


Leakage Reduces Device Coherence Demands for Pulse-Level Molecular Simulations

Ayush Asthana¹,² Chenxu Liu²,³ Oinam Romesh Meitei³,² Sophia E. Economou,² Edwin Barnes,² and Nicholas J. Mayhall^{1,*}

¹*Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

²*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*

³*Department of Chemistry, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

 (Received 2 May 2022; revised 27 February 2023; accepted 4 May 2023; published 26 June 2023)

Quantum simulation on noisy intermediate-scale quantum devices is severely limited by short qubit coherence times. A variational pulse-shaping algorithm known as ctrl-VQE was recently proposed to address this issue by eliminating the need for parameterized quantum circuits, which lead to long state preparation times. Here, we find the fastest possible pulses that prepare target molecular wave functions for a given device Hamiltonian describing coupled transmon qubits using simulations. We find that the resulting time-optimal pulses develop a bang-bang form consistent with Pontryagin’s maximum principle. We further investigate how the minimal state preparation time is impacted by the number of energy levels active in the transmon simulations. We find that leakage outside the computational subspace (something that is usually considered problematic) speeds up the state preparation, further reducing device coherence-time demands. Our analysis reveals that this speedup is due to both an enlarged solution space of target wave functions and the appearance of additional channels connecting initial and target states.

DOI: [10.1103/PhysRevApplied.19.064071](https://doi.org/10.1103/PhysRevApplied.19.064071)

I. INTRODUCTION

Noisy intermediate-scale quantum (NISQ) devices have qubits with low coherence times and frequent gate errors. This constrains today’s quantum algorithms to short circuit depths and has motivated the development of algorithms that minimize the quantum resources needed to obtain accurate results [1–7]. Despite rapid advances on both the algorithmic and hardware fronts, the path to useful quantum advantage remains elusive [8–15]. Thus, generating unitary transformations in a short time on a quantum device is of fundamental and practical importance in quantum computing.

Because most qubit platforms contain more than two levels, one must typically designate two of the states (i.e., $|0\rangle$ and $|1\rangle$) as the “computational basis” (the states to which the computational problem is mapped). Any uncontrolled population of states outside of the computational space (referred to as “leakage”) is generally harmful from a computational perspective. For example, superconducting transmon qubits are one of the leading approaches for realizing scalable [16] quantum devices with high fidelity and low noise in gates [17]. However, because the anharmonicity of transmon qubits is typically quite modest, higher-energy levels are easily populated. While these

higher-energy states can sometimes be leveraged as an additional quantum resource [18–26] if the quantum logic gates can be extended to explicitly act on the increased computational space, leakage is generally harmful and leads to difficulties in error mitigation and measurement [27–29]. To minimize leakage, one typically attempts to constrain the power of the drive or to use optimal control techniques [30,31] to implement fast quantum gates with high fidelity on transmon qubits [32,33].

On the algorithms side, variational quantum algorithms (VQAs) [34] are among the most promising approaches for solving a range of both optimization and simulation problems on NISQ-era devices due to their resilience to noise [35]. Although VQAs have been developed for applications, including quantum chemistry [12,14,36], optimization problems [37,38], quantum neural networks, [39–42], and many more, our current focus is on the electronic structure problem, which deals with finding the low-lying energy states of a molecule [12,12,36].

The electronic structure problem is the cornerstone of computational chemistry, which is believed to be one of the first applications in which quantum computers can demonstrate an advantage over classical computers. Molecular simulations on a quantum computer have already reproduced some known results for small molecules using variational quantum eigensolver (VQE) algorithms [1]. VQEs, a subset of VQAs, are classical-quantum algorithms that

*nmayhall@vt.edu

approximate the ground state of a many-body Hamiltonian by optimizing a parameterized quantum circuit that prepares a trial state wave function $[|\Psi(\theta)\rangle]$ to minimize the molecular energy. The optimal parameter values are given by

$$\theta^* = \arg \min_{\theta} (\langle \Psi(\theta) | \hat{H}_{\text{mol}} | \Psi(\theta) \rangle), \quad (1)$$

where \hat{H}_{mol} is the molecular Hamiltonian.

It has recently been recognized that VQAs can be both better understood and significantly improved by establishing connections with quantum optimal control theory (QOC) [43–47]. VQAs involve the optimization of parameters in discrete quantum circuit logic elements, whereas QOC functions more broadly to optimize a time-dependent drive Hamiltonian, typically by determining optimal pulse shapes of applied fields. Every parameterized quantum circuit corresponds to a pulse, but not all pulses have an associated parameterized quantum circuit. Although the two approaches share the same objective, to reach the final state that minimizes a cost function (e.g., molecular energy), QOC-based VQAs are inherently more powerful than circuit-based VQAs.

This connection is directly exploited by the recently developed ctrl-VQE algorithm [43], which defines a QOC-based ansatz at the device level, taking pulse parameters as optimization parameters and performing VQA optimization steps to minimize the molecular energy. Ctrl-VQE has been shown to reduce the coherence-time requirements for molecular simulations by several orders of magnitude, therefore showcasing the potential to simulate strongly correlated systems on NISQ-era quantum computers.

In addition to the reduction in coherence-time demands demonstrated in Ref. [43], connections between QOC and VQAs also provide critical conceptual insights regarding the complexity of numerical optimization. For controllable quantum systems with unconstrained quantum resources, the control landscape has been shown to be free from local minima [48]. The absence of local minima is important for the optimization process to reach the target state without getting stuck in traps. On the other hand, such traps often occur in parameterized gate-based VQAs unless more than a critical number of parameters are included (referred to as overparameterization) [49]. A conceptual understanding of numerical optimization is important for the development of scalable quantum computing algorithms.

The minimum evolution time (MET) required to rotate one quantum state to another has been studied using both constrained and unconstrained controls [50–52]. The field required to drive a MET transformation forms the “time-optimal” control, and the parameters for this field are the time-optimal control parameters. Coming from classical control theory, Pontryagin’s maximum principle has been applied to study optimal control for VQAs

with a continuous-time evolution, using an amplitude-constrained drive [45,53,54]. Yang *et al.* [45] found that the time-optimal controls exhibit a bang-bang form (control parameters saturate the bounds) due to the linear dependence on control waveforms in the drive Hamiltonian. Brady *et al.* [53] showed that more general optimal waveform shapes may exist (depending on the problem), as a key assumption in earlier work is not valid universally.

In this work, we show through numerical simulations that ctrl-VQE yields time-optimal control fields that prepare target molecular ground states in the shortest time possible for a given transmon device. These optimal pulses have a bang-bang form that is fully consistent with the Pontryagin maximum principle, as we confirm with analytical calculations. Using these time-optimal solutions, we investigate the effect of additional states outside the computational subspace, which naturally occur in transmon systems. We find that, surprisingly and contrary to typical scenarios, leakage outside of the computational space can significantly improve the performance of ctrl-VQE by further reducing the state preparation time, at the cost of additional measurements needed to renormalize the objective function. We show that this speedup can be attributed to two factors: (i) an expansion of the solution space of target wave functions, and (ii) the appearance of shorter quantum trajectories that are made accessible via the leakage states. These results constitute an important step toward overcoming coherence-time limitations of quantum simulation on NISQ devices.

This paper is organized in the following way. We first discuss the ctrl-VQE algorithm (Sec. II A) and time-optimal controls (Sec. II B). We then provide the computational details for the simulations in this paper in Sec. III. Results related to time-optimal controls in ctrl-VQE (Sec. IV A) and faster state evolution using leakage states (Sec. IV B) are discussed in Sec. IV. Finally, we present the summary and outlook in Sec. V.

II. ALGORITHM AND TIME-OPTIMAL CONTROLS

A. Ctrl-VQE

The ctrl-VQE approach uses a QOC-based wavefunction ansatz at the device level, instead of parameterized quantum circuits, for VQE simulations [43]. Because ctrl-VQE creates a device-specific ansatz built from parameterized control fields, the Hamiltonian of the device needs to be defined upfront. In our simulations, we use superconducting transmon qubits with constant, always-on, interqubit couplings. The device Hamiltonian is given by

$$\hat{H}_D = \sum_q \omega_q a_q^\dagger a_q - \sum_q \frac{\delta_q}{2} a_q^\dagger a_q^\dagger a_q a_q + \sum_{\langle pq \rangle} g_{pq} a_p^\dagger a_q, \quad (2)$$

where ω_q is the transition frequency of the $|0\rangle \rightarrow |1\rangle$ transition (resonance frequency) for qubit q , δ_q is the qubit anharmonicity, g_{pq} is the constant coupling rate for a pair of transmons, and the brackets in $\langle pq \rangle$ indicate a restricted sum over directly connected transmons. The control Hamiltonian that drives the system under the rotating-wave approximation is given by

$$\hat{H}_C(t) = \sum_q \Omega_q(t) (e^{i\nu_q t} a_q + e^{-i\nu_q t} a_q^\dagger), \quad (3)$$

where $\Omega_q(t)$ and ν_q are the time-dependent amplitude and (time-independent) frequency of the drive on qubit q . The state evolves according to the time-dependent Schrödinger equation in the interaction picture,

$$\frac{\partial}{\partial t} |\Psi_I(t)\rangle = -i\hat{H}_{I,C}(t) |\Psi_I(t)\rangle, \quad (4)$$

where $|\Psi_I(t)\rangle$ is the wave function, and the interaction picture control Hamiltonian is given by

$$\hat{H}_{I,C}(t) = e^{i\hat{H}_D t} \hat{H}_C(t) e^{-i\hat{H}_D t}. \quad (5)$$

To obtain the final state at time T , the time-dependent Schrödinger equation is solved, while the optimization is carried out on the drive amplitudes and frequency [$\Omega(t)_n$, ν_n for each qubit], which are varied to minimize the molecular energy given by

$$E_{\text{mol}} = \langle \Psi_I(T) | \hat{H}_{\text{mol}} | \Psi_I(T) \rangle, \quad (6)$$

where \hat{H}_{mol} is the molecular Hamiltonian mapped to the qubit system, and $|\Psi_I(T)\rangle$ is the wave function at the final evolution time.

However, because $|\Psi_I(T)\rangle$ could potentially populate higher-energy levels outside of the computational subspace (the first two levels $|0\rangle$ and $|1\rangle$ on each qubit), we first project onto the computational subspace and then renormalize. By discarding any measurements from leakage states, we end up with a measured expectation value for the wave function projected onto the computational basis, $|\tilde{\Psi}(T)\rangle$. The number of discarded shots gives us a measure of the norm of the wave function discarded, allowing us to recover the variational principle:

$$E_{\text{mol}} = \frac{\langle \tilde{\Psi}_I(T) | \hat{H}_{\text{mol}} | \tilde{\Psi}_I(T) \rangle}{\langle \tilde{\Psi}_I(T) | \tilde{\Psi}_I(T) \rangle} \quad (7)$$

with T the total time of evolution that is kept fixed during a simulation. The ctrl-VQE ansatz can be summarized as

$$|\Psi_I^{\text{trial}}(\Omega_n(t), \nu_n)\rangle = \hat{T} e^{-i \int_{t=0}^T dt \hat{H}_{I,C}(t, \Omega(t)_n, \nu_n)} |\psi_I(0)\rangle, \quad (8)$$

where \hat{T} is the time-ordering operator, and $|\psi_I(0)\rangle$ is the initial state, usually taken as the Hartree-Fock (HF)

wave function. This trial wave function is evaluated at each iteration, and the optimization parameters are varied to minimize the molecular energy. In what follows, we suppress subscript I on the states for notational simplicity.

B. Minimum evolution time and time-optimal controls

Control problems aim to find a target state that minimizes the cost function, which in the case of molecular simulation is the molecular energy [Eq. (6)]. In the qualitative depiction in Fig. 1, the target state is represented by $|\Psi(T)\rangle$. Starting from the initial state, $|\Psi(0)\rangle$, the state evolves in time through the action of driving pulses until it reaches the target state. There exists a minimum pulse duration (the MET) that is required to carry out this state evolution. Control pulses with durations that exceed the MET ($T > T^*$ in Fig. 1, where T^* is the MET for qubits) can drive the system along one of the infinitely many possible trajectories in Hilbert space to reach the target state, while time-optimal pulses drive the initial state to the target state along special trajectories in a time equal to the MET, which depends on the constraints imposed on the driving pulses. Ctrl-VQE generates controllable evolutions in the control theory sense, as the entire Hilbert space can be spanned to arbitrary precision by varying unconstrained parameters of the driving pulses. In ctrl-VQE, as the total pulse duration (T) increases, the accessible (or reachable) Hilbert space increases as well.

In the present study, we numerically find T^* and the associated time-optimal controls by starting at a large enough pulse duration and reducing the pulse duration until we find the shortest pulse at which the system can reach the target state to the desired accuracy. The pulse




FIG. 1. Qualitative depiction of state evolution in Hilbert space of qubits with two or more levels. The time-optimal path in the computational subspace (yellow surface) is shown as the path with evolution time $T = T^*$. Many paths are possible to reach the solution $|\Psi(T)\rangle$ when $T > T^*$ (gray curly line). Access to leakage states provides new paths (shown by darker red paths) that reach the solution faster than T^* when the final state is projected onto the computational subspace.

is divided into many stepwise constant segments with practical bounds on amplitudes and frequencies. Dividing the pulse into many adjustable segments provides a large amount of freedom in shaping the pulse, essentially allowing the ctrl-VQE algorithm to choose any possible optimal pulse shape in the optimization process. Because the numerical optimization of time-optimal controls is quite challenging, it is difficult to ensure that one has converged to T^* . As a guarantee, in what follows we compare our numerical solutions to analytical results obtained from Pontryagin's theory to confirm that our numerically obtained METs are correct.

III. COMPUTATIONAL DETAILS

The calculations are carried out using a local version of the ctrlq code [43,55]. As a test system for our simulations, we have chosen the problem of finding the ground-state energy of the H_2 molecule at a bond distance of 1.5 Å, which has relatively strong correlations. Molecular integrals are generated using PySCF [56] and the STO-3G basis set is used in this work. We use the parity mapping with Z2 symmetry reduction to map the molecular Hamiltonian on the qubit basis. The initial state is taken as the unentangled HF state for this problem, and the target state is the exact ground state [i.e., full configuration interaction (FCI)]. The numerical optimization of pulse amplitudes and frequencies is carried out using the limited-memory Broyden-Fletcher-Goldfarb-Shanno bound-constrained optimization algorithm, and the pulses are taken to be piecewise constant with 100 segments (unless otherwise stated). The time-discretized amplitude, $\Omega(t)$, is given by

$$\Omega_k(t) = \begin{cases} c_1, & 0 < t \leq t_1, \\ c_2, & t_1 < t \leq t_2, \\ \vdots & \\ c_n, & t_{n-1} < t \leq T, \end{cases} \quad (9)$$

for qubit k , where c_1, c_2, \dots, c_n represent amplitude parameters for each pulse segment, and T is the final time. The pulse amplitude is constrained between ± 20 MHz unless otherwise mentioned, which is based on common values for control signals in transmon qubits [57]. The carrier signal frequency ν_k given in Eq. (3) is constrained such that its difference from the qubit frequency, $\nu_k - \omega_k$, is in the range ± 1 GHz, where ω_k is the qubit frequency. A convergence error threshold of 10^{-8} hartree in molecular energy is used in the optimization process, which corresponds to close to unit fidelity with the target FCI wave function. Since we encounter many local minima when the pulse duration is close to its minimum value, we consider MET to be the minimum pulse duration at which a solution could be reached in 1000 random initializations. The results are minimally affected by increasing the

TABLE I. Device parameters used in Eq. (3) for a two-transmon system. All values are in gigahertz.

	Transmon	
	1	2
$\omega/(2\pi)$	4.8080	4.8333
$\delta/(2\pi)$	0.3102	0.2916
	1 \rightarrow 2	
$g/(2\pi)$	0.018 31	

number of pulse segments or Trotter steps (for numerical tests, see the Supplementary Material [58]). The parameters in the device Hamiltonian, ω_q , η_q , and g_{pq} , used in the simulations are given in Table I.

IV. RESULTS AND DISCUSSION

In the following subsections, we discuss our main results related to time-optimal controls in the ctrl-VQE algorithm applied to transmon qubits. We first discuss the time-optimal pulse shape for ctrl-VQE through the application of Pontryagin's maximum principle and numerical simulations. Next, we qualitatively analyze the near-MET evolution of the two-qudit state, with a particular focus on the role of leakage to higher transmon levels outside the computational subspace.

A. Time-optimal control in ctrl-VQE

Time-optimal controls for quantum algorithms have been previously studied using quantum optimal control theory with continuous-time controls [45,53,54,59]. Pontryagin's principle is applied to derive conditions under which an associated cost function is minimized at the MET. We follow the same procedure and use Pontryagin's principle to find conditions for time-optimal controls for the ctrl-VQE algorithm. The dynamics of the multi-qudit state in ctrl-VQE is governed by the time-dependent Schrödinger equation given in Eqs. (4) and (5). The cost function in the case of the electronic structure problem takes the form of Eq. (6). This can be formulated into a time-optimal control problem, with the modified cost function to be minimized as

$$J = \langle \Psi(T) | \hat{H}_{\text{mol}} | \Psi(T) \rangle + \left(\int_0^T \langle \lambda(t) | (-i\hat{H}_{I,C}(t) | \Psi(t) \rangle - |\dot{\Psi}(t)\rangle) dt + \text{H.c} \right) + \eta T. \quad (10)$$

This corresponds to minimizing the molecular energy at the final time T , under the constraint that the wave function is obtained as a proper solution to the device Schrödinger equation. Here, $\langle \lambda(t) |$ and $|\lambda(t)\rangle$ are treated as Lagrange

multipliers. The control function [60] may be defined following Pontryagin's maximum principle [53,54,61] and Eq. (4), taking the form

$$\mathbb{H} = \langle \lambda(t) | i\hat{H}_{I,C}(t) | \Psi(t) \rangle - \langle \Psi(t) | i\hat{H}_{I,C}(t) | \lambda(t) \rangle \quad (11)$$

with the time evolution of state $|\Psi(t)\rangle$ and co-state $|\lambda(t)\rangle$, and the end point of co-state $|\lambda(T)\rangle$ given by

$$|\dot{\Psi}(t)\rangle = -i\hat{H}_{I,C}|\Psi(t)\rangle, \quad (12)$$

$$|\dot{\lambda}(t)\rangle = -i\hat{H}_{I,C}|\lambda(t)\rangle, \quad (13)$$

$$|\lambda(T)\rangle = -\hat{H}_{\text{mol}}|\Psi(T)\rangle. \quad (14)$$

The derivation of the control function \mathbb{H} and details of the application of Pontryagin's principle are given in the Supplementary Material [58]. At the optimal control $[\Omega^*(t), \Psi^*(t), \lambda^*(t)]$, where the asterisk denotes the optimal path, Pontryagin's maximum principle states that

$$\mathbb{H}(t, \Omega_n^*(t), \Psi^*(t), \lambda^*(t)) \geq \mathbb{H}(t, \Omega_n(t), \Psi^*(t), \lambda^*(t)). \quad (15)$$

This means that the modified cost function \mathbb{H} should be maximized at each point in time along the evolution. Note that the \mathbb{H} formed in ctrl-VQE is an explicit function of time. The modified cost function \mathbb{H} may be written as a sum of \mathbb{H} for each qubit as

$$\mathbb{H} = \sum_q \mathbb{H}_q, \quad (16)$$

where q is the qubit index. For ctrl-VQE, the control function for each qubit may be written as (for a derivation, see the Supplementary Material [58])

$$\begin{aligned} \mathbb{H}_q &= 2\text{Re}\langle \lambda(t) | i\Omega_q(t) e^{i\hat{H}_D t} (e^{iv_q t} a_q + e^{-iv_q t} a_q^\dagger) e^{-i\hat{H}_D t} | \Psi(t) \rangle \\ &= \Omega_q(t) \phi_q(t), \end{aligned} \quad (17)$$

where the ‘‘switching function’’ $\phi_q(t)$ is given by

$$\phi_q(t) = 2\text{Re}\langle \lambda(t) | i e^{i\hat{H}_D t} (e^{iv_q t} a_q + e^{-iv_q t} a_q^\dagger) e^{-i\hat{H}_D t} | \Psi(t) \rangle. \quad (18)$$

Through Eqs. (15) to (18), it can be shown that the optimal controls $\Omega_q^*(t)$ obey the condition

$$\Omega_q^*(t) = \begin{cases} \text{max bound} & \text{if } \phi_q(t) > 0, \\ \text{min bound} & \text{if } \phi_q(t) < 0, \\ \text{undefined} & \text{if } \phi_q(t) = 0. \end{cases} \quad (19)$$

The result of this is that optimal controls follow a ‘‘bang-bang’’ structure when the switching function has a nonzero value; specifically, the controls are maximized at time t

in the case of positive $\phi_q(t)$ and vice versa. The optimal solution is undefined in the case of a singular switching function $\phi_q(t) = 0$ as the control function becomes independent of control parameters in that case. This condition, $\phi_q(t) = 0$, is well studied and is known as the ‘‘singularity’’ condition [61]. For a single-qubit system, Eq. (18) is zero for an extended period of time only when $\lambda(t) = \Psi(t)$ for an extended period. This is not possible in the electronic structure problem due to Eq. (14); therefore, the optimal protocol will follow a bang-bang form for a single qubit. To study the case of two qubits, as in the case of the H_2 molecular calculation, we investigate the optimal pulse shape by carrying out ctrl-VQE transmon simulations and limiting the pulse duration to the requisite MET.

In agreement with the analytic results, our simulations reveal the emergence of bang-bang pulse controls for time-optimal simulations of the H_2 ground state at a bond length of 1.5 Å. Figure 2 presents the pulse shapes obtained after optimization at different pulse durations for both two-level qubits [Fig. 2(a)] and three-level qutrits [Fig. 2(b)]. Similar to the sketch depiction in Fig. 1, there are many ways of reaching the final state when the evolution time exceeds the MET (i.e., at $T > T^*$). In the top two panels, we present one of the many solutions found, noting that both qubits and qutrits exhibit a highly unstructured pulse shape. In the middle two panels (16 and 10 ns, respectively), both qubits and qutrits begin to develop a bang-bang-like structure. As the pulse duration is further reduced, the pulse shape becomes increasingly more uniform until it reaches a well-organized bang-bang shape near METs of 15 and 8.94 ns for qubits and qutrits, respectively.

As the evolution time is decreased (such that we reach our VQE trial state faster), the computational task of optimizing a pulse becomes increasingly difficult due to the quickly rising number of local parameter traps. This makes it difficult to confirm that one has actually converged to T^* . In order to verify that we are near T^* , in Fig. 3, we compare the switching function from Eq. (18), $\phi_q(t)$, with the optimized pulses to check the agreement with results obtained using Pontryagin's maximum principle. By plotting the function $\phi_q(t)$ [62] atop the optimized pulses in Fig. 3, we can readily see that the pulse switching times coincide with the roots of the switching function and that at each time, the pulse is maximized in the direction determined by the sign of $\phi_q(t)$, in agreement with Pontryagin's maximum principle, i.e., when $\phi_q(t)$ is positive (negative), $\Omega_q(t)$ should assume the largest positive (negative) value allowed by the constraints. Since we were not able to numerically find pulses that would reach the desired solution and are shorter than the discussed pulses, along with the fact that these optimized pulses agree with results from Pontryagin's principle indicate that these are indeed time-optimal pulses.

For the calculation of $\phi_q(t)$, we first carry out a ctrl-VQE simulation using time-optimal controls to evaluate $|\Psi(t)\rangle$

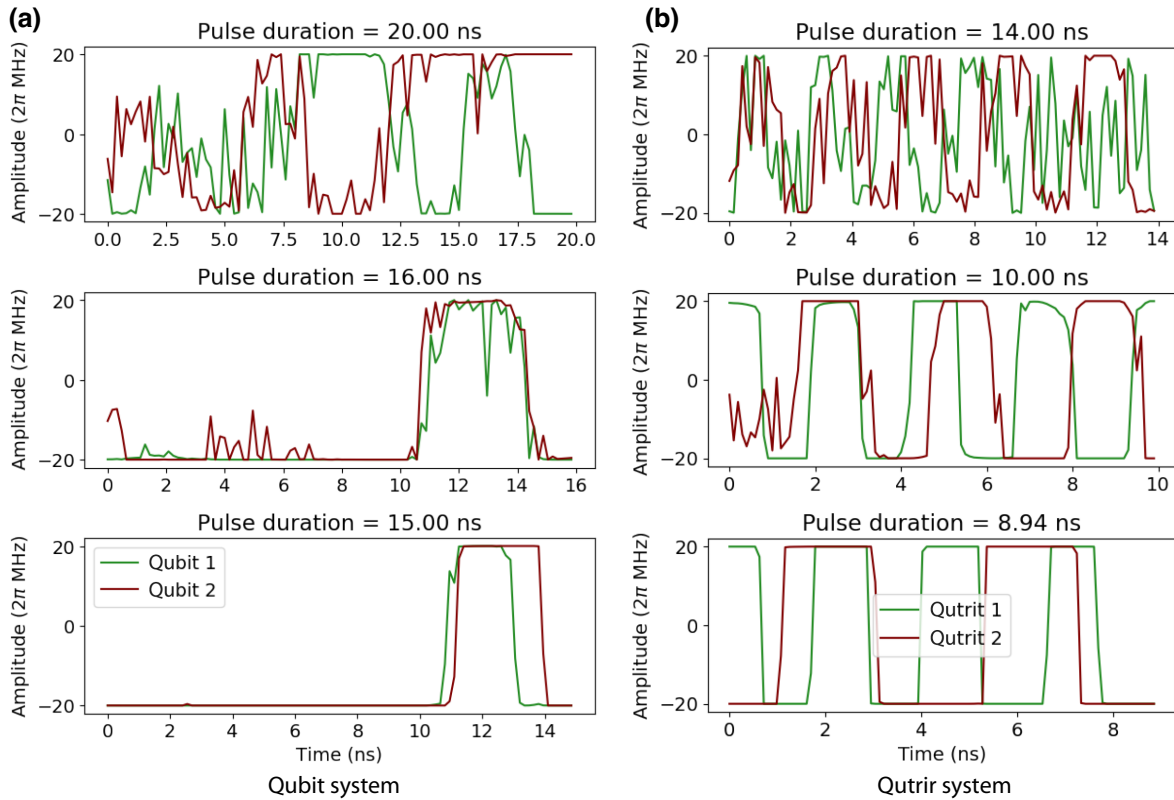


FIG. 2. (a) How the optimized pulse shape changes for a two-qubit system as the pulse duration is reduced from 20.00 to 15.00 ns. (b) How the optimized pulse shape changes for a two-qutrit system as the pulse duration is reduced from 14.00 to 8.94 ns. The H_2 ground-state energy is calculated in these simulations at a bond distance of 1.5 Å. The pulse shape evolves into bang-bang form as the pulse duration is restricted towards the MET needed to produce the ground state.

at each time step. We then form $|\lambda(T)\rangle$, and backpropagate it using Eqs. (13) and (14). Finally, we evaluate $\phi_q(t)$ through Eq. (18). It can be observed that there is no region in time where the function $\phi_q(t)$ is 0 for a continuous period of time. Furthermore, the times at which $\phi_q(t)$ vanishes coincide with the times at which the pulse changes sign, consistent with the analytical prediction of Eq. (19). This can also be understood from the fact that $\phi_q(t)$ is proportional to the analytical derivative derived in Ref. [43].

B. Minimum evolution time for a system of transmon qutrits

1. Leakage speeds up evolution

The ctrl-VQE optimized pulses for H_2 shown in Fig. 2 demonstrate that the ground state is obtained significantly faster when higher transmon levels are available. We find that the MET for a system of qubits for this problem is 15.00 ns while the MET for a system of qutrits is 8.94 ns. This amounts to about a 40% reduction in evolution time when a qutrit system is used. Considering that short coherence times are one of the most restrictive bottlenecks on NISQ devices, this reduction in state preparation

time constitutes a significant practical advantage of this approach.

Of course, this decreased evolution time is not entirely free of cost. By populating higher-energy states (as depicted in Fig. 1), the final molecular energy can only be computed by first projecting onto the computational space and renormalizing the expectation value. However, assuming one can resolve the difference between computational states and leakage states in a measurement, one can simply postselect, discarding measurements of leakage states, but keeping track of how often such measurement outcomes are obtained. The percentage of leakage state measurements then directly informs one how to normalize the objective function. With increasing weight in the leakage space, the number of discarded measurements increases, which requires one to perform an increasing number of shots. However, this increased shot-count overhead is directly determined by the extent of the leakage, which is also partially controlled by the amplitude constraints on the pulse.

In addition to the decreased time of evolution, access to leakage states also improves the ability to optimize the pulses themselves. Figure 4 shows the probability of finding the target solution plotted against pulse duration for

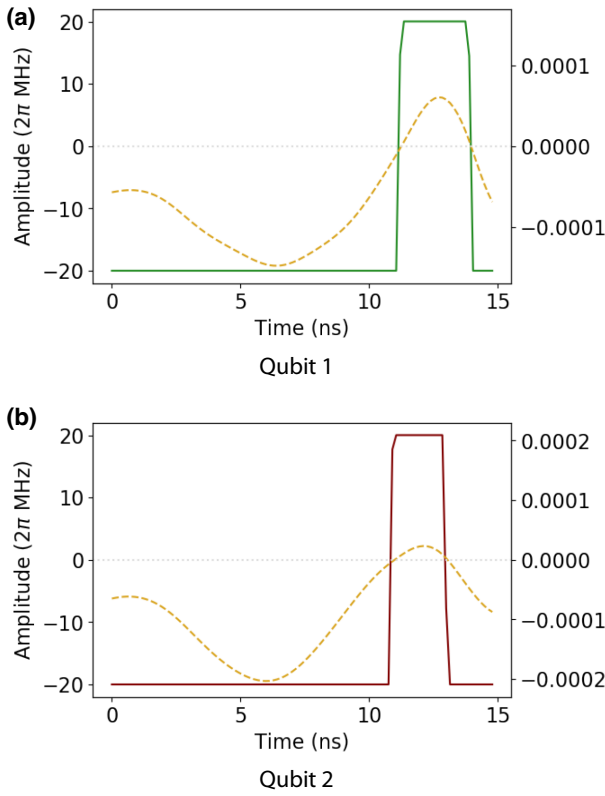


FIG. 3. (a) The switching function $\phi_1(t)$ (yellow dashed) and pulse shape $\Omega_1^*(t)$ (green) for the optimal drive on qubit 1 in two-level transmon qubits. (b) The switching function $\phi_2(t)$ (yellow dashed) and pulse shape $\Omega_2^*(t)$ (red) for the optimal drive on qubit 2 in two-level transmon qubits. The plots are made using 100 pulse segments, 1000 Trotter steps in the simulation, and a pulse duration of $T = 14.93$ ns, reaching an accuracy of 10^{-6} hartree relative to the FCI energy.

transmons with two, three, four, and five levels retained. The probability is determined by using 100 random initializations and counting how many of them succeed in preparing the ground state of the H_2 molecule (to within 10^{-8} hartree). The results show that the probability to reach the target state drops abruptly at about 9 ns for qutrits and at about 19 ns for qubits. It is evident that the computational difficulty to reach the target state in the case of qubits is significantly higher, and it takes more optimization steps for the simulation to converge in the case of qubits as compared to qutrits (increase in optimization steps near the MET was also observed in Ref. [52]). This indicates that a computational phase transition occurs at the MET. The speedup obtained using qutrits is not further enhanced by using more than three levels in our two-transmon calculations. This is likely due to the fact that the higher levels are further detuned from the microwave drive, and are thus not easily populated. Therefore, the contribution of these states to preparing the final state is negligible compared to that of the third transmon level.




FIG. 4. Success probability of the simulation to reach the target state versus pulse duration. The probability is computed using 100 random initializations and with an energy error of less than 10^{-8} hartree compared to the FCI solution. The plot shows that transmons truncated to two levels require a longer evolution time as compared to transmons with more than two levels.

We find that the leakage at pulse durations close to the MET for qutrits is about 50%, indicating a necessary twofold increase in the number of shots required. We have further seen that this leakage can be optimized to a desired fraction of the total population by introducing a penalty term in our cost function to penalize leakage over the desired amount. For instance, using a 0.01 hartree penalty per 1% leakage for any amount of leakage over 10% of the total population (such that a leakage of 11% will result in an energy penalty of 0.01 hartree, while any amount of leakage below 10% will have no energy penalty) reduces the MET from 15 ns to about 12.5 ns. So a penalty can be introduced in the cost function to keep the leakage below a target threshold, while still taking advantage of faster state evolution. However, we find that a penalty term that reduces the leakage to zero does not result in any speedup in our simulations (despite not penalizing leakage at earlier times, $t < T$). The no-leakage case is the same as performing the simulation without projecting onto the computational space (the space spanned by the $|0\rangle$ and $|1\rangle$ levels of the qudits), which has also not resulted in a faster state evolution in our simulations. This indicates that some population in the higher-lying qudit levels is necessary at the final time to achieve this faster state evolution, evidence that an enlarged solution space is at least partly responsible for the decreased MET. Another important observation is that most of the leakage lies in specific states, namely, $|02\rangle$ and $|20\rangle$ in the present simulations.

It is also worth noting that, for pulse durations T significantly greater than the minimum time T^* , we do not observe any traps (getting stuck in local minima) for either qubits or qutrits. This can be understood through results in quantum optimal control that state that a controllable system is free from traps unless the drive is severely constrained [52]. Once the drive has a duration significantly

greater than the minimum duration needed for the state evolution, the system reaches the target state without getting hindered by traps. Local traps are also often encountered in gate-based VQEs unless a minimum number of parameters is included in the circuit (overparameterization). A major benefit of the reduced MET is that it allows this trap-free (analogous to overparameterization in the gate-based VQE) region to be reached with shorter pulses in ctrl-VQE.

2. Mechanism of evolution speedup in qutrits

We investigate the mechanism of the observed speedup using a numerical analysis based on time-dependent perturbation theory to determine which processes contribute the most toward reaching the target wave function in the case of qubits and qutrits. In the H_2 molecular problem, the initial state is the HF state, $|01\rangle$. The FCI target state is a linear combination of $|01\rangle$ and $|10\rangle$, where the population in the former is 6.92 times that of the latter. Thus, some population is transferred from $|01\rangle$ to $|10\rangle$ during the evolution. The population of each state as a function of time is shown in Fig. 5, both for transmon qubits (a) and qutrits (b). It is evident that the evolution is very different between the two cases. In the case of qubits, states $|00\rangle$ and $|11\rangle$ get populated during the evolution and depopulate completely by the end of the evolution. These two states assist with the transfer of population from $|01\rangle$ to $|10\rangle$. In the case of qutrits, states $|00\rangle$ and $|11\rangle$ receive very little population throughout the evolution, while states $|02\rangle$ and $|20\rangle$ exhibit a significant population buildup that remains at the end of the evolution. This striking disparity indicates that higher transmon levels provide alternate channels for reaching the target state. To determine the exact transition channels at work, we analyze the evolution using time-dependent perturbation theory. The unitary evolution operator $U(0, T)$ that carries out the state evolution, $|\Psi(T)\rangle = U(0, T)|\Psi(0)\rangle$, can be expanded in a Dyson series:

$$\begin{aligned}
 U(0, t_f) = & I + (-i) \int_0^{t_f} dt H_{I,C}(t) \\
 & + (-i)^2 \int_0^{t_f} dt_1 \int_0^{t_1} dt_2 H_{I,C}(t_1) H_{I,C}(t_2) + \dots
 \end{aligned}
 \quad (20)$$

We consider the transition from the initial state $|01\rangle$ to the final state $|10\rangle$. The zeroth-order transition amplitude is zero, $\langle 01|10\rangle = 0$, while the first- and second-order transition amplitudes from state $|01\rangle$ to $|10\rangle$ are

$$A_{01 \rightarrow 10}^{(1)} = \int_0^{t_f} \langle 10|(-i)H_{I,C}(t)|01\rangle dt, \quad (21)$$

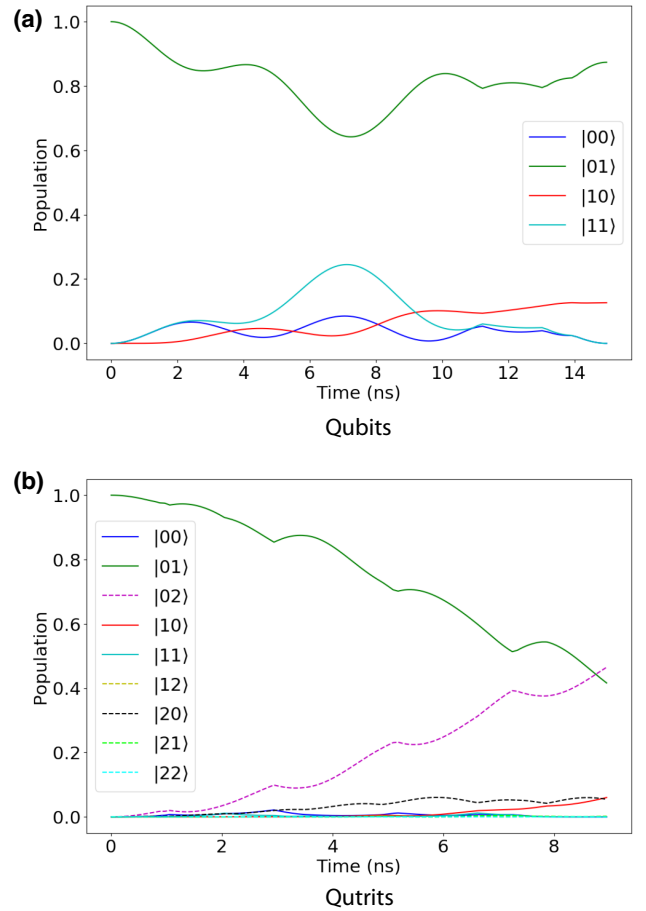


FIG. 5. Population in various basis states versus evolution time for (a) qubits and (b) qutrits. The solid lines are states in the computational space, while the dashed lines are states that involve the extra transmon level. The initial state is the HF state, while the final state is the FCI wave function of H_2 at a bond distance of 1.5 Å.

$$A_{01 \rightarrow 10}^{(2)} = \int_0^{t_f} dt_1 \int_0^{t_1} dt_2 \langle 10|(-i)^2 H_{I,C}(t_1) H_{I,C}(t_2)|01\rangle. \quad (22)$$

The second-order transition amplitude can be rewritten as a sum over transitions through intermediate states:

$$\begin{aligned}
 A_{01 \rightarrow 10}^{(2)} = & \sum_m \int_0^{t_f} dt_1 \int_0^{t_1} dt_2 \langle 10|(-i)H_{I,C}(t_1)|m\rangle \\
 & \times \langle m|(-i)H_{I,C}(t_2)|10\rangle \\
 = & \sum_m A_{01 \rightarrow m \rightarrow 10}
 \end{aligned}
 \quad (23)$$

with the summation taken over all the eigenstates of the transmons. We have checked that evolving the wave function using the second-order evolution operator $U^{(2)}(0, T)$ yields the target state up to an infidelity of about 0.01, and hence second order is sufficient for a qualitative study of the state evolution. The first-order transition probability,

$A_{01 \rightarrow 10}^{(1)} = 0$, vanishes for both qubits and qutrits, as the static coupling between transmon qubits is treated as a part of the device Hamiltonian, which is diagonal in the computational basis. Thus, the transition from initial to final state is entirely through higher-order processes.

We show the second-order transition probability from state $|01\rangle$ to $|10\rangle$ for two transmon qubits and two transmon qutrits in Figs. 6(a) and 6(b), respectively. The second-order transition probability is defined as

$$P_{01 \rightarrow 10} \equiv |A_{01 \rightarrow 10}^{(2)}|^2 = \left| \sum_m A_{01 \rightarrow m \rightarrow 10}^{(2)} \right|^2. \quad (24)$$

To investigate which transition path is most important in the time-optimal evolution, we compare the contributions to the transition probability from two dominant paths. In the qubit case, the paths are those that go through states $|00\rangle$ and $|11\rangle$ [see Fig. 6(a), orange line], while for the qutrit case, the dominant paths are those through $|02\rangle$ and $|20\rangle$ [see Fig. 6(b), orange line]. Note that in both

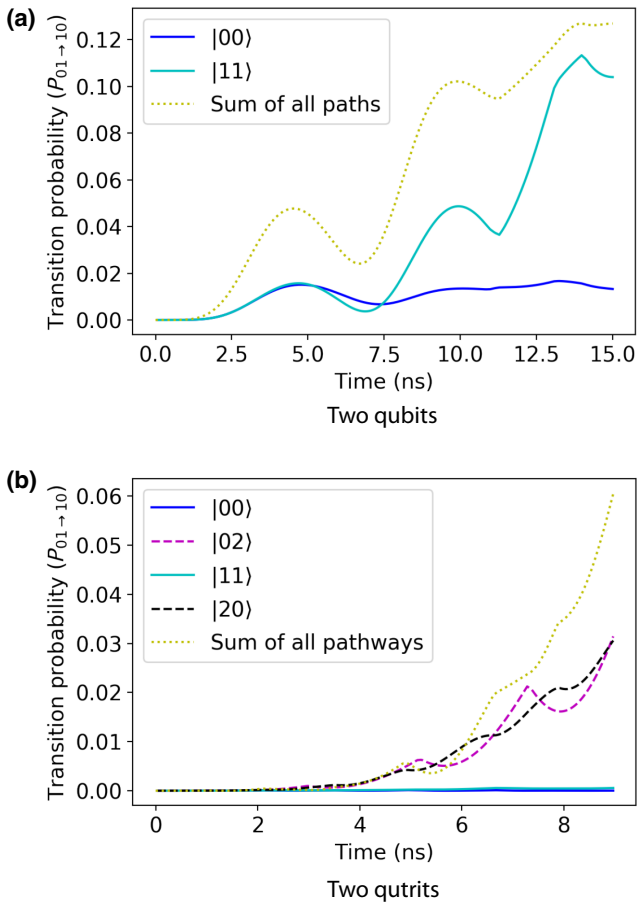


FIG. 6. Transition probability from state $|01\rangle$ to $|10\rangle$ through various intermediate states $[m$ in Eq. (23)] for (a) two qubits and (b) two qutrits. The dotted yellow line shows the sum of contributions from all channels to state $|10\rangle$.

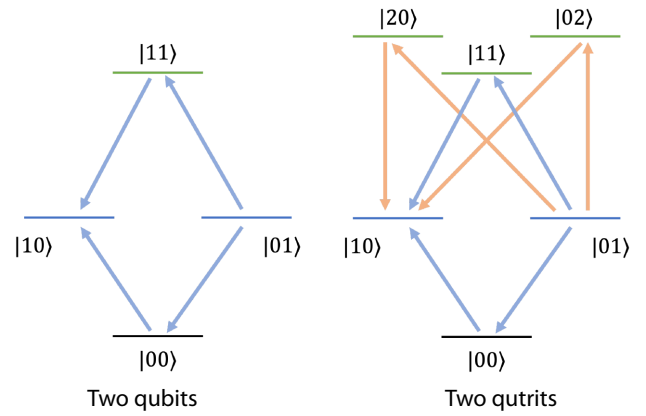


FIG. 7. Major channels involved at second order in evolution from state $|01\rangle$ to $|10\rangle$ in the case of (a) two qubits and (b) two qutrits. The orange channels are the new channels that arise in the qutrit case.

cases, the truncated transition probability is close to the second-order one, which shows that the transition is predominantly happening through those channels. Thus, the higher-lying leakage states are participating and providing alternate channels for state evolution. This can be visualized with the help of the schematic level diagram shown in Fig. 7. Basis states $|02\rangle$ and $|20\rangle$ do not exist in qubits, while this channel becomes available in qutrits. Furthermore, the population of basis states $|00\rangle$ and $|11\rangle$ is zero in the target state. This puts a constraint on the dynamics of qubits as it has to depopulate these basis states at the end of the evolution, while these basis states do not have any population in the case of qutrits. Furthermore, in Figs. 8(a) and 8(b), we compare the truncated transition probability with the squares of the individual transition amplitudes. If the two transition paths constructively interfere with each other, we expect

$$|A_{01 \rightarrow m_1 \rightarrow 10} + A_{01 \rightarrow m_2 \rightarrow 10}|^2 > |A_{01 \rightarrow m_1 \rightarrow 10}|^2 + |A_{01 \rightarrow m_2 \rightarrow 10}|^2. \quad (25)$$

This is observed in the qubit case, while it is not true for the qutrit case (see Fig. 8).

Finally, an additional contribution to the speedup evident in the qutrit case comes from the classical step of projecting the final qutrit state onto the computational subspace and normalizing it. This can be seen from Fig. 9, where the population in state $|10\rangle$ is plotted as a function of the evolution time in the case of qubits and qutrits (unnormalized and normalized cases). It can be seen that the population in qutrits increases faster than qubits and reaches the target population faster when the final state on qutrits is projected onto the computational basis and normalized. This is because the population required by the target state in the computational subspace is smaller when we finally normalize the state. The presence of additional

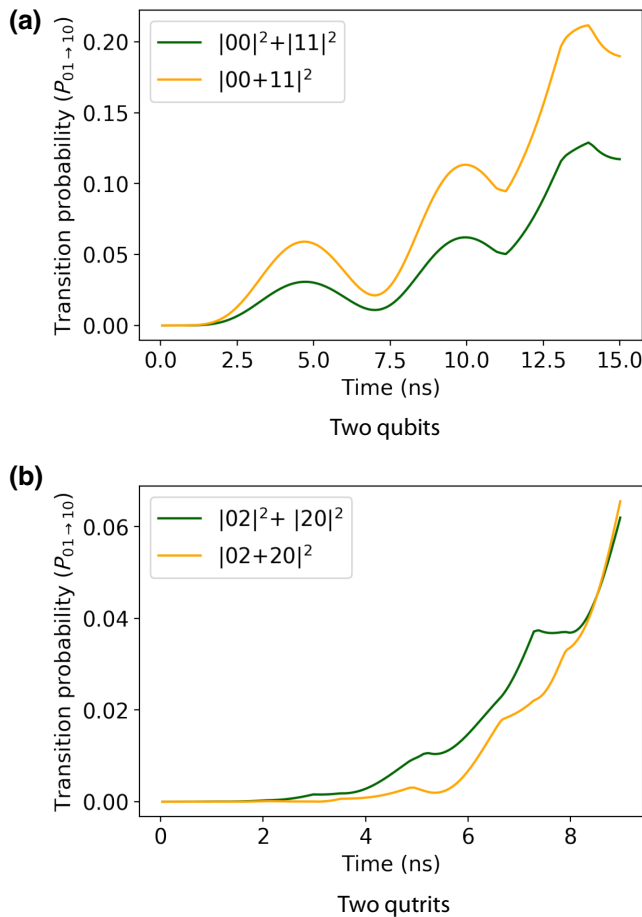


FIG. 8. Role of constructive interference between the main transition channels in the case of (a) two qubits and (b) two qutrits. The constructive interference (difference between the orange and green lines) is more significant in the qubit case.

channels for state evolution, along with reduced population demand in the desired target state in the computational subspace, leads to a faster state evolution in qutrits compared to qubits.

Although we have focused on the electronic structure problem in the case of the H_2 molecule ground state in our current work, the key results are expected to generalize. This is because any problem can be mapped to the task of finding trajectories that connect an initial and target state in Hilbert space. Changing the problem would correspond to changing these initial and target states. Including higher levels will generally provide new channels for state evolution that provide shorter paths to the target state.

It should be noted that Ref. [26] looks somewhat similar in the fact that leakage is shown to reduce device coherence time requirements in that paper as well. But the effect is coming from an entirely different physical idea in the two papers. The faster state propagation in Ref. [26] is coming from a higher bosonic amplification factor, whereas the




FIG. 9. Population buildup in state $|10\rangle$ versus evolution time in the case of two qubits (2 level) and two qutrits (3 level). The basis state $|10\rangle$ reaches the target population faster in the case of qutrits when the final state is projected onto the computational subspace and normalized (3-level normalized).

effect in the current manuscript is primarily by a simplification of the cost function provided by leakage. This leads to a much milder condition to reach the effect in this paper compared to what is proposed in Ref. [26] as it requires strong pulses to pump qubit states to higher levels.

V. SUMMARY

In this paper, we present time-optimal control fields that prepare target molecular ground states on transmon quantum processors. These optimal controls are obtained using the pulse-level VQE algorithm known as ctrl-VQE simulations. We find that the optimal controls converge to a bang-bang form, i.e., they saturate pulse amplitude constraints, as the total evolution time is reduced. We show that this behavior is consistent with analytical results derived from Pontryagin's maximum principle. These results suggest that ctrl-VQE prepares target states on quantum computers in the shortest possible time allowed by quantum speed limits.

In addition, we investigate how the minimal state preparation time is impacted by truncating the transmons to two or more levels using ctrl-VQE simulations. We find that the inclusion of higher transmon levels can lead to a substantial reduction in the time needed to reach the target state when the final state is projected on the computational subspace. We show that this speedup is the result of two effects caused by the additional levels: (i) an enlargement of the space of wave functions that are parallel to the target state in the computational subspace, and (ii) the appearance of new transition paths that connect the initial and target states. Together, these effects cut down the evolution time by about 40% for the simulations performed in this work.

Given the limited coherence times in NISQ devices, it is crucial that we minimize the time it takes to prepare target wave functions. Our results indicate that ctrl-VQE does this, by searching over a wider range of unitaries, not being restricted to those that can be decomposed into a circuit of native gates. Our work also shows that leakage to higher transmon levels is not always problematic. It can potentially be used to further reduce the coherence time requirements needed for VQEs. Another important benefit of faster state evolution is that the overparameterized region, where the VQE algorithms become free from local minima (trap-free), is reached faster in qudits as compared with qubits. Local traps have been shown to prevent gate-based VQE algorithms from reaching the solution to the desired accuracy unless a significantly large gate depth is achieved [49].

ACKNOWLEDGMENTS

This research is supported by the Department of Energy (DOE). E.B. and N.J.M. acknowledge Award No. DE-SC0019199, and S.E.E. acknowledges the DOE Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA), Contract No. DE-SC0012704. A.A. would like to thank Dr. Luke Bertels, Dr. Sheril Avikkal Kunhippurayil, and Dr. Pablo Poggi for helpful discussions. The authors thank the Advanced Research Computing (ARC) facility at Virginia Tech for the computational infrastructure.

-
- [1] 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).
 - [2] W. J. Huggins, J. Lee, U. Baek, B. O’Gorman, and K. B. Whaley, A non-orthogonal variational quantum eigensolver, *New J. Phys.* **22**, 073009 (2020).
 - [3] 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**, 1 (2019).
 - [4] H. L. Tang, V. 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, *PRX Quantum* **2**, 020310 (2021).
 - [5] J. Lee, W. J. Huggins, M. Head-Gordon, and K. B. Whaley, Generalized unitary coupled cluster wave functions for quantum computation, *J. Chem. Theory Comput.* **15**, 311 (2018).
 - [6] I. G. Ryabinkin, T.-C. Yen, S. N. Genin, and A. F. Izmaylov, Qubit coupled cluster method: A systematic approach to quantum chemistry on a quantum computer, *J. Chem. Theory Comput.* **14**, 6317 (2018).
 - [7] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, and E. Barnes, Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm, *npj Quantum Inf.* **6**, 1 (2020).
 - [8] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, and D. A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
 - [9] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, and Y. Hu, *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
 - [10] D. Maslov, J.-S. Kim, S. Bravyi, T. J. Yoder, and S. Sheldon, Quantum advantage for computations with limited space, *Nat. Phys.* **17**, 894 (2021).
 - [11] Y. Nam and D. Maslov, Low-cost quantum circuits for classically intractable instances of the Hamiltonian dynamics simulation problem, *npj Quantum Inf.* **5**, 1 (2019).
 - [12] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, and L. Cincio, *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
 - [13] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, and N. P. Sawaya, *et al.*, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
 - [14] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, and G. H. Booth, *et al.*, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
 - [15] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, VQE method: A short survey and recent developments, *Mater. Theory* **6**, 1 (2022).
 - [16] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, and F. Jelezko, *et al.*, The quantum technologies roadmap: A European community view, *New J. Phys.* **20**, 080201 (2018).
 - [17] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
 - [18] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, E. Lucero, A. D. O’Connell, D. Sank, H. Wang, J. Wenner, and A. N. Cleland, *et al.*, Emulation of a quantum spin with a superconducting phase qudit, *Science* **325**, 722 (2009).
 - [19] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O’Brien, A. Gilchrist, and A. G. White, Simplifying quantum logic using higher-dimensional Hilbert spaces, *Nat. Phys.* **5**, 134 (2009).
 - [20] H. Bechmann-Pasquinucci and A. Peres, Quantum Cryptography with 3-State Systems, *Phys. Rev. Lett.* **85**, 3313 (2000).
 - [21] N. J. Cerf, M. Bourennane, A. Karlsson, and N. Gisin, Security of Quantum Key Distribution Using d -Level Systems, *Phys. Rev. Lett.* **88**, 127902 (2002).
 - [22] M. S. Blok, V. Ramasesh, T. Schuster, K. O’Brien, J. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Y. Yao, and I. Siddiqi, Quantum Information Scrambling on a

- Superconducting Qutrit Processor, *Phys. Rev. X* **11**, 021010 (2021).
- [23] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, Fault-tolerant detection of a quantum error, *Science* **361**, 266 (2018).
- [24] J. M. Chow, J. M. Gambetta, A. W. Cross, S. T. Merkel, C. Rigetti, and M. Steffen, Microwave-activated conditional-phase gate for superconducting qubits, *New J. Phys.* **15**, 115012 (2013).
- [25] X. Wu, S. L. Tomarken, N. A. Petersson, L. A. Martinez, Y. J. Rosen, and J. L. DuBois, High-Fidelity Software-Defined Quantum Logic on a Superconducting Qudit, *Phys. Rev. Lett.* **125**, 170502 (2020).
- [26] S. Ashhab, F. Yoshihara, T. Fuse, N. Yamamoto, A. Lupascu, and K. Semba, Speed limits for two-qubit gates with weakly anharmonic qubits, *Phys. Rev. A* **105**, 042614 (2022).
- [27] B. M. Varbanov, F. Battistel, B. M. Tarasinski, V. P. Ostroukh, T. E. O'Brien, L. DiCarlo, and B. M. Terhal, Leakage detection for a transmon-based surface code, *npj Quantum Inf.* **6**, 1 (2020).
- [28] J. Ghosh, A. G. Fowler, J. M. Martinis, and M. R. Geller, Understanding the effects of leakage in superconducting quantum-error-detection circuits, *Phys. Rev. A* **88**, 062329 (2013).
- [29] A. G. Fowler, Coping with qubit leakage in topological codes, *Phys. Rev. A* **88**, 042308 (2013).
- [30] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [31] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quantum Inf.* **7**, 1 (2021).
- [32] S. E. Economou and E. Barnes, Analytical approach to swift nonleaky entangling gates in superconducting qubits, *Phys. Rev. B* **91**, 161405 (2015).
- [33] X.-H. Deng, E. Barnes, and S. E. Economou, Robustness of error-suppressing entangling gates in cavity-coupled transmon qubits, *Phys. Rev. B* **96**, 035441 (2017).
- [34] 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**, 1 (2014).
- [35] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, *New J. Phys.* **22**, 043006 (2020).
- [36] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [37] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm. Preprint [ArXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).
- [38] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm. Preprint [ArXiv:1602.07674](https://arxiv.org/abs/1602.07674) (2016).
- [39] E. Farhi and H. Neven, Classification with quantum neural networks on near term processors. Preprint [ArXiv:1802.06002](https://arxiv.org/abs/1802.06002) (2018).
- [40] X. Hou, G. Zhou, Q. Li, S. Jin, and X. Wang, A universal duplication-free quantum neural network. Preprint [ArXiv:2106.13211](https://arxiv.org/abs/2106.13211) (2021).
- [41] M. Schuld, I. Sinayskiy, and F. Petruccione, The quest for a quantum neural network, *Quantum Inf. Process.* **13**, 2567 (2014).
- [42] S. Jeswal and S. Chakraverty, Recent developments and applications in quantum neural network: A review, *Arch. Comput. Methods Eng.* **26**, 793 (2019).
- [43] 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, *npj Quantum Inf.* **7**, 155 (2021).
- [44] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, From Pulses to Circuits and Back Again: A Quantum Optimal Control Perspective on Variational Quantum Algorithms, *PRX Quantum* **2**, 010101 (2021).
- [45] 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).
- [46] A. Choquette, A. Di Paolo, P. K. Barkoutsos, D. Sénéchal, I. Tavernelli, and A. Blais, Quantum-optimal-control-inspired ansatz for variational quantum algorithms, *Phys. Rev. Res.* **3**, 023092 (2021).
- [47] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control, *Quantum* **6**, 824 (2022).
- [48] B. Russell, H. Rabitz, and R.-B. Wu, Control landscapes are almost always trap free: A geometric assessment, *J. Phys. A: Math. Theor.* **50**, 205302 (2017).
- [49] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, Theory of overparametrization in quantum neural networks. Preprint [ArXiv:2109.11676](https://arxiv.org/abs/2109.11676) (2021).
- [50] S. Deffner and S. Campbell, Quantum speed limits: From Heisenberg's uncertainty principle to optimal quantum control, *J. Phys. A: Math. Theor.* **50**, 453001 (2017).
- [51] P. M. Poggi, Analysis of lower bounds for quantum control times and their relation to the quantum speed limit. Preprint [ArXiv:2002.11147](https://arxiv.org/abs/2002.11147) (2020).
- [52] K. W. M. 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).
- [53] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal Protocols in Quantum Annealing and Quantum Approximate Optimization Algorithm Problems, *Phys. Rev. Lett.* **126**, 070505 (2021).
- [54] C. Lin, Y. Wang, G. Kolesov, and U. Kalabić, Application of Pontryagin's minimum principle to Grover's quantum search problem, *Phys. Rev. A* **100**, 022327 (2019).
- [55] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, Ctrl-VQE code, <https://github.com/mayhallgroup/ctrlq> (2021).
- [56] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, and S. Sharma, *et al.*, PySCF: The Python-based simulations

- of chemistry framework, Wiley Interdiscip. Rev. Comput. Mol. Sci. **8**, e1340 (2018).
- [57] H. S. Ku, J. L. Long, X. Wu, M. Bal, R. E. Lake, E. Barnes, S. E. Economou, and D. P. Pappas, Single qubit operations using microwave hyperbolic secant pulses, *Phys. Rev. A* **96**, 042339 (2017).
- [58] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevApplied.19.064071> for a derivation of the control Hamiltonian used and a numerical study to study the impact of Trotter steps and pulse windows.
- [59] U. Boscain, M. Sigalotti, and D. Sugny, Introduction to the Pontryagin Maximum Principle for Quantum Optimal Control, *PRX Quantum* **2**, 030203 (2021).
- [60] Conventionally known as the control Hamiltonian, but we use the edited name to avoid confusion with the drive Hamiltonian (\bar{H}_C).
- [61] D. Liberzon, *Calculus of Variations and Optimal Control Theory* (Princeton University Press, Princeton, NJ, 2012).
- [62] Here we use 100 pulse segments and 1000 Trotter steps. We tested smaller and larger Trotter steps and found no appreciable change in results.