ACCELERATING FEEDBACK-BASED ALGORITHMS FOR QUANTUM OPTIMIZATION USING GRADIENT DESCENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Feedback-based methods, built on Quantum Lyapunov Control (QLC), have gained significant attention as an alternative training paradigm for Variational Quantum Algorithms (VQAs) in solving combinatorial optimization problems such as MAX-CUT. These methods aim to reduce the training overhead associated with VQAs, and can mitigate barren plateaus. QLC methods often require long control sequences, and may show sensitivity to the evolution timestep, which can result in slow convergence compared to gradient-based VQA approaches. In this work, we address these challenges by proposing a hybrid method that incorporates per-layer gradient estimation to accelerate the convergence of QLC while preserving its low training overhead and stability guarantees. By leveraging layer-wise gradient information, our approach selects improved control parameters and achieves faster convergence rates compared to common QLC methods such as FALQON. We validate our algorithm through various empirical studies across a range of problem instances and optimization settings. We empirically show our algorithm's robustness against noise and variations in the timestep parameter. We also demonstrate a parameter initialization scheme for QAOA using our method that can lead to improved solution quality.

1 INTRODUCTION

Variational quantum algorithms (VQAs), and the Quantum Approximate Optimization Algorithm (QAOA) in particular, are among the leading candidates for solving combinatorial optimization problems on near-term quantum hardware (Farhi et al., 2014; Cerezo et al., 2021; Blekos et al., 2024). Their practical utility, however, is limited by the cost of training: the circuit parameters are tuned by a classical optimizer in an outer loop that repeatedly evaluates the circuit, which incurs a high measurement overhead and must navigate non-convex landscapes that can suffer from barren plateaus (Cerezo et al., 2021; Larocca et al., 2022). Quantum control theory offers a different route to the same goal (Gross et al., 1993; Sugawara, 2003; Doherty et al., 2000; Zhang et al., 2017; Brif et al., 2010; Cong & Meng, 2013). In *Quantum Lyapunov Control* (QLC), the control parameters are not searched for but computed: a feedback law derived from a Lyapunov function maps measurement outcomes on the current state directly to the next control value, guaranteeing that the objective is non-increasing along the trajectory without any classical optimization loop (Grivopoulos & Bamieh, 2003; Jamalinia et al., 2021; Magann et al., 2022b). Feedback-based algorithms built on this principle have been applied to finding ground states of quantum Hamiltonians (Larsen et al., 2023), quadratic unconstrained binary optimization (QUBO) problems (Blekos et al., 2024), and to combinatorial optimization, where the feedback-based algorithm for quantum optimization (FALQON) is the prominent example. FALQON grows a QAOA-like circuit one layer at a time, setting each layer's parameter from measurements on the state produced by the previous layers (Magann et al., 2022a;b).

One drawback of such feedback-based methods is (circuit) depth. Because the feedback law is derived from a first-order approximation of the dynamics over a short evolution time Δt, feedback-based algorithms require long and finely discretized control sequences, and the resulting circuits are considerably deeper than those of gradient-trained VQAs. For weighted MAX-CUT on 3-regular graphs, for instance, FALQON needs up to two orders of magnitude more layers than QAOA to reach

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

a comparable approximation ratio (Magann et al., 2022b). Increasing Δt to shorten the sequence is not a reliable solution, as the approximation underlying the feedback law breaks down and the control parameters become unstable (Arai et al., 2025). A number of extensions of FALQON have been proposed to address these issues, including second-order feedback laws (Arai et al., 2025), perturbative refinement of earlier layers (Brady & Hadfield, 2024), and counterdiabatic, imaginary-time, and other modified drivers (Malla et al., 2024; Rattighieri et al., 2025; Chandarana et al., 2024; Long et al., 2025). These methods improve on FALQON in specific regimes, but they share its basic design: each control parameter is fixed by a single closed-form evaluation of a feedback law, so the rate at which the Lyapunov function decreases is dictated by the accuracy of that approximation rather than optimized.

In this work we take a different approach and treat the per-layer control parameter as an intermediate optimization problem. We propose *GD-QLC*, a hybrid scheme that, at every layer, runs a small number L of stochastic gradient descent steps on the energy function using the same measured quantities that the feedback law is built from. The feedback rule of FALQON is recovered as the special case $L = 1$. We show that GD-QLC satisfies the Lyapunov condition and hence inherits the monotone-descent guarantee of feedback-based methods (Theorem 1) and that the per-layer optimization converges to a stationary point at an $O(\log L/\sqrt{L})$ rate under unbiased gradient estimates (Theorem 2). Since L is a fixed small constant, the measurement overhead per layer remains within a constant factor of FALQON's, and, as we show empirically, it is repaid by the significant reduction in the number of layers. Unlike gradient-trained VQAs, GD-QLC never revisits earlier layers, so the monotonicity guarantee of feedback control is preserved and the training cost does not grow with the depth of the circuit.

Through numerical experiments on MAX-CUT, weighted MAX-CUT, MAX-CLIQUE, and MIN-COVER, we show that GD-QLC consistently converges in fewer layers than FALQON and matches or outperforms its second-order variant (Arai et al., 2025) and FOCQS (Brady & Hadfield, 2024), while producing markedly better-behaved control trajectories and remaining stable across a wide range of timesteps Δt where the baselines degrade. We evaluate GD-QLC under finite-shot measurement noise and realistic hardware noise models and find that it retains its advantage in the low-shot regime relevant to Noisy Intermediate-Scale Quantum (NISQ) devices. Finally, we show that the control parameters produced by GD-QLC provide an effective warm start for QAOA, improving solution quality and substantially reducing the variance relative to random initialization.

2 PRELIMINARIES

Lyapunov control (LC) is a framework for steering the dynamics of a system toward a desired target by manipulating control parameters according to a feedback law derived from a *control Lyapunov function*. Consider a controlled dynamical system of the form

$$\dot{x} = f(x, u),$$

where $x \in \mathbb{R}^n$ denotes the system state and $u \in \mathbb{R}^m$ is the control input. Our objective is to drive the system to a desired equilibrium point x^* from any initial state.

Definition 1 (Lyapunov function)**.** *A continuously differentiable function $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is called a* Lyapunov function *if it satisfies: (i) $V(x)$ is positive definite with respect to x^*, i.e., $V(x) > 0$ for all $x \neq x^*$ and $V(x^*) = 0$; and (ii) for every $x \neq x^*$, there exists a control input u such that V is non-increasing along the system trajectories,*

$$\dot{V}(x) = \nabla V(x) \cdot f(x, u) \leq 0. \tag{1}$$

This condition guarantees that, for each state x, an appropriate choice of control input can be made to decrease the Lyapunov function, thereby driving the system asymptotically toward the equilibrium x^*. Consider a quantum dynamic governed by the Schrödinger equation with $\hbar = 1$:

$$i\frac{d}{dt}\left|\psi(t)\right\rangle = \left(H_p + \beta(t)H_d\right)\left|\psi(t)\right\rangle, \left|\psi(0)\right\rangle = \left|\psi_0\right\rangle, \tag{2}$$

where $|\psi(t)\rangle$ denotes the system state, and H_p and H_d are the drift and control Hamiltonians, respectively, and $\beta(t)$ acts as a control function. Systems of this form are frequently studied in *bilinear*

control theory (see for example (Panos M. Pardalos, 2008)). Typically, the problem is encoded in H_p such that the desired solution is the ground state, and the goal is to minimize

$$E_p(t) = \langle\psi(t)|\, H_p\, |\psi(t)\rangle\,. \tag{3}$$

We consider $E_p(t)$ as a Lyapunov *candidate* for monotonic objective improvement. Let $E_0 = \lambda_{\min}(H_p)$ be the minimum eigenvalue and define the shifted function

$$V(|\psi\rangle) := \langle\psi|\, H_p\, |\psi\rangle - E_0 \geq 0,$$

which equals 0 on the ground state of H_p. Since $\dot{V} = \frac{d}{dt}E_p(t)$, a feedback law that enforces $\frac{d}{dt}E_p(t) \leq 0$ guarantees monotonic non-increasing of the objective. We shall refer to the quantum system in equation 2 coupled with an admissible feedback control law that satisfies equation 1 as a *quantum Lyapunov control (QLC) scheme*. While QLC schemes guarantee non-increasing changes in the objective function, convergence to the ground state generally requires additional invariance/spectral conditions: for instance, if H_p has non-degenerate eigenvalues (i.e., a unique ground state) with nonzero eigenvalue gaps, then QLC is guaranteed to converge to the ground state (Magann et al., 2022a). For a system evolving as in equation 2, the time derivative of $E_p(t)$ is

$$\frac{d}{dt}E_p = \langle\psi(t)|i[H_d, H_p]|\psi(t)\rangle\,\beta(t) = A(t)\beta(t), \tag{4}$$

where $A(t) := \langle\psi(t)|i[H_d, H_p]|\psi(t)\rangle$. Choosing a feedback control law $\beta(t) = -F(A(t))$ where $F(\cdot)$ is a function respecting the sign of $A(t)$ yields

$$\frac{d}{dt}E_p = -A(t)F(A(t)) \leq 0,$$

which satisfies the QLC condition. As an example, $F(A(t)) = -A(t)$ leads to the control scheme known as FALQON (see Appendix E for further details). In order to realize these QLC schemes, quantum measurements must be performed to estimate the feedback. For this purpose, the system is evolved for a small timestep Δt before the resulting state is measured. For K iterations totaling $2K\Delta t$ time, the dynamics in equation 2 are approximated by

$$|\psi_K\rangle = \prod_{k=0}^{K} U_d(\beta_k)U_p\,|\psi_0\rangle\,,\; U_p = e^{-i\Delta t H_p},\; U_d(\beta) = e^{-i\Delta t\beta H_d}. \tag{5}$$

The state $|\psi_k\rangle$ can then be measured to estimate the feedback $A(2k\Delta t)$ which is then used to set the control parameter for the following iteration:

$$E_{p,k} = \langle\psi_k|\, H_p\, |\psi_k\rangle\,,\; A_k = \langle\psi_k|i[H_d, H_p]|\psi_k\rangle\,,\; \beta_{k+1} = -F(A_k). \tag{6}$$

For small enough Δt such that the approximations hold, $\{\beta_k\}$ provides a piecewise approximation of $\beta(t)$. The circuit in equation 5 is similar to the ansatz of QAOA (see Appendix D).

3 Main results

We propose a framework to optimize the control parameter β using gradient descent. Since the optimization loop itself can be viewed as a dynamic system, we propose a dynamical system with two time axes: the physical time t, and optimization time τ. Let $|\psi(t,\tau)\rangle$ be the state of the physical quantum system at time t followed by optimization times for τ units. The dynamic of the system is then described by the following differential equations:

$$\begin{cases} i\frac{\partial}{\partial t}\,|\psi(t,\,\tau)\rangle = \big(H_p + \beta(t,\,\tau)H_d\big)\,|\psi(t,\,\tau)\rangle & |\psi(0,\,\tau)\rangle = |\psi_0\rangle\,, \\ \frac{\partial}{\partial\tau}\beta(t,\,\tau) = -\eta(t,\,\tau)A(t,\,\tau) & \beta(t,0) = \beta(0,\tau) = 0 \end{cases} \tag{7}$$

where

$$A(t,\,\tau) := \langle\psi(t,\,\tau)|i[H_d,\,H_p]|\psi(t,\,\tau)\rangle$$

is the first order feedback and $\eta(t,\tau)$ is the learning rate. One can connect t and τ through a parameterized curve $t' \mapsto (t,\tau)$ and consider

$$E_p(t') := \langle\psi(t,\tau)|H_p|\psi(t,\tau)\rangle - E_0$$

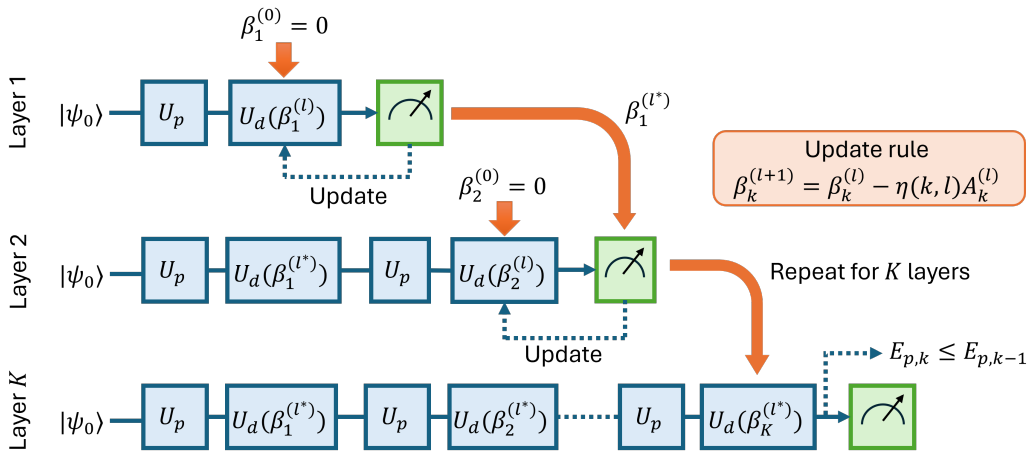

Figure 1: In our method, for each layer GD is used to update $\beta_k^{(l)}$. This process is repeated K times to construct all the layers.

as the Lyapunov function along the trajectory of t'. Different parameterized curves and/or dynamics for $\beta(t, \tau)$ result in a variety of optimization schemes.

In order to realize the continuous system in equation 7 we discretize both time variables. The physical time t is discritized by Δt as in equation 5 and the optimization time is discritized by time steps $\Delta \tau = 1$. Then, GD-QLC feedback law is obtained from the update rule

$$\beta_k^{(l+1)} = \beta_k^{(l)} - \eta(k,\, l)\, A_k^{(l)}, \qquad \beta_k^{(0)} = 0, \quad l = 0, \ldots, L-1, \tag{8}$$

where $A_k^{(l)} = \left\langle \psi_k^{(l)} \middle| i[H_d, H_p] \middle| \psi_k^{(l)} \right\rangle$ is evaluated at $\left|\psi_k^{(l)}\right\rangle = U_d(\beta_k^{(l)}) U_p \, |\psi_{k-1}\rangle$, and $\eta(k,\, l)$ is the learning rate. With this method, the discretized trajectory t' is given by increasing τ linearly with L steps, then increasing t by Δt and repeating the process.

The GD-QLC process is summarized in Algorithm 1: at each layer k, we perform L such updates, choose the best $\beta_k^{(l^*)}$, and proceed to the next layer. We use a learning rate that scales as $c_k/\sqrt{l+1}$ at iteration l of layer k. Intuitively, GD-QLC combines layerwise gradient descent update rule and FALQON feedback law as demonstrated in Figure 1. A single step ($L = 1$) with $\eta = 1$ reduces to the FALQON update $\beta_k = -A_k^{(0)}$. The following theorem, with the proof given in Appendix A, showes that GD-QLC satisfies the Lyapunov condition when Δt is small enough.

Theorem 1. *Let $E_{k,l} := \left\langle \psi_k^{(l)} \middle| H_p \middle| \psi_k^{(l)} \right\rangle$ be the energy at iteration l of layer k of GD-QLC, let $\dot{E}_{k,l} := A_k^{(l)} \beta_k^{(l)}$ be the corresponding value of the Lyapunov rate equation 4, and let l^* be the iterate selected by Algorithm 1, so that E_{k,l^*} is the energy after layer k (with $E_{0,l^*} := \langle \psi_0 | H_p | \psi_0 \rangle$). Suppose the learning rates satisfy $0 < \eta(k,\, l) \leq 1/\mathcal{L}_{GD}$, where $\mathcal{L}_{GD} := 4\Delta t \, \|H_d\|^2 \|H_p\|$. Then, for every layer k: (i) $\dot{E}_{k,l} \leq 0$ for all l, with strict inequality for $l \geq 1$ whenever $A_k^{(0)} \neq 0$ and $\eta(k,\, l) < 1/\mathcal{L}_{GD}$; and (ii) $E_{k,l}$ is non-increasing in l and $E_{k,l^*} \leq E_{k-1,l^*}$, i.e., the energy does not increase from one layer to the next. Moreover, for the selected parameter $\beta_k = \beta_k^{(l^*)}$,*

$$E_{k,l^*} - E_{k-1,l^*} \leq \Delta t \, \dot{E}_{k,l^*} + 2\|H_d\|^2 \|H_p\| \, \Delta t^2 \, \beta_k^2. \tag{9}$$

In particular, GD-QLC is a valid QLC scheme.

The bound in equation 9 shows that the per-layer descent is driven by the Lyapunov rate $\dot{E}_{k,l^*} \leq 0$ up to an $O(\Delta t^2)$ discretization term; its right-hand side is itself non-positive whenever $\Delta t \, |\beta_k| \leq |A_k^{(l^*)}| / \left(2\|H_d\|^2 \|H_p\|\right)$, which is the discrete-time analogue of the timestep condition of FALQON (Magann et al., 2022a).

Note that we do not necessarily have global convergence, since $E_{p,k}(\beta)$ is not necessarily convex and may not satisfy standard assumptions such as the PŁ condition, as is standard in the machine

learning and optimization literature (Bottou et al., 2018). Standard results in stochastic optimization, however, allow us to still guarantee convergence to a stationary point in expectation:

Theorem 2. *Fix a layer k and let the assumptions of Theorem 1 hold. Suppose the update rule equation 8 is run for L iterations with $A_k^{(l)}$ replaced by a measurement-based estimate $g_k^{(l)}$ that is unbiased given $\beta_k^{(l)}$, i.e., $\mathbb{E}\left[g_k^{(l)} \mid \beta_k^{(l)}\right] = A_k^{(l)}$, and has conditional variance at most σ^2. Let the learning rate be $\eta(k,\, l) = \frac{c_k}{\sqrt{l+1}}$ with $0 < c_k \leq 1/\mathcal{L}_{GD}$, and let $\hat{l} \in \{0, \ldots, L-1\}$ be drawn with probability proportional to $\eta(k,\, \hat{l})$. Then*

$$\mathbb{E}\Big[\big|A_k^{(\hat{l})}\big|^2\Big] \leq \frac{4\|H_p\|/\Delta t + \mathcal{L}_{GD}\,\sigma^2 c_k^2\,(1+\log L)}{c_k\sqrt{L}} = O\Big(\frac{\log L}{\sqrt{L}}\Big), \tag{10}$$

where the expectation is taken over the measurement randomness and the draw of $\hat{l}$. Since $\Delta t\, A_k^{(l)}$ is the derivative of the energy $E_{k,l}$ with respect to the control parameter $\beta_k^{(l)}$, the inner loop converges in expectation to a stationary point of the per-layer energy.

This result proved in Appendix A.3 gives a bound on the number of iterations required to reach a neighborhood of a stationary point. Note that L controls the trade-off between convergence speed and measurement overhead, as each gradient update requires additional measurements to estimate the necessary observables. Our empirical experiments suggest that $L < 7$ is sufficient for most scenarios.

Algorithm 1 GD-QLC

Input: problem Hamiltonian H_p, control Hamiltonian H_d, initial state $|\psi_0\rangle$, timestep Δt, iteration limit K, gradient descent iterations L, learning rate constant c

1: Define $U_p = e^{-i\Delta t H_p}$ and $U_d(\beta) = e^{-i\Delta t \beta H_d}$
2: **for** $k = 1$ to K **do**
3: Initialize $\beta_k^{(0)} = 0$
4: **for** $l = 1$ to L **do**
5: Prepare the state $\left|\psi_k^{(l)}\right\rangle := U_d(\beta_k^{(l-1)})U_p\,|\psi_{k-1}\rangle$
6: Estimate $A_k^{(l)} = \left\langle\psi_k^{(l)}\middle|i[H_d,\, H_p]\middle|\psi_k^{(l)}\right\rangle$.
7: $\beta_k^{(l)} \leftarrow \beta_k^{(l-1)} - \frac{c}{\sqrt{l+1}\log k} A_k^{(l)}$
8: Choose $\beta_k^{(l^*)}$ that minimizes $E_{p,k}$ among $\beta_k^{(l)}$'s.
9: Set $|\psi_k\rangle = U_d(\beta_k^{(l^*)})U_p\,|\psi_{k-1}\rangle$.
10: Return $\beta_k^{(l^*)}$ and the state $|\psi_K\rangle$

Overhead comparison: GD-QLC measures $A_k^{(l)}, l \in [L]$. Hence, its overhead has a factor L compared to that of FALQON. For illustration purposes, assume that H_p and H_d consist of m_p and m_d Pauli terms, respectively. Estimating A_k is done with $O(m_p m_d)$ measurements. Hence, at each layer, GD-QLC with L iterations needs $O(L\, m_p\, m_d)$ measurements. We show in the next section that this additional overhead is justified by the significant reduction in the number of layers K needed to achieve comparable performance.

4 Numerical simulation

To evaluate GD-QLC, we consider four well-studied graph-based combinatorial optimization problems: MAX-CUT, weighted MAX-CUT, MAX-CLIQUE, and MIN-COVER[1]. These problems are standard benchmarks for quantum optimization, as they admit QUBO formulations and can therefore be mapped to Ising Hamiltonians (Hadfield et al., 2019; Lucas, 2014). We compare GD-QLC against FALQON (Magann et al., 2022b), its second-order variant SO-FALQON (Arai et al., 2025),

[1] All experiments were performed using classical simulations.

and Feedback Optimally Controlled Quantum States (FOCQS) (Brady & Hadfield, 2024). FOCQS extends QLC by using the collected first-order feedback $\{A_k\}$ to perturbatively refine the control parameters of all previous layers. Unless stated otherwise, GD-QLC uses $L = 5$ gradient iterations per layer; Δt and the number of qubits vary across experiments and are reported with each figure. Shaded regions indicate $\pm 1\sigma$ across problem instances. Further details are given in Appendix B.

Performance is evaluated as a function of the number of layers k through the approximation ratio and the success probability,

$$r_A(k) = \frac{E_{p,k}}{E_{p,\min}}, \qquad p(k) = \sum_i |\langle \psi_k | q_i \rangle|^2 ,$$

where $E_{p,\min}$ is the minimum eigenvalue of H_p and $\{|q_i\rangle\}$ are the computational basis states encoding the (possibly degenerate) optimal solutions; that is, $p(k)$ is the probability that measuring $|\psi_k\rangle$ yields a globally optimal solution of the original problem. For problems cast as maximization, $r_A(k) = E_{p,k}/E_{p,\max}$ with $E_{p,\max}$ the maximum eigenvalue of H_p. We do not include QAOA as a baseline in the convergence comparisons, since FALQON and SO-FALQON have already been compared directly against QAOA (Magann et al., 2022b; Arai et al., 2025); QAOA appears instead in the warm-starting study at the end of this section.

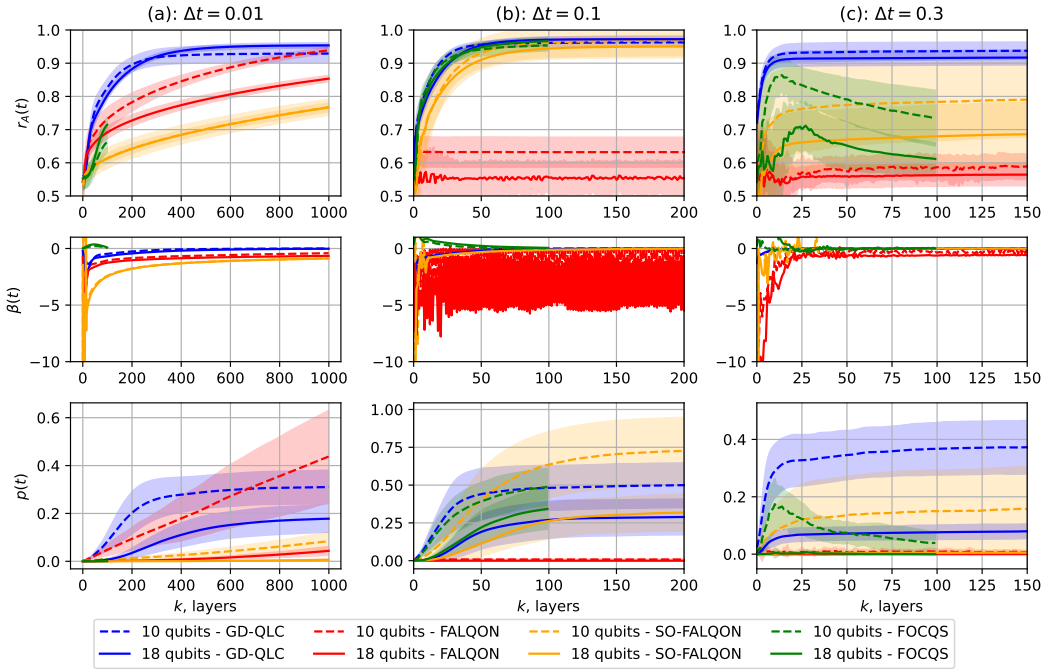

Figure 2: Weighted MAX-CUT on 10 and 18 qubits using GD-QLC, FALQON, SO-FALQON, and FOCQS with $\Delta t = 0.01$ (column (a)), 0.1 (b), and 0.3 (c). Rows show the approximation ratio r_A (top), the control parameters β_k (middle), and the success probability p (bottom). FOCQS was simulated only up to $K = 100$ layers due to computational limitations. The ranges of the x-axis and of β are clipped where possible for visualization.

Comparison with feedback-based baselines: Figure 2 compares GD-QLC with FALQON, SO-FALQON, and FOCQS on weighted MAX-CUT for $\Delta t \in \{0.01, 0.1, 0.3\}$, representative of small, medium, and large timesteps in practice. In all three regimes GD-QLC reaches a higher r_A than FALQON and does so in substantially fewer layers. Its r_A is comparable to or better than that of SO-FALQON and FOCQS, and its success probability is competitive with both; the main exception is $\Delta t = 0.1$, where SO-FALQON eventually reaches a higher p on the 10-qubit instances, albeit with the unstable control trajectories discussed next.

Stability of the control trajectories: The middle row of Figure 2 shows that GD-QLC produces smooth, bounded control sequences, whereas FALQON and SO-FALQON exhibit large spikes in β_k. Although SO-FALQON is designed to tolerate larger Δt through its second-order approximation, its feedback law frequently selects very large $|\beta_k|$ in the early layers, which is problematic for practical

implementation, as noted in the original work (Arai et al., 2025). For $\Delta t = 0.1$ we observed spikes in β_k exceeding 4000 in magnitude for FALQON and SO-FALQON (clipped in the figure), while GD-QLC avoids such values entirely and still converges rapidly in both r_A and p.

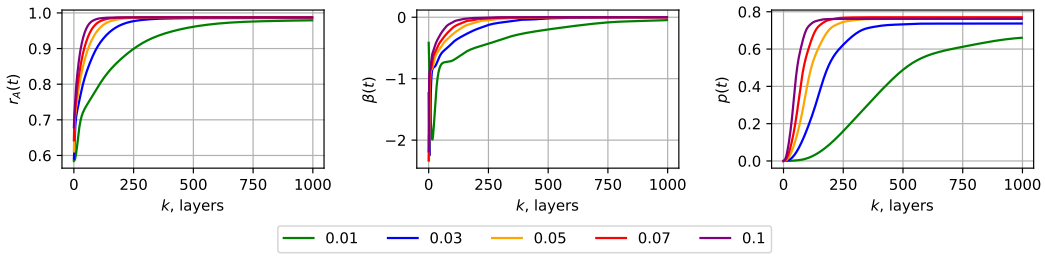

Figure 3: GD-QLC on weighted MAX-CUT over 18 vertices (qubits) for different values of Δt.

Robustness to the timestep Δt: The choice of Δt involves a trade-off: when Δt is too small, each layer induces only a weak evolution and convergence is slow; when it is too large, the first-order approximations underlying the feedback laws break down. Since tuning Δt per instance is costly, robustness to this choice is valuable in practice. Figure 2 already indicates that GD-QLC tolerates increasing Δt well: its performance degrades only gracefully from $\Delta t = 0.01$ to 0.3, whereas FALQON stalls near $r_A \approx 0.6$ for $\Delta t \geq 0.1$. Figure 3 isolates this effect for GD-QLC on 18-qubit weighted MAX-CUT with $\Delta t \in \{0.01, 0.03, 0.05, 0.07\}$. For every timestep tested, r_A plateaus at essentially the same value within roughly 500 layers, with larger timesteps converging faster, and moderate values ($\Delta t = 0.03$–0.05) also yield the highest success probability.

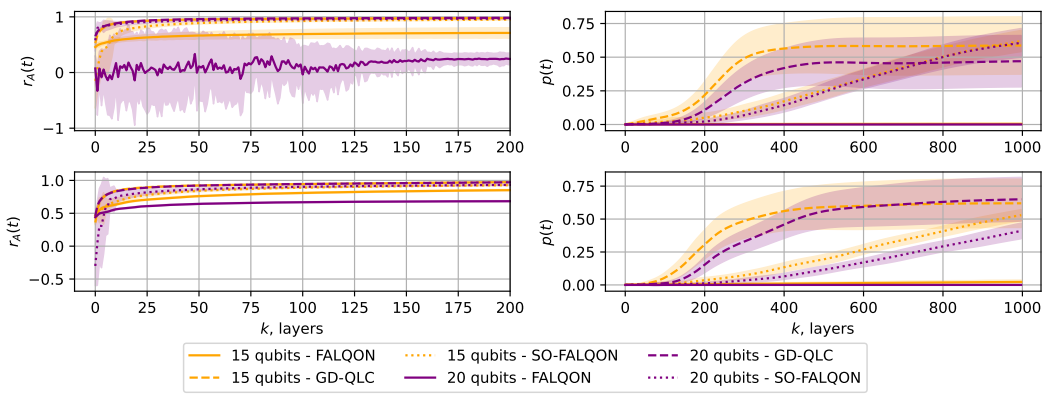

Figure 4: GD-QLC against FALQON and SO-FALQON for MAX-CLIQUE (top row) and MIN-COVER (bottom row) on 15 and 20 qubits with $\Delta t = 0.005$. FOCQS was omitted due to computational limitations.

Other combinatorial problems: Figure 4 reports analogous comparisons for MAX-CLIQUE and MIN-COVER on 15 and 20 qubits. Here a smaller timestep, $\Delta t = 0.005$, is used for all methods: typically $\|H_{p,\,\text{MIN-COVER}}\| > \|H_{p,\,\text{MAX-CUT}}\|$, and similarly for MAX-CLIQUE, so FALQON requires a smaller Δt to remain stable. Both GD-QLC and SO-FALQON clearly outperform FALQON on both problems, and GD-QLC matches the final r_A and p of SO-FALQON in fewer layers. Larger problem sizes were not tractable within our classical simulation budget.

Effect of the number of GD iterations L: Figure 5 shows the effect of the number of per-layer GD iterations on weighted MAX-CUT with $\Delta t = 0.1$; we expect similar behavior on the other problems. We show one curve that is computed with $L = 1$ and the GD-QLC learning rate, which translates to a gain scaled version of FALQON. We also include a curve with the FALQON optimization rule. Already a small L yields substantial gains over FALQON, consistent with the sharp initial progress typical of GD, and the improvement in r_A saturates as L grows, indicating diminishing returns beyond a moderate number of gradient steps. Two further effects are noteworthy. First, larger L continues to improve the success probability p even after r_A has saturated. Second, larger L regularizes the control amplitudes, suppressing the sharp early transients in β_k observed for small

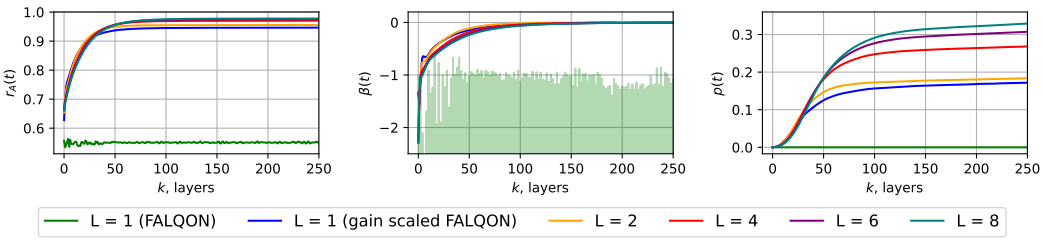

Figure 5: GD-QLC on weighted MAX-CUT with $\Delta t = 0.1$ for different numbers of per-layer GD iterations L.

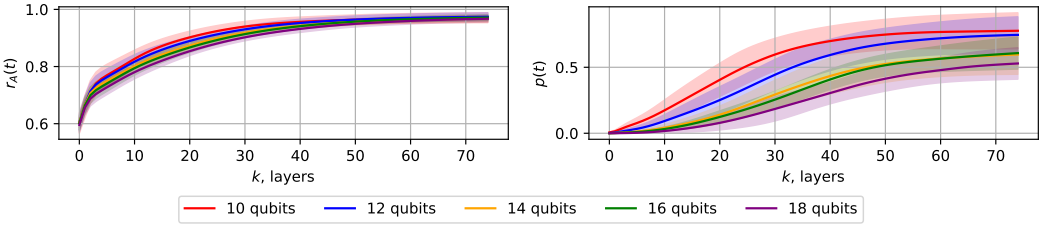

Figure 6: GD-QLC on MAX-CUT with $\Delta t = 0.1$ for different numbers of qubits.

L (middle panel). Per-layer GD thus accelerates convergence while also improving control stability and solution quality, at a measurement overhead that is linear in L (Section 3).

Effect of problem size: Figure 6 shows GD-QLC on MAX-CUT with $\Delta t = 0.1$ for 10 to 18 qubits. The approximation ratio is essentially unaffected by the system size over this range, while the success probability decreases with the number of qubits, reflecting the shrinking fraction of optimal states in the search space.

Effect of the shot budget: We next account for finite-shot estimation of the feedback. Figure 7 shows MAX-CUT on 14 qubits with $\Delta t = 0.3$, averaged over 50 random graphs, under total per-iteration budgets of 100, 250, and 1000 shots. The budget is divided evenly among the distinct circuits each algorithm requires per iteration (qubit-wise commuting observables are measured from a single circuit), so methods needing more circuits per iteration, such as GD-QLC, receive fewer shots per circuit. Under all three budgets GD-QLC attains the highest r_A and p and is largely insensitive to measurement noise, whereas the baselines degrade markedly; SO-FALQON and FOCQS in particular lose much of their noiseless advantage.

Hardware noise: To assess robustness to realistic device noise, we simulated the methods using Qiskit's hardware noise profiles (Javadi-Abhari et al., 2024), which emulate the noise profiles of real-world quantum processors. We solved MAX-CUT on 50 random 10-vertex graphs with 10-layer circuits, $\Delta t = 0.3$, and 100 shots per observable (Figure 8 (a)). GD-QLC reaches a lower energy E_p than other methods even at low depth and exhibits a smoother control trajectory, indicating that its advantage persists under hardware noise.

GD-QLC as an initialization for QAOA: Finally, GD-QLC can serve as a warm start for QAOA. After running GD-QLC (Algorithm 1), we initialize a K-layer QAOA circuit with $\beta_k \leftarrow \Delta t\, \beta_k^{(l^*)}$ and $\gamma_k \leftarrow \Delta t$ for all k, and then optimize the QAOA parameters as usual (see Appendix D). We call this scheme QAOA+GD-QLC. QAOA+FALQON is defined analogously and was introduced in (Magann et al., 2022a). Figure 8 (b) compares QAOA+GD-QLC with randomly initialized QAOA, GD-QLC, FALQON, and QAOA+FALQON on 10-qubit MAX-CUT with $\Delta t = 0.3$ and 10-layer circuits, simulated with Qiskit's hardware noise profiles under 100 and 1000 shots per observable; the violins show the distribution of the final r_A over 50 random instances. Randomly initialized QAOA exhibits a wide spread with a long tail of poor solutions, GD-QLC alone is already markedly better and more consistent, and QAOA+GD-QLC attains the highest and most concentrated r_A. GD-QLC thus provides a reliable initialization that improves both the quality and consistency.

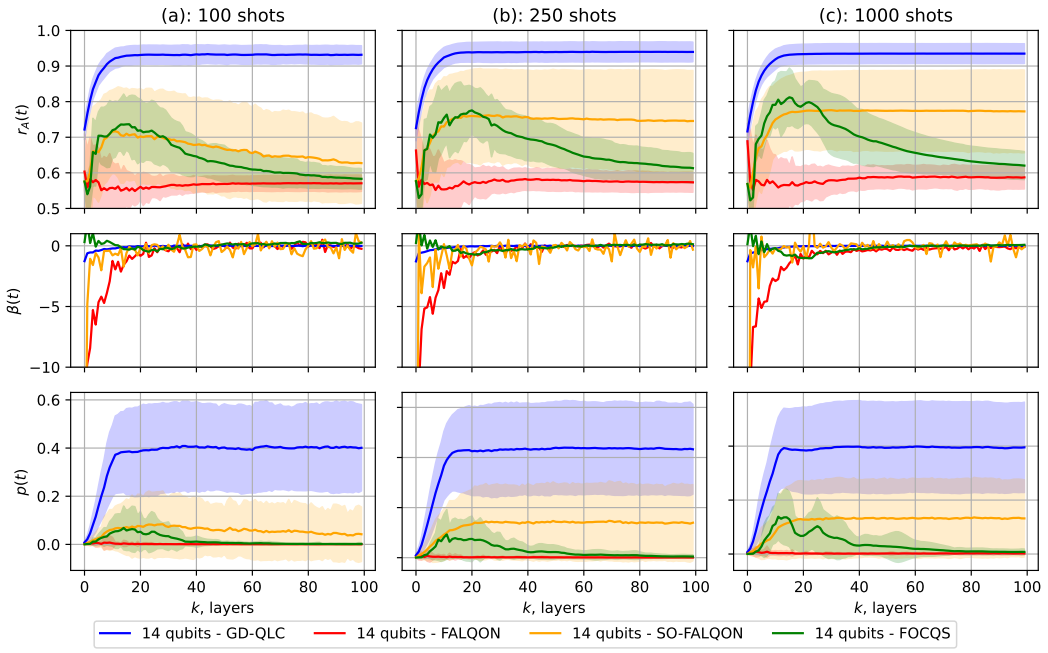

Figure 7: MAX-CUT on 14 qubits with $\Delta t = 0.3$ using GD-QLC (with $L = 5$), FALQON, SO-FALQON, and FOCQS under three per-iteration shot budgets: 100 (column (a)), 250 (b), and 1000 (c). The range of β is clipped to $[-10, 0]$ for visualization.

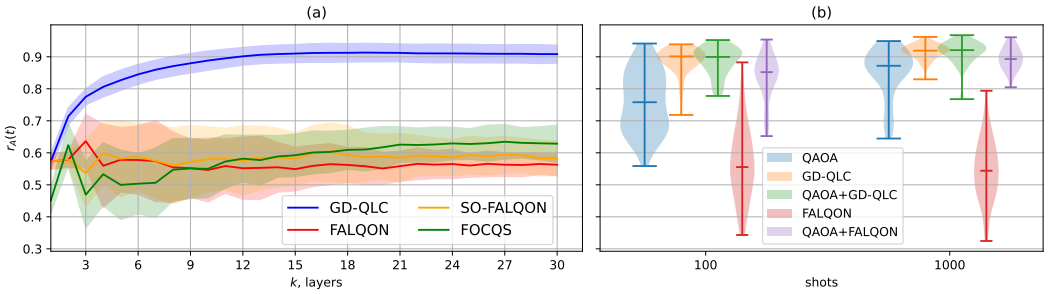

Figure 8: Final approximation ratio on 10-qubit MAX-CUT for randomly initialized QAOA, GD-QLC, and QAOA warm-started with GD-QLC (QAOA+GD-QLC), with 50 and 500 shots per observable; distributions over 50 random instances.

5 Conclusion and Future Work

We introduced GD-QLC, a hybrid optimization framework that integrates per-layer GD updates into feedback-based QLC methods to accelerate convergence and reduce circuit depth for combinatorial optimization problems. We established the this method satisfies the Lyapunov condition and has per-layer convergence guarantees. Through extensive empirical studies, we show the effectiveness of our method compared to common existing works. We further showed that GD-QLC exhibits strong robustness to variations in the timestep and hardware noise compared to existing methods.

Limitations: In terms of theory, the main limitation of this work is the absence of overall convergence guarantees as a function of number of layers. Such guarantees might require study of *bilinear control theory*, which is out of the scope of this paper, and it is not clear whether a straightforward analysis exists. Another limitation is per-layer convergence to global minima that seems out of reach with the current literature for non-convex optimization. Regarding the empirical results, (1) scalability to larger system sizes and (2) implementation on real hardware are the primary limitations. To address these one needs implementation on large scale quantum hardware (with possibly hundreds of qubits) that would require extensive computational resources. Nevertheless, our theoretical results seem to align with out empirical studies.

AI USE STATEMENT

In this work, we used generative AI tools for help with polishing the text of the main paper and for improving the presentation of proofs to be more concise and readable. The main ideas, and theoretical arguments were mostly done by the authors except Step 4 in the proof of Theorem 1 where generative AI is used to derive equation 25. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We have included all data and procedures required for reproduction of our work in the body of the main text (particularly Section 4 and Appendix B).

REFERENCES

Don Arai, Ken N. Okada, Yuichiro Nakano, Kosuke Mitarai, and Keisuke Fujii. Scalable circuit depth reduction in feedback-based quantum optimization with a quadratic approximation. *Phys. Rev. Res.*, 7:013035, Jan 2025. doi: 10.1103/PhysRevResearch.7.013035. URL https://link.aps.org/doi/10.1103/PhysRevResearch.7.013035.

Kostas Blekos, Dean Brand, Andrea Ceschini, Chiao-Hui Chou, Rui-Hao Li, Komal Pandya, and Alessandro Summer. A review on quantum approximate optimization algorithm and its variants. *Physics Reports*, 1068:1–66, 2024. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2024.03.002. URL https://www.sciencedirect.com/science/article/pii/S0370157324001078.

Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning, 2018. URL https://arxiv.org/abs/1606.04838.

Lucas T. Brady and Stuart Hadfield. Focqs: Feedback optimally controlled quantum states, 2024. URL https://arxiv.org/abs/2409.15426.

Constantin Brif, Raj Chakrabarti, and Herschel Rabitz. Control of quantum phenomena: past, present and future. *New Journal of Physics*, 12(7):075008, jul 2010. doi: 10.1088/1367-2630/12/7/075008. URL https://dx.doi.org/10.1088/1367-2630/12/7/075008.

M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, August 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00348-9. URL http://dx.doi.org/10.1038/s42254-021-00348-9.

Pranav Chandarana, Koushik Paul, Kasturi Ranjan Swain, Xi Chen, and Adolfo del Campo. Lyapunov controlled counterdiabatic quantum optimization, 2024. URL https://arxiv.org/abs/2409.12525.

Shuang Cong and Fangfang Meng. A survey of quantum lyapunov control methods. *The Scientific World Journal*, 2013(1), January 2013. ISSN 1537-744X. doi: 10.1155/2013/967529. URL http://dx.doi.org/10.1155/2013/967529.

Andrew C. Doherty, Salman Habib, Kurt Jacobs, Hideo Mabuchi, and Sze M. Tan. Quantum feedback control and classical control theory. *Physical Review A*, 62(1), June 2000. ISSN 1094-1622. doi: 10.1103/physreva.62.012105. URL http://dx.doi.org/10.1103/PhysRevA.62.012105.

Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm, 2014. URL https://arxiv.org/abs/1411.4028.

Guillaume Garrigos and Robert M. Gower. Handbook of convergence theorems for (stochastic) gradient methods, 2024. URL https://arxiv.org/abs/2301.11235.

Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.

S. Grivopoulos and B. Bamieh. Lyapunov-based control of quantum systems. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No.03CH37475)*, volume 1, pp. 434–438 Vol.1, 2003. doi: 10.1109/CDC.2003.1272601.

Peter Gross, Harjinder Singh, Herschel Rabitz, Kenneth Mease, and G. M. Huang. Inverse quantum-mechanical control: A means for design and a test of intuition. *Phys. Rev. A*, 47:4593–4604, Jun 1993. doi: 10.1103/PhysRevA.47.4593. URL https://link.aps.org/doi/10.1103/PhysRevA.47.4593.

Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2), 2019. ISSN 1999-4893. doi: 10.3390/a12020034. URL https://www.mdpi.com/1999-4893/12/2/34.

Elham Jamalinia, Peyman Azodi, Alireza Khayatian, and Peyman Setoodeh. Lyapunov-based stabilization and control of closed quantum systems, 2021. URL https://arxiv.org/abs/2102.00344.

Ali Javadi-Abhari, Matthew Treinish, Kevin Krsulich, Christopher J. Wood, Jake Lishman, Julien Gacon, Simon Martiel, Paul D. Nation, Lev S. Bishop, Andrew W. Cross, Blake R. Johnson, and Jay M. Gambetta. Quantum computing with Qiskit, 2024.

Martin Larocca, Piotr Czarnik, Kunal Sharma, Gopikrishnan Muraleedharan, Patrick J. Coles, and M. Cerezo. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6:824, September 2022. ISSN 2521-327X. doi: 10.22331/q-2022-09-29-824. URL http://dx.doi.org/10.22331/q-2022-09-29-824.

James B. Larsen, Matthew D. Grace, Andrew D. Baczewski, and Alicia B. Magann. Feedback-based quantum algorithms for ground state preparation, 2023. URL https://arxiv.org/abs/2303.02917.

Thanh Nguyen Van Long, Lan Nguyen Tran, and Le Bin Ho. Imaginary-time-enhanced feedback-based quantum algorithms for universal ground-state preparation, 2025. URL https://arxiv.org/abs/2512.13044.

Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2, 2014. ISSN 2296-424X. doi: 10.3389/fphy.2014.00005. URL http://dx.doi.org/10.3389/fphy.2014.00005.

Alicia B. Magann, Kenneth M. Rudinger, Matthew D. Grace, and Mohan Sarovar. Lyapunov-control-inspired strategies for quantum combinatorial optimization. *Physical Review A*, 106 (6), December 2022a. ISSN 2469-9934. doi: 10.1103/physreva.106.062414. URL http://dx.doi.org/10.1103/PhysRevA.106.062414.

Alicia B. Magann, Kenneth M. Rudinger, Matthew D. Grace, and Mohan Sarovar. Feedback-based quantum optimization. *Physical Review Letters*, 129(25), December 2022b. ISSN 1079-7114. doi: 10.1103/physrevlett.129.250502. URL http://dx.doi.org/10.1103/PhysRevLett.129.250502.

Rajesh K. Malla, Hiroki Sukeno, Hongye Yu, Tzu-Chieh Wei, Andreas Weichselbaum, and Robert M. Konik. Feedback-based quantum algorithm inspired by counterdiabatic driving, 2024. URL https://arxiv.org/abs/2401.15303.

Markus Meringer. Fast generation of regular graphs and construction of cages. *Journal of Graph Theory*, 30(2):137–146, 1999. doi: https://doi.org/10.1002/(SICI)1097-0118(199902)30:2⟨137::AID-JGT7⟩3.0.CO;2-G. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-0118%28199902%2930%3A2%3C137%3A%3AAID-JGT7%3E3.0.CO%3B2-G.

Vitaliy Yatsenko Panos M. Pardalos. *Optimization and Control of Bilinear Systems*. Springer New York, NY, 2008. ISBN 978-0-387-73669-3. doi: 10.1007/978-0-387-73669-3.

L. A. M. Rattighieri, G. E. L. Pexe, B. L. Bernado, and F. F. Fanchini. Accelerating feedback-based quantum algorithms through time rescaling, 2025. URL https://arxiv.org/abs/2504.01256.

M. Sugawara. General formulation of locally designed coherent control theory for quantum system. *The Journal of Chemical Physics*, 118(15):6784–6800, 04 2003. ISSN 0021-9606. doi: 10.1063/1.1559680. URL https://doi.org/10.1063/1.1559680.

Jing Zhang, Yu xi Liu, Re-Bing Wu, Kurt Jacobs, and Franco Nori. Quantum feedback: Theory, experiments, and applications. *Physics Reports*, 679:1–60, 2017. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2017.02.003. URL https://www.sciencedirect.com/science/article/pii/S0370157317300479. Quantum feedback: theory, experiments, and applications.

A PROOFS

This section contains the proofs of Theorems 1 and 2. We first fix notation and collect the elementary identities on which both proofs rely.

A.1 NOTATION AND PRELIMINARY IDENTITIES

Fix a layer $k \geq 1$ and the state $|\psi_{k-1}\rangle$ produced by the previous layers (with $|\psi_0\rangle$ the initial state). For a control parameter $\beta \in \mathbb{R}$ let

$$|\psi_k(\beta)\rangle := U_d(\beta)\, U_p\, |\psi_{k-1}\rangle = e^{-i\Delta t\, \beta H_d}\, e^{-i\Delta t H_p}\, |\psi_{k-1}\rangle \tag{11}$$

be the state produced by layer k when the control β is applied, so that $\left|\psi_k^{(l)}\right\rangle = \left|\psi_k(\beta_k^{(l)})\right\rangle$. Define the real-valued functions

$$\begin{aligned} E_k(\beta) &:= \langle\psi_k(\beta)|H_p|\psi_k(\beta)\rangle\,, \\ A_k(\beta) &:= \langle\psi_k(\beta)|i[H_d, H_p]|\psi_k(\beta)\rangle\,, \\ B_k(\beta) &:= \langle\psi_k(\beta)|[H_d,[H_d, H_p]]|\psi_k(\beta)\rangle\,, \end{aligned}$$

so that $E_{k,l} = E_k(\beta_k^{(l)})$, $A_k^{(l)} = A_k(\beta_k^{(l)})$ and $\dot{E}_{k,l} = \beta_k^{(l)} A_k^{(l)}$. Two facts are used repeatedly. First, since $U_p = e^{-i\Delta t H_p}$ commutes with H_p and $U_d(0) = I$,

$$E_k(0) = \langle\psi_{k-1}|U_p^\dagger H_p U_p|\psi_{k-1}\rangle = \langle\psi_{k-1}|H_p|\psi_{k-1}\rangle\,, \tag{12}$$

i.e., the energy at the start of layer k equals the energy at the end of layer $k-1$: $E_{k,0} = E_{k-1,l^*}$ (for $k = 1$, $E_{1,0} = \langle\psi_0|H_p|\psi_0\rangle$). Second, the state depends on β only through the single exponential $e^{-i\Delta t\, \beta H_d}$, which yields exact derivative formulas.

Lemma 1. *For any Hermitian operator O,*

$$\frac{\mathrm{d}}{\mathrm{d}\beta}\, \langle\psi_k(\beta)|O|\psi_k(\beta)\rangle = \Delta t\, \langle\psi_k(\beta)|i[H_d, O]|\psi_k(\beta)\rangle\,.$$

Consequently,

$$E_k'(\beta) = \Delta t\, A_k(\beta), \qquad A_k'(\beta) = -\Delta t\, B_k(\beta), \qquad E_k''(\beta) = -\Delta t^2\, B_k(\beta), \tag{13}$$

and

$$|A_k(\beta)| \leq 2\|H_d\|\, \|H_p\|, \qquad |B_k(\beta)| \leq 4\|H_d\|^2\, \|H_p\|, \tag{14}$$

where $\|\cdot\|$ denotes the operator norm.

Proof. Since $\frac{\mathrm{d}}{\mathrm{d}\beta}\, |\psi_k(\beta)\rangle = -i\Delta t\, H_d\, |\psi_k(\beta)\rangle$,

$$\frac{\mathrm{d}}{\mathrm{d}\beta}\, \langle\psi_k(\beta)|O|\psi_k(\beta)\rangle = \Delta t\, \langle\psi_k(\beta)|iH_dO - iOH_d|\psi_k(\beta)\rangle = \Delta t\, \langle\psi_k(\beta)|i[H_d, O]|\psi_k(\beta)\rangle\,.$$

Applying this with $O = H_p$ gives the first identity in equation 13; applying it with $O = i[H_d, H_p]$ gives $A_k'(\beta) = \Delta t\, \langle\psi_k(\beta)|i[H_d, i[H_d, H_p]]|\psi_k(\beta)\rangle = -\Delta t\, B_k(\beta)$, and the expression for E_k'' follows by combining the two. Finally, $|\,\langle\psi|O|\psi\rangle\,| \leq \|O\|$ for a unit vector $|\psi\rangle$ and $\|[X, Y]\| \leq 2\|X\|\, \|Y\|$, which gives equation 14. □

We stress that Lemma 1 is exact: no Trotter or small-Δt approximation is involved.

We will also use the following standard fact.

Lemma 2 (Descent lemma, see e.g. Garrigos & Gower (2024)). *If $f : \mathbb{R} \to \mathbb{R}$ is $\mathcal{L}$-smooth, i.e., differentiable with $|f'(x) - f'(y)| \le \mathcal{L}|x - y|$ for all x, y, then for every x and every $\gamma > 0$*

$$f\big(x - \gamma f'(x)\big) - f(x) \le -\gamma \left(1 - \frac{\gamma \mathcal{L}}{2}\right) |f'(x)|^2. \tag{15}$$

In particular, if $0 < \gamma \le 1/\mathcal{L}$ then $f(x - \gamma f'(x)) \le f(x) - \frac{\gamma}{2}|f'(x)|^2$.

A.2 Proof of Theorem 1

Theorem 1 makes two claims: (i) *within a layer*, every gradient iterate satisfies the Lyapunov sign condition $\dot{E}_{k,l} \le 0$; and (ii) *across layers*, the energy does not increase, with an explicit bound on the per-layer change. The proof proceeds in four steps: we identify the update rule equation 8 as gradient descent on a smooth function; we apply the descent lemma to obtain monotone descent within a layer; we show that the gradient iterates never change the sign of the feedback A_k, which gives (i); and we derive (ii) from equation 12 together with an exact second-order expansion of the energy in the control.

Proof of Theorem 1. **Step 1: smoothness of the objective.** Let

$$f_k(\beta) := \frac{E_k(\beta)}{\Delta t}, \tag{16}$$

so that, by Lemma 1, $f_k'(\beta) = A_k(\beta)$ and the update rule equation 8 reads

$$\beta_k^{(l+1)} = \beta_k^{(l)} - \eta(k,l)\, f_k'\big(\beta_k^{(l)}\big),$$

i.e., it is exactly gradient descent on f_k (equivalently, on the energy E_k) with step size $\eta(k,l)$, started at $\beta_k^{(0)} = 0$. By the mean value theorem, a twice differentiable function is $\mathcal{L}$-smooth with $\mathcal{L} = \sup_\beta |f_k''(\beta)|$, and by Lemma 1

$$|f_k''(\beta)| = \Delta t\, |B_k(\beta)| \le 4\Delta t\, \|H_d\|^2\, \|H_p\| = \mathcal{L}_{GD} \qquad \text{for all } \beta. \tag{17}$$

Hence f_k is $\mathcal{L}_{GD}$-smooth for every layer k, with a constant that depends only on Δt, $\|H_d\|$ and $\|H_p\|$. Equivalently, $A_k = f_k'$ is $\mathcal{L}_{GD}$-Lipschitz:

$$|A_k(\beta) - A_k(\beta')| \le \mathcal{L}_{GD}\, |\beta - \beta'| \qquad \text{for all } \beta, \beta'. \tag{18}$$

Step 2: monotone descent within a layer. Applying Lemma 2 to f_k with $\gamma = \eta(k,l) \le 1/\mathcal{L}_{GD}$ and multiplying by Δt gives, for every $l \ge 0$,

$$E_{k,l+1} \;\le\; E_{k,l} - \frac{\Delta t\, \eta(k,l)}{2} \big(A_k^{(l)}\big)^2 \;\le\; E_{k,l}. \tag{19}$$

Thus $E_{k,l}$ is non-increasing in l; in particular $E_{k,l} \le E_{k,0} = E_k(0)$ for every l, with strict inequality for every $l \ge 1$ whenever $A_k^{(0)} \ne 0$.

Step 3: sign condition (claim (i)). We claim that the iterates never change the sign of the feedback: for every $l \ge 0$,

$$A_k^{(l+1)} A_k^{(l)} \ge 0 \qquad \text{and} \qquad \Big|A_k^{(l+1)}\Big| \ge \big(1 - \mathcal{L}_{GD}\, \eta(k,l)\big) \Big|A_k^{(l)}\Big|. \tag{20}$$

Indeed, if $A_k^{(l)} = 0$ then $\beta_k^{(l+1)} = \beta_k^{(l)}$ and $A_k^{(l+1)} = 0$. If $A_k^{(l)} > 0$, then $\beta_k^{(l+1)} - \beta_k^{(l)} = -\eta(k,l) A_k^{(l)} < 0$ and by equation 18

$$A_k^{(l+1)} \ge A_k^{(l)} - \mathcal{L}_{GD} \Big|\beta_k^{(l+1)} - \beta_k^{(l)}\Big| = \big(1 - \mathcal{L}_{GD}\, \eta(k,l)\big)\, A_k^{(l)} \ge 0,$$

because $\eta(k,l) \le 1/\mathcal{L}_{GD}$. The case $A_k^{(l)} < 0$ is symmetric, which proves equation 20. Consequently all of $A_k^{(0)}, A_k^{(1)}, \dots$ have the same (weak) sign as $A_k^{(0)}$. Since $\beta_k^{(0)} = 0$, unrolling the update rule gives

$$\beta_k^{(l)} = -\sum_{j=0}^{l-1} \eta(k,j)\, A_k^{(j)},$$

and therefore

$$\dot{E}_{k,l} = A_k^{(l)}\, \beta_k^{(l)} = -\sum_{j=0}^{l-1} \eta(k,j)\, A_k^{(j)} A_k^{(l)} \le 0 \qquad \text{for all } l \ge 0, \tag{21}$$

because every product $A_k^{(j)} A_k^{(l)}$ is non-negative. This is the Lyapunov sign condition equation 4 evaluated at the state $\left|\psi_k^{(l)}\right\rangle$ with the control $\beta_k^{(l)}$, and it holds for every iterate, in particular for the parameter $\beta_k = \beta_k^{(l^*)}$ selected by Algorithm 1. If moreover $A_k^{(0)} \neq 0$ and $\eta(k,l) < 1/\mathcal{L}_{GD}$, then the second part of equation 20 shows by induction that $A_k^{(l)} \neq 0$ for all l, so the term $j = 0$ in equation 21 is strictly negative and $\dot{E}_{k,l} < 0$ for all $l \ge 1$. This proves claim (i).

Step 4: monotone descent across layers (claim (ii)). Let $\beta_k^* = \beta_k^{(l^*)}$ be the parameter selected at layer k, so that $|\psi_k\rangle = |\psi_k(\beta_k^*)\rangle$ and $E_{k,l^*} = E_k(\beta_k^*)$. By Step 2 and equation 12,

$$E_{k,l^*} \le E_{k,0} = \langle \psi_{k-1} | H_p | \psi_{k-1} \rangle = E_{k-1,l^*}, \tag{22}$$

so the energy after layer k does not exceed the energy after layer $k-1$. This holds for every $\Delta t > 0$ satisfying the step-size condition $\eta(k,l) \le 1/\mathcal{L}_{GD}$, because Lemma 1 is exact. To quantify the per-layer change, we use Lemma 1 and integrate by parts:

$$\begin{aligned} E_k(\beta_k^*) - E_k(0) &= \Delta t \int_0^{\beta_k^*} A_k(s)\, ds = \Delta t \left[s A_k(s) \right]_{s=0}^{s=\beta_k^*} - \Delta t \int_0^{\beta_k^*} s\, A_k'(s)\, ds \\ &= \Delta t\, \beta_k^* A_k(\beta_k^*) + \Delta t^2 \int_0^{\beta_k^*} s\, B_k(s)\, ds. \end{aligned} \tag{23}$$

Using $|B_k| \le 4\|H_d\|^2 \|H_p\|$ from equation 14, the last integral is at most $2\|H_d\|^2\|H_p\|\,(\beta_k^*)^2$ in absolute value, so

$$E_{k,l^*} - E_{k-1,l^*} \;\le\; \Delta t\, \dot{E}_{k,l^*} + 2\|H_d\|^2 \|H_p\|\, \Delta t^2\, (\beta_k^*)^2, \tag{24}$$

which is equation 9. The first term is non-positive by equation 21 and of order Δt, while the correction is of order Δt^2: the per-layer descent is driven by the Lyapunov rate $\dot{E}_{k,l^*}$ up to an $O(\Delta t^2)$ discretization term, exactly as for continuous-time QLC discretized with step Δt (Magann et al., 2022a). By equation 21, $\Delta t\, \dot{E}_{k,l^*} = -\Delta t\, |\beta_k^*|\, |A_k^{(l^*)}|$, so the right-hand side of equation 24 is itself non-positive as soon as

$$\Delta t\, |\beta_k^*| \;\le\; \frac{|A_k^{(l^*)}|}{2\, \|H_d\|^2\, \|H_p\|}. \tag{25}$$

Condition equation 25 is the discrete-time counterpart of the timestep condition of Magann et al. (2022a) for FALQON and, like it, requires $A_k^{(l^*)} \neq 0$, i.e., that the state has not yet reached a stationary point of the Lyapunov function, where the control has no first-order effect. Together with equation 22, this proves claim (ii). Since $\dot{E}_{k,l} \le 0$ for every iterate and the energy is non-increasing both within and across layers, GD-QLC is a valid QLC scheme, which completes the proof. □

Remark 1. *Identity equation 23 is the discrete analogue of $\frac{d}{dt}E_p = A\beta$ in equation 4: the change of the energy over a layer equals Δt times the Lyapunov rate at the end of the layer, plus a curvature correction of order Δt^2 controlled by B_k. Because the update rule performs gradient descent on the energy itself, only the first-order observable A_k has to be measured; B_k enters the analysis only through the smoothness constant $\mathcal{L}_{GD}$ and the size of the correction term.*

A.3 Proof of Theorem 2

In practice $A_k^{(l)}$ is replaced by a finite-shot estimate $\hat{A}_k^{(l)}$, and equation 8 becomes stochastic gradient descent (SGD) on $f_k = E_k/\Delta t$ with the stochastic gradient

$$g_k^{(l)} := \hat{A}_k^{(l)}. \tag{26}$$

If $\hat{A}_k^{(l)}$ is unbiased given $\beta_k^{(l)}$ (as is the case for an empirical average of projective measurements of the Hermitian observable $i[H_d, H_p]$), then $\mathbb{E}\Big[g_k^{(l)} \mid \beta_k^{(l)}\Big] = A_k^{(l)} = f_k'(\beta_k^{(l)})$. Since $i[H_d, H_p]$ is a bounded operator, the estimator has finite conditional variance for a fixed number of shots, so $\mathrm{Var}\Big(g_k^{(l)} \mid \beta_k^{(l)}\Big) \le \sigma^2$ for some σ^2 that decreases inversely with the number of shots. These are the two oracle assumptions in Theorem 2. The rest of the argument is the standard analysis of SGD for smooth non-convex objectives (Ghadimi & Lan, 2013; Bottou et al., 2018).

Proof of Theorem 2. Fix k, write $f = f_k$ and $\mathcal{L} = \mathcal{L}_{GD}$, and abbreviate $\beta^{(l)} = \beta_k^{(l)}$, $g^{(l)} = g_k^{(l)}$, $\eta_l = \eta(k,l) = c_k/\sqrt{l+1}$. By Step 1 of the proof of Theorem 1, f is $\mathcal{L}$-smooth, so for $\beta^{(l+1)} = \beta^{(l)} - \eta_l g^{(l)}$,

$$f\big(\beta^{(l+1)}\big) \le f\big(\beta^{(l)}\big) - \eta_l\, f'\big(\beta^{(l)}\big)\, g^{(l)} + \frac{\mathcal{L}\eta_l^2}{2}\Big|g^{(l)}\Big|^2.$$

Taking the conditional expectation given $\beta^{(l)}$, using unbiasedness together with $\mathbb{E}\big[|g^{(l)}|^2 \mid \beta^{(l)}\big] \le |f'(\beta^{(l)})|^2 + \sigma^2$, and then taking the full expectation,

$$\mathbb{E}\Big[f\big(\beta^{(l+1)}\big)\Big] \le \mathbb{E}\Big[f\big(\beta^{(l)}\big)\Big] - \eta_l\Big(1 - \frac{\mathcal{L}\eta_l}{2}\Big)\,\mathbb{E}\Big|f'\big(\beta^{(l)}\big)\Big|^2 + \frac{\mathcal{L}\sigma^2}{2}\,\eta_l^2.$$

Because $\eta_l \le \eta_0 = c_k \le 1/\mathcal{L}$, we have $1 - \mathcal{L}\eta_l/2 \ge \frac{1}{2}$. Summing over $l = 0, \dots, L-1$ and telescoping,

$$\frac{1}{2}\sum_{l=0}^{L-1}\eta_l\,\mathbb{E}\Big|f'\big(\beta^{(l)}\big)\Big|^2 \le f\big(\beta^{(0)}\big) - \mathbb{E}\Big[f\big(\beta^{(L)}\big)\Big] + \frac{\mathcal{L}\sigma^2}{2}\sum_{l=0}^{L-1}\eta_l^2 \le \frac{2\|H_p\|}{\Delta t} + \frac{\mathcal{L}\sigma^2}{2}\sum_{l=0}^{L-1}\eta_l^2,$$

where we used $f = E_k/\Delta t$ and $|E_k(\beta)| \le \|H_p\|$ for all β. Now $\sum_{l=0}^{L-1}\eta_l = c_k\sum_{j=1}^{L} j^{-1/2} \ge c_k\sqrt{L}$ (each summand is at least $L^{-1/2}$) and $\sum_{l=0}^{L-1}\eta_l^2 = c_k^2\sum_{j=1}^{L} j^{-1} \le c_k^2(1+\log L)$. If $\hat{l} \in \{0,\dots,L-1\}$ is drawn with probability $\eta_{\hat{l}}/\sum_l \eta_l$, independently of the measurement randomness, then $\mathbb{E}|f'(\beta^{(\hat{l})})|^2 = \sum_l \eta_l\,\mathbb{E}|f'(\beta^{(l)})|^2 \big/ \sum_l \eta_l$, and therefore, recalling that $f' = A_k$,

$$\mathbb{E}\Big[\,\big|A_k^{(\hat{l})}\big|^2\,\Big] \le \frac{4\|H_p\|/\Delta t + \mathcal{L}_{GD}\,\sigma^2 c_k^2\,(1+\log L)}{c_k\sqrt{L}} = O\Big(\frac{\log L}{\sqrt{L}}\Big),$$

which is the claim. □

Remark 2. *As is standard for non-convex SGD, the guarantee is stated for a randomly selected iterate $\hat{l}$. Algorithm 1 instead keeps the iterate l^* with the smallest estimated energy; this is the choice that the Lyapunov guarantee of Theorem 1 refers to, while Theorem 2 quantifies how close the inner loop gets to a stationary point of the per-layer energy after L steps. Since $\mathcal{L}_{GD}$ does not depend on k, a single constant $c_k = c \le 1/\mathcal{L}_{GD}$ is admissible for all layers. For the largest admissible choice $c_k = 1/\mathcal{L}_{GD}$ the bound becomes $\big(16\|H_d\|^2\|H_p\|^2 + \sigma^2(1+\log L)\big)/\sqrt{L}$, which is independent of Δt; the trade-off between the number of gradient steps L and the number of measurement shots (which controls σ^2) is captured by the term $\sigma^2(1+\log L)$.*

B Experimental details

In this section we will expand on the details of the quantum simulations performed throughout this work. All experiments were performed on a compute node with 512 GB of memory, a single 64-core, 2.0 GHz, 225-watt AMD EPYC 7713 processor, and four NVIDIA A100 GPUs. The data for the figures in the paper required approximately 50 hours of compute time to generate, not including

preliminary studies. The corresponding problem Hamiltonians H_p for each task are summarized in Table 1. For all problem instances, the initial state $|\psi_0\rangle$ is chosen to be the uniform superposition, and the driver Hamiltonian is fixed to $H_d = \sum_i X_i$, where X_i is the Pauli-X on the ith register.

We consider 3-regular graphs for MAX-CUT, scale-free graphs for MIN-COVER, and random graphs for MAX-CLIQUE. Scale-free graphs are generated using the Barabási–Albert model with parameter $m = 3$, while random graphs are drawn from the Erdős–Rényi model with edge probability $p = 0.5$. For MAX-CUT, we include all non-isomorphic 3-regular graphs (see, e.g., (Meringer, 1999)). For each qubit count and each method, we evaluate 40 randomly generated problem instances to compute averaged performance metrics unless otherwise noted. In the case of 10-vertex 3-regular graphs, only 19 distinct graphs exist and therefore averages were taken over all graphs. For weighted MAX-CUT, edge weights were randomly chosen uniformly from the interval $[0,\ 2]$, leading to an average expected weight of one. As for the choice of learning rate, we consider $\eta(k, l) = \frac{c}{\sqrt{l+1}\log k}$, where c is a tunable constant. Our intuition is that as k gets larger, we expect $|\psi_k\rangle$ to approach the desired solution implying that β approaches zero. We specifically set $\eta(k, l) = \frac{0.1}{\sqrt{l+1}\log k}$ as the learning rate, and $L = 5$ for the number of iterations in GD-QLC. We choose $K = 1000$ as the limit for the total number of layers for all the algorithms except for FOCQS, which computational constraints limited our simulation to $K = 100$ layers.

Table 1: Problem Hamiltonian for the studied graph problems

PROBLEM	H_p
MAX-CUT	$\frac{1}{2}\sum_{(i,\,j)\in E(G)} (Z_iZ_j - I)$
WEIGHTED MAX-CUT	$\frac{1}{2}\sum_{(i,\,j)\in E(G)} (w_{ij}Z_iZ_j - I)$
MAX-CLIQUE	$3\sum_{(i,\,j)\in E(\bar{G})} Z_iZ_j - Z_i - Z_j + \sum_{i\in V(G)} Z_i$
MIN-COVER	$3\sum_{(i,\,j)\in E(G)} Z_iZ_j + Z_i + Z_j - \sum_{i\in V(G)} Z_i$

C EXTENSION TO MULTIPLE CONTROL HAMILTONIANS

GD-QLC can be easily extended to multiple controls, as the derivations do not require a fixed number of controls *a priori*. A multi-control version of equation 2 can be written as

$$i\frac{d}{dt}|\psi(t)\rangle = \left(H_p + \langle\phi(\vec{\beta}(t)),\, \vec{H}_d\rangle\right)|\psi(t)\rangle\,,\ |\psi(0)\rangle = |\psi_0\rangle\,, \tag{27}$$

and we have

$$\frac{d}{dt}E_p = \langle\phi(\vec{\beta}(t)),\, \vec{A}(t)\rangle \tag{28}$$

where $\langle\cdot,\,\cdot\rangle$ denotes the standard vector inner product, $\vec{\beta}(t) = (\beta_0(t),\,\ldots,\beta_J(t))$ is the vector of control inputs, $\vec{H}_d = (H_{d,0},\,\ldots,H_{d,J})$ are the corresponding control Hamiltonians, and $\vec{A}(t) = (A_0(t),\,\ldots,A_J(t))$ are the first-order commutation observables defined as $A_j(t) = \langle\psi(t)|i[H_{d,j},\,H_p|\psi(t)\rangle$. This system can be discretized in the same manner as equation 5 to yield the following circuit

$$|\psi_K\rangle = \prod_{k=0}^{K}\left(\prod_{j=0}^{J} U_{d,j}(\beta_j)U_p\right)|\psi_0\rangle\,,\ U_p = e^{-i\Delta t H_p},\ U_{d,j}(\beta_j) = e^{-i\Delta t\beta_j H_{d,j}}. \tag{29}$$

Following the same analysis as in Appendix A, the gradient at the k'th layer with respect to the control input is

$$\nabla E_{p,k} = \vec{A}(t). \tag{30}$$

and the update rule

$$\vec{\beta}_k^{(l+1)} = \beta_k^{(l)} - \eta(k,\,l)\nabla E_{p,k},\quad \vec{\beta}_k^{(0)} = \vec{0} \tag{31}$$

is the GD-QLC update rule for the multi-parameter case.

D QAOA

QAOA is a hybrid quantum–classical algorithm for solving combinatorial optimization problems (Farhi et al., 2014). This framework works by alternating applications of a problem-specific cost Hamiltonian H_p and a driver (mixing) Hamiltonian H_d, with the goal of concentrating probability on high-quality solutions. The Hamiltonians evolutions characterize the following unitary operations:

$$U_p(\gamma) = e^{-i\gamma H_p},\, U_d(\beta) = e^{-i\beta H_d}, \tag{32}$$

with tunable parameters γ and β. Typically H_p is designed so that its ground state encodes the solution to the optimization problem. Starting with an initial state $|s\rangle$, p alternating applications of U_p and U_d with parameters $\boldsymbol{\gamma} = (\gamma_0,\, \gamma_1,\, \ldots,\, \gamma_p)$ and $\boldsymbol{\beta} = (\beta_0,\, \beta_1,\, \ldots,\, \beta_p)$ results in the state

$$|\psi(\boldsymbol{\gamma},\, \boldsymbol{\beta})\rangle = U_d(\beta_p)U_p(\gamma_p)\cdots U_d(\beta_1)U_p(\gamma_1)\,|s\rangle\,, \tag{33}$$

A classical computer is utilized to iteratively tune $\boldsymbol{\gamma}$, $\boldsymbol{\beta}$ so that the following objective function is minimized:

$$E(\boldsymbol{\gamma},\, \boldsymbol{\beta}) = \langle\psi(\boldsymbol{\gamma},\, \boldsymbol{\beta})|\, H_p\, |\psi(\boldsymbol{\gamma},\, \boldsymbol{\beta})\rangle\,. \tag{34}$$

The performance of QAOA is typically measured by the approximation ratio $E_{\mathrm{QAOA}}/E_{min}$, which compares the expected cost value obtained from the QAOA output state to that of the true optimal solution.

E Overview of FALQON

Magann et al. (2022a) introduced a time-discretized realization of QLC for solving quantum optimization tasks known as the *Feedback-based Algorithm for Quantum Optimization* (FALQON). In FALQON, the continuous-time dynamics in equation 2 are approximated by a sequence of discrete control steps of duration Δt, implemented via alternating unitary evolutions under the control and problem Hamiltonians. Specifically, defining

$$U_{d,j} = e^{-i\Delta t\beta_j H_d}, \qquad U_p = e^{-i\Delta t H_p}, \tag{35}$$

the state of the system after k steps, starting from an initial state $|\psi_0\rangle$, is given by

$$|\psi_k\rangle = \prod_{j=0}^{k} U_{d,j}\, U_p\, |\psi_0\rangle\,. \tag{36}$$

The resulting circuit is parameterized by the vector $\boldsymbol{\beta} = (\beta_0, \ldots, \beta_k)$, which provides a piecewise-constant approximation to the continuous control field $\beta(t)$. By iteratively appending layers and updating β_{j+1} using measurement-based estimates of $A_j := A(j\Delta t) = \langle\psi_k|\, i[H_d,\, H_p]\, |\psi_k\rangle$, FALQON constructs increasingly accurate approximations to the continuous-time Lyapunov control trajectory. Provided that Δt is chosen sufficiently small, this discretization preserves the monotonic improvement guarantees of the underlying QLC framework.