



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Lyapunov-control-inspired strategies for quantum combinatorial optimization

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

Phys. Rev. A **106**, 062414 — Published 13 December 2022

DOI: [10.1103/PhysRevA.106.062414](https://doi.org/10.1103/PhysRevA.106.062414)

Lyapunov control-inspired strategies for quantum combinatorial optimization

Alicia B. Magann,^{1,2,3} Kenneth M. Rudinger,² Matthew D. Grace,¹ and Mohan Sarovar¹

¹*Quantum Algorithms and Applications Collaboratory,
Sandia National Laboratories, Livermore, California 94550, USA*

²*Quantum Algorithms and Applications Collaboratory,
Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*

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

(Dated: October 6, 2022)

The prospect of using quantum computers to solve combinatorial optimization problems via the quantum approximate optimization algorithm (QAOA) has attracted considerable interest in recent years. However, a key limitation associated with QAOA is the need to classically optimize over a set of quantum circuit parameters. This classical optimization can have significant associated costs and challenges. Here, we provide an expanded description of Lyapunov control-inspired strategies for quantum optimization, as first presented in [arXiv:2103.08619](https://arxiv.org/abs/2103.08619), that do not require any classical optimization effort. Instead, these strategies utilize feedback from qubit measurements to assign values to the quantum circuit parameters in a deterministic manner, such that the combinatorial optimization problem solution improves monotonically with the quantum circuit depth. Numerical analyses are presented that investigate the utility of these strategies towards MaxCut on weighted and unweighted 3-regular graphs, both in ideal implementations and also in the presence of measurement noise. We also discuss how these strategies compare with QAOA, how they may be used to seed QAOA optimizations in order to improve performance for near-term applications, and explore connections to quantum annealing.

I. INTRODUCTION

Combinatorial optimization problems have a variety of broad and high-value applications, including in routing and scheduling problems [1, 2]. The desire to use quantum resources to aid in solving them has a long history, spanning the development of adiabatic and annealing-based strategies [3–5], as well as the development of early quantum algorithms [6, 7]. More recently, the quantum approximate optimization algorithm (QAOA) [8] was proposed in 2014, as a method for leveraging quantum computers to solve combinatorial optimization problems. In particular, QAOA is a method for determining an approximate solution to a combinatorial optimization problem by using a hybrid quantum-classical framework; a classical computer is utilized to iteratively minimize the value of the cost function, and the cost function is evaluated on a quantum computer using a parameterized quantum circuit. Since its development, QAOA has captured the attention of numerous theoretical and experimental groups, e.g., [9–14], particularly as a potential application of noisy, intermediate-scale quantum (NISQ) [15] devices.

Recently, numerous connections have been made between QAOA and quantum optimal control (QOC) [16], which is a strategy for identifying the controls needed to steer the dynamics of a quantum system in a desired manner by iteratively optimizing over a set of control functions or parameters [17, 18]. Certain connections have rested on the control-theoretic notion of controllability, which implies that QOC solutions can be found for driving the dynamics of a system under consideration towards arbitrary objectives [19–26]. For instance,

controllability considerations have recently been applied to show that QAOA can be computationally universal in certain circumstances [27, 28], and to assess the number of QAOA quantum circuit parameters needed to achieve controllability [29].

A key challenge in identifying QAOA and QOC solutions is the difficulty of searching for the optimal QAOA and QOC parameters, respectively. If the number of QAOA layers, and correspondingly variational parameters, can be limited to $O(\text{poly}(\log(n)))$, then the complexity scaling of the optimization can be polynomial in the number of qubits, n . However, this scaling belies the difficulty of such optimizations – the fact is that QAOA, and most variational algorithms, require optimization over non-linear, stochastic cost functions with derivative information that is noisy and hard to obtain. Scaling classical optimization to thousands of parameters is challenging in this context, e.g., [30].

Even in the absence of noise, the difficulty of identifying the optimal parameters is determined in large part by the structure of the optimization landscape, and landscape features such as local optima, saddle points, and barren plateaus can complicate and hinder the optimization process [31–38]. However, we note that a variety of alternate quantum control frameworks, including quantum tracking control [39–43], quantum Lyapunov control [44–52], and quantum feedback control [53–56], have been developed that do not rely on the iterative classical optimization procedure inherent to QOC and thus do not share these challenges.

Here, we explore a new connection between quantum algorithms and quantum control theory, and develop strategies for quantum combinatorial optimization.

tion inspired by the theory of *quantum Lyapunov control* (QLC). In particular, this article provides an expanded discussion of the details of these strategies, beyond that contained in [57]. Importantly, these QLC-inspired strategies do not involve any classical optimization. Instead, they use measurement-based feedback to assign values to the quantum circuit parameters. We show that this feedback-based procedure yields a monotonically improving solution to the original combinatorial optimization problem, with respect to the depth of the quantum circuit.

The remainder of this article is organized as follows. We begin by providing background on QAOA, as motivation for this work. This is followed by a description of certain aspects of QLC. We then introduce a Feedback-based ALgorithm for Quantum Optimization (FALQON) inspired by QAOA and QLC, and discuss extensions that can be used to improve performance, including the addition of a reference perturbation, the implementation of an iterative procedure, and the introduction of additional control functions. We then discuss applications of FALQON towards solving the MaxCut problem. To this end, we provide numerical illustrations of the ideal performance of FALQON towards solving MaxCut on 3-regular graphs, and we also explore how it performs in the presence of measurement noise, and how its performance compares to that of QAOA. We go on to discuss how FALQON can be used to boost the performance of QAOA in NISQ applications, and explore connections to quantum annealing. We conclude with an outlook.

II. QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

Combinatorial optimization problems are concerned with identifying configurations of discrete optimization variables that best achieve one or multiple goals, as quantified by an associated cost function C . The quantum approximate optimization algorithm (QAOA) [8] is an approach for finding or approximating solutions to combinatorial optimization problems using quantum computers. It operates by first encoding the cost function C into an Ising Hamiltonian, H_p , which is diagonal in the quantum computational basis such that each eigenstate of H_p corresponds to a single spin configuration, which encodes a configuration of the associated optimization variables [58]. The encoding is done such that the best solution to the combinatorial optimization problem is encoded in the ground state of H_p .

Using this encoding, QAOA is a hybrid quantum-classical algorithm for solving

$$\min_{\{\gamma_k\}, \{\beta_k\}} \langle \psi(\{\gamma_k\}, \{\beta_k\}) | H_p | \psi(\{\gamma_k\}, \{\beta_k\}) \rangle. \quad (1)$$

To do so, a classical computer is utilized to iteratively minimize the value of the objective function, evaluated

as $\langle \psi(\{\gamma_k\}, \{\beta_k\}) | H_p | \psi(\{\gamma_k\}, \{\beta_k\}) \rangle$, and optimized over the set of 2ℓ parameters $\{\gamma_k\}_{k=1}^{\ell}$ and $\{\beta_k\}_{k=1}^{\ell}$.

The objective function value is determined at each iteration using a quantum computer, which prepares the multiqubit state $|\psi(\{\gamma_k\}, \{\beta_k\})\rangle$ using a parameterized quantum circuit of the form

$$U_{\text{QAOA}} = U_d(\beta_\ell) U_p(\gamma_\ell) \cdots U_d(\beta_1) U_p(\gamma_1), \quad (2)$$

such that $|\psi(\{\gamma_k\}, \{\beta_k\})\rangle = U_{\text{QAOA}} |\psi_0\rangle$ for an initial state $|\psi_0\rangle$. The elements of the QAOA circuit, $U_p(\cdot)$ and $U_d(\cdot)$, are created by simulating an evolution under H_p and under a *driver* Hamiltonian H_d , which is chosen to not commute with H_p . Following the implementation of U_{QAOA} , measurements of H_p in the multiqubit state $|\psi(\{\gamma_k\}, \{\beta_k\})\rangle$ allow for estimating the value of the objective function at each iteration of the classical optimization algorithm.

III. QUANTUM LYAPUNOV CONTROL

Quantum Lyapunov control (QLC) is a local-in-time method for identifying controls to asymptotically steer the dynamics of a quantum system towards a desired objective [44–52]. The controls are identified utilizing a feedback law, which is derived from a suitable control Lyapunov function [59], chosen to capture the target objective. In this section, we describe the theory of QLC and outline certain results from the literature pertaining to its asymptotic convergence behavior. We begin by considering a quantum system whose dynamics are governed by

$$i \frac{d}{dt} |\psi(t)\rangle = (H_p + H_d \beta(t)) |\psi(t)\rangle \quad (3)$$

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, and the latter couples a scalar, time-dependent control function $\beta(t)$ to the system. In this article, we choose our QLC objective to be the minimization of $\langle H_p \rangle = \langle \psi(t) | H_p | \psi(t) \rangle$, and thus seek a QLC strategy for designing $\beta(t)$ to accomplish this. We proceed by defining a Lyapunov function

$$E_p(|\psi(t)\rangle) = \langle \psi(t) | H_p | \psi(t) \rangle \quad (4)$$

to capture our QLC objective. Then, to minimize E_p we seek to design $\beta(t)$ such that the QLC condition

$$\frac{d}{dt} E_p \leq 0, \quad \forall t \geq 0 \quad (5)$$

is satisfied. There is significant flexibility in choosing $\beta(t)$ to satisfy Eq. (5). Namely, given that

$$\begin{aligned} \frac{dE_p}{dt} &= \langle \psi(t) | i[H_p + \beta(t)H_d, H_p] | \psi(t) \rangle \\ &= \langle \psi(t) | i[\beta(t)H_d, H_p] | \psi(t) \rangle \\ &= \langle \psi(t) | i[H_d, H_p] | \psi(t) \rangle \beta(t) \\ &= A(t)\beta(t), \end{aligned} \quad (6)$$

where

$$A(t) \equiv \langle \psi(t) | i[H_d, H_p] | \psi(t) \rangle, \quad (7)$$

we may take

$$\beta(t) = -w f(t, A(t)), \quad (8)$$

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 $A(t) \neq 0$ [60]. This choice of $\beta(t)$ guarantees that E_p will decrease monotonically over time. When $\beta(t)$ is chosen according to Eq. (8), the system dynamics are governed by

$$i \frac{d}{dt} |\psi(t)\rangle = (H_p + H_d \beta(t, A(t))) |\psi(t)\rangle, \quad (9)$$

which are highly nonlinear, due to the dependence of $\beta(t)$ on the state $|\psi(t)\rangle$ via $A(t)$.

We refer to Eq. (8) as a *feedback law*, as it relies on feedback in order to evaluate the observable expectation value $A(t)$. Conventionally, QLC laws like Eq. (8) are used in simulations to design open-loop control laws; that is, they cannot be applied directly in experiments as-is, as the destructive, real-time measurements required to estimate $A(t)$ would lead to a collapse of the system state. This distinguishes QLC from real-time feedback control.

Ideally, designing β per Eq. (8) would result in asymptotic convergence to the global minimum of E_p , and it has been shown that this behavior can be guaranteed when a set of sufficient conditions are met [50, 60–62]. However, these conditions are very stringent (see Appendix A). When the protocol we construct in Sec. IV is applied to the MaxCut problem, per Sec. V, they are not satisfied. Nonetheless, asymptotic convergence to the global minimum can still be obtained in such settings (e.g., as illustrated in Fig. 1), and in situations where convergence is not obtained, a variety of techniques can be employed to improve control performance, as discussed in the following subsections.

A. Inclusion of reference perturbation in $\beta(t)$

The inclusion of a reference perturbation in $\beta(t)$ can improve the likelihood of asymptotic convergence to the global minimum of E_p [60–62]. Here, we consider the inclusion of a reference perturbation $\lambda(t)$ such that the time-dependent system Hamiltonian is given by

$$H(t) = H_p + (\lambda(t) + \beta(t))H_d. \quad (10)$$

Inspecting Eq. (10), we may define system (a) as a system with drift Hamiltonian H_p , control Hamiltonian H_d , and control function $(\lambda(t) + \beta(t))$. Meanwhile, we may define system (b) as a perturbed system with time-dependent drift Hamiltonian

$$H_{p,(b)}(t) \equiv H_p + \lambda(t)H_d, \quad (11)$$

control Hamiltonian H_d , and control function $\beta(t)$. Within (b), we may define the perturbed Lyapunov function

$$E_{p,(b)}(|\psi(t)\rangle) = \langle \psi(t) | H_{p,(b)}(t) | \psi(t) \rangle, \quad (12)$$

and seek a control law that will ensure

$$\frac{dE_{p,(b)}}{dt} \leq 0, \quad (13)$$

while at the same time, ideally improving convergence to the minimum of our original objective E_p . At this stage, it is important to note that in practice, $\lambda(t)$ can be defined explicitly as a desired function of t , or implicitly as a function of $|\psi(t)\rangle$, E_p , or $E_{p,(b)}$. For practical reasons, we restrict our attention to the former case; for further details on the latter, we refer to refs. [61, 62].

When $\lambda(t)$ is defined as an explicit function of time, the left-side of Eq. (13) is given, conveniently, by,

$$\begin{aligned} \frac{d}{dt} E_{p,(b)} &= \langle \psi(t) | i[H(t), H_{p,(b)}(t)] | \psi(t) \rangle \\ &= \langle \psi(t) | i[H_{p,(b)}(t) + \beta(t)H_d, H_{p,(b)}(t)] | \psi(t) \rangle \\ &= \langle \psi(t) | i[\beta(t)H_d, H_{p,(b)}(t)] | \psi(t) \rangle \\ &= \langle \psi(t) | i[H_d, H_p] | \psi(t) \rangle \beta(t) \\ &= A(t)\beta(t) \end{aligned} \quad (14)$$

which implies that even after the inclusion of the reference perturbation $\lambda(t)$ and the definition of a perturbed Lyapunov function $E_{p,(b)}$, we may nonetheless define the control law for $\beta(t)$ as before, as $\beta(t) = -w f(t, A(t))$, to ensure Eq. (13) is satisfied.

Within this framework, if system (b) converges to the ground state of $H_{p,(b)}(t)$ at a terminal time $t = T$, and $\lambda(T) = 0$, then system (b) becomes system (a), such that $H_{p,(b)}(T) = H_p$, meaning that the ground state of $H_{p,(b)}(t)$ is also the ground state of H_p , and convergence to the desired state has been obtained. As such, it is often practical to select $\lambda(t)$ to be a slowly-varying function that tends to 0 as $t \rightarrow T$.

B. Iterative quantum Lyapunov control

Another technique to improve the likelihood of asymptotic convergence to the minimum of E_p is to use an iterative procedure for refining the QLC control function $\beta(t)$ [51]. We emphasize that the iterative QLC procedure outlined in this section is conceptually distinct from the iterative optimization procedure utilized in QOC, as the iterations involved do not involve updates determined by an optimization routine. Instead, $\beta(t)$ is updated using a QLC-derived control law, in the following manner.

We begin by considering a system initialized as $|\psi(t=0)\rangle = |\psi_0\rangle$, and then design a control field $\beta^{(0)}(t)$ to control E_p using QLC, as per the control law of Eq. (8), over some fixed time interval $t \in [0, T]$. We denote the trajectory of E_p over this time interval by $E_p^{(0)}(t)$, and denote

the associated state by $|\psi^{(0)}(t)\rangle$. Subsequent steps are then carried out as follows. For iterations $j \geq 1$, $\beta^{(j-1)}(t)$ serves as a reference perturbation, as denoted by $\lambda(t)$ in Sec. III A. Then, $\beta^{(j-1)}(t)$, $E_p^{(j-1)}(t)$, and $|\psi^{(j-1)}(t)\rangle$ all describe the dynamics of a perturbed system (b), whose time-dependent Hamiltonian is

$$H_{p,(b)}^{(j-1)}(t) = H_p + \beta^{(j-1)}(t)H_d. \quad (15)$$

A new QLC field $\tilde{\beta}^{(j)}(t)$ is then determined for $t \in [0, T]$ using the framework in Sec. III A, where a perturbed Lyapunov function based on $H_{p,(b)}^{(j-1)}$ is utilized, and $\tilde{\beta}^{(j)}(t)$ is chosen according to Eq. (8). After $\tilde{\beta}^{(j)}(t)$ has been computed for $t \in [0, T]$, the update rule is given by

$$\beta^{(j)}(t) = \beta^{(j-1)}(t) + \tilde{\beta}^{(j)}(t). \quad (16)$$

For T chosen to be large enough for the perturbed system to converge to the unperturbed system at each iteration, i.e., such that $\beta^{(j)}(T) = 0$, causing $H_{p,(b)}^{(j)}(T) = H_p$, $\forall j$, this procedure guarantees a monotonic improvement of $E_p(T)$ with respect to iteration, as per,

$$\begin{aligned} \langle \psi^{(j)}(T) | H_p | \psi^{(j)}(T) \rangle &= \langle \psi^{(j)}(T) | H_{p,(b)}^{(j-1)}(T) | \psi^{(j)}(T) \rangle \\ &\leq \langle \psi^{(j)}(0) | H_{p,(b)}^{(j-1)}(0) | \psi^{(j)}(0) \rangle \\ &= \langle \psi_0 | H_{p,(b)}^{(j-1)}(0) | \psi_0 \rangle \\ &= \langle \psi^{(j-1)}(0) | H_{p,(b)}^{(j-1)}(0) | \psi^{(j-1)}(0) \rangle \\ &= \langle \psi^{(j-1)}(T) | H_{p,(b)}^{(j-1)}(T) | \psi^{(j-1)}(T) \rangle \\ &= \langle \psi^{(j-1)}(T) | H_p | \psi^{(j-1)}(T) \rangle, \end{aligned} \quad (17)$$

such that $E_p^{(j)}(T) \leq E_p^{(j-1)}(T)$ [51]. We note that in line 5 of the above, we have utilized the fact that $\frac{d}{dt} \langle \psi^{(j-1)}(t) | H_{p,(b)}^{(j-1)}(t) | \psi^{(j-1)}(t) \rangle = 0$, due to the fact that $|\psi^{(j-1)}(t)\rangle$ evolves under $H_{p,(b)}^{(j-1)}(t)$.

C. Extensions to multiple control functions

The framework of QLC can be extended in a straightforward manner to settings with multiple control functions, i.e., where the system Hamiltonian is given by

$$H(t) = H_p + \sum_j \beta(j, t) H_{d,j}, \quad (18)$$

where $\beta(j, t)$ denotes the value of the control function that scales the j -th control Hamiltonian $H_{d,j}$ at time t . Then, in order to satisfy the QLC condition that $\frac{d}{dt} E_p \leq$

0, we see that

$$\begin{aligned} \frac{d}{dt} \langle \psi(t) | H_p | \psi(t) \rangle &= \langle \psi(t) | i [H_p + \sum_j \beta(j, t) H_{d,j}, H_p] | \psi(t) \rangle \\ &= \langle \psi(t) | i [\sum_j \beta(j, t) H_{d,j}, H_p] | \psi(t) \rangle \\ &= \sum_j \langle \psi(t) | i [H_{d,j}, H_p] | \psi(t) \rangle \beta(j, t) \\ &= \sum_j A(j, t) \beta(j, t), \end{aligned} \quad (19)$$

where

$$A(j, t) \equiv \langle \psi(t) | i [H_{d,j}, H_p] | \psi(t) \rangle. \quad (20)$$

As such, the following control laws may be used:

$$\beta(j, t) = -w_j f_j(t, A(j, t)), \quad \forall j \quad (21)$$

to ensure that E_p decreases monotonically over time. We remark that in cases with multiple control functions, reference perturbations may be included for any $\beta(j, t)$, following the framework outlined in Sec. III A, and iterative QLC schemes can also be used, following the framework of Sec. III B.

IV. FEEDBACK-BASED ALGORITHM FOR QUANTUM OPTIMIZATION

We now consider how the QLC framework outlined in Sec. III (eqs. (5) – (8)) can be translated into FALQON, a feedback-based algorithm for minimizing the expectation value of a Hamiltonian, that can be implemented on quantum devices. To this end, we now assume that H_p is the problem Hamiltonian that encodes a combinatorial optimization problem of interest, noting that when defined this way,

$$E_p = \langle \psi(t) | H_p | \psi(t) \rangle \quad (22)$$

may not meet all of the requirements of being a true Lyapunov function, such as positive definiteness. Now, without loss of generality, we consider alternating, rather than concurrent, applications of H_p and H_d , such that the state $|\psi(t)\rangle$ undergoes a time evolution of the form

$$U = U_d(\beta_\ell) U_p \cdots U_d(\beta_1) U_p, \quad (23)$$

where

$$U_d(\beta_k) \equiv e^{-i\beta_k H_d \Delta t} \quad (24)$$

and

$$U_p \equiv e^{-iH_p \Delta t} \quad (25)$$

and $\beta_k = \beta((2k-1)\Delta t)$ for $k = 1, 2, \dots, \ell$, such that after each period of Δt , the Hamiltonian that is applied

alternates between H_p and H_d . For small Δt , this yields a Trotterized approximation to the time evolution that would be achieved in Eq. (3). To ensure that Eq. (5) is satisfied, we may again define β from Eq. (8). In this work, we use

$$w = 1, \quad f(t, A(t)) = A(t), \quad (26)$$

such that in the alternating framework

$$\beta_{k+1} = -A_k, \quad (27)$$

where $A_k = \langle \psi_k | i[H_d, H_p] | \psi_k \rangle$, and $|\psi_k\rangle = |\psi(2k\Delta t)\rangle$. Importantly, we note that it is always possible to select Δt small enough such that Eq. (5) is satisfied when the control law in Eq. (27) is used (see Sec. IV A). However, if Δt is chosen to be too large, the condition in Eq. (5) can be violated.

The implementation of this alternating procedure on a qubit device can be accomplished according to the steps in Algorithm 1. The preliminary step is to seed the procedure by setting $\beta_1 = \beta_{\text{init}}$, and we use $\beta_{\text{init}} = 0$. Then, a set of qubits are initialized in a fixed initial state $|\psi_0\rangle$, and a single circuit “layer” is implemented to prepare the state

$$|\psi_1\rangle = U_d(\beta_1)U_p|\psi_0\rangle. \quad (28)$$

Next, the qubits are then measured in order to estimate A_1 . This can be accomplished by expanding A_1 in the Pauli operator basis as

$$A_1 = \langle \psi_1 | i[H_d, H_p] | \psi_1 \rangle = \sum_{j=1}^N \alpha_j \langle \psi_1 | P_j | \psi_1 \rangle, \quad (29)$$

where α_j are scalar coefficients and P_j are Pauli basis operators. We note that the number of Pauli basis operators N in the expansion depends on the structure of H_p and H_d (see Eq. (48) below). Each P_j can then be measured, and the measurements can be repeated to collect sufficiently many samples to estimate the associated expectation values. Finally, the resultant expectation values of each P_j can then be used to evaluate the weighted sum in Eq. (29) to estimate A_1 . Following this, the result is “fed back” to set $\beta_2 = -A_1$ (or, more precisely, β_2 is set to be the negative of the approximation of A_1).

For subsequent steps $k = 2, \dots, \ell$, the same procedure is repeated: the qubits are initialized in the state $|\psi_0\rangle$, after which k layers are applied to obtain

$$|\psi_k\rangle = U_d(\beta_k)U_p \cdots U_d(\beta_1)U_p|\psi_0\rangle. \quad (30)$$

Then, the qubits are measured to estimate A_k using the same procedure described above, and the result is fed back to set the value of β_{k+1} . By design, this procedure causes $\langle H_p \rangle$ to decrease layer-by-layer as per

$$\langle \psi_1 | H_p | \psi_1 \rangle \geq \langle \psi_2 | H_p | \psi_2 \rangle \geq \cdots \geq \langle \psi_\ell | H_p | \psi_\ell \rangle, \quad (31)$$

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 $\langle H_p \rangle$ converges (i.e., stops decreasing), as determined via measurements, or when a threshold number of layers ℓ is reached. At that point, the set of β values $\{\beta_k\}_{k=1}^\ell$ is recorded as the output.

After Algorithm 1 concludes, the set $\{\beta_k\}_{k=1}^\ell$ can subsequently be used to prepare the state $|\psi_\ell\rangle$ in post-processing steps as needed, e.g., in order to estimate the value of $\langle H_p \rangle$, by measuring $|\psi_\ell\rangle$ and repeating the experiment enough times to ensure reliable statistics. In addition, the associated ℓ -layer quantum circuit can also be implemented in order to estimate the bit string $z = z_1 z_2 \cdots z_n$ associated with the best candidate solution to the underlying combinatorial optimization problem. This latter task can be accomplished by sampling the bit string $z_1 z_2 \cdots z_n$ from the output distribution associated with the output state $|\psi_\ell\rangle$, i.e., by measuring $z_j = \langle \psi_\ell | Z_j | \psi_\ell \rangle$ for $j = 1, \dots, n$ and then concatenating the results to form

$$z = z_1 z_2 \cdots z_n, \quad (32)$$

where Z_j denotes the Pauli operator acting on qubit j [63]. After collecting a set of samples, the bit string associated with the best solution to the combinatorial optimization problem, i.e., the bit string z that returns the minimum value of the associated cost function $C(z)$, should be saved as the best approximate solution to the combinatorial optimization problem of interest.

Algorithm 1 FALQON

- 1: **set** $H_p, H_d, \Delta t, \ell, |\psi_0\rangle$
 - 2: Seed the procedure
 $\beta_1 \leftarrow 0$
 - 3: Initialize the qubits
 $|\psi\rangle \leftarrow |\psi_0\rangle$
 - 4: Implement 1 layer
 $|\psi_1\rangle \leftarrow U_d(\beta_1)U_p|\psi_0\rangle$
 - 5: Estimate the value of A_1 by measuring the qubits in the state $|\psi_1\rangle$ and repeating the experiment enough times to ensure reliable statistics.
 - 6: $\beta_2 \leftarrow -A_1$
 - 7: $k \leftarrow 1$
 - 8: **while** $k < \ell$ **do**
 - 9: $k \leftarrow k + 1$
 - 10: Initialize the qubits
 $|\psi\rangle \leftarrow |\psi_0\rangle$
 - 11: Implement k layers
 $|\psi_k\rangle \leftarrow U_d(\beta_k)U_p \cdots U_d(\beta_1)U_p|\psi_0\rangle$
 - 12: Estimate the value of A_k by measuring the qubits in the state $|\psi_k\rangle$ and repeating the experiment enough times to ensure reliable statistics.
 - 13: $\beta_{k+1} \leftarrow -A_k$
 - 14: **end while**
 - 15: **output** $\{\beta_k\}_{k=1}^\ell$
-

We note that FALQON has similarities to other quantum circuit parameter-setting protocols that involve “greedy”, layer-by-layer optimization, e.g., where a classical optimization routine is used to sequentially optimize quantum circuit parameters in order to minimize

a cost function in a layer-wise manner [64–67]. In fact, the parameter-setting rule given in Eq. (27) corresponds to taking a step “down” in the direction of the local gradient $\frac{d}{d\beta_k} E_p$ with a step size of Δt , thus suggesting that there is a natural connection between FALQON and layer-wise circuit optimization methods that proceed by gradient descent [68]. We also remark that the ADAPT-QAOA approach developed in ref. [69] has certain similarities to FALQON, e.g., it also utilizes information about $\frac{d}{dt} \langle \psi(t) | H_p | \psi(t) \rangle$ to step forward from layer to layer. However, their stepping procedure involves selecting from a set of driver Hamiltonians, while still containing a classical optimization loop.

Having outlined FALQON, we now turn to the prospect of boosting its performance using the techniques outlined in Secs. (III A), (III B), and (III C). We begin by discussing how a reference perturbation may be introduced, as per Eq. (10) and how the framework outlined in Sec. (III A) may be adapted to the quantum device setting. As before, this can be accomplished by simply “Trotterizing” Eq. (10), and implementing a quantum circuit with the form

$$U_d(\nu_\ell) U_p \cdots U_d(\nu_1) U_p U_d(\nu_0) U_p \quad (33)$$

where $\nu_k = \lambda_k + \beta_k$, where λ_k is the value of the reference perturbation at the k -th layer and β_k is the value of the control parameter at the k -th layer, determined via $\beta_k = A_{k-1}$. Numerical illustrations showing how the performance of FALQON can be improved with the inclusion of a reference perturbation can be found in [57]. In a similar fashion, the iterative QLC procedure discussed in Sec. III B can also be adapted the context of quantum optimization, in order to successively improve the quality of the solutions obtained. After applying a first-order Trotter decomposition, the circuits at the j -th iteration will have the structure

$$U_d(\beta_\ell^{(j)}) U_p \cdots U_d(\beta_1^{(j)}) U_p U_d(\beta_0^{(j)}) U_p, \quad (34)$$

where $\beta_k^{(j)} = \beta_k^{(j-1)}(t) + \tilde{\beta}_k^{(j)}(t)$ for $k = 0, \dots, \ell$, and the iterative procedure is seeded by determining $\beta^{(0)}$ via Algorithm 1. Finally, the approach discussed in Sec. (III C) may also be extended to the quantum device setting, by modifying the layered quantum circuit structure to include evolutions under additional driver Hamiltonians.

A. Selecting Δt

We now consider the selection of the time step Δt in order to ensure that Eq. (5) will hold. To this end, we consider a single layer of FALQON, such that

$$\begin{aligned} E_{p,t+2\Delta t} &= \langle \psi(t+2\Delta t) | H_p | \psi(t+2\Delta t) \rangle \\ &= \langle \psi(t) | e^{iH_p \Delta t} e^{iH_d \beta_t \Delta t} H_p e^{-iH_d \beta_t \Delta t} e^{-iH_p \Delta t} | \psi(t) \rangle. \end{aligned} \quad (35)$$

For the following, we adopt the notation $\langle \cdot \rangle_t \equiv \langle \psi(t) | \cdot | \psi(t) \rangle$. We express each of the exponentials above using a Taylor series expansion:

$$\begin{aligned} E_{p,t+2\Delta t} &= E_{p,t} + i \langle [H_d, H_p] \rangle_t \beta_t \Delta t - \langle [H_p, [H_d, H_p]] \rangle_t \beta_t \Delta t^2 \\ &\quad - \langle [H_d, [H_d, H_p]] \rangle_t \beta_t^2 \Delta t^2 + O(\Delta t^3) \\ &= E_{p,t+2\Delta t}^{(0)} + E_{p,t+2\Delta t}^{(1)} + E_{p,t+2\Delta t}^{(2)} + \cdots \end{aligned} \quad (36)$$

where the superscripts label the orders of Δt . Since we only use β to design the first order term, we would like this to be the dominant term in the expansion such that

$$|E_{p,t+2\Delta t}^{(1)}| > \left| \sum_{k=2}^{\infty} E_{p,t+2\Delta t}^{(k)} \right|. \quad (37)$$

This way, designing $E_{p,t+2\Delta t}^{(1)}$ appropriately will enforce that $E_{p,t+2\Delta t}$ decreases. The left-side of Eq. (37) is given by

$$\begin{aligned} |E_{p,t+2\Delta t}^{(1)}| &= |\langle [H_d, H_p] \rangle_t| |\beta_t| \Delta t \\ &= |A_\ell| |\beta_t| \Delta t. \end{aligned} \quad (38)$$

Meanwhile, the magnitude of higher-order terms such as $E_{p,t+2\Delta t}^{(2)}$ can be bounded as,

$$\begin{aligned} |E_{p,t+2\Delta t}^{(2)}| &= |\langle [H_p, [H_d, H_p]] \rangle_t \beta_t + \langle [H_d, [H_d, H_p]] \rangle_t \beta_t^2| \Delta t^2 \\ &\leq (|\langle [H_p, [H_d, H_p]] \rangle_t| \\ &\quad + |\langle [H_d, [H_d, H_p]] \rangle_t| |\beta_t|) |\beta_t| \Delta t^2 \\ &\leq (\| [H_p, [H_d, H_p]] \| + \| [H_d, [H_d, H_p]] \| |\beta_t|) |\beta_t| \Delta t^2 \\ &\leq (2 \| H_p H_d H_p \| + \| H_p H_p H_d \| + \| H_d H_p H_p \| \\ &\quad + \| H_d H_d H_p \| + \| H_p H_d H_d \| \\ &\quad + 2 \| H_d H_p H_d \| |\beta_t|) |\beta_t| \Delta t^2 \\ &\leq (2 \| H_p \|^2 \| H_d \| + 2 \| H_d \|^2 \| H_p \| |\beta_t|) 2 |\beta_t| \Delta t^2 \\ &= 2 n_p n_d |\beta_t| (2 n_p + 2 n_d |\beta_t|) \Delta t^2, \end{aligned} \quad (39)$$

where in the last line we introduce the abbreviated notation $n_d \equiv \| H_d \|$ and $n_p \equiv \| H_p \|$. Expressions for the magnitude of any higher-order (i.e., $k \geq 2$) term can be found following the same procedure, which results in the following general expression at k -th order:

$$|E_{p,t+2\Delta t}^{(k)}| = 2 n_p n_d |\beta_t| (2 n_p + 2 n_d |\beta_t|)^{k-1} \Delta t^k. \quad (40)$$

Given Eq. (40), the right side of Eq. (37) can be

bounded by

$$\begin{aligned}
\left| \sum_{k=2}^{\infty} E_{\mathbf{p},t+2\Delta t}^{(k)} \right| &\leq \sum_{k=2}^{\infty} |E_{\mathbf{p},t+2\Delta t}^{(k)}| \\
&\leq 2n_{\mathbf{p}}n_{\mathbf{d}}|\beta_t| \sum_{k=2}^{\infty} (2n_{\mathbf{p}} + 2n_{\mathbf{d}}|\beta_t|)^{k-1} \Delta t^k \\
&= \frac{n_{\mathbf{p}}n_{\mathbf{d}}|\beta_t|}{n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|} \left(\sum_{k=0}^{\infty} (2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|))^k \right. \\
&\quad \left. - 1 - 2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|) \right). \tag{41}
\end{aligned}$$

For $2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|) < 1$ the geometric series converges. Under this assumption, we can rewrite the condition in Eq. (37) as

$$\begin{aligned}
|A_t||\beta_t| \Delta t &> \frac{n_{\mathbf{p}}n_{\mathbf{d}}|\beta_t|}{n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|} \left(\frac{1}{1 - 2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|)} \right. \\
&\quad \left. - 1 - 2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|) \right). \tag{42}
\end{aligned}$$

We rearrange this equation to obtain a bound for Δt :

$$|\Delta t| < \frac{|A_t|}{2(2n_{\mathbf{d}}n_{\mathbf{p}} + |A_t|)(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|)}. \tag{43}$$

For these values of Δt , we can confirm that the geometric series in Eq. (41) does converge, because

$$\begin{aligned}
1 &> 2\Delta t(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|) \\
&= \frac{2|A_t|(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|)}{2(2n_{\mathbf{d}}n_{\mathbf{p}} + |A_t|)(n_{\mathbf{p}} + n_{\mathbf{d}}|\beta_t|)} \\
&= \frac{|A_t|}{2n_{\mathbf{d}}n_{\mathbf{p}} + |A_t|}
\end{aligned} \tag{44}$$

is always satisfied. Therefore, if Δt is selected according to Eq. (43), it is ensured that the QLC condition in Eq. (5) will hold, and thus, that $E_{\mathbf{p}}$ will decrease monotonically as a function of layer as desired. In practice we find that Δt can be chosen to be much larger than the value in Eq. (43) due to the looseness of the bound.

V. APPLICATIONS TO MAXCUT

We now consider applications of FALQON towards the combinatorial optimization problem MaxCut, which aims to identify a graph partition that maximizes the number of edges that are cut. For a graph \mathcal{G} , with n nodes and edge set \mathcal{E} , the MaxCut problem Hamiltonian is defined on n qubits as

$$H_{\mathbf{p}} = - \sum_{i,j \in \mathcal{E}} \frac{1}{2} (1 - w_{ij} Z_i Z_j), \tag{45}$$

where w_{ij} denote the edge weights, and for unweighted graphs, $w_{ij} = 1$ for all $i, j \in \mathcal{E}$. In our analyses involving weighted graphs, we consider random edge weights w_{ij} drawn from a uniform distribution between 0 and 2, such that the average edge weight is $\bar{w} = 1$, matching the unweighted case. Furthermore, we consider $H_{\mathbf{d}}$ to have the standard form

$$H_{\mathbf{d}} = \sum_{j=1}^n X_j. \tag{46}$$

Given these choices for $H_{\mathbf{p}}$ and $H_{\mathbf{d}}$,

$$\begin{aligned}
A_k &= \langle \psi_k | i[H_{\mathbf{d}}, H_{\mathbf{p}}] | \psi_k \rangle \\
&= \sum_{i,j \in \mathcal{E}} w_{ij} (\langle \psi_k | Y_i Z_j | \psi_k \rangle + \langle \psi_k | Z_i Y_j | \psi_k \rangle), \tag{47}
\end{aligned}$$

where X_j and Y_j denote the Pauli operators acting on qubit j . As such, evaluating the feedback law $\beta_{k+1} = -A_k = -\langle \psi_k | i[H_{\mathbf{d}}, H_{\mathbf{p}}] | \psi_k \rangle$ requires measuring the expectation values of

$$N \leq n(n-1) \tag{48}$$

Pauli basis operators (i.e., in this case, two-qubit Pauli strings), where the exact value of N depends on the structure of the graph under consideration.

In order to assess the performance of FALQON towards MaxCut, we consider two figures of merit: the approximation ratio,

$$r_A = \frac{\langle H_{\mathbf{p}} \rangle}{\langle H_{\mathbf{p}} \rangle_{\min}}, \tag{49}$$

which is proportional to the original Lyapunov function $E_{\mathbf{p}}$, and the success probability of measuring the (potentially) degenerate ground state,

$$\phi = \sum_i |\langle \psi | q_{0,i} \rangle|^2, \tag{50}$$

which gives the probability of obtaining the global minimum solution to the original combinatorial optimization problem. Each of these two figures of merit can take on values between 0 and 1, where $r_A = \phi = 1$ corresponds to the optimal (i.e., ground state) solution.

A. Numerical illustrations on 3-regular graphs

We now examine the performance of FALQON towards MaxCut on 3-regular graphs via a series of numerical illustrations (for a demonstration in quantum hardware, see our companion article [57]). We consider both weighted and unweighted 3-regular graphs with $n \in \{8, 10, \dots, 20\}$ vertices. For weighted graphs, the edge weights are randomly sampled from a uniform distribution over $(0, 2)$. For graphs with $n = 8, 10$ vertices, we consider the set of all nonisomorphic, connected 3-regular graphs. For each value of $n = 12, \dots, 20$, we

consider a set of 50 randomly-generated nonisomorphic graphs. The qubits are initialized in the ground state of H_d , and the performance of FALQON is quantified using the mean (over the set of graphs) of r_A , and ϕ . We relate the performance to two reference values: $r_A = 0.932$, corresponding to the highest approximation ratio that can currently be guaranteed using a classical approximation algorithm for MaxCut on unweighted, 3-regular graphs (i.e., the algorithm of Goemans and Williamson [70]), and $\phi = 0.25$, which implies that on average, 4 circuit repetitions will be needed in order to obtain a sample bit string corresponding to the ground state. The results are collected in Fig. 1.

In Fig. 1(a) and (b), the r_A and ϕ results are shown for graphs with $n = 8, 10, \dots, 16$ vertices, and the associated reference values are plotted in black. The solid curves show the results for unweighted graphs, while the dotted curves of corresponding color show the results for weighted graphs. We find that FALQON has superior performance on unweighted graphs, given that r_A and ϕ appear to converge to higher values. Nonetheless, FALQON does lead to monotonic convergence towards high r_A values as a function of layer for weighted graphs as well. However, for weighted graphs, we find many instances where ϕ fails to converge to 1, as shown in Fig. 1(b). Like behavior has been found in numerical studies of QAOA, where the inclusion of edge weights in MaxCut leads to the appearance of additional poor-quality local minima in the optimization landscape [71]. Meanwhile, in Fig. 1(c), we plot the associated values of β up to layer $k = 50$, where the solid and dotted curves correspond to the results for unweighted and weighted 3-regular graphs, respectively. We find that there is strong agreement between the β curves for unweighted and weighted graphs at each problem size n . Furthermore, all β curves exhibit a very consistent shape.

In these illustrations, the only free parameter is the time step Δt . This is tuned to be as large as possible for each value of n , a value we call the critical time step and denote by Δt_c , as long as the condition in Eq. (5) is met for all (unweighted) problem instances considered up to 1,000 layers. We then utilize the same value of Δt for the weighted graphs at each problem size, noting that in some weighted instances, this leads to a violation of Eq. (5), and subsequent non-monotonic behavior of r_A and ϕ . For a closer look at how violations of Eq. (5) can manifest in individual problem instances, we refer to Appendix B).

Noting that convergence can be more challenging for weighted instances of MaxCut, we next explore how the performance of FALQON can be improved when it is modified using the iterative QLC heuristic introduced in Sec. III B, with results presented in Fig. 2. We apply this iterative QLC heuristic to a weighted instance of a 3-regular graph with $n = 8$ vertices, where the base FALQON algorithm displays good convergence with respect to r_A , but where ϕ fails to reach high values, and asymptotes to only $\phi \approx 0.57$. In Fig. 2(c), we show the β

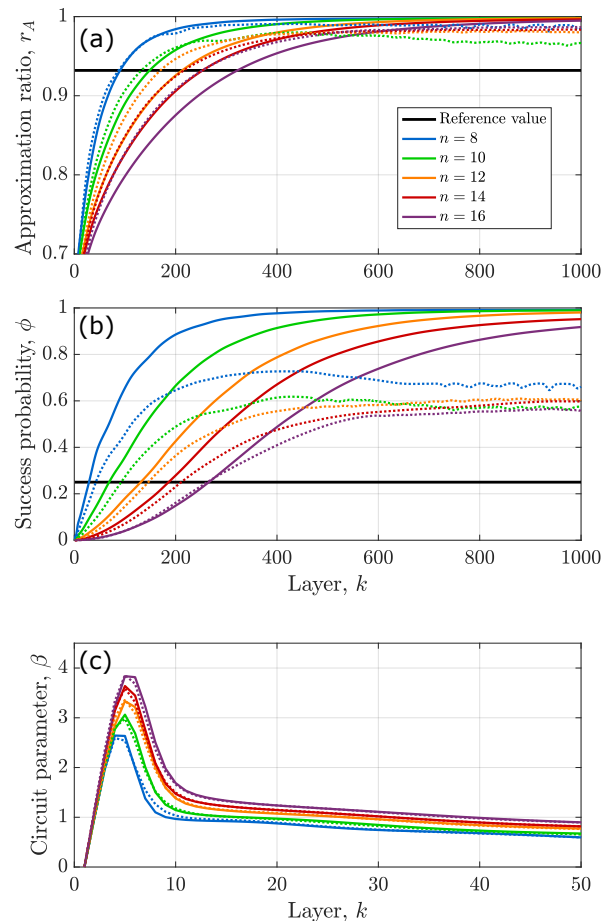


Figure 1. The performance of FALQON, as quantified by the approximation ratio r_A in (a) and the success probability ϕ in (b) is shown for different values of n as a function of layer k , where each layer is formed by an application of the problem Hamiltonian, H_p , and a driver Hamiltonian, H_d , defined according to Eqs. (45) and (46), respectively. This leads to full circuits with an alternating structure, as per Eq. (30). The solid curves of different colors show the mean results for MaxCut on unweighted 3-regular graphs over a set of different graphs at each problem size n , where the latter denotes the number of vertices (equivalently, the number of qubits) of the graph. For $n = 8, 10$, all possible graphs are considered; for $n = 12, 14, 16$, 50 graphs are considered at each problem size. The dotted curves show mean results for weighted 3-regular graphs, where edge weights are assigned to each graph, selected from a uniform distribution between 0 and 2. The solid black lines show the reference values $r_A = 0.932$ and $\phi = 0.25$. In (c), the mean values of β are plotted as a function of layer.

curves that result from three iterations of the procedure, while panel (b) shows how these iterations serve to improve the convergence of ϕ . We refer to Ref. [57] for an illustration of the improvement provided by the reference perturbation heuristic and also another random perturbation heuristic motivated by simulated annealing.

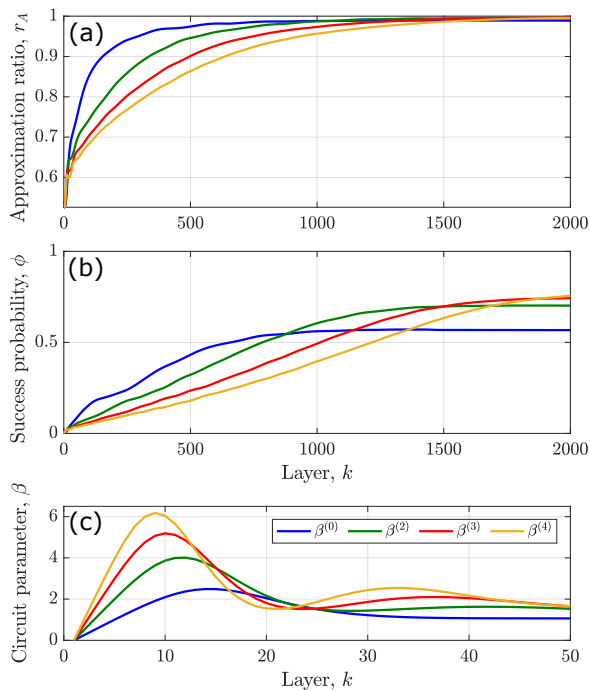


Figure 2. Performance of FALQON using the iterative QLC approach, applied to a weighted 3-regular graph with $n = 8$ vertices. Three iterations are shown, using a time step of $\Delta t = 0.012$. In (a) and (b), the approximation ratio, r_A and the success probability, ϕ are plotted as a function of layer, respectively. Panel (c) shows how the associated β curves are refined at each iteration of the procedure.

B. Behavior under measurement noise

Here, we analyze the performance of FALQON under measurement noise, which affects each A_k value, and consequently, each β_k value as well. This type of noise enters due to the fact that in practice, a finite number of measurement samples m are used to estimate each A_k . We simulate this effect by sampling measurement outcomes from a multinomial distribution, defined at layer k as the probability distribution over the set of eigenvalues of $i[H_d, H_p]$ when in the state $|\psi_k\rangle$ [72]. The results are collected in Fig. 3, which shows the performance of FALQON when $m = 2, 5, 20$ and 50 samples are used to estimate A_k , for an instance of MaxCut on a 3-regular graph with 8 vertices. The results shown are representative of the behavior across other instances we studied. Our findings suggest that FALQON is robust to the effects of sampling noise, and can be effective even in the presence of significant measurement noise. We also find that as the number of samples m that are used decreases, performance improves if Δt is selected to decrease as well.

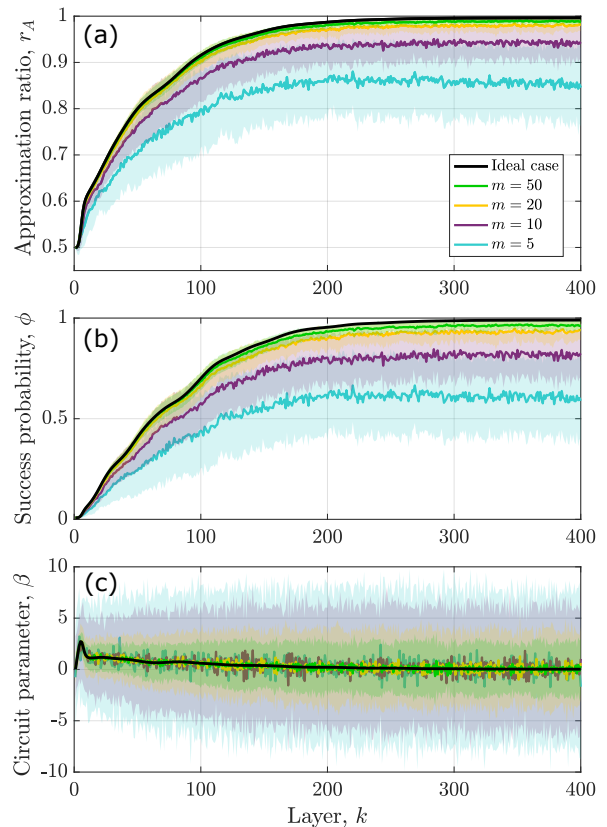


Figure 3. Performance of FALQON on an instance of Max-Cut on an unweighted, 3-regular graph with $n = 8$ vertices in the presence of sampling noise. The time step is set to $\Delta t = 0.034$. In panel (a), the approximation ratio r_A is plotted as a function of layer when $m = 2, 5, 20$, and 50 measurement samples are used to evaluate each of the expectation values A_k , $k = 1, \dots, 400$. The solid black curves show the ideal, noiseless reference case. The remaining solid curves represent the mean taken over 100 realizations of this sampling process using different numbers of measurement samples, m ; the shading shows the associated standard deviation over these realizations. Panels (b) and (c) show analogous results for the success probability, ϕ , of measuring the 2-degenerate ground state, and the circuit parameter, β , respectively.

C. Comparison with QAOA

In this section, we consider how the performance of FALQON can be expected to compare with that of QAOA. We recall that a key feature of FALQON is that it does not require any classical optimization, and as a result, the resources required for FALQON are substantially different, compared with the resources required for QAOA. In particular, in Fig. 1 we found that when FALQON is applied to MaxCut on regular graphs, it is able to achieve high values of the approximation ratio, r_A , and relatively high values of the success probability, ϕ , with no classical optimization. Furthermore, it can also achieve high values of r_A and ϕ in the presence of measurement noise, i.e., when only a small number of




Figure 4. Numerical analysis of the performance of QAOA for an instance of MaxCut on an unweighted, 3-regular graph with $n = 8$ vertices; the instance of MaxCut considered here is the same that is considered in the analysis presented in Fig. 3. For this analysis, we simulate implementations of QAOA that scan over $k = 1, 2, \dots, 25$ layers. A series of 100 realizations are performed at each value of k , in which the initial circuit parameter values are selected at random, and $q = 1000$ iterations of SPSA are performed in order to optimize the circuit parameters. SPSA makes two calls to the objective function per iteration, and for each of these calls, we consider estimating the value of $\langle H_p \rangle$ using $m = 10$, $m = 1000$, and $m = \infty$ samples, where the latter corresponds to the case of ideal measurements and perfect resolution of the expectation value $\langle H_p \rangle$. The performance of QAOA for each pair of (m, k) values is then quantified by the approximation ratio, r_A , and the success probability, ϕ , in panels (a) and (b), respectively. In each panel, the dotted curves and shaded regions show the mean and standard deviations computed over the realizations, respectively, as a function of k for each value of m . Meanwhile, the gray horizontal lines denote the reference values of $r_A = 0.932$ and $\phi = 0.25$.

measurement samples, m , are used to estimate the expectation values A_k at each layer, as shown in Fig. 3. However, it is evident from these figures that FALQON does require relatively deep circuits in order to achieve this good performance.

We now turn to QAOA. Applications of QAOA as a hybrid quantum-classical algorithm have mostly focused on shallow circuits. In this regime, QAOA can be expected to find better solutions than FALQON, through the aid of classical optimization. The premise is that the classical optimization will allow for extracting the best attainable solution from the quantum processor within a limited circuit depth. Beyond shallow circuits, in principle QAOA is capable of achieving solutions that improve

monotonically with respect to the depth of the circuit. However in practice, seeing a monotonic improvement in solution quality as the number of QAOA layers $k \rightarrow \infty$ would require resources that are not practically feasible (i.e., the ability to identify globally optimal solutions at each value of k through classical optimization). In practice, scaling up QAOA to larger problem sizes and deeper circuits will cause the classical optimization cost to rise, due to the fact that optimization is harder in higher dimensions, and in fact, formally scales exponentially with increasing number of layers [73].

Because QAOA and FALQON require different resources it is difficult to compare them directly, especially with one figure of merit. Instead, in the following we analyze the performance of each separately on the same problem instance. Fig. 4 presents the performance of QAOA on the same instance of MaxCut as is considered in the FALQON analysis presented in Fig. 3. In particular, Fig. 4 presents the results of QAOA simulations for circuits with $k = 1, 2, \dots, 25$ layers. A series of 100 QAOA realizations are performed at each value of k . For each realization, the initial parameter values are chosen uniformly randomly from $\beta_k \in [0, \pi)$, $\gamma_k \in [0, 2\pi)$. Then, the parameter optimization is performed using 1,000 iterations of the Simultaneous Perturbation Stochastic Approximation (SPSA) algorithm [74, 75], which involves perturbing $\langle H_p \rangle$ in order to estimate an approximate gradient at each optimization iteration, and utilizing this to perform gradient descent. The gradient is approximated at each iteration by evaluating the objective function $\langle H_p \rangle$ twice, regardless of how many optimization parameters are involved. For each of these evaluations, we consider estimating the value of $\langle H_p \rangle$ using $m = 10$, $m = 1000$, and $m = \infty$ samples, where the latter corresponds to the case of ideal measurements and perfect resolution of the expectation value $\langle H_p \rangle$. The performance of QAOA is then quantified by the approximation ratio, r_A , and the success probability, ϕ , with results plotted in Fig. 4(a) and (b), respectively.

We now contrast the results in Fig. 4 with the earlier FALQON results in Fig. 3. Recalling first the FALQON results, we find monotonic improvements in r_A and ϕ with respect to layer, k . We also find that increasing m leads to faster convergence. Turning to QAOA, we find that the QAOA results for different values of m are not significantly different, indicating that SPSA performs comparably in the presence of different levels of sampling noise for this problem instance. We also find that the behavior of r_A and ϕ with respect to k is non-monotonic. That is, both of these figures of merit first increase as a function of k . This is likely due to the increased expressivity of circuits as more layers are added, i.e., better solutions become reachable, and 1,000 SPSA iterations is sufficient for exploring the parameter space and identifying these solutions. As k increases further, the advantages of this increasing expressivity are subsequently counterbalanced by the fact that adding layers also adds more optimization variables, increasing the dimension of

the optimization space and the difficulty of the optimization problem. The behavior of r_A and ϕ then rolls over and begins to deteriorate as a consequence of the increasing difficulty of the optimization problem, once the limited number of optimization iterations allowed becomes insufficient for exploring the space and finding good solutions.

We emphasize that these findings are not specific to the MaxCut problem instance analyzed here, and hold generically. Also, the choice of optimization algorithm does not significantly affect the conclusions, the same overall behavior is seen with other optimization algorithms.

These findings support the notion that QAOA is likely favorable in settings where classical optimization resources are sufficiently available and quantum resources are limited to the regime of shallow circuits. On the other hand, FALQON demonstrates strong performance for deep circuits, and does not require classical optimization resources, meaning that it has the advantage of not incurring a rising classical cost as the circuit depth is scaled up. This suggests that in cases where it is feasible to implement deep circuits, FALQON could offer a considerable advantage.

We conclude this section with another comparison of the resources required by FALQON and QAOA, now in the context of their sampling complexity for MaxCut, as quantified by the total number of samples (i.e., circuit repetitions) that are required, denoted N_s . We denote by m 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 . We first consider QAOA: given that all terms in H_p commute, for $q(\ell)$ classical optimization iterations of QAOA, $N_s = mq(\ell)$, assuming one evaluation of $\langle H_p \rangle$ per optimization iteration. We note that for any reasonable convergence, the number of optimization iterations, $q(\ell)$ depends at least linearly on ℓ . However, if a gradient algorithm is used for QAOA, additional samples will be needed. We assume that for ℓ layers of QAOA, 2ℓ gradient elements are required, for each of the 2ℓ circuit parameters. Assuming that at least m samples are needed to estimate each gradient element (i.e., to evaluate $\langle H_p \rangle$ for at least one perturbation of each circuit parameter), then for q iterations,

$$N_s^{QAOA} \geq mq(\ell)(1 + 2\ell) = \mathcal{O}(mq(\ell)\ell). \quad (51)$$

In FALQON, additional measurements are needed to evaluate A_1, \dots, A_ℓ . This involves measuring each of the terms in $i[H_d, H_p]$, which contains a set of YZ terms and a set of ZY terms. In principle, these two sets can be combined to form a single set, given that each $Y_j Z_k$ term commutes with each $Z_j Y_k$ term and can thus be measured together [76–78]; however, we consider the scenario that commuting $Y_i Z_j$ and $Z_i Y_j$ terms are measured separately, as per current convention, although terms such as $Y_i Z_j$ and $Z_i Y_m$, which act nontrivially on disjoint pairs of qubits, may be measured simultaneously. Then, for a graph \mathcal{G} with maximum degree d , the expectation value

of $i[H_d, H_p]$ can be estimated in maximally $2m(d + 1)$ repetitions [79]. For ℓ layers, this yields

$$N_s^{FALQON} \leq 2m\ell(d + 1) = \mathcal{O}(md\ell). \quad (52)$$

This comparison 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.

VI. COMBINING FALQON AND QAOA

As we described previously, FALQON has flexibility in the choice of a control law, e.g., Eq. (27) and the value of β_1 , the introduction of a reference perturbation (Sec. III A), and the choice of driver Hamiltonian. Once these features are selected, FALQON is a deterministic, constructive procedure, i.e., in the limit of perfect measurements, the resulting set of parameters $\{\beta_k\}$ is uniquely specified for a problem Hamiltonian H_p . On the other hand, QAOA has flexibility in the choice of the driver Hamiltonian(s), as well as the classical optimization method and *all* initial values of the parameter set elements β_k and γ_k .

The numerical results presented in Sec. V suggest that solving MaxCut using FALQON alone can require many layers and therefore may not be suitable for NISQ devices with limited circuit depths. In this section, we explore how FALQON results from a smaller number of layers can be used as a seed to initialize a QAOA circuit, thereby aiding in the subsequent search for optimal circuit parameters. Related work from Egger et al. proposes a somewhat similar idea for “warm-starting” low-depth QAOA using the solution from a relaxation of the original combinatorial optimization problem [80]. While Sack and Serbyn introduce a “Trotterized quantum annealing protocol” to initialize QAOA [81], parametrized by the time step Δt . In our work, we consider MaxCut on ensembles of unweighted 3-regular graphs with $n \in \{8, 10, 12, 14\}$ vertices and 10-layer circuits implementing FALQON and QAOA. All of the simulations in this section are performed using pyQAOA, a Python-based simulator of QAOA circuits [82]. As before, for graphs with $n = 8, 10$ vertices, we consider all nonisomorphic, connected 3-regular graphs. For graphs with $n = 12$ or $n = 14$ vertices, we consider a set of 50 randomly-generated nonisomorphic graphs for each value of n .

For each graph, the set of parameters $\{\beta_k\}$ is generated for a 10-layer circuit using FALQON as described in Algorithm 1 with H_d as specified in Eq. (46). To use the results of FALQON to initialize QAOA, the product $\beta_k \Delta t$ obtained from FALQON becomes the initial value for β_k in QAOA; analogously, Δt from FALQON becomes the initial value of *all* γ_k in QAOA, i.e.,

$$\beta_k \Delta t \rightarrow \beta_k \quad \text{and} \quad \Delta t \rightarrow \gamma_k \quad \text{for all } k. \quad (53)$$

This follows from the relationship between FALQON unitary operations and QAOA unitary operations, i.e., compare U_p and $U_d(\beta_k)$ in Eqs. (24)–(25) to the corresponding QAOA operations.

Within the context of a 10-layer QAOA circuit, the two sets of parameters $\{\beta_k\}$ and $\{\gamma_k\}$ are then optimized using a quasi-Newton optimization method (BFGS). For simplicity, we refer to this sequential FALQON + QAOA procedure as “FALQON+”. In addition, we compare the performance of FALQON+ to QAOA with multiple random initializations, i.e., we perform QAOA using multistart BFGS with 20 randomly-selected initial values for $\{\beta_k\}$ and $\{\gamma_k\}$.

Approximation ratios r_A and success probabilities ϕ for FALQON, corresponding FALQON+, and multistart QAOA are reported in Fig. 5. In both figures, the color-shaded boxes and encompassed horizontal line represent the interquartile range and the median value of the data, respectively, while the whiskers extend out to 1.5 of the interquartile range; points beyond this range are identified as “outliers” (represented as grey diamond symbols). Comparing the results of FALQON and FALQON+, note that in addition to improved approximation ratios, FALQON+ also substantially improves the success probabilities, at the cost of only one application of QAOA with BFGS. Since the final state measured at the end of the circuit corresponds to an actual (approximate) solution of the MaxCut problem, increasing success probabilities is more important than increasing approximation ratios. For our multistart QAOA simulations, we present distributions of maximum, median, and minimum (with respect to the randomly-selected then optimized sets of initial parameters $\{\beta_k\}$ and $\{\gamma_k\}$) approximation ratios and corresponding success probabilities for each ensemble of graphs. In the legend of Fig. 5, these distributions are denoted as QAOA_{\max} , QAOA_{med} , and QAOA_{\min} , respectively. In principle, QAOA can perform better than FALQON overall since QAOA parameters can be optimized globally *and* these parameters include the additional set $\{\gamma_k\}$. However, in practice, achieving this improved performance may require multiple optimizations. For our simulations, FALQON+ performs comparably to the maximum and median cases of multistart QAOA.

In Fig. 6, we present example instances of FALQON ($\beta_k \Delta t$ and Δt) and corresponding FALQON+ (β_k and γ_k) parameters for $n \in \{8, 10, 12, 14\}$ vertices [83]. Combined with results presented in Fig. 5, these examples illustrate that substantial changes and improvements can occur between FALQON initialization and subsequent FALQON+ convergence, indicating that the parameters generated by FALQON, i.e., $\beta_k \Delta t$ and Δt , do not correspond to local optima for the QAOA landscape. In this sense, FALQON can prepare parameters for a successful application of QAOA, thereby reducing the expense of the optimization effort for QAOA. Although not presented here, the FALQON-QAOA parameter differences presented in Fig. 6 are typical of all of our simulation results.

Based on results and analysis presented here, FALQON+ may provide a tractable solution to the challenge of identifying optimal parameters for QAOA. Overall, our results demonstrate that FALQON can be used to enhance the performance of depth-limited QAOA, with minimal additional cost. See ref. [57] for estimates of FALQON and QAOA sampling complexity.

VII. QUANTUM ANNEALING APPLICATIONS

Quantum annealing [84] is an approach for preparing the ground state of a problem Hamiltonian H_p that proceeds by initializing a quantum system in the ground state of another Hamiltonian H_d , and then evolving the system via the time-dependent Hamiltonian

$$H(t) = u(t)H_d + (1 - u(t))H_p \quad (54)$$

for $t \in [0, T]$, where $u(t)$ is the quantum annealing schedule, with $u(0) = 1$ and $u(T) = 0$. Without known structure in H_p to exploit, often the simplest annealing schedule is linear, where $u(t) = 1 - t/T$, and we consider this in the following. Then, the aim is to choose T to be large enough such that the system remains in the instantaneous ground state of $H(t)$ at all times, so that as H_p is slowly turned on, this will evolve the system into the ground state of H_p at time T .

In this section, we compare FALQON to linear quantum annealing because of numerical evidence suggesting that FALQON may proceed via a similar adiabatic mechanism, i.e., by slowly switching on the problem Hamiltonian H_p , such that the system remains in the instantaneous ground state. Evidence of this potential adiabatic behavior is shown in Fig. 7, for a representative instance of unweighted 3-regular MaxCut on $n = 8$ vertices with $\Delta t = \Delta t_c$. The population in the instantaneous ground state, ϕ_{inst} , is computed as $\phi_{\text{inst}} = \sum_j |\langle \psi | \tilde{q}_{0,j} \rangle|^2$, where the sum is taken over j degenerate instantaneous ground states (found numerically, as the eigenstates whose associated eigenvalues are within 0.01 of the lowest eigenvalue); the set of instantaneous eigenstates $\{\tilde{q}\}$ is computed by numerically diagonalizing $H_p + \beta H_d$ at each layer. The consistently high values of ϕ_{inst} in Fig. 7, which is representative of the behavior seen across other MaxCut instances, suggest that FALQON may give rise to a Trotterized version of an adiabatic process, i.e., in which strong rotations are applied initially to transfer $|\psi\rangle$ into the instantaneous ground state, and then, the system remains primarily in the instantaneous ground state for the remaining evolution. In order to achieve this behavior, we see in Fig. 1(c) that β initially has large values, then decreases monotonically as a function of layer, similar to the behavior of an annealing schedule $u(t)$. Particularly notable from Fig. 1(c) is that the β curves appear to concentrate around a single average curve for each value of n , indicating that there may be a universal FALQON solution for this class of problems. As such, we

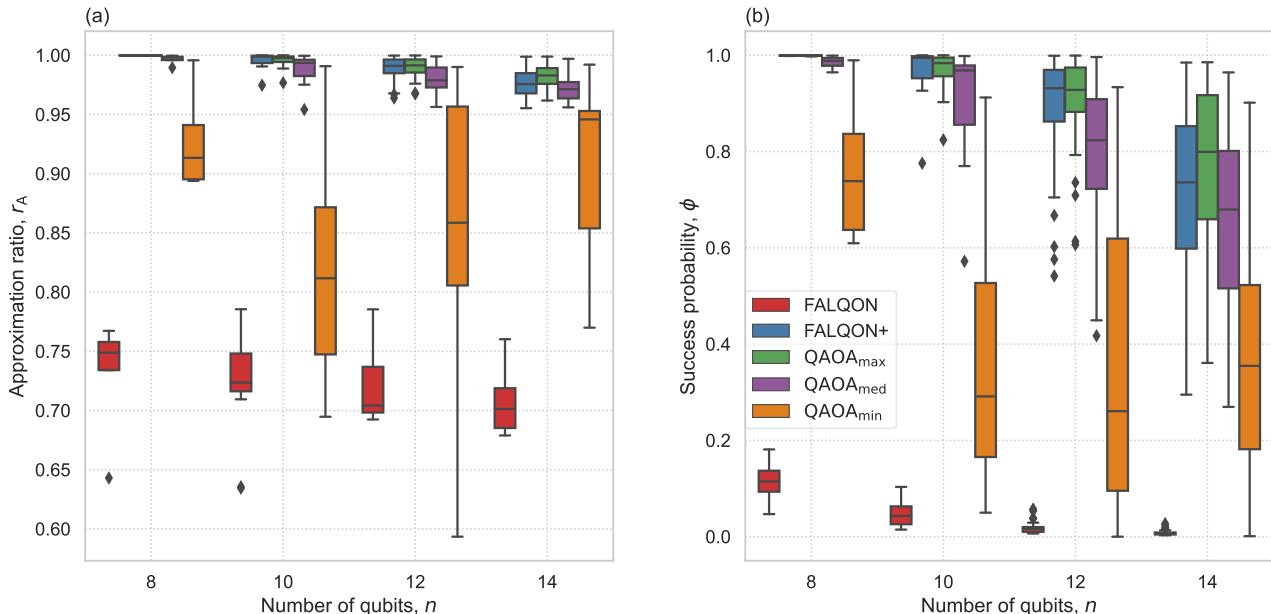


Figure 5. Performance of MaxCut on ensembles of 3-regular graphs with n vertices for $n \in \{8, 10, 12, 14\}$ for FALQON, FALQON+, and multistart QAOA for 10-layer circuits, quantified by (a) the approximation ratio r_A and (b) the success probability ϕ , where QAOA_{max}, QAOA_{med}, and QAOA_{min} denote distributions of maximum, median, and minimum QAOA results, respectively. In both figures, the color-shaded boxes and encompassed horizontal line represent the interquartile range and the median value of the data, respectively, while the whiskers extend out to 1.5 of the interquartile range; points beyond this range are identified as “outliers” (represented as grey diamond symbols).

relate these curves to digitized quantum annealing schedules, and consider the digitized time $T = 2k\Delta t$ needed to achieve $r_A = 0.932$ or $\phi = 0.25$ using FALQON. The results are shown in Fig. 8, which shows that T scales favorably with respect to n , with an a linear scaling at the problem sizes evaluated.

We then compare the performance of FALQON against that of a digitized linear quantum annealing schedule. We present the results of our numerical comparison in Fig. 9 for the same MaxCut problem instances considered in Fig. 1. Our findings indicate that for the same value of T , FALQON consistently shows stronger performance, as quantified by both r_A and ϕ .

It is of course important to note that this comparison is limited insofar as we compare only to a linear annealing schedule. This restriction was chosen for simplicity; it remains to be seen how FALQON compares relative to quantum annealing with various optimized schedules [85–87]. Nevertheless, these results suggest that feedback-based protocols could be useful for improving performance of analog annealing devices as well. For example, the control schedule determined by FALQON, $\beta(t)$, could be used as the basis for an adiabatic annealing schedule. Alternatively, an annealing schedule could be derived through execution of an analogous Lyapunov-control inspired feedback strategy on an analog annealer, assuming the required measurements for determining $A(t)$ could be performed.

VIII. OUTLOOK

We have introduced FALQON as a constructive, feedback-based algorithm for solving combinatorial optimization problems using quantum computers, and explored its utility towards the MaxCut problem via a series of numerical experiments. Crucially, FALQON does not require classical optimization, unlike other quantum optimization frameworks such as QAOA. However, this advantage comes at a cost. As we found in our numerical illustrations, the quantum circuits needed tend to be much deeper than those conventionally considered in QAOA, suggesting that there is a tradeoff between the classical and quantum costs.

Our numerical demonstrations utilized the feedback law given in Eq. (27), although a much broader class of functions could be considered, as per Eq. (8), and the performance for different choices of w and f should be explored. Furthermore, the use of bang-bang control laws, e.g., where $\beta \in \{\pm\beta_{\max}\}$ switches between $\pm\beta_{\max}$, for a value of β_{\max} chosen according to the sign of A , could also be considered in the future. Furthermore, we remark that the performance of FALQON depends on the choice of Δt , suggesting that it may be possible to design methods to optimally or adaptively choose Δt , e.g., for a given problem, or in a layer-by-layer manner based on measurement data, perhaps informed by Eq. (43), in order to enhance the algorithm performance.

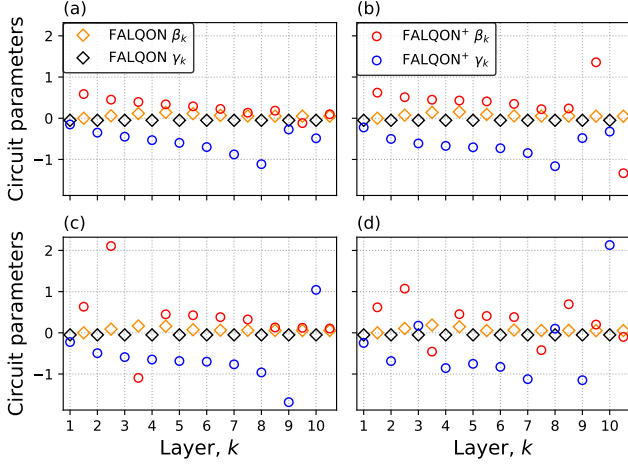


Figure 6. Comparison of FALQON and corresponding FALQON+ parameters for example instances of 3-regular graphs with n vertices for $n = 8$ (a), $n = 10$ (b), $n = 12$ (c), and $n = 14$ (d). Orange and black diamonds denote FALQON $\beta_k \Delta t$ and $\gamma_k = \Delta t$ parameters; red and blue circles denote FALQON+ β_k and γ_k parameters, as described in Eq. (53). Because each FALQON and QAOA circuit layer contains a sequence of unitary operations, i.e., $U_d(\beta_k)U_p(\gamma_k)$ in Eqs. (2) and (23), we plot γ_k at k and β_k at $k + 1/2$ for clarity.




Figure 7. The population in the ground state, ϕ , and the population in the instantaneous ground state, ϕ_{inst} , are plotted as a function of layer, for an application of FALQON to MaxCut on an unweighted 3-regular graph with $n = 8$ vertices.

In future implementations, FALQON could be used alone as a substitute for conventional QAOA, or it could be used in combination with QAOA (e.g., by taking $|\psi_0\rangle$ to be the terminal state from an already-optimized QAOA circuit). Similarly, in this work we have explored how FALQON could be used to seed QAOA, by identifying a set of initial QAOA parameters that can serve as the starting point for subsequent classical optimization. We expect that this seeding procedure may have particular benefit 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 classical optimization procedure.




Figure 8. The mean digitized time $T = 2k\Delta t$ needed to achieve the reference values of $r_A = 0.932$ (brown) and $\phi = 0.25$ (black) using a FALQON-inspired annealing schedule is shown for the unweighted MaxCut problems considered in Fig. 1. The mean is taken over the set of unweighted, 3-regular graphs that are considered at each problem size n , and the error bars show the associated standard deviations.




Figure 9. (a) The mean approximation ratios, r_A , obtained by FALQON (solid curves) and a linear quantum annealing schedule (dotted curves) for MaxCut on unweighted, 3-regular graphs with $n = 8, 10, 12, 14$ vertices. The mean is taken over the set of different graphs that are considered at each problem size n . For the linear quantum annealing schedule, T is chosen to be the time when FALQON reaches the reference value $r_A = 0.932$ for each value of n . This reference value is plotted as a horizontal black line. Panel (b) shows the corresponding results for the success probabilities, ϕ .

We further remark that in situations where circuit depth is limited, FALQON can be extended to incorporate additional driver Hamiltonians, drawing on the framework outlined in Sec. III C, and it could also be modified to use a hardware-inspired ansatz, where the circuit is formed by alternating rounds of an Ising Hamiltonian defined by the hardware connectivity, denoted as

H_h , and a driver Hamiltonian H_d , while the objective remains determined by the Ising problem Hamiltonian, denoted as H_p . In this scenario, despite changes in the structure of the quantum circuits, the measurements of $i[H_d, H_p]$ needed to assign values to the β parameters would remain unchanged.

ACKNOWLEDGMENTS

We acknowledge discussions with C. Arenz, L. Brady, L. Cincio, T.S. Ho, L. Kocia, O. Parekh, H. Rabitz, and K. Young. MDG gratefully acknowledges helpful discussions about pyQAOA usage with G. von Winckel [82]. 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.

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all

right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan <https://www.energy.gov/downloads/doe-public-access-plan>. 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.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

-
- [1] B. L. Golden, S. Raghavan, E. A. Wasil, *et al.*, *The vehicle routing problem: latest advances and new challenges*, Vol. 43 (Springer, 2008).
- [2] J. Blażewicz, W. Domschke, and E. Pesch, The job shop scheduling problem: Conventional and new solution techniques, *Eur. J. Oper. Res.* **93**, 1 (1996).
- [3] A. B. Finnila, M. Gomez, C. Sebenik, C. Stenson, and J. D. Doll, Quantum annealing: A new method for minimizing multidimensional functions, *Chem. Phys. Lett.* **219**, 343 (1994).
- [4] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [5] J. Brooke, D. Bitko, G. Aeppli, *et al.*, Quantum annealing of a disordered magnet, *Science* **284**, 779 (1999).
- [6] C. Durr and P. Hoyer, A quantum algorithm for finding the minimum (1996), [arXiv:9607014 \[quant-ph\]](https://arxiv.org/abs/9607014).
- [7] C. Dürr, M. Heiligman, P. Hoyer, and M. Mhalla, Quantum query complexity of some graph problems, *SIAM J. Comput.* **35**, 1310 (2006).
- [8] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm (2014), [arXiv:1411.4028 \[quant-ph\]](https://arxiv.org/abs/1411.4028).
- [9] J. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, *et al.*, Unsupervised machine learning on a hybrid quantum computer (2017), [arXiv:1712.05771 \[quant-ph\]](https://arxiv.org/abs/1712.05771).
- [10] M. Willsch, D. Willsch, F. Jin, H. De Raedt, and K. Michielsen, Benchmarking the quantum approximate optimization algorithm, *Quantum Inf. Process.* **19**, 197 (2020).
- [11] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, Implementation of the XY interaction family with calibration of a single pulse, *Nat. Electron.* **3**, 744 (2020).
- [12] A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Krizan, D. Shiri, I.-M. Svensson, G. Tancredi, G. Johansson, P. Delsing, G. Ferrini, and J. Bylander, Improved success probability with greater circuit depth for the quantum approximate optimization algorithm, *Phys. Rev. Applied* **14**, 034010 (2020).
- [13] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo, and *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [14] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De,

- P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, and et al., Quantum approximate optimization of the long-range ising model with a trapped-ion quantum simulator, *Proc. Natl. Acad. Sci. U.S.A* **117**, 25396 (2020).
- [15] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [16] 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).
- [17] C. Brif, R. Chakrabarti, and H. Rabitz, Control of quantum phenomena: past, present and future, *New J. Phys.* **12**, 075008 (2010).
- [18] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Training schrödinger's cat: quantum optimal control, *Eur. Phys. J. D* **69**, 279 (2015).
- [19] C. Altafini, Controllability of quantum mechanical systems by root space decomposition of $su(N)$, *J. Math. Phys.* **43**, 2051 (2002).
- [20] F. Albertini and D. D'Alessandro, The Lie algebra structure and controllability of spin systems, *Linear Algebra Appl.* **350**, 213 (2002).
- [21] S. G. Schirmer, H. Fu, and A. I. Solomon, Complete controllability of quantum systems, *Phys. Rev. A* **63**, 063410 (2001).
- [22] H. Fu, S. G. Schirmer, and A. I. Solomon, Complete controllability of finite-level quantum systems, *J. Phys. A* **34**, 1679 (2001).
- [23] G. Turinici and H. Rabitz, Wavefunction controllability for finite-dimensional bilinear quantum systems, *J. Phys. A* **36**, 2565 (2003).
- [24] D. Burgarth, D. D'Alessandro, L. Hogben, S. Severini, and M. Young, Zero forcing, linear and quantum controllability for systems evolving on networks, *IEEE Trans. Automat. Contr.* **58**, 2349 (2013).
- [25] C. Arenz, D. Burgarth, P. Facchi, V. Giovannetti, H. Nakazato, S. Pascazio, and K. Yuasa, Universal control induced by noise, *Phys. Rev. A* **93**, 062308 (2016).
- [26] C. Arenz and H. Rabitz, Controlling qubit networks in polynomial time, *Phys. Rev. Lett.* **120**, 220503 (2018).
- [27] S. Lloyd, Quantum approximate optimization is computationally universal (2018), [arXiv:1812.11075 \[quant-ph\]](https://arxiv.org/abs/1812.11075).
- [28] M. E. Morales, J. Biamonte, and Z. Zimborás, On the universality of the quantum approximate optimization algorithm, *Quantum Inf. Process.* **19**, 1 (2020).
- [29] G. Bigan Mbeng, R. Fazio, and G. Santoro, Quantum Annealing: a journey through Digitalization, Control, and hybrid Quantum Variational schemes (2019), [arXiv:1906.08948 \[quant-ph\]](https://arxiv.org/abs/1906.08948).
- [30] A. S. Berahas, R. H. Byrd, and J. Nocedal, Derivative-free optimization of noisy functions via quasi-newton methods, *SIAM J. Optim.* **29**, 965 (2019).
- [31] R. Chakrabarti and H. Rabitz, Quantum control landscapes, *Int. Rev. Phys. Chem.* **26**, 671 (2007).
- [32] B. Russell, H. Rabitz, and R.-B. Wu, Control landscapes are almost always trap free: A geometric assessment, *J. Phys. A* **50**, 205302 (2017).
- [33] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 1 (2018).
- [34] C. Arenz and H. Rabitz, Drawing together control landscape and tomography principles, *Phys. Rev. A* **102**, 042207 (2020).
- [35] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [36] L. Bittel and M. Kliesch, Training variational quantum algorithms is np-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [37] J. Lee, A. B. Magann, H. A. Rabitz, and C. Arenz, Progress toward favorable landscapes in quantum combinatorial optimization, *Phys. Rev. A* **104**, 032401 (2021).
- [38] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control (2021), [arXiv:2105.14377 \[quant-ph\]](https://arxiv.org/abs/2105.14377).
- [39] P. Gross, H. Singh, H. Rabitz, K. Mease, and G. M. Huang, Inverse quantum-mechanical control: A means for design and a test of intuition, *Phys. Rev. A* **47**, 4593 (1993).
- [40] Y. Chen, P. Gross, V. Ramakrishna, H. Rabitz, and K. Mease, Competitive tracking of molecular objectives described by quantum mechanics, *J. Chem. Phys.* **102** (1995).
- [41] A. G. Campos, D. I. Bondar, R. Cabrera, and H. A. Rabitz, How to make distinct dynamical systems appear spectrally identical, *Phys. Rev. Lett.* **118**, 083201 (2017).
- [42] G. McCaul, C. Orthodoxou, K. Jacobs, G. H. Booth, and D. I. Bondar, Controlling arbitrary observables in correlated many-body systems, *Phys. Rev. A* **101**, 053408 (2020).
- [43] G. McCaul, C. Orthodoxou, K. Jacobs, G. H. Booth, and D. I. Bondar, Driven imposters: Controlling expectations in many-body systems, *Phys. Rev. Lett.* **124**, 183201 (2020).
- [44] 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).
- [45] 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).
- [46] 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).
- [47] 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).
- [48] 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).
- [49] M. Sugawara, General formulation of locally designed coherent control theory for quantum system, *J. Chem. Phys.* **118**, 6784 (2003).
- [50] 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.
- [51] M. Mirrahimi, G. Turinici, and P. Rouchon, Reference

- trajectory tracking for locally designed coherent quantum controls, *J. Phys. Chem. A* **109** **11**, 2631 (2005).
- [52] 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.
- [53] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan, Quantum feedback control and classical control theory, *Phys. Rev. A* **62**, 012105 (2000).
- [54] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, 2009).
- [55] J. Combes, J. Kerckhoff, and M. Sarovar, The slh framework for modeling quantum input-output networks, *Adv. Phys.: X* **2**, 784 (2017).
- [56] J. Zhang, Y.-x. Liu, R.-B. Wu, K. Jacobs, and F. Nori, Quantum feedback: theory, experiments, and applications, *Phys. Rep.* **679**, 1 (2017).
- [57] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Feedback-based quantum optimization (2021), [arXiv:2103.08619 \[quant-ph\]](https://arxiv.org/abs/2103.08619).
- [58] A. Lucas, Ising formulations of many np problems, *Front. Phys.*, **5** (2014).
- [59] A. Isidori, M. Thoma, E. D. Sontag, B. W. Dickinson, A. Fettweis, J. L. Massey, and J. W. Modestino, *Nonlinear Control Systems*, 3rd ed. (Springer-Verlag, Berlin, Heidelberg, 1995).
- [60] S. Cong and F. Meng, A survey of quantum lyapunov control methods, *Sci. World J.* **2013** (2013).
- [61] K. Beauchard, J. M. Coron, M. Mirrahimi, and P. Rouchon, Implicit lyapunov control of finite dimensional schrödinger equations, *Syst. Control. Lett.* **56**, 388 (2007).
- [62] 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).
- [63] For applications of FALQON to MaxCut on regular graphs, as studied in Sec. V, the results of this procedure for measuring the bit string $z_1 z_2 \cdots z_n$ will be concentrated around the mean when shallow circuits are used. The proof for concentration of $\langle H_p \rangle$ for fixed ℓ follows directly from the proof in Sec. III of [8]. Given that H_p and $Z_1 Z_2 \cdots Z_n$ are both diagonal in the measured, computational basis, it follows that concentration holds for the latter as well.
- [64] J. Carolan, M. Mohseni, J. P. Olson, M. Prabhu, C. Chen, D. Bunandar, M. Y. Niu, N. C. Harris, F. N. Wong, M. Hochberg, *et al.*, Variational quantum unsampling on a quantum photonic processor, *Nat. Phys.* **16**, 322 (2020).
- [65] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, *Quantum Mach. Intell.* **3**, 1 (2021).
- [66] E. Campos, A. Nasrallah, and J. Biamonte, Abrupt transitions in variational quantum circuit training, *Phys. Rev. A* **103**, 032607 (2021).
- [67] E. Campos, D. Rabinovich, V. Akshay, and J. Biamonte, Training saturation in layerwise quantum approximate optimization, *Phys. Rev. A* **104**, L030401 (2021).
- [68] G. Verdon, J. M. Arrazola, K. Brádler, and N. Killoran, A quantum approximate optimization algorithm for continuous problems (2019), [arXiv:1902.00409 \[quant-ph\]](https://arxiv.org/abs/1902.00409).
- [69] L. Zhu, H. L. Tang, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer (2020), [arXiv:2005.10258 \[quant-ph\]](https://arxiv.org/abs/2005.10258).
- [70] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [71] 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 \[quant-ph\]](https://arxiv.org/abs/2201.11785).
- [72] We remark that in practice, it may not be possible to sample from the set of eigenvalues of $i[H_d, H_p]$ by measuring the observable $i[H_d, H_p]$ directly. Instead, the observable $i[H_d, H_p]$ may be expanded as a weighted sum of terms as per Eqs. (29) and (47). Then, the terms can then be grouped into sets of observables that can be measured jointly.
- [73] L. Bittel and M. Kliesch, Training variational quantum algorithms is np-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [74] J. C. Spall, An overview of the simultaneous perturbation method for efficient optimization, Johns Hopkins apl technical digest **19**, 482 (1998).
- [75] J. C. Spall, Implementation of the simultaneous perturbation algorithm for stochastic optimization, *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817 (1998).
- [76] O. Crawford, B. v. Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [77] T.-C. Yen, V. Verteletskyi, and A. F. Izmaylov, Measuring all compatible operators in one series of single-qubit measurements using unitary transformations, *J. Chem. Theory Comput.* **16**, 2400 (2020).
- [78] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).
- [79] We can see this by mapping the problem of estimating the number of repetitions to the problem of edge coloring, which asks for an assignment of colors to the edges of a graph such that no edges sharing a common vertex have the same color. The number of colors corresponds to the number of repetitions, as the same rule applies (i.e., that no non-commuting 2-qubit terms sharing a common qubit can be measured simultaneously). From Vizing’s theorem, at most $d + 1$ colors are needed for edge coloring a (simple) graph whose maximum degree is d [88]. Although the task of determining the optimal edge coloring is in general NP-hard, classical, polynomial-time algorithms exist for assigning at most $d + 1$ colors to graphs with maximum degree d [89].
- [80] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [81] S. H. Sack and M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *Quantum* **5**, 491 (2021).
- [82] G. von Winckel, pyQAOA: Simulation of Quantum Approximate Optimization Algorithms in Python, <https://github.com/gregvw/pyQAOA> (2021), accessed: 2021-07-27.
- [83] In our version of pyQAOA, unitary operations are implemented as the conjugate transpose of Eqs. (24)–(25), i.e., $-i \rightarrow +i$ in the exponent. To account for this difference,

we plot $-\beta_k$ and $-\gamma_k$ rather than β_k and γ_k in Fig. 6, hence the negative values for the FALQON circuit parameters γ_k , corresponding to $-\Delta t = -0.025$.

- [84] A. Das and B. K. Chakrabarti, Colloquium: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* **80**, 1061 (2008).
- [85] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Phys. Rev. A* **65**, 042308 (2002).
- [86] C. Brif, M. D. Grace, M. Sarovar, and K. C. Young, Exploring adiabatic quantum trajectories via optimal control, *New J. Phys.* **16**, 065013 (2014).
- [87] L. Zeng, J. Zhang, and M. Sarovar, Schedule path optimization for adiabatic quantum computing and optimization, *J. Phys. A* **49**, 1 (2016).
- [88] V. G. Vizing, On an estimate of the chromatic class of a p-graph, *Discret Analiz* **3**, 25 (1964).
- [89] J. Misra and D. Gries, A constructive proof of Vizing's theorem, *Inf. Process. Lett* **41** (1992).
- [90] J. P. La Salle, *The stability of dynamical systems* (SIAM, 1976).

Appendix A: Convergence of QLC

Within the QLC framework outlined in Sec. III, it has been shown that asymptotic convergence to the ground state of H_p can be guaranteed when the following sufficient criteria are met [50, 60–62]:

1. H_p has no degenerate eigenvalues, i.e., $q_i \neq q_j$ for $i \neq j$ where q_i and q_j are the i -th and j -th eigenvalues of H_p
2. H_p has no degenerate eigenvalue gaps, i.e., $\omega_{ji} \neq \omega_{lk}$ for $(i, j) \neq (k, l)$, where $\omega_{ji} = q_j - q_i$ is the gap between the i -th and j -th eigenvalues of H_p
3. $\langle q_j | H_d | q_i \rangle \neq 0$ for all $i \neq j$
4. $E_p(|q_0\rangle) < E_p(|\psi(t=0)\rangle) < E_p(|q_1\rangle)$

In particular, if criteria 1-3 are met, then the LaSalle invariance principle [90] can be used to show that any initial state $|\psi(t=0)\rangle$ will converge asymptotically to the largest invariant set, i.e., the largest set of states where $\frac{d}{dt}E_p = 0$. When E_p is chosen per Eq. (4), it can be shown that the largest invariant set is the set of eigenstates of H_p . Within this set, the eigenstate $|q_0\rangle$ with the smallest eigenvalue is the minimum, the eigenstate with the largest eigenvalue is the maximum, and all other eigenstates with intermediate eigenvalues are saddle points. In order to ensure convergence to the desired critical point $|q_0\rangle$, criterion (4) stipulates that the value of $E_p(|\psi(0)\rangle)$ at time $t=0$ must be strictly lower than $E_p(|q_1\rangle) = q_1$, such that the only critical point inside $E_p(|\psi(t)\rangle) \leq E_p(|\psi(0)\rangle)$ is the desired target, $|q_0\rangle$. Thus, the satisfaction of criteria (1)-(4) is sufficient to ensure that the system state will converge asymptotically to the desired target $|q_0\rangle$ [50, 60–62].

Appendix B: Selecting $\Delta t > \Delta t_c$

Fig. 10 illustrates the effects of selecting a time step that is too large, i.e., $\Delta t > \Delta t_c$, for an instance of unweighted, 3-regular MaxCut on $n=8$ vertices, with $\Delta t = 0.065$. The behavior in Fig. 10 is representative of the behavior seen across other instances when $\Delta t > \Delta t_c$. In general, there is a balance between selecting a large Δt for improving convergence *and* satisfying $\Delta t \leq \Delta t_c$ for ensuring monotonic improvement in E_p .

Appendix C: Log-log plot of MaxCut results

In Fig. 11, we plot the results presented in the main text in Fig. 1 using a log-log scale.

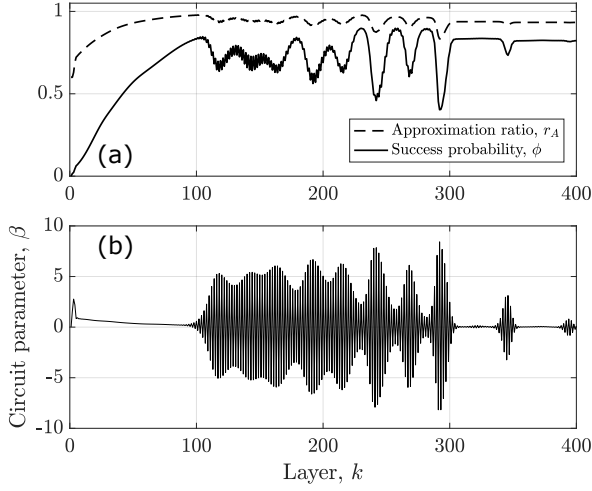


Figure 10. Typical behavior when the time step Δt is chosen to be too large, leading to a violation of the QLC criterion that $\langle H_D \rangle$ decreases monotonically with respect to layer, k . In (a), the behavior of the approximation ratio, r_A , and the success probability, ϕ , are shown, with the violation in monotonicity occurring around 100 layers. In (b), the associated behavior of the circuit parameter β is plotted, indicating that the violation of the QLC criterion corresponds to the creation of rapid oscillations in β .




Figure 11. The performance of FALQON, as quantified by the approximation ratio r_A in (a) and the success probability ϕ in (b), is shown as a function of the layer k for different qubit counts n on a log-log scale. The same results are plotted on a linear scale in Fig. 1. The solid curves show the average results for unweighted 3-regular graphs, and the dotted curves show average results for weighted 3-regular graphs.