

Robust quantum control for adiabatic quantum computation

Gregory Quiroz

The Johns Hopkins University Applied Physics Laboratory, Laurel, Maryland 20723, USA



(Received 11 April 2019; published 7 June 2019)

Properly designed control has been shown to be particularly advantageous for improving adiabatic quantum computation (AQC) accuracy and time complexity scaling. Here, an *in situ* quantum control optimization protocol is developed to indirectly optimize state fidelity without knowledge of the instantaneous spectral gap or the computational solution. The protocol is shown to converge to analytically derived time-optimal controls for Grover's search algorithm. Furthermore, the protocol is utilized to explore optimized control trajectories for the maximum two-bit satisfiability problem, where appreciable improvement in fidelity and the minimum spectral gap over a linear schedule is observed. The approach is also shown to be robust against system model uncertainties (unitary control errors). This method is designed to enable robust control optimization on existing quantum annealing hardware and future AQC processors.

DOI: [10.1103/PhysRevA.99.062306](https://doi.org/10.1103/PhysRevA.99.062306)

I. INTRODUCTION

Adiabatic quantum computation (AQC) utilizes controlled adiabatic evolution of a many-body quantum system to implement a quantum algorithm. The quantum system is described by a Hamiltonian $H_{\text{ad}}[\mathbf{x}(t)] = \sum_{l=1}^L x_l(t)H_l$, where the ground state of the initial Hamiltonian $H_0 = H_{\text{ad}}[\mathbf{x}(0)]$ is assumed to be easily prepared and the ground state of the problem Hamiltonian $H_P = H_{\text{ad}}[\mathbf{x}(T)]$ represents the solution to the computational problem [1–3]. The system evolution is dictated by the control schedules $\mathbf{x}(t) = \{x_l(t)\}_{l=1}^L$, which effectively controls the amplitude of each noncommuting, linearly independent, primitive Hamiltonian H_l .

The accuracy of AQC is determined by the adiabatic theorem, which asserts that the system will remain in an instantaneous eigenstate of $H_{\text{ad}}(t)$ provided the dynamics are sufficiently slow. In the noise-free case, the adiabatic theorem yields the rigorous bound on the trace-norm distance [4] between $|\Phi_0(t)\rangle$, the instantaneous ground state of $H_{\text{ad}}(t)$, and the time-evolved state $|\psi(t)\rangle$: $D[|\Phi_0(T)\rangle, |\psi(T)\rangle] \lesssim q^a$, provided

$$T \gtrsim \frac{a \max_{s \in (0,1)} \left\| \frac{d}{ds} H_{\text{ad}} \right\|^{b-1}}{q \Delta_{\min}^b}, \quad (1)$$

where $s = t/T$ is the normalized time, $\|A\|$ denotes the operator norm, and Δ_{\min} is the *minimum spectral gap* between the instantaneous ground state and first excited state of $H_{\text{ad}}(t)$. The parameter $q \in (0, 1)$, while the integer exponents a and b depend upon the differentiability and analyticity properties of $H_{\text{ad}}(t)$ and the boundary conditions satisfied by its derivatives [5–7].

The adiabatic theorem has been the basis for a number of studies focused on properly designing $\mathbf{x}(t)$ to minimize the adiabatic error D and reduce the lower bound on T by modifying Δ_{\min} via a local adiabatic condition that seeks to minimize ground-state transitions $\forall t$ [8]. These approaches have employed variational time-optimal strategies [9], optimal control theory [10], and convex optimization [11] that exploit

accurate system models and either knowledge of the computational solution, i.e., $|\Phi_0(T)\rangle$, or the instantaneous spectral gap $\Delta(s)$. In this paper, an optimization technique referred to as closed-loop optimized adiabatic quantum control (CLOAQC) is developed to indirectly optimize the adiabatic error using only the time-evolved system state at $t = T$ measured in the computational basis. The method is shown to converge towards known time-optimal solutions for Grover's search algorithm (GSA) [8,9] and substantially improve adiabatic error and enhance Δ_{\min} for the maximum two-bit satisfiability (MAX 2-SAT) problem relative to a linear schedule. The protocol is shown to exhibit robustness to unitary control errors and it is argued that due to the form of the objective function the method can be readily extended to more generic noise models.

II. CLOAQC

A. The algorithm

Closed-loop quantum control learning is an iterative optimization method that relies on information from previous experiments to update control parameters and effectively optimize system performance with respect to a given objective function. In the case of AQC control optimization, the learning procedure includes three main steps: (1) the generation of a set of control parameters, (2) the implementation of the AQC algorithm and subsequent quantification of performance, and (3) a learning algorithm that incorporates prior performance information to provide updated control parameters. Observing the impact of varying the control parameters via prior experiments is a key aspect of closed-loop learning that affords inherent robustness to system uncertainty. It is exploited here to consider a quantum processor designed to implement quantum annealing or, more generally, AQC in a blackbox framework where one has limited knowledge of the intrinsic noise processes, systematic errors, and the underlying structure of the energy spectrum of $H_{\text{ad}}(t)$. To this end, it is assumed that one only has knowledge of the Hamiltonian

one *believes* is being implemented on the hardware and the state of the system at the *end* of the computation measured in the computational basis. These assumptions fit well within the confines of currently available quantum annealing based hardware, such as the D-wave processor [12], and future AQC processors.

A function space is used to parametrize the control functions for each constituent Hamiltonian H_i in $H_{\text{ad}}[\mathbf{x}(s)]$. Each control function is defined as

$$x_i(s) = \sum_{j=1}^{d+1} \lambda_{ij} \phi_j(s), \quad (2)$$

where λ_{ij} denote weights for each of the j th basis functions $\phi_j(s)$. Thus, the optimization parameter space is defined by $\Lambda \equiv \{\lambda_{ij}\}_{i,j=1}^{L,d+1}$, which includes the weights for the d basis elements in the function expansion for all L controls. Note that the function representation has advantages over piecewise control in that the parameter space is drastically reduced and variations in control parameters result in global rather than local changes in control functions. Here, the polynomial basis $\phi_k(s) = s^{k-1}$ is chosen for simplicity; however, one can readily consider any alternative. Intrinsically band-limited functions such as the discrete prolate spheroidal sequences may be attractive for imposing intrinsic bandwidth constraints on control profiles [13].

Traditional state fidelity metrics, such as the trace-norm distance,

$$D[|\psi(T, \Lambda)\rangle, |\Phi_0(T)\rangle] = \sqrt{1 - |\langle \psi(T, \Lambda) | \Phi_0(T) \rangle|^2}, \quad (3)$$

require knowledge of $|\psi(T, \Lambda)\rangle$ and $|\Phi_0(T)\rangle$. While one may estimate $|\psi(T, \Lambda)\rangle$ by sampling the AQC algorithm, knowledge of $|\Phi_0(T)\rangle$ implies knowledge of the computational solution. One may envision control protocols that exploit partial or approximate solutions obtained from classical algorithms; however, the focus of this paper is the case where the computational solution is unknown and additional classical preprocessing (i.e., approximate solution optimization algorithm) can be avoided.

State fidelity metrics are circumvented by employing the energy of the system with respect to H_p , $E(\Lambda) = \langle \psi(T, \Lambda) | H_p | \psi(T, \Lambda) \rangle$ to evaluate the performance of a given set of control parameters Λ . The expectation value is estimated by accumulating samples of the final state measured in the computational basis $\{|\tilde{\psi}_i(T, \Lambda)\rangle\}_{i=1}^M$ from M implementations of the AQC algorithm; this estimate is denoted as $\hat{E}(\Lambda)$. Note that $E(\Lambda)$ serves as a viable surrogate objective function that achieves a minimum value when $|\psi(T, \Lambda)\rangle = |\Phi_0(T)\rangle$. Alternative objective functions based on the average energy have been discussed as a means for optimizing control schedules for quantum annealing. In Ref. [14], the energy difference at each point in time is employed as a figure of merit, while in Ref. [15] the average energy appears via the fluctuation-dissipation theorem. In both studies, the control schedules rely on information about the system at times $t \in [0, T]$. CLOAQC relaxes this constraint, requiring only an estimation of the energy of the system at $t = T$.

Due to sampling statistics, function calls to the $\hat{E}(\Lambda)$ are stochastic. Therefore, a stochastic optimization technique must be employed to perform the control optimization.

Simultaneous perturbative stochastic approximation (SPSA), an iterative gradient-based optimization technique that requires two function calls per iteration to estimate the gradient, is selected for this task [16]. Note that SPSA has been previously used for quantum information applications [17–19]. For each iteration k , the first step is to generate a random search vector Δ_k , where each of the $i = 1, 2, \dots, |\Lambda|$ elements is a Bernoulli distributed variable, i.e. $\Delta_{kj} = \pm 1$. The estimated gradient is

$$g_k = \frac{\hat{E}(\Lambda_k + \beta_k \Delta_k) - \hat{E}(\Lambda_k - \beta_k \Delta_k)}{2\beta_k} \Delta_k, \quad (4)$$

where the control parameter update is given by

$$\Lambda_{k+1} = \Lambda_k + \alpha_k [g_k + \lambda_{\text{ad}} \nabla J_{\text{ad}}(\Lambda_k)] \quad (5)$$

and $\lambda_{\text{ad}} = 0.005$ in all subsequent simulations. The update includes an additional analytically calculated gradient for the objective function

$$J_{\text{ad}} = \sum_{\mu \in \mathbf{x}} \int_0^1 \|\dot{\mu}(s)\| ds, \quad (6)$$

which seeks to enforce adiabaticity by minimizing the derivative of each control field over the total time interval [10].

The functions α_k and β_k are convergence parameters and typically defined as

$$\alpha_k = \frac{\alpha_0}{(k+1+R)^\delta}, \quad \beta_k = \frac{\beta_0}{(k+1)^\zeta}, \quad (7)$$

where α_0 , β_0 , and R are chosen following the procedure outlined in Ref. [20]. The remaining parameters $\delta = 0.602$ and $\zeta = 0.101$ are typically good values [16] and appear to be good choices for the problem presented in this paper. It is also found that the asymptotically optimal values $\delta = 1$ and $\zeta = 1/6$ yield similar results [21].

B. Convergence and performance

SPSA convergence analyses [16,21] establish convergence in part by imposing conditions on the convergence parameters α_k and β_k , the distribution of Δ_k , and the statistical relationship between Δ_k and measurements of $E(\Lambda)$. Essentially, the main conditions stipulate that (1) α_k and β_k both tend to zero at rates that are neither too fast nor too slow, (2) the objective function $E(\Lambda)$ is sufficiently smooth near Λ_{opt} , and (3) Δ_k defines a vector of elements of independent and symmetrically distributed values about zero with finite inverse moments $\langle |\Delta_{ki}^{-1}| \rangle$. By selecting α_k and β_k according to Eq. (7) and Ref. [20], condition 1 is satisfied. Condition 2 is inherently satisfied, while condition 3 is satisfied by selecting Δ_k to follow the symmetric Bernoulli ± 1 distribution.

Provided that the above conditions are satisfied, the convergence in the objective function can be related to convergence in the control space. According to SPSA convergence analyses, the error in the control space reduces asymptotically as $\|\Lambda_k - \Lambda_{\text{opt}}\| \sim O(k^{-\gamma/2})$, where in typical situations

$\gamma \in [1, 0]$. Utilizing this result, one obtains the bound

$$\begin{aligned} |f(\Lambda_k) - f(\Lambda_{\text{opt}})| &\leq \frac{C}{2} \|H[f(\Lambda_{\text{opt}})]\| \|\Lambda_k - \Lambda_{\text{opt}}\|^2 \quad (8) \\ &\leq \frac{C}{2} \|H[f(\Lambda_{\text{opt}})]\| k^{-\gamma}, \quad (9) \end{aligned}$$

where C is a constant and $\|H(f(\Lambda_{\text{opt}}))\|$ is the spectral norm of the Hessian. This bound holds for any differentiable function f satisfying $\nabla f(\Lambda_{\text{opt}}) = 0$. Since both the surrogate $E(\Lambda)$ and true objective $D(\Lambda)$ satisfy these properties, the scaling is expected to hold for both functions. In the subsequent section, this claim is corroborated by numerical simulations that convey the expected scaling $|D(\Lambda_k) - D(\Lambda_{\text{opt}})| \sim O(1/k^\gamma)$. The value of γ is extracted for different control scenarios to determine convergence rates. It is important to note that while asymptotic results suggest $\gamma \approx 1$ the actual convergence rate is highly problem dependent.

CLOAQC performance is dictated by the typical overhead requirements of SPSA. The algorithm requires a total number of $2MK$ experiments, where K is the number of iterations of the algorithm. The factor of 2 arises from SPSA's finite difference gradient [Eq. (4)] estimation, which differs from the $2d$ estimates required for standard finite difference gradient techniques.

III. NUMERICAL EXPERIMENTS

A. Grover's search algorithm

The efficacy of CLOAQC is explored via GSA for the identification of a marked element in an unsorted database of N elements [22]. GSA requires a minimum of $O(\sqrt{N})$ oracle queries to identify the marked element, a quadratic improvement over the best possible classical algorithm [23]. Recast in the language of AQC [8,9], Grover's algorithm is generically defined by an n -qubit Hamiltonian

$$H_G(s) = x_1(s)[I - |+\rangle\langle +|] + x_2(s)[I - |m\rangle\langle m|], \quad (10)$$

where $x_{1,2}(s)$ are the control functions, I is the identity operator, $|+\rangle$ represents the uniform superposition over all $N = 2^n$ computational basis states, and $|m\rangle$ is the marked state. Time-optimal controls can be designed such that $\Delta_{\text{min}} \sim O(1/\sqrt{N})$, and the total runtime required to reach the ground state $|\Phi_0(T)\rangle = |m\rangle$ is $T \sim O(\sqrt{N})$, thus achieving the well-known quadratic speedup [8,9].

The optimized control functions obtained from CLOAQC are compared to the time-optimal GSA controls for one independent control (IC) [$x_1(s) = 1 - x_2(s)$] and two ICs, where $x_1(s)$ and $x_2(s)$ are linearly independent. The CLOAQC algorithm is initialized such that Λ_0 describes a linear ramping control schedule. The boundary conditions $x_1(0) = x_2(1) = 1$ and $x_1(1) = x_2(0) = 0$ are enforced on the control profiles throughout the optimization procedure. Each control function is expanded into five basis functions, and thus the control parameter space is described by five and ten parameters in the one and two IC case, respectively. The total runtime $T = 5/\Delta_{\text{min}}$ is chosen for all simulations (see the Appendix B for a further analysis of CLOAQC performance as a function of runtime).




FIG. 1. Relative difference between trace-norm distances for time-optimal GSA controls and CLOAQC for one and two control degrees of freedom. Comparison illustrates the dependence of CLOAQC performance on the number of algorithm implementations M and number of qubits n .

A comparison of performance indicates a convergence in CLOAQC solutions toward the time-optimal GSA solutions with increasing iteration. In Fig. 1, CLOAQC is compared to the Roland-Cerf (RC) [8] and the quantum adiabatic Brachistochrone (QAB) [9] GSA solutions for various values of sampling parameter M and number of qubits n using 100 realizations of CLOAQC. The top and bottom rows illustrate the relative difference in adiabatic error D between CLOAQC and RC and QAB, respectively, as a function of iteration k . Median CLOAQC performance is denoted by the solid colors, while the shaded region denotes the interquartile range. Although both RC and QAB solutions require knowledge of the spectral gap $\Delta(s)$, CLOAQC is capable of converging toward the equivalent time-optimal solutions using only the state of the system at the end of the computation. In panels (a) and (c), the convergence rate of CLOAQC is shown to be approximately independent of the sampling parameter M . The most considerable improvement in performance is observed between $M = 10$ and 100 for both control scenarios. CLOAQC convergence does not convey a compelling dependence on n [see Figs. 1(b) and 1(d)]. Therefore, CLOAQC performance is dependent on the dimension of Λ and not on the underlying hardness of the AQC algorithm. Note that each panel includes a fit for $n = 4$ with $M = 100$, along with the corresponding convergence parameter γ . Note that γ appears to be consistent with the typical SPSA scaling for both control scenarios.

CLOAQC optimized control profiles are qualitatively similar to the time-optimal GSA solutions. In particular, for one IC, optimized controls closely resemble the RC solution, while deviations from the QAB control profiles are more significant for two ICs. Higher-order polynomial expansions are needed to more accurately reproduce the QAB profiles. Note that despite such a distinction between CLOAQC and QAB paths their adiabatic errors only differ by less than 10^{-3} for the n values considered here (see the Appendix A for a more detailed discussion).

B. MAX 2-SAT

The utility of CLOAQC is further illustrated via the MAX 2-SAT problem, where optimized controls are shown to offer substantial reductions in adiabatic error and amplification of Δ_{\min} relative to a linear control schedule. MAX 2-SAT offers variability in the energy spectrum, including location and magnitude of Δ_{\min} , that is highly problem instance dependent. By focusing on unique satisfying assignment (USA) instances (i.e., instances with nondegenerate ground-state manifolds), CLOAQC's ability to maintain adiabaticity and effectively navigate the ground-state manifold without knowledge of the specifications of Δ_{\min} is demonstrated. CLOAQC's ability to enhance adiabatic error is supplemented by gap amplifications that generally become more pronounced with increasing ICs.

CLOAQC performance is assessed with respect to an ensemble of 100 USA 2-SAT instances. Each 2-sat instance is a logical AND of M_c clauses, where each clause C_j itself is a logical OR of exactly two Boolean variables from the set $\{x_i\}_{i=1}^n$. Each 2-SAT problem Hamiltonian is constructed by associating the binary values of each Boolean variable x_j with the ± 1 eigenstates of the Pauli spin operator σ_j^z for the j th qubit and summing the M_c clause Hamiltonians [1]. After rescaling and dropping the constant term, the resulting problem Hamiltonian is

$$H_P = \sum_j h_j \sigma_j^z + \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z, \quad (11)$$

where $h_j = -\sum_m v_j^m$ and $J_{ij} = \sum_m v_i^m v_j^m$ [24]. The variables $v_j^m \in \{-1, 0, 1\}$, where $j = \{1, 2, \dots, n\}$ and $m = \{1, 2, \dots, M_c\}$ label the variables and clauses, respectively, and encode the specifications of C_m . Namely, if x_j appears negated (un-negated) in the m th clause then $v_j^m = -1(+1)$. $v_j^m = 0$ for all clauses where x_j does not appear. The USA ensemble is generated using the approach in Ref. [2] and found to possess an average clause density $\hat{\alpha} = n/M_c \approx 1.2$. Clause density is known to play a role in discerning problem hardness. $\hat{\alpha}$ is found to be sufficiently close to the critical clause density $\alpha_c = 1$, where the most (classically) difficult MAX 2-SAT problems lay [25].

A variety of control scenarios are considered by defining the MAX-2 SAT algorithm as

$$H_{2S}(s) = x_1(s)H_0 + x_2(s)H_I + x_3(s)H_{P,1} + x_4(s)H_{P,2}. \quad (12)$$

The initial Hamiltonian $H_0 = \sum_j \sigma_j^x$ represents a transverse field on each qubit, thus defining the initial ground state to be the uniform superposition state over all computational basis states. The intermediate Hamiltonian $H_I = \sum_{i \neq j} \sigma_i^x \sigma_j^x$ is only present for $t \in (0, T)$ and defines a nonstochastic contribution. Nonstochastic Hamiltonians have been studied for the MAX 2-SAT problem [26], and they have been shown to benefit algorithmic performance for certain algorithms [27–29]. Here, their advantages in the presence of optimized control are investigated. The last two terms in Eq. (12), $H_{P,1} = \sum_j h_j \sigma_j^z$ and $H_{P,2} = \sum_{i,j} J_{ij} \sigma_i^z \sigma_j^z$, denote the 1- and 2-local terms of the problem Hamiltonian $H_P = H_{P,1} + H_{P,2}$. The number of independent controls is varied by imposing constraints on $x_i(s)$. Four scenarios are considered here:




FIG. 2. Comparison of CLOAQC vs a linear schedule for 100 instances of the MAX 2-SAT problem for $n = 8$ qubits and four different control scenarios.

(1) one IC, where $x_1(s)$ and $x_4(s) = 1 - x_1(s)$ are nonzero; (2) two ICs, where $x_2(s) = 0$ and $x_3(s) = x_4(s)$; (3) three ICs, where $x_3(s) = x_4(s)$; and (4) four ICs, where all control functions vary independently. In all cases, $x_1(0) = x_3(1) = x_4(1) = 1$, $x_1(1) = x_3(0) = x_4(0) = 0$, and $x_2(0) = x_2(1) = 0$ are imposed on the control functions.

CLOAQC is shown to outperform a linear control schedule for the ensemble of 100 USA instances in Fig. 2. Median adiabatic error \tilde{D} is obtained from an ensemble of 25 realizations of CLOAQC for each problem instance, with Δ_{CLOAQC} corresponding to the minimum gap for the control profile that produces \tilde{D} . CLOAQC is implemented for 1000 iterations with $M = 100$ and $T = 10/\Delta_{\min}$.

Performance features between different control scenarios are remarkably distinct. Letting r_D and r_Δ denote the median ratios $\tilde{D}_{\text{CLOAQC}}/D_{\text{lin}}$ and $\Delta_{\text{CLOAQC}}/\Delta_{\text{lin}}$ with respect to the distribution of problem instances, control scenarios are compared against each other by their median improvement in adiabatic error and gap enhancement. In the one control case, $r_D \approx 0.066$ and no gap amplification is observed. A minimum of two ICs are required to achieve gap amplifications, where $r_\Delta \approx 2$ and the median adiabatic error ratio improves to $r_D \approx 0.011$. Further improvements in adiabatic error and minimum gap size are achievable by including H_I with optimized control; however, the degree of improvement is strongly dependent upon the choice of H_I and problem instance; this is consistent with previous findings [28]. The distribution for the three IC case is fairly localized in adiabatic error and broad in gap enhancement, where the median performance ratio $r_D \approx 0.011$ and $r_\Delta \approx 2.97$. Note that in some cases the gap enhancement is as large as approximately $4.4 \times \Delta_{\text{lin}}$. While a subset of the instances do benefit from four ICs, a majority of the instances do not. The median ratio of adiabatic error is $r_D \approx 0.062$ and the median gap enhancement is $r_\Delta \approx 2.538$ for the distribution of instances. The degradation in performance is due to the fixed number of CLOAQC iterations and increasing dimension of the search space. Increasing the number of iterations to $K = 2000$, CLOAQC performance improves to $r_D \approx 0.020$ and a modest median gap enhancement of $r_\Delta \approx 2.89$. Further improvements in D likely require an increase in M and K , specifically in the high fidelity regime where achieving nonzero $\hat{E}(\Delta)$ values may require $M \gg 1$.




FIG. 3. Comparison of CLOAQC vs a linear schedule for 100 instances of the MAX 2-SAT problem for $n = 8$ and four different control scenarios. Panels (a) and (b) convey results for $H_I = H^{(y)}$ and $H_I = H^{(xz)}$, respectively.

C. More on MAX 2-SAT: Alternative intermediate Hamiltonians

The benefits of nonstochastic Hamiltonians combined with optimized controls provided by CLOAQC are further explored here for the MAX 2-SAT problem using alternative Hamiltonians to the two-local XX interaction Hamiltonian presented in the main text. In Figs. 3(a) and 3(b), the intermediate Hamiltonian is given by $H_I^{(y)} = \sum_j \sigma_j^y$ and $H_I^{(xz)} = \sum_{ij} \sigma_i^x \sigma_j^z + \sigma_i^z \sigma_j^x$, respectively. The former has been previously studied as an additional driving term for quantum annealing, where techniques from “shortcuts to adiabaticity” were employed to produce optimized control for particular Ising-type problem Hamiltonians [30]. A Hamiltonian similar to $H_I^{(xz)}$ has been previously employed as a catalyst Hamiltonian that improves the success probability and runtime scaling for a specific choice of H_P [27]. This paper focuses on potential improvements provided by each H_I in conjunction with CLOAQC optimized control for the 100 USA problem instances discussed in the main text.

The addition of $H_I^{(y)}$ leads to improvements in adiabatic error and no substantial improvements in the minimum spectral gap. In Fig. 3(a), the median performance of CLOAQC is compared to the one control linear ramping schedule. The one and two control schedules are included for reference, with the results for three and four controls constituting the H_I -dependent distributions. Significantly more localized than the XX interaction Hamiltonian, $H_I^{(y)}$ with optimized control generally offers improvements in adiabatic error without the need for gap amplification for the 100 USA problems considered here. Degradations in adiabatic error for the four control case are still observed and again attributed to fixing the number of iterations and energy samples used for the CLOAQC optimization experiments.

In contrast $H^{(y)}$, improvements in adiabatic error and minimum gap size are strongly instance dependent when utilizing $H_I^{(xz)}$. Improvements in adiabatic error up to a factor of approximately 10^{-3} and 10^{-4} are observed for three and four independent controls, respectively, with gap amplifications up to approximately $4.5 \times \Delta_{lin}$. While substantial improvements are observed for a subset of instances, approximately 1/3 of

TABLE I. Summary of median improvements in adiabatic error and gap amplifications for the distribution of 100 MAX 2-SAT USA problem instances for each nonstochastic Hamiltonian considered. Data are formatted as (r_D, r_Δ) . SAA denotes “same as above” for equivalent results among different intermediate Hamiltonians.

H_I	No. of controls			
	1	2	3	4
$H_I^{(xx)}$	(0.065, 1)	(0.014, 1.735)	(0.012, 2.973)	(0.062, 2.538)
$H_I^{(y)}$	SAA	SAA	(0.009, 1.673)	(0.011, 1.658)
$H_I^{(xz)}$	SAA	SAA	(0.023, 2.604)	(0.031, 2.722)

the instances do not benefit from the addition of $H_I^{(xz)}$ with optimized control. Tracking the adiabatic error as a function of iteration, the addition of $H_I^{(xz)}$ appears to induce stability issues in CLOAQC that lead to limited success in optimization over the 25 realizations. By altering the convergence parameters, it is possible to improve median performance and overcome such issues at the cost of an increase in the number of iterations.

Median improvements in adiabatic errors and gap amplification taken over the distribution of problem instances are summarized in Table I for all choices of H_I .

D. Robustness to noise

CLOAQC possesses a degree of inherent robustness to uncertainty in $H_{ad}(s)$ due to the fact that it is a closed-loop protocol and it relies on the minimization of the average energy with respect to H_P , the problem Hamiltonian one wishes to encode on the AQC hardware. CLOAQC’s robustness is assessed here by including an additive, unitary control error with three different types of ramping schedules. The Hamiltonian

$$H'_{ad}(s) = H_{ad}(s) + H_E(s) \quad (13)$$

is used to describe the faulty AQC algorithm, with the additional additive term $H_E(s) = \Gamma(s) \sum_i^n \hat{m}_i \cdot \vec{\sigma}_i$ contributing to the deformation of the ground-state manifold. $\hat{m}_i = (m_{i,x}, m_{i,y}, m_{i,z})$ is a unit vector where $m_{i,\mu}$ is generated from a zero mean normal distribution with unit standard deviation. The ramping schedule $\Gamma(s)$ is chosen in accordance with Ref. [31] and takes three forms: (a) $\Gamma(s) = Cs$, (b) $\Gamma(s) = C \sin(\pi s)$, and (c) $\Gamma(s) = 1/2 \sin(C\pi s)$, where $C \in \mathbb{R}$. Figure 4 compares CLOAQC to a linear, RC, and QAB control profile for GSA, with each panel corresponding to the three $\Gamma(s)$ schedules, respectively. Median performance is shown for a distribution of 25 realizations of CLOAQC, using a realization of $H_E(t)$ that does *not* exhibit favorable recurrences in D with increasing C [31]. CLOAQC conveys considerable improvements in adiabatic error and robustness for sufficiently small and slow-oscillating unitary control errors, most notably for ramping schedules (b) and (c), where one control CLOAQC outperforms (two control) QAB. CLOAQC performance exhibits an abrupt degradation in performance at critical C values where the dynamics are dominated by




FIG. 4. Comparison of CLOAQC against linear, RC, and QAB control schedules for the three unitary control error models discussed in the main text. Median CLOAQC performance is shown for a distribution of 25 realizations of CLOAQC using $n = 4$, $K = 1000$, and $M = 100$.

$H_E(s)$ and insufficient sampling of the H_p ground state exists.

While the focus here is unitary control errors, given the form of the objective function, CLOAQC is expected to be robust against more generic noise sources. For example, the algorithm also exhibits robustness to time-dependent semiclassical stochastic noise processes as well (results not shown here). Generically, CLOAQC can be expected to outperform a linear schedule when $\|H_{ad}(s)\| \gg \|H_E(s)\|$, where $\|H_E(s)\| < \infty$ and $\|\cdot\|$ is the operator norm, and the dynamics generated by $H_E(s)$ are sufficiently slow $\forall s$. The development of rigorous bounds and potential relaxations of this local condition on $H_E(s)$ related to CLOAQC performance in the open quantum system setting is left for future work.

IV. CONCLUSIONS

A blackbox AQC control optimization protocol (referred to as CLOAQC) is presented and shown to drastically improve algorithmic performance solely using the state of the quantum system at the end of the computation and no knowledge of the minimum spectral gap. CLOAQC is shown to converge towards Grover's search algorithm's time-optimal control solutions with increasing iteration and achieve robustness to uncertainties in the Hamiltonian, specifically unitary control errors. Further exploring the utility of the approach, CLOAQC is used to optimize control profiles for the MAX 2-SAT problem and achieve improvements in computational accuracy and the size of the minimum energy gap relative to a linear control schedule. CLOAQC is designed to be conducive to current quantum annealing hardware and future AQC processors. Furthermore, the algorithm is quite versatile and can be adapted to alternative control scenarios, e.g., inhomogeneous driving [32].

ACKNOWLEDGMENT

The author thanks B. D. Clader, D. Lucarelli, and C. Ferrie for useful discussions.




FIG. 5. Comparison between CLOAQC optimized control paths and RC (top) and QAB (bottom) control paths for $n = 4$. CLOAQC is implemented with $K = 10000$ and $M = 100$. The median paths are shown by dark lines while shaded regions correspond to interquartile range for 100 realizations of CLOAQC. CLOAQC control paths correspond to the results presented in the main text (Fig. 1).

APPENDIX A: CONTROL SCHEDULES

1. Grover's search algorithm

A comparison between CLOAQC and the time-optimal GSA solutions corresponding to the main text results (Fig. 1) for $n = 4$ and $M = 100$ is shown in Fig. 5. The one control case yields optimized controls that closely resemble the RC control path, as seen from the top panel of Fig. 5. Deviations from the RC solution are attributed to the truncated polynomial expansion used to define the CLOAQC control field. In contrast, the complexity of the QAB profile is far more difficult for CLOAQC to reproduce under the fifth-order polynomial expansion; hence, a greater distinction between the CLOAQC optimized path and the QAB profile is observed (see Fig. 5, bottom panel). Interestingly, even with such a distinction, the relative difference in performance between CLOAQC and the QAB path is roughly only 10^{-3} for $n = 4$ and less than that for $n = 6$ and 8.

2. MAX 2-SAT

The CLOAQC optimized control paths for one MAX 2-SAT instance are shown in Figs. 6 and 7. The main text results for $H_I^{(xx)}$ are shown in Fig. 6 for one to four independent controls. Lines correspond to the median path, while shaded regions denote the interquartile range for 25 CLOAQC realizations. Figure 7 displays similar results for $H_I^{(y)}$ and $H_I^{(xz)}$ for three and four controls only. The minimum spectral gap for this particular problem instance lies between $s = 0.6$ and 0.8 for a majority of the control schedules. Note that the most significant ramping of the control schedules typically occurs within this range for $x_j(s)$, $j = 2, 3, 4$.




FIG. 6. CLOAQC control paths resulting from 25 realizations of CLOAQC implemented with $n = 8$, $K = 1000$, and $M = 100$. Results for one to four independent controls are shown in panels (a)–(d), respectively. The intermediate Hamiltonian is the two-local XX interaction discussed in the main text.

APPENDIX B: PERFORMANCE AS A FUNCTION OF RUNTIME

While CLOAQC performance varies as a function of runtime, numerical results indicate that the algorithm typically performs no worse than a linear schedule and typically converges towards known optimal solutions with increasing T . In this section, the performance of CLOAQC is examined as a function of runtime for both GSA and MAX 2-SAT using one independent control as an example of the typical convergence behavior of the algorithm. In Fig. 8, CLOAQC performance is assessed as a function of T for GSA. CLOAQC optimizations are performed with $K = 5000$ and $M = 100$ for 100 realizations of the protocol. The relative difference




FIG. 7. CLOAQC control paths resulting from 25 realizations of CLOAQC implemented with $K = 1000$ and $M = 100$. Panels (a) and (b) correspond to the $H_I^{(y)}$ intermediate Hamiltonian for three and four independent controls, respectively. Panels (c) and (d) correspond to the $H_I^{(xz)}$ intermediate Hamiltonian for three and four independent controls, respectively.

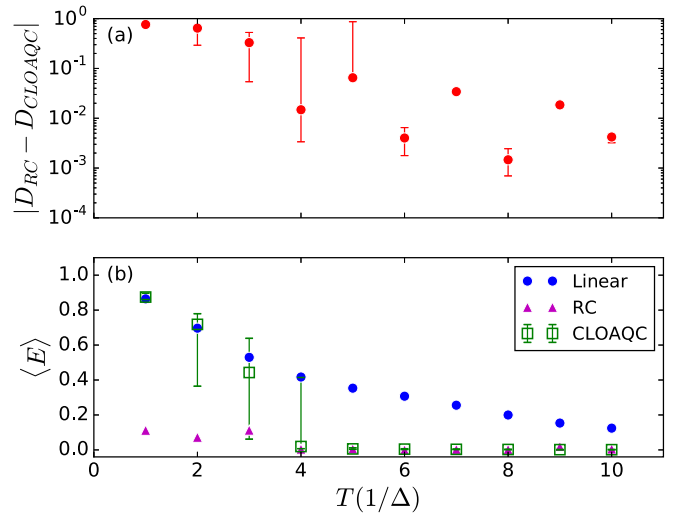


FIG. 8. CLOAQC performance as a function of the adiabatic runtime T for GSA using $n = 4$ qubits and one independent control. (a) Relative difference in adiabatic error between CLOAQC optimized path and RC solution. (b) Average energy for the initial (linear) control schedule, RC, and CLOAQC. Distributions include 100 realizations of CLOAQC with markers and whiskers denoting medians and interquartile ranges, respectively. CLOAQC is implemented using $K = 5000$ and $M = 100$.

between CLOAQC adiabatic error and the RC solution adiabatic error is displayed in panel (a) along with average energy for both the initial (linear) schedule, RC, and CLOAQC optimized control schedule in panel (b). In Fig. 9, the ratio of adiabatic errors and average energy for the linear and CLOAQC optimized controls are compared in panels (a) and (b), respectively, for the MAX 2-SAT problem with one




FIG. 9. CLOAQC performance as a function of the adiabatic runtime T for one representative $n = 8$ MAX 2-SAT USA instance using one independent control. (a) Ratio of adiabatic errors for the CLOAQC optimized path and the linear control schedule. (b) Average energy for the initial (linear) control schedule and CLOAQC. Distributions include 25 realizations of CLOAQC with markers and whiskers denoting medians and interquartile ranges, respectively. CLOAQC is implemented using $K = 1000$ and $M = 100$.

independent control. CLOAQC is implemented with $K = 1000$ and $M = 100$ for 25 realizations. Figure 9 focuses on one USA instance explored in the main text; however, the results shown here capture the typical behavior for the remaining 99 instances considered. Note that for both the GSA and MAX 2-SAT results markers denote medians, while whiskers denote the interquartile range.

CLOAQC successfully optimizes control schedules when initialized with controls paths that yield relatively low probabilities of sampling the ground state of H_P . As T increases, CLOAQC is able to more accurately reproduce the optimal RC path for GSA and offer improvements in adiabatic error over the linear schedule for MAX 2-SAT. This observation is evident from both the trace-norm distance and the average energy, where the ground-state energy E_0 is sampled with near unit median probability for the CLOAQC optimized path at $T \approx 7/\Delta$ and $\approx 20/\Delta$ for GSA and MAX 2-SAT,

respectively. In the case of GSA, the linear schedule only achieves an average energy of $\langle E \rangle \approx 0.3$, which corresponds to a ground-state sampling probability of $P(E_0) \approx 50\%$. Similar results are found for the MAX 2-SAT problem, where CLOAQC is shown to double the ground-state probability. Substantial improvements in $\langle E \rangle$ and $P(E_0)$ are also observed for shorter T values. For example, for GSA, CLOAQC can achieve a median $\langle E \rangle \approx 0.04$ and $P(E_0) \approx 96\%$ at $T \approx 4/\Delta$, where the linear control schedule yields $\langle E \rangle \approx 0.42$ and $P(E_0) \approx 38\%$. Similarly, for MAX-2SAT, at $T \approx 15/\Delta$, CLOAQC achieves $\langle E \rangle \approx 0.01$ and $P(E_0) \approx 96\%$ using an initial linear schedule that yields $\langle E \rangle \approx 0.1$ and $P(E_0) \approx 35\text{--}40\%$ for all problem instances considered. Lastly, it is important to note that CLOAQC is also capable of attaining nonzero $P(E_0)$ when initialized at T values where $P(E_0) = 0$. Such is the case at $T \approx 1/\Delta$ for a majority of the MAX 2-SAT USA instances discussed here.

-
- [1] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, [arXiv:quant-ph/0001106](#).
- [2] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001).
- [3] T. Albash and D. A. Lidar, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [4] The trace-norm distance $D[\rho_1, \rho_2] := \frac{1}{2} \|\rho_1 - \rho_2\|_1$, where $\|A\|_1 := \text{Tr}\sqrt{A^\dagger A}$, is a standard distance measure between quantum states [33]. When ρ_i is a pure state $\rho_i = |\psi_i\rangle\langle\psi_i|$ we replace the corresponding argument of D by $|\psi_i\rangle$.
- [5] S. Jansen, M.-B. Ruskai, and R. Seiler, *J. Math. Phys.* **48**, 102111 (2007).
- [6] D. A. Lidar, A. T. Rezakhani, and A. Hama, *J. Math. Phys.* **50**, 102106 (2009).
- [7] N. Wiebe and N. S. Babcock, *New J. Phys.* **14**, 013024 (2012).
- [8] J. Roland and N. J. Cerf, *Phys. Rev. A* **65**, 042308 (2002).
- [9] A. T. Rezakhani, W.-J. Kuo, A. Hama, D. A. Lidar, and P. Zanardi, *Phys. Rev. Lett.* **103**, 080502 (2009).
- [10] C. Brif, M. D. Grace, M. Sarovar, and K. C. Young, *New J. Phys.* **16**, 065013 (2014).
- [11] L. Zeng, J. Zhang, and M. Sarovar, *J. Phys. A* **49**, 165305 (2016).
- [12] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, *Nature (London)* **473**, 194 (2011).
- [13] D. Lucarelli, *Phys. Rev. A* **97**, 062346 (2018).
- [14] J. Nehrkorn, S. Montangero, A. Ekert, A. Smerzi, R. Fazio, and T. Calarco, [arXiv:1105.1707](#).
- [15] D. Herr, E. Brown, B. Heim, M. Konz, G. Mazzola, and M. Troyer, [arXiv:1705.00420](#).
- [16] J. C. Spall, *IEEE Trans. Autom. Control* **37**, 332 (1992).
- [17] C. Ferrie, *Phys. Rev. Lett.* **113**, 190404 (2014).
- [18] C. Ferrie and O. Moussa, *Phys. Rev. A* **91**, 052306 (2015).
- [19] C. E. Granade, Characterization, verification, and control for large quantum systems, Ph.D. thesis, University of Waterloo, 2015.
- [20] J. C. Spall, *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817 (1998).
- [21] P. Sadegh and J. C. Spall, *IEEE Trans. Autom. Control* **43**, 1480 (1998).
- [22] L. K. Grover, *Phys. Rev. Lett.* **79**, 325 (1997).
- [23] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, *SIAM J. Comput.* **26**, 1510 (1997).
- [24] S. Santra, G. Quiroz, G. V. Steeg, and D. A. Lidar, *New J. Phys.* **16**, 045006 (2014).
- [25] D. Coppersmith, D. Gamarnik, M. Hajiaghayi, and G. B. Sorkin, *Random Struct. Algorithms* **24**, 502 (2004).
- [26] E. Crosson, E. Farhi, C. Yen-Yu Lin, H.-H. Lin, and P. Shor, [arXiv:1401.7320](#).
- [27] E. Farhi, J. Goldstone, and S. Gutmann, [arXiv:quant-ph/0208135](#).
- [28] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, *Phys. Rev. B* **95**, 184416 (2017).
- [29] H. Nishimori and K. Takada, *Front. ICT* **4**, 2 (2017).
- [30] K. Takahashi, *Phys. Rev. A* **95**, 012309 (2017).
- [31] A. M. Childs, E. Farhi, and J. Preskill, *Phys. Rev. A* **65**, 012322 (2001).
- [32] J. I. Adame and P. L. McMahon, [arXiv:1806.11091](#).
- [33] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University, New York, 2000).