

From Pulses to Circuits and Back Again: A Quantum Optimal Control Perspective on Variational Quantum Algorithms

Alicia B. Magann^{1,2,†}, Christian Arenz^{3,†}, Matthew D. Grace², Tak-San Ho,³
Robert L. Kosut^{3,4}, Jarrod R. McClean,⁵ Herschel A. Rabitz³ and Mohan Sarovar^{2,*}

¹Department of Chemical & Biological Engineering, Princeton University, Princeton, New Jersey 08544, USA

²Extreme-Scale Data Science & Analytics, Sandia National Laboratories, Livermore, California 94550, USA

³Department of Chemistry, Princeton University, Princeton, New Jersey 08544, USA

⁴SC Solutions, Sunnyvale, California 94085, USA

⁵Google Research, 340 Main Street, Venice, California 90291, USA



(Received 11 September 2020; published 12 January 2021)

The last decade has witnessed remarkable progress in the development of quantum technologies. Although fault-tolerant devices likely remain years away, the noisy intermediate-scale quantum devices of today may be leveraged for other purposes. Leading candidates are variational quantum algorithms (VQAs), which have been developed for applications including chemistry, optimization, and machine learning, but whose implementations on quantum devices have yet to demonstrate improvements over classical capabilities. In this Perspective, we propose a variety of ways that the performance of VQAs could be informed by quantum optimal control theory. A major theme throughout is the need for sufficient control resources in VQA implementations; we discuss different ways this need can manifest, outline a variety of open questions, and look to the future.

DOI: [10.1103/PRXQuantum.2.010101](https://doi.org/10.1103/PRXQuantum.2.010101)

I. INTRODUCTION

The development of large-scale, fault-tolerant quantum computers would enable diverse and disruptive applications, such as the ability to break modern encryption protocols using Shor's factoring algorithm [1] and to efficiently simulate the dynamics of complex quantum systems [2]. Although significant progress has been made [3], persistent technological challenges in current quantum devices means they cannot yet serve as platforms for implementing these landmark algorithms at scale. As such, a major goal is to identify classically difficult problems that could be solved with these noisy intermediate-scale quantum (NISQ) [4] devices.

This goal has motivated the development of variational quantum algorithms (VQAs) for a variety of applications, including ground-state chemistry [5], optimization [6], and machine learning [7]. In VQAs, the problem to be solved is reformulated as an optimization problem, whose solution

is sought using quantum hardware and classical optimization in concert [8]. The quantum device is used to evaluate the objective function, which is accomplished via a relatively shallow, parametrized quantum circuit applied to an appropriately initialized register of qubits, after which the value of the objective function can be determined by measuring the register. Meanwhile, the classical coprocessor iteratively optimizes the parameters of the shallow quantum circuit. To date, hardware implementations of VQAs have not yet demonstrated improvements over the capabilities of classical computers, and the aim of this Perspective is to examine how progress can be made towards meeting this milestone in the future. In particular, this Perspective considers VQAs, their associated challenges, and potential paths forward, through the lens of quantum optimal control. To motivate this choice, we first look back and review certain aspects of the research efforts that have led us to the NISQ era of today.

We begin by recalling early efforts to create quantum computers, which focused on developing methods to control their components, and involved one or a few qubits [9,10]. In these experiments, control was typically realized using electromagnetic fields or "pulses" designed to drive the dynamics of the qubits in a desired fashion. Techniques for qubit control were studied extensively, especially in the context of implementing high-fidelity entangling gates between qubits, as this is a necessary

*mnsarov@sandia.gov

†These authors contributed equally to this work.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

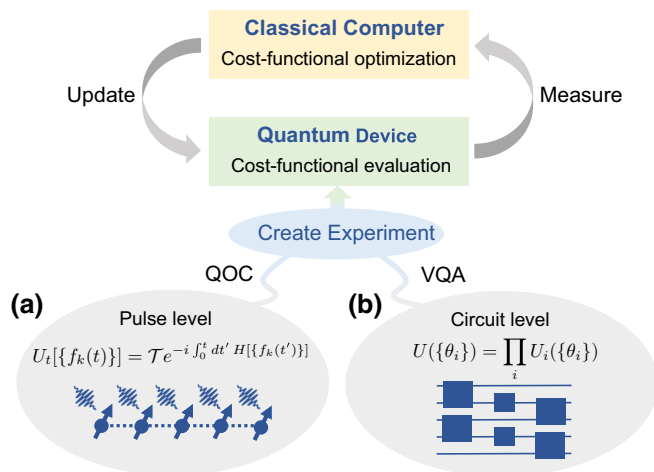


FIG. 1. The quantum-classical optimization loop used in VQA and QOC experiments is shown. The quantum device (green) evaluates the objective functional, whose value is determined via measurements and input into a classical computer (yellow), which updates the optimization parameters. In conventional QOC experiments, these parameters are defined at the pulse level (a), and serve to parametrize a set of control fields $\{f_k(t)\}$ entering in a Hamiltonian describing the physics of the system. In VQAs, the parameters typically enter at the circuit level (b), via a set of gates parametrized through a set of e.g., gate angles $\{\theta_i\}$. In this Perspective, we explore how VQAs may be informed by returning to a more physical, pulse-level description inspired by QOC.

ingredient for quantum computation in the gate model. One method, quantum optimal control (QOC), stands out for its ability to improve gate fidelities beyond what other techniques could offer [11,12]. In QOC, an objective functional is defined that quantifies how well a desired control task is achieved; then, the pulses to minimize the objective functional are sought using iterative optimization methods [13].

Following these early demonstrations of qubit control, devices began to scale up to higher qubit counts, which has led to the advent of the NISQ era today, and in tandem, to the development of VQAs. These concurrent developments have inspired significant research on circuit compilation and optimization [14–16], similar in spirit to many of the earlier efforts that studied pulse optimization in the context of QOC. Thus, the reach of technology and the focus of the community have broadened “*from pulses to circuits.*” In this Perspective, we explore how the development of VQAs can be informed by going from the circuit level “*back again*” to the pulse level to strengthen ties to QOC and leverage results and tools from this well-developed field.

As illustrated in Fig. 1, we explore this prospect by framing VQA implementations and QOC experiments as quantum-classical optimization loops; in the former (b), the optimization is done over quantum circuit elements, in

contrast to a conventional QOC experiment (a), where the optimization is performed over a set of continuous pulses. We remark now that although the parametrized quantum circuits in (b) are formed by gates, which are in turn implemented using pulses, a user seeking to implement a VQA typically has no need for any knowledge of what happens at the pulse level, which offers a layer of abstraction separating the user from the underlying hardware physics.

With this unifying picture in mind, we now describe the remainder of this article. We begin by introducing the concept of VQAs and discussing their current state in Sec. II. We then discuss QOC theory in Sec. III and review connections that have been made to VQAs to date in Sec. IV. In Sec. V we present a hierarchy of abstractions in variational optimization that serves to provide a common framework for VQA and QOC experiments. This is followed by an in-depth discussion of QOC-motivated future research directions aimed at addressing four important challenges associated with VQAs: ansatz selection, optimization landscapes, noise, and robustness. In each of these cases, we emphasize the need for appropriate control resources to enhance the performance of VQAs. Finally, we conclude in Sec. VI with a look ahead.

II. VARIATIONAL QUANTUM ALGORITHMS

Variational quantum algorithms seek to solve problems by leveraging the dynamical and representational power of quantum computers in conjunction with classical computers. They do so by reformulating problems of interest as the minimization of some objective or loss function $J[\{\theta_i\}]$ over a set of parameters $\{\theta_i\}$, as

$$\min_{\{\theta_i\}} J[\{\theta_i\}], \quad (1)$$

where $\{\theta_i\}$ parametrizes a quantum circuit

$$U(\{\theta_i\}) = \prod_i U_i(\theta_i), \quad (2)$$

on n qubits. The objective function can be nonlinear in the general case, such as cross-entropy for machine learning. However, for reasons of simplicity, it is often formulated as the minimization of the expectation of a linear operator H_p , referred to as the “problem” Hamiltonian [17]:

$$J[\{\theta_i\}] = \langle \psi(\{\theta_i\}) | H_p | \psi(\{\theta_i\}) \rangle, \quad (3)$$

where the state of the qubits at the culmination of the circuit is $|\psi(\{\theta_i\})\rangle = U(\{\theta_i\})|\psi_0\rangle$, with $|\psi_0\rangle$ denoting their fixed initial state. Exact minimization of J corresponds to the preparation of the ground state of H_p .

The set of variational parameters $\{\theta_i\}$ that minimize $J[\{\theta_i\}]$ are sought iteratively, where at every iteration, $J[\{\theta_i\}]$ is evaluated by preparing the qubits in the

state $|\psi_0\rangle$, applying a circuit $U(\{\theta_i\})$ with a particular parametrization $\{\theta_i\}$, and then measuring a set of qubit observables to estimate $J[\{\theta_i\}]$. This can be accomplished by expanding H_p in the Pauli operator basis as per $H_p = \sum_{j=1}^N \alpha_j P_j$, where α_j are scalar coefficients, P_j are tensor products of Pauli operators that are easy to measure on a quantum device, and $N = \mathcal{O}[\text{poly}(n)]$ [18], and then measuring the expectations of each of the N Pauli basis operators in the expansion. Due to the stochastic nature of measurement in quantum mechanics, repeated measurements on an identically prepared state $|\psi(\{\theta_i\})\rangle$ are needed to estimate these expectations. Then, the expectation value of H_p can be computed to determine J by classically evaluating the weighted sum $\langle \psi(\{\theta_i\}) | H_p | \psi(\{\theta_i\}) \rangle = \sum_{j=1}^N \alpha_j \langle \psi(\{\theta_i\}) | P_j | \psi(\{\theta_i\}) \rangle$. After the value of $J[\{\theta_i\}]$ has been evaluated in this manner, a classical optimization routine is then used to iteratively update the values of $\{\theta_i\}$ until convergence. Using this method to estimate the value of $J[\{\theta_i\}]$ to a specified precision ϵ requires a number of repetitions and measurements of the state that scales as $N_s \propto \lambda^2/\epsilon^2$, where $\lambda = \sum_j |\alpha_j|$ [19]. Recent work has shown that clever grouping of terms and other techniques can be used to reduce the naive scaling of these measurements by orders of magnitude, even with techniques in the near term [20]. As quantum computers advance, it is possible to improve the scaling of this estimation even further using techniques that leverage phase estimation, but the increased resource costs for such approaches can be prohibitive.

For practical reasons, the circuit represented by $U(\{\theta_i\})$ is usually assumed to be formed by a sequence of elementary one- and two-qubit gates drawn from a specified gate set, which is constructed and parametrized according to a particular ansatz. However, we remark here that knowledge of how the parameters $\{\theta_i\}$ explicitly enter into the circuit is not always required, as $J[\{\theta_i\}]$ is evaluated via measurements using the quantum device. That is, the method is robust to many types of labeling errors, since properties of the quantum system rather than specific parameter values are of interest. This means that VQAs possess some degree of robustness to drifts, crosstalk, and other systematic errors that can occur during the implementation of the circuit [21], which is a primary reason why VQAs are believed to be a way to derive practical algorithmic use from NISQ devices.

Furthermore, we note that the choice of ansatz for any VQA is a crucial step. Although there is no general approach for developing good ansätze, they are often derived from physical intuition (e.g., the QAOA ansatz, see below), knowledge of states generated by a particular ansatz (e.g., the coupled-cluster ansatz, see below), or practical convenience (e.g., hardware-efficient ansätze [5,22]). A critical feature of an ansatz is that it should be scalable, i.e., having a circuit depth that scales as $\mathcal{O}[\text{poly}(n)]$ for an n -qubit quantum computer.

Furthermore, the number of variational parameters $\{\theta_i\}$ should also scale as $\mathcal{O}[\text{poly}(n)]$.

VQAs have been developed for numerous application areas including machine learning [7], linear systems [23], and the compilation of quantum circuits [24]. However, the first application area for VQAs was the ground-state problem in quantum chemistry [8,25]. In this context, the variational quantum eigensolver (VQE) was developed as a VQA for seeking the electronic ground state of a chemical system in a field of fixed nuclear charges. The solution of this chemistry problem has a variety of applications, including in chemical reaction prediction, the determination of molecular properties, etc. One common ansatz is the unitary coupled-cluster ansatz [5,26], which is a norm-preserving variant of the common coupled-cluster ansatz used in quantum chemistry, which constructs a size-extensive ansatz through an exponential parametrization. Due to its unitary formulation, it naturally preserves physical properties of the state, but it is not efficient to evaluate classically. It represents an example of a structured ansatz that carries the fermionic structure into its translation into gates after the use of a Jordan-Wigner transformation and Suzuki-Trotter splitting. Strictly speaking, this formulation represents a slight deviation from the formal construction [27], but the construction remains unitary, independent of the choice of parameters, and upon repetition it can be used to express arbitrary states within the manifold of fixed particle number.

With the evidence that NISQ devices can achieve classically intractable, but perhaps not useful, tasks [28], there is a belief that variational preparation of ground states of correlated systems may represent one of the first classically intractable and useful roles for NISQ devices. As such, the study of correlated quantum systems remains a major objective for NISQ devices, and is a focus of many experimental and theoretical efforts at present. VQE demonstrations have been shown experimentally on a variety of photonic, ion trap, and superconducting qubit setups in combination with error-mitigation techniques [21,22,29–33]. Indeed the discrepancy between available qubits in current quantum devices ($\gtrsim 50$ qubits) and the number used in VQE experiments ($\lesssim 10$ qubits) is that the impact of noise makes the experiments incompatible with the high accuracy necessary to claim an application advantage. With the help of error mitigation from symmetries in the reduced subspace, the largest variational calculation performed to chemical accuracy on a quantum computer utilized 12 qubits and approximately 200 quantum gates to simulate a 12-atom hydrogen chain [34]. Using more qubits will require advances in both hardware and error-reduction techniques.

Another major application of VQAs is combinatorial optimization. In this domain, the quantum approximate optimization algorithm (QAOA) has been developed as a

variational method for determining approximate solutions to combinatorial optimization problems, by encoding them into diagonal Ising Hamiltonians, such that the solution of the problem is encoded in the ground state of the Hamiltonian [6]. QAOA seeks to find the solution by variationally minimizing the expectation value of the Ising Hamiltonian. Unlike the VQE, the QAOA ansatz is typically fixed, and consists of an alternating sequence of unitary operations generated by the problem Hamiltonian H_p , and a so-called “driver” Hamiltonian, which does not commute with the problem Hamiltonian, and is denoted by H_d . Explicitly, the QAOA ansatz is formed by p rounds of alternating applications of these two Hamiltonians,

$$U(\{\theta_i\}) = \prod_{j=1}^p \exp(-i\beta_j H_d) \exp(-i\gamma_j H_p), \quad (4)$$

where $\{\theta_i\}$ is the set of variational parameters $\{\beta_j, \gamma_j\}$, and the true optimum for the original combinatorial optimization problem can be achieved as $p \rightarrow \infty$ [6].

QAOA has been implemented experimentally using superconducting circuits for up to 23 qubits [35–39], a photonic system with two qubits [40], and trapped ions with up to 40 qubits [41]. Due to limited qubit coherence times, most of these implementations (with some exceptions [37,39]) considered combinatorial optimization problems defined on the connectivity graph of the hardware only, keeping circuit depths at a minimum. Recently, it has been shown that success in quantum optimization can be related to increased flexibility in the training landscape of a continuous extension of the original problem [42], which is connected to the condition of full controllability in quantum control theory. Furthermore, although in theory, the ansatz with higher p should improve on the quality of the solutions achievable at lower p , they also require deeper circuits. So in practice, the effects of noise and decoherence may negate any improvements in quality. For this reason, most experiments implemented a single round of QAOA only, i.e., $p = 1$. These cases involved only two parameters whose optimal values could be obtained analytically, meaning that a variational optimization loop was not necessary. However, in recent superconducting circuit experiments it was shown that going to $p = 2$ [38] or $p = 3$ [39] and performing the optimization could improve the solution quality beyond $p = 1$.

Although these recent hardware demonstrations show a promising trend, the ultimate goal of implementing VQAs to solve problems that are intractable classically has yet to be reached. The path towards meeting this goal will involve a confluence of theoretical and experimental progress. In the following, we propose a few ways that this progress could be informed by QOC. In particular, we focus in on how methods and results from QOC could be leveraged to address VQA challenges associated

with ansatz selection, optimization landscapes, noise and robustness. Before getting to this, we first introduce QOC and discuss its connections to VQAs.

III. QUANTUM OPTIMAL CONTROL

The aim of QOC is to design one or more electromagnetic fields or “pulses” to steer the dynamics of a quantum system towards a desired control target, which can be a state, observable expectation value, or evolution operator, at some terminal time T . A standard formulation in QOC seeks to minimize a control objective functional $J[\{f_k(t)\}]$ over $\{f_k(t)\}$, as [43–47]

$$\min_{\{f_k(t)\}} J[\{f_k(t)\}], \quad (5)$$

where $J[\{f_k(t)\}]$ includes the control target and physical constraints, often along with other criteria, which can be defined to represent available laboratory resources or quantify robustness to errors or uncertainties [48].

The set of pulses $\{f_k(t)\}$ used in QOC are typically considered to be classical fields in the semiclassical approximation [49], in contrast to fully quantized fields. In addition, the wavelengths of the fields are typically assumed to be much greater than the size of the controlled quantum system in the dipole approximation [49], such that the control fields $\{f_k(t)\}$ typically enter the time-dependent Hamiltonian $H[\{f_k(t)\}]$, as follows [50,51]:

$$H[\{f_k(t)\}] = H_0 + \sum_k f_k(t) H_k, \quad (6)$$

where H_0 is the drift Hamiltonian describing the time-independent system and $\{H_k\}$ is the set of control Hamiltonians that couple the fields to the system, e.g., via dipole interactions. The dynamical equation for the system time-evolution operator U_t is given by the Schrödinger equation $\dot{U}_t = -iH[\{f_k(t)\}]U_t$, with $U_0 = \mathbb{1}$ (throughout this article, we set $\hbar = 1$). This is a bilinear control system, making an analytical formulation of QOC solutions intractable in general [50–52]. Its formal solution reads

$$U_t[\{f_k(t)\}] = \mathcal{T} e^{-i \int_0^t dt' H[\{f_k(t')\}]}, \quad (7)$$

where \mathcal{T} indicates time ordering, such that the system state at time t is given by $|\psi(t)\rangle = U_t|\psi_0\rangle$ where $|\psi_0\rangle$ is the initial state.

In analogy to the formulation of VQAs in Eqs. (1) and (3), we now turn our attention to the QOC problem in Eq. (5), where we define

$$J[\{f_k(t)\}] = \langle \psi(T) | H_p | \psi(T) \rangle, \quad (8)$$

such that (unconstrained) $\{f_k(t)\}$ are sought to minimize the expectation value of H_p at a designated time T . Solutions

of this QOC problem, subject to the dynamical constraint that $|\psi(t)\rangle$ evolves according to the Schrödinger equation, are stationary points of the control objective functional J , given the initial condition $|\psi(0)\rangle = |\psi_0\rangle$. Optimal control fields can be constructed iteratively via the corresponding Euler-Lagrange equations [53]. To this end, a plethora of methods have been developed for updating the QOC solutions, including GRAPE [54], Krotov [55–57], TBQCP [58–60], D-MORPH [61–63], and SCP [64], until specified optimality or convergence criteria are satisfied.

In practice, the minimization of J is usually accomplished by first parametrizing one or more continuous control fields by a set of variables $\{c_i\}$, such that the Hamiltonian describing the controlled system becomes

$$H(t, \{c_i\}) = H_0 + \sum_k f_k(t, \{c_{i,k}\}) H_k. \quad (9)$$

Common parametrizations include setting $\{c_i\}$ to be the amplitudes and phases of a set of frequency components of the field, or setting $\{c_i\}$ to be piecewise-constant field amplitudes in the time domain. Then, the objective is to optimize $\{c_i\}$ to generate $U_T(\{c_i\}) = \mathcal{T} e^{-i \int_0^T H(t, \{c_i\}) dt}$ such that the control objective functional $J[\{c_i\}]$ is minimized at the terminal time T . In practice, this minimization is typically performed in an iterative fashion. If a tractable and accurate model is available, this iteration can be performed numerically. However, it can also be carried out experimentally via learning control [65–70], which does not require knowledge of the underlying system model. Instead, at each iteration of such QOC experiments, $J[\{c_i\}]$ is evaluated by first preparing the system in a specified initial state, then evolving it in the presence of applied fields with parametrization $\{c_i\}$, and finally measuring the observable expectation value(s) needed to estimate $J[\{c_i\}]$. A classical optimization routine is used to update the values of $\{c_i\}$ from one iteration to the next, until $J[\{c_i\}]$ converges [65]. In this manner, objective functional evaluations are performed by the quantum system directly, inherently accounting for parameter uncertainties and other systematic errors, limitations, etc.

This quantum-classical optimization loop associated with QOC at the pulse level is directly analogous to the procedure used in VQA implementations at the circuit level, as shown in Fig. 1. In fact, if one has access to the quantum circuit at the pulse level, then Eq. (1) can instead be solved by optimizing over the set of continuous pulses $\{f_k(t)\}$ that are available. As such, we consider VQA implementations to be a form of *digital* QOC experiment on qubits, where the quantum circuit generating the unitary transformation $U(\{\theta_i\})$ is designed directly, through the selection and optimization of a parametrized unitary ansatz. In general, the depth, dimension, and structure of the ansatz, as well as the continuous and discrete parameters within it, are all tunable, giving the resulting

optimization space both continuous and discrete degrees of freedom.

In the past few years, this fundamental relationship between QOC and VQA has been exploited to derive a deeper understanding of VQAs and novel variational strategies for quantum computing problems. In the following section, we review some of this work at the intersection of VQAs and QOC.

IV. PRIOR WORK CONNECTING QOC WITH VQAS

As argued above, standard parametrized quantum circuit ansätze can be viewed as examples of digitized QOC implementations. There are many benefits to relaxing, or embedding, such digitized ansätze into a continuously (in time) parametrized framework, similar to the typical setting in QOC. Such a relaxation often allows one to eliminate any discrete optimization component of the problem, and more importantly, by formulating VQAs within a standard control theory setting, this enables one to apply many powerful methods and results of optimal control theory. Of course, any relaxation of a parametrized ansatz into a continuous-in-time ansatz must be done in a way that is consistent with the capabilities of the available control hardware—e.g., since most quantum computing platforms have native one- and two-qubit gates, the natural continuously parametrized ansätze will be composed of pulses that address individual qubits or the coupling between two qubits.

An early example of work with this reformulation is by Yang *et al.* [71], who used a continuously parametrized formulation of variational quantum optimization to demonstrate that a bang-bang approach (similar to the alternating structure of QAOA) is optimal for preparing the state encoding the optimization solution, given amplitude-constrained control fields and a finite time to solution. This is achieved by applying Pontryagin’s minimum principle [53] to the continuously parametrized formulation of the problem. Similarly, Lin *et al.* apply Pontryagin’s minimum principle associated with time-optimal quantum control to Grover’s quantum search problem, and find that the time-optimal control solution has a bang-singular-bang structure [72].

Another direction has considered connections between the QOC concept of *controllability* and the quantum computing concept of *computational universality* [73]. In brief, controllability is the study of which control objectives can be realized with a given set of controls and constraints. For unconstrained control fields, the dynamical Lie algebra \mathcal{L} , formed by iterated commutators of the drift and the control Hamiltonians and their real linear combinations, is a powerful tool for deciding these matters. In particular, the dynamical Lie algebra gives rise to the Lie rank criterion [50,51], which states that if \mathcal{L} spans the full space

[i.e., the special unitary algebra $\mathfrak{su}(2^n)$ for a n qubit system], then every unitary transformation $V \in \text{SU}(2^n)$ can be created to arbitrary precision in finite time by shaping the control fields, and the system is said to be *fully operator controllable*. A vast literature characterizing the controllability of quantum systems has been developed in recent decades [74–86]. More recently, the relationship between controllability and computational universality was utilized in Refs. [87,88] to show that the QAOA ansatz is universal for quantum computing for specific choices of the problem Hamiltonian. In addition, Mbeng *et al.* made connections between digitized quantum annealing, QAOA, and QOC, e.g., analyzing the number of angles that are needed for controllability [89], while Akshay *et al.* examined reachability deficits in QAOA, providing strategies for improving reachability [90].

Some groups have investigated using the QAOA ansatz for bang-bang control of state transitions in quantum spin systems, e.g., Refs. [91–93], exploring robustness and reachability as a function of the ansatz depth. In addition, Bapat and Jordan analyzed the performance of bang-bang control protocols for optimization algorithms, showing that on certain problem instances, these protocols can yield an exponential speedup for both classical and quantum optimization, compared with quasistatic scheduling [94]. Using QOC, Brady *et al.* [95] investigate the optimal continuous-time control solution for preparing an approximation of a ground state with a time-dependent Hamiltonian that is the sum of the driver and problem Hamiltonians found in QAOA, and the total evolution time is fixed. They demonstrate that the optimal solution can in fact be more general than the bang-bang solution identified by Yang *et al.* [71], because an assumption made in that work (the absence of “singular” intervals) is not generally valid. With a careful variational analysis, along with numerical simulations of various transverse-field Ising models, Brady *et al.* argue that the more general form of the optimal QAOA variational schedule is a *bang-anneal-bang* schedule, whereby the control parameters take extreme values at the beginning and end of the time interval, but take on intermediate values that smoothly vary during intermediate times. With a fermionic representation, Wang *et al.* [96] show that the evolution of a quantum system implementing QAOA on the so-called “ring of disagrees” problem translates into QOC of an ensemble of independent spins, thereby simplifying the determination of the optimal angle vectors. On a related note, Wu *et al.* propose a scheme to machine learning tasks into corresponding QOC problems on NISQ devices [97].

Recently, several groups have proposed adaptive, variational, or QOC-inspired approaches to design improved VQA ansätze [98–103]. For example, the approach developed in Refs. [98–100] uses derivative information to adaptively modify the circuit depth and ansätze structure using a “pool” of predetermined single- or multiqubit

Hamiltonian operators. Whereas, in the context of quantum chemistry, the approach presented in Ref. [101] uses a set of QOC-informed driving Hamiltonians to generate VQA ansätze with symmetry-breaking features that can decrease the circuit depth required for convergence. Similarly, Ref. [102] proposes ansätze determined by QOC at the device level, rather than parametrized quantum circuits, to perform VQE simulations. These last works especially strengthen the connections between QOC and VQAs, and provide a natural segue to the next section, where we outline some promising new directions of research at this intersection.

V. NEW DIRECTIONS FOR VQAS INFORMED BY QOC

A unifying view of VQAs and QOC can be obtained by viewing both as formulations of variational optimization at different levels within an abstraction hierarchy. This is illustrated in Fig. 2, which we now discuss. We first assume that the objective function J , determined by the application, is shared between the two approaches. Then, one can think of experimental ways to evaluate J within a hierarchy of abstractions modeling the experimental hardware. At each level (i)–(iii) of this hierarchy, there is a natural parametrization of the control one has over the hardware and this defines a natural variational ansatz at that level.

At the bottom of this hierarchy is a pulse-level abstraction (i), where we are closest to a first-principles model of the hardware and think in terms of Hamiltonians for

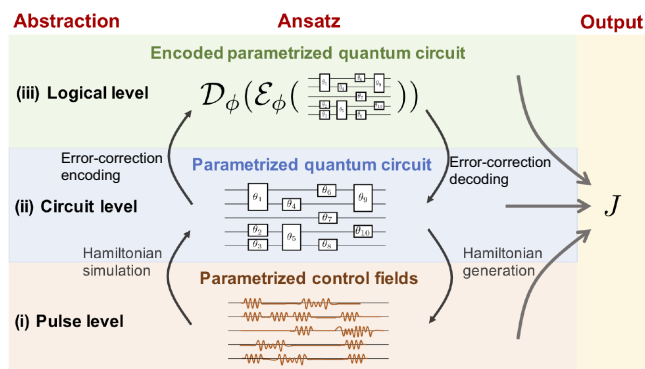


FIG. 2. A representation of the hierarchy of abstractions used to model quantum hardware, and how variational optimization enters into each level of the hierarchy. Output produced by the hardware is the desired objective function, J , which is parametrized in different ways by each level of the hierarchy—by control-field parameters c_i at the pulse level (i), by circuit parameters (usually angles θ_i) at the quantum circuit level (ii), and by the structure of the circuit and the encoding and decoding maps (\mathcal{E}_ϕ and \mathcal{D}_ϕ) at the logical circuit level (iii). By understanding the relationship between the natural variational parametrizations at each of the levels, we can develop a rich family of variational ansätze.

the localized computing elements (e.g., qubits, qudits) and fields coupling them. At this level, the parametrization of control is in terms of a continuously parametrized control field or control Hamiltonian (with parameters c_i) that is often realized by a set of electromagnetic fields, which are coupled to the computing elements. QOC typically operates at this level of the modeling hierarchy.

In the middle of the hierarchy (ii), we abstract away first-principles descriptions of the hardware and think in terms of universal circuit elements, or gates, that (ideally) perform well-defined maps on the computing elements. Although there might be a discrete set of types of gates, they can be continuously parametrized by some set of gate parameters $\{\theta_i\}$. VQAs typically operate at this level of the modeling hierarchy.

Finally, at the highest level of the modeling hierarchy, the logical circuit level (iii), we think in terms of circuits operating on quantum states encoded within an error-correcting code. Any physical circuit at the circuit level of abstraction can be converted into a logical circuit given an error-correcting code and its associated logical gates, with an error rate that is determined by the code and the hardware. Strictly speaking, with conventional error-correcting codes, only random errors in discrete gates are thought to be error correctable to arbitrary precision, thus, arbitrary rotation gates depending on θ_i are synthesized as a sequence of discrete gates, which performs this rotation to a specified precision. Hence, the rotation angle θ_i is still present in the logical circuit, but only up to the precision that is given by the synthesis and code procedure. In practice, one may optimize the angle as if it is continuous, so long as the synthesis map is performed after, and the precision is great enough to impact the optimization in practice. Error-corrected quantum computing experiments operate at this level of the modeling hierarchy, and benefit from decreased susceptibility to hardware noise due to the encoding and careful implementation of fault-tolerant operations.

Importantly, an error-correction procedure could be parametrized and considered as variational parameters (these are denoted by the collective parameter ϕ in Fig. 2). This is not done conventionally, but is a new avenue that opens up when one considers performing VQAs at the logical level of abstraction. For example, if (even partial) characterization of the noise sources in a NISQ device is possible, as shown in Refs. [104–108], efficient quantum error correction (QEC), which can significantly outperform standard schemes is possible. Unitary encoding and recovery operations tuned to the range and “character” of noise-induced uncertainties can be calculated via a biconvex optimization. The circuit representations of these channel-optimized unitaries can then be parametrized. By adding these parameters to the VQA parameters, the QEC system becomes another part of the “learning control” environment. Furthermore, continuous

parameters like error likelihoods in decoding could be adjusted to attempt to improve performance, and even more sophisticated schemes that allow discrete code modifications could be another potential area of research. While maintaining fault tolerance under variation of the error-correction procedure can be challenging, doing such variational optimization of small encodings could be useful in the intermediate term where fully fault-tolerant operation is not feasible, but access to limited error correction is possible.

A. Ansatz selection


With this hierarchy of modeling abstractions and a description of variational optimization at each level, we can exploit the connections between the levels in the hierarchy to define new and richer variational ansätze, and in fact, a family of ansätze, built from paths on this diagram. To see this, we explore a few sample paths across this diagram and understand the implications of following particular trajectories.

As a first example, we consider a path from the pulse level (i) to the circuit level (ii). We begin at the pulse level (i), with a continuous control perspective of fields acting on n qubits, which is then discretized as a parametrized control. We now have a time-dependent Hamiltonian $H(t, \{c_i\})$ acting on a system of n qubits. The field of quantum algorithms of simulation of time-dependent Hamiltonians is well developed, and a range of methods exist for simulating the time evolution generated by $H(t)$ over some desired time interval $[0, T]$ to arbitrary accuracy. Among the simplest is an operator splitting, also known as “Trotter factorization” [109], but a host of methods with more accurate implementations without direct classical simulation analogs, including quantum walk and so-called linear combinations of unitary approaches, exist that may also be used [110–112]. Hence, we may understand this step as implying that the parametrized control can be combined with Hamiltonian simulation to yield a quantum circuit, whose parametrization is naturally understood at the level of $H(t)$. As such, optimizations proceed on a landscape determined by a parametrization at the pulse level (i), despite being implemented at the circuit level (ii). That is, while arbitrary rotations depending on some θ_i exist in the circuit, they are entirely determined by the composition $\theta_i(\{c_i\})$ in conjunction with the chosen Hamiltonian simulation map. This has the advantage that this circuit can be converted to a logical circuit by discretization into error-correctable gates, mitigating the effects of decoherence while retaining the essential prescription of control. In this case, we retain all of the power of QOC machinery in manipulating and understanding the ansatz, while leveraging the power of quantum algorithms and error correction to ensure theoretical guarantees of implementation.

We can take this same path farther, by appending to it additional gate parametrizations. That is, if Hamiltonian simulation also maps parametrized control to a circuit at the quantum circuit level (ii), e.g., Ref. [113], additional parameters may also be added to the generated circuits, creating a hybrid ansatz. To explore the full power of these connections though, let us explore the other direction that one can take in this perspective.

Consider starting at the quantum circuit level (ii), with descriptions of hardware in terms of quantum circuits. To make these circuits realizable, a structure and parametrization is selected. Once this circuit has been determined, one can, in principle, map this back to a QOC problem at the pulse level (i) in a number of ways, where now the control parametrization depends on the angles from the circuit such that $c_i(\{\theta_i\})$ is determined through a map that we denote as ‘‘Hamiltonian generation.’’ This mapping is typically nonunique in a more severe way than in the other direction, especially if one considers mapping the entire circuit to a generating Hamiltonian, whether it be time independent or time dependent. However, much better formed mappings can be used and are related to existing strategies for gate design. For example, one map can perform a gate-wise mapping for each parametrized gate back to the Hamiltonian control parameters $c_i(\{\theta_i\})$ that are used to accurately implement this gate. Similar to before, a hybrid control ansatz can then be created by dropping the dependence on θ_i and allowing free variation of the parameter c_i . This can also be done for fixed gates without parametrizations in the gate model, increasing the overall expressiveness of the resulting circuit in a systematic and new way. We see then, that the gate formulation also naturally provides a family of control ansätze.

To give a specific example, we consider a system and ansatz that has received considerable attention in the VQA literature, which is the preparation of the H_2 ground state in a minimal, molecular orbital basis. Without symmetry reduction, the ground state of this problem is given by $|\Psi_{\text{GS}}\rangle = \cos \theta |0011\rangle + \sin \theta |1100\rangle$. A simplified version of a circuit that contains this state within its parametrization is given by



Where, in this diagram, H is the Hadamard gate, R_y is a rotation about the Pauli Y axis, and the connecting line is a two-qubit controlled- Z gate, $CZ = \text{diag}(1, 1, 1, -1)$. In this case, variation over the parameter θ provides the variational freedom required to prepare the ground state.

In addition, as outlined above, we can map back to control implementations of gates, including those that do not yet have parameters. For example, a controlled-phase gates may be written as a time-independent evolution under a Hamiltonian of the form $H_z = c_{12}Z_1Z_2 + c_1Z_2 + c_2Z_2$ for a time τ , up to an unimportant global phase. If this is how the gate is implemented in our physical system, then the parameter τ can be added to the list of parameters to create a hybrid control ansatz. The ansatz in this case includes the beginning of the circuit, as well as analog evolution under H_z , which may be implemented by physical means.

Finally, taking the mapping in the other direction, the evolution under the Hamiltonian H for a time τ in this case, could be redigitized into a Trotter factorization of $\exp(-iH_z\tau) = \exp(-i\tau c_{12}Z_1Z_2) \exp(-i\tau c_1Z_1) \exp(-i\tau c_2Z_2)$, where each term in the product can be translated back into a digital sequence, and now we can vary the set of parameters $\{\theta, c_1, c_2, c_{12}\}$, where we retain the θ parametrization and the rest of the circuit, but add more controls. Note that for this simple problem, if the implementation is error free, this additional freedom may be unnecessary. However, if some errors occur in the gates, this freedom can be used to improve the result. Moreover, problems that are similar, but not identical to the hydrogen molecule, can then leverage this base ansatz and natural extensions from it. This simple example is designed to illustrate the ways that one can move from the circuit level (ii) to a hybrid ansatz that combines tools from the circuit and pulse levels (ii) and (i) and back again, while changing the nature of the ansatz.

Finally, we see that beginning at any level in the hierarchy in Fig. 2, it is possible to iterate on the connections between these approaches, choosing to carry forward parametrizations or develop new ones, and map back to the other corresponding formalism. For example, if one takes a variational parametrization at the control-field level, maps it to a quantum circuit with a Hamiltonian simulation map, then selects a particular gate and maps it back to a control parametrization, and uses the combined circuit, we get a new hybrid parametrization that can be manipulated closer to a device-level description. It is easy to see that these choices can be mixed, matched, and iterated on to create an entire family of parametrizations. Here, we envision that the tools of QOC are bolstered by concepts from digital simulation, including error correction and advanced time-propagation algorithms, and vice versa.

In addition to expanding the set of options and tools of both areas, we hope that this formulation and perspective of connections allows more direct cross-pollination of ideas. For example, by formulating a circuit ansatz based on control theory, we more directly understand the consequences of controllability in a circuit model. In turn, by developing a control-theoretic ansatz from a circuit family known to be universal, we may find new insights in relation

to entanglement-generating power or expressiveness of quantum circuits.

B. Optimization landscapes

A critical consideration in VQAs is the difficulty of optimizing over the circuit parameters. The ease of finding the global minimum of J during this optimization process is dictated by the structure of the underlying optimization landscape as well as available prior information about the location of good optima. A central theme of QOC is the study of such landscapes, and translating these tools into the context of VQAs will allow for new strategies in ansatz and problem design.

At their core, VQAs represent the mapping of a convex optimization problem in an exponentially large linear space, to a nonconvex optimization problem over parametrized quantum circuits. This can be seen by expanding the variational state $|\psi(\{\theta_i\})\rangle$ in the eigenbasis $\{|n\rangle\}$ of H_p . The objective functional then takes the form $J[\{\theta_i\}] = \sum_{n=1}^d \lambda_n(\{\theta_i\}) E_n$, where E_n are the eigenvalues of H_p and $\lambda_n = |\langle n | \psi(\{\theta_i\}) \rangle|^2$. When J is optimized over the set $\{\lambda_n\}$, the optimization problem is convex, which implies that the corresponding optimization landscape is free from local optima. In contrast, the parameters θ_i typically enter in a nonlinear fashion in λ_n , often via an exponential map in $U_i(\theta_i)$. As such, when J is optimized over the set $\{\theta_i\}$ the optimization problem is nonconvex and in general, local optima can appear.

In the context of QOC, to address these and other landscape considerations, a sizeable body of research has analyzed the topological features and properties of the (dynamic) QOC landscape, defined by J as a functional of the control fields [114–119]. To characterize these landscapes, recall that J depends on the time-evolution operator up to time T , such that we have $J = J\{U_T[f(t)]\}$. Then, the functional derivative of J with respect to a single field $\delta J/\delta f(t)$ is given by the composition $\delta J/\delta f(t) = \partial J/\partial U_T \circ \delta U_T/\delta f(t)$. The first term in the composition captures the properties of the kinematic control landscape, i.e., J defined as a function of U_T . The composition above implies that if the variation of U_T with respect to the control field $\delta U_T/\delta f(t)$ is assumed to be full rank, i.e., equal to d^2 , then the dynamic and kinematic critical points coincide. This result is referred to as *local surjectivity*. In this scenario, the topology of the dynamic control landscape is fully characterized by the critical point structure of the kinematic control landscape. Results from QOC theory have shown that the kinematic control landscapes of typical objective functionals consist of global extrema and saddles [116,117,120–122]. As such, if we are able to create every unitary transformation such that J as a function of U_T can be varied arbitrarily, and if U_T can be varied arbitrarily also by varying the control field, then the landscape of J consists of saddles and global extrema. More

precisely, under the assumptions of (1) a full dynamical Lie algebra, (2) unconstrained control fields of arbitrary length and shape, and (3) local surjectivity, the control landscapes for commonly employed objective functionals are free from local extrema. Although debated in the literature [118,119,123–129], extensive QOC simulation and laboratory studies indicate the relative ease of satisfying all three assumptions [49]. Furthermore, it has been shown that assumptions (1) and (3) are almost always satisfied in a measure theoretic sense [74,127,130], thereby implying that QOC landscapes are almost always free of local minima under the premise of sufficient control-field resources [127]. While the precise meaning of “sufficient” is application dependent and remains an open challenge to systematically assess, it has recently been shown that local surjectivity is almost always met when the control fields allow for approximating Haar random evolutions [130,131] within the time interval of interest: $[0, T]$. As such, even though performing the optimization over a set of parametrized control fields does not avoid the nonconvexity of the problem, provided the assumptions (1)–(3) hold, the interior of the optimization landscape provably contains saddles only.

Looking ahead, we hope that the theory of QOC landscapes can provide insight into optimization landscapes of VQAs, which could include better tools to understand different ansätze. Most directly, the optimization landscape of a variational ansatz formulated at the control-field level in Fig. 2 could be analyzed in terms of existing QOC landscape theory. Here, in addition to the theoretical foundations, numerical tools, such as D-MORPH (a flexible continuation-method developed in the context of QOC) [132,133] and FLACCO (software for feature-based landscape analysis of continuous and constrained optimization problems) [134], could be employed to explore these landscapes. However, there are a number of caveats to consider for this direct approach; that is, although QOC has many well-developed tools for the characterization of QOC landscapes, there remain challenges for applying these tools to scalable and practical implementations of VQAs. For instance, characterizing the landscape topology in the presence of constraints is challenging; although these landscapes are expected to be free from local extrema when sufficiently many unconstrained variables $\{c_i\}$ are used to parametrize the control fields, there is evidence that an insufficient number of parameters (i.e., less than d^2), leading to a violation of the surjectivity assumption (3), and control-field constraints are among the reasons for local extrema to appear [135]. In general, further research is needed when the control resources do not scale with the dimension d of the quantum system being controlled, but rather only scale as $\mathcal{O}\{\text{poly}[\log(d)]\}$. To this end, it may be interesting to study systems that are not fully controllable, but where all states within a certain subspace defined by the symmetries of the (controlled) system are reachable.

One may then ask whether a (significantly) smaller number of control parameters could be used to obtain a trap-free landscape when moving only in a restricted subspace (containing the ground state) whose dimension does not scale exponentially.

We also remark that as the dimension d of the quantum system increases the landscape flattens out, such that for certain objective functionals, gradients become exponentially small as a function of the number of qubits, making these regions difficult to leave with local optimization algorithms. These so-called barren plateaus [25] are a consequence of a concentration of measure phenomenon, and occur in both VQA implementations and QOC [130]. Understanding how to avoid them in each context independently could yield useful and transferable techniques benefiting both areas of study, and will likely involve careful design of ansatz, initialization, and training methods, e.g., Ref. [136].

C. Noise and time-optimal control

Today, errors in NISQ devices severely restrict the circuit depths that are achievable for VQAs. For example, certain errors can arise from stochastic fluctuations in the gate implementations, leading to errors that accumulate as circuit depth increases. In QOC, robust control strategies have been developed to suppress random errors like this, by seeking pulses that are robust in the presence of finite control noise; the condition of the Hessian of J can be used to determine such properties [133]. We believe more direct translation of robust control tools from QOC into variational algorithms will help improve their robustness against general errors.

Other, often more insidious, errors stem from interactions of the qubits with their surrounding environment that can cause the system to decohere over time. Thus, the timescale associated with performing a quantum circuit should ideally be restricted to the coherence time of the system. When coherence times are limited and gates are dominated by stochastic errors, it is desirable to seek VQA circuits that drive the qubits to the minimum of J as quickly as possible, i.e., using circuits with minimum depth. In the context of QOC, the theory of time-optimal control, and the associated theory of quantum speed limits, provide a powerful framework for considering this issue at the pulse level, and may offer valuable tools to enhance VQA performance. In general, identifying time-optimal fields is a challenging task, as it corresponds to finding the associated geodesic on the unitary group [125,137–141]. Exact results are only known for certain systems consisting of 1–3 qubits [142–146]. However, upper and lower bounds on T can be found [147–150]. In addition, for n -qubit networks, the upper bounds in Ref. [149,150] allow for characterizing the unitary transformations that can be created efficiently with $2n$ local

fields, i.e., where $T = \mathcal{O}[\text{poly}(n)]$. While some progress has recently been made to characterize efficiently controllable qubit graphs [151], in general, it remains an open challenge to systematically determine the set of unitary transformations that are reachable in polynomial time with fewer controls. Finally, we remark that additional control resources not only offer faster control strategies and in general, richer ansätze, but can also allow for counteracting decoherence by enabling decoupling schemes [152–155] that can suppress interactions with the environment.

D. Robust control through digitization

All the directions covered by the previous subsections are fairly natural in the VQA setting. However, if we think more broadly, the scope of interactions between VQAs and QOC could be even more fruitful. In particular, consider how VQAs could impact traditional QOC experiments. Until this point, we have considered a setting natural to VQAs, where one typically starts from a known initial state, e.g., $|00\dots 0\rangle$ and builds up the state of interest using a set of parametrized gates or controls. However, many QOC experiments deal with situations where one is given an unknown quantum state $|\psi\rangle$, and wishes to manipulate it. For example, this quantum state may be the state of reactants in a chemical system and the control goal is to steer the reaction products. Alternatively, in a quantum sensing context, the quantum state may be the state of the sensed environment and the control goal is to distinguish features in the environment.

If we now consider such a QOC experiment from the point of view of VQAs, and think of the applied controls as a parametrized unitary transformation, one can translate these controls into a parametrized circuit ansatz [113] and further, into an error-corrected circuit. This would have the effect of ensuring that the controls are applied to an arbitrarily high degree of accuracy, an advantage of taking the digital point of view. Of course, to be fully compatible with quantum error correction, the input quantum state may also need to be encoded, and effective digital encoding of quantum states from nature is an experimental challenge that has not yet been realized, but this perspective has the potential to enable QOC experiments of arbitrary accuracy on many-body systems.

The increasing difficulty of precisely controlling a quantum system with nonideal analog controls as its dimension increases can be understood in various ways. Conceptually, this is a consequence of the orthogonality catastrophe—where small perturbations from the ideal Hamiltonian (in this setting, caused by small errors in the analog controls) can lead to an exponential decay (with system dimension) of fidelity with the state under ideal evolution [156,157]. More operationally, consider an N -body quantum system with state vector of dimension $d \sim \exp(N)$.

We can estimate the requirements on control precision with a simple model. Assume we want to prepare the state with 1 in the first entry of the state vector, and 0 on the remaining $(d - 1)$ entries. Now for each dimension we are able to achieve a precision δ , and consider for simplicity the normalized erred state that is off by δ in each of the entries except the first. In this case, the fidelity with the target state is given by $1/[1 + (d - 1)\delta^2]$, which means to maintain a constant fidelity, one must have $\delta \propto d^{-1/2}$. This is an exponential (in N) precision requirement for the control.

In contrast, if the state of this N -body system can be *transduced* into $n = \log_2(d)$ qubits, the control of the state can be encoded in a VQA circuit [113]. Then, assuming the fidelity of qubit gates are sufficiently high and using the tools of quantum error correction, the fidelity of controlling this state can be brought arbitrarily close to 1 with logarithmic overhead in physical resources.

At a broader application level, the manipulation of unknown quantum states in a VQA manner encompasses what is sometimes called “quantum machine learning” for training data provided in the form of quantum states rather than classical inputs, or quantum data. This idea allows for the QOC tools to come to bear in this community. Moreover, this permits a closer connection to query-based proofs in quantum computer science, where stronger proofs are possible in the setting where only a limited number of quantum states are available [158–162].

VI. OUTLOOK

In this Perspective, we have explored the connections between VQAs and the theory of QOC. While VQA applications constitute some of the most promising applications of near-term quantum computers, more progress bridging the gap between theoretical VQAs and NISQ computing hardware is necessary to realize this promise. We have argued that concepts from QOC, and the more general framework outlined in Sec. V, which integrates traditional models of QOC with VQAs and extends both, are critical to bridging this gap. Both fields have something to gain from exploring this fertile area of overlap. From the perspective of VQAs, the mature theory and powerful tools of QOC can provide richer variational structures and offer a deeper understanding of variational experiments. Conversely, from the perspective of QOC, VQAs present an exciting frontier of many-body quantum control, pushing this established field in new directions. For these reasons, we expect bountiful fruits at the intersection of these two areas of research in the years to come.

ACKNOWLEDGMENTS

A.B.M., M.D.G., and M.S. are supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under the Quantum Computing Application Teams program. A.B.M.

also acknowledges support from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award No. DE-FG02-97ER25308. M.D.G. also acknowledges support from the U.S. Department of Energy, Office of Advanced Scientific Computing Research, under the Quantum Algorithm Teams program. C.A. and H.A.R. are supported by the U.S. Army Research Office, Grants No. W911NF-16-1-0014 and No. W911NF-19-1-0382, respectively. T.S.H. acknowledges support from the U.S. Department of Energy, Grant No. DE-FG02-02ER15344. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under Contract No. DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

-
- [1] P. W. Shor, Algorithms for Quantum Computation: Discrete Logarithms and Factoring, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, SFCS '94 (IEEE Computer Society, USA, 1994), p. 124.
 - [2] S. Lloyd, Universal Quantum Simulators, *Science* **273**, 1073 (1996).
 - [3] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, and C. Monroe, arXiv:1906.02700 (2020).
 - [4] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [5] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2013).

- [6] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm, arXiv:1411.4028 [quant-ph] (2014).
- [7] V. Dunjko and P. Wittek, A non-review of Quantum Machine Learning: trends and explorations, *Quantum Views* **4**, 32 (2020).
- [8] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [9] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, Coherent control of macroscopic quantum states in a single-Cooper-pair box, *Nature* **398**, 786 (1999).
- [10] T. Yamamoto, Y. A. Pashkin, O. Astafiev, Y. Nakamura, and J.-S. Tsai, Demonstration of conditional gate operation using superconducting charge qubits, *Nature* **425**, 941 (2003).
- [11] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. Du, P. Neumann, *et al.*, Quantum error correction in a solid-state hybrid spin register, *Nature* **506**, 204 (2014).
- [12] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, *et al.*, High-fidelity spin entanglement using optimal control, *Nat. Commun.* **5**, 3371 (2014).
- [13] We note that we use the term *quantum optimal control* to refer solely to the closed-loop control scenario just described, and do not consider other quantum control settings such as real-time feedback control in this Perspective.
- [14] D. Venturelli, M. Do, E. Rieffel, and J. Frank, *Quantum Sci. Technol.* **3**, 025004 (2018).
- [15] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, *Quantum* **3**, 140 (2019).
- [16] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (2019), p. 1015.
- [17] The term “problem Hamiltonian” for this quantity is standard in optimization applications of VQAs. It is often called “system Hamiltonian” in other applications.
- [18] We remark here that this is a common property of optimization problems, like the ones we consider in this Perspective, but does not always hold.
- [19] N. C. Rubin, R. Babbush, and J. McClean, Application of fermionic marginal constraints to hybrid quantum algorithms, *New J. Phys.* **20**, 053020 (2018).
- [20] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, arXiv:1907.13117 (2019).
- [21] P. J. J. O’Malley *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 31007 (2016).
- [22] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
- [23] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Scaling of variational quantum circuit depth for condensed matter systems, arXiv:1909.05820 [quant-ph] (2019).
- [24] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, *Quantum* **3**, 140 (2019).
- [25] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 1 (2018).
- [26] A. G. Taube and R. J. Bartlett, New perspectives on unitary coupled-cluster theory, *Int. J. Quantum Chem.* **106**, 3393 (2006).
- [27] H. R. Grimsley, D. Claudino, S. E. Economou, E. Barnes, and N. J. Mayhall, Is the Trotterized UCCSD Ansatz Chemically Well-Defined?, *J. Chem. Theory Comput.* **16**, 1 (2019).
- [28] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [29] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [30] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, Error mitigation extends the computational reach of a noisy quantum processor, *Nature* **567**, 491 (2019).
- [31] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Robust Determination of Molecular Spectra on a Quantum Processor, *Phys. Rev. X* **8**, 011021 (2018).
- [32] S. E. Smart and D. A. Mazziotti, Quantum-classical hybrid algorithm using an error-mitigating N -representability condition to compute the Mott metal-insulator transition, *Phys. Rev. A* **100**, 022517 (2019).
- [33] R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. C. Bultink, X. Fu, C. H. Price, V. P. Ostroukh, N. Muthusubramanian, A. Bruno, M. Beekman, N. Haider, T. E. O’Brien, and L. DiCarlo, Experimental error mitigation via symmetry verification in a variational quantum eigensolver, *Phys. Rev. A* **100**, 010302 (2019).
- [34] G. A. Quantum *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [35] J. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, *et al.*, Unsupervised machine learning on a hybrid quantum computer, arXiv:1712.05771 (2017).
- [36] M. Willsch, D. Willsch, F. Jin, H. De Raedt, and K. Michielsen, Benchmarking the quantum approximate optimization algorithm, *Quantum Inf. Process.* **19**, 197 (2020).
- [37] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, Implementation of the XY interaction family with calibration of a single pulse, arXiv:1912.04424 (2019).
- [38] A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Križan, D. Shiri, I.-M. Svensson *et al.*, Improved success probability with greater

- circuit depth for the quantum approximate optimization algorithm, arXiv:1912.10495 (2019).
- [39] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, et al., Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, arXiv:2004.04197 (2020).
- [40] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O’Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph, *et al.*, Large-scale silicon quantum photonics implementing arbitrary two-qubit processing, *Nat. Photonics* **12**, 534 (2018).
- [41] G. Pagano, A. Bapat, P. Becker, K. Collins, A. De, P. Hess, H. Kaplan, A. Kyprianidis, W. Tan, C. Baldwin, *et al.*, Quantum Approximate Optimization of the Long-Range Ising Model with a Trapped-Ion Quantum Simulator, arXiv:1906.02700 (2019).
- [42] J. R. McClean, M. P. Harrigan, M. Mohseni, N. C. Rubin, Z. Jiang, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Low depth mechanisms for quantum optimization, arXiv:2008.08615 (2020).
- [43] R. Kosloff, S. A. Rice, P. Gaspard, S. Tersigni, and D. Tanner, Wavepacket dancing: Achieving chemical selectivity by shaping light pulses, *Chem. Phys.* **139**, 201 (1989).
- [44] S. Shi and H. Rabitz, Quantum mechanical optimal control of physical observables in microsystems, *J. Chem. Phys.* **92**, 364 (1990).
- [45] K. Sundermann and R. de Vivie-Riedle, Extensions to quantum optimal control algorithms and applications to special problems in state selective molecular dynamics, *J. Chem. Phys.* **110**, 1896 (1999).
- [46] D. Dong and I. R. Petersen, Quantum control theory and applications: a survey, *IET Control. Theory Appl.* **4**, 2651 (2010).
- [47] S. Lloyd and S. Montangero, Information Theoretical Analysis of Quantum Optimal Control, *Phys. Rev. Lett.* **113**, 010502 (2014).
- [48] A. P. Peirce, M. A. Dahleh, and H. Rabitz, Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications, *Phys. Rev. A* **37**, 4950 (1988).
- [49] C. Brif, R. Chakrabarti, and H. Rabitz, Control of quantum phenomena: past, present and future, *New J. Phys.* **12**, 075008 (2010).
- [50] D. Elliott, *Bilinear control systems: Matrices in action, Bilinear Control Systems: Matrices in Action* (Springer Science & Business Media, New York, NY, 2009), Vol. 169.
- [51] D. d’Alessandro, *Introduction to Quantum Control and Dynamics* (CRC Press, Boca Raton, FL, 2007).
- [52] V. Jurdjevic and H. J. Sussmann, Control systems on Lie groups, *J. Differ. Eq.* **12**, 313 (1972).
- [53] R. F. Stengel, *Optimal Control and Estimation* (Springer, New York, NY, 1994).
- [54] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [55] J. P. Palao and R. Kosloff, Optimal control theory for unitary transformations, *Phys. Rev. A* **68**, 062308 (2003).
- [56] D. M. Reich, M. Ndong, and C. P. Koch, Monotonically convergent optimization in quantum control using Krotov’s method, *J. Chem. Phys.* **136**, 104103 (2012).
- [57] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. P. Koch, Krotov: A Python implementation of Krotov’s method for quantum optimal control, *SciPost Phys.* **7**, 80 (2019).
- [58] T.-S. Ho and H. Rabitz, Accelerated monotonic convergence of optimal control over quantum dynamics, *Phys. Rev. E* **82**, 026703 (2010).
- [59] T.-S. Ho, H. Rabitz, and S.-I. Chu, A general formulation of monotonically convergent algorithms in the control of quantum dynamics beyond the linear dipole interaction, *Comput. Phys. Commun.* **182**, 14 (2011).
- [60] S.-L. Liao, T.-S. Ho, S.-I. Chu, and H. Rabitz, Fast-kick-off monotonically convergent algorithm for searching optimal control fields, *Phys. Rev. A* **84**, 031401 (2011).
- [61] A. Rothman, T.-S. Ho, and H. Rabitz, Observable-preserving control of quantum dynamics over a family of related systems, *Phys. Rev. A* **72**, 023416 (2005).
- [62] A. Rothman, T.-S. Ho, and H. Rabitz, Quantum observable homotopy tracking control, *J. Chem. Phys.* **123**, 134104 (2005).
- [63] A. Rothman, T.-S. Ho, and H. Rabitz, Exploring the level sets of quantum control landscapes, *Phys. Rev. A* **73**, 053401 (2006).
- [64] R. L. Kosut, M. D. Grace, and C. Brif, *Phys. Rev. A* **88**, 052326 (2013).
- [65] R. S. Judson and H. Rabitz, Teaching Lasers to Control Molecules, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [66] D. J. Egger and F. K. Wilhelm, Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems, *Phys. Rev. Lett.* **112**, 240503 (2014).
- [67] D. Lu, K. Li, J. Li, H. Katiyar, A. J. Park, G. Feng, T. Xin, H. Li, G. Long, A. Brodutch, *et al.*, Enhancing quantum control by bootstrapping a quantum processor of 12 qubits, *NPJ Quantum Inf.* **3**, 45 (2017).
- [68] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid Quantum-Classical Approach to Quantum Optimal Control, *Phys. Rev. Lett.* **118**, 150503 (2017).
- [69] Q.-M. Chen, X. Yang, C. Arenz, R.-B. Wu, X. Peng, I. Pelcer, and H. Rabitz, Combining the synergistic control capabilities of modeling and experiments: Illustration of finding a minimum-time quantum objective, *Phys. Rev. A* **101**, 032313 (2020).
- [70] X.-D. Yang, C. Arenz, I. Pelcer, Q.-M. Chen, R.-B. Wu, X.-h. Peng, and H. Rabitz, Assessing three closed-loop learning algorithms by searching for high-quality quantum control pulses, arXiv:2008.03874 [quant-ph] (2020).
- [71] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing Variational Quantum Algorithms Using Pontryagin’s Minimum Principle, *Phys. Rev. X* **7**, 021027 (2017).
- [72] C. Lin, Y. Wang, G. Kolesov, and U. c. v. Kalabić, Application of Pontryagin’s minimum principle to Grover’s quantum search problem, *Phys. Rev. A* **100**, 022327 (2019).
- [73] V. Ramakrishna and H. Rabitz, Relation between quantum computing and quantum controllability, *Phys. Rev. A* **54**, 1715 (1996).

- [74] C. Altafini, Controllability of quantum mechanical systems by root space decomposition of $\mathfrak{su}(N)$, *J. Math. Phys.* **43**, 2051 (2002).
- [75] F. Albertini and D. D'Alessandro, The Lie algebra structure and controllability of spin systems, *Linear Algebra Appl.* **350**, 213 (2002).
- [76] S. G. Schirmer, H. Fu, and A. I. Solomon, Complete controllability of quantum systems, *Phys. Rev. A* **63**, 063410 (2001).
- [77] G. Turinici, in *Mathematical Models and Methods for Ab Initio Quantum Chemistry* (Springer, New York, NY, 2000), p. 75.
- [78] H. Fu, S. G. Schirmer, and A. I. Solomon, Complete controllability of finite-level quantum systems, *J. Phys. A* **34**, 1679 (2001).
- [79] G. Turinici and H. Rabitz, Wavefunction controllability for finite-dimensional bilinear quantum systems, *J. Phys. A* **36**, 2565 (2003).
- [80] D. Burgarth, D. D'Alessandro, L. Hogben, S. Severini, and M. Young, Zero forcing, linear and quantum controllability for systems evolving on networks, *IEEE Trans. Automat. Contr.* **58**, 2349 (2013).
- [81] R. Heule, C. Bruder, D. Burgarth, and V. M. Stojanović, Local quantum control of Heisenberg spin chains, *Phys. Rev. A* **82**, 052333 (2010).
- [82] Z. Zimborás, R. Zeier, M. Keyl, and T. Schulte-Herbrüggen, A dynamic systems approach to fermions and their relation to spins, *EPJ Quantum Technol.* **1**, 11 (2014).
- [83] M. G. Genoni, A. Serafini, M. Kim, and D. Burgarth, Dynamical Recurrence and the Quantum Control of Coupled Oscillators, *Phys. Rev. Lett.* **108**, 150501 (2012).
- [84] G. Dirr, U. Helmke, I. Kurniawan, and T. Schulte-Herbrüggen, Lie-semigroup structures for reachability and control of open quantum systems: Kossakowski-Lindblad generators form Lie wedge to Markovian channels, *Rep. Math. Phys.* **64**, 93 (2009).
- [85] D. K. Burgarth, P. Facchi, V. Giovannetti, H. Nakazato, S. Pascazio, and K. Yuasa, Exponential rise of dynamical complexity in quantum computing through projections, *Nat. Commun.* **5**, 1 (2014).
- [86] C. Arenz, D. Burgarth, P. Facchi, V. Giovannetti, H. Nakazato, S. Pascazio, and K. Yuasa, Universal control induced by noise, *Phys. Rev. A* **93**, 062308 (2016).
- [87] S. Lloyd, Quantum approximate optimization is computationally universal, arXiv:1812.11075 [quant-ph] (2018).
- [88] M. E. Morales, J. Biamonte, and Z. Zimborás, On the universality of the quantum approximate optimization algorithm, *Quantum Inf. Process.* **19**, 1 (2020).
- [89] G. Bigan Mbeng, R. Fazio, and G. Santoro, Quantum Annealing: a journey through Digitalization, Control, and hybrid Quantum Variational schemes, arXiv:1906.08948 [quant-ph] (2019).
- [90] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, Reachability Deficits in Quantum Approximate Optimization, *Phys. Rev. Lett.* **124**, 090504 (2020).
- [91] Y. Dong, X. Meng, L. Lin, R. Kosut, and K. B. Whaley, Robust Control Optimization for Quantum Approximate Optimization Algorithm, arXiv:1911.00789 [quant-ph] (2019).
- [92] M. Y. Niu, S. Lu, and I. L. Chuang, Optimizing QAOA: Success Probability and Runtime Dependence on Circuit Depth, arXiv:1905.12134 [quant-ph] (2019).
- [93] J. Yao, M. Bukov, and L. Lin, Policy Gradient based Quantum Approximate Optimization Algorithm, arXiv:2002.01068 [quant-ph] (2020).
- [94] A. Bapat and S. Jordan, Bang-bang control as a design principle for classical and quantum optimization algorithms, *Quantum Inf. Comput.* **19**, 424 (2019).
- [95] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal Protocols in Quantum Annealing and QAOA Problems, arXiv:2003.08952 [quant-ph] (2020).
- [96] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for MaxCut: A fermionic view, *Phys. Rev. A* **97**, 022304 (2018).
- [97] R.-B. Wu, X. Cao, P. Xie, and Y.-X. Liu, End-to-End Quantum Machine Learning with Quantum Control Systems, arXiv:2003.13658 [quant-ph] (2020).
- [98] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).
- [99] H. L. Tang, V. O. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor, arXiv:1911.10205 [quant-ph] (2020).
- [100] L. Zhu, H. L. Tang, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, arXiv:2005.10258 [quant-ph] (2020).
- [101] A. Choquette, A. Di Paolo, P. K. Barkoutsos, D. Sénéchal, I. Tavernelli, and A. Blais, Quantum-optimal-control-inspired ansatz for variational quantum algorithms, arXiv:2008.01098 [quant-ph] (2020).
- [102] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, Gate-free state preparation for fast variational quantum eigensolver simulations: ctrl-VQE, arXiv:2008.04302 (2020).
- [103] Z.-J. Zhang, T. H. Kyaw, J. S. Kottmann, M. Degroote, and A. Aspuru-Guzik, Mutual information-assisted Adaptive Variational Quantum Eigensolver Ansatz Construction, arXiv:2008.07553 [quant-ph] (2020).
- [104] M. Reimpell and R. F. Werner, Iterative Optimization of Quantum Error Correcting Codes, *Phys. Rev. Lett.* **94**, 080501 (2005).
- [105] A. S. Fletcher, P. W. Shor, and M. Z. Win, Optimum Quantum Error Recovery using Semidefinite Programming, *Phys. Rev. A* **75**, 012338 (2007).
- [106] R. L. Kosut, A. Shabani, and D. A. Lidar, Robust Quantum Error Correction via Convex Optimization, *Phys. Rev. Lett.* **100**, 020502 (2008).
- [107] R. L. Kosut and D. A. Lidar, Quantum error correction via convex optimization, *Quantum Inf. Process.* **8**, 443 (2009).
- [108] S. Taghavi, R. L. Kosut, and D. A. Lidar, Channel-Optimized Quantum Error Correction, *IEEE Trans. Inf. Theory* **56**, 1461 (2010).

- [109] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, *J. Math. Phys.* **32**, 400 (1991).
- [110] D. W. Berry and A. M. Childs, Black-box Hamiltonian simulation and unitary implementation, arXiv:0910.4157 (2009).
- [111] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [112] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [113] A. B. Magann, M. D. Grace, H. A. Rabitz, and M. Sarovar, Digital quantum simulation of molecular dynamics and control, arXiv:2002.12497 [quant-ph] (2020).
- [114] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Quantum optimally controlled transition landscapes, *Science* **303**, 1998 (2004).
- [115] H. Rabitz, T.-S. Ho, M. Hsieh, R. Kosut, and M. Demiralp, Topology of optimally controlled quantum mechanical transition probability landscapes, *Phys. Rev. A* **74**, 012721 (2006).
- [116] T.-S. Ho and H. Rabitz, Why do effective quantum controls appear easy to find? *J. Photochem. Photobiol. A* **180**, 226 (2006).
- [117] T.-S. Ho, J. Dominy, and H. Rabitz, Landscape of unitary transformations in controlled quantum dynamics, *Phys. Rev. A* **79**, 013422 (2009).
- [118] A. N. Pechen and D. J. Tannor, Are There Traps in Quantum Control Landscapes?, *Phys. Rev. Lett.* **106**, 120402 (2011).
- [119] P. De Fouquieres and S. G. Schirmer, A closer look at quantum control landscapes and their implication for control optimization, *Infin. Dimens. Anal. Quantum Probab. Relat. Top.* **16**, 1350021 (2013).
- [120] R. Chakrabarti and H. Rabitz, Quantum control landscapes, *Int. Rev. Phys. Chem.* **26**, 671 (2007).
- [121] H. Rabitz, M. Hsieh, and C. Rosenthal, Optimal control landscapes for quantum observables, *J. Chem. Phys.* **124**, 204107 (2006).
- [122] M. Hsieh and H. Rabitz, Optimal control landscape for the generation of unitary transformations, *Phys. Rev. A* **77**, 042306 (2008).
- [123] H. Rabitz, T.-S. Ho, R. Long, R. Wu, and C. Brif, Comment on “are there traps in quantum control landscapes?”, *Phys. Rev. Lett.* **108**, 198901 (2012).
- [124] A. N. Pechen and D. J. Tannor, *Phys. Rev. Lett.* **108**, 198902 (2012).
- [125] K. W. Moore Tibbetts, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T.-S. Ho, R.-B. Wu, and H. Rabitz, Exploring the tradeoff between fidelity and time optimal control of quantum unitary transformations, *Phys. Rev. A* **86**, 062309 (2012).
- [126] G. Riviello, C. Brif, R. Long, R.-B. Wu, K. M. Tibbetts, T.-S. Ho, and H. Rabitz, Searching for quantum optimal control fields in the presence of singular critical points, *Phys. Rev. A* **90**, 013404 (2014).
- [127] B. Russell, H. Rabitz, and R.-B. Wu, Control landscapes are almost always trap free: A geometric assessment, *J. Phys. A* **50**, 205302 (2017).
- [128] D. V. Zhdanov, *J. Phys. A: Math. Theor.* **51**, 508001 (2018).
- [129] B. Russell, R. Wu, and H. Rabitz, *J. Phys. A: Math. Theor.* **51**, 508002 (2018).
- [130] C. Arenz and H. Rabitz, Drawing together quantum control and quantum tomography principles, arXiv:2004.02729 [quant-ph] (2020).
- [131] L. Banchi, D. Burgarth, and M. J. Kastoryano, Driven quantum dynamics: Will it blend?, *Phys. Rev. X* **7**, 041015 (2017).
- [132] K. W. Moore and H. Rabitz, Exploring quantum control landscapes: Topology, features, and optimization scaling, *Phys. Rev. A* **84**, 012109 (2011).
- [133] D. Hocker, C. Brif, M. D. Grace, A. Donovan, T.-S. Ho, K. M. Tibbetts, R. Wu, and H. Rabitz, Characterization of control noise effects in optimal quantum unitary dynamics, *Phys. Rev. A* **90**, 062309 (2014).
- [134] P. Kerschke, Comprehensive Feature-Based Landscape Analysis of Continuous and Constrained Optimization Problems Using the R-Package FLACCO, arXiv:1708.05258 [stat.ML] (2017).
- [135] G. Riviello, K. Moore-Tibbetts, C. Brif, R. Long, R.-B. Wu, T.-S. Ho, and H. Rabitz, Searching for quantum optimal controls under severe constraints, *Phys. Rev. A* **91**, 043401 (2015).
- [136] T. Volkoff and P. J. Coles, Large gradients via correlation in random parameterized quantum circuits, arXiv:2005.12200 [quant-ph] (2020).
- [137] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, Time-optimal Quantum Evolution, *Phys. Rev. Lett.* **96**, 060503 (2006).
- [138] M. A. Nielsen, M. R. Dowling, M. Gu, and A. C. Doherty, Quantum computation as geometry, *Science* **311**, 1133 (2006).
- [139] M. A. Nielsen, M. R. Dowling, M. Gu, and A. C. Doherty, Optimal control, geometry, and quantum computing, *Phys. Rev. A* **73**, 062323 (2006).
- [140] A. Carlini, A. Hosoya, T. Koike, and Y. Okudaira, Time-optimal unitary operations, *Phys. Rev. A* **75**, 042308 (2007).
- [141] X. Wang, M. Allegra, K. Jacobs, S. Lloyd, C. Lupo, and M. Mohseni, Quantum Brachistochrone Curves as Geodesics: Obtaining Accurate Minimum-time Protocols for the Control of Quantum Systems, *Phys. Rev. Lett.* **114**, 170501 (2015).
- [142] F. Albertini and D. D’Alessandro, Minimum time optimal synthesis for two level quantum systems, *J. Math. Phys.* **56**, 012106 (2015).
- [143] R. Romano, Geometric analysis of minimum-time trajectories for a two-level quantum system, *Phys. Rev. A* **90**, 062302 (2014).
- [144] G. C. Hegerfeldt, Driving at the Quantum Speed Limit: Optimal Control of a Two-Level System, *Phys. Rev. Lett.* **111**, 260501 (2013).
- [145] N. Khaneja, R. Brockett, and S. J. Glaser, Time optimal control in spin systems, *Phys. Rev. A* **63**, 032308 (2001).
- [146] N. Khaneja, S. J. Glaser, and R. Brockett, Sub-Riemannian geometry and time optimal control of three spin systems: quantum gates and coherence transfer, *Phys. Rev. A* **65**, 032301 (2002).

- [147] C. Arenz, B. Russell, D. Burgarth, and H. Rabitz, The roles of drift and control field constraints upon quantum control speed limits, *New J. Phys.* **19**, 103015 (2017).
- [148] J. Lee, C. Arenz, H. Rabitz, and B. Russell, Dependence of the quantum speed limit on system size and control complexity, *New J. Phys.* **20**, 063002 (2018).
- [149] C. Arenz and H. Rabitz, Controlling Qubit Networks in Polynomial Time, *Phys. Rev. Lett.* **120**, 220503 (2018).
- [150] J. Lee, C. Arenz, D. Burgarth, and H. Rabitz, An upper bound on the time required to implement unitary operations, *J. Phys. A* **53**, 125304 (2020).
- [151] C. Gokler, S. Lloyd, P. Shor, and K. Thompson, Efficiently Controllable Graphs, *Phys. Rev. Lett.* **118**, 260501 (2017).
- [152] L. Viola, Quantum control via encoded dynamical decoupling, *Phys. Rev. A* **66**, 012307 (2002).
- [153] J. R. West, D. A. Lidar, B. H. Fong, and M. F. Gyure, High Fidelity Quantum Gates via Dynamical Decoupling, *Phys. Rev. Lett.* **105**, 230503 (2010).
- [154] C. Arenz, G. Gualdi, and D. Burgarth, Control of open quantum systems: case study of the central spin model, *New J. Phys.* **16**, 065023 (2014).
- [155] R. L. Kosut, C. Arenz, and H. Rabitz, Quantum control landscape of bipartite systems, *J. Phys. A* **52**, 165305 (2019).
- [156] T. Dowdall, A. Benseny, T. Busch, and A. Ruschhaupt, Fast and robust quantum control based on Pauli blocking, *Phys. Rev. A* **96**, 043601 (2017).
- [157] T. Fogarty, S. Deffner, T. Busch, and S. Campbell, Orthogonality Catastrophe as a Consequence of the Quantum Speed Limit, *Phys. Rev. Lett.* **124**, 110601 (2020).
- [158] E. Knill, Quantum Randomness and Nondeterminism, arXiv:quant-ph/9610012 [quant-ph] (1996).
- [159] R. Cleve, P. Hoyer, B. Toner, and J. Watrous, Consequences and limits of nonlocal strategies, in *Proceedings 19th IEEE Annual Conference on Computational Complexity, 2004* (IEEE, Amherst, MA, 2004), p. 236.
- [160] A. W. Harrow, C. Y.-Y. Lin, and A. Montanaro, Sequential measurements, disturbance and property testing, in *Proceedings of the 2017 Annual ACM-SIAM Symposium on Discrete Algorithms* (Society for Industrial and Applied Mathematics, Barcelona, Spain, 2017), p. 1598.
- [161] S. Bravyi, D. Gosset, and R. König, Quantum advantage with shallow circuits, *Science* **362**, 308 (2018).
- [162] F. Centrone, N. Kumar, E. Diamanti, and I. Kerenidis, Experimental demonstration of quantum advantage for NP verification, arXiv:2007.15876 (2020).