

Preparation of many-body states for quantum simulation

Nicholas J. Ward, Ivan Kassal, and Alán Aspuru-Guzik^{a)}*Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts 02138, USA*

(Received 19 December 2008; accepted 16 March 2009; published online 18 May 2009)

While quantum computers are capable of simulating many quantum systems efficiently, the simulation algorithms must begin with the preparation of an appropriate initial state. We present a method for generating physically relevant quantum states on a lattice in real space. In particular, the present algorithm is able to prepare general pure and mixed many-particle states of any number of particles. It relies on a procedure for converting from a second-quantized state to its first-quantized counterpart. The algorithm is efficient in that it operates in time that is polynomial in all the essential descriptors of the system, the number of particles, the resolution of the lattice, and the inverse of the maximum final error. This scaling holds under the assumption that the wave function to be prepared is bounded or its indefinite integral is known and that the Fock operator of the system is efficiently simulatable. © 2009 American Institute of Physics. [DOI: 10.1063/1.3115177]

I. INTRODUCTION

Simulating quantum systems on a conventional computer requires resources that generally scale exponentially with the size of the system. Feynman¹ proposed to solve this problem using a quantum machine that would be able to mimic the properties of the quantum system. Subsequently, it has been demonstrated that quantum computers would be able to simulate the time-dependent Schrödinger equation for many systems of interest using resources that scale polynomially with the size of the system.²⁻⁹ However, all such simulations require the preparation of an appropriate initial state, which must be preparable to within a chosen error.

In this work, we focus on the preparation of states on a gate-model quantum computer. Our techniques can therefore complement those developed for the preparation of states in other models of quantum computation, such as adiabatic quantum computing.^{6,10}

In general, we will call a state on n qubits “efficiently preparable” if it can be prepared to within error ϵ using $\text{poly}(n, \epsilon^{-1})$ elementary (one- and two-qubit) quantum gates. Unfortunately, the efficiently preparable states form only a small subset of all quantum states. This is because a general state on n qubits contains 2^n amplitudes and therefore one needs $O(2^n)$ gates to prepare it.¹¹ Indeed, state-preparation algorithms are known that almost reach this lower bound.^{12,13}

In this work, we show that if wave functions are represented on a grid in real space, then most quantum states of physical interest are efficiently preparable. This is of interest because efficient, grid-based simulation algorithms are known for physically realistic systems.^{3-5,8}

Our work extends that of Zalka,⁴ who, in introducing real-space quantum simulation, also provided the first state preparation algorithm. However, his procedure is able to prepare only states of single particles or uncorrelated many-particle systems. We show how to use Zalka’s single-particle

wave functions as building blocks, permitting the preparation of general superpositions and mixed states of an arbitrary number of particles. Our approach is motivated by electronic-structure theory, in that we choose particularly convenient single-particle bases in which to expand more complicated states. We use the single-particle eigenstates to form Slater determinants (configurations), superpositions of which are used to express general many-particle states.

Our scheme is essentially a method for translating states in second quantization to the corresponding states in first quantization. This has two advantages. First, many useful states that might be needed in first quantization are easily prepared in second quantization.¹⁴ In particular, we can prepare eigenstates of operators if our scheme is combined with full configuration interaction (FCI),¹⁵ an exact diagonalization method. FCI is classically an exponentially hard problem due to the exponential growth of the number of configurations with system size, but it can be computed on a quantum computer in polynomial time.⁶ The quantum FCI operates in second quantization and can compute, for example, the ground state wave function of a molecular system. The second benefit of our method is that it is often easier to simulate time evolution in real space than in Fock space. For instance, every simulation in second quantization would require a separate set of basis-set-dependent operators and there might be some processes, such as ionization, which could not be adequately described using a small, localized basis set. In first quantization, however, all problems of chemical interest can be efficiently simulated by direct simulation of the molecular Hamiltonian in real space.⁸

This paper is organized as follows. We first consider the preparation of many-particle states in which all the particles are the same. There are three steps: the preparation of single-particle eigenstates in a chosen basis, the preparation of many-particle configurations, and finally the preparation of superpositions of configurations. We discuss the preparation of mixed states, after which we turn to systems with many different types of particles. Again, we consider the prepara-

^{a)}Electronic mail: aspuru@chemistry.harvard.edu.

tion of configurations, their superpositions, and mixed states. We close by showing that the algorithm is efficient in that its run time is polynomial in the size of the system, the number of qubits used to encode the wave function, and the inverse of the maximum allowed error.

II. ONE TYPE OF PARTICLE

Our algorithm translates from second to first quantization and therefore depends on the basis which is chosen for the representation of the second-quantized states. We require a finite orthonormal basis of functions $\{\phi_1, \dots, \phi_M\}$, which are the eigenstates of a known operator \hat{F} on an M -dimensional, single-particle Hilbert space. In our analogy with electronic-structure theory, \hat{F} would be the Fock operator for a single particle,¹⁵ and indeed we expect that the algorithm would be at its most useful if \hat{F} is chosen as the Fock operator of an actual system. Although the form of \hat{F} can be arbitrary, subject to a few restrictions below, we will take advantage of the analogy and refer to \hat{F} as the *Fock operator*. We will also say, for example, that two eigenstates of \hat{F} are *degenerate* if their *energies* (the corresponding eigenvalues) are the same.

To ensure that the overall state-preparation algorithm scales efficiently, we require that \hat{F} can be efficiently simulated on a quantum computer, i.e., that the simulation time scales polynomially with the size of the system. More precisely, if there are m particles occupying the M orbitals and the simulation is done on a grid of 2^l sites (see below), then, for any t and any $\varepsilon > 0$, there exist a unitary \hat{U} , composed of $\text{poly}(m, M, l, t, \varepsilon^{-1})$ elementary (one- and two-qubit) quantum gates, such that $\|\hat{U} - e^{-i\hat{F}t}\| \leq \varepsilon$. Intuitively, this means that given an initial state, the final state generated by the action of \hat{F} for time t can be calculated with reasonable effort and reasonable error.

Several classes of Hamiltonians are known to be efficiently simulatable and together they ensure that most physically relevant Fock operators will also be efficiently simulatable. Very generally, an operator can be efficiently simulated if its matrix in a given basis is sparse.^{10,16,17} In particular, this includes Hamiltonians that are sums of local operators, each of which acts on only a few degrees of freedom.^{2,11} In addition, many physically realistic real-space Hamiltonians (such as those for chemical systems) can be efficiently simulated.^{3-5,8}

We finally note that the requirement that the basis be orthonormal may exclude certain commonly used basis sets. Many of the usually encountered bases are appropriate, such as plane waves or molecular orbitals, which diagonalize the molecular Hartree–Fock Hamiltonian. However, nonorthogonal bases, such as Gaussian wavepackets or atomic orbitals on more than one atomic center, are not suitable for state preparation using our procedure.

A. Single-particle eigenstates

A single-particle basis function ϕ can be prepared on a grid by the state preparation method first proposed by Zalka⁴

and rediscovered independently by both Grover and Rudolph¹⁸ and Kaye and Mosca.¹⁹ The algorithm first prepares the absolute value of the function, followed by the addition of the phases. Specifically, given a register of l qubits, representing a grid of 2^l points, and a basis state $\phi(x)$ normalized over a length L , the algorithm first performs the transformation

$$|0\rangle \rightarrow |\phi\rangle = \sum_{x=0}^{2^l-1} \left| \phi\left(x\frac{L}{2^l}\right) \right| |x\rangle,$$

where each integer-valued state $|x\rangle$ is a position on the suitably scaled grid. This state is generated from the state $|000\dots\rangle$ by redistributing its amplitude l times across the eigenstates $|x\rangle$. To perform the redistribution correctly, we calculate the integrals

$$I_{i,k} = \frac{\int_{k(L/2^i)}^{(k+1)(L/2^i)} |\phi(x)|^2 dx}{\int_{k(L/2^i)}^{(k+2)(L/2^i)} |\phi(x)|^2 dx}, \quad (1)$$

for $k=0, \dots, 2^i-2$ and $i=1, \dots, l$. The fraction $I_{i,k}$ is simply the probability that a particle in the $(k+1)$ th subdivision of size $L/2^i$ is also in its left half. If the denominator in $I_{i,k}$ is zero, there is no amplitude to redistribute, so we can skip this step. The first split is realized by performing a rotation on the first qubit by $\arccos(\sqrt{I_{1,0}})$, corresponding to the transformation

$$|0, \dots\rangle \rightarrow \sqrt{I_{1,0}}|0, \dots\rangle + \sqrt{1-I_{1,0}}|1, \dots\rangle.$$

This splits the norm of the initial state so that the appropriate proportion is present on each half of the grid. The subsequent finer splits are carried out in superposition using controlled rotations on each qubit. For example, after the second iteration, the correct proportion of the norm is present in each quarter of the grid. After l iterations, one obtains the desired state. Note that adding a single qubit and the corresponding rotation doubles the precision of the grid. Consequently, the absolute value of the wave function can be efficiently approximated to any desired accuracy. Phases can be added where necessary through phase kickback.²⁰ Given a procedure that can transform $|x\rangle \rightarrow e^{i \arg \phi(x)} |x\rangle$, we can complete the preparation of $|\phi\rangle$ as

$$\begin{aligned} \sum_{x=0}^{2^l-1} \left| \phi\left(x\frac{L}{2^l}\right) \right| |x\rangle &\rightarrow \sum_{x=0}^{2^l-1} e^{i \arg \phi(xL/2^l)} \left| \phi\left(x\frac{L}{2^l}\right) \right| |x\rangle \\ &= \sum_{x=0}^{2^l-1} \phi\left(x\frac{L}{2^l}\right) |x\rangle = |\phi\rangle. \end{aligned}$$

The same algorithm can be straightforwardly generalized to a three-dimensional grid, where the position eigenstates are in Cartesian coordinates and the corresponding three-dimensional integrals are used. In addition, particle spin can be represented using additional qubits. A particle with spin S requires $\lceil \log_2(2S+1) \rceil$ qubits to store its z -projection m_S . In particular, there is a natural mapping between the spin of $\text{spin-}\frac{1}{2}$ particles and the states of a single qubit. If the Fock operator is spin-free, the eigenstates will have separable spatial and spin degrees of freedom, making the complete

single-particle state $|\phi\rangle|m_S\rangle$. Preparing the spin part of this wave function is relatively easy, for it suffices to initialize the spin register in an integer state. If, however, the eigenstate has correlation between the spatial and spin degrees of freedom, it can be prepared using the techniques in Secs. II D and III. That is, we treat the particle as if it were a composite system—composed of a spinless, spatial part and a spin—and prepare its eigenstate using the techniques below. In what follows, we will assume that our particles are fermions and we will note where the algorithm would need to be modified for bosons.

B. Computational complexity of integration

The preceding method for preparing single-particle states requires the evaluation of integrals [Eq. (1)]. Since this must be performed in superposition, the integrals must be computed on the quantum computer: precomputing them classically would require an exponentially large look-up table. Consequently, the computational complexity of the state preparation procedure will depend on the cost of computing the integrals.^{4,18} An integration procedure will, given a function $\phi: V \subset \mathbb{R}^d \rightarrow \mathbb{R}$ (where V is a bounded region), supply an estimate \tilde{I} of the integral $I = \int_V \phi(\mathbf{x}) d\mathbf{x}$ such that $|\tilde{I} - I| \leq \varepsilon_I$ with a certain fixed probability δ [we will call this condition the (ε_I, δ) absolute error].

Integrals can be evaluated either analytically or numerically. If the indefinite integrals of the basis functions are known, the definite integrals over any box on the Cartesian grid can be computed. The values of the indefinite integrals themselves can usually be computed efficiently (i.e., with polynomial cost in the desired accuracy) because they usually contain simple mathematical functions. In particular, the time it takes to retrieve n digits of any elementary function is a polynomial in n (Ref. 21) and likewise for compositions of elementary functions.

If the indefinite integrals are either unknown or impractical to compute, numerical techniques can be used. In particular, any classical numerical technique can, in principle, be implemented on a quantum computer. For example, computing \tilde{I} by Monte Carlo requires, in the worst case,²²

$$[(\Phi^{-1}(1 - \delta/2))^2 \sigma^2 / \varepsilon_I^2]$$

samples of ϕ for an (ε_I, δ) absolute error, where σ^2 is an estimate of the variance of ϕ over V and $\Phi(z) = (2\pi)^{-1/2} \int_{-\infty}^z \exp(-u^2/2) du$ is the standard normal cumulative distribution function. In particular, if ϕ is bounded so that $\phi_L \leq \phi(\mathbf{x}) \leq \phi_U$ for all $\mathbf{x} \in V$, the number of required samples is limited²² to

$$[(\Phi^{-1}(1 - \delta/2)(\phi_U - \phi_L) / 2\varepsilon_I)^2].$$

That is, Monte Carlo integration of any finite-variance function requires time that scales as $O(\varepsilon_I^{-2})$. Acceptable wave functions need not be continuous or even finite²³ (and hence may have infinite variance), but such examples are rather contrived and rarely encountered in practice (but see below for δ -functions).

Furthermore, it is known that quantum computers are able to offer a quadratic speed-up over conventional proba-

bilistic methods of integral evaluation. Quantum integration techniques^{24,25} rely on amplitude amplification²⁶ to achieve a computational complexity of $O(\varepsilon_I^{-1})$. This has been proven optimal by Nayak and Wu²⁷ and Novak.²⁸ These techniques have the same general applicability as classical Monte Carlo and will likewise succeed for any bounded function. Furthermore, the state preparation scheme of Soklakov and Schack,²⁹ which relies on amplitude amplification, also succeeds in $O(\varepsilon_I^{-1})$ time.

The preceding assumes that the function that we seek to prepare does not depend substantially on grid spacing. We would expect that of realistic wave functions, assuming that the grid spacing is smaller than the smallest wavelength of the system. A useful exception are Kronecker δ -functions, defined on a grid of 2^l points as $\phi(\mathbf{x}) = 2^{l/2} \delta_{\mathbf{x}, \mathbf{x}_0}$, where \mathbf{x}_0 is a constant vector. The variance of δ -functions grows exponentially in l and therefore they cannot be integrated efficiently by Monte Carlo or prepared efficiently using the method of Soklakov and Schack.²⁹ However, they can still be prepared efficiently using our techniques because their indefinite integral, the Heaviside function, can be easily computed in time independent of l .

It remains to be shown that an error in the evaluated integral translates to a comparable error in the prepared function. If the integrals (1) have a maximum error ε_I , that is, $|\tilde{I}_{i,k} - I_{i,k}| \leq \varepsilon_I$, and the error in the final prepared state $|\tilde{\phi}\rangle$ is $\varepsilon_\phi = 1 - |\langle \tilde{\phi} | \phi \rangle|$, we find that $\varepsilon_\phi \leq l\varepsilon_I/2$. In the case $l=1$, $|\phi\rangle = \sqrt{I}|0\rangle + \sqrt{1-I}|1\rangle$ and $|\tilde{\phi}\rangle = \sqrt{\tilde{I}}|0\rangle + \sqrt{1-\tilde{I}}|1\rangle$. Then, assuming $0 \leq \tilde{I} \leq 1$, which is necessary for $|\tilde{\phi}\rangle$ to be an acceptable state,

$$\varepsilon_\phi = 1 - \sqrt{\tilde{I}} - \sqrt{(1-I)(1-\tilde{I})} \leq 1 - \sqrt{1-\varepsilon_I} \approx \frac{\varepsilon_I}{2},$$

where we assumed that $\varepsilon_I \ll 1$. For larger l , a similar analysis applies qubit-wise: one finds that $\langle \tilde{\phi} | \phi \rangle \geq (1 - \varepsilon_I)^{l/2} \geq 1 - l\varepsilon_I/2$ (the last inequality holds for all ε_I if $l \geq 2$), whence $\varepsilon_\phi \leq l\varepsilon_I/2$. That is to say, the error in the prepared state grows only polynomially with the error in the evaluated integral, a fact that we will use later on to establish the computational cost of the state preparation algorithm.

C. Many-particle eigenstates

The next step is to use Zalka's algorithm to prepare multipartite configurations. That is, we wish to prepare the position-space representation of a second-quantization state $|n_1 n_2 \cdots n_M\rangle_{2\text{nd}}$ (a Fock eigenstate), where n_i is the occupation number of the basis orbital ϕ_i ($1 \leq i \leq M$). The position-space representation of $|n_1 n_2 \cdots n_M\rangle_{2\text{nd}}$ will be a Slater determinant of the occupied orbitals, and it will be an eigenstate of the many-body *Hartree-Fock Hamiltonian*,

$$\hat{H} = \sum_{i=1}^m \hat{F}_i, \quad (2)$$

where m is the total number of particles and \hat{F}_i is the Fock operator \hat{F} acting on the particle i .¹⁵

We assume that the state $|n_1 n_2 \cdots n_M\rangle_{2\text{nd}}$ has already been prepared by some previous algorithm. The M basis orbitals are occupied by m particles and we let j_1, \dots, j_m be the indices of the occupied orbitals. We therefore wish to perform the transformation

$$|n_1 n_2 \cdots n_M\rangle_{2\text{nd}} \otimes |0 \cdots\rangle_{1\text{st}} \rightarrow |n_1 n_2 \cdots n_M\rangle_{2\text{nd}} \otimes \frac{1}{\sqrt{m!}} \sum_{\sigma \in S_m} \text{sgn}(\sigma) |\sigma(\phi_{j_1} \phi_{j_2} \cdots \phi_{j_m})\rangle_{1\text{st}} \quad (3)$$

which takes the input state and prepares the appropriate first-quantized Slater determinant $|\Phi\rangle_{1\text{st}}$, a superposition of all the permutations on the m occupied orbitals (S_m is the symmetric group on m elements and sgn denotes the signature). Here $|0 \cdots\rangle_{1\text{st}}$ contains m registers for the m first-quantized occupied orbitals $|\phi_{j_1}\rangle_{1\text{st}}, \dots, |\phi_{j_m}\rangle_{1\text{st}}$. Note that Eq. (3) is not in general a reversible operation, as multiple input states would be mapped to the same antisymmetrized result. To ensure the algorithm is reversible, we additionally require³⁰ that $j_1 < j_2 < \cdots < j_m$. The procedure can be slightly modified if bosons are in question, then $\text{sgn}(\sigma)$ is to be omitted, and the j_i must satisfy $j_1 \leq j_2 \leq \cdots \leq j_m$.

The transformation (3) is accomplished in two steps. First, the occupied single-particle basis orbitals are each prepared in a separate register, forming a Hartree product,

$$|n_1 n_2 \cdots n_M\rangle_{2\text{nd}} \otimes |0 \cdots\rangle_{1\text{st}} \rightarrow |n_1 n_2 \cdots n_M\rangle_{2\text{nd}} \otimes |\phi_{j_1} \phi_{j_2} \cdots \phi_{j_m}\rangle_{1\text{st}}.$$

The procedure can be modified in the case of bosons by counting the occupation of each orbital and preparing that many copies in separate registers.

In the next step, the Hartree product is antisymmetrized, which produces the desired Slater determinant. To complete this step, we introduce an improved form of the antisymmetrization algorithm developed by Abrams and Lloyd.³⁰ The algorithm begins with the m wave functions $|\phi_{j_i}\rangle$ to be antisymmetrized in register A , and $m \lceil \log_2 m \rceil$ qubits in register B (where each grouping of $\log_2 m$ qubits constitutes a “quword”) initialized to $|0\rangle$. Using a series of controlled rotations, B is converted to the state

$$\frac{1}{\sqrt{m!}} \sum_{i=1}^m |i\rangle_{B[1]} \otimes \sum_{i=1}^{m-1} |i\rangle_{B[2]} \otimes \cdots \otimes |1\rangle_{B[m]},$$

which is a superposition of $m!$ unique states consisting of m quwords each and $B[i]$ denotes the i th quword in register B . Next we will transform this state into the superposition

$$\frac{1}{\sqrt{m!}} \sum_{\sigma \in S_m} |\sigma(1, \dots, m)\rangle_B$$

as follows. First let $B'[1]=B[1]$. Then assign to $B'[i]$ the $B[i]$ th natural number not present in the set $\{B'[1], B'[2], \dots, B'[i-1]\}$. This leaves the quantum computer in the state

$$\frac{1}{\sqrt{m!}} |\phi_{j_1} \phi_{j_2} \cdots \phi_{j_m}\rangle_A \otimes \sum_{\sigma \in S_m} |\sigma(1, \dots, m)\rangle_B. \quad (4)$$

Register B now contains a symmetrized state and this symmetry can be transferred to register A by sorting B while performing the same swaps on the wave functions in A . This yields the symmetrized state

$$\frac{1}{\sqrt{m!}} \sum_{\sigma \in S_m} |\sigma(\phi_{j_1} \phi_{j_2} \cdots \phi_{j_m})\rangle_A \otimes |1, 2, \dots, m\rangle_B, \quad (5)$$

which is what we would keep if we were interested in preparing bosonic states. To instead obtain an antisymmetrized state, we need only count the number of exchanges made in the sort, and reverse the sign of the wave function if it is odd. If we now eliminate register B , A contains the desired multiparticle state $|\Phi\rangle_{1\text{st}}$.

The original algorithm, introduced by Abrams and Lloyd, included an additional auxiliary register C , which would then be used as an intermediate for the sorting of A and B . We eliminate this step by sorting A and B together directly.

D. Superpositions

We now generalize the algorithm to the preparation of superpositions of many-particle states. Given a superposition of second-quantization states $|\mathbf{n}_i\rangle_{2\text{nd}} = |n_{1i} n_{2i} \cdots n_{Mi}\rangle_{2\text{nd}}$, with amplitudes α_i , we wish to perform the transformation

$$\left(\sum_i \alpha_i |\mathbf{n}_i\rangle_{2\text{nd}} \right) \otimes |0 \cdots\rangle_{1\text{st}} \rightarrow |0 \cdots\rangle_{2\text{nd}} \otimes \left(\sum_i \alpha_i |\Phi_i\rangle_{1\text{st}} \right).$$

The superposition on the left might come from a variety of sources. For example, an easily prepared equal superposition of Fock states would result in an equal superposition of real-space wave functions. Wang *et al.*¹⁴ provided an algorithm for preparing general superpositions of Fock states on a quantum computer. Alternatively, a quantum electronic-structure algorithm could be used to efficiently produce a physically relevant superposition. For example, an FCI algorithm could specify the ground state of a chemical or other many-body system in terms of a superposition of Fock states.⁶

As before, we begin by applying Zalka’s state preparation algorithm to the input state. Because this linear operation is carried out in superposition, it accomplishes the transformation

$$\left(\sum_i \alpha_i |\mathbf{n}_i\rangle_{2\text{nd}} \right) \otimes |0 \cdots\rangle_{1\text{st}} \rightarrow \sum_i \alpha_i |\mathbf{n}_i\rangle_{2\text{nd}} \otimes |\phi_{j_1 i} \phi_{j_2 i} \cdots \phi_{j_m i}\rangle_{1\text{st}}.$$

Note that the single-particle wave functions are now entangled with the input state. For a multiparticle eigenstate, the situation was different because the resulting state was separable. Hence, to separate the first-quantized wave functions from the second-quantized ones, we must “uncompute” the second-quantized states. This must be accomplished using only manipulations on the register containing the first-quantized wave functions $|\phi_i\rangle_{1\text{st}}$; if we can regenerate the

input state from the wave functions, the input register can be reset to $|0\rangle_{2\text{nd}}$ as desired. Given the one-to-one correspondence between a second-quantization state and the corresponding first quantization wave functions, regenerating the input state amounts to the problem of identifying the wave functions $|\phi_i\rangle_{1\text{st}}$ given only the information contained in their first-quantized representation.

For nondegenerate eigenstates, each state $|\phi_i\rangle_{1\text{st}}$ can be uniquely identified using its energy, which can be obtained through the phase estimation procedure.^{20,31,32} In general, given a unitary \hat{U} and its eigenstate $|\psi\rangle$, the phase estimation algorithm finds the eigenvalue of $|\psi\rangle$. Specifically, since $\hat{U}|\psi\rangle = e^{2\pi i\theta}|\psi\rangle$, we have $\hat{U}^k|\psi\rangle = e^{2\pi i k\theta}|\psi\rangle$. By controlled applications of the powers of \hat{U} to $|\psi\rangle$, controlled on the state $(1/2^{q/2})\sum_{k=0}^{2^q-1}|k\rangle$, one gets $(1/2^{q/2})\sum_{k=0}^{2^q-1}|k\rangle\hat{U}^k|\psi\rangle = (1/2^{q/2})\sum_{k=0}^{2^q-1}e^{2\pi i k\theta}|k\rangle|\psi\rangle$. An efficient quantum Fourier transform (QFT) on the control qubits will now yield the first q digits of the binary expansion of θ . If we choose \hat{U} such that $\hat{U}|\phi_i\rangle_{1\text{st}} = e^{2\pi i E_i}|\phi_i\rangle_{1\text{st}}$, we can use phase estimation with enough control qubits to obtain an approximation of the energies E_i . In particular, the natural choice $\hat{U} = e^{-i\hat{H}t}$, where \hat{H} is the Hartree–Fock Hamiltonian (2), supplies the appropriate unitary for a suitable choice of the time t . Note that \hat{U} can be simulated efficiently because \hat{H} is a sum of Fock operators, which are efficiently simulatable by assumption. The energy eigenvalues are stored in an additional register containing enough qubits to provide precision that distinguishes between nearby energies.

In the case that the spectrum of \hat{H} is degenerate, properties other than the energy of the states need to be used to distinguish them. If the degeneracy is caused by a symmetry of the Hamiltonian, the elements of the symmetry group can be used for this discrimination, as we outline in Sec. II E. If the degeneracies are accidental, other techniques are required, and we give some suggestions in Sec. II F. In addition, the techniques in Sec. II F can be used for distinguishing eigenvalues that are exponentially close and therefore cannot be distinguished efficiently by phase estimation.

Phase estimation using both \hat{U} to find energy eigenvalues and appropriate symmetry operations to distinguish degenerate states will provide us with a unique combination of eigenvalues for each state in the superposition. These eigenvalues can then be used (for example in conjunction with a look-up table) to uniquely identify the wave function and subtract 1 from the corresponding occupation number vector of the second-quantization state. Because this is done in superposition for every single-particle wave function, the input state is converted to $|0\rangle_{2\text{nd}}$. This accomplishes the total transformation

$$\left(\sum_i \alpha_i |\mathbf{n}_i\rangle_{2\text{nd}}\right) \otimes |0\cdots\rangle_{1\text{st}} \rightarrow |0\cdots\rangle_{2\text{nd}} \otimes \left(\sum_i \alpha_i |\phi_{j_1 i} \phi_{j_2 i} \cdots \phi_{j_m i}\rangle_{1\text{st}}\right), \quad (6)$$

which is a separable state. The antisymmetrization step can now proceed in superposition as usual, resulting in the final

state $|\Psi\rangle_{1\text{st}} = \sum_i \alpha_i |\Phi_i\rangle_{1\text{st}}$, as desired. This completes the state-preparation algorithm for a given superposition of multiparticle states.

E. Resolving degeneracies caused by symmetry

The procedure in Sec. II D assumes that it is possible to distinguish eigenstates based on their energy. If there are degenerate states, additional operations are required to distinguish them. Degeneracies in quantum states usually arise as a result of symmetry—degeneracies that do not are called “accidental” and we treat them separately in Sec. II F. For symmetry-caused degeneracy, distinguishing degenerate states requires an understanding of how they transform under the symmetry operations of the system. All of the wave functions $|\phi_i\rangle_{1\text{st}}$ are eigenstates of each symmetry operation within the point group, but degenerate wave functions will always have different eigenvalues for at least one of the operations. Phase estimation can still be used to obtain a unique set of eigenvalues, but in addition to finding the energies, we can distinguish the wave functions by symmetry. By applying phase estimation using an appropriate symmetry operation as the unitary operator, we obtain additional eigenvalues to distinguish degenerate states.

Because there are only a limited number of symmetries that are possible in physical systems, it will rarely be necessary to use more than a few readout qubits to retrieve all the distinguishing eigenvalues. With the exception of systems with spherical, cubic, or icosahedral symmetries, which we treat below, all systems in three-dimensional space have a symmetry point group all of whose irreducible representations are one- or two-dimensional.³³ Wave functions transforming as the one-dimensional irreducible representations are nondegenerate, while the ones transforming as the two-dimensional irreducible representations come in degenerate pairs. Distinguishing them, therefore, requires the determination of only one symmetry eigenvalue, which is different for the two wave functions.

This is most easily done in the case of point groups C_{nv} , $C_{\infty v}$, D_n , D_{nh} , $D_{\infty h}$, and D_{nd} , all of which contain a C_2 axis or a reflection plane that has character zero in all of the two-dimensional irreducible representations. In this case, one of the two degenerate wave functions is invariant under the reflection or C_2 rotation, while the other acquires a phase of -1 . To distinguish them, one would use the reflection or the C_2 rotation as the unitary of phase estimation with one readout qubit (note that these operations are easy to implement, being simple linear transformations). The readout qubit, initialized in the state $(|0\rangle + |1\rangle)/(\sqrt{2})$, would, under the action of the symmetry operation, be converted to $(|0\rangle \pm |1\rangle)/(\sqrt{2})$, depending on the acquired phase. A Hadamard gate would then return $|0\rangle$ or $|1\rangle$, perfectly discriminating between the two eigenfunctions.

Symmetry groups C_n , C_{nh} , and S_{2n} have, strictly speaking, only one-dimensional irreducible representations. However, there are pairs of representations that are complex conjugates of each other, meaning that the corresponding energy levels are degenerate due to time-reversal symmetry. These pairs of conjugate representations are called “separably

degenerate,"³⁴ and the corresponding wave functions can be distinguished using the principal symmetry axis C_n (or S_{2n} in the S_{2n} groups). In each case, under the action of C_n , one of the wave functions acquires a phase ω and the other ω^* , where $\omega = e^{2\pi i/n}$ (there are also cases where the pairs acquire phases such as $-\omega$ and $-\omega^*$, ω^2 and $(\omega^2)^*$, and so on, but these do not change the procedure outlined here). Phase estimation can, as usual, measure this phase up to a certain precision. However, since $1/n$ usually does not have a finite binary expansion, there will be an associated error in the phase estimation. This can be reduced below an arbitrary threshold by the addition of more readout qubits, as discussed in Sec. IV. This is especially true since real physical systems almost never have C_n axes with $n > 8$, meaning that only several qubits will be required for readout.

The cubic and icosahedral groups, T , T_h , T_d , O , O_h , I , and I_h , all have three-dimensional irreducible representations (and I and I_h also have four- and five-dimensional ones). Fortunately, there are plenty of reflection planes and C_2 axes which can be used for discrimination just as was done in the simpler groups above. Distinguishing three or four degenerate states requires two symmetry eigenvalue comparisons (and three in the case of fivefold degeneracy). Consequently, two readout qubits are required in these cases, one for each comparison (or three qubits in the fivefold degenerate case).

Degenerate states of spherically symmetric systems, such as atoms, can be distinguished by energy and by their angular momentum quantum numbers ℓ and m_ℓ . The maximally symmetric case is the $1/r$ potential, where the conservation of the Laplace–Runge–Lenz vector implies that all states with equal principal quantum number n are degenerate. If our basis contains states with $n \leq n_{\max}$, we would require $O(\log_2 n_{\max})$ qubits for the discrimination of the angular momentum states [that is, $O(\log_2 n_{\max})$ qubits each for ℓ and m_ℓ]. While circumstances where one encounters states of extremely high angular momentum are rare, we can see that the discrimination can be performed efficiently. The phase estimation in this case would use discrete rotations as its unitary operator. A similar approach was suggested by Zalka for the related problem of implementing unitary representations of $SU(2)$.³⁵

F. Resolving accidental degeneracies and exponentially close eigenstates

In Sec. II E, we outlined a procedure for distinguishing states that are degenerate because of symmetry. However, the eigenstates might also be accidentally degenerate or exponentially close in energy so that they cannot be efficiently distinguished by phase estimation. In those cases, it is not possible to distinguish between the (near-)degenerate states using the symmetry-based procedure.

One way around these problems is to transform to another basis where the (near-) degeneracy does not arise. A way of accomplishing this is to use a perturbed Fock operator $\hat{F}' = \hat{F} + \hat{V}$, where \hat{V} is a small, efficiently simulatable perturbation that breaks the (near-) degeneracies. In a finite basis, \hat{V} must also be small to ensure that the new basis can adequately describe the target state. The new eigenfunctions

are obtained from the old using perturbation theory, as are the new coefficients of the state that we wish to prepare. This change of basis can be done efficiently on a classical computer, before proceeding as normal with the state preparation algorithm. For the purposes of phase estimation, the new Fock operator can be efficiently simulated by operator splitting because both \hat{F} and \hat{V} are efficiently simulatable.

A drawback of this procedure is that the perturbation may destroy certain desirable symmetries of the system. In some cases, this can be avoided if we choose $\hat{V} \propto |\phi_i\rangle\langle\phi_i|$, where $|\phi_i\rangle$ is one of the (near-) degenerate eigenstates. In that case, \hat{F}' and \hat{F} would have the same eigenstates and no change of basis would be needed. Of course, it is possible that \hat{V} in this form is not efficiently simulatable, in which case this scheme would not be efficient.

G. Mixed states

The previous sections outline the procedure for preparing general pure states, which in the chosen basis read

$$|\Psi\rangle_{1st} = \sum_i \alpha_i |\Phi_i\rangle_{1st}. \quad (7)$$

From now on, we will drop the subscript 1st for clarity. We now wish to prepare a mixed state with density operator

$$\hat{\rho} = \sum_i p_i |\Psi_i\rangle\langle\Psi_i|,$$

where $|\Psi_i\rangle$ are arbitrary pure states of the form (7) and the probabilities p_i add up to 1. This scheme could be used for the preparation of thermal states, in which case one would choose $|\Psi_i\rangle$ to be the Hamiltonian eigenstates and $p_i = e^{-\beta E_i}/Z$, where $\beta = 1/k_B T$ and Z is the partition function. Our approach to the thermalization problem is therefore different from that of Terhal and DiVincenzo, who prepared thermal states by simulating an external bath.³⁶

We assume that each $|\Psi_i\rangle$ can be efficiently specified using some specification $|\xi_i\rangle$ (for example, $|\Psi_i\rangle$ is the ξ_i th eigenstate of the Hamiltonian). We begin by preparing the state $\sum_i \sqrt{p_i} |\xi_i\rangle$. This can be done using the procedure in Sec. II A if we order the ξ_i so that they may be thought of as a function on a one-dimensional grid. We then run the entire state-preparation algorithm in superposition, preparing the appropriate $|\Psi_i\rangle$ conditional on the value of the $|\xi_i\rangle$. This yields the state

$$|\Xi\rangle = \sum_i \sqrt{p_i} |\xi_i\rangle |\Psi_i\rangle,$$

the density operator of which is

$$\hat{\rho}_\Xi = \sum_{i,i'} \sqrt{p_i p_{i'}} |\xi_i\rangle\langle\xi_{i'}| \otimes |\Psi_i\rangle\langle\Psi_{i'}|.$$

Tracing out the specification register, we get the desired density operator

$$\hat{\rho} = \text{Tr}_\xi \hat{\rho}_\Xi = \sum_{i,i'} \sqrt{p_i p_{i'}} |\Psi_i\rangle\langle\Psi_{i'}| \text{Tr} |\xi_i\rangle\langle\xi_{i'}| = \sum_i p_i |\Psi_i\rangle\langle\Psi_i|.$$

In practical terms, tracing out the specification register amounts to doing nothing at all. That is, each $|\Psi_i\rangle$ is en-

tangled to a different $|\xi_i\rangle$, meaning that the $|\Psi_i\rangle$ evolve separately under time evolution, as they would if they were independent members of an ensemble.

One can notice that density operators diagonal in the $|\Phi_i\rangle$ basis can be prepared more directly. In the previous Sec. II D, we had to “disentangle” the first- and second-quantized states. If we had instead simply traced out the input register, we would have obtained a mixed state diagonal in the $|\Phi_i\rangle$ basis.

III. MANY TYPES OF PARTICLES

In Sec. II, we outlined an algorithm for the preparation of arbitrary many-particle states (pure or mixed) of a system of identical particles. However, one often wants to consider systems of more than one type of particle, or treat particles of the same kind, but separated in space, as different (the latter approach might be useful, for example, in computing electron transfer matrix elements for large molecules).³⁷ We consider the case of two types of particles, with the generalization to more types being clear.

One wants to prepare an arbitrary two-particle state

$$|\Theta\rangle = \sum_{i,j} \alpha_{i,j} |\Phi_{A,i}\rangle |\Phi_{B,j}\rangle,$$

where $|\Phi_{A,i}\rangle$ is a many-particle eigenstate of particles of type A and $|\Phi_{B,j}\rangle$ is an eigenstate of particles of type B . Each element $|\Phi_{A,i}\rangle |\Phi_{B,j}\rangle$ of this superposition is easily created by preparing the appropriate state in separate registers as was done in Sec. II C. Creating $|\Theta\rangle$ itself can be done in analogy to the preparation of superpositions in Sec. II D. We start by efficiently specifying $|\Theta\rangle$ using occupation number vectors of the $|\Phi_{A,i}\rangle$ and the $|\Phi_{B,i}\rangle$, namely,

$$\sum_{i,j} \alpha_{i,j} |\mathbf{n}_{A,i}\rangle_{2\text{nd}} |\mathbf{n}_{B,j}\rangle_{2\text{nd}} \otimes |0_A\rangle_{1\text{st}} |0_B\rangle_{1\text{st}}.$$

We then complete the state preparation, in superposition, as we did in Sec. II D, treating each register separately. Doing so produces $|\Theta\rangle$.

There are many circumstances in which the ability to prepare states such as these would be valuable. For instance, in chemical dynamics it is necessary to treat the nuclei and the electrons separately. If we restricted our state preparation to simple product states such as $|\Phi_{A,i}\rangle |\Phi_{B,j}\rangle$, we would get a state in the Born–Oppenheimer approximation, which is often a good approximation to the initial states of reactants participating in chemical reactions. However, as the procedure for preparing $|\Theta\rangle$ shows, quantum computers could just as easily prepare non-Born–Oppenheimer states in which there is correlation between electronic and nuclear degrees of freedom. Many-particle mixed states can likewise be prepared by following the procedure in Sec. II G separately for each type of particle.

IV. ERRORS AND THE COMPUTATIONAL COST

For the state preparation algorithm to be considered efficient, the time it takes to execute it must scale as a polynomial in the sizes of the input. More precisely, it should scale as a polynomial in l , the number of qubits used to store the

wave function and m , the number of occupied single-particle orbitals, which is the best descriptor of the total size of the system.

In this section, we first show that pre-existing errors are amplified at most linearly by subsequent steps of the algorithm. We then use this fact to obtain the total computational cost of preparing an arbitrary quantum state.

A. Errors

Assuming that the quantum gates are executed perfectly—or that the gate errors are corrected using efficient error correction algorithms—there are five sources of error in the state preparation algorithm.

1. Preparation of single-particle eigenstates

Zalka’s method that we adopted in Sec. II A requires evaluation of the integrals (1). We addressed the computational cost of integral evaluation in Sec. II B, where we show that the procedure can be accomplished in time polynomial in ε_l^{-1} if the wave function’s indefinite integral is known or, more generally, if the wave function is bounded. The resulting error in the prepared single-particle eigenstate is $\varepsilon_\phi \leq l\varepsilon_l/2$.

2. Assembly of many-particle eigenstates

Many-particle eigenstates (3) inherit the errors present in the single-particle eigenstates $|\phi_i\rangle$ that are used to assemble them. Supposing that the prepared states $|\tilde{\phi}_i\rangle$ approximate the true states $|\phi_i\rangle$ with error $\varepsilon_{\phi,i} = 1 - |\langle \tilde{\phi}_i | \phi_i \rangle|$, then the prepared Hartree product $|\tilde{\phi}_{j_1} \cdots \tilde{\phi}_{j_m}\rangle$ suffers an error

$$\begin{aligned} \varepsilon_\Phi &= 1 - |\langle \tilde{\phi}_{j_1} \cdots \tilde{\phi}_{j_m} | \phi_{j_1} \cdots \phi_{j_m} \rangle| \\ &= 1 - \prod_{i=1}^m (1 - \varepsilon_{\phi,j_i}) \leq \sum_{i=1}^m \varepsilon_{\phi,j_i} \leq m\varepsilon_\phi \leq ml\varepsilon_l/2, \end{aligned}$$

where $\varepsilon_\phi = \max \varepsilon_{\phi,i}$. Since the total error grows as a polynomial in both the single-state error and the number of occupied states, the assembly of Hartree products amplifies the pre-existing errors only linearly in m . The remaining step, the antisymmetrization of the Hartree product, does not introduce additional errors.

3. Preparation of superpositions

The parallel state preparation that is used to perform transformation (6) does not introduce any additional errors with the exception of the possible failures of phase estimation, discussed below. Nevertheless, we should see how pre-existing errors propagate through this step. If the prepared state is $|\tilde{\Psi}\rangle = \sum_i \alpha_i |\tilde{\Phi}_i\rangle$, we see that it suffers an error with respect to the target state

$$\begin{aligned} \varepsilon_\Psi &= 1 - \sum_i |\alpha_i|^2 |\langle \tilde{\Phi}_i | \Phi_i \rangle| \\ &= 1 - \sum_i |\alpha_i|^2 (1 - \varepsilon_{\Phi,i}) \leq \varepsilon_\Phi \leq ml\varepsilon_l/2, \end{aligned}$$

where $\varepsilon_\Phi = \max \varepsilon_{\Phi,i}$ and we assumed that $\langle \tilde{\Phi}_i | \Phi_j \rangle = 0$ for

$i \neq j$. In other words, the error in $|\tilde{\Psi}\rangle$ is limited by the error of its components.

4. Discrimination of states in a superposition

The preparation of superpositions described in Secs. II D–II F and III relies on phase estimation as a means of distinguishing states. Since the eigenenergies will rarely have finite binary expansions, there will be errors introduced at this step. If two phases differ at the n th bit and we perform phase estimation with $q=n+p$ qubits, the probability of an incorrect identification is $1/2(2^p-2)$, meaning that the success probability will be $1-\varepsilon_{\text{PE}}$ provided we implement phase estimation with $p=\lceil \log(2+1/2\varepsilon_{\text{PE}}) \rceil$ additional qubits.¹¹ The additional overhead, logarithmic in $\varepsilon_{\text{PE}}^{-1}$, does not compromise the efficiency. The same arguments apply to the phase estimation of eigenvalues of C_n belonging to states in separably degenerate irreducible representations of groups C_n , C_{nh} , and S_{2n} (see Sec. II E). The symmetry eigenvalues that are useful for states in the other point groups are always ± 1 and can be perfectly resolved using phase estimation with a single readout qubit.

In addition, failures of state discrimination based on phase estimation can be detected after the state preparation is complete. The second-quantized register, which should be uncomputed during the procedure, should be measured at the end. If $|0 \cdots 0\rangle$ is observed, phase estimation will have succeeded. Otherwise, a misidentification will have occurred, and the procedure ought to be repeated. This simple, classical error correction introduces only a constant overhead and ensures that phase estimation does not contribute to the error in the final prepared state.

5. Assembly of mixed states

In the notation of Sec. II G, if the prepared states $|\tilde{\Psi}_i\rangle$ approximate the true states $|\Psi_i\rangle$ with an error $\varepsilon_{\Psi,i}=1-|\langle \tilde{\Psi}_i | \Psi_i \rangle|$, and assuming perfect preparation of the state $\sum_i \sqrt{p_i} |\xi_i\rangle$, the final prepared mixed state will be $\hat{\rho}=\sum_i p_i |\tilde{\Psi}_i\rangle \langle \tilde{\Psi}_i|$. If we assume that $\langle \tilde{\Psi}_i | \Psi_j \rangle=0$ for $i \neq j$, then $\hat{\rho}$ suffers an error¹¹

$$\begin{aligned} \varepsilon_\rho &= 1 - \text{Tr} \left(\sqrt{\hat{\rho}} \hat{\rho} \sqrt{\hat{\rho}} \right)^{1/2} \\ &= 1 - \text{Tr} \left(\sum_i p_i^2 |\tilde{\Psi}_i\rangle \langle \tilde{\Psi}_i| \langle \tilde{\Psi}_i | \Psi_i \rangle \langle \Psi_i | \tilde{\Psi}_i \rangle \langle \tilde{\Psi}_i | \right)^{1/2} \\ &= 1 - \sum_i p_i (1 - \varepsilon_{\Psi,i}) \leq \varepsilon_\Psi \leq ml\varepsilon_l/2, \end{aligned}$$

where $\varepsilon_\Psi = \max \varepsilon_{\Psi,i}$. That is, the assembly of mixed states does not magnify the pre-existing errors.

Overall, we see that errors introduced in any stage of the state preparation algorithm are not amplified more than polynomially by subsequent stages. The final error in the prepared state is $\varepsilon = \varepsilon_\rho \leq ml\varepsilon_l/2$, meaning that the error scales linearly with the size of the system m and the error of the integration procedure, as well as logarithmically with the grid size 2^l .

B. Computational cost

There are three time-consuming steps in the state preparation algorithm. The first is the evaluation of the integrals (1) and the resulting single-qubit rotations, the second is the phase estimation used to distinguish states in the superposition (see Sec. II D), and the final is the antisymmetrization procedure described in Sec. II C. We characterize the cost of each step in turn.

In the previous section, we have seen that the total error of the prepared state will be $\varepsilon \leq ml\varepsilon_l/2$. Therefore, if we want to ensure a maximum error ε , we must choose $\varepsilon_l = 2\varepsilon/ml$, implying that $O(ml\varepsilon^{-1})$ time is required for each integration (see Sec. II B). The integration procedure itself is called ml times: for each of the m occupied orbitals, l qubits have to be rotated correctly. Therefore, the total time required for all the qubit rotations is $O(m^2l^2\varepsilon^{-1})$.

The cost of the phase-estimation procedure that is used to distinguish the eigenstates cannot be given precisely because we have not made any assumptions about the nature of the Fock operator \hat{F} other than that it is efficiently simulatable, that is, running in time $\text{poly}(m, M, l, \Delta^{-1})$ (here, Δ is the precision at which the simulation needs to be run, i.e., it is half the gap between the closest two eigenstates, which we assumed is not exponentially small). Simulating the entire Hartree–Fock Hamiltonian requires the simulation of the Fock operator acting separately on each particle, meaning that the total simulation requires $m \text{poly}(m, M, l, \Delta^{-1})$ time. In addition to this, two QFTs are required on the readout register of the phase estimation. If q qubits are used for the readout (see Sec. IV A), the QFTs require $O(q^2)$ time. It should be noted that the required q is determined only by needed precision in the phase estimation and that it does not depend strongly on m , l , or M . Therefore, the cost of the QFTs can be treated as essentially a constant overhead. Furthermore, there is the cost of looking up the state's energy in the look-up table; a simple binary search requires $O(\log_2 M)$ time per register, for a total cost of $O(m \log_2 M)$. However this, too, is a negligible cost in comparison to $m \text{poly}(m, M, l, \Delta^{-1})$, which we conclude is the asymptotic cost of the eigenstate discrimination portion of the state preparation algorithm.

The bottleneck of the antisymmetrization procedure used to produce fermionic states (or the symmetrization for bosonic ones) is the sort that takes state (4) to state (5). Sorting register B by a comparison sort requires $\Omega(m \log m)$ swaps. These swaps must also be performed on each of the corresponding l qubits of register A , for a total cost of $\Omega(lm \log m)$. For large systems, this expression will be dominated by the scalings of the integral evaluation and the phase estimation.

Based on the foregoing, the total computational cost of the state preparation algorithm is $O(m^2l^2\varepsilon^{-1} + m \text{poly}(m, M, l, \Delta^{-1})) = \text{poly}(m, M, l, \varepsilon^{-1}, \Delta^{-1})$, an expression polynomial in all the basic descriptors of the system. This allows us to conclude that the algorithm, as described above, is efficient.

V. CONCLUSION

We outlined a quantum algorithm for the preparation of physically realistic quantum states on a lattice. In particular, we have gone beyond previous proposals by describing a method for preparing any pure or mixed state of any number of particles. This is achieved by using Zalka's method for preparing single-particle states and then combining those into many-particle states. The assembly of many-particle states requires that we be able to distinguish them on a quantum computer, a task that we address using phase estimation. We also provided symmetry-based solutions for degenerate cases, where phase estimation using a single operator is insufficient to distinguish the states. Accidentally degenerate states can be distinguished by adding a perturbation to the system Hamiltonian. Our algorithm is efficient, with a run time of $\text{poly}(m, M, l, \epsilon^{-1}, \Delta^{-1})$, subject only to the requirements that the wave function be bounded or that its indefinite integral be known and that the Fock operator be efficiently simulatable.

ACKNOWLEDGMENTS

We acknowledge support from the Army Research Office under Contract No. W911NF-07-0304. N.J.W. thanks the Harvard College Research Program and I.K. thanks the Joyce and Zlatko Baloković scholarship.

¹R. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).

²S. Lloyd, *Science* **273**, 1073 (1996).

³S. Wiesner, e-print arXiv:quant-ph/9603028, 1996.

⁴C. Zalka, *Proc. R. Soc. London, Ser. A* **454**, 313 (1998).

⁵D. A. Lidar and H. Wang, *Phys. Rev. E* **59**, 2429 (1999).

⁶A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).

⁷H. Wang, S. Kais, A. Aspuru-Guzik, and M. R. Hoffmann, *Phys. Chem. Chem. Phys.* **10**, 5388 (2008).

⁸I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 18681 (2008).

⁹B. P. Lanyon, J. D. Whitfield, G. G. Gillet, M. E. Goggin, M. P. Almeida,

I. Kassal, M. Mohseni, J. D. Biamonte, B. J. Powell, M. Barbieri *et al.* (submitted).

¹⁰D. Aharonov and A. Ta-Shma, *SIAM J. Comput.* **37**, 47 (2007).

¹¹M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

¹²M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, *Quantum Inf. Comput.* **5**, 467 (2005).

¹³V. Bergholm, J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, *Phys. Rev. A* **71**, 052330 (2005).

¹⁴H. Wang, S. Ashhab, and F. Nori, e-print arXiv:0902.1419, 2009.

¹⁵A. Szabo and N. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover, New York, 1996).

¹⁶D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Commun. Math. Phys.* **270**, 359 (2007).

¹⁷A. Childs, Ph.D. thesis, Massachusetts Institute of Technology, 2004.

¹⁸L. Grover and T. Rudolph, e-print arXiv:quant-ph/0208112, 2002.

¹⁹P. Kaye and M. Mosca, e-print arXiv:quant-ph/0407102, 2004.

²⁰R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, *Proc. R. Soc. London, Ser. A* **454**, 339 (1998).

²¹J. M. Borwein and P. B. Borwein, *Pi and the AGM: A Study in Analytic Number Theory and Computational Complexity* (Wiley-Interscience, New York, 1998).

²²G. S. Fishman, *Monte Carlo: Concepts, Algorithms, and Applications* (Springer, New York, 1995).

²³A. Peres, *Quantum Theory: Concepts and Methods* (Springer, New York, 1995).

²⁴L. K. Grover, *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing* (ACM, New York, 1998), pp. 53–62.

²⁵D. S. Abrams and C. P. Williams, e-print arXiv:quant-ph/9908083, 1999.

²⁶G. Brassard, P. Høyer, M. Mosca, and A. Tapp, e-print arXiv:quant-ph/0005055, 2000.

²⁷A. Nayak and F. Wu, *Proceedings of the 31st Annual ACM Symposium on the Theory of Computing* (ACM, New York, 1999), pp. 384–393.

²⁸E. Novak, *J. Complex.* **17**, 2 (2001).

²⁹A. N. Soklakov and R. Schack, *Phys. Rev. A* **73**, 012307 (2006).

³⁰D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **79**, 2586 (1997).

³¹A. Y. Kitaev, e-print arXiv:quant-ph/9511026, 1995.

³²D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **83**, 5162 (1999).

³³F. A. Cotton, *Chemical Applications of Group Theory*, 3rd ed. (Wiley, New York, 1990).

³⁴P. R. Bunker and P. Jensen, *Molecular Symmetry and Spectroscopy*, 2nd ed. (NRC Research, Ottawa, 2006).

³⁵C. Zalka, e-print arXiv:quant-ph/0407140, 2004.

³⁶B. M. Terhal and D. P. DiVincenzo, *Phys. Rev. A* **61**, 022301 (2000).

³⁷L. Y. Zhang, R. A. Friesner, and R. B. Murphy, *J. Chem. Phys.* **107**, 450 (1997).