Using Device Physics and Error Mitigation to Improve the Performance of Quantum Computers

Samantha V. Barron

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Physics

Sophia E. Economou, Chair
Edwin F. Barnes
Kyungwha Park
Nguyen Q. Vinh

October 17th, 2022
Blacksburg, Virginia

Keywords: Quantum Computing, Quantum Simulation, Near Term Intermediate Scale Quantum Computing, Quantum Control, Variational Quantum Algorithms
Copyright 2023, Samantha V. Barron

Using Device Physics and Error Mitigation to Improve the Performance of Quantum Computers

Samantha V. Barron

(ABSTRACT)

Quantum computers have seen rapid development over the last two decades. Despite this, they are not yet scalable or fault-tolerant (i.e. we cannot address arbitrarily many error-corrected qubits). Therefore, improvements that include consideration of the underlying physics are paramount. To do this, we must reduce existing errors and understand how algorithms perform without error correction. In this dissertation, we provide contributions toward these goals. We organize these efforts into three groups.

Firstly, we focus on quantum control. We introduce a novel scheme for performing entangling gates on superconducting qubits. We create fast, high-fidelity entangling operations and single-qubit gates to implement arbitrary quantum operations. Then, we implement entangling gates on real transmon qubits. Finally, we develop new techniques for entangling gates on spin qubits. In total, we improve low-level device performance with high-fidelity entangling operations.

Secondly, we focus on quantum simulation algorithms. First, we apply error mitigation techniques to a quantum simulation algorithm while simultaneously performing device characterization. Then we take advantage of known symmetries of the input Hamiltonian to improve the same algorithm. Then, we demonstrate that this reduces resources compared to other approaches in the presence of noise. Then we compare this technique with state-of-the-art approaches. Then, we improve this algorithm with approaches from quantum control.

Finally, we develop a novel algorithm to simulate spin chains on a quantum processor with improved resources compared to other techniques. In total, we improve quantum simulation algorithms, with the aim of better utilizing current devices.

Thirdly, we consider the ADAPT-VQE algorithm, which is used to construct quantum circuits for preparing trial states in quantum simulation. In total, we improve gate counts for the algorithm, improve a separate algorithm by utilizing the gradient criterion, and leverage the repeating structure of an input Hamiltonian to improve performance. Finally, we provide a deeper understanding of ADAPT-VQE and demonstrate its robustness to scaling issues of competing algorithms. In total, we improve the algorithm and its applicability. Thus, we improve quantum simulation algorithms that can be run in the near term.

Using Device Physics and Error Mitigation to Improve the Performance of Quantum Computers

Samantha V. Barron

(GENERAL AUDIENCE ABSTRACT)

The computers that we interact with every day rely on the processing of bits, represented as 1's or 0's. The rules that govern how they operate mostly rely on classical physics (i.e. discovered before quantum physics), which does not include any quantum effects. If we instead allow for quantum rules and quantum bits ("qubits"), new types of algorithms are possible. That is to say, quantum computers can do some things more efficiently than classical computers. As such, there is a massive effort to build these devices. Because these devices are so delicate and in their early stages, this requires an understanding of the algorithm and the physical device performing it. Therefore, improving the overall performance requires taking high and low-level aspects of this design into consideration.

In this dissertation, we provide three groups of contributions to achieving this goal. In the first group, we improve the device performance by considering how operations are performed on qubits, primarily in terms of producing quantum operations that have no classical analog. In the second group, we improve the simulation of quantum systems on quantum devices with a focus on how existing imperfections in the device impact the results. In the third group, we make improvements to an algorithm used to simulate quantum systems like molecules, while also developing a deeper understanding of how the algorithm functions. In each of these parts, we develop novel techniques to improve device and algorithm performance, contributing to the applicability and utility of current and future quantum devices.

Dedication

To my high school friend, who was the first person to believe I could do this.

Acknowledgments

I would like to thank my advisor, Sophia Economou for her endless support throughout my time at Virginia Tech. Without her, this work would not have been possible, and I have been extremely fortunate to have been her student.

I would like to thank Ed Barnes and Nick Mayhall for their support on various projects, and endless hours of discussion in meetings. I would also like to thank my mentor Chris Wood for his support and guidance during my internships.

I would like to thank all of the members of Nick, Ed, and Sophia's groups (past and current) for their discussion and support. Particularly Harper Grimsley, Bryan Gard, Sona Najafi, Xiuhao Deng, Fernando Calderon-Vargas, John Van Dyke, and Linghua Zhu for debugging code and discussing ideas for projects.

I would like to thank the friends I met in Blacksburg for helping me be a complete person: Lauren Malhotra, Dan Ideozu, Pat Hicks, Molly Graham, Jess Nguyen, Collin Miller, Gretchen Goeke Dee, Avni Trasi, Riley DeHority, Erin Milligan, and Shelby Greenburg. I have never felt as strong a shared connection to a group of people, and you all make my life more vibrant than I could have ever imagined. I would also like to thank Charli XCX in general.

I would like to thank my family, including my mom, dad, and brother for their love, for instilling in me the importance of education when I was very young, and for supporting me at every step in my education.

And of course, I would like to thank my cat, Hoss Qubit Barron, for his endless love and affection. Like a qubit, you are so small and hard to control.

Contents

I Introduction 1

1 Qubits and Quantum Gates 2

1.1 Co-Authorship Acknowledgement 2

1.2 States of Qubits 2

1.3 Operations on Qubits 3

1.4 Single Qubit Operations 4

1.5 Two Qubit Gates 5

1.6 Quantum Control for Superconducting Qubits 6

1.7 Quantum Circuits 7

2 Variational Quantum Algorithms 9

II Quantum Control 12

3 Microwave-based Arbitrary CPHASE Gates for Transmon Qubits 13

4 A universal quantum gate set for transmon qubits with strong ZZ interactions 25

5 Fast high-fidelity entangling gates for spin qubits in Si double quantum

dots 41

III Quantum Simulation 52

6 Measurement Error Mitigation for Variational Quantum Algorithms 53

7 Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm 61

8 Preserving Symmetries for Variational Quantum Eigensolvers in the Presence of Noise 71

9 Benchmarking Quantum Chemistry Computations with Variational, Imaginary Time Evolution, and Krylov Space Solver Algorithms 87

10 Gate-free state preparation for fast variational quantum eigensolver simulations 103

11 Preparing Bethe Ansatz Eigenstates on a Quantum Computer 115

IV ADAPT-VQE 131

12 Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor 132

13 Adaptive quantum approximate optimization algorithm for solving com-

binatorial problems on a quantum computer 149

14 Scaling adaptive quantum simulation algorithms via operator pool tiling 159

15 ADAPT-VQE is insensitive to rough parameter landscapes and barren plateaus 166

V Conclusions 180

16 Conclusions 181

Bibliography 184

Part I

Introduction

Chapter 1

Qubits and Quantum Gates

1.1 Co-Authorship Acknowledgement

I have had the great pleasure and fortune of collaborating with researchers at several institutions in different areas of quantum computing. In some of the following works, I am listed as the primary author. In works where this is not the case, this is highlighted along with individual author contributions at the beginning of each chapter. To briefly summarize here, the works where I am the primary author are Chapter 3 (Ref [1]), Chapter 6 (Ref [2]), Chapter 8 (Ref [3]), and the works where I am not the primary author are: Chapter 4 (Ref [4]), Chapter 5 (Ref [5]), Chapter 7 (Ref [6]), Chapter 9 (Ref [7]), Chapter 10 (Ref [8]), Chapter 11 (Ref [9]), Chapter 12 (Ref [10]), Chapter 13 (Ref [11]), Chapter 14 (Ref [12]), Chapter 15 (Ref [13]),

1.2 States of Qubits

This section and those following introduce basic concepts, notations, and techniques in quantum computing. Although this introduction does not follow it, each of these topics and more are covered in the well-known textbook by Nielsen and Chuang [14].

The state of a qubit is often represented with a statevector. The so-called "computational

states" that a qubit takes are often $|0\rangle$ and $|1\rangle$, which form an orthonormal basis for the Hilbert space on one qubit. Hence, we can represent the arbitrary state of a qubit with $|\psi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$ where $\alpha_i \in \mathbb{C}$ and $\langle\psi|\psi\rangle = 1$, hence $|\alpha_0|^2 + |\alpha_1|^2 = 1$. We will later discuss generalizations of this for applications to systems including other states (generally known as qudits, having d possible states). In particular, we will discuss transmon qubits in Chapters 3, 4, 10. Note that given the possible constraints on α_i here, we could parameterize $|\psi\rangle$ by $|\psi\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)e^{i\phi}|1\rangle$ for the azimuthal and polar angles, ϕ and θ, respectively. This mapping allows us to view single qubit (pure) states as points on a sphere of radius 1. This is known as the Bloch sphere.

The combined state of a multi-qubit system is expressed using tensor products. Correspondingly, a system of qubits may exist in a superposition of all possible combinations of the states of the individual qubits. We oftentimes abbreviate tensor products as $|i\rangle \otimes |j\rangle = |i, j\rangle = |ij\rangle$. Hence, an arbitrary state on two qubits can be written as $|\psi\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle$. Note that for a system of n qubits, there are 2^n computational basis states. In this case, the statevector for the system can be expressed as $|\psi\rangle = \sum_{b\in B^n} \alpha_b |b\rangle$ where B^n is the set of bitstrings acting on n qubits, e.g. $b = 101011$ for $n = 6$.

1.3 Operations on Qubits

Quantum gates on qubits are performed by time-evolving the state of the qubits. In the case of unitary dynamics, the evolution is governed by the Schrödinger equation and the Hamiltonian of the system, given by $i\hbar\frac{d}{dt}|\psi(t)\rangle = H(t)|\psi(t)\rangle$ where $\hbar$ is the reduced Planck constant, $|\psi(t)\rangle$ is the time-dependent state of the qubits, and $H(t)$ is the (Hermitian) Hamiltonian of the system. The solution to this differential equation is the time-evolved

state, given by $|\psi(t)\rangle$ over some time interval, usually $t \in [t_i, t_f] = [0, T]$ for simplicity. For constructing gates on qubits, this scheme is somewhat limiting, as we do not simply want to evolve one specific state to another, but instead we want to evolve all possible input states according to some gate. Therefore for quantum control applications, we typically seek the unitary evolution operator $U(t)$ such that $|\psi(t)\rangle = U(t)\,|\psi(0)\rangle$. The corresponding Schrödinger equation for unitary evolution operators is then $i\hbar\frac{d}{dt}U(t) = H(t)U(t)$. Note that in some applications we are interested in the time dynamics of $U(t)$ over some time interval, but at other times we mainly interested in $U(T)$, oftentimes abbreviated by U.

1.4 Single Qubit Operations

Some examples of these operations are Pauli operators (all unitary) (plus the identity, I), given by

$$\begin{aligned} I &= |0\rangle\langle 0| + |1\rangle\langle 1| \\ X &= |0\rangle\langle 1| + |1\rangle\langle 0| \\ Y &= i\,|1\rangle\langle 0| - i\,|0\rangle\langle 1| \\ Z &= |0\rangle\langle 0| - |1\rangle\langle 1|\,. \end{aligned}$$

Each of these can be seen as generators for a unitary operator, or unitary rotations. For $A \in \{X, Y, Z\}$ we define $R_A(\theta) = e^{-i\theta A/2}$. The sense in which these constitute rotations on the Bloch sphere are as follows: Each $R_A(\theta)$ is a rotation by θ about the axis spanned by the eigenstates of A on a statevector on the Bloch sphere. The eigenstates of the X and Y operators are usually denoted by $|\pm\rangle$ and $|\pm i\rangle$ respectively, expressed in the computational

basis as

$$|\pm\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}}$$
$$|\pm i\rangle = \frac{|0\rangle \pm i\,|1\rangle}{\sqrt{2}}.$$

The X, Y, and Z axes on the Bloch sphere are then the $+1$ eigenstates of the X, Y, and Z operators, namely $|0\rangle$, $|+\rangle$, and $|+i\rangle$.

One way of constructing arbitrary operations on single qubits is by Euler angles. For example, we could use only $R_Z(\theta)$ and $R_X(\pi/2)$ operations to construct an arbitrary single qubit gate. We discuss this in more detail and apply it to the construction of arbitrary single qubit gates in Chapters 3, 4.

1.5 Two Qubit Gates

In order to perform universal quantum operations on qubits, we also need to be able to perform two qubit operations. Some common examples of two qubit operations are

$$
\begin{aligned}
CNOT &= |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X \\
CZ &= |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes Z \\
CPHASE(\theta) &= |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes R_Z(\theta) \\
C(U) &= |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U \\
SWAP &= |00\rangle\langle 00| + |01\rangle\langle 10| + |10\rangle\langle 01| + |11\rangle\langle 11|
\end{aligned}
$$

among others. Here, U is an arbitrary single qubit operation. We remark here that even though these gates act on two qubits, all except the $SWAP$ gate can produce entanglement. In the case of $CNOT$, CZ, $CPHASE(\theta)$, and $C(U)$, note that the leftmost indexed qubit is the control qubit and the target qubit is on the right. This reflects the fact that the operation on the target qubit is applied conditionally on the state of the control qubit. The space of possible operations is much larger, since unitary operations on 4 dimensions can be described with 15 real parameters (i.e. $|SU(N)| = N^2 - 1$). Later in Chapters 3, 5 we will discuss more sophisticated methods of analyzing the effect of two qubit gates, known as the Makhlin Invariants.

1.6 Quantum Control for Superconducting Qubits

Thus far, we have not said much about the physical implementation of qubits. That is to say, what quantum system will we use to represent our qubit, what will $|0\rangle$ and $|1\rangle$ mean when implemented in a lab, and are there characteristics of this setup that make it more or less favorable than others? It is generally still an open question as to what physical implementation will be successful in providing scalable, fault tolerant quantum processors.

One of the leading types of qubits are superconducting qubits. Broadly speaking, these qubits exist on physical circuits that are externally controlled and measured with electrical components, and need to be kept at very low temperatures. In the case of transmons, the circuit is a capacitor and a Josephson junction. Superconducting qubits and competing systems are discussed in Chapters 3, 4.

Transmon superconducting qubits will be the focus of Chapters 3, 4. Although the allowable states in the system are not just $|0\rangle$ and $|1\rangle$ (i.e. there are higher excited states), the Josephson junction provides nonlinearity in the potential for the circuit that, when quantized, allows the $|0\rangle \leftrightarrow |1\rangle$ transition to be individually addressed. This is in contrast with a harmonic oscillator where the transition frequency between neighboring states is constant. The Hamiltonian and details of this system are described in more detail in Chapters 3, 4.

1.7 Quantum Circuits

A common visual representation for quantum operations is quantum circuits. In this representation, qubits are denoted by horizontal lines. Quantum operations are represented by blocks and symbols on and/or connecting these lines. Quantum circuits are typically read from left to right (i.e. gates on the left are performed before gates on the right), and the initial state of the register is represented on the left. An example of a simple quantum circuit is provided in the figure below.

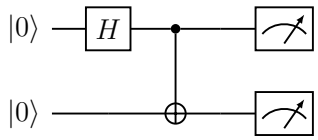

Figure 1.1: Simple quantum circuit containing an H and a $CNOT$. The state of the two qubits is initially $|00\rangle$. After the gates are applied, measurement is performed on each of the qubits. The final state before measurement is represented by $CNOT_{1,2}H_1 |00\rangle$, where the subscripts indicate the qubits on which the gate act.

Chapter 2

Variational Quantum Algorithms

A Variational Quantum Algorithm (VQA) is an algorithm that involves some sort of optimization where a quantum circuit is being varied in order to produce results for processing on a classical computer. Since the algorithm involves a classical and a quantum computer communicating, they are considered examples of hybrid quantum-classical algorithms. One can choose the cost function in a VQA to be the expectation value of a Hamiltonian, resulting in the Variational Quantum Eigensolver (VQE), as we discuss in more detail later. These algorithms have been developed over roughly the last decade, and have numerous extensions, variants, and applications. For more information about these algorithms, Ref. [15] is an extensive review of VQAs, and Refs. [16],[17] are concerned with VQEs specifically. Additionally, Ref. [18] discusses VQEs in the context of quantum chemistry. More broadly, VQAs and VQEs fall under the topic of noisy intermediate scale quantum (NISQ) algorithms for quantum simulation. NISQ topics are reviewed in Ref. [19] and quantum simulation is reviewed in Ref. [20].

A more concrete description of VQAs are as follows. The quantum computer is initialized in some state $|\psi_0\rangle$, then a circuit described by the unitary $U(\theta)$ is applied, where θ are parameters that determine the circuit applied (e.g. angles for gates). Following the circuit, some set of observables $\{O_i\}_i$ is measured. These expectation values are combined in some way according to g to produce the cost function $f(\theta)$. In total, this looks like

$$f(\theta) = g(\langle\psi_0| U(\theta)^\dagger O_0 U(\theta) |\psi_0\rangle, \langle\psi_0| U(\theta)^\dagger O_1 U(\theta) |\psi_0\rangle, ...). \tag{2.1}$$

Once $f(\theta)$ can be computed using data from the quantum computer, the job of the classical computer is to suggest values of θ to try in order to minimize (or maximize) the cost function θ. This formulation is fairly general, and in practice g and $\{O_i\}_i$ are determined by the algorithm. Typically $U(\theta)$ and $|\psi_0\rangle$ are chosen with some domain knowledge as discussed in Chapters 7, 8, 9, or adaptively constructed as discussed in Chapters 12, 13, 14, 15.

Throughout this work, we will often be interested in the Variational Quantum Eigensolver (VQE). In this case, the cost function is given by $f_{\text{VQE}}(\theta) = \langle\psi_0| U(\theta)^\dagger H U(\theta) |\psi_0\rangle$ where H is a Hamiltonian. The goal of the VQE is to prepare the ground state of the given Hamiltonian, and measure its energy. This relies on the variational principle, which states that the expectation value of a Hamiltonian H by any state in the hilbert space $|\phi\rangle$ is bounded below by the smallest eigenvalue of H, E_0, that is $\langle\phi| H |\phi\rangle \geq E_0$. In the present situation, this means that the cost function obeys $f_{\text{VQE}}(\theta) \geq E_0$ where E_0 is the ground state energy of the Hamiltonian H. The idea is that if there exists θ^* such that $U(\theta^*) |\psi_0\rangle = |E_0\rangle$, then globally minimizing the cost function $f_{\text{VQE}}(\theta)$ is equivalent to finding the ground state.

There are several rather limiting assumptions in this setup, which we describe in more detail in Chapters 7, 8, but to briefly summarize some key points here: The most immediate is that we can indeed prepare the ground state with the given circuit (i.e. that such θ^*) exists in the first place. Additionally, it is not clear whether globally minimizing $f_{\text{VQE}}(\theta)$ is viable for an arbitrary Hamiltonian. In fact, it has been shown that this problem is in a complexity class that is considered difficult. Furthermore, in the presence of noise, the cost function $f_{\text{VQE}}(\theta)$ may not reliably or consistently be evaluated. Therefore, special consideration is needed if this algorithm is performed on current hardware. These limitations and others are central

to the work in Chapters 7, 8. Notably, we consider circuits which are known to satisfy the preparation criteria (involving θ^*) for the given problem, and we analyze VQE performance in the presence of simulated and real device noise.

Part II

Quantum Control

Chapter 3

Microwave-based Arbitrary CPHASE Gates for Transmon Qubits

This chapter includes pages from our publication [1]

G. S. Barron, F. A. Calderon-Vargas, J. L. Long, D. P. Pappas, and S. E. Economou, "Microwave-based arbitrary CPHASE gates for transmon qubits", Physical Review B **101**, ARTN05450810.1103/PhysRevB.101.054508 (2020)

Copyright (2020) of the American Physical Society.

The contributions to this work are as follows: My contributions were carrying out numerical simulations, analytical derivations, preparation of figures, primary writing of the manuscript, bibliography preparation and conducting background material search. F.A.C.V. assisted in theoretical developments. J.L. assisted in choosing Hamiltonian parameters and evaluating experimental considerations. D.P.P. and S.E.E. supervised.

Microwave-based arbitrary CPHASE gates for transmon qubits

George S. Barron,[1,*] F. A. Calderon-Vargas,[1] Junling Long,[2] David P. Pappas,[2] and Sophia E. Economou[1,†]
[1]*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2]*National Institute of Standards and Technology, Boulder, Colorado 80305–3328, USA*

(Received 31 March 2019; revised manuscript received 4 January 2020; accepted 17 January 2020; published 20 February 2020)

Superconducting transmon qubits are of great interest for quantum computing and quantum simulation. A key component of quantum chemistry simulation algorithms is breaking up the evolution into small steps, which naturally leads to the need for nonmaximally entangling, arbitrary CPHASE gates. Here we design such microwave-based gates using an analytically solvable approach leading to smooth, simple pulses. We use the local invariants of the evolution operator in SU(4) to develop a method of constructing pulse protocols, which allows for the continuous tuning of the phase. We find CPHASE fidelities of more than 0.999 and gate times as low as 100 ns.

DOI: 10.1103/PhysRevB.101.054508

I. INTRODUCTION

Quantum computing promises solutions to a number of problems in computing, chemistry, and material science. Superconducting qubits are a promising candidate for qubits because their fabrication relies on existing techniques [1,2], and they can also have their characteristics tailored for specific applications.

Superconducting qubits have been recently used in the implementation of quantum algorithms for molecular problems [3–5], reinforcing the idea that quantum chemistry is one of the most appealing applications of quantum computing [6]. In many quantum simulation algorithms, gate decompositions of Trotterized Hamiltonians often include CPHASE gates, which are then written in terms of two maximally entangling CNOT gates [7]. This decomposition is shown in Fig. 1. Clearly, using CPHASE gates instead of CNOTs would reduce circuit depth and potentially improve resource use in terms of time and fidelity.

Fast high-fidelity two-qubit gates remain challenging in superconducting qubits [8]. Spectral crowding makes accurately addressing an individual transition to produce a controlled operation difficult over short times because the bandwidth required to resolve differences between nearby transitions becomes very small, increasing the time required for each gate [9]. The trade-off is then that either gate times are long or the gate fidelity is low.

One approach to implementing two-qubit gates in superconducting qubits is to dynamically tune elements of the circuit. For example, one can tune the qubit frequency [2,10–14], resonator frequency [15,16], or the coupling strength [17,18]. Unfortunately, tunable elements introduce charge noise, leading to decoherence and low fidelity. An alternative method is to apply microwave pulses to the qubits to drive transitions that implement unitary rotations [19–33]. Typically, microwave-based control selects a single transition to implement a two-qubit gate. However, spectral crowding is a generic issue for systems controlled exclusively by microwave pulses since, without tuning, their spectra are fixed (to Stark shift effects) and this usually forces the gate time to be very long to spectrally select the target transition. Moreover, the always-on coupling in these systems makes single-qubit gates nontrivial, especially for strongly coupled qubits.

In this work, we develop a collection of microwave-based CPHASE gates using the speeding up wave forms by inducing phases to harmful transitions (SWIPHT) [22] protocol, which overcomes spectral crowding. This protocol was recently used in experiment to produce CNOT gates between two transmon qubits [33]. Here we make use of hyperbolic secant (sech) pulse envelopes [34] which are smooth and simple to implement and produce high fidelities with low gate times for a variety of angles [35,36]. These type of pulses were recently used on transmons in experimental demonstrations of Z gates [37] and as part of a two-qubit gate [28]. We use the local invariants [38,39] of the two-qubit analytic evolution operator with control sech pulses to find conditions on the pulse parameters that achieve the desired two-qubit operation. Through simulations of transmons with typical parameters, we show that our CPHASE gates produce high fidelities for low gate times. These CPHASE gates are applicable in either an all-microwave context or a microwave-tuning hybrid context. Regarding the latter, our CPHASE gates are applicable in the sense that they only rely on a weak effective $Z \otimes Z$ coupling compared to methods that dynamically tune circuit elements. This reliance on only a weak amount of dynamical tuning of the circuit parameters allows these gates to be performed in a variety of parameter regimes. To address the generic challenge of implementing single-qubit gates with fixed-frequency, always-coupled transmons, we design a composite pulse protocol that gives high-fidelity X rotations, which along with our two-qubit gates and previously available Z gates [37,40] form a universal set. These single-qubit gates all take less than 50 ns each and have fidelities in excess of 0.992.

*gbarron@vt.edu
†economou@vt.edu

 ©2020 American Physical Society

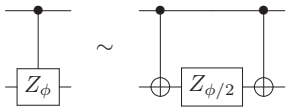

FIG. 1. Nonmaximally entangling CPHASE gate decomposed into two maximally entangling CNOT gates.

This paper is organized as follows. In Sec. II we introduce the two-qubit Hamiltonian for the system of transmons coupled by a resonator. In Sec. III we present the results of the analytical CPHASE protocols and numerical performance, as well as their robustness in other coupling strength regimes. In Sec. IV we present our single-qubit gates along with their fidelities. We conclude in Sec. V.

II. TRANSMON HAMILTONIAN

We focus on two superconducting transmon qubits coupled by a cavity [41]. The transmons are modeled as weakly anharmonic oscillators and the cavity as a harmonic oscillator. The Hamiltonian for this system is given by

$$H_0 = \omega_c a^\dagger a + \sum_{j=1,2} \epsilon_{j,1} a_j^\dagger a_j - \frac{\eta_j}{2} a_j^\dagger a_j (a_j^\dagger a_j - 1) + g_j (a_j^\dagger a + a^\dagger a_j), \tag{1}$$

where ω_c is the frequency of the cavity connecting the two qubits, $\epsilon_{j,1}$ is the transition frequency between the ground and first excited states for the jth qubit, η_j is the anharmonicity of the jth qubit, g_j is the coupling strength between the cavity and the jth qubit, a ($a^\dagger$) is the annihilation (creation) operator for the cavity, and a_j ($a_j^\dagger$) is the annihilation (creation) operator for the jth qubit. The Hamiltonian describing the coupling to the external microwave electric field is given by

$$H_p(t) = \sum_{j=1,2} E_j(t) e^{i\omega_{p,j} t} a_j + \text{H.c.}, \tag{2}$$

where $E_j(t)$ and $\omega_{p,j}$ are the pulse envelope and frequency driving the jth qubit, respectively. For the design of our gates we only drive (without loss of generality) the second qubit so that $E_1(t) = 0$, $E_2(t) = E(t)$, and $\omega_{p,2} = \omega_p$.

The states in the system are $|i; j, k\rangle = |i\rangle_c |j, k\rangle$, where $|i\rangle_c$ is the ith cavity level and the jth (kth) index denotes the level of the first (second) transmon. It is advantageous to write out the Hamiltonian in the dressed basis [42], which diagonalizes H_0, and the indices of each element of the dressed basis is determined by the state in the bare basis that has the largest overlap with the dressed state. For example, for indices s_i we write an element of the dressed basis as an eigenstate of H_0 with $|\widetilde{s_1}\rangle = \sum_i \alpha_i |s_i\rangle$, where $|\alpha_1| > |\alpha_i|$ with $i \neq 1$. We encode each qubit into the lowest two levels of each transmon. Consequently, the projection operator for the two-qubit subspace is $P_{\text{QSS}} = |\widetilde{0;0,0}\rangle \langle\widetilde{0;0,0}| + |\widetilde{0;0,1}\rangle \langle\widetilde{0;0,1}| + |\widetilde{0;1,0}\rangle \langle\widetilde{0;1,0}| + |\widetilde{0;1,1}\rangle \langle\widetilde{0;1,1}|$. Going to the dressed basis and projecting into the qubit subspace spanned by the basis $|\widetilde{0;0,0}\rangle$, $|\widetilde{0;0,1}\rangle$, $|\widetilde{0;1,0}\rangle$, $|\widetilde{0;1,1}\rangle$, the approximate two-qubit Hamiltonian when only one qubit is driven is given by

$$H_{\text{QSS}} \approx \begin{bmatrix} -\omega_{I,1}/2 & \Omega_1(t) e^{i\omega_p t} & 0 & 0 \\ \Omega_1(t)^* e^{-i\omega_p t} & +\omega_{I,1}/2 & 0 & 0 \\ 0 & 0 & -\omega_{I,2}/2 & \Omega_2(t) e^{i\omega_p t} \\ 0 & 0 & \Omega_2(t)^* e^{-i\omega_p t} & +\omega_{I,2}/2 \end{bmatrix}. \tag{3}$$

The block-diagonal approximation here corresponds to neglecting the presence of the cross-resonance interaction that is introduced when strongly driving qubits at or near the frequency of neighboring qubits. Our pulses are limited to comparatively small driving fields so we expect the cross-resonance effect to be negligible. Indeed, this is confirmed by our numerical simulations. We define $\delta\omega_I$ as the difference between the transition frequencies of the two subspaces $\omega_{I,1}$ and $\omega_{I,2}$ each corresponding with subspace 1 (upper left block) and subspace 2 (lower right block) of the Hamiltonian, respectively, as well as $\Omega_i(t) = E(t) d_i$ for the dipole moment d_i of each transition. Here we have made the approximation that terms in the Hamiltonian that couple states with a different number of excitations on the first qubit will vanish. This is due to the fact that in the dressed basis, since our off-diagonal coupling terms in H_0 are small compared to the diagonal terms, $|\langle\widetilde{i; j, k}\rangle i; j, k|$ is large compared to contributions from other states.

To design fast gates, we avoid spectrally selecting one of the two subspaces and allow the pulse to drive both transitions. Because in general $d_1 \neq d_2$ and $\omega_{I,1} \neq \omega_{I,2}$, the same $E(t)$ on each block will produce different evolutions. Our goal is to design control pulses $E(t)$ that generate two-qubit gates of the form $|0\rangle\langle 0| \otimes \mathbb{I}_2 + |1\rangle\langle 1| \otimes U$, and other control pulses that generate single-qubit gates of the form $\mathbb{I}_2 \otimes U$.

III. CPHASE GATES

For each of the following CPHASE gates, we use hyperbolic secant pulses of the form $\Omega(t) = \Omega_0 \text{sech}(\sigma t)$ with bandwidth σ, amplitude Ω_0, and pulse frequency ω_p. This pulse is chosen because it gives an analytically solvable time-dependent Schrödinger equation for a two-level system [34], is smooth, and has nice analytic properties for rotations about the Z axis [35] (see Appendix A for the derivation of the evolution operator and discussion of its properties). Specifically, for detuning Δ and bandwidth σ, a 2π hyperbolic secant pulse will induce a phase $2\arctan(\sigma/\Delta)$ and a 4π pulse will induce a phase $2\arctan[\frac{4\Delta/\sigma}{(\Delta/\sigma)^2-3}]$ [36]. A plot of two examples of hyperbolic secant pulses is shown in Fig. 2. The main idea is that the same sech pulse acts on both (target and harmful) transitions, causing a cyclic evolution to each subspace. This assumes that the dipoles of the two transitions are the same, which is not strictly the case. Nevertheless, approximately equal dipoles, as is the case for the parameters here, suffice for high fidelities. Due to the different detunings of the two transitions from the pulse, each acquires a different phase. The choice of phases for the two transitions, which we can control through the bandwidth and frequency of the pulse, determines the specific CPHASE gate. Since we focus on CPHASE gates, we use 2π and 4π pulses, which only implement cyclic transitions between energy levels. Our pulses generate generalized CPHASE gates,

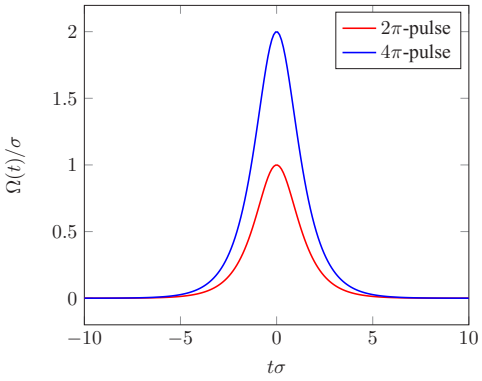

FIG. 2. Hyperbolic secant 2π (lower curve) and 4π (upper curve) pulses. These two different pulse areas have different algebraic properties, which lead to different types of protocols.

defined as $\text{CPHASE}' = \text{diag}(e^{i\phi_{00}}, e^{i\phi_{01}}, e^{i\phi_{10}}, e^{i\phi_{11}})$, which is equivalent to a regular CPHASE gate, $\text{CPHASE} = (1, 1, 1, e^{i\theta})$, up to local Z rotations. The phases in both the generalized and regular CPHASE gates satisfy $\theta = \phi_{00} - \phi_{01} - \phi_{10} + \phi_{11}$. In systems of transmons, it has been shown that zero-duration single-qubit Z rotations may be accomplished by shifting the phase of the microwave pulse [40], so this generalization does not affect our gate times or fidelities. Moreover, as discussed in Appendix A, the pulse areas considered here produce no transfer of population and hence only perform rotations about the Z axis. For this reason, although the local invariants allow us to consider arbitrary evolutions in SU(4) up to arbitrary rotations in SU(2) (see Appendix B), our pulses only require that we consider local operations of the form $R_Z(\phi)$, which, as discussed above, do not affect gate times or fidelities.

In the following results, we denote protocols that use transitions that exist inside the qubit subspace as "IQSS" and protocols that use transitions partially outside the qubit subspace as "OQSS." These two sets of transitions are illustrated in Fig. 3. In particular, when we refer to a protocol that is "IQSS," the transitions and their respective frequencies that we consider are $\omega_{I,1} : |\widetilde{0;00}\rangle \leftrightarrow |\widetilde{0;01}\rangle$, $\omega_{I,2} : |\widetilde{0;10}\rangle \leftrightarrow |\widetilde{0;11}\rangle$. On the other hand, if the protocol is "OQSS," then the transitions and their respective frequencies that we consider are $\omega_{O,1} : |\widetilde{0;01}\rangle \leftrightarrow |\widetilde{0;02}\rangle$, $\omega_{O,2} : |\widetilde{0;11}\rangle \leftrightarrow |\widetilde{0;12}\rangle$. As per the SWIPHT protocol, in either of these cases we designate either the IQSS or OQSS transitions with either the harmful or target transitions with transition frequencies $\omega_{x,h}$ and $\omega_{x,t}$, respectively. From these we define the difference $\delta\omega_x = \omega_{x,t} - \omega_{x,h}$ with $x \in \{I, O\}$ depending on the transitions chosen.

When evaluating the performance of the derived protocols, we numerically solve the Schrödinger equation to obtain the evolution operator at the end of each pulse. In our simulations we keep 3 states for the cavity and 4 states for each of the qubits, so that the Hilbert space simulated is 48 dimensional. This sufficiently simulates the full dynamics of the system in that adding more available states does not change our resulting fidelities. To compare the final evolution operator we obtain from the simulation with the target one, we calculate the fidelity given by $F = \frac{1}{n(n+1)}[\text{Tr}(MM^\dagger) + |\text{Tr}(M)|^2]$ [43], where $M = U_0^\dagger U$, with U_0 being the desired gate and U being the actual gate from simulations. Each U and U_0 are truncated so that they act only on the qubit subspace. In our numerical simulations we use $\omega_c = 7.15$ GHz, $\epsilon_{1,1} = 6.2$ GHz, $\epsilon_{2,1} = 6.8$ GHz, $\eta_1 = \eta_2 = \eta = 350$ MHz, and $g_1 = g_2 = g = 130$ MHz as the fixed parameters, except in Sec. III B, where we evaluate the performance of the gates when varying the coupling strength g. From these, we find that $\delta\omega_I = 3.23$ MHz and $\delta\omega_O = -11.07$ MHz. Bell-shaped pulses such as Gaussians or hyperbolic secants used here theoretically extend to infinity. In practice, of course, experimentalists switch on and off the pulses at finite times. Similarly, simulations require truncation of the pulses. We choose the gate time such that the simulation runtime is reasonable without sacrificing accuracy. We find that $T = 10/\sigma$ [which corresponds to $\Omega(T)/\Omega_0 \approx 0.01347$] suffices to achieve this. Moreover, we numerically optimize around the analytically predicted solution to compensate for errors such as a difference in the dipoles of the two transitions. Below we describe each of the protocols and provide results from numerical simulations quantifying their performance.

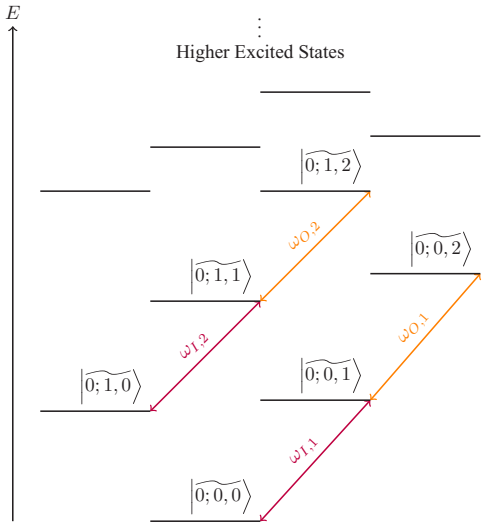

FIG. 3. Transitions between two-qubit states. The purple transitions (labeled $\omega_{I,1}$ and $\omega_{I,2}$) correspond to the IQSS protocols, and the orange transitions (labeled $\omega_{O,1}$ and $\omega_{O,2}$) correspond to the OQSS protocols.

A. CPHASE gate via off-resonant 2π-pulse OQSS

Our strategy here is to find conditions on the bandwidth and pulse frequency of a hyperbolic secant pulse that performs a generalized CPHASE gate on the two subspaces defined above.

To this end, we use the local invariants (Appendix B) of the analytical evolution operator for the two-qubit system driven by sech pulses (Appendix A).

First, we recall that the definition of the generalized CPHASE gate is $\text{CPHASE}' = \text{diag}(e^{i\phi_{00}}, e^{i\phi_{01}}, e^{i\phi_{10}}, e^{i\phi_{11}})$, where the phase imparted is $\theta = \phi_{00} - \phi_{01} - \phi_{10} + \phi_{11}$. If we consider the two block-diagonal portions of the Hamiltonian H_{QSS} (3), then we can define two detunings between the pulse and the desired transitions, $\Delta_1 = \omega_p - \omega_{O,1}$ and $\Delta_2 = \omega_p - \omega_{O,2}$.

Now for the OQSS protocol the block-diagonal form of the Hamiltonian as well as the analytical solution for the unitary operator of a two-level system discussed in Appendix A allows us to write the evolution operator as $U = \text{diag}(1, f_1, 1, f_2)$. Here $f_j = {}_2F_1(-a, +a, \frac{-i\Delta_j+\sigma}{2\sigma}, 1)$ is the Gaussian hypergeometric function with $a \equiv \frac{\Omega}{\sigma} \in \mathbb{Z}$. Using Eq. (B3) We can compute the local invariants of the two-qubit evolution operator U, yielding

$$\begin{aligned} G_1(U) &= \tfrac{1}{8}(4 + \gamma + \gamma^*), \\ G_2(U) &= 0, \\ G_3(U) &= 2 + \frac{\gamma}{2}, \end{aligned} \tag{4}$$

where

$$\begin{aligned} \gamma = &\frac{\Gamma^2\left(\frac{\sigma-i\Delta_1}{2\sigma}\right)\Gamma\left(-a+\frac{1}{2}-\frac{i\Delta_2}{2\sigma}\right)\Gamma\left(a+\frac{1}{2}-\frac{i\Delta_2}{2\sigma}\right)}{\Gamma^2\left(\frac{\sigma-i\Delta_2}{2\sigma}\right)\Gamma\left(-a+\frac{1}{2}-\frac{i\Delta_1}{2\sigma}\right)\Gamma\left(a+\frac{1}{2}-\frac{i\Delta_1}{2\sigma}\right)} \\ &+ \frac{\Gamma^2\left(\frac{\sigma-i\Delta_2}{2\sigma}\right)\Gamma\left(-a+\frac{1}{2}-\frac{i\Delta_1}{2\sigma}\right)\Gamma\left(a+\frac{1}{2}-\frac{i\Delta_1}{2\sigma}\right)}{\Gamma^2\left(\frac{\sigma-i\Delta_1}{2\sigma}\right)\Gamma\left(-a+\frac{1}{2}-\frac{i\Delta_2}{2\sigma}\right)\Gamma\left(a+\frac{1}{2}-\frac{i\Delta_2}{2\sigma}\right)}. \end{aligned} \tag{5}$$

We also compute the local invariants for the target CPHASE gate, yielding

$$\mathbf{G}[\text{diag}(1, 1, 1, e^{i\theta})] = [\cos(\theta/2)^2, 0, 2 + \cos(\theta)], \tag{6}$$

where $\mathbf{G}(U) = (G_1(U), G_2(U), G_3(U))$.

In order to find the conditions on the pulse parameters, we demand that the local invariants of U [Eq. (4)] be equal to those of the CPHASE gate, and thus we arrive at (assuming for simplicity that $a = 1$)

$$\cos\theta = \cos[\phi_1(\Delta_1) - \phi_1(\Delta_2)], \tag{7}$$

where $e^{i\phi_1(\Delta_j)} = \frac{\sigma - i\Delta_j}{\sigma + i\Delta_j}$.

For the IQSS protocol, we follow the same procedure, except that in that case the IQSS evolution operator is $U = \text{diag}(f_1, f_1^*, f_2, f_2^*)$ After imposing that the respective local invariants of U and the CPHASE gate be equal, for a 2π pulse we arrive at

$$\cos\theta = \cos\{2[\phi_1(\Delta_1) - \phi_1(\Delta_2)]\}. \tag{8}$$

In our parameter regime, the IQSS protocol has similar or lower performance than the OQSS protocol, and thus in the remainder of this section we focus on the OQSS protocol only.

In the SWIPHT protocol, there is a notion of a "harmful" and "target" transition. The difference between these two transitions is that we select the "harmful" transition to be the transition that we want to drive to obtain a trivial phase. The "target" transition then corresponds to the transition that looks like the target portion of a controlled unitary operation. So here we see that there is some freedom in defining which of the two blocks (1 or 2) involves the target and which the harmful transition. We define the following choice of sign:

$$\lambda = \begin{cases} +1 & \text{Target is block 1} \\ -1 & \text{Target is block 2} \end{cases}. \tag{9}$$

We also define Δ_t and Δ_h based on this choice. Then noting that $\omega_p = \Delta_h + \omega_h = \Delta_t + \omega_t$ and defining $\delta\omega_O = \omega_t - \omega_h$, we can use the definitions for θ, ϕ_1 as well as trigonometric identities to find $\delta\omega_O/\sigma = \cot[\phi_1(\Delta_h)/2] - \cot[\phi_1(\Delta_t)/2]$.

Now if we specify an angle for the two-qubit gate, then all of the restrictions up until now allow us to find a pulse frequency and bandwidth that perform two different rotations on each block, but together they combine to form a CPHASE operation. We define $\theta_1 = \phi_{00} - \phi_{01}$, $\theta_2 = \phi_{10} - \phi_{11}$ as well as $\theta = \theta_1 + \theta_2$, $\delta\theta = \theta_1 - \theta_2$. These can be written in terms of the harmful/target detunings and λ as follows: $\theta = \lambda[\phi_1(\Delta_t) - \phi_1(\Delta_h)]$ and $\delta\theta = \phi_1(\Delta_t) + \phi_1(\Delta_h)$. With these expressions we can then find an expression for the bandwidth in terms of the desired angle,

$$\sigma = \lambda\delta\omega_O \frac{\cos(\theta/2) - \cos(\delta\theta/2)}{2\sin(\theta/2)}. \tag{10}$$

To successfully generate an arbitrary CPHASE gate we also need to express the control pulse frequency ω_p in terms of the desired θ angle. In this line, using previous definitions for ω_p, we can write $\omega_p = \frac{\omega_t+\omega_h}{2} + \frac{1}{2}(\Delta_t + \Delta_h)$, which can be easily rewritten in terms of θ and $\delta\theta$:

$$\begin{aligned} \omega_p &= \frac{\omega_t + \omega_h}{2} + \frac{1}{2}(\Delta_t + \Delta_h) \\ &= \frac{\omega_t + \omega_h}{2} + \sigma\frac{\sin(\delta\theta/2)}{\cos(\theta/2) - \cos(\delta\theta/2)}. \end{aligned}$$

Using Eq. (10) to further simplify the previous equation, we find that the pulse frequency in terms of the desired angle θ is

$$\begin{aligned} \omega_p = &\frac{\omega_{O,t} + \omega_{O,h}}{2} \\ &+ \lambda\frac{\delta\omega_O}{2\sin(\theta/2)}\sqrt{1 - [\cos(\theta/2) - 2\sin(\theta/2)\sigma/\delta\omega_O]^2}, \end{aligned} \tag{11}$$

where, in order to ensure that the resulting pulse has finite frequency, we require that the angle of the CPHASE gate is within the range $\theta \in (0, \pi]$. Moreover, to make the pulse frequency real, this expression also provides a maximum allowable bandwidth for a given angle θ, $\sigma_{\max} = \frac{|\delta\omega_O|}{2}\cot(\theta/4)$.

Using this protocol, we find fidelities in excess of 0.999 and gate times as low as 60 ns. The numerical evaluations of the fidelity in the simulation for this protocol are shown in Fig. 4. From the figure we see that the fidelity is consistently above 0.992 for all angles and gate times. By choosing smaller bandwidths, one is able to increase the fidelity. The infidelity at small gate times is due to leakage outside the qubit subspace.

B. CPHASE gate via resonant pulses

The construction of the protocols presented in this section is similar to that of the off-resonant protocols in Sec. III A, with the difference that now we require the pulses to be

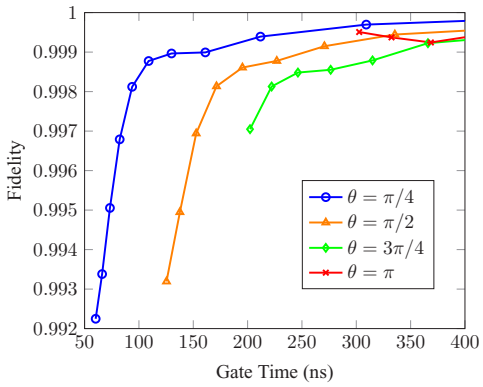

FIG. 4. Fidelity of numerical simulations for the CPHASE gate with the OQSS protocol. For every angle, there is a maximum allowable bandwidth and hence a minimum allowable gate time.

resonant with one of the two transitions in the system. In this family of protocols, we will first restrict the value $a = \Omega/\sigma \in \mathbb{Z}$ so that each subspace sees a rotation about the Z axis, independent of the detuning with the frequency of each transition.

For the IQSS case, this means that the evolution operator is diagonal and has the form $U = \mathrm{diag}(f_1, f_1^*, f_2, f_2^*)$, where $f_j = {}_2F_1(-a, +a, \frac{-i\Delta_j+\sigma}{2\sigma}, 1)$. Because we are constructing resonant protocols, without loss of generality, we let the pulse be resonant with the first transition so that $\Delta_1 = 0$ and $\Delta_2 = \delta\omega_I$. We may then compute the local invariants, yielding

$$\mathbf{G}(U) = \left[\frac{\mathrm{Re}(\alpha)^2}{|\alpha|^2}, 0, 2 + \frac{1}{2}\left(\frac{\alpha}{\alpha^*} + \frac{\alpha^*}{\alpha}\right)\right], \tag{12}$$

where $\alpha = {}_2F_1(-a, +a, \frac{-i\delta\omega_I+\sigma}{2\sigma}, 1)$ is the Gaussian hypergeometric function. Then, demanding that the local invariants of the CPHASE gate [Eq. (6)] be equal to the evolution operator for the pulsed system yields a single equation for this protocol:

$$\cos\theta = \frac{\alpha^2 + \alpha^{*2}}{2|\alpha|^2}. \tag{13}$$

As an example, for a 2π pulse we let $a = 1$. Then this equation reduces to

$$\frac{\delta\omega_I^4 - 6\delta\omega_I^2\sigma^2 + \sigma^4}{\left(\delta\omega_I^2 + \sigma^2\right)^2} = \cos\theta. \tag{14}$$

Solving this equation for the bandwidth yields four solutions, and because the bandwidth must be positive, only two are physical based on the sign of $\delta\omega_I$. For the $a = 1$ case, the final bandwidths are

$$\sigma_1 = |\delta\omega_I| \cot(\theta/4), \tag{15}$$

$$\sigma_2 = |\delta\omega_I| \tan(\theta/4). \tag{16}$$

This procedure may be repeated for pulses with $a \in \mathbb{N}$, though the resulting equations are more complicated and may be treated numerically. Additionally, one may repeat this protocol while using the OQSS transitions. The setup is essentially the same, except the OQSS evolution operator will be

$$U = \mathrm{diag}(1, f_1, 1, f_2). \tag{17}$$

In this case, the analogous single equation for these protocols will be

$$\cos\theta = (-1)^a \frac{1+\alpha^2}{2\alpha}, \tag{18}$$

where now $\alpha = {}_2F_1(-a, +a, \frac{+i\delta\omega_O+\sigma}{2\sigma}, 1)$ For OQSS transitions, where the control sech pulse is resonant with a transition partially out of the qubit space, the derivation of the protocols rely on solving Eq. (18) for the bandwidth of the pulse. In particular, for a sech 2π pulse ($a = 1$) we find that the associated bandwidth that produces a CPHASE gate for a given angle is $\sigma = |\delta\omega_O| \cot(\theta/2)$, where again we require that $\theta \in (0, \pi]$. On the other hand, if we instead use a 4π pulse, then the solution to the evolution operator has different properties compared to the resonant 2π case, and we find that the bandwidth for a specific angle $\theta \in (0, \pi)$ is given by

$$\sigma = |\delta\omega_O| \frac{\tan(\theta/2)}{\sqrt{4 + 3\tan(\theta/2)^2} \pm 2}. \tag{19}$$

In this case, the choice of sign is arbitrary and the bandwidth does not depend on which transition is designated as the harmful or target. However, the choice of sign determines the range of the bandwidth. We find that if the sign choice is positive, then $0 < \frac{\sigma}{|\delta\omega_O|} < 1/\sqrt{3}$ and if the choice of sign is negative, then $1/\sqrt{3} < \frac{\sigma}{|\delta\omega_O|}$. In some protocols derived here, there are multiple ranges of allowed bandwidths. These ranges result from the fact that multiple bandwidths satisfy Eq. (18) for a given angle and $\delta\omega$. In our parameter regime, this protocol has comparable or lower performance from the others simulated here, so we do not show numerical results in this case.

If we repeat this procedure but now choosing transitions corresponding to the IQSS case, then the different bandwidths are obtained by solving Eq. (13). For example, in the case of a 2π pulse ($a = 1$), the bandwidth for this CPHASE gate of angle $\theta \in (0, \pi]$ is $\sigma = |\delta\omega_I| \cot(\theta/4)$. Here we find gates with fidelities as high as 0.999999 and gate times as low as 24 ns for angles in the range of $\pi/16$ to $\pi/2$. To construct this protocol, the pulse is driven on resonance with one of the transitions inside the qubit subspace. We evaluate the performance of this protocol in simulation by calculating the gate fidelity, shown in Fig. 5. The two curves correspond to the two different choices of resonant transitions. The upper (blue) curve corresponds to the lower right block being the target, and the bottom (red) curve corresponds to the upper left block being the target. We find that the fidelity using subspace 2 as the target is above 0.9998 for angles from $\pi/8$ to $\pi/2$, and using the other transition as the target produces lower fidelities of ∼0.9995. In either case, we find reasonable gate times for this range of angles. The infidelity at smaller angles is due to leakage as a result of larger pulse amplitudes. In contrast with the other numerical results, in this protocol the desired angle of the gate fixes the bandwidth and hence the gate time.

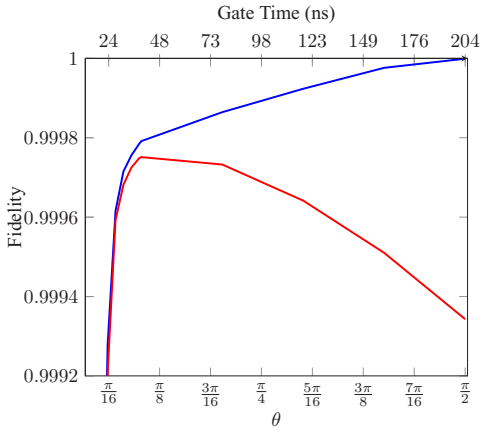

FIG. 5. Performance of CPHASE gate from simulation using the resonant IQSS 2π protocol. The upper curve corresponds to using subspace 2 as the target, $\omega_{l,2} = \omega_t$, and the bottom curve corresponds to using subspace 1 as the target, $\omega_{l,1} = \omega_t$.

When we repeat this procedure for a 4π pulse, the CPHASE gate of angle $\theta \in (0, \pi)$, has bandwidth $\sigma = |\delta\omega_l| \frac{\tan(\theta/4)}{\sqrt{4+3\tan(\theta/4)^2 \pm 2}}$. Again, the pulse is driven on resonance with one of the transitions inside the qubit subspace. We also find that the range on the bandwidth in the case when the choice of sign is positive becomes $0 < \frac{\sigma}{|\delta\omega_l|} < \frac{1}{2+\sqrt{7}}$ and when the choice of sign is negative, we have $\frac{1}{\sqrt{7}-2} < \frac{\sigma}{|\delta\omega_l|}$.

So far we have fixed the coupling to $g = 130$ MHz. Now we focus on the IQSS 2π protocol and evaluate its performance as a function of the coupling strength. We determine two primary features as we vary the coupling strength. First, weakly coupled systems produce gate times that increase rapidly as a function of the desired angle, as shown in Fig. 6. Second, increasing the coupling strength decreases the fidelity, as shown in Fig. 6. Overall, we find that for a range of coupling strengths we are able to find high fidelities exceeding 0.998. In some cases the fidelity is as high as 0.999999. In all cases, the fidelity drops for smaller angles due to leakage as a result of larger pulse amplitudes. We limit these simulations to 200 ns gate durations to compare the different coupling strengths because this protocol has no upper bound on the gate time.

C. CPHASE protocols comparison

In Table I we provide a summary of the results of the various protocols. Overall, we find that there is flexibility in the way of constructing CPHASE gates. For instance, one does not necessarily need to drive on resonance with one of the transitions. Additionally, one may choose various pulse areas or bandwidths for different implementations. For instance, one may choose to derive protocols with $a = 3, 4, \ldots$. This has the potential to reduce the duration of the gate at the cost of potentially introducing more leakage, which one could possibly address by incorporating Derivative Removal by Adiabatic Gate (DRAG) [44–46] into the pulse design, but this is beyond the scope of this work. In terms of performance, the two best protocols are the OQSS arbitrary frequency via 2π-pulse and IQSS resonant 2π-pulse protocols. Comparing the fidelities and gate times for a range of angles, if a small angle is desired one should choose the IQSS resonant 2π-pulse protocol because at small angles it provides consistently higher fidelities ($\sim$0.9999 compared to $\sim$0.998) at comparable gate times and sometimes fidelities as high as 0.999999. On the other hand, if a larger angle is desired, then the OQSS arbitrary frequency via 2π-pulse protocol is preferable due to its flexibility in the bandwidth, yielding potentially lower gate times ($\sim$120 ns compared to $\sim$200 ns). The other protocols produce fidelities on the order of $\sim$0.98 generally due to their higher bandwidths, which result in more leakage. In systems that do not have higher available states, these protocols may be more useful as they can produce smaller gate times for a range of angles.

IV. SINGLE QUBIT GATES

Now we turn our attention to single qubit operations. Developing single qubit operations in an always-coupled two-qubit system is inherently challenging. This is because the interaction between the two qubits dresses the energy eigen-

TABLE I. Summary of results for various CPHASE protocols. These results are derived in Secs. III A and III B. Each row denotes the analytical results of a particular protocol producing a CPHASE gate by an angle θ. The way to read this column is from left to right. First, choose a pair of transitions corresponding to the transitions in Fig. 3 and then choose a pulse area. The properties of the selected pulse will then be the bandwidth for a particular angle and the range on such bandwidths. In some cases, there are two disconnected ranges of allowable bandwidths. The OQSS 2π protocol has an arbitrary bandwidth in the sense that it is not a function of the desired angle. However, the bandwidth must satisfy the constraint in the corresponding "Bandwidth range" column.

Transitions	Pulse area	Frequency	Bandwidth	Bandwidth range
IQSS	2π	ω_t	$\sigma = \lvert\delta\omega_l\rvert \cot(\theta/4)$	$0 < \sigma < +\infty$
IQSS	4π	ω_t	$\sigma_\pm = \lvert\delta\omega_l\rvert \frac{\tan(\theta/4)}{\sqrt{4+3\tan(\theta/4)^2 \pm 2}}$	$0 < \frac{\sigma_+}{\lvert\delta\omega_l\rvert} < \frac{1}{2+\sqrt{7}} < \frac{1}{-2+\sqrt{7}} < \frac{\sigma_-}{\lvert\delta\omega_l\rvert} < +\infty$
OQSS	2π	ω_t	$\sigma = \lvert\delta\omega_O\rvert \cot(\theta/2)$	$0 < \sigma < +\infty$
OQSS	4π	ω_t	$\sigma_\pm = \lvert\delta\omega_O\rvert \frac{\tan(\theta/2)}{\sqrt{4+3\tan(\theta/2)^2 \pm 2}}$	$0 < \frac{\sigma_+}{\lvert\delta\omega_O\rvert} < \frac{1}{\sqrt{3}} < \frac{\sigma_-}{\lvert\delta\omega_O\rvert} < +\infty$
OQSS	2π	Eq. (11)	Arbitrary	$0 < \sigma < \frac{\lvert\delta\omega_O\rvert}{2} \cot(\theta/4) < +\infty$

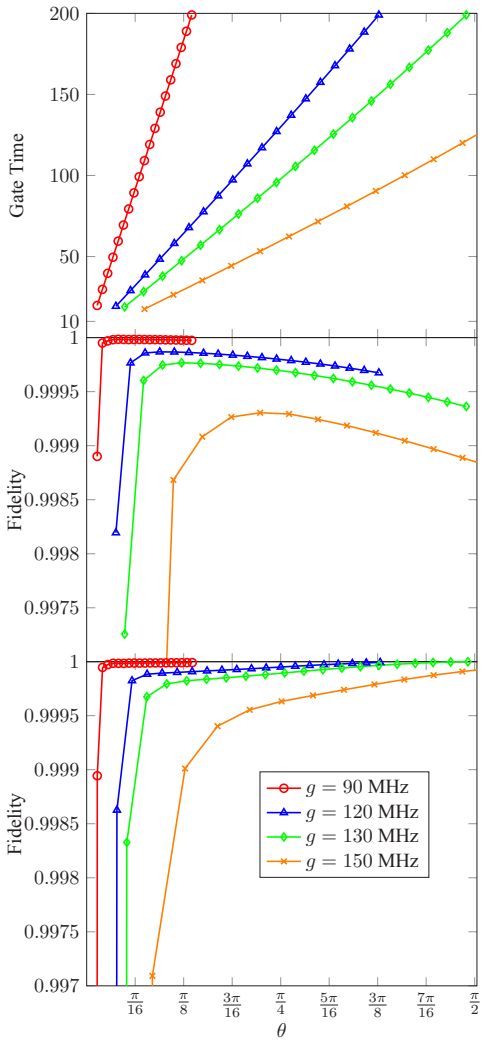

FIG. 6. Properties and performance of the CPHASE gates using the IQSS 2π protocol over a range of coupling strengths. The upper panel shows the gate time for each coupling strength g as a function of the angle of the CPHASE gate. The middle panel shows the fidelity at each angle for the various coupling strengths using subspace 1 as the target, $\omega_{l,1} = \omega_t$. The lower panel shows the fidelity at each angle for the various coupling strengths using subspace 2 as the target, $\omega_{l,2} = \omega_t$. For all panels, gate times between 10 ns and 200 ns are considered.

states, such that the frequency of the qubit transition in one qubit depends on the state of the other qubit. As a result, the problem is to find a pulse which implements the same operation in two distinct two-level systems. In the following subsections, we design single-qubit gates by providing two solutions to this problem. One is based on individual Gaussian pulses (Sec. IV A) and the other on sequences of square pulses (Sec. IV B).

A. Gaussian pulse protocol

Here we will design $R_{\hat{z}}(\theta)$ and $R_{\hat{x}}(\pi/2)$ rotations, which can be combined with rotations about Z to give arbitrary single-qubit gates. Since rotations about the Z axis may be produced by shifts in the frequency of the microwave pulse [40] with zero gate time and no loss in fidelity, we only consider the development of the rotations about the x axis.

For this protocol, we use a single Gaussian pulse, $\Omega(t) = \Omega_0 e^{-(\sigma t)^2}$, that is resonant with the transition frequency of the target qubit. Since the qubits are strongly coupled, the $|0\rangle \leftrightarrow |1\rangle$ transition frequency is conditional on the state of the other qubit. Here we choose the frequency when the qubit not being driven is in the $|0\rangle$ state. This pulse has two free parameters, the amplitude Ω_0 and the bandwidth σ. Treating this problem numerically, we will optimize the final fidelity of the evolution operator with the desired evolution operator, $\mathbb{I} \otimes R_{\hat{x}}(\pi/2)$. Accordingly, the evolution operator of the whole system (i.e., including higher levels of the transmons and cavity) is projected into the qubit subspace before calculating the fidelity. We use a Bayesian optimization algorithm, which is useful in cases where calls to the objective function are costly and gradients are not easily be evaluated. This technique has been used in quantum control before [47]. The result of this procedure is a Gaussian pulse with $\Omega_0 = 60.25$ MHz and $\sigma = 67.57$ MHz. Plots of the relevant transition probabilities and the pulse envelope are shown in Figs. 7(a) and 7(b), respectively. The pulse duration is 11.75 ns and the gate has a final fidelity of $F = 0.99563$. Interleaving two copies of this pulse with five rotations about the z axis allows us to perform any single qubit rotation.

B. Square pulse protocol

We also develop $R_{\hat{x}}(\theta)$ rotations via a sequence of square pulses. In this case, the Hamiltonian is simply the Hamiltonian in Eq. (3) for piecewise constant $\Omega_i(t)$ and $\omega_{p,i}(t)$. The evolution operator for the qubit subspace can be written as

$$U(\vec{\tau}, \vec{E}, \vec{\omega}) = \begin{bmatrix} U_1(\vec{\tau}, \vec{E}, \vec{\omega}) & 0 \\ 0 & U_2(\vec{\tau}, \vec{E}, \vec{\omega}) \end{bmatrix}, \quad (20)$$

where $U_j(\vec{\tau}, \vec{E}, \vec{\omega}) = \prod_{i=1}^{N} U_{j,i}(\tau_i, E_i, \omega_{p,i})$ and $U_{j,i}(\tau_i, E_i, \omega_{p,i})$ is the evolution operator for the jth block over the duration of the ith square pulse. The ith square pulse has duration τ_i, pulse amplitude E_i and frequency $\omega_{p,i}$.

Instead of solving exactly for parameters of each pulse that perform the desired evolution on each subspace, we define an objective function to optimize which is $f_{\hat{n},\theta}(\vec{\tau}, \vec{E}, \vec{\omega}) = F(\mathbb{I} \otimes R_{\hat{n}}(\theta), U(\vec{\tau}, \vec{E}, \vec{\omega}))$, where $F(U, V)$ is the fidelity between two unitary operators. We do this for several reasons. Primarily, there is no guarantee that such solutions exist, and even if they did, they would likely not be simple. Moreover, even if we solve for a sequence of pulses

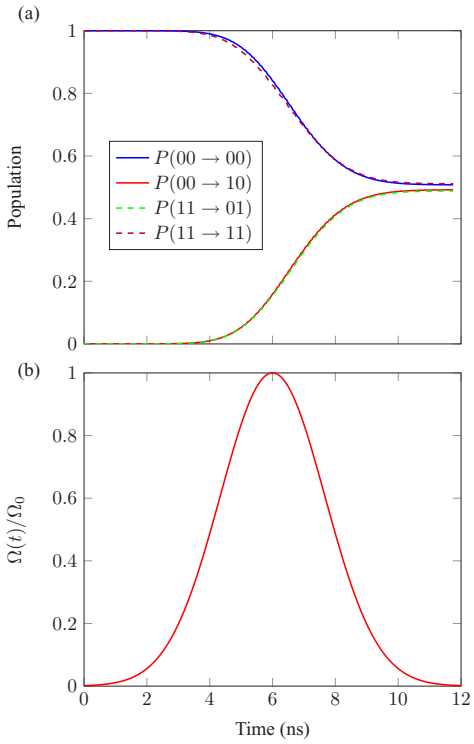

FIG. 7. (a) Plot of the probabilities as a function of time for various transfers of population to and from elements of the computational basis. Since we are performing the rotation $\mathbb{I} \otimes R_{\hat{x}}(\pi/2)$, we expect that the final probabilities shown end up at $1/2$. The dashed/solid lines are simply so that the plots do not overlap. (b) Plot of the pulse envelope $\Omega(t)$ as a function of time.

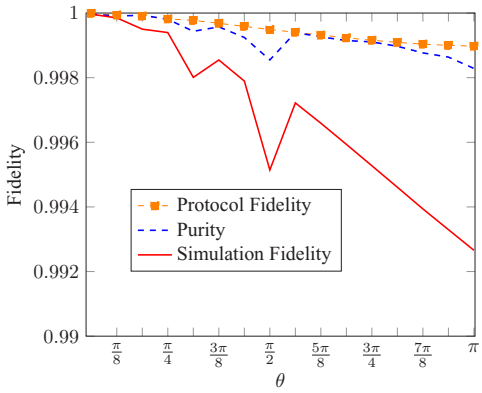

FIG. 8. Performance of single qubit rotation protocol. The upper (orange) curve is the protocol fidelity, determined by global optimization over the parameters of a sequence of resonant square pulses for the four-level system. The middle (blue) curve is the purity as a function of the angle. The bottom (red) curve is the fidelity found from simulation of the full time dynamics of the system after using the result from the protocol as an initial condition in a local optimization.

that exactly implements the desired evolution, in simulation and experiment the fidelity will not be exactly 1 due to decoherence. Here we use global, constrained optimization algorithms over the $3N$ parameters to find such sequences of pulses. The region in which the optimization is performed is determined by experimental limitations such as ramp-up times for the microwave pulses on the order of 1 ns and maximum possible amplitudes of each pulse based on the microwave pulse generators of about 20 MHz.

The desired evolution operator for the qubit subspace here is $\mathbb{I} \otimes R_{\hat{x}}(\theta)$ so that $U_j(\vec{\tau}, \vec{E}, \vec{\omega}) = R_{\hat{x}}(\theta)$ for each $j = 1, 2$. Without loss of generality, we choose a sequence of pulses resonant with the first subspace so that $\omega_{p,i} = \omega_{I,1}$. Then, with $E_i \in \mathbb{R}$, this sequence of square pulses naturally produces rotations about the x axis for the first subspace, $R_{\hat{x}}(\theta) = U_1(\vec{\tau}, \vec{E}, \vec{\omega})$. This provides the constraint $\theta/2 = d_1 \sum_{i=1}^{N} \tau_i |E_i|$. Now the optimization is over $2N$ parameters with one constraint.

We evaluate the performance of the single qubit X rotation protocol. This involves two steps: The first step is to determine the parameters on some sequence of square pulses by the optimization of $f_{\hat{n},\theta}(\vec{\tau}, \vec{E}, \vec{\omega})$, which yields what we define as the "protocol fidelity," see Fig. 8. The second step is to take the resulting sequence of square pulses and simulate the full time dynamics of the system, using a local optimization to improve the results of the protocol in the simulation. This is done by using the parameters of each pulse sequence from the protocol as initial conditions to a local optimization algorithm that improves the fidelity. We refer to this as the "simulation fidelity" in Fig. 8. We find "simulation fidelities" above 0.992 for all angles $0 < \theta \leqslant \pi$. All of these gates have durations from $\sim$15 ns to $\sim$25 ns. Because there is a gap between the "simulation fidelity" and purity in the figure, we see that there is some coherent error occurring. This is due to coupling to higher excited states which are not included in the four-level system and is the primary cause of the infidelity. The dip at $\theta = \pi/2$ is due to the fact that we use a local optimizer for the "simulation Ffidelity" and the curve is not guaranteed to be smooth.

V. CONCLUSIONS

Using the analytical evolution operator for the hyperbolic secant pulse acting on a two-level system, we have derived a collection of CPHASE gates for transmon qubits. We have demonstrated that these gates produce high fidelities typically in excess of $\sim$0.999 and in some cases as high as 0.999999 and typical gate times less than $\sim$100 ns. Moreover, we show that one of these protocols is robust in the fidelity for a range of angles and coupling strengths g. Finally, we

demonstrate that arbitrary single qubit gates may be achieved via microwave control in this parameter regime using either single Gaussian pulses or sequences of square pulses. Our high-fidelity parameterized entangling gates may be applied in realistic systems for use in quantum simulation algorithms.

ACKNOWLEDGMENTS

S.E.E. acknowledges support from the Department of Energy, Award No. de-sc0019318. D.P. acknowledges support by the Army Research Office, Award No. ARO W911NF1810114.

APPENDIX A: HYPERBOLIC SECANT PULSE SOLUTION

The basis for the CPHASE gate is the analytic solution for the evolution operator of a two-level system driven by a hyperbolic secant pulse. The pulse is defined as $\Omega(t) = \Omega_0 \mathrm{sech}(\sigma t)$, where σ is the pulse bandwidth and Ω_0 is the pulse strength. One can show that the form of the Hamiltonian for a given transition in the interaction frame is

$$H_2(t) = \begin{bmatrix} 0 & \Omega(t)e^{i\Delta t} \\ \Omega(t)^* e^{-i\Delta t} & 0 \end{bmatrix}, \tag{A1}$$

where Δ is the detuning of the pulse with the $|0\rangle \leftrightarrow |1\rangle$ transition, i.e., $\Delta = \omega_p - \epsilon_1$. By following a previous discussion of this problem [35], we define $c = \frac{1}{2}(1 + i\frac{\Delta}{\sigma})$, $a = \frac{\Omega}{\sigma}$, $\zeta = \frac{1}{2}[1 + \tanh(\sigma t)]$, $\alpha(a, c, \zeta) = {}_2F_1(a, -a, c^*, \zeta)$, and $\beta(a, c, \zeta) = {}_2F_1(a + 1 - c, 1 - a - c, 2 - c, \zeta)$ where ${}_2F_1$ is one of Gauss's hypergeometric functions. Then the evolution operator is

$$U(t, -\infty) = \begin{bmatrix} \alpha(a, c, \zeta) & -\frac{ia}{c}\zeta^c \beta(a, c, \zeta) \\ -\frac{ia}{c^*}\zeta^{c^*} \beta(a, c, \zeta)^* & \alpha(a, c, \zeta)^* \end{bmatrix}. \tag{A2}$$

Here the initial condition is $U(-\infty, -\infty) = \mathbb{I}$, though in practice we take the initial time to be some finite value that is sufficiently large for our results to converge. Since we are only interested in the end result of the pulse, we consider the evolution operator at $t = +\infty$,

$$\begin{aligned} U &= U(+\infty, -\infty) \\ &= \begin{bmatrix} {}_2F_1\left(-a, a; \frac{\sigma - i\Delta}{2\sigma}; 1\right) & -i\mathrm{sech}\left(\frac{\pi\Delta}{2\sigma}\right)\sin(a\pi) \\ -i\mathrm{sech}\left(\frac{\pi\Delta}{2\sigma}\right)\sin(a\pi) & {}_2F_1\left(-a, a; \frac{i\Delta + \sigma}{2\sigma}; 1\right) \end{bmatrix}. \end{aligned} \tag{A3}$$

Then it is clear that for $a \in \mathbb{Z}$, the evolution operator is diagonal. In this instance, we can express the evolution operator with $U = \mathrm{diag}\{e^{-i\phi_a}, e^{+i\phi_a}\}$. If we use a 2π pulse (i.e., $a = 1$), then $\phi_1(\Delta) = 2\arctan(\frac{\sigma}{\Delta})$. If instead we consider a 4π pulse, then $a = 2$ and $\phi_2(\Delta) = 2\arctan[\frac{4\Delta/\sigma}{(\Delta/\sigma)^2 - 3}]$.

APPENDIX B: LOCAL INVARIANTS

To develop our CPHASE gates, we will use the local invariants [38,39] of unitary operations in SU(4), denoted here as G_i for $i \in \{1, 2, 3\}$. These are three quantities that may be computed from any element of SU(4) and are invariant under operations in SU(2). That is, for $\mathcal{U} \in \mathrm{SU}(4)$ and $V_i \in \mathrm{SU}(2)$,

$$G_i(V_1 \otimes V_2\, \mathcal{U}\, V_3 \otimes V_4) = G_i(\mathcal{U}). \tag{B1}$$

Therefore, the local invariants convey the nonlocal properties of the operator $\mathcal{U}$ and give a unique representation of any class of two-qubit gates that are equivalent up to local operations. Note that single-qubit rotations about the Z axis may be efficiently performed for transmons [40] and that hyperbolic secant pulses can produce no population transfer, as discussed in Appendix A and in Ref. [36]. These two facts allow us to develop protocols for CPHASE operations that only consider their nonlocal characteristics and have no overhead in terms of the fidelity or time required to perform the single-qubit gates associated with their local-equivalence classes. Hence, the local invariants of the analytical unitary evolution, U, for our four-level system driven by a hyperbolic secant pulse will be the starting place for constructing our protocols for a CPHASE gate. The quantities G_i are obtained by first placing U in the magic basis [48] defined by the unitary transformation [38,39]

$$Q = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & i \\ 0 & i & 1 & 0 \\ 0 & i & -1 & 0 \\ 1 & 0 & 0 & -i \end{pmatrix}. \tag{B2}$$

The local invariants are the coefficients of the characteristic polynomial of the matrix $M(U) = (Q^\dagger U Q)^T (Q^\dagger U Q)$, and they are given by the following expressions:

$$\begin{aligned} G_1 &= \mathrm{Re}\frac{\mathrm{tr}[M(U)]^2}{16 \det U}, \\ G_2 &= \mathrm{Im}\frac{\mathrm{tr}[M(U)]^2}{16 \det U}, \\ G_3 &= \frac{\mathrm{tr}[M(U)]^2 - \mathrm{tr}[M^2(U)]}{4 \det U}. \end{aligned} \tag{B3}$$

[1] L. Frunzio, A. Wallraff, D. Schuster, J. Majer, and R. Schoelkopf, Fabrication and characterization of superconducting circuit QED devices for quantum computation, IEEE Trans. Appl. Supercond. **15**, 860 (2005).

[2] J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Coupling superconducting qubits via a cavity bus, Nature **449**, 443 (2007).

[3] 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).

[4] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, Phys. Rev. X **8**, 011021 (2018).

[5] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, Scalable Quantum Simulation of Molecular Energies, Phys. Rev. X **6**, 031007 (2016).

[6] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[7] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, Rev. Mod. Phys. **86**, 153 (2014).

[8] X. Gu, A. F. Kockum, A. Miranowicz, Y.-x. Liu, and F. Nori, Microwave photonics with superconducting quantum circuits, Phys. Rep. **718-719**, 1 (2017).

[9] L. S. Theis, F. Motzoi, and F. K. Wilhelm, Simultaneous gates in frequency-crowded multilevel systems using fast, robust, analytic control shapes, Phys. Rev. A **93**, 012324 (2016).

[10] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Demonstration of two-qubit algorithms with a superconducting quantum processor, Nature **460**, 240 (2009).

[11] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Preparation and measurement of three-qubit entanglement in a superconducting circuit, Nature **467**, 574 (2010).

[12] J. Ghosh, A. Galiautdinov, Z. Zhou, A. N. Korotkov, J. M. Martinis, and M. R. Geller, High-fidelity controlled-σ^z gate for resonator-based superconducting quantum computers, Phys. Rev. A **87**, 022309 (2013).

[13] J. M. Martinis and M. R. Geller, Fast adiabatic qubit gates using only σ_z control, Phys. Rev. A **90**, 022307 (2014).

[14] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, Superconducting quantum circuits at the surface code threshold for fault tolerance, Nature **508**, 500 (2014).

[15] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, Universal Gate for Fixed-Frequency Qubits Via a Tunable Bus, Phys. Rev. Appl. **6**, 064007 (2016).

[16] M. Roth, M. Ganzhorn, N. Moll, S. Filipp, G. Salis, and S. Schmidt, Analysis of a parametrically driven exchange-type gate and a two-photon excitation gate between superconducting qubits, Phys. Rev. A **96**, 062323 (2017).

[17] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, Qubit Architecture with High Coherence and Fast Tunable Coupling, Phys. Rev. Lett. **113**, 220502 (2014).

[18] B. Royer, A. L. Grimsmo, N. Didier, and A. Blais, Fast and high-fidelity entangling gate through parametrically modulated longitudinal coupling, Quantum **1**, 11 (2017).

[19] A. Galiautdinov, Generation of high-fidelity controlled-NOT logic gates by coupled superconducting qubits, Phys. Rev. A **75**, 052303 (2007).

[20] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits, Phys. Rev. Lett. **107**, 080502 (2011).

[21] 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).

[22] S. E. Economou and E. Barnes, Analytical approach to swift nonleaky entangling gates in superconducting qubits, Phys. Rev. B **91**, 161405(R) (2015).

[23] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, Phys. Rev. A **93**, 060302(R) (2016).

[24] H. Paik, A. Mezzacapo, M. Sandberg, D. T. McClure, B. Abdo, A. D. Córcoles, O. Dial, D. F. Bogorin, B. L. T. Plourde, M. Steffen, A. W. Cross, J. M. Gambetta, and J. M. Chow, Experimental Demonstration of a Resonator-Induced Phase Gate in a Multiqubit Circuit-Qed System, Phys. Rev. Lett. **117**, 250502 (2016).

[25] 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).

[26] J. L. Allen, R. Kosut, J. Joo, P. Leek, and E. Ginossar, Optimal control of two qubits via a single cavity drive in circuit quantum electrodynamics, Phys. Rev. A **95**, 042325 (2017).

[27] E. Barnes, C. Arenz, A. Pitchford, and S. E. Economou, Fast microwave-driven three-qubit gates for cavity-coupled superconducting qubits, Phys. Rev. B **96**, 024504 (2017).

[28] T. Noh, G. Park, S.-G. Lee, W. Song, and Y. Chong, Construction of controlled-NOT gate based on microwave-activated phase (MAP) gate in two transmon system, Sci. Rep. **8**, 13598 (2018).

[29] E. Magesan and J. M. Gambetta, Effective Hamiltonian models of the cross-resonance gate, arXiv:1804.04073.

[30] D. J. Egger, M. Ganzhorn, G. Salis, A. Fuhrer, P. Müller, P. K. Barkoutsos, N. Moll, I. Tavernelli, and S. Filipp, Entanglement Generation in Superconducting Qubits using Holonomic Operations, Phys. Rev. Appl. **11**, 014017 (2019).

[31] V. Tripathi, M. Khezri, and A. N. Korotkov, Operation and intrinsic error budget of two-qubit cross-resonance gate, Phys. Rev. A **100**, 012301 (2019).

[32] J. L. Allen, R. Kosut, and E. Ginossar, Minimal time robust two qubit gates in circuit QED, arXiv:1902.08056.

[33] S. P. Premaratne, J.-H. Yeh, F. C. Wellstood, and B. S. Palmer, Implementation of a generalized controlled-NOT gate between fixed-frequency transmons, Phys. Rev. A **99**, 012317 (2019).

[34] N. Rosen and C. Zener, Double Stern-Gerlach experiment and related collision phenomena, Phys. Rev. **40**, 502 (1932).

[35] S. E. Economou, L. J. Sham, Y. Wu, and D. G. Steel, Proposal for optical u(1) rotations of electron spin trapped in a quantum dot, Phys. Rev. B **74**, 205415 (2006).

[36] S. E. Economou, High-fidelity quantum gates via analytically solvable pulses, Phys. Rev. B **85**, 241401(R) (2012).

[37] 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).

[38] Y. Makhlin, Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations, Quant. Info. Proc. **1**, 243 (2002).

[39] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Geometric theory of nonlocal two-qubit operations, Phys. Rev. A **67**, 042313 (2003).

[40] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient Z gates for quantum computing, Phys. Rev. A **96**, 022330 (2017).

[41] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, Phys. Rev. A **76**, 042319 (2007).

[42] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions: Basic Processes and Applications* (J. Wiley, New York, 1992).

[43] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, Phys. Lett. A **367**, 47 (2007).

[44] 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).

[45] J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm, Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator, Phys. Rev. A **83**, 012308 (2011).

[46] F. Motzoi and F. K. Wilhelm, Improving frequency selection of driven pulses using derivative-based transition suppression, Phys. Rev. A **88**, 062318 (2013).

[47] F. Sauvage and F. Mintert, Optimal quantum control with poor statistics, arXiv:1909.01229.

[48] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction.Phys. Rev. A **54**, 3824 (1996).

Chapter 4

A universal quantum gate set for transmon qubits with strong ZZ interactions

This chapter includes pages from our publication [4]

J. Long, T. Zhao, M. Bal, R. Zhao, G. S. Barron, H.-s. Ku, J. A. Howard, X. Wu, C. R. H. McRae, X.-H. Deng, G. J. Ribeill, M. Singh, T. A. Ohki, E. Barnes, S. E. Economou, and D. P. Pappas, "A universal quantum gate set for transmon qubits with strong ZZ interactions", arXiv:2103.12305 (2021)

This work has been posted to arXiv.org.

The contributions to this work are as follows: My contributions were assisting in the theoretical work relating to the SWIPHT protocol. J.L. and T.Z. primarily carried out this work. M.B., R.Z., H.K., J.A.H., X.W., C.R.H.M., G.J.R., M.S., and T.A.O. performed experimental fabrication, characterization, and control. E.B., S.E.E. and D.P.P. supervised.

arXiv:2103.12305v1 [quant-ph] 23 Mar 2021

A universal quantum gate set for transmon qubits with strong ZZ interactions

Junling Long,[1,*] Tongyu Zhao,[1,†] Mustafa Bal,[1] Ruichen Zhao,[1] George S. Barron,[2] Hsiang-sheng Ku,[3] Joel A. Howard,[4] Xian Wu,[1] Corey Rae H. McRae,[1] Xiu-Hao Deng,[5] Guilhem J. Ribeill,[6] Meenakshi Singh,[4] Thomas A. Ohki,[6] Edwin Barnes,[2] Sophia E. Economou,[2] and David P. Pappas[3]

[1] *Department of Physics, University of Colorado, Boulder, Colorado 80309, USA*
[2] *Department of Physics, Virginia Tech, Blacksburg, VA 24061, USA*
[3] *National Institute of Standards and Technology, Boulder, Colorado 80305, USA*
[4] *Department of Physics, Colorado School of Mines, Golden, Colorado 80401, USA*
[5] *Shenzhen Institute of Quantum Science and Engineering, Southern University of Science and Technology, Shenzhen, Guangdong 518055, China*
[6] *Quantum Engineering and Computing, Raytheon BBN Technologies, Cambridge, MA 02138, USA.*

(Dated: March 24, 2021)

High-fidelity single- and two-qubit gates are essential building blocks for a fault-tolerant quantum computer. While there has been much progress in suppressing single-qubit gate errors in superconducting qubit systems, two-qubit gates still suffer from error rates that are orders of magnitude higher. One limiting factor is the residual ZZ-interaction, which originates from a coupling between computational states and higher-energy states. While this interaction is usually viewed as a nuisance, here we experimentally demonstrate that it can be exploited to produce a universal set of fast single- and two-qubit entangling gates in a coupled transmon qubit system. To implement arbitrary single-qubit rotations, we design a new protocol called the two-axis gate that is based on a three-part composite pulse. It rotates a single qubit independently of the state of the other qubit despite the strong ZZ-coupling. We achieve single-qubit gate fidelities as high as 99.1% (for $|\widetilde{01}\rangle \leftrightarrow |\widetilde{11}\rangle$ transition) from randomized benchmarking measurements. We then demonstrate both a CZ gate and a CNOT gate. Because the system has a strong ZZ-interaction, a CZ gate can be achieved by letting the system freely evolve for a gate time $t_g = 53.8\,\text{ns}$. To design the CNOT gate, we utilize an analytical microwave pulse shape based on the SWIPHT protocol for realizing fast, low-leakage gates. We obtain fidelities of 94.6% and 97.8% for the CNOT and CZ gates respectively from quantum progress tomography.

PACS numbers: 03.67.Lx, 42.50.-p, 42.50.Gy, 42.50.Pq

I. INTRODUCTION

Superconducting qubits hold promise as the main building blocks of a future fault-tolerant quantum computer [1–3]. In recent years, advances in both coherence times [4] and quantum gate fidelities [5–7], together with the inherent scalability of superconducting circuit architectures, have enabled demonstrations of superconducting quantum computers comprised of tens of qubits [8–10]. While single-qubit gate errors are routinely on the order of 0.01% [11–13], two-qubit gates across different qubit architectures have exhibited higher error rates, ranging from 0.1% up to the few-percent level [1, 4–6, 14–19]. Thus, two-qubit gates remain a bottleneck in the performance of existing devices.

One significant source of two-qubit gate errors is the residual ZZ interaction between coupled qubits [6]. The ZZ interaction, also known as ZZ coupling or crosstalk, originates from a mixing between computational states and higher-energy states of the coupled qubits. It shifts the resonance frequency of one qubit conditionally on the state of another qubit. As a result, spurious phases are accumulated during gates, reducing fidelities. Some efforts have been devoted to suppressing the ZZ coupling using tunable couplers [20, 21] or hybrid superconducting qubits [22].

With the recent progress in tunable couplers [18, 20, 21], new opportunities open up for scalable architectures. In this context, it has been realized and demonstrated that the ZZ interaction can in fact be exploited for the implementation of a maximally entangling CZ gate [18]. This approach is very promising for scalability, as the qubits can be brought into the strong ZZ coupling regime to generate fast gates, while in the idling regime crosstalk is suppressed. Since tuning can be costly in terms of time and possibly coherence, it would be beneficial for the performance of the device to maximize the use of this strongly coupled regime by developing additional quantum gates that can be implemented within it. For example, a CNOT gate, despite being locally equivalent to a CZ, can be directly implemented using microwave fields [23, 24], eliminating the need for additional single-qubit rotations. The development of a broader toolset of gates applicable in the nonzero ZZ coupling regime is also important for cases where this coupling cannot be completely switched off, which is the generic situation in superconducting circuits.

In this work, we develop and demonstrate a universal set of fast gates that operate in the regime of strong

* Junling.Long@Colorado.edu
† Tongyu.Zhao@Colorado.edu

ZZ interactions between capacitively coupled superconducting transmon qubits. To implement arbitrary single-qubit gates in this regime, we introduce a family of three-part composite pulses that effectively freeze the ZZ coupling so that one qubit is rotated independently of the state of another. We call these two-axis gates (TAGs). Randomized benchmarking is performed on each qubit to characterize the TAG fidelity, which we find to be 99.1%. We also demonstrate two types of maximally entangling gates: CZ and CNOT. The CZ gate is implemented via free evolution under the ZZ interaction, yielding a gate time of 53.8 ns. For the CNOT gate, we employ the SWIPHT protocol [23–25], which amounts to cleverly designing the pulse such that it acts on two transitions (the target and the nearest competing transition) in such a way that the net effect is the implementation of the target gate. This allows for a much faster gate compared to resolving the transitions. We create a maximally entangled state with the CNOT gate and obtain a fidelity of 98.2%, confirmed using quantum state tomography. Finally, we characterize the CZ and CNOT gates using quantum process tomography. The measured average gate fidelity is about 94.6% (97.8%) for the CNOT (CZ) gate.

This paper is organized as follows. In Sec. II, we provide details about the transmon device and the effective ZZ Hamiltonian. Sec. III presents our two-axis gate designs for implementing arbitrary single-qubit rotations in the strong ZZ coupling regime. In Sec. IV, we describe how we implement our CZ and CNOT gates, including a review of the SWIPHT protocol we use for the latter. Our randomized benchmarking measurements for the two-axis gates are presented in Sec. V, while our quantum state tomography and process tomography results for our two-qubit entangling gates are shown in Secs. VI and VII, respectively. We conclude in Sec. VIII. Three appendices contain additional information about the spectrum of the system and relaxation and dephasing times, along with a detailed derivation of the effective ZZ interaction for both capacitively coupled and resonator-coupled transmons.

II. DEVICE

The sample chip is shown in Fig. 1 (a) and (b), where optical images of the full device and a zoomed-in view of the two qubits (red and blue) are shown, respectively. The device-under-test (DUT) consists of two floating transmon qubits [26] capacitively coupled to each other. Each transmon qubit is coupled to a coplanar waveguide readout resonator. The two readout resonators are coupled to the feedline in a hanger-style. Readout and control signals are all sent through the feedline. The capacitive components of each transmon are comprised of two identical rectangular pads [red (blue) for the first (second) qubit, as shown in Fig. 1(b)]. Each pair of pads is connected by an Al/AlO_x/Al Josephson junction that is fabricated with an overlap technique [27]. The full chip (except for the Josephson junctions) is made of 100 nm thick Niobium superconducting films. In this particular DUT, the two transmon qubits have frequencies of $\omega_1/2\pi = 5.075\,\mathrm{GHz}$ and $\omega_2/2\pi = 5.310\,\mathrm{GHz}$, and anharmonicities of $\alpha_1/2\pi = -260\,\mathrm{MHz}$ and $\alpha_2/2\pi = -340\,\mathrm{MHz}$ (Appendix A). The relaxation and coherence times of the two qubits are measured to be $T_1^{(1)} = 76.98\,\mu\mathrm{s}$, $T_2^{*(1)} = 50.65\,\mu\mathrm{s}$, $T_1^{(2)} = 79.71\,\mu\mathrm{s}$, and $T_2^{*(2)} = 17.09\,\mu\mathrm{s}$ (see Appendix B). The equivalent grounded circuit model of the two capacitively coupled transmon qubits is shown in Fig. 1(c). The Gaussian elimination method described in Ref. [28] was applied to convert the full floating circuit model to this simplified grounded circuit model. Note that the circuit of the readout resonators is not shown.

The level diagram of the DUT is shown in Fig. 1(d). Two-qubit states are labeled as $|mn\rangle$, where m and n in the ket represent the mth state of the first qubit and nth state of the second qubit, respectively. The gray dashed lines in Fig. 1(d) are the uncoupled two-qubit states, while the solid lines are the dressed two-qubit states with the always-on capacitive coupling. We only retain states involving up to two total excitations. The coupling between the $|10\rangle$ and $|01\rangle$ states is given by g_1. This coupling causes the two states to repel each other. As a result, the dressed qubit frequency shifts down (up) compared to the uncoupled qubit frequency of the first (second) qubit.

The couplings in the two-excitation manifold are between the $|20\rangle \leftrightarrow |11\rangle$, and $|11\rangle \leftrightarrow |02\rangle$ states, with the same strength, $g_2 \approx \sqrt{2}g_1$. As a result of the two couplings, the dressed energy level $\widetilde{|11\rangle}$ is shifted up from the uncoupled $|11\rangle$ state. This makes the frequencies of the transitions $\widetilde{|01\rangle} \leftrightarrow \widetilde{|11\rangle}$ and $\widetilde{|10\rangle} \leftrightarrow \widetilde{|11\rangle}$ higher than those of $\widetilde{|00\rangle} \leftrightarrow \widetilde{|10\rangle}$ and $\widetilde{|00\rangle} \leftrightarrow \widetilde{|01\rangle}$, respectively, i.e., this is a state-dependent shift in qubit frequencies. In the computational subspace, defined by the dressed qubit states $\widetilde{|00\rangle}$, $\widetilde{|01\rangle}$, $\widetilde{|10\rangle}$, and $\widetilde{|11\rangle}$, the two-qubit system is described by the following Hamiltonian (see Appendix C),

$$H_q/\hbar = (\omega_1 + \frac{g_z}{2})\frac{\sigma_z^{(1)}}{2} + (\omega_2 + \frac{g_z}{2})\frac{\sigma_z^{(2)}}{2} + g_z\frac{\sigma_z^{(1)}}{2}\frac{\sigma_z^{(2)}}{2}, \tag{1}$$

where $\sigma_z^{(i)} = (\widetilde{|1\rangle}\widetilde{\langle 1|} - \widetilde{|0\rangle}\widetilde{\langle 0|})$ is the atomic inversion operator for the i^{th} qubit; ω_1 (ω_2) is the frequency of the first (second) qubit while the other qubit is in its ground state, i.e., it is the transition frequency of $\widetilde{|00\rangle} \leftrightarrow \widetilde{|10\rangle}$ ($\widetilde{|00\rangle} \leftrightarrow \widetilde{|01\rangle}$); g_z is the effective coupling strength determined by the amount of the up-shift of state $\widetilde{|11\rangle}$. We obtain a value of $g_z/2\pi = 9.29$ MHz from spectroscopy measurements (see Appendix A). As shown in Eq. (1), the coupling between the two qubits in the computational subspace (dressed basis) is ZZ, but it originates from the

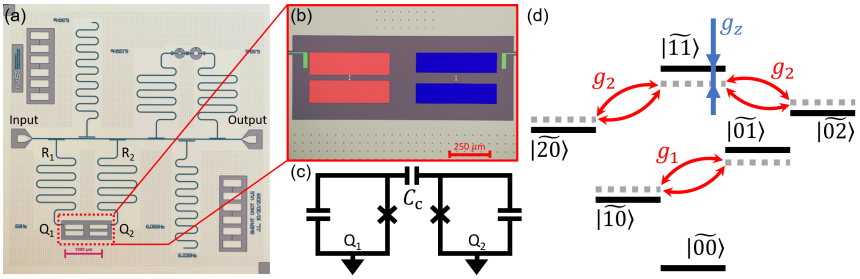

FIG. 1. (a) Optical micrograph of the full chip including qubits, readout resonators, test Josephson junctions, and test resonators. (b) Zoomed-in view of the two floating qubits (DUT) used in this experiment. Each qubit consists of two identical pads (red for the left qubit and blue for the right qubit) and a Josephson junction connecting the two pads. Each qubit is coupled to its own readout resonator (green). (c) Grounded circuit model of the capacitively coupled DUT. See Ref. [28] for how to obtain the equivalent grounded circuit model of floating qubits. (d) Energy level diagram of the two-qubit system. Gray dashed lines are uncoupled two-qubit energy levels. Solid lines represent the diagonalized energy levels of the two-qubit system with the coupling on. The red bidirectional curved arrows denote the swapping couplings between excitation-preserved states, and the coupling strengths are given by g_1 and g_2. g_z denotes the shift of $|11\rangle$ due to the couplings.

transverse XX coupling between level $|11\rangle$ and the higher levels $|20\rangle$ and $|02\rangle$. As shown in Appendix C, Eq. 1 can originate from either a direct, capacitive coupling or a resonator-mediated coupling between two transmons.

III. ARBITRARY SINGLE-QUBIT ROTATIONS

Performing arbitrary single-qubit rotations is not a trivial task in a two-qubit system with strong ZZ coupling because the frequency of one qubit is conditional on the state of the other qubit. Here, we design a new type of single-qubit gate, called the two-axis gate (TAG), that unconditionally drives arbitrary single-qubit rotations in this ZZ-coupled, two-qubit system.

To explain how TAGs work, we take a $X(\pi/2)$ rotation on the first qubit as an example. As shown in Fig. 2(a), TAGs are implemented with three-part composite pulses. Note that the TAG protocol can be implemented using any pulse shapes for each of the three pieces. Here, we use square pulses in our experiments to facilitate pulse tune up. The first and third pulses (red) are identical, and we refer to them as Bloch sphere rotation (BR) pulses. The second pulse (blue) is called the qubit rotation (QR) pulse. For our $X(\pi/2)$ rotation example, the center frequency of the TAG drive is set to be on resonance with the $|\widetilde{00}\rangle \leftrightarrow |\widetilde{10}\rangle$ transition. In the frame rotating at this frequency, the rotation axes for the three pulses of the TAG are all along the x-axis, as shown in Fig. 2(b). The pulse strengths and durations are chosen such that all three pulses combine to rotate the state vector around the x-axis by $\pi/2$ and drive the original state $|\widetilde{00}\rangle$ to the final state $|-y\rangle = (|\widetilde{00}\rangle - i|\widetilde{10}\rangle)/\sqrt{2}$.

Since the center frequency is detuned from the $|\widetilde{01}\rangle \leftrightarrow |\widetilde{11}\rangle$ transition by g_z as described by Eq. (1), the rotation axes of the BR and QR are both tilted from the x-axis in the xz plane by ϕ_{BR} and ϕ_{QR}, respectively, as shown in Fig. 2(c). The BR pulses can be treated as a rotation of the Bloch sphere around the blue axis in the opposite direction. This is the reason why we call it a Bloch sphere rotation. When the first BR pulse is applied, we require it to rotate the Bloch sphere such that the x-axis after the rotation coincides with the QR axis, i.e., the blue axis in Fig. 2(c). To achieve this, two conditions need to be satisfied. First, the red axis needs to be the bisector of the angle formed between the blue- and x-axis, i.e., $\phi_{QR} = 2\phi_{BR}$. Second, the rotation angle of the BR pulse needs to be π or $-\pi$. In the rotated Bloch sphere, since the QR (blue) axis is on the x-axis, we can perform a rotation around the x-axis by $\pi/2$. After the QR pulse, we apply the BR pulse again to rotate the Bloch sphere around the BR (red) axis by another π angle to get back to the original Bloch sphere. Treating a BR pulse as a rotation of the Bloch sphere is a good way to understand the TAG. If we treat all the pulses in the TAG as rotations of a state vector, we can see the trajectories on both Bloch spheres. This is shown in Fig. 2(c). When the first BR pulse is applied, the rotated state vector becomes perpendicular to the blue axis. Then, the QR pulse is applied to rotate the state vector by $\Theta = \pi/2$. Finally, the BR pulse is applied again to bring the state to $|-y\rangle$.

Similarly, we can design TAGs that rotate the state of one qubit around the x-axis by other arbitrary angles. To perform a y-rotation with the TAG, one can just add a 90° global phase shift to the TAG. Arbitrary rotations can be performed by combinations of x and y-rotations.

IV. TWO-QUBIT ENTANGLING GATES

In this section, we describe how we realize two different two-qubit entangling gates, a CZ gate and a generalized CNOT gate, on our device.

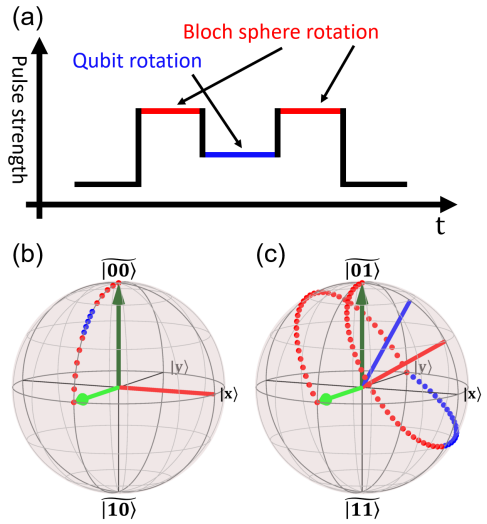

FIG. 2. (a) A square shaped two-axis gate that implements a $X(\pi/2)$ on the first qubit. (b) and (c) show the trajectories of the state vector for the $|\widetilde{00}\rangle \leftrightarrow |\widetilde{10}\rangle$ and $|\widetilde{01}\rangle \leftrightarrow |\widetilde{11}\rangle$ transitions, respectively. In (b) and (c) the initial state is $|\widetilde{00}\rangle$, as denoted by the dark green arrow, while the final state is $|-y\rangle = (|\widetilde{00}\rangle - i|\widetilde{10}\rangle)/\sqrt{2}$, which is denoted by the light green arrow. The red (blue) dots in (a) and (c) denote the trajectories corresponding to the BR (QR) pulses.

A. A controlled-Z gate by free evolution

With the ZZ coupling, free evolution of the system is, in fact, a controlled-phase gate. By transforming the Hamiltonian in Eq. (1) to the rotating frame of ω_1 and ω_2 for the first and second qubit, respectively, and adding a constant energy shift $g_z/4$, one obtains

$$H_q^{rot}/\hbar = g_z|\widetilde{11}\rangle\langle\widetilde{11}|. \tag{2}$$

Note that we used the fact that $\sigma_z^{(i)}$ is defined in the dressed two-qubit basis to obtain this result. Also note that adding a constant energy shift to the Hamiltonian only produces a global phase; it does not change the dynamics of the two qubits. The time evolution operator of this Hamiltonian above is given by

$$\begin{aligned} U(t) &= e^{-iH_q^{rot}t/\hbar} = e^{-ig_z t|\widetilde{11}\rangle\langle\widetilde{11}|} \\ &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-ig_z t} \end{pmatrix}. \end{aligned} \tag{3}$$

This is a standard controlled-phase gate with the phase given by $-g_z t$. To realize a CZ gate, we can just let the system evolve for time $t = \frac{\pi}{g_z}$, in which case the evolution operator becomes

$$U_{\rm CZ} = U\left(\frac{\pi}{g_z}\right) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \tag{4}$$

The CZ gate time is given by $\frac{\pi}{g_z} = 53.8$ ns. We analyze the performance of this gate using quantum process tomography in Sec. VII. We note that a similar CZ gate based on free evolution in the strong ZZ coupling regime, achieved via a tunable coupler, was recently demonstrated [18].

B. A generalized controlled-NOT gate based on the SWIPHT protocol

In addition to the free-evolution CZ gate, it is also possible to implement a CNOT gate in the strong ZZ coupling regime using a shaped microwave pulse. A protocol known as SWIPHT (speeding up wave forms by inducing phases to harmful transitions), provides a general framework for designing fast, high-fidelity two-qubit entangling gates [23–25]. The SWIPHT protocol is based on reverse-engineering analytical pulse shapes that implement two-qubit entangling gates [29]. SWIPHT was first demonstrated experimentally for a two-transmon system in the weak ZZ coupling regime [30]. The SWIPHT protocol was originally proposed for two transmon qubits coupled with a bus resonator, but it applies to any two-qubit system with ZZ coupling in the dressed qubit basis. Here, we employ SWIPHT to implement a generalized CNOT gate for our two-qubit system.

Fig. 3(a) shows the energy levels in the computational subspace of the two-qubit system. A CNOT gate can be viewed as effectively swapping the $|\widetilde{00}\rangle$ and $|\widetilde{01}\rangle$ states while leaving the other logical states alone. A natural way to implement this is to apply a resonant microwave drive to perform a π rotation on the target transition $|\widetilde{00}\rangle \leftrightarrow |\widetilde{01}\rangle$. However, there exists a nearby harmful transition $|\widetilde{10}\rangle \leftrightarrow |\widetilde{11}\rangle$ [red in Fig. 3(a)] that is detuned from the target transition by g_z, so it would be driven off-resonantly. As a result, one would have to make the gate time much longer than $2\pi/g_z$ to selectively drive the target transition. This is not optimal, given the limited coherence times of the qubits. Instead of attempting to avoid the harmful transition, the SWIPHT CNOT gate is designed to purposely drive it through a trivial cyclic evolution while driving a π rotation on the target transition. This is shown in Fig. 3(b). A cyclic evolution on a two-level system can be described by a Z rotation, $e^{-i\varphi\sigma_z^{(2)}/2}$, where φ is the phase accumulated during the cyclic evolution. We can express the SWIPHT CNOT

gate in the dressed two-qubit basis as

$$U_{\rm CNOT} = \widetilde{|0\rangle}\widetilde{\langle 0|} \otimes e^{-i\pi\sigma_x^{(2)}/2} + \widetilde{|1\rangle}\widetilde{\langle 1|} \otimes e^{-i\varphi\sigma_z^{(2)}/2} = \begin{pmatrix} 0 & -i & 0 & 0 \\ -i & 0 & 0 & 0 \\ 0 & 0 & e^{-i\frac{\varphi}{2}} & 0 \\ 0 & 0 & 0 & e^{i\frac{\varphi}{2}} \end{pmatrix}. \quad (5)$$

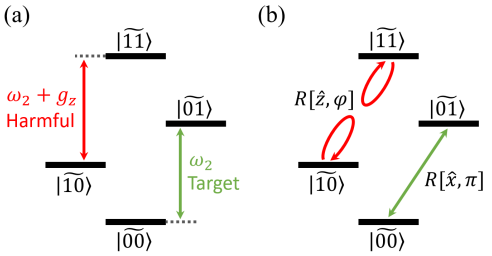

FIG. 3. (a) The computational subspace of the two transmon qubit system with strong ZZ coupling. The frequency of the target transition (green) $\widetilde{|00\rangle} \leftrightarrow \widetilde{|01\rangle}$ is ω_2. The harmful transition (red) $\widetilde{|10\rangle} \leftrightarrow \widetilde{|11\rangle}$ is detuned from the target transition by g_z. (b) The SWIPHT pulse (see Fig. 4) is on resonance with the target transition and induces a π rotation around the x-axis in the two-level target subspace. At the same time, the harmful transition is off-resonantly driven by the SWIPHT pulse and undergoes cyclic evolution, acquiring trivial phases.

Compared to a standard CNOT gate, the SWIPHT CNOT gate described above has some extra phases, which is why we refer to it as a generalized CNOT gate [23]. This generalized CNOT gate is maximally entangling and is related to the standard CNOT by single-qubit Z rotations and a global phase.

According to the SWIPHT protocol [23], analytical pulse shapes that implement the operation given in Eq. (5) can be generated from the formula

$$\Omega(t) = \frac{\ddot{\chi}}{\sqrt{\frac{g_z^2}{4} - \dot{\chi}^2}} - 2\sqrt{\frac{g_z^2}{4} - \dot{\chi}^2}\cot(2\chi). \quad (6)$$

Any real function $\chi(t)$ obeying the constraint $|\dot{\chi}| \leq g_z/2$ and the boundary conditions $\chi(0) = \chi(t_g) = \pi/4$, $\dot{\chi}(0) = \dot{\chi}(t_g) = 0$, where t_g is the gate time, produces a pulse wave form $\Omega(t)$ that performs a generalized CNOT. One may notice a factor of two difference from the original expression in Ref. [23]; This is due to different definitions of the Rabi strength. Here, we choose $\chi(t)$ as in Ref. [23]:

$$\chi(t) = A\left(\frac{t}{t_g}\right)^4\left(1 - \frac{t}{t_g}\right)^4 + \frac{\pi}{4}, \quad (7)$$

where A and t_g are given by

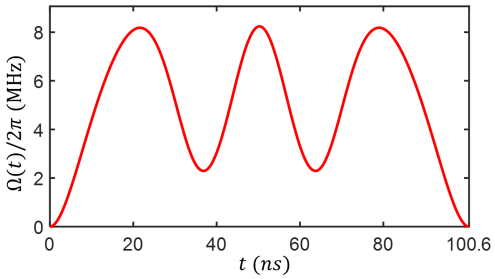

FIG. 4. Pulse shape that implements the SWIPHT CNOT gate. The gate time is about 100.6 ns.

$$t_g = \frac{5.87}{g_z}, \quad A = 138.9. \quad (8)$$

Given $g_z = 9.29$ MHz, the resulted gate time is $t_g = 100.6$ ns. The pulse shape of the SWIPHT gate is plotted in Fig. 4. The performance of our CNOT gate is evaluated in Secs. VI and VII through quantum state and process tomography.

V. RANDOMIZED BENCHMARKING OF THE TWO-AXIS, SINGLE-QUBIT GATE

To characterize the TAG, we perform the randomized benchmarking (RB) protocol described in Ref. [31] with assumption that errors are gate independent. The RB experiment starts with the qubit initialized in the ground state. Various $\pi/2$ gates and π gates are randomly chosen and concatenated together to form a sequence. A comparison of the measured result of these operations with what would be expected from an ideal system can then be used to evaluate performance.

The $\pi/2$ gates are Clifford group generators $e^{\pm i\sigma_u\pi/4}$, with $u = x, y$. The π gates are chosen from $e^{\pm i\sigma_b\pi/2}$ with $b = 0, x, y, z$, where we have defined σ_0 to be the identity operator. The RB sequence is truncated at different numbers of $\pi/2 - \pi$ gate pairs and applied to the qubit. Each truncation is followed by a $\pi/2$ or π gate to always bring the qubit to a certain eigenstate of σ_z. For convenience, we typically use the ground state in our experiments. Then, the probability of the ground state, alternatively referred to as the sequence fidelity, is measured for each truncation of the RB sequence. The above process is repeated for many different RB sequences. Finally, the sequence fidelity $\mathcal{F}$ is averaged over all of the sequences, thereby becoming a function of only the truncation, which is quantified by the number of $\pi/2$ gates. $\mathcal{F}$ decays exponentially, and the decay rate is determined by the average gate infidelity [31]. Since the TAG is designed to address two transitions for each qubit (one where the other qubit is in the ground state, and one

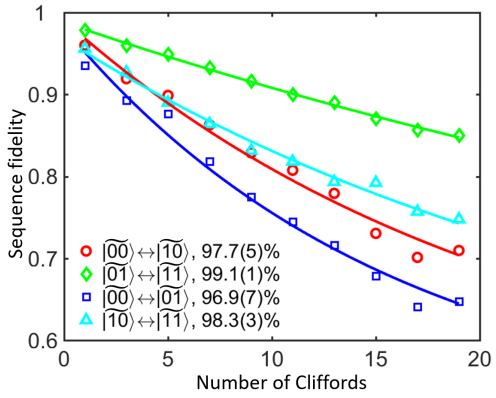

FIG. 5. Randomized benchmarking of the two-axis, single-qubit gate. Average sequency fidelity vs. the number of Cliffords/$\frac{\pi}{2}$ gates for the transition $|\widetilde{00}\rangle \leftrightarrow |\widetilde{10}\rangle$ (red circle), $|\widetilde{01}\rangle \leftrightarrow |\widetilde{11}\rangle$ (green diamond), $|\widetilde{00}\rangle \leftrightarrow |\widetilde{01}\rangle$ (blue square), and $|\widetilde{10}\rangle \leftrightarrow |\widetilde{11}\rangle$ (cyan triangle). Solid curves are the exponential fittings to the data to extract an average gate fidelity for each transition.

where it is in the excited state), we run the RB experiment on each transition separately for each qubit. This measures the average fidelity of the TAG as it acts on the corresponding transition. For each RB experiment, we choose 5 random $\pi/2$ gate sequences and 8 random π gate sequences, yielding a total of 40 different RB sequences. Each sequence is truncated at every other pair of π and $\pi/2$ gates. For each truncation, 1000 copies are measured to obtain the sequence fidelity $\mathcal{F}$. The experimental results are plotted in Fig. 5 for all four transitions of the two qubits. The average gate fidelities of the TAG are 97.7(5)%, 99.1(1)%, 96.9(7)%, and 98.3(3)% for the transitions $|\widetilde{00}\rangle \leftrightarrow |\widetilde{10}\rangle$, $|\widetilde{01}\rangle \leftrightarrow |\widetilde{11}\rangle$, $|\widetilde{00}\rangle \leftrightarrow |\widetilde{01}\rangle$, and $|\widetilde{10}\rangle \leftrightarrow |\widetilde{11}\rangle$, respectively. Since each TAG has to be calibrated simultaneously for two transitions of a qubit, we believe there are some coherence errors in the TAGs due to imperfect calibration. The gate fidelities of the two transitions of the second qubit are lower than those of the first qubit. This is consistent with the fact that T_2^* of the second qubit is only half that of the first qubit.

VI. QUANTUM STATE TOMOGRAPHY OF A MAXIMALLY ENTANGLED STATE

The core functionality of two-qubit gates is to entangle two qubits. We first verify this functionality by preparing a maximally entangled state and performing quantum state tomography on it. Since the CNOT gate is more straightforward in terms of generating entangled states, we present experimental results of an entangled state generated with the SWIPHT CNOT gate.

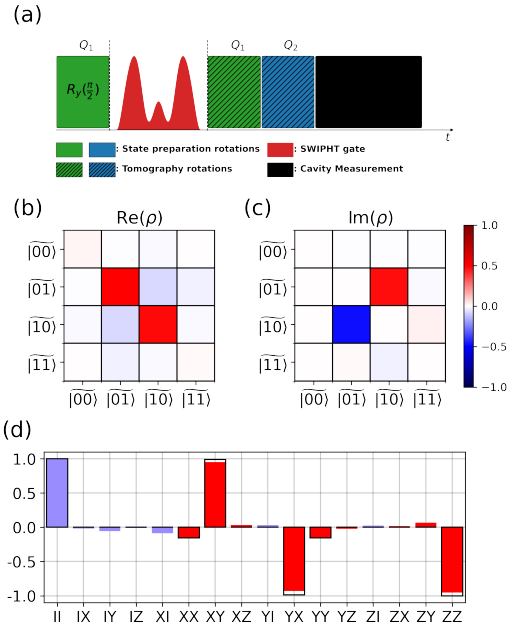

FIG. 6. Results of quantum state tomography on a maximally entangled state prepared with the SWIPHT CNOT gate. (a) Pulse sequence for the quantum state tomography experiment (not to scale). (b) and (c) are the real and imaginary parts of the reconstructed density matrix. (d) The vector representation of the state in the Pauli basis. The light blue bars are the single-qubit components, and the red bars are the two-qubit components. The outlined bars represent the ideal entangled state $\frac{1}{\sqrt{2}}(e^{-i\frac{\varphi}{2}}|10\rangle - i|01\rangle)$. By comparing to the ideal entangled state, we obtained an experimental fidelity of 98.2%.

As shown in Fig. 6(a), the experiment starts from a $\pi/2$ rotation around the y-axis of the first qubit, which is the control qubit. This prepares the system in the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|0\rangle$. We then apply the SWIPHT CNOT gate, which brings the system to a maximally entangled state $|\psi\rangle = \frac{1}{\sqrt{2}}(e^{-i\frac{\varphi}{2}}|10\rangle - i|01\rangle)$. The extra phase $\varphi/2$ on $|10\rangle$ is due to the fact that the SWIPHT CNOT is a generalized CNOT gate. Finally, tomography rotations are applied, followed by projective measurements on both qubits. Each run consists of $N = 1000$ measurements in order to obtain good statistics on the final state populations. Noise and imperfections during the experiment can lead to an unphysical density matrix. To remedy this, a maximum likelihood estimation is applied to search for a physical density matrix such that the predicted probability of obtaining the observed data with the physical density matrix is maximized [32]. In addition, a wide-band, quantum limited Josephson parametric amplifier was used that allowed us to read out both qubits at the

same time.

The real and imaginary parts of the reconstructed density matrix are shown in Fig. 6(b) and (c), respectively. As we can see, the largest contributions are all within the $|10\rangle$, $|01\rangle$ subspace, as expected. To compare to the ideal entangled state, the vector representation in the Pauli basis of the entangled state is shown in Fig. 6(d). A fidelity of 98.2 % is obtained using $\mathcal{F} = \langle\psi|\rho|\psi\rangle$, where $|\psi\rangle$ and ρ are the ideal entangled state and the reconstructed density matrix from experiment, respectively.

VII. QUANTUM PROCESS TOMOGRAPHY OF TWO-QUBIT GATES

Finally, we perform a full characterization of our two-qubit entangling gates using quantum process tomography. The process tomography starts from single-qubit rotations implemented with TAGs on both qubits to prepare the system. There are 36 different possible initial states for the two qubits: $\{|\pm x\rangle, |\pm y\rangle, |\pm z\rangle\}^{\otimes 2}$. A separate run for each initial state is conducted. In each run, a two-qubit entangling gate is applied after the initial state is prepared. For the SWIPHT CNOT gate, the gate time is about 100.6 ns. However, to remove a residual conditional phase due to the ZZ interaction, right after the SWIPHT gate we let the system idle (evolve freely) for a time such that the sum of the SWIPHT gate time and the idle time is equal to a full period of the ZZ interaction, $2\pi/g_z$. For the free-evolution CZ gate, no entangling microwave pulse is applied, and the system instead undergoes free evolution for 53.8 ns. Subsequently, tomography rotations and projective measurements are performed on both qubits. Each run consists of $N = 1000$ measurements in order to obtain good statistics on the final state populations. We apply a maximum likelihood estimation method to reconstruct the Pauli transfer matrix R for the process.

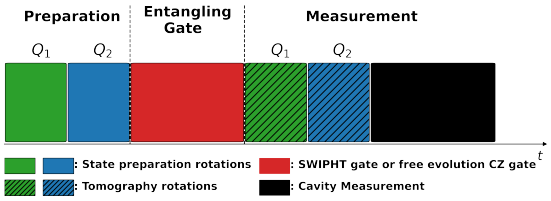

FIG. 7. Pulse sequence for the quantum process tomography experiment (not to scale).

The experimental Pauli transfer matrices, $R^{\rm exp}_{\rm CNOT}$ for the SWIPHT CNOT gate and $R^{\rm exp}_{\rm CZ}$ for the free-evolution CZ gate, are shown in Fig. 8(a) and (d), respectively. Figure 8(c) and (f) illustrate the ideal R matrices, $R^{\rm ideal}_{\rm CNOT}$ and $R^{\rm ideal}_{\rm CNOT}$, for the SWIPHT CNOT and CZ gate, respectively. The average gate fidelity in each case is computed by comparing $R^{\rm exp}$ to $R^{\rm ideal}$,

$$\mathcal{F}_{ave} = \frac{\mathrm{Tr}[(R^{\rm exp})^T R^{\rm ideal}]/4 + 1}{5}. \quad (9)$$

For the SWIPHT CNOT (CZ), we obtain an average gate fidelity of 94.6% (97.8%). To understand how much the decoherence contributes to the infidelity of the gates, we performed master equation simulations for both the SWIPHT CNOT and CZ gates with state preparation rotations also included, incorporating decoherence at levels consistent with our measured T_1 and T_2^* times. We then run the quantum process tomography on the simulated data to extract the simulated R matrices, shown in Fig. 8(b) and (e) for the SWIPHT CNOT and CZ gates, respectively. The average gate fidelities from the master equation simulations are 98.92% for the SWIPHT CNOT gate and 99.25% for the CZ gate. In another simulation, perfect state preparation is assumed. The outcome gate fidelities are 99.4% for the SWIPHT CNOT gate and 99.7% for the CZ gate. This indicates that the SPAM errors [33] contributes about 0.5% to the infidelity. The experiments report about 4% lower fidelities compared to the simulations. We attribute this difference to coherent errors [34] and leakage errors [19]. Note that our master equation simulation only includes the computational subspace and therefore does not capture leakage errors. Purity of the SWIPHT CNOT Pauli transfer matrix is measured to be 95.54%, which indicates the presence of leakage errors. Extra nonzero terms in Fig. 8(a) and (d) compared to the simulation results in Fig. 8(b) and (e) indicates that there are indeed coherent errors. On the other hand, the maximally entangled state in the quantum state tomography described in the last section reports a fidelity of 98.2%, which is close to the decoherence limit from the master equation simulation. This means that other states in the quantum process tomography would have fidelities lower than the average gate fidelity. The varying state fidelities of different states in the quantum process tomography also indicates that there are coherent errors in the experiment.

VIII. CONCLUSIONS

In this paper, we experimentally demonstrated a universal quantum gate set for the regime of strong ZZ coupling between superconducting transmon qubits. We introduced the concept of two-axis gates to implement arbitrary single-qubit rotations in this regime, and we tested their performance in a two-qubit system using randomized benchmarking. The average gate fidelity was found to be about 99% for the first qubit and 98% for the second qubit. In addition, we demonstrated two types of two-qubit entangling gates: a SWIPHT CNOT gate implemented with a shaped microwave pulse and a CZ gate based on free evolution under the ZZ interaction. The SWIPHT CNOT was found to produce a maximally entangled state with a measured fidelity of 98.2% from

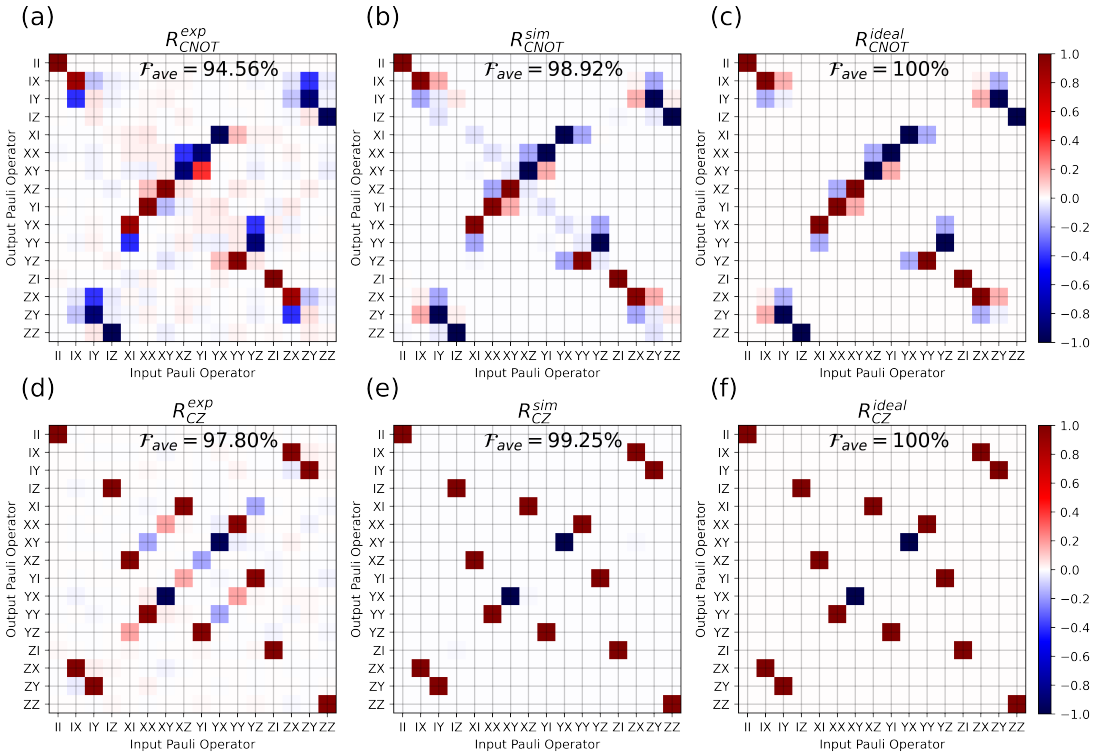

FIG. 8. (a)-(c) [(d)-(f)] are the experimental, simulated, ideal Pauli transfer matrices for the SWIPHT CNOT gate [free evolution CZ gate]. Note that (e) and (f) are different although no significant difference can be noticed under this color scale.

quantum state tomography. We further characterized the performance of both the SWIPHT CNOT and the free-evolution CZ gates using quantum process tomography. We obtained average gate fidelities of 94.6% and 97.8% for the SWIPHT CNOT and the CZ, respectively. We presented evidence that the infidelities are likely due to coherent, SPAM, and leakage errors. Advanced pulse shaping techniques [11] and calibration improvements would likely reduce these errors significantly. The universal gate set demonstrated here offers additional control flexibility that can help to optimize the use of tunable couplers and to mitigate the adverse effects of residual ZZ interactions in superconducting qubit processors.

IX. ACKNOWLEDGMENTS

This research was supported by the National Science Foundation (Award No. 1839136), DOE (Grant No. DE-SC0019199) and the DOE through FNAL as well as the new and emerging qubit science and technology (NEQST) program initiated by the US Army Research Office (ARO) in collaboration with the Laboratory of Physical Sciences (LPS) under Grant No. W911NF-18-1-0114. NIST authors acknowledge support of the NIST Quantum Based Metrology Initiative, NQI, and partial support from Google. This work is property of the US Government and not subject to copyright.

Junling Long and Tongyu Zhao contributed equally to this work.

Appendix A: Spectrum of the two-qubit system measured with Ramsey oscillations

The spectrum of the coupled two-qubit transmon system is shown in Table I.

Appendix B: Relaxation and Ramsey times

The Relaxation and Ramsey times of the coupled two-qubit system are shown in Table II.

Qubit Transition	Frequency(GHz)
$\widetilde{\lvert 00\rangle} \leftrightarrow \widetilde{\lvert 10\rangle}$	5.07478658
$\widetilde{\lvert 00\rangle} \leftrightarrow \widetilde{\lvert 01\rangle}$	5.30990762
$\widetilde{\lvert 01\rangle} \leftrightarrow \widetilde{\lvert 11\rangle}$	5.08406906
$\widetilde{\lvert 10\rangle} \leftrightarrow \widetilde{\lvert 11\rangle}$	5.31920716
$\widetilde{\lvert 10\rangle} \leftrightarrow \widetilde{\lvert 20\rangle}$	4.81503094
$\widetilde{\lvert 10\rangle} \leftrightarrow \widetilde{\lvert 02\rangle}$	4.96857665

TABLE I. Spectrum of the two-qubit system.

	T_1(µs)	T_2^*(µs)
Q_1	76.98	50.65
Q_2	79.71	17.09

TABLE II. Relaxation and Ramsey times of the qubits.

Appendix C: Derivation of the ZZ coupling

In this appendix, we derive effective ZZ-coupling Hamiltonians for two coupled transmons. We first consider the case of a direct capacitive coupling, and then we consider an indirect, cavity-mediated coupling. In both cases, an effective Hamiltonian of the form shown in Eq. 1 is obtained. Similar derivations have been given in prior works [6, 26], but we include them here for the sake of completeness.

1. Direct capacitive coupling

Two transmons that have a direct capacitive coupling are described by the Hamiltonian

$$H = \sum_j \omega_{1,j} \left|j\right\rangle \left\langle j\right| + \sum_\alpha \omega_{2,\alpha} \left|\alpha\right\rangle \left\langle \alpha\right| + \sum_{j\alpha} g_{j,\alpha} \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right| + \text{H.c.} \quad \text{(C1)}$$

The energy of the jth level of the first transmon is $\omega_{1,j}$, while the energy of the αth level of the second transmon is $\omega_{2,\alpha}$. We use Latin indices to index transmon 1 states and Greek indices to index transmon 2 states. The coupling strengths between these levels are denoted by $g_{j,\alpha}$. We assume these are real.

We will perform a Schrieffer-Wolff transformation to eliminate the coupling terms in Eq. (C1). This will be done using a similarity transformation $D = e^{S-S^\dagger}$, where

$$S = \sum_{j\alpha} \beta_{j,\alpha} \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|. \quad \text{(C2)}$$

Under this transformation, the Hamiltonian becomes

$$H' = DHD^\dagger = H + [S-S^\dagger, H] + \frac{1}{2}[S-S^\dagger, [S-S^\dagger, H]] + \ldots \quad \text{(C3)}$$

Our goal is to find the values of $\beta_{j,\alpha}$ that make the terms in $[S-S^\dagger, H]$ that are linear in $g_{j,\alpha}$ cancel the coupling terms in H.

For each Hermitian term A in Eq. (C1), it suffices to compute $[S, A]$, because $[S-S^\dagger, A] = [S, A] + [S, A]^\dagger$. Keeping only terms up to second order in $g_{j,\alpha}$ and $\beta_{j,\alpha}$, we have

$$\left[S, \sum_j \omega_{1,j} \left|j\right\rangle \left\langle j\right|\right] = \sum_{j\alpha} \beta_{j,\alpha}(\omega_{1,j} - \omega_{1,j+1}) \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|, \quad \text{(C4)}$$

$$\left[S, \sum_\alpha \omega_{2,\alpha} \left|\alpha\right\rangle \left\langle \alpha\right|\right] = \sum_{j\alpha} \beta_{j,\alpha}(\omega_{2,\alpha+1} - \omega_{2,\alpha}) \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|, \quad \text{(C5)}$$

$$\left[S, \sum_{j\alpha} g_{j,\alpha} \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right|\right] = \sum_{j\alpha} \beta_{j,\alpha} g_{j,\alpha} \left|j+1\right\rangle \left\langle j+1\right| \otimes \left|\alpha\right\rangle \left\langle \alpha\right| - \sum_{j\alpha} \beta_{j,\alpha} g_{j,\alpha} \left|j\right\rangle \left\langle j\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha+1\right|, \quad \text{(C6)}$$

$$\left[S, \sum_{j\alpha} g_{j,\alpha} \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|\right] = \sum_{j\alpha} \beta_{j+1,\alpha} g_{j,\alpha+1} \left|j+2\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+2\right| - \sum_{j\alpha} \beta_{j,\alpha+1} g_{j+1,\alpha} \left|j+2\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+2\right|, \quad \text{(C7)}$$

$$\begin{aligned}\left[S, \left[S-S^\dagger, \sum_j \omega_{1,j} \left|j\right\rangle \left\langle j\right|\right]\right] &= \sum_{j\alpha} \beta_{j,\alpha}^2(\omega_{1,j} - \omega_{1,j+1}) \left|j+1\right\rangle \left\langle j+1\right| \otimes \left|\alpha\right\rangle \left\langle \alpha\right| \\ &- \sum_{j\alpha} \beta_{j,\alpha}^2(\omega_{1,j} - \omega_{1,j+1}) \left|j\right\rangle \left\langle j\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha+1\right| \\ &+ \sum_{j\alpha} \beta_{j+1,\alpha}\beta_{j,\alpha+1}(\omega_{1,j} - \omega_{1,j+1}) \left|j+2\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+2\right| \\ &- \sum_{j\alpha} \beta_{j,\alpha+1}\beta_{j+1,\alpha}(\omega_{1,j+1} - \omega_{1,j+2}) \left|j+2\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+2\right|, \end{aligned} \quad \text{(C8)}$$

$$\begin{aligned}
&\left[S,\left[S-S^\dagger,\sum_\alpha \omega_{2,\alpha}\left|\alpha\right\rangle\left\langle\alpha\right|\right]\right]\\
&=\sum_{j\alpha}\beta_{j,\alpha}^2(\omega_{2,\alpha+1}-\omega_{2,\alpha})\left|j+1\right\rangle\left\langle j+1\right|\otimes\left|\alpha\right\rangle\left\langle\alpha\right|\\
&\quad-\sum_{j\alpha}\beta_{j,\alpha}^2(\omega_{2,\alpha+1}-\omega_{2,\alpha})\left|j\right\rangle\left\langle j\right|\otimes\left|\alpha+1\right\rangle\left\langle\alpha+1\right|\\
&+\sum_{j\alpha}\beta_{j+1,\alpha}\beta_{j,\alpha+1}(\omega_{2,\alpha+2}-\omega_{2,\alpha+1})\left|j+2\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+2\right|\\
&\quad-\sum_{j\alpha}\beta_{j,\alpha+1}\beta_{j+1,\alpha}(\omega_{2,\alpha+1}-\omega_{2,\alpha})\left|j+2\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+2\right|,
\end{aligned}\tag{C9}$$

From Eqs. (C4) and (C5), we see that we need

$$\beta_{j,\alpha}=-\frac{g_{j,\alpha}}{\omega_{1,j}-\omega_{1,j+1}+\omega_{2,\alpha+1}-\omega_{2,\alpha}}.\tag{C10}$$

Eq. (C6) then gives rise to shifts in the transmon energy levels and to an effective ZZ coupling. Restricting to the logical subspace, the full Hamiltonian then becomes

$$\begin{aligned}
H'_{logical}&=\frac{1}{2}\left(\omega_{1,0}+\omega_{1,1}+\omega_{2,0}+\omega_{2,1}+\frac{1}{2}\lambda_{ZZ}\right)II\\
&+\frac{1}{2}\left(\omega_{2,0}-\omega_{2,1}-\eta-\frac{1}{2}\lambda_{ZZ}\right)IZ\\
&+\frac{1}{2}\left(\omega_{1,0}-\omega_{1,1}+\eta-\frac{1}{2}\lambda_{ZZ}\right)ZI+\frac{1}{4}\lambda_{ZZ}ZZ,
\end{aligned}\tag{C11}$$

where

$$\eta=\frac{g_{0,0}^2}{\omega_{1,0}+\omega_{2,1}-\omega_{1,1}-\omega_{2,0}},\tag{C12}$$

$$\lambda_{ZZ}=\frac{g_{1,0}^2}{\omega_{1,1}+\omega_{2,1}-\omega_{1,2}-\omega_{2,0}}+\frac{g_{0,1}^2}{\omega_{1,1}+\omega_{2,1}-\omega_{1,0}-\omega_{2,2}}.\tag{C13}$$

We can see that the frequencies of the transitions $\widetilde{|00\rangle}\leftrightarrow\widetilde{|01\rangle}$ and $\widetilde{|10\rangle}\leftrightarrow\widetilde{|11\rangle}$ in the dressed basis are split by λ_{ZZ}. Notice that λ_{ZZ} itself depends quadratically on the bare couplings and inversely on the bare energy splittings between $|11\rangle$ and $|20\rangle$ and between $|11\rangle$ and $|02\rangle$. If we define $\omega_1=\omega_{1,1}-\omega_{1,0}-\eta$, $\omega_2=\omega_{2,1}-\omega_{2,0}+\eta$, and $g_z=\lambda_{ZZ}$, we obtain Eq. 1.

a. *Driving terms*

We now consider the inclusion of driving terms on each transmon:

$$H_d=\Omega_1(t)\sum_j d_{1,j}\left|j\right\rangle\left\langle j+1\right|+\Omega_2(t)\sum_\alpha d_{2,\alpha}\left|\alpha\right\rangle\left\langle\alpha+1\right|+\text{H.c.}\tag{C14}$$

Under the Schrieffer-Wolff transformation, this becomes

$$H'_d=DH_dD^\dagger=H_d+[S-S^\dagger,H_d]+\frac{1}{2}[S-S^\dagger,[S-S^\dagger,H_d]]+\ldots\tag{C15}$$

where

$$\begin{aligned}
\left[S,\Omega_1(t)\sum_j d_{1,j}\left|j\right\rangle\left\langle j+1\right|\right]&=\Omega_1(t)\sum_{j\alpha}\beta_{j,\alpha}d_{1,j}\left|j+1\right\rangle\left\langle j+1\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|-\Omega_1(t)\sum_{j,\alpha}\beta_{j,\alpha}d_{1,j}\left|j\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|\\
&=\Omega_1(t)\sum_{j\alpha}\left(\beta_{j,\alpha}d_{1,j}-\beta_{j+1,\alpha}d_{1,j+1}\right)\left|j+1\right\rangle\left\langle j+1\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|\\
&\quad-\Omega_1(t)d_{1,0}\left|0\right\rangle\left\langle 0\right|\otimes\sum_\alpha\beta_{0,\alpha}\left|\alpha\right\rangle\left\langle\alpha+1\right|.
\end{aligned}\tag{C16}$$

$$\begin{aligned}
\left[S,\Omega_1^*(t)\sum_j d_{1,j}^*\left|j+1\right\rangle\left\langle j\right|\right]&=\Omega_1^*(t)\sum_{j\alpha}\beta_{j+1,\alpha}d_{1,j}^*\left|j+2\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|\\
&\quad-\Omega_1^*(t)\sum_{j,\alpha}\beta_{j,\alpha}d_{1,j+1}^*\left|j+2\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|\\
&=\Omega_1^*(t)\sum_{j\alpha}\left(\beta_{j+1,\alpha}d_{1,j}^*-\beta_{j,\alpha}d_{1,j+1}^*\right)\left|j+2\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+1\right|.
\end{aligned}\tag{C17}$$

$$\left[S, \Omega_2(t)\sum_\alpha d_{2,\alpha}\left|\alpha\right\rangle\left\langle\alpha+1\right|\right] = \Omega_2(t)\sum_{j\alpha}\beta_{j,\alpha}d_{2,\alpha+1}\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+2\right| - \Omega_2(t)\sum_{j,\alpha}\beta_{j,\alpha+1}d_{2,\alpha}\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+2\right| = \Omega_2(t)\sum_{j\alpha}\left(\beta_{j,\alpha}d_{2,\alpha+1}-\beta_{j,\alpha+1}d_{2,\alpha}\right)\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha+2\right|. \tag{C18}$$

$$\left[S, \Omega_2^*(t)\sum_\alpha d_{2,\alpha}^*\left|\alpha+1\right\rangle\left\langle\alpha\right|\right] = \Omega_2^*(t)\sum_{j\alpha}\beta_{j,\alpha}d_{2,\alpha}^*\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha\right\rangle\left\langle\alpha\right| - \Omega_2^*(t)\sum_{j,\alpha}\beta_{j,\alpha}d_{2,\alpha}^*\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha+1\right\rangle\left\langle\alpha+1\right| = \Omega_2^*(t)\sum_{j\alpha}\left(\beta_{j,\alpha+1}d_{2,\alpha+1}^*-\beta_{j,\alpha}d_{2,\alpha}^*\right)\left|j+1\right\rangle\left\langle j\right|\otimes\left|\alpha+1\right\rangle\left\langle\alpha+1\right| + \Omega_2^*(t)d_{2,0}^*\sum_j\beta_{j,0}\left|j+1\right\rangle\left\langle j\right|\otimes\left|0\right\rangle\left\langle 0\right|. \tag{C19}$$

In the logical subspace, we then have

$$\begin{aligned}[S-S^\dagger, H_d]_{logical} &= \left|1\right\rangle\left\langle 1\right|\otimes\left[\Omega_1(\beta_{0,0}d_{1,0}-\beta_{1,0}d_{1,1})\left|0\right\rangle\left\langle 1\right| + \Omega_1^*(\beta_{0,0}d_{1,0}^*-\beta_{1,0}d_{1,1}^*)\left|1\right\rangle\left\langle 0\right|\right]\\ &- \beta_{0,0}\left|0\right\rangle\left\langle 0\right|\otimes\left[\Omega_1 d_{1,0}\left|0\right\rangle\left\langle 1\right| + \Omega_1^* d_{1,0}^*\left|1\right\rangle\left\langle 0\right|\right]\\ &+ \left[\Omega_2(\beta_{0,1}d_{2,1}-\beta_{0,0}d_{2,0})\left|0\right\rangle\left\langle 1\right| + \Omega_2^*(\beta_{0,1}d_{2,1}^*-\beta_{0,0}d_{2,0}^*)\left|1\right\rangle\left\langle 0\right|\right]\otimes\left|1\right\rangle\left\langle 1\right|\\ &+ \beta_{0,0}\left[\Omega_2 d_{2,0}\left|0\right\rangle\left\langle 1\right| + \Omega_2^* d_{2,0}^*\left|1\right\rangle\left\langle 0\right|\right]\otimes\left|0\right\rangle\left\langle 0\right|. \end{aligned} \tag{C20}$$

The full effective Hamiltonian in the logical subspace is

$$H'_{logical} = \begin{pmatrix} \omega_{1,0}+\omega_{2,0} & d_{2,0}\Omega_2 - d_{1,0}\beta_{0,0}\Omega_1 & d_{1,0}\Omega_1 + d_{2,0}\beta_{0,0}\Omega_2 & 0\\ d_{2,0}^*\Omega_2^* - d_{1,0}^*\beta_{0,0}\Omega_1^* & \omega_{1,0}+\omega_{2,1}+\eta & 0 & d_{1,0}\Omega_1 + (d_{2,1}\beta_{0,1}-d_{2,0}\beta_{0,0})\Omega_2\\ d_{1,0}^*\Omega_1^* + d_{2,0}^*\beta_{0,0}\Omega_2^* & 0 & \omega_{1,1}+\omega_{2,0}-\eta & (d_{1,0}\beta_{0,0}-d_{1,1}\beta_{1,0})\Omega_1 + d_{2,0}\Omega_2\\ 0 & d_{1,0}^*\Omega_1^* + (d_{2,1}^*\beta_{0,1}-d_{2,0}^*\beta_{0,0})\Omega_2^* & (d_{1,0}^*\beta_{0,0}-d_{1,1}^*\beta_{1,0})\Omega_1^* + d_{2,0}^*\Omega_2^* & \omega_{1,1}+\omega_{2,1}+\lambda_{ZZ}\end{pmatrix}.$$

2. Resonator-mediated coupling

Next, we consider two transmons coupled to the same resonator:

$$H = \omega_c a^\dagger a + \sum_j \epsilon_{1,j}\left|j\right\rangle\left\langle j\right| + \sum_\alpha \epsilon_{2,\alpha}\left|\alpha\right\rangle\left\langle\alpha\right| + \sum_j h_{1,j}a^\dagger\left|j\right\rangle\left\langle j+1\right| + \sum_\alpha h_{2,\alpha}a^\dagger\left|\alpha\right\rangle\left\langle\alpha+1\right| + \text{H.c.} \tag{C21}$$

ω_c is the cavity frequency. The energy levels of the first transmon are $\epsilon_{1,j}$, while those of the second are $\epsilon_{2,\alpha}$. $h_{1,j}$ and $h_{2,\alpha}$ are the transmon-cavity coupling strengths, which we assume are real. We again employ a Schrieffer-Wolff transformation $D = e^{S-S^\dagger}$ to transform to a new basis in which the Hamiltonian, $H' = DHD^\dagger$, does not have transmon-cavity couplings to first order in $h_{1,j}$ and $h_{2,\alpha}$. This is accomplished by choosing

$$S = \sum_j \gamma_{1,j}a^\dagger\left|j\right\rangle\left\langle j+1\right| + \sum_\alpha \gamma_{2,\alpha}a^\dagger\left|\alpha\right\rangle\left\langle\alpha+1\right|. \tag{C22}$$

To determine $\gamma_{1,j}$ and $\gamma_{2,\alpha}$, we compute commutators up to second order in the transmon-cavity couplings:

$$\left[S, \omega_c a^\dagger a\right] = -\omega_c\sum_j\gamma_{1,j}a^\dagger\left|j\right\rangle\left\langle j+1\right| - \omega_c\sum_\alpha\gamma_{2,\alpha}a^\dagger\left|\alpha\right\rangle\left\langle\alpha+1\right|, \tag{C23}$$

$$\left[S, \sum_j\epsilon_{1,j}\left|j\right\rangle\left\langle j\right|\right] = \sum_j\gamma_{1,j}(\epsilon_{1,j+1}-\epsilon_{1,j})a^\dagger\left|j\right\rangle\left\langle j+1\right|, \tag{C24}$$

$$\left[S, \sum_\alpha\epsilon_{2,\alpha}\left|\alpha\right\rangle\left\langle\alpha\right|\right] = \sum_\alpha\gamma_{2,\alpha}(\epsilon_{2,\alpha+1}-\epsilon_{2,\alpha})a^\dagger\left|\alpha\right\rangle\left\langle\alpha+1\right|, \tag{C25}$$

$$\left[S, \sum_j h_{1,j}a^\dagger\left|j\right\rangle\left\langle j+1\right|\right] = \sum_j(\gamma_{1,j}h_{1,j+1}-\gamma_{1,j+1}h_{1,j})(a^\dagger)^2\left|j\right\rangle\left\langle j+2\right|, \tag{C26}$$

$$\left[S, \sum_j h_{1,j} a \left|j+1\right\rangle \left\langle j\right|\right]$$
$$= \sum_j \gamma_{1,j} h_{1,j} \left(a^\dagger a \left|j\right\rangle \left\langle j\right| - a a^\dagger \left|j+1\right\rangle \left\langle j+1\right|\right)$$
$$- \sum_{j\alpha} \gamma_{2,\alpha} h_{1,j} \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|, \tag{C27}$$

$$\left[S, \sum_\alpha h_{2,\alpha} a \left|\alpha+1\right\rangle \left\langle \alpha\right|\right]$$
$$= \sum_\alpha \gamma_{2,\alpha} h_{2,\alpha} \left(a^\dagger a \left|\alpha\right\rangle \left\langle \alpha\right| - a a^\dagger \left|\alpha+1\right\rangle \left\langle \alpha+1\right|\right)$$
$$- \sum_{j\alpha} \gamma_{1,j} h_{2,\alpha} \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right|, \tag{C29}$$

$$\left[S, \sum_\alpha h_{2,\alpha} a^\dagger \left|\alpha\right\rangle \left\langle \alpha+1\right|\right]$$
$$= \sum_\alpha (\gamma_{2,\alpha} h_{2,\alpha+1} - \gamma_{2,\alpha+1} h_{2,\alpha}) (a^\dagger)^2 \left|\alpha\right\rangle \left\langle \alpha+2\right| \tag{C28}$$

$$\left[S, \left[S - S^\dagger, \omega_c a^\dagger a\right]\right] = \left[S, -\omega_c \sum_j \gamma_{1,j} \left(a^\dagger \left|j\right\rangle \left\langle j+1\right| + a \left|j+1\right\rangle \left\langle j\right|\right)\right] + \left[S, -\omega_c \sum_\alpha \gamma_{2,\alpha} \left(a^\dagger \left|\alpha\right\rangle \left\langle \alpha+1\right| + a \left|\alpha+1\right\rangle \left\langle \alpha\right|\right)\right]$$
$$= -\omega_c \sum_j \gamma_{1,j}^2 \left(a^\dagger a \left|j\right\rangle \left\langle j\right| - a a^\dagger \left|j+1\right\rangle \left\langle j+1\right|\right) - \omega_c \sum_\alpha \gamma_{2,\alpha}^2 \left(a^\dagger a \left|\alpha\right\rangle \left\langle \alpha\right| - a a^\dagger \left|\alpha+1\right\rangle \left\langle \alpha+1\right|\right)$$
$$+ \omega_c \sum_{j\alpha} \gamma_{1,j} \gamma_{2,\alpha} \left(\left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right| + \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right|\right), \tag{C30}$$

$$\left[S, \left[S - S^\dagger, \sum_j \epsilon_{1,j} \left|j\right\rangle \left\langle j\right|\right]\right] = \left[S, \sum_j \gamma_{1,j} (\epsilon_{1,j+1} - \epsilon_{1,j}) \left(a^\dagger \left|j\right\rangle \left\langle j+1\right| + a \left|j+1\right\rangle \left\langle j\right|\right)\right]$$
$$= \sum_j \gamma_{1,j} \gamma_{1,j+1} (\epsilon_{1,j+2} - 2\epsilon_{1,j+1} + \epsilon_{1,j}) (a^\dagger)^2 \left|j\right\rangle \left\langle j+2\right|$$
$$+ \sum_j \gamma_{1,j}^2 (\epsilon_{1,j+1} - \epsilon_{1,j}) \left(a^\dagger a \left|j\right\rangle \left\langle j\right| - a a^\dagger \left|j+1\right\rangle \left\langle j+1\right|\right)$$
$$- \sum_{j\alpha} \gamma_{1,j} \gamma_{2,\alpha} (\epsilon_{1,j+1} - \epsilon_{1,j}) \left|j+1\right\rangle \left\langle j\right| \otimes \left|\alpha\right\rangle \left\langle \alpha+1\right|, \tag{C31}$$

$$\left[S, \left[S - S^\dagger, \sum_\alpha \epsilon_{2,\alpha} \left|\alpha\right\rangle \left\langle \alpha\right|\right]\right] = \left[S, \sum_\alpha \gamma_{2,\alpha} (\epsilon_{2,\alpha+1} - \epsilon_{2,\alpha}) \left(a^\dagger \left|\alpha\right\rangle \left\langle \alpha+1\right| + a \left|\alpha+1\right\rangle \left\langle \alpha\right|\right)\right]$$
$$= \sum_\alpha \gamma_{2,\alpha} \gamma_{2,\alpha+1} (\epsilon_{2,\alpha+2} - 2\epsilon_{2,\alpha+1} + \epsilon_{2,\alpha}) (a^\dagger)^2 \left|\alpha\right\rangle \left\langle \alpha+2\right|$$
$$+ \sum_\alpha \gamma_{2,\alpha}^2 (\epsilon_{2,\alpha+1} - \epsilon_{2,\alpha}) \left(a^\dagger a \left|\alpha\right\rangle \left\langle \alpha\right| - a a^\dagger \left|\alpha+1\right\rangle \left\langle \alpha+1\right|\right)$$
$$- \sum_{j\alpha} \gamma_{1,j} \gamma_{2,\alpha} (\epsilon_{2,\alpha+1} - \epsilon_{2,\alpha}) \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right|. \tag{C32}$$

Eqs. (C23)-(C25) yield the desired values of $\gamma_{1,j}$ and $\gamma_{2,\alpha}$:

$$\gamma_{1,j} = -\frac{h_{1,j}}{\epsilon_{1,j+1} - \epsilon_{1,j} - \omega_c}, \quad \gamma_{2,\alpha} = -\frac{h_{2,\alpha}}{\epsilon_{2,\alpha+1} - \epsilon_{2,\alpha} - \omega_c}. \tag{C33}$$

If we assume zero cavity photons, then we obtain the following effective Hamiltonian:

$$H' \approx \sum_j \tilde{\epsilon}_{1,j} \left|j\right\rangle \left\langle j\right| + \sum_\alpha \tilde{\epsilon}_{2,\alpha} \left|\alpha\right\rangle \left\langle \alpha\right|$$
$$+ \sum_{j\alpha} \tilde{g}_{j,\alpha} \left|j\right\rangle \left\langle j+1\right| \otimes \left|\alpha+1\right\rangle \left\langle \alpha\right| + \text{H.c.} \tag{C34}$$

where $\tilde{\epsilon}_{1,0} = \epsilon_{1,0}$, $\tilde{\epsilon}_{2,0} = \epsilon_{2,0}$, and

$$\begin{aligned}\tilde{\epsilon}_{1,j} &= \epsilon_{1,j} - 2\gamma_{1,j-1}h_{1,j-1} + (\omega_c - \epsilon_{1,j} + \epsilon_{1,j-1})\gamma_{1,j-1}^2 \\ &= \epsilon_{1,j} + \frac{h_{1,j-1}^2}{\epsilon_{1,j} - \epsilon_{1,j-1} - \omega_c}, \quad \text{for} \quad j > 0, \end{aligned} \tag{C35}$$

$$\begin{aligned}\tilde{\epsilon}_{2,\alpha} &= \epsilon_{2,\alpha} - 2\gamma_{2,\alpha-1}h_{2,\alpha-1} + (\omega_c - \epsilon_{2,\alpha} + \epsilon_{2,\alpha-1})\gamma_{2,\alpha-1}^2 \\ &= \epsilon_{2,\alpha} + \frac{h_{2,\alpha-1}^2}{\epsilon_{2,\alpha} - \epsilon_{2,\alpha-1} - \omega_c}, \quad \text{for} \quad \alpha > 0, \end{aligned} \tag{C36}$$

$$\begin{aligned}\tilde{g}_{j,\alpha} &= -\gamma_{2,\alpha}h_{1,j} - \gamma_{1,j}h_{2,\alpha} \\ &\quad + \frac{1}{2}(2\omega_c - \epsilon_{1,j+1} + \epsilon_{1,j} - \epsilon_{2,\alpha+1} + \epsilon_{2,\alpha})\gamma_{1,j}\gamma_{2,\alpha} \\ &= \frac{h_{1,j}h_{2,\alpha}(\epsilon_{1,j+1} - \epsilon_{1,j} + \epsilon_{2,\alpha+1} - \epsilon_{2,\alpha} - 2\omega_c)}{2(\epsilon_{1,j+1} - \epsilon_{1,j} - \omega_c)(\epsilon_{2,\alpha+1} - \epsilon_{2,\alpha} - \omega_c)}. \end{aligned} \tag{C37}$$

Eq. (C34) has the same form as the Hamiltonian describing the direct capacitive coupling between two transmons, Eq. (C1). We can therefore use the results from the previous section to write down the effective ZZ-coupling Hamiltonian for two cavity-coupled transmons:

$$H_{eff} = \begin{pmatrix} \tilde{\epsilon}_{1,0} + \tilde{\epsilon}_{2,0} & 0 & 0 & 0 \\ 0 & \tilde{\epsilon}_{1,0} + \tilde{\epsilon}_{2,1} + \tilde{\eta} & 0 & 0 \\ 0 & 0 & \tilde{\epsilon}_{1,1} + \tilde{\epsilon}_{2,0} - \tilde{\eta} & 0 \\ 0 & 0 & 0 & \tilde{\epsilon}_{1,1} + \tilde{\epsilon}_{2,1} + \tilde{\lambda}_{ZZ} \end{pmatrix}, \tag{C38}$$

where

$$\tilde{\eta} = \frac{\tilde{g}_{0,0}^2}{\tilde{\epsilon}_{1,0} + \tilde{\epsilon}_{2,1} - \tilde{\epsilon}_{1,1} - \tilde{\epsilon}_{2,0}}, \tag{C39}$$

$$\tilde{\lambda}_{ZZ} = \frac{\tilde{g}_{1,0}^2}{\tilde{\epsilon}_{1,1} + \tilde{\epsilon}_{2,1} - \tilde{\epsilon}_{1,2} - \tilde{\epsilon}_{2,0}} + \frac{\tilde{g}_{0,1}^2}{\tilde{\epsilon}_{1,1} + \tilde{\epsilon}_{2,1} - \tilde{\epsilon}_{1,0} - \tilde{\epsilon}_{2,2}}. \tag{C40}$$

If we define $\omega_1 = \tilde{\epsilon}_{1,1} - \tilde{\epsilon}_{1,0} - \tilde{\eta}$, $\omega_2 = \tilde{\epsilon}_{2,1} - \tilde{\epsilon}_{2,0} + \tilde{\eta}$, and $g_z = \tilde{\lambda}_{ZZ}$, then (after a constant overall energy shift) we obtain Eq. 1.

$$H'_{logical} = \begin{pmatrix} \tilde{\epsilon}_{1,0} + \tilde{\epsilon}_{2,0} & d_{2,0}\Omega_2 - d_{1,0}\tilde{\beta}_{0,0}\Omega_1 & d_{1,0}\Omega_1 + d_{2,0}\tilde{\beta}_{0,0}\Omega_2 & 0 \\ d_{2,0}^*\Omega_2^* - d_{1,0}^*\tilde{\beta}_{0,0}\Omega_1^* & \tilde{\epsilon}_{1,0} + \tilde{\epsilon}_{2,1} + \tilde{\eta} & 0 & d_{1,0}\Omega_1 + (d_{2,1}\tilde{\beta}_{0,1} - d_{2,0}\tilde{\beta}_{0,0})\Omega_2 \\ d_{1,0}^*\Omega_1^* + d_{2,0}^*\tilde{\beta}_{0,0}\Omega_2^* & 0 & \tilde{\epsilon}_{1,1} + \tilde{\epsilon}_{2,0} - \tilde{\eta} & (d_{1,0}\tilde{\beta}_{0,0} - d_{1,1}\tilde{\beta}_{1,0})\Omega_1 + d_{2,0}\Omega_2 \\ 0 & d_{1,0}^*\Omega_1^* + (d_{2,1}^*\tilde{\beta}_{0,1} - d_{2,0}^*\tilde{\beta}_{0,0})\Omega_2^* & (d_{1,0}^*\tilde{\beta}_{0,0} - d_{1,1}^*\tilde{\beta}_{1,0})\Omega_1^* + d_{2,0}^*\Omega_2^* & \tilde{\epsilon}_{tr1,1} + \tilde{\epsilon}_{2,1} + \tilde{\lambda}_{ZZ} \end{pmatrix}, \tag{C43}$$

where

$$\tilde{\beta}_{j,\alpha} = -\frac{\tilde{g}_{j,\alpha}}{\tilde{\epsilon}_{1,j} - \tilde{\epsilon}_{1,j+1} + \tilde{\epsilon}_{2,\alpha+1} - \tilde{\epsilon}_{2,\alpha}}. \tag{C43}$$

a. Driving terms

We now consider the inclusion of driving terms on each transmon:

$$\begin{aligned} H_d = \Omega_1(t)\sum_j d_{1,j}\left|j\right\rangle\left\langle j+1\right| \\ + \Omega_2(t)\sum_\alpha d_{2,\alpha}\left|\alpha\right\rangle\left\langle \alpha+1\right| + \text{H.c.} \end{aligned} \tag{C41}$$

Under the Schrieffer-Wolff transformation, this becomes

$$\begin{aligned} H'_d = DH_dD^\dagger = H_d + [S - S^\dagger, H_d] \\ + \frac{1}{2}[S - S^\dagger, [S - S^\dagger, H_d]] + \ldots \end{aligned} \tag{C42}$$

The first-order corrections vanish in the zero-cavity-photon subspace because S is proportional to $a^\dagger$, while H_d does not contain any cavity-photon operators. Therefore, the driving terms after the SW transformation are still given by H_d. The full effective Hamiltonian in the logical subspace is

[1] Jay M Gambetta, Jerry M Chow, and Matthias Steffen, "Building logical qubits in a superconducting quantum computing system," npj Quantum Information **3**,

1–7 (2017).
[2] Philip Krantz, Morten Kjaergaard, Fei Yan, Terry P Orlando, Simon Gustavsson, and William D Oliver, "A quantum engineer's guide to superconducting qubits," Applied Physics Reviews **6**, 021318 (2019).
[3] Michel H Devoret and Robert J Schoelkopf, "Superconducting circuits for quantum information: an outlook," Science **339**, 1169–1174 (2013).
[4] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver, "Superconducting qubits: Current state of play," Annual Review of Condensed Matter Physics **11** (2019).
[5] Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C White, Josh Mutus, Austin G Fowler, Brooks Campbell, *et al.*, "Superconducting quantum circuits at the surface code threshold for fault tolerance," Nature **508**, 500–503 (2014).
[6] David C McKay, Sarah Sheldon, John A Smolin, Jerry M Chow, and Jay M Gambetta, "Three-qubit randomized benchmarking," Physical review letters **122**, 200502 (2019).
[7] Sabrina S Hong, Alexander T Papageorge, Prasahnt Sivarajah, Genya Crossman, Nicolas Didier, Anthony M Polloreno, Eyob A Sete, Stefan W Turkowski, Marcus P da Silva, and Blake R Johnson, "Demonstration of a parametrically activated entangling gate protected from flux noise," Physical Review A **101**, 012302 (2020).
[8] Andrew W Cross, Lev S Bishop, Sarah Sheldon, Paul D Nation, and Jay M Gambetta, "Validating quantum computers using randomized model circuits," Physical Review A **100**, 032328 (2019).
[9] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, *et al.*, "Quantum supremacy using a programmable superconducting processor," Nature **574**, 505–510 (2019).
[10] JS Otterbach, R Manenti, N Alidoust, A Bestwick, M Block, B Bloom, S Caldwell, N Didier, E Schuyler Fried, S Hong, *et al.*, "Unsupervised machine learning on a hybrid quantum computer," arXiv preprint arXiv:1712.05771 (2017).
[11] Felix Motzoi, Jay M Gambetta, Patrick Rebentrost, and Frank K Wilhelm, "Simple pulses for elimination of leakage in weakly nonlinear qubits," Physical review letters **103**, 110501 (2009).
[12] Sarah Sheldon, Lev S Bishop, Easwar Magesan, Stefan Filipp, Jerry M Chow, and Jay M Gambetta, "Characterizing errors on qubit operations via iterative randomized benchmarking," Physical Review A **93**, 012301 (2016).
[13] Julian Kelly, R Barends, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, Austin G Fowler, I-C Hoi, E Jeffrey, *et al.*, "Optimal quantum control using randomized benchmarking," Physical review letters **112**, 240504 (2014).
[14] Brooks Foxen, Charles Neill, Andrew Dunsworth, Pedram Roushan, Ben Chiaro, Anthony Megrant, Julian Kelly, Zijun Chen, Kevin Satzinger, Rami Barends, *et al.*, "Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms," arXiv preprint arXiv:2001.08343 (2020).
[15] M Kjaergaard, ME Schwartz, A Greene, GO Samach, A Bengtsson, M O'Keeffe, CM McNally, J Braumüller, DK Kim, P Krantz, *et al.*, "A quantum instruction set implemented on a superconducting quantum processor," arXiv preprint arXiv:2001.08838 (2020).
[16] S Krinner, P Kurpiers, B Royer, P Magnard, I Tsitsilin, J-C Besse, A Remm, A Blais, and A Wallraff, "Demonstration of an all-microwave controlled-phase gate between far detuned qubits," arXiv preprint arXiv:2006.10639 (2020).
[17] Yuan Xu, Ji Chu, Jiahao Yuan, Jiawei Qiu, Yuxuan Zhou, Libo Zhang, Xinsheng Tan, Yang Yu, Song Liu, Jian Li, *et al.*, "High-fidelity, high-scalability two-qubit gate scheme for superconducting qubits," arXiv preprint arXiv:2006.11860 (2020).
[18] Michele C Collodo, Johannes Herrmann, Nathan Lacroix, Christian Kraglund Andersen, Ants Remm, Stefania Lazar, Jean-Claude Besse, Theo Walter, Andreas Wallraff, and Christopher Eichler, "Implementation of conditional-phase gates based on tunable zz-interactions," arXiv preprint arXiv:2005.08863 (2020).
[19] David C McKay, Stefan Filipp, Antonio Mezzacapo, Easwar Magesan, Jerry M Chow, and Jay M Gambetta, "Universal gate for fixed-frequency qubits via a tunable bus," Physical Review Applied **6**, 064007 (2016).
[20] Pranav Mundada, Gengyan Zhang, Thomas Hazard, and Andrew Houck, "Suppression of qubit crosstalk in a tunable coupling superconducting circuit," Physical Review Applied **12**, 054023 (2019).
[21] X. Li, T. Cai, H. Yan, Z. Wang, X. Pan, Y. Ma, W. Cai, J. Han, Z. Hua, X. Han, Y. Wu, H. Zhang, H. Wang, Yipu Song, Luming Duan, and Luyan Sun, "Tunable coupler for realizing a controlled-phase gate with dynamically decoupled regime in a superconducting circuit," Phys. Rev. Applied **14**, 024070 (2020).
[22] Jaseung Ku, Xuexin Xu, Markus Brink, David C McKay, Jared B Hertzberg, Mohammad H Ansari, and BLT Plourde, "Suppression of unwanted zz interactions in a hybrid two-qubit system," arXiv preprint arXiv:2003.02775 (2020).
[23] Sophia E Economou and Edwin Barnes, "Analytical approach to swift nonleaky entangling gates in superconducting qubits," Physical Review B **91**, 161405 (2015).
[24] Xiu-Hao Deng, Edwin Barnes, and Sophia E Economou, "Robustness of error-suppressing entangling gates in cavity-coupled transmon qubits," Physical Review B **96**, 035441 (2017).
[25] George S. Barron, F. A. Calderon-Vargas, Junling Long, David P. Pappas, and Sophia E. Economou, "Microwave-based arbitrary cphase gates for transmon qubits," Phys. Rev. B **101**, 054508 (2020).
[26] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Charge-insensitive qubit design derived from the cooper pair box," Phys. Rev. A **76**, 042319 (2007).
[27] X Wu, JL Long, HS Ku, RE Lake, M Bal, and DP Pappas, "Overlap junctions for high coherence superconducting qubits," Applied Physics Letters **111**, 032602 (2017).
[28] Junling Long, *Superconducting Quantum Circuits for QuantumInformation Processing* (University of Colorado Boulder, 2020).
[29] Edwin Barnes and S Das Sarma, "Analytically solvable driven time-dependent two-level quantum systems," Physical review letters **109**, 060401 (2012).
[30] Shavindra P. Premaratne, Jen-Hao Yeh, F. C. Well-

stood, and B. S. Palmer, "Implementation of a generalized controlled-not gate between fixed-frequency transmons," Phys. Rev. A **99**, 012317 (2019).
[31] Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R Brad Blakestad, John D Jost, Chris Langer, Roee Ozeri, Signe Seidelin, and David J Wineland, "Randomized benchmarking of quantum gates," Physical Review A **77**, 012307 (2008).
[32] Jeremy L O'Brien, GJ Pryde, Alexei Gilchrist, DFV James, Nathan K Langford, TC Ralph, and AG White, "Quantum process tomography of a controlled-not gate," Physical review letters **93**, 080502 (2004).
[33] A Dewes, FR Ong, V Schmitt, R Lauro, N Boulant, P Bertet, D Vion, and D Esteve, "Characterization of a two-transmon processor with individual single-shot qubit readout," Physical review letters **108**, 057002 (2012).
[34] Phillip Kaye, Raymond Laflamme, Michele Mosca, *et al.*, *An introduction to quantum computing* (Oxford university press, 2007).

Chapter 5

Fast high-fidelity entangling gates for spin qubits in Si double quantum dots

This chapter includes pages from our publication [5]

F. A. Calderon-Vargas, G. S. Barron, X. H. Deng, A. J. Sigillito, E. Barnes, and S. E. Economou, "Fast high-fidelity entangling gates for spin qubits in Si double quantum dots", Physical Review B **100**, ARTN03530410.1103/PhysRevB.100.035304 (2019)

Copyright (2019) of the American Physical Society.

The contributions to this work are as follows: My contributions were assisting in the implementation of the SWIPHT protocol, assisting with numerical simulations. F.A.C.V. primarily carried out this work. X.D. assisted in theoretical work and simulations. A.J.S. contributed experimental expertise. E.B., and S.E.E. supervised.

Fast high-fidelity entangling gates for spin qubits in Si double quantum dots

F. A. Calderon-Vargas,[1,*] George S. Barron,[1] Xiu-Hao Deng,[1] A. J. Sigillito,[2] Edwin Barnes,[1,†] and Sophia E. Economou[1,‡]
[1]*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2]*Department of Physics, Princeton University, Princeton, New Jersey 08544, USA*

(Received 13 February 2019; revised manuscript received 21 June 2019; published 9 July 2019)

Implementing high-fidelity two-qubit gates in single-electron spin qubits in silicon double quantum dots is still a major challenge. In this work we employ analytical methods to design control pulses that generate high-fidelity entangling gates for quantum computers based on this platform. Using realistic parameters and initially assuming a noise-free environment, we present simple control pulses that generate CNOT, CPHASE, and CZ gates with average fidelities greater than 99.99% and gate times as short as 45 ns. Moreover, using the local invariants of the system's evolution operator, we show that a simple square pulse generates a CNOT gate in less than 27 ns and with a fidelity greater than 99.99%. Last, we use the same analytical methods to generate two-qubit gates locally equivalent to $\sqrt{\text{CNOT}}$ and $\sqrt{\text{CZ}}$ that are used to implement simple two-piece pulse sequences that produce high-fidelity CNOT and CZ gates in the presence of low-frequency noise.

DOI: 10.1103/PhysRevB.100.035304

I. INTRODUCTION

Among the many candidates to become the hardware of the first quantum computer, single electron spin qubits in silicon quantum dots (QDs) present a distinctive potential for scalability due to the well-established fabrication technologies and all-electrical control [1,2]. In the last five years, several groups have demonstrated high-fidelity single-qubit gates [3–7] with fidelities higher than 99.9% [6]. Two-qubit gates—a second crucial component of any quantum logic circuit—have also been demonstrated in these systems [7–12], however the fidelities in this case have not yet exceeded 98% [11]. This is mostly due to charge noise that causes fluctuations in the tunneling and detuning energies [13,14], which directly affects the exchange coupling and single-qubit resonance frequencies [6,15]. A second source of decoherence comes from fluctuations in the nuclear spin bath of the surrounding substrate and its hyperfine interaction with the electron spins, but for natural silicon, in contrast to other semiconductor platforms, the abundance of the nuclear-spin isotope ^{29}Si is only 4.7%. Moreover, the spinful isotopes can be removed through isotopic purification [16], largely eliminating this type of noise.

Therefore, charge noise remains the main obstacle in the path towards high-fidelity gates. Several methods have been proposed to mitigate its effect on qubit operations. In particular, it has been experimentally shown that the exchange energy's sensitivity to charge noise can be reduced to second order by symmetric operations [7,17–19]. However, the remaining noise still limits gate fidelities, necessitating the use of additional techniques. One approach is dynamically corrected gates (DCGs), where pulse shaping or carefully designed pulse sequences are used to suppress the effects of noise during operations. Most DCG schemes work only if the noise is concentrated at low frequencies. This is true of charge noise in Si quantum dots, where the power spectrum has been measured to have the form $1/f^{\alpha}$, with $\alpha \approx 1$ [6,15]. In fact, DCGs have been proposed for both single- and two-qubit gates [20–22]. However, it is important to make the pulses as fast as possible in these schemes, because the assumption of slow noise becomes invalid for longer gate times, weakening the power of DCGs. Of course, speeding up gates is generally desirable as it leads to faster quantum information processing and improves fidelities in the presence of finite coherence times.

The speed of two-qubit gates in exchange-coupled quantum dots is ultimately set by the strength of the exchange interaction. However, most proposed two-qubit gates take longer than what one would expect just based on the interaction strength. Recently, it has been experimentally demonstrated that entangling gates can be generated in a single shot using electric dipole spin resonance (EDSR) in conjunction with an exchange pulse [7,21]. However, this approach requires the pulse to be slow compared to the exchange time scale in order to avoid unwanted transitions in the spectrum and to accumulate the correct phases for the desired entangling gate, reducing the effectiveness of dynamical noise suppression and leading to fidelity loss.

In this work we show that fast, high-fidelity two-qubit entangling gates in Si quantum dots can be achieved with analytical pulse-shaping techniques. Our gate designs are based on single-shot exchange pulses combined with EDSR control. We overcome naive gate speed limits by purposely driving spectrally close unwanted transitions within the logical space and incorporating them into our gate design rather than simply avoiding them [23–25]. In this way we achieve entangling operations with fidelities exceeding 99.99% in the absence of noise in times as low as 45 ns, a factor of 2 improvement over previous methods. Moreover, we calculate the local

*f.calderon@vt.edu
†efbarnes@vt.edu
‡economou@vt.edu

 ©2019 American Physical Society

invariants [26] of the system's evolution operator and find that a simple square pulse generates, up to local operations, a CNOT gate in less than 27 ns and with a fidelity larger than 99.99% in the absence of noise. Finally, we combine our pulse-shaping methods with dynamical gate correction to further improve fidelities. We design fast, robust pulses for the three most commonly used entangling operations: CNOT, CZ, and CPHASE gates, which arise in numerous algorithms pertaining to quantum computing and simulation.

The paper is organized as follows. In Sec. II we introduce the two-qubit Hamiltonian. Then, in Sec. III, we present a simple analytical method, based on a general analytical solution for the unitary evolution of an arbitrary two-level system, to design control pulses that generate high-fidelity CNOT gates. In Sec. IV we compute the local invariants of the evolution operator and show that a short square pulse generates a high-fidelity CNOT gate. In Sec. V we show that with a simple hyperbolic secant (sech) function as the control pulse we can generate high-fidelity θ-CPHASE gates, with arbitrary phase angle, including the CZ (or π-CPHASE) gate. We conclude in Sec. VI.

II. TWO-QUBIT HAMILTONIAN

The system being studied comprises two electrons confined in a silicon double quantum dot (DQD). The two-electron system is under an external magnetic field, applied on the z axis, and an inhomogeneous magnetic field, created by a nearby micromagnet, with large gradients along the z and x axes (the two-dimensional electron gas is in the z-x plane and the two quantum dots are spatially separated along the x direction). In particular, the component of the magnetic field out of the plane B_y varies along the z direction, while B_z varies along the x direction. The spatial dependence of B_y enables EDSR control in which the application of a microwave pulse to a nearby metal gate rapidly oscillates the position of the electrons in the z direction, causing a transverse oscillating magnetic field B_y in the electron's frame of reference. The gradient of B_z along the x direction gives the two electrons distinct resonance frequencies, allowing individual addressability. Following the experimental setup reported in Ref. [7], we assume that the system is operated at a symmetric operation point where the exchange coupling $J(t)$ is first-order insensitive to electric noise [17,18]. The exchange coupling is controlled by a middle metal gate that modulates the energy barrier between dots and, therefore, the interdot tunneling amplitude.

We use the Heisenberg Hamiltonian to model the system,

$$H(t) = J(t)(\mathbf{S}_L \cdot \mathbf{S}_R - 1/4) + \mathbf{S}_L \cdot \mathbf{B}_L + \mathbf{S}_R \cdot \mathbf{B}_R, \quad (1)$$

where $\mathbf{S}_L$ ($\mathbf{S}_R$) is the spin of the electron in the left (right) quantum dot, $\mathbf{B}_L = (0, B_{y,L}(t), B_z^{\text{ext}} + B_{z,L}(t))^T$ and $\mathbf{B}_R = (0, B_{y,R}(t), B_z^{\text{ext}} + B_{z,R}(t))^T$ are the magnetic fields (in energy units) on the spins in the left and right dots, respectively. Here B_z^{ext} is the homogeneous external magnetic field in the z direction, $B_{z,L}(t)$ [$B_{z,R}(t)$] is the local magnetic field in the left (right) dot induced by the micromagnet's inhomogeneous field, and $B_{y,q}(t) = B_{y,q}^0 + B_{y,q}^1 \cos(\omega t + \phi)$ ($q = L, R$) is the transverse oscillating field composed of a static part $B_{y,q}^0$ and a variable amplitude $B_{y,q}^1$ controlled by a nearby metal gate. In similar fashion, the z component of the local magnetic field has a time-dependent component, i.e., $B_{z,q}(t) = B_{z,q}^0 + B_{z,q}^1(t)$, caused by the change in the position of the electrons in the magnetic field gradient when the energy barrier between dots is changed by the middle metal gate. Therefore, when the exchange coupling is turned on, the two parameters $B_{z,L}^1$ and $B_{z,R}^1$ cause shifts in the single-qubit resonance frequencies. Following Ref. [7], our analysis is focused on the regime where the Zeeman splitting of the two qubits is much larger than the strength of the exchange interaction $J \ll (B_{z,R} - B_{z,L})$. In particular, in the regime where the exchange interaction is negligible, $J \approx 0$ (see Fig. 1), single-qubit gates can be performed by matching the drive frequency ω with the resonance frequency of either dot, and by adjusting the drive phase ϕ one can control the single-qubit rotation axis in the Bloch sphere's equatorial plane.

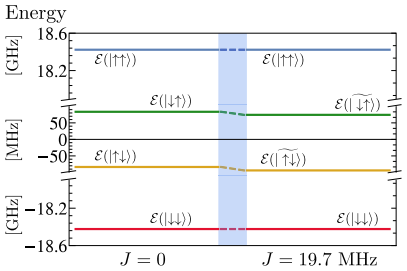

FIG. 1. Eigenenergies of the system's Hamiltonian, in the absence of AC driving, for zero exchange coupling ($J = 0$, left side), and for nonzero exchange coupling ($J = 19.7$ MHz, right side). We use the parameters reported in Ref. [7] to calculate the magnitude of the four eigenenergies. Note that when the exchange coupling is turned on the energies of the states with antiparallel spin are lowered by $\sim J/2$ for $J \ll (B_{z,R} - B_{z,L})$, and thus the conditional spin flips (i.e., the ESR frequency of one spin depends on the other spin state) are now distinguishable from each other.

In the absence of AC driving, the Hamiltonian (1) in the diabatic two-qubit basis $\{|\uparrow\uparrow\rangle, |\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\downarrow\rangle\}$ is

$$H_0(t) = \begin{pmatrix} E_z + E_z^1[J] & 0 & 0 & 0 \\ 0 & \frac{1}{2}\big(\Delta E_z + \Delta E_z^1[J] - J(t)\big) & \frac{J(t)}{2} & 0 \\ 0 & \frac{J(t)}{2} & \frac{1}{2}\big(-\Delta E_z - \Delta E_z^1[J] - J(t)\big) & 0 \\ 0 & 0 & 0 & -E_z - E_z^1[J] \end{pmatrix}, \quad (2)$$

where $E_z = B_z^{\rm ext} + (B_{z,R}^0 + B_{z,L}^0)/2$ is the average Zeeman splitting, $\Delta E_z = B_{z,R}^0 - B_{z,L}^0$ is the Zeeman splitting between dots, and $E_z^1[J] = (B_{z,R}^1 + B_{z,L}^1)/2$ and $\Delta E_z^1[J] = B_{z,R}^1 - B_{z,L}^1$ represent the Zeeman shifts that occur when the exchange J is turned on. Here $E_z^1[J]$ and $\Delta E_z^1[J]$ would vanish if J vanishes too. If the exchange is adiabatically turned on (relative to the Zeeman splitting of the two qubits), in the absence of AC driving, we can use the Hamiltonian's instantaneous adiabatic eigenstates as the new basis. Accordingly, the Hamiltonian's instantaneous eigenvalues are

$$\mathcal{E}(|\uparrow\uparrow\rangle) = E_z + E_z^1, \tag{3}$$

$$\mathcal{E}(\widetilde{|\downarrow\uparrow\rangle}) = \tfrac{1}{2}\Big[-J + \sqrt{\left(\Delta E_z + \Delta E_z^1\right)^2 + J^2}\Big], \tag{4}$$

$$\mathcal{E}(\widetilde{|\uparrow\downarrow\rangle}) = \tfrac{1}{2}\Big[-J - \sqrt{\left(\Delta E_z + \Delta E_z^1\right)^2 + J^2}\Big], \tag{5}$$

$$\mathcal{E}(|\downarrow\downarrow\rangle) = -E_z - E_z^1, \tag{6}$$

$$H_{\rm int} = \begin{pmatrix} E_z^1 & -i\frac{B_{y,L}\Delta_+^{(1)}}{2}e^{-\frac{i}{2}(\Delta E_z - 2E_z)t} & -i\frac{B_{y,R}\Delta_-^{(2)}}{2}e^{\frac{i}{2}(\Delta E_z+2E_z)t} & 0 \\ i\frac{B_{y,L}\Delta_+^{(1)}}{2}e^{\frac{i}{2}(\Delta E_z-2E_z)t} & \frac{1}{2}\left(\Delta E_z^1 - J + \frac{J^2}{2\left(\Delta E_z+\Delta E_z^1\right)}\right) & 0 & -i\frac{B_{y,R}\Delta_+^{(2)}}{2}e^{\frac{i}{2}(\Delta E_z+2E_z)t} \\ i\frac{B_{y,R}\Delta_-^{(2)}}{2}e^{-\frac{i}{2}(\Delta E_z+2E_z)t} & 0 & -\frac{1}{2}\left(\Delta E_z^1 + J + \frac{J^2}{2\left(\Delta E_z+\Delta E_z^1\right)}\right) & -i\frac{B_{y,L}\Delta_-^{(1)}}{2}e^{-\frac{i}{2}(\Delta E_z-2E_z)t} \\ 0 & i\frac{B_{y,R}\Delta_+^{(2)}}{2}e^{-\frac{i}{2}(\Delta E_z+2E_z)t} & i\frac{B_{y,L}\Delta_-^{(1)}}{2}e^{\frac{i}{2}(\Delta E_z-2E_z)t} & -E_z^1 \end{pmatrix}, \tag{7}$$

where $\Delta_\pm^{(1)}(t) = 1 \pm JB_{y,R}(t)/2B_{y,L}(t)(\Delta E_z + \Delta E_z^1)$ and $\Delta_\pm^{(2)}(t) = 1 \pm JB_{y,L}(t)/2B_{y,R}(t)(\Delta E_z + \Delta E_z^1)$. Our goal is to find pulse envelopes $B_{y,L}^1(t)$, $B_{y,R}^1(t)$ that render the corresponding evolution operator $U_{\rm int}$ locally equivalent to a two-qubit entangling gate at the end of the pulse. Note that although we have expressed $H_{\rm int}$ in the adiabatic basis, the gates we design apply equally well to the original diabatic basis since the two bases coincide at the end of the gate when the exchange coupling is switched off.

A Hamiltonian similar to the one in Eq. (7) has been addressed in a previous work on superconducting qubits [25], where the SWIPHT (speeding up wave forms by inducing phases to harmful transitions) protocol was introduced and control pulses were designed for the CNOT and CZ gates. The analytical methods developed in the aforementioned work and earlier works [23,24,27,28] are the basis for the control pulses we present in the following sections.

For the numerical results hereafter, we use the parameters reported in Ref. [7]: $B_z^{\rm ext}/2\pi = 14$ GHz, $\Delta E_z/2\pi = 214$ MHz, $B_{z,L}^0/2\pi = 4.287$ GHz, $B_{y,L}^0/2\pi = 5$ MHz, $B_{y,R}^0/2\pi = 55$ MHz, $J/2\pi = 19.7$ MHz, $B_{y,R}^1(t) = B_{y,L}^1(t)$, $\phi = 3\pi/2$, $E_z^1/2\pi = 29.23$ MHz, $\Delta E_z^1/2\pi = -46.94$ MHz.

$$H_{\rm rot} = \begin{pmatrix} E_z + E_z^1 - \omega & \frac{B_{y,L}^1\delta_+^{(1)}}{4} & \frac{B_{y,R}^1\delta_-^{(2)}}{4} & 0 \\ \frac{B_{y,L}^1\delta_+^{(1)}}{4} & \frac{1}{2}\left(\Delta E_z + \Delta E_z^1 - J + \frac{J^2}{2\left(\Delta E_z+\Delta E_z^1\right)}\right) & 0 & \frac{B_{y,R}^1\delta_+^{(2)}}{4} \\ \frac{B_{y,R}^1\delta_-^{(2)}}{4} & 0 & -\frac{1}{2}\left(\Delta E_z + \Delta E_z^1 + J + \frac{J^2}{2\left(\Delta E_z+\Delta E_z^1\right)}\right) & \frac{B_{y,L}^1\delta_-^{(1)}}{4} \\ 0 & \frac{B_{y,R}^1\delta_+^{(2)}}{4} & \frac{B_{y,L}^1\delta_-^{(1)}}{4} & \omega - E_z - E_z^1 \end{pmatrix}, \tag{8}$$

where $\widetilde{|\downarrow\uparrow\rangle}, \widetilde{|\uparrow\downarrow\rangle}$ are hybridizations of the $|\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle$ spin states.

The fast high-fidelity two-qubit gates proposed in this work require microwave AC driving along with a nonzero constant exchange interaction. Since we are assuming that $J \ll \Delta E_z + \Delta E_z^1$, the Hamiltonian (1) in the instantaneous adiabatic basis $\{|\uparrow\uparrow\rangle, \widetilde{|\downarrow\uparrow\rangle}, \widetilde{|\uparrow\downarrow\rangle}, |\downarrow\downarrow\rangle\}$ can be simplified by expanding $\sqrt{(\Delta E_z + \Delta E_z^1)^2 + J^2} \approx \Delta E_z + \Delta E_z^1 + J^2/2(\Delta E_z + \Delta E_z^1)$. Moreover, given that the Zeeman splitting is much larger than the exchange coupling, we only keep terms up to leading order in $J/(\Delta E_z + \Delta E_z^1)$. At this point it is convenient to switch to the interaction picture given by $H_0(J = 0)$, Eq. (2), because we do not include the free, uncoupled evolution in the definition of our gates—they are generated purely by AC driving. Our results can easily be adapted if one wishes to define logical gates with respect to a different frame. Consequently, in the adiabatic basis, the interaction picture Hamiltonian is

III. CNOT GATE FROM PARTIAL REVERSE ENGINEERING

In order to facilitate the analytical derivation of our shaped pulses, it is convenient to momentarily move the Hamiltonian (1) (in the adiabatic basis) to a rotating frame and apply the rotating wave approximation. Since the interaction and rotating frames are related by local unitaries, if we design a pulse that implements an entangling gate in the rotating frame, it will implement the same gate in the interaction frame up to these local unitaries. We stress that although we make the rotating wave approximation when designing pulses, we do not make this approximation when we numerically calculate the fidelity of the corresponding gates. Accordingly, moving the Hamiltonian (1) to the rotating frame, $H_{\rm rot}(t) = U_\omega H(t) U_\omega^\dagger - iU_\omega \dot{U}_\omega^\dagger$ with $U_\omega = \exp[i\omega t(S_{z,R} + S_{z,L})/\hbar]$, and applying the rotating wave approximation, we get in the adiabatic basis for $\phi = 3\pi/2$:

where $\delta_{\pm}^{(1)} = 1 \pm JB_{y,R}^1/2B_{y,L}^1(\Delta E_z + \Delta E_z^1)$ and $\delta_{\pm}^{(2)} = 1 \pm JB_{y,L}^1/2B_{y,R}^1(\Delta E_z + \Delta E_z^1)$. If the driving frequency ω is near to the left qubit's resonance frequency $|\mathcal{E}(|\uparrow\uparrow\rangle) - \mathcal{E}(|\widetilde{\downarrow\uparrow}\rangle)|$, the terms $B_{y,R}^1\delta_{\pm}^{(2)}/4$ in $H_{\rm rot}$ (8) can be dropped because $B_{y,R}^1\delta_{\pm}^{(2)}/4 \ll \Delta E_z + \Delta E_z^1$. As a consequence, we have two blocks along the diagonal, S_1 and S_2, spanned by $\{|\uparrow\uparrow\rangle, |\widetilde{\downarrow\uparrow}\rangle\}$ and $\{|\widetilde{\uparrow\downarrow}\rangle, |\downarrow\downarrow\rangle\}$, respectively, which are effectively decoupled in the rotating Hamiltonian (8) and are separated in energy by J. The resonance frequency of each block is $\omega_1 = E_z + E_z^1 - [\Delta E_z + \Delta E_z^1 - J + J^2/2(\Delta E_z + \Delta E_z^1)]/2$ and $\omega_2 = E_z + E_z^1 - [\Delta E_z + \Delta E_z^1 + J + J^2/2(\Delta E_z + \Delta E_z^1)]/2$, respectively. We start by matching the driving frequency ω with the resonance frequency of S_1, i.e., inducing a resonant spin flip between the states $|\uparrow\uparrow\rangle$ and $|\widetilde{\downarrow\uparrow}\rangle$. Moreover, since $J \ll \Delta E_z + \Delta E_z^1$ and $B_{y,R}^1(t) = B_{y,L}^1(t)$ we set $\delta_{\pm}^{(1)} \approx 1$. Consequently, the Hamiltonian (8) becomes, up to a local phase in S_2,

$$H_{\rm rot} = \begin{pmatrix} 0 & \frac{B_{y,L}^1(t)}{4} & 0 & 0 \\ \frac{B_{y,L}^1(t)}{4} & 0 & 0 & 0 \\ 0 & 0 & -\frac{J}{2} & \frac{B_{y,L}^1(t)}{4} \\ 0 & 0 & \frac{B_{y,L}^1(t)}{4} & \frac{J}{2} \end{pmatrix}. \tag{9}$$

The block S_1 can be readily integrated to give the evolution operator $U_{S_1}(t) = \exp\{-\frac{i}{4}\int_0^t dt' B_{y,L}(t')\sigma_x\}$. For the block S_2, in Refs. [25,27] it has been shown that for a two-level system with detuning Δ and driving field $\Omega(t)$ the system's evolution operator $\mathcal{U}(t)$ and pulse $\Omega(t)$ can be expressed in terms of a single real function $\chi(t)$:

$$\mathcal{U}(t) = e^{-i\frac{\pi}{4}\sigma_y}\begin{pmatrix} \cos\chi e^{i\psi_-} & \sin\chi e^{-i\psi_+} \\ -\sin\chi e^{i\psi_+} & \cos\chi e^{-i\psi_-} \end{pmatrix}, \tag{10}$$

$$\Omega(t) = \frac{\ddot{\chi}}{2\sqrt{\frac{\Delta^2}{4} - \dot{\chi}^2}} - \sqrt{\frac{\Delta^2}{4} - \dot{\chi}^2}\cot(2\chi), \tag{11}$$

where $\psi_{\pm}(t) = \int_0^t dt'\sqrt{\Delta^2/4 - \dot{\chi}^2(t')}\csc[2\chi(t')] \pm \arcsin[2\dot{\chi}(t)/\Delta]/2$, and any choice of $\chi(t)$ that fulfills $|\dot{\chi}| \leqslant |\Delta/2|$ gives an exact solution to the system's Schrödinger equation. We have included a $\pi/2$ rotation about y in Eq. (10) to account for the fact that, unlike in the present case, the driving field appears in the diagonal entries of the Hamiltonian in Ref. [27]. We can apply this method to the block S_2 in Eq. (9), where $B_{y,L}^1(t) = 4\Omega(t)$ and $J = \Delta$. Following Ref. [25], in order to generate a π x rotation in S_1 and an identity operation in S_2, which would correspond to a CNOT gate, the following conditions must be satisfied: $\int_0^\tau B_{y,L}^1(t) = 2\pi$, $\chi(0) = \chi(\tau) = \pi/4$, and $\dot{\chi}(0) = \dot{\chi}(\tau) = 0$, where τ is the pulse duration (gate time). For this particular gate, the phases $\psi_{\pm}$ would only affect the local unitaries that accompany the entangling operation, and thus no constraint is imposed. An ansatz for $\chi(t)$ that automatically satisfies the previous conditions is

$$\chi(t) = A\left(\frac{t}{\tau}\right)^4\left(1 - \frac{t}{\tau}\right)^4 + \frac{\pi}{4}. \tag{12}$$

The parameters A and τ in Eq. (12) can be adjusted to generate CNOT gates with different gate times and pulse amplitudes. Figures 2(a) and 2(b) show a pair of control pulses, $B_{y,L}^1(t) = 4\Omega(t)$ with $\Omega(t)$ given by Eq. (11), that generate CNOT gates. The parameters that produce the control pulse in Fig. 2(a) are $A_{(a)} = 139.2947$ and $\tau_{(a)} = 5.54498/J$, and the parameters that produce the second pulse, Fig. 2(b), are $A_{(b)} = 61.4617$ and $\tau_{(b)} = 15.38016/J$. Therefore we get a CNOT gate in the interaction picture:

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \approx K_1^{(j)}U_{\rm int}(\tau_{(j)})K_2^{(j)}, \tag{13}$$

where $j = a, b$ identifies the set of χ parameters used $\{A_{(j)}, \tau_{(j)}\}$, $\omega = \omega_1$, and $K_i^{(j)}$ are Kronecker products of single-qubit gates (see the Appendix), where the superscript in $K_i^{(j)}$ refers to which pulse we are using to generate the CNOT (see Fig. 2). This distinction is needed since different pulses will produce different sets of single-qubit gates. Using the gate fidelity given by [29]

$$F = \frac{1}{n(n+1)}[{\rm Tr}(U^\dagger U) + |{\rm Tr}(U_0^\dagger U)|^2], \tag{14}$$

where n is the Hilbert space dimension, U is the generated gate, and U_0 is the desired gate, we compute that the CNOT gate generated in Eq. (13) with the parameters $\{A_{(a)}, \tau_{(a)}\}$ has a 99.994% fidelity and a gate time $\tau_{(a)} \approx 45$ ns (this does not take into account the gate time for the single-qubit gates $K_i^{(a)}$). On the other hand, the second set of parameters $\{A_{(b)}, \tau_{(b)}\}$ generates a CNOT gate with 99.997% fidelity and a gate time $\tau_{(b)} \approx 124$ ns. Even though this second CNOT gate has a longer gate time, the control pulse amplitude is relatively small, corresponding to a maximum Rabi frequency approximately equal to 9 MHz [at resonance, the Rabi frequency for the left qubit is equal to $B_{y,L}^1(t)/2$], and thus easier to implement in the laboratory—the up-to-date Rabi frequencies for single-spin qubits in Si quantum dots reported in the literature range from 0.5 [9] up to 35 MHz [5,6]. This shows the versatility of

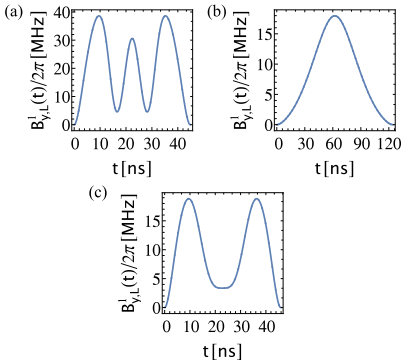

FIG. 2. (a) and (b) Control pulse that generates a maximally entangling gate locally equivalent to CNOT in (a) 45 ns with 99.994% fidelity, and (b) 124 ns with 99.997% fidelity. (c) Control pulse that generates a two-qubit gate locally equivalent to $\sqrt{\rm CNOT}$ in 46 ns. The two-piece pulse is shown in Eq. (15) and generates a CNOT gate with 99.9999% fidelity.

our scheme which, by tweaking parameters, produces two-qubit gates that conform with the experimental constraints; for example, for systems with short coherence times, shorter gates can be obtained to increase circuit depth, for long coherence times, where gate duration is less of a concern, our approach also offers gates with high fidelities.

Now we simulate the gate-fidelity loss caused by the low-frequency component of $1/f$ charge noise, which is approximated as quasistatic during the operation time. Comparisons between theory and experiment have shown that the quasistatic approximation works well, provided the pulses are not too long [17,18]. This quasistatic noise drives fluctuations in the exchange coupling, represented by δJ, as well as in the Zeeman shift terms: δE_z^1 and $\delta\Delta E_z^1$, respectively. Consequently, we form the noisy Hamiltonian by replacing the exchange coupling and Zeeman shift terms (J, E_z^1, ΔE_z^1) with their respective perturbed forms: $J+\delta J$, $E_z^1+\delta E_z^1$, and $\Delta E_z^1+\delta\Delta E_z^1$. Then the infidelity is averaged over the quasistatic noise by independently sampling the three random variables δJ, δE_z^1, and $\delta\Delta E_z^1$ over a normal distribution of standard deviation σ_δ and zero mean, with the average being taken over 500 samples for each value of σ_δ. The infidelities for both CNOT gates generated with Eq. (13) are shown in Fig. 3. Moreover, we can generate a gate locally equivalent to $\sqrt{\text{CNOT}}$ that can be used as part of a simple two-piece pulse [30,31] that can correct some of the leading error and generate a CNOT with higher fidelity. Using the same ansatz for χ given in Eq. (12) with the parameters $A_{(c)}=75.95269$ and $\tau_{(c)}=5.67638/J$, we generate a gate equivalent to $\sqrt{\text{CNOT}}$ up to single-qubit unitaries with a gate time of 46 ns. The control pulse that generates such a gate is shown in Fig. 2(c). Furthermore, the maximum amplitude of the control pulse is also relatively small in this case and within the range accessible in recent experiments [7]. The two-piece pulse is [31]

$$\begin{pmatrix}0&1&0&0\\1&0&0&0\\0&0&1&0\\0&0&0&1\end{pmatrix}\approx K_1^{(c)}U_{\text{int}}(\tau_{(c)})\kappa_2^{(c)}\times\sigma_X^R\otimes\sigma_Y^L\times\kappa_1^{(c)}U_{\text{int}}(\tau_{(c)})K_2^{(c)},\quad(15)$$

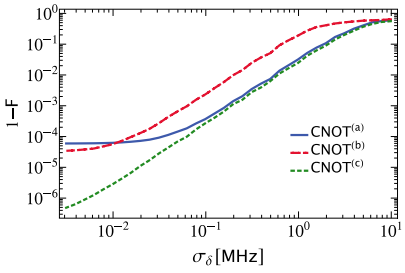

FIG. 3. Infidelity of the three CNOT$^{(j)}$ gates, where $j=a,b,c$ corresponds to the control pulses depicted in Fig. 2. The blue (solid) curve corresponds to a gate of duration 45 ns, the red (long-dashed) curve corresponds to a gate of duration 124 ns, and the green (short-dashed) curve corresponds to the composite gate, with duration 92 ns (it does not take into account the single-qubit gate times).

where $K_i^{(c)}$, $\kappa_i^{(c)}$ are Kronecker products of single-qubit gates (see the Appendix), and $\sigma_X^R\otimes\sigma_Y^L$ are π rotations around the x and y axes for the right (R) and left (L) qubits, respectively [32]. The total duration of the two-piece pulse CNOT will depend on the gate times of each of the single-qubit gates that accompany the two two-qubit gates. Figure 3 shows the infidelity of each of the three CNOT gates as functions of the charge noise strength. For noise with standard deviation less than 100 kHz, the two-piece pulse CNOT outperforms the other gates, but for higher noise strength its performance matches the control pulse shown in Fig. 2(a). In the absence of noise, the CNOT in Eq. (15) has a 99.9999% fidelity. Moreover, for a typical 1% noise level (≈200 kHz) all three CNOT gates have fidelities of at least 99%.

IV. CNOT GATE GENERATED BY A SQUARE PULSE

In the previous section we have shown that perfect entangling gates can be generated with smooth pulses in a relatively short time. Now we change our focus to simpler square pulses, which, due to their larger area per gate time, should generate entangling gates in much shorter times. In particular, we are interested in determining what type of entangling gates can be generated with a single square pulse. To this end, we consider the nonlocal, i.e., entangling, properties of a general two-qubit gate U, which are uniquely characterized by a set of three real numbers termed "local invariants" [26,33]. These invariants are related to the coefficients of the characteristic polynomial of the unitary and symmetric matrix $M(U)$, defined as $M(U)=(Q^\dagger UQ)^TQ^\dagger UQ$ where Q denotes the transformation matrix from the logical basis to the magic basis [34]. In the magic basis any single-qubit operation is represented by a real orthogonal matrix, and thus the spectrum of M is invariant under local operations. The set of local invariants are

$$G_1=\text{Re}\left[\frac{\text{tr}^2[M(U)]}{16\det U}\right],\quad G_2=\text{Im}\left[\frac{\text{tr}^2[M(U)]}{16\det U}\right],\quad G_3=\frac{\text{tr}^2[M(U)]-\text{tr}[M^2(U)]}{4\det U}.\quad(16)$$

If two two-qubit gates are equivalent up to single-qubit operations, then they have the same set of local invariants, e.g., CNOT and CZ are equivalent up to local operations and, indeed, they share the same set of local invariants: $\{G_1=0,G_2=0,G_3=1\}$.

Using the Hamiltonian in the rotating frame, Eq. (8), with constant driving amplitude $B_{y,L}^1=B_{y,R}^1=J/2$ and driving at resonance frequency $\omega=\omega_1=E_z+E_z^1-[\Delta E_z+\Delta E_z^1-J+J^2/2(\Delta E_z+\Delta E_z^1)]/2$, we calculate the local invariants of the evolution operator. Figure 4 shows the set $\{G_1,G_2,G_3\}$ versus time, where it is evident that at $t\approx 26$ ns the evolution operator is locally equivalent to a CNOT gate. Therefore, using the interaction picture Hamiltonian (7) with a square pulse of amplitude $B_{y,L}^1/(2\pi)=9.85$ MHz and gate time

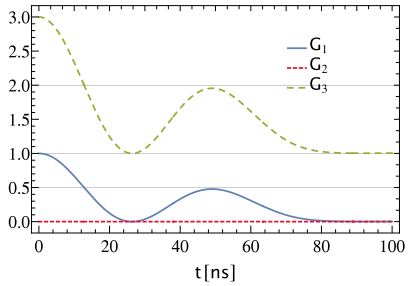

FIG. 4. Local invariants $\{G_1, G_2, G_3\}$ of the system's evolution operator $U_{\rm int}(t)$ (with square control pulse) as functions of time. At $\tau_{(1)} = 26.445$ ns ($\{G_1 = 0, G_2 = 0, G_3 = 1\}$) the evolution operator is locally equivalent to a CNOT gate, and at $\tau_{(2)} = 12.8$ ns ($\{G_1 = 0.5, G_2 = 0, G_3 = 2\}$) the evolution operator is locally equivalent to a $\sqrt{\rm CNOT}$ gate.

$\tau_{(1)} = 26.445$ ns we get a CNOT gate:

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \approx K_1^{(1)} U_{\rm int}(\tau_{(1)}) K_2^{(1)}, \tag{17}$$

where, again, $K_i^{(1)}$ are Kronecker products of single-qubit gates (see the Appendix). The CNOT gate generated in Eq. (17) has a fidelity of 99.999%. Although we find that considerably faster gates are possible when square pulses are used, the experimental challenges with creating such ideal waveforms may lead to additional pulse generation errors, and the smooth pulses of the previous section may ultimately yield better results in practice.

In Sec. III we use a gate locally equivalent to $\sqrt{\rm CNOT}$ as part of a two-piece pulse sequence, Eq. (15), that corrects some of the leading error and generates a high-fidelity CNOT gate. Similarly, using the local invariants of the $\sqrt{\rm CNOT}$ gate (i.e., $\{G_1 = 0.5, G_2 = 0, G_3 = 2\}$), we see in Fig. 4 that the system's evolution operator also generates a gate locally equivalent to $\sqrt{\rm CNOT}$ at $\tau_{(2)} = 12.8$ ns. Accordingly, the two-piece pulse sequence is

$$\begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \approx K_1^{(2)} U_{\rm int}(\tau_{(2)}) \kappa_2^{(2)} \times \sigma_X^R \otimes \sigma_X^L \times \kappa_1^{(2)} U_{\rm int}(\tau_{(2)}) K_2^{(2)}, \tag{18}$$

where $K_i^{(2)}$, $\kappa_i^{(2)}$ are Kronecker products of single-qubit gates (see the Appendix), and $\sigma_X^R \otimes \sigma_X^L$ are π rotations around the x axis for both qubits. This pulse sequence generates a CNOT gate with 99.9999% fidelity in the absence of noise.

Next, we calculate the average gate infidelity of the CNOT gates, Eqs. (17) and (18), versus noise strength, as shown in Fig. 5. The noise is assumed to be quasistatic and causing fluctuations in the exchange coupling δJ and both Zeeman shift terms δE_z^1 and $\delta \Delta E_z^1$. The infidelity is calculated by independently sampling the three random variables over a

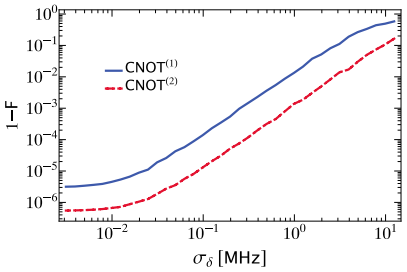

FIG. 5. Infidelity of the CNOT gates vs noise strength. The blue (solid) curve corresponds to a single-shot CNOT gate, Eq. (17), while the red (dashed) curve corresponds to a composite pulse gate, Eq. (18).

normal distribution of standard deviation σ_δ and zero mean. Its average is taken over 500 samples for each value of σ_δ. Figure 5 shows that the two-piece pulse increases the CNOT fidelity by at least an order of magnitude. Furthermore, for a typical 1% noise level (≈200 kHz) both CNOT gates have fidelities larger than 99.9%.

V. θ-CPHASE AND CZ GATES BASED ON SECH PULSES

CZ and CPHASE gates naturally appear in certain quantum algorithms such as the quantum Fourier transform and in quantum simulation. To achieve shallower circuit depth, it is therefore of interest to be able to directly implement these gates instead of decomposing them in terms of CNOTs. In previous works [23–25] it has been demonstrated that by using a hyperbolic secant (sech) control pulse, $\Omega(t) = \Omega_0 \mathrm{sech}(\sigma t)$ with $\Omega_0/\sigma =$ integer, a cyclic evolution can be induced, and thus local phase gates are easily implemented. However, the method used in those works involves an auxiliary excited state and, therefore, it is not directly applicable to the current case. Here we modify the aforementioned method such that the evolution is within the logical space.

To illustrate the basics of the method, we apply a sech pulse in a general two-level system, whose Hamiltonian in the rotating frame is given by

$$\mathcal{H}(t) = \begin{pmatrix} -\Delta/2 & \Omega(t) \\ \Omega(t) & \Delta/2 \end{pmatrix}. \tag{19}$$

The control pulse $\Omega(t) = \sigma \mathrm{sech}(\sigma t - n\pi/2)$, with n being a positive real number, generates a general phase gate $\mathcal{U} = \mathrm{diag}\{e^{-i\theta}, e^{i\theta}\}$ with gate time $\tau = n\pi/\sigma$. The induced phase is given by

$$\theta = -2 \arctan\left[\frac{\sigma}{\Delta}\right] - \frac{\Delta\tau}{2}. \tag{20}$$

Here the precision of the induced phase θ is proportional to n.

We can easily extend this method to a two-qubit system with a Hamiltonian

$$\mathcal{H}(t) = \begin{pmatrix} -\Delta_1/2 & \Omega(t) & 0 & 0 \\ \Omega(t) & \Delta_1/2 & 0 & 0 \\ 0 & 0 & -\Delta_2/2 & \Omega(t) \\ 0 & 0 & \Omega(t) & \Delta_2/2 \end{pmatrix}, \tag{21}$$

which is equivalent to the one in Eq. (8) if we again drop the terms involving $B^1_{y,R}\delta^{(2)}_{\pm}$ and use $\delta^{(1)}_{\pm}\approx 1$. In order to generate a generalized θ-CPHASE gate, defined as $\theta-\widetilde{\text{CPHASE}}=\text{diag}\{e^{i\theta_1},e^{i\theta_2},e^{i\theta_3},e^{i\theta_4}\}$ with $\theta_1-\theta_2-\theta_3+\theta_4=\pm\theta$, where θ is the induced phase on the target qubit, we use the same sech pulse $\Omega(t)=\sigma\,\text{sech}(\sigma t-n\pi/2)$. The resulting phase θ is now given by

$$\frac{\theta}{2}=-2\arctan\left[\frac{\sigma}{\Delta_1}\right]+2\arctan\left[\frac{\sigma}{\Delta_2}\right]-\frac{(\Delta_1-\Delta_2)\tau}{2}, \tag{22}$$

with a gate time $\tau=n\pi/\sigma$.

In order to generate a generalized θ-CPHASE gate with the spin Hamiltonian, we compare Eqs. (8) and (21), and define $\alpha=\Delta_1/J$, where $\Delta_1=\omega-\omega_1$ and ω_1 is the resonance frequency of the diagonal block S_1. Moreover, we define m to be σ/J and, since the difference between the S_1 and S_2 detunings is equal to $-J$, we have that $\Delta_2/J=\alpha+1$. Consequently, Eq. (22) becomes

$$\frac{\theta}{2}=-2\arctan\left[\frac{m}{\alpha}\right]+2\arctan\left[\frac{m}{\alpha+1}\right]+\frac{n\pi}{2m}. \tag{23}$$

Solving Eq. (23) for α gives us four solutions:

$$\alpha^{(1)}_{\pm}=\frac{1}{2}\left\{-1\pm\sqrt{1-4m^2+4m\cot\left[\frac{1}{4}\left(\frac{n\pi}{m}-\theta\right)\right]}\right\}, \tag{24}$$

$$\alpha^{(2)}_{\pm}=\frac{1}{2}\left\{-1\pm\sqrt{1-4m^2-4m\tan\left[\frac{1}{4}\left(\frac{n\pi}{m}-\theta\right)\right]}\right\}. \tag{25}$$

These four equations, owing to the presence of the trigonometric functions tangent and cotangent, have real solutions only when m is within the following respective ranges:

$$x^{(1)}_r\leqslant m<\frac{n\pi}{4\pi r+\theta}, \tag{26}$$

$$x^{(2)}_r\leqslant m<\frac{n\pi}{2\pi+4\pi r+\theta}, \tag{27}$$

where r is a positive integer and $x^{(1)}_r$ ($x^{(2)}_r$) is the root of the function $1-4x^2+4x\cot[\frac{1}{4}(\frac{n\pi}{x}-\theta)]$ $(1-4m^2-4m\tan[\frac{1}{4}(\frac{n\pi}{m}-\theta)])$ that is smaller than and closest in magnitude to $\frac{n\pi}{4\pi r+\theta}$ ($\frac{n\pi}{2\pi+4\pi r+\theta}$). To show the efficacy of the method, Fig. 6 depicts the infidelities and gate times for several θ-CPHASE gates that were obtained with the interaction picture Hamiltonian (7) and the parameters listed at the end of Sec. II, along with $n=3$, $\sigma=(0.9999)\frac{n\pi}{6\pi+\theta}J$, $\omega=\omega_1+\alpha^{(2)}_{+}J$, where $\alpha^{(2)}_{+}$ is given by Eq. (25), and $B^1_{y,L}(t)=4\sigma\,\text{sech}(\sigma t-n\pi/2)$. All the θ-CPHASE gates in Fig. 6 have fidelities exceeding 99.9999%.

Note that the regular θ-CPHASE$=\text{diag}\{1,e^{i\theta},1,1\}$ can be obtained from θ-CPHASE by applying subsequent single-qubit z rotations and including a global phase. Moreover, after calculating the local invariants, Eq. (16), for θ-CPHASE we see that $\theta-\widetilde{\text{cphase}}$ is maximally entangling only when $\theta_1-\theta_2-\theta_3+\theta_4=\pm\pi$, which corresponds to a generalized control-Z gate, $\tilde{C}\tilde{Z}$ [25]. Accordingly, using the same parameters employed in the calculation of the θ-CPHASE gates shown in Fig. 6, we generate a regular CZ gate with the interaction picture Hamiltonian:

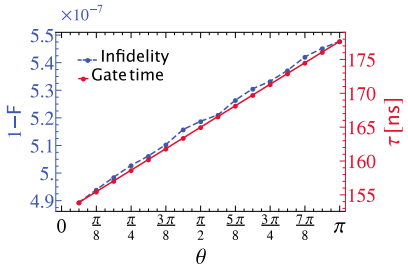

FIG. 6. Infidelities and gate times for several θ-CPHASE gates. The abscissa gives the magnitude of the phase angle of the CPHASE gate. The peak amplitudes of the control pulses, see Fig. 7(a), of these gates vary between 39 MHz (for $\theta=\pi/16$) and 34 MHz (for $\theta=\pi$).

$$\begin{pmatrix}1&0&0&0\\0&-1&0&0\\0&0&1&0\\0&0&0&1\end{pmatrix}\approx K^{(\alpha)}_1U_{\text{int}}(\tau_{(\alpha)})K^{(\alpha)}_2, \tag{28}$$

where $K^{(\alpha)}_i$ are Kronecker products of single-qubit gates (see the Appendix). The CZ gate generated in Eq. (28) has a 99.9999% fidelity and a gate time $\tau\approx 178$ ns (this does not take into account the gate time for the single-qubit gates K^{CZ}_i). Moreover, Fig. 7(a) shows the control pulse $B^1_{y,L}(t)=4\sigma\,\text{sech}(\sigma t-n\pi/2)$ versus time, which has a peak amplitude of 33.8 MHz. Furthermore, Fig. 7(b) shows that the amplitude of the sech pulse is, as expected, inversely proportional to the gate time and that the gate fidelity (99.9999%) is slightly larger in longer gates. Therefore, a shorter gate would require a larger peak amplitude which is currently unattainable due to experimental limitations: The largest Rabi frequency reported in the literature is 35 MHz [5,6] which corresponds to a pulse peak amplitude of 70 MHz.

In similar fashion to the use of $\sqrt{\text{CNOT}}$ in Eq. (15), we use the $\pi/2$-CPHASE gate as part of a two-piece pulse sequence that corrects some of the leading error and generates a CZ gate. The $\pi/2$-CPHASE pulse is similar to the one depicted in Fig. 7(a) with a slightly larger peak amplitude $B^1_{y,L}(\tau_{(\beta)}/2)=36.4$ MHz and shorter gate time $\tau_{(\beta)}=165$ ns. The two-piece pulse sequence is

$$\begin{pmatrix}1&0&0&0\\0&-1&0&0\\0&0&1&0\\0&0&0&1\end{pmatrix}\approx K^{(\beta)}_1U_{\text{int}}(\tau_{(\beta)})\kappa^{(\beta)}_2\times\sigma^R_X\otimes\sigma^L_Y\times\kappa^{(\beta)}_1U_{\text{int}}(\tau_{(\beta)})K^{(\beta)}_2, \tag{29}$$

where $K^{(\beta)}_i$ and $\kappa^{(\beta)}_i$ are Kronecker products of single-qubit gates (see the Appendix). This pulse sequence generates a CZ gate with 99.9999% fidelity in the absence of noise.

Finally, we calculate the average gate infidelity of the previous CZ gates, this is depicted in Fig. 7(c). The noise is assumed to be quasistatic and affecting the exchange coupling

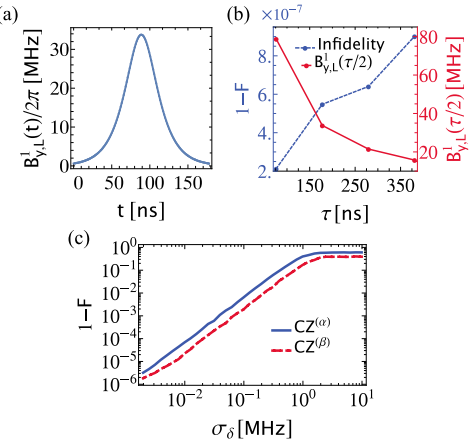

FIG. 7. (a) Control pulse that generates a maximally entangling gate locally equivalent to CZ in 178 ns with 99.9999% fidelity. Its peak amplitude is equal to 33.8 MHz. (b) Infidelity (left ordinate, blue dashed curve) and pulse peak amplitude (right ordinate, red solid curve) versus gate time. (c) Infidelity of the CZ gates, Eqs. (28) and (29), versus noise strength. The blue (solid) curve corresponds to a single-shot CZ gate (28), while the red (dashed) curve corresponds to a composite pulse gate (29).

δJ and both Zeeman shift terms δE_z^1 and $\delta\Delta E_z^1$. We independently sample the three random variables $(\{\delta J, \delta E_z^1, \delta\Delta E_z^1\})$ over a normal distribution of standard deviation σ_δ and zero mean. Its average is taken over 500 samples for each value of σ_δ. As seen in Fig. 7(c), the two-piece CZ gate has slightly better fidelity than the single-shot CZ gate, for a typical 1% noise level (≈200 kHz) the two-piece pulse produces a CZ gate with an average fidelity larger than 99%.

VI. CONCLUSIONS

Using a protocol to analytically design the unitary evolution of an arbitrary two-level system, we have presented a set of control pulses that generate high-fidelity CNOT, CPHASE, and CZ gates in the system formed by two electrons confined in a silicon double quantum dot. With these smooth control pulses, and using real parameters from the experiment reported in Ref. [7], we have predicted, in the absence of noise, single-shot CNOT gates with more than 99.99% fidelity in times as short as 45 ns, and a CZ gate with 99.9999% fidelity and gate time of 178 ns. Moreover, after calculating the local invariants of the evolution operator we have shown that a simple square pulse can generate a CNOT gate in less than 27 ns and with a fidelity of 99.999%. In addition, we have shown that high-fidelity (∼99.9999%) CPHASE gates with arbitrary phase angles can be implemented with simple sech pulses. Furthermore, using a simple two-piece pulse sequence interrupted by single-qubit gates, we have appreciably improved the fidelity of CNOT and CZ gates in the presence of quasistatic noise. In conclusion, we have presented a set of control pulses for fast high-fidelity two-qubit gates, whose predicted performances do not have parallel in other theoretical and experimental works, and thus they are immediately relevant for current experiments on exchange-coupled single-spin qubits in silicon quantum dots.

ACKNOWLEDGMENTS

We are grateful to J. R. Petta for useful discussions and feedback on the manuscript. This work is supported by the Army Research Office (W911NF-17-0287) and the U.S. Office of Naval Research (Grant No. N00014-17-1-2971).

APPENDIX A: SINGLE-QUBIT GATES FOR CNOT and CZ GATES

Single-qubit gates for the CNOT gates, Eqs. (13) and (15), in Sec. III are

$$K_1^{(a)} = \exp[i(0.00564\sigma_X - 0.006156\sigma_Y + 0.507546\sigma_Z)]_R \otimes \exp[i(-0.107076\sigma_X + 0.091219\sigma_Y - 0.663155\sigma_Z)]_L,$$
$$K_2^{(a)} = \exp[i(0.005831\sigma_X + 0.006493\sigma_Y - 0.320524\sigma_Z)]_R \otimes \exp[i(0.104122\sigma_X - 0.008626\sigma_Y + 0.034753\sigma_Z)]_L,$$
$$K_1^{(b)} = \exp[i(0.000429\sigma_X - 0.001655\sigma_Y - 0.836667\sigma_Z)]_R \otimes \exp[i(-0.101941\sigma_X - 0.204093\sigma_Y + 1.37345\sigma_Z)]_L,$$
$$K_2^{(b)} = \exp[i(0.00087\sigma_X + 0.001517\sigma_Y + 1.018818\sigma_Z)]_R \otimes \exp[i(0.153317\sigma_X + 0.090985\sigma_Y - 0.816242\sigma_Z)]_L,$$
$$K_1^{(c)} = \exp[i(0.003253\sigma_X - 0.000834\sigma_Y - 1.120708\sigma_Z)]_R \otimes \exp[i(0.356721\sigma_X - 0.003073\sigma_Y - 0.400727\sigma_Z)]_L,$$
$$K_2^{(c)} = \exp[i(0.04293\sigma_X + 1.57321\sigma_Y - 0.000016\sigma_Z)]_R \otimes \exp[i(-0.174612\sigma_X + 0.106183\sigma_Y - 1.42782\sigma_Z)]_L,$$
$$\kappa_1^{(c)} = \exp[i(0.746971\sigma_X + 1.379465\sigma_Y + 0.002712\sigma_Z)]_R \otimes \exp[i(-0.248607\sigma_X - 0.622127\sigma_Y - 0.313053\sigma_Z)]_L,$$
$$\kappa_2^{(c)} = \exp[i(0.208987\sigma_X - 1.553846\sigma_Y + 0.000769\sigma_Z)]_R \otimes \exp[i(0.328546\sigma_X + 0.577004\sigma_Y - 0.23246\sigma_Z)]_L, \quad \text{(A1)}$$

where the subindices L and R refer to the left and right qubits, respectively.

The single-qubit gates for the CNOT gate, Eqs. (17) and (18), in Sec. IV are

$$K_1^{(1)} = \exp[i(0.031741\sigma_X - 0.024265\sigma_Y + 2.143595\sigma_Z)]_R \otimes \exp[i(-0.157713\sigma_X - 0.835226\sigma_Y + 0.22501\sigma_Z)]_L,$$
$$K_2^{(1)} = \exp[i(-0.010813\sigma_X - 0.011452\sigma_Y - 0.818712\sigma_Z)]_R \otimes \exp[i(0.402313\sigma_X + 0.941847\sigma_Y + 0.138859\sigma_Z)]_L,$$
$$K_1^{(2)} = \exp[i(1.23768\sigma_X + 0.97176\sigma_Y - 0.02182\sigma_Z)]_R \otimes \exp[i(-0.0331\sigma_X - 0.88781\sigma_Y - 0.06691\sigma_Z)]_L,$$

$$K_2^{(2)} = \exp\left[i(-0.00418\sigma_X - 0.01482\sigma_Y - 0.3255\sigma_Z)\right]_R \otimes \exp\left[i(0.64823\sigma_X - 0.68884\sigma_Y - 0.49805\sigma_Z)\right]_L,$$
$$\kappa_1^{(2)} = \exp\left[i(-0.003159\sigma_X + 0.013948\sigma_Y + 0.233523\sigma_Z)\right]_R \otimes \exp\left[i(-0.007219\sigma_X - 0.106195\sigma_Y + 0.204111\sigma_Z)\right]_L,$$
$$\kappa_2^{(2)} = \exp\left[i(0.001186\sigma_X - 0.015079\sigma_Y + 0.027787\sigma_Z)\right]_R \otimes \exp\left[i(0.060973\sigma_X - 0.124889\sigma_Y - 1.244681\sigma_Z)\right]_L. \quad \text{(A2)}$$

The single-qubit gates for the CZ gate, Eqs. (28) and (29), in Sec. V are

$$K_1^{(\alpha)} = \exp\left[i(-0.036888\sigma_X + 1.571355\sigma_Y + 0.00069\sigma_Z)\right]_R \otimes \exp\left[i(0.000045\sigma_X - 0.000142\sigma_Y + 0.34968\sigma_Z)\right]_L,$$
$$K_2^{(\alpha)} = \exp\left[i(-0.626514\sigma_X + 1.440206\sigma_Y - 0.002372\sigma_Z)\right]_R \otimes \exp\left[i(0.000262\sigma_X + 0.000279\sigma_Y - 0.344756\sigma_Z)\right]_L,$$
$$K_1^{(\beta)} = \exp\left[i(0.001121\sigma_X - 0.00114\sigma_Y + 0.529903\sigma_Z)\right]_R \otimes \exp\left[i(1.567642\sigma_X - 0.09954\sigma_Y - 0.000507\sigma_Z)\right]_L,$$
$$K_2^{(\beta)} = \exp\left[i(-0.222578\sigma_X + 1.555129\sigma_Y - 0.002324\sigma_Z)\right]_R \otimes \exp\left[i(0.000321\sigma_X - 0.000072\sigma_Y - 0.692129\sigma_Z)\right]_L,$$
$$\kappa_1^{(\beta)} = \exp\left[i(0.852858\sigma_X + 1.319655\sigma_Y + 0.002282\sigma_Z)\right]_R \otimes \exp\left[i(0.000107\sigma_X + 0.000308\sigma_Y + 0.259482\sigma_Z)\right]_L,$$
$$\kappa_2^{(\beta)} = \exp\left[i(-0.697808\sigma_X + 1.406964\sigma_Y - 0.002295\sigma_Z)\right]_R \otimes \exp\left[i(0.000197\sigma_X + 0.000242\sigma_Y + 0.416163\sigma_Z)\right]_L. \quad \text{(A3)}$$

[1] D. Loss and D. P. DiVincenzo, Quantum computation with quantum dots, Phys. Rev. A **57**, 120 (1998).

[2] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. L. Hollenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson, Silicon quantum electronics, Rev. Mod. Phys. **85**, 961 (2013).

[3] M. Veldhorst, J. C. C. Hwang, C. H. Yang, A. W. Leenstra, B. de Ronde, J. P. Dehollain, J. T. Muhonen, F. E. Hudson, K. M. Itoh, A Morello, and A. S. Dzurak, An addressable quantum dot qubit with fault-tolerant control-fidelity, Nat. Nanotechnol. **9**, 981 (2014).

[4] E. Kawakami, T. Jullien, P. Scarlino, D. R. Ward, D. E. Savage, M. G. Lagally, V. V. Dobrovitski, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen, Gate fidelity and coherence of an electron spin in an Si/SiGe quantum dot with micromagnet, Proc. Natl. Acad. Sci. **113**, 11738 (2016).

[5] K. Takeda, J. Kamioka, T. Otsuka, J. Yoneda, T. Nakajima, M. R. Delbecq, S. Amaha, G. Allison, T. Kodera, S. Oda, and S. Tarucha, A fault-tolerant addressable spin qubit in a natural silicon quantum dot, Sci. Adv. **2**, e1600694 (2016).

[6] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, N. Usami, K. M. Itoh, and S. Tarucha, A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%, Nat. Nanotechnol. **13**, 102 (2018).

[7] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard, and J. R. Petta, Resonantly driven CNOT gate for electron spins, Science **359**, 439 (2018).

[8] R. Brunner, Y.-S. Shin, T. Obata, M. Pioro-Ladrière, T. Kubo, K. Yoshida, T. Taniyama, Y. Tokura, and S. Tarucha, Two-Qubit Gate of Combined Single-Spin Rotation and Interdot Spin Exchange in a Double Quantum Dot, Phys. Rev. Lett. **107**, 146801 (2011).

[9] M. Veldhorst, C. H. Yang, J. C. C. C Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, A two-qubit logic gate in silicon, Nature (London) **526**, 410 (2015).

[10] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen, A programmable two-qubit quantum processor in silicon, Nature (London) **555**, 633 (2018).

[11] W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R. C. C. C. Leon, M. A. Fogarty, J. C. C. C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak, Fidelity benchmarks for two-qubit gates in silicon, Nature **569**, 532 (2019).

[12] X. Xue, T. F. Watson, J. Helsen, D. R. Ward, D. E. Savage, M. G. Lagally, S. N. Coppersmith, M. A. Eriksson, S. Wehner, and L. M. K. Vandersypen, Benchmarking Gate Fidelities in a Si/SiGe Two-Qubit Device, Phys. Rev. X **9**, 021011 (2019).

[13] P. Huang, N. M. Zimmerman, and G. W. Bryant, Spin decoherence in a two-qubit CPHASE gate: The critical role of tunneling noise, npj Quantum Inf. **4**, 62 (2018).

[14] J. P. G. van Dijk, E. Kawakami, R. N Schouten, M. Veldhorst, L. M K Vandersypen, M. Babaie, E. Charbon, and F. Sebastiano, The impact of classical control electronics on qubit fidelity, arXiv:1803.06176.

[15] K. W. Chan, W. Huang, C. H. Yang, J. C. C. Hwang, B. Hensen, T. Tanttu, F. E. Hudson, K. M. Itoh, A. Laucht, A. Morello, and A. S. Dzurak, Assessment of a Silicon Quantum Dot Spin Qubit Environment via Noise Spectroscopy, Phys. Rev. Appl. **10**, 044017 (2018).

[16] K. M. Itoh and H. Watanabe, Isotope engineering of silicon and diamond for quantum computing and sensing applications, MRS Commun. **4**, 143 (2014).

[17] M. D. Reed, B. M. Maune, R. W. Andrews, M. G. Borselli, K. Eng, M. P. Jura, A. A. Kiselev, T. D. Ladd, S. T. Merkel, I. Milosavljevic, E. J. Pritchett, M. T. Rakher, R. S. Ross, A. E. Schmitz, A. Smith, J. A. Wright, M. F. Gyure, and A. T. Hunter, Reduced Sensitivity to Charge Noise in Semiconductor Spin Qubits via Symmetric Operation, Phys. Rev. Lett. **116**, 110402 (2016).

[18] F. Martins, F. K. Malinowski, P. D. Nissen, E. Barnes, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuemmeth, Noise Suppression Using Symmetric Exchange Gates in Spin Qubits, Phys. Rev. Lett. **116**, 116801 (2016).

[19] C. Zhang, R. E. Throckmorton, X.-C. Yang, X. Wang, E. Barnes, and S. Das Sarma, Randomized Benchmarking of Barrier versus Tilt Control of a Singlet-Triplet Qubit, Phys. Rev. Lett. **118**, 216802 (2017).

[20] J. Zeng, C. H. Yang, A. S. Dzurak, and E. Barnes, Geometric formalism for constructing arbitrary single-qubit dynamically corrected gates, Phys. Rev. A **99**, 052321 (2019).

[21] M. Russ, D. M. Zajac, A. J. Sigillito, F. Borjans, J. M. Taylor, J. R. Petta, and G. Burkard, High-fidelity quantum gates in Si/SiGe double quantum dots, Phys. Rev. B **97**, 085421 (2018).

[22] U. Güngördü and J. P. Kestner, Pulse sequence designed for robust C-phase gates in SiMOS and Si/SiGe double quantum dots, Phys. Rev. B **98**, 165301 (2018).

[23] S. E. Economou, L. J. Sham, Y. Wu, and D. G. Steel, Proposal for optical $U(1)$ rotations of electron spin trapped in a quantum dot, Phys. Rev. B **74**, 205415 (2006).

[24] S. E. Economou, High-fidelity quantum gates via analytically solvable pulses, Phys. Rev. B **85**, 241401(R) (2012).

[25] S. E. Economou and E. Barnes, Analytical approach to swift nonleaky entangling gates in superconducting qubits, Phys. Rev. B **91**, 161405(R) (2015).

[26] Y. Makhlin, Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations, Quantum Inf. Process. **1**, 243 (2002).

[27] E. Barnes, Analytically solvable two-level quantum systems and Landau-Zener interferometry, Phys. Rev. A **88**, 013818 (2013).

[28] E. Barnes and S. Das Sarma, Analytically Solvable Driven Time-Dependent Two-Level Quantum Systems, Phys. Rev. Lett. **109**, 060401 (2012).

[29] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, Phys. Lett. A **367**, 47 (2007).

[30] C. D. Hill, Robust Controlled-NOT Gates from Almost Any Interaction, Phys. Rev. Lett. **98**, 180501 (2007).

[31] F. A. Calderon-Vargas and J. P. Kestner, Dynamically Correcting a CNOT Gate for any Systematic Logical Error, Phys. Rev. Lett. **118**, 150502 (2017).

[32] We choose to use an $\sigma_X \otimes \sigma_Y$ operation instead, let us say, an $\sigma_X \otimes \sigma_X$ operation because the former corrects more leading errors.

[33] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Geometric theory of nonlocal two-qubit operations, Phys. Rev. A **67**, 042313 (2003).

[34] S. Hill and W. K. Wootters, Entanglement of a Pair of Quantum Bits, Phys. Rev. Lett. **78**, 5022 (1997).

Part III

Quantum Simulation

Chapter 6

Measurement Error Mitigation for Variational Quantum Algorithms

This chapter includes pages from our publication [2]

G. S. Barron and C. J. Wood, "Measurement Error Mitigation for Variational Quantum Algorithms", arXiv:2010.08520 (2020)

This work has been posted to arXiv.org.

The contributions to this work are as follows: My contributions were carrying out numerical simulations, preparation of figures, primary writing of the manuscript, bibliography preparation and conducting background material search. C.J.W. supervised and assisted with the development of the code.

arXiv:2010.08520v1 [quant-ph] 16 Oct 2020

Measurement Error Mitigation for Variational Quantum Algorithms

George S. Barron[1,*] and Christopher J. Wood[2,†]

[1] *Department of Physics, Virginia Tech, Blacksburg, VA 24061, U.S.A*
[2] *IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA*

(Dated: October 19, 2020)

Variational Quantum Algorithms (VQAs) are a promising application for near-term quantum processors, however the quality of their results is greatly limited by noise. For this reason, various error mitigation techniques have emerged to deal with noise that can be applied to these algorithms. Recent work introduced a technique for mitigating expectation values against correlated measurement errors that can be applied to measurements of 10s of qubits. We apply these techniques to VQAs and demonstrate its effectiveness in improving estimates to the cost function. Moreover, we use the data resulting from this technique to experimentally characterize measurement errors in terms of the device connectivity on devices of up to 20 qubits. These results should be useful for better understanding the near-term potential of VQAs as well as understanding the correlations in measurement errors on large, near-term devices.

I. INTRODUCTION

The development of quantum computers and their applications have rapidly accelerated over the last few years. Several different hardware platforms have been experimentally realized at varying scales [1–4], and there has been an increased focus on studying algorithms and applications that can be run on these noisy near-term devices. Some of the most promising algorithms for near-term devices are Variational Quantum Algorithms (VQAs) [5–12], which use a quantum device to evaluate an objective function that is minimized using a classical optimizer. Instances of this algorithm include the Variational Quantum Eigensolver (VQE) [6–10, 13, 14] and Quantum Approximate Optimization Algorithm (QAOA) [5, 15, 16]. Recent process in this field includes, for example, improvements to the measurement process [17–20], selection of variational ansätze [21–24], and optimization. Moreover, they have been demonstrated experimentally on a variety of physical systems [9, 25–27].

The ability of current experimental implementaions of VQAs to produce accurate results is limited by noise on the device, despite these algorithms not explicitly requiring error correction. All proposed platforms for quantum computation experience some combination of different errors including decoherence, calibration errors, leakage, cross-talk, and measurement errors. With superconducting systems, cross-talk and measurement are among the largest sources of errors [28]. Error mitigation techniques have been developed for VQAs to reduce the amount of error on near-term devices in the absence of error correction. For example, *extrapolation to the zero-noise limit* [29, 30] uses pulse-level control to mitigate expectation values against decoherence, requiring only a constant factor of overhead in the number of circuits executed. Techniques have also been developed to characterize and mitigate against measurement errors [31–33]. Moreover, Ref. [34] analyzes the underlying model and provides rigorous improvements to correction techniques. Related techniques have also been used in VQE experiments [9] and are implemented in the IBM Qiskit package [35]. Recently, Ref. [36] has introduced a readout error mitigation technique, which we will call Continuous-Time Markov-Process Error Mitigation (CTMP-EM), that mitigates expectation values against correlated measurement errors. The n-qubit calibration procedure for CTMP-EM requires as few as $n+2$ circuits to execute and $O(n^2)$ parameters to fit. In this work, they used it to mitigate estimates of the fidelity of graph states using stabilizer measurements, as well as expectation values of stabilizers with respect to Clifford circuits.

In this paper we apply the CTMP-EM technique to experimentally characterize long-range correlations in measurement errors, and to calibrate error mitigated measurements in several quantum computers. In addition to demonstrating the presence of long-range correlations in these devices, we are also able to show that these long-range correlations can be as strong between distant qubits as they are between neighboring qubits. Moreover, rather than only considering the global minimum for the VQE objective function, we consider the objective function holistically. Evaluating the objective function at other parameter values is important for various tasks, for example the application of the ubiquitous parameter shift rule [37] for analytic gradient computation in variational experiments. For the Fermi-Hubbard model, we demonstrate that CTMP-EM can fundamentally improve the shape of the objective function for the VQE not only at its minimum, but globally as well.

This article is structured as follows. In Section II we review VQAs and the CTMP-EM technique. In Section III we demonstrate that applying CTMP-EM to the VQE algorithm changes the shape of the objective function. In Section IV we use the calibration data from CTMP-EM to analyze the long-range correlations in readout er-

* gbarron@vt.edu
† cjwood@us.ibm.com

rors on devices. We also compare several different IBM Quantum superconducting devices with the CTMP-EM calibration data. In Section V we conclude.

II. BACKGROUND

A. Variational Quantum Algorithms

A VQA is an optimization problem $\min_\theta f(\theta)$ where the objective function $f(\theta)$ is evaluated using a quantum device within the classical optimization loop. Two examples of VQAs are the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA). In each case, the objective function is $f(\theta) = \langle\psi(\theta)| H |\psi(\theta)\rangle$ for a Hamiltonian H and parameterized state $|\psi(\theta)\rangle$. In the case of a VQE, H corresponds to the Hamiltonian of some physical system. In the case of QAOA, H corresponds to a cost function with binary variables. In QAOA, the state $|\psi(\theta)\rangle$ is prepared by alternating unitary gates corresponding to evolution of the cost Hamiltonian H and some mixer operator(s). In VQE, the state $|\psi(\theta)\rangle$ takes the form of a specific trial-state ansatz. The objective function $f(\theta)$ is computed by performing measurements to estimate the expectation value $\langle\psi(\theta)| H |\psi(\theta)\rangle$ on a quantum device.

VQAs are appealing for near-term devices as they are agnostic to errors in state preparation so long as the true minimum expecation value can be reached. This has been demonstrated on different hardware platforms, with different instances of VQAs, and varying degrees of accuracy. For the case of VQE it has been shown that, in special cases, the parameters that minimize the objective function are resilient to certain types of noise [38]. In general, however, device noise significantly impacts the value of the objective function at that minimum and other points in parameter space.

B. Continuous Time Markov Process Error Mitigation

Error mitigation techniques generally aim to improve the accuracy of the results obtained from using a *noisy* quantum device. Typically, each technique mitigates against a certain kind of error. Measurement errors are a large source of error in VQA experiments on near-term devices. Measurement errors are modeled by a stochastic *assignment matrix* A acting on the state before readout. The elements of the matrix $A_{y,x} = P(y|x)$ are the probability of reading out the basis state y where x was prepared. A general n-qubit stocastic matrix has $2^n(2^n-1)$ independent real parameters and thus can only be completely characterized for a small number of qubits. Once this matrix is determined, applying the inverse matrix to a given probability distribution undoes the effect of readout error, however since the inverse is not in general a stochastic matrix the resulting output is not a valid probability distribution.

The CTMP-EM algorithm of Ref. [36] circumvents dealing with the A matrix directly, and is used to mitigate expectation values computed with a quantum device. This is done by modeling $A = e^G$, where $G = \sum_i r_i G_i$ with rates $0 \leq r_i \in \mathbb{R}$ and operators G_i that generate different readout errors. In particular, for multi-qubit bitstrings a, b, the readout error $a \to b$ corresponds to the generator $G_i = |b\rangle\langle a| - |a\rangle\langle a|$. For readout errors $0 \leftrightarrow 1$, $01 \leftrightarrow 10$, and $11 \leftrightarrow 00$ on subsets of n-qubits, CTMP-EM can determine the corresponding r_i with as few as $n+2$ circuits. Once r_i have been determined, Algorithm 1 of Ref. [36] can be used to estimate measurement error mitigated expectation values. The idea of this algorithm is to use the measurement counts collected from an expectation value experiment and classically post-process them by simulating a Markov process that applies A^{-1} to the resulting distribution while simultaneously computing the expectation value of the desired operator. This algorithm runs in $n^2\gamma e^{4\gamma}/\delta^2$ time, where n is the number of qubits, δ is the desired additive error in the expectation value of the Pauli string considered, and $\gamma = -\max_x \langle x| G |x\rangle$ for bitstrings x. Ref. [36] shows experimentally that $\gamma \approx 0.05n$ for up to $n = 20$. Hence, the algorithm requires exponential post-processing time, but in practical situations this technique is applicable up to about 50 qubits.

To study the effects of measurement error mitigation on VQAs we perform simulations of noisy measurements which simulate the device including only readout errors present on individual qubits and all qubit pairs. We choose readout error rates to reflect the *ibmq_boeblingen* device. However for device characterization in Section IV we use the real IBM Quantum devices.

III. MITIGATING VQA OBJECTIVE FUNCTIONS

A. Ground State Energy Mitigation

One of the main applications of VQE is to estimate the ground state energy of a Hamiltonian by minimizing the measured operator expectation value over the variational parameters. Both the expectation value and gradient estimates are particularly senstive to measurement errors which can greatly effect the the performance of the classical optimizer that depends on these values. This makes measurement error mitigation essential for improving the accuracy of VQE and other VQAs. To investigate the impact of CTMP-EM on variational algorithms, we choose the Fermi-Hubbard model which describes a system of Fermions interacting on a lattice [39]. The Hamiltonian

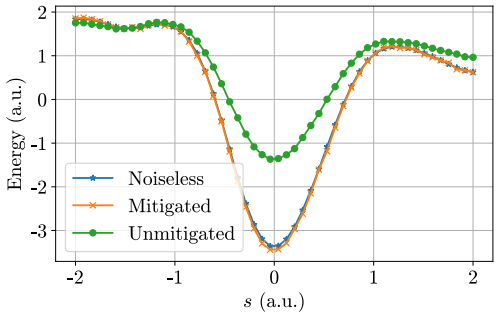

FIG. 1. Sweep of the objective function f through the global minimum ($s = 0$) computed several ways. Here, f corresponds to the energy of the Fermi-Hubbard model for the given state. In the "Noiseless" case, there is no error in the simulation and no mitigation. In the "Unmitigated" case, we include readout error, but no mitigation. In the "Mitigated" case, we include readout error and CTMP-EM. Estimates of the ground state energy ($s = 0$) and surrounding points are significantly improved by applying CTMP-EM.

is

$$H = -t \sum_{\langle j,k \rangle} \sum_{\sigma} \left(a^\dagger_{j,\sigma} a_{k,\sigma} + a^\dagger_{k,\sigma} a_{j,\sigma} \right) + U \sum_k n_{k,\uparrow} n_{k,\downarrow}, \tag{1}$$

where t is the tunneling parameter, and U is the interaction parameter between Fermions on the same site, $a^\dagger_{k,\sigma}$ is the raising operator for site k with spin $\sigma \in \{\uparrow, \downarrow\}$, and $n_{k,\sigma} = a^\dagger_{k,\sigma} a_{k,\sigma}$ is the number operator. The Fermionic operators are mapped to qubit operators using the Bravyi-Kitaev mapping [40]. For all calculations we will assume $U = 2t$. Energies are expressed in units of t. We assume that the lattice is a 1-D chain coupled by nearest neighbors with periodic boundary conditions. We chose an n-qubit variational ansatz $|\psi(\theta)\rangle$ consisting of an initial state $|+\rangle^{\otimes n}$ followed by six repititions of a layer of parameterized single-qubit Y-rotations followed by a layer of CZ gates between neighbouring qubits.

In Fig. 1 we plot the objective function $f(\theta_0 + s\ \phi)$ around the global minimum, where s is the parameter of the sweep, ϕ is a randomly chosen vector with $\dim(\theta_0) = \dim(\phi)$, and θ_0 are the parameters that globally minimize f. Evaluations of f are repeated three times: in the absence of any noise ("Noiseless"), including measurement error without any mitigation ("Unmitigated"), and including measurement error with CTMP-EM applied ("Mitigated"). We observe that CTMP-EM is able to significantly improve both the estimate of the ground state energy, as well as objective function values around the ground state. Nevertheless, arriving at the correct ground state depends not only on the point containing the ground state itself, but the objective function as a whole.

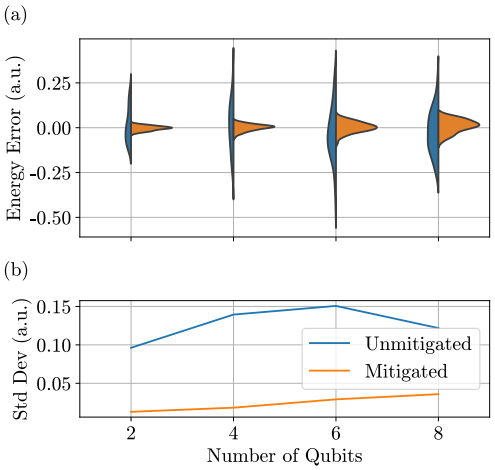

FIG. 2. Comparison of samples of the objective function $f(\theta)$ for random θ and different numbers of qubits. The energy error in panel (a) is the difference between the noisy (mitigated or unmitigated) energy and the exact result. The lefthand distributions are using the unmitigated objective function, and the righthand distributions are using the objective function mitigated with CTMP-EM. The standard deviations in panel (b) are those of the distributions in panel (a). For all system sizes considered, adding CTMP-EM significantly improves the estimate of the objective function. We use $8192n$ shots (for n qubits) to compensate for the overhead in applying CTMP-EM. The remaining source of error in the mitigated case is due to undersampling in the number of shots required to perform CTMP-EM. This emphasizes the importance of the scaling of measurements needed for CTMP-EM with the number of qubits.

B. Objective Function Sampling

In VQE the quantum computer is treated as a black box evaluation of the objective function f, hence it is critical that evaluations of f are accurate. To investigate the effectivness of measurement error mitigation for black box evaluations of $f(\theta)$ we sample points in parameter space θ and compare the noiseless, unmitigated, and mitigated cases for Fermi-Hubbard models with 1, 2, 3, and 4 sites (with 2, 4, 6, and 8 qubits respectively). We evaluate the energy with and without CTMP-EM for randomly sampled parameters of the objective function and compare the distribution of values with the noiseless result as shown in Fig. 2.

We find that the standard deviation of the error distribution is significantly reduced when mitigation is applied. Specifically, for 2, 4, 6, and 8 qubits respectively,

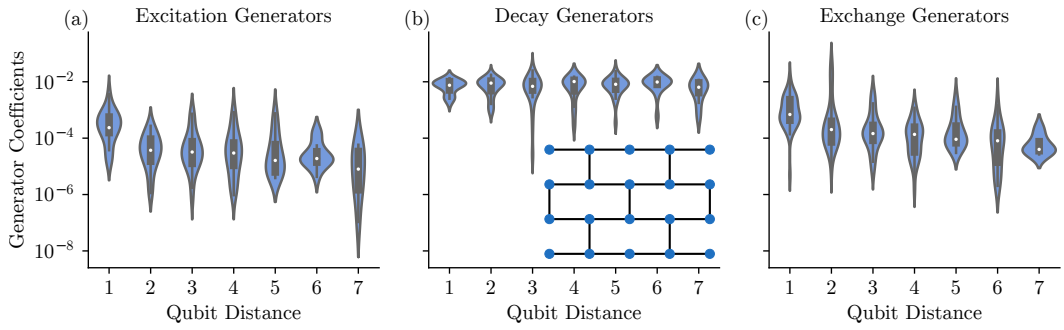

FIG. 3. Histograms and quartiles of coefficients r_i for various generators G_i. The "Excitation" (a), "Decay" (b), and "Exchange" (c) generators are associated with readout errors $00 \to 11$, $11 \to 00$, and $01 \to 10$, respectively. The "Qubit Distance" is the length of the shortest path between two qubits on *ibmq_boeblingen* device. Higher generator coefficients correspond to higher readout error rates. In all cases, the long-range correlations in readout error can be non-trivial. Cases where $r_i = 0$ due to shot limitations are elided for the purposes of plotting with logarithmic axes. In panel (b) we include the connectivity for the *ibmq_boeblingen* 20 qubit device. The vertices represent qubits, and the edges represent connections between qubits. Despite qubits being physically separated on the device, long-range correlations between qubits can occur.

the standard deviation is reduced by factors of approximately 7.46, 7.64, 5.18, and 3.40. This demonstrates that in the absence of CTMP-EM, the noisy objective function deviates significantly from its noiseless form which can greatly limit the effectivness of the VQE algorithm even for a small number of qubits.

IV. CHARACTERIZING READOUT ERRORS

The CTMP-EM method can also be used as a characterization protocol for correlated measurement errors in a quantum device, which we will demonstrate by using it to characterize correlated readout errors in several experimental devices. One method for measurement error characterization involves computing the full A matrix (which has 2^{2n} elements) as a form of measurement tomography [41], however this is not possible to do past a small number of qubits. Instead we use the set of rates $\{r_i\}$ of the CTMP-EM generator G to characterize correlated measurement errors. This is scalable in the sense that it requires as few as $n+2$ circuits, and G is parameterized by the $O(n^2)$ rates of its generator components.

A. Characterizing Correlated Errors

Calibration of G in the CTMP-EM model described in Ref. [36] is done by preparing a input set of computational basis states $\{|a_i\rangle\}$, labelled by bitstrings a_i, and performing measurements in the computational basis to estimate the assignment probabilities $P(x|a_i)$ for $x = 0, ..., 2^n - 1$. These probabilities are then processed to compute the CTMP-EM generator rates $\{r_i\}$ for each of the 1 and 2-qubit generator terms. The set of input labels a_i is *complete* if the set of all measurement outcomes contains all 1 and 2-qubit transitions for the CTMP-EM generators, if only 2-qubit correlations are present. This requires at least $n+2$ generators, though more may be used to provide a more uniform distribution across generator terms. We use the set of all bitstrings that have Hamming weight ≤ 2, of which there are $(n^2+n+2)/2$. For the purposes of characterizing correlated errors, we focus on the 2 qubit generators, as these generate correlations that cannot be captured in the single-qubit tensor product error model. We compare the distributions of the 2-qubit of rates r_i grouped by the *qubit distance*, which we define as the shortest path between two qubits in terms of the device connectivity.

We apply this technique to the 20 qubit *ibmq_boeblingen* device, which has a planar qubit connectivity graph as shown inset in Fig. 3. Here, for example, neighboring qubits have distance 1, and some pairs of qubits in the corners of the layout have distance 7. The histograms of measured 2-qubit generator rates r_i vs qubit distance for the *ibmq_boeblingen* device are shown inset in Fig. 3. Here we further group the generators into three types: excitation generators ($00 \to 11$), decay generators ($11 \to 00$), and exchange generators ($01 \leftrightarrow 10$).

A natural expectation is that the correlated errors between qubits is dependent on their connectivity, and that correlated errors will be largest on neighbouring qubits. However, our results indicate that the correlation in measurement errors between distant qubits is non-trivial, and in some cases comparable to neighboring qubits.

As one may expect, the decay generators have the highest associated generator coefficient. This is to be

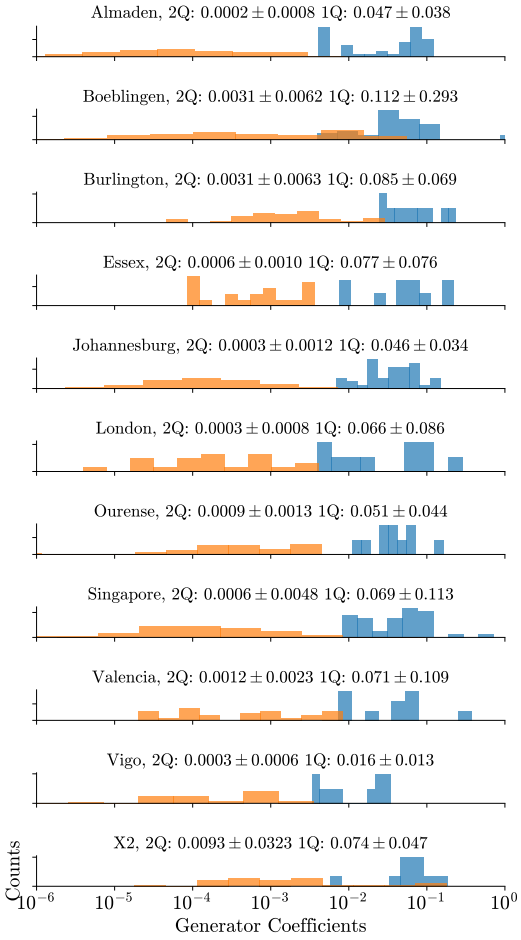

FIG. 4. Generator coefficients for 1 (blue) and 2-qubit (orange) readout errors on several IBM Quantum devices. In all cases, the rates are relatively consistent. Some error rates are calculated as $r_i = 0$ due to limitations in the number of shots. These cases are elided for the purposes of plotting with logarithmic axes.

expected as thermal relaxation to the ground state is a significant contribution to measurement errors. However, it is surprising that the distribution median does not decay appreciably with qubit distance and is non-negligible even for distantly connected qubits. The excitation and exchange generators on the other hand show reduction with qubit distance on average, however in some cases certain rates can be as large as those between neighboring qubits.

B. Comparing Devices with CTMP-EM Data

To illustrate the usefulness of CTMP-EM for characterization, we use this method to characterize readout errors on several IBM Quantum devices and compare their local 1-qubit and correlated 2-qubit generator coefficients. The distribution in generator values is shown in Fig. 4. For all measured devices the 1-qubit error rates are generally higher than the 2-qubit error rates as expected, and all devices give relatively consistent error rates. The devices shown range from 5-qubit to 20-qubits with the calibrations run using the minimum number of $n+2$ calibration circuits, using the maximum number of shots available for each device. Moreover, CTMP-EM as a characterization technique does not depend on device connectivity, and includes information about long-range correlations. This is advantageous since we saw before that the long-range correlations in readout errors can be non-trivial.

V. CONCLUSION

In this paper, we have shown that CTMP-EM is vital for improving the performance of VQAs on near-term devices, and that it can fundamentally change the shape of the objective function computed by the quantum device. Moreover, we demonstrate that CTMP-EM can efficiently be used to characterize long-range correlations in readout error on near-term devices in terms of the device connectivity, and that these long-range correlations are present on current devices. Nevertheless, an interesting topic for future work would be to expand the CTMP model to generators that act on more than two qubits, or analyze how the calibration parameters drift over time. Additionally, it would be interesting to investigate the impact on performance for QAOA Hamiltonians. For now, our work emphasizes the importance of using CTMP-EM for characterization, because it only requires computing the terms of G, of which there are only $O(n^2)$ many, instead of $A = e^G$, which is dense. Moreover, the calibration technique is efficient in the number of circuits. We believe that our results will be useful to understanding the objective function in VQAs on near-term devices, as well as characterizing near-term devices in terms of readout error.

ACKNOWLEDGEMENTS

The authors thank Sergey Bravyi, Sophia Economou, and Sarah Sheldon for helpful discussions. We also thank the IBM Quantum team for providing access to the devices used in this work. This work was done during G.S.B internship at IBM Quantum during the summer of 2020. G.S.B. thanks the IBM Quantum team for a very enriching internship experience.

[1] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta, "Validating quantum computers using randomized model circuits," Phys. Rev. A **100**, 032328 (2019).

[2] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and et al., "Demonstration of the qccd trapped-ion quantum computer architecture," arXiv e-prints , 2003.01293 (2020).

[3] Peter J Karalekas, Nikolas A Tezak, Eric C Peterson, Colm A Ryan, Marcus P da Silva, and Robert S Smith, "A quantum-classical cloud platform optimized for variational hybrid algorithms," Quantum Science and Technology **5**, 024003 (2020).

[4] Petar Jurcevic, Ali Javadi-Abhari, Lev S. Bishop, Isaac Lauer, Daniela F. Bogorin, Markus Brink, Lauren Capelluto, Oktay Günlük, Toshinaro Itoko, Naoki Kanazawa, Abhinav Kandala, George A. Keefe, Kevin Kruslich, William Landers, Eric P. Lewandowski, Douglas T. McClure, Giacomo Nannicini, Adinath Narasgond, Hasan M. Nayfeh, Emily Pritchett, Mary Beth Rothwell, Srikanth Srinivasan, Neereja Sundaresan, Cindy Wang, Ken X. Wei, Christopher J. Wood, Jeng-Bang Yau, Eric J. Zhang, Oliver E. Dial, Jerry M. Chow, and Jay M. Gambetta, "Demonstration of quantum volume 64 on a superconducting quantum computing system," (2020), arXiv:2008.08571 [quant-ph].

[5] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, "A quantum approximate optimization algorithm," arXiv preprint arXiv:1411.4028 (2014).

[6] 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," Nature Commun. **5**, 4213 (2014).

[7] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, "Scalable quantum simulation of molecular energies," Phys. Rev. X **6**, 031007 (2016).

[8] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," New J. Phys **18**, 023023 (2016).

[9] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," Nature **549**, 242–246 (2017).

[10] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, "Computation of molecular spectra on a quantum processor with an error-resilient algorithm," Phys. Rev. X **8**, 011021 (2018).

[11] G Pagano, A Bapat, P Becker, KS Collins, A De, PW Hess, HB Kaplan, A Kyprianidis, WL Tan, C Baldwin, *et al.*, "Quantum approximate optimization with a trapped-ion quantum simulator," arXiv preprint arXiv:1906.02700 (2019).

[12] Jonathan Romero and Alan Aspuru-Guzik, "Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions," arXiv preprint arXiv:1901.00848 (2019).

[13] Pauline J Ollitrault, Abhinav Kandala, Chun-Fu Chen, Panagiotis Kl Barkoutsos, Antonio Mezzacapo, Marco Pistoia, Sarah Sheldon, Stefan Woerner, Jay Gambetta, and Ivano Tavernelli, "Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor," arXiv preprint arXiv:1910.12890 (2019), arXiv:1910.12890 [quant-ph].

[14] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Benjamin Chiaro, Roberto Collins, William Courtney, Sean Demura, Andrew Dunsworth, Daniel Eppens, Edward Farhi, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Steve Habegger, Matthew P. Harrigan, Alan Ho, Sabrina Hong, Trent Huang, William J. Huggins, Lev Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Erik Lucero, Orion Martin, John M. Martinis, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Hartmut Neven, Murphy Yuezhen Niu, Thomas E. O'Brien, Eric Ostby, Andre Petukhov, Harald Putterman, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Doug Strain, Kevin J. Sung, Marco Szalay, Tyler Y. Takeshita, Amit Vainsencher, Theodore White, Nathan Wiebe, Z. Jamie Yao, Ping Yeh, and Adam Zalcman, "Hartree-Fock on a superconducting qubit quantum computer," arXiv e-prints , arXiv:2004.04174 (2020), arXiv:2004.04174 [quant-ph].

[15] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. Schuyler Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, Colm A. Ryan, D. Scarabelli, M. Scheer, E. A. Sete, P. Sivarajah, Robert S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, Blake R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, "Unsupervised Machine Learning on a Hybrid Quantum Computer," arXiv e-prints , arXiv:1712.05771 (2017), arXiv:1712.05771 [quant-ph].

[16] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas, "From the quantum approximate optimization algorithm to a quantum alternating operator ansatz," Algorithms **12**, 34 (2019).

[17] Ryan Babbush, Nathan Wiebe, Jarrod McClean, James McClain, Hartmut Neven, and Garnet Kin-Lic Chan, "Low-depth quantum simulation of materials," Physical Review X **8**, 011044 (2018).

[18] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F. Izmaylov, "Measurement optimization in the variational quantum eigensolver using a minimum clique cover," The Journal of Chemical Physics **152**, 124114 (2020), https://doi.org/10.1063/1.5141458.

[19] William J. Huggins, Jarrod McClean, Nicholas Rubin, Zhang Jiang, Nathan Wiebe, K. Birgitta Whaley, and Ryan Babbush, "Efficient and Noise Resilient Measurements for Quantum Chemistry on Near-Term Quantum Computers," arXiv:1907.13117 [physics, physics:quant-ph] (2019), arXiv:1907.13117 [physics, physics:quant-ph].

[20] Andrew Zhao, Andrew Tranter, William M. Kirby, Shu Fay Ung, Akimasa Miyake, and Peter J. Love, "Measurement reduction in variational quantum algorithms," Phys. Rev. A **101**, 062322 (2020).

[21] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall, "An adaptive variational algorithm for exact molecular simulations on a quantum computer," Nat. Commun. **10**, 3007 (2019).

[22] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, Stefan Filipp, and Ivano Tavernelli, "Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wavefunction expansions," Phys. Rev. A **98**, 022322 (2018).

[23] Ho Lun Tang, V. O. Shkolnikov, George S. Barron, Harper R. Grimsley, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou, "qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansatze on a quantum processor," (2019), arXiv:1911.10205 [quant-ph].

[24] Bryan T Gard, Linghua Zhu, George S Barron, Nicholas J Mayhall, Sophia E Economou, and Edwin Barnes, "Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm," arXiv preprint arXiv:1904.10910 (2019).

[25] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, "Training of quantum circuits on a hybrid quantum computer," Science Advances **5** (2019), 10.1126/sciadv.aaw9918, https://advances.sciencemag.org/content/5/10/eaaw9918.full.pdf

[26] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, "Cloud quantum computing of an atomic nucleus," Phys. Rev. Lett. **120**, 210501 (2018).

[27] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage, "Quantum-classical computation of schwinger model dynamics using quantum computers," Phys. Rev. A **98**, 032331 (2018).

[28] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei, "Detector tomography on ibm quantum computers and mitigation of an imperfect measurement," Phys. Rev. A **100**, 052315 (2019).

[29] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta, "Error mitigation for short-depth quantum circuits," Phys. Rev. Lett. **119**, 180509 (2017).

[30] Abhinav Kandala, Kristan Temme, Antonio D Córcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta, "Error mitigation extends the computational reach of a noisy quantum processor," Nature **567**, 491–495 (2019).

[31] Kathleen E. Hamilton and Raphael C. Pooser, "Error-mitigated data-driven circuit learning on noisy quantum hardware," Quantum Machine Intelligence **2**, 1–15 (2020).

[32] Michael R. Geller and Mingyu Sun, "Efficient correction of multiqubit measurement errors," arXiv preprint arXiv:2001.09980 (2020), arXiv:2001.09980 [quant-ph].

[33] Kathleen E. Hamilton, Tyler Kharazi, Titus Morris, Alexander J. McCaskey, Ryan S. Bennink, and Raphael C. Pooser, "Scalable quantum processor noise characterization," arXiv preprint arXiv:2006.01805 (2020), arXiv:2006.01805 [quant-ph].

[34] Michael R Geller, "Rigorous measurement error correction," Quantum Science and Technology **5**, 03LT01 (2020).

[35] https://github.com/qiskit/qiskit, "Qiskit: An open-source framework for quantum computing," (2019).

[36] Sergey Bravyi, Sarah Sheldon, Abhinav Kandala, David C. Mckay, and Jay M. Gambetta, "Mitigating measurement errors in multi-qubit experiments," arXiv preprint arXiv:2006.14044 (2020), arXiv:2006.14044 [quant-ph].

[37] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran, "Evaluating analytic gradients on quantum hardware," Phys. Rev. A **99**, 032331 (2019).

[38] Kunal Sharma, Sumeet Khatri, M Cerezo, and Patrick J Coles, "Noise resilience of variational quantum compiling," New Journal of Physics **22**, 043006 (2020).

[39] J. Hubbard and Brian Hilton Flowers, "Electron correlations in narrow energy bands," Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences **276**, 238–257 (1963), https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1963.02

[40] Sergey B. Bravyi and Alexei Yu. Kitaev, "Fermionic quantum computation," Annals of Physics **298**, 210 – 226 (2002).

[41] R. C. Bialczak, M. Ansmann, M. Hofheinz, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, H. Wang, J. Wenner, M. Steffen, and et al., "Quantum process tomography of a universal entangling gate implemented with josephson phase qubits," Nature Physics **6**, 409–413 (2010).

Chapter 7

Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm

This chapter includes pages from our publication [6]

Material from: B. T. Gard, L. H. 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 Information **6**, 10.1038/s41534-019-0240-1 (2020)

Publisher: Springer Nature.

The contributions to this work are as follows: My contributions were assisting in development of theoretical results, assisting in numerical simulations, assisting in preparation of figures, and assisting in manuscript preparation. B.T.G. primarily carried out this work. L.Z. assisted with theoretical developments, numerical simulations, and manuscript preparation. N.J.M., S.E.E. and E.B. supervised.

npj | Quantum Information

www.nature.com/npjqi

ARTICLE OPEN

Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm

Bryan T. Gard[1*], Linghua Zhu[1], George S. Barron[1], Nicholas J. Mayhall[2], Sophia E. Economou[1] and Edwin Barnes[1]

The variational quantum eigensolver is one of the most promising approaches for performing chemistry simulations using noisy intermediate-scale quantum (NISQ) processors. The efficiency of this algorithm depends crucially on the ability to prepare multi-qubit trial states on the quantum processor that either include, or at least closely approximate, the actual energy eigenstates of the problem being simulated while avoiding states that have little overlap with them. Symmetries play a central role in determining the best trial states. Here, we present efficient state preparation circuits that respect particle number, total spin, spin projection, and time-reversal symmetries. These circuits contain the minimal number of variational parameters needed to fully span the appropriate symmetry subspace dictated by the chemistry problem while avoiding all irrelevant sectors of Hilbert space. We show how to construct these circuits for arbitrary numbers of orbitals, electrons, and spin quantum numbers, and we provide explicit decompositions and gate counts in terms of standard gate sets in each case. We test our circuits in quantum simulations of the H_2 and *LiH* molecules and find that they outperform standard state preparation methods in terms of both accuracy and circuit depth.

npj Quantum Information (2020)6:10 ; https://doi.org/10.1038/s41534-019-0240-1

INTRODUCTION

Quantum simulation of fermionic systems, such as molecules, is one of the first envisioned applications of quantum computers, as famously proposed by Feynman.[1] The first protocol introduced for quantum chemistry simulations is based on the so-called phase estimation algorithm.[2,3] This algorithm however requires a large number of quantum gates, leading to long quantum circuits that are challenging for existing and near-term noisy intermediate scale quantum (NISQ) devices.[4,5] For such devices, alternative, hybrid algorithms are instead envisioned. In such algorithms, the work is shared between a quantum processor and a classical computer. In particular, the variational quantum eigensolver (VQE), first introduced and demonstrated experimentally by Peruzzo et al.,[6] has become the prevailing algorithm for chemistry simulations with NISQ devices, with several milestone papers demonstrating the calculation of molecular energies and wavefunctions.[7–12]

The VQE algorithm relies on preparing and measuring multi-qubit states based on a variational ansatz, and using the classical computer to optimize and update the variational parameters in this ansatz. Some of the advantages of VQE are that its variational character can provide some degree of error mitigation in the gates,[7,13–15] and that it features shallower circuits compared to the phase estimation algorithm. The form of the ansatz is a crucial ingredient of VQE and one that can determine its success on NISQ devices. There are two main approaches in determining the ansatz. One approach is based on a technique from chemistry, the unitary coupled cluster method,[16–18] translated into quantum gates by Trotterization.[18] This approach tends to lead to deeper circuits than what is currently feasible on hardware, and is generally not exact. To address these issues, a new, iterative algorithm termed ADAPT-VQE, was recently put forward and was shown to enable a much more compact ansatz while simultaneously exhibiting higher accuracy.[19] An alternative approach is to base the ansatz on the capabilities of the hardware and prepare states by combining parameterized gates available on the processor.[8,20] Such an ansatz has the advantage of compatibility with the capabilities of the hardware, and as such is NISQ-friendly. On the other hand, in its simplest form it is an ad hoc ansatz that can cause the algorithm to get stuck on 'barren plateaus'[21] as the number of qubits increases and the Hilbert space correspondingly grows exponentially. Therefore, for hardware-based ansätze to be a viable approach for problems of interest, they must be selected in a way that guarantees they span the part of the Hilbert space where the solution lives, while avoiding generating unphysical states.

Two ways to guarantee that the desired part of the Hilbert space is accessed include adding terms in the VQE energy function that penalize symmetry violations[22,23] or carefully designing state preparation circuits so that they only produce states with the appropriate symmetries regardless of how their variational parameters are chosen. An early step toward the latter direction was taken by Wang et al.,[24] who focused on the preparation of states with a well-defined number of occupied spin orbitals and showed that the number of CNOT gates required for this scales polynomially with the number of qubits in the limit where the number of electrons is much smaller than the number of qubits. More recently, Barkoutsos et al.[15] enforced particle number conservation by using the particle-hole representation in conjunction with a parametrized particle-conserving exchange-type gate,[25–28] which we also make use of here. However, important open questions remain, including how other symmetries can also be built into the circuits and whether more efficient circuits containing the minimal number of parameters necessary to span the symmetry subspace exist.

In this paper, we address these questions by introducing state preparation circuits that provide a systematic, economical way to generate states with well-defined symmetries, including particle number, total spin, spin projection, and time-reversal. Our circuits incorporate the minimal number of parameters needed to fully span the appropriate symmetry subspace while avoiding all states outside this subspace. This general approach has two key

[1]Department of Physics, Virginia Tech, Blacksburg, VA 24061, USA. [2]Department of Chemistry, Virginia Tech, Blacksburg, VA 24061, USA. *email: bgard1@vt.edu

advantages: the first is that the true ground state is guaranteed to be contained within the space of states spanned by the circuit, and the second is that resources are not spent on generating irrelevant parts of the Hilbert space, reducing the complexity of the classical optimization step of the VQE algorithm. Eliminating extraneous parameters can dramatically speed up the optimization process and suppress the probability of getting trapped in local extrema or barren plateaus. We present circuits for arbitrary numbers of single-particle orbitals and electrons and for arbitrary spin quantum numbers. Our circuits are constructed with hardware constraint considerations, including a reduced number of CNOT gates that need only be applied between adjacent qubits in a linear array, making our work particularly suited for NISQ devices. In addition, since our construction conserves number and spin symmetries, symmetry verification techniques can be used to mitigate any errors which violate these symmetries.[13,14,27] Our most general circuits, which conserve particle number, time-reversal, and spin symmetries, are guaranteed to span the appropriate symmetry subspace by construction. We also present circuits that conserve particle number and time-reversal symmetries, but not total spin; the symmetry-preserving properties of these circuits are verified with numerical calculations. We test the performance of our circuits against standard state preparation ansätze by running VQE simulations of the H_2 and *LiH* molecules. We find that our circuits outperform the standard methods in terms of both accuracy and circuit depth.

The paper is organized as follows. In the Results section, we define the basic gate set we use to manipulate particle number on the quantum processor and show how these gates can be systematically assembled into circuits that preserve particle number symmetry. We also discuss how to respect time-reversal symmetry as well. In addition, in the Results section, we show how to create circuits that respect total spin and spin projection symmetries in addition to particle number. We give some concluding remarks in the Discussion section. An appendix contains additional details about gate decompositions.

RESULTS

Particle number and time-reversal symmetries

In this work, we focus on mapping chemistry problems onto quantum processors using the Jordan-Wigner mapping,[3,29] in which each qubit in the quantum processor corresponds to a particular spin-orbital, and the qubit states $|0\rangle$ and $|1\rangle$ encode the occupation of that spin-orbital. Any multi-electron state involving n spin-orbitals on the chemistry side can be mapped to a corresponding state of n qubits on the quantum processor. In this mapping, fixing the total number of electrons is tantamount to fixing the total number of qubits that are in the excited state $|1\rangle$. Thus, the Jordan-Wigner mapping relates fixed-particle-number subspaces to fixed-excitation subspaces in the qubit Hilbert space.

Formally, we can define the qubit subspace corresponding to m electrons occupying n spin-orbitals as

$$H_{n,m} = \text{span}\left\{|s_1, s_2, ..., s_n\rangle | s_i \in \{0,1\}, \sum_{i=1}^{n} s_i = m\right\}. \quad (1)$$

This is the subspace of multi-qubit states containing m qubits in the $|1\rangle$ state and $n-m$ qubits in the $|0\rangle$ state, so that $\dim(H_{n,m}) = \binom{n}{m}$. A general state in this subspace can be represented as an arbitrary superposition of these basis states with complex coefficients and is thus characterized by $2\dim(H_{n,m}) - 2$ real parameters, where we have removed two parameters by fixing the normalization and neglecting a global phase. In the absence of any other symmetries, $2\dim(H_{n,m}) - 2$ is the minimal number of real variational parameters needed to prepare arbitrary trial states describing the correct number of electrons.

An additional symmetry that often arises in chemistry problems is time-reversal symmetry. This symmetry is typically present, for example, when one wants to solve the stationary Schrödinger equation in the absence of any applied magnetic field. In this case, one can always choose the energy eigenstates to be strictly real functions. Under the Jordan-Wigner mapping, this means that the coefficients appearing in the multi-qubit superposition states we prepare as trial states for the VQE algorithm should be restricted to real values. This will reduce the dimensionality of the target symmetry subspace by a factor of two down to $\dim(H_{n,m}) - 1$. Imposing this restriction on the trial states will prevent the classical optimizer from wasting time exploring a large portion of Hilbert space that does not contain any of the desired energy eigenstates. When we introduce our state preparation circuits below, we will see that time-reversal symmetry can be imposed easily after other symmetries are already built into the circuits, essentially just by fixing half of the variational parameters in the circuits in such a way that the resulting states are strictly real.

Before we introduce our general scheme for constructing particle-number-conserving state preparation circuits, we first present a few simple examples that may provide some intuition about the general structure of such circuits. First note that the cases $m = 0$ and $m = n$ are trivial since in each case, there is only a single state spanning the subspace. Therefore, we restrict our attention to $0 < m < n$ throughout this work. The simplest nontrivial example is the case of $m = 1$ electron in $n = 2$ orbitals. (Since we are presently only concerned with particle-number symmetry, these could be spin-orbitals or spatial orbitals in the case of spinless fermions. Spin symmetries will be incorporated in the next section). A circuit that spans the corresponding subspace $H_{2,1}$ is shown in Fig. 1.

This circuit requires only two parameters to span the single-excitation (one-electron) subspace, which is comprised of states of the form $\alpha|01\rangle + \beta|10\rangle$. Here, although α and β are complex, they contribute only two real parameters after we impose normalization $|\alpha|^2 + |\beta|^2 = 1$ and discard a global phase. The key ingredient in this circuit is a two-qubit entangling gate that we have denoted as $A(\theta, \phi)$. In the basis $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$, it is defined as[15]

$$A(\theta,\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & e^{i\phi}\sin\theta & 0 \\ 0 & e^{-i\phi}\sin\theta & -\cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2)$$

It is clear from the form of this exchange-type gate that it preserves particle number since it mixes $|01\rangle$ and $|10\rangle$ but does nothing to the $|00\rangle, |11\rangle$ subspace. The initial X gate on the first qubit in Fig. 1 brings the two-qubit state into the one-excitation subspace, while the subsequent A gate generates all possible superpositions within this subspace, as can be seen by inspection in this case. If we wish to also impose time-reversal symmetry, then it suffices to set $\phi = 0$ in each of the A gates. This removes the phase from each coefficient in the resulting superposition state without restricting the magnitude, thus ensuring that the resulting state is an arbitrary real state. The A gate plays a central role in our state preparation circuits. It can be decomposed into a sequence of two single-qubit gates and three CNOT gates, as

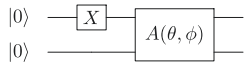

Fig. 1 2-Qubit *A* circuit. A simple example of a 2 qubit circuit which exactly spans the subspace defined by 1 excitation, $\alpha|01\rangle + \beta|10\rangle$, with two parameters (θ, ϕ). This circuit saturates the lower bound on the number of real parameters required to construct arbitrary states in $H_{2,1}$.

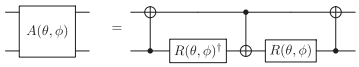

Fig. 2 *A* gate decomposition. Decomposition of the A gate in terms of elementary single and two-qubit gates. $R(\theta,\phi) = R_z(\phi+\pi)\,R_y(\theta+\pi/2)$, where $R_z(\theta) = \exp(-i\theta\sigma_z/2)$, $R_y(\phi) = \exp(-i\phi\sigma_y/2)$.

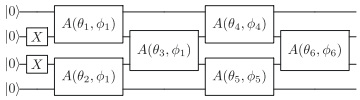

Fig. 3 4-Qubit *A* circuit. An example circuit for the case of $n = 4$, $m = 2$ which exactly spans the subspace defined by six basis states using the minimal number (10) of parameters. Note that ϕ_1 is used three times.

shown in Fig. 2. This decomposition is minimal in the number of CNOT gates. In operator form, the A gate is

$$A = e^{-\frac{i}{2}(\frac{\pi}{2}-\phi)Z_2} e^{-\frac{i}{2}(\theta X_1 X_2 + \theta Y_1 Y_2 + \frac{\pi}{2} Z_1 Z_2)} e^{-\frac{i}{2}(\frac{\pi}{2} Z_1 - \phi Z_2)},$$

up to an irrelevant global phase.[30]

Building upon the example of Fig. 1, we find that we can generate trial states corresponding to other particle and orbital numbers using a similar construction.

Figure 3 shows another example circuit for the case of $m = 2$ electrons in $n = 4$ orbitals. In this case, we start with two X gates to bring the state into the two-excitation subspace, and we then apply a series of A gates between neighboring qubits to create arbitrary superpositions within this subspace. In this example, an arbitrary complex two-electron state is specified by $2\binom{4}{2} - 2 = 10$ real parameters. Since each A gate introduces two parameters, one might expect that only five A gates would be needed to generate an arbitrary state. However, we find that at least six A gates are needed to do this, although two of the parameters can be fixed to reduce the total parameter count back down to 10. We can again impose time-reversal symmetry by setting all the ϕ_i parameters to zero in the A gates, which reduces the parameter count down to 6—one more than the minimal number of 5. To remove this extra parameter, we find numerically that it works to set either θ_4 or θ_5 to zero. Note that setting both θ_i and ϕ_i to zero does not remove the ith A gate completely but instead reduces it to a CZ gate, as is evident from Eq. (2).

To confirm that this and all other circuits presented in this section indeed span the target subspace, we compute the fidelity $F = \frac{1}{N}\sum_{i=1}^{N} |\langle\Psi_i|\psi_i\rangle|^2$, where $|\Psi_i\rangle$ is a random state within the chosen subspace, and $|\psi_i\rangle$ is the state output by our circuit after we maximize $|\langle\Psi_i|\psi_i\rangle|^2$ with respect to the variational parameters. We check numerically that $F = 1$ can be achieved using the minimal number of parameters, $2\dim(\mathrm{H_{n,m}}) - 2$. We choose enough random states, $N \gg \dim(H_{n,m})$, to ensure that the subspace is adequately represented. This is how we determine that, in the case of Fig. 3 with time-reversal symmetry, it works to set θ_4 or θ_5 to zero, but setting other θ parameters to zero instead does not achieve unit fidelity (although the fidelity still remains very high). Further details about the numerical verification of our circuits are given in the supplementary information.[31]

Although we have found a circuit that prepares all states in the two-excitation subspace using the minimal number of variational parameters (Fig. 3), this solution is neither unique nor optimal in terms of the number of CNOT gates. To illustrate these points, we present another circuit that accomplishes the same task in Fig. 4. In addition to the A gate, this circuit also makes use of single-qubit gates beyond just X gates. This example is quite different from the one shown in Fig. 3 in that it does not first apply X gates to two

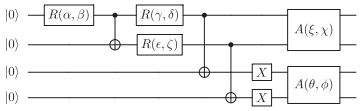

Fig. 4 Alternate 4-qubit circuit. Another example circuit for the case of $n = 4, m = 2$, which also spans the desired subspace with the minimal 10 parameters, but with fewer two-qubit gates than the circuit shown in Fig. 3. The single-qubit gates $R(\theta,\phi)$ are as defined in Fig. 2.

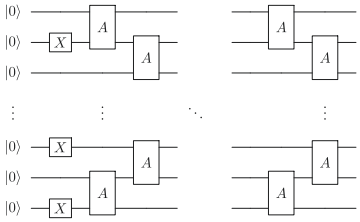

Fig. 5 General *A* circuit. General construction of an efficient circuit which also enforces number symmetry for any number of qubits (orbitals) n and excitations (electrons) m and is constructed using the logic discussed in the text. Each A gate contributes two variational parameters θ_i, ϕ_i to the ansatz, except the last two A gates, which each contribute one (see text). This general structure only requires single-qubit X gates and a cascade of two-qubit A gates and always generates circuits with the minimal number of required parameters $2\binom{n}{m} - 2$.

qubits in order to bring the quantum processor into the appropriate particle-number subspace. Instead, the appropriate subspace is approached gradually as the circuit is performed, making it more challenging to understand and generalize the circuit. Since each A requires three CNOT gates to implement, we see that the circuit in Fig. 4 requires only 9 CNOT gates, while the one in Fig. 3 requires twice as many. However, it should be noted that the circuit of Fig. 3 requires only nearest-neighbor qubit coupling, which is not true of the one in Fig. 4. This example highlights the fact that further reductions in circuit depth are possible even if the circuit contains the minimal number of parameters, although this may require an increase in the qubit connectivity.

An efficient circuit for the general case of n orbitals and m fermions is shown in Fig. 5. This circuit builds on the approach of Fig. 3, where X gates are applied to m qubits to bring the system into the correct particle-number subspace, and then a sequence of A gates is performed to create different superpositions. Through trial and error, we find that the total number of A gates needed to produce all possible superpositions is $\binom{n}{m}$. Since each A gate contributes two parameters, we have a total of $2\binom{n}{m}$ parameters; the ϕ parameters in the last two A gates can be fixed to reduce the number of parameters down to the minimal number, $2\binom{n}{m} - 2$, while still fully spanning the subspace. Our general recipe for constructing circuits for any n, m like that shown in Fig. 5 can be summarized into the following steps:

1. Apply X gates to m qubits. For an efficient circuit, avoid placing X gates on neighboring qubits.
2. Apply a "first layer" of A gates on all adjacent pairs of qubits on which either $X \otimes 1$ or $1 \otimes X$ has been applied.
3. Apply a "second layer" of A gates on adjacent pairs of qubits, where each pair includes one qubit acted on by an A gate from the previous step and a qubit free of A gates. Continue

to place A gates on adjacent qubits as necessary until all neighboring qubits are connected with A gates. The first and second layers define a primitive pattern.

4. Repeat the primitive pattern until $\binom{n}{m}A$ gates are placed. Any two A gates have ϕ as a free parameter and therefore the full circuit contains exactly $2\binom{n}{m}-2$ parameters, the minimum required to span the subspace $H_{n,m}$.

In Step 1 above, we can see that this step simply places the full system into the proper m excitation subspace. The following two steps, which only involve applications of A gates, do not change the excitation number of the system but instead mix the $|01\rangle$ and $|10\rangle$ basis states of the two qubits on which they act, producing the desired subspace spanned by m, dependent on the parameters $(\vec{\theta},\vec{\phi})$. Although we do not have an analytical argument for why this particular arrangement of A gates works, we have confirmed this through extensive numerical testing. We conjecture that this circuit pattern will continue to work for arbitrarily many fermions and orbitals. Another key point is contained in Step 4, specifically that this general construction always uses the minimal number of required parameters to span the desired space; increasing the gate depth any further is unnecessary, as unit fidelity is achievable at this gate depth. Time-reversal symmetry can be imposed by setting all the ϕ parameters and one θ parameter in the A gates to zero as before. This yields a purely real ansatz with the minimal number of parameters needed. Numerical calculations of the fidelity as a function of the number of variational parameters for several different values of n and m are given in Fig. S4 of the supplement.[31] We have numerically verified all cases for which (n) and all permutations of $n=\{2,3,4,5,6\}, m=\{1,2,3,4,5\}$. In all cases, the fidelity increases monotonically with the number of parameters and saturates at unity once this number equals the dimension of the symmetry subspace, confirming that the circuits fully span the subspace.

While this construction is not necessarily the most resource efficient in terms of the required number of gates, it is straightforward to extend to any desired subspace defined by n and m and only requires nearest-neighbor qubit coupling, which is typically more straightforward to engineer. This should be contrasted with the circuit of Fig. 4, where two-qubit gates between non-neighboring qubits are required. We also note that our general construction naturally exhibits the symmetries of binomial functions. For example, the number of gates used for the case of $\binom{n}{p}$ and $\binom{n}{n-p}$ for $0<p<n$ are identical, which is a reflection of particle-hole symmetry. This in turn means that we can focus on the case where $m \leq n/2$ without loss of generality, so that in Step 1 above, we can always avoid applying X gates on two adjacent qubits.

It is worth comparing our general construction with existing state preparation algorithms in terms of gate counts. If we wished to span the full Hilbert space, then this would require $\mathcal{O}(2^n)$ CNOT gates.[32,33] Some previous state preparation algorithms involve transforming one arbitrary $n-$qubit state into another arbitrary state, which requires $2^{n+1}-2n-2$ CNOT gates.[34] However, since we are interested in only spanning a subspace of the full Hilbert space, namely $H_{n,m}$, we can span this subspace with significantly fewer CNOT gates. Wang et al.[24] and Ortiz et al.[35] also considered this restricted subspace, requiring no more than $2^{m+1}n^m/m!$ and $\binom{n}{m}^2 n^2$ CNOT gates, respectively. In the case of our general construction for n qubits and m excitations, we find that our algorithm requires at most $N_{\mathrm{CNOT}}(n,m)=3\binom{n}{m}$ CNOT gates. Since we always consider a fixed input state (the state with all qubits in $|0\rangle$), simplifications of our required gates are always possible, which reduces the number of required CNOT gates. Specifically, if we eliminate the unnecessary CNOT gates, then the actual number of CNOT gates in the general circuit is

$$N_{\mathrm{CNOT}}(n,m)=\begin{cases}3\binom{n}{m}-3m+1 & 0/2\\ 3\binom{n}{m}-2m-2 & m=n/2\\ 3\binom{n}{m}-3n+3m+1 & n/2.\end{cases} \quad (3)$$

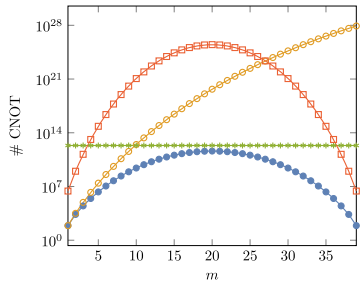

Fig. 6 CNOT scaling. Number of required CNOT gates as a function of particle number m for a fixed number of qubits $n=40$. Our general state preparation circuits that respect particle-number symmetry are shown in blue, while the results of Wang et al.[24] are shown in yellow, Bergholm et al.[34] in green and Ortiz et al.[35] in red.

Figure 6 shows how our approach to constructing state preparation circuits compares to existing works. We see that our scheme significantly decreases the required number of CNOT gates.

We noted earlier that our example of Fig. 4 contains only 9 CNOT gates, while our general construction uses $N_{\mathrm{CNOT}}(4,2)=18$ CNOT gates in this case. However, our general construction in terms of A gates is relatively straightforward for arbitrary $H_{n,m}$, while it is not clear how to generalize the construction of Fig. 4. We stress that on a case by case basis, it may be possible to find circuits that more efficiently span a reduced Hilbert space, but finding a general procedure for constructing such circuits for arbitrary numbers of orbitals and electrons can be challenging, and it may require a more complicated qubit connectivity beyond just nearest-neighbor coupling. Further investigation of these differences and tradeoffs is an interesting topic for future work but beyond the scope of this paper.

Spin symmetries

Many systems that we are interested in simulating possess not only particle-number and time-reversal symmetries, but also spin symmetries. This can include both the net spin magnetization s_z and also total spin s. To our knowledge, it remains an open problem to find state preparation circuits that respect all these symmetries at the same time. Here, we introduce a protocol for constructing circuits that achieve this.

We begin by first showing how our general scheme for conserving particle number described in the previous section can be extended to conserve s_z as well with only minor modifications. The first step is to choose our fermion-qubit mapping such that the first $n/2$ qubits represent spin-up orbitals while the remaining $n/2$ represent spin-down orbitals. We follow the same general steps of forming a cascade of A gates as before, but now with the added constraint that the parameters of any A gate that bridges the two spin subspaces (those that entangle the $n/2$ qubit with the $n/2+1$ qubit) are set to zero ($\theta_i=\phi_i=0$). This prevents any mixing of the two spin subspaces. Therefore, if we start with the proper number of spin-up and spin-down orbitals occupied, then this sequence of gates guarantees that the final state also has the correct spin occupation numbers. We also require that these parameter-free A gates do not appear in the

first layer. This can effectively swap the first and second layers as outlined in the previous general protocol. For example, in the case of Fig. 3 and with our spin assignment, we only need to set the four parameters contained within the A gates that act between qubits $2,3$ to zero along with $\phi_5 = \phi_4$. This assignment then generates all states with $n = 4, m = 2, s_z = 0$ and with the minimal number of parameters (6) for this space.

We now move on to the case where total spin s is also a good quantum number. If we impose conservation of s, s_z and particle number, then the size of the relevant symmetry subspace for a chosen n, m, s, s_z is given by

$$dim(H_{n,m,s,s_z}) = \sum_{k=0}^{m/2-s} \binom{n/2}{k}\binom{n/2-k}{m-2k} \times \frac{(2s+1)(m-2k)!}{(m/2-k-s)!(m/2-k+s+1)!}. \quad (4)$$

Here, we maintain the notation from the previous section where n is the number of qubits and m is the number of particles, but now we focus on the case in which each qubit encodes the occupancy of a particular spin-orbital. (In the previous section each qubit could correspond to either a spin-orbital or a spatial orbital in the case of spinless fermions.) Therefore, the n qubits encode $n/2$ spatial orbitals, and we omit the case of odd values for n. The first binomial coefficient in Eq. (4) counts the number of ways to assign doubly occupied spatial orbitals in the system, the second counts the number of ways to assign singly occupied orbitals, and the last factor is the number of s irreducible representations in a tensor product of m spin-$\frac{1}{2}$ particles. We show in Fig. 7 that exploiting all the symmetries may significantly reduce the dimension of the Hilbert space. For example, already for $n = 28$ spin-orbitals, the relevant subspace when all symmetries are imposed is at least two orders of magnitude smaller than the full Hilbert space of 28 qubits. Taking advantage of this reduction can significantly reduce the demands on the quantum processor and improve the speed and accuracy of the classical optimization step of the VQE algorithm. We note that, while this reduction is significant, this scaling always remains exponential even when symmetries are imposed. Nevertheless, use of symmetries can simplify ansätze and thus reduce the required CNOT and parameter counts.

The problem of constructing circuits that also conserve total spin s is much more challenging compared to the symmetries we have discussed thus far. In the remainder of this section, we address this problem by developing a completely different approach to obtaining symmetry-preserving circuits. In addition to allowing for conservation of s, this approach also provides an alternative method for building circuits that respect particle-number symmetry. In this new approach, we adopt the convention that the qubits alternate between spin-up and spin-down orbitals such that the first, third, fifth, etc. qubits encode spin up, while the second, fourth, sixth, etc. encode spin down. This choice of spin encoding leads to an efficient scaling in terms of number of gates, as we will see.

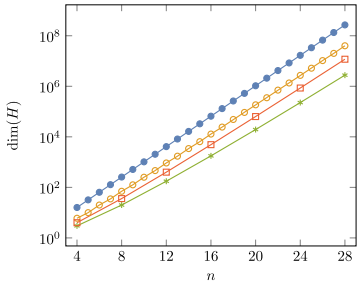

Fig. 7 Hilbert space scaling. Hilbert space dimension as a function of number of qubits, n, when relevant symmetries are enforced. We show the dimension of the full Hilbert space (blue), the largest particle-number subspace with $m = n/2$ (yellow), the subspace with $m = n/2$ and $s_z = 0$ (red), and the subspace with $m = n/2$ and $s = 0$ (green). Lines are included only as a guide. All cases remain an exponential scaling Hilbert space, where use of symmetries reduces the exponential factor.

The starting point for this alternative circuit-building scheme is to find an n-qubit unitary that transforms the state $\otimes_{i=1}^{n}|0\rangle$ into an arbitrary superposition of states that share the same value of m. This is of course exactly the problem of finding a circuit that respects particle-number symmetry that we discussed extensively in the previous section. The key idea here is that by starting with the explicit n-qubit unitary upfront instead of a decomposition of it in terms of A gates, we can control how circuit parameters appear in the coefficients of the final superposition states. If the circuit parameters appear in a sufficiently simple way, then it would be straightforward to impose constraints on them so that spin symmetries are also respected. We could attempt to use the particle-number circuits from the previous section and impose constraints on the circuit parameters to enforce total spin symmetry, but this would lead to complicated, highly non-linear conditions on the parameters that would be difficult to solve.

It is straightforward to construct an n-qubit unitary that transforms $\otimes_{i=1}^{n}|0\rangle$ into an arbitrary state in a particular particle-number subspace. In the simplest case of only two spin-orbitals and a single fermion, this unitary is just the A gate multiplied by a single X gate (see Fig. 1). In the case of four spin-orbitals and two fermions, the desired unitary should generate arbitrary superposition states formed from a total of $\binom{4}{2} = 6$ basis states. While many different unitaries transform $\otimes_{i=1}^{n}|0\rangle$ into such superposition states, here we desire a unitary which generates these states with *specific* coefficients that are easily adjusted to respect the appropriate spin symmetries. The unitary we choose to use, referred to as the E_4 gate, is based on hyperspherical coordinates:

$$\begin{aligned} E_4|0000\rangle = {} & \sin u_1 \sin u_2 \sin u_3 \sin u_4 \sin u_5|0101\rangle \\ & +\sin u_1 \sin u_2 \sin u_3 \sin u_4 \cos u_5|1001\rangle \\ & +\sin u_1 \sin u_2 \sin u_3 \cos u_4|0011\rangle \\ & +\sin u_1 \sin u_2 \cos u_3|0110\rangle \\ & +\sin u_1 \cos u_2|1010\rangle \\ & +\cos u_1|1100\rangle. \end{aligned} \quad (5)$$

The E_4 gate clearly generates any two-particle state in a system with four spin-orbitals. Notice that we have also imposed time-reversal symmetry by purposely choosing the coefficients to be real. This can be easily generalized to problems without time-reversal symmetry by inserting additional arbitrary phase factors on any five of the six terms in Eq. (5). The use of hyperspherical coordinates provides a simple parameterization that automatically ensures normalization and facilitates the incorporation of spin symmetries. For example, if we want to restrict to the subspace with $s = 1, s_z = 0$, then there is only one state:

$$|s = 1, s_z = 0\rangle = \frac{1}{\sqrt{2}}(|1001\rangle + |0110\rangle).$$

This state can be created from the E_4 gate by setting $u_1 = u_2 = u_4 = \pi/2, u_3 = \pi/4, u_5 = 0$. We summarize this case and the other spin subspace of interest $s = 0, s_z = 0$, which is given by the general superposition

$$|s = 0, s_z = 0\rangle = \frac{\gamma}{\sqrt{2}}(|1001\rangle - |0110\rangle) + \alpha|1100\rangle + \beta|0011\rangle, \quad (6)$$

in Table 1.

The advantage of this construction is that we can easily fix parameters in the E_4 gate to generate any desired spin subspace in this $n = 4, m = 2$ space. There are of course many unitaries that satisfy Eq. (5). Once we have settled on a particular choice for E_4, we can then decompose it into more elementary gates. Note that in Eq. (5) we have chosen a basis ordering which leads to an efficient Gray code decomposition[36] (there are only two bit changes from term to term). A specific choice for E_4 and its corresponding decomposition are shown in the appendix.

The case of $n > 4$ spin-orbitals can be treated in a similar fashion. We begin by constructing a unitary E_n that transforms $\otimes_{i=1}^{n}|0\rangle$ into an arbitrary superposition of $\binom{n}{m}$ fixed-particle-number basis states. For example, in the case $n = 6, m = 3$, we have

$$E_6|000000\rangle = \sum_{j=20}^{1} \cos\,\mathrm{u_j} \prod_{i=1}^{j-1} \sin\,\mathrm{u_i}\big|\mathrm{p_{j1}}, \mathrm{p_{j2}}, \ldots, \mathrm{p_{j6}}\big\rangle, \tag{7}$$

where $p_{jk} \in \{0, 1\}$, $\sum_{k=1}^{6} p_{jk} = 3$, $u_{20} = 0$, and the order of the basis states is given by

$$\begin{aligned}&|000111\rangle, |001011\rangle, |001110\rangle, |001101\rangle, |011001\rangle, |011010\rangle\\ &|011100\rangle, |010110\rangle, |010101\rangle, |010011\rangle, |100011\rangle, |100101\rangle\\ &|100110\rangle, |101010\rangle, |101100\rangle, |101001\rangle, |110001\rangle, |110010\rangle\\ &|110100\rangle, |111000\rangle.\end{aligned} \tag{8}$$

This ordering results in minimal (2) bit changes per step of the Gray code, which facilitates the decomposition of E_6. We list the specific spin subspaces and summarize how to generate them by fixing the parameters of E_6 in Table 2.

An example of a gate decomposition for a specific choice of E_6 is given in the appendix along with explicit gate counts for E_4, E_6, and E_8.

Extrapolating from these examples, we can then form a general procedure for constructing symmetry-preserving state preparation circuits for any valid choice of the quantum numbers n, m, s, s_z for a time-reversal-symmetric system as follows:

1. For a given choice of n, m, there are $\binom{n}{m}$ basis states $(|p_{j1}, p_{j2}, \ldots, p_{jn}\rangle)$ that span the corresponding particle-number subspace, such that $\sum_{k=1}^{n} p_{jk} = m$. Assign hyperspherical coefficients to these basis states according to $\sum_{j=\binom{n}{m}}^{1} \cos u_j \prod_{i=1}^{j-1} \sin u_i \big|p_{j1}, p_{j2}, \ldots, p_{jn}\big\rangle$, where $u_{\binom{n}{m}} = 0$, and the basis states are ordered such that there are only two bit changes per step in the Gray code.
2. Determine the constraints that must be imposed on the u_i such that the appropriate spin subspace labled by s and s_z is obtained.
3. Construct a unitary E_n such that the first column contains the coefficients in Step 1. The remaining columns can be chosen as desired so long as they respect unitarity.
4. Decompose E_n using a Gray code scheme or alternative gate decomposition technique.

Following this procedure, we can enforce arbitrary spin and particle-number symmetries for any choice of n, m, s, s_z. Step 2 is facilitated by the use of hyperspherical coordinates because various basis states are easily eliminated from the final superposition by setting the corresponding u_j to $\pi/2$, yielding another hypersphere parameterization of lower dimension. This recursive structure in the coefficients maintains regularity and simplicity as the effective Hilbert space dimension is reduced. Note that in cases where it is necessary to fix some of the u_j in terms of the others (see e.g., the case of $s = s_z = 0$ in Table 1 or $s = 1/2 = \pm s_z$ in Table 2), use of inverse functions without a restricted real-domain (e.g. $\tan^{-1}, \cot^{-1}$) are desirable for optimization stability and enforcement of time-reversal symmetry, while the use of other inverse trigonometric functions could yield complex-valued coefficients in the trial states. Finally, we emphasize that here it is not necessary to numerically compute fidelities to confirm that the resulting circuits indeed span the appropriate symmetry subspace as in the previous section because this is guaranteed by construction in the present approach.

To demonstrate the efficiency of our state preparation circuits, we use them in a VQE simulation to compute the ground state energy of the H_2 molecule. We work in the STO-3G basis and map to qubits using the Jordan-Wigner transform.

We note that there are many choices of qubit mappings and unitary transformations that can be used to reduce the complexity of molecular simulations.[37] However, we have formulated our ansätze to preserve symmetries in the natural basis where each qubit represents a physical spin orbital. There are also several techniques to tailor an ansatz to efficiently describe the allowed transitions in H_2, but here we are concerned with constructing general ansätze, applicable to any molecule, which only utilize particle number and spin symmetries, to reduce the complexity of the resulting circuits. Incorporation of particle number, spin projection and total spin symmetries in conjuction with other qubit mappings and reduction methods is an interesting topic, but beyond the scope of the present work. The results are given in Fig. 8, which shows the difference between the computed ground state energy and the exact ground state energy. For comparison, we also show the results obtained from standard ansätze (SWAPRZ, RY, RYRZ, and UCCSD) included in IBM's Qiskit software package.[38]

Table 1. Two particle spin symmetry.

s, s_z	u_1	u_2	u_3	u_4	u_5
1, 0	$\pi/2$	$\pi/2$	$\pi/4$	$\pi/2$	0
0, 0	u_1	$\pi/2$	$-\tan^{-1}(\csc\,u_4)$	u_4	0

By fixing the coefficients of Eq. 5, we can generate the two spin subspaces defined by their quantum numbers for the $n = 4, m = 2$ space. In the case of $s = 1$, $s_z = 0$, the subspace is spanned by a single state, so all the u_i are fixed. The subspace with $s = s_z = 0$ is three-dimensional, and so two of the u_i are left unspecified

Table 2. Three particle Spin Symmetry.

s, s_z	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}	u_{11}	u_{12}	u_{13}	u_{14}	u_{15}	u_{16}	u_{17}	u_{18}	u_{19}
$\frac{3}{2}, \frac{1}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\cot^{-1}(\frac{1}{\sqrt{2}})$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	0	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$
$\frac{3}{2}, -\frac{1}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\cot^{-1}(\frac{1}{\sqrt{2}})$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{4}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	0	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$
$\frac{1}{2}, \frac{1}{2}$	u_1	$\frac{\pi}{2}$	u_3	$\frac{\pi}{2}$	$-\cot^{-1}(\delta)$	u_6	$\frac{\pi}{2}$	u_8	$\frac{\pi}{2}$	u_{10}	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	u_{15}	$\frac{\pi}{2}$	$\frac{\pi}{2}$	u_{18}	0
$\frac{1}{2}, -\frac{1}{2}$	$\frac{\pi}{2}$	u_2	$\frac{\pi}{2}$	u_4	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$\frac{\pi}{2}$	$-\cot^{-1}(\kappa)$	$\frac{\pi}{2}$	u_{11}	$\frac{\pi}{2}$	u_{13}	u_{14}	$\frac{\pi}{2}$	u_{16}	u_{17}	$\frac{\pi}{2}$	$\frac{\pi}{2}$

By fixing the coefficients of Eq. 7, we can generate the four spin subspaces defined by their spin quantum numbers for the $n = 6$, $m = 3$ space. Here $\delta = \{\sin(u_{10})\cos(u_{15})\sin(u_6)\sin(u_8) + \sin(u_6)\cos(u_8)\}$, $\kappa = \{\sin(u_{11})\sin(u_{13})\sin(u_{14})\cos(u_{16}) + \sin(u_{11})\cos(u_{13})\}$. The order of basis states for this space is defined in the text

Furthermore, we show the variational parameter and CNOT gate counts in Table 3, accompanying Fig. 8 for all ansätze considered. It is clear that our proposed circuits obtain the ground state energy with high accuracy using the minimal number of variational parameters and a low number of CNOT gates. For the RY, RYRZ, and SWAPRZ ansätze, we choose a circuit depth of three, while for UCCSD, we use a circuit depth of one. In all cases, we use the BFGS optimizer (run several times with random initial conditions) included in Qiskit and consider an ideal, noiseless simulation.

Figure 9 shows the results of a similar simulation, but this time with noise included. This noise is characterized by the physical parameters of IBM's Poughkeepsie device. Here we only take the standard ansatz that has reasonable parameter and CNOT scaling, RY, and compare it to our ansätze. Contrary to the noiseless case, the RY ansatz here is chosen to have linear connectivity (nearest-neighbor CNOT connectivity) and a circuit depth of one as this configuration performs best under these noisy conditions. With these settings, the RY ansatz has 8 parameters and 3 CNOT gates. The CNOT and parameter counts for our ansätze are the same as in Fig. 8 since these are determined by the ground state symmetries of the molecule. Even though the RY ansatz has fewer costly CNOT gates than our *A* gate ansatz and is over parameterized for the target ground state, it is outperformed by the *A* gate ansatz. This is due to the fact that at this low depth, the RY ansatz does not have sufficient coverage of the target Hilbert space. Increasing the depth of the RY ansatz will increase its coverage of the target space, but also increases the number of CNOT gates. This is not an equivalent trade-off and actually decreases the ability of the ansatz to approximate the ground state. In the case of both *E* gate ansätze, their relatively larger number of CNOT gates can be seen to hinder their performance. However, as these gates are constructed in terms of arbitrary state generation, these gate counts could be improved with more efficient state generation algorithms. In addition, mitigation of CNOT errors may be possible with the use of Richardson extrapolation,[39,40] a technique we will pursue in future work.

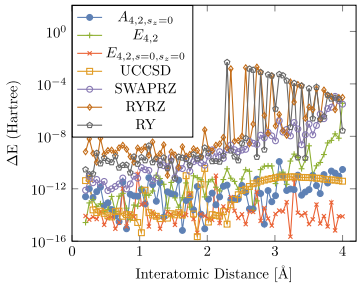

Fig. 8 H_2 Noiseless VQE. Energy difference from the exact ground state of H_2 for various ansätze. All ansätze perform well below chemical accuracy, but vary significantly in their number of required parameters and CNOT gates (shown in the accompanying table). Our proposed ansätze ($A_{4,2}$ and $E_{4,2}$) achieve very small energy differences and have low to modest CNOT counts. The UCCSD ansatz also performs very well with low parameters (3), but has a fairly large CNOT count in this example.

Table 3. 4-qubit ansätze comparison.

Ansätze	# CNOT	# Parameter
$A_{4,2,s_z=0}$	6	3
$E_{4,2}$	14	5
$E_{4,2,s=0,s_z=0}$	20	2
UCCSD	56	3
SWAPRZ	34	72
RYRZ	18	32
RY	18	16

In the ideal, noiseless case, we compare many standard ansätze to our new proposed ansätze and find that our proposed ansätze are more compact in terms of CNOT count and parameter count

An interesting feature in Fig. 9a) can be seen where the $E_{4,2}$ ansatz crosses the $E_{4,2,s=0,s_z=0}$ ansatz. In principle, the spin symmetry ansatz should always have a better ability to target the ground state in a VQE algorithm. This is explained in Fig. 9b), where as we see that the $E_{4,2}$ ansatz abruptly changes the total spin value of its lowest energy state. Indeed, the triplet states ($s = 1$) are nearly degenerate with the singlet state ($s = 0$, Eq. 6) at large bound distances. At an interatomic distance near two Angstroms, the energy difference between the lowest energy singlet and triplet states is $\mathcal{O}(10^{-2})$ Hartree. The optimization landscape becomes much more difficult to navigate when there are many nearly equal local minima, as is the case here at large bond distances. Therefore since the algorithm attempts to find the lowest energy eigenvalue and the singlet and triplet states are nearly degenerate at longer bond distances, the optimizer will frequently return a minimum (triplet state) rather than the global minimum (singlet state). In the presence of noise, this issue is blurred further as both energies are not exactly attainable, but a triplet state is trivial to produce (e.g. $s = 1, s_z = 1$, $|1010\rangle$), so finding the triplet state as the lowest energy state is frequently encountered. Additionally, this issue grows more difficult to navigate with larger molecules and strongly correlated systems, which have many more nearly degenerate ground states. Naturally one could also enforce a spin eigenvalue constraint on the optimizer by modifying the objective function, but this places stricter requirements on the optimizer and requires a physical measurement of spin eigenvalues, through tomography or additional circuit elements. However, enforcing symmetry in the ansätze lessens the burden on the optimizer and does not directly require spin measurements.

We also perform similar noisy simulations on *LiH*. *LiH* in the STO-3G basis maps to 12 qubits but we follow the same methods as ref., [8] removing two non-interacting orbitals and freezing the core orbital, reducing *LiH* to 6 qubits. We show in Fig. 10 and Table 4 that our ansätze perform comparably for *LiH* as they do for H_2 with notably more error for the *E* gate ansatz as it requires significantly more CNOTs. As in the noisy H_2 case, the RY ansatz is chosen to have a fixed depth (one) and connectivity (nearest-neighbor). These settings give the best performance under these noisy conditions but the same coverage-CNOT trade off issue arises as before. From both the H_2 and *LiH* results we can see that the *A*-gate circuit finds a middle-ground in terms of CNOT and parameter count and has good performance, while the *E*-gate circuit exactly minimizes parameter count based on particle number, spin projection, and total spin, but its performance suffers due to a large CNOT count. Thus, in the presence of noise, we find that it can be favorable to relax some of the symmetry constraints to reduce the circuit depth at the expense of more variational parameters.

Since our constructions maintain number and spin symmetries, potentially any noise sources which violate these symmetries can be mitigated through post-selection or symmetry verification.[13,14,27] This will be a topic of future work. Another interesting question for future work is whether it is possible to find circuits that minimize the number of CNOT gates while respecting spin symmetries and while using the minimal number of parameters necessary to span the symmetry subspace.

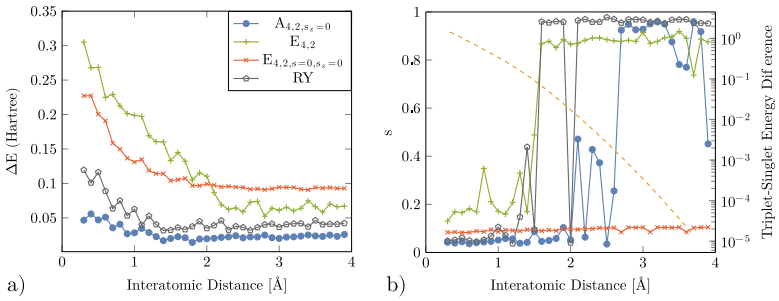

Fig. 9 H_2 Noisy VQE. a) Energy difference from the exact ground state of H_2 in the presence of noise characterized by IBM's Poughkeepsie device. Included is Qiskit's state preperation and measurement error mitigation for all ansätze. **b**) Total spin eigenvalues for the same dissociation curve (left axis). The total spin values indicate that some ansätze drift outside the singlet subspace ($s = 0$) and begin to instead find a (nearly degenerate) triplet state ($s = 1$) at large bond distances. Since the $E_{4,2,s=0,s_z=0}$ ansatz is restricted to a fixed total spin, it is the only displayed ansätze which always attempts to find the true ground state across the full dissociation curve. The energy difference between this triplet state and the true singlet ground state as a function of interatomic distance is also shown for comparison (orange dashed line, right axis).

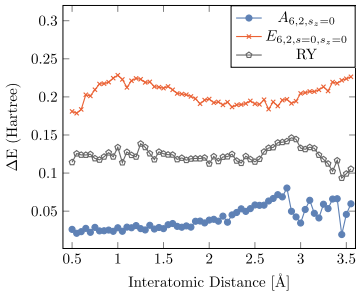

Fig. 10 *LiH* Noisy VQE. Energy difference from the exact ground state of *LiH* in the presence of noise characterized by IBM's Poughkeepsie device. With appropriate orbital reductions, *LiH* is run on 6 qubits and has two excitations. The A-gate ansatz continues to perform well utilizing particle number and spin projection symmetries and relatively low CNOT count for this dimension.

Table 4. 6-qubit ansätze comparison.

Ansätze	# CNOT	# Parameter
$A_{6,2,s_z=0}$	25	8
$E_{6,2,s=0,s_z=0}$	48	5
RY	5	12

CNOT and parameter count for the case of estimating the ground state of *LiH* in the presence of typical device noise

DISCUSSION

In this work, we presented general schemes to construct state preparation circuits for quantum simulation that respect a number of symmetries that commonly arise in systems of interest, including particle number, time-reversal, total spin, and spin magnetization. In each case, we provide general construction procedures, explicit examples of circuits, and gate counts. In the case of particle-number symmetry, for which state preparation circuits have been proposed previously by other authors, our circuits outperform existing methods in terms of the number of two-qubit entangling gates they contain. Enforcing spin symmetries in addition to particle number can significantly enhance the performance of the VQE by preventing the algorithm from wasting time exploring vast, irrelevant regions of Hilbert space. Also, our ansätze are *guaranteed* to contain the ground state of interest, with the minimal number of parameters defined by the symmetry subspace. It is important to emphasize, however, that the number of variational parameters still grow exponentially with system size even when all symmetries are imposed. Thus, spanning the entire symmetry subspace will become impractical as the size of the quantum processor is increased beyond a few tens of qubits. However, we expect that our symmetry preservation techniques will continue to play an important role for larger NISQ devices as systematic methods to further reduce the search to still smaller regions of Hilbert space are developed. This could be done, for example, by looking for ways to freeze large sets of the variational parameters appearing in our circuits without sacrificing the accuracy of the ansatz. In addition, as a general state creation circuit with a particular symmetry, one can imagine applications for these circuits in the Phase Estimation Algorithm (PEA).

METHODS

Confirmation of circuit compilation was performed in Mathematica software as well as numerical optimization to confirm that the proposed circuits span any state in the desired Hilbert space. In numerical optimization testing, we find that Mathematica's NMaximize method "Simulated Annealing" performs best for our purposes, when run over many samples of a randomly chosen state in the Hilbert space. Gate decomposition was performed using the method of Gray codes, was done "by hand" and confirmed in Mathematica.

Our VQE simulations were done using IBM's Qiskit software, coding in our own proposed ansätze, which enforce relevant symmetries. Noisy simulations were simulated using the noise parameters of IBM's Poughkeepsie device.

DATA AVAILABILITY

The data that support the findings of this study are available from the authors upon request.

CODE AVAILABILITY

A custom Mathematica code to reproduce our results is available on GitHub.[41]

Received: 25 April 2019; Accepted: 10 November 2019;
Published online: 28 January 2020

REFERENCES

1. Feynman, R. P. Simulating physics with computers. *Int. J. Theoretical Phys.* **21**, 467–488 (1982).
2. Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
3. Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707 (2005).
4. Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
5. Kühn, M., Zanker, S., Deglmann, P., Marthaler, M. & Weiß, H. Accuracy and resource estimations for quantum chemistry on a near-term quantum computer. *J. Chem. Theory Comput.* **15**, 4764 (2019).
6. Peruzzo, A. et al. A variational eigenvalue solver on a quantum processor. *Nat. Commun.* **5**, 4213 (2014).
7. O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
8. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242 (2017).
9. Colless, J. I. et al. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X* **8**, 011021 (2018).
10. Shen, Y. et al. Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys. Rev. A* **95**, 020501 (2017).
11. Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* **8**, 031022 (2018).
12. Nam, Y. et al. Ground-state energy estimation of the water molecule on a trapped ion quantum computer. Preprint at https://arxiv.org/abs/1902.10171 (2019).
13. McArdle, S., Yuan, X. & Benjamin, S. Error-mitigated digital quantum simulation. *Phys. Rev. Lett.* **122**, 180501 (2019).
14. Bonet-Monroig, X., Sagastizabal, R., Singh, M. & O'Brien, T. E. Low-cost error mitigation by symmetry verification. *Phys. Rev. A* **98**, 062339 (2018).
15. Barkoutsos, P. K. et al. Quantum algorithms for electronic structure calculations: particle-hole hamiltonian and optimized wave-function expansions. *Phys. Rev. A* **98**, 022322 (2018).
16. Bartlett, R. J., Kucharski, S. A. & Noga, J. Alternative coupled-cluster ansätze ii. the unitary coupled-cluster method. *Chemical Physics Letters* **155**, 133–140 (1989).
17. Kutzelnigg, W. Error analysis and improvements of coupled-cluster theory. *Theoretica Chimica Acta* **80**, 349–386 (1991).
18. Ryabinkin, I. G., Yen, T.-C., Genin, S. N. & Izmaylov, A. F. Qubit coupled cluster method: a systematic approach to quantum chemistry on a quantum computer. *J. Chem. Theory Comput.* **14**, 6317–6326 (2018).
19. Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 3007 (2019).
20. Bian, T., Murphy, D., Xia, R., Daskin, A. & Kais, S. Quantum computing methods for electronic states of the water molecule. *Mol. Phys.* **117**, 2069–2082 (2019).
21. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
22. McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023 (2016).
23. Ryabinkin, I. G., Genin, S. N. & Izmaylov, A. F. Constrained variational quantum eigensolver: quantum computer search engine in the fock space. *J. Chem. Theory Comput.* **15**, 249–255 (2019).
24. Wang, H., Ashhab, S. & Nori, F. Efficient quantum algorithm for preparing molecular-system-like states on a quantum computer. *Phys. Rev. A* **79**, 042335 (2009).
25. Roth, M. et al. Analysis of a parametrically driven exchange-type gate and a two-photon excitation gate between superconducting qubits. *Phys. Rev. A* **96**, 062323 (2017).
26. Egger, D. et al. Entanglement generation in superconducting qubits using holonomic operations. *Phys. Rev. Applied* **11**, 014017 (2019).
27. Sagastizabal, R. et al. Error mitigation by symmetry verification on a variational quantum eigensolver. *Phys. Rev. A* **100**, 010302 (2019).
28. Ganzhorn, M. et al. Gate-efficient simulation of molecular eigenstates on a quantum computer. *Phys. Rev. Applied* **11**, 044092 (2019).
29. Jordan, P. & Wigner, E. über das paulische äquivalenzverbot. *Z. Phys.* **47**, 631 (1928).
30. Zhang, J., Vala, J., Sastry, S. & Whaley, K. B. Geometric theory of nonlocal two-qubit operations. *Phys. Rev. A* **67**, 042313 (2003).
31. Gard, B., Zhu, L., Barron, G., Barnes, E. & Economou, S. Supplementary Information for: Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. https://doi.org/10.1038/s41534-019-0240-1.
32. Shende, V. V. & Markov, I. L. Quantum circuits for incompletely specified two-qubit operators. *Quantum Inf. Comput.* **5**, 49 (2005).
33. Möttönen, M., Vartiainen, J. J., Bergholm, V. & Salomaa, M. M. Quantum circuits for general multiqubit gates. *Phys. Rev. Lett.* **93**, 130502 (2004).
34. Bergholm, V., Vartiainen, J. J., Möttönen, M. & Salomaa, M. M. Quantum circuits with uniformly controlled one-qubit gates. *Phys. Rev. A* **71**, 052330 (2005).
35. Ortiz, G., Gubernatis, J. E., Knill, E. & Laflamme, R. Quantum algorithms for fermionic simulations. *Phys. Rev. A* **64**, 022319 (2001).
36. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. (Cambridge University Press, New York, 2011).
37. Bravyi, S., Gambetta, J.M., Mezzacapo, A. & Temme, K. Tapering off qubits to simulate fermionic hamiltonians. Preprint at https://arxiv.org/abs/1905.01534 (2017).
38. Abraham, H. et al. Qiskit: An open-source framework for quantum computing. https://qiskit.org/. (2019).
39. Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
40. McCaskey, A.J. et al. Quantum chemistry as a benchmark for near-term quantum computers. Preprint at https://arxiv.org/abs/1905.01534 (2019).
41. Gard, B., Zhu, L., Barron, G., Barnes, E. & Economou, S. https://github.com/bgard1/Efficient-Symmetry-Preserving-State-Preparation-Circuits-for-the-VQE-algorithm. (2019).

ACKNOWLEDGEMENTS

This research was supported by the US Department of Energy (Award No. de-sc 0019199) and the National Science Foundation (Award No. 1839136). S.E.E. also acknowledges support from Award No. de-sc 0019318 from the Department of Energy. This research used quantum computing system resources supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research program office. Oak Ridge National Laboratory manages access to the IBM Q System as part of the IBM Q Network.

AUTHOR CONTRIBUTIONS

S.E.E., E.B., and N.J.M. structured and supervised the project. E.B. proposed methods for symmetry encoding and minimal parameterization. B.T.G. performed simulations, formulated general circuits and drafted the manuscript. L.Z. performed circuit decomposition, CNOT gate counts and suggested example circuits. G.S.B. drafted figures and confirmed circuit compilations. All authors contributed to the revisions of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information is available for this paper at https://doi.org/10.1038/s41534-019-0240-1.

Correspondence and requests for materials should be addressed to B.T.G.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2020

Chapter 8

Preserving Symmetries for Variational Quantum Eigensolvers in the Presence of Noise

This chapter includes pages from our publication [3]

G. S. Barron, B. T. Gard, O. J. Altman, N. J. Mayhall, E. Barnes, and S. E. Economou, "Preserving symmetries for variational quantum eigensolvers in the presence of noise", Physical Review Applied **16**, 034003 (2021)

Copyright (2021) of the American Physical Society.

The contributions to this work are as follows: My contributions were carrying out numerical simulations, maintenance, writing, testing, and distribution of most of the supporting source code, preparation of figures, primary writing of the manuscript, bibliography preparation and conducting background material search. B.T.G. assisted with theoretical work, code development, numerical simulations and manuscript preparation. O.J.A. assisted with numerical simulations. N.J.M., E.B., and S.E.E. supervised.

Preserving Symmetries for Variational Quantum Eigensolvers in the Presence of Noise

George S. Barron,[1,*] Bryan T. Gard,[1] Orien J. Altman,[1] Nicholas J. Mayhall,[2] Edwin Barnes,[1] and Sophia E. Economou[1]

[1] *Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2] *Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

(Received 8 March 2021; accepted 30 July 2021; published 1 September 2021)

One of the most promising applications of noisy intermediate-scale quantum computers is the simulation of molecular Hamiltonians using the variational quantum eigensolver (VQE). We show that encoding symmetries of the simulated Hamiltonian in the VQE ansatz reduces both classical and quantum resources compared to other widely available ansatze. Through simulations of the H_2 molecule and of a Heisenberg model on a two-dimensional lattice, we verify that these improvements persist in the presence of noise. This is done using both real IBM devices and classical simulations. We also demonstrate how these techniques can be used to find molecular excited states of various symmetries using a noisy processor. We use error-mitigation techniques to further improve the quality of our results.

DOI: 10.1103/PhysRevApplied.16.034003

I. INTRODUCTION

Quantum computers are believed to be one of the most promising technologies currently being developed that will help extend the reach of scientific discovery. This may be achieved through quantum simulation [1], which leverages the properties of a quantum processing unit (QPU) to simulate naturally occurring quantum-mechanical systems. One of the most popular algorithms within the reach of near-term devices is the variational quantum eigensolver (VQE) [2–7].

This algorithm falls under a more general class of algorithms known as hybrid variational quantum algorithms [8–14]. The general principle of these algorithms is to use a feedback loop between the quantum and classical computers to minimize a predefined cost function. This method has been applied to a variety of quantum systems in both theoretical [15–27] and experimental [2,3,5,12,28–31] contexts. In the case of the VQE, the predefined function is the expectation value of the simulated Hamiltonian with respect to the state of the QPU. Additionally, a variety of techniques are available that use the VQE to find higher excited states of such systems [6,32,33]. These variational algorithms in general are of particular interest for near-term devices due to their proposed noise resilience. Notably, Ref. [4] has demonstrated noise resilience against coherent errors and Ref. [34] has proved noise resilience against incoherent errors for a special class of variational algorithms.

The advantage of hybrid variational algorithms lies in their ability to exploit advanced classical computational resources without having to store the wave function on a classical computer. Wave functions are instead prepared and measured on the QPU. The success of these algorithms depends sensitively on challenges involving optimizing rapidly varying functions with many parameters, limited computational resources, and the presence of noise [24,35–45]. Work has been done to manipulate the input and output of the QPU so as to mitigate and characterize the error [46–49]. Another, perhaps more fundamental, issue is related to the complexity of the quantum circuits used to prepare the wave functions. This complexity in turn is fundamentally related to the number of parameters needed to describe arbitrary states. Hence, a fundamental question is "how likely is it that a given parametrized quantum circuit (a variational ansatz) will represent the targeted state with sufficient accuracy?" Typically, there are no accuracy guarantees for a given ansatz. One approach to designing ansatze with tunable accuracy is to iteratively build a customized variational ansatz in a "problem-tailored" fashion, as proposed with the adaptive derivative assembled problem-tailored-VQE (ADAPT VQE) algorithm [50,51]. This is accomplished by growing the ansatz operator by operator, each time selecting the operator that creates the largest change to the objective function, using the VQE as a subroutine at each step. This is in contrast to alternative approaches that fix the number of parameters from the start. In Ref. [22], we have shown how to create

*gbarron@vt.edu

 © 2021 American Physical Society

quantum circuits that have exactly the necessary number of variational parameters to describe any state in the relevant symmetry subspace of the Hilbert space.

One of the advantages of Ref. [22] is that the ansatz is guaranteed to find the correct state in the absence of noise, both for ground and for excited states. Other algorithms that accomplish excited-state simulation are sometimes iterative and rely on previous results of a VQE, thus introducing computational overheads. Despite these advantages of symmetry-enforcing circuits, however, it is not clear how the presence of noise might affect their performance. The operators that create noise clearly will not commute with the same symmetries as the target Hamiltonian and, as such, might drive the system into undesirable symmetry subspaces. Will a circuit that preserves these symmetries perform worse due to a lack of relevant error-mitigating degrees of freedom?

In this work, we consider the effect that noise has on the ability of a VQE to preserve symmetries. The simulations here are done using IBM's QISKIT [52] software, using noise models derived from real superconducting devices. We also test our methods on real IBM devices. We perform these simulations and experiments using our previously developed techniques [22] for encoding symmetries of the problem Hamiltonians directly into the state-preparation circuit. We test our approach on H_2 and on a two-dimensional Heisenberg model, finding in both cases that our ansatze can outperform standard *ad hoc* ansatze. We adapt this approach to solving for excited states that obey certain symmetries. This allows us to guarantee the ability of the ansatz to find the desired state and, in fact, to accomplish this with a minimal number of variational parameters. Additionally, our approach for producing excited states does not rely on previous runs of VQEs and is parallelizable for the different excited states calculated.

This paper is structured as follows. In Sec. II, we provide the necessary background to make this paper self-contained. Specifically, in Sec. II A, we review the theory behind VQE and its application to quantum chemistry problems. In Sec. II B, we review our earlier work on symmetry-preserving VQE ansatze. In Secs. III and IV, we perform simulations on the H_2 molecule and analyze the performance of these algorithms in noiseless and noisy simulation contexts, respectively. In Sec. IV, we use, benchmark, and compare a variety of error-mitigation techniques, as well as discuss challenges associated with noisy optimization. In Sec. V, we compare the ability of the symmetry-preserving and *ad hoc* ansatze to preserve desired symmetries with both noisy simulations and experiments on IBM devices. In Sec. VI, we present the results of this procedure for excited-state simulation. In Sec. VII, we discuss potential future work incorporating device connectivity. In Sec. VIII, we give a summary of our conclusions.

II. BACKGROUND

A. Variational quantum eigensolver

The goal of a VQE is to minimize an objective function of the form $f(\vec{\theta}) = \langle\psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle$, where $|\psi(\vec{\theta})\rangle = U(\vec{\theta})|\psi_0\rangle$ is the state of the QPU after initializing the register to some initial state $|\psi_0\rangle$ and applying a circuit with unitary $U(\vec{\theta})$, where $\vec{\theta}$ is a vector of variational parameters chosen by an optimization algorithm. Per the variational principle, $f(\vec{\theta}) \geq E_0$, where E_0 is the ground-state energy of the Hamiltonian H.

We primarily focus on the second quantized molecular Hamiltonian

$$\hat{H}_f = \sum_{ij} h_{ij}\,\hat{a}_i^\dagger \hat{a}_j + \sum_{ijkl} g_{ijkl}\hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l, \tag{1}$$

where $\hat{a}$ ($\hat{a}^\dagger$) is the fermionic annihilation (creation) operator. The quantities h_{ij} and g_{ijkl} are the single and double electron integrals, respectively, which can be computed efficiently on a classical computer with existing software. For these calculations, we use the STO-3G basis set. To evaluate the objective function $f(\vec{\theta})$ on the QPU, one must map the fermionic operators to a set of qubit operators in a way that accounts for the fermionic anticommutation relations.

This can be accomplished with several different methods, including the Jordan-Wigner (JW) [53], parity [54], Bravyi-Kitaev (BK) [55], or Bravyi-Kitaev Super Fast (BKSF) [56] mappings. In this work, we choose the JW mapping, as it directly maps occupations of fermionic spin orbitals to excitations of qubits. The result is that our ansatz consists of only nearest-neighbor two-qubit gates, which we choose to preserve particle-number and spin projection. This choice is not necessary but other maps would require gates that act on more than two qubits to attain the same effect. To achieve this, one would need to find mapping-specific gates that preserve the desired expectation values of the states on which they act but this is beyond the scope of the current work.

The resulting Hamiltonian appearing in the objective function $f(\vec{\theta})$ is a weighted sum of Pauli strings $\hat{H} = \sum_i \alpha_i T_i$, where $\alpha_i \in \mathbb{R}$ are the weights of Pauli strings and $T_i \in \{I, X, Y, Z\}^{\otimes n}$ for n qubits (with the JW mapping, this is the number of spin orbitals). The role of the QPU is to compute expectation values of this Hamiltonian. In particular, for a given set of parameters $\vec{\theta}$, the QPU calculates $\{\langle T_i\rangle_{\vec{\theta}}\}$, which may then be weighted on the classical computer according to the coefficients α_i to form $\langle H\rangle_{\vec{\theta}}$. Work has been done to reduce the number of measurements at this step based on the properties of $\{T_i\}$ [41]. Additionally, Ref. [57] has developed techniques that use symmetry to reduce the number of qubits for certain systems including H_2; however, the transformations used are not compatible with our framework.

Despite substantial progress, accurately performing VQEs on current and near-term hardware is still challenging for a number of reasons. First, calls to the objective function $f(\vec{\theta})$ are inherently probabilistic, meaning that one must use an optimization algorithm to minimize $f(\vec{\theta})$ that is resistant to noise. Additionally, the number of variational parameters used to construct the variational form can potentially limit the performance of an optimization algorithm in finding a suitable minimum for the estimated ground-state energy. Furthermore, the depth of circuits and number of controlled-NOT (CNOT) gates are costly resources when executing circuits on a real QPU, so creating ansatze that use these resources efficiently is paramount.

B. Symmetry-preserving ansatze

The minimal number of parameters for a general Hilbert space of n qubits (i.e., n spin orbitals) is given by $2(2^n - 1)$. However, if symmetries are present, they can be used to reduce this parameter count, sometimes dramatically. For example, if particle number is conserved, then the number of parameters reduces to $2\left(\binom{n}{m} - 1\right)$, where $m < n$ is the number of fermions. If the Hamiltonian also respects time-reversal symmetry, as is often the case in molecular problems, then the energy eigenstates can be chosen to have only real coefficients, which further reduces the parameter count by a factor of 2. Total spin s and spin projection s_z are also commonly conserved, in which case the parameter count becomes [22]

$$\sum_{k=0}^{m/2-s} \binom{n/2}{k} \binom{n/2-k}{m-2k} \times \frac{(2s+1)(m-2k)!}{(m/2-k-s)!(m/2-k+s+1)!}. \quad (2)$$

Choosing a particular symmetry-preserving ansatz amounts to restricting the search space of the optimizer to a smaller region of the Hilbert space. This in turn reduces the number of variational parameters, lessening the burden on the classical optimizer. We show below that it also lessens the demands on the QPU by cutting down the number of function calls in the optimization process and reducing the circuit depth.

The use of a symmetry-preserving ansatz contrasts sharply with using ansatze that rely on *ad hoc* arguments [5,58]. These ansatze usually attempt to express all possible states in the Hilbert space, including the ground state of the Hamiltonian in question. The conventional wisdom here is that more variational parameters result in more flexibility in producing the target states. However, too many variational parameters can result in large circuit depths and challenging conditions for the optimizer. It may also not be possible to exactly represent the ground state with these types of ansatze if they are not designed to uniformly cover the relevant part of the Hilbert space.

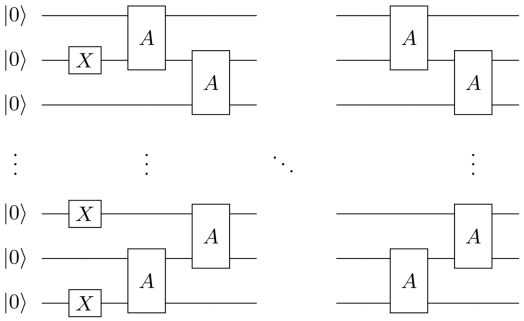

FIG. 1. The circuit used to produce the "ASWAP" ansatz. Each A gate has two separate parameters, θ and ϕ. Each A gate preserves the particle number of the state it acts on, so the entirety of the ansatz does too. The initial X gates are used to place the register in the desired particle-number subspace.

In this work, we primarily focus on the "ASWAP" ansatz, which is constructed using gates that preserve the number of excitations in a state. In the case of a molecular Hamiltonian, this corresponds to fixing the number of electrons occupying some number of orbitals. The gates that we consider [21] are of the form

$$A(\theta,\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & e^{i\phi}\sin\theta & 0 \\ 0 & e^{-i\phi}\sin\theta & -\cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3)$$

Reference [22] has shown how to construct circuits that preserve particle-number and spin projection S_z by tiling these A gates in a regular pattern like the one shown in Fig. 1. The preservation of any value of $\langle S_z \rangle$ is accomplished by dividing the register of qubits into two halves (the upper half represents spin-up orbitals and the lower half spin-down). The parameters for the A gates that join the two spin subspaces are then set to 0. The desired S_z subspace can be selected by changing which subset of qubits are acted on by the initial X gates. In this work, we primarily use the ASWAP ansatz that fixes $\langle S_z \rangle = 0$, except when considering excited states in Sec. VI.

The advantages of the symmetry-preserving ansatze are thus twofold. First, restricting the space of states prepared by the ansatz always produces states that are consistent with physical constraints imposed by the Hamiltonian. Second, both the classical and quantum resources required to run the VQE are reduced. The latter is especially important when considering near-term devices.

III. PERFORMANCE COMPARISONS OF VARIOUS ANSATZE WITHOUT NOISE

Before testing the performance of the symmetry-preserving ansatze in the presence of noise, we first demonstrate that the correct ground state can be prepared in a noiseless context. To do this, we run different instances of the VQE algorithm for H_2 under the JW mapping on four qubits using a noiseless simulator. The noiseless simulator used is the "state-vector simulator" implemented in QISKIT. To perform the optimization of the variational parameters in the VQE, we use the limited-memory Broyden-Fletcher-Goldfarb-Shannon (BFGS) optimization algorithm. For noiseless simulations, the number of CNOT gates in an ansatz will not affect the quality of the results. Thus, if a given ansatz is able to reach chemical accuracy, the relevant benchmark is how many function evaluations are required to do so.

Since our ansatze are symmetry preserving, the reference state ultimately determines the targeted values of $\langle N \rangle$ and $\langle S_z \rangle$. Therefore, in the case of H_2, we expect a singlet ground state, so we have $\langle N \rangle = 2$ and $\langle S_z \rangle = 0$. The ansatze that we compare against are standard *ad hoc* ansatze available in QISKIT [52]. Generally, the structure of these ansatze involves interleaving layers of single-qubit rotations with entangling operations. In the case of the RY (RYRZ) ansatze, layers of single-qubit $R_Y(\theta)$ [$R_Y(\theta)R_Z(\phi)$] operations on all qubits are interleaved with CZ gates. However, since CZ and CNOT are locally equivalent, we instead count the number of CNOT operations for comparison. In all cases when executing circuits, we use the highest level of optimization available in QISKIT to reduce gate counts (i.e., transpilation). The number of such layers in the ansatz is referred to as the depth. The SwapRZ ansatz interleaves $R_Z(\theta)$ operations with parametrized SWAP operations, meaning that this ansatz preserves particle-number symmetry. In contrast with our ASWAP ansatz, the SwapRZ ansatz does not, in general, span the full symmetry subspace and so does not necessarily capture the ground state exactly. Additionally, the ASWAP ansatz generally has fewer parameters and CNOT gates. In these *ad hoc* ansatze, the entangling operations performed at each layer can be chosen for the appropriate problem. For the ansatze considered here, we perform entangling operations between all pairs of qubits unless specified otherwise.

In Fig. 2, we illustrate the performance of several *ad hoc* ansatze, along with our ASWAP ansatz. All the ansatze considered are able to find the correct ground-state energy to well below chemical accuracy (approximately 1.5 millihartree). One distinction between our ansatz and the *ad hoc* ansatze, however, is that the ASWAP ansatz requires at least 5 times fewer function evaluations to converge compared to any of the other considered ansatze, while simultaneously having lower error rates. For the example shown in Fig. 2, ASWAP requires 414 function calls, while RY uses 2142. This improvement is expected since the number of variational parameters is smaller than the other considered ansatze. This illustrates the key point that enforcing known symmetries of the simulated Hamiltonian can significantly reduce the computational load on the QPU, namely the number of calls to the objective function.

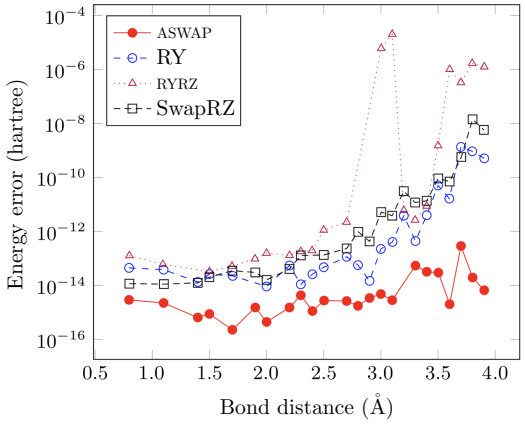

FIG. 2. Results for H_2 in a noiseless simulator with the limited-memory BFGS algorithm. All results are well below chemical accuracy and sufficiently converged. To account for bad initial variational parameters, each result is obtained by running additional optimizations with randomized initial variational parameters. The lowest energy found with this method is then chosen as the final result. The median number of iterations of the BFGS optimizer for each ansatz are as follows: ASWAP, 414; RY, 2142; SwapRZ, 3955; RYRZ, 4851.

The metrics that we use to evaluate the resource "efficiency" of a particular ansatz are the number of variational parameters and the number of CNOT operations it contains. We summarize these resource requirements for ASWAP and for several of the *ad hoc* ansatze described above in Table I for four qubits. Generally, as expected, the number of variational parameters in the ASWAP ansatze can be small, despite being able to quickly and accurately prepare the correct ground state.

TABLE I. A comparison of ansatz resources on four qubits.

	N	S_z	Number of parameters	Number of CNOT gates
ASWAP	0, 4	None	0	0
	1, 3	None	3	9
	2	None	5	16
	2	0	3	3
RY(d)	None	None	$4(d+1)$	$3d$
RYRZ(d)	None	None	$8(d+1)$	$3d$
SwapRZ(d)	N	None	$10d+4$	$24d$

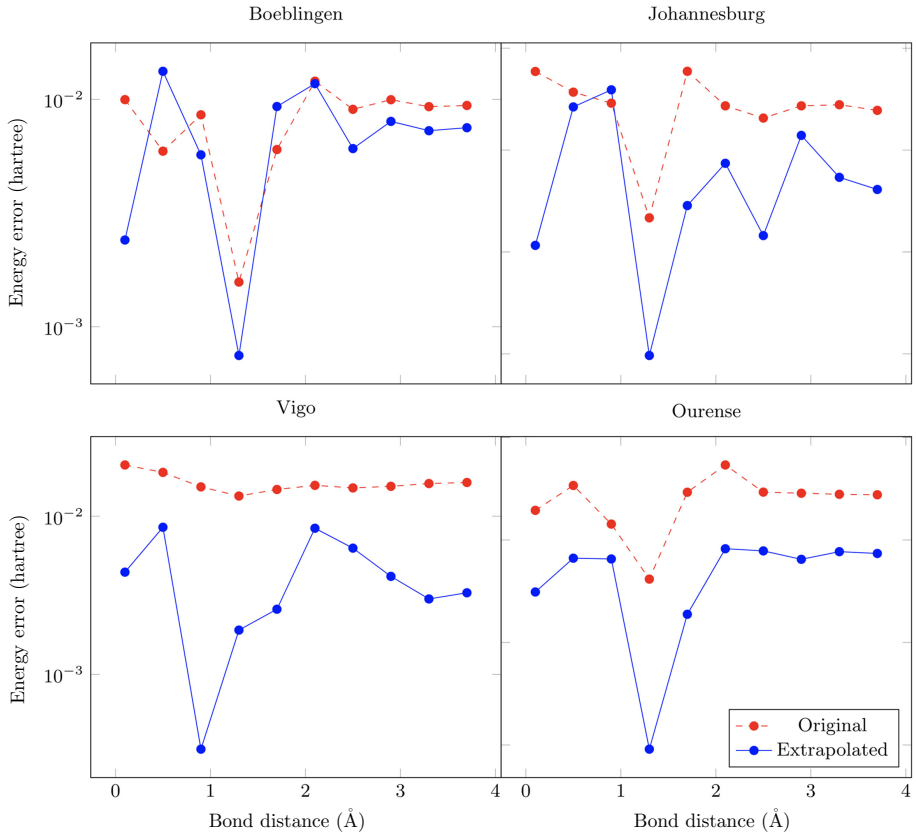

FIG. 3. A comparison of results with and without Richardson extrapolation for H_2 at varying interatomic distances for the ASWAP ansatz. Some points on the extrapolated curve are within chemical accuracy (approximately 1.5 millihartree). Richardson extrapolation provides as much as a 2-order-of-magnitude improvement over a standard VQE run.

Ansatze with more parameters require more optimization steps and hence calls to the objective function. Ansatze with more CNOT gates suffer from higher error rates and are more limited by decoherence. Here, the ASWAP ansatze enforce time-reversal symmetry. The combinations of symmetries for ASWAP are not exhaustive, as more are possible. Some ansatze are parametrized by the "depth" $d \geq 1$, which is the number of times the layer is repeated. Ansatze for which a symmetry is "None" means that the ansatz does not enforce this symmetry. The RY (RZ) ansatze apply CZ operations between nearest neighbors at each layer.

This comparison of resources illustrates the potential of enforcing symmetries at the level of the ansatz. Not only is the final prepared state guaranteed to have certain desired properties but enforcing this at the level of the ansatz allows for a reduced number of parameters as well as CNOT gates, compared to a more generic ansatz.

IV. EFFECTS OF NOISE ON PERFORMANCE OF VQE WITH VARIOUS ANSATZE

A. Hardware noise and noise model simulation

All of the ansatze used in this paper are applicable to any general-purpose quantum processor. Nevertheless, for the purposes of noisy simulation, we consider a noise model that is derived from information about the IBM Q processors. In particular, we include information from the devices about measurement errors, single- and two-qubit gate errors, depolarization, and thermal-relaxation errors that are dependent on individual gate times.

To establish a vocabulary, the different simulators in QISKIT that we use here are as follows. We refer to the noiseless simulator that keeps track of the state vector of the register under unitary operations as the "state vector simulator." We refer to the noisy simulator that includes the above information as the "QASM simulator" (quantum

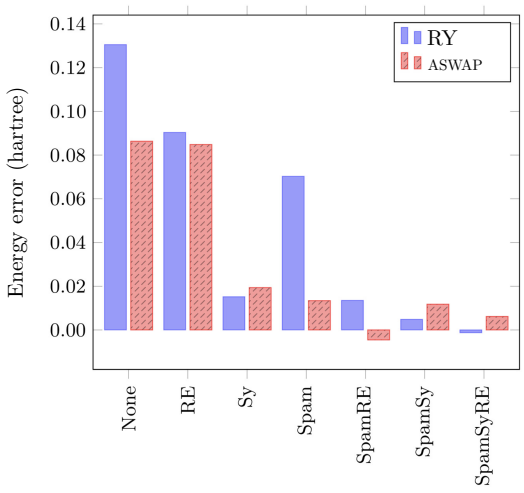

FIG. 4. The performance of two competing ansatze with different combinations of error-mitigation strategies for H_2 at equilibrium. Each of these strategies is introduced in Sec. IV B. The abbreviations for the strategies are RE (Richardson extrapolation), Sy (partial-symmetry enforcement), Spam (SPAM), or combinations of any of these. We note here that the ASWAP ansatz performs better even in the absence of any error-mitigation strategies. This underscores the fact that the reduction in quantum and classical resources due to imposing symmetries at the ansatz level leads to improvements in the energy error. Additionally, some values of the energy error are below zero, which violates the variational principle. This is due to our implementation of Richardson extrapolation, which can potentially undershoot the value of the ground-state energy. This is also potentially due to shot noise, since our simulations rely on a finite number of samples from the QPU.

assembly language) [52]. When using the QASM simulator, we use noise models constructed from information about the Vigo, Boeblingen, Ourense, and Johannesburg IBM Q devices.

B. Error-mitigation techniques

When running a VQE in the presence of noise, we use three error-mitigation techniques. The first of these addresses state-preparation-and-measurement (SPAM) errors by computing and inverting a matrix T of dimension $2^n \times 2^n$ the entries of which represent the probability of preparing one state and immediately measuring another. Inversion of this matrix allows one to perform error mitigation on population counts of a particular experiment, thereby accounting for SPAM errors. This and similar techniques have been previously explored [59–64]. Though there is not yet a fully scalable solution for mitigating SPAM errors, each of these techniques have applicability in different contexts for near-term devices. Here, we use the implementation in QISKIT.

The second technique, known as Richardson extrapolation [29,65,66], involves systematically increasing the amount of error in a given computation in a controlled way so that one can then extrapolate an objective function to a point at which those errors would be reduced. In this case, the extrapolation occurs by taking each CNOT gate in a particular circuit and inserting an additional even number of CNOT gates immediately after it. For instance, each CNOT gate in a given circuit would then become a sequence of an odd number of CNOT gates. In this way, the logic of the circuit is unaffected but in the presence of noise and faulty gates, more redundant CNOT gates will increase the error in the results. In Fig. 3, this procedure is shown for a variety of bond distances. In this case, we find that the extrapolated minimum energy obtained by the VQE is more accurate than the original results by as much as an order of magnitude in the energy error. Moreover, the extrapolated results are within a standard deviation of chemical accuracy for most of the interatomic distances.

The third error mitigation that we use is partial-symmetry enforcement during measurement. This noise mitigation relies on the fact that the true ground state has a known definite particle number and therefore the measurement results should also have this symmetry, even in the presence of noise. This mitigation is only partially possible, since the measurement process itself requires postrotation gates, which do not themselves conserve particle number. However, for the chemical Hamiltonians we are interested in, we can always group together Pauli strings that do preserve particle number. For instance, the H_2 Hamiltonian, after the JW mapping onto four qubits, can be divided into five sets of commuting Pauli strings; one of these sets is made up of only Pauli strings that commute with the total-particle-number operator (terms such as $IIIZ$, $ZZII$, $IZII$, etc.). During the measurement of this set, we can enforce that the particle number must match that of the ground state, discarding results that do not respect this symmetry. We also note that this procedure is not limited to the JW mapping. Under the BK mapping, for instance, the $a_i^\dagger a_i$ and $a_i^\dagger a_j^\dagger a_j a_i$ terms in a molecular Hamiltonian map to operators that are diagonal in the Pauli basis [67], as does the total-particle-number operator. Hence, this procedure extends to other mappings; however, we cannot presently compare their performance, since the ASWAP ansatz currently only applies to the JW mapping.

This process simply postselects data that we know to positively contribute to our goal of finding the ground state. The measurements of Pauli strings that do not commute with the total-particle-number operator are not modified. This technique may be more applicable when the state produced on the QPU during the VQE produces expectation values that are close to the total particle number, which is the case with the ASWAP ansatze. Nevertheless,

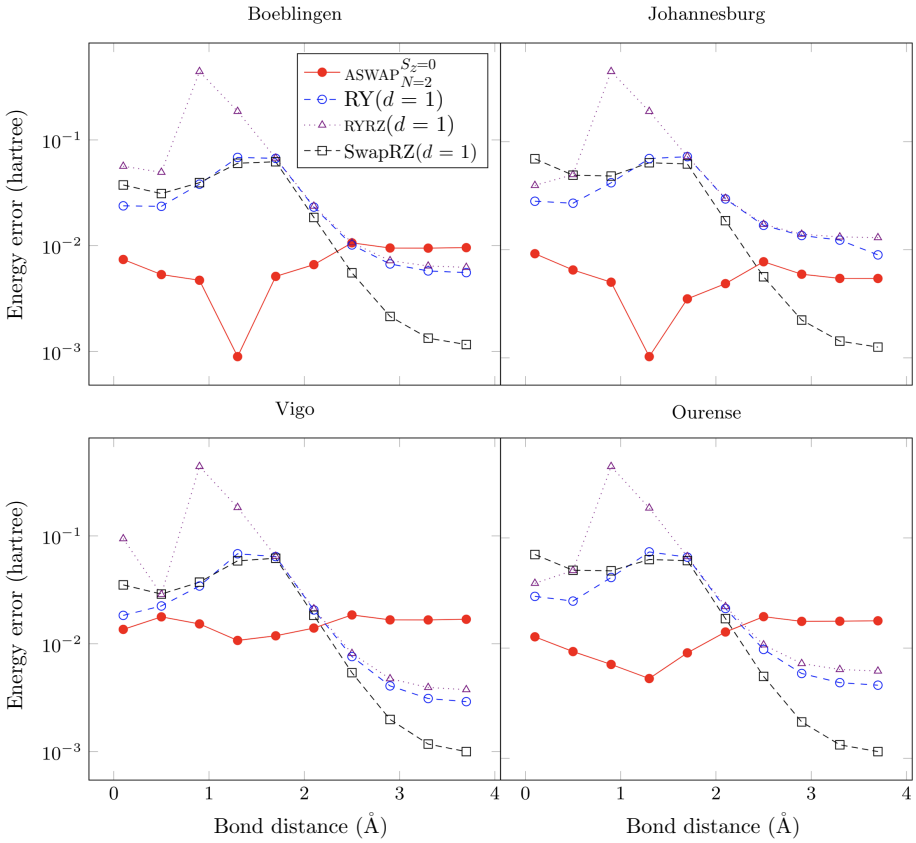

FIG. 5. A comparison of the symmetry-preserving ASWAP ansatz with several of the best competing *ad hoc* ansatze, for a variety of noise models corresponding to real devices. In all cases, the ASWAP ansatz performs better at smaller bond distances. This is because for large bond distances, the *ad hoc* ansatze find a degenerate separable ground state, whereas the ASWAP ansatz consistently finds the nontrivial ground state for all bond distances. All results in this figure have SPAM error mitigation applied.

this may be applied to any ansatz, as we later demonstrate. One potential drawback of this approach is the overhead introduced in the number of shots needed to accurately take expectation values. In discarding shots, it may be necessary to perform additional measurements to compensate for the reduced shots after postselection. In this case, as we later discuss, we are limited in the number of shots by the IBM Q devices.

We employ all of these strategies in various combinations in a noisy H_2 simulation, using both the RY and ASWAP ansatze. The results are shown in Fig. 4. We find that the largest improvement in performance for the RY and ASWAP ansatze come from partial-symmetry enforcement and SPAM strategies, respectively. In the case of RY, the large improvement under partial-symmetry enforcement is likely due to the fact that RY can prepare states that do not obey the desired symmetry but partial-symmetry enforcement corrects this by classical postprocessing. In the case of ASWAP, the large improvement from SPAM is likely due to the structure of the ansatz. Specifically, the ASWAP ansatz uses logic that assumes a fixed initial state of the register. Moreover, measurement errors can in principle violate the expected value of the particle number. Both of these effects can be suppressed by SPAM error mitigation.

C. Noisy optimization

Minimization of the objective function $f(\vec{\theta})$ requires an optimization algorithm that is resistant to noise. In previous works, variational quantum algorithms have been implemented using the simultaneous perturbation

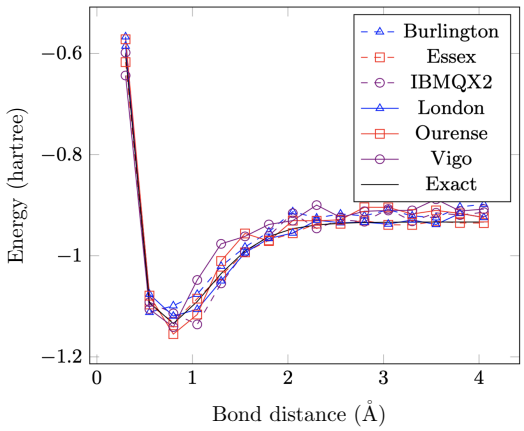

FIG. 6. A comparison of experimental VQE results for the dissociation of H_2 using IBM Q devices using the ASWAP ansatz with $S_z = 0$, $N = 2$, and time-reversal symmetry enforced. In each case, the shape of the results from hardware match the general shape of the exact result, even at longer bond distances. All points (except the exact results) have SPAM and Richardson-extrapolation error mitigation applied. This highlights the robustness of the ASWAP ansatz to a variety of noise levels in an experimental setting.

stochastic approximation (SPSA) [68], adaptive moment estimation (Adam) [69], constrained optimization by linear approximation (COBYLA) [70], and dividing rectangles (DIRECT) [71] algorithms. For this work, we choose the DIRECT and COBYLA algorithms for three primary reasons. First, the DIRECT algorithm has been demonstrated to be resistant to substantial amounts of noise in variational quantum contexts [31]. Second, DIRECT is a global-optimization algorithm that rarely gets trapped in local minima. In this sense, DIRECT sidesteps the need to consider carefully chosen (and potentially repeated) initial conditions for the variational parameters. Finally, though DIRECT does not converge quickly, it gets within the vicinity of the correct solution quickly. This trade-off is beneficial for the case of a noisy objective function, since high levels of precision in the objective function are limited by statistical and hardware errors. We also choose the COBYLA optimization algorithm for similar reasons, namely its ability to find an accurate estimate of the ground state in a short number of function calls and its resilience to noise. In previous experiments, it has been shown that the number of shots limits the resulting accuracy of the final energy estimate [5]. For this reason, we perform the maximum number of shots allowed by IBM devices (8192).

For our noisy simulations, the back end chosen is the QASM simulator, which uses a noise model constructed to mimic a variety of IBM Q processors. We use the DIRECT and COBYLA optimization algorithms. In all simulations using DIRECT, we limit the optimization to a budget of 100 or 300 calls to the objective function. For the COBYLA optimizer, on average the optimizer converges with about 400 calls to the objective function.

D. Results for noisy VQE simulations

In Fig. 5 we show the results of running the VQE for a variety of bond distances of H_2 for the best-performing symmetry-preserving and *ad hoc* ansatze. We find that for smaller bond distances, the ASWAP ansatz performs better in terms of finding the correct ground-state energy than the *ad hoc* ansatze. At larger bond distances, it appears that the *ad hoc* ansatze perform better; however, as we show in the next section, this is actually not the case. The reason is because in this regime, the *ad hoc* ansatze are actually converging to an excited state that is nearly degenerate with the true ground state. All the ansatze have greater difficulty in identifying the true ground state at large bond distance because correlations are stronger here (the true ground state is a singlet). On the other hand, the excited state found by the *ad hoc* ansatze is a (triplet) product state, which is easy for the *ad hoc* algorithms to identify. Because ASWAP has the correct symmetry quantum numbers built in, it does not misidentify the ground state and the error at large bond distance is larger because of the correlations in this state.

E. Experimental simulation of H_2

To more rigorously demonstrate the effectiveness of these techniques in the presence of noise, we use the physical IBM superconducting qubits to validate the results. In Fig. 6 we show the dissociation of H_2 using several different devices. In each case, the resulting energy shows good agreement with the exact result and the shape of the dissociation curve is generally preserved.

V. EFFECT OF NOISE ON SYMMETRY PRESERVATION

In addition to reducing computational resources and error levels, the ASWAP ansatz also correctly preserves the desired symmetry throughout the entire dissociation curve. This is illustrated in Fig. 7 for the case of H_2. While all ansatze considered remain reasonably close to the correct particle number, $\langle N \rangle = 2$, only ASWAP succeeds in finding the right value of $\langle S_z \rangle$ at all bond distances. In contrast, the RY ansatze find another, nearly degenerate, solution that has an incorrect value of S_z. This is a very important advantage of our ansatze when considering quantum chemistry problems, because with other ansatze there is no way to control which solution is found when there are approximate degeneracies aside from adding penalty terms to the objective function to find different symmetry states. Even

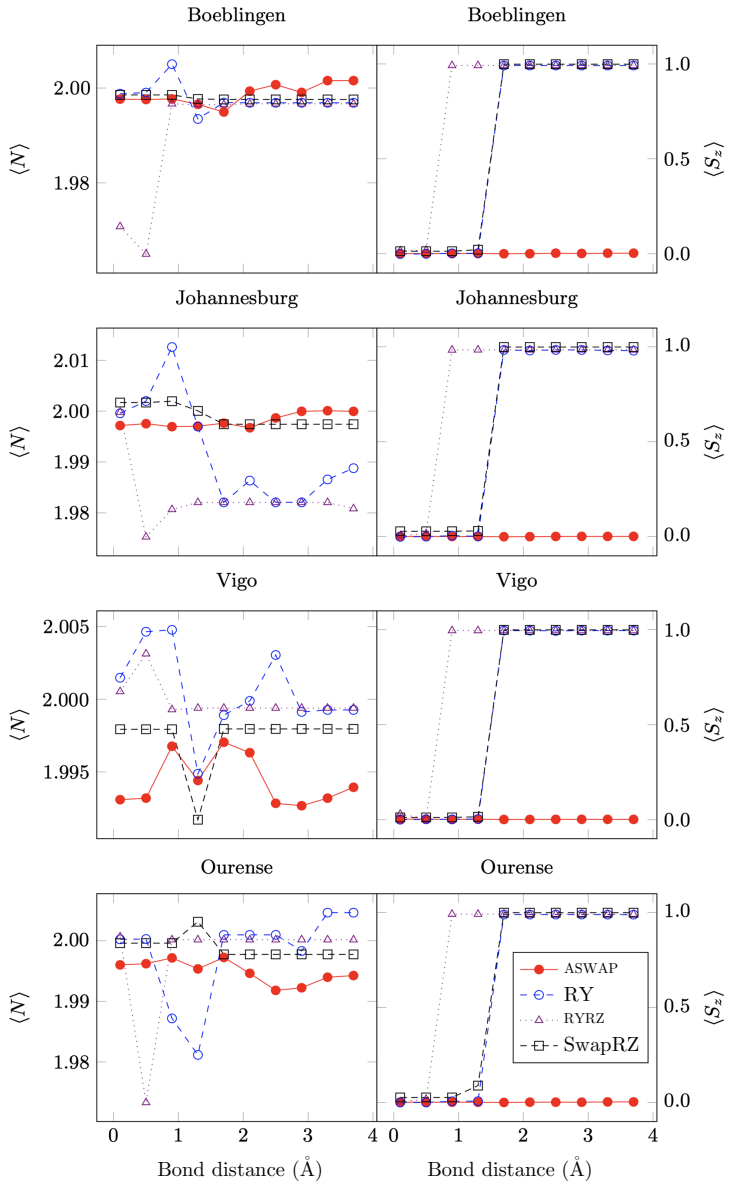

FIG. 7. Comparisons of preserved symmetries ($\langle N\rangle, \langle S_z\rangle$) in the ASWAP, RY, and RYRZ ansatze. The ASWAP ansatz consistently preserves each desired symmetry throughout the dissociation and any deviation, in this case, is due to the noise model. On the other hand, deviations in the RY and RYRZ cases can also be attributed to the inability of the ansatz to target a particular subspace. This is most prominently shown in the case of the $\langle S_z\rangle$ operator (right column). For the RY ansatze, the change in the resulting expectation values around 1 Å is due to the algorithm finding another, nearly degenerate, solution for the ground state that has an incorrect value of $\langle S_z\rangle$. The fact that the *ad hoc* ansatze produce values of $\langle S_z\rangle = +1$ and not -1 is dependent on both the optimization algorithm used, as well as the depth and connectivity used in the ansatz. Results using the DIRECT algorithm (shown here) consistently produce $\langle S_z\rangle = +1$, whereas COBYLA occasionally finds solutions with $\langle S_z\rangle = -1$. When using COBYLA, the value of $\langle S_z\rangle$ depends on the depth and connectivity of the ansatz used. This underscores the importance of using an ansatz that preserves known symmetries, so that one can guarantee consistent results. All results in this figure have SPAM error mitigation applied.

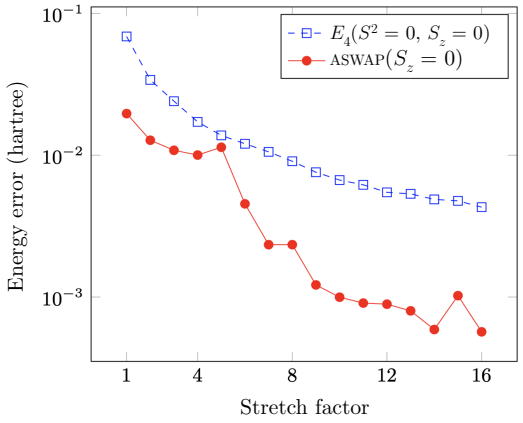

FIG. 8. The performance of the E_4 and ASWAP ansatze for H_2 at equilibrium for a range of stretch factors. The stretch factors here are the ratios of the original and simulated T_1 and T_2 times for the noise model. A stretch factor of 1 corresponds to the original noise model. Here, we choose the noise model for the IBM Q Vigo device. We find that improving the T_1 and T_2 times of the device by a factor of roughly 4 makes the E_4 ansatz perform as well as the ASWAP ansatz does today. Since the only error-mitigation technique used here is SPAM error mitigation, we find that roughly an order-of-magnitude improvement of the T_1 and T_2 times will allow the ASWAP ansatz to reach chemical accuracy without Richardson extrapolation.

with that approach, it is not clear how one would guarantee consistent eigenvalues for the symmetries throughout the whole curve with another ansatz. The ASWAP ansatz allows us to specify exactly which of these solutions we want to find.

A. Total spin ansatz performance

So far, this work has only focused on the ASWAP ansatz from Ref. [22]. That work also presented another ansatz, referred to as the E_n ansatze, where n is the number of qubits. Unlike any of the other ansatze in the present work, it has the ability to preserve the total spin of the trial state. The E_n ansatz is constructed by first writing out a general state vector in the appropriate total spin subspace, with the coefficients parametrized by hyperspherical coordinates. Then a unitary that produces these general states of fixed total spin starting from $|0\rangle^{\otimes n}$ is constructed and decomposed in terms of Toffoli gates using Gray codes [72]. The resulting ansatze can then be used to prepare arbitrary states with the appropriate spin projection, total spin, and particle-number symmetries. The details of this construction can be found in Ref. [22].

The E_n ansatze have larger gate depths and they are generally more challenging to implement on noisy hardware. In Fig. 8 we compare the performance of the E_4 and ASWAP ansatze for simulating H_2 at equilibrium. To accomplish this, we take the noise model from the IBM Q Vigo device and uniformly stretch the T_1 and T_2 times for all of the qubits, so as to simulate a device with less noise. T_1 is the energy-relaxation time, associated with decay from the $|1\rangle$ to the $|0\rangle$ state. T_2 is the decoherence time, associated with the loss of the phase of the quantum state. From these simulations, we confirm that the performance of the E_4 gate is limited by the $T_{1/2}$ times. Moreover, we find that these times would need to be approximately 4 times longer than they currently are to perform as well as the ASWAP ansatz with current $T_{1/2}$ times. Hence, we expect that the E_4 ansatz will be useful for preserving total spin, but is not yet currently viable. Nevertheless, there may be more complex molecules for which the ability to preserve S^2 will have a greater impact

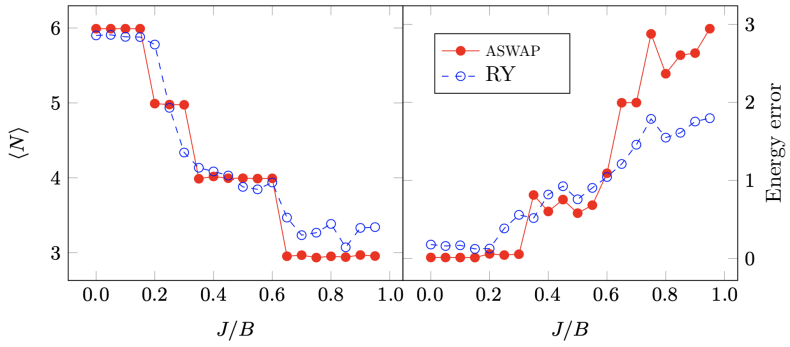

FIG. 9. A comparison of VQE results using ASWAP and RY ansatze for the Heisenberg model on a 2×3 lattice. Since the particle number cannot be fixed to a constant value with the RY ansatz, it only sometimes produces the correct result. On the other hand, by fixing the particle number in each case to the correct value, the ASWAP ansatz is able to obtain better results. To improve the RY results, we use three layers of entanglers and rotations. Here, we use the IBM Q Vigo noise model and include SPAM error mitigation.

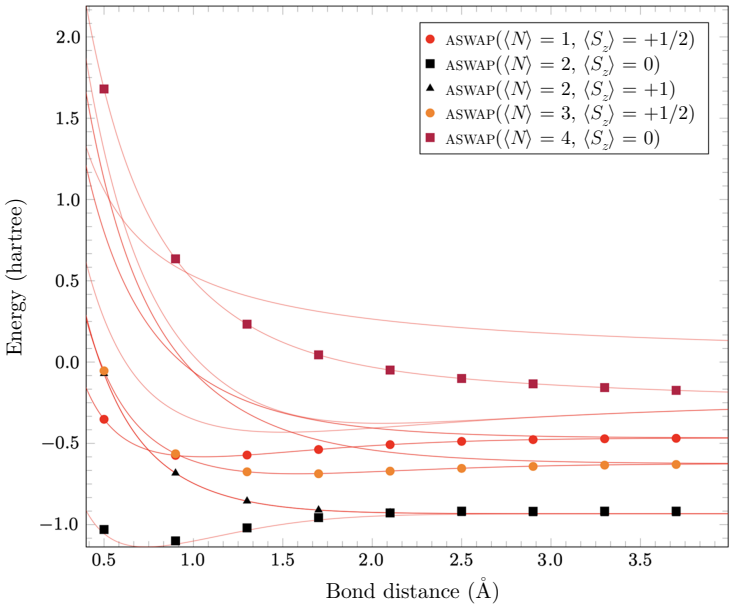

FIG. 10. The dissociation curve for a variety of excited states for H_2 calculated using the ASWAP ansatz. Dotted points correspond to results obtained using simulations of the IBM Q Vigo quantum processor, where different colors correspond to different symmetries enforced by the ASWAP ansatz (here, the particle number $\langle N\rangle$ and spin projection $\langle S_z\rangle$). The lines correspond to the different excited states as a function of the bond distance calculated using exact diagonalization, where different colors correspond to different excited states. In all cases, we find strong agreement between the results from noisy simulation and exact diagonalization. All results in this figure have SPAM error mitigation applied.

on the results and offset the challenges of the larger circuit depth.

B. Application to Heisenberg model

So far, we have only focused on molecular systems but now we turn our attention to the antiferromagnetic Heisenberg model (with $J > 0$). The Hamiltonian is given by

$$H = J \sum_{\{i,j\}\in E(L)} \left(X_iX_j + Y_iY_j + Z_iZ_j\right) + B \sum_{i\in V(L)} Z_i, \quad (4)$$

where J is the interaction strength, B is the magnetic field strength, and $E(L)$, $V(L)$ are the edge and vertex sets of the lattice L, respectively. Here, we choose the 2×3 lattice. For this system, the ground state has a well-defined particle number that varies with different values of J/B. Here, the particle number is defined as $\sum_{i\in V(L)}(1 - \langle Z_i\rangle)/2$. We choose this definition of the particle number since (in the computational basis) its diagonal entries correspond to the Hamming weight of the corresponding basis state. This is more amenable to the application of ASWAP operations. In Fig. 9, we show that again the ASWAP ansatz is able to successfully find the ground state with the correct particle number, whereas the RY ansatz produces results that deviate from the correct particle number. Moreover, in some regions, the ASWAP ansatz produces better results in terms of the energy error. These results show that if the symmetry of the ground state is known *a priori*, then enforcing it at the level of the ansatz can lead to a reduction in resource requirements and error rates. It is also important to note that even if the symmetry of the ground state is not known beforehand, one can still try multiple ansatz symmetries in independent VQEs run on different QPUs in parallel. It should also be noted that the *ad hoc* ansatze can explore different symmetries during a single VQE implementation run on one QPU. This potentially makes such ansatze appealing if nothing about the symmetries is known and access to multiple QPUs is limited.

VI. EXCITED STATES

An important aspect of our symmetry-preserving ansatze is the ability to find excited states by targeting subspaces of the Hilbert space that have the appropriate symmetry, provided that this symmetry differs from that of the ground state. Other methods of calculating excited states using VQEs exist but they require either iteratively running multiple VQEs to determine excited states [32] or executing additional circuits [6,33]. Nevertheless, each of these approaches has its own advantages. For example, Ref. [33] does not require multiple optimizations and is robust to noise. Hence, in terms of optimization, no additional effort is required to obtain excited states.

A notable feature of our method is that the search for excited states does not depend on having already obtained the ground state. In fact, we could compute both ground and excited states simultaneously by running separate VQEs in parallel or skip the ground state if we are only interested in excited states. Since our method amounts to

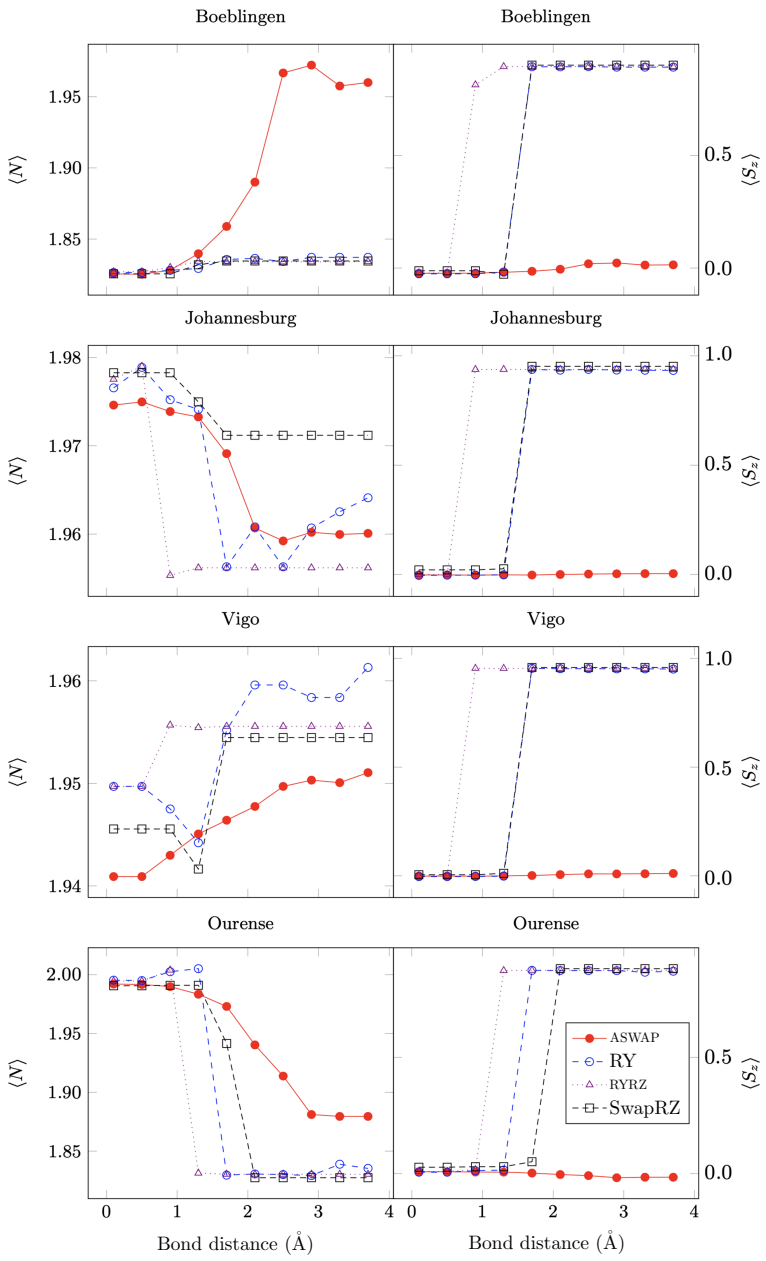

FIG. 11. Identical to Fig. 7, except with no SPAM error mitigation. Comparing with Fig. 7, SPAM error mitigation seems to consistently help preserve $\langle N \rangle$ for all ansatze.

simply changing the ansatze for different excited states, finding an excited state is as easy as finding the ground state. One possible disadvantage of our approach is that depending on the system chosen, it may not be the case that one can uniquely specify all excited states by changing symmetries alone. Therefore, the optimal approach may be a combination of our techniques and some of the previously introduced ones.

A demonstration of excited-state calculations using our ASWAP ansatz is shown for the H_2 molecule for several excited states in Fig. 10. Note that some of the exact results do not have corresponding simulation results. This

is because our method does not, in general, allow one to produce all of the higher excited states, only those directly accessible by partitioning the Hilbert space according to symmetries. In all cases considered, we are able to prepare five excited states (in addition to the ground state) out of the possible ten (not counting degeneracies). The error rates here are sufficiently low so that one may correctly order the eigenstates according to their energies.

VII. OUTLOOK

The key development of Ref. [22] is that the ASWAP tiling procedure produces ansatze that can be used to prepare any state in the given symmetry subspace. However, the tiling procedure used has the requirement of using nearest-neighbor connectivity. This raises the possibility of other connectivities having this same property, namely that they span symmetry subspaces. Though beyond the scope of the current work, future implementations of ASWAP ansatze could be hardware tailored by tiling A gates according to the connectivity of the device. Such a result would alleviate the need to insert additional SWAP operations to account for non-nearest-neighbor operations. In this case, one would need to check that the ansatz is able to capture all states in the symmetry subspace. The resulting ansatz could be considered a hybrid between hardware-efficient (taking into account the device connectivity) and symmetry-preserving (only utilizing A gates) ansatze. Such considerations may become more advantageous as device architectures become more complex.

VIII. CONCLUSIONS

In this work, we demonstrate that the deleterious effects of noise on a VQE can be mitigated by exploiting the symmetries of the Hamiltonian being simulated. We achieve this by performing simulations of the noiseless and noisy systems for a variety of ansatze, utilizing the most current error-mitigation techniques. We also test our methods on real IBM devices. The results of these simulations and experiments indicate that, with or without noise, the use of state-preparation circuits that preserve the relevant symmetries of the problem reduces computational resources and error rates compared to more *ad hoc* approaches. In addition, we show that the built-in symmetry preservation of our ansatze allows us to find excited states without having to first find the ground state and without using longer state-preparation circuits.

ACKNOWLEDGMENTS

S.E.E. acknowledges support from the U.S. Department of Energy (Award No. DE-SC0019318). E.B. and N.J.M. acknowledge support from the U.S. Department of Energy (Award No. DE-SC0019199). This research used quantum computing system resources supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research program office. The Oak Ridge National Laboratory manages access to the IBM Q System as part of the IBM Q Network.

APPENDIX: OPERATORS WITHOUT SPAM MITIGATION

In Fig. 11, we show the same information as in Fig. 7, except without using SPAM error mitigation. In terms of preserving symmetries, SPAM error mitigation helps to reduce the errors in estimating $\langle N \rangle$ that are due to state preparation and measurement. Generally, we see that without SPAM error mitigation, $\langle N \rangle$ is underestimated. This is likely due to T_1 relaxation during state preparation and measurement.

[1] R. P. Feynman, Simulating physics with computers, Int. J. Theor. Phys. **21**, 467 (1982).
[2] 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 (2014).
[3] P. J. J. O'Malley, *et al.*, Scalable Quantum Simulation of Molecular Energies, Phys. Rev. X **6**, 031007 (2016).
[4] 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).
[5] 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).
[6] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, Phys. Rev. X **8**, 011021 (2018).
[7] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
[8] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, ArXiv:1411.4028 (2014).
[9] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, and C. Baldwin, *et al.*, Quantum approximate optimization with a trapped-ion quantum simulator, ArXiv:1906.02700 (2019).
[10] J. Romero and A. Aspuru-Guzik, Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions, ArXiv:1901.00848 (2019).
[11] V. Leyton-Ortega, A. Perdomo-Ortiz, and O. Perdomo, Robust implementation of generative modeling with parametrized quantum circuits, ArXiv:1901.08047 (2019).
[12] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A.

Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, Training of quantum circuits on a hybrid quantum computer, Sci. Adv. **5**, eaaw9918 (2019).
[13] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, npj Quantum Inf. **5**, 1 (2019).
[14] G. Verdon, M. Broughton, J. R. McClean, K. J. Sung, R. Babbush, Z. Jiang, H. Neven, and M. Mohseni, Learning to learn with quantum neural networks via classical neural networks, ArXiv:1907.05415 (2019).
[15] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. Kin-Lic Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, Phys. Rev. Lett. **120**, 110501 (2018).
[16] T. Jones, S. Endo, S. McArdle, X. Yuan, and S. C. Benjamin, Variational quantum algorithms for discovering Hamiltonian spectra, Phys. Rev. A **99**, 062304 (2019).
[17] K. Mitarai, T. Yan, and K. Fujii, Generalization of the Output of a Variational Quantum Eigensolver by Parameter Interpolation with a Low-Depth Ansatz, Phys. Rev. Appl. **11**, 044087 (2019).
[18] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martínez, Quantum Computation of Electronic Transitions Using a Variational Quantum Eigensolver, Phys. Rev. Lett. **122**, 230401 (2019).
[19] E. Anschuetz, J. Olson, A. Aspuru-Guzik, and Y. Cao, in *International Workshop on Quantum Technology and Optimization Problems* (Springer, Munich, Germany, 2019), p. 74.
[20] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Müller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, and K. Temme, Quantum optimization using variational algorithms on near-term quantum devices, Quantum Sci. Technol. **3**, 030503 (2018).
[21] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions, Phys. Rev. A **98**, 022322 (2018).
[22] B. T. Gard, L. H. 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**, 10 (2020).
[23] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, Adv. Quantum Technol. **2**, 1900070 (2019).
[24] P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, Improving variational quantum optimization using cvar, ArXiv:1907.04769 (2019).
[25] A. Robert, P. K. Barkoutsos, S. Woerner, and I. Tavernelli, Resource-efficient quantum algorithm for protein folding, ArXiv:1908.02163 (2019).
[26] S. Endo, Q. Zhao, Y. Li, S. Benjamin, and X. Yuan, Mitigating algorithmic errors in a Hamiltonian simulation, Phys. Rev. A **99**, 012334 (2019).
[27] Y. S. Yordanov and C. H. W. Barnes, Efficient quantum circuits for quantum computational chemistry, arXiv:2005.14475 (2020).
[28] R. LaRose, A. Tikku, E. O'Neel-Judy, L. Cincio, and P. J. Coles, Variational quantum state diagonalization, npj Quantum Inf. **5**, 1 (2019).
[29] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, Cloud Quantum Computing of an Atomic Nucleus, Phys. Rev. Lett. **120**, 210501 (2018).
[30] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, and M. J. Savage, Quantum-classical computation of Schwinger model dynamics using quantum computers, Phys. Rev. A **98**, 032331 (2018).
[31] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, *et al.*, Self-verifying variational quantum simulation of lattice models, Nature **569**, 355 (2019).
[32] O. Higgott, D. Wang, and S. Brierley, Variational quantum computation of excited states, Quantum **3**, 156 (2019).
[33] P. J. Ollitrault, A. Kandala, C.-F. Chen, P. K. Barkoutsos, A. Mezzacapo, M. Pistoia, S. Sheldon, S. Woerner, J. Gambetta, and I. Tavernelli, Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor, arXiv:1910.12890 (2019).
[34] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, arXiv:1908.04416 (2019).
[35] M. Kühn, S. Zanker, P. Deglmann, M. Marthaler, and H. Weiß, Accuracy and resource estimations for quantum chemistry on a near-term quantum computer, J. Chem. Theory Comput. **15**, 4764 (2019).
[36] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).
[37] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, Quantum **3**, 214 (2019).
[38] R. M. Parrish, J. T. Iosue, A. Ozaeta, and P. L. McMahon, A Jacobi diagonalization and Anderson acceleration algorithm for variational quantum algorithm parameter optimization, ArXiv:1904.03206 (2019).
[39] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martinez, Hybrid quantum/classical derivative theory: Analytical gradients and excited-state dynamics for the multistate contracted variational quantum eigensolver, ArXiv:1906.08728 (2019).
[40] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, ArXiv:1907.03358 (2019).
[41] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskyi, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, J. Chem. Theory Comput. **16**, 190 (2020).
[42] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families, ArXiv:1907.13623 (2019).

[43] M. Wilson, S. Stromswold, F. Wudarski, S. Hadfield, N. M. Tubman, and E. Rieffel, Optimizing quantum heuristics with meta-learning, ArXiv:1908.03185 (2019).
[44] X. Bonet-Monroig, R. Babbush, and T. E. O'Brien, Nearly optimal measurement scheduling for partial tomography of quantum states, ArXiv:1908.05628 (2019).
[45] A. F. Izmaylov, T.-C. Yen, and I. G. Ryabinkin, Revising the measurement process in the variational quantum eigensolver: Is it possible to reduce the number of separately measured operators? Chem. Sci. **10**, 3746 (2019).
[46] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, Low-cost error mitigation by symmetry verification, Phys. Rev. A **98**, 062339 (2018).
[47] S. McArdle, X. Yuan, and S. Benjamin, Error-Mitigated Digital Quantum Simulation, Phys. Rev. Lett. **122**, 180501 (2019).
[48] 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).
[49] T. Takeshita, N. C. Rubin, Z. Jiang, E. Lee, R. Babbush, and J. R. McClean, Increasing the Representation Accuracy of Quantum Simulations of Chemistry without Extra Quantum Resources, Phys. Rev. X **10**, 011004 (2020).
[50] 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).
[51] H. Lun 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 ansatze on a quantum processor, ArXiv:1911.10205 (2019).
[52] QISKIT: An open-source framework for quantum computing (2019), https://github.com/qiskit/qiskit.
[53] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, Z. Phys. **47**, 631 (1928).
[54] J. T. Seeley, M. J. Richard, and P. J. Love, The Bravyi-Kitaev transformation for quantum computation of electronic structure, J. Chem. Phys. **137**, 224109 (2012).
[55] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, Ann. Phys. (NY) **298**, 210 (2002).
[56] K. Setia, S. Bravyi, A. Mezzacapo, and J. D. Whitfield, Superfast encodings for fermionic quantum simulation, Phys. Rev. Res. **1**, 033033 (2019).
[57] N. Moll, A. Fuhrer, P. Staar, and I. Tavernelli, Optimizing qubit resources for quantum chemistry simulations in second quantization on a quantum computer, J. Phys. A: Math. Theor. **49**, 295301 (2016).
[58] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp, Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer, Phys. Rev. Appl. **11**, 044092 (2019).
[59] M. R. Geller and M. Sun, Efficient correction of multiqubit measurement errors, ArXiv:2001.09980 (2020).
[60] K. E. Hamilton, T. Kharazi, T. Morris, A. J. McCaskey, R. S. Bennink, and R. C. Pooser, Scalable quantum processor noise characterization, ArXiv:2006.01805 (2020).
[61] S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay, and J. M. Gambetta, Mitigating measurement errors in multi-qubit experiments, ArXiv:2006.14044 (2020).
[62] K. E. Hamilton and R. C. Pooser, Error-mitigated data-driven circuit learning on noisy quantum hardware, Quantum Mach. Intell. **2**, 1 (2020).
[63] M. R. Geller, Rigorous measurement error correction, Quantum Sci. Technol. **5**, 03LT01 (2020).
[64] G. S. Barron and C. J. Wood, Measurement error mitigation for variational quantum algorithms, ArXiv:2010.08520 (2020).
[65] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, Phys. Rev. Lett. **119**, 180509 (2017).
[66] 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).
[67] J. T. Seeley, M. J. Richard, and P. J. Love, The Bravyi-Kitaev transformation for quantum computation of electronic structure, J. Chem. Phys. **137**, 224109 (2012).
[68] J. C. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, IEEE Trans. Automat. Contr. **37**, 332 (1992).
[69] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980 (2014).
[70] M. J. D. Powell, *Advances in Optimization and Numerical Analysis*, edited by Susana Gomez and Jean-Pierre Hennart (Springer Netherlands, Dordrecht, 1994), p. 51.
[71] D. R. Jones, C. D. Perttunen, and B. E. Stuckman, Lipschitzian optimization without the Lipschitz constant, J. Optim. Theory Appl. **79**, 157 (1993).
[72] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, 2011).

Chapter 9

Benchmarking Quantum Chemistry Computations with Variational, Imaginary Time Evolution, and Krylov Space Solver Algorithms

This chapter includes pages from our publication [7]

K. Yeter-Aydeniz, B. T. Gard, J. Jakowski, S. Majumder, G. S. Barron, G. Siopsis, T. S. Humble, and R. C. Pooser, "Benchmarking quantum chemistry computations with variational, imaginary time evolution, and krylov space solver algorithms", Advanced Quantum Technologies **4**, 2100012 (2021)

This work is copyright 2021 Wiley-VCH GmbH.

The contributions to this work are as follows: My contributions were assisting in the testing of the symmetry preserving circuits, assisting in writing of the portions relevant to symmetry preserving circuits. K.Y.A. primarily carried out this work. B.T.G. carried out numerical simulations in relation to symmetry preserving circuits. J.J. assisted with theoretical work and simulations. S.M. assisted with work relating to hidden inverses and preparation of the manuscript. G.S., T.S.H. and R.C.P. supervised.

ADVANCED QUANTUM TECHNOLOGIES
www.advquantumtech.com

Benchmarking Quantum Chemistry Computations with Variational, Imaginary Time Evolution, and Krylov Space Solver Algorithms

*Kübra Yeter-Aydeniz, Bryan T. Gard, Jacek Jakowski, Swarnadeep Majumder, George S. Barron, George Siopsis, Travis S. Humble, and Raphael C. Pooser**

Quantum chemistry is a key application area for noisy-intermediate scale quantum (NISQ) devices, and therefore serves as an important benchmark for current and future quantum computer performance. Previous benchmarks in this field have focused on variational methods for computing ground and excited states of various molecules, including a benchmarking suite focused on the performance of computing ground states for alkali-hydrides under an array of error mitigation methods. State-of-the-art methods to reach chemical accuracy in hybrid quantum-classical electronic structure calculations of alkali hydride molecules on NISQ devices from IBM are outlined here. It is demonstrated how to extend the reach of variational eigensolvers with symmetry preserving Ansätze. Next, it is outlined how to use quantum imaginary time evolution and Lanczos as a complementary method to variational techniques, highlighting the advantages of each approach. Finally, a new error mitigation method is demonstrated which uses systematic error cancellation via hidden inverse gate constructions, improving the performance of typical variational algorithms. These results show that electronic structure calculations have advanced rapidly, to routine chemical accuracy for simple molecules, from their inception on quantum computers a few short years ago, and they point to further rapid progress to larger molecules as the power of NISQ devices grows.

1. Introduction

Electronic structure calculation in quantum chemical systems is one of the most important applications of noisy intermediate scale quantum (NISQ) computers. Because of its important role as a key application, it has recently filled the gap in application-level benchmarks of quantum computer performance. Computational chemistry has long played a major role in benchmarking classical computers, starting with interpretation of molecular spectra against quantum chemistry calculations. Its importance has been recognized in Gerhard Hertzberg's 1971 Nobel lecture awarded "for his contributions to the knowledge of electronic structure and geometry of molecules". Today, a wide range recent electronic structure calculations using variational methods in the literature represent a promising paradigm within which to evaluate NISQ computers as well.[1–6] Recently these ideas were combined into a benchmark suite that aimed to enumerate the potential of cloud NISQ devices under an array of configurations including active orbitals, Ansatz construction, and error mitigation methods.[7] Further results in preparing and measuring quantum states for simple molecules indicate a growing capability that may eventually demonstrate quantum computational advantage,[8] which makes tracking and validating performance of NISQ devices in this field all the more important.

Notably, recent methods for benchmarking, including quantum volume benchmarking,[9] cross-entropy benchmarking,[10]

Dr. K. Yeter-Aydeniz, Dr. J. Jakowski, Dr. R. C. Pooser
Computational Sciences and Engineering Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA
E-mail: pooserrc@ornl.gov
Dr. K. Yeter-Aydeniz
Physics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA

The ORCID identification number(s) for the author(s) of this article can be found under https://doi.org/10.1002/qute.202100012

DOI: 10.1002/qute.202100012

Dr. B. T. Gard, G. S. Barron
Department of Physics
Virginia Tech
Blacksburg, VA 24061, USA
S. Majumder
Department of Electrical and Computer Engineering
Duke University
Durham, NC 27708, USA
Prof. G. Siopsis
Department of Physics and Astronomy
The University of Tennessee
Knoxville, TN 37996, USA
Dr. T. S. Humble
Quantum Science Center
Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA

 © 2021 Wiley-VCH GmbH

and randomized benchmarking,[11] provide good low-level detail about quantum gate performance, but they have lacked an application-specific context that relates to the role of computation as a tool for discovery. This gap is compounded by the fact that performance metrics such as single and two qubit gate fidelities are not always good predictors of higher level application performance. This gap further motivates the use of quantum chemistry as a key performance indicator in the NISQ era. Here, we present the current progress in developing this application area as a performance benchmark, and we outline the accuracy and performance that is achievable for electronic structure calculations on quantum computers today with the state-of-the-art algorithms and error mitigation techniques. We subsequently estimate where progress in this field can move to in the near term in light of these developments.

2. Quantum Chemistry Algorithms

One of the most well-known NISQ algorithms is the variational quantum eigensolver (VQE), which has proven to be especially versatile for testing and tuning applications of quantum chemistry on NISQ devices. This algorithm is part of the broader growth in variational quantum algorithms over the past decade.[12–18] In particular, the VQE method offered the first experimental example of an electronic structure calculation by quantum computing on quantum photonic hardware.[17] Upon its first implementation on superconducting hardware, the algorithm was shown to outperform time evolution via Trotterization and subsequent phase estimation.[19] The advantage of variational methods is that low-depth, parameterized circuits produce simple trial wavefunctions, minimizing errors due to decoherence, while classical optimization proved robust to the systematic and depolarization errors found in NISQ devices. Several implementations that followed focused on the utility of unitary coupled cluster methods and variations thereof.[7,20–23] Other implementations demonstrated that the trial Ansatz need not be in UCC form to obtain chemical accuracy.[3] Eventually, variational methods were modified to be able to compute excited states as well.[4]

The key limitation in time-evolution computation outlined in ref. [19] motivated a search other techniques to compute time-evolved quantities such as full eigenspectra. Quantum imaginary time evolution (QITE) algorithm to find the ground-state energy of many-particle systems is one such method.[18] In principle, QITE is capable of preparing exact quantum states without variational optimization. Direct computation of the eigenspectra without ancilla qubits using this algorithm was outlined in ref. [24], and the method was further extended to a practical application on NISQ devices in ref. [25]. The quantum Lanczos method was used in tandem with QITE in these demonstrations in order to recover higher excited states of molecules and nuclei, avoiding the problem of high circuit depth associated with Trotterization and real time evolution. Further demonstrations of the QITE algorithm applied to chemistry problems are included in refs. [26, 27]. These early examples were notably limited to two-qubit computations due to hardware noise. Details on the inner workings of these algorithms were included in a recent review.[28] Concurrently with the present work, a quantum algorithm called quantum solver of the anti-Hermitian contracted eigenvalue equation (QACSE) was demonstrated in ref. [29]. This quantum algorithm gives the energies and reduced density matrices of ground and excited states with low depth quantum circuits without a requirement for complex classical optimization. It also shares the advantages of short depth and noise resilience with VQE and QLanczos algorithms. Although not included in our original benchmark suite, future iterations will include the QACSE solver.

Here, we start from this current state of the field and produce new benchmark techniques that extend the performance of both variational methods and QITE/QLanczos algorithms beyond the state of the art in electronic structure calculations on NISQ devices. In particular, we apply and obtain experimental data for an Ansatz constructed with a symmetry preserving circuit (SPC) introduced in refs. [30–32]. This Ansatz preserves relevant chemical symmetries and has the benefit of minimal parameter count and low depth. Since the main limitation of coupled cluster methods is noise associated with large circuit depth, particularly in the number of CNOTs required, we find that VQE combined with SPCs vastly outperforms coupled cluster methods due to the commensurately lower noise associated with short depth. Further, since these circuits identify target symmetry spaces, they are also capable of probing select excited states for each unique combination of symmetry eigenvalues. We demonstrate this capability by directly calculating eigenspectra for a range of alkali-hydride molecules. Importantly, we show that VQE+SPC enables chemical accuracy in all of the test molecules in our suite over the cloud for the first time. We also compare and contrast the strengths of VQE+SPC with QITE/QLanczos and demonstrate the use of QITE on 4-qubit depths for the first time in chemistry problems. We make extensive use of error mitigation in our benchmarks, and here we also outline a new error mitigation technique via the use of hidden inverse gate constructions and demonstrate its use in variational algorithms. We find that this error mitigation technique substantially reduces the effects of coherent errors in simulation, further extending the reach of both variational and QITE algorithms by mitigating circuit noise.

The rest of this report proceeds as follows. In Section 3, we discuss the chemical Hamiltonians of the alkali hydride molecules LiH, NaH, KH, RbH which are used in this paper as a suite of model chemical systems for benchmarking. In Section 4, we discuss three hybrid quantum-classical algorithms that are used to calculate the eigenvalues of the molecules of interest on a quantum computer. Specifically, in Section 4.1, we discuss SPCs in VQE and present our data obtained using IBM's quantum computers. We then discuss QITE and QLanczos algorithms in Section 4.2 and present our experimental data collected from IBM's quantum hardware. In Section 4.3, we introduce and demonstrate the hidden inverse as an error mitigation strategy for variational algorithms and demonstrate its application on a noisy quantum computer simulator for LiH and NaH molecules. Finally, we conclude and present further predictions for the state of the field in Section 5.

3. Chemistry Model

In this section, we introduce the underlying model of chemistry used throughout this paper. We start by discussing electronic

© 2021 Wiley-VCH GmbH

structure of alkali hydrides molecules LiH, NaH, KH, RbH. First, we discuss the model of Hamiltonian in the second quantization and sketch the embedding of the quantum computation kernel in Hartree–Fock which separates small active set of orbitals from the remaining orbital treated with a mean-field approach. Our model of the quantum region includes four spin orbitals and requires four qubits when using the Jordan–Wigner transformation. Finally, we discuss eigenstates that correspond to our model Hamiltonian.

The test molecular systems used in this work are alkali hydride molecules LiH, NaH, KH, RbH which, respectively, have 4, 12, 20, 38 electrons. The experimentally measured bond lengths of gas phase LiH, NaH, and KH are, respectively, 1.595, 1.887, and 2.243 Å. For comparison, the theoretical calculation with frozen core CCSD(T)/6-311G(d,p) theory estimate the bond lengths for these species as equal to 1.595, 1.913, and 2.257 Å. The discrepancy between theoretical and experimentally measured bond lengths is attributed mainly to basis set superposition error (BSSE) and it is a manifestation of incompleteness of basis set used for electronic structure. In general, for the atomic basis sets localized on nuclei the BSSE depends on the interatomic distances and leads to shortening of the bond lengths and overestimation of binding energy for short distances. We also note that the proper description of the bond breaking requires a proper accounting for a multireference character and is not a subject of this study. In the current work, the interatomic distances range between 1.0 and 3.0 Å, covering the bond range and not the whole range up to the dissociation limit. We employ the minimal basis set as it is sufficient for the current studies. The minimal STO-3G basis set is used and the corresponding Hilbert space is spanned by, respectively, 6, 10, 14, 23 spherical atomic orbitals and twice as many molecular spin-orbitals. Even though we use only a minimal basis set, the size of the Hilbert space is too large for a direct computation of electronic structure on NISQ devices with the customary Jordan–Wigner transformation,[33] as it requires that each spin orbital is represented by an individual qubit. Alternative choices for this mapping exist.[34,35] Each technique has its own advantages and disadvantages. Most notably, the alternatives encode information non-locally in the qubits. This can come with the advantage of requiring fewer qubits, but could also make the results more dependent on correlated errors on the device. In this case, we choose the customary Jordan–Wigner transform for simplicity, and since the SPCs are currently tailored for this specific mapping. Though it is beyond the scope of this work, the relative performance of each of these mappings is a potentially interesting topic of study.

Fortuitously, the requirement for large number of qubits can be alleviated. In the majority of chemical processes, such as bond breaking or formation, only the highest energy, valence electrons play a significant (active) role, whereas the tightly bound core electrons are largely (inactive) spectators whose role is limited to a mean-field screening of the Coulomb electrostatic field from the nuclei. The resulting effective Hamiltonian in the second quantization language is given by ref. [7]:

$$H = H_0 + H_1 + H_2 \tag{1}$$

with H_0 given by

$$H_0 = E_{\text{nucl}} + \sum_a \left(h_a^a + \tfrac{1}{2} \sum_b \bar{g}_{ab}^{ab} \right) \tag{2}$$

where a and b run over inactive occupied spin-orbitals of frozen-core. The E_{nucl} describes Coulomb repulsion between bare nuclei cores and the second term describes the effect of screening by the core electrons. Similarly, the 1-body term is given by

$$H_1 = \sum_{p,q} \hat{p}^\dagger \hat{q} \cdot \left(h_q^p + \tfrac{1}{2} \sum_a \bar{g}_{aq}^{ap} \right) \tag{3}$$

where the term h_q^p represents interaction of valence electrons with all core ions and the second term in the parenthesis describes its screening by core electrons. Finally, the two-body part is

$$H_2 = \tfrac{1}{4} \sum_{p,q,r,s} \bar{g}_{sr}^{pq} \cdot \hat{p}^\dagger \hat{q}^\dagger \hat{r} \hat{s} \tag{4}$$

where the indices p, q, r, s run over active spin orbitals whereas indices a and b run over inactive spin-orbitals of a frozen core. The symbols h_q^p and $\bar{g}_{sr}^{pq}$ denote, respectively, matrix elements of core Hamiltonian (kinetic energy of electrons plus Coulomb interaction with core ions) and an anti-symmetrized repulsion integral:

$$\bar{g}_{sr}^{pq} = g_{sr}^{pq} - g_{r,s}^{pq} = \langle p, q | s, r \rangle - \langle p, q | r, s \rangle \tag{5}$$

This Hamiltonian provides a powerful recipe for embedding of the quantum electronic structure calculation into a classically computed environment obtained via Hartree–Fock (HF).[7,36]

For the model alkali hydride molecules studied here, only the two highest energy electrons are involved in chemical bonding. The selected active space includes only the highest occupied and lowest unoccupied molecular orbitals (HOMO and LUMO) and four molecular spin orbitals. The resulting Hilbert space spanned by the HOMO and LUMO is of O(2^4) size and can describe up to four electrons. The neutral case is described by placing two electrons within the active space. The Hamiltonian matrix has 16×16 elements and is block diagonal with each block corresponding to a different number electrons, with block size of 1, 4, 6, 4, and 1 for, respectively, 0, 1, 2, 3, and 4 electrons cases. In summary, the effective Hamiltonian applied allows for a significant reduction of quantum resources requiring only 4 qubits for the Jordan–Wigner transformed Hamiltonian. This construction is effectively the state of the art in quantum chemistry computations to date. While higher basis sets have been demonstrated, our benchmark basis set and subsequent qubit numbers are well-matched to today's hardware. Subsequent implementations of these techniques underway will use larger basis sets and enable benchmarking of larger numbers of qubits as the coherence time of the larger cloud devices increases.

For the 4-qubit problems we investigate here, there are 2^4 = 16 eigenstates which we can group into distinct symmetry subspaces. Denoting the number of electrons within active space as N_e and corresponding spin projection as s_z, one can write parametrized vectors associated with each Hamiltonian block as

© 2021 Wiley-VCH GmbH

labeled by their distinguishable quantum numbers in general form:

$$\begin{aligned}
|N_e = 1, s_z = 0.5\rangle &= \alpha_1|1000\rangle + \alpha_2|0100\rangle \\
|N_e = 2, s_z = 0\rangle &= \beta_1|1010\rangle + \beta_2|1001\rangle \\
&\quad +\beta_3|0110\rangle + \beta_4|0101\rangle \\
|N_e = 3, s_z = 0.5\rangle &= \gamma_1|1110\rangle + \gamma_2|1101\rangle
\end{aligned} \tag{6}$$

where we used spin block notation such that the first two bits in the ket occupation vector correspond to spin-up orbital whereas the last two bits in the ket correspond to spin-down orbitals. Next, we note that kets corresponding to $s_z = -0.5$ can be easily obtained by symmetry relations. The trivial states for $N_e = 0$ ($|0000\rangle$) and $N_e = 4$ ($|1111\rangle$) have been skipped. The second formula in Equation (6) for $N_e = 2$ parametrizes three singlet states ($s = 0$) and one triplet state ($s = 1$) mixed. For the $N_e = 2$ case, the singlet states can be written as

$$\begin{aligned}
|N_e = 2, s = 0, s_z = 0\rangle &= \beta_1'|1010\rangle + \beta_2'|0101\rangle \\
&\quad +\beta_3' \cdot \frac{1}{\sqrt{2}}(|1001\rangle + |0110\rangle)
\end{aligned} \tag{7}$$

whereas the triplet states are

$$\begin{aligned}
|N_e = 2, s = 1, s_z = 1\rangle &= |1100\rangle \\
|N_e = 2, s = 1, s_z = 0\rangle &= \frac{1}{\sqrt{2}}(|1001\rangle - |0110\rangle) \\
|N_e = 2, s = 1, s_z = -1\rangle &= |0011\rangle
\end{aligned} \tag{8}$$

Note that of these, only the one with $s_z = 0$ is an entangled state, the remaining states are separable. The ground state of our model chemistry is a singlet state model described by Equation (7).

4. Quantum Variational Electronic Structure Computations

In this section, we discuss application of various variational quantum algorithms toward the estimation of ground and excited state electronic structure for model chemical Hamiltonians discussed in Section 3. First we explain the symmetry preserving circuits in which trial state vectors are constructed to preserve spin projection during the variational search. As an illustration, these circuits are used to explore ground and excited states of our model systems on IBM Q hardware. Next, we discuss imaginary time evolution and quantum Lanczos eigensolvers and its application to alkali hydrides on IBM Q hardware. Finally, we illustrate the importance of noise mitigation and introduce the hidden inverse method as an approach to reduce noise in quantum computing.

4.1. Symmetry Preserving Circuits

In a standard VQE, parameterized Ansätze are chosen that, ideally, are accurate, have a small amount of parameters and a low enough circuit depth for NISQ hardware. Symmetry preserving circuits (SPC)[30–32] have very efficient implementation on NISQ hardware, which we demonstrate here on metal alkali hydrides, experimentally for the first time. Since all of the molecules we consider can be mapped using the Jordan–Wigner mapping onto 4 qubits, the methodology of constructing the SPC results in a fairly simple, compact set of circuits. We note that the qubit tapering methods using symmetry information can also be used to improve performance by reducing circuit width.[37] The SPC method is complementary in its focus on circuit depth. Moreover, SPC uses the minimum number of parameters needed to describe arbitrary states with constraints on symmetries. These circuits are constructed following specific rules to preserve natural symmetries which exist in the problem Hamiltonian and are maintained through the Jordan–Wigner mapping. Specifically, in a system of $2n$ qubits, each qubit represents the occupation of a spin-orbital, and if we assume a block assignment of spins, then the n "top" qubits correspond to spin-up orbitals, while the n remaining "bottom" qubits correspond to spin-down orbitals. The SPC methodology constructs a circuit with gates that preserve the number of electrons (N_e) and spin projection (s_z). The primitive gate used in the SPC construction is a parameterized SWAP-type gate (denoted as ASWAP), given in the computational basis by

$$A(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta & e^{i\phi}\sin\theta & 0 \\ 0 & e^{-i\phi}\sin\theta & -\cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{9}$$

Since a chemical Hamiltonian is real-valued and symmetric (Hermitian) hence its energy eigenvalues are real (time-reversal symmetry). Therefore, we can safely set $\phi = 0$ in Equation (9). By inspection, we can also see that Equation (9) naturally preserves particle number as it only acts non-trivially on the subspace spanned by $\{|01\rangle, |10\rangle\}$ as controlled by the parameter θ. Similarly, ASWAP also preserves spin projection if the gate is not placed in a way that mixes the up-spin and down-spin subspaces, which is achieved by the construction described in ref. [31]. Under the Jordan–Wigner mapping, total spin is a non-local quantity and therefore it is not simple to conserve using local (two-qubit) gates. Consequently, we focus only on preserving the local quantities of particle number and spin projection, using circuits constructed using Equation (9).

In principle, SPC are parameterized circuits constructed of ASWAP gates that operate within a small subspace while preserving the number of particles, N_e, and spin projection, s_z, but not the total spin. Moreover, the triplet states in our 4-qubits test case are not parameterized. When we consider full VQE experiments, we therefore are only concerned with generating the singlet state with $s_z = 0$. For completeness, the numerical parameters which generate the triplet state with $s_z = 0$ using the circuit in **Figure 1** are $\theta_1 = -\pi/4$, $\theta_2 = -\pi/4$, $\theta_3 = 3\pi/4$.

In Figure 1 we show the simplified SPC for constructing any state with $N_e = 2$ in Equation (6). For clarity we discuss how to simplify this circuit in the appendix and only present the result here. In total this circuit is composed of 3 CNOT gates, 21 single qubit gates (when expressed in terms of R_z, R_y, X, H, though other choices are possible), and the minimal number of parameters for this symmetry subspace, which is 3. The resulting

© 2021 Wiley-VCH GmbH

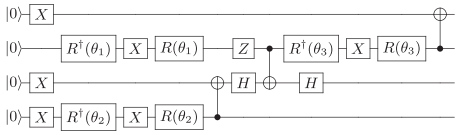

Figure 1. A simplified sample circuit for the case of time-reversal symmetry with $n = 4$, $N_e = 2$, $s_z = 0$ spanning the 4D subspace with $N_e = 2$ in Equation (6) using a minimal number of parameters (3). Here we define $R(\theta) \equiv R_z(\pi)R_y(\theta + \frac{\pi}{2})$ with $R_z(\phi) = e^{-i\phi\sigma_z/2}$, $R_y(\theta) = e^{-i\theta\sigma_y/2}$.

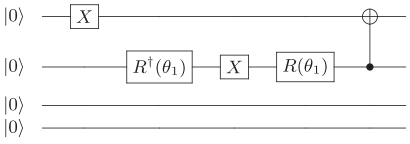

Figure 2. A circuit which generates exactly any state with $N_e = 1$ in Equation (6). As explained in the text, by using particle-hole symmetry we obtain a similar circuit for any state with $N_e = 3$ in Equation (6).

quantum circuit is very similar to the recently proposed UCC-3 circuit in ref. [7]. Since the ground state of the molecules we investigate is of the form of a state with $N_e = 2$ in Equation (6), the circuit given by Figure 1 can efficiently find these ground states.

In **Figure 2**, we similarly show an SPC which finds any state with $N_e = 1$ in Equation (6). Since this symmetry space is smaller, the resulting circuit is simpler. In this case, the total number of CNOT gates is 1, single qubit gates is 6 and the minimal parameter count is 1. Note that the states with $N_e = 1$ in Equation (6) are degenerate with those with $N_e = 1, s_z = -0.5$. However, we could also build an SPC which finds this degenerate state, since it has a different s_z eigenvalue. We omit finding these degenerate states in practice, but note that they are in principle, distinguishable with SPCs. By particle-hole symmetry we can also view Figure 2 as the case of a single hole being exchanged (rather than a particle), therefore naturally representing the $N_e = 3, s_z = -0.5$ case. In this sense, the N_e and $N - N_e$ particles cases only differ in the initial single qubit gates. Therefore, we have two primary configurations of SPC, Figures 1 and 2, for the case of chemical Hamiltonians on 4-qubits.

By employing direct search of energy minimum and maximum within each SPC subspace one can find all of the 16 eigenstates except one of three singlet states which is neither a minimum nor maximum energy within its own symmetry subspace. That is, energy minimization (maximization) within the singlet state subspace leads to the lowest (highest) energy state within this manifold. The third singlet state can be extracted by employing constraint optimization in which the orthogonality to the ground singlet state is enforced. We summarize the total number of states within each symmetry subspace in **Table 1**. As discussed, the SPC allows to efficiently reduce search space while at the same time employing relatively simple circuits.

In **Figure 3** we show the dissociation of LiH, NaH, KH, and RbH molecules at five fixed interatomic distances. We plot the exact eigenenergies (found through diagonalization) in solid lines, while results from IBM quantum hardware are shown with markers. Specifically, we ran experiments on IBM Q devices Valencia, Bogota, Santiago, and Vigo but note that results are consistent across these devices. When executing SPC with three CNOT gates (blue, green and purple), we utilize Richardson extrapolation[38] and readout error mitigation strategies. For SPCs with only a single CNOT (red, orange, brown, and black), we chose to only use readout error mitigation as the benefit from Richardson extrapolation is not significant. As discussed above for each symmetry subspace, the SPCs can target two eigenenergies, one as a minimization and one as a maximization optimization within a target subspace.

Table 1. Table of symmetry spaces along with the number of states within each symmetry subspace and the CNOT and parameter count for the corresponding SPC. Note that we omit four trivial, separable states.

Symmetry space	# of states	CNOTs	Parameters
$N_e = 1, s = 0.5, s_z = 0.5$	2	1	1
$N_e = 1, s = 0.5, s_z = -0.5$	2	1	1
$N_e = 2, s = 0, \ s_z = 0$	3	3	3
$N_e = 2, s = 1, \ s_z = 0$	1	3	0
$N_e = 3, s = 0.5, s_z = 0.5$	2	1	1
$N_e = 3, s = 0.5, s_z = -0.5$	2	1	1

In **Figure 4** we show the energy difference between the exact target eigenenergy and the energy obtained using IBM hardware. Each marker is offset to reduce crowding but the same fixed interatomic distances are present as shown in Figure 3. Error bars are included, representing one standard deviation. We note that many of the standard deviations are of the same order as the energy error mean values. Even though we use the maximum allowed number of shots (8192) this indicates that shot noise is a major source of noise. We can also see that building our SPCs for these three molecules leads to similar results, as well as many measurements within chemical accuracy.

4.2. Quantum Imaginary Time Evolution and Quantum Lanczos Algorithms

In this section, we discuss the QITE and QLanczos algorithms which we used to calculate the energy levels for the same suite of molecules as above. The QITE algorithm offers the quantum version of imaginary-time evolution for the non-unitary evolution operator $\mathcal{U}^* = e^{-\beta H}$. Starting with an appropriately chosen initial state $|\Psi_0\rangle$, after n steps of imaginary-time evolution we obtain

$$|\Psi(\beta)\rangle = c_n(e^{-\Delta\tau H})^n|\Psi_0\rangle \tag{10}$$

where $\Delta\tau = \frac{\beta}{n}$ and c_n is the normalization constant such that $c_n^{-2} = \langle\Psi_0|\mathcal{U}^{*2}|\Psi_0\rangle$. β can be thought of as the inverse temperature. As the temperature goes to zero, the system cools down to its ground state as long as the initial state has a non-zero overlap with it. To represent this non-unitary imaginary-time evolution on a quantum computer without the requirement of any ancilla qubits the QITE algorithm proposes to approximate

© 2021 Wiley-VCH GmbH

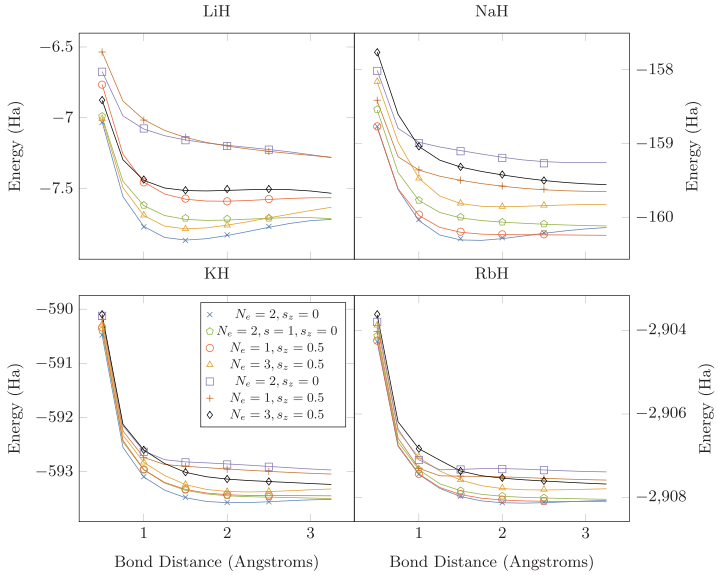

Figure 3. Dissociation curves of four molecules run on IBM Hardware (markers) at fixed interatomic distances of {0.5, 1.0, 1.5, 2.0, 2.5} Angstroms. For each molecule considered, the SPC can accurately probe seven of the eight interesting eigenstates (discussed in text). In all cases, we can see excellent agreement with the exact energies within each symmetry subspace (solid lines). Since the SPC distinguishes between different symmetry spaces, we label each curve by its symmetry eigenvalues throughout dissociation.

these non-unitary operators with unitary updates so that the sth step of evolution is

$$|\Psi_s\rangle = \frac{c_s}{c_{s-1}} e^{-\Delta\tau H}|\Psi_{s-1}\rangle \approx e^{-i\Delta\tau A[s]}|\Psi_{s-1}\rangle \tag{11}$$

with $c_0 = 1$ and the unitary update operator given in terms of the Hermitian operator

$$A[s] = \sum_{i_1,\ldots,i_{N_q}} a[s]_{i_1,\ldots,i_{N_q}} \sigma_{i_1} \ldots \sigma_{i_{N_q}} \tag{12}$$

where N_q is the number of qubits and $\sigma \in \{X, Y, Z\}$ are Pauli matrices. Our next step is to calculate the coefficients $a[s]$ by solving Equation (11) to second order, $\mathcal{O}(\Delta\tau^2)$. This amounts to solving a linear system of equations, $(\boldsymbol{S} + \boldsymbol{S}^T) \cdot \boldsymbol{a} = \boldsymbol{b}$ where

$$S_{I,I'} = \langle \sigma_{i_1} \ldots \sigma_{i_{N_q}} \sigma_{i'_1} \ldots \sigma_{i'_{N_q}} \rangle \tag{13}$$

and

$$b_I = -i\sqrt{\frac{c_{s-1}}{c_s}} \langle \sigma_{i_1} \ldots \sigma_{i_{N_q}} H \rangle \tag{14}$$

with $I = \{i_1, \ldots, i_{N_q}\}$ and similarly for I'. The expectation values at a given step s are calculated with respect to the state in the previous QITE step, $|\Psi_{s-1}\rangle$. The bottleneck in this algorithm is due to fact that the number of measurements needed to calculate $\boldsymbol{S}$ (a $3^{N_q} \times 3^{N_q}$ matrix) and $\boldsymbol{b}$ (a 3^{N_q}-dimensional vector) scales exponentially with the number of qubits N_q in the system. Having a real Hamiltonian eliminates the terms with an odd number of Y Pauli matrices in $\boldsymbol{b}$. Also, the fact that the matrix $\boldsymbol{S} + \boldsymbol{S}^T$ is symmetric reduces the required number of measurements significantly. Nevertheless, the problem of complexity persists, and reducing QITE complexity is an active field of research, just as it has been for reducing circuit depth in VQE (see, e.g., ref. [39]).

The cloud access to NISQ devices limits the number of measurements that can be done in a reasonable amount of time. Additionally, the depth of quantum circuits is limited by the decoherence time of the qubits. Because of these limitations, we perform QITE algorithms on a NISQ device by following these steps:

- To decide on the convergence of the energy expectation value, we use exact QITE calculations. This allows us to validate the hardware results even though the noise in quantum hardware might result in deviations of energy values that are greater than the value that we choose for our convergence criterion. We set the latter to $\epsilon = 0.001$ for energy expectation values, and find the number of steps required for convergence.
- We then calculate the unitary updates in (12) at every QITE step from (13) and (14) by performing the energy expectation value calculations on a noisy quantum simulator which includes the noise model of the quantum hardware to be used.

© 2021 Wiley-VCH GmbH

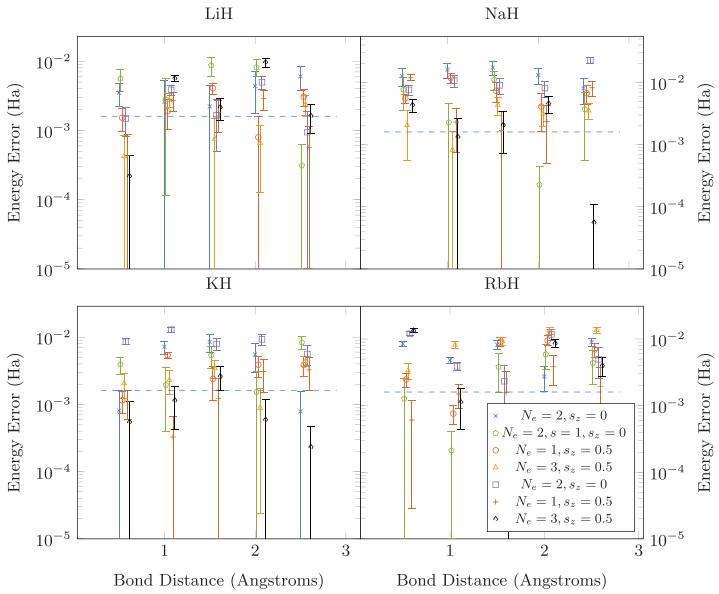

Figure 4. Energy differences from exact eigenenergies of each of the four molecules we consider. Markers are results from IBM hardware with error bars representing one standard deviation. Different symmetry spaces are artificially offset for easier viewing but still have the same fixed interatomic distances shown in Figure 3. We show several results within chemical accuracy (dashed line).

- After choosing the initial state $|\Psi_0\rangle$ using the symmetry considerations for each of the reduced Hamiltonian blocks, we find the state

$$|\Psi_s\rangle = e^{-i\Delta\tau A[s]} e^{-i\Delta\tau A[s-1]} \cdots e^{-i\Delta\tau A[1]} |\Psi_0\rangle \tag{15}$$

  after application of s unitary updates.
- We obtain the quantum circuit corresponding to the state $|\Psi_s\rangle$ by using the isometry function in IBM Qiskit library. This function is based on ref. [40] which decomposes a given state into single-qubit and Controlled-NOT (CNOT) gates with the aim of having the least number of CNOT gates. An example of this quantum circuit can be seen in Figure B2. This quantum circuit has the same single-qubit $U3$ and two-qubit CNOT gates at every QITE step and only the angles of $U3$ gates are updated at every QITE step, where

$$U3(\theta, \phi, \lambda) = \begin{pmatrix} \cos\frac{\theta}{2} & -e^{i\lambda}\sin\frac{\theta}{2} \\ e^{i\phi}\sin\frac{\theta}{2} & e^{i(\phi+\lambda)}\cos\frac{\theta}{2} \end{pmatrix} \tag{16}$$

- We then run the quantum circuits corresponding to the states $|\Psi_s\rangle$ on quantum hardware and obtain the energy expectation values at every QITE step. (The experimental results of these calculations can be seen in **Figure 5**).

Next, we discuss how to implement the QLanczos algorithm.

The QLanczos algorithm makes use of the QITE algorithm to obtain the states in Krylov space $\mathcal{K}$ which is spanned by $\{|\Phi_0\rangle, |\Phi_2\rangle, \dots\}$, where $|\Phi_l\rangle \in \{|\Psi_0\rangle, |\Psi_2\rangle, \dots\}$. After choosing the vectors in Krylov space, we calculate the overlap ($\mathcal{T}$) and Hamiltonian ($\mathcal{H}$) matrices from the QITE algorithm measurements on quantum hardware. The normalization constants are calculated recursively using

$$\frac{1}{c_{r+1}^2} = \frac{\langle\Phi_r|e^{-2\Delta\tau H}|\Phi_r\rangle}{c_r^2} \tag{17}$$

Since we are using a small $\Delta\tau$ approximation, we calculate the expectation value by expanding to first order in $\Delta\tau$, $\langle\Phi_r|e^{-2\Delta\tau H}|\Phi_r\rangle = 1 - 2\Delta\tau\langle\Phi_r|H|\Phi_r\rangle + \mathcal{O}(\Delta\tau^2)$. Thus, we obtain the normalization constants using the energy expectation values obtained from quantum hardware in the implementation of the QITE algorithm. To increase the accuracy of the calculations, one can use higher-order approximations.

The overlap and Hamiltonian matrix elements can be expressed in terms of these constants and energy expectation values,

$$\mathcal{T}_{l,l'} = \langle\Phi_l|\Phi_{l'}\rangle = \frac{c_l c_{l'}}{c_r^2} \tag{18}$$

and

$$\mathcal{H}_{l,l'} = \langle\Phi_l|H|\Phi_{l'}\rangle = \mathcal{T}_{l,l'}\langle\Phi_r|H|\Phi_r\rangle \tag{19}$$

© 2021 Wiley-VCH GmbH

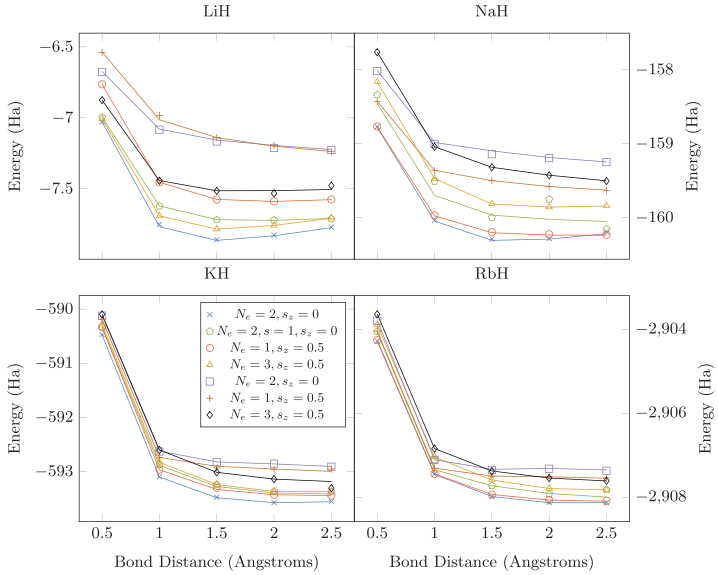

Figure 5. Energy versus bond distance obtained using the reduced Hamiltonian blocks for molecules LiH, KH, RbH, NaH. The experimental data, notated with markers in the figure, were obtained using QITE and QLanczos algorithms. The quantum circuits were run on IBM Q's several quantum computers such as Casablanca, Manhattan, Vigo, Bogota, Rome. The experimental results are in very good agreement with the exact values shown with straight lines in figure.

where $r = \frac{l+l'}{2}$, and l, l' are even integers.

The energy eigenvalues E are found by solving the generalized eigenvalue equation

$$\boldsymbol{\mathcal{H}}\boldsymbol{x} = E\boldsymbol{\mathcal{T}}\boldsymbol{x} \tag{20}$$

The corresponding eigenvectors $\boldsymbol{x} = (x_0, x_1, \dots)$ determine the eigenstates of the system Hamiltonian,

$$|\Psi[E]\rangle = c_E\left(x_0|\Phi_0\rangle + x_1|\Phi_2\rangle + \dots\right) \tag{21}$$

where $c_E^{-1} = \| \sum_i x_i|\Phi_i\rangle\|$. To decide which energy levels are non-spurious, we use the minimum total uncertainty criterion. In this method, we scan all possible combinations of QITE states and calculate classically the uncertainty in energy as $\Delta E = ||H|\Psi[E]\rangle - E|\Psi[E]\rangle||$ for each vector in the Krylov space, and we add these uncertainties up. We choose the vectors in Krylov space that give the minimum total uncertainty below a certain value, $\Delta E \leq \delta$. We discard the eigenvalues and eigenvectors that do not satisfy this criterion. In the next subsection, we outline the case with $\delta \leq 0.5$. A detailed description of the method can be found in ref. [41].

As explained in the previous section, the symmetry of each of these molecules leads to a simplification of the system which makes them good candidates for implementation on NISQ devices. In the previous section we used this symmetry to simplify the quantum circuits. As a second perspective, we present a simplification of the system by reducing the Hamiltonian into smaller blocks so that we can study these molecules using a smaller number of qubits on quantum hardware. For benchmarking purposes, we also present our QITE algorithm results for the full LiH molecule Hamiltonian for comparison with the results presented using the SPC method. In what follows, we explain the reduction of the system size using the symmetry of the Hamiltonian of various molecules.

In Section 3, we discuss chemical Hamiltonian for our benchmark molecules. Some of the solutions are given in Equations (6)–(8), and also listed in Table 1. In a 4-qubit molecular Hamiltonian, the matrix elements corresponding to a no-electron state in the active space ($N_e = 0, s_z = 0$) and a 4-electron state ($N_e = 4, s_z = 0$) are diagonal so no further computation of eigenvalues and corresponding eigenstates is needed. The Hamiltonian blocks formed by single ($N_e = 1$) and three ($N_e = 3$) electron states are 4×4 matrices. However, each 4×4 Hamiltonian block is itself block-diagonal consisting of two 2×2 blocks corresponding to a different value of spin projection ($s_z = +\frac{1}{2}$ and $s_z = -\frac{1}{2}$). Each of the four 2×2 blocks can be handled by a single-qubit circuit. Similarly, two-electron states ($N_e = 2$) form a 6×6 matrix which includes a diagonal 4×4 Hamiltonian block for $s_z = 0$ (three singlet states and one triplet state) and can be represented using a 2-qubit circuit. The remaining two states represent triplet states with $s_z = 1+$ and $s_z = -1$. In conclusion, we can compute the energy levels of each $2^4 \times 2^4$ molecular Hamiltonian using only two-qubit and single-qubit circuits. This analysis shows that

© 2021 Wiley-VCH GmbH

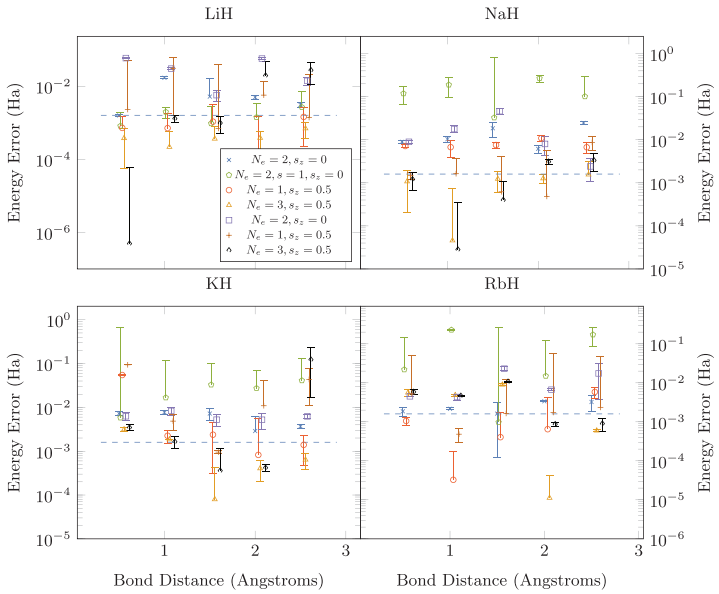

Figure 6. Error in energy versus bond distance obtained using the reduced Hamiltonian blocks for molecules LiH, KH, RbH, NaH. The dashed line in each panel represents the chemical accuracy (1.5×10^{-3}). As seen in the figures we were able to obtain several energy eigenvalues within the chemical accuracy.

decomposition of the full Hamiltonian into diagonal blocks based on the number of electron and spin symmetries is essential to making quantum computation compatible with NISQ devices.

We start by calculating the energy expectation values of the system using the QITE algorithm which will give us the ground state energies of the symmetry sector that the initial state belongs to. To this end, as mentioned earlier, the Hamiltonian of the molecules are divided into smaller blocks and the initial states are chosen accordingly using symmetry which results in a smaller number of imaginary-time steps required for convergence. To be able to represent these reduced Hamiltonian blocks on a quantum computer we need to express them in terms of Pauli matrices. This can be done by writing the block Hamiltonian H_{block} as

$$H_{\text{block}} = \sum_I c_I \sigma_I \tag{22}$$

where $I = \{i_0, \ldots, i_{n_q-1}\}$ and $\sigma_I = \sigma_{i_0}\sigma_{i_1}\ldots\sigma_{n_q-1}$, with $\sigma \in \{I, X, Y, Z\}$ and n_q the number of qubits used for this block Hamiltonian. The coefficients c_I are found from

$$c_I = \frac{1}{2^{n_q}} \text{Tr}[H_{\text{block}} \cdot \sigma_I] \tag{23}$$

Once the coefficients c_I are determined, the Hamiltonian H_{block} expressed in terms of Pauli matrices can be implemented on a quantum computer. We use H_{block} in QITE and QLanczos algorithms to calculate the energy eigenvalues of each molecular Hamiltonian with states shown in Table 1.

We use the information from symmetry considerations of the reduced block Hamiltonians to choose different initial states, $|\Psi_0\rangle$, so as we obtain all energy levels of the system Hamiltonian.

4.2.1 QITE and QLanczos Results

To calculate the energy expectation values at each QITE step for the molecules that we study, we first ran the quantum circuits obtained using the isometry function for each QITE step on IBM's cloud accessible quantum computers. We did this for the reduced Hamiltonian blocks first and then for the 4-qubit full LiH molecule Hamiltonian. An example of 2-qubit quantum circuit for reduced block Hamiltonian of $N_e = 2$ states can be found in Figure B2. At each QITE step, the angles θ_1, θ_2, and θ_3 change but the quantum circuit depth remains fixed. We then used the measured energy expectation values to find the energy eigenvalues of the systems of interest.

The calculation of the energy eigenvalue from the QLanczos algorithm using the method described above is often numerically unstable due to the noise introduced by NISQ devices. To stabilize the results, we worked as follows. We only used two-dimensional Krylov spaces for the $\boldsymbol{\mathcal{T}}$ and $\boldsymbol{\mathcal{H}}$ matrices in (20), which have been shown to yield more accurate results in NISQ devices.[41] We also used a total minimum uncertainty criterion of

© 2021 Wiley-VCH GmbH

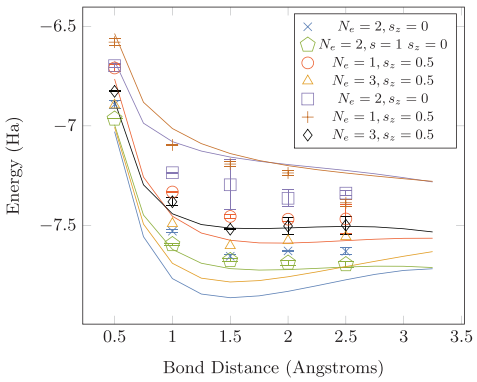

Figure 7. Energy versus bond distance obtained using the 4-qubit LiH molecule Hamiltonian. The Richardson extrapolation applied to experimental data, notated with markers in the figure, were obtained using QITE algorithm. The quantum circuits were run on IBM Q's Yorktown and Melbourne device. The experimental results are in agreement with the exact values shown with straight lines in figure but most of them are not in chemical accuracy. The error bars represent $\pm\sigma$.

$\delta \leq 0.5$; if the total uncertainty for the vectors in Krylov space was higher than this value, we increased l and m, which corresponds to higher QITE steps, which theoretically reduced the total uncertainty. The Hamiltonian matrices are diagonally dominant with the diagonal elements being a large negative number since due to the electron–nuclei Coulomb interaction. Each reduced Hamiltonian block was shifted by appropriate constant value, to ensure that all matrix elements are of the same order. While this did not affect the eigenstates of the system, it improved the accuracy of the results significantly. The accuracy of the eigenvalues can be improved further if one includes higher-order terms in the expansion of $e^{-\Delta\tau H}$, but this would require more measurements of higher moments of the Hamiltonian resulting in longer runs on quantum hardware.

See Appendix B for an example application of QITE algorithm on IBM Q hardware for 2-qubit reduced Hamiltonian blocks of molecules LiH and RbH where we also demonstrate the 2-qubit circuit obtained from isometry function of IBM's Qiskit library.

In Figure 5, we present our experimental results obtained using QITE and QLanczos algorithms for the energy of molecules LiH, KH, RbH, NaH as a function of bond length. We see an excellent agreement between the experimental results and the values obtained from exact diagonalization. In these plots, we present 7 of the 8 eigenvalues that could not be readily obtained from the Hamiltonian. The remaining eighth eigenvalue is the first excited-state energy of the $N_e = 2$ symmetry subspace for the molecules. It turns out to be exactly given by the analytic expression $\frac{1}{\sqrt{2}}(|10\rangle - |01\rangle)$. This is confirmed by implementing the QITE algorithm on hardware starting with the initial state $|\Psi_0\rangle = \frac{1}{\sqrt{2}}(|10\rangle - |01\rangle)$. We show the experimental results for this state in upper left panel of Figure 5 for the LiH molecule.

In the results presented here we used the QITE algorithm to find the minimum and maximum state energies corresponding to $N_e = 1$, $N_e = 2$, and $N_e = 3$ symmetry subspaces and the QLanczos algorithm to find the remaining excited-state energies.

We also present the errors in the energy eigenvalues as a function of the bond length as a result of using the QITE and QLanczos algorithms from reduced Hamiltonian blocks of molecules LiH, RbH, NaH, KH with respect to chemical accuracy in **Figure 6**. Although we only used readout error mitigation, we were able to obtain several energy eigenvalues within chemical accuracy.

Although the Hamiltonian can be separated into smaller blocks and these molecules can be studied in smaller quantum hardware sizes, for benchmarking purposes we also ran the QITE algorithm for the full 4-qubit LiH molecule Hamiltonian. Compared to the reduced system, the quantum circuit obtained using the IBM Q library isometry function is deeper than before and requires connections between all qubits in the system. The isometry function does include the minimum required number of CNOTs in the quantum circuit, but the increased depth arises from the need to include SWAP gates to provide the needed connectivity. On the other hand, if one wants to study system Hamiltonians which result in encodings that require strong interactions between specific qubits, then correlation-informed permutation method proposed in ref. [42] would reduce the CNOT depth compared to isometry. In this method, the qubits used to encode strongly correlated spin orbitals during the algorithm are placed in close proximity, reducing the SWAP depth on linear qubit layouts. Many of the IBM Q devices have linear layout, with IBM Q Yorktown device having the most connections between qubits. Consequently, we ran the QITE algorithm on the IBM Q Yorktown and Melbourne device and the results can be seen in **Figure 7**. The Richardson extrapolation error mitigation strategy was applied to experimental energy values seen in Figure 7. Although the experimental data are in agreement with exact values, most of them are not chemically accurate. Due to this noise, the QLanczos algorithm is not numerically stable and cannot be used to improve on QITE algorithm results in this case.

4.3. Reducing Noise with Hidden Inverse

Variational algorithms are designed to work in near-term quantum computers with imperfect gates and in a nascent quantum error correction regime. It is known that the variational formalism has inherent robustness to certain types of time independent systematic errors and this was discussed in ref. [43]. In addition, as demonstrated in the previous two sections, the "workhorse" error mitigation methods, readout error mitigation and Richardson extrapolation, can improve the expectation value accuracy in many instances. These two methods have key limitations, however. Beyond a certain circuit depth determined by overall hardware noise, Richardson extrapolation will not yield useful data.[44,45] Meanwhile, exact readout error mitigation is in principle unscalable, though scalable approximate methods are under development. On the other hand, recent advancements in quantum characterization procedures[46] have demonstrated systematic error parameters to be time-dependent stochastic in the time scale of multiple circuits run within a given circuit list. One way to model this type of error (in the context of variational algorithms) is with systematic errors that are drawn independently

© 2021 Wiley-VCH GmbH

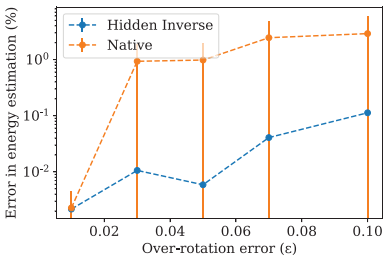

Figure 8. Average error in estimated ground state energy of LiH at 0.5 interatomic distance plotted against over-rotation error ϵ. Hidden inverse optimized Ansatz outperforms native UCC-3 Ansatz for the entire range of ϵ. Error bars represent one standard deviation.

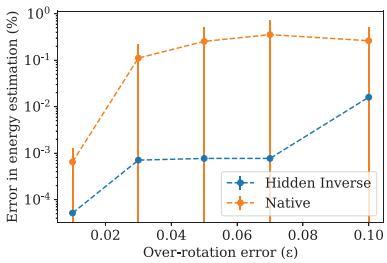

Figure 9. Average error in estimated ground state energy of NaH at 0.5 interatomic distance plotted against over-rotation error ϵ. Hidden inverse optimized Ansatz outperforms native UCC-3 Ansatz for the entire range of ϵ. Error bars represent one standard deviation.

from a stationary distribution in every iteration of the classical optimizer. This has a direct effect on the performance of variational algorithms. In particular, mitigating such an error does not require exponential resources, and it can be done while maintaining constant circuit depth, unlike extrapolation methods.

In this section, we test the effectiveness of a circuit level error mitigation technique called hidden inverse[47,48] to mitigate precisely this type of error. The key idea behind this error mitigation technique is that each self-inverse gate (such as CNOT) can be experimentally implemented in standard or inverted configuration. CNOT gates are not native gates in most hardware but are implemented by a combination of single qubit and two-qubit native rotations. By carefully choosing either the standard gate or the inverted gate within a circuit, it is possible to mitigate some of the errors of native rotations. This technique can be easily understood with a simple example circuit. Let us imagine a two qubit circuit composed of two CNOT gates and a RZ rotation sandwiched between the two in the target qubit. We compare the circuit error for two versions of the circuit 1) both CNOTs are performed in the standard configuration, 2) first CNOT is performed in the standard configuration and the second CNOT is performed in the inverted configuration. Under a coherent error noise model, for small values of RZ rotation, error in circuit (1) will add up while error in circuit (2) will cancel out (degree of cancellation will depend on error correlation time). In the absence of noise, both circuits should result in identical final states but in the presence of coherent error, circuit (2) is clearly the better choice. Here we simulate VQE circuits with UCC-3 Ansatz[7] for LiH and NaH molecules under a two qubit over-rotation error noise model (see Appendix C for more information on hidden inverse technique and details on simulation such as noise model, Ansatz and classical optimization). Essentially, we identify subcircuits within our Ansatz that can benefit from hidden inverse noise cancellation and replace them with the optimized circuit.

Figures 8 and **9** display the comparison between errors in calculated ground state energy computed with the native Ansatz and the hidden inverse optimized Ansatz as a function of the over-rotation error ϵ. These plots are for LiH and NaH for a fixed interatomic distance of 0.5 Angstroms. The range of ϵ is chosen to match realistic noise parameters of current hardware. We can see from these two plots that when the over-rotation error increases, the average error in the calculated ground state energy with the native UCC-3 Ansatz (orange curve) tends to get worse. We also see large variability in the estimated energy as expected from the stochastic modeling of the error parameters. The hidden inverse optimized Ansatz (blue curve) is found to be robust against this stochastic over-rotation error as we see the error stay low (with reduced variability) across the error range. Hidden inverse offers one to two orders of magnitude improvement for the entire range.

5. Conclusion

In this report, we have outlined the state of the art of the field devoted to benchmarking quantum chemistry computations on NISQ devices. This application is a key indicator of NISQ capabilities, and it fills a performance assessment gap left open by low level benchmarks to date. We outlined how to extend these benchmarks with variational methods plus short depth Ansätze on one hand, and imaginary time evolution plus Lanczos on the other. We also outlined how to improve all of our algorithms with scalable error mitigation techniques via the hidden inverse.

In doing so, we have added QITE, Lanczos, and SPCs to our electronic structure calculation benchmark suite. We used these algorithms to demonstrate that NISQ devices are now capable of chemical accuracy over the cloud for minimal basis set computations. Notably, only 4 years ago the first chemically accurate computations were performed by researchers with direct access to the hardware with intricate precision in analog quantum control, and chemical accuracy over the cloud was impossible. Meanwhile, the benchmarks presented here are fully specified in QASM. When this suite of molecules was simulated on quantum computers 2 years ago, it was found that chemical accuracy was just out of reach. Here, many of the same computations resulted in chemical accuracy, highlighting both the immense progress of quantum computers to date and the efficacy of application-level benchmarking in tracking this progress.

While the algorithms we tested were quite different in their approach and circuit constructions, we note that both were able to achieve chemically accurate results for certain configurations. While all NISQ era algorithms necessarily suffer drawbacks, notably the requirement of very short depths, we note here that VQE+SPC and QITE+Lanczos have complementary properties and thus are optimal for different scenarios.

© 2021 Wiley-VCH GmbH

The SPC Ansätze are primarily designed to recover ground state and are optimal in circuit depth for ground state energy computation. However, VQE+SPC can also be applied to extract some excited states, particularly those which correspond to a minimum or maximum of energy within given symmetry Hamiltonian block. On the other hand, the QITE+Lanczos method is suitable not only for ground state calculations but is also intrinsically designed to recover excited states which are not accessible to VQE+SPC. However, the scalability of the algorithm is at an earlier state of development compared to VQE+SPC. We look forward to further refining these algorithms, and combining them with advanced error mitigation via the hidden inverse in hardware in the near future. Such advanced error mitigation techniques will be critical for enabling simulation of chemical systems beyond 4-qubits.

Finally, we note that given the rapid progress of NISQ devices in chemistry, as outlined by performance in this benchmark suite, we expect expanded basis sets to be usable for chemical accuracy in the very near future. While quantum computers with more than 50 qubits are now available, large molecules with many active orbitals will require a consistent and concerted effort in coherence time improvement in order to be simulatable. This is mostly due to the fact that the Hilbert space covered by active orbitals spanning an encoding across 50 qubits cannot be fully explored within the coherence times of today's devices. Even linear-depth circuits will require a 5 to 10× improvement in coherence time over today's to run to completion before full decoherence. Nonetheless, given the current rate of progress, 10 qubit computations with more advanced bases will be routine within a year or two, and computations involving more than 30 qubits will soon follow. In this mesoscale era, just at the border of a potential quantum advantage, the pace will be brisk and the potential for new science and new discovery cannot be overestimated.

Appendix A: Details of SPC Simplifications

The SPC presented in the main text are constructed using the ASWAP gate as a primitive. Each ASWAP gate can be decomposed into three CNOTs and four single qubit gates according to **Figure A1**.

Building the SPC for the case of $N_e = 2, s_z = 0$ immediately results in the circuit shown in **Figure A2**. Naively, this circuit contains 12 CNOT gates, but with further investigation we can see that we can significantly simplify the circuit. The first such simplification can be seen since the input state is fixed by placing an X on qubits q_0, q_3. In this case, the first two CNOT gates in both of the first two ASWAP gates are unnecessary and can be replaced by either an I (Identity) or X gate. Furthermore, by inspection of Equation (9), we can see that the ASWAP gate with zero arguments is locally equivalent to a controlled-Z gate. Additionally, another simplification can be performed since the final ASWAP(θ_3) gate has a "first" CNOT gate which can be commuted through the controlled-Z gate and cancelled with the "last" CNOT gate of ASWAP(θ_1). We can then also replace the "second" CNOT gate of ASWAP(θ_3) with a single qubit X gate. In total, this then simplifies the SPC to only require three CNOT gates. Similar simplifications can be applied to arrive at the simplified Figure 2.

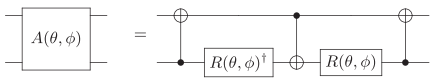

Figure A1. Decomposition of the A gate in terms of elementary single and two-qubit gates. $R(\theta, \phi) = R_z(\phi + \pi)R_y(\theta + \pi/2)$, where $R_z(\theta) = \exp(-i\theta\sigma_z/2)$, $R_y(\phi) = \exp(-i\phi\sigma_y/2)$.

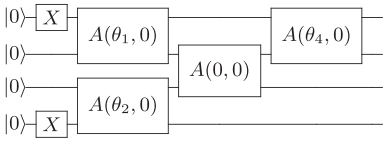

Figure A2. Following the SPC construction rules, we can directly build the SPC for the $n = 4$, $m = 2$, $s_z = 0$ symmetry eigenvalues. This circuit can be significantly simplified and the result is presented in the main text Figure 1.

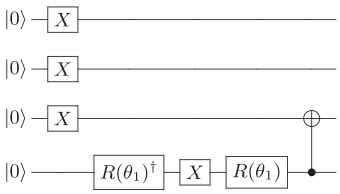

Figure A3. A circuit which exactly generates any state with $N_e = 3$ in Equation (6). This circuit can be thought of as swapping a single hole, while Figure 2 swaps a single particle.

We can slightly modify the $N_e = 1, s_z = 0.5$ circuit shown Figure 2 in order to generate a similar circuit which works for the $N_e = 3, s_z = 0.5$ subspace. If we place initial X gates on the q_1, q_2 qubits and we also move the ASWAP gate to act on the bottom two qubits, then we immediately have a circuit which targets $N_e = 3, s_z = 0.5$ subspace with a similar form as Figure 2. For completeness, we show the resulting circuit in **Figure A3**.

Appendix B: Example QITE Results

In **Figure B1**, we present two examples of the measured energy expectation values as a function of imaginary time obtained from the QITE algorithm. Since the QITE algorithm always converges to the ground-state energy of the symmetry subspace that the initial state belongs to, in order to access some of the excited-state energies, we ran the QITE algorithm using $-H_{\text{block}}$, instead. Figure B1a shows an example in which the energy expectation value converges to the third excited-state energy of the $N_e = 2$ symmetry subspace block Hamiltonian of the LiH molecule where we started with the initial state $|\Psi_0\rangle = |11\rangle$. Figure B1b shows the convergence to the ground state energy of the $N_e = 2$ symmetry subspace of RbH starting with the initial state $|\Psi_0\rangle = |00\rangle$ and using H_{block} in the QITE algorithm.

The readout-error-mitigated (ROEM) experimental data in Figure B1 were obtained by running the quantum circuits of the form shown in **Figure B2** on IBM Q 7-qubit Casablanca quantum hardware using the number of shots $N_{\text{shots}} = 8192$ at each measurement. The experiments in Figure B1a (Figure B1b) were

© 2021 Wiley-VCH GmbH

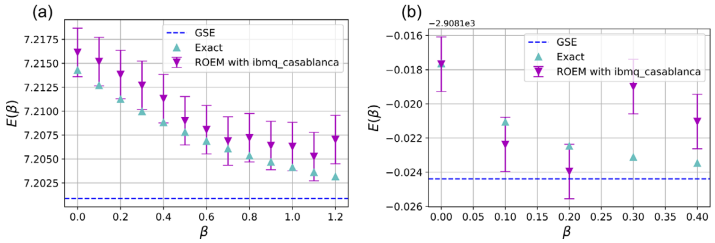

Figure B1. The convergence of the energy expectation values as a function of imaginary-time, β, using QITE algorithm. a) Convergence to the ground state of $N_e = 2$ symmetry subspace for LiH molecule with initial state $|\Psi_0\rangle = |11\rangle$ and $-H$. The converged energy corresponds to the negative of the maximum (3rd excited state) energy of that symmetry subspace. b) Convergence to the ground state of $N_e = 2$ symmetry subspace for RbH molecule with initial state $|\Psi_0\rangle = |00\rangle$ and H. The readout error mitigated experimental data were collected on IBM Q 7-qubit Casablanca device and compared to the data calculated using exact diagonalization. Error bars represent $\pm\sigma$.

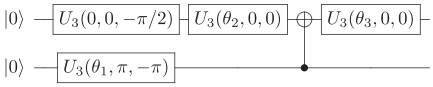

Figure B2. The quantum circuit of each QITE step for 2-qubit reduced Hamiltonian blocks. At every QITE step the angles θ_1, θ_2, and θ_3 change but the circuit depth is fixed.

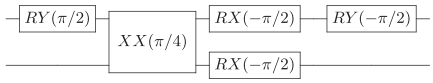

Figure C1. Decomposition of CNOT gate into native ion trap gates.

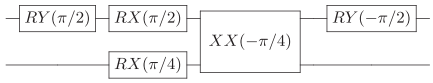

Figure C2. Decomposition of CNOT† into native ion trap gates.

run $N_{\text{runs}} = 3$ ($N_{\text{runs}} = 2$) times and the error bars represent one standard deviation ($\pm\sigma$). Although in these two examples we used IBM Q's Casablanca device, for the rest of the measurements on quantum hardware we used various devices such as Manhattan, Vigo, Bogota, and Rome depending on their availabilities.

Appendix C: Details of Hidden Inverse Error Mitigation

In order to demonstrate the effectiveness of the hidden inverse protocol, we have decided to perform simulation under a ion trap noise model. A CNOT gate in trapped ion systems is synthesized using XX interaction padded with single qubit rotations as seen in **Figure C1** while the CNOT inverse is implemented using the configuration found in **Figure C2**. **Figure C3** shows the standard UCC-3 Ansatz with three unknown parameters $(\theta_1, \theta_2, \theta_3)$. We then modified this Ansatz (shown in **Figure C4**) where we carefully replaced some of the CNOT gates with CNOT inverse gates to cancel out systematic errors. For our simulation, we assumed single qubit gates are error free and XX gates

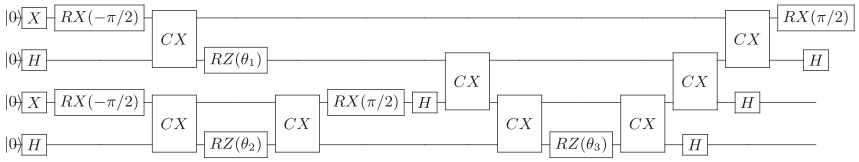

Figure C3. UCC-3 Native Ansatz. The top qubit is control and the bottom qubit is target for all the CX gates in the Ansatz.

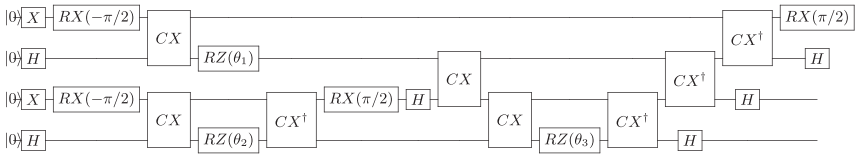

Figure C4. UCC-3 Hidden inverse Ansatz. The top qubit is control and the bottom qubit is target for all the CX and $CX^\dagger$ gates in the Ansatz.

© 2021 Wiley-VCH GmbH

contain multiplicative over-rotation type errors as such we apply $XX(\pm(1+\text{err})\times\frac{\pi}{4})$. Error amounts are drawn independently from $\text{err}\approx N(0,\epsilon)$ at each round of the classical optimizer. We used BOBYQA (Bound Optimization BY Quadratic Approximation) as implemented in SKQuant-Opt,[49] a standard optimizer package for near term hybrid quantum classical algorithms. Although we have simulated the results under a ion trap model, hidden inverse can be applied to other systems too(as an example to reduce errors in ZX gates for a superconducting hardware).

Acknowledgements

The quantum circuits were drawn using Q-circuit package.[50] This work was supported by the ASCR Quantum Testbed Pathfinder program at Oak Ridge National Laboratory under FWP number ERKJ332. This research used quantum computing system resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. S.M. was supported through US Department of Energy grant DE-SC0019294 awarded to Duke and is funded in part by an NSF QISE-NET fellowship (DMR-1747426). G.B. and B.G. were supported through US Department of Energy grants awarded to Virginia Tech (Awards DE-SC0019318 and DE-SC0019199, respectively). G.S. acknowledges the Army Research Office award W911NF-19-1-0397 and the National Science Foundation award OMA-1937008.

Conflict of Interest

The authors declare no conflict of interest.

Data Availability Statement

Research data are not shared.

This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan.

Keywords

quantum benchmarks, quantum chemistry, quantum computing, quantum imaginary time evolution, variational algorithms

Received: January 26, 2021
Revised: March 14, 2021
Published online: May 7, 2021

[1] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya, S. Sim, L. Veis, A. Aspuru-Guzik, *Chem. Rev.* **2019**, *119*, 10856.
[2] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, X. Yuan, *Rev. Mod. Phys.* **2020**, *92*, 015003.
[3] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, J. M. Gambetta, *Nature* **2017**, *549*, 242.
[4] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, I. Siddiqi, *Phys. Rev. X* **2018**, *8*, 011021.
[5] M. Urbanek, D. Camps, R. V. Beeumen, W. A. de Jong, *J. Chem. Theory Comput.* **2020**, *16*, 5425.
[6] 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, C. F. Roos, *Phys. Rev. X* **2018**, *8*, 031022.
[7] A. J. McCaskey, Z. P. Parks, J. Jakowski, S. V. Moore, T. D. Morris, T. S. Humble, R. C. Pooser, *npj Quantum Inf.* **2019**, *5*, 99.
[8] Google AI Quantum and Collaborators, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, et al., *Science* **2020**, *369*, 1084.
[9] R. Blume-Kohout, K. C. Young, *Quantum* **2020**, *4*, 362.
[10] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, H. Neven, *Nat. Phys.* **2018**, *14*, 595.
[11] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, D. J. Wineland, *Phys. Rev. A* **2008**, *77*, 012307.
[12] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, P. J. Coles, arXiv:2012.09265 [quant-ph], **2020**.
[13] O. Higgott, D. Wang, S. Brierley, *Quantum* **2019**, *3*, 156.
[14] T. Jones, S. Endo, S. McArdle, X. Yuan, S. C. Benjamin, *Phys. Rev. A* **2019**, *99*, 062304.
[15] S. C. Benjamin, Y. Li, *Phys. Rev. X* **2017**, *7*, 021050.
[16] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, A. Kandala, A. Mezzacapo, P. Müller, W. Riess, G. Salis, J. Smolin, I. Tavernelli, K. Temme, *Quantum Sci. Technol.* **2018**, *3*, 030503.
[17] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, J. L. O'Brien, *Nat. Commun.* **2014**, *5*, 4213.
[18] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, X. Yuan, *npj Quantum Inf.* **2019**, *5*, 75.
[19] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, J. M. Martinis, *Phys. Rev. X* **2016**, *6*, 031007.
[20] J. Romero, R. Babbush, J. R McClean, C. Hempel, P. J. Love, A. Aspuru-Guzik, *Quantum Sci. Technol.* **2019**, *4*, 014008.
[21] H. R. Grimsley, S. E. Economou, E. Barnes, N. J. Mayhall, *Nat. Commun.* **2019**, *10*, 1.
[22] D. Claudino, J. Wright, A. J. McCaskey, T. S. Humble, arXiv:2011.01279 [quant-ph], **2020**.
[23] I. O. Sokolov, P. K. Barkoutsos, P. J. Ollitrault, D. Greenberg, J. Rice, M. Pistoia, I. Tavernelli, *J. Chem. Phys.* **2020**, *152*, 124107.
[24] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, G. K.-L. Chan, *Nat. Phys.* **2020**, *16*, 205.
[25] K. Yeter-Aydeniz, R. C. Pooser, G. Siopsis, *npj Quantum Inf.* **2020**, *6*, 63.
[26] N. Gomes, F. Zhang, N. F. Berthusen, C.-Z. Wang, K.-M. Ho, P. P. Orth, Y. Yao, *J. Chem. Theory Comput.* **2020**, *16*, 6256.
[27] S. Barison, D. E. Galli, M. Motta, arXiv:2011.08137 [quant-ph], **2020**.
[28] B. Bauer, S. Bravyi, M. Motta, G. K.-L. Chan, *Chem. Rev.* **2020**, *120*, 12685.

 © 2021 Wiley-VCH GmbH

[29] S. E. Smart, D. A. Mazziotti, *Phys. Rev. Lett.* **2021**, *126*, 070504.
[30] G. S. Barron, B. T. Gard, O. J. Altman, N. J. Mayhall, E. Barnes, S. E. Economou, arXiv:2003.00171 [quant-ph], **2020**.
[31] B. T. Gard, L. Zhu, G. S. Barron, N. J. Mayhall, S. E. Economou, E. Barnes, *npj Quantum Inf.* **2020**, *6*, 10.
[32] P. Kl. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, I. Tavernelli, *Phys. Rev. A* **2018**, *98*, 022322.
[33] P. Jordan, E. Wigner, *Z. Phys.* **1928**, *47*, 631.
[34] S. B. Bravyi, A. Y. Kitaev, *Ann. Phys.* **2002**, *298*, 210.
[35] K. Setia, S. Bravyi, A. Mezzacapo, J. D. Whitfield, *Phys. Rev. Res.* **2019**, *1*, 033033.
[36] M. Rossmannek, P. K. Barkoutsos, P. J. Ollitrault, I. Tavernelli, arXiv:2009.01872 [quant-ph], **2020**.
[37] K. Setia, R. Chen, J. E. Rice, A. Mezzacapo, M. Pistoia, J. D. Whitfield, *J. Chem. Theory Comput.* **2020**, *16*, 6091.
[38] Y. Li, S. C. Benjamin, *Phys. Rev. X* **2017**, *7*, 021050.
[39] K. C. Han, arXiv:2009.12239 [quant-ph], **2020**.
[40] R. Iten, R. Colbeck, I. Kukuljan, J. Home, M. Christandl, *Phys. Rev. A* **2016**, *93*, 032318.
[41] K. Yeter-Aydeniz, G. Siopsis, R. C. Pooser, *New J. Phys.*, in press, **2021**, https://doi.org/10.1088/1367-2630/abe63d.
[42] N. V. Tkachenko, J. Sud, Y. Zhang, S. Tretiak, P. M. Anisimov, A. T. Arrasmith, P. J. Coles, L. Cincio, P. A. Dub, arXiv:2009.04996 [quant-ph], **2020**.
[43] K. Sharma, S. Khatri, M. Cerezo, P. J. Coles, *New J. Phys.* **2020**, *22*, 043006.
[44] N. Klco, E. F. Dumitrescu, A. J. McCaskey, T. D. Morris, R. C. Pooser, M. Sanz, E. Solano, P. Lougovski, M. J. Savage, *Phys. Rev. A* **2018**, *98*, 032331.
[45] P. Czarnik, A. Arrasmith, P. J. Coles, L. Cincio, arXiv:2005.10189v2 [quant-ph], **2021**.
[46] T. Proctor, M. Revelle, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Kohout, K. Young, *Nat. Commun.* **2020**, *11*, 5396.
[47] P. H. Leung, *Dissertation*, Duke University, **2020**.
[48] B. Zhang, S. Majumder, P. H. Leung, S. Crain, Y. Wang, C. Fang, D. Debroy, J. Kim, K. R. Brown, arXiv:2104.01119 [quant-ph], **2021**.
[49] W. Lavrijsen, A. Tudor, J. Müller, C. Iancu, W. Jong, arXiv:2004.03004 [quant-ph], **2020**.
[50] B. Eastin, S. T. Flammia, arXiv:quant-ph/0406003v2, **2004**.

© 2021 Wiley-VCH GmbH

Chapter 10

Gate-free state preparation for fast variational quantum eigensolver simulations

This chapter includes pages from our publication [8]

Material from: 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 Information **7**, 155 (2021)

Publisher: Springer Nature.

The contributions to this work are as follows: My contributions were assisting in theoretical development of the results, performing simulations for comparison to gate-based ansatze. O.R.M. and B.T.G. primarily carried out this work. D.P.P., S.E.E., E.B., and N.M. supervised.

npj | Quantum Information

www.nature.com/npjqi

ARTICLE OPEN

Gate-free state preparation for fast variational quantum eigensolver simulations

Oinam Romesh Meitei[1], Bryan T. Gard[2], George S. Barron[2], David P. Pappas[3,4], Sophia E. Economou[2], Edwin Barnes[2] and Nicholas J. Mayhall[1]✉

The variational quantum eigensolver is currently the flagship algorithm for solving electronic structure problems on near-term quantum computers. The algorithm involves implementing a sequence of parameterized gates on quantum hardware to generate a target quantum state, and then measuring the molecular energy. Due to finite coherence times and gate errors, the number of gates that can be implemented remains limited. In this work, we propose an alternative algorithm where device-level pulse shapes are variationally optimized for the state preparation rather than using an abstract-level quantum circuit. In doing so, the coherence time required for the state preparation is drastically reduced. We numerically demonstrate this by directly optimizing pulse shapes which accurately model the dissociation of H_2 and HeH^+, and we compute the ground state energy for LiH with four transmons where we see reductions in state preparation times of roughly three orders of magnitude compared to gate-based strategies.

npj Quantum Information (2021)7:155 ; https://doi.org/10.1038/s41534-021-00493-0

INTRODUCTION

Molecular modeling stands in the juncture of key advances in many important fields including and not limited to energy storage, material designs, and drug discovery. For more than half a century, many seminal works have been reported on the development of theories and methods to enable molecular modeling with high accuracies. Approximate numerical methods which are built on a single Slater determinant reference state, such as density functional theory (DFT), perturbation theory, or coupled cluster, perform well when the amount of electron correlation, ranges from minimal to moderate. However, for systems which are qualitatively governed by electron correlation effects (strongly correlated systems), such approximate methods fail to be sufficiently accurate. While alternative strategies exist, such as density matrix renormalization group (DMRG)[1–3] or selected configurational interaction (SCI) methods[4–13], which can handle strong correlation, these approaches assume that the correlations are either low-dimensional, or limited in scale. Currently, no polynomially scaling classical algorithm exists which can solve for arbitrary molecular ground states.

High dimensionality of the wavefunction is ultimately the core reason for the exponential cost of modeling electronic structure. Even before optimization, simply storing the wavefunction on classical hardware quickly becomes a bottleneck. This is because one typically represents the wavefunction in a basis of "classical" states, or basis states which have a direct product structure (Slater determinants, occupation number vectors, or even tensor product states[12]). In this classical basis, the exact (and generally entangled) state of the system is represented as an exponentially (or factorially, if a basis of Slater determinants is used to span a target projected spin subspace) large vector of coefficients weighting the corresponding classical basis states.

Quantum computing offers a radical departure from this computational strategy. As quantum systems themselves, the state of a quantum processing unit (QPU) is also a vector in a Hilbert space of identical dimension to the molecular problem. This ability to perform a one-to-one mapping between vectors in the Hilbert space containing the molecule's electronic wavefunction and those in the state space accessible to a QPU means that with enough control over the QPU, it should be possible to take the vector corresponding to the molecular wavefunction and realize it on the QPU, avoiding altogether the requirement to work with an exponentially large vector of coefficients. Once the QPU is prepared into the state corresponding to the target molecule, any molecular observable (energy, dipole moment, etc) can be obtained by measuring the corresponding operator on the QPU.

In order to turn this strategy into an algorithm, one needs a procedure for realizing the target molecular wavefunction on the QPU. As the leading quantum algorithm for molecular simulation on near-term devices, the variational quantum eigensolver (VQE)[14] provides an efficient procedure for this purpose. In VQE, one defines a parameterized quantum circuit comprised of tunable gates, and then optimizes these gates using the variational principle, minimizing the energy of the molecular Hamiltonian. This parameterized quantum circuit (referred to as an ansatz) defines the variational flexibility (and thus the subspace reachable on the QPU) of the algorithm.

State-preparation circuits with more parameters generally have more variational flexibility, but come with the cost of having deeper quantum circuits and more difficult optimization. This cost can be significant. Current and near-term quantum computers are classified as noisy intermediate scale quantum (NISQ) devices due to the presence of short coherence times, system noise, and frequent gate errors[15]. Because each gate has limited fidelity, the success probability for a sequence of gates decreases exponentially with circuit depth. Even if gates could reach unit fidelity, the finite coherence times of a NISQ device still limits the number of gates that one can apply in a circuit which, in turn, limits the accuracy of the molecular VQE simulation. Although VQE is relatively robust in the presence of noise and errors in certain cases[16], the critical limitation preventing larger scale experiments is the accurate implementation of deep circuits. The goal of

[1]Department of Chemistry, Virginia Tech, Blacksburg, VA 24061, USA. [2]Department of Physics, Virginia Tech, Blacksburg, VA 24061, USA. [3]Physics Department, University of Colorado, Boulder, CO 80309, USA. [4]National Institute of Standards and Technology, Boulder, CO 80305, USA. ✉email: nmayhall@vt.edu

finding parameterized circuits which minimize the circuit depth and maximize the accuracy has led to a number of approaches such as hardware efficient ansätze[17], physically motivated fixed ansätze[18–22], and adaptive ansätze [23–25].

Each of these mentioned approaches essentially aims to find the shortest quantum circuit for preparing a state to within an acceptable error. However, even if one succeeds at finding the "best" circuit, it is likely that its execution time will still exceed the coherence time of a NISQ device. In order to make meaningful progress toward larger problems, one can imagine three strategies: (i) device improved error mitigation techniques, (ii) significantly increase the ratio of coherence to gate time, or (iii) simply abandon gates and directly optimize the electronic controls. Of these, the latter strategy appears to be the most feasible on near-term NISQ devices.

In this paper, we explore the possibility of performing gate-free VQE simulations by replacing the parameterized quantum circuit with a direct optimization of the laboratory-frame analog control settings. In the following sections, we argue that quantum control techniques are likely to be better suited for fast VQE state preparation than the more conventional circuit-based approaches on NISQ devices. We first provide a detailed overview of circuit-based VQE, then introduce our proposed strategy, ctrl-VQE, then discuss initial results along with a strategy to avoid over-parameterization, and finally compare to gate-based ansätze. Several technical aspects, numerical results, and additional discussion are provided in the Supplementary Methods and Supplementary Discussion.

RESULTS AND DISCUSSION

Dissociation of diatomic molecules

In this section, we explore the ability of ctrl-VQE to reproduce full configuration interaction (FCI) (exact diagonalization) bond dissociation energy curves for two small example diatomic molecules. Here, we demonstrate the performance of our approach by computing the ground state molecular electronic energy along the bond dissociation of the H_2 molecule and the HeH^+ molecular ion. Although small, these molecules have been used as benchmarks for quantum algorithms in recent years[17,20,21,26–35]. These two molecular examples are chosen because they exhibit different behaviors during bond dissocation in terms of their electronic structure. As two-orbital, two-electron problems, these would naturally be modeled on four qubits following the commonly used Jordan-Wigner transformation[36]. However, to make the pulse simulations more computationally efficient, we have used the parity mapping instead in which two qubits are diagonal and can be removed from consideration. Details are given in Bravyi et al.[37], and we use the implementation for mapping and pruning in Qiskit software [38].

As a prototypical example of a homolytic dissociation,

$$\begin{gathered} H_2 \rightarrow H^{\bullet} + H^{\bullet} \\ \text{Dynamic correlation} \rightarrow \text{Static correlation} \end{gathered} \tag{1}$$

as H_2 dissociates, the ground state moves from being dominated by a single Slater determinant to increasingly multiconfigurational, due to the shrinking HOMO-LUMO gap. As a result, the accuracy of a mean-field treatment (e.g., HF) diminishes as the bond length is stretched.

In stark contrast, HeH^+ becomes easier to model as the bond distance increases. The reason is that, being the strongest acid and, interestingly, the first molecule formed in the universe[39], dissociation is a heterolytic deprotonation,

$$\begin{gathered} HeH^+ \rightarrow He + H^+ \\ \text{Dynamic correlation} \rightarrow \text{No correlation}* \end{gathered} \tag{2}$$

such that both the reactants and the products are both closed shell and well represented by a single Slater determinant. If fact, in a minimal basis set (i.e., the STO-3G basis set used in this paper), the products have exactly zero electron correlation energy because the H^+ ion has no electrons, while the He atom has no empty virtual orbitals into which electrons can be promoted. As a result, HF becomes exact at dissociation. Of course, in larger basis sets, the He atom would have some dynamic correlation, hence the * in Eq. (2).

The dissociation curves of the H_2 and HeH^+ molecular systems produced with ctrl-VQE using the simple square pulses are presented in Figs. 1 and 2 respectively. Consistent with the physical descriptions above, the HF state moves quickly away from the FCI ground state with bond dissociation for H_2, while the opposite is true for HeH^+, where the HF state gradually converges to the exact FCI ground state with increasing bond distance. ctrl-VQE reproduces the FCI bond dissociation curve of H_2 and HeH^+ with high accuracies. The maximum difference from the FCI energy along the dissociation curve is 0.03 mHa for both H_2 and HeH^+, and the average error is 0.002 mH in both the cases. More detailed information including the pulse parameters, characteristics, and molecular energies along the dissociation curves are provided in the SI. The resulting states have a good overlap with the exact FCI ground states (99% in all cases). For H_2, it is possible for the overlap to deviate at long bond distances due to degenerate singlet and triplet states. As the bond distance increases, the singlet ground state of H_2 becomes degenerate with the lowest triplet state, making it possible to generate a superposition of a singlet and triplet. It is possible to converge to the exact ground singlet state by supplementing the objective function with a penalty term proportional to the total spin operator (see SI for details).

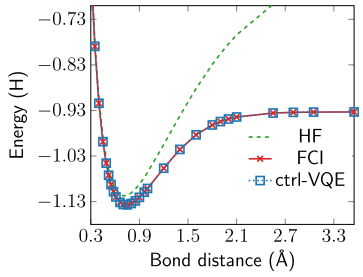

Fig. 1 Bond dissociation curve of H_2 molecule. ctrl-VQE energies are computed using square pulses with two time segments.

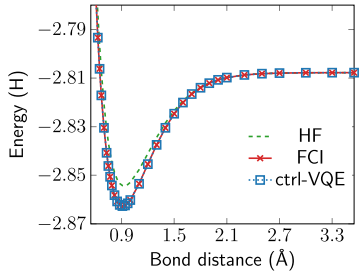

Fig. 2 Bond dissociation curve of HeH^+ molecule. ctrl-VQE energies are computed using square pulses with two time segments.

Effect of electron correlation on pulse duration

In any VQE, one is typically interested in finding a useful balance between accuracy (needing deep circuits) and noise (needing shallow circuits). As such, molecular simulations of strongly correlated molecules are intrinsically more difficult as deeper circuits are required, which increases problems from noise and gate errors. An analogous balance is targeted in ctrl-VQE, where one would hope to obtain sufficiently accurate results with as short of a pulse time as possible. Because entanglement cannot be created instantly[40], we expect molecules with strong correlation to require longer pulses than simpler molecules.

In order to examine this relationship, in Fig. 3 we plot the duration of the shortest pulses we were able to obtain at each molecular geometry, with H_2 (HeH^+) shown in red (purple). Referring back to Eqs. (1) and (2), one would expect that as the bond distance is increased, H_2, needing more entanglement, would in turn require increasing pulse durations, whereas pulses for HeH^+ would get increasingly fast as the bond is stretched. This trend is indeed observed.

The total pulse duration significantly decreases for the dissociation of the HeH^+ molecular ion, whereas it significantly increases for the H_2 molecule. Note that the initial state and the final target state for HeH^+ become degenerate with increasing bond distance (above 2.0 Å). Thus, the Hartree Fock state ($|01\rangle$) is a good approximation to the exact FCI ground state, and only a slight modification to the initial state is required to well approximate the ground state. With ctrl-VQE, the total pulse duration at the far bond distances of HeH^+ are only 1.0 ns. This directly reflects the efficiency of the method presented in this work. In a gate-based compilation method, generally one would still construct an ansatz comprised of costly two-qubit gates, even though the target state is very nearly approximated by the initial state.

The total pulse duration at the dissociation limit for H_2 is significantly longer than near the equilibrium bond distance. Although the HF molecular energy increases monotonically beyond the equilibrium point, the same is not observed for the pulse duration. This is suggestive of the pulse durations reflecting the different dynamics along the bond dissociation of the two molecular systems (see SI for a more detailed analysis).

Leakage to higher energy states

In Figs. 4 and 5, we show the energy and leakage as a function of evolution time for various choices of the hyper-parameter, T, (the total evolution time). In each of the two figures, the top panel displays the behavior of the molecular energy as the state evolves from the reference Hartree-Fock state to the converged trial state, and the bottom panel displays the total population outside of the $|0\rangle$ and $|1\rangle$ computational states for each qubit (leakage).

We see that for each value of T, the molecular energy of both H_2 and HeH^+ never decreases toward the exact energy monotonically, but rather increases initially, and then either decreases or oscillates before rapidly converging to the exact energy. While each of the optimized pulses do indeed generate suitable transmon dynamics which accurately produce the molecular ground state at the specified time T, the different pulses are not each equally favorable. From the bottom panels of Figs. 4 and 5, we see that some pulses create much more leakage than other pulses. High levels of leakage will likely require a larger number of shots (pulse executions) to get precise determinations of the expectation values, since a higher portion of the shots end up in excited states which are either discarded, or collected for normalizing the results. However, with leakage of only 10%, we anticipate only needing a commensurate increase in the shot count to compensate for post-selection, and so we do not expect this to be a fundamental limitation of the approach. Alternatively, one can use an unnormalized energy in the objective function. This approach naturally constrains the solutions to simultaneously minimize leakage, as any leakage necessarily penalizes the associated pulse. Results using an unnormalized energy in the objective function are comparable to those shown here, and are given in the SI.

Noise analysis: imprecise control

All the analysis thus far has assumed perfect control over the driving parameters used to generate approximate ground states of chemical Hamiltonians. However, there are always uncertainties

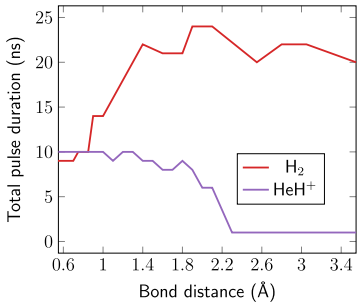

Fig. 3 Total pulse duration. Shown as a function of bond distance of H_2 and HeH^+ molecular systems. Square pulses were used to generate the ctrl-VQE molecular energies in Figs. 1 and 2. The shortest pulse duration at each geometry is plotted here.

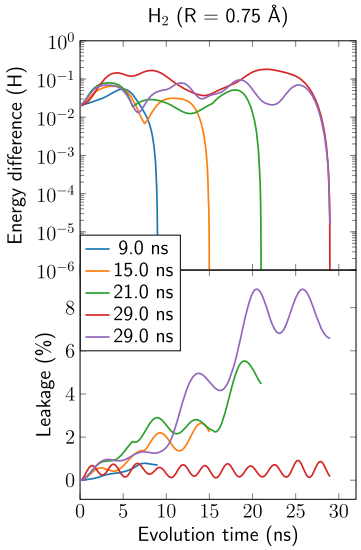

Fig. 4 Error in energy and leakage for H^2. Energy difference shown between ctrl-VQE and FCI (top) and leakage (bottom) along the time evolution steps with optimal pulses of different durations for H_2 with a bond distance of 0.75 Å. The red and purple lines with $T =$ 29 ns are distinct solutions to the same optimization. Optimized pulse parameters are provided in the SI.

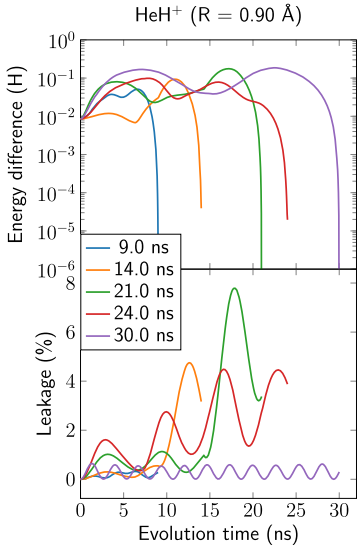

Fig. 5 Error in energy and leakage for HeH^+. Energy difference shown between ctrl-VQE and FCI (top) and leakage (bottom) along the time evolution steps with optimal pulses of different durations for HeH^+ with a bond distance of 0.90 Å. Optimized pulse parameters are provided in the SI.

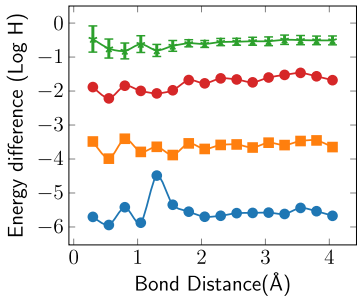

Fig. 6 Effect of imprecise control parameters. Energy error for the dissociation of H_2 is shown. Instead of using the exact, optimized parameters, we use parameters which have added Gaussian noise uncertainty. For each optimal parameter, we randomly sample 100 points each from a Gaussian distribution centered on its optimal setting and with a standard deviation of $\sigma = \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ which correspond to the blue, orange, red and green curves respectively. The resulting energies for these 100 trials are then averaged. Error bars are given for each point but are smaller than some markers.

when attempting to implement the optimal driving parameters. In Fig. 6 we simulate the effect of imprecisely implementing these parameters. For each set of optimal drive parameters found by the ctrl-VQE protocol, we produce 100 samples of noisy parameters taken from Gaussian distributions that are centered about the optimal parameter values:

$$\theta_n \leftarrow \exp\left(-\frac{(\theta-\theta_o)^2}{2\sigma^2}\right), \tag{3}$$

where θ_n is a noisy version of the optimal parameter θ_o. These in turn are used to produce 100 noisy energy samples. We then average these 100 samples and find the resulting energy difference relative to the target ground state energy. With this applied noise model, we can see that even imprecise settings can still be used to achieve accuracies below a 10^{-4} energy error. As shown in the red and green curves in Fig. 6, only when the noise becomes higher than $\sigma = 0.01$ (where errors reach a magnitude of 10 MHz or 1 ns) do the energy differences become significant enough to affect the performance of ctrl-VQE. Since these errors are of the same order as the parameters themselves, they do not hinder the realistic implementation of ctrl-VQE. We do not include any explicit noise due to finite decoherence (T_1) or dephasing (T_2^*) since the pulses we find are many orders of magnitude shorter than the typical time scales for these effects[41]. To confirm this, we ran simulations using a Lindblad master equation that includes these effects. Our objective function did not change significantly due to our short time scales.

Comparison with circuit compilation techniques

Although several ansätze have been proposed to achieve shorter circuits, even the most compact approaches involve too many gates to implement on current hardware for all but the smallest molecules. In order to reduce the time spent during the state-preparation stage, and thus the coherence time demands on the hardware, circuit compilation techniques have been designed to take a given quantum circuit and execute it either with fewer gates or by reoptimizing groups of gates for faster execution.

To execute the gate-based VQE (described in Section III A) experimentally, the gates in a circuit are compiled to sequences of analog control pulses using a look-up table that maps each elementary gate to the associated predefined analog pulse. The sequence of control pulses corresponding to the elementary gates in the quantum circuit are then simply collected and appropriately scheduled to be executed on the hardware. As such the compilation is essentially instantaneous, making the gate-based compilation technique well suited for VQE algorithms where numerous iterations are performed.

From a compactness perspective however, this gate-based compilation is far from ideal, resulting in total pulse durations which are much longer than what might be obtained with optimized compilation techniques. As is obvious from the one-to-one translation of gates to pulses, the overall circuit structure is typically not considered in gate-based compilations. Thus one may naturally be inclined to seek an optimal pulse sequence for the entire circuit. This has motivated compilation algorithms where control pulses are optimized for the target circuit, using numerical optimal control techniques such as gradient ascent pulse engineering (GRAPE)[42] for partially optimal compilation[43]. However, because the GRAPE algorithm itself is highly non-trivial, the compilation latency for each iteration is of critical concern. Two of the GRAPE-based techniques are briefly described below; see Gokhale et al.[43] for a detailed discussion.

The GRAPE compilation technique employs an optimal control routine which compiles to machine-level sequences of analog pulses for a target quantum circuit. This is achieved by manipulating the sets of time-discrete control fields that are executed on the quantum system. The control fields in the optimal routine are updated using gradients (see Boutin et al.[44] for use of analytical gradients) of the cost function with respect to the control fields. GRAPE compilation achieves up to 5-fold speedups compared to standard gate-based compilation. However, such

speedup comes with a substantial computational cost. This amounts to long compilation latency, making this approach impractical for VQE algorithms in which multiple iterations of circuit-parameter optimizations are performed. The GRAPE-based compilation also suffers from limitations in the circuit sizes it can handle[45–47].

On the other hand, partial compilation techniques achieve significant pulse speedups by leveraging the fast compilation of standard gate-based techniques together with the pulse speedups of GRAPE-based compilations. Two flavors of such an approach were reported in Gokhale et al.[43]. Both divide the whole circuit into blocks of subcircuits. In the so-called strict partial compilation, the structure of quantum circuits used in quantum variational algorithms are exploited to only perform GRAPE-based compilation on fixed subcircuits that are independent of the circuit parametrization. The optimal pulses using the GRAPE compilation techniques for the fixed blocks are pre-computed and simply concatenated with the control pulses from the gate-based compilations for the remaining blocks of the circuit. Thus, the compilation speed is comparable to the gate-based compilations in each iteration, but this method also takes advantage of pulse speedups from GRAPE-based compilations. As one may expect, the pulse speedups heavily depend on the circuit depth of the fixed blocks.

In the other flavor, flexible partial compilation, circuits are blocked into subcircuits, each with a single variational parameter, ensuring the subcircuit blocks have low depth. Hyperparameter optimizations are performed for each subcircuit, and they are utilized to find optimal pulses for the circuit blocks. It is noteworthy to mention that in the flexible partial compilation, compilation latency is reduced significantly (~80x, see Gokhale et al.[43]) by tuning the hyperparameters of the circuit blocks to speed up the convergence of optimal control in GRAPE. Flexible partial compilations achieve significant pulse speedups as compared to strict partial compilations. However, in spite of the high pulse speedups, the flexible partial compilation technique still diverges from the extremely fast compilation time of the standard gate-based methods. Thus, the algorithm still suffers from compilation latency, albeit to a lesser degree compared to GRAPE-based compilation.

Although both of these GRAPE-based compilation techniques and ctrl-VQE share a direct pulse-shaping step, they fundamentally differ because ctrl-VQE is not a compiler. In contrast to full or partial compiling, we make no reference to a circuit or to implementing any specific unitary. Of course, the control pulses we find do implement some equivalent circuit or unitary (and we analyze these in Sec. II F), but we have no need to ever characterize the unitary. In fact, because ctrl-VQE only targets a single state, a large family of unitaries (defined by the behavior in the spaces orthogonal to the initial and final states) exist which minimize our objective function. As such, many possible solutions exist, with no immediate preference given to any one.

Decompiled control pulses

In the previous sections, the efficiency of ctrl-VQE using optimized control pulses at the device level was demonstrated. The short pulse durations imply that applications to larger molecules might have even more significant speedup since the number of CNOTs in most VQE ansätze increases quickly.

Although ctrl-VQE is performed using no state preparation circuit, the unitary created by the time dynamics of the applied pulse can be decomposed into gates, allowing one to analyze the circuit. This is essentially running a compiler in reverse, or "decompiling" the pulse to yield a state preparation circuit. With this decompiled circuit at hand, we can evaluate the time it would take to execute the optimized pulse as a traditional circuit. By comparing this time to that of the pulse duration, one has a clean benchmark for quantifying the overhead associated with gate-based state preparation. This decompilation can be done in two ways. In the first approach, we simply compute the matrix representation of the evolution operator generated by the optimized ctrl-VQE pulse. For the two-qubit case in focus here, a quantum circuit is then constructed using the KAK decomposition technique. For a detailed description of the technique, we refer the reader to Williams[48].

In the second approach, an arbitrary circuit corresponding to the state vector from ctrl-VQE is constructed and then transpiled to obtain a shorter circuit depth. In this way we obtained gate durations on IBMQ hardware (mock Johannesburg device available in Qiskit software). The KAK decomposition, arbitrary circuit construction, transpilation, and the mock circuit compilations were performed using Qiskit software[38]. The quantum circuits are illustrated in Fig. 7. The corresponding pulse durations are 1202 ns for the circuit obtained using the KAK decomposition and 825 ns for the circuit obtained from transpilation. The state vector used in this study was for the H_2 molecule at 0.75 Å, and the corresponding pulse duration was 9 ns with ctrl-VQE. This clearly demonstrates that state preparation circuits can be much deeper than is necessary.

Adaptive update of pulse parameterization

The manner in which a pulse is parameterized (in particular, the number of variational parameters) will significantly impact

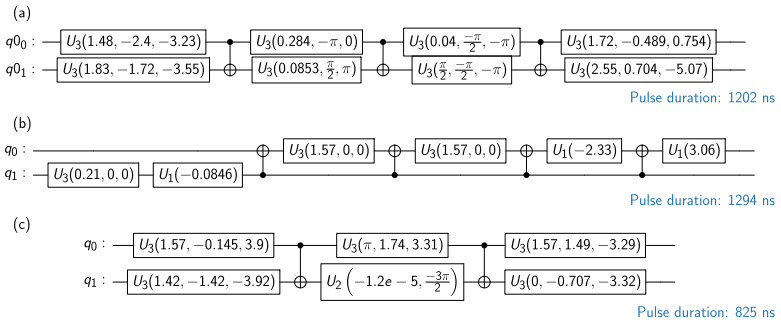

Fig. 7 Circuits from decompiled control pulses. Illustration of circuits constructed using the unitary and state vectors from ctrl-VQE. **a** Circuit corresponding to unitary obtained from a KAK decomposition. **b** An arbitary circuit corresponding to the state vector. **c** Transpiled version of circuit (**b**).

experimental performance. An over-parameterization of the pulse leads to difficulties in experimental optimization. In response, we seek a means to limit the number of parameters. In this section, we describe a scheme to avoid over-parameterizations and arrive at the minimal number of pulse parameters to achieve a target accuracy.

Unlike in the previous sections where we chose a fixed number of time segments, in this section we propose an adaptive algorithm which slowly grows the number of segments (and thus parameters). We begin with a single time-segment square pulse (constant amplitude throughout the time evolution), the pulse is then iteratively sliced at random intervals such that the number of time segments systematically increases. This adaptive scheme is outlined below:

1. Initialize a parameterized square pulse with $n = 1$ time segment and perform the variational pulse optimization.
2. Divide the pulse obtained from the previous step at a randomly choosen time. The square pulse now has $n = 2$ time segments.
3. Perform pulse optimization using the pulse shape from the previous step as an initial guess.
4. Divide the largest time segment in the square pulse obtained from the previous step into two at a randomly chosen time. This increases the number of segments from n to $n + 1$.
5. Perform pulse optimization using the pulse shape from the previous step as an initial guess.
6. Repeat Steps 4 and 5 until the energy obtained is sufficiently converged.

Note that drive frequencies for the pulses are also optimized in the above procedure. Irrespective of the pulse shape used, only a single drive frequency is used for each qubit. The switching times for the square pulses with more than one time segment are not optimized and remain fixed in Steps 3 and 5. For N transmons with n pulse time segments each, the total number of parameters to optimize is $N(n + 1)$. In the following simulations, we constrain the amplitudes and drive frequencies within the ranges ±40 MHz and $\omega_k \pm 3\pi$ GHz, respectively.

Following the above strategy of adaptively increasing the number of time segments in the square pulses, we obtained optimal pulses for H_2, HeH^+ and LiH with total pulse durations of 9, 9, and 40 ns respectively. Figure 8 plots the energy difference relative to FCI with increasing number of time segments in the optimal pulse. For H_2 and HeH^+, chemical accuracy is obtained with a single-time-segment square pulse on each qubit, while for LiH, three time segments are required to achieve chemical accuracy. The molecular energy converges for square pulses with two time segments for H_2 and HeH^+ and with five time segments for LiH molecules. The error in ground state energy of LiH molecule (Li–H R = 1.5 Å) computed with ctrl-VQE along with pulse duration, leakage, and overlap with the exact FCI state is tabulated in Table 1. To illustrate the sequential pulse refinement from iteration to iteration, the square pulse shapes at each adaptive step for LiH are shown in Fig. 9.

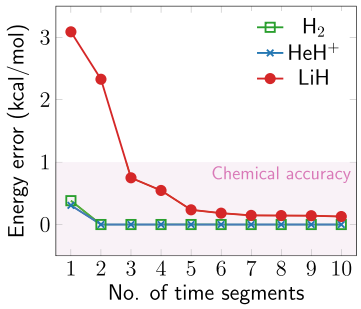

Fig. 8 Convergence of pulse parameterization. Energy difference of ctrl-VQE relative to FCI with increasing number of time-segments in the square pulse used.

In the case of LiH, we see that the accuracy is not of the same order as in the other two diatomic molecules. To further examine this, we perform over 100 pulse optimizations on the square pulse with 10 time segments obtained from the adaptive strategy by retaining only the switching times and starting with random amplitudes and drive frequencies. We use a large number of initial guesses to reduce the likelihood of getting trapped in a local minimum. We find that the highest accuracy is achieved when we use the pulse obtained from the adaptive strategy above. This suggests that we are nearing the accuracy limit of LiH using a 40 ns pulse.

While the above adaptive strategy aims at practical experimental implementation, the cost of classical simulation remains the same irrespective of the number of time-segments when analytical gradients for the pulse amplitudes are used. The cost of evaluating an analytical gradient is only about 2.5 times the cost of an energy evaluation. Here, the analytical gradient is obtained at each trotter step of time evolution and adapted for the pulse shaped considered. Details along with the derivation are provided in the SI.

Comparison to gate-based ansätze

Now we directly compare the results of our pulse-based technique with gate-based variational ansätze. We use calibration data from an IBMQ device (mock Johannesburg device available in Qiskit Terra version 0.14.2) to compute the duration of the circuits used to prepare trial wavefunctions. We consider the RY and UCCSD ansätze, which are capable of producing the exact ground state.

For both H_2 and HeH^+, the RY ansatz requires 1 CNOT, and the UCCSD requires 2 CNOTs. The RY ansatz has four parameters, and the UCCSD ansatz has three parameters. The total pulse time of the RY ansatz is 519 ns, and the UCCSD ansatz is 825 ns. For the LiH molecule, the RY and UCCSD ansätz require 9 and 245 CNOTs respectively. The RY ansatz has 16 parameters whereas the UCCSD ansatz has 8 parameters. The total pulse duration of the RY ansatz is 3485 ns, and that of the UCCSD ansatz is 82,169 ns for LiH. In each case, the time to execute the circuit is significantly longer than the time required to apply the pulses using ctrl-VQE, which is 9 ns for H_2 and HeH^+, and 40 ns for LiH. The T_1 and T_2 times of the IBMQ device used range from 45,150 to 115,217 ns (average 71,262 ns) and from 48,173 to 107,025 ns (average 74,490 ns) respectively.

Table 1. Performance of ctrl-VQE in computing the ground state energy of the LiH molecule (Li–H R = 1.5 Å). The energy error relative to FCI, the leakage to higher energy states, the total pulse duration, and the overlap with the exact FCI state are shown. The data correspond to the pulse with five time segments from Fig. 8.

Method	Energy	Error
Hartree-Fock	−7.8633576	0.0177
ctrl-VQE	−7.8806399	0.0004
FCI	−7.8810157	–
Pulse duration	40 ns	–
Leakage	0.59%	–
Overlap	99.38%	–

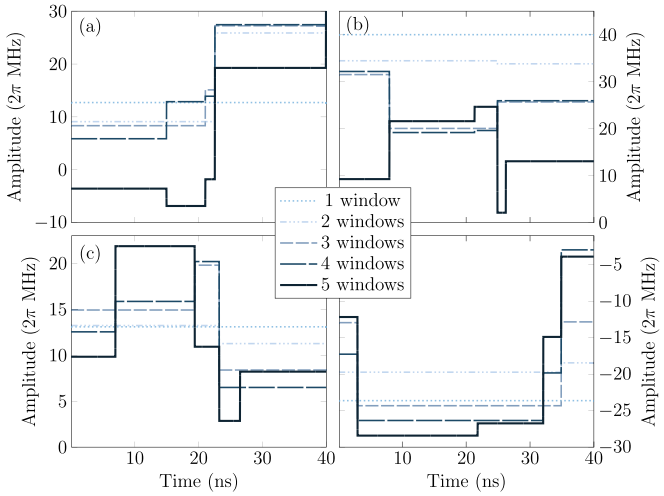

Fig. 9 Adaptive pulse sequences for LiH molecule. Square pulse shapes obtained at each adaptive step, starting from one time window and adaptively dividing the time window up to five-segments. The accuracy increases with each division of the pulse. Pulse shapes are shown for the pulses on **a** qubit 1, **b** qubit 2, **c** qubit 3, **d** qubit 4.

This means that the pulse duration for LiH using the UCCSD ansatz is already reaching the limit set by decoherence times of the device, while the pulse duration obtained with ctrl-VQE is well within the decoherence time.

We note here that this comparison needs to be approached with a bit of caution, as our simulated device has slightly different parameters than the IBM-Q devices used for the circuit timings. However, we do not expect that our results would change significantly if we had access to the full set of exact parameters of the device (including anharmonicities, bus frequencies, and couplings), because previous works using the present device parameter regime[49,50] have demonstrated universal gate sets with single and two-qubit operations commensurate with those of the IBM-Q devices. As such, we expect the ctrl-VQE pulse times reported in this study to be close to what one would get by directly running the calculations on IBM-Q devices.

Summary and future outlook

In summary, we have presented an alternative new variational quantum algorithm which is fundamentally different from the existing quantum algorithms for molecular simulation. The quantum circuit used for state preparation in standard variational algorithms is entirely replaced by a hardware-level control routine with optimized pulse shapes to directly drive the qubit system. As such, we demonstrate that VQE applications do not need to be confined to two level systems (qubits), which are required for general quantum computing. This opens a possibility of faster ansatz preparation within a time window defined by coherence times of NISQ devices.

The efficacy of the presented method is numerically demonstrated by modeling the bond dissociation of two diatomic molecules, H_2 and HeH^+. The maximum error from the exact FCI energy is well within chemical accuracy (0.02 kcal/mol) for both molecular systems throughout the bond dissociation. ctrl-VQE captures the important electron correlation effects involved in the bond dissociation of the two molecular systems, as reflected in the pulse durations along the bond dissociation.

To demonstrate the application of ctrl-VQE to a larger molecular system, we have also computed the energy of the LiH molecule at a single bond distance. An adaptive scheme which slowly increases the variational flexibility of the pulse parameterization was used to determine the minimal number of pulse parameters required to acheive a target accuracy, ultimately simplifying experimental implementation and avoiding over-parameterization. The results demonstrate that pulse shapes with relatively modest numbers of parameters achieve convergence in the energy even as the problem size increased. The approach yields significant state-preparation speedups for VQE as compared to standard gate-based ansätze. The short pulse durations minimize losses due to decoherence and dephasing, which is a step toward enabling more accurate VQE simulations on larger, strongly correlated systems. ctrl-VQE can be viewed as a lowest-level quantum variational algorithm.

Because ctrl-VQE operates directly on the hardware, classical simulations of the algorithm are even more computationally expensive than typical VQE emulations, as one needs to not only solve the time-dependent Schrödinger equation, but also optimize the driving Hamiltonian. As such, the systems we have been able to study so far are very small. In future work, we will develop an improved implementation to study the behavior of larger systems and with more sophisticated constraints on the pulse shape ($\Omega_n(t)$). Improved software will also make it possible to numerically study the impact that device controllability limitations[51] have on the amount or type of electron correlation able to be described.

METHOD

Variational quantum eigensolver

The VQE algorithm aims to leverage classical resources to reduce the circuit depth required for molecular simulation. The algorithm finds optimal rotation angles for a parameterized quantum circuit of fixed depth by variationally minimizing the energy of a target molecule, which is obtained by repeated state preparation and measurement cycles. In order to account for the distinguishability of qubits, we start by transforming the

second quantized electronic Hamiltonian into an equivalent form involving non-local strings of Pauli spin operators, $\hat{o}_i$:

$$\begin{aligned}\hat{H}^{\text{molecule}} &= \sum_{pq} h_{pq}\hat{p}^\dagger\hat{q} + \tfrac{1}{2}\sum_{pqrs}\langle pq|rs\rangle\hat{p}^\dagger\hat{q}^\dagger\hat{s}\hat{r} \\ &= \sum_i \hat{o}_i h_i,\end{aligned} \tag{4}$$

where h_i is a sum of molecular one or two-electron integrals, and $\hat{p}$ operators are are fermionic annihilation operators. Several such transformations exist, such as the Jordan–Wigner[36], Bravyi–Kitaev[52], or parity transformation[53]. The main steps in VQE are defined as follows:

1. Precompute all h_i values, and transform terms in the Hamiltonian operator into the desired qubit representation.
2. Choose an ansatz for the problem which defines the variables $(\vec{\theta})$ to optimize. Assuming one starts from the Hartree-Fock (HF) reference state, this involves predefining a parameterized circuit which provides enough variational flexibility to describe the molecule's electron correlation effects. Many ansätze have been proposed, several of which are variants of the original proposal using the unitary coupled-cluster (UCC) ansatz[54,55].
3. Choose an initial set of parameter values, $\vec{\theta} = \vec{\theta}_0$. These can be initialized to zero, chosen randomly, or if appropriate, chosen based on some classical precomputation such as using MP2 parameters to start a UCCSD optimization.
4. Using current parameters, $\vec{\theta}$, repeatedly execute the circuit, each time performing an individual measurement of one of the operators, $\hat{o}_i$, in $\hat{H}$ or a set of mutually commuting operators. After a sufficient number of circuit executions (shots), the averages of the resulting data converge to the expectation values of the operators such that the average molecular energy can be obtained by multiplication with the one and two-electron integrals,

$$E(\vec{\theta}) = \sum_i h_i \left\langle \psi(\vec{\theta}) \middle| \hat{o}_i \middle| \psi(\vec{\theta}) \right\rangle. \tag{5}$$

5. Check convergence. If the average energy has decreased by a small enough value determined to be converged, exit. If the energy is not yet converged, update $\vec{\theta}$, and return to step 4.

Various approaches have been proposed, each differing in the details of state-preparation[18,24,56–59], and ways to reduce the number of circuits required to compute expectation values[60–63].

Control variational quantum eigensolver

In this section, we present an alternative to the gate-based VQE algorithm, replacing the parameterized state-preparation circuit with a parameterized laboratory-frame pulse representation, which is optimized in an analogous manner, but with the benefit of a much faster state preparation, opening up the possibility of more accurate simulations on NISQ devices. All other aspects of VQE (i.e., measurement protocols) are essentially the same. Using the molecular energy as the objective function to minimize, the pulse parameters are optimized using the variational principle. This general strategy, which we refer to as control variational quantum eigensolver (ctrl-VQE), is outlined as follows:

1. As done in any regular VQE, compute the one- and two-electron integrals and transform the molecular Hamiltonian into a qubit representation, for example using Jordan–Wigner, Parity or Bravyi–Kitaev mappings. This defines the objective function to minimize, $\langle \hat{H}^{\text{molecule}} \rangle$.
2. Define a fixed pulse representation (e.g. square pulses, sum of Gaussian pulses, etc.). Parametrize the chosen pulse representation, and initialize parameters.
3. Choose an initial state for the qubit(s) system. HF is a good choice for the molecular problems studied here. Controls are assumed to be in the form of direct drives on each qubit.
4. Measure the expectation value of the objective function $\langle \hat{H}^{\text{molecule}} \rangle$ on the quantum device.
5. Using a classical optimization routine, determine the pulse parameters for the next measurement on the quantum device.
6. Repeat until a target convergence threshold is met. If the chosen parameterized pulse can fully specify the target Hilbert-space, then the optimal pulse finds the minimum energy state.

We note here that the optimization used in this work excludes the total pulse duration. This is fixed throughout the optimization routine, and only the pulse parameters such as the amplitudes or the frequencies are directly optimized. As such, the total pulse time enters the algorithm as a "hyper-parameter", which can be optimized in an outerloop if desired. In the square pulses considered in this work, the time segments are also optimized. Unless stated otherwise, optimal pulses correspond to pulse shapes that are optimal with a given fixed total pulse duration.

Unlike universal quantum computing algorithms, ctrl-VQE occurs at the hardware-level, and any simulation must refer to a specific physical platform. For this work, we choose a well-established transmon platform with the following 1D device Hamiltonian[41]:

$$\begin{aligned}\hat{H}_D = &\sum_{k=1}^{N}\left(\omega_k \hat{a}_k^\dagger\hat{a}_k - \tfrac{\delta_k}{2}\hat{a}_k^\dagger\hat{a}_k^\dagger\hat{a}_k\hat{a}_k\right) \\ &+ \sum_{\langle kl\rangle} g(\hat{a}_k^\dagger\hat{a}_l + \hat{a}_l^\dagger\hat{a}_k).\end{aligned} \tag{6}$$

where $\langle kl\rangle$ indicates that the summation is over pairs of coupled qubits. Here, we assume that the qubits are arranged in a closed linear chain and have always-on direct couplings, where $\hat{a}_k$ is the bosonic annihilation operator for the kth transmon, and ω_k, δ_k, and g are the resonant frequency, anharmonicity, and constant coupling rate, respectively. Furthermore, each transmon in Eq. (6) formally supports an infinite number of states. However, in our simulations we necessarily approximate this system by a finite number of levels (three unless otherwise stated) per transmon. We tested the accuracy of this by adding more levels and found that the results did not change significantly. The parameters used in this work are chosen to be typical of values found in current superconducting transmons, and are provided in Table 2. We find that our results below do not qualitatively depend on the frequency difference between the qubits, and in the SI we provide a comparison of this current device to one with a larger detuning between the transmons. In order to drive the device, we apply a time-dependent field, with separate controls on each qubit such that within the rotating wave approximation, the control Hamiltonian is expressed as:

$$\hat{H}_C = \sum_{k=1}^{N}\Omega_k(t)(e^{i\nu_k t}\hat{a}_k + e^{-i\nu_k t}\hat{a}_k^\dagger). \tag{7}$$

where $\Omega_k(t)$ is the real-valued, time-dependent amplitude of the drive, and ν_k is the frequency of the field. The system therefore evolves under the total Hamiltonian:

$$\hat{H} = \hat{H}_D + \hat{H}_C(t, \Omega_n(t), \nu_n). \tag{8}$$

By moving to the interacting frame, the final ansatz has the following form:

$$\left|\psi^{\text{trial}}(\Omega_n(t), \nu_n)\right\rangle = \mathcal{T}e^{-i\int_0^T dt\hat{H}_C(t,\Omega_n(t),\nu_n)}|\psi_0\rangle. \tag{9}$$

where $|\psi_0\rangle$ is the VQE reference state (e.g., the HF state), $\{\Omega_n(t), \nu_n\}$ are the variational parameters for the ansatz, T is the total pulse time, and $\mathcal{T}$ is the time-ordering operator. Note that although the control Hamiltonian above only has single-qubit terms, the device itself has inter-qubit couplings (Eq. (6) with strength g), which create an entangling control Hamiltonian in the interacting frame:

$$\hat{H}(t)_C = e^{i\hat{H}_D t}\hat{H}_C(t)e^{-i\hat{H}_D t}. \tag{10}$$

Table 2. Device parameters appearing in Eq. (6) for a four-transmon system.

	Transmon			
	1	2	3	4
ω	4.8080	4.8333	4.9400	4.7960
δ	0.3102	0.2916	0.3302	0.2616
	$1 \leftrightarrow 2$	$2 \leftrightarrow 3$	$3 \leftrightarrow 4$	$4 \leftrightarrow 1$
g	0.01831	0.02131	0.01931	0.02031

For the simulations of H_2 and HeH^+, which only involve two transmons, device parameters corresponding to transmons 1 and 2 were used. All units are in 2π GHz.

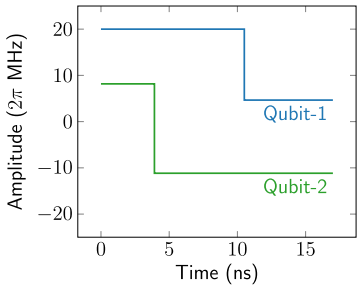

Fig. 10 Illustration of square pulses for a two transmon system. Simple square pulse with two time segments.

As such, the coupling strength g is ultimately responsible for describing electron correlation in the target molecule.

Using this ansatz in Eq. (9), the energy to be minimized in the ctrl-VQE objective function is,

$$E(\Omega_n(t), \nu_n) = \left\langle \psi^{\text{trial}'} \middle| H^{\text{molecule}} \middle| \psi^{\text{trial}'} \right\rangle. \tag{11}$$

where $\left|\psi^{\text{trial}'}\right\rangle$ is just the $\left|\psi^{\text{trial}}\right\rangle$ state above, projected onto the computational basis and normalized. Note that an unnormalized state can also be used; this yields similar results, as shown in the SI.

The ansatz above is completely (and solely) determined by the device and controls, granting enormous flexibility to the ansatz. In fact, any digital quantum circuit ansatz can be compiled into the form above. As such, the ansatz in Eq. (9) does not intrinsically possess any limitation on its potential accuracy beyond the fundamental limitations imposed by quantum speed limits[40]. However, this additional flexibility can make optimization more difficult.

In this work, we have chosen to impose simple constraints on the form of $\Omega_n(t)$ to minimize the number of optimization parameters and simplify experimental implementation. We have considered two examples: (i) piecewise-constant pulses, and (ii) sum of Gaussian pulses. Because these two examples yield similar results, we present only the square pulse data in the main text, and provide the Gaussian pulse data in the SI. An example of a pair of two-time-segment square pulses for a two transmon system is shown in Fig. 10. With a single time segment square pulse, the amplitude is taken as constant throughout the pulse duration. For n time segments, the parameterized square pulse is given by,

$$\Omega_k(t) = \begin{cases} c_1 & 0 \le t < t_1 \\ c_2 & t_1 \le t < t_2 \\ \vdots & \\ c_n & t_{n-1} \le t < T \end{cases}, \tag{12}$$

where c_i are amplitudes of the pulse, t_i are the switching times, and T is the total pulse duration. The amplitudes are constrained to ±20 MHz unless otherwise mentioned; this is a typical amplitude for RF control signals in superconducting qubit systems[64]. Each transmon drive term also has a frequency modulation of the form $\exp(i\nu_k t)$, with a driving frequency ν_k constrained to $(\omega_k - 2\pi\,\text{GHz}) \le \nu_k \le (\omega_k + 2\pi\,\text{GHz})$, where ω_k, ν_k are in units of 2π GHz. Therefore, the pulse parameters that will be optimized include c_i, t_i, and ν_k. With N transmons and n square pulses on each transmon, we then have $2Nn$ parameters to optimize. The pulse parameter optimizations are performed using l-BFGS-b.

We note that a very recent preprint uses similar quantum control considerations to improve VQE[65]. However, in that work, these considerations are used only to define an improved gate-based ansatz, unlike the direct variational pulse shaping described in this work. Our work is also somewhat related to the paper from Chamon and coworkers[66] in which they find that the optimal pulse for preparing the ground state of a Hamiltonian is of a bang-bang form. Although both ctrl-VQE and the work of Yang et al.[66] both combine variational quantum algorithms with quantum control, our effort aims to find ground states of non-diagonal Hamiltonians (such as the molecular Hamiltonian in Eq. (4)), preventing the efficient implementation of the time-evolution operator needed for realizing the bang-bang protocol. Because our control Hamiltonian is not determined by our problem Hamiltonian (Eq. (4)), we should not expect such bang-bang protocols to be optimal. We also note that a preprint was posted soon after this paper which suggests that optimal control would be helpful in VQE applications[67]. While pulse level control presents a challenge for front-end development which has focused on circuits, IBM has already begun making direct pulse access available[68], and we hope this work will help encourage other platforms to do the same.

Computational details

The numerical results presented in this work were generated with a locally developed program: CtrlQ. The source code for the program is available on Github, codebase: https://github.com/oimeitei/ctrlq under a Apache License 2.0. Functionalities from Qiskit[38] and Qutip[69] were used to check and compare the results presented in this work. Molecular integrals were generated using PySCF[70], and an STO-3G basis set was used throughout. Below, we consider simulations that involve either two-transmon or four-transmon devices depending on the molecule. The parameters ω, δ and g used in the simulations are explicitly given in Table 2. To demonstrate qualitative insensitivity to the device parameters, we provide some results using a device with a larger detuning in the SI.

Although the simulations presented in the following sections use highly restricted forms of control pulses, our in-house code (CtrlQ) supports various forms of pulses. In fact, we have implemented efficient analytic gradients of the molecular energy with respect to pulse amplitudes (see SI), making it possible to optimize pulses with arbitrary shapes. As mentioned in the previous section, here we focus primarily on simple pulse waveforms with few parameters to simplify experimental implementations.

DATA AVAILABILITY

The data for the numerical simulations is available upon reasonable request.

CODE AVAILABILITY

The code for the numerical simulations is available at https://github.com/mayhallgroup/ctrlq.

Received: 4 March 2021; Accepted: 28 September 2021;
Published online: 27 October 2021

REFERENCES

1. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
2. Chan, G. K.-L. & Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *J. Chem. Phys.* **116**, 4462–4476 (2002).
3. Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**, 96–192 (2011).
4. Huron, B., Malrieu, J. P. & Rancurel, P. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *J. Chem. Phys.* **58**, 5745 (1973).
5. Bender, C. F. & Davidson, E. R. Studies in configuration interaction: The first-row diatomic hydrides. *Phys. Rev.* **183**, 23–30 (1969).
6. Buenker, R. J. Ab initio SCF MO and CI studies of the electronic states of butadiene. *J. Chem. Phys.* **49**, 5381 (1968).
7. Evangelisti, S., Daudey, J.-P. & Malrieu, J.-P. Convergence of an improved CIPSI algorithm. *Chem. Phys.* **75**, 91–102 (1983).
8. Tubman, N. M., Lee, J., Takeshita, T. Y., Head-Gordon, M. & Whaley, K. B. A deterministic alternative to the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **145**, 044112 (2016).
9. Schriber, J. B. & Evangelista, F. A. Communication: an adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *J. Chem. Phys.* **144**, 161106–161106 (2016).
10. Holmes, A. A., Tubman, N. M. & Umrigar, C. J. Heat-bath configuration interaction: an efficient selected configuration interaction algorithm inspired by heat-bath sampling. *J. Chem. Theory Comput.* **12**, 3674–3680 (2016).
11. Levine, D. S. et al. CASSCF with extremely large active spaces using the adaptive sampling configuration interaction method. *J. Chem. Theory Comput.* **16**, 2340–2354 (2020).
12. Abraham, V. & Mayhall, N. J. Selected configuration interaction in a basis of cluster state tensor products. *J. Chem. Theory Comput.* **16**, 6098–6113 (2020).

13. Caffarel, M., Applencourt, T., Giner, E. & Scemama, A. *Recent Progress in Quantum Monte Carlo*, vol. 1234 of *ACS Symposium Series*, Chap. 2, p. 15–46 (American Chemical Society, 2016).
14. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213–4213 (2014).
15. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
16. Sharma, K., Khatri, S., Cerezo, M. & Coles, P. J. Noise resilience of variational quantum compiling. *New J. Phys.* **22**, 043006 (2020).
17. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
18. Huggins, W. J., Lee, J., Baek, U., O'Gorman, B. & Whaley, K. B. A non-orthogonal variational quantum eigensolver. *New J. Phys.* **22**, 073009 (2020).
19. Lee, J., Huggins, W. J., Head-Gordon, M. & Whaley, K. B. Generalized unitary coupled cluster wave functions for quantum computation. *J. Chem. Theory Comput.* **15**, 311–324 (2019).
20. Gard, B. T. et al. Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. *npj Quant. Inf.* **6**, 1–9 (2020).
21. Barron, G. S. et al. Preserving symmetries for variational quantum eigensolvers in the presence of noise. *Phys. Rev. Applied* **16**, 034002 (2021).
22. Ryabinkin, I. G., Yen, T.-C., Genin, S. N. & Izmaylov, A. F. Qubit coupled cluster method: A systematic approach to quantum chemistry on a quantum computer. *J. Chem. Theory Comput.* **14**, 6317–6326 (2018).
23. Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 3007 (2019).
24. Tang, H. L. et al. Qubit-adapt-vqe: an adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *PRX Quantum* **2**, 020310 (2021).
25. Ryabinkin, I. G., Lang, R. A., Genin, S. N. & Izmaylov, A. F. Iterative qubit coupled cluster approach with efficient screening of generators. *J. Chem. Theory Comput.* **16**, 1055–1063 (2020).
26. Wang, Y. et al. Quantum simulation of helium hydride cation in a solid-state spin register. *ACS Nano* **9**, 7769–7774 (2015).
27. Shen, Y. et al. Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure. *Phys. Rev. A* **95**, 020501 (2017).
28. Colless, J. I. et al. Computation of molecular spectra on a quantum processor with an error-resilient algorithm. *Phys. Rev. X* **8**, 011021 (2018).
29. McCaskey, A. J. et al. Quantum chemistry as a benchmark for near-term quantum computers. *npj Quant. Inf.* **5**, 1–8 (2019).
30. Xia, R., Bian, T. & Kais, S. Electronic structure calculations and the ising hamiltonian. *J. Phys. Chem. B* **122**, 3384–3395 (2018).
31. Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* **8**, 031022 (2018).
32. Sugisaki, K. et al. Quantum chemistry on quantum computers: a method for preparation of multiconfigurational wave functions on quantum computers without performing post-hartree-fock calculations. *ACS Cent. Sci.* **5**, 167–175 (2019).
33. Ritter, M. B. Near-term quantum algorithms for quantum many-body systems. *J. Phys. Conf. Ser.* **1290**, 012003 (2019).
34. Sagastizabal, R. et al. Experimental error mitigation via symmetry verification in a variational quantum eigensolver. *Phys. Rev. A* **100**, 010302 (2019).
35. Armaos, V., Badounas, D. A., Deligiannis, P. & Lianos, K. Computational chemistry on quantum computers. *Appl. Phys. A* **126**, 625 (2020).
36. Jordan, W. P. über das paulische äquivalenzverbot. *Z. Phys.* **47**, 631 (1928).
37. Bravyi, S., Gambetta, J. M., Mezzacapo, A. & Temme, K. Tapering off qubits to simulate fermionic hamiltonians (2017). Preprint at https://arxiv.org/abs/1701.08213 (2017).
38. Abraham, H. et al. *Qiskit: An Open-source Framework For Quantum Computing* (Zenodo, 2019).
39. Güsten, R. et al. Astrophysical detection of the helium hydride ion HeH+. *Nature* **568**, 357–359 (2019).
40. Deffner, S. & Campbell, S. Quantum speed limits: from heisenberg's uncertainty principle to optimal quantum control. *J. Phys. A Math. Theor.* **50**, 453001 (2017).
41. Koch, J. et al. Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A* **76**, 042319 (2007).
42. Khaneja, N. et al. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *J. Magn. Reson.* **172**, 296–305 (2005).
43. Gokhale, P. et al. Partial compilation of variational algorithms for noisy intermediate-scale quantum machines. *MICRO '52: Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 266. https://doi.org/10.1145/3352460.3358313 (2019).
44. Boutin, S., Andersen, C. K., Venkatraman, J., Ferris, A. J. & Blais, A. Resonator reset in circuit qed by optimal control for large open quantum systems. *Phys. Rev. A* **96**, 042315 (2017).
45. Lu, D. et al. Enhancing quantum control by bootstrapping a quantum processor of 12 qubits. *npj Quant. Inf.* **3**, 45 (2017).
46. Gradl, T., Spörl, A., Huckle, T., Glaser, S. J. & Schulte-Herbrüggen, T. *Euro-Par 2006 Parallel Processing* (eds. Nagel, W. E., Walter, W. V. & Lehner, W.) p. 751–762 (Springer Berlin Heidelberg, 2006).
47. Cheng, J., Deng, H. & Qia, X. Accqoc: Accelerating quantum optimal control based pulse generation. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, p. 543–555 (IEEE, 2020).
48. Williams, C. P. *Explorations in Quantum Computing* (Springer, 1998).
49. Economou, S. E. & Barnes, E. Analytical approach to swift nonleaky entangling gates in superconducting qubits. *Phys. Rev. B* **91**, 161405 (2015).
50. Barron, G. S., Calderon-Vargas, F. A., Long, J., Pappas, D. P. & Economou, S. E. Microwave-based arbitrary cphase gates for transmon qubits. *Phys. Rev. B* **101**, 054508 (2020).
51. Ball, H. et al. Software tools for quantum control: Improving quantum computer performance through noise and error suppression. *Quantum Sci. technol.* **6**, 044011 (2021).
52. Bravyi, S. B. & Kitaev, A. Y. Fermionic quantum computation. *Ann. Phys.* **298**, 210 (2002).
53. Seeley, J. T., Richard, M. J. & Love, P. J. The bravyi-kitaev transformation for quantum computation of electronic structure. *J. Chem. Phys.* **137**, 224109 (2012).
54. Bartlett, R. J. & Musiał, M. Coupled-cluster theory in quantum chemistry. *Rev. Mod. Phys.* **79**, 291–352 (2007).
55. Yung, M.-H. et al. From transistor to trapped-ion computers for quantum chemistry. *Sci. Rep.* **4**, 3589 (2014).
56. Ostaszewski, M., Grant, E. & Benedetti, M. Quantum circuit structure learning. *arXiv:1905.09692 [quant-ph]* (2019).
57. Chivilikhin, D. et al. MoG-VQE: multiobjective genetic variational quantum eigensolver. *arXiv:2007.04424 [cond-mat, physics:quant-ph]* (2020).
58. Matsuzawa, Y. & Kurashige, Y. Jastrow-type decomposition in quantum chemistry for low-depth quantum circuits. *J. Chem. Theory Comput.* **16**, 944–952 (2020).
59. Lee, J., Huggins, W. J., Head-Gordon, M. & Whaley, K. B. Generalized unitary coupled cluster wavefunctions for quantum computation. *J. Chem. Theory Comput.* **15**, 311 (2019).
60. Babbush, R. et al. Low-depth quantum simulation of materials. *Phys. Rev. X* **8**, 011044 (2018).
61. Verteletskyi, V., Yen, T.-C. & Izmaylov, A. F. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *J. Chem. Phys.* **152**, 124114 (2020).
62. Huggins, W. J. et al. Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers. *arXiv:1907.13117 [physics, physics: quant-ph]* (2019).
63. Zhao, A. et al. Measurement reduction in variational quantum algorithms. *Phys. Rev. A* **101**, 062322 (2020).
64. Ku, H. S. et al. Single qubit operations using microwave hyperbolic secant pulses. *Phys. Rev. A* **96**, 042339 (2017).
65. Choquette, A. et al. Quantum-optimal-control-inspired ansatz for variational quantum algorithms. *arXiv:2008.01098 [quant-ph]* (2020).
66. Yang, Z.-C., Rahmani, A., Shabani, A., Neven, H. & Chamon, C. Optimizing variational quantum algorithms using pontryagin's minimum principle. *Phys. Rev. X* **7**, 021027 (2017).
67. Magann, A. B. et al. From pulses to circuits and back again: a quantum optimal control perspective on variational quantum algorithms. *PRX Quant.* **2**, 010101 (2021).
68. Alexander, T. et al. Qiskit pulse: programming quantum computers through the cloud with pulses. *Quant. Sci. Technol.* **5**, 044006 (2020).
69. Johansson, J. R., Nation, P. D. & Nori, F. Qutip 2: A python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.* **184**, 1234–1240 (2013).
70. Sun, Q. et al. Pyscf: the python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.* **8**, e1340 (2018).

ACKNOWLEDGEMENTS

N.J.M., S.E.E., E.B. and D.P.P. are grateful for financial support provided by the U.S. Department of Energy (Award No. DE-SC0019199). S.E.E. acknowledges support provided by the U.S. Department of Energy (Award No. DE-SC0019318). This work was done in part while N.J.M. was visiting the Simons Institute for the Theory of Computing. The authors thanks the Advanced Research Computing at Virginia Tech for the computational infrastructure.

AUTHOR CONTRIBUTIONS

N.J.M., S.E.E., E.B. and D.P.P. conceived the project. O.R.M. wrote the python/C++ code and helped design the adaptive pulse parameterization algorithm. B.T.G. wrote mathematica code. B.T.G. and O.R.M performed numerical experiments. All authors contributed to the conceptual analysis of the results and to ideas that refined the algorithm. All authors wrote the paper. O.R.M. and B.T.G. contributed equally.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41534-021-00493-0.

Correspondence and requests for materials should be addressed to Nicholas J. Mayhall.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2021

Chapter 11

Preparing Bethe Ansatz Eigenstates on a Quantum Computer

This chapter includes pages from our publication [9]

J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, "Preparing bethe ansatz eigenstates on a quantum computer", PRX Quantum **2**, 040329 (2021)

Copyright (2021) of the American Physical Society.

The contributions to this work are as follows: My contributions were assisting in numerical simulations, development of the underlying simulator. J.S.V.D. primarily developed this work. N.J.M., E.B., and S.E.E. supervised.

Preparing Bethe Ansatz Eigenstates on a Quantum Computer

John S. Van Dyke,[1,*] George S. Barron,[1] Nicholas J. Mayhall,[2] Edwin Barnes,[1] and Sophia E. Economou[1]

[1] *Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2] *Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

(Received 29 April 2021; revised 22 August 2021; accepted 4 October 2021; published 9 November 2021)

Several quantum many-body models in one dimension possess exact solutions via the Bethe ansatz method, which has been highly successful for understanding their behavior. Nevertheless, there remain physical properties of such models for which analytic results are unavailable and which are also not well described by approximate numerical methods. Preparing Bethe ansatz eigenstates directly on a quantum computer would allow straightforward extraction of these quantities via measurement. We present a quantum algorithm for preparing Bethe ansatz eigenstates of the spin-1/2 *XXZ* spin chain that correspond to real-valued solutions of the Bethe equations. The algorithm is polynomial in the number of T gates and the circuit depth, with modest constant prefactors. Although the algorithm is probabilistic, with a success rate that decreases with increasing eigenstate energy, we employ amplitude amplification to boost the success probability. The resource requirements for our approach are lower than for other state-of-the-art quantum simulation algorithms for small error-corrected devices and thus may offer an alternative and computationally less demanding demonstration of quantum advantage for physically relevant problems.

DOI: 10.1103/PRXQuantum.2.040329

I. INTRODUCTION

Quantum computers hold the promise of transformative applications in a variety of fields, including cryptanalysis [1], quantum chemistry [2,3], materials science [4,5], and, potentially, combinatorial optimization [6,7]. To realize the full potential of quantum computing, large-scale fault-tolerant devices will ultimately be necessary. As these do not yet exist, much recent work has studied possible near-term applications in the present era of noisy intermediate-scale quantum computers (NISQs) [8,9]. In this context, a key question concerns the demonstration of "quantum advantage"—that is, the ability to perform computations that cannot be done efficiently with classical methods. Recently, quantum advantage has been shown for a superconducting processor sampling random quantum circuits [10] and photonic-based Gaussian boson sampling [11]. Although these are important achievements, the specific tasks performed have not been closely related to the practical applications mentioned above but have essentially been designed for the purpose of demonstrating advantage. Thus, the realization of quantum advantage for a problem of practical interest remains open.

A question of increasing interest is what applications become feasible with small-scale error-corrected devices, i.e., in the intermediate era between NISQs and fault-tolerant quantum computers with many logical qubits and a high clock speed for non-Clifford gates. Recent estimates suggest that algorithms that provide only a quadratic speed-up over classical methods may have difficulty achieving quantum advantage on problem sizes accessible with small devices [12–14]. The simulation of quantum systems, on the other hand, can yield exponential improvement over conventional approaches. These still require formidable resources, despite recent algorithmic advances [15–18]. This motivates the search for physically interesting problems and algorithms that can lead to quantum advantage with fewer resources in the near future.

We propose the study of Bethe ansatz (BA) states on a quantum computer as a computationally less demanding route to the demonstration of quantum advantage for problems relevant to physics, including quantum magnetism [19,20], ultracold atoms [21], and unconventional superconductivity [22]. BA methods allow for deep insight into the static and dynamic properties of these many-body systems and are able to explore not only ground states but also interactions between complex collective excitations, such as magnons, and the response to quantum quench experiments [23].

*jsvandyke86@gmail.com

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

More specifically, the BA technique yields exact solutions to a class of one-dimensional quantum many-body models, including the spin-1/2 Heisenberg and Hubbard models, among others [24–27]. The resulting wave functions depend on algebraic equations that can be efficiently solved classically. The exponential growth of the Hilbert space with the system size L has historically limited the direct computational studies of the eigenstates to small systems. Instead, various mathematical techniques have been extensively developed to access physical quantities in the thermodynamic limit $L \to \infty$, bypassing the calculation of the wave function itself. While many different quantities can be determined, the difficulty of their calculation varies widely. In particular, arbitrary-range and higher-order correlation functions have been very challenging to access and remain an active area of research [28–30]. Quantum computers, however, can compute such correlation functions [5,31] straightforwardly, thus suggesting the possibility of quantum advantage for this task. The importance of higher-order correlation functions for strongly correlated systems has recently been emphasized [32]. The calculation of such observables using a quantum computer in turn hinges on the possibility of efficiently preparing the Bethe ansatz states.

To this end, we demonstrate an efficient quantum algorithm that can prepare a subset of the Bethe ansatz eigenstates of the one-dimensional *XXZ* chain, a model that is fundamental to the study of quantum magnetism. Our algorithm uses the so-called coordinate Bethe ansatz method and has polynomial scaling in the circuit depth and T-gate complexity, along with low constant prefactors. As we show with explicit gate counts for the corresponding circuits, the approach scales to large enough systems for the calculation of classically inaccessible quantities in near-term error-corrected devices. While the algorithm we present is probabilistic, we also show that amplitude amplification can be used to increase the success rate [33].

Algorithms have been previously given (and also implemented) that prepare exact eigenstates of quantum many-body models [34–37]. However, these have been largely limited to cases that are equivalent to noninteracting fermions, for instance, under the Jordan-Wigner mapping. In contrast, the *XXZ* model corresponds to an interacting fermionic problem, which is computationally much more difficult. We note that the possibility of constructing circuits to diagonalize Bethe ansatz-solvable models and measure challenging correlation functions has previously been suggested, though without an indication of how this could be done [34].

The importance of Bethe ansatz-solvable models for benchmarking NISQ devices has been previously recognized [37–39], as they provide exact values for quantities (such as the energy) to compare against the results of noisy quantum computations. On the other hand, the direct preparation of Bethe ansatz states has been relatively unexplored. This question has recently been studied in Ref. [40], which considered treating Bethe ansatz states variationally (using the algebraic Bethe ansatz) and concluded that the approach was not scalable. Furthermore, that work did not make a connection to the possibility of quantum advantage. Apart from direct preparation of Bethe ansatz eigenstates, other works have considered variational approaches using generic ansatzes [38,39,41–43]. The comparison of the computational complexity of these methods to that of the direct construction is an interesting question for future studies, as are probabilistic algorithms for preparing other strongly correlated states [44].

The paper is organized as follows. Section II introduces the *XXZ* model and the elements of the Bethe ansatz solution needed for the construction of the algorithm. Section III describes the Bethe ansatz state-preparation algorithm. Section IV presents numerical results that validate the method and studies its success probability. This section also includes resource estimates for classically intractable problem sizes. Section V describes the amplitude amplification procedure for our algorithm and presents numerical calculations that confirm its success. Section VI compares our algorithm with conceptually simpler but less efficient approaches to the same task, explicitly verifying the enormous speed-up of our method. Section VII argues that quantum advantage can be achieved with Bethe-state preparation by comparison with classical computational methods, and presents additional applications of the algorithm. Finally, we conclude in Sec. VIII.

II. MODEL AND SOLUTION

We consider the one-dimensional spin-1/2 *XXZ* chain on L sites with periodic boundary conditions, the Hamiltonian of which is given by

$$H = \sum_{i=0}^{L-1} J_{xy}\left(S_i^x S_{i+1}^x + S_i^y S_{i+1}^y\right) + J_z S_i^z S_{i+1}^z, \qquad (1)$$

with $S_L^\alpha \equiv S_0^\alpha$ ($\alpha = x, y, z$). Here, the S_j^α are the spin operators with eigenvalues $\pm 1/2$. The exact solution of this model via the Bethe ansatz method has been presented in Ref. [45] and many introductions to the problem (in both the coordinate and algebraic formulations) exist [46–50]. We follow the account of the Bethe ansatz method given in Ref. [48]. The eigenstates of the above Hamiltonian are given by

$$\psi(x_1, \ldots, x_M) = \sum_P A_P \exp\left[i \sum_{j=1}^{M} k_{Pj} x_j\right], \qquad (2)$$

where $x_1, \ldots, x_M$ label the positions of the M down spins in the chain (the Hamiltonian conserves the z component of the total spin, $S_{\text{tot}}^z = \sum_i S_i^z$) and the momenta k_i label

the different states. The wave function of Eq. (2) gives the amplitude for the M down spins to occur on the sites $x_j = 0, \ldots, L-1$. The summation here is over the $M!$ permutations of the down-spin sites. These permutations arise from the fact that, within Bethe ansatz models, scattering processes exchange momenta between particles but do not alter their magnitudes. The coefficients A_P are related by

$$\frac{A_P}{A_{P'}} = -\frac{1 + e^{i(k_{Pl} + k_{P'l})} - 2\Delta e^{ik_{Pl}}}{1 + e^{i(k_{Pl} + k_{P'l})} - 2\Delta e^{ik_{P'l}}} \equiv -e^{-i\Theta(k_{Pl}, k_{P'l})}. \tag{3}$$

To fix the coefficients, we take $A_I = 1$, where I is the identity permutation. Here, $\Delta = J_z/J_{xy}$ is the anisotropy in the interactions and P and P' are permutations that differ by a single transposition between adjacent elements, $P(l+1) = P'(l)$ and $P(l) = P'(l+1)$. The momenta k_i are constrained by the quantization conditions,

$$Lk_i = 2\pi I_i + \sum_j \Theta(k_i, k_j), \tag{4}$$

where I_i is an integer (half-integer) for M odd (even). Physically, these constraints on k_i arise from imposing periodic boundary conditions on the model. Equations (3) and (4) are a set of algebraic equations (the Bethe equations) for the quantum numbers $\{k_i\}$. In general, these equations admit complex solutions but for the special case when all $\{k_i\}$ are real, $\Theta(k_i, k_j)$ is also real and is given by

$$\Theta(k_i, k_j) = 2 \arctan\left[\frac{\Delta \sin(k_{ij}^-/2)}{\Delta \cos(k_{ij}^-/2) - \cos(k_{ij}^+/2)}\right]. \tag{5}$$

where $k_{ij}^{\pm} = k_i \pm k_j$.

In the spirit of hybrid quantum-classical algorithms, we solve the Bethe equations classically to obtain the momenta $\{k_i\}$ and phases $\Theta(k_i, k_j)$. These values are then used as input to our quantum algorithm for generating the corresponding eigenstate. Our algorithm allows for the preparation of Bethe ansatz eigenstates for which the $\{k_i\}$ are real, such that A_P and $e^{ik_{Pj}x_j}$ amount to complex phases applied to the second-quantized basis states of the system.

III. BETHE ANSATZ STATE-PREPARATION ALGORITHM

The quantum algorithm for preparing a Bethe ansatz state consists of several steps and is summarized in Algorithm 1. The general structure of our approach is based on the linear-combination-of-unitaries (LCU) method (which also finds application in the Taylor-series approach to Hamiltonian simulation) [51,52]. However, a key difference is that our algorithm aims to generate specific quantum states starting from a particular initial state, rather than compiling a generic unitary evolution operator. In addition to the L qubits representing the system, a register of M^2 ancilla qubits are used to label the different permutation terms in Eq. (2). By preparing a superposition of the allowed label values on these ancillas, using these to apply controlled operations on the system, and finally disentangling the label and system registers, we perform the summation over all permutations present in Eq. (2). This process is facilitated by introducing a second ancillary register of M qubits that we call the "faucet register," along with one additional ancilla work qubit. Thus, the algorithm requires a total of $M^2 + M + 1$ ancilla qubits.

1: Prepare the Dicke state $|D_{L,M}\rangle$ on the system qubits
2: Create permutation labels while applying pieces of A_P
3: Apply $e^{ik_{Pj}x_j}$ using the "faucet" method
4: Reverse permutation label (without phases)
5: Measure permutation label, with success on $|00\cdots0\rangle$

Algorithm 1. Bethe-state preparation.

The algorithm begins by preparing the Dicke state on L sites with M down spins, $|D_{L,M}\rangle$. Relabeling $|\uparrow\rangle \equiv |0\rangle$, $|\downarrow\rangle \equiv |1\rangle$, $|D_{L,M}\rangle$ is the equal superposition (that is, without relative phases) of all basis states on L qubits with Hamming weight M. This state forms the underlying "canvas" on which the phases in Eq. (2) are applied. Dicke-state preparation can be accomplished using the recent deterministic algorithm of Ref. [53], for which the gate count has been improved in Ref. [54]. This algorithm uses an inductive method to prepare smaller Dicke states that are subsequently combined to yield the desired one. We use this algorithm in our explicit circuit constructions, though any other deterministic method of preparing $|D_{L,M}\rangle$ would also work.

As discussed above, the amplitudes that must be applied to $|D_{L,M}\rangle$ to generate a Bethe ansatz state depend on the permutations $\{P\}$ of M objects. We use the permutation-label register to create the different permutations and their associated phases A_P. Naively, one could use an integer encoded in a binary representation to label each of the permutations. The difficulty with this approach is that the number of permutations is $M!$, so that imprinting the phases A_P and $e^{ik_{Pj}x_j}$ onto $|D_{L,M}\rangle$ would require combinatorially many operations. This leads to circuit depths and complexities that are superexponential in M, quickly becoming unfeasible as M grows (we explore a concrete realization of this approach in Sec. VI). To overcome this fundamental limitation of this method and design an efficient algorithm, we introduce a conceptually distinct approach for labeling the permutations. Rather than assigning an arbitrary number to a given permutation, we implement its explicit action on the string consisting of the numbers $1, \ldots, M$. As described below, this allows for an

efficient generation of the permutation labels, while also generating the distinct A_P simultaneously.

The permutation-label register consists of M subregisters, each of which can store an integer value $k \in \{1, \ldots, M\}$. To represent a valid permutation, the subregisters must contain distinct values (for instance, $|213\rangle$ is valid whereas $|233\rangle$ is not). We use a one-hot encoding such that each subregister consists of M qubits and the number k is represented by a 1 on the kth qubit and 0s on the rest. Thus, for $M = 3$, the allowed states on each subregister are $|1\rangle \equiv |001\rangle$, $|2\rangle \equiv |010\rangle$, and $|3\rangle \equiv |100\rangle$. This one-hot encoding requires M^2 qubits to represent the complete label. The use of the one-hot encoding facilitates a trade-off between time and space resources [15,55,56] by reducing the number of controls required to implement the necessary phase gates.

The goal of step 2 of Algorithm 1 is to create the state $(1/\sqrt{M!}) \sum_P A_P |P\rangle$ on the permutation-label register. The phases A_P are kicked back onto the system qubits, while the $|P\rangle$ are used to apply the conditional gates needed in step 3, as explained below. For clarity, we first describe the construction of the equal superposition of all permutation labels. We then show how to slightly modify this procedure to simultaneously generate the phases A_P for all $M!$ permutations. We use an iterative method to construct the permutation-label state, starting from the vacuum state $|00\ldots0\rangle$ on M^2 qubits. The complete label superposition state is built up sequentially from the first (rightmost) subregister to the last (leftmost) using a series of exchange-type gates. We describe the method inductively as follows. The zeroth sublabel is prepared by setting the zeroth qubit of the zeroth subregister to 1 (in the following, the index k is enumerated starting from 0). Assume that the kth sublabel (i.e., an equal superposition of permutations of integers 1 through $k+1$) is constructed on the $k+1$ rightmost subregisters. Set the $(k+1)$th qubit of the $(k+1)$th subregister to 1, thus introducing the next integer value to be included in the permutation-label state. Perform the exchange-type ASWAP gate [57,58],

$$A(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & e^{i\phi}\sin(\theta) & 0 \\ 0 & e^{-i\phi}\sin(\theta) & -\cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (6)$$

between the $(k+1)$th qubits of subregisters $k+1$ and k, with $\theta = \arccos(1/\sqrt{k+2})$, $\phi = 0$. This generates a superposition state consisting of two sets of terms, those in which the 1 remains in the $(k+1)$th subregister and those in which it is transferred to the kth. In the latter case, the $(k+1)$th subregister now contains all zeros, while the kth has two qubits with 1, which is not valid. This is fixed by applying controlled-SWAP gates between all qubits $l < k+1$ in subregisters k and $k+1$, controlled on the state of qubit $k+1$ in subregister k. Taken together, these operations produce a partial SWAP between subregisters $k+1$ and k,

$$|k+1\rangle_{k+1} |m\rangle_k \to \frac{1}{\sqrt{k+2}} |k+1\rangle_{k+1} |m\rangle_k + \sqrt{\frac{k+1}{k+2}} |m\rangle_{k+1} |k+1\rangle_k , \quad (7)$$

where $|i\rangle_j$ is the one-hot encoded state for i on subregister j, and $m < k$ by construction. One repeats this partial swapping process, now between subregisters k and $k-1$, then between $k-1$ and $k-2$, and so on, until the last register is swapped. By implementing the inductive process up to the $(M-1)$th subregister, the complete equally weighted superposition of permutation labels is formed. As an example, for $M = 3$, the above algorithm generates the following sequence of state transformations:

$$\begin{aligned}
&|000\rangle\ |000\rangle\ |000\rangle \xrightarrow{X} |000\rangle\ |000\rangle\ |001\rangle \\
&\xrightarrow{X} |000\rangle\ |010\rangle\ |001\rangle \\
&\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{2}} |000\rangle \left(|010\rangle\ |001\rangle + |000\rangle\ |011\rangle \right) \\
&\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{2}} |000\rangle \left(|010\rangle\ |001\rangle + |001\rangle\ |010\rangle \right) \\
&\xrightarrow{X} \frac{1}{\sqrt{2}} |100\rangle \left(|010\rangle\ |001\rangle + |001\rangle\ |010\rangle \right) \\
&\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{6}} |100\rangle\ |010\rangle\ |001\rangle + \frac{1}{\sqrt{3}} |000\rangle\ |110\rangle\ |001\rangle \\
&\quad + \frac{1}{\sqrt{6}} |100\rangle\ |001\rangle\ |010\rangle + \frac{1}{\sqrt{3}} |000\rangle\ |101\rangle\ |010\rangle \\
&\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{6}} |100\rangle\ |010\rangle\ |001\rangle + \frac{1}{\sqrt{3}} |010\rangle\ |100\rangle\ |001\rangle \\
&\quad + \frac{1}{\sqrt{6}} |100\rangle\ |001\rangle\ |010\rangle + \frac{1}{\sqrt{3}} |001\rangle\ |100\rangle\ |010\rangle \\
&\xrightarrow{\text{ASWAP}} \frac{1}{\sqrt{6}} |100\rangle\ |010\rangle\ |001\rangle + \frac{1}{\sqrt{6}} |010\rangle\ |100\rangle\ |001\rangle \\
&\quad + \frac{1}{\sqrt{6}} |010\rangle\ |000\rangle\ |101\rangle + \frac{1}{\sqrt{6}} |100\rangle\ |001\rangle\ |010\rangle \\
&\quad + \frac{1}{\sqrt{6}} |001\rangle\ |100\rangle\ |010\rangle + \frac{1}{\sqrt{6}} |001\rangle\ |000\rangle\ |110\rangle \\
&\xrightarrow{\text{CSWAPs}} \frac{1}{\sqrt{6}} |100\rangle\ |010\rangle\ |001\rangle + \frac{1}{\sqrt{6}} |010\rangle\ |100\rangle\ |001\rangle \\
&\quad + \frac{1}{\sqrt{6}} |010\rangle\ |001\rangle\ |100\rangle + \frac{1}{\sqrt{6}} |100\rangle\ |001\rangle\ |010\rangle \\
&\quad + \frac{1}{\sqrt{6}} |001\rangle\ |100\rangle\ |010\rangle + \frac{1}{\sqrt{6}} |001\rangle\ |010\rangle\ |100\rangle . \quad (8)
\end{aligned}$$

The explicit circuit for $M = 3$ is shown in Fig. 1, where the gates inside the red dashed rectangles are excluded at this point. Furthermore, gates acting on distinct qubits are pushed to the left, thereby reducing the circuit depth. This produces a different sequence of intermediate states than above, but the final state is the same.

A slight modification of the above procedure allows one to simultaneously apply the phases A_P to the appropriate terms in the superposition. After each partial subregister swap, one applies a controlled-phase gate with angle $\Theta(k_i, k_j) + \pi$, where i,j are the integer values that are swapped and where the additional phase π implements the signature of the permutation (red dashed rectangles in Fig. 1). Suppose, as in Eq. (7), that the value $(k+1)$ is swapped to the right. Then the $(k+1)$th qubit of the right subregister can serve as the target qubit in the required controlled phase. Since the left subregister can store any value $m < k+1$, a separate controlled phase is used for each possibility, where the control bits are given by the values m. This leads to a total of $M^3/3 - M^2/2 + M/6$ controlled-phase gates for this part of the algorithm.

By the end of the construction, the phases A_P have been applied to the corresponding permutation-label state, having been successively built up from the elementary transpositions of which they are composed. The explicit circuit for $M = 3$ is displayed in Fig. 1, where now the gates in the red dashed rectangles are included, in order to produce the phases A_P. At this point in the construction, the total state of the physical system and the permutation label is

$$\left(\frac{1}{\sqrt{M!}} \sum_P A_P \otimes_{j=M}^{1} |Pj\rangle_p\right) |D_{L,M}\rangle_s , \qquad (9)$$

where $|\cdots\rangle_p$ is a state of the permutation-label qubits and $|\cdots\rangle_s$ is a state of the system qubits.

In the next step, one applies the position-dependent phase factors $e^{ik_{Pj}x_j}$ to the relevant basis states on the system qubits (step 3 of Algorithm 1). To do this, we introduce an efficient method that acts on all the $\binom{L}{M}$ basis states in $|D_{L,M}\rangle$ simultaneously, yielding an enormous speed-up over classical approaches. The technique, which we call the "faucet" method, is based on the observation that the positions x_j take integer values $x_j = 0, \ldots, L-1$, so that the total phase $e^{ik_{Pj}x_j}$ can be produced by x_j repetitions of the phase $e^{ik_{Pj}}$.

For this part of the algorithm, we use the M additional ancilla qubits comprising the faucet register. Each of these new qubits is initialized to $|1\rangle$. In the outer loop of the faucet subroutine, one traverses the register of the system qubits site by site from $x = 0$ to $x = L-1$. At each site, if it is occupied by a down spin (i.e., the bit is 1), one turns off the next faucet ancilla qubit, $|1\rangle \to |0\rangle$. This is achieved through a sequence of multicontrolled X gates, which are controlled on the previous ancilla being in the state $|0\rangle$ and the next one being in the state $|1\rangle$, along with the additional control that the current system site is a $|1\rangle$. Since the meaning of the "next faucet ancilla" at a given site depends on the bit string, one must generically apply a multicontrolled X gate for each ancilla at every step. We note that one can decrease the number of gates for sites near the edges of the chain. For instance, at site 2 at most two faucets could have been turned off, so that the later ones do not need to be checked.

Next, the phases $e^{ik_{Pj}}$ $(j = 1, \ldots, M)$ are applied to the system qubits, each being controlled on the state of one of the faucet ancillas. For the j th ancilla, this gate is also controlled on the state of the permutation-label subregister j (since the value of k_{Pj} is permutation dependent). By the

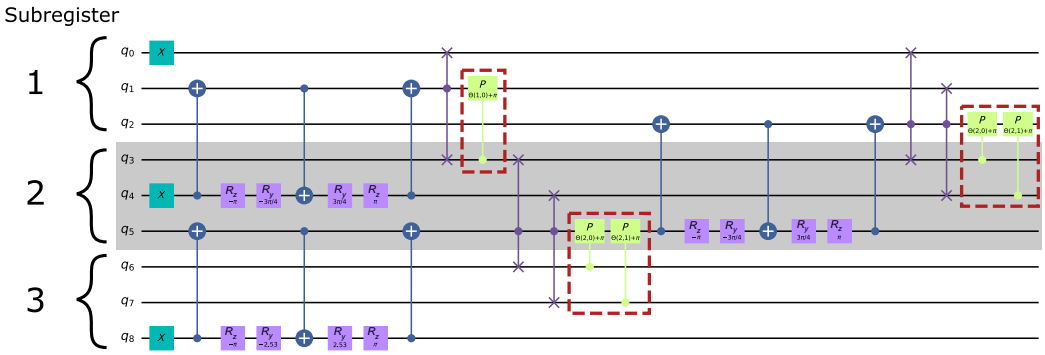

FIG. 1. The circuit to prepare the permutation-label ancilla state for $M = 3$. When the gates inside the red dashed rectangles are excluded, the circuit realizes the equally weighted superposition encoding the permutations of three objects. With these gates included, the circuit applies the additional permutation-dependent phases A_P. The three subregisters are indicated by shading and the initial state of each qubit is $|0\rangle$. The alternating CNOT gates interleaved with R_y and R_z rotations realize the ASWAP gates.

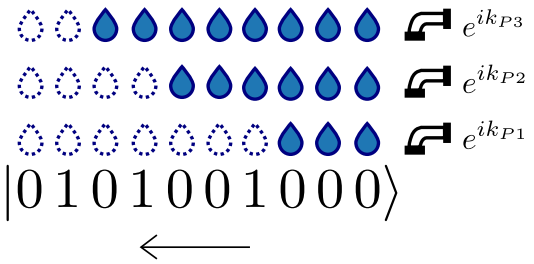

FIG. 2. A schematic diagram illustrating the idea of the faucet method. Phases $e^{ik_{Pj}}$ are applied for each system qubit while traversing the bit string. When a "1" is encountered, the next faucet in the list is turned off, so that no more phases with the given k_{Pj} value are applied.

end of the system bit string, all the faucet-register qubits are in state $|0\rangle$. Thus, the subroutine can be compared to a set of M running faucets (that correspond to applying the phases $e^{ik_{Pj}}$), which are turned off at the appropriate times (upon encountering a "1" in the traversal of the system register). This analogy is illustrated in Fig. 2. In total, this step requires M^2L doubly controlled-phase gates and a number of Toffoli gates that scales as ML (as mentioned above, some gates can be skipped near the edge of the system). In our numerical implementation of the method, we introduce additional work qubits to facilitate the construction of these gates using chains of Toffolis [59].

After the relevant phases have been applied, it remains to disentangle the system from the permutation label (the entanglement having been generated during the faucet method, since the phases there are permutation dependent). This is accomplished by applying the inverse of the circuit that generates the permutation-label superposition, without the additional controlled-phase gates that are used to produce the A_P phases. The fact that the phases A_P and $e^{ik_{Pj}x_j}$ have been applied to the initial Dicke state implies that the permutation-label reversal will not completely disentangle the system qubits from the permutation-label register. However, it turns out that the $|00\ldots0\rangle$ component of the permutation state corresponds precisely with the occurrence of the target Bethe ansatz state on the system qubits. That is, the full state vector takes the form

$$|\psi\rangle = \alpha|00\ldots0\rangle_p|\psi_B\rangle_s + \beta|\phi_j\rangle, \tag{10}$$

where $|\psi_B\rangle_s$ is the normalized target Bethe ansatz state on the system qubits $|\phi_j\rangle$ is a junk state with $(|00\ldots0\rangle\langle00\ldots0|_p \otimes \mathbb{1}_s)|\phi_j\rangle = 0$. Thus, by measuring the permutation-label qubits, the target Bethe ansatz state is successfully prepared on the outcome $|00\ldots0\rangle$. This result is essentially that obtained from LCU methods [52], though here we have the additional construction of A_P during the label-preparation step, which is not present in the standard LCU. To illustrate the full algorithm, the complete circuit to construct a BA eigenstate with $L = 4$, $M = 2$ is given in Appendix A. As discussed in our numerical simulations below, the success probability $|\alpha|^2$ depends on the system parameters and also varies between different eigenstates. Thus, in general one must repeat the procedure multiple times to obtain the correct state, which can then be used to calculate physical quantities or in other applications, as we discuss in Sec. VII. We also show in Sec. V that amplitude amplification can be used to boost the success rate, thereby reducing the overall resource requirements.

IV. NUMERICAL SIMULATIONS

To calculate the momenta $\{k_i\}$ defining the Bethe eigenstates, we solve the Bethe equations iteratively using the approach presented in Ref. [48]. We then perform numerical simulations of Algorithm 1 using the IBM QISKIT library's state-vector simulator [60]. These calculations verify the correctness of our algorithm and reveal its success probabilities for the sufficiently small systems that can be studied on a classical computer. However, we can also explicitly construct the circuits that would need to be run for much larger instances, far beyond what is classically tractable. The corresponding circuit depths and gate counts in these cases indicate that our algorithm is feasible for near-term error-corrected quantum computers. We stress that our analysis does not rely on asymptotic resource scaling arguments but, rather, provides exact gate counts, since the corresponding circuits are precisely known. Although we do not compile our algorithm down to an error-correcting code such as the surface code, we estimate the required number of T gates below.

In Fig. 3(a) we present the numerically calculated success probability of the algorithm for preparing selected eigenstates when $L = 2M$ and $M = 2, 3, 4$, with $J_{xy} = 1$, $J_z = -1/2$. The interaction strength in this case corresponds to the critical regime of the ferromagnetic model. We note the eigenstates included in Fig. 3 are not meant to comprise the complete set of real-valued solutions but, rather, are simply those for which we obtain numerical solutions of the Bethe equations by using the algorithm presented in Ref. [48]. Two general trends are apparent: a significant suppression of the success rate with increasing M and a more moderate suppression as a function of the eigenstate energy, within each set of system parameters L, M. The worst-case probabilities are roughly consistent with $1/M!$, although we have only limited values of M to support this (larger M being outside of our computational resources for classical simulation). This value is further supported by Fig. 3(b), which shows the success

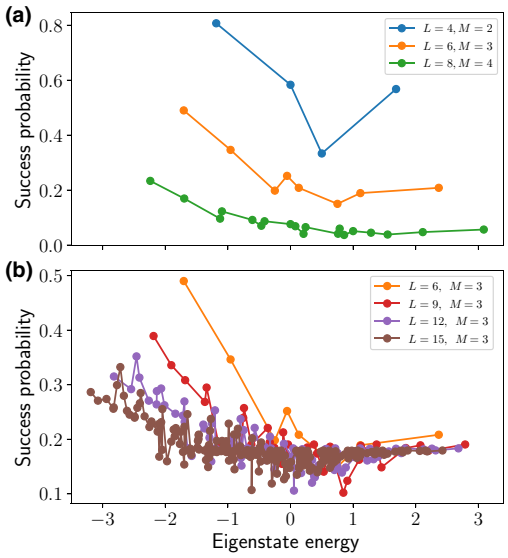

FIG. 3. (a) The success probabilities for the preparation of selected eigenstates as a function of their energies, when $L = 2M$ and $M = 2, 3, 4$. (b) The success probabilities of selected eigenstates for $M = 3$ and varying L. The Hamiltonian parameter values are $J_{xy} = 1$ and $J_z = -1/2$.

probabilities for various eigenstates when $M = 3$ and for different values of L, with the clear trend that increasing L tends to flatten the success probability across the spectrum. Although the low-energy states enjoy less of an advantage over the higher-energy ones in this case, the lowest probabilities are still around $1/M!$ on average. However, the behavior of the success rate changes significantly depending on the value of J_z. These effects are considered in Appendix B.

These results suggest that one can go to very large system sizes L while still preserving a reasonably large success probability, if M is sufficiently small. We note that the case of small M is particularly interesting for physics applications in the ferromagnetic regime of the model. In this case, the all-up state $|0\rangle^{\otimes L}$ is a ground state of the model, while small M states are low-lying excited states of interacting magnons (spin waves). The ability to study these states as a function of M, as enabled by our algorithm, would yield deeper insight into the development of strong correlations in these systems as the number of interacting quasiparticles grows. While the ferromagnetic regime is especially natural for our algorithm, we note that interesting physics in the paramagnetic and antiferromagnetic regimes can also be explored at small values of M. These correspond to highly excited eigenstates, which are relevant, for instance, in the study of quantum thermalization [61,62].

For these applications (and others discussed below), it appears feasible to access values of L and M that would not be classically simulable (even by approximate methods), while maintaining relatively modest resource requirements for the algorithm. This is demonstrated in Fig. 4, which provides circuit depths [Fig. 4(a)], Toffoli-gate counts [Fig. 4(b)], controlled-phase gate counts [Fig. 4(c)], and the number of qubits required [Fig. 4(d)] for the Bethe-state-preparation algorithm. The linear scaling of all these metrics in L is immediately apparent. The slopes of these lines for the case $M = 5$ are (a) 39, (b) 11, (c) 25, and (d) 1, respectively. The results for the number of controlled-phases and qubits are in exact agreement with the analytical results of Sec. III.

Apart from the asymptotic scaling behavior, the absolute values of the circuit depths and gate counts are seen to be very low, on the order of 10^3–10^4, even for large systems of $L \sim 100$ sites. Furthermore, the total number of qubits required (approximately 10^2) is also quite reasonable for small error-corrected devices. In Fig. 5, we show the total gate and measurement counts for the case $M = 5$ as L is varied. This indicates that the controlled-phase, R_y, and Toffoli gates are the most prevalent non-Clifford operations in the algorithm. To estimate the fault-tolerant resources needed, we therefore convert the counts for these gates into the corresponding numbers of T gates. Following Ref. [63], we assume a worst-case scenario for the number of T gates needed to realize an arbitrary z rotation to be given by $4\log_2(1/\epsilon) + 11$, where ϵ is the rotation synthesis error [64]. Similarly, arbitrary y rotations can be performed by conjugation with Clifford operations. As in Ref. [16], we replace each Toffoli gate with two T gates [65]. For $L = 100$, $M = 5$, this leads to a T count of approximately 6.2×10^5 for a single run of the state-preparation algorithm, with $\epsilon = 10^{-10}$. Assuming a worst-case success probability of $1/M!$, approximately 120 attempts would need to be performed on average to correctly generate the target eigenstate. This yields approximately 7.4×10^7 T gates overall, which is comparable to the estimates for simulating the Hubbard model using the methods of Ref. [16]. We note that the estimates in that work involve optimizing an error budget between multiple sources (Trotterization, phase estimation, and rotation synthesis) and do not appear to include the cost of preparing a good initial state for the phase estimation routine. Furthermore, the above estimate for our algorithm assumes the seemingly worst-case scenario in the number of repetitions (approximately $M!$), whereas the results in Fig. 3 suggest that lower-energy states require fewer repetitions in general. To reduce the number of repetitions required for our algorithm, we implement amplitude amplification in the following section.

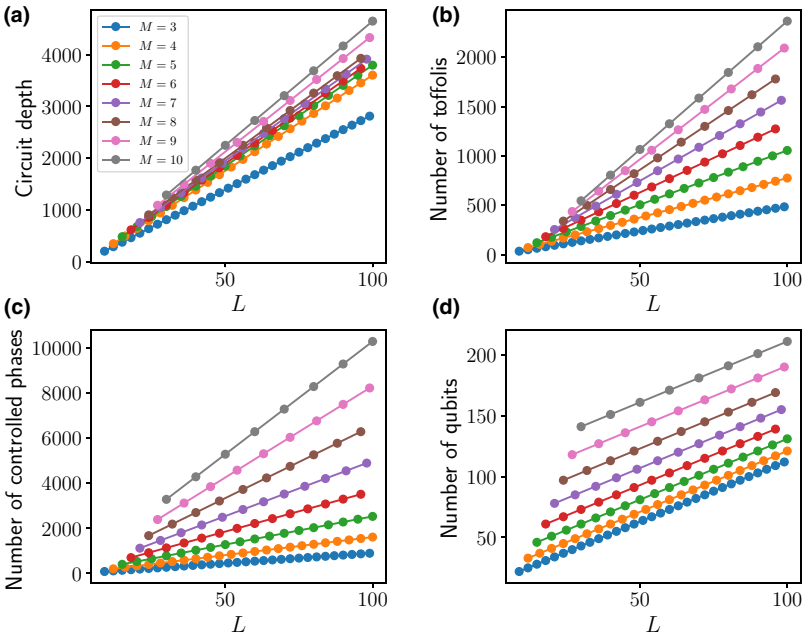

FIG. 4. The Bethe-state-preparation (a) circuit depth, (b) number of Toffoli gates, (c) number of controlled-phase gates, and (d) number of qubits versus the system size L, for different numbers of down spins M.

V. AMPLITUDE AMPLIFICATION

Amplitude amplification, a generalization of the well-known Grover search algorithm, is a quantum subroutine that can boost the probability of a desired measurement outcome, leading in general to a square-root improvement in the number of repetitions required for the success of a probabilistic algorithm [33]. For our problem, we use $\mathcal{B}$ to denote Algorithm 1 with the measurement step removed.

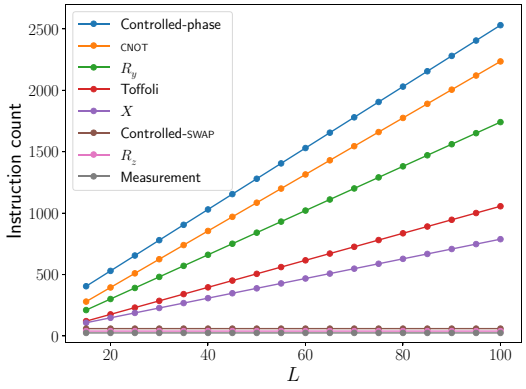

FIG. 5. The Bethe-state-preparation gate and measurement counts for $M = 5$ as a function of L.

Amplitude amplification defines an operator

$$\mathcal{Q} = -\mathcal{B}S_0\mathcal{B}^{-1}S_B, \tag{11}$$

where S_B changes the relative sign of the "good" states in the Hilbert space, while S_0 changes the relative sign of the vacuum state $|00\ldots0\rangle$. In the present case, the good states are the components of the Bethe ansatz state, which correspond to $|00\ldots0\rangle_p$ on the permutation-label qubits. S_B can therefore be implemented using a OR circuit on the permutation label, followed by Z on the work qubit that stores the result, after which the OR is uncomputed. In our numerical calculations of the success probability, we use the ancilla-free implementation of OR in the standard QISKIT circuit library. The ancilla-free approach is used here to decrease the number of qubits needed for the simulation, allowing us to study larger system sizes. Below, we examine an ancilla-based method for which the gate counts at large sizes are reduced. To produce S_0, we use the same approach, with the OR circuit extended to include the system qubits (we do not implement the -1 in Eq. (11), as it is an overall phase). We present numerical results for amplitude amplification in Fig. 6(a). This confirms the clear enhancement of the algorithm success probability using this method. Although only one round of amplification is applied here, the protocol can be repeated in the standard way to further increase the success rate. Applying this improvement to the resource estimate of the previous section, the $M = 5$ worst-case eigenstates should

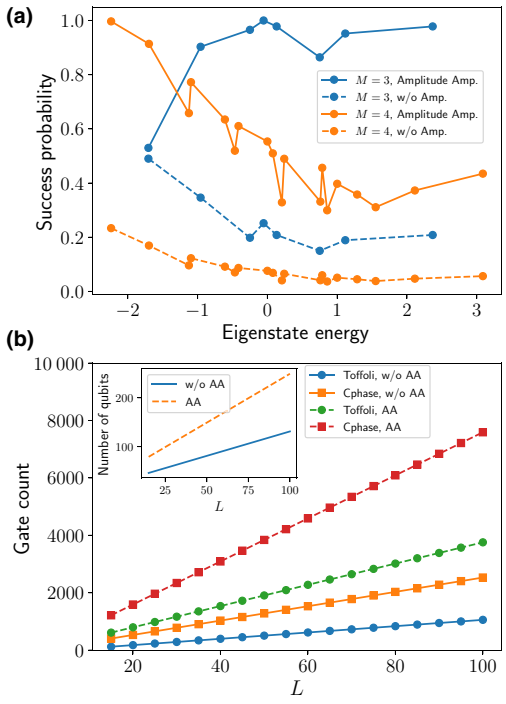

FIG. 6. (a) The success probability for Bethe-state preparation with (solid lines) and without (dashed lines) amplitude amplification ("AA"). In the former case, a single round of amplification is applied. The system parameters are $J_{xy} = 1$, $J_z = -1/2$, $L = 2M$. (b) The gate counts as a function of the system size L for Toffoli (circles) and controlled-phase (squares) gates ("Cphase") without (solid lines) and with one round of amplitude amplification (dashed lines). The inset shows the total number of qubits without (solid) and with (dashed) amplitude amplification.

require on average $\sqrt{120} \approx 11$ repetitions of the algorithm $\mathcal{B}$ to achieve success, leading to an overall T count of approximately 4.1×10^6 (neglecting the costs of S_B and S_0).

To estimate the resources needed for amplitude amplification more precisely, we implement the multicontrolled OR operation using elementary gates and ancillas [59]. The Toffoli count, controlled-phase count, and number of qubits are shown in Fig. 6(b) for a single round of amplification when $M = 5$. Since the S_0 reflection depends on the state of both the system and the permutation label, the implementation requires an additional $L + M^2 - M$ ancillas (we can reuse the M faucet ancillas for the amplification). The additional gates required are dominated by the cost of executing $\mathcal{B}$ three times. Although the total number of qubits is approximately doubled in this approach, we note that it may be possible to reduce this number by acting with the OR operation on only a subset of the qubits that are nominally necessary to identify the state. For example, although S_B is an OR circuit acting on the M^2 permutation-label qubits, in practice the computational basis states appearing in the junk state $|\phi_j\rangle$ can be distinguished from $|00\ldots0\rangle_p$ by only acting on a reduced number of qubits. While the particular subset of qubits needed will depend on the state under consideration, it is possible to confirm the success of this approach by measuring the energy or other quantities that can be compared to exact analytic expressions.

We also implement a different version of amplitude amplification, which is a modified form of the oblivious amplitude amplification of Ref. [52]. Unfortunately, this method leads to a reduced fidelity of the actually prepared state with the exact target state, though in some cases the overlap remains quite high (> 0.99). We attribute this reduced fidelity to the nonunitarity of summing exponentials with unit modulus, since $|e^{ia} + e^{ib}| \neq 1$ in general. We note, however, that nearly deterministic success of the oblivious amplitude amplification procedure was obtained for the application of Ref. [52] (simulation of Hamiltonian dynamics with Taylor-series expansions). For this reason, it is less clear how the approach will fare for the Bethe-state-preparation problem at larger values of M. In addition, further modification to the method of Ref. [52] may ameliorate some of the difficulties with applying it to Bethe-state preparation in its present form.

VI. COMPARISON WITH ALTERNATIVE ALGORITHMS

To highlight the advantages of Algorithm 1, we compare it with conceptually simpler, but much less efficient, methods of preparing Bethe ansatz states on a quantum computer. First, one could imagine applying controlled-phase gates directly to each term in the Dicke-state superposition to generate the desired eigenstate. This approach still requires permutation-label ancillas to generate the linear combination of phases needed in Eq. (2). However, it has the seeming advantage of allowing one to combine the phases A_P and $e^{ik_{Pj}x_j}$ into a single controlled-phase rotation, whereas the former term requires order M^3 and the latter one order M^2 controlled phases to implement using Algorithm 1. Nevertheless, it is clear that this benefit is vastly outweighed by the large number of terms in the superposition that must be separately addressed, $M!\binom{L}{M}$. For the case $L = 100$, $M = 5$ this amounts to approximately 9.0×10^9 controlled phases, compared to the 2530 of Algorithm 1.

A more promising approach is to use the "faucet" method of Algorithm 1 to handle the $e^{ik_{Pj}x_j}$ phases whilestill applying the full A_P in a single multicontrolled

phase gate, rather than decomposing it into its elementary transpositions. This replaces the $M!\binom{L}{M}$ dependence above with $M!LM$. While the scaling is still inferior to that of Algorithm 1 for large M, it is conceivable that for small M this alternative method may be competitive. In particular, one can replace the complicated permutation label of Algorithm 1, which requires M^2 qubits to construct A_P in terms of individual transpositions, with a compressed label that simply assigns a number to each permutation. This approach uses significantly fewer qubits, at the expense of requiring more controls for the relevant phase gates. Since the permutation-label construction still needs to be reversed to disentangle the system from the ancillas, it is important that it can still be executed in a unitary fashion. This in turn requires an efficient method for generating an equal superposition of $M!$ states. We implement such states using the prime factorization $M! = 2^{n_2}3^{n_3}\cdots$. We then construct the permutation label as the tensor product of the binary representation of 2^{n_2} (using n_2 qubits) and W_n states for the odd prime factors. The equal superposition for the binary part of the label is easily generated by applying Hadamard gates to the relevant qubits, while various efficient algorithms exist for W_n state preparation [66–68]. We implement this algorithm numerically and verify that it successfully prepares Bethe ansatz eigenstates. Since the fundamental approach for creating the linear combination of phases is the same between this method and Algorithm 1, their success probabilities are equal. Unfortunately, explicit construction of the circuits for the alternative method indicates that the resource requirements are significantly higher, even for small M. This is shown in Fig. 7, which reveals that the number of controlled-phase and Toffoli gates required for the alternative method vastly exceeds those of Algorithm 1, even for $M = 5$.

It is also useful to compare Algorithm 1 with other more standard techniques, such as adiabatic state preparation [2,69,70]. As is well known, the evolution time to produce a large overlap with the desired eigenstate is expected to scale as the inverse square of the gap between the energy of the given state and that of the next closest one. This time can be short, for instance, in the ordered regime of the XXZ chain, where the gap between the ground and first excited state remains finite in the large system limit. However, in the critical regime of the model, the spectrum is gapless and so the evolution time is expected to diverge. Furthermore, even in the ordered regime, the spectrum at higher energies above the ground state generically has dense regions, preventing an efficient preparation of those eigenstates by the adiabatic approach. Our algorithm does not suffer from such complications, as it implements the exact analytic expression for the wave function, irrespective of the gap between the target eigenstate and the other ones.

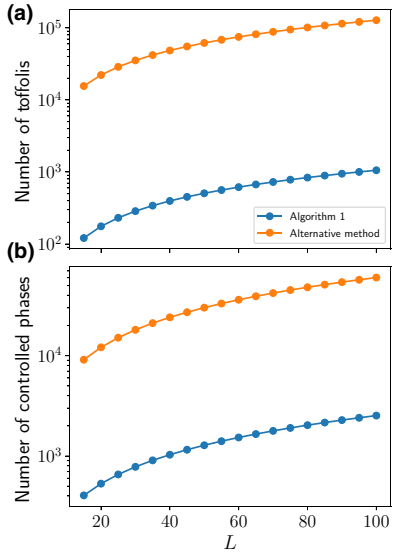

FIG. 7. A comparison between Algorithm 1 and the alternative method of the number of (a) Toffoli and (b) controlled-phase gates used, for $M = 5$.

VII. DISCUSSION

To achieve quantum advantage for a physically relevant problem, it should be the case that no classical method can deliver results of a comparable accuracy. Since the present quantum algorithm exactly prepares eigenstates of the XXZ chain, it is reasonable to compare it to the numerical exact diagonalization of finite-size systems (i.e., using classical computers). In a recent study, a matrix-free approach has been used to investigate Heisenberg spin chains up to length $L = 26$ [71]. In this work, the memory requirements have been vastly reduced compared to conventional methods, though the scaling remains exponential with the system size. Specifically the $S^z = 0$ subspace ($M = 13$) has been considered, for which the dimension is approximately 10^7. In contrast, the $L = 100$, $M = 5$ subspace is roughly seven times larger (dimension approximately 7.5×10^7). Although state vectors of this size can still be stored in memory, we note that the computation time is also exponential in the system size, ultimately limiting the practicality of exact diagonalization.

In addition to numerically exact calculations, approximate tensor-network methods have been highly successful for studying one-dimensional quantum many-body systems with the matrix product state (MPS) ansatz [72,73]. However, these approaches are best suited for states with a relatively low amount of entanglement, such as gapped ground states obeying an area law for the entanglement entropy. This makes simulation of long-time dynamics

challenging, due to the growth of entanglement from, for instance, an initial product state. Our Bethe ansatz algorithm can prepare eigenstates throughout the spectrum at the same computational cost, including highly excited states the entanglement entropy of which grows more quickly than logarithmically, even in the small-M limit [74]. This yields an advantage over MPS methods for large systems when targeting these strongly entangled states. An explicit link between the Bethe ansatz and MPS has been developed in Refs. [75,76], which uses the algebraic Bethe ansatz to produce exact tensor-network representations for generic eigenstates. These networks have a PEPS-like structure (but with fewer physical indices), which underscores the computational intractability of such states for large systems.

In addition to computing arbitrary-range and higher-order correlation functions that are inaccessible with traditional Bethe ansatz methods, our algorithm has a number of other applications. Simulation of the real-time dynamics of many-body systems is widely recognized as a task allowing for quantum advantage. Such simulations often take the form of quench experiments, for which the system is initialized in an easy-to-prepare product state, then allowed to evolve under the influence of an interacting many-body Hamiltonian. Our algorithm would enable interesting variations on this approach, for instance by initializing the system in an eigenstate of a given value of the interaction strength, then subsequently evolving it with a different value. The evolution here can be performed using any of the known algorithms for quantum simulation, whether by Trotterization [16,77,78], Taylor expansions [52,79], or other approaches [51,80].

In a different direction, one could use our algorithm as a starting point to explore integrability-breaking perturbations. Thus, we consider a Hamiltonian of the form $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_p$, where $\mathcal{H}_0$ is solvable by the Bethe ansatz and $\mathcal{H}_p$ includes perturbations that break the integrability of the total Hamiltonian $\mathcal{H}$ (such as disorder in the coupling strengths). In this case, the Bethe-state-preparation algorithm is used to prepare an eigenstate of $\mathcal{H}_0$, which then serves as an initial state for quantum annealing or phase estimation on $\mathcal{H}$. For sufficiently weak perturbations, the overlap of this state with the corresponding exact eigenstate of $\mathcal{H}$ should be much greater than that of a mean-field or noninteracting trial state.

Although we focus on deploying our algorithm on small error-corrected quantum computers, one may also consider implementing it on present-day or near-term NISQ devices. Many of the controlled-phase gates in our algorithm involve repetitions of the same basic rotation angles, of which there are only $(M^2 + M)/2$ distinct values. This suggests replacing the exact values with variational parameters, similarly to Ref. [40]. The resulting variational form can then be optimized under the cost function $|E - E_B|$, where E is the energy calculated on the quantum computer and E_B is the exact value, known analytically from the Bethe ansatz solution. We note that the present optimization problem should be significantly easier than that of a standard VQE, since the ideal values of the rotation angles can serve as a good initial guess. Updates to the parameters then serve to directly mitigate systematic errors due to over- or under-rotation in the controlled-phase gates.

VIII. CONCLUSIONS

We present a quantum algorithm for the efficient preparation of Bethe ansatz eigenstates of the XXZ model. To our knowledge, this is the first quantum algorithm for the direct preparation of eigenstates of an interacting many-body problem. The circuit depth and gate counts of the algorithm scale linearly in the system size, for a fixed number of down spins. Our algorithm is feasible to perform on small error-corrected devices of order 100 qubits, provided that the number of down spins is small. The usefulness of the approach can be extended through amplitude amplification. In particular, quantum advantage over classical computational methods appears to be achievable, with resource estimates that are comparable to the most efficient known quantum simulation algorithms. Our work suggests directions for future research, including the modification of the algorithm for the case of complex-valued $\{k_i\}$, and its generalization to other Bethe ansatz-solvable models, such as the one-dimensional Hubbard model.

ACKNOWLEDGMENTS

The numerical simulations were performed with IBM QISKIT and the QuSpin PYTHON library [81]. We would like to thank Chandra Sekhar Mukherjee, Jesko Sirker, and Rafael Nepomechie for helpful discussions. This work was supported by the U.S. 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.

APPENDIX A: FULL CIRCUIT FOR $L = 4, M = 2$

In Fig. 8, we present the full quantum circuit to prepare a Bethe ansatz eigenstate with $L = 4$ and $M = 2$. We note that the work qubit is not needed in the $M = 2$ case but we include it here as a reminder that it is used to implement multicontrolled gates for $M > 2$.

APPENDIX B: SUCCESS PROBABILITY FOR DIFFERENT J_Z

The success rate of the Bethe ansatz state preparation changes qualitatively with both the sign of J_z and whether

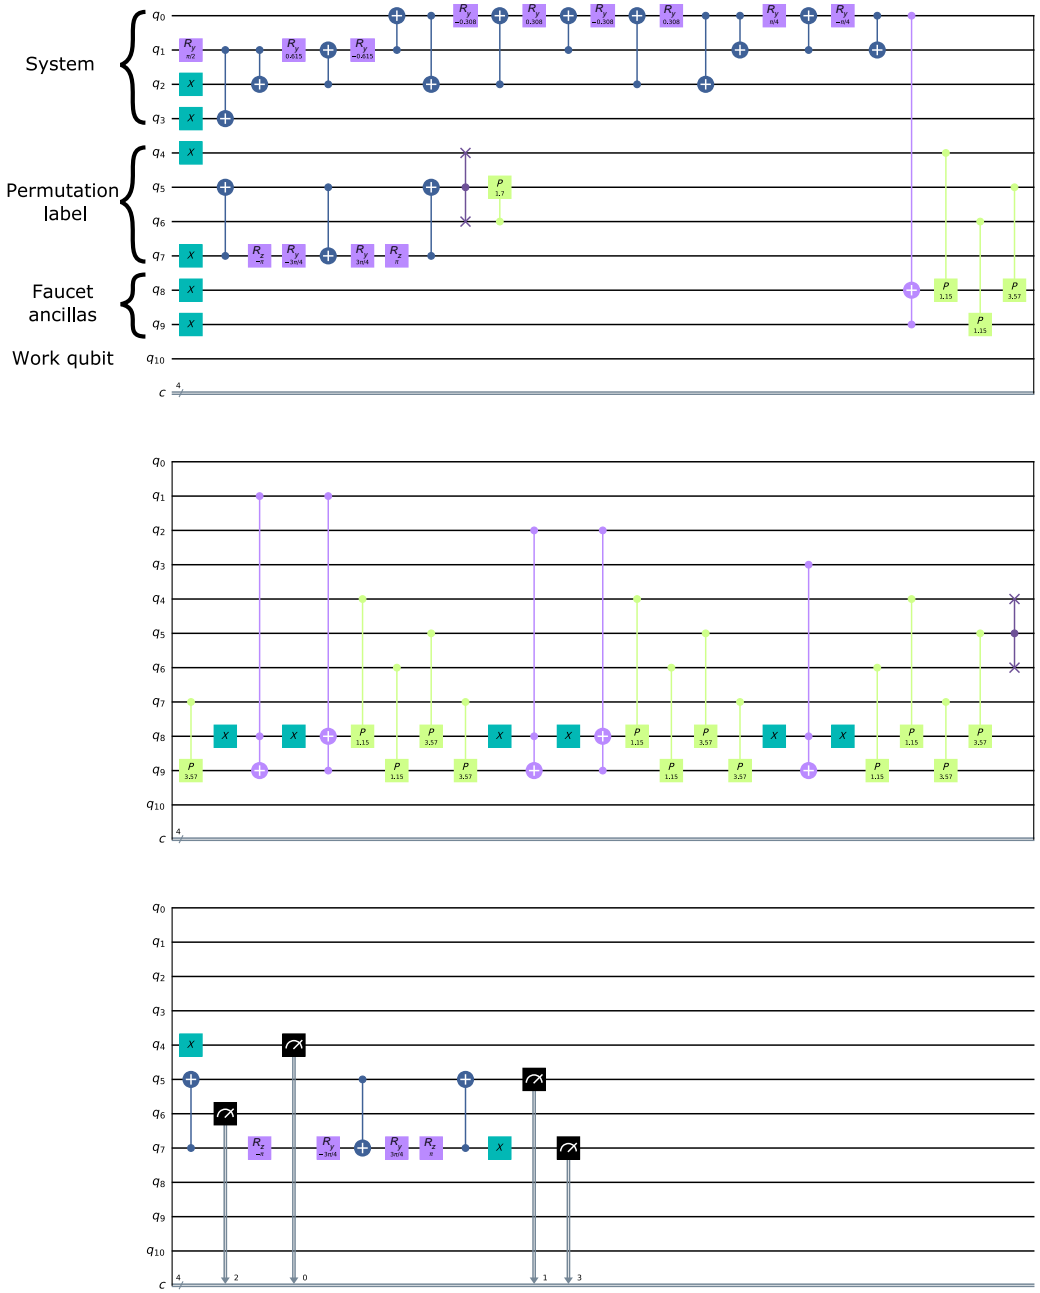

FIG. 8. The quantum circuit to prepare the Bethe ansatz eigenstate with $L = 4$, $M = 2$, $k_1 = 1.14676529$, and $k_2 = 3.56562369$.

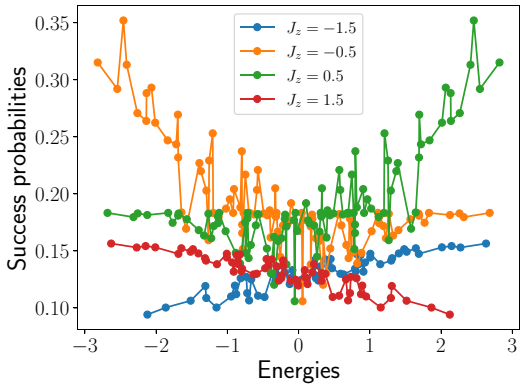

FIG. 9. The success probabilities for the preparation of selected eigenstates as a function of their energies, for $L = 12$, $M = 3$, $J_{xy} = 1$, with varying interaction strengths J_z.

it lies in the critical ($|J_z| < 1$) or noncritical ($|J_z| \geq 1$) regimes. These effects are illustrated in Fig. 9. Whereas the ferromagnetic model in the critical regime has higher success probabilities for lower-energy states, the antiferromagnetic case shows the opposite behavior. In fact, the success probabilities for different states are exactly mirrored across the $E = 0$ axis. This is especially interesting since the corresponding eigenstates at $E = -E_i$ for the FM case and $E = E_i$ for the AFM one are different in general. It is unclear at present why these distinct states should have the same success probability. When J_z is outside the critical regime, we find that the success probability can be lower than the $1/M!$ value (approximately 0.17 for $M = 3$) that appears to bound the critical case.

[1] P. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc. Press, 1994), p. 124.

[2] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, Science **309**, 1704 (2005).

[3] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, Rev. Mod. Phys. **92**, 015003 (2020).

[4] D. S. Abrams and S. Lloyd, Simulation of Many-Body Fermi Systems on a Universal Quantum Computer, Phys. Rev. Lett. **79**, 2586 (1997).

[5] D. Wecker, M. B. Hastings, N. Wiebe, B. K. Clark, C. Nayak, and M. Troyer, Solving strongly correlated electron models on a quantum computer, Phys. Rev. A **92**, 062318 (2015).

[6] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, Phys. Rev. Lett. **79**, 325 (1997).

[7] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028 (2014).

[8] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[9] I. H. Deutsch, Harnessing the power of the second quantum revolution, PRX Quantum **1**, 020101 (2020).

[10] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, and B. Burkett, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[11] H.-S. Zhong, *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[12] E. Campbell, A. Khurana, and A. Montanaro, Applying quantum algorithms to constraint satisfaction problems, Quantum **3**, 167 (2019).

[13] Y. R. Sanders, D. W. Berry, P. C. S. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, Compilation of fault-tolerant quantum heuristics for combinatorial optimization, PRX Quantum **1**, 020312 (2020).

[14] R. Babbush, J. McClean, C. Gidney, S. Boixo, and H. Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, arXiv:2011.04149 (2020).

[15] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity, Phys. Rev. X **8**, 041015 (2018).

[16] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik, H. Neven, and R. Babbush, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, Quantum **4**, 296 (2020).

[17] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, arXiv:2007.14460 (2020).

[18] J. Lee, D. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, arXiv:2011.03494 (2020).

[19] S. E. Nagler, D. A. Tennant, R. A. Cowley, T. G. Perring, and S. K. Satija, Spin dynamics in the quantum antiferromagnetic chain compound $KCuF_3$, Phys. Rev. B **44**, 12361 (1991).

[20] Z. Wang, J. Wu, W. Yang, A. K. Bera, D. Kamenskyi, A. T. M. N. Islam, S. Xu, J. M. Law, B. Lake, C. Wu, and A. Loidl, Experimental observation of Bethe strings, Nature **554**, 219 (2018).

[21] T. Fukuhara, P. Schauß, M. Endres, S. Hild, M. Cheneau, I. Bloch, and C. Gross, Microscopic observation of magnon bound states and their dynamics, Nature **502**, 76 (2013).

[22] P. R. Pasnoori, N. Andrei, and P. Azaria, Edge modes in one-dimensional topological charge conserving spin-triplet superconductors: Exact results from Bethe ansatz, Phys. Rev. B **102**, 214511 (2020).

[23] P. Calabrese and J. Cardy, Time Dependence of Correlation Functions Following a Quantum Quench, Phys. Rev. Lett. **96**, 136801 (2006).

[24] H. Bethe, Zur Theorie der Metalle i. Eigenwerte und Eigenfunktionen der linearen Atomkette, Z. Phys. **71**, 205 (1931).

[25] E. H. Lieb and F. Y. Wu, Absence of Mott Transition in an Exact Solution of the Short-Range, One-Band Model in One Dimension, Phys. Rev. Lett. **20**, 1445 (1968).

[26] D. C. Mattis, *The Many-Body Problem: An Encyclopedia of Exactly Solved Models in One Dimension (3rd Printing with Revisions and Corrections)* (World Scientific, Singapore, 1993).

[27] F. H. L. Essler, H. Frahm, F. Göhmann, A. Klümper, and V. E. Korepin, *The One-Dimensional Hubbard Model* (Cambridge University Press, Cambridge, 2005).

[28] F. Göhmann, M. Karbach, A. Klümper, K. K. Kozlowski, and J. Suzuki, Thermal form-factor approach to dynamical correlation functions of integrable lattice models, J. Stat. Mech.: Theory Exp. **2017**, 113106 (2017).

[29] C. Babenko, F. Göhmann, K. K. Kozlowski, and J. Suzuki, A thermal form factor series for the longitudinal two-point function of the Heisenberg-Ising chain in the antiferromagnetic massive regime, J. Math. Phys. **62**, 041901 (2021).

[30] C. Babenko, F. Göhmann, K. K. Kozlowski, J. Sirker, and J. Suzuki, Exact real-time longitudinal correlation functions of the massive *XXZ* chain, arXiv:2012.07378 (2020).

[31] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Simulating physical phenomena by quantum networks, Phys. Rev. A **65**, 042323 (2002).

[32] A. Bohrdt, Y. Wang, J. Koepsell, M. Kánasz-Nagy, E. Demler, and F. Grusdt, Dominant Fifth-Order Correlations in Doped Quantum Antiferromagnets, Phys. Rev. Lett. **126**, 026401 (2021).

[33] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, Quantum Comput. Quantum Inf. 305, 53 (2002). ArXiv:aquant-ph/0005055

[34] F. Verstraete, J. I. Cirac, and J. I. Latorre, Quantum circuits for strongly correlated quantum systems, Phys. Rev. A **79**, 032316 (2009).

[35] P. Schmoll and R. Orús, Kitaev honeycomb tensor networks: Exact unitary circuits and applications, Phys. Rev. B **95**, 045112 (2017).

[36] A. Cervera-Lierta, Exact Ising model simulation on a quantum computer, Quantum **2**, 114 (2018).

[37] K. Robbins and P. J. Love, Benchmarking near-term quantum devices with the variational quantum eigensolver and the Lipkin-Meshkov-Glick model, Phys. Rev. A **104**, 022412 (2021).

[38] P.-L. Dallaire-Demers, M. Stęchły, J. F. Gonthier, N. T. Bashige, J. Romero, and Y. Cao, An application benchmark for fermionic quantum simulations, arXiv:2003.01862 (2020).

[39] M. J. Cervia, A. B. Balantekin, S. N. Coppersmith, C. W. Johnson, P. J. Love, C. Poole, K. Robbins, and M. Saffman, Exactly solvable model as a testbed for quantum-enhanced dark matter detection, arXiv:2011.04097 (2020).

[40] R. I. Nepomechie, Bethe ansatz on a quantum computer? Quantum Inf. Comput. **32**, 255 (2021).

[41] W. W. Ho and T. H. Hsieh, Efficient variational simulation of non-trivial quantum states, SciPost Phys. **6**, 029 (2019).

[42] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, PRX Quantum **1**, 020319 (2020).

[43] A. Cervera-Lierta, J. S. Kottmann, and A. Aspuru-Guzik, The meta-variational quantum eigensolver (Meta-VQE): Learning energy profiles of parameterized Hamiltonians for quantum simulation, arXiv:2009.13545 (2020).

[44] B. Murta and J. Fernández-Rossier, Gutzwiller wave function on a digital quantum computer, arXiv:2103.15445 (2021).

[45] R. Orbach, Linear antiferromagnetic chain with anisotropic coupling, Phys. Rev. **112**, 309 (1958).

[46] V. E. Korepin, N. M. Bogoliubov, and A. G. Izergin, *Quantum Inverse Scattering Method and Correlation Functions* (Cambridge University Press, Cambridge, 1993).

[47] M. Takahashi, *Thermodynamics of One-Dimensional Solvable Models* (Cambridge University Press, Cambridge, 1999), 1st ed.

[48] T. Giamarchi, *Quantum Physics in One Dimension*, The International Series of Monographs on Physics No. 121 (Clarendon, Oxford University Press, Oxford, New York, 2004), oCLC: ocm52784724.

[49] B. Sutherland, *Beautiful Models: 70 Years of Exactly Solved Quantum Many-Body Problems* (World Scientific, Singapore, 2004).

[50] M. Gaudin and J.-S. Caux, *The Bethe Wavefunction* (Cambridge University Press, Cambridge, 2014).

[51] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, Quantum Inf. Comput. **12**, 901 (2012).

[52] 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).

[53] A. Bärtschi and S. Eidenbenz, in *Fundamentals of Computation Theory* (Springer, Cham, 2019) p. 126.

[54] C. S. Mukherjee, S. Maitra, V. Gaurav, and D. Roy, Preparing Dicke states on a quantum computer, IEEE Trans. Quantum Eng. **1**, 3102517 (2020).

[55] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, Proc. NatL. Acad. Sci. **115**, 9456 (2018).

[56] K. Wan, Exponentially faster implementations of Select(H) for fermionic Hamiltonians, Quantum **5**, 380 (2021).

[57] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions, Phys. Rev. A **98**, 022322 (2018).

[58] 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**, 00 (2020).

[59] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2019), oCLC: 1129119093.

[60] H. Abraham, *et al.*, QISKIT: An open-source framework for quantum computing (2019).

[61] F. H. L. Essler and M. Fagotti, Quench dynamics and relaxation in isolated integrable quantum spin chains, J. Stat. Mech.: Theory Exp. **2016**, 064002 (2016).

[62] L. Vidmar and M. Rigol, Generalized Gibbs ensemble in integrable lattice models, J. Stat. Mech.: Theory Exp. **2016**, 064007 (2016).
[63] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, Proc. Natl. Acad. Sci. **114**, 7555 (2017).
[64] P. Selinger, Efficient Clifford+T approximation of single-qubit operators, Quantum Inf. Comput. **15**, 159 (2015).
[65] C. Gidney and A. G. Fowler, Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation, Quantum **3**, 135 (2019).
[66] C. Yesilyurt, S. Bugu, F. Ozaydin, A. A. Altintas, M. Tame, L. Yang, and Ş. K. Özdemir, Deterministic local doubling of W states, JOSA B **33**, 2313 (2016).
[67] D. Cruz, R. Fournier, F. Gremion, A. Jeannerot, K. Komagata, T. Tosic, J. Thiesbrummel, C. L. Chan, N. Macris, M.-A. Dupertuis, and C. Javerzac-Galy, Efficient quantum algorithms for GHZ and W states, and implementation on the IBM quantum computer, Adv. Quantum Technol. **2**, 1900015 (2019).
[68] C. Gidney, General construction of W_n state, https://quantumcomputing.stackexchange.com/questions/4350/general-construction-of-w-n-state.
[69] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv:quant-ph/0001106v1 (2000).
[70] B. W. Reichardt, In *Proc. 36th ACM Symp. on Theory of Computing (STOC)*, (2004), p. 502.
[71] R. Van Beeumen, G. D. Kahanamoku-Meyer, N. Y. Yao, and C. Yang, in *HPCAsia2020: Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region* (ACM, New York, 2020), p. 179.
[72] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, Phys. Rev. B **73**, 094423 (2006).
[73] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, Adv. Phys. **57**, 143 (2008).
[74] J. Mölter, T. Barthel, U. Schollwöck, and V. Alba, Bound states and entanglement in the excited states of quantum spin chains, J. Stat. Mech.: Theory Exp. **2014**, P10029 (2014).
[75] V. Murg, V. E. Korepin, and F. Verstraete, Algebraic Bethe ansatz and tensor networks, Phys. Rev. B **86**, 045125 (2012).
[76] Y. Q. Chong, V. Murg, V. E. Korepin, and F. Verstraete, Nested algebraic Bethe ansatz for the supersymmetric $t-j$ model and tensor networks, Phys. Rev. B **91**, 195132 (2015).
[77] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Simulation of electronic structure Hamiltonians using quantum computers, Mol. Phys. **109**, 735 (2011).
[78] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer, Improving quantum algorithms for quantum chemistry, Quantum Inf. Comput. **15**, 1 (2015).
[79] R. Babbush, D. W. Berry, I. D. Kivlichan, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, Exponentially more precise quantum simulation of fermions in second quantization, New J. Phys. **18**, 033032 (2016).
[80] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, Quantum **3**, 163 (2019).
[81] P. Weinberg and M. Bukov, QuSpin: A PYTHON package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains, SciPost Phys. **2**, 003 (2017).

Part IV

ADAPT-VQE

Chapter 12

Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor

This chapter includes pages from our publication [10]

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)

Copyright (2021) of the American Physical Society.

The contributions to this work are as follows: My contributions were assisting in theoretical developments, assisting in the development of minimal complete pools. H.L.T. primarily developed this work. V.O.S. assisted in developing minimal complete pools and simulations. H.R.G. assisted in numerical simulations. N.J.M., E.B., and S.E.E. supervised.

Qubit-ADAPT-VQE: An Adaptive Algorithm for Constructing Hardware-Efficient Ansätze on a Quantum Processor

Ho Lun Tang,[1,*] V.O. Shkolnikov,[1] George S. Barron,[1] Harper R. Grimsley,[2] Nicholas J. Mayhall,[2] Edwin Barnes,[1] and Sophia E. Economou[1,†]

[1] *Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2] *Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

(Received 5 February 2020; revised 6 June 2020; accepted 30 March 2021; published 28 April 2021)

Quantum simulation, one of the most promising applications of a quantum computer, is currently being explored intensely using the variational quantum eigensolver. The feasibility and performance of this algorithm depend critically on the form of the wave-function ansatz. Recently in Ref. [Nat. Commun. 10, 3007 (2019)], an algorithm termed ADAPT-VQE was introduced to build system-adapted ansätze with substantially fewer variational parameters compared to other approaches. This algorithm relies heavily on a predefined operator pool with which it builds the ansatz. However, Ref. [Nat. Commun. 10, 3007 (2019)] did not provide a prescription for how to select the pool, how many operators it must contain, or whether the resulting ansatz will succeed in converging to the ground state. In addition, the pool used in that work leads to state-preparation circuits that are too deep for a practical application on near-term devices. Here, we address all these key outstanding issues of the algorithm. We present a hardware-efficient variant of ADAPT-VQE that drastically reduces circuit depths using an operator pool that is guaranteed to contain the operators necessary to construct exact ansätze. Moreover, we show that the minimal pool size that achieves this scales linearly with the number of qubits. Through numerical simulations on H_4, LiH and H_6, we show that our algorithm ("qubit-ADAPT") reduces the circuit depth by an order of magnitude while maintaining the same accuracy as the original ADAPT-VQE. A central result of our approach is that the additional measurement overhead of qubit-ADAPT compared to fixed-ansatz variational algorithms scales only linearly with the number of qubits. Our work provides a crucial step forward in running algorithms on near-term quantum devices.

DOI: 10.1103/PRXQuantum.2.020310

I. INTRODUCTION

Finding the ground state of a many-body interacting electronic Hamiltonian is one of the most important problems in modern quantum chemistry and physics. As the dimension of the Hamiltonian scales exponentially with the number of particles, accurate classical simulations can only be performed for systems with few electrons. Although many classical computational techniques have been developed to approximate the ground electronic state, no classical method is available that can perform accurately for arbitrary systems with polynomial effort. While density-functional theory (DFT) [1,2] has been tremendously useful for unraveling the microscopic details of weakly correlated molecules and materials, the most accurate form (Kohn-Sham DFT [1]) relies on a single Slater determinant wave function, which fails to describe strong correlation with current functionals. Unlike DFT, whose accuracy relies on a density functional, which is not able to be systematically improved, wave-function-based methods such as configuration interaction (CI) or coupled cluster (CC) can be used to describe many-body systems with clear paths to arbitrary accuracy. However, in the presence of strong correlation such methods incur an exponential computational cost due to the large number of Slater determinants involved. Alternatively, methods based on tensor network states [3–7] (most notably density matrix renormalization group [8–10]) can exploit rank sparsity in low-energy states of one-dimensional systems to simulate strongly correlated systems, with polynomial cost. This polynomial scaling is lost in higher dimensions, and the computational cost again grows exponentially. A radically different approach is Feynman's proposal to study quantum systems using quantum computers [11]. Recent

*holuntang@vt.edu
†economou@vt.edu

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

reviews provide a comprehensive background and discuss new developments in quantum simulation with quantum computers [12–14].

The long-term method for simulating chemical systems with quantum computers is the quantum phase estimation algorithm (PEA) [15,16], which shows an exponential speedup over classical algorithms [17,18]. Since the number of gates, i.e., unitary operations, involved in this algorithm is very large, it requires a long coherent evolution, which can only be realized in fault-tolerant quantum computers [19]. These scalable, error-correcting devices may take decades to realize experimentally. In the meantime, the community is exploring algorithms that can be applied to existing and near-term processors, namely noisy intermediate-scale quantum (NISQ) devices [19].

A promising algorithm for NISQ hardware is the variational quantum eigensolver (VQE) [20,21]. VQE is a hybrid method that combines classical computational power with a quantum processor. Based on the variational principle in quantum physics, the VQE algorithm constructs a trial wave function by applying gates on a quantum device and estimates the average energy by measuring the Hamiltonian on that device. This measured energy is then minimized by tuning the quantum circuit. The circuit parameter optimization is performed by a classical computer, and quantum resources are only used for the classically intractable parts (state preparation and energy evaluation) of the calculation. Compared to PEA, VQE has much more modest requirements on the coherence times of the quantum processor, and it has already been realized on NISQ devices, such as superconducting qubits [22,23], photons [20], and trapped ions [24,25]. The accuracy of VQE is highly dependent on the explicit form of the wavefunction ansatz and can only obtain the exact ground-state energy if the ansatz is capable of representing any state in the subspace that contains the ground state. For example, if symmetries of the ground state are known, then one can construct an ansatz that represents any state in the corresponding symmetry subspace [26], guaranteeing that the exact ground state can be expressed in terms of the ansatz.

One of the commonly used ansätze for VQE is unitary coupled-cluster singles and doubles (UCCSD) [27, 28]. Stemming from the coupled cluster theory used in chemistry, this unitary version is more suitable for quantum circuit implementation. The “singles and doubles” in UCCSD means only single and double excitation operators are included in the ansatz, each carrying its own variational parameter. One drawback of UCCSD is that including all the singles and doubles operators leads to a (potentially unnecessarily) deep circuit and a large number of parameters to optimize. To reduce this complexity, several alternatives have been proposed, which attempt to keep only the most important operators [29–36]. A second drawback is that UCCSD is also generally not exact and suffers from ambiguities in operator ordering upon factorization into a product of exponentiated operators (Trotterization), which is a necessary step in converting the ansatz to a state-preparation circuit [37,38]. Another approach to building the ansatz is to use the most accessible gates in the quantum device, alternating single-qubit gates and two-qubit gates layer by layer. This is referred to as a hardware-efficient ansatz [23]. Although this approach lessens the demands on the quantum processor, the resulting wavefunction ansatz can lead to difficulties in parameter optimization [39]. This problem was addressed by constructing particle-conserving entangling gates instead of ordinary two-qubit gates [40,41] and symmetry-preserving circuits [26]. While these fixed-ansatz approaches can be applied to any problem and can reduce the number of variational parameters and circuit depths, further improvement may still be possible by tailoring the ansatz to a given simulation problem.

Recently, a new algorithm that provides a systematic method to build an ansatz dynamically was introduced. This algorithm, termed adaptive derivative assembled pseudo-Trotter (ADAPT) VQE [29], employs a predetermined pool of operators from which the ansatz is dynamically constructed. ADAPT-VQE differs substantially from approaches based on compiling fixed circuits. Instead, the ansatz is grown iteratively, such that at each step, the operator that affects the energy the most is added to the ansatz. Using fermionic operators as a pool, Ref. [29] demonstrated that ADAPT-VQE substantially outperforms UCCSD, in terms of both number of variational parameters and accuracy. This result demonstrates the promise of the ADAPT algorithm. However, due to the gate overhead of the fermion-to-spin mapping, the operators considered in Ref. [29] translate to a fairly large number of quantum gates, and, therefore, while the number of parameters is very low, the circuit depth (which is significantly reduced compared to UCCSD) may still be impractically large, limiting the applicability of ADAPT-VQE to NISQ devices. Even more importantly, it is not clear (i) how the operator pool should be chosen in general, (ii) how many operators it should contain, and (iii) what guarantees that the pool is complete, i.e., that it enables convergence to the ground state.

In this paper, we address these issues by introducing qubit-ADAPT-VQE, an algorithm that substantially reduces both the number of measurements and the circuit depths needed to achieve convergence. We term this algorithm qubit-adaptive derivative assembled problem-tailored (qubit-ADAPT) VQE, in contrast to the implementation in Ref. [29], which we refer to as fermionic-ADAPT in this paper. Through classical simulations of several different molecules, we demonstrate that compared to fermionic-ADAPT, qubit-ADAPT reduces the circuit depth by an order of magnitude while maintaining the same accuracy. This compactness of qubit-ADAPT and its iteratively grown ansatz are expected to reduce the

risk of barren plateaus. Moreover, we introduce a pool completeness criterion that determines whether a given pool will generate an exact ADAPT ansatz. We prove that the minimal number of pool operators that satisfy this condition grows linearly with the number of qubits. This is much smaller than the quartic scaling originally assumed in fermionic-ADAPT, and it demonstrates that the additional measurement overhead of ADAPT-VQE remains modest for larger systems (increasing only linearly over conventional, fixed-ansatz VQEs). Our results pave the way toward both practical and accurate VQE algorithms on NISQ devices.

This paper is organized as follows. First, we briefly review fermionic-ADAPT and estimate the corresponding circuit depth in Sec. II. In Sec. III, we introduce qubit-ADAPT and provide a detailed description of the operator pool. In Sec. III A, we compare the performance of qubit-ADAPT and fermionic-ADAPT through numerical simulations of H_4, LiH, and H_6 molecules. We show that the minimal size of the operator pool for qubit-ADAPT scales linearly in the number of qubits in Sec. III B. We conclude in Sec. IV. Details of the estimate for the circuit depth of fermionic-ADAPT and a constructive proof of the linear scaling of minimal complete pools are included in appendices.

II. CIRCUIT DEPTH ESTIMATE FOR FERMIONIC-ADAPT

The ADAPT ansatz is grown by one operator $\hat{\tau}_i = -\hat{\tau}_i^\dagger$ at each iteration, and after the nth iteration is given by

$$\left|\psi^{\mathrm{ADAPT}}(\vec{\theta})\right\rangle = e^{\theta_k \hat{\tau}_k} \dots e^{\theta_2 \hat{\tau}_2} e^{\theta_1 \hat{\tau}_1} \left|\psi^{\mathrm{HF}}\right\rangle, \tag{1}$$

where $\left|\psi^{\mathrm{HF}}\right\rangle$ is the HF state. These operators are selected from an operator pool defined upfront. At each iteration, the operator, which induces the maximum change to the energy, is selected. This energy response is represented by the gradient of the energy with respect to the corresponding parameter, i.e.,

$$\frac{\partial}{\partial \theta_i}\langle E\rangle = \langle\psi|\,[\hat{H}, \hat{\tau}_i]\,|\psi\rangle\,, \tag{2}$$

which can be measured on the quantum device. The ansatz keeps growing until the norm of the gradient vector,

$$||\vec{g}|| = \sqrt{\sum_i \left(\frac{\partial}{\partial \theta_i}\langle E\rangle\right)^2}, \tag{3}$$

is zero or, in practice, smaller than a chosen threshold, ϵ. Compared to ordinary VQE, ADAPT-VQE requires additional measurements to obtain the gradient at each iteration; the number of these measurements is roughly equal to the size of the pool times the number of terms in the Hamiltonian [42], although the number of measurements needed to compute the mean energy can be decreased using recently developed techniques [43].

In order to attain a compact circuit for the ADAPT ansatz, we want to minimize both the number of parameters and the circuit depth. While the first requirement can be satisfied by the general structure of the ADAPT-VQE algorithm, the second one is not guaranteed and depends on the chosen pool. Using fermionic operators, a *parameter-efficient* operator pool can be constructed from spin-adapted single excitation operators:

$$\hat{\tau}_1 \propto |\uparrow\rangle_a \langle\uparrow|_b + |\downarrow\rangle_a \langle\downarrow|_b - \text{h.c.}, \tag{4}$$

and double excitation operators:

$$\begin{aligned}\hat{\tau}_{2,T} &\propto |T,1\rangle_{pq} \langle T,1|_{rs} + |T,-1\rangle_{pq} \langle T,-1|_{rs} \\ &\quad + |T,0\rangle_{pq} \langle T,0|_{rs} - \text{h.c.}, \\ \hat{\tau}_{2,S} &\propto |S,0\rangle_{pq} \langle S,0|_{rs} - \text{h.c.},\end{aligned} \tag{5}$$

where a, b, p, q, r, s are spatial orbitals and T, S refer to triplets and singlets formed by p, q or r, s. To implement ADAPT-VQE on qubits, we have to map these fermionic operators to Pauli operators, which has the consequence that a *parameter-efficient* pool is most likely not *gate-efficient*. In this paper, we employ the Jordan-Wigner (JW) mapping, i.e., $\hat{a}_i \to \prod_{j=0}^{i-1} Z_j (X_i - iY_i)$, although other mappings [44–46] could be used instead.

In the JW mapping, a double excitation operator may contain more than one fermionic operator $a_p^\dagger a_q^\dagger a_r a_s - \text{h.c.}$, which is transformed into at most eight Pauli strings. (A product of four fermionic operators gives 16 terms in total, but each symmetric term is cancelled by its Hermitian conjugate [35].) These excitation operators conserve S^2, S_z and particle number by summing a number of Pauli strings together, which results in a high gate count per excitation operator.

In order to make a rough estimate of the number of CNOT gates involved in one operator from the fermionic pool, we consider the generalized singles-doubles excitation fermionic pool. "Generalized" indicates that the excitations are not restricted to occur only from occupied to virtual orbitals. Rather, all combinations of excitations are included. To obtain the gate count, we perform first-order Trotterization of each unitary, i.e., $e^{\theta_i \hat{\tau}_i} \approx \prod_j e^{\theta_j \hat{P}_j}$, where P_j are the Pauli strings appearing in $\hat{\tau}_i$ after the JW mapping. For simplicity, we assume only double excitation operators are picked by the algorithm, which makes this a conservative estimate as any single excitation included would result in a smaller CNOT-to-parameter count ratio. The number of CNOTs needed for a single Pauli string $\hat{P}_i$ with length q is $2(q-1)$ [48]. Here, the length q is the number of nonidentity Pauli operators in

the string. The average number of CNOTs involved in a spin-adapted double excitation operator is approximately $\bar{N}_{\text{Pauli}}(6 + 2\bar{N}_Z)\bar{N}_{\text{spin}}$, where $\bar{N}_{\text{Pauli}}$ is the average number of Pauli strings in a doubles operator $a_p^\dagger a_q^\dagger a_r a_s - \text{h.c.}$, $\bar{N}_Z$ is the average number of Pauli Z's in a Pauli string due to the anticommutation relation of fermionic operators, and $\bar{N}_{\text{spin}}$ is the average number of doubles operators summed in a spin-adapted operator [Eq. (5)]. By Appendix A, these quantities are given by

$$\bar{N}_{\text{Pauli}} = \frac{8m(5m-7)}{m(5m-1)-8}, \tag{6}$$

$$\bar{N}_Z = \frac{12m^3 - 30m^2 + 34m - 20}{3(5m^2 - m - 8)}, \tag{7}$$

$$\bar{N}_{\text{spin}} = \frac{5m^2 - m - 8}{m^2 + m}, \tag{8}$$

where m is the number of spatial orbitals. For large m, the number of CNOTs in a spin-adapted doubles operator is approximately $64m$.

III. QUBIT-ADAPT

Because the spin-adapted fermionic operators each introduce a large number of CNOTs into the state-preparation circuit, we are motivated to construct a new pool consisting of operators that involve fewer CNOTs. One way is to break down the spin-adapted fermionic operators after the JW mapping and choose the individual Pauli strings as the operator pool $\hat{\tau} = \hat{P} = i\prod_i p_i$, $p_i \in \{X, Y, Z\}$. This more hardware-efficient choice effectively reduces $\bar{N}_{\text{spin}}$ and $\bar{N}_{\text{Pauli}}$ to 1, while it contains the same basic elements as the spin-adapted fermionic pool. An important property of this qubit pool is that it contains only Pauli strings with odd numbers of Y's because the fermionic operators are real, hence $\hat{P} = i\prod_i p_i$ has to be real. We refer to these as "odd" Pauli strings. The remaining "even" Pauli strings have no effect on the energy. This is because $\hat{H}$ is symmetric (time-reversal symmetry is preserved), and so the expectation value of the commutator in Eq. (2) will vanish for real $|\psi\rangle$ if $\hat{\tau}_i$ is symmetric (i.e., even). We can therefore restrict $\hat{\tau}_i$ to the odd Pauli strings, which are antisymmetric. Using such operators in the pool also ensures that the ansatz remains real throughout the qubit-ADAPT algorithm, which should be the case when $\hat{H}$ is time-reversal symmetric. The length of these strings ranges from 2 to n due to the Pauli-Z chain responsible for the fermionic anticommutation relation. To further reduce gate depth, we remove these Pauli-Z chains from the operators. Numerically, we find that these two pools (with and without Z chains) perform similarly. The fact that qubit-ADAPT does not require the Z chains in order to achieve convergence suggests that this algorithm avoids issues related to the long-range correlations inherent in the JW mapping. The pool without Z chains gives Pauli strings with maximum length 4. The size of this pool is much smaller than the full set of Pauli strings with maximum length 4, as we only pick operators that already appear in the fermionic pool, which are capable of transforming $|\psi^{\text{HF}}\rangle$ to the ground state. We refer to this reduced pool as the "qubit pool". Below, we demonstrate that this pool produces significantly shallower state preparation circuits compared to fermionic-ADAPT. We also show that the size of the qubit pool can be reduced dramatically down to a size that scales only linearly in the number of qubits, which substantially cuts down on the number of measurements needed to run qubit-ADAPT.

A. Numerical simulations

We compare the performance of the qubit pool to the fermionic pool in terms of the number of parameters and the number of CNOTs for different molecules, H_4, LiH, and H_6 (Fig. 1). In each case, we choose a bond distance such that correlation effects are significant: $r = 1.5$ Å for H_4 and H_6, and $r = 2$ Å for LiH. All the calculations are performed using an STO-3G basis and start from restricted Hartree-Fock orbitals without using a frozen-core approximation such that we have eight spin orbitals for H_4 and 12 spin orbitals for each of the other two molecules. All CNOT gate counts are obtained using the qiskit command `count_ops`, where we consider the case of all-to-all qubit connectivity for concreteness. Although these counts will increase as the connectivity is reduced, we do not expect the relative performance of the two algorithms to change significantly. In the case of fermionic-ADAPT, we show gate counts before and after transpilation is used. The counts before transpilation agree well with the estimates obtained in Appendix A.

It is evident from Fig. 1 that in the case of the qubit pool, more parameters are used compared to the fermionic pool. On the other hand, the number of CNOTs is reduced significantly, by about an order of magnitude in the case of H_6. Switching from the fermionic pool to the qubit pool increases the number of parameters in the ansatz, which is the price for compressing the circuit depth. However, the increase in parameter number increases only the required classical computational power during the classical optimization, while the decrease in circuit depth reduces the demands on the quantum processor. For NISQ devices, the number of CNOTs that can be implemented within the coherence time is very limited, so the ability of qubit-ADAPT to divert more of the computational cost away from the quantum processor and onto the classical optimizer should be advantageous overall.

To evaluate the effectiveness of using the gradient as a means to grow the ADAPT ansatz, we compare the performance of qubit-ADAPT with random operator orderings drawn from the same qubit pool. In Fig. 2, we

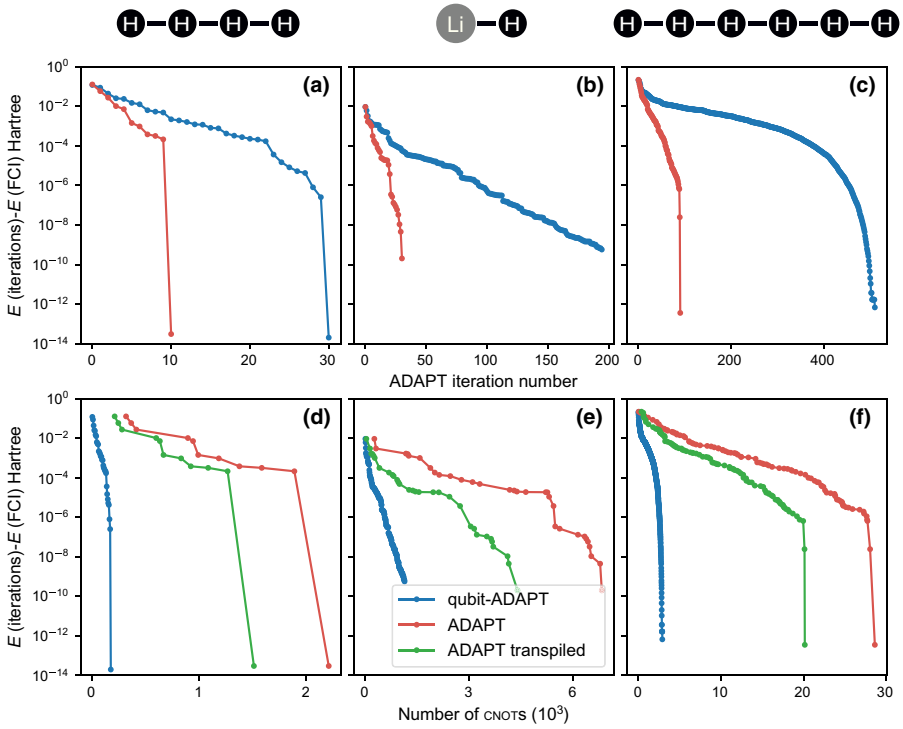

FIG. 1. Ground-state energy difference between ADAPT and the exact (FCI) result for (a),(d) H_4 at bond distance 1.5 Å, (b),(e) LiH at bond distance 2 Å, and (c),(f) H_6 at bond distance 1.5 Å. The bond distances are chosen to ensure that correlation effects are substantial. All the calculations are performed using an STO-3G basis and start from restricted Hartree-Fock orbitals without using a frozen-core approximation such that we have eight spin orbitals for H_4, 12 spin orbitals for LiH, and 12 spin orbitals for H_6. Results are shown for qubit-ADAPT (blue), fermionic-ADAPT (orange), and transpiled fermionic-ADAPT (green). (a)–(c) Energy difference as a function of ADAPT iterations (which is equal to the number of variational parameters). (d)–(f) Energy difference as a function of the number of CNOTs in the corresponding state-preparation circuit. All CNOT gate counts are obtained using the qiskit command `count_ops`. The transpiled counts are obtained using `qiskit.transpile` [47] with heavy optimization, which includes canceling back-to-back CNOTs and "commutative cancellation" (see Supplementary Material [56] for the code). Although qubit-ADAPT requires more variational parameters, it entails significantly fewer CNOTs compared to fermionic-ADAPT.

see that qubit-ADAPT always converges much faster than the random orderings for both H_4 and LiH. When random orderings are used for H_4, convergence to the ground state requires 46–78 parameters, compared to only 30 parameters for qubit-ADAPT. In the case of LiH, the random orderings require more than 3 times as many parameters as qubit-ADAPT to converge. We can provide only a lower bound on the number of parameters needed to converge the random-ordering results in Fig. 2(b) due to the long computational times needed in this case. These findings suggest that the role of the gradient selection in qubit-ADAPT is crucial for larger problems.

These findings suggest that qubit-ADAPT will likely avoid issues with barren plateaus. Studies of barren plateaus have revealed that circuit depth plays a key role [39,49], indicating that qubit-ADAPT-VQE may avoid this issue due to the compactness of the ansätze it produces. In fact, there are two recent papers that follow similar strategies as ADAPT-VQE, and which demonstrate that variable ansätze are much more resilient to barren plateaus. One of these works, the so-called layerwise learning, grows the ansatz by adding layers with fixed operator structure, thus optimizing the parameters gradually [50]. Similar conclusions were also drawn in Ref. [51]. The features that protect against barren plateaus in these works are also naturally present in the qubit-ADAPT-VQE ansatz-growth strategy. Additionally, it has been demonstrated that correlating random parameters in the ansatz can also prevent barren plateaus [52]. Again, we expect that qubit-ADAPT-VQE may also have this feature since the parameters of the ansatz at each layer are related to the parameters in the previous iteration. More

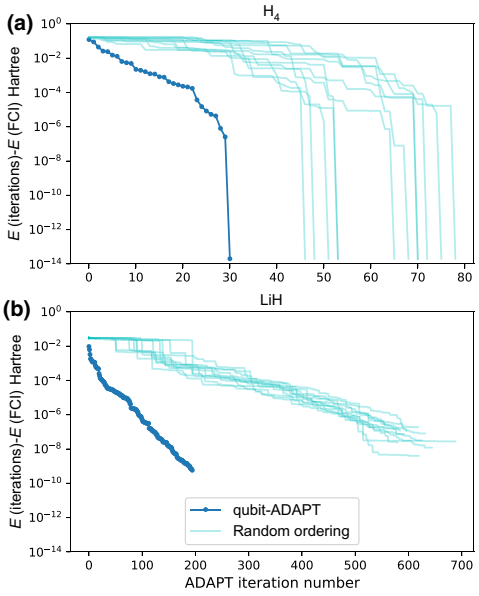

FIG. 2. Energy error of qubit-ADAPT versus ansätze with random operator orderings as a function of the number of iterations for (a) H_4 and (b) LiH. The qubit pool is used in all cases, and the orbital bases and bond distances are chosen as in Fig. 1.

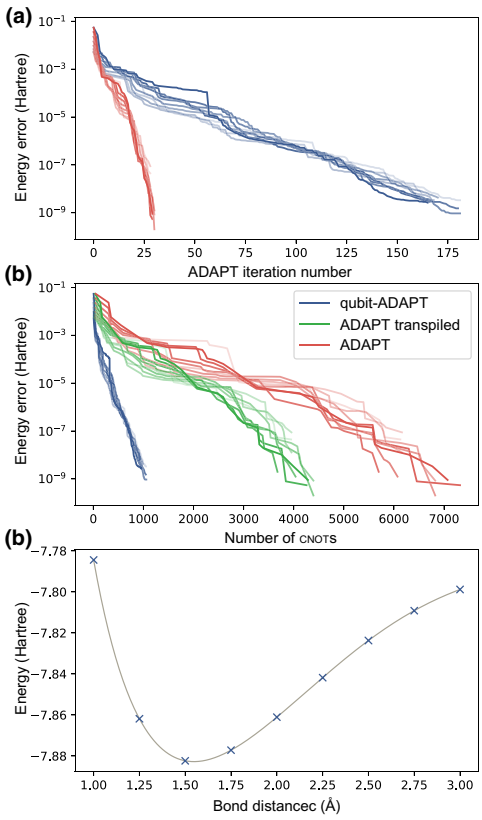

FIG. 3. qubit-ADAPT and fermionic-ADAPT energy error of LiH for various bond distances. The bond distance varies from 1 to 3 Å with darker curves in (a), (b) corresponding to larger bond distance. The energy error is plotted against (a) the number of parameters and (b) the number of CNOTs. The performance remains similar across different bond lengths and hence across different amounts of correlation in the ground state. The orbital basis used is the same as in Fig. 1. (c) The final energies obtained by qubit-ADAPT (points marked by ×) for each bond distance considered are plotted along with the exact ground-state energies (solid curve).

generally, barren plateaus are believed to arise because the parameterized ansatz has too much expressivity in the space of unitaries. Since qubit-ADAPT-VQE is problem tailored, each ansatz is only intended to be used for a single Hamiltonian. This is unlike the "hardware-efficient" approach where one ansatz is intended to solve multiple Hamiltonians. In this way, the ansätze resulting from qubit-ADAPT-VQE have lower expressivity and thus less susceptibility to barren plateaus.

Figure 3 shows that the performance of qubit-ADAPT remains essentially the same across different bond distances. Results are shown for LiH, where the bond distance is varied from 1 to 3 Å. The fact that the curves corresponding to different bond lengths largely overlap one another shows that the rate of convergence does not change significantly. Because the amount of correlation in the ground state depends on the bond distance (more correlations tend to arise as the bonds are stretched), this suggests that the convergence of qubit-ADAPT is not sensitive to the strength of correlations. This finding indicates that qubit-ADAPT is a promising approach to studying strongly correlated systems.

B. Operator pool reduction

So far, a drawback of the qubit pool is its large size compared to the fermionic pool, which in turn leads to proportionally more measurements at each iteration. Even though the qubit pool is defined from the fermionic pool, which is a small subset of the full Pauli group, the pool size grows quickly with the number of orbitals. However, many of these operators are redundant in the sense that eliminating them has no effect on the convergence of the algorithm. For example, there are pairs of operators that are related by a global rotation, e.g., $X_0Y_1Y_2Y_3$ and $Y_0X_1X_2X_3$, so we need only to keep one of them in the pool; discarding the other does not have any significant effect on the performance of qubit-ADAPT.

The redundancy in the pool can be illustrated by removing randomly selected operators from the pool and monitoring the impact on convergence. First we randomly remove 3/4 of the operators in the pool. As shown in Fig. 4, despite this large reduction of the pool size, the performance of the algorithm is similar to that with the original pool. However, as is also evident in Fig. 4, if we further remove more operators, the pool is sometimes incomplete, and the energy may not converge to the ground-state energy. We can understand the tolerance of the algorithm to the drastic reduction of the pool by studying the Hilbert space spanned by the pool operators. Starting from the expression for $\left|\psi^{\mathrm{ADAPT}}(\vec{\theta})\right\rangle$ in Eq. 1 and using the Baker-Campbell-Hausdorff formula repeatedly, we obtain

$$\left|\psi^{\mathrm{ADAPT}}(\vec{\theta})\right\rangle = e^{\sum_i \phi_i A_i} \left|\psi^{\mathrm{HF}}\right\rangle, \tag{9}$$

where the A_i include all the pool operators and their commutators, and the ϕ_i are functions of $\vec{\theta}$. Therefore, the Hilbert space spanned by the pool is determined by the set $\{A_i\}$. Note that if the original pool is comprised of odd Pauli strings, then so is $\{A_i\}$. If the operators in $\{A_i\}$ can transform the reference state to any real state in the n-qubit Hilbert space, then the qubit-ADAPT ansatz is guaranteed to be exact, and it is capable of converging to the ground state. (Here, the only symmetry we impose is time reversal.) Note that if $\{A_i\}$ includes all the odd Pauli strings [of which there are $2^{n-1}(2^n-1)$, which scales exponentially with the number of qubits], then we could create an arbitrary orthogonal transformation in Eq. (9). However, spanning the Hilbert space requires only a subset of the odd Pauli strings, because only 2^n-1 real parameters are required to create an arbitrary real state. In particular, we need $\{A_i\}$ to be such that for an arbitrary state $|\psi\rangle$, the states $A_i|\psi\rangle$ form a complete basis. In this case, we refer to $\{A_i\}$ as a complete basis of operators.

The problem then is to determine the minimal pools that produce a complete basis of operators. For a given pool, we define the overlap matrix as $M_{ij} = \langle\psi|A_i^\dagger A_j|\psi\rangle$ where $|\psi\rangle$ is an arbitrary real state. If the rank of M satisfies $r(M) \geq 2^n - 1$, this pool is called complete. To determine the smallest complete pools, we randomly generate many different pools of increasing size and compute $r(M)$ in each case to test for completeness. We did this for up to seven qubits. Our numerical investigations reveal that, surprisingly, the minimal pool size required for the overlap matrix to have the required rank of 2^n-1 is only $2n-2$. This is evident in Fig. 5, which shows the fraction of pools that are complete for various pool sizes and numbers of qubits. In each case, complete pools are found for pool sizes that contain at least $2n-2$ operators. Note that not only is this much smaller than the Hilbert space, it is also much smaller than the size of the fermionic pool, which scales like n^4.

Randomly selecting pools of size $2n-2$ and testing for completeness by computing $r(M)$ is a numerically intensive process that quickly becomes infeasible as the number of qubits increases. This is further exacerbated by the fact that the fraction of complete pools of size $2n-2$ becomes smaller as n increases, as is evident in Fig. 5, which raises the question of whether complete pools of size $2n-2$ even exist for large values of n. In Appendix B, we prove analytically using induction that complete pools of size $2n-2$ exist for any n. In fact, the proof is constructive: We present two families of minimal pools that are provably complete for any number of qubits. One of these pools, which we call $\{V_j\}_n$ where $j = 1, \ldots, 2n-2$, is defined recursively as $\{V_j\}_n = \{Z_n\{V_k\}_{n-1}, iY_n, iY_{n-1}\}$, starting from the pool for $n=2$ qubits: $\{V_j\}_2 = \{iZ_2Y_1, iY_2\}$. This pool is comprised of generators for single-qubit Y rotations and conditional Y

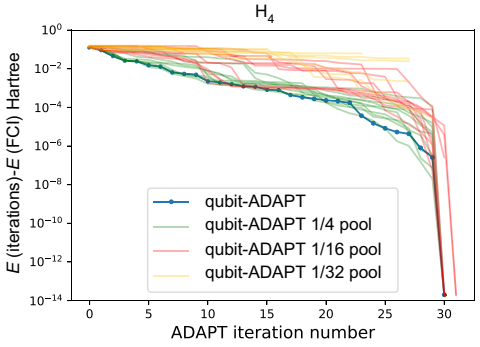

FIG. 4. qubit-ADAPT energy error of H_4 using the full qubit pool, and randomly chosen 1/4, 1/16, and 1/32 pools are plotted against the number of iterations. For the 1/4 and 1/16 pools, the algorithm can converge for almost every run. For the 1/32 pool, most of the runs can reach only an energy error of 10^{-3}. The orbital basis and bond distance are chosen as in Fig. 1.

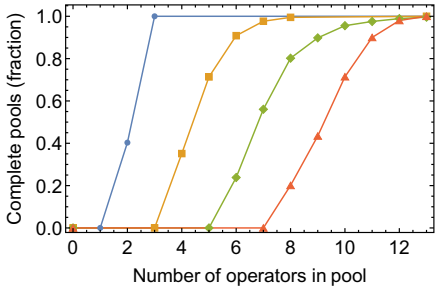

FIG. 5. The fraction of pools that are complete as a function of pool size for two (blue circles), three (yellow squares), four (green diamonds), and five (red triangles) qubits.

rotations, and it contains operators that act on up to all n qubits. In Appendix B, we prove that these operators are sufficient to rotate any real state to any other real state in the Hilbert space. As shown in Appendix C, $\{V_j\}_n$ can be mapped to a second family of minimal complete pools, which we call $\{G_j\}_n$, that has a very local structure. This pool contains all two-qubit operators of the form $iZ_{k+1}Y_k$ that act on two neighboring qubits labeled by k and $k+1$. There are $n-1$ such operators. $\{G_j\}_n$ also contains all single-qubit Pauli iY operators except on the first qubit. Therefore, this pool contains $2n-2$ operators. In Appendix C, we show that the G_j can be obtained from commutators of the V_j for any n, and so the completeness of the former follows from that of the latter.

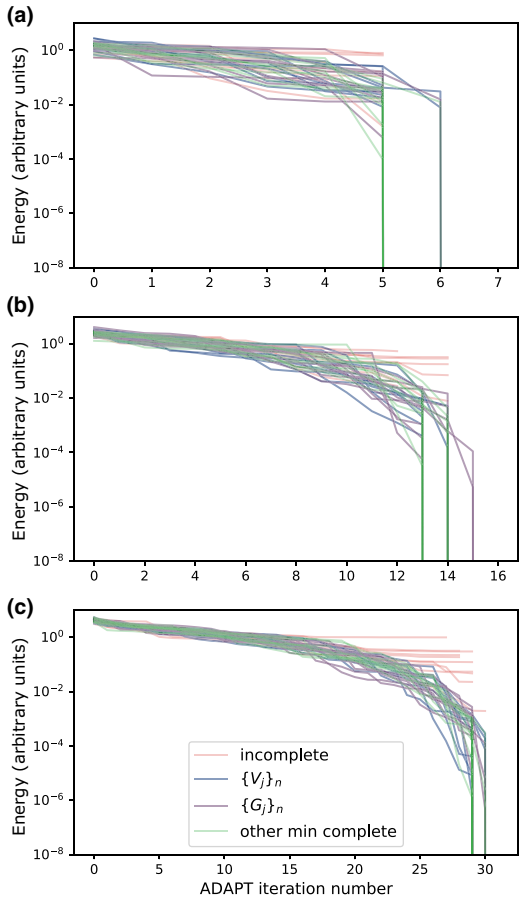

FIG. 6. Energy error curves for three different kinds of minimal complete pools versus incomplete pools. Results for three, four, and five-qubit random Hamiltonians are shown in (a), (b), and (c), respectively.

To investigate the importance of the pool being complete, we run qubit-ADAPT for random real Hamiltonians of three, four, and five qubits with random initial states and for pools consisting of $2n-2$ operators randomly chosen from the set of odd Pauli strings. We also run simulations using the minimal complete pools $\{V_j\}_n$ and $\{G_j\}_n$. The operator coefficients in each Hamiltonian (which is taken to be real and symmetric) are obtained by sampling uniformly in the range $[-2,2]$ with ten samples for each of these three cases. The corresponding energy error curves are illustrated in Fig. 6. Each curve is the result for a different random Hamiltonian. All of the curves that fail to converge correspond to incomplete pools. For these cases, even though the gradient goes to zero the ground state is not reached because important operators are never generated. On the other hand, the runs with complete pools always converge, highlighting the importance of this criterion. For the cases considered, we find that 20%–40% of pools containing $2n-2$ operators are complete. These findings shed light on the question of what constitutes a good operator pool for ADAPT-VQE. These points were not appreciated in the original paper [29], and moreover, they suggest that the fermionic pool used in that work is overcomplete. Furthermore, the fact that the minimal pool size is linearly proportional to n means that the number of additional measurements needed for each step of qubit-ADAPT is also linear in n. Thus, the extra measurement overhead of qubit-ADAPT remains modest as the problem size increases.

IV. CONCLUSIONS

In conclusion, we introduce a more efficient and NISQ-compatible version of the ADAPT-VQE algorithm, called qubit-ADAPT. The basic idea of qubit-ADAPT is to use a pool consisting of Pauli strings (rather than fermionic operators) so that the number of CNOT gates associated with each pool operator is reduced. We establish a completeness condition that guarantees that a pool will generate an exact ADAPT ansatz, and we prove that the smallest possible pool that obeys this criterion scales linearly with the number of qubits. We also provide a constructive approach to generating minimal complete pools for an arbitrary number of qubits. These results remove the ad hoc elements of the ADAPT algorithm and lead to a substantial reduction in the depths of state-preparation circuits and in the number of measurements needed to run ADAPT-VQE on realistic hardware.

ACKNOWLEDGMENTS

This research is supported by the National Science Foundation (Award No. 1839136) and the US Department of Energy (Award No. DE-SC0019199).

APPENDIX A: CNOT ESTIMATION

From the structure of the pool operators, we can try estimating the number of CNOT gates as a function of the number of variational parameters. These estimates are described below.

1. Qubit-ADAPT

In order to get a conservative estimate on the CNOT count, we assume that all operators picked are four-qubit Pauli strings, since exponentiating an n-qubit Pauli string requires $2(n-1)$ CNOT gates, the number of CNOT gates is

$$\bar{N}_{\text{CNOT}} = 6. \tag{A1}$$

2. Fermionic-ADAPT

In the case of fermionic-ADAPT, the number of CNOTs used can be estimated as

$$\bar{N}_{\text{CNOT}} = (6 + 2 \times \bar{N}_Z) \times \bar{N}_{\text{Pauli}} \times \bar{N}_{\text{spin}}, \tag{A2}$$

where $\bar{N}_Z$ is the average number of Pauli Z used for fixing the antimsymmetry of fermionic operators, $\bar{N}_{\text{spin}}$ is the average number of fermionic operators in each spin-adapted group and $\bar{N}_{\text{Pauli}}$ is the average number of Pauli strings from each fermionic operator.

$$\begin{aligned}
\bar{N}_{\text{spin}} &= \frac{\text{Totalno.offermionicterms}(a_p^\dagger a_q^\dagger a_r a_s)}{\text{Total spin adapted double ops in pool}} = \text{no. of fermionic terms in a spin adapted op,} \\
\bar{N}_{\text{Pauli}} &= \frac{\text{Total number of Pauli strings after JW}}{\text{Total fermionic terms}} = \text{no. of Pauli strings from each fermionic term,} \\
\bar{N}_Z &= \frac{\text{TotallengthofPauli}Z\text{fromallfermionic}}{\text{Total no. of fermionic terms}} = \text{increases the length of Pauli strings.}
\end{aligned} \tag{A3}$$

To compute these total counts, we can first classify the combination of spatial orbitals into five groups:

1. $p \neq q \neq r \neq s$.
2. One of p, q = one of r, s.
3. $p = q \neq r \neq s$.
4. $p = q = r \neq s$.
5. $p = q \neq r = s$.

Since some of them can have both triplet and singlet while the others can only have singlet. And the number of fermionic terms in spin-adapted grouping also depends on the combination of spatial orbitals.

For group 1, the triplet operator reads as

$$\begin{aligned}
\hat{\tau}_{2,T} &\propto |T,1\rangle_{pq}\langle T,1|_{rs} + |T,-1\rangle_{pq}\langle T,-1|_{rs} \\
&\quad + |T,0\rangle_{pq}\langle T,0|_{rs} - \text{h.c.} \\
&= |pq\rangle\langle rs| + \frac{1}{2}(|p\bar{q}\rangle + |\bar{p}q\rangle)(\langle r\bar{s}| + \langle \bar{r}s|) \\
&\quad + |\bar{p}\bar{q}\rangle\langle \bar{r}\bar{s}| - \text{h.c.} \\
&\to a_p^\dagger a_q^\dagger a_r a_s + \frac{1}{2}\Big(a_p^\dagger a_{\bar{q}}^\dagger a_r a_{\bar{s}} + a_p^\dagger a_{\bar{q}}^\dagger a_{\bar{r}} a_s \\
&\quad + a_{\bar{p}}^\dagger a_q^\dagger a_r a_{\bar{s}} + a_{\bar{p}}^\dagger a_q^\dagger a_{\bar{r}} a_s\Big) + a_{\bar{p}}^\dagger a_{\bar{q}}^\dagger a_{\bar{r}} a_{\bar{s}} - \text{h.c.},
\end{aligned}$$

so there are six terms in total (12 terms including Hermitian conjugate). For the singlet,

$$\begin{aligned}
\hat{\tau}_{2,S} &\propto |S,0\rangle_{pq}\langle S,0|_{rs} - \text{h.c.} \\
&= \frac{1}{2}(|p\bar{q}\rangle - |\bar{p}q\rangle)(\langle r\bar{s}| - \langle \bar{r}s|) - \text{h.c.} \\
&\to \frac{1}{2}\Big(a_p^\dagger a_{\bar{q}}^\dagger a_r a_{\bar{s}} - a_p^\dagger a_{\bar{q}}^\dagger a_{\bar{r}} a_s - a_{\bar{p}}^\dagger a_q^\dagger a_r a_{\bar{s}} + a_{\bar{p}}^\dagger a_q^\dagger a_{\bar{r}} a_s\Big) \\
&\quad - \text{h.c.},
\end{aligned}$$

so there are four terms in total (eight terms including Hermitian conjugate). For group 2, the triplet operator reads as (for $q = r$)

$$\begin{aligned}
\hat{\tau}_{2,T} \to\ & a_p^\dagger a_q^\dagger a_q a_s + \frac{1}{2}\Big(a_p^\dagger a_{\bar{q}}^\dagger a_q a_{\bar{s}} + a_p^\dagger a_{\bar{q}}^\dagger a_{\bar{q}} a_s \\
& + a_{\bar{p}}^\dagger a_q^\dagger a_q a_{\bar{s}} + a_{\bar{p}}^\dagger a_q^\dagger a_{\bar{q}} a_s\Big) \\
& + a_{\bar{p}}^\dagger a_{\bar{q}}^\dagger a_{\bar{q}} a_{\bar{s}} - \text{h.c.},
\end{aligned}$$

while the singlet operator is

$$\begin{aligned}
\hat{\tau}_{2,S} \to\ & \frac{1}{2}\Big(a_p^\dagger a_{\bar{q}}^\dagger a_q a_{\bar{s}} - a_p^\dagger a_{\bar{q}}^\dagger a_{\bar{q}} a_s - a_{\bar{p}}^\dagger a_q^\dagger a_q a_{\bar{s}} \\
& + a_{\bar{p}}^\dagger a_q^\dagger a_{\bar{q}} a_s\Big) - \text{h.c.}
\end{aligned}$$

For group 3, there is only singlet

$$\begin{aligned}\hat{\tau}_{2,S} \to \frac{1}{2}\Big(&a_p^\dagger a_{\bar{p}}^\dagger a_r a_{\bar{s}} - a_p^\dagger a_{\bar{p}}^\dagger a_{\bar{r}} a_s - a_{\bar{p}}^\dagger a_p^\dagger a_r a_{\bar{s}} \\ &+ a_{\bar{p}}^\dagger a_p^\dagger a_{\bar{r}} a_s\Big) - \text{h.c.} \\ = a_p^\dagger a_{\bar{p}}^\dagger a_r a_{\bar{s}} &+ a_p^\dagger a_{\bar{p}}^\dagger a_{\bar{r}} a_s - \text{h.c.}\end{aligned}$$

For group 4, there is only singlet

$$\begin{aligned}\hat{\tau}_{2,S} \to \frac{1}{2}\Big(&a_p^\dagger a_{\bar{p}}^\dagger a_p a_{\bar{s}} - a_p^\dagger a_{\bar{p}}^\dagger a_{\bar{p}} a_s - a_{\bar{p}}^\dagger a_p^\dagger a_p a_{\bar{s}} \\ &+ a_{\bar{p}}^\dagger a_p^\dagger a_{\bar{p}} a_s\Big) - \text{h.c.} \\ = a_{\bar{p}}^\dagger a_p^\dagger a_p a_{\bar{s}} &+ a_p^\dagger a_{\bar{p}}^\dagger a_{\bar{p}} a_s - \text{h.c.}\end{aligned}$$

For group 5, there is only singlet

$$\begin{aligned}\hat{\tau}_{2,S} \to \frac{1}{2}\Big(&a_p^\dagger a_{\bar{p}}^\dagger a_r a_{\bar{r}} - a_p^\dagger a_{\bar{p}}^\dagger a_{\bar{r}} a_r - a_{\bar{p}}^\dagger a_p^\dagger a_r a_{\bar{r}} \\ &+ a_{\bar{p}}^\dagger a_p^\dagger a_{\bar{r}} a_r\Big) - \text{h.c.} \\ = 2a_p^\dagger a_{\bar{p}}^\dagger a_r a_{\bar{r}} &- \text{h.c.}.\end{aligned}$$

Then we can perform Jordan-Wigner transformation to each of these terms. Each anti-Hermitian pair of fermionic operators having four different indices, i.e., $a_i^\dagger a_j^\dagger a_k a_l -$ h.c., would result in eight Pauli strings, i.e.,

$$\begin{aligned}&a_i^\dagger a_j^\dagger a_k a_l - \text{h.c.} \\ &= 2i \prod_n Z_n \Big(Y_i X_j X_k X_l + X_i Y_j X_k X_l - X_i X_j Y_k X_l - X_i X_j X_k Y_l \\ &\quad - X_i Y_j Y_k Y_l - Y_i X_j Y_k Y_l + Y_i Y_j X_k Y_l + Y_i Y_j Y_k X_l \Big).\end{aligned}$$

The other case is two of the four indices are the same,

$$\begin{aligned}a_i^\dagger a_j^\dagger a_j a_l - \text{h.c.} &= \prod_n Z_n (X_i + iY_i) Z_j (X_l - iY_l) - \text{h.c.} \\ &= 4i \prod_n Z_n \Big(Y_i Z_j X_l - X_i Z_j Y_l \Big),\end{aligned}$$

only two Pauli strings left. From the expression of the five cases' fermionic operators, we can then determine the average number of Pauli strings per anti-Hermitian pair for each case.

For the length of Pauli Z, it is equivalent to count the interval between i and j and between k and l where $i \le j \le k \le l$, we calculate them for separate cases, for m spatial orbitals,

1. $p \neq q \neq r \neq s$:
The total number of Pauli Z is given by the sum

$$\begin{aligned}&\frac{1}{2}\binom{4}{2} \times (6+4) \\ &\quad \times \sum_{m_1<m_2<m_3<m_4} \Big[(2m_4 - 2m_3 - 1)(2m_2 - 2m_1 - 1) \Big] \\ &= \frac{m}{2}(2m^4 - 15m^3 + 40m^2 - 45m + 18).\end{aligned}$$

The prefactor $\frac{1}{2}\binom{4}{2}$ is counting the ways we choose two of $\{p,q,r,s\}$ to be in the dagger side, while the prefactor $(6+4)$ is from the number of fermionic operators for each combination. If we have $\{i,j,k,l\} = \{1,2,3,4\}$, which is in group 5, the Jordan-Wigner transformation would bring a Pauli-Z chain $Z_1 Z_2 Z_3$ for the creation and annihilation operator on index 4 and a Pauli-Z chain $Z_1 Z_2$ for operator on index 3, so only Z_3 can survive, but it would combine with the $X_3 \pm iY_3$ from the operator on index 3, so there is no Pauli Zs effectively. Therefore, we should count the number of integers can be fitted in the intervals $2m_4 - 2m_3$ and $2m_2 - 2m_1$, that explains the terms -1 in the calculation.

In this counting, we treat all the spin orbitals to be spin up (even number indices) even though each term in $|T,0\rangle\langle T,0|$ and $|S,0\rangle\langle S,0|$ has two spin-down and two spin-up operators. It is because whenever there are terms with both spin-up and spin-down spin orbitals, there would be another term with the spin totally flipped, so the counting of these two terms would cancel each other. For example, consider $p < q < r < s$, the $|T,0\rangle\langle T,0|$ terms read as $\frac{1}{2}\left(a_p^\dagger a_{\bar{q}}^\dagger a_r a_{\bar{s}} + a_p^\dagger a_{\bar{q}}^\dagger a_{\bar{r}} a_s + a_{\bar{p}}^\dagger a_q^\dagger a_r a_{\bar{s}} + a_{\bar{p}}^\dagger a_q^\dagger a_{\bar{r}} a_s\right) - \text{h.c.}$, the first term and the last term are different by a total spin flip, the number of Pauli Z would be $2s+1-2r-1+2q+1-2p-1$ and $2s-(2r+1)-1+2q-(2p+1)-1$, so the $+1$ due to the odd index switches sign by the spin flip. Therefore, the average number of Pauli Z of this pair is the $2s-2r-1+2q-2p-1$, which is the number of Pauli Z in the all spin-up term $a_p^\dagger a_q^\dagger a_r a_s$ (or the all spin-down term).

2. One of p,q = one of r,s.
The total number of Pauli Z is given by the sum

$$\begin{aligned}10 \times \Big[&\sum_{m_1=m_2<m_3<m_4} (2m_4 - 2m_3 - 1) \\ &+ \sum_{m_1<m_2=m_3<m_4} (2m_4 - 2m_3 - 1 + 2m_3 - 2m_1 - 1) \\ &+ \sum_{m_1<m_2<m_3=m_4} (2m_2 - 2m_1 - 1) \Big] \\ &= \frac{10m}{3}(m^3 - 4m^2 + 5m - 2).\end{aligned}$$

Here, we do not have the counting prefactor in group 1 as we already decide that one of the two identical indices is in the dagger side.

3. $p = q \neq r \neq s$.

The total number of Pauli Z is given by the sum

$$\begin{aligned} 2 \times \Big[& \sum_{m_1=m_2<m_3<m_4} (2m_4 - 2m_3 - 1) \\ & + \sum_{m_1<m_2=m_3<m_4} (2m_4 - 2m_3 - 1 + 2m_3 - 2m_1 - 1) \\ & + \sum_{m_1<m_2<m_3=m_4} (2m_2 - 2m_1 - 1) \Big] \\ & = \frac{2m}{3}(m^3 - 4m^2 + 5m - 2). \end{aligned}$$

It is only different from case 2 by the prefactor 2 instead of 10 as it has two fermionic operators in total.

4. $p = q = r \neq s$.

The total number of Pauli Z is given by the sum

$$\begin{aligned} 2 \times \Big[& \sum_{m_1=m_2=m_3<m_4} (2m_4 - 2m_1 - 1) \\ & + \sum_{m_1<m_2=m_3=m_4} (2m_4 - 2m_1 - 1) \Big] \\ & = \frac{2m}{3}(2m^2 - 3m + 1). \end{aligned}$$

5. $p = q \neq r = s$.

Since the operators have a form $a_p^\dagger a_{\bar{p}}^\dagger a_r a_{\bar{r}}$, they do not have Pauli Z.

Using Eq. (A3) and Table I, the average number can then be calculated as follows:

$$\bar{N}_{\text{spin}} = \frac{5m^2 - m - 8}{m^2 + m}, \tag{A4}$$

TABLE I. Top: Average number of Pauli strings per anti-Hermitian pair of spin-orbital operators. Middle and bottom: For each group, the number of combinations are calculated. For each combination, it has two spin group if it has both triplet and singlet operators, only one spin group if it only has singlet. For each singlet and triplet operator, the number of fermionic terms are calculated. And then we calculate the number of Paulis in each term. Finally, we count the length of Pauli Z.

Group	1	2	3	4	5
No. of Fermi ops (T)	6	6	NA	NA	NA
No. of Fermi ops (S)	4	4	2	2	1
Average no. of Paulis (T)	8	$(4 \times 2 + 2 \times 8)/6 = 4$	NA	NA	NA
Average no. of Paulis (S)	8	$(2 \times 2 + 2 \times 8)/4 = 5$	8	2	8

Group	1	2	3
No. of combinations	$\binom{m}{2}\binom{m-2}{2}/2$	$m\binom{m-1}{2}$	$m\binom{m-1}{2}$
No. of spin groups	$\binom{m}{2}\binom{m-2}{2}/2 \times 2$	$m\binom{m-1}{2} \times 2$	$m\binom{m-1}{2}$
No. of Fermi ops	$\binom{m}{2}\binom{m-2}{2}/2 \times (6+4)$	$m\binom{m-1}{2} \times (6+4)$	$m\binom{m-1}{2} \times 2$
No. of Pauli strings	$\binom{m}{2}\binom{m-2}{2}/2 \times (6+4) \times 8$	$m\binom{m-1}{2} \times (6 \times 4 + 4 \times 5)$	$m\binom{m-1}{2} \times 2 \times 8$
No. of Pauli Z	$\frac{m}{2}(2m^4 - 15m^3 + 40m^2 - 45m + 18)$	$\frac{10m}{3}(m^3 - 4m^2 + 5m - 2)$	$\frac{2m}{3}(m^3 - 4m^2 + 5m - 2)$

Group	4	5	Total
No. of combinations	$2\binom{m}{2}$	$\binom{m}{2}$	$m(m^3 + 2m^2 - m - 2)/8$
No. of spin groups	$2\binom{m}{2}$	$\binom{m}{2}$	$m^2(m^2 - 1)/4$
No. of Fermi ops	$2\binom{m}{2} \times 2$	$\binom{m}{2}$	$m(5m^3 - 6m^2 - 7m + 8)/4$
No. of Pauli strings	$2\binom{m}{2} \times 2 \times 2$	$\binom{m}{2} \times 8$	$2m(5m^3 - 15m^2 + 14m - 4)$
No. of Pauli Z	$\frac{2m}{3}(2m^2 - 3m + 1)$	0	$m(6m^4 - 21m^3 + 32m^2 - 27m + 10)/6$

$$\bar{N}_{\text{Pauli}} = \frac{8(5m^3 - 15m^2 + 14m - 4)}{5m^3 - 6m^2 - 7m + 8}, \tag{A5}$$

$$\bar{N}_Z = \frac{12m^3 - 30m^2 + 34m - 20}{3(8 + m - 5m^2)}. \tag{A6}$$

We can estimate the CNOTs required for the H4 molecule ($m = 4$) is

$$\begin{aligned} N_{\text{CNOT}} &\approx N_{\text{pars}} \times N_{\text{Pauli}} \times (6 + 2 \times N_Z) \times N_{\text{spin}} \\ &= 11 \times N_{\text{Pauli}} \times (6 + 2 \times N_Z) \times N_{\text{spin}} \\ &= 1928.41, \end{aligned}$$

which is very close to the number we obtain.

	H4	LiH	H6
Estimated no. of CNOTS per parameters	175	303	303
No. of parameters	11	30	91
Total estimated no. of CNOTS	1928	9078	27536
No. of CNOTS	2208	6824	28632

This estimation is accurate when all the spin orbitals are evenly explored (H6 and H4), i.e., no "core" orbitals.

APPENDIX B: CONSTRUCTIVE PROOF OF MINIMAL COMPLETE POOLS

In this appendix, we present a constructive proof that there exist complete pools for qubit-ADAPT containing only $2n - 2$ operators, which we refer to as minimal complete pools. We first prove that the set of operators $\{V_i\}_{n=3} = \{iZ_3Z_2Y_1, iZ_3Y_2, iY_3, iY_2\}$ forms a complete pool for $n = 3$ qubits, i.e., using operators of the form $\prod_k e^{\theta_k V_k}$, we can rotate any real state into any other real state. We then prove that the recursively defined pool, $\{V_i\}_n = \{Z_n\{V_i\}_{n-1}, iY_n, iY_{n-1}\}$, is a complete pool for n qubits, such that $\prod_k e^{\theta_k V_k} |\psi\rangle = |\varphi\rangle$ for any two real states $|\psi\rangle$ and $|\varphi\rangle$ in the Hilbert space.

1. Complete pool for three qubits

Here, we prove that $\{V_i\}_{n=3}=\{iZ_3Z_2Y_1, iZ_3Y_2, iY_3, iY_2\}$ is a minimal complete pool for three qubits. We do this by showing that we can map any state $|\psi\rangle$ to $|000\rangle$ using only operators of the form $\prod_k e^{\theta_k V_k}$.

We begin by decomposing an arbitrary real state $|\psi\rangle$ as follows:

$$|\psi\rangle = |0\rangle_3 |\psi_0\rangle_{21} + |1\rangle_3 |\psi_1\rangle_{21}. \tag{B1}$$

The subscripts outside the kets indicate which qubit(s) the ket belongs to. In general, the two-qubit states $|\psi_0\rangle$ and $|\psi_1\rangle$ do not have the same norm: $\langle\psi_0|\psi_0\rangle \neq \langle\psi_1|\psi_1\rangle$. However, we can always perform a rotation $e^{i\theta Y_3}$ on qubit 3 (the leftmost qubit) to transform $|\psi\rangle$ to the form

$$|\psi'\rangle = e^{i\theta Y_3} |\psi\rangle = |0\rangle_3 |\psi'_0\rangle_{21} + |1\rangle_3 |\psi'_1\rangle_{21}, \tag{B2}$$

where

$$|\psi'_0\rangle = \cos\theta |\psi_0\rangle + \sin\theta |\psi_1\rangle, \tag{B3}$$

$$|\psi'_1\rangle = \cos\theta |\psi_1\rangle - \sin\theta |\psi_0\rangle. \tag{B4}$$

The norms of these states are

$$\begin{aligned} \langle\psi'_0|\psi'_0\rangle &= \cos^2\theta \langle\psi_0|\psi_0\rangle + \sin^2\theta \langle\psi_1|\psi_1\rangle \\ &\quad + \sin(2\theta) \langle\psi_0|\psi_1\rangle, \end{aligned} \tag{B5}$$

$$\begin{aligned} \langle\psi'_1|\psi'_1\rangle &= \cos^2\theta \langle\psi_1|\psi_1\rangle + \sin^2\theta \langle\psi_0|\psi_0\rangle \\ &\quad - \sin(2\theta) \langle\psi_0|\psi_1\rangle. \end{aligned} \tag{B6}$$

We can make these two norms equal by choosing θ such that

$$\tan(2\theta) = \frac{\langle\psi_1|\psi_1\rangle - \langle\psi_0|\psi_0\rangle}{2\langle\psi_0|\psi_1\rangle}. \tag{B7}$$

Note that this procedure works for arbitrary states $|\psi_0\rangle$ and $|\psi_1\rangle$, not just for two-qubit states.

The heart of the proof amounts to showing that a state of the form in Eq. (B2), where $\langle\psi'_0|\psi'_0\rangle = \langle\psi'_1|\psi'_1\rangle = 1/2$, can be transformed into a form where qubit 3 has been factored out:

$$|\psi'\rangle \to (|0\rangle + |1\rangle)_3 |\chi\rangle_{21}, \tag{B8}$$

where $|\chi\rangle$ is some two-qubit state. We do this using the (slightly reduced) set of pool operators $\{V_i^{\text{red}}\}_3 = \{iZ_3Z_2Y_1, iZ_3Y_2, iY_2\}$. These are the same as $\{V_i\}_3$, except that iY_3 is not included. In the next subsection, we generalize Eq. (B8) to n qubits and then use it to prove that $\{V_i\}_n$ is a complete pool.

To prove Eq. (B8), let us begin by expanding $|\psi'\rangle$ as follows:

$$|\psi'\rangle = |00\rangle |\psi'_{0a}\rangle + |01\rangle |\psi'_{0b}\rangle + |10\rangle |\psi'_{1a}\rangle + |11\rangle |\psi'_{1b}\rangle. \tag{B9}$$

From this point onward, we suppress the subscripts on the kets for notational simplicity. Similarly to how we made the states $|\psi'_0\rangle$ and $|\psi'_1\rangle$ have the same norm, here we can apply the conditional operations $e^{i\theta_0[(1+Z_3)/2]Y_2}$ and $e^{i\theta_1[(1-Z_3)/2]Y_2}$ to make

$$\langle\psi'_{0a}|\psi'_{0a}\rangle = \langle\psi'_{0b}|\psi'_{0b}\rangle = \langle\psi'_{1a}|\psi'_{1a}\rangle = \langle\psi'_{1b}|\psi'_{1b}\rangle = 1/4. \tag{B10}$$

We assume that this has already been done in Eq. (B9), and that $|\psi'\rangle$ has been redefined accordingly.

Next, we apply the conditional rotation $e^{i\phi_1 Z_3 Z_2 Y_1}$ on qubit 1 to bring the state to

$$e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = (|00\rangle + |01\rangle)\,|\chi_0\rangle + |10\rangle\,|\psi''_{1a}\rangle + |11\rangle\,|\psi''_{1b}\rangle. \tag{B11}$$

Here, we use the fact that $|\psi'_{0a}\rangle$ and $|\psi'_{0b}\rangle$ are real single-qubit states, and thus can be viewed as two-dimensional vectors lying in the same plane. The conditional rotation $e^{i\phi_1 Z_3 Z_2 Y_1}$ rotates these two vectors in opposite directions because one is conditioned on the state $|00\rangle$ and the other on $|01\rangle$, and eventually the two vectors coincide at $|\chi_0\rangle$ for a particular value of the rotation angle ϕ_1. This operation also affects the other terms in Eq. (B11) (as indicated with the extra primes), but their form remains the same. This conditional operation will be used repeatedly in what follows.

The next step is to apply a conditional rotation on qubit 2 to bring it to the state $|0\rangle$ in the first term:

$$e^{i\phi_2[(1+Z_3)/2]Y_2} e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = \sqrt{2}\,|00\rangle\,|\chi_0\rangle + |10\rangle\,|\psi''_{1a}\rangle + |11\rangle\,|\psi''_{1b}\rangle. \tag{B12}$$

We can then apply another conditional operation on qubit 1 to rotate the states $|\psi''_{1a}\rangle$ and $|\psi''_{1b}\rangle$ into each other:

$$e^{i\phi_3 Z_3 Z_2 Y_1} e^{i\phi_2[(1+Z_3)/2]Y_2} e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = \sqrt{2}\,|00\rangle\,|\chi'_0\rangle + (|10\rangle + |11\rangle)\,|\chi_1\rangle\,. \tag{B13}$$

This operation brings the two states to some state $|\chi_1\rangle$. It also changes $|\chi_0\rangle$ to $|\chi'_0\rangle$. Another conditional rotation on qubit 2 simplifies the second term:

$$e^{i\phi_4[(1-Z_3)/2]Y_2} e^{i\phi_3 Z_3 Z_2 Y_1} e^{i\phi_2[(1+Z_3)/2]Y_2} e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = \sqrt{2}\,|00\rangle\,|\chi'_0\rangle + \sqrt{2}\,|10\rangle\,|\chi_1\rangle\,. \tag{B14}$$

We need one more conditional operation on qubit 1 to bring this to the desired form:

$$e^{i\phi_5 Z_3 Z_2 Y_1} e^{i\phi_4[(1-Z_3)/2]Y_2} e^{i\phi_3 Z_3 Z_2 Y_1} e^{i\phi_2[(1+Z_3)/2]Y_2} e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = \sqrt{2}\,(|00\rangle + |10\rangle)\,|\chi\rangle = (|0\rangle + |1\rangle)\,|0\rangle\,|\chi\rangle\,\sqrt{2}. \tag{B15}$$

We thus prove Eq. (B8). For reasons that will become clear in the next subsection, it is important to note that we did this using only the reduced pool $\{V_i^{\rm red}\}_3$.

All that remains is to show that Eq. (B15) can be mapped to the state $|000\rangle$. This can be done using a Y_3 rotation followed by a conditional operation on qubit 1:

$$e^{i\phi_6 Z_3 Z_2 Y_1} e^{i(\pi/4)Y_3} e^{i\phi_5 Z_3 Z_2 Y_1} e^{i\phi_4[(1-Z_3)/2]Y_2} \times e^{i\phi_3 Z_3 Z_2 Y_1} e^{i\phi_2[(1+Z_3)/2]Y_2} \times e^{i\phi_1 Z_3 Z_2 Y_1}|\psi'\rangle = |000\rangle\,. \tag{B16}$$

Since we can map any state to $|000\rangle$, it follows that we can map any three-qubit state to any other three-qubit state using only the pool operators. We thus show that $\{V_i\}_3 = \{iZ_3Z_2Y_1, iZ_3Y_2, iY_3, iY_2\}$ is a minimal complete pool for three qubits.

2. Proof of the complete pool for *n* qubits

In the previous subsection, we prove that the pool $\{V_i\}_3$ is a complete pool for three qubits. A key piece of this proof is the statement that we can use the reduced pool $\{V_i^{\rm red}\}_3 = \{iZ_3Z_2Y_1, iZ_3Y_2, iY_2\}$ to factor out the third qubit Eq. (B8):

$$|0\rangle\,|\psi'_0\rangle + |1\rangle\,|\psi'_1\rangle \to (|0\rangle + |1\rangle)\,|\chi\rangle\,, \tag{B17}$$

where we assume $\langle\psi'_0|\psi'_0\rangle = \langle\psi'_1|\psi'_1\rangle$, and $|\chi\rangle$ is some two-qubit state. We prove that a similar statement holds for n qubits.

Theorem: We can factor out the nth qubit from a generic n-qubit state:

$$|0\rangle\,|\psi'_0\rangle + |1\rangle\,|\psi'_1\rangle \to (|0\rangle + |1\rangle)\,|\chi\rangle\,, \tag{B18}$$

where $\langle\psi'_0|\psi'_0\rangle = \langle\psi'_1|\psi'_1\rangle$, and $|\chi\rangle$ is some $(n-1)$-qubit state, using only the reduced n-qubit pool:

$$\{V_i^{\rm red}\}_n = \{Z_n\{V_i\}_{n-1}, iY_{n-1}\} = \{Z_n\{V_i^{\rm red}\}_{n-1}, iZ_nY_{n-1}, iY_{n-1}\}. \tag{B19}$$

This pool is defined recursively starting from $\{V_i^{\rm red}\}_3$, and it differs from the full pool by only one operator (iY_n):

$$\{V_i\}_n = \{Z_n\{V_i\}_{n-1}, iY_n, iY_{n-1}\}. \tag{B20}$$

We use induction to prove this theorem. We assume it holds for n qubits and then show that it also holds for $n+1$ qubits. This proof will closely follow what we did for three qubits above. Thus, we start with the $(n+1)$-qubit state:

$$|0\rangle_{n+1}\,|\psi'_0\rangle + |1\rangle_{n+1}\,|\psi'_1\rangle, \tag{B21}$$

where $|\psi'_0\rangle$ and $|\psi'_1\rangle$ are n-qubit states that have the same norm. We want to factor out the leftmost $(n+1)$th qubit using only the operators $\{V_i^{\rm red}\}_{n+1} = \{Z_{n+1}\{V_i^{\rm red}\}_n, iZ_{n+1}Y_n, iY_n\}$. We decompose the state further:

$$|00\rangle\,|\psi'_{0a}\rangle + |01\rangle\,|\psi'_{0b}\rangle + |10\rangle\,|\psi'_{1a}\rangle + |11\rangle\,|\psi'_{1b}\rangle. \tag{B22}$$

As before, we can assume that all four states $|\psi'_{0a}\rangle$, $|\psi'_{0b}\rangle$, $|\psi'_{1a}\rangle$, $|\psi'_{1b}\rangle$, have the same norm of $1/4$, since this can be achieved by applying two conditional operations $e^{i\theta_0[(1+Z_{n+1})/2]Y_n}$ and $e^{i\theta_1[(1-Z_{n+1})/2]Y_n}$.

We next apply a conditional operation built from the operators $e^{\phi_1 Z_{n+1}\{V_i^{\text{red}}\}_n}$ to map the state to

$$(|00\rangle + |01\rangle)\,|\chi_0\rangle + |10\rangle\,|\psi''_{1a}\rangle + |11\rangle\,|\psi''_{1b}\rangle. \tag{B23}$$

We use the induction hypothesis for n qubits to obtain this result. Since the leftmost qubit is in the state $|0\rangle$ in the first two terms, the effect of $e^{\phi_1 Z_{n+1}\{V_i^{\text{red}}\}_n}$ is the same as $e^{\phi_1\{V_i^{\text{red}}\}_n}$ on these terms, and the hypothesis for n qubits then allows us to bring these two terms into a factorized form as above. The last two terms also evolve under this operation, but their general form remains the same since the $n+1$ qubit remains in the state $|1\rangle$ under this operation. We follow this with a conditional operation $e^{i\phi_2[(1+Z_{n+1})/2]Y_n}$ on the nth qubit to rotate it to $|0\rangle$ in the first term:

$$\sqrt{2}\,|00\rangle\,|\chi_0\rangle + |10\rangle\,|\psi''_{1a}\rangle + |11\rangle\,|\psi''_{1b}\rangle. \tag{B24}$$

Another application of $e^{\phi_3 Z_{n+1}\{V_i^{\text{red}}\}_n}$ can be used to similarly factorize the last two terms:

$$\sqrt{2}\,|00\rangle\,\left|\chi'_0\right\rangle + (|10\rangle + |11\rangle)\,|\chi_1\rangle\,. \tag{B25}$$

Note that the form of the first term remains the same since this conditional operation does not rotate the $(n+1)$th or nth qubits. This is because all the operators in $\{V_i^{\text{red}}\}_n$ have either a Z or a I on the nth qubit. Next apply $e^{i\phi_4[(1-Z_{n+1})/2]Y_n}$ to rotate the nth qubit in the second term to $|0\rangle$:

$$\sqrt{2}\,|00\rangle\,\left|\chi'_0\right\rangle + \sqrt{2}\,|10\rangle\,|\chi_1\rangle\,. \tag{B26}$$

Additional applications of $e^{\phi Z_{n+1}\{V_i^{\text{red}}\}_n}$ bring this to

$$\sqrt{2}\,(|0\rangle + |1\rangle)\,|0\rangle\,|\chi\rangle\,, \tag{B27}$$

as desired. In order to make this final step, we need to use a slight extension of the theorem in which we factor out the $(n+1)$th qubit instead of the nth qubit. Do we have all the necessary operators to make this extension work? Yes. To see this, recall that in order to factor out the nth qubit, we need $\{V_i^{\text{red}}\}_n = \{Z_n\{V_i^{\text{red}}\}_{n-1}, iZ_nY_{n-1}, iY_{n-1}\}$. However, here we have

$$Z_{n+1}\{V_i^{\text{red}}\}_n = \{Z_{n+1}Z_n\{V_i^{\text{red}}\}_{n-1}, iZ_{n+1}Z_nY_{n-1}, iZ_{n+1}Y_{n-1}\}. \tag{B28}$$

The operators $Z_{n+1}Z_n\{V_i^{\text{red}}\}_{n-1}$ and $iZ_{n+1}Z_nY_{n-1}$ produce exactly the same conditional rotations on the state in Eq. (B26) as $Z_n\{V_i^{\text{red}}\}_{n-1}$ and iZ_nY_{n-1} produce on $\sqrt{2}\,|0\rangle\,\left|\chi'_0\right\rangle + \sqrt{2}\,|1\rangle\,|\chi_1\rangle$. However, the theorem also requires unconditional Y_{n-1} rotations on $\left|\chi'_0\right\rangle$ and $|\chi_1\rangle$. These can be implemented using the composite operator

$$e^{i(\pi/4)[(1-Z_{n+1})/2]Y_n} e^{i\phi Z_{n+1}Z_nY_{n-1}} e^{i(\pi/4)[(1-Z_{n+1})/2]Y_n}, \tag{B29}$$

which first rotates Eq. (B26) into a state in which the $(n+1)$th and nth qubit states in the two terms have the same parity before performing the $Z_{n+1}Z_nY_{n-1}$ rotation, so that $\left|\chi'_0\right\rangle$ and $|\chi_1\rangle$ undergo the same rotation. The $(n+1)$th and nth qubits are then restored to their original states by the final rotation in Eq. (B29). Given that we already show above that this procedure works for $n=3$ qubits, this completes the proof of the theorem.

With the theorem proven, it is easy to prove that $\{V_i\}_{n+1}$ is a complete pool. We already know that a rotation generated by Y_{n+1} is enough to bring an arbitrary $(n+1)$-qubit state to the form needed to apply the theorem. As we just saw, we can then use $Z_{n+1}\{V_i^{\text{red}}\}_n$, $iZ_{n+1}Y_n$, and iY_n to factor out the $(n+1)$th qubit and arrive at Eq. (B27). If we then apply another Y_{n+1} rotation to bring the $(n+1)$th qubit to $|0\rangle$, we can perform the conditional operation $e^{i\phi_6 Z_{n+1}\{V_i^{\text{red}}\}_n}$ to rotate $\sqrt{2}\,|\chi\rangle$ to $|0\rangle^{\otimes n-1}$. Thus, we can rotate an arbitrary $(n+1)$-qubit state to $|0\rangle^{\otimes n+1}$ using only the generators $\{V_i\}_{n+1} = \{Z_{n+1}\{V_i^{\text{red}}\}_n, iZ_{n+1}Y_n, iY_{n+1}, iY_n\}$. This in turn implies that we can rotate any $(n+1)$-qubit state to any other $(n+1)$-qubit state using the same set of operators. Therefore, $\{V_i\}_{n+1}$ is a complete pool, and it contains $2n-2$ operators.

APPENDIX C: MAPPING BETWEEN MINIMAL COMPLETE POOLS

In this section, we show that the two examples of minimal complete pools discussed in the main text can be mapped into one another. The first example is the pool $\{V_i\}$ that is constructed iteratively in Appendix B. The second example is the pool comprised of operators that act only on adjacent qubits in a linear array, $\{G_i\}$. To facilitate the following analysis, we order the elements in these two pools as follows (where we ignore the factors of i in the definition of the pool operators for simplicity):

$$\begin{aligned}
&V_1 = ZZ\ldots ZY, \quad V_2 = ZZ\ldots ZYI,\\
&\quad V_3 = ZZ\ldots ZYII, \quad \ldots, \quad V_{n-1} = ZYII\ldots I, \quad V_n = YII\ldots I,\\
&V_{n+1} = ZZ\ldots ZIYI, \quad V_{n+2} = ZZ\ldots ZIYII,\\
&\quad \ldots, \quad V_{2n-3} = ZIYII\ldots I, \quad V_{2n-2} = IYII\ldots I.
\end{aligned} \tag{C1}$$

$$
\begin{aligned}
&G_1 = ZYII\ldots I, \quad G_2 = IZYII\ldots I,\\
&\quad G_3 = IIZYII\ldots I, \quad \ldots, \quad G_{n-2} = II\ldots IZYI, \quad G_{n-1} = II\ldots IZY,\\
&G_n = YII\ldots I, \quad G_{n+1} = IYII\ldots I,\\
&\quad G_{n+2} = IIYII\ldots I, \quad \ldots, \quad G_{2n-3} = II\ldots IYII, \quad G_{2n-2} = II\ldots IYI.
\end{aligned}
\tag{C2}
$$

We show that the V_i can be obtained from commutators of the G_i. This means that the two pools share the same operator basis, and so completeness of the one pool implies completeness of the other. Since we already prove above that the V_i are complete in Appendix B, this demonstrates that the G_i also form a minimal complete pool for any number n of qubits.

Recall that two Pauli strings do not commute if they differ by an odd number of Pauli operators, ignoring any qubits with identity operators in at least one of the two strings. In this case, the result of the commutator is proportional to the product of the two Pauli strings. Now notice that we can write the V_i as products of the G_i:

$$
\begin{aligned}
V_{n-k} &\sim G_1 \prod_{j=2}^{k} G_j G_{j+n-1}, \quad k = 2,\ldots,n-1,\\
V_{n-1} &= G_1,\\
V_n &= G_n,\\
V_{2n-k} &\sim G_{k-1} G_k V_{n-k}, \quad k = 3,\ldots,n-1,\\
V_{2n-2} &= G_{n+1}.
\end{aligned}
\tag{C3}
$$

The operators in the product in the first line should be ordered such that the operators with the smallest values of j are on the left and the ones with the largest values are on the right. Notice that no two adjacent operators in these products commute. Therefore, we can rewrite them as nested commutators of the G_i. This proves that the V_i and G_i produce the same operator basis when all possible commutators of each pool are computed. We already show in Appendix B that the V_i are complete. Therefore, the G_i also constitute a minimal complete pool. This pool is particularly useful for quantum processors containing a linear array of qubits with nearest-neighbor coupling only. Interestingly, the entangling operators in the G_i pool are similar to the cross-resonance interaction in fixed-frequency superconducting qubits, such as those in the IBM quantum processors [53–55]. This similarity could be exploited to implement qubit-ADAPT with native hardware operations.

[1] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, Phys. Rev. **140**, A1133 (1965).

[2] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, Phys. Rev. **136**, B864 (1964).

[3] A. Kovyrshin and M. Reiher, Self-adaptive tensor network states with multi-site correlators, J. Chem. Phys. **147**, 214111 (2017).

[4] K. H. Marti, B. Bauer, M. Reiher, M. Troyer, and F. Verstraete, Complete-graph tensor network states: A new fermionic wave function ansatz for molecules, New J. Phys. **12**, 103008 (2010).

[5] V. Murg, F. Verstraete, R. Schneider, P. R. Nagy, and Ö. Legeza, Tree tensor network state with variable tensor order: An efficient multireference method for strongly correlated systems, J. Chem. Theory Comput. **11**, 1027 (2015).

[6] S. Szalay, M. Pfeffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider, and r. Legeza, Tensor product methods and entanglement optimization for *ab initio* quantum chemistry, Int. J. Quantum Chem. **115**, 1342 (2015).

[7] K. Choo, A. Mezzacapo, and G. Carleo, arXiv:1909.12852 [cond-mat, physics:physics, physics:quant-ph] (2019).

[8] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, Phys. Rev. Lett. **69**, 2863 (1992).

[9] U. Schollwöck, The density-matrix renormalization group, Rev. Mod. Phys. **77**, 259 (2005).

[10] G. K.-L. Chan and S. Sharma, The density matrix renormalization group in quantum chemistry, Ann. Rev. Phys. Chem. **62**, 465 (2011).

[11] R. P. Feynman, Simulating physics with computers, Int. J. Theor. Phys. **21**, 467 (1982).

[12] S. McArdle, S. Endo, A. Aspuru-Guzik, S. Benjamin, and X. Yuan, arXiv:1808.10402 (2018).

[13] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, quantum chemistry in the age of quantum computing, Chem. Rev. **119**, 10856 (2019).

[14] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, arXiv:2001.03685 (2020).

[15] A. Y. Kitaev, arXiv:quant-ph/9511026 (1995).

[16] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Simulation of electronic structure hamiltonians using quantum computers, Mol. Phys. **109**, 735 (2011).

[17] D. S. Abrams and S. Lloyd, Simulation of Many-Body Fermi Systems on a Universal Quantum Computer, Phys. Rev. Lett. **79**, 2586 (1997).

[18] D. S. Abrams and S. Lloyd, Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors, Phys. Rev. Lett. **83**, 5162 (1999).

[19] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[20] 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).

[21] 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).
[22] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of Molecular Spectra on a Quantum Processor with an Error-Resilient Algorithm, Phys. Rev. X **8**, 011021 (2018).
[23] 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 EP (2017).
[24] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure, Phys. Rev. A **95**, 020501 (2017).
[25] 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).
[26] 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**, 10 (2020).
[27] R. J. Bartlett and M. Musiał, Coupled-cluster theory in quantum chemistry, Rev. Mod. Phys. **79**, 291 (2007).
[28] M.-H. Yung, J. Casanova, A. Mezzacapo, J. McClean, L. Lamata, A. Aspuru-Guzik, and E. Solano, From transistor to trapped-ion computers for quantum chemistry, Sci. Rep. **4**, 3589 EP (2014).
[29] 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).
[30] 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).
[31] I. G. Ryabinkin, R. A. Lang, S. N. Genin, and A. F. Izmaylov, arXiv:1906.11192 (2019).
[32] 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 (2019).
[33] W. J. Huggins, J. Lee, U. Baek, B. O'Gorman, and K. B. Whaley, arXiv:1909.09114 (2019).
[34] P.-L. Dallaire-Demers, J. Romero, L. Veis, S. Sim, and A. Aspuru-Guzik, arXiv:1801.01053 (2018).
[35] J. Romero, R. Babbush, J. R. McClean, C. Hempel, P. J. Love, and A. Aspuru-Guzik, Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz, Quantum Sci. Technol. **4**, 014008 (2018).
[36] Y. Matsuzawa and Y. Kurashige, arXiv:1909.12410 (2019).
[37] F. A. Evangelista, G. Kin-Lic Chan, and G. E. Scuseria, arXiv:1910.10130 (2019).
[38] 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 (2020).
[39] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).
[40] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions, Phys. Rev. A **98**, 022322 (2018).
[41] M. Ganzhorn, D. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner, I. Tavernelli, and S. Filipp, Gate-Efficient Simulation of Molecular Eigenstates on a Quantum Computer, Phys. Rev. Appl. **11**, 044092 (2019).
[42] Because the gradient depends on the commutator with $\hat{H}$ such that measurements associated with terms in $\hat{H}$, which commute with a specific operator in the pool, can be skipped, the actual number of measurements needed is difficult to predict, as it will depend on the operator pool chosen. Getting estimates for this number for various pools will be the focus of future work.
[43] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, arXiv e-prints (2019).
[44] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, Ann. Phys. (N. Y.) **298**, 210 (2002).
[45] V. Havlíček, M. Troyer, and J. D. Whitfield, Operator locality in the quantum simulation of fermionic models, Phys. Rev. A **95**, 032332 (2017).
[46] K. Setia and J. D. Whitfield, Bravyi-kitaev superfast simulation of electronic structure on a quantum computer, J. Chem. Phys. **148**, 164104 (2018).
[47] H. Abraham, *et al.*, Qiskit: An open-source framework for quantum computing (2019), https://doi.org/10.5281/zenodo.2562110.
[48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, NY, USA, 2011), 2nd ed.
[49] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, arXiv:2001.00550 [quant-ph] (2020).
[50] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, Quantum Mach. Intell. **3**, 5 (2021).
[51] M. Bilkis, M. Cerezo, G. Verdon, P. J. Coles, and L. Cincio, arXiv:2103.06712 [quant-ph] (2021).
[52] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, Quantum **3**, 214 (2019).
[53] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, Phys. Rev. B **81**, 134507 (2010).
[54] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits, Phys. Rev. Lett. **107**, 080502 (2011).
[55] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, Phys. Rev. A **93**, 060302 (2016).
[56] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PRXQuantum.2.020310 for the qiskit code for circuit transpilation.

Chapter 13

Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer

This chapter includes pages from our publication [11]

L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, "Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer", Phys. Rev. Research **4**, 033029 (2022)

Copyright (2022) of the American Physical Society.

The contributions to this work are as follows: My contributions were assisting in theoretical developments, manuscript preparation, and assisting in numerical simulations. L.Z. primarily developed this work. H.L.T. and F.A.C.V. assisted with numerical simulations, manuscript preparation, and theoretical work. N.J.M., E.B., and S.E.E. supervised.

Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer

Linghua Zhu [1,*] Ho Lun Tang,[1] George S. Barron [1] F. A. Calderon-Vargas,[1]
Nicholas J. Mayhall,[2] Edwin Barnes,[1] and Sophia E. Economou[1,†]
[1]*Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*
[2]*Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

(Received 25 May 2020; revised 22 December 2020; accepted 17 June 2022; published 11 July 2022)

The quantum approximate optimization algorithm (QAOA) is a hybrid variational quantum-classical algorithm that solves combinatorial optimization problems. While there is evidence suggesting that the fixed form of the standard QAOA *Ansatz* is not optimal, there is no systematic approach for finding better *Ansätze*. We address this problem by developing an iterative version of QAOA that is problem tailored, and which can also be adapted to specific hardware constraints. We simulate the algorithm on a class of Max-Cut graph problems and show that it converges much faster than the standard QAOA, while simultaneously reducing the required number of CNOT gates and optimization parameters. We provide evidence that this speedup is connected to the concept of shortcuts to adiabaticity.

DOI: 10.1103/PhysRevResearch.4.033029

I. INTRODUCTION

Many important computationally hard combinatorial optimization problems such as Max-Cut, graph coloring, traveling salesman, and scheduling management [1–3] can be mapped to Ising Hamiltonians whose ground states provide the solutions. One can in principle solve these optimization problems on a quantum computer by initializing the quantum device in the ground state of a Hamiltonian that is easy to prepare and adiabatically tuning the latter into the problem Hamiltonian. In a digital quantum computer, this translates into a Trotterized version of the adiabatic evolution operator, which is the alternating product of the evolution operators corresponding to the initial mixer and the problem (Ising) Hamiltonians. In the limit of an infinite product, this Trotterized form becomes exact.

QAOA is a hybrid quantum-classical variational algorithm that uses a finite-order version of this evolution operator to prepare wave-function *Ansätze* on a quantum processor [4–7]. QAOA is performed by variationally minimizing the expectation value of the Ising Hamiltonian with respect to the parameters in the *Ansatz*. The quantum processor is also used to measure energy expectation values, while the optimization is done on a classical computer. There has been a lot of progress on QAOA recently on both the experimental and theoretical fronts [1,8–14]. There is evidence suggesting that QAOA may provide a significant quantum advantage over classical algorithms [14–16], and that it is computationally universal [17,18].

Despite these advances, there are limitations and potential issues with this algorithm. The performance improves with the number of alternating layers in the *Ansatz*, but the latter is limited by coherence times in existing and near-term quantum processors. Moreover, more layers implies more variational parameters, which introduces challenges for the classical optimizer [19]. Furthermore, Ref. [20] points out that the locality and symmetry of QAOA can also severely limit its performance. These issues can be attributed to, or are at least exacerbated by, the form of the QAOA *Ansatz*. In particular, short-depth QAOA is not really the digitized version of the adiabatic problem, but rather an ad hoc *Ansatz*, and as a result should not be expected to perform optimally, or even well. A short-depth *Ansatz* that is further tailored to a given combinatorial problem could therefore address the issues with the standard QAOA *Ansatz*. However, identifying such an alternative is a highly nontrivial problem given the vast space of possible *Ansätze*.

In this work, we propose an iterative version of QAOA termed adaptive derivative assembled problem tailored-quantum approximate optimization algorithm (ADAPT-QAOA). Our algorithm grows the *Ansatz* two operators at a time by using a gradient criterion to systematically select the QAOA mixer from a predefined operator pool. While ADAPT-QAOA is general and can be applied to any optimization problem, we focus on Max-Cut to quantify its performance. We find that when entangling gates are included in the operator pool, there is a dramatically faster convergence compared to standard QAOA. Surprisingly, despite the introduction of entangling gates as mixers, these improvements come with a reduction in the numbers of both the optimization parameters and the CNOT gates by approximately 50% each compared to

*zlinghua18@vt.edu
†economou@vt.edu

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

standard QAOA. We summarize the details of our algorithm and its performance in Sec. II and provide evidence that this improved performance is related to the concept of shortcuts to adiabaticity [21–24] in Sec. III. Finally, we conclude in Sec. IV with a look ahead.

II. ADAPT-QAOA

A. Framework

In QAOA [4,5], the variational *Ansatz* consists of p layers, each containing the cost Hamiltonian H_C and a mixer, H_M:

$$|\psi_p(\vec{\gamma},\vec{\beta})\rangle = \left(\prod_{k=1}^{p}[e^{-iH_M\beta_k}e^{-iH_C\gamma_k}]\right)|\psi_{\text{ref}}\rangle\,, \tag{1}$$

where $|\psi_{\text{ref}}\rangle = |+\rangle^{\otimes n}$, n is the number of qubits, and $\vec{\gamma}$ and $\vec{\beta}$ are sets of variational parameters. If these parameters are chosen such that $\langle\psi_p(\vec{\gamma},\vec{\beta})|\,H_C\,|\psi_p(\vec{\gamma},\vec{\beta})\rangle$ is minimized, then the resulting energy and state provide an approximate solution to the optimization problem encoded in H_C. The accuracy of the result and the efficiency with which it can be obtained depend sensitively on H_M. In the standard QAOA *Ansatz*, the mixer is chosen to be a single-qubit X rotation applied to all qubits. A few papers have suggested modifications to the standard QAOA *Ansatz* for specific problems and hardware architectures [25–27]. These interesting results reveal the potential advantages of the QAOA *Ansatz* but do not provide a universal strategy for choosing mixers that works across a broad range of optimization problems.

In this work, we replace the single, fixed mixer H_M by a set of mixers A_k that change from one layer to the next:

$$|\psi_p(\vec{\gamma},\vec{\beta})\rangle = \left(\prod_{k=1}^{p}[e^{-iA_k\beta_k}e^{-iH_C\gamma_k}]\right)|\psi_{\text{ref}}\rangle\,. \tag{2}$$

We build up this *Ansatz* iteratively, one layer at a time, in a way that is determined by H_C. This iterative process is inspired by the variational quantum eigensolver algorithm, ADAPT-VQE [28,29]. It can be summarized by three basic steps: First, define the operator set $\{A_j\}$ (called the mixer pool, and where $A_j = A_j^\dagger$) and select a suitable reference state to be the initial state: $|\psi^{(0)}\rangle = |\psi_{\text{ref}}\rangle$. Here, we choose $|\psi_{\text{ref}}\rangle = |+\rangle^{\otimes n}$ as in the standard QAOA. We will return shortly to the question of how to choose the pool. Second, prepare the current *Ansatz* $|\psi^{(k-1)}\rangle$ on the quantum processor and measure the energy gradient with respect to the pool, the jth component of which is given by $-i\,\langle\psi^{(k-1)}|\,e^{iH_C\gamma_k}[H_C,A_j]e^{-iH_C\gamma_k}\,|\psi^{(k-1)}\rangle$, where the new variational parameter γ_k is set to a predefined value γ_0. For the measurement, we can decompose the commutator into linear combinations of Pauli strings and measure the expectation values of the strings using general variational quantum algorithm methods [30]. If the norm of the gradient is below a predefined threshold, then the algorithm stops, and the current state and energy estimate approximate the desired solution. If the gradient threshold is not met, modify the *Ansatz* by adding the operator, $A_{\text{max}}^{(k)}$, associated with the largest component of the gradient: $|\psi^{(k)}\rangle = e^{-iA_{\text{max}}^{(k)}\beta_k}e^{-iH_C\gamma_k}\,|\psi^{(k-1)}\rangle$, where β_k is a new variational parameter. Third, optimize all parameters currently in the *Ansatz*, β_m, γ_m, $m = 1,\ldots,k$, such that $\langle\psi^{(k)}|H_C|\psi^{(k)}\rangle$ is minimized, and return to the second step. This algorithm, which we call ADAPT-QAOA, lies somewhere between standard QAOA and ADAPT-VQE in the sense that it possesses the alternating-operator structure of QAOA but enjoys additional flexibility by allowing the mixers to vary over the course of the iterative construction.

B. Operator pool

The first step in running this algorithm is to define the mixer pool. Define Q to be the set of qubits. The pool corresponding to the standard QAOA contains only one operator, $P_{\text{QAOA}} = \{\sum_{i\in Q} X_i\}$. Here, we introduce two qualitatively different pools: one consisting entirely of single-qubit mixers, and one with both single-qubit and multiqubit entangling gates: $P_{\text{single}} = \cup_{i\in Q}\{X_i, Y_i\} \cup \{\sum_{i\in Q} Y_i\} \cup P_{\text{QAOA}}$, $P_{\text{multi}} = \cup_{i,j\in Q\times Q}\{B_iC_j|B_i, C_j \in \{X,Y,Z\}\} \cup P_{\text{single}}$. Because $P_{\text{QAOA}} \subset P_{\text{single}} \subset P_{\text{multi}}$, we expect that P_{multi} will provide the best performance. The QAOA, single-qubit, and multiqubit pools have $O(1)$, $O(n)$, and $O(n^2)$ elements, respectively.

Max-Cut is a classic (NP-hard) quadratic unconstrained binary optimization problem, and it can be used to solve other optimization problems. Thus, it is a useful benchmarking problem for QAOA and has been used as such in prior works [5,7,11]. It is defined as follows: Given a graph $G = (V, E)$, with weight $w_{i,j}$ for edge (i, j), find a cut $S \subseteq V$ such that $S \cup \bar{S} = V$, and $\sum_{i\in S, j\in\bar{S}, i,j\in E} w_{i,j}$ is maximized. This problem can be encoded in the Ising Hamiltonian

$$H_C = -\frac{1}{2}\sum_{i,j} w_{i,j}(I - Z_iZ_j), \tag{3}$$

where the couplings are given by the edge weights. Each classical state (i.e., tensor product of Z eigenstates) represents a possible cut. H_C counts the sum of the weights of the edges connecting one subgraph to the other, and its ground state corresponds to the maximum cut. H_C has a Z_2 symmetry generated by $F = \otimes_i X_i$. Only the A_j that commute with F have a nonzero gradient (see Appendix A), so we retain only these Pauli strings (which have an even number of Y or Z operators) in our mixer pool.

C. Performance and resource comparison

We use the Max-Cut problem on regular graphs with $n = 6$ vertices and degrees $D = 3$ and $D = 5$ to benchmark the performance of ADAPT-QAOA. For each type of graph, we analyze 20 instances of random edge weights, which are drawn from the uniform distribution $U(0, 1)$ [31]. We use Nelder-Mead for the optimization of the variational parameters $\vec{\beta}$ and $\vec{\gamma}$. The gradients used to select new operators are sensitive to the initial values for $\vec{\gamma}$. It is natural to initialize these parameters at $\gamma_0 = 0$ to avoid biasing the optimization. However, as we show in Appendix B, $\gamma_0 = 0$ is a critical point of the cost function [32]. Moreover, the minimum of the energy in the first layer of ADAPT-QAOA never occurs at $\gamma_0 = 0$. Therefore, we shift the initial value γ_0 slightly away from zero ($\gamma_0 = 0.01$) to avoid these issues.

In Fig. 1 we show the error as a function of the number of *Ansatz* layers for the standard QAOA and for ADAPT-QAOA using single-qubit and multiqubit mixer pools. For

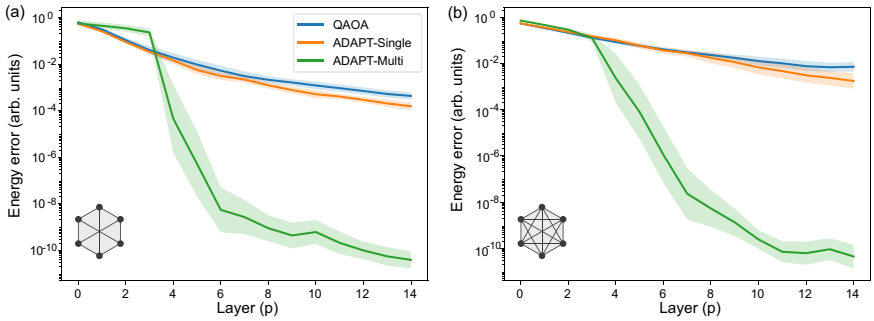

FIG. 1. Comparison of the performance of standard QAOA (blue) with ADAPT-QAOA for the single-qubit (orange) and multiqubit (green) pools. The algorithms are run on the Max-Cut problem for the regular graphs shown in the figure, which have $n = 6$ vertices and are of degree (a) $D = 3$ and (b) $D = 5$. The energy error (the difference between the energy estimate obtained by the algorithm and the exact ground-state energy of H_C) is shown as a function of the number of layers in the *Ansatz*. Results are shown for 20 different instances of edge weights, which are randomly sampled from the uniform distribution $U(0, 1)$. The shaded regions indicate 95% confidence intervals.

both three- and five-regular graphs, we find that using the single-qubit mixer pool provides a modest improvement over standard QAOA. On the other hand, the multiqubit pool performs dramatically better, leading to a rapid convergence to the exact solution after only ∼3 layers. We also find that for the degree-five graphs, standard QAOA and ADAPT-QAOA with single-qubit mixers converge slower than the degree-three case, whereas the performance of ADAPT-QAOA with the multiqubit operator pool remains approximately the same. Note that the particular form of the two-qubit operators in the pool was chosen for its simplicity. In general, one can choose a hardware-tailored operator pool, in the spirit of Ref. [33]. In Appendix C, we show similar results for $n = 8$ and $n = 10$ graphs of degree $D = 2$, where ADAPT-QAOA with the multiqubit pool substantially outperforms the standard QAOA again. Going to larger values of D or n is made challenging by a sharp increase in the number of layers needed to reach convergence, as reported for standard QAOA in Ref. [34].

It is interesting to ask how much the ADAPT-QAOA *Ansätze* differ from the standard QAOA *Ansatz*. We find that when the single-qubit mixer pool is used, the single-qubit operators X_i are chosen instead of the standard mixer approximately 36.6% of the time for $n = 6, D = 3$ graphs and 25% of the time for $n = 6, D = 5$ graphs. For the multiqubit mixer pool, the algorithm chooses operators other than the standard mixer approximately 75% of the time for $n = 6, D = 3$ graphs and 80% of the time for $n = 6, D = 5$ graphs (see Appendix D). This trend supports the intuitive idea that a more connected graph requires more entanglement for a rapid convergence to the solution.

A crucial question, especially for near-term platforms, is how the different mixer pools compare with respect to resource overhead. Figure 2 shows the number of CNOTs and number of parameters for the three algorithms. The CNOT counts are determined by decomposing each two-qubit operator into two CNOT gates and one or two single-qubit gates. Surprisingly, we find that both the standard QAOA and the single-qubit mixer *Ansätze* in fact have more CNOTs compared to that constructed from the entangling multiqubit mixer pool. Moreover, on average, the standard QAOA algorithm uses more parameters and CNOTs to reach the same convergence threshold than either version of ADAPT-QAOA. About half as many CNOTs are required for the ADAPT-QAOA multiqubit pool case, despite the fact that the mixers in the multiqubit pool themselves introduce additional CNOT gates on top of those coming from H_C. Reference [25] proposed using a restricted form of entangling gates in the *Ansatz* to obtain better performance in combinatorial problems at the cost of introducing more variational parameters. In contrast, ADAPT-QAOA provides a systematic way to both improve performance and reduce the number of parameters and CNOTs.

III. SHORTCUTS TO ADIABATICITY

One may wonder whether there is a physically intuitive way to understand the strikingly better performance of ADAPT-QAOA. Considering that the standard QAOA *Ansatz* has a structure dictated by the adiabatic theorem, a possible explanation is that the ADAPT algorithm is related to shortcuts to adiabaticity (STA). STA, also known as counterdiabatic or transitionless driving, was introduced for quantum systems by Demirplak and Rice [21] and later, independently, by Berry [22,23]. STA has also been explored in the classical context [35,36], including a recent application in biology [37]. The idea is that if we want to drive a system such that it remains in the instantaneous ground state at all times, then by adding a certain term $\mathcal{H}_{CD}$ to the Hamiltonian, we can achieve this without paying the price of a slow evolution. Although the instantaneous eigenstates of the original Hamiltonian only solve the time-dependent Schrödinger equation in the adiabatic limit, they become exact solutions when the Hamiltonian is updated to include $\mathcal{H}_{CD}$. The advantage of STA is that the evolution can be achieved nonadiabatically. Below, we provide evidence that ADAPT-QAOA is indeed related to STA, a likely explanation for why it converges to

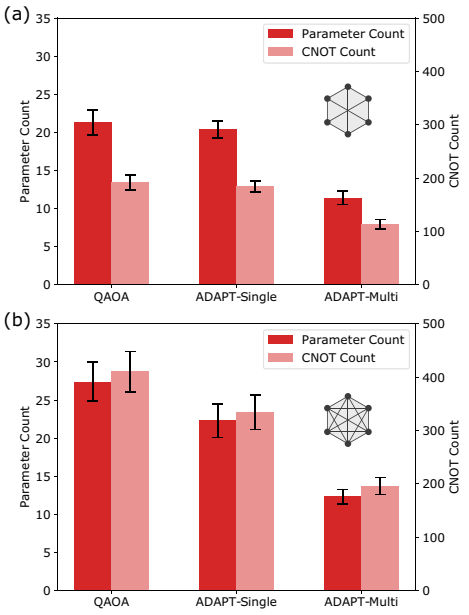

FIG. 2. Resource comparison of the standard QAOA, ADAPT-QAOA with the single-qubit mixer pool, and ADAPT-QAOA with the multiqubit mixer pool for the Max-Cut problem on regular graphs with $n = 6$ vertices and random edge weights. (a) and (b) show the comparison for graphs of degree $D = 3$ and $D = 5$, respectively. For all cases except the standard QAOA applied to $D = 5$ graphs, we count the number of parameters and CNOTs needed to reach an energy error of $\delta E = 10^{-3}$. As standard QAOA for $D = 5$ graphs never reaches this error threshold, we instead count the CNOT gates and parameters at the end of the simulation (15 layers). The dark (light) red bars show variational parameter (CNOT gate) counts. The error bars show variances obtained by sampling over 20 different instances of edge weights.

the solution much faster than its adiabatic counterpart, the standard QAOA. Before we present this evidence, we must first explain how $\mathcal{H}_{CD}$ can be constructed using the concept of adiabatic gauge potentials.

A. Approximate adiabatic gauge potentials

Here we briefly review the mathematical machinery of STA and adiabatic gauge potentials [38–40]. Let $|\psi\rangle$ be a state evolving under $\mathcal{H}(\theta(t))$, $i\partial_t |\psi\rangle = \mathcal{H}(\theta(t))|\psi\rangle$, where θ is a continuous variable that parameterizes the Hamiltonian. A unitary transformation $U(\theta(t))$ can be applied to move the Hamiltonian $\mathcal{H}(\theta(t))$ from the initial basis to its instantaneous eigenbasis, where $\tilde{\mathcal{H}}(\theta) = U^\dagger(\theta)\mathcal{H}(\theta)U(\theta)$ is diagonal at all times. The Schrödinger equation in the instantaneous eigenbasis is $i\partial_t |*\rangle \tilde{\psi} = [\tilde{\mathcal{H}} - \dot{\theta}\tilde{\mathcal{A}}_\theta]|*\rangle \tilde{\psi}$, where $|*\rangle \tilde{\psi} = U^\dagger |\psi\rangle$, $\dot{\theta} = d\theta/dt$, and $\tilde{\mathcal{A}}_\theta = iU^\dagger \partial_\theta U$ is the adiabatic gauge potential in the rotated frame. It is evident that the term $-\dot{\theta}\tilde{\mathcal{A}}_\theta$ drives transitions between the energy levels of the original Hamiltonian $\mathcal{H}$. Therefore, one can add the counterdiabatic term $\mathcal{H}_{CD} = \dot{\theta}\mathcal{A}_\theta$ to $\mathcal{H}(\theta)$, with $\mathcal{A}_\theta = U\tilde{\mathcal{A}}_\theta U^\dagger$, to eliminate such transitions in the rotated frame. This is the core of transitionless driving protocols.

Now, the matrix elements of the adiabatic gauge potential in the instantaneous eigenbasis are

$$\begin{aligned}\langle m(\theta)|\,\mathcal{A}_\theta\,|n(\theta)\rangle &= \langle m(\theta)|\,U\tilde{\mathcal{A}}_\theta U^\dagger\,|n(\theta)\rangle \\ &= i\,\langle m(\theta)|\,\partial_\theta U U^\dagger\,|n(\theta)\rangle \\ &= i\,\langle m(\theta)|\,|\partial_\theta n(\theta)\rangle\,, \end{aligned} \tag{4}$$

where we used $\tilde{\mathcal{A}}_\theta = iU^\dagger\partial_\theta U$ and $|n(\theta)\rangle = U(\theta)|n_0\rangle$ with $|n_0\rangle$ being independent of θ. Moreover, the adiabatic gauge potential $\mathcal{A}_\theta$ satisfies [22,38]

$$\langle m|\,\mathcal{A}_\theta\,|n\rangle = i\,\langle m|\,|\partial_\theta n\rangle = i\frac{\langle m|\,\partial_\theta\mathcal{H}\,|n\rangle}{E_n - E_m}, \tag{5}$$

which is obtained by differentiating the eigenfunction $\mathcal{H}(\theta)|n(\theta)\rangle = E_n(\theta)|n(\theta)\rangle$ with respect to θ. Note that increasing the size of the system can lead to divergent matrix elements due to exponentially small denominators $(E_n - E_m)$. In this regard, Ref. [40] proposes an approximate gauge potential

$$\mathcal{A}_\theta^{(p)} = i\sum_{k=1}^{p} a_k[\mathcal{H}, \partial_\theta\mathcal{H}]_{2k-1}, \tag{6}$$

where $[X, Y]_{k+1} = [X, [X, Y]]_k$ and $\{a_1, a_2, \ldots, a_p\}$ is a set of coefficients with p being the order of the expansion. This set of coefficients is found by minimizing $\mathrm{Tr}[G^2(\mathcal{A}_\theta^{(p)})]$, where $G(\mathcal{A}_\theta^{(p)}) = \partial_\theta\mathcal{H} - i[\mathcal{H}, \mathcal{A}_\theta^{(p)}]$ [40]. In fact, $\mathrm{Tr}[G^2(\mathcal{X})]$ is minimized when $\mathcal{X}$ is equal to the exact adiabatic gauge potential $\mathcal{A}_\theta$ [38,39]. Using matrix calculus identities and properties of the trace, it is straightforward to show that

$$\frac{\partial\mathrm{Tr}[G^2(\mathcal{X})]}{\partial\mathcal{X}} = 2[\mathcal{H}, i\partial_\theta\mathcal{H} - [\mathcal{X}, \mathcal{H}]]. \tag{7}$$

Only adiabatic gauge potentials satisfy $[\mathcal{H}, i\partial_\theta\mathcal{H} - [\mathcal{A}_\theta, \mathcal{H}]] = 0$. This is easily proven by differentiating $\tilde{\mathcal{H}}(\theta) = U^\dagger(\theta)\mathcal{H}(\theta)U(\theta)$ with respect to θ,

$$\partial_\theta\tilde{\mathcal{H}} = \partial_\theta U^\dagger U U^\dagger\mathcal{H}U + U^\dagger\partial_\theta\mathcal{H}U + U^\dagger\mathcal{H}UU^\dagger\partial_\theta U, \tag{8}$$

and noting that $\tilde{\mathcal{A}}_\theta = -i\partial_\theta U^\dagger U = iU^\dagger\partial_\theta U$ and $\mathcal{H} = \sum_n E_n(\theta)|n(\theta)\rangle\langle n(\theta)|$, then

$$[\mathcal{A}_\theta, \mathcal{H}] = i(\partial_\theta\mathcal{H} - \sum_n \partial_\theta E_n(\theta)|n(\theta)\rangle\langle n(\theta)|). \tag{9}$$

Given that $[\mathcal{H}, \sum_n \partial_\theta E_n(\theta)|n(\theta)\rangle\langle n(\theta)|] = 0$, adiabatic gauge potentials clearly satisfy

$$[\mathcal{H}, i\partial_\theta\mathcal{H} - [\mathcal{A}_\theta, \mathcal{H}]] = 0. \tag{10}$$

B. Connection between ADAPT-QAOA and STA

To investigate the connection between ADAPT-QAOA and STA, we apply the above formalism using the Hamiltonian $\mathcal{H} = f(t)H_C + [1 - f(t)]\sum_i^n X_i$ with $f(t) = t/T$ and, since there is no other continuous variable that parameterizes the

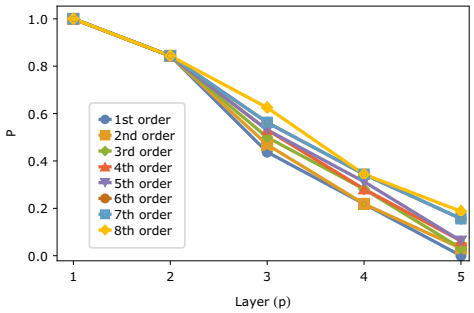

FIG. 3. Probability P of the operator at layer p of the ADAPT-QAOA *Ansatz* to be among the Pauli strings with the largest coefficient in $\mathcal{H}_{CD}$ averaged over 32 graphs with $n = 6$, $D = 3$. The different curves correspond to different orders of the approximation.

Hamiltonian, we simply set $\theta = t$ in the equations above. T is the duration of the evolution from the initial state $|\psi_{ref}\rangle = |+\rangle^{\otimes n}$ to the ground state of the cost Hamiltonian H_C, which is given by Eq. (3). The counterdiabatic Hamiltonian $\mathcal{H}_{CD}$ is approximated using Eq. (6), where p is the order of the approximation.

As a concrete example, we study the Max-Cut problem on 32 instances of regular graphs ($n = 6$, $D = 3$) with random edge weights. Figure 3 shows the probability that an operator in the ADAPT-QAOA *Ansatz* is also one of the dominant operators in $\mathcal{H}_{CD}$. For each of the 32 cases, we define a set $\mathcal{O}_{CD}^{(i)}$ (with $i = 1, \ldots, 32$) comprised of the five operators with the largest coefficient in the time-averaged $\mathcal{H}_{CD}$.[1] The probability P in Fig. 3 is constructed by taking the total number of times the mixer operator at layer p is also an element of the corresponding set $\mathcal{O}_{CD}^{(i)}$ and dividing it by the total number of cases. In all cases, the mixer operator at the first layer is also an element of the set $\mathcal{O}_{CD}^{(i)}$. For higher layers, the probability of the mixer operator to be in $\mathcal{O}_{CD}^{(i)}$ is inversely proportional to the layer number. We attribute this to the fact that $\mathcal{H}_{CD}$ is computed for a specific mixer Hamiltonian ($\sum_i^n X_i$), while information about this choice does not enter into ADAPT-QAOA, which only relies on the initial state $|+\rangle^{\otimes n}$.[2] Interestingly, from Fig. 3 we see that going to higher order in the $\mathcal{H}_{CD}$ approximation increases the probability of finding the mixers in the set $\mathcal{O}_{CD}^{(i)}$. It therefore appears that ADAPT-QAOA finds the appropriate rotation axes in Hilbert space for faster convergence to the solution, and that these axes may in some sense be universal across all possible choices of $\mathcal{H}(t)$ that interpolate between the initial and target states. This suggests that STA can be used as a tool to construct operator pools for ADAPT-QAOA.

[1]We have seen in our simulations that only Pauli string operators are chosen by ADAPT-QAOA until deep into the layer number. Since we stop at layer five, only Pauli strings are chosen in all 32 cases.

[2]In principle, the mixer Hamiltonian in $\mathcal{H}(t)$ could be replaced by any other Hamiltonian that has $|+\rangle^{\otimes n}$ as the ground state, but this would modify the counterdiabatic Hamiltonian and any resemblance to the ADAPT-QAOA *Ansatz* would be reduced or even lost. Interestingly, using the mixer Hamiltonian $\sum_i^n X_i$ involves the lowest possible energetic cost [41,42] of implementing the shortcut to adiabaticity.

IV. CONCLUSION

In conclusion, we introduced ADAPT-QAOA, an optimization algorithm that grows the *Ansatz* iteratively in a way that is naturally tailored to a given problem. We tested several instances of random diagonal Hamiltonians and found that ADAPT-QAOA always outperforms the standard QAOA. Given its flexibility with the choice of mixer pool, the algorithm can be tailored to the native gates, connectivities, and experimental constraints of hardware. It would also be fruitful to employ ADAPT-QAOA for optimization problems that use higher-dimensional Hilbert spaces, such as graph coloring [26,27]. Finally, more work into the connection to STA would be of both fundamental and practical interest.

ACKNOWLEDGMENTS

We thank Bryan T. Gard and Ada Warren for helpful discussions. S.E.E. acknowledges support from the US Department of Energy (Award No. DE-SC0019318). E.B. and N.J.M. acknowledge support from the US Department of Energy (Award No. DE-SC0019199).

APPENDIX A: ISING SYMMETRY AND MIXER POOL OPERATORS

In this work, we focus on Ising Hamiltonians of the form

$$H_C = -\frac{1}{2}\sum_{i,j} w_{i,j}(I - Z_i Z_j), \tag{A1}$$

which have a Z_2 symmetry associated with the operator $F = \otimes_i X_i$. Since $[F, H_C] = 0$ and $F|\psi_{\text{ref}}\rangle = |\psi_{\text{ref}}\rangle$, we can rewrite the gradient in the first iteration as

$$\begin{aligned}&\langle\psi_{\text{ref}}|e^{iH_C\gamma_1}[H_C, A_j]e^{-iH_C\gamma_1}|\psi_{\text{ref}}\rangle \\ &\quad= \langle\psi_{\text{ref}}|e^{iH_C\gamma_1}F[H_C, A_j]Fe^{-iH_C\gamma_1}|\psi_{\text{ref}}\rangle,\end{aligned} \tag{A2}$$

where A_j is an operator from the mixer pool. However, we also know that $FA_j = \pm A_j F$ because F and A_j are Pauli strings (except when A_j is the standard QAOA mixer $\sum_{i\in Q} X_i$ or $\sum_{i\in Q} Y_i$, but the former commutes and the latter anticommutes with F), so $F[H_C, A_j]F = \pm[H_C, A_j]$. Comparing this to Eq. (A2), we see that to have a nonzero gradient, we need $[F, A_j] = 0$. This holds for all steps of the algorithm, because only operators that commute with F appear in the *Ansätze*, so a formula like Eq. (A2) holds at every iteration. The A_j that commute with F are Pauli strings that have an even number of Y or Z operators, so we retain only these Pauli strings in our mixer pool.

APPENDIX B: FIRST LAYER OF ADAPT-QAOA *ANSATZ*

Here we analyze the ADAPT-QAOA cost function in the first layer, and the results show that the minimum of the energy in the first layer of ADAPT-QAOA never occurs at $\gamma = 0$ for

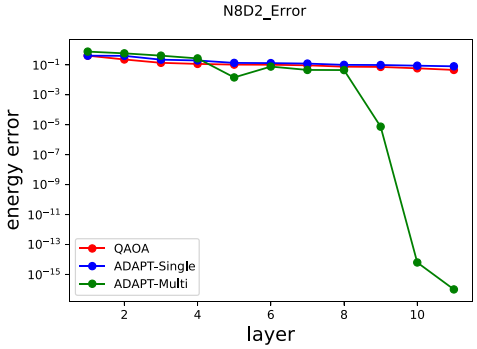

FIG. 4. The performance of three algorithms for the Max-Cut problem on a $n = 8$ and $D = 2$ graph with randomly chosen edge weights.

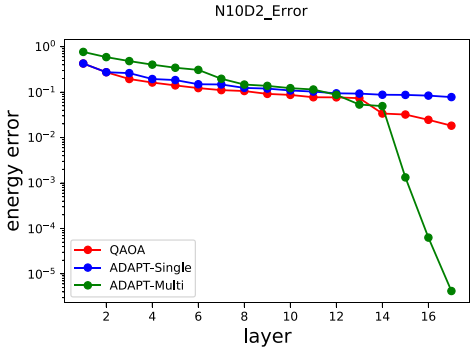

FIG. 5. The performance of three algorithms for the Max-Cut problem on a $n = 10$ and $D = 2$ graph with randomly chosen edge weights.

any operator included in the pool. We also show that $\gamma = 0$ is a critical point of the cost function.

At level p of ADAPT-QAOA, the cost function is

$$E_p(\beta,\gamma) = \langle\psi^{(p-1)}|e^{-i\gamma H}e^{-i\beta M}He^{i\beta M}e^{i\gamma H}|\psi^{(p-1)}\rangle. \quad \text{(B1)}$$

Where H is a linear combination of Pauli strings that are tensor products of the identity and Z. All the terms in H commute with each other. M is the mixer. If the mixers M are single-Pauli strings, we have

$$e^{-i\beta M}He^{i\beta M} = H_c + \cos(2\beta)H_a - i\sin(2\beta)MH_a, \quad \text{(B2)}$$

where H_c is the part of H that commutes with M, and H_a is the part of H that anticommutes with M. We then have

$$\begin{aligned} &e^{-i\gamma H}e^{-i\beta M}He^{i\beta M}e^{i\gamma H} \\ &\quad = H_c + \cos(2\beta)H_a - i\sin(2\beta)e^{-i\gamma H}Me^{i\gamma H}H_a \\ &\quad = H_c + \cos(2\beta)H_a - i\sin(2\beta)MH_ae^{2i\gamma H_a}. \end{aligned} \quad \text{(B3)}$$

We know that $E_p(\beta,\gamma)$ is periodic in β with period π. Therefore, we can restrict β to the range $\beta \in [-\pi/2, \pi/2]$ without loss of generality. Let us define

$$G(\gamma) \equiv -iMH_ae^{2i\gamma H_a}. \quad \text{(B4)}$$

The cost function is then

$$E_p(\beta,\gamma) = \langle H_c\rangle + \cos(2\beta)\langle H_a\rangle + \sin(2\beta)\langle G(\gamma)\rangle, \quad \text{(B5)}$$

where the expectation values are taken with respect to $|\psi^{(p-1)}\rangle$. Therefore,

$$\begin{aligned} &\frac{\partial E_p}{\partial\beta} = -2\sin(2\beta)\langle H_a\rangle + 2\cos(2\beta)\langle G(\gamma)\rangle = 0 \\ &\Rightarrow \tan(2\beta) = \frac{\langle G(\gamma)\rangle}{\langle H_a\rangle}, \end{aligned} \quad \text{(B6)}$$

and

$$\begin{aligned} &\frac{\partial E_p}{\partial\gamma} = \sin(2\beta)\langle G'(\gamma)\rangle = 0 \\ &\Rightarrow \beta = 0, \pm\pi/2 \quad \text{or} \quad \langle G'(\gamma)\rangle = 0. \end{aligned} \quad \text{(B7)}$$

For the first layer, $p = 1$, the state is $|\psi^{(0)}\rangle = |+\rangle^{\otimes n}$, and so $\langle H_a\rangle = 0$. From Eq. (B6), we see that $\beta = \pm\pi/4$ assuming $\langle 0|G(\gamma)|0\rangle \neq 0$. Eq. (B7) then requires $\langle G'(\gamma)\rangle = 0$. Using Eq. (B4), for $p = 1$ this is

$$\begin{aligned} \langle 0|G'(\gamma)|0\rangle &= 2\langle 0|MH_a^2e^{2i\gamma H_a}|0\rangle \\ &= 2\langle 0|H_a^2e^{2i\gamma H_a}\,0\rangle = 0. \end{aligned} \quad \text{(B8)}$$

Notice that $\gamma = 0$ is not a solution of this equation because $\langle 0|H_a^2|0\rangle > 0$, which follows from the fact that this is the norm of a nonzero state, $H_a|+\rangle^{\otimes n}$. Numerics are needed to determine if there is a nonzero value of γ that satisfies $\langle 0|G'(\gamma)|0\rangle = 0$.

On the other hand, if $\langle 0|G(\gamma)|0\rangle = 0$, then there is no constraint on β from $\frac{\partial E_p}{\partial\beta} = 0$. Notice that this happens when $\gamma = 0$ since $\langle 0|G(0)|0\rangle = -i\langle 0|MH_a|0\rangle = 0$, which is true for regular graphs where H_a is a sum of terms like Z_jZ_k, and M is a Pauli string that commutes with $F = \otimes_\ell X_\ell$. In this case, we must have $\beta = 0$ or $\pm\pi/2$. There could be other values of γ that satisfy $\langle 0|G(\gamma)|0\rangle = 0$, but it seems unlikely that both this condition and $\langle 0|G'(\gamma)|0\rangle = 0$ can be satisfied simultaneously, so it is probably still true that β must be 0 or $\pm\pi/2$ in this case.

In summary, there are two classes of possible solutions for $p = 1$: (i) $\beta = \pm\pi/4$ and $\gamma = \gamma^* \neq 0$ where $\langle 0|G'(\gamma^*)|0\rangle = 0$; (ii) $\gamma = \gamma^*$ where $\langle 0|G(\gamma^*)|\rangle = 0$ and either $\langle 0|G'(\gamma^*)|0\rangle = 0$ or $\beta = 0, \pm\pi/2$. To clarify whether these solutions are maxima or minima, we calculated the second derivatives:

$$\frac{\partial^2 E_p}{\partial\beta^2} = -4\cos(2\beta)\langle 0|H_a|0\rangle - 4\sin(2\beta)\langle 0|G(\gamma)|0\rangle, \quad \text{(B9)}$$

$$\frac{\partial^2 E_p}{\partial\gamma^2} = \sin(2\beta)\langle 0|G''(\gamma)|0\rangle, \quad \text{(B10)}$$

$$\frac{\partial^2 E_p}{\partial\beta\partial\gamma} = 2\cos(2\beta)\langle 0|G'(\gamma)|0\rangle. \quad \text{(B11)}$$

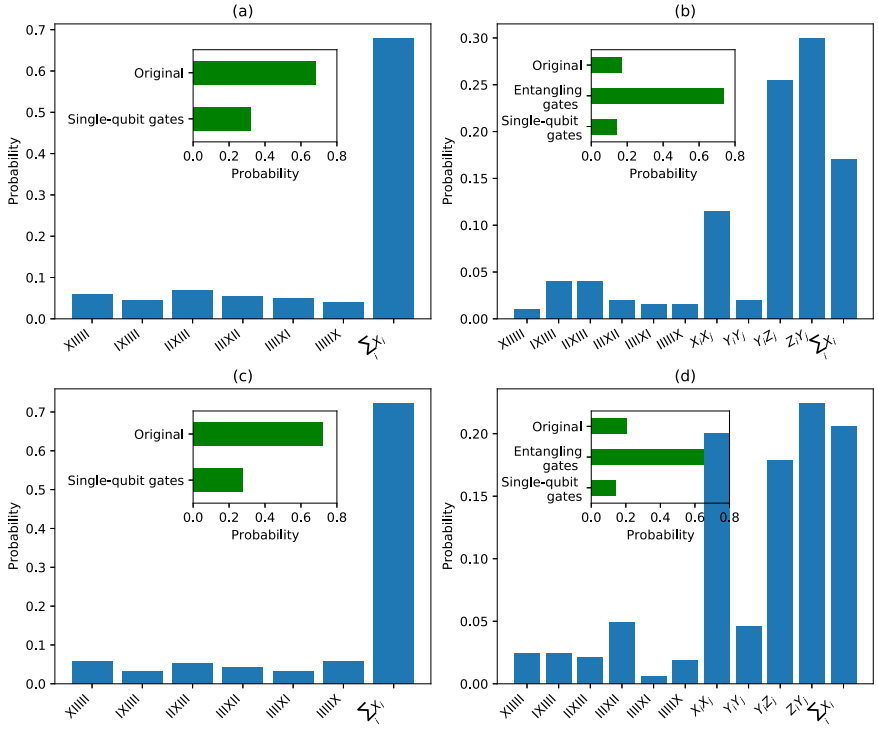

FIG. 6. Probability of operators picked by the original QAOA, ADAPT-QAOA with the single-qubit mixer and ADAPT-QAOA with multiqubit pool for the Max-Cut problem on regular graphs with $n = 6$ vertices with degree (a), (b) $D = 3$ and (c), (d) $D = 5$ with random edge weights sampled from a uniform distribution $U(0, 1)$. The blue bars show the probability of each particular operator used for *Ansatz*, and green bars show the probability of the original mixer, sum over all single-qubit gates and sum over all entangling gates used in *Ansatz*. The results from 20 instances of random edge weights.

Let us consider the class (i) extrema first. In this case

$$\left.\frac{\partial^2 E_p}{\partial\beta^2}\right|_{\beta=\pm\pi/4,\gamma=\gamma^*} = \mp 4\langle 0|G(\gamma^*)|0\rangle, \tag{B12}$$

and the Hessian is

$$\left\{\frac{\partial^2 E_p}{\partial\beta^2}\frac{\partial^2 E_p}{\partial\gamma^2} - \left(\frac{\partial^2 E_p}{\partial\beta\partial\gamma}\right)^2\right\}\Bigg|_{\beta=\pm\pi/4,\gamma=\gamma^*} = -4\langle 0|G(\gamma^*)|0\rangle\langle 0|G''(\gamma^*)|0\rangle. \tag{B13}$$

Thus, we need $\langle 0|G(\gamma^*)|0\rangle\langle 0|G''(\gamma^*)|0\rangle < 0$, in which case $\beta = \pi/4$ is a maximum and $\beta = -\pi/4$ is a minimum if $\langle 0|G(\gamma^*)|0\rangle > 0$, while $\beta = -\pi/4$ is a maximum and $\beta = \pi/4$ is a minimum if $\langle 0|G(\gamma^*)|0\rangle < 0$.

For type (ii) extrema, we have

$$\left.\frac{\partial^2 E_p}{\partial\beta^2}\right|_{\beta=0,\pm\pi/2,\gamma=\gamma^*} = \pm 4\langle 0|H_a|0\rangle = 0, \tag{B14}$$

and the Hessian is

$$\left\{\frac{\partial^2 E_p}{\partial\beta^2}\frac{\partial^2 E_p}{\partial\gamma^2} - \left(\frac{\partial^2 E_p}{\partial\beta\partial\gamma}\right)^2\right\}\Bigg|_{\beta=0,\pm\pi/2,\gamma=\gamma^*} = -4\langle 0|G'(\gamma^*)|0\rangle^2, \tag{B15}$$

these extrema are saddle points. Notice that $\gamma = 0$ is one of these saddle points.

We can see that the minimum of the energy in the first layer of ADAPT-QAOA never occurs at $\gamma = 0$, and the only minima occur at $\beta = \pm\pi/4$ and $\gamma = \gamma^* \neq 0$, which is consistent with our numerical simulation results. The fact that $\gamma = 0$ is a saddle point (at least in the first layer) makes this a problematic choice for the initial value of γ in the classical optimization subroutine of ADAPT-QAOA. For this reason, we instead choose the initial value, γ_0, to be a small nonzero value, e.g., $\gamma_0 = 0.01$. Similar performance is achieved for any value of $\gamma_0 \lesssim 0.1$. Keeping γ_0 close to zero is preferable to avoid biasing the optimization.

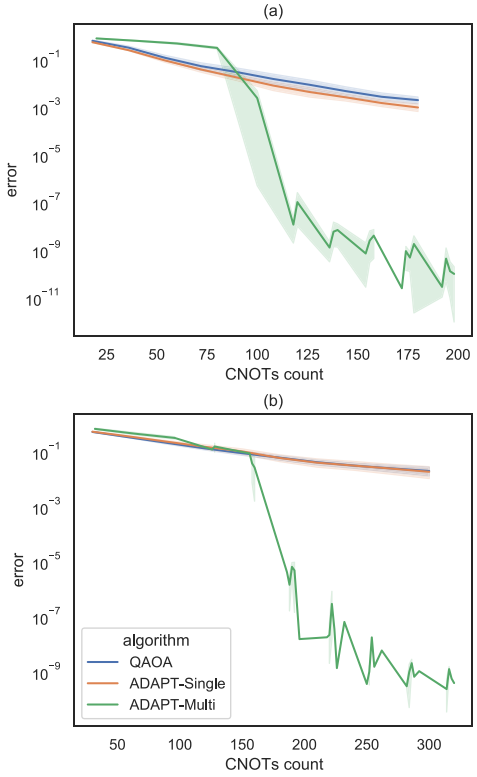

FIG. 7. Comparison of the performance of the original QAOA algorithm (blue) with the ADAPT-QAOA algorithm for the single-qubit (orange) and multiqubit (green) pools run for the Max-Cut problem on graphs with $n = 6$ vertices with degree (a) $D = 3$ and (b) $D = 5$. The energy error is shown as a function of the number of total CNOTs used in the *Ansatz*.

APPENDIX C: SCALABILITY OF ADAPT-QAOA

Here, we investigate the scalability of our algorithm. We perform simulations for degree $D = 2$ graphs with $n = 8$ and $n = 10$ qubits, with results shown in Figs. 4 and 5. Clearly, even for larger system sizes and less connected graphs where $D/n < 1/2$, our algorithm exhibits substantially better performance compared to standard QAOA.

APPENDIX D: ROLE OF ENTANGLING MIXERS VERSUS ENTANGLING GATES

In Fig. 2, we compared the resources used by three different algorithms for the Max-Cut problem on regular graphs. These resources include the number of CNOT gates and the number of optimization parameters. From the comparison we can see that including entangling mixers in the *Ansatz* produces a dramatically faster convergence to the exact solution compared to the original QAOA. Surprisingly, despite the inclusion of these entangling mixers, the improvement in convergence comes with a simultaneous reduction in both the number of entangling gates in the compiled *Ansatz*.

To further investigate the role of entangling mixers, we consider the $n = 6$, $D = 3$ [in Figs. 6(a), 6(b)] and $n = 6$, $D = 5$ graphs [in Figs. 6(c), 6(d)] and show the probability for an operator to be picked by the original QAOA and by ADAPT-QAOA with a single-qubit mixer pool in Fig. S1. Similar results for ADAPT-QAOA with a multiqubit mixer pool are also shown in Figs. 6(b), 6(d). We find that when only the single-qubit mixer pool is used, the single-qubit operators X_i are chosen instead of the original mixer approximately 25% of the time. For the multiqubit mixer pool, the algorithm chooses two-qubit entangling operators approximately 70% of the time. Clearly, entangling mixers play a central role in the improved performance of ADAPT-QAOA.

Additionally, to understand the importance of CNOT gates in the compiled *Ansatz*, in Fig. 7 we show the error for the solution determined by each algorithm as a function of the number of CNOTs used in the *Ansatz* for both $n = 6$, $D = 3$, and $n = 6$, $D = 5$ graphs. Based on the figure, we can see that ADAPT-QAOA with the multiqubit mixer pool converges using a much smaller number of CNOTs compared to the original QAOA or to ADAPT-QAOA with the single-qubit mixer pool. On the other hand, the latter only provides a modest reduction in CNOT count compared to the original QAOA. Interestingly, we see that while entangling mixers are crucial for the dramatic improvement in the performance afforded by ADAPT-QAOA with the multiqubit mixer pool, at the same time far fewer entangling gates are needed in the *Ansatz*. Thus, the role of entanglement in this algorithm is rather subtle.

[1] Y.-H. Oh, H. Mohammadbagherpoor, P. Dreher, A. Singh, X. Yu, and A. J. Rindos, Solving multi-coloring combinatorial optimization problems using hybrid quantum algorithms, arXiv:1911.00595.

[2] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, The unconstrained binary quadratic programming problem: A survey, J. Comb. Opt. **28**, 58 (2014).

[3] A. Lucas, Ising formulations of many NP problems, Front. Phys. **2**, 5 (2014).

[4] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.

[5] E. Farhi, J. Goldstone, and P. Shor, Design of quantum algorithms using physics tools, Tech. Rep. 56291-PH-QC.10 (Massachusetts Institute of Technology, Cambridge, 2014).

[6] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, arXiv:1602.07674.

[7] R. Shaydulin, I. Safro, and J. Larson, Multistart methods for quantum approximate optimization, in *Proceedings of the IEEE*

High Performance Extreme Computing Conference (HPEC), Waltham, MA, 2019 (IEEE, Piscataway, NJ, 2019).

[8] 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, Proc. Natl. Acad. Sci. **117**, 25396 (2020).

[9] 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).

[10] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, Phys. Rev. X **10**, 021067 (2020).

[11] G. E. Crooks, Performance of the quantum approximate optimization algorithm on the maximum cut problem, arXiv:1811.08419.

[12] 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).

[13] Z. Jiang, E. G. Rieffel, and Z. Wang, Near-optimal quantum circuit for Grover's unstructured search using a transverse field, Phys. Rev. A **95**, 062317 (2017).

[14] P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, Improving variational quantum optimization using CVaR, Quantum **4**, 256 (2020).

[15] G. G. Guerreschi and A. Matsuura, QAOA for Max-Cut requires hundreds of qubits for quantum speed-up, Sci. Rep. **9**, 6903 (2019).

[16] M. Y. Niu, S. Lu, and I. L. Chuang, Optimizing QAOA: Success probability and runtime dependence on circuit depth, arXiv:1905.12134.

[17] M. E. S. Morales, J. D. Biamonte, and Z. Zimborás, On the universality of the quantum approximate optimization algorithm, Quantum Info. Proc. **19**, 291 (2020).

[18] S. Lloyd, Quantum approximate optimization is computationally universal, arXiv:1812.11075.

[19] D. Wierichs, C. Gogolin, and M. Kastoryano, Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer, Phys. Rev. Research **2**, 043246 (2020).

[20] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to Variational Quantum Optimization from Symmetry Protection, Phys. Rev. Lett. **125**, 260505 (2020).

[21] M. Demirplak and S. A. Rice, Adiabatic population transfer with control fields, J. Phys. Chem. A **107**, 9937 (2003).

[22] M. V. Berry, Transitionless quantum driving, J. Phys. A: Math. Theor. **42**, 365303 (2009).

[23] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, Rev. Mod. Phys. **91**, 045001 (2019).

[24] N. N. Hegade, K. Paul, Y. Ding, M. Sanz, F. Albarrán-Arriagada, E. Solano, and X. Chen, Shortcuts to Adiabaticity in Digitized Adiabatic Quantum Computing, Phys. Rev. Appl. **15**, 024038 (2021).

[25] E. Farhi, J. Goldstone, S. Gutmann, and H. Neven, Quantum algorithms for fixed qubit architectures, arXiv:1703.06199.

[26] S. Hadfield, Z. Wang, B. O'Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator *Ansatz*, Algorithms **12**, 34 (2019).

[27] Z. Wang, N. C. Rubin, J. M. Dominy, and E. G. Rieffel, *XY* mixers: Analytical and numerical results for the quantum alternating operator *Ansatz*, Phys. Rev. A **101**, 012320 (2020).

[28] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, Nature Commun. **10**, 3007 (2019).

[29] 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, PRX Quantum **2**, 020310 (2021).

[30] W. J. Huggins, J. R. McClean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, npj Quantum Info. **7**, 23 (2021).

[31] L. Li, M. Fan, M. Coram, P. Riley, and S. Leichenauer, Quantum optimization with a novel Gibbs objective function and *Ansatz* architecture search, Phys. Rev. Research **2**, 023074 (2020).

[32] R. Shaydulin, I. Safro, and J. Larson, Multistart methods for quantum approximate optimization, in *2019 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, New York, 2019), pp. 1–8.

[33] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, Nat. Phys. **17**, 332 (2021).

[34] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, Reachability Deficits in Quantum Approximate Optimization, Phys. Rev. Lett. **124**, 090504 (2020).

[35] C. Jarzynski, Generating shortcuts to adiabaticity in quantum and classical dynamics, Phys. Rev. A **88**, 040101(R) (2013).

[36] J. Deng, Q.-h. Wang, Z. Liu, P. Hänggi, and J. Gong, Boosting work characteristics and overall heat-engine performance via shortcuts to adiabaticity: Quantum and classical systems, Phys. Rev. E **88**, 062122 (2013).

[37] S. Iram, E. Dolson, J. Chiel, J. Pelesko, N. Krishnan, Ö. Güngör, B. Kuznets-Speck, S. Deffner, E. Ilker, J. G. Scott, and M. Hinczewski, Controlling the speed and trajectory of evolution with counterdiabatic driving, Nature Phys. **17**, 135 (2020).

[38] M. Kolodrubetz, D. Sels, P. Mehta, and A. Polkovnikov, Geometry and non-adiabatic response in quantum and classical systems, Phys. Rep. **697**, 1 (2017).

[39] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, Proc. Natl. Acad. Sci. USA **114**, E3909 (2017).

[40] P. W. Claeys, M. Pandey, D. Sels, and A. Polkovnikov, Floquet-Engineering Counterdiabatic Protocols in Quantum Many-Body Systems, Phys. Rev. Lett. **123**, 090602 (2019).

[41] A. C. Santos and M. S. Sarandy, Superadiabatic controlled evolutions and universal quantum computation, Sci. Rep. **5**, 15775 (2015).

[42] Y. Zheng, S. Campbell, G. De Chiara, and D. Poletti, Cost of counterdiabatic driving and work output, Phys. Rev. A **94**, 042132 (2016).

Chapter 14

Scaling adaptive quantum simulation algorithms via operator pool tiling

This chapter includes pages from our publication [12]

J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, "Scaling adaptive quantum simulation algorithms via operator pool tiling", arXiv:2206.14215 (2022)

This work has been posted to arXiv.org.

The contributions to this work are as follows: My contributions were assisting in numerical simulations, development of the underlying simulator. J.S.V.D. primarily developed this work. N.J.M., E.B., and S.E.E. supervised.

Scaling adaptive quantum simulation algorithms via operator pool tiling

John S. Van Dyke,[1,*] George S. Barron,[1] Nicholas J. Mayhall,[2] Edwin Barnes,[1] and Sophia E. Economou[1]

[1] *Department of Physics, Virginia Tech, Blacksburg, VA 24061*
[2] *Department of Chemistry, Virginia Tech, Blacksburg, VA 24061*

(Dated: June 30, 2022)

Adaptive variational quantum simulation algorithms use information from the quantum computer to dynamically create optimal trial wavefunctions for a given problem Hamiltonian. A key ingredient in these algorithms is a predefined operator pool from which trial wavefunctions are constructed. Finding suitable pools is critical for the efficiency of the algorithm as the problem size increases. Here, we present a technique called operator pool tiling that facilitates the construction of problem-tailored pools for arbitrarily large problem instances. By first performing an ADAPT-VQE calculation on a smaller instance of the problem using a large, but computationally inefficient operator pool, we extract the most relevant operators and use them to design more efficient pools for larger instances. Given that many problems, such as those arising in condensed matter physics, have a naturally repeating (lattice) structure, we expect the pool tiling method to be widely applicable. We demonstrate the technique here on strongly correlated quantum spin models in one and two dimensions, finding that the resulting state preparation circuits are significantly shorter compared to existing methods.

arXiv:2206.14215v1 [quant-ph] 28 Jun 2022

Variational quantum eigensolvers (VQEs) comprise an important and much-studied class of algorithms for near-term quantum computers [1–4]. As examples of hybrid quantum-classical algorithms, they divide the computational task of optimizing an objective function $f(\boldsymbol{\theta})$ (for instance, the ground state energy of a physical system) between a quantum and a classical processor, in which the former efficiently queries the objective function, while the latter determines updated guesses for the variational parameters $\boldsymbol{\theta}$. In a VQE, the objective function takes the form of an expectation value $f(\boldsymbol{\theta}) = \langle H \rangle_{\boldsymbol{\theta}}$ of an observable H in the state $|\psi(\boldsymbol{\theta})\rangle$. The latter is prepared by a parameterized quantum circuit, and is subsequently measured to determine the relevant expectation value.

While the choice of H is typically fixed by the problem at hand, the success of the algorithm depends strongly on the form of the circuit that prepares $|\psi(\boldsymbol{\theta})\rangle$ (also known as the "ansatz") and the optimization method used. Many different ansatze have been proposed, including hardware-efficient [5] and physics- or chemistry-inspired [1, 6–8] ones. In general, a good ansatz ought to be expressive enough to closely approximate the solution for suitable values of $\boldsymbol{\theta}$, while maintaining a low circuit depth, for compatibility with the limited coherence times of noisy quantum devices.

An important development in this regard was the introduction of the ADAPT-VQE method [9], which put forward the idea of adaptive VQEs. This approach iteratively constructs an ansatz

$$|\psi_{n+1}(\boldsymbol{\theta}^{(n+1)})\rangle = e^{-i\theta_{n+1}A_{n+1}}|\psi_n(\boldsymbol{\theta}^{(n)})\rangle \tag{1}$$

where $\boldsymbol{\theta}^{(n)} = (\theta_1, \ldots, \theta_n)$ are the variational parameters at step n in the algorithm. In adaptive VQEs, the operators A_n are selected from a fixed, predefined "operator pool" according to a selection criterion. In the ADAPT-VQE appraoch, the operator chosen at each step is the one that maximizes the gradient of the objective function, i.e., the A_k for which $\left|\partial\langle H\rangle/\partial\theta_{n+1}|_{\theta_{n+1}=0}\right|$ is maximized, where H is the objective function operator. Intuitively, this criterion seeks to identify the operators A_k which are most effective for decreasing the objective function at a given step.

The ADAPT-VQE algorithm was first applied to finding the ground state energies of small molecules, where it obtained high accuracy with shallower circuit depths compared to other commonly used ansatze. The approach was subsequently generalized to a number of other applications, including optimization problems [10, 11], real and imaginary time evolution [12, 13], the preparation of excited states [14], and strongly-correlated lattice models [15]. The adaptive philosophy towards ansatz construction has also found other expressions, as in evolutionary algorithms and related methods [16, 17].

As one might expect, the success of ADAPT is crucially dependent on the choice of operator pool. The first application of the method used pools consisting of fermionic operators that correspond to single and double excitations of electrons. It was then discovered that qubit-based pools, consisting of individual Pauli strings, can also reach the ground states of molecular systems but with considerably reduced counts of two-qubit gates [18]. Subsequent work [19] demonstrated that even more savings can be achieved by symmetrizing the pool of Ref. [18].

In general, the optimal choice of pool for a given problem is subject to a tradeoff: on the one hand, pools with more operators are more likely to lead to convergence (and more quickly), due to the greater flexibility of operator selection at each step in the algorithm. On the other hand, since $|\partial\langle H\rangle/\partial\theta|$ is calculated for each operator in the pool at every step, this can lead to considerable measurement overhead when estimating the expectation val-

ues on a quantum computer. In seeking to reduce this, it was proven that "minimally complete" pools can be constructed that contain only $2L-2$ operators (where L is the number of qubits) [18]. These pools have the property that all states in the Hilbert space can be obtained from an ansatz using only the operators of the given pool. However, due to the local nature of the ansatz updates, it is not guaranteed that the sequence of operators required can be found by the algorithm. One method for alleviating some of these issues is to explicitly encode information about symmetries of the problem into the choice of operator pool [20].

Despite these advances, the space of possible choices for operator pools remains staggeringly large, growing exponentially in the number of qubits. At present, identifying ultimately successful pools is somewhat of an art, often based on a trial-and-error approach in which various candidates are investigated and compared, slowing the pace of discovery.

In this work, we present a general systematic procedure for constructing pools that are both small in size and also facilitate rapid convergence to the ground state. This is achieved by tiling successful pools from small problem instances to construct pools that are effective for larger problems. We then illustrate the method on the example of the spin-1/2 XXZ model, a prototypical correlated spin model. We find that our approach leads to shorter state preparation ansatze compared to previous methods. Finally, we discuss other variations of the underlying idea and their connection to many-body expansions used in quantum chemistry.

ADAPT instances that use larger pools are expected to converge more often and with fewer iterations, owing to the greater number of states that can be accessed at each step. Among qubit-based pools, the largest possible one is the "full Pauli pool", consisting of all Pauli strings on L qubits, excluding the identity. As this pool includes all other qubit pools as proper subsets, it performs at least as well as each of them, in the sense of making the locally optimal choice. However, the exponential scaling of the size, 4^L-1, clearly prevents its application beyond small problem instances.

Many problems, such as those arising in condensed matter physics, have a simple repeated (lattice) structure that relates instances defined for different numbers of qubits, up to finite-size effects. Exploiting the lattice structure of these systems on near-term quantum devices has attracted increasing attention [21–23]. Example problems include spin models (as used in the field of quantum magnetism), as well as fermionic ones like the Hubbard model, well-known from the study of high-temperature superconductivity. The lattice structure of these models suggests the following approach to designing effective pools: (1) Start with a small instance of the particular model or problem of interest on L_1 qubits and run ADAPT-VQE with a large, highly expressive pool that can quickly converge to the desired solution. (2) Take the operators chosen by ADAPT during the execution of the algorithm as the basis for a new pool defined on a larger instance of the same problem on $L_2 > L_1$ qubits. This requires taking the tensor product of each operator with the single-qubit identity $L_2 - L_1$ times, such that operators in the new pool have the correct dimension. Since the products can be applied either on the left or the right, a single operator of the original pool yields $L_2 - L_1 + 1$ operators for the new one, corresponding to "tiling" the operator across the system in the one-dimensional case (with straightforward generalizations to higher dimensions). Importantly, while the original pool size can be exponential in the number of qubits L_1 (as with the full Pauli pool), the number of operators in the new pool scales linearly with the system size L_2.

We now demonstrate the viability of this strategy through numerical simulations on a concrete example. We consider both one-dimensional and two-dimensional spin-1/2 XXZ models. The first case consists of a chain of L sites with open boundary conditions, whose Hamiltonian is given by

$$H_{XXZ}^{(1D)} = \sum_{i=1}^{L-1} J_{xy}\left(X_iX_{i+1} + Y_iY_{i+1}\right) + J_zZ_iZ_{i+1}, \quad (2)$$

where X_i, Y_i, and Z_i are Pauli operators acting on site i. We set $J_{xy} = 1$ throughout this work. For the first part of the tiling method, we solve the $L = 4$ case of the XXZ chain using the full Pauli pool. ADAPT converges in five steps for this simple problem, choosing the operators: `YXII`, `IIYX`, `ZYXI`, `XZYI`, and `YXZI` in sequential order (here the site indices have been suppressed and the identity is written explicitly). Convergence of the algorithm is defined to occur when the maximum gradient for a given step drops below a threshold ε. We note that in obtaining this result, we used a modified operator selection compared to the original version of ADAPT: If multiple operators have the same maximal gradient at a given step, we select the one that has the lowest Pauli weight (i.e. the number of non-identity Pauli operators in the string). This is useful both for reducing the number of two-qubit gates (which is better suited for noisy devices) and for identifying particular motifs that can be tiled when going to a larger system. Indeed, given the simple structure of the set of operators chosen by ADAPT, we first generalize the starting $L = 4$ pool to include other related Pauli strings, before tiling these to produce pools for larger systems. Thus we start from a "Z-decorated XY pool" consisting of the operators: {`YX`, `XY`, `ZYX`, `ZXY`, `YZX`, `XZY`, `YXZ`, `XYZ`}. We then tile these operators to create pools for systems with $L > 4$ and run ADAPT. Fig. 1(a) shows the relative error in the ground state energy compared to the exact solution for this model, for system sizes $4 \leq L \leq 16$ and several values of the interaction strength J_z. In each case, the error

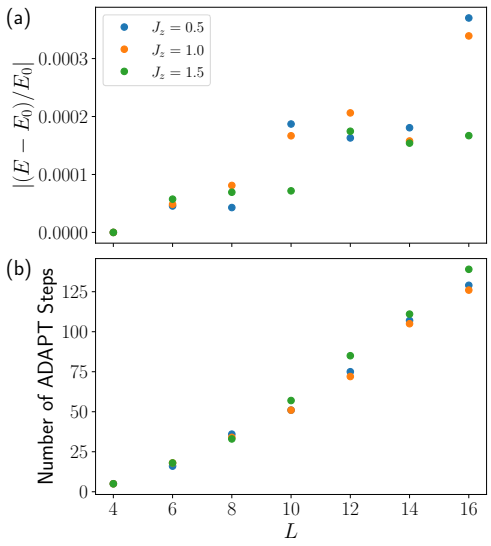

FIG. 1. (a) Relative error in the ground state energy versus system size L for the XXZ chain, using the Z-decorated XY pool. E is the energy found by ADAPT after convergence, while E_0 is the exact ground state energy. (b) Number of ADAPT steps required to converge the algorithm versus system size L, using the Z-decorated XY pool. The initial state is the Néel state $|\uparrow\downarrow\uparrow\downarrow\cdots\rangle$, and we set $\varepsilon = 0.01$. Variational parameters are optimized using the LBFGS algorithm, with $\theta_{n+1} = 0$ and $\{\theta_j\}_{j=1,\ldots,n}$ set to the optimal values from the previous ADAPT iteration.

is less than 0.04%, indicating good convergence across a range of J_z that includes both the paramagnetic and ordered regimes.

In Fig. 1(b) we present the number of ADAPT iterations required to find the ground state as a function of L. We note that this value is also equal to the number of optimization parameters in the final ansatz. The behavior over the simulated range appears to be close to linear, such that a simple extrapolation would suggest that 10^3–10^4 parameters are needed for a system with $L \approx 100$ sites. To compare with another popular VQE ansatz, we consider Ref. [24], which studied the XXZ chain for $J_z = 1$ and closed boundary conditions using the Hamiltonian variational ansatz (HVA). For a system with $L = 12$ sites, 136 variational parameters were required to converge to the ground state. In contrast, the present simulations with ADAPT use only 72 parameters in the final ansatz for the $L = 12$ open chain, illustrating the ability of ADAPT with the tiling pool to obtain highly compact representations of the ground state.

We now turn to the details of how ADAPT-VQE approaches convergence in the present example. The energy and the absolute value of the maximum gradient at each iteration of the algorithm are shown in Fig. 2(a),(b) respectively. Three distinct regimes are apparent in both quantities. The energy decreases sharply during the first $L/2$ iterations, while the maximum gradient remains at its initial value for $L/2 - 1$ iterations. This is followed by a second region in which the energy decreases roughly linearly with iteration number, while the gradient shows a plateau with oscillations between several values. The final regime shows a slower decay both in the energy and the maximum gradient. One can obtain some physical insight into this behavior by inspecting the operators chosen by ADAPT and the corresponding optimal parameter values. The first $L/2$ operators in the ansatz are of the form $e^{-i\theta_k X_{2k-1} Y_{2k}}$ with $k = 1, \ldots, L/2$ and $\theta_k \approx -\pi/4$. Acting on the initial Néel state, this sequence produces a superposition of basis states in which pairs of up and down spins have been swapped. The terms in this superposition all have the same magnitude, but with a nontrivial relative sign structure between them. The other two regimes in the ADAPT convergence process involve the Z-decorated three-local operators, and possess different optimal parameter values which do not form an easily discernible pattern. Nevertheless, it is clear that the operations performed during the middle regime are more effective for reducing the energy overall, while the final regime involves finer tuning of the wave function.

For the two-dimensional XXZ model we study the case of open boundary conditions and nearest-neighbor interactions on rectangular lattices of $N = L_x L_y$ sites:

$$H_{XXZ}^{(2D)} = \sum_{\langle i,j \rangle} J_{xy} \left(X_i X_j + Y_i Y_j \right) + J_z Z_i Z_j, \qquad (3)$$

where $\langle i, j \rangle$ denotes nearest neighbors. We first consider the two-leg ladder with $L_y = 2$ and variable L_x. In generating the ADAPT operator pool, it is natural to consider tiles consisting of 2×2 blocks, given the two-dimensional nature of the problem. Inspired by the success of the Z-decorated XY pool in the one-dimensional case, we consider two types of 2×2 blocks: (1) those having X and Y on two of the four sites, the other two sites being I; (2) those having X, Y, and Z on three of the four sites, the remaining site being I. In each case, we use all permutations of the locations of the nontrivial operators on the block. These blocks are then tiled across the ladder (in overlapping fashion) to create the operator pool. The procedure is illustrated in Fig. 3.

In Fig. 4 we show the convergence properties of the two-leg XXZ ladder ($L_y = 2$) as a function of L_x. In each case, the algorithm achieves a ground state energy close to the exact value [Fig. 4(a)], with a number of ADAPT steps that is roughly linear in the system length. This behavior is consistent with that of the XXZ chain, as expected from the quasi-1D nature of the model.

A more challenging problem to consider is the 2D XXZ model with $L_x = L_y$. In Fig. 5 we present the en-

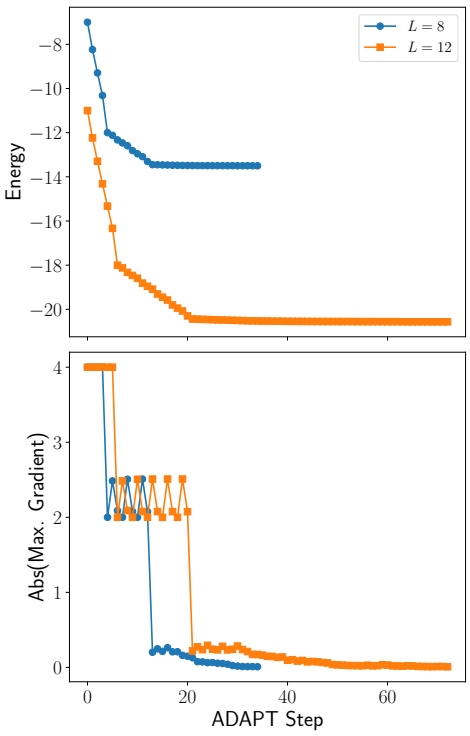

FIG. 2. (a) Energy and (b) absolute value of the maximum gradient as a function of ADAPT iteration number for the XXZ chain, using the Z-decorated XY pool. Simulation parameters are the same as for Fig. 1 with $J_z = 1$.

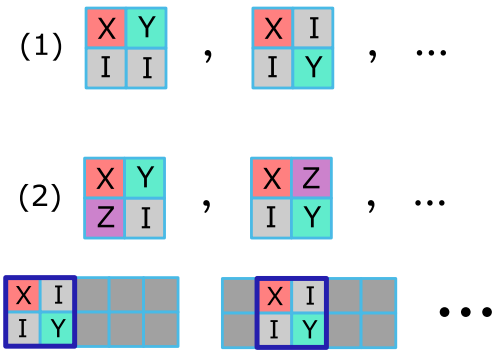

FIG. 3. Illustration of the construction of 2x2 tiles for the XXZ two-leg ladder. The operators {X,Y} and {X,Y,Z} are distributed across the block in all possible arrangements, and then tiled across the full system.

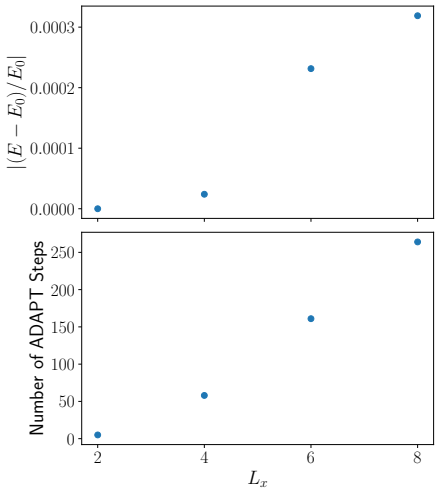

FIG. 4. (a) Relative error in the ground state energy and (b) number of ADAPT steps required for convergence for the XXZ two-leg ladder, using the pool described in Fig. 3. Simulation parameters are $L_y = 2$, $J_z = 1$.

ergy error [Fig. 5(a)] and number of steps to convergence [Fig. 5(b)] for this case. While the ground state is accurately approximated here as well, the computational difficulty associated with the L_x^2 growth in the total number of sites limits our calculations to only three values ($L_x = 2, 3, 4$). Although a quantitative extrapolation to large systems would not be warranted, it is clear that the required number of ADAPT steps grows more quickly with system size than in the 1D case. Notably, this nonlinear growth remains when the number of steps is plotted as a function of L_x^2, suggesting that the computational difficulty is at least partially due to the intrinsic connectivity of the 2D lattice, and not just the total number of sites. In any case, it appears that the 2D XXZ model presents a significant challenge for VQE methods in the quantum advantage regime, due to the large number of optimization steps needed. Nevertheless, the rapid progress in variational quantum algorithms lends hope that such difficulties may be surmounted.

Having demonstrated the power of the "tiling" approach to designing operator pools for ADAPT, we turn to possible generalizations of the method. While the above examples focused on 1D and 2D systems, it is clear how to generalize to higher dimensions (though likely to be computationally challenging). Generalizations to other popular models in condensed matter physics, such as the Hubbard model, are also straightforward. While such constructions are quite natural for lattice systems with short-range interactions, it is less clear that a simple tiling procedure will work for more complex geome-

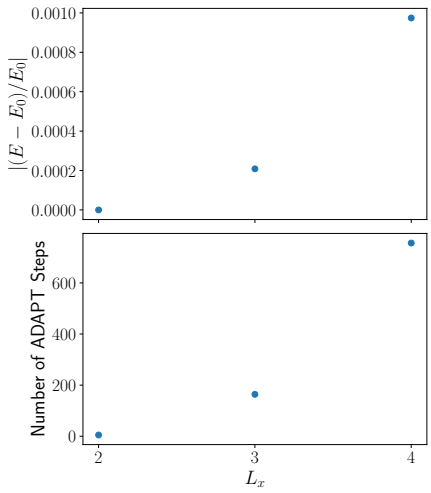

FIG. 5. (a) Relative error in the ground state energy and (b) number of ADAPT steps required for convergence for the 2D XXZ model with a square geometry, using the pool described in Fig. 3. Simulation parameters are $L_y = 2$, $J_z = 1$.

tries and/or long-range interactions (as found in quantum chemistry applications, for instance). In this case, one can consider the more general approach of starting from the full problem, selecting out various subsets of the degrees of freedom, and solving these reduced problems using ADAPT with a large operator pool. The resulting operators chosen from each subproblem can then be collected to generate a pool for executing ADAPT on the original problem. This method bears a close resemblance to many-body expansion techniques used in quantum chemistry. We leave further development of these ideas for future work.

To conclude, we have proposed and demonstrated a systematic operator tiling method for designing small but accurate operator pools for use with the ADAPT-VQE algorithm. The ability of these pools to produce highly compact ansatzes is a significant advantage for using ADAPT on noisy quantum devices. We have illustrated the approach on the open XXZ spin chain and found that it performs favorably compared to the popular Hamiltonian variational ansatz, requiring a significantly reduced number of variational parameters to reach convergence. Confirming the usefulness of the operator tiling method for other models in strongly correlated physics and chemistry is an interesting direction for future work.

Numerical calculations of the exact ground state energies were performed with the QuSpin Python library [25]. This work was supported by the Department of Energy. 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 number DE-SC0012704.

* Present Address: Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel, Maryland 20723, USA

[1] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nature Communications **5**, 1 (2014).
[2] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Nature Reviews Physics **3**, 625 (2021).
[3] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, arXiv.org (2021).
[4] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, arXiv:2111.05176 [quant-ph] (2021).
[5] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Nature **549**, 242 (2017).
[6] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, New Journal of Physics **18**, 023023 (2016).
[7] D. Wecker, M. B. Hastings, and M. Troyer, Physical Review A **92**, 042303 (2015).
[8] H. R. Grimsley, D. Claudino, S. E. Economou, E. Barnes, and N. J. Mayhall, Journal of Chemical Theory and Computation (2019), 10.1021/acs.jctc.9b01083.
[9] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, Nature Communications **10**, 1 (2019).
[10] L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, (2020).
[11] Y. Chen, L. Zhu, C. Liu, N. J. Mayhall, E. Barnes, and S. E. Economou, arXiv:2205.12283 [quant-ph] (2022).
[12] Y.-X. Yao, N. Gomes, F. Zhang, C.-Z. Wang, K.-M. Ho, T. Iadecola, and P. P. Orth, PRX Quantum **2**, 030307 (2021).
[13] N. Gomes, A. Mukherjee, F. Zhang, T. Iadecola, C.-Z. Wang, K.-M. Ho, P. P. Orth, and Y.-X. Yao, arXiv.org (2021).
[14] F. Zhang, N. Gomes, Y. Yao, P. P. Orth, and T. Iadecola, Physical Review B **104**, 075159 (2021).
[15] G. Gyawali and M. J. Lawler, Physical Review A **105**, 012413 (2022).
[16] A. G. Rattew, S. Hu, M. Pistoia, R. Chen, and S. Wood, "A Domain-agnostic, Noise-resistant, Hardware-efficient Evolutionary Variational Quantum Eigensolver," (2019), arXiv:1910.09694.
[17] D. Chivilikhin, A. Samarin, V. Ulyantsev, I. Iorsh, A. R. Oganov, and O. Kyriienko, "MoG-VQE: Multiobjective genetic variational quantum eigensolver," (2020), arXiv:2007.04424.
[18] H. L. Tang, V. O. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, PRX Quantum **2**, 020310 (2021).
[19] Y. S. Yordanov, V. Armaos, C. H. W. Barnes, and D. R. M. Arvidsson-Shukur, Communications Physics **4**,

1 (2021).
[20] V. O. Shkolnikov, N. J. Mayhall, S. E. Economou, and E. Barnes, (2021).
[21] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Physical Review Research **1**, 023025 (2019).
[22] A. Eddins, M. Motta, T. P. Gujarati, S. Bravyi, A. Mezzacapo, C. Hadfield, and S. Sheldon, PRX Quantum **3**, 010309 (2022).
[23] K. Fujii, K. Mizuta, H. Ueda, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, PRX Quantum **3**, 010346 (2022).
[24] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, PRX Quantum **1**, 020319 (2020).
[25] P. Weinberg and M. Bukov, SciPost Physics **2**, 003 (2017).

Chapter 15

ADAPT-VQE is insensitive to rough parameter landscapes and barren plateaus

This chapter includes pages from our publication [13]

H. R. Grimsley, G. S. Barron, E. Barnes, S. E. Economou, and N. J. Mayhall, "ADAPT-VQE is insensitive to rough parameter landscapes and barren plateaus", arXiv:2204.07179 (2022)

This work has been posted to arXiv.org.

The contributions to this work are as follows: My contributions were assisting in theoretical development of techniques, assisting in the preparation of the manuscript. H.R.G. primarily developed this work. N.J.M., E.B., and S.E.E. supervised.

arXiv:2204.07179v1 [quant-ph] 14 Apr 2022

ADAPT-VQE is insensitive to rough parameter landscapes and barren plateaus

Harper R. Grimsley and Nicholas J. Mayhall*
Chemistry Department, Virginia Tech

George S. Barron, Edwin Barnes, and Sophia E. Economou
Physics Department, Virginia Tech

(Dated: April 18, 2022)

Variational quantum eigensolvers (VQEs) represent a powerful class of hybrid quantum-classical algorithms for computing molecular energies. Various numerical issues exist for these methods, however, including barren plateaus and large numbers of local minima. In this work, we consider Adaptive, Problem-Tailored (ADAPT)-VQE ansätze, and examine how they are impacted by these local minima. We find that while ADAPT-VQE does not remove local minima, the gradient-informed, one-operator-at-a-time circuit construction seems to accomplish two things: First, it provides an initialization strategy that is dramatically better than random initialization, and which is applicable in situations where chemical intuition cannot help with initialization, i.e., when Hartree-Fock is a poor approximation to the ground state. Second, even if an ADAPT-VQE iteration converges to a local trap at one step, it can still "burrow" toward the exact solution by adding more operators, which preferentially deepens the occupied trap. This same mechanism helps highlight a surprising feature of ADAPT-VQE: It should not suffer optimization problems due to "barren plateaus." Even if barren plateaus appear in the parameter landscape, our analysis and simulations reveal that ADAPT-VQE avoids such regions by design.

I. INTRODUCTION

Quantum computers have long been viewed as a promising technology for quantum simulation.[1] However, the limited capabilities of Noisy, Intermediate-Scale Quantum (NISQ) devices restrict the types of algorithms that can be implemented at present.[2] While quantum phase estimation (QPE) provides a route to efficient molecular simulation,[3] the presence of both noise and errors on NISQ devices make near-term implementation of large-scale phase estimation intractable.

In response to the intractability of QPE, the variational quantum eigensolver (VQE) was introduced by Peruzzo et. al.[4] as a hybrid quantum-classical approach to finding approximate eigenvalues of a Hamiltonian, $\mathcal{H}$. In VQE, a quantum processor is used to apply a parameterized unitary transformation expressed as a quantum circuit (or even a direct pulse[5–7]), $\mathcal{U}(\boldsymbol{\theta})$, to some easily prepared reference state, $|0\rangle$.[4,8–11] The target Hamiltonian is then measured with the prepared state to obtain the energy as a function of circuit parameters:

$$E(\boldsymbol{\theta}) = \langle 0|\mathcal{U}^{\dagger}(\boldsymbol{\theta})\,\mathcal{H}\mathcal{U}(\boldsymbol{\theta})\,|0\rangle\,. \tag{1}$$

Using such quantum resources to prepare states and measure observables, a VQE will classically optimize $\boldsymbol{\theta}$ in order to minimize $E(\boldsymbol{\theta})$. The quality of the optimal energy for a given VQE is naturally dependent on the quality of the parameterization $\mathcal{U}(\boldsymbol{\theta})$, but because unitary operators are norm-preserving, the energy in Eq. (1) is variationally bounded from below by the ground-state energy of $\mathcal{H}$. The main advantage of VQEs is relatively low circuit depth,[4] avoiding the long, coherent evolutions of QPE.[12] This makes VQEs more appealing in the absence of fault-tolerant quantum computers. The circuit depth of a VQE is defined by the choice of $\mathcal{U}$, so that there is generally a trade-off between accuracy and circuit depth.

An outstanding challenge with many VQE ansätze is that the cost function, Eq. (1), creates a rough parameter landscape full of local minima, complicating the parameter optimization. Bittel and Kliesch have identified situations where there are so many far-from-optimal local minima that VQEs must be NP-hard in general.[13] The problem of local minima can be ameliorated through overparametrization in both quantum optimal control[5,14] and classical neural network settings.[15,16] This idea of overparametrization avoiding local minima has since been applied to VQEs: Rivera-Dean et. al. used this philosophy by employing a neural network to distort their cost function landscape mid-VQE.[17] This enabled them, in some cases, to escape from local minima.[18] Alternative strategies for avoiding local minima include collectively optimizing an ansatz for several Hamiltonians at the same time with a "snake" algorithm[19] and a "sweeping" approach to energy minimization called Unitary Block Optimization.[20]

A recent theoretical analysis by Larocca et. al. suggests that quantum neural networks (of which VQEs are a special case) undergo a sort of phase transition where local minima cease to be a problem.[21] This transition tends to occur when the number of parameters surpasses the dimension of the associated ansatz's dynamical Lie algebra, or DLA. The DLA for an ansatz of the form $|\Psi\rangle = e^{\theta_1 A_1} e^{\theta_2 A_2} \ldots e^{\theta_M A_M} |\phi_0\rangle$ is defined as the span of the set of repeated commutators of $\{\hat{A}_i\}$. As the authors point out, their results imply that this desirable overparametrization is likely to be unachievable for ansätze due to the exponential scaling of the DLA dimension with ansatz length. Perhaps even more alarmingly, Wierichs et. al. were able to identify situations where adding additional parameters actually hurts the performance of

gradient descent methods.[22]

In addition to the problems with local traps, it has recently been recognized that VQEs might also become impossible to optimize (even to a local mininum) as the system size increases. For sufficiently flexible or expressive VQE ansätze[23], it has been found that the energy landscape flattens (as quantified by the variance in the parameter gradients) exponentially fast as the system size increases.[24] The exponential growth of these flat landscapes (so-called "barren plateaus"), means that only a vanishingly small region of parameter space exists which has gradients large enough to measure with high enough precision to perform gradient descent. This region of concentrated cost has been termed a "narrow gorge."[25] As a result, initializing the optimization from a random point in parameter space is bound to land in a barren plateau, meaning that the number of circuit executions (shots) needed to resolve the search direction increases exponentially with the number of qubits, preventing any opportunity for quantum advantage. While intelligent heuristics for parameter initializations might help protect an optimization from getting stuck in a barren plateau (e.g., starting from a Hartree-Fock solution in molecular VQEs), the success is largely determined on a case-by-case basis.[24]

In this work, we show that our recently introduced adaptive variational algorithm, ADAPT-VQE,[26] is largely immune to local minima and barren plateaus in the parameter landscape. Both issues are avoided because the algorithm systematically "burrows" a deep well in the landscape until the global minimum is reached. In other words, ADAPT-VQE dynamically modifies its parameter landscape in such a way that problematic regions are never explored. This phenomenon can be understood directly from the gradient criterion used to iteratively update the wavefunction ansatz. We illustrate this behavior with simulations of several different molecules. In Appendix B, we also show that the smoothness of the landscape can be controlled by intentionally overparameterizing the ansatz.

II. ADAPT-VQE ALGORITHM AND NUMERICAL SIMULATIONS

A. ADAPT-VQE

In recent work, we developed a dynamic framework for constructing ansätze that have much faster energy convergence with respect to circuit depth. This approach, referred to as ADAPT-VQE,[26,27] uses measurements of the molecular energy gradient to dynamically grow an ansatz, operator by operator, creating a highly compact ansatz that quickly converges to the exact solution. Defining a pool of anti-Hermitian operators, $\mathcal{A} = \{A_i\}$, we outline the steps in Algorithm 1.

At each ADAPT-VQE iteration, the gradient, $\frac{\partial E}{\partial \theta_i}$, is measured with respect to all operators in the pool. The operator with the largest gradient magnitude is then added to the ansatz with the associated parameter initialized to zero. The other parameters in the ansatz are initialized using the optimal values from the previous step (we refer to this as parameter "recycling"). At this point, an ordinary VQE is performed using some classical optimization algorithm. In this work, we exclusively use the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method,[28] a quasi-Newton strategy, because we are explicitly seeking information about local minima, and because we are not including any noise models in our simulations. Because we initialize the new parameter added during each ADAPT-VQE iteration to zero, the new trial circuit is equivalent to the previous one during the first VQE iteration. Consequently, the energy can only improve during this VQE, i.e., the energy decreases monotonically. Parameters are added one-by-one in this fashion until some convergence criteria are achieved. Reasonable choices include the norm (either l^2 or l^∞) of the vector of gradients, $\mathbf{g}$, or the number of operators in the ansatz.

Algorithm 1: ADAPT-VQE Algorithm

$k \leftarrow 0$;
$\mathcal{U}_0(\boldsymbol{\theta}) \leftarrow \mathbf{1}$;
while *Converged = False* **do**
  for $A_i \in \mathcal{A}$ **do**
    $\mathbf{g}_i \leftarrow \left| \frac{\partial}{\partial \theta_i} \langle 0| \mathcal{U}_k^\dagger(\boldsymbol{\theta}) e^{-\theta_i A_i} \mathcal{H} e^{\theta_i A_i} \mathcal{U}_k(\boldsymbol{\theta}) |0\rangle \right|_{\theta_i=0,\boldsymbol{\theta}}$;
  end
  $i \leftarrow$ index of largest element of $\mathbf{g}$;
  $k \leftarrow k+1$;
  $\mathcal{U}_k(\boldsymbol{\theta}) \leftarrow e^{\theta_i A_i} \mathcal{U}_k(\boldsymbol{\theta})$;
  $\boldsymbol{\theta} \leftarrow \underset{\boldsymbol{\theta}}{argmin} \langle 0| \mathcal{U}_k^\dagger(\boldsymbol{\theta}) \mathcal{H} \mathcal{U}_k(\boldsymbol{\theta}) |0\rangle$;
  $E_k \leftarrow \langle 0| \mathcal{U}_k^\dagger(\boldsymbol{\theta}) \mathcal{H} \mathcal{U}_k(\boldsymbol{\theta}) |0\rangle$;
end
Return E_k, $\boldsymbol{\theta}$;

All simulations were conducted using a locally developed code which can be found on GitHub at `https://github.com/hrgrimsl/adapt`. OpenFermion[29] was used to construct matrix representations of operators under the Jordan-Wigner transformation and PySCF[30] was used to obtain molecular integrals. Because our focus in this work is to first understand the noise-free parameter landscapes associated with ADAPT-VQE, all simulations are performed without any noise models. Future work will explore how the presence of noise affects the landscapes. For all the ADAPT-VQE calculations in this work, the unitary coupled cluster with singles and doubles (UCCSD) operator pool is used,[26] without spin-complemented or spin-adapted operators. Details of this pool are provided in Appendix A.

B. Prevalence and Distribution of Local Minima

In this section, we numerically explore the parameter landscapes of several example systems using ADAPT-VQE. Our aim is to characterize the way in which the number and distribution of local minima change as ADAPT-VQE gradually increases the length of the ansatz (and thus the depth of the circuit). For each molecule and bond distance considered, we first run ADAPT-VQE normally, where the initial parameter values used in the VQE at each iteration of the algorithm are chosen to be the "recycled" parameters, i.e., the optimal values obtained from the previous iteration. This yields an ansatz that reproduces the target ground state with high accuracy.

After using ADAPT-VQE to define the ansatz, we then use this ansatz to search for local minima by repeatedly reinitializing each VQE with randomly chosen parameters, and reoptimizing.[31] In this work, we performed 300 such random initializations for each ansatz considered. For each layer of the ansatz and each random initialization, we record the minimum energy obtained by the VQE subroutine. These values correspond to the energies of local minima in the landscape associated with each ansatz.

In addition to these random initializations, we also include both the "recycled" parameters from the previous VQE (the default initialization in ADAPT-VQE[26]) and the $\mathbf{0}$ parameter vector associated with the Hartree-Fock (HF) reference. All 302 initializations of a given ansatz are then optimized with BFGS, and the resulting energy errors are shown with rainbow-colored bars in each figure. The colors indicate relative energy ordering at a given ansatz, such that red corresponds to the highest energy and violet to the lowest energy. The recycled initialization's outcome is of particular interest since this is the default, deterministic initialization for ADAPT-VQE, and the approach used when growing the ansätze used in the data. These conventions will be used throughout this work.

We consider linear H_4 and H_6 at 1 and 3 Å as toy models exhibiting varying degrees of electron correlation (and entanglement in the target wavefunction). In addition, we study LiH at 1.62 Å and BeH_2 at 1.33 Å as examples of real molecules at equilibrium geometries. These geometries were obtained through optimization at the B3LYP[32]/6-31G*[33–36] level of theory in PySCF,[30] and are given in the Supplementary Information. All ADAPT-VQE calculations were performed in the STO-3G[37,38] basis.

H_4 molecule In Fig. 1 we show the energies (relative to the global minimum obtained from a full configuration interaction (FCI) calculation) of the various local minima as a function of ansatz length (as defined by the ADAPT-VQE algorithm). After a short period without local minima, the random initializations begin to converge on an increasing number of distinct local minima as the number of parameters increases. In contrast to the random initialization, both the HF and the recycled initializations converge to the same minimum for H_4 at 1 Å, which is consistently better than the average random initialization. This is our first indication that good initializations can reliably avoid high-energy traps. Interestingly, even though ADAPT-VQE doesn't always find the lowest energy trap, it does eventually converge.

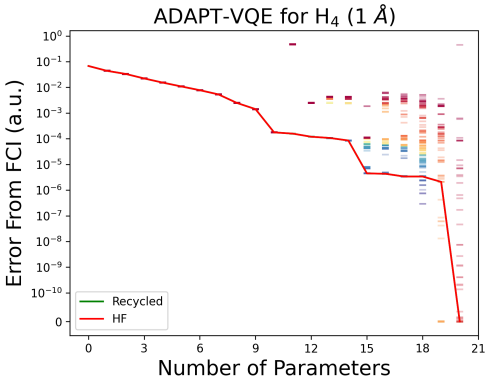

FIG. 1. ADAPT-VQE Results for H_4 at 1 Å. The x-axis corresponds to the number of ADAPT-VQE iterations, i.e. the number of operators in the ansatz at a given step. The y-axis corresponds to the error from the exact FCI energy. The red curve corresponds to the energy obtained through BFGS minimization using an HF guess, i.e. one where all parameters are zero. The green curve corresponds to the energy obtained through BFGS minimization using the standard ADAPT-VQE in which optimal parameter values in one iteration are recycled as initial guesses in the next iteration, and with the new parameter initialized to zero. The colored dots correspond to all the energies obtained through BFGS optimizations, with red being the highest energy and violet the lowest.

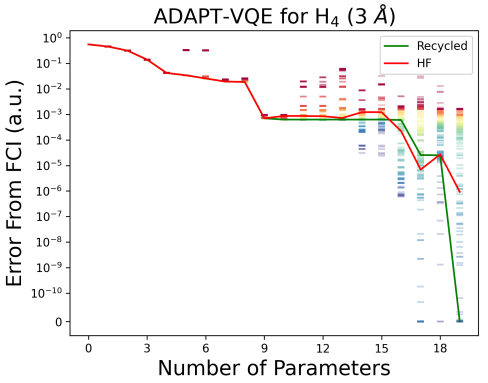

FIG. 2. ADAPT-VQE Results for H_4 at 3 Å. The axes and colors are as in Fig. 1.

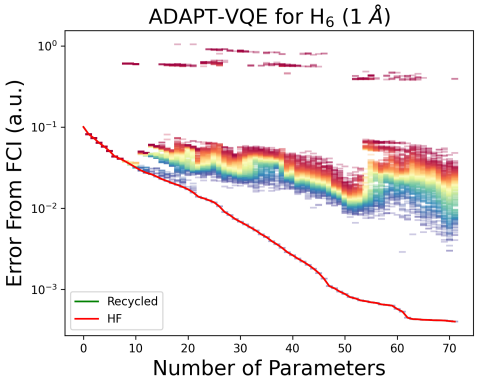

FIG. 3. ADAPT-VQE Results for H_6 at 1 Å. The axes and colors are as in Fig. 1

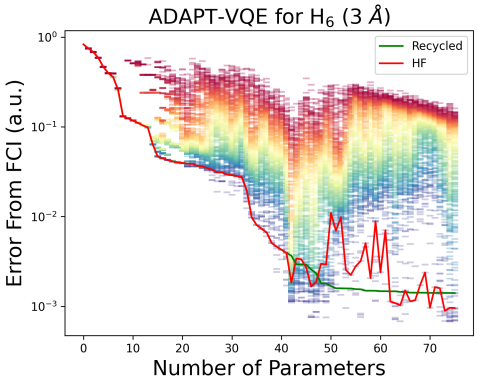

FIG. 4. ADAPT-VQE Results for H_6 at 3 Å. The axes and colors are as in Fig. 1

Additionally, we observe that there are still many local minima even after these "chemically informed" guesses are able to reach the exact ground state. In Appendix B, we consider the prospect of removing local minima through systematic overparameterization for H_4 at 1 Å. While we are successful in removing local minima using a novel "ADAPT^N" approach, deeper circuits are actually required to achieve the overparameterization than to simply add operators until ADAPT-VQE reaches the ground state in spite of local minima.

In Fig. 2, we see that for the more strongly correlated 3 Å bond distance, the HF and recycled initializations differ. The recycled initialization is able to reach the ground state with fewer parameters than the HF initialization, suggesting that it is generally a more robust choice than HF for finding the global minimum. Again, we see ADAPT-VQE converging to the exact solution far faster than a typical (yellow-green) random initialization.

H_6 molecule In Fig. 3, we begin to see the true power of an intelligent guess by simulating H_6 at 1Å. As the ansatz grows longer, a massive gap opens up between the random guesses and the HF/recycled ones. This gap implies that in practice, it is very difficult to do better than simply recycling the previous parameters in ADAPT. This gap is further numerical evidence of a "narrow gorge", in which the exact solution is hypothesized to exist.[25] Although such a landscape is often associated with optimization difficulties, here we see that ADAPT-VQE is able to stay very close to the narrow gorge, avoiding such issues. We emphasize that this feature is not only a result of good initialization,[39] but rather a cooperative effect between initialization and the gradient-guided ansatz construction. In Appendix D, we demonstrate this explicitly by performing simulations using the recycled initialization, but on randomized (not gradient-guided) ansätze.

In Fig. 4, the energy distribution of the local traps significantly increases. While we no longer see such an obvious gap for H_6 at 3 Å, the HF and recycled initializations are still far better than random ones. As the ansatz grows in depth (i.e., around 45 parameters), we notice a quick rise in the median energy of the traps found. This indicates that as the number of parameters increases, so too does the number of local traps. Furthermore, these new traps are preferentially high in energy, thus moving the median solution to higher energies. This further implies that as the system grows in size, the overwhelming number of solutions will be high in energy, making random sampling of VQE initializations intractable. Focusing on the comparison between HF and the recycled initialization, we see that they begin to differ significantly (around 50-60 parameters) where the recycled guess outperforms the zero initialization (HF). The improvement over HF is unsurprising given the strong correlation in this system, which ADAPT-VQE can handle.

LiH molecule In Fig. 5 we see similar behavior for LiH to that of H_6 at 1 Å. While the solution gap is less pronounced, both HF and the recycled initialization are always significantly better than nearly every random initialization.

BeH_2 molecule We observe similar behavior once again in Fig. 6 for BeH_2. In this relatively weakly correlated system, the HF and recycled curves overlap completely, and once again outperform the random initialization strategy.

In all cases, we observe that for more than a few parameters, local minima emerge, and for large numbers of parameters, these minima often dominate the energy landscape. In many cases initializing all parameters to 0 (HF) is a reasonable choice that leads to low energy minima. Unsurprisingly, however, the quality of HF initialization diminishes in strongly correlated systems like H_4 and H_6 at 3 Å. In such cases, we expect the optimal parameters to be further from 0 since the mean-field

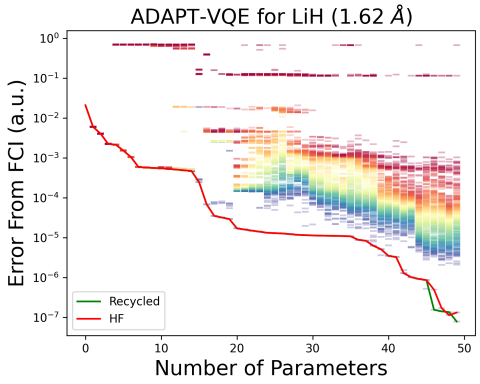

FIG. 5. ADAPT-VQE Results for LiH at 1.62 Å. The axes and colors are as in Fig. 1.

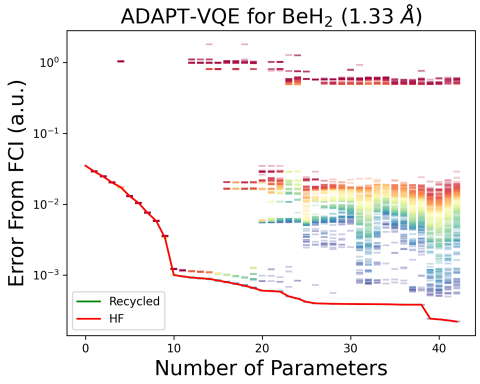

FIG. 6. ADAPT-VQE Results for BeH_2 at 1.33Å. The axes and colors are as in Fig. 1.

HF solution does not accurately represent the wavefunction. For these two strongly correlated systems, the recycled ADAPT-VQE parameters are more reliable as initial guesses, and we expect the advantage to become more pronounced as system sizes increase for strongly correlated systems.

Trap "Burrowing" The problem of local minima seems to be partially mitigated by ADAPT-VQE itself. Even in cases where the recycled initialization converges to a high-energy trap, ADAPT-VQE progresses by adding an operator which is chosen to preferentially deepen the current trap (via the gradient criterion). As such, over a sequence of ADAPT-VQE iterations, the current trap becomes increasingly deep relative to the other parameter traps, such that a gap can open up between the current minimum (which approaches the global minimum) and all other local minima. Thus ADAPT-VQE appears to "burrow" into the parameter landscape, creating a single deep well as opposed to stabilizing all local minima (i.e., reaching overparameterization). This burrowing effect is depicted graphically in Fig. 7.

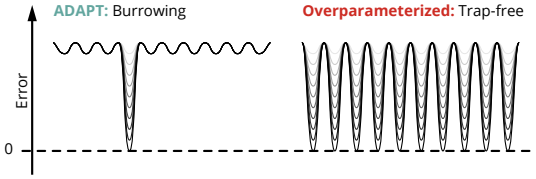

FIG. 7. Schematic illustration of how the parameter landscapes change as parameters are added using a) ADAPT-VQE, and b) a "controllable" or overparameterized ansatz which has a greater number of parameters than the rank of the DLA.

C. Sensitivity to Parameter Landscape Gradients

1. *Insensitivity to Barren Plateaus*

In the previous section, we demonstrated that while the parameter landscapes exhibit a large number of local traps that are high in energy, ADAPT-VQE is robust due to the fact that any local minimum in early stages of the algorithm can be deepened into a global minimum at later stages. This same mechanism implies a similar robustness to the presence of barren plateaus. As mentioned above, the barren plateau phenomenon has been recently recognized as a serious obstacle to the use of VQEs in practical settings. The problem arises from the observation that highly expressive ansätze (more specifically, circuits which form a 2-design), which are attractive from an accuracy perspective, exhibit an exponentially decreasing gradient variance with increasing system size. This means that the vast majority of parameter space becomes essentially flat. In the course of optimizing the parameters of such an expressive ansatz, a randomly chosen initialization will (with overwhelming probability) correspond to a point in parameter space where the gradient of the cost function is so small that an exponentially large number of measurements are needed to resolve a meaningful search direction in the presence of noise. As a result, the ability to optimize or train such expressive circuits is suspect at best. While a physically inspired parameter initialization can be effective (e.g., HF initialization), difficult cases (like those exhibiting strong correlation) may prevent efficient initialization.

Unlike the non-adaptive situation in which a static ansatz is first defined and then optimized, ADAPT-VQE slowly brings a given stationary point (initially the reference state) to the exact solution, via this burrowing mechanism. As such, each VQE performed along the way is "warm-started", in that one already has a decent ini-

tialization coming from the previous optimization. Using this recycled initialization, we have a clear characterization of the parameter landscape about the initial point: all previous parameters are optimized, and thus have zero gradients, and the newly added operator has a large gradient by design, since we specifically add the operator with the largest gradient. This means that each VQE subroutine in the ADAPT-VQE algorithm is initialized with a single parameter which is guaranteed to be greater than ϵ (the ADAPT-VQE convergence threshold). Based on this argument, we do not expect difficulty due to barren plateaus when training ADAPT-VQE ansätze as system sizes are scaled up. We emphasize that this argument does not suggest that the ansätze constructed by ADAPT-VQE are free from barren plateaus, only that our algorithm remains localized to a region in parameter space with significant gradients.

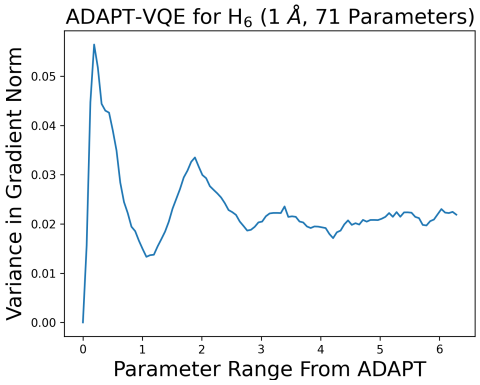

FIG. 8. Variance of the gradients for the 1 Å H_6 ansatz at 71 parameters as a function of parameter search region. Parameter vectors were selected from a 71-dimensional hypercube with each parameter randomly chosen within a certain adjustable range (x-axis), i.e. at $x = 2\pi$, each parameter is chosen from the interval $[\theta^* - 2\pi, \theta^* + 2\pi)$, where θ^* is the recycled ADAPT-VQE solution parameter.

In Fig. 8, we show that the landscape is steeper near the minimum compared to points further away in parameter space. Here, the variance of the gradient is obtained by sampling the parameter space within a hypercube of increasing volume, centered at the ADAPT-VQE default initialization. As the side of this hypercube increases (moving right on the x-axis), we plot the gradient variance of an increasing reach in parameter space. We find that near the default initialization defined by ADAPT-VQE, the landscape has large gradients. However, as the search space is increased such that we randomly consider arbitrary points in parameter space, the gradient variance flattens out.

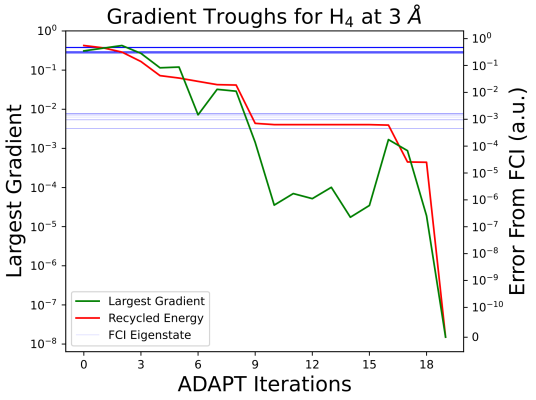

FIG. 9. ADAPT-VQE with recycled parameters for H_4 at 3 Å. The x-axis corresponds to the ADAPT-VQE iteration. The green curve depicts the gradient associated with the operator to be added at each step, while the red curve depicts the energy at each ADAPT-VQE step. The blue lines depict the excited FCI eigenstates which are lower than the HF energy.

2. *"Gradient Troughs"*

Although barren plateaus seem to pose no threat to the ability to scale up ADAPT-VQE based on the arguments in the previous section, there is still a related issue that might prevent ADAPT-VQE from converging to accurate solutions. As described above, at each ADAPT-VQE step, the ansatz is extended using the operator with the largest gradient:

$$\frac{\partial E}{\partial \theta_i} = \langle \psi(\theta) | \left[\hat{H}, \hat{A}_i \right] | \psi(\theta) \rangle \,. \tag{2}$$

The ansatz is then repeatedly extended until the largest gradient in the operator pool[40] is smaller than some threshold, ϵ. Noise on a NISQ device, however, defines some lowest possible threshold, ϵ_{min}, that can be resolved using a given shot allowance. In our earlier work,[26] we sometimes observed non-monotonic convergence of the gradients as a function of ansatz length (although the energy convergence is guaranteed to be monotonic), such that as the ansatz is extended, the pool gradients might first decrease, then increase again before finally converging. This "gradient trough", therefore presents a challenge in the presence of noise. If a gradient trough appears and drops below the NISQ resolvable threshold, ϵ_{min}, then the ADAPT-VQE algorithm may halt prematurely.

How do these gradient troughs grow with system size? If we were to find that they grow exponentially fast, meaning that the largest gradient in the operator pool is exponentially suppressed as the number of qubits increases, then this would suggest concern for the scalability of ADAPT-VQE. However, this does not need to be

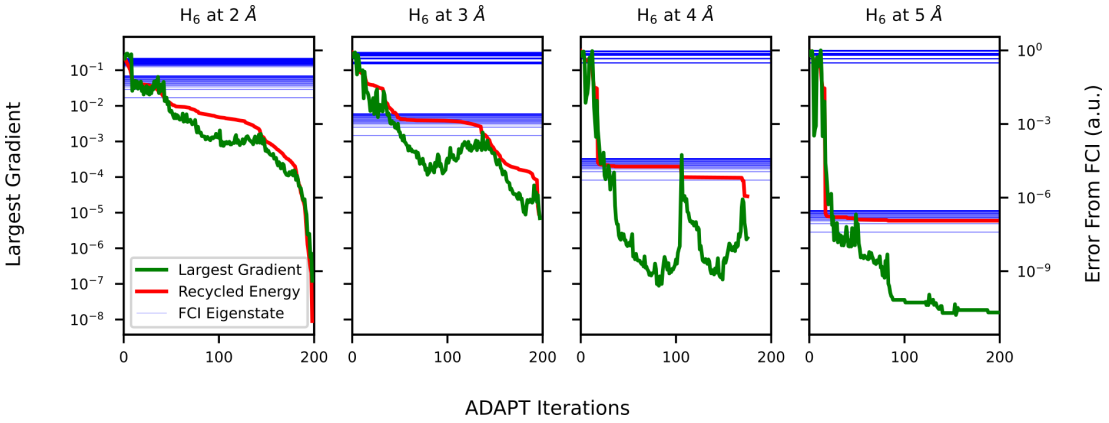

FIG. 10. ADAPT-VQE with recycled parameters for H_6 at 3 Å. The x-axis corresponds to the ADAPT-VQE iteration. The green curve depicts the gradient associated with the operator to be added at each step, while the red curve depicts the energy at each ADAPT-VQE step. The blue lines depict the excited FCI eigenstates which are lower than the HF energy.

the case. Choosing a local orbital basis[41] one can imagine trivial situations where the gradients not only avoid exponential suppression, but any suppression at all. Consider the n^{th} iteration of an ADAPT-VQE calculation of a molecular wavefunction, $|\psi_n\rangle$. If one were to double the number of qubits by adding another molecule (at infinite distance so as to remove interactions between the systems), the total wavefunction at iteration $2n$ would have a product form, $|\psi_{2n}^{AB}\rangle = |\psi_n^A\rangle\,|\psi_n^B\rangle$. Any pool operator $\hat{O}_i$ that is local to either subsystem has the exact same gradient in the supersystem, $|\psi_{2n}^{AB}\rangle$, as it does in the subsystem, $|\psi_n^A\rangle$. For example, consider an operator, $\hat{O}_i^A$, local to subsystem A:

$$\begin{aligned}\frac{\partial E^{AB}}{\partial \theta_i} &= \langle\psi_{2n}^{AB}|[\hat{H},\hat{O}_i^A]|\psi_{2n}^{AB}\rangle \\ &= \langle\psi_{2n}^{AB}|[\hat{H}^A+\hat{H}^B,\hat{O}_i^A]|\psi_{2n}^{AB}\rangle \\ &= \langle\psi_n^A|[\hat{H}^A,\hat{O}_i^A]|\psi_n^A\rangle\,\langle\psi_n^B|\psi_n^B\rangle \\ &= \frac{\partial E^A}{\partial \theta_i}. \end{aligned} \tag{3}$$

The additive separability of non-interacting subsystems is referred to as "size-consistency" in the chemistry literature. However, in addition to additive separability of the energy, size-extensive wavefunctions (like UCCSD) also demonstrate "size-intensivity" for intensive properties (e.g., density, optical gaps, etc). As shown in Eq. 3, the gradient with respect to a local rotation is not affected by the presence of an additional non-interacting system, thus demonstrating size-intensivity. In the limit of a large system, any further additions to the system size will necessarily be too far away from a given subsystem to interact. Based on this argument, we don't expect gradient troughs to deepen asymptotically with system size. However, more work is needed to characterize the behavior of gradient troughs as the system size increases in the presence of interactions.

3. *Effect of Low-Lying FCI Eigenstates*

In order to understand the nature of the "gradient troughs" discussed in Sec. II C 2, and shown in Figures 9 and 10, we superimposed the low-lying FCI energies with the ADAPT-VQE energies computed. The FCI spectrum is plotted as a set of blue horizontal lines. We only plot H_4 and H_6, as the other systems studied have no nearby excited states, nor do they exhibit any gradient troughs. In the region of the gradient trough, the energy also becomes very flat, (i.e., consider operators 9-16 in Fig. 9 and operators 50-100 in Fig. 10).

By plotting the exact eigenstates on top of these curves, one readily sees that the gradient troughs occur when ADAPT-VQE falls inside of a nearly degenerate manifold of FCI excited states. Should the ADAPT-VQE threshold be chosen loose enough (or if there is too much device noise to measure the gradient below this value) that the algorithm is aborted in this region, then ADAPT-VQE will be unable to advance further toward the ground state, remaining stuck as an approximation to an excited state (or in general some arbitrary superposition of the nearly degenerate eigenstates). This appearance of gradient troughs was first noticed in the paper that introduced ADAPT-VQE,[26] however the origin of the onset and the interpretation was not clear at that time.

As a consequence, although ADAPT-VQE isn't expected to suffer from the more general problem of barren plateaus, more work is needed to understand how to escape any gradient troughs to ensure smooth convergence to the exact solution, particularly when noise is

included. This remains an outstanding problem associated with ADAPT-VQE, warranting more research.

III. CONCLUSIONS

Underparameterized ansätze are difficult to optimize due to large numbers of local minima, while highly expressive ansätze are difficult to optimize due to barren plateaus. In this paper, we find that ADAPT-VQE does not suffer from these challenges. We have studied the parameter landscapes arising from various ADAPT-VQE generated ansätze and have arrived at the following conclusions:

1. **Recycled initialization helps avoid traps**: ADAPT-VQE's process of re-using parameters at each step focuses the search space on a local region, keeping the algorithm relatively easy to train despite the rough overall landscape. The parameter vector from the previous iteration tends to be a relatively good initial guess for the following ADAPT-VQE iteration. This means that by simply "recycling" the parameters from one ADAPT-VQE iteration to the next, the vast majority of parameter traps are entirely avoided.

2. **Trap burrowing corrects local minima**: Even if the early iterations get stuck in a trap, the adaptive construction iteratively extends the ansatz in a direction that is guaranteed to improve the cost function near the current stationary point. By continually focusing on a local point in parameter space, ADAPT-VQE can "burrow" into a given local minimum, even if the vast number of traps remain high in energy.

3. **Barren plateau avoidance**: The nature of the ADAPT-VQE algorithm suggests that barren plateaus should not prove problematic in the parameter optimization step. This originates from the fact that ADAPT-VQE specifically adds a large gradient operator, generating a steep landscape, such that a search direction is resolvable without an exponential number of shots.

4. **Gradient troughs**: ADAPT-VQE can still exhibit numerical challenges. An exponentially vanishing *pool* operator gradient could potentially arise, resulting in ADAPT-VQE becoming stuck during the operator addition step (in contrast to the parameter optimization step). Numerical evidence suggests that these gradient troughs appear when the ADAPT-VQE energy starts to converge near one or more excited states. Heuristics for diagnosing and addressing such issues will be the focus of future work.

Despite the presence of local minima and the possibility of barren plateaus in standard ADAPT-VQE ansatze, we conclude that ADAPT-VQE can be optimized reasonably well through parameter recycling. Consequently, in addition to being parameter- and gate- efficient, ADAPT-VQE appears to be relatively immune to the problems of both local minima and barren plateaus in VQEs.

IV. ACKNOWLEDGEMENTS

N.J.M., S.E.E., and E.B. are grateful for financial support provided by the U.S. Department of Energy. N.J.M. and E.B. acknowledge Award No. DE-SC0019199. 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. H.R.G. acknowledges support provided by the Institute for Critical Technology and Applied Science at Virginia Tech. The authors thank the Advanced Research Computing at Virginia Tech for the computational infrastructure.

[1] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).
[2] J. Preskill, Quantum **2**, 79 (2018).
[3] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309** (2005).
[4] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nature communications **5**, 4213 (2014).
[5] A. Asthana, C. Liu, O. R. Meitei, S. E. Economou, E. Barnes, and N. J. Mayhall, arXiv:2203.06818 [physics, physics:quant-ph] (2022), arXiv: 2203.06818.
[6] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, npj Quantum Inf **7**, 1 (2021), number: 1 Publisher: Nature Publishing Group.
[7] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, PRX Quantum **2**, 010101 (2021), publisher: American Physical Society.
[8] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Chemical Reviews **119**, 10856 (2019), publisher: American Chemical Society.
[9] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Nature Reviews Physics **3**, 625 (2021).
[10] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and

J. Tennyson, arXiv:2111.05176 [quant-ph] (2021), arXiv: 2111.05176.
[11] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, Materials Theory **6**, 2 (2022).
[12] A. Y. Kitaev, arXiv e-prints , quant (1995).
[13] L. Bittel and M. Kliesch, Physical Review Letters **127**, 120502 (2021), publisher: American Physical Society.
[14] G. Riviello, K. M. Tibbetts, C. Brif, R. Long, R.-B. Wu, T.-S. Ho, and H. Rabitz, Physical Review A **91**, 043401 (2015), publisher: American Physical Society.
[15] D. Lopez-Paz and L. Sagun, in *International Conference on Learning Representations* (2018).
[16] S. S. Du and X. Zhai, in *International Conference on Learning Representations* (2019).
[17] J. Rivera-Dean, P. Huembeli, A. Acín, and J. Bowles, arXiv:2104.02955 [quant-ph] (2021), arXiv: 2104.02955.
[18] The neural network temporarily adds additional "weight" parameters to the optimization. Even when the neural network is then reset to the identity, a better set of parameters $\boldsymbol{\theta}$ was sometimes found for the undistorted cost function.
[19] D.-B. Zhang and T. Yin, Physical Review A **101**, 032311 (2020).
[20] L. Slattery, B. Villalonga, and B. K. Clark, arXiv:2102.08403 [cond-mat, physics:quant-ph] (2021), arXiv: 2102.08403.
[21] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, arXiv:2109.11676 [quant-ph, stat] (2021), arXiv: 2109.11676.
[22] D. Wierichs, C. Gogolin, and M. Kastoryano, Physical Review Research **2**, 043246 (2020), arXiv: 2004.14666.
[23] Formal arguments have largely been restricted to 2-design structures.
[24] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Nat Commun **9**, 4812 (2018), number: 1 Publisher: Nature Publishing Group.
[25] A. Arrasmith, Z. Holmes, M. Cerezo, and P. J. Coles, arXiv:2104.05868 [quant-ph] (2021), arXiv: 2104.05868.
[26] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, Nat. Commun. **10**, 3007 (2019), tex.ids: grimsleyAdaptiveVariationalAlgorithm2019a.
[27] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, PRX Quantum **2**, 020310 (2021), publisher: American Physical Society.
[28] R. Fletcher, *Practical Methods of Optimization*, 2nd ed. (Wiley, Chichester, 2000).
[29] J. R. McClean, N. C. Rubin, K. J. Sung, I. D. Kivlichan, X. Bonet-Monroig, Y. Cao, C. Dai, E. S. Fried, C. Gidney, B. Gimby, P. Gokhale, T. Häner, T. Hardikar, V. Havlíček, O. Higgott, C. Huang, J. Izaac, Z. Jiang, X. Liu, S. McArdle, M. Neeley, T. O'Brien, B. O'Gorman, I. Ozfidan, M. D. Radin, J. Romero, N. P. D. Sawaya, B. Senjean, K. Setia, S. Sim, D. S. Steiger, M. Steudtner, Q. Sun, W. Sun, D. Wang, F. Zhang, and R. Babbush, Quantum Science and Technology **5**, 034014 (2020), publisher: IOP Publishing.
[30] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan, WIREs Computational Molecular Science **8**, e1340 (2018).
[31] Each parameter was randomly initialized on the $[0, 2\pi)$ interval in order to coincide with the period of $e^{A_i\theta_i}$ for the chosen pool.
[32] A. D. Becke, The Journal of Chemical Physics **98**, 5648 (1993), publisher: American Institute of Physics.
[33] J. D. Dill and J. A. Pople, The Journal of Chemical Physics **62**, 2921 (1975), publisher: American Institute of Physics.
[34] R. Ditchfield, W. J. Hehre, and J. A. Pople, The Journal of Chemical Physics **54**, 724 (1971), publisher: American Institute of Physics.
[35] P. C. Hariharan and J. A. Pople, Theoretica chimica acta **28**, 213 (1973).
[36] W. J. Hehre, R. Ditchfield, and J. A. Pople, The Journal of Chemical Physics **56**, 2257 (1972), publisher: American Institute of Physics.
[37] W. J. Hehre, R. F. Stewart, and J. A. Pople, The Journal of Chemical Physics **51**, 2657 (1969), publisher: American Institute of Physics.
[38] J. B. Collins, P. von R. Schleyer, J. S. Binkley, and J. A. Pople, The Journal of Chemical Physics **64**, 5142 (1976), publisher: American Institute of Physics.
[39] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Quantum Mach. Intell. **3**, 5 (2021).
[40] In the first paper the convergence criterion was taken to be the norm of the gradients in the pool, rather than the maximum.
[41] One is always free to rotate occupied or virtual orbitals without changing the associated Slater determinant, due to orbital subspace rotational invariance.
[42] F. A. Evangelista, G. K.-L. Chan, and G. E. Scuseria, arXiv:1910.10130 [cond-mat, physics:physics, physics:quant-ph] (2019), arXiv: 1910.10130.
[43] M. D. Sapova and A. K. Fedorov, arXiv:2108.11167 [physics, physics:quant-ph] (2021), arXiv: 2108.11167.

Appendix A: Operator Pool

In this work, all experiments were performed with a pool of particle-hole (i.e., not generalized) Fermionic operators. These operators were restricted to single excitations of the form $\hat{a}_i^a - \hat{a}_a^i$ and double excitations of the form $\hat{a}_{ij}^{ab} - \hat{a}_{ab}^{ij}$, where i and j are occupied spin orbitals and a and b are virtual spin orbitals. While these operators are not required to respect $\hat{S}^2$ symmetry, the pool does not include "spin flips," operators which change the total numbers of alpha and beta electrons. This pool grows quartically with system size. If these operators are used as generators for an ansatz of the form $|\Psi\rangle = e^{\theta_1 A_1} e^{\theta_2 A_2} \dots e^{\theta_M A_M} |\phi_0\rangle$ with infinite repetition of the operators and correct ordering, the FCI solution should be reachable.[42] However, we have not yet proven that the ADAPT-VQE convergence criterion goes to zero *only* when converged to an exact eigenstate.

Appendix B: ADAPT$^{\mathbf{N}}$

In this section we turn our attention to the following question: can we remove all local traps from ADAPT-VQE without necessarily reaching "controllability" where all local minima are exact? It follows from

Ref. 21 that one can systematically overparametrize a given ansatz of the form $|\Psi\rangle = e^{\theta_1 A_1} e^{\theta_2 A_2} \dots e^{\theta_M A_M} |\phi_0\rangle$ by simply repeating all the operators N times such that the number of parameters approaches or exceeds the dimension of the associated DLA. This ansatz repetition can be defined as:

$$|\Psi\rangle = \prod_{i=1}^{N} \left(\prod_{j=1}^{M} e^{\theta_j^i A_j} \right) |\phi_0\rangle \,. \tag{B1}$$

We use the superscript on θ_j^i to indicate the ith independent parameter associated with the same operator A_j. Because the same generators are being used, the DLA remains the same, but the number of parameters has now been multiplied by N. We apply this idea to ADAPT-VQE in what we call ADAPT^N. As with normal ADAPT, we begin with a reference $|\phi_0\rangle$. We choose the first operator A_1 to add as before, but apply it to the reference N times with N distinct parameters: B2

$$|\Psi_1\rangle = e^{\theta_1^1 A_1} e^{\theta_1^2 A_1} \dots e^{\theta_1^N A_1} |\phi_0\rangle \,. \tag{B2}$$

The energy of $|\Psi_1\rangle$ is variationally minimized with respect to $\boldsymbol{\theta}$, beginning from $\boldsymbol{\theta} = \mathbf{0}$. For each operator A_i in the operator pool, we now measure the energy gradient associated with adding it to the front of the ansatz with a parameter of 0. As with the original ADAPT-VQE, we choose the next operator in our ansatz, A_2, to be the operator with largest gradient magnitude. We then add the operator in a collated fashion:

$$|\Psi_2\rangle = \left(e^{\theta_2^1 A_2} e^{\theta_1^1 A_1} \right) \left(e^{\theta_2^2 A_2} e^{\theta_1^2 A_1} \right) \dots \left(e^{\theta_2^N A_2} e^{\theta_1^N A_1} \right) |\phi_0\rangle \,. \tag{B3}$$

This procedure of adding a new operator to the "core" ansatz, and replicating N times, is then repeated until convergence. The intuitive way to understand ADAPT^N is that ADAPT-VQE is adding N parameters at a time, rather than the usual 1. However, unlike the "batched" approach of Ref. 43, our method does not add N *different* operators, a choice made to ensure that the DLA dimension does not increase with N. Because the ratio of the number of parameters to DLA dimension increases with N, we hypothesize that ADAPT^N should exhibit fewer local minima with increased N.

To test this hypothesis, we consider first the numerically challenging linear H_4 at 1Å bond distance in the STO-3G basis. In Figure 11, we show the convergence of ADAPT^N for $N \leq 4$ as a function of the number of parameters. Like before, we consider 300 random initializations at each ansatz length in addition to the recycled initialization from the previous VQE subroutine, and the HF initialization. We then plot the optimized results of all 302 initializations. For ADAPT^1 (i.e. standard ADAPT) and ADAPT^2, we see similar behavior. Random guesses are frequently better than recycled ones, with random initializations of $\boldsymbol{\theta}$ landing in numerous local minima. For ADAPT^1 and ADAPT^2, the landscape is dominated by local minima except when there are very few operators or when FCI has almost been reached. For ADAPT^3, most of the local minima seem to be gone, and local minima never outperform the recycled initialization. For ADAPT^4, every VQE subroutine appears to be free of local minima. This gives us a simple way to make ADAPT-VQE (or similar VQEs) trap-free. The downside is that more parameters and deeper circuits are required to introduce the same number of distinct operators. Based on these simulations, the lower values of N seem to be more depth-efficient at achieving any given accuracy at any given circuit depth, even if only the recycled initialization is tried. Furthermore, based on Ref. 21, the necessary N to keep ADAPT-VQE trap-free will increase with the number of operators needed to describe a system, which will generally grow with system size. We therefore expect this overparametrization strategy to scale poorly. Contrary to the prediction of Ref. 22, however, we have shown that an ADAPT-like method can avoid local traps via overparametrization.

Appendix C: Difficulty of Exhaustive Sampling

While sampling over parameter initializations provides a path for enumerating different local traps in the landscape, the computational cost of such sampling is high, and any such study will necessarily exhibit artifacts arising from insufficient sampling. While our best efforts were made to avoid such artifacts, we have noticed a few. Take for example, Figure 4, where the best solution at 60 parameters is very slightly worse than the best solution at 59 parameters. We tried using 1000 random initializations at 60 parameters (Fig. 12), and were still unable to improve on the recycled initialization at 60 parameters, let alone the best initialization at 59 parameters. While these small artifacts persist, the general trends seem converged. For instance, in Fig. 12, increasing the number of initializations increases the number of solutions found. The distribution of the solutions is essentially unchanged, with the median values (yellow-green) staying roughly in the same place.

Appendix D: Importance of gradient-guided ansatz

In Figs. 1-6, we demonstrated that the recycled (and HF) initialization performed far better than the average local minimum. In order to ascertain the role that the ADAPT-VQE ansatz construction has on the ability of the ansatz to burrow, we carry out similar simulations to those presented in Fig. 3, except that before performing the VQE optimizations, we shuffle the operators in the ansatz. As a result, the ansatz is no longer constructed by ADAPT-VQE. The results are shown in Figs. 13-15. In each of the 3 random reorderings, we find that the large gap that appeared in Fig. 3 between the recycled (and HF) minimum and the local traps, largely disappears when considering a randomly shuffled ansatz. This

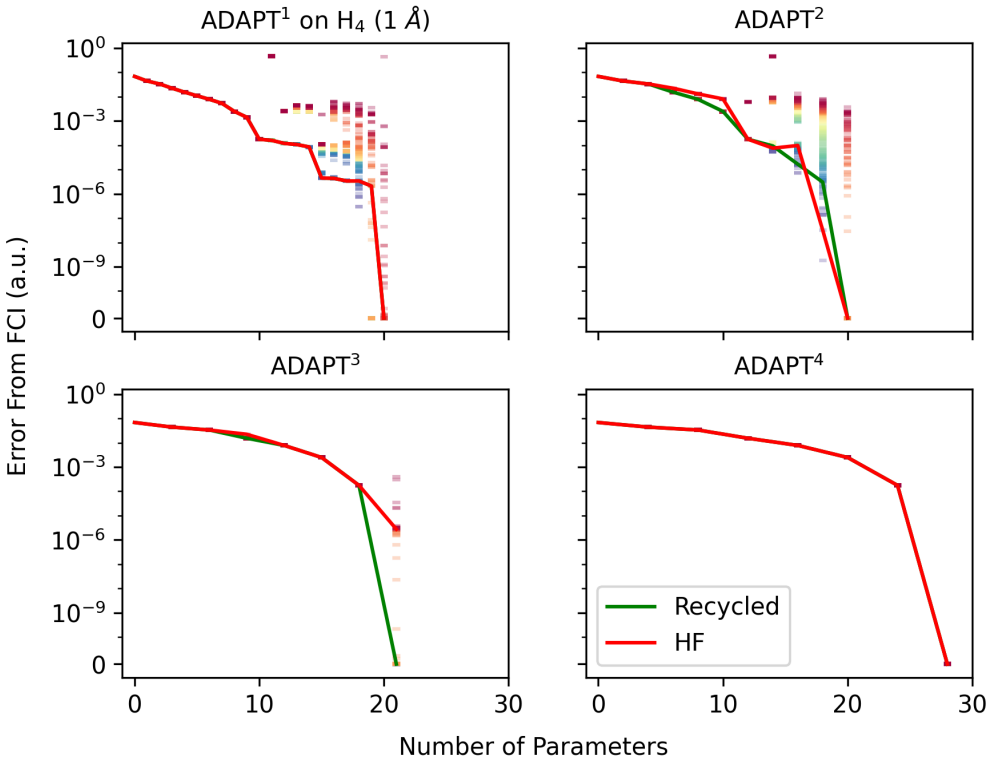

FIG. 11. ADAPT^N for $N = 1, 2, 3, 4$ on H_4 at 1 Å separation. The axes and colors are as in Fig. 1

is clear evidence that it is not sufficient to only slowly optimize the parameters one-at-a-time via recycled initialization, but that the ansatz must be grown in an efficient manner as well if one is to observe the full burrowing effect.

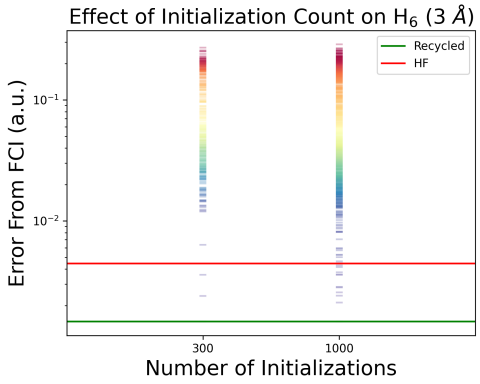

FIG. 12. Results of various initializations for the 60-parameter ADAPT-VQE ansatz for H_6 at 3Å. The original 300 random initializations from Figure 4 are shown on the left. 1000 random initializations are depicted on the right.

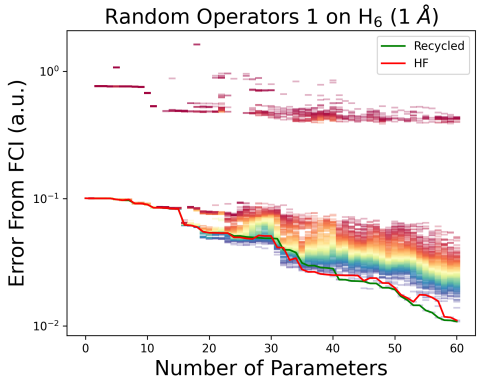

FIG. 13. Random Reordering 1 of the 71 operators identified in Fig. 3. (Computer time was exhausted before all 71 operators were added in this case.) The x-axis corresponds to the number of operators in the ansatz at a given step. The y-axis corresponds to the error from the exact FCI energy. The red curve (HF) corresponds to the energy obtained through BFGS minimization using an all-zero initialization. The green curve corresponds to the energy obtained through BFGS minimization using the standard ADAPT-VQE recycled initialization. The colored dashes correspond to all the energies obtained through BFGS optimizations, with red being the highest energy and violet the lowest.

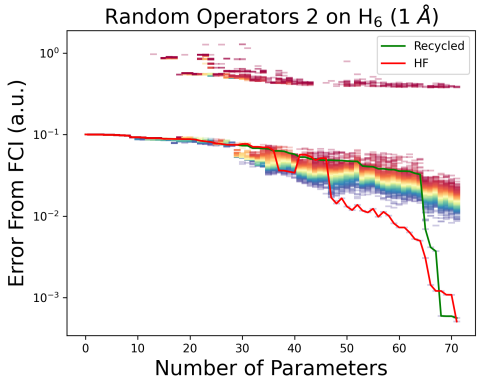

FIG. 14. Random Reordering 2 of the 71 operators identified in Fig. 3. Labels same as in Fig. 13

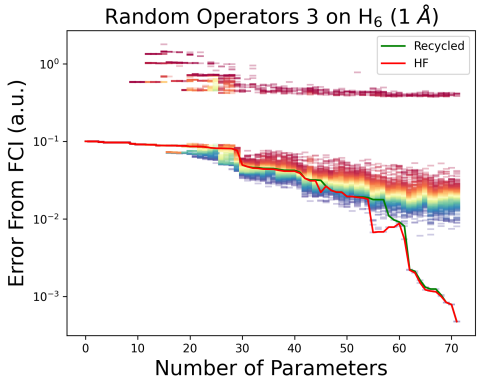

FIG. 15. Random Reordering 3 of the 71 operators identified in Fig. 3. Labels same as in Fig. 13

Part V

Conclusions

Chapter 16

Conclusions

In Chapter 1 of this dissertation, we introduce foundational concepts in quantum computing, including common notation, conventions, and concepts necessary for the discussion that follows. This includes the discussion of quantum states, quantum gates, quantum circuits, VQAs, and quantum control for superconducting qubits.

In Part II we use and develop techniques in quantum control to improve the low-level performance of algorithms and gate performance. This comes in several primary contributions. In Chapters 3, 4, and 5 we develop novel entangling gates by using the SWIPHT protocol. In each case, we use the physics of the device to our advantage, overcoming existing limitations of existing hardware.

In Part III we improve and analyze the performance of quantum simulation algorithms, both in the presence and absence of noise. Firstly, we demonstrate the advantages of using an existing readout error mitigation scheme on evaluating VQE cost functions, and simultaneously use the resulting calibration data for device characterization. In Chapters 7 and 8 we focus on our novel technique of producing symmetry ansatze for use in VQEs. In the former chapter, we develop these techniques in a primarily theoretical context, and numerically verify statements about their generality. In the latter chapter, we analyze the performance of these ansatze in the presence of noise, and compare with existing techniques. We find that the reduction in classical and quantum resources results in overall improved algorithm performance. We also consider applications to other systems, in some cases using real quantum

devices. Then, in Chapter 9 we compare this technique and other state-of-the-art techniques in quantum simulation. In Chapter 10 we use quantum control to improve the VQE algorithm compared to existing methods. Finally, in Chapter 11 develop a novel technique to preparing Bethe Ansatz eigenstates. We show that the algorithm has improved gate counts compared to existing techniques, and demonstrate it numerically. In each case, we improve the applicability of quantum simulation algorithms for both current and future devices.

In Part IV we focus on a collection of works on the ADAPT-VQE algorithm. In Chapter 12 we develop the qubit-ADAPT-VQE algorithm, in which we generalize the ADAPT-VQE algorithm to other operator pools, resulting in improved gate counts. We also derive operator pools which can be used to prepare arbitrary states. In Chapter 13 we develop the ADAPT-QAOA algorithm, which improves the existing QAOA algorithm by using gradient-selected mixers which improve algorithm convergence. In Chapter 14 construct operator pools for the ADAPT-VQE algorithm in a way that is tailored to Hamiltonians with lattice structures, a task which is otherwise computationally demanding. We numerically demonstrate that this provides performance improvements compared to existing methods. Finally, we numerically demonstrate the "burrowing" phenomenon in the ADAPT-VQE algorithm and its connection and insensitivity to barren plateaus. In each case, we improve the applicability and performance of the ADAPT-VQE algorithm and its variants, in turn reducing the amount of quantum resources necessary to perform certain tasks.

Quantum computing as a field is still very much in its early stages. There have been massive developments and improvements to algorithms, hardware, and applications since its creation. Nevertheless, inquiries into these systems and their application will require a deep understanding of not only the application at a high level, or the implementation at a low level, but instead the relationship between these two points in the quantum computing "stack". As such, it is imperative to study the interaction and relationship between the components

of a quantum computer at each level. To this end, this work contains a collection of novel contributions towards this effort, in order to use device physics and error mitigation in order to improve the performance of quantum computers overall.

Bibliography

[1]G. S. Barron, F. A. Calderon-Vargas, J. L. Long, D. P. Pappas, and S. E. Economou, "Microwave-based arbitrary CPHASE gates for transmon qubits", Physical Review B **101**, ARTN05450810.1103/PhysRevB.101.054508 (2020).

[2]G. S. Barron and C. J. Wood, "Measurement Error Mitigation for Variational Quantum Algorithms", arXiv:2010.08520 (2020).

[3]G. S. Barron, B. T. Gard, O. J. Altman, N. J. Mayhall, E. Barnes, and S. E. Economou, "Preserving symmetries for variational quantum eigensolvers in the presence of noise", Physical Review Applied **16**, 034003 (2021).

[4]J. Long, T. Zhao, M. Bal, R. Zhao, G. S. Barron, H.-s. Ku, J. A. Howard, X. Wu, C. R. H. McRae, X.-H. Deng, G. J. Ribeill, M. Singh, T. A. Ohki, E. Barnes, S. E. Economou, and D. P. Pappas, "A universal quantum gate set for transmon qubits with strong ZZ interactions", arXiv:2103.12305 (2021).

[5]F. A. Calderon-Vargas, G. S. Barron, X. H. Deng, A. J. Sigillito, E. Barnes, and S. E. Economou, "Fast high-fidelity entangling gates for spin qubits in Si double quantum dots", Physical Review B **100**, ARTN03530410.1103/PhysRevB.100.035304 (2019).

[6]B. T. Gard, L. H. 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 Information **6**, 10.1038/s41534-019-0240-1 (2020).

[7]K. Yeter-Aydeniz, B. T. Gard, J. Jakowski, S. Majumder, G. S. Barron, G. Siopsis, T. S. Humble, and R. C. Pooser, "Benchmarking quantum chemistry computations with varia-

tional, imaginary time evolution, and krylov space solver algorithms", Advanced Quantum Technologies **4**, 2100012 (2021).

[8]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 Information **7**, 155 (2021).

[9]J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, "Preparing bethe ansatz eigenstates on a quantum computer", PRX Quantum **2**, 040329 (2021).

[10]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).

[11]L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, "Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer", Phys. Rev. Research **4**, 033029 (2022).

[12]J. S. Van Dyke, G. S. Barron, N. J. Mayhall, E. Barnes, and S. E. Economou, "Scaling adaptive quantum simulation algorithms via operator pool tiling", arXiv:2206.14215 (2022).

[13]H. R. Grimsley, G. S. Barron, E. Barnes, S. E. Economou, and N. J. Mayhall, "ADAPT-VQE is insensitive to rough parameter landscapes and barren plateaus", arXiv:2204.07179 (2022).

[14]M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information: 10th anniversary edition* (Cambridge University Press, 2010).

[15]M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, "Variational quantum algorithms", Nature Reviews Physics **3**, 625–644 (2021).

[16]D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, "VQE method: a short survey and recent developments", Materials Theory **6**, 2 (2022).

[17]J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, "The Variational Quantum Eigensolver: A review of methods and best practices", Physics Reports **986**, 1–128 (2022).

[18]Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, "Quantum Chemistry in the Age of Quantum Computing", Chemical Reviews **119**, 10856–10915 (2019).

[19]J. Preskill, "Quantum Computing in the NISQ era and beyond", Quantum **2**, 79 (2018).

[20]E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, M. Greiner, K. R. Hazzard, R. G. Hulet, A. J. Kollár, B. L. Lev, M. D. Lukin, R. Ma, X. Mi, S. Misra, C. Monroe, K. Murch, Z. Nazario, K.-K. Ni, A. C. Potter, P. Roushan, M. Saffman, M. Schleier-Smith, I. Siddiqi, R. Simmonds, M. Singh, I. Spielman, K. Temme, D. S. Weiss, J. Vu čkovi ć, V. Vuleti ć, J. Ye, and M. Zwierlein, "Quantum simulators: architectures and opportunities", PRX Quantum **2**, 017003 (2021).