Estimating exact energies in quantum simulation without Toffoli gates

Mark Steudtner^{1,2} and Stephanie Wehner²

¹*Instituut-Lorentz, Universiteit Leiden, P.O. Box 9506, 2300 RA Leiden, The Netherlands*

²*QuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands*

(Received 26 January 2020; accepted 25 March 2020; published 18 May 2020)

Qubitization is a modern approach to estimate Hamiltonian eigenvalues without simulating its time evolution. While in this way approximation errors are avoided, its resource and gate requirements are more extensive: qubitization requires additional qubits to store information about the Hamiltonian, and Toffoli gates to probe them throughout the routine. Recently, it was shown that storing the Hamiltonian in a unary representation can alleviate the need for such gates in one of the qubitization subroutines. Building on that principle, we develop an alternative decomposition of the entire algorithm: without Toffoli gates, we can encode the Hamiltonian into qubits within logarithmic depth.

DOI: [10.1103/PhysRevA.101.052329](https://doi.org/10.1103/PhysRevA.101.052329)

I. INTRODUCTION

In digital quantum simulation, a quantum computer acts as a universal simulator for systems difficult to predict with classical methods. However, the goal of this field is beyond simple imitation of one system by another: after the Hamiltonian of the model is mapped onto qubits, quantum algorithms are employed to extract its spectrum and eigenstates. The perhaps most sophisticated algorithm of such kind is quantum phase estimation, which allows one to project into spectral eigenstates by a Fourier analysis of the simulated time evolution (under the model Hamiltonian). However, despite being conceptually simple, quantum phase estimation is challenging on a technical level. Not only are its requirements beyond the abilities of current hardware, but it is likely to even present technical challenges in the future. Part of the problem is that the time evolution cannot be simulated exactly, but must generally be approximated. As originally suggested in [1], this can be achieved with Trotterization, which means that the simulator is made to evolve in stroboscopic slices of the exact time evolution. The shorter the time period of the evolution, the more accurate is the approximation, but quantum phase estimation has a better resolution for longer time evolutions [2,3]. The algorithm also requires an additional register of estimator qubits to couple to every slice in the trotterized time evolution, which is likely to demand nonlocal operations inside the quantum computer.

However, there are more advanced methods that could replace Trotterization in the phase estimation algorithm. In *qubitization* [4], the simulator is extended by a certain number of qubits. The time evolution is then replaced by a unitary that, in a certain subspace of the extension, acts as the Hamiltonian on the simulator qubits. As the unitary describes a rotation out of that subspace, the rotation angles—functions of the Hamiltonian eigenvalues—can be read out by the phase estimation routine. The appeal of qubitization is that it does not involve any approximation of the Hamiltonian; however, it generally requires higher-level quantum operations, such as the Toffoli gate [5]. This is in particular the case for when one tries to keep the number of additional qubits to a

minimum. For qubitization, Toffoli gates can be regarded as a symptom of a compressed data structure, meaning that bit strings are encoded nonlinearly. For the unitary entangling the simulator qubits to the extension, Toffoli gates signify a drastic increase in circuit depth. Unfortunately, the depth of a quantum algorithm (its theoretical runtime) is a bottleneck of every quantum computation.

Recently, a qubitization version with a decompressed storage structure was introduced in [5], opening an important subroutine of the algorithm up for parallelization. Here, we want to build upon this work, and obtain a qubitization version that is entirely free of Toffoli gates. For that purpose, we apply recently popular state preparation techniques [6,7], and so parallelize the remaining subroutines of qubitization. The new routines also fit more naturally into the framework of quantum phase estimation: rather than being in permanent contact with the estimator qubits like the Trotter circuits, the qubitization routines have relay points, at which singular controlled gates can act as switches. Without any compression, we will present a version of qubitization in which all its components are of low algorithmic depth. For local Hamiltonians, this phase estimation version has a total runtime scaling logarithmically with the number of terms.

II. BACKGROUND

The goal of quantum simulation is to extract the eigenstates and eigenvalues of a Hamiltonian acting on qubits. Let that Hamiltonian have Λ terms and be of the form

$$H = \sum_{k=1}^{\Lambda} \alpha_k p_{\text{sim}}^k, \quad (1)$$

where p_{sim}^k are Pauli strings—signed products of Pauli operators $\{X, Y, Z\}$ acting on different qubits inside the simulator, a quantum register we will call “sim.” By shifting all the minus signs into the Pauli strings, their coefficients α_k are positive and for the moment we will normalize them to $\sum_{j=1}^{\Lambda} \alpha_j \leq 1$, such that $-1 \leq E \leq +1$ for all energy eigenvalues E . Trotterized time evolution and other product formulas

FIG. 1. Quantum phase estimation circuits. H is the Hadamard gate and $|\varphi\rangle_{\text{sim}}$ a trial state of the simulator. (a) Kitaev's phase estimation circuit [14]. U is an approximation to the simulated time evolution $\exp(iH)$. (b) Qubitization circuit [5] with one estimator qubit, featuring the three unitaries $\hat{V}$, $\hat{S}$, and $\hat{G}$, where only the latter is applied conditionally. All qubitization qubits are initialized in the zero state, such that the entire register equals $|0\rangle = \bigotimes_{j=1}^n |0\rangle_j$.

[2,3,8–11] approximate $\exp(iH)$ by products of Pauli string rotations $\exp(i\alpha_k p_{\text{sim}}^k \Delta)$ for different time slices $\Delta \leq 1$. Rotations of Pauli strings are generally straightforward to implement [12]. Using those textbook circuits, quantum phase estimation [13] can discern states with different eigenvalues E . The projective power of the phase estimation algorithm is generally bigger the more qubits it has at its disposal. In its minimal version [14], Kitaev's phase estimation [shown in Fig. 1(a)], only one estimator qubit is required: Hamiltonian eigenvalues are inferred from its measurement histogram. The routine in Fig. 1(a) features the component U , an approximation to the simulated time evolution, $U \approx \exp(iH)$.

An alternative to Trotterized time evolution is qubitization [4]. In qubitization without quantum signal processing [15], not the energy E is estimated, but $\pm \arccos(E)$, where the sign $\pm$ occurs for every run of the circuit at random. While Trotterization only approximates the time evolution e^{iH} , qubitization employs a sequence of two unitaries, $\hat{V}$ and $\hat{S}$, to estimate the phase factors $\exp(i \arccos E) = E \pm i\sqrt{1-E^2}$ for the exact energy eigenvalues E .

However, qubitization has a major drawback: it requires an additional register of at least $n \geq \log \Lambda$ qubits. We will refer to it as the qubitization register. In those qubits, information about the Hamiltonian must be encoded with a unitary $\hat{G}$. This operation prepares a state $|G\rangle$, which features the weights of the Hamiltonian $\{\sqrt{\alpha_k}\}$. Besides $\hat{S}$, a reflection of $\hat{G}$, $\hat{S}|G\rangle = -|G\rangle$, an entangling operation $\hat{V}$ is required to facilitate the action of the Hamiltonian (on the simulation register) within the subspace of $|G\rangle$, i.e., $\langle G|\hat{V}|G\rangle = H$. All three unitaries have the form [4]

$$\hat{G}|0\rangle = |G\rangle = \sum_{k=1}^{\Lambda} \sqrt{\alpha_k} |\mu_k\rangle, \quad (2)$$

$$\hat{S} = 2|G\rangle\langle G| - \mathbb{I}, \quad (3)$$

$$\hat{V} = \sum_{k=1}^{\Lambda} |\mu_k\rangle\langle\mu_k| \otimes p_{\text{sim}}^k + |0\rangle\langle 0| \otimes \mathbb{I}_{\text{sim}} + \dots, \quad (4)$$

where $|0\rangle$ is the state in which all qubits are in $|0\rangle$, and $\{|\mu_k\rangle\}$ is a subset of the computational basis not containing $|0\rangle$. The state $|G\rangle$ in (2) superposes the configurations $|\mu_k\rangle$ according to the weights $\{\alpha_k\}$, which from now on shall be normalized to $\sum_k \alpha_k = 1$. In (4), only the relevant part of $\hat{V}$ is presented: its action on the sim-register given it finds the states $\{|\mu_k\rangle\}$ in the qubitization register, and trivial action if the qubitization register is in $|0\rangle$. The action of $\hat{V}$ outside of this subspace does not play a role in the qubitization process.

Note that throughout the literature, the unitaries $\hat{G}$ and $\hat{V}$ are also referred to as PREPARE and SELECT [16].

A naive implementation of the qubitization routine would now replace the controlled time evolution, U , in Fig. 1(a) with the sequence of controlled $\hat{V}$ and $\hat{S}$. However, $\hat{V}$ features controlled gates already, and so the amount of Toffoli gates would increase if $\hat{V}$ was applied conditionally. Fortunately, as was noticed in [5], $\hat{V}$ acts trivially if $|G\rangle$ is not initialized, and the same applies to $\hat{S}$. For the benefit of the gate complexity, the application of $\hat{G}$ was made conditional, and $\hat{G}$ was reintroduced after $\hat{S} \cdot \hat{V}$ for only the $|0\rangle$ subspace of the estimator qubit; see Fig. 1(b). Note that using qubitization, the phase estimation does no longer project into eigenstates of the Hamiltonian (in the sim-register). When estimating $\pm \arccos E$, the eigenstate of E must be retrieved probabilistically. Again we learn from [5] that the unitary $\hat{G}$ can be applied once more after the circuit in Fig. 1(b). A following measurement of every qubit in the qubitization register has then a 50% chance of all outcomes being $+1$, which is a flag that the sim-register has been projected into the eigenstate of E .

Within this work, we are particularly interested in the basis of the qubitization register. Here, $\mu_k = \mu_{k,1}\mu_{k,2}\mu_{k,3}\dots\mu_{k,n}$ is a string of n bits, $\mu_{k,j} \in \{0, 1\}$, describing a computational basis state by $|\mu_k\rangle = \bigotimes_{j=1}^n |\mu_{k,j}\rangle_j$. The set of bit strings $\{\mu_k\}$ is determined by the encoding (i.e., the mapping: $k \mapsto \mu_k$), which has important implications for the depth of $\hat{G}$, $\hat{S}$, and $\hat{V}$. Prior works have considered two encoding strategies.

Binary encoding [4,16]. To reach the minimal qubit requirements of $n = \log \Lambda$, all possible bit strings $\{0, 1\}^{\otimes n}$ must be considered as basis configurations $\{\mu_k\}$. For the construction of $\hat{V}$, sequences of multiqubit Toffoli gates must probe the qubitization register for configurations μ_k and apply p_{sim}^k for all k from 1 to Λ ; see Fig. 2(a) or [16]. A comparable effort is required when constructing the state $|G\rangle$, where $O(\Lambda)$ rotations have to be conditioned on specific basis configurations, probed again with Toffoli gates. The algorithmic depth of $\hat{V}$, $\hat{S}$, and $\hat{G}$ is therefore fixed to be at least proportional to Λ .

Unary encoding [5]. To eliminate the need for Toffoli gates in the construction of $\hat{V}$ (but not directly $\hat{S}$), we can encode information in $|G\rangle$ with only the strings $\mu_k = e_k$, where e_k is the bit string corresponding to the k th unit vector with $e_{k,j} = \delta_{jk}$. While this increases the qubit count to $n = \Lambda$, it allows one to define $\hat{V}$ in a simpler way

$$\hat{V} = \prod_{k=1}^{\Lambda} (|0\rangle\langle 0|_k + |1\rangle\langle 1|_k \otimes p_{\text{sim}}^k), \quad (5)$$

FIG. 2. Implementations of $\hat{V}$. (a) Binary representation requiring $n = \log \Lambda$ qubits. Sequences of multi-Toffoli gates probing the state $|G\rangle$ for the conditioned application of singular p^k . To benefit from cancellations between adjacent gates [16], μ_k is the k th word of the Gray code [23]. However, even with those cancellations, the depth of $\hat{V}$ is $O(\Lambda)$. (b) Unary encoding requiring $n = \Lambda$ qubits. The application of the string p^k is only controlled on the k th qubit in the qubitization register. Since nonoverlapping p strings can be applied in parallel, $\hat{V}$ can in the best case be implemented within $O(1)$ time.

which means that instead of checking the states of the entire qubitization register, each string p_{sim}^k is applied conditioned on the k th qubit for all k between 1 and Λ ; see Fig. 2(b). Since each of those gates is controlled by a different qubit, the only thing hindering a massive parallelization of $\hat{V}$ is the structure of the p strings themselves: the complexity of (5) depends on their individual Pauli weight and whether they overlap with one another. Fortunately, there is a considerable body of theoretical work on how to produce local strings when simulating fermions [17–22] and bosons [23]. While the decompressed data structure of the unary encoding benefits the complexity of $\hat{V}$, benefits for the other components have not been shown: regardless of the encoding, $\hat{S}$ is suggested to be implemented by the sequence $\hat{G} \cdot (2|0\rangle\langle 0| - \mathbb{I}) \cdot \hat{G}^\dagger$ [5,16], with the reflection $(2|0\rangle\langle 0| - \mathbb{I})$ requiring a Toffoli-type gate across the entire qubitization register; see Fig. 3. For the unary encoding, this gate would have particularly many controls—the Λ -fold Toffoli would most certainly dominate the time complexity of the entire algorithm. Here, adding more qubits would help: a depth of $\log \Lambda$ could theoretically be achieved by doubling the size of the qubitization register; however, the high resource requirements are a critical downside of the

FIG. 3. Implementation of $\hat{S}$ with a Toffoli-type gate in a five-qubit qubitization register [5,16]. When using a unary encoding, the size of this register will be proportional to the number of Hamiltonian terms, resulting in a large gate.

unary encoding already. Λ must typically be regarded as a large number, such that adding another Λ qubits is a serious repercussion. Other strategies with which the reflection could be approximated can be found in [24]. Fortunately, the unary encoding does not actually require $\hat{S}$ a high-level gate and the following sections will replace $\hat{S}$ and $\hat{G}$ with a series of Pauli string rotations.

III. RESULTS

In this section, we summarize the techniques and findings of this work, and give an overview of its further organization. This work studies generalized unary encodings in quantum simulation with qubitization: for these encodings, the unitaries (2)–(4) are defined with a set of configurations $\{\mu_k\}$ that are linearly independent. We here eliminate Toffoli gates from the qubitization circuits of these encodings. This means in particular that we simplify $\hat{S}$, for which a Toffoli-type gate like in Fig. 3 is not necessary. Instead, we decompose the routines $\hat{G}$ and $\hat{S}$ into rotations of Pauli strings, such that the implementation of a qubitization routine becomes as straightforward as Trotterizing a time evolution. Pauli string rotations can be further decomposed into a sequence of two-qubit gates and single-qubit rotations with an up to logarithmic scaling in the number of qubits involved [25].

We will now give a high level overview of our results. The decomposition is achieved in two stages: first $\hat{G}$ and $\hat{S}$ are expressed as linear combinations of Pauli strings—the correctness of these decompositions is proven in Sec. V. Then, we review gadgets with which these linear combinations can be created. This is done in Sec. IV. The simplification of the qubitization circuit is however not the only benefit of the decomposition: in Sec. VI we introduce a particular encoding with $\mu_k \neq e_k$, for which we can implement $\hat{G}$ and $\hat{S}$ within $O(\log \Lambda)$ depth. Discussions about the qubit connectivity are shifted to the conclusion of this work (Sec. VII). For the remainder of this section, we will describe our findings analytically, and integrate them into the phase estimation routine before giving an overview of the entire protocol.

Before we start, however, we would like to introduce a few shorthands used throughout this paper: let us write the set of integers from 1 to N as $[N]$ and $[0] = \emptyset$. We define sums over an empty set to yield zero and products over an empty set to yield one. Also, we will index Pauli operators with the qubits they act on, e.g., $Z_1, X_2 \otimes Y_4$.

Rather than Toffoli gates, $\hat{G}$ and $\hat{S}$ are going to be built from a set of Pauli strings γ_j^x, γ_j^y , for all $j \in [\Lambda]$. All of these 2Λ Pauli strings anticommute pairwise—they are in fact qubit representations of Majorana fermions. For $\mu_k = e_k$, the strings are of the form

$$\begin{aligned} \gamma_j^x &= \left(\bigotimes_{k \in [j-1]} Z_k \right) \otimes X_j, \\ \gamma_j^y &= \left(\bigotimes_{k \in [j-1]} Z_k \right) \otimes Y_j, \end{aligned} \quad (6)$$

which would correspond to Jordan-Wigner transforms of their fermionic selves. Note, however, that these strings are

FIG. 4. Quantum phase estimation circuit using qubitization with a unary encoding. This circuit, estimating $\pm \arccos(H)$ up to a constant phase shift, features the subroutines g^y and g^x that can be parallelized to the runtime of $O(\log \Lambda)$.

only tools, which, living in the qubitization register, have no physical meaning in the simulated model. The strings are chosen such that they can create the basis states by $\gamma_k^x|\mathbf{0}\rangle = -i\gamma_k^y|\mathbf{0}\rangle = |\mu_k\rangle$, and their anticommutation properties are desired such that the linear combinations

$$g^x = \sum_{j \in [\Lambda]} \sqrt{\alpha_j} \gamma_j^x, \quad (7)$$

$$g^y = \sum_{j \in [\Lambda]} \sqrt{\alpha_j} \gamma_j^y \quad (8)$$

are unitary. The operators (7) and (8) can then be applied directly to the system through gadgets based on Pauli string rotations. This means we can construct the operators $\widehat{G}$ and $\widehat{S}$ as we find them to be made up by g^x and g^y :

$$\widehat{G} = g^x, \quad \widehat{S} = -ig^y g^x. \quad (9)$$

While $\widehat{G}$ can be prepared by a single instance of the aforementioned gadgets, the product within $\widehat{S}$ means that it is implemented by two consecutive gadgets. Not only can we build the linear combinations g^x and g^y , but also their rotations and conditional versions, e.g., $\exp(i\theta g^x) = \cos \theta + i \sin \theta \cdot g^x$ and $|0\rangle\langle 0|_{\text{est}} \otimes \mathbb{I} + |1\rangle\langle 1|_{\text{est}} \otimes g^y$, such that the gadgets can be integrated easily into the phase estimation routine.

Using the fine structure of $\widehat{S}$ and assuming we additionally apply the gadget for g^x (just like in the trick of [5], reclaiming the Hamiltonian eigenstates), we present the qubitization circuit featuring only gadgets for g^x and g^y in Fig. 4. One can easily verify that, up to a phase shift, this circuit acts as $\widehat{G}\widehat{S}\widehat{V}\widehat{G}$ in the $|1\rangle$ subspace of the estimator qubit and as the identity in its $|0\rangle$ subspace.

Protocol

The entire process of simulating a Hamiltonian (1) can be summarized with the following instructions.

- (1) Extract the coefficients $\{\alpha_k\}$ and Pauli strings $\{p_{\text{sim}}^k\}$ from the Hamiltonian.
- (2) Choose an encoding: $k \mapsto \mu_k$ and compute the sets of Pauli strings $\{\gamma_k^x\}$ and $\{\gamma_k^y\}$.
- (3) Decide the type of gadgets and obtain their parametric angles from the Hamiltonian coefficients: $\{\alpha_k\} \mapsto \{\phi_k\}$.
- (4) Compile the quantum circuits for the gadgets of g^y and g^x according to $\{\gamma_k^x\}$, $\{\gamma_k^y\}$, and $\{\phi_k\}$ as well as the entangling routine $\widehat{V}$ with respect to the basis $\{\mu_k\}$ and the set of Pauli strings $\{p_{\text{sim}}^k\}$.

- (5) Run the circuit in Fig. 4 on a quantum computer to estimate energy eigenvalues. Measure the qubitization register for a 50% chance to recover Hamiltonian eigenstates.

IV. TECHNIQUES

In this section we will present gadgets for the operators (7) and (8) that are therefore essential for a Toffoli-free decomposition of the operators $\widehat{G}$ and $\widehat{S}$. Rather than Toffoli gates, the decomposition relies on sequences of Pauli string rotations, rendering the hardware requirements basically identical to those of Trotterized time evolution, for which Pauli string rotations are also the elementary building blocks. In the most general use case, these gadgets receive a set of L pairwise-anticommuting Pauli strings $\{h^k\}$, as well as a set of L real coefficients $\{\beta_k\}$ and realize (controlled versions of) the operators

$$\exp\left(i \sum_{k=1}^L \beta_k h^k\right) = \cos \theta + i \sin \theta \left(\frac{1}{\theta} \sum_{k=1}^L \beta_k h^k\right), \quad (10)$$

where $\theta = (\sum_k |\beta_k|^2)^{1/2}$. The gadgets contain $2L - 1$ Pauli-string rotations featuring some angles $\phi_1, \phi_2, \dots, \phi_L$, which have to be matched with the coefficients $\{\beta_k\}$ in classical preprocessing. Here we consider two different gadgets that we will refer to as “symmetric” and “antisymmetric,” found in [6] and [7], yielding circuits of depth as short as $O(L)$ and $O(\log L)$, respectively. In both gadgets we identify relay points to control their application on an external qubit without adding a control to every gate. For the implementation of g^x or g^y , we find $L = \Lambda$, and $\{h^j\}$ equal to $\{\gamma_j^x\}$ or $\{\gamma_j^y\}$. However, while we equate these sets, we do not necessarily want to equate all individual elements with the same index, so $\{h^j\} = \{\gamma_j^x\}$ does not necessarily imply that for all j we have $h^j = \gamma_j^x$. Our purposes require $\theta = \pi/2$, so the set of parameters $\{\beta_k\}$ must be equal to $\{\sqrt{\frac{\pi}{2}} \alpha_k\}$.

A. Symmetric gadget

The first version of this gadget works in a very intuitive way. The symmetric gadget is based on second order Trotterization of the h strings. Since all those strings anticommute, this Trotter approximation does not yield (10) when setting $\phi_k = \beta_k$ for all $k \in [L]$. However, the result of the Trotterization is still quite predictable; in fact, arbitrary $\{\phi_k\}$ turn out to be the generalized Euler angles in the resulting superposition [26], i.e., when $\{\sqrt{\frac{\pi}{2}} \beta_k\}$ describe the Cartesian coordinates on an L -dimensional unit sphere, its spherical coordinates are $\{\phi_k\}$:

$$\begin{aligned} & \exp\left(\frac{i}{2} \phi_L h^L\right) \exp\left(\frac{i}{2} \phi_{L-1} h^{L-1}\right) \cdots \exp\left(\frac{i}{2} \phi_2 h^2\right) \exp(i \phi_1 h^1) \\ & \times \exp\left(\frac{i}{2} \phi_2 h^2\right) \cdots \exp\left(\frac{i}{2} \phi_{L-1} h^{L-1}\right) \exp\left(\frac{i}{2} \phi_L h^L\right) \\ & = \prod_{m \in [L]} \cos \phi_m + i \sum_{k \in [L]} h^k \sin \phi_k \prod_{j \in [k-1]} \cos \phi_j. \end{aligned} \quad (11)$$

One can now arrive at (10) by setting $\phi_1 = \arcsin(\frac{\beta_1}{\theta} \sin \theta)$ and then progress iteratively with

FIG. 5. Symmetric gadget for the linear combination of pairwise anticommuting strings $\{h^k\}$, controlled on a single qubit. The gates labeled $ih^k(\phi)$ denote rotations $\exp(ih^k\phi)$ around angles ϕ . (a) The main circuit featuring a h^1 rotation, in between the “right arm” and “left arm” subroutines. The control relay is built around the fact that h^1 (right arm) $h^1 = (\text{left arm})^\dagger$. (b) Right and left arm, which are sequences of rotations of the strings $\{h^k\}_{k=2}^L$. Comparing the two arms, the order of their rotations is reversed.

$\phi_{m+1} = \arcsin(\frac{\beta_{m+1}}{\beta_m} \tan(\phi_m))$. As shown in Fig. 5, the gadget can be switched off from an external control qubit by completing the innermost h^1 rotation to $\pi/2$. Canceling the right and left arm of the Trotterization, the circuit becomes trivial with one more controlled application of $(-ih^1)$ on the outside.

B. Antisymmetric gadget

While the symmetric gadget has a depth of $O(L)$, the antisymmetric one can achieve a depth of up to $O(\log L)$. We here present it in its utmost parallel version. The gadget consists of a h^1 rotation sandwiched by $L - 1$ rotations on one side and the same amount of counterrotations on the other; see Fig. 6. We are going to group those rotations into layers, where rotations inside the same layer commute. The first layer is a rotation of the string $h^1 h^2$ about the angle $\frac{1}{2}\phi_2$. Note that the product $h^1 h^2$ is just another Pauli string with imaginary coefficient. Embracing the h^1 rotation with the first layer and its inverse yields

$$\begin{aligned} & \exp(-\tfrac{1}{2}\phi_2 h^1 h^2) \exp(i\phi_1 h^1) \exp(\tfrac{1}{2}\phi_2 h^1 h^2) \\ &= \cos \phi_1 + i \sin \phi_1 h^1 \exp(\phi_2 h^1 h^2) \\ &= \cos \phi_1 + i \sin \phi_1 \cos \phi_2 h^1 + i \sin \phi_1 \sin \phi_2 h^2, \quad (12) \end{aligned}$$

where the second line results from the $(h^1 h^2)$ rotation being annihilated on the constant part of the inner rotation, but the counterrotation being flipped on the h^1 part, creating a conditional $(h^1 h^2)$ rotation around the angle ϕ_2 . The outcome is a superposition of the identity, ih^1 and ih^2 . The second layer will feature rotations of $h^1 h^3$ and $h^2 h^4$ around new angles $\frac{1}{2}\phi_3$ and $\frac{1}{2}\phi_4$. There, the counterrotation of $h^1 h^3$ will be reversed only upon encountering h^1 , and $h^2 h^4$ rotations add up to ϕ_4 when sandwiching h^2 . Note that the $h^1 h^3$ and $h^2 h^4$ rotations commute and can therefore potentially be executed at the same time. The result is again a superposition ih^k (for all $k \in [4]$) and $\mathbb{I}$. The procedure is repeated with every new layer employing rotations of $h^j h^k$ of all different j and k , where all j are drawn from the pool of strings already used in an inner layer, and h^k is a completely new string not yet in the superposition. The third layer would for instance feature rotations of $h^1 h^5$, $h^2 h^6$, $h^3 h^7$, and $h^4 h^8$. Since the pool of already used rotations doubles with every layer, we only need $\log L$ layers in total. A classical procedure matching the angles $\{\phi_k\}$ with the coefficients $\{\beta_k\}$ is sketched in the Appendix.

The resulting circuit has an easy relay point for the controlled application: as can be inspected in Fig. 6, rotations and counterrotations cancel each other when pulling the h^1 rotation from the center.

FIG. 6. Antisymmetric gadget to implement the linear combination of pairwise anticommuting strings controlled on a single qubit. The gate labeled $ih^1(\phi_1)$ denotes a rotation $\exp(ih^1\phi_1)$ and the gates $h^j h^k(\phi)$ with the argument ϕ signify rotations $\exp(h^j h^k \phi)$. (a) Main circuit featuring $\log L$ subroutines referred to as “layers” and their inverses, sandwiching the initial rotation around h^1 . The control relay is build around the fact that, when the initial rotation is removed, the remaining circuit acts trivially. (b) Layer subroutines featuring more and more rotations $h^j h^k$, that all pairwise commute.

V. PROOF

In this section, we will prove the results (9), for which we will first take a deeper look into the qubitization mechanism as defined by the unitaries (2)–(4) in [4].

We start with the state $|G\rangle \otimes |\varphi\rangle_{\text{sim}}$, where $|\varphi\rangle_{\text{sim}}$ is an ansatz for an energy eigenstate $|E\rangle_{\text{sim}}$. Let $-1 < E < 1$ be the corresponding eigenvalue $H|E\rangle_{\text{sim}} = E|E\rangle_{\text{sim}}$. Using $\hat{V} \cdot \hat{V} = \mathbb{I}$, the operator $\hat{V}$ spans a subspace for every state $|E\rangle_{\text{sim}}$:

$$|\tilde{E}\rangle = |G\rangle \otimes |E\rangle_{\text{sim}}, \quad (13)$$

$$|\tilde{E}^\perp\rangle = \frac{\hat{V} - E}{\sqrt{1-E^2}} |G\rangle \otimes |E\rangle_{\text{sim}}. \quad (14)$$

In the basis of $|\tilde{E}\rangle$ and $|\tilde{E}^\perp\rangle$, $\hat{V}$ has the matrix representation

$$\begin{bmatrix} E & \sqrt{1-E^2} \\ \sqrt{1-E^2} & -E \end{bmatrix}, \quad (15)$$

with which alone we cannot extract information about E yet, since it has eigenvalues ± 1 . However, the operator $\hat{S}$ acts on the basis of this subspace as $\begin{bmatrix} 1 & \\ & -1 \end{bmatrix}$, and so the sequence of $\hat{S}\hat{V}$ is represented by the matrix

$$\begin{bmatrix} E & \sqrt{1-E^2} \\ -\sqrt{1-E^2} & E \end{bmatrix}. \quad (16)$$

The eigenvalues of this matrix are $E \pm i\sqrt{1-E^2}$, and so we have completely described the qubitization routine. We now have to show that the same is true when $\hat{S}$ and $\hat{G}$ are defined as (9). In particular, we want to do this independently of the representation chosen for the strings γ_j^x , γ_j^y and therefore only assume $\gamma_k^x|0\rangle = -i\gamma_k^y|0\rangle = |\mu_k\rangle$. It immediately follows that $\hat{G}$ of (9) prepares the state $|G\rangle$, and so we only have to show that $\hat{S}$ of (9) fulfills $\hat{S}|\tilde{E}\rangle = |\tilde{E}\rangle$ and $\hat{S}|\tilde{E}^\perp\rangle = -|\tilde{E}^\perp\rangle$.

We first examine the product $g^y g^x = i\hat{S}$:

$$\begin{aligned} g^y g^x &= \sum_{m \in [\Lambda]} \alpha_m \gamma_m^y \gamma_m^x \\ &+ \frac{1}{2} \sum_{k \in [\Lambda]} \sum_{j \in [k-1]} \sqrt{\alpha_j \alpha_k} (\gamma_j^y \gamma_k^x + \gamma_k^y \gamma_j^x). \end{aligned} \quad (17)$$

Furthermore, using the anticommutation relations of the Majoranas we find $i\gamma_k^y \gamma_k^x |\mu_j\rangle = (-1)^{\delta_{jk}} |\mu_j\rangle$ and

$$\begin{aligned} & -\frac{i}{2} (\gamma_j^y \gamma_k^x + \gamma_k^y \gamma_j^x) |\mu_m\rangle \\ &= -\frac{i}{2} \gamma_j^y \gamma_k^x [1 - (-1)^{\delta_{mj} + \delta_{mk}}] |\mu_m\rangle \\ &= \begin{cases} |\mu_k\rangle & \text{if } m = j, \\ |\mu_j\rangle & \text{if } m = k, \\ 0 & \text{else.} \end{cases} \end{aligned} \quad (18)$$

Therefore, we can split the product $\hat{S}\hat{V}|\tilde{E}\rangle$ into two parts:

$$\hat{S}\hat{V}|\tilde{E}\rangle = -ig^y g^x \sum_k \sqrt{\alpha_k} |\mu_k\rangle \otimes p_{\text{sim}}^k |E\rangle_{\text{sim}} \quad (19)$$

$$\begin{aligned} &= - \sum_k \left[\sum_m \overbrace{(-1)^{\delta_{mk}} \alpha_m}^{1-2\alpha_k} \right] \sqrt{\alpha_k} |\mu_k\rangle \otimes p_{\text{sim}}^k |E\rangle_{\text{sim}} \\ &+ 2 \sum_k \left[\sum_{j \neq k} \underbrace{\sqrt{\alpha_j} |\mu_j\rangle}_{|G\rangle - \sqrt{\alpha_k} |\mu_k\rangle} \right] \otimes \alpha_k p_{\text{sim}}^k |E\rangle_{\text{sim}} \end{aligned} \quad (20)$$

$$= (2E - \hat{V})|\tilde{E}\rangle. \quad (21)$$

Plugging (21) into (14) and considering also that $\hat{S}|\tilde{E}\rangle = -ig^y|0\rangle \otimes |E\rangle_{\text{sim}} = |\tilde{E}\rangle$, we verify that $\hat{S}$ acts as Z operator onto the subspace spanned by $|\tilde{E}\rangle$ and $|\tilde{E}^\perp\rangle$.

VI. LOGARITHMIC DEPTH

In this section, we will change the representation of the Majorana strings, and so be able to show that g^y and g^x (and therefore $\hat{S}$ and $\hat{G}$) can be implemented in depth $O(\log \Lambda)$. Let us start by properly generalizing unary encodings.

Generalized unary encodings can be defined indirectly by considering the sets $F(j), U(j) \subseteq [\Lambda]$ for every $j \in [\Lambda]$, that we will refer to as flip and update sets of j , respectively. These sets are chosen such that $(\bigotimes_{l \in U(k)} X_l)$ and $(\bigotimes_{m \in F(j)} Z_m)$ only anticommute if $j = k$, and the γ strings are defined as

$$\begin{aligned} \gamma_k^x &= \left(\bigotimes_{l \in U(k)} X_l \right) \prod_{j \in [k-1]} \bigotimes_{m \in F(j)} Z_m, \\ \gamma_k^y &= i \left(\bigotimes_{l \in U(k)} X_l \right) \prod_{j \in [k-1]} \bigotimes_{m \in F(j)} Z_m. \end{aligned} \quad (22)$$

This is an idea originating from fermionic encodings [18,27]. In our case, we want the flip and update sets to be defined on a graph. Let us consider a tree graph with Λ nodes, that are uniquely labeled with the integers $[\Lambda]$ in a certain way. We define the following.

(i) $F(j)$ includes j and all the integers given to the direct children of node j in the tree.

(ii) $U(k)$ includes k and all the integers given to the ancestors of node k , including the root of the tree.

Regarding the labeling of the nodes we will at first assume Λ to be equal to $2^N - 1$ for some integer N , and consider a perfect binary tree as a graph. See Fig. 7(a) for an example with $\Lambda = 15$, in which two of the sets are highlighted. The next bigger tree with $\Lambda(N+1)$ nodes can always be obtained in taking two trees of size $\Lambda(N)$, adding the number $2^N - 1$ to the labels of each node in one of the trees, and then connecting the roots of both trees with a new node labeled $2^{N+1} - 1$. In this way, a tree of every size $\Lambda(N)$ can iteratively be reached from the one in Fig. 7(a). In case Λ is not a number $2^N - 1$ (for an integer N), the next biggest tree can be considered and superfluous nodes deleted.

Note that if the tree was a Fenwick tree, we would call the encoding a Bravyi-Kitaev transform [17–19], but we have made a different choice here to keep the number of children on each node constant. This is an important ingredient of the complexity proof that we will now commence.

FIG. 7. Binary-tree encoding of γ_j^x and γ_j^y . (a) Perfect binary tree with $\Lambda = 15$ nodes labeled systematically. We highlighted examples of flip and update sets: $F(14) = \{10, 13, 14\}$ and $U(6) = \{6, 7, 15\}$. (b) Clipping of a binary tree, showing only qubits (nodes) the products $\gamma_j^x \gamma_k^x$ and $\gamma_j^y \gamma_k^y$ have support on: the union of sets $F(j)$ and $F(k)$, with j and k highlighted.

We will show that, with (22), the linear combinations g^x and g^y can be implemented in $O(\log \Lambda)$ time using antisymmetric gadgets. In the case of g^x , the innermost rotation is about the string γ_Λ^x , which happens to be the root of the tree. From there on, pairs of layers are defined with the following scheme: each layer consist of rotations $\gamma_j^x \gamma_k^x$, with j being on the level last finished, and k being a direct child of j on an unfinished level. The first layer of every pair features all nodes j of the previous level and is followed by a layer of rotations $\gamma_j^x \gamma_\ell^x$, with the same j , but where ℓ is the respective other child of each j . After these two layers one level is completed and the procedure is repeated until the leaves are included. The implementation of g^y is analogous. Note that this procedure is slightly different from the one outlined in Sec. IV B, as here a layer only uses roughly half of the pool of strings already used in inner layers. The reason for this is that only those strings are ensured not to overlap: Bravyi-Kitaev-like encodings are defined such that most of the Z operators in (22) cancel and due to the product of $\gamma_j^x \gamma_k^x$ (respectively $\gamma_j^y \gamma_k^y$) the X strings cancel almost completely since $U(k) = U(j) \cup \{k\}$. The only qubits on which the product is supported are in the union of the flips sets $F(j)$ and $F(k)$; see Fig. 7(b) for an illustration. Since the rotations of the same layer have no overlap, they can be performed in parallel. Furthermore, as the Pauli weight of those operators is bounded, every layer is finished in constant time and since the number of layers is proportional to the number of levels, of which there are $O(\log \Lambda)$, the algorithmic depth exhibits the same scaling.

With $\mu_k \neq e_k$, $\hat{V}$ has to be adapted, as (5) is no longer valid. Since the basis of the binary-tree encoding has the form $\mu_k = \sum_{j \in U(k)} e_j \pmod 2$, we find that, in a binary-tree version of $\hat{V}$, all strings p_{sim}^k are applied conditioned on the joint parity of the qubits in the flip set $F(k)$. Although a bit longer, the depth of this operation can achieve the same scaling as (5).

VII. CONCLUSION

We have attempted to make the qubitization more elegant and tractable by eliminating higher-level gates in interpreting its unitaries as fermionic operators: for the preparation subroutine a delocalized fermion is created in an external register of qubits, and the reflection is implemented by the parity operator of that fermionic mode. For a logarithmic time complexity, no additional qubits are needed, but only a change in the representation of those fermions is required. However, to harness this time scaling without using swap or sorting networks, a tree-shaped qubit connectivity would be required. More realistically, one could employ different encodings for a scaling of at least $O(\sqrt{\Lambda})$ that is, however, resulting in local gates on a square lattice. In fact, it might be beneficial to use custom tree graphs which can be embedded into the qubit connectivity.

We are aware that all unary encodings are resource intensive; for problems defined on lattices, however, the qubit requirements could be feasible. For the simulation of different models, the gadgets for the linear combination of anticommuting strings could be used to relax the qubit requirements by banding anticommuting Hamiltonian terms together. In fact, this has been the original idea behind [6,7], where the authors set out to reduce the number of measurements in VQE experiments. Note, however, that we only expect substantial reductions in the number of terms when considering problems with highly delocalized fermions, as in those Hamiltonians we expect large sets of pairwise anticommuting strings. Alternatively, the resource requirements could be relaxed by combining the routine $\hat{V}$ with circuits that (nearly) diagonalize parts of a Hamiltonian [28–30]. The linear combination gadgets could also be used for the preparation of trial states in the simulator. For the benefit of both, qubitization and state preparation, the relationship between the presented gadgets and Givens rotations [31,32] should be studied further.

ACKNOWLEDGMENTS

We would like to thank T. E. O’Brien for some interesting discussions about that topic and C. W. J. Beenakker for his support. M.S. was supported by the Netherlands Organization for Scientific Research (NWO/OCW) and an ERC Synergy Grant. S.W. is supported by an NWO VIDI Grant, the NWO Zwaartekracht QSC, and an ERC Starting Grant.

APPENDIX: MATCHING ANGLES IN THE ANTISYMMETRIC GADGET

In this Appendix we outline the classical routine to obtain the set of angles $\{\phi_k\}$ such that the antisymmetric gadget in Sec. IV B outputs (10), given a set of parameters $\{\beta_k\}$. For this purpose we place all $h^j h^k$ rotations in a tree: the tree of the

FIG. 8. Matching the sets of angles $\{\phi_k\}$ to the parameters $\{\beta_k\}$ using a tree describing the circuit in Fig. 6, where each forking stands for a rotation (and counterrotation) where the edges are labeled with the corresponding coefficients. The edges, with which the angle ϕ_2 is inferred, are highlighted, since the ratio of the coefficients on the right and left flank must be equal to β_2/β_5 .

process described in Sec. IV B can be found in Fig. 8. The first level of the tree considers the initial rotation, and, from the inside out, every rotation and counterrotation of one of the operators $h^j h^k$ is denoted by a fork of the current node ih^j into new nodes ih^j and ih^k , where the edges are labeled with the corresponding rotation angles. When the gadget is completed, the tree has nodes labeled with all $\{ih^k\}$ for $k \in [L]$ and with $\mathbb{I}$. The coefficient of one of those operators is on the one hand found in (10) and, on the other hand, it can be read out from the tree: one needs to find the corresponding operator on the

leaf level, and then multiply all the sines and cosines found on edges that connect its ancestors with each other. For example, in the case of ih^3 , we need to match the coefficients

$$\frac{\sin \theta}{\theta} \beta_3 = \cos(\phi_7) \sin(\phi_3) \cos(\phi_2) \sin(\phi_1). \quad (\text{A1})$$

Considering a pair of operators, the ratio of coefficients in (10) must match the ratio of coefficients obtained from the tree, and so we have a means to infer $\{\phi_k\}$. We start by comparing the ratios of siblings on the leaf level: ϕ_5 , for instance, is inferred by considering the coefficient ratios of ih^5 and ih^1 ,

$$\frac{\beta_5}{\beta_1} = \tan \phi_5, \quad (\text{A2})$$

as all other sines and cosines involving the angles ϕ_3 , ϕ_2 , and ϕ_1 cancel. As we have inferred some angles $\{\phi_k\}$ we can consider the coefficient ratios of pairs with the first common ancestor one level lower in the tree and so, level by level, move towards the root inferring all angles. In Fig. 8, for instance, we have highlighted the paths used for the ratio of coefficients of ih^5 and ih^2 ,

$$\frac{\beta_2}{\beta_5} = \frac{\cos(\phi_6) \cos(\phi_4)}{\cos(\phi_5) \sin(\phi_3)} \tan \phi_2, \quad (\text{A3})$$

from which ϕ_2 can be inferred since, at that point, the angles ϕ_3 , ϕ_4 , ϕ_5 , and ϕ_6 are already known. The last angle to be matched is ϕ_1 , for which we could make use of the constant coefficient's ratio with ih^8 ,

$$\frac{\theta}{\beta_8} \tan \theta = \tan(\phi_1) \sin(\phi_2) \sin(\phi_4) \sin(\phi_8). \quad (\text{A4})$$

- [1] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [2] M. Suzuki, Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations, *Phys. Lett. A* **146**, 319 (1990).
- [3] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, *J. Math. Phys.* **32**, 400 (1991).
- [4] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [5] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, Quantum Algorithm for Spectral Measurement with a Lower Gate Count, *Phys. Rev. Lett.* **121**, 010501 (2018).
- [6] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskiy, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, *J. Chem. Theory Comput.* **16**, 190 (2020).
- [7] X. Bonet-Monroig, R. Babbush, and T. E. O'Brien, Nearly optimal measurement scheduling for partial tomography of quantum states, [arXiv:1908.05628](https://arxiv.org/abs/1908.05628).
- [8] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doberty, and M. Troyer, The Trotter step size required for accurate quantum simulation of quantum chemistry, *Quantum Inf. Comput.* **15**, 361 (2015).
- [9] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Gate-count estimates for performing quantum chemistry on small quantum computers, *Phys. Rev. A* **90**, 022305 (2014).
- [10] A. M. Childs, A. Ostrander, and Y. Su, Faster quantum simulation by randomization, *Quantum* **3**, 182 (2019).
- [11] I. D. Kivlichan, C. E. Granade, and N. Wiebe, Phase estimation with randomized Hamiltonians, [arXiv:1907.10070](https://arxiv.org/abs/1907.10070).
- [12] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [13] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proc. R. Soc. London, Ser. A: Math. Phys. Eng. Sci.* **454**, 339 (1998).
- [14] A. Yu. Kitaev, Quantum measurements and the Abelian stabilizer problem, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [15] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [16] 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).
- [17] S. B. Bravyi and A. Yu. Kitaev, Fermionic quantum computation, *Ann. Phys. (NY)* **298**, 210 (2002).

- [18] 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).
- [19] V. Havlíček, M. Troyer, and J. D. Whitfield, Operator locality in the quantum simulation of fermionic models, *Phys. Rev. A* **95**, 032332 (2017).
- [20] K. Setia, S. Bravyi, A. Mezzacapo, and J. D. Whitfield, Superfast encodings for fermionic quantum simulation, *Phys. Rev. Research* **1**, 033033 (2019).
- [21] M. Steudtner and S. Wehner, Quantum codes for quantum simulation of fermions on a square lattice of qubits, *Phys. Rev. A* **99**, 022308 (2019).
- [22] Z. Jiang, J. McClean, R. Babbush, and H. Neven, Majorana Loop Stabilizer Codes for Error Correction of Fermionic Quantum Simulations, *Phys. Rev. Applied* **12**, 064041 (2019).
- [23] N. P. D. Sawaya, T. Menke, T. Ha. Kyaw, S. Johri, A. Aspuru-Guzik, and G. G. Guerreschi, Resource-efficient digital quantum simulation of d -level systems for photonic, vibrational, and spin- s Hamiltonians, [arXiv:1909.12847](https://arxiv.org/abs/1909.12847).
- [24] A. N. Chowdhury, Y. Subasi, and R. D. Somma, Improved implementation of reflection operators, [arXiv:1803.02466](https://arxiv.org/abs/1803.02466).
- [25] F. Motzoi, M. P. Kaicher, and F. K. Wilhelm, Linear and Logarithmic Time Compositions of Quantum Many-body Operators, *Phys. Rev. Lett.* **119**, 160503 (2017).
- [26] D. K. Hoffman, R. C. Raffinetti, and K. Ruedenberg, Generalization of Euler angles to N -dimensional orthogonal matrices, *J. Math. Phys.* **13**, 528 (1972).
- [27] M. Steudtner and S. Wehner, Fermion-to-qubit mappings with varying resource requirements for quantum simulation, *New J. Phys.* **20**, 063010 (2018).
- [28] S. Wiesner, Simulations of many-body quantum systems by a quantum computer, [arXiv:quant-ph/9603028](https://arxiv.org/abs/quant-ph/9603028).
- [29] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K.-L. Chan, Low-depth Quantum Simulation of Materials, *Phys. Rev. X* **8**, 011044 (2018).
- [30] W. J. Huggins, J. McClean, N. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, [arXiv:1907.13117](https://arxiv.org/abs/1907.13117).
- [31] 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).
- [32] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).