

Eigenvalue estimation of differential operators with a quantum algorithm

Thomas Szkopek,* Vwani Roychowdhury, and Eli Yablonovitch

Department of Electrical Engineering, University of California Los Angeles, Los Angeles, California 90095, USA

Daniel S. Abrams

Luminescent Technologies, Inc., Mountain View, California 94041, USA

(Received 6 January 2005; revised manuscript received 30 June 2005; published 13 December 2005)

We demonstrate how linear differential operators could be emulated by a quantum processor, should one ever be built, using the Abrams-Lloyd algorithm. Given a linear differential operator of order $2S$, acting on functions $\psi(x_1, x_2, \dots, x_D)$ with D arguments, the computational cost required to estimate a low order eigenvalue to accuracy $\Theta(1/N^2)$ is $\Theta((2(S+1)(1+1/\nu)+D)\ln N)$ qubits and $O(N^{2(S+1)(1+1/\nu)} \ln^c N^D)$ gate operations, where N is the number of points to which each argument is discretized, ν and c are implementation dependent constants of $O(1)$. Optimal classical methods require $\Theta(N^D)$ bits and $\Omega(N^D)$ gate operations to perform the same eigenvalue estimation. The Abrams-Lloyd algorithm thereby leads to exponential reduction in memory and polynomial reduction in gate operations, provided the domain has sufficiently large dimension $D > 2(S+1)(1+1/\nu)$. In the case of Schrödinger's equation, ground state energy estimation of two or more particles can in principle be performed with fewer quantum mechanical gates than classical gates.

DOI: [10.1103/PhysRevA.72.062318](https://doi.org/10.1103/PhysRevA.72.062318)

PACS number(s): 03.67.Lx, 02.60.Lj

I. INTRODUCTION

An early motivation for research in quantum information processing has been the simulation of quantum mechanical systems [1]. The Abrams-Lloyd algorithm [2–4] is an instance of quantum mechanical simulation (followed by variations [5–7]). We describe in this paper the application of the Abrams-Lloyd algorithm to estimating low order eigenvalues of linear partial differential equations with homogeneous boundary conditions (more precisely, Hermitian boundary value problems). The significance of our analysis is twofold. First, we generalize the Abrams-Lloyd algorithm to boundary value problems other than Schrödinger's equation, which may find application to classical problems. Second, we quantify computational cost and determine under what conditions we may expect the Abrams-Lloyd algorithm to give a reduction in computational work compared to optimal classical techniques in order to achieve the same eigenvalue accuracy.

Very briefly, the Abrams-Lloyd algorithm as originally envisaged for the many-body Schrödinger equation is structured as follows. An initial estimate $|\psi(0)\rangle$ of the target eigenstate is loaded into a multiple qubit register. Controlled application of a unitary operation, chosen to correspond to the time evolution operator $\exp(-i\mathcal{H}\tau)$ of the many-body Hamiltonian \mathcal{H} under study for time step τ , allows one to generate a sequence of time evolved states originating from the initial guess, $\{|\psi(0)\rangle, |\psi(\tau)\rangle, |\psi(2\tau)\rangle, \dots\}$. A spectral analysis of the sequence of time evolved states recovers the frequency (energy) of the target eigenstate (provided the initial guess was “close enough”). The Abrams-Lloyd algorithm is akin to a stroboscope for quantum states evolving under a many-body Hamiltonian. If the total time of evolution is sufficiently long, while the individual time steps are sufficiently

small, a high frequency (energy) resolution can be achieved. Following the determination of the eigenvalue, the corresponding eigenstate remains in the qubit register. Although the full amplitude description of an eigenstate is inaccessible, some information about the state can be extracted to a precision ultimately limited by the number of qubits used to represent the eigenstate (so for instance, one can test symmetries of the eigenstate).

The algorithm can be extended to more general partial differential equations rather easily. So long as the boundary value problem is Hermitian, we can map our mathematical problem to a fictional quantum system and apply the algorithm without change. The partial differential operator, \mathcal{L} , will correspond to a (possibly) fictional Hamiltonian \mathcal{H} , and an initial guess $\psi(0)$ will correspond to an initial wave function $|\psi(0)\rangle$. In other words, quantum mechanical amplitudes represent function values. Less obviously, the sequence of time evolved states, $\{|\psi(0)\rangle, |\psi(\tau)\rangle, |\psi(2\tau)\rangle, \dots\}$ has a mathematical analogue of great use in classical matrix eigenvalue analysis, known as the Krylov subspace, $\text{span}\{\psi(0), \exp(-i\mathcal{L}\tau)\psi(0), \exp(-i2\mathcal{L}\tau)\psi(0), \dots\}$. The subspace is generated by repeated application of $\exp(-i\mathcal{L}\tau)$, although in classical techniques one more typically uses rational functions of \mathcal{L} . Here, τ no longer has the physical meaning of time. Rather, τ sets the scale for how much phase is applied per application of $\exp(-i\mathcal{L}\tau)$. As in the quantum simulation, a large total phase applied one small phase step at a time allows a high resolution estimate of eigenvalues. We quantify these notions now.

The computational cost of the Abrams-Lloyd algorithm for a specified eigenvalue accuracy is limited as a consequence of three sources of error, expressed here in the language of quantum mechanical simulation.

(i) *Truncation error*: Discretization is necessary for a computational model based on qubits. However, discretization of the continuous problem to N points per coordinate

*Electronic address: szkopek@ee.ucla.edu

results in $\Theta(1/N^2)$ relative error in low order energy eigenvalues due to truncation of high spatial frequency contributions. The choice of N must be made appropriate to the accuracy that is desired.

(ii) *Splitting error*: The full many-body evolution $\exp(-i\mathcal{H}\tau)$ over time step τ can be implemented with universal gates by splitting the full evolution into a sequence of efficiently implementable unitaries $\exp(-i\mathcal{H}^k\tau)$, where $\mathcal{H} = \sum_k \mathcal{H}^k$. The approximation results in an absolute eigenvalue error $O(\|\mathcal{H}\|_2^{\nu+1}\tau^\nu)$, where $\|\mathcal{H}\|_2$ is the maximum eigenvalue of the discretized Hamiltonian \mathcal{H} and ν is a constant of $O(1)$ determined by the precise sequence of local operators chosen. Splitting error requires us to use small time steps τ .

(iii) *Frequency resolution*: A quantum Fourier transform, like any discrete Fourier transform, can resolve absolute phase to accuracy at best $\pm\pi$. For a sequence of M samples, the relative error in an energy eigenvalue E will be $\pm\pi/(ME\tau)$. Frequency resolution requires us to simulate over a large total time $M\tau$.

The optimal way to balance these errors is as follows. Since we are interested in the continuous problem, we first choose a discretization of N points per sample so that the discrete problem eigenvalue approximates the continuous eigenvalue problem to some desired accuracy $\Theta(1/N^2)$. We wish to solve the discretized problem to an accuracy determined by the truncation error limit; solving the discrete problem to greater accuracy leads to wasted effort since we are interested in the continuous problem, while solving the discrete problem to lesser accuracy implies we have wasted effort by choosing too many discrete points N per coordinate. We can thereby determine the maximum time step τ to keep splitting error no greater than truncation error. Next, we can determine the number of time steps M required to resolve eigenvalues with the quantum Fourier transform at the truncation error limit. In the case of Hermitian boundary value problems, we show in this paper that the resulting computational cost is $\Theta(D \ln N)$ qubits and $O(N^{2(S+1)(1+1/\nu)} \ln^c N)$ gate operations, where $2S$ is the differential order of \mathcal{L} and c is a constant $O(1)$. This can be compared with the optimal classical cost of $\Theta(N^D)$ bits and $\Omega(N^D)$ gate operations. Near optimal classical methods approaching these costs do in fact exist ¹.

We emphasize that in our analysis, we take a constructive approach wherein we account for all the logical operations required to implement the algorithm without recourse to oracles that may or may not have physically efficient implementations. This is in contrast to previous work including the simulation of spin glass physics [8], and Sturm-Liouville problems (Ref. [9] and references therein). As stated, our motivation is to compare the computational cost of eigenvalue estimation by the Abrams-Lloyd algorithm and optimal classical methods.

Our paper is organized as follows. In Sec. II, we introduce the one-dimensional eigenvalue problem, which will serve as

¹A near optimal classical method can be constructed using a combination of Krylov subspace iteration, matrix preconditioning and multigrid solutions, as in Ref. [20]. See Refs. [13,19] for a sampling of the vast array of classical numerical techniques available.

a useful example with which the principles of the algorithm can be illustrated. We derive the truncation error for low order eigenvalues in a way suitable for extension to higher dimensional problems. The algorithm itself is described in Sec. III, followed by an analysis of computational cost as it is applied to the one-dimensional problem in Sec. IV. A circuit suitable for a second order differential equation is given as a concrete example. Generalization of the algorithm to higher dimensional problems is given in Sec. V along with an analysis of computational cost, where we show a reduction in computational work polynomial in N over classical techniques. Concluding remarks about the computational efficiency of the Abrams-Lloyd algorithm are given in Sec. VI.

II. ONE-DIMENSIONAL PROBLEM

To illustrate the essential features of eigenvalue estimation of differential operators, it is instructive to consider a Hermitian one-dimensional problem, which we introduce here in some detail. The primary result of this section is a derivation of truncation error in low order eigenvalues as a function of discretization. Much of the notation used throughout this paper are defined in this section. We begin with a linear, $2S$ -order differential operator \mathcal{D} that maps a complex valued function $\psi(x)$, $x \in [0, 1]$ to a function according to the rule,

$$\mathcal{D}\psi(x) = \sum_{s=0}^S \frac{\partial^s}{\partial x^s} \left(a_s(x) \frac{\partial^s \psi(x)}{\partial x^s} \right) = a_0(x)\psi(x) + \frac{\partial}{\partial x} \left(a_1(x) \frac{\partial \psi(x)}{\partial x} \right) + \cdots + \frac{\partial^S}{\partial x^S} \left(a_S \frac{\partial^S \psi(x)}{\partial x^S} \right), \quad (1)$$

where we assume $\psi(x)$ has finite derivatives up to order $2S$. The coefficients $a_s(x)$, $s=0, 1, 2, \dots, S$ are finite, real valued functions on the domain $x \in [0, 1]$ with finite derivatives to order s and satisfy periodic boundary conditions,

$$\frac{\partial^t a_s}{\partial x^t}(0) = \frac{\partial^t a_s}{\partial x^t}(1), \quad t = 0, 1, \dots, s. \quad (2)$$

The minimal smoothness assumed of $a_0(x)$ is continuity on $x \in [0, 1]$. For concreteness, we impose periodic boundary conditions upon $\psi(x)$ itself,

$$\frac{\partial^t \psi}{\partial x^t}(0) = \frac{\partial^t \psi}{\partial x^t}(1), \quad t = 0, 1, \dots, 2S \quad (3)$$

although more general homogeneous boundary conditions could be insisted upon. Given the above definitions, a set of eigenfunctions $\phi_f(x)$ with corresponding real eigenvalues λ_f is defined through

$$\mathcal{D}\phi_f(x) = \lambda_f \phi_f(x), \quad (4)$$

and we order the eigenvalues λ_f , $f=1, 2, 3, \dots$ in ascending order $\lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$. The definition and boundary conditions in Eqs. (1)–(3) guarantee a Hermitian \mathcal{D} , meaning $\int_0^1 dx (\phi_f \mathcal{D} \phi_{f'} - \phi_{f'} \mathcal{D} \phi_f) = 0$ for any pair of eigenfunctions $\phi_f, \phi_{f'}$. All the usual eigenvalue/eigenfunction properties of Hermitian operators follow. Our task is to estimate a low order [$f=O(1)$] eigenvalue λ_f .

The most useful expression of the eigenvalue is the Rayleigh quotient,

$$\lambda_f = \int_0^1 dx \phi_f^*(x) \mathcal{D} \phi_f(x) = \int_0^1 dx \phi_f^*(x) \mathcal{L} \phi_f(x), \quad (5)$$

where we impose unity L_2 norm on the eigenfunctions

$$\|\phi_f\|_{L_2} = \left(\int_0^1 dx \phi_f^*(x) \phi_f(x) \right)^{1/2} \quad (6)$$

in anticipation of the quantum algorithm and the operator \mathcal{L} , derived from \mathcal{D} by simple integration by parts, is a more convenient (bilinear) operator to work with due to its symmetric form,

$$\varphi^*(x) \mathcal{L} \psi(x) = \sum_{s=0}^S \frac{\partial^s \varphi^*(x)}{\partial x^s} a_s(x) \frac{\partial^s \psi(x)}{\partial x^s} \quad (7)$$

for any two functions $\varphi(x)$ and $\psi(x)$.

It is useful to work not only in the “space” domain $x \in [0, 1]$, but also in the “reciprocal space” domain of integers, $k \in \mathbb{Z}$. The connection between the two representations is defined by the Fourier transforms

$$\begin{aligned} \tilde{\psi}_k &= \int_0^1 dx \exp(-2\pi i k x) \psi(x), \\ \psi(x) &= \sum_{k=-\infty}^{\infty} \exp(2\pi i k x) \tilde{\psi}_k, \end{aligned} \quad (8)$$

where a tilde will indicate a reciprocal space representation throughout the paper. Our eigenvalue Eq. (4) is Fourier transformed to

$$\sum_{k'=-\infty}^{\infty} \tilde{\mathcal{L}}_{k,k'} \tilde{\phi}_{f,k'} = \lambda_f \tilde{\phi}_{f,k}, \quad (9)$$

where

$$\tilde{\mathcal{L}}_{k,k'} = \sum_{s=0}^S (2\pi i k)^s \tilde{a}_{s,k-k'} (2\pi i k')^s \quad (10)$$

is the reciprocal space matrix representation of the operator \mathcal{L} (and \mathcal{D}). The Rayleigh quotient of Eq. (5) is Fourier transformed to

$$\lambda_f = \sum_{k,k'=-\infty}^{\infty} \tilde{\phi}_{f,k} \tilde{\mathcal{L}}_{k,k'} \tilde{\phi}_{f,k'}, \quad (11)$$

where we now have the Euclidean normalization,

$$\|\tilde{\phi}_f\|_2 = \left(\sum_{k=-\infty}^{\infty} \tilde{\phi}_{f,k}^* \tilde{\phi}_{f,k} \right)^{1/2} = 1, \quad (12)$$

consistent with $\|\phi_f\|_{L_2} = 1$ and our Fourier transform definition.

In a classical digital computer, discretization of the domain $x \in [0, 1]$ is required so that values x can be represented with a finite number of bits. For the quantum algorithm we

will be discussing, discretization of the domain will also be required so that values x can be identified with a finite number of qubits. We can then sample the spatial domain at the points $x=0, 1/N, 2/N, \dots, (N-1)/N$, where $N=2^n$ requires n qubits. It will be more convenient to work with the integers $\bar{x}=Nx=0, 1, 2, \dots, N-1$. The discrete spatial domain of N points allows us to approximate a function $\psi(x)$ by a vector,

$$\psi^{(N)} = (\psi_0^{(N)}, \psi_1^{(N)}, \dots, \psi_{N-1}^{(N)}) \quad (13)$$

for computational purposes, where we shall impose the Euclidean norm $\|\psi^{(N)}\|_2 = 1$. In particular, we wish to generate discretized approximations $\phi_f^{(N)}$ that approach the continuous problem eigenvector $\phi_f(x)$ such that taking an ever greater number of discretization points N gives us the limit $\lim_{N \rightarrow \infty} \sqrt{N} \phi_{f,\bar{x}}^{(N)} = \phi_f(x)$, the factor \sqrt{N} accounting for Euclidean normalization of the vector $\phi_f^{(N)}$ and L_2 normalization of the function $\phi_f(x)$. We discuss how we generate $\phi_f^{(N)}$ and how to quantify the quality of our discrete approximations as a function of N further below.

In addition to having a discrete approximation to functions $\psi(x)$, we shall require discrete approximations of the differential operator \mathcal{L} of Eq. (7) in the form of an $N \times N$ matrix $\mathcal{L}^{(N)}$ acting on vectors $\psi^{(N)}$. Hence, we will need $N \times N$ finite difference matrices, which we shall denote $\Delta_{(s)}^{(N)}$, to approximate derivatives $\partial^s / \partial x^s$. There is freedom in choosing finite differences to approximate derivatives, here we (arbitrarily) choose the forward difference for a concrete example,

$$(\Delta_{(1)}^{(N)} \psi^{(N)})_{\bar{x}} = N(\psi_{\bar{x}+1}^{(N)} - \psi_{\bar{x}}^{(N)}), \quad (14)$$

and higher order finite differences can be generated by $\Delta_{(s)}^{(N)} = (\Delta_{(1)}^{(N)})^s$ for integer s . From the very definition of derivatives, we have $\lim_{N \rightarrow \infty} \sqrt{N} (\Delta_{(s)}^{(N)} \psi^{(N)})_{\bar{x}} = \partial^s \psi(x) / \partial x^s$ if $\lim_{N \rightarrow \infty} \sqrt{N} \psi_{\bar{x}}^{(N)} = \psi(x)$, the factor \sqrt{N} again accounting for Euclidean normalization of the vector $\psi^{(N)}$ and L_2 normalization of the function $\psi(x)$. The subscript arithmetic $\bar{x} \pm 1$ in the definition of finite differences is to be performed modulo- N , consistent with the boundary conditions of Eqs. (2) and (3). The resulting matrix operator $\mathcal{L}^{(N)}$ is

$$\mathcal{L}^{(N)} = \sum_{s=0}^S (\Delta_{(s)}^{(N)})^T \text{Diag}(a_s^{(N)}) \Delta_{(s)}^{(N)}, \quad (15)$$

where $(\cdot)^T$ indicates matrix transpose and $\text{Diag}(\cdot)$ indicates a diagonal matrix with the vector argument along the diagonal. With the above construction for $\mathcal{L}^{(N)}$, we can now pose a Hermitian matrix eigenvalue problem,

$$\sum_{\bar{x}=0}^{N-1} \mathcal{L}_{\bar{x},\bar{x}}^{(N)} \phi_{f,\bar{x}}^{(N)} = \lambda_f \phi_{f,\bar{x}}^{(N)} \quad (16)$$

whose solutions will have the desired properties $\lim_{N \rightarrow \infty} \sqrt{N} \phi_{f,\bar{x}}^{(N)} = \phi_f(x)$ and $\lim_{N \rightarrow \infty} \lambda_f^{(N)} = \lambda_f$, with the obvious restriction $f \leq N$. A reciprocal space description is useful, for which we introduce the discrete Fourier transforms,

$$\begin{aligned}\tilde{\psi}_k^{(N)} &= \frac{1}{\sqrt{N}} \sum_{\bar{x}=0}^{N-1} \omega^{-k\bar{x}} \psi_{\bar{x}}^{(N)}, \\ \psi_{\bar{x}}^{(N)} &= \frac{1}{\sqrt{N}} \sum_{k=-N/2}^{N/2-1} \omega^{k\bar{x}} \tilde{\psi}_k^{(N)},\end{aligned}\quad (17)$$

where $\omega = \exp(2\pi i/N)$ and reciprocal space has been truncated to the set $\mathcal{N} = \{k \in \mathbb{Z}; -N/2 \leq k \leq N/2 - 1\}$. The eigenvectors are assigned unit Euclidean norm in both \bar{x} and k space representations,

$$\|\phi_f^{(N)}\|_2 = \left(\sum_{\bar{x}=0}^{N-1} \phi_{f,\bar{x}}^{*(N)} \phi_{f,\bar{x}}^{(N)} \right)^{1/2} = \left(\sum_{k \in \mathcal{N}} \phi_{f,k}^{*(N)} \phi_{f,k}^{(N)} \right)^{1/2} = 1, \quad (18)$$

so that the discrete analogs of Eqs. (5) and (11) are

$$\lambda_f^{(N)} = \sum_{\bar{x}, \bar{x}'=0}^{N-1} \phi_{f,\bar{x}}^{*(N)} \mathcal{L}_{\bar{x}', \bar{x}}^{(N)} \phi_{f,\bar{x}'}^{(N)} = \sum_{k, k' \in \mathcal{N}} \tilde{\phi}_{f,k'}^{*(N)} \tilde{\mathcal{L}}_{k', k}^{(N)} \tilde{\phi}_{f,k}^{(N)} \quad (19)$$

which we shall find useful below.

We shall call $|\lambda_f^{(N)} - \lambda_f|$ the truncation error, alluding to the fact that we wish to approximate λ_f with $\lambda_f^{(N)}$ while truncating reciprocal space from all integers \mathbb{Z} to the subset \mathcal{N} . We now proceed to show the well-known fact that replacing derivatives by finite differences ultimately limits the convergence of $\lambda_f^{(N)}$ to λ_f as the number of sampling points N increases. Straightforward application of previously stated definitions gives

$$\begin{aligned}(\Delta_{(1)}^{(N)} \tilde{\psi}_k^{(N)}) &= N[\exp(2\pi i k/N) - 1] \tilde{\psi}_k^{(N)} \\ &= 2\pi i k \tilde{\psi}_k^{(N)} \left[1 + \Theta\left(\frac{k^2}{N^2}\right) \right],\end{aligned}\quad (20)$$

where we have made use of series expansions and the fact that $|k| \leq N/2$ to arrive at the contribution $\Theta(k^2/N^2)$. The result holds for higher order derivatives.

An important parameter in characterizing truncation error is a reciprocal space cutoff $k(\phi_f)$, which can be defined for every ϕ_f . There will always exist a number $k(\phi_f)$ such that $|\tilde{\phi}_{f,k}|^2 = O(k^{-(4S+1+\epsilon)})$ for all $|k| > k(\phi_f)$ and some infinitesimal ϵ . This follows simply because ϕ_f must be differentiable up to order $2S$, and therefore the series $\sum_k (2\pi k)^{4S} |\tilde{\phi}_{f,k}|^2$ giving the norm of the $2S$ th derivative of ϕ_f must converge. The eigenvalue spectrum of the continuous domain operator \mathcal{L} is unbounded, and it can be shown that $\sup\{k(\phi_f)\}$ does not exist. However, since we restrict ourselves to $f=O(1)$, we can specify a finite $k(\phi_f)$ independent of N . For $N/2 > k(\phi_f)$, a reciprocal space cutoff $k(\phi_f^{(N)})$ must also exist since $\lim_{N \rightarrow \infty} \tilde{\phi}_{f,k}^{(N)} = \tilde{\phi}_{f,k}$. From here on, we shall not distinguish between $k(\phi_f^{(N)})$ and $k(\phi_f)$ as the precise value of the reciprocal space cutoff is not needed, but simply its existence. We thus define another subset of reciprocal space $\mathcal{M} = \{k \in \mathbb{Z}; |k| < k(\phi_f)\}$.

We have now collected enough ingredients to find the truncation error $|\lambda_f^{(N)} - \lambda_f|$. We assume that $N/2 > k(\phi_f)$, so

that a ‘‘reasonable’’ representation of ϕ_f can be made on the discretized domain. By ‘‘reasonable,’’ we mean the eigenvalue λ_f can be estimated using Eq. (11) and the truncated reciprocal space \mathcal{N} to give,

$$\lambda_f = \sum_{k, k' \in \mathcal{N}} \tilde{\phi}_{f,k'}^* \tilde{\mathcal{L}}_{k', k} \tilde{\phi}_{f,k} + O(N^{-(2S+\epsilon)}), \quad (21)$$

where the above result arises from the least convergent (highest order derivative) contribution to λ_f in the region $k, k' \notin \mathcal{N}$,

$$\begin{aligned}\sum_{k, k' \notin \mathcal{N}} \tilde{\phi}_{f,k'}^* (2\pi i k)^s \tilde{a}_{s, k-k'} (2\pi i k')^s \tilde{\phi}_{f,k} \\ = \sum_{k, k' \notin \mathcal{N}} O(k'^{-(S+1/2+\epsilon)} k^{-(S+1/2+\epsilon)}) = O(N^{-(2S+\epsilon)}),\end{aligned}\quad (22)$$

where we have made use of $|\tilde{\phi}_{f,k}| = O(k^{-(2S+1/2+\epsilon)})$ for $k > k(\phi_f)$. Thus, for $N/2 > k(\phi_f)$, truncation of the reciprocal space sum in Eq. (11) gives $O(N^{-(2S+\epsilon)})$ error.

Using the finite difference error of Eq. (20), the reciprocal space matrix $\tilde{\mathcal{L}}^{(N)}$ can be written

$$\tilde{\mathcal{L}}_{k, k'}^{(N)} = \sum_{s=0}^S \left\{ (2\pi i k)^s \tilde{a}_{s, k-k'} (2\pi i k')^s \left[1 + \Theta\left(\frac{k^2 + k'^2}{N^2}\right) \right] \right\}, \quad (23)$$

where we have used the fact that there is some freedom in approximating $a_s(x)$ by $a_{s, \bar{x}}^{(N)}$. We choose to match spectral components $\tilde{a}_{s, k}^{(N)} = \tilde{a}_{s, k}$, and accept that $a_{s, \bar{x}}^{(N)}$ may exhibit oscillation artifacts (Gibb’s phenomenon) due to discarding the contributions $\tilde{a}_{s, k}$ for $k \in \mathbb{Z} - \mathcal{N}$. Note that the smoothness of $a_s(x)$, meaning continuity and finite s order derivatives for $x \in [0, 1]$, implies the existence of reciprocal space cutoffs $k(a_s)$. We use Eq. (19) to decompose,

$$\begin{aligned}\lambda_f^{(N)} &= \sum_{k, k' \in \mathcal{M}} \tilde{\phi}_{f,k'}^* \tilde{\mathcal{L}}_{k', k} \tilde{\phi}_{f,k} \left[1 + \Theta\left(\frac{k(\phi_f)^2}{N^2}\right) \right] \\ &+ \sum_{k, k' \in (\mathcal{N} - \mathcal{M})} \tilde{\phi}_{f,k'}^* \tilde{\mathcal{L}}_{k', k} \tilde{\phi}_{f,k} \left[1 + \Theta\left(\frac{k^2 + k'^2}{N^2}\right) \right],\end{aligned}\quad (24)$$

where the error summed over $\mathcal{N} - \mathcal{M}$ is of order,

$$\begin{aligned}\sum_{k, k' \in (\mathcal{N} - \mathcal{M})} \Theta\left(\frac{1}{N^2}\right) O\left(\frac{k^2 + k'^2}{k'^{S+1/2+\epsilon} k^{S+1/2+\epsilon}}\right) \\ = \Theta\left(\frac{1}{N^2}\right) O\left(\frac{1}{k(\phi_f)^{2S-3}}\right).\end{aligned}\quad (25)$$

The diminishing contribution of the region $\mathcal{N} - \mathcal{M}$ to $\lambda_f^{(N)}$ ensures that the relative finite difference error $\Theta(k^2/N^2)$ does not approach unity but remains $\Theta(1/N^2)$. Collecting the results of Eqs. (21), (24), and (25), we can express $\lambda_f^{(N)} - \lambda_f$ as

$$\begin{aligned} \lambda_f^{(N)} - \lambda_f = & \sum_{k,k' \in \mathcal{N}} \tilde{\phi}_{f,k'}^{*(N)} \tilde{\mathcal{L}}_{k',k} \tilde{\phi}_{f,k}^{(N)} \left[1 + \Theta \left(\frac{1}{N^2} \right) \right] \\ & - \sum_{k,k' \in \mathcal{N}} \tilde{\phi}_{f,k'}^* \tilde{\mathcal{L}}_{k',k} \tilde{\phi}_{f,k}, \end{aligned} \quad (26)$$

where we have dropped the dependence upon $k(\phi_f)$ as it shall be of no further use. We note that $\tilde{\mathcal{L}}^{(N)} - \tilde{\mathcal{L}} = \Theta(1/N^2)$ in the reciprocal space \mathcal{M} , so we can consider $1/N^2$ a parameter of expansion in perturbation theory. The lowest order perturbation gives $\|\delta\tilde{\phi}_f^{(N)}\|_2 = \|\tilde{\phi}_f^{(N)} - \tilde{\phi}_f\|_2 = O(1/N^2)$ for a non-degenerate ϕ_f . Degenerate eigenvectors might be perturbed substantially, but this is merely the result of there being no preferred basis for the span of the degenerate eigenvectors. The same bounds on truncation error can be shown to apply to the degenerate case. Noting that Eq. (26) is second order in eigenvector and $\delta\tilde{\phi}_f$ is orthogonal to $\tilde{\phi}_f$, the contribution of $\delta\tilde{\phi}_f^{(N)}$ to the eigenvalue error is $O(1/N^4)$ and can therefore be ignored. The relative truncation error is

$$\left| \frac{\lambda_f^{(N)} - \lambda_f}{\lambda_f} \right| = \Theta \left(\frac{1}{N^2} \right) \quad (27)$$

which is the final result of this section. We emphasize that truncation error arises solely from the uniform discretization of the domain $x \in [0, 1]$.

III. QUANTUM ALGORITHM—ONE-DIMENSION

We present now the quantum algorithm as it applies to the Hermitian, one-dimensional boundary value problem discussed in the preceding section. We will show the various computational steps, and the rationale behind them.

First, we set forth some preliminaries. We will represent a vector $\psi^{(N)}$ with a quantum state composed of $n = \ln_2 N$ qubits whose probability amplitudes are encoded as follows:

$$|\psi^{(N)}\rangle = \frac{1}{\sqrt{N}} \sum_{\bar{x}=0}^{N-1} \psi_{\bar{x}}^{(N)} |\bar{x}\rangle, \quad (28)$$

where $|\bar{x}\rangle$ is an n qubit state storing the binary representation of \bar{x} . Similarly, the finite difference matrix $\mathcal{L}^{(N)}$ is mapped to an operator,

$$\Lambda^{(N)} = \sum_{\bar{x}, \bar{x}'=0}^{N-1} |\bar{x}\rangle \mathcal{L}_{\bar{x}, \bar{x}'}^{(N)} \langle \bar{x}'|, \quad (29)$$

and we define the unitary exponential,

$$U = \exp(i\Lambda^{(N)}\tau) = \sum_{q=0}^{\infty} \frac{(i\Lambda^{(N)}\tau)^q}{q!}, \quad (30)$$

where τ is a dimensionless constant whose value is chosen in advance of the simulation and where the unitarity of U follows from the Hermitian nature of $\Lambda^{(N)}$. The constant τ must be carefully chosen to arrive at a desired accuracy in eigenvalue $\lambda_f^{(N)}$ without an unnecessarily large number of operations. The prescription for choosing τ is described

further below in Sec. III. Note that τ is now an abstract scaling parameter rather than the time step of a quantum simulation.

We shall call the register of $n = \ln_2 N$ qubits the accumulator register. In addition, a register of $m = \ln_2 M$ qubits will be required to count phase steps, which we shall call the index register. Several ancilla qubits will be required, their number depending on the desired precision for the coefficients $a_s^{(N)}$ that specify Λ . The first steps are to load an initial state $|\psi^{(N)}\rangle$ into the accumulator and to form an equal superposition of all index qubit states, giving a complete state,

$$|\Psi\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} |\psi^{(N)}\rangle |j\rangle. \quad (31)$$

The state $|\psi^{(N)}\rangle$ is an initial estimate of the field eigenvector of interest. To determine the required computational work to arrive at a suitable initial estimate $|\psi^{(N)}\rangle$, it is useful to decompose the accumulator state in terms of the initially unknown eigenstates $|\phi_f^{(N)}\rangle$,

$$|\Psi\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \sum_{x=0}^{N-1} \alpha_f |\phi_f^{(N)}\rangle |j\rangle = \sum_{f=0}^{N-1} \alpha_f |\Psi_f\rangle, \quad (32)$$

where $\alpha_f = \langle \phi_f^{(N)} | \psi^{(N)} \rangle$. As will be shown, the probability that the Abrams-Lloyd algorithm will give an estimate of eigenvalue λ_f in a single iteration is $|\alpha_f|^2$. To obtain an estimate of λ_f with probability approaching unity, approximately $1/|\alpha_f|^2$ iterations will be required. It is thus necessary for $|\psi^{(N)}\rangle$ to have a large overlap with $|\phi_f^{(N)}\rangle$ in order to avoid numerous iterations of the algorithm. The best technique proposed thus far is that of Jaksch and Papageorgiou [10], where a more coarsely defined $\phi_f^{(N_0)}$ is determined first (i.e., $N_0 < N$). According to the analysis of the preceding section, a coarse approximation limited by truncation error will allow one to achieve

$$|\alpha_f|^2 \leq \|\tilde{\phi}_f - \tilde{\phi}_f^{(N_0)}\|_2^2 = \|\delta\tilde{\phi}_f^{(N_0)}\|_2^2 = 1 - O(1/N_0^2) \quad (33)$$

for $f = O(1)$. Thus, one might solve for a desired $\phi_f^{(N_0)}$ classically (with cost that we will discuss later), and load the state $|\phi_f^{(N_0)}\rangle$ into the accumulator with $\Theta(N_0)$ operations.

We shall now follow the linear portion of the algorithm as it operates on a particular component $|\Psi_f\rangle$, reintroducing the full superposition over all f in the final (nonlinear) measurement step. The next stage of the algorithm is to apply the unitary U to the accumulator conditional upon the index to produce the superposition,

$$|\Psi'_f\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} U^j |\phi_f^{(N)}\rangle |j\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \exp(ij\lambda_f^{(N)}\tau) |\phi_f^{(N)}\rangle |j\rangle. \quad (34)$$

Only M conditional applications of U are in fact required to form $|\Psi'_f\rangle$ from $|\Psi_f\rangle$. One applies U conditional on $j > 1$, then one applies U conditional on $j > 2$ and so forth until the $(M-1)$ th conditional U is applied for $j = M-1$. The conditional applications of U can be performed with a single ad-

ditional ancilla qubit as follows. With at most $\ln M$ logical operations, one can entangle the index register with an ancilla to form the state $\sum_j |j\rangle |C_{j,j'}\rangle$ where the ancilla $C_{j,j'}=1$ for $j \geq j'$ and $C_{j,j'}=0$ otherwise. The j' 'th application of U can be implemented as a U conditional on the ancilla $C_{j,j'}$. The ancilla is then disentangled from the quantum register by running the initial entangling operation once again.

The operator U acts in the full n qubit Hilbert space of $|\psi^{(N)}\rangle$, which will in general be prohibitively large, but it is

nevertheless possible to efficiently generate an approximation to U using operations in a few qubit Hilbert space. The structure of $\Lambda^{(N)}$ is a band diagonal matrix resulting from local operations, and thus it has a block diagonal representation in the qubit basis of the accumulator.

To illustrate explicitly some of the key features of the algorithm at work, it is useful to consider the simple example where $\mathcal{D} = \partial/\partial x \{a(x)(\partial/\partial x)\}$. The following decomposition is appropriate:

$$\Lambda^{(N)} = N^2 \underbrace{\begin{pmatrix} d_0 & & & & & \\ & d_1 & & & & \\ & & d_2 & & & \\ & & & d_3 & & \\ & & & & \ddots & \\ & & & & & \ddots \end{pmatrix}}_{\Lambda^{(N,1)}} - N^2 \underbrace{\begin{pmatrix} 0 & a_1 & & & & \\ a_1 & 0 & & & & \\ & & 0 & a_3 & & \\ & & a_3 & 0 & & \\ & & & & \ddots & \end{pmatrix}}_{\Lambda^{(N,2)}} - N^2 \underbrace{\begin{pmatrix} 0 & & & & & a_0 \\ & 0 & a_2 & & & \\ & a_2 & 0 & & & \\ & & & 0 & a_4 & \\ & & & a_4 & 0 & \\ a_0 & & & & & \ddots \end{pmatrix}}_{\Lambda^{(N,3)}}, \quad (35)$$

where $d_{\bar{x}} = a_{\bar{x}} + a_{\bar{x}+1}$. The operators $\Lambda^{(N,p)}$ can be written more compactly,

$$\Lambda^{(N,1)} = N^2 \sum_{\bar{x}} d_{\bar{x}} |\bar{x}\rangle \langle \bar{x}|,$$

$$\Lambda^{(N,2)} = -N^2 \sum_{\bar{x} \text{ even}} a_{\bar{x}+1} \{ |\bar{x}\rangle \langle \bar{x}+1| + |\bar{x}+1\rangle \langle \bar{x}| \},$$

$$\Lambda^{(N,3)} = -N^2 \sum_{\bar{x} \text{ odd}} a_{\bar{x}+1} \{ |\bar{x}\rangle \langle \bar{x}+1| + |\bar{x}+1\rangle \langle \bar{x}| \},$$

where $\Lambda^{(N,1)}$ is diagonal, and $\Lambda^{(N,2)}$, $\Lambda^{(N,3)}$ act in one qubit subspaces (conditional upon $[\bar{x}/2]$) in lieu of the full Hilbert space of $\Lambda^{(N)}$.

The unitary U can be approximated to take advantage of the above decomposition in several ways. For a general decomposition,

$$\Lambda^{(N)} = \sum_{p=1}^R \Lambda^{(N,p)}, \quad (36)$$

where for our simple example $R=3$, the Baker-Campbell-Hausdorff formulas can be used to show

$$\begin{aligned} U_{\Pi} &= \prod_{p=1}^R \exp(i\Lambda^{(N,p)}\tau/2) \prod_{p=R}^1 \exp(i\Lambda^{(N,p)}\tau/2) \\ &= \exp\left(i\Lambda^{(N)}\tau - \frac{i}{3!} \sum_{p,q=1}^R [\Lambda^{(N,p)}, [\Lambda^{(N,q)}, \Lambda^{(N,R)}]]\tau^3 \right. \\ &\quad \left. + O(\|\Lambda^{(N)}\|_2^4 \tau^4) \right) \end{aligned} \quad (37)$$

where terms bilinear in $\Lambda^{(N,p)}\tau$ are suppressed by the symmetry of the product formula shown. One may approximate U by U_{Π} to take advantage of the efficient implementation of $\exp(i\Lambda^{(N,p)}\tau)$ at the cost of introducing error. The quantum circuit for implementing U_{Π} for our simple example $\mathcal{D} = \partial/\partial x \{a(x)(\partial/\partial x)\}$ is shown in Fig. 1. The reason for the ease of implementing U_{Π} is apparent in Fig. 1, one applies single qubit unitaries conditional upon the evaluation of $a_{\bar{x}}$. The Solovay-Kitaev theorem guarantees that the single qubit unitaries can be implemented to an accuracy $\Theta(1/N^2)$ with $\Theta(\ln^c N)$ universal quantum gates [11]. We also assume that evaluation of $a_{\bar{x}}$ requires $O(\ln N)$ operations. Roughly speaking, the differentiability of $a(x)$ rules out pathological functions that have greater complexity.

Approximating U by U_{Π} implies that the algorithm will give an estimate of the eigenvalue $\lambda_{f,\Pi}^{(N)}$ of the operator

$$\Lambda_{\Pi}^{(N)} = \Lambda^{(N)} + O(\|\Lambda^{(N)}\|_2^3 \tau^2) \quad (38)$$

instead of the desired eigenvalue $\lambda_f^{(N)}$. We call the error




FIG. 1. The quantum circuits for applying (a) $\exp(i\Lambda^{(N,2)}\tau/2)$, (b) $\exp(i\Lambda^{(N,2)}\tau)$, and (c) $\exp(i\Lambda^{(N,3)}\tau/2)$, to the accumulator qubits $|\bar{x}\rangle = |\bar{x}_3\bar{x}_2\bar{x}_1\bar{x}_0\rangle$ for the decomposition of Eq. (54) with $N=2^4$. The ancillae initialized to state $|0\rangle$ are used to store the coefficients $d_{\bar{x}}$ or $a_{\bar{x}+1}$, to three bit precision with resolution δ , $d_{\bar{x}} = (\mathbf{d}_{\bar{x},2}2^2 + \mathbf{d}_{\bar{x},1}2^1 + \mathbf{d}_{\bar{x},0}2^0) \times \delta$ with $\mathbf{d}_{\bar{x},j} \in \{0,1\}$, and a similar binary description for $a_{\bar{x}+1}$. The circuit “ $d_{\bar{x}}$ ” maps $|\bar{x}\rangle|\text{anc}\rangle \rightarrow |\bar{x}\rangle|\mathbf{d}_{\bar{x}} \oplus \text{anc}\rangle$, while “ $a_{\bar{x}+1}$ (even/odd)” maps $|\bar{x}\rangle|\text{anc}\rangle \rightarrow |\bar{x}\rangle|\mathbf{a}_{\bar{x}+1} \oplus \text{anc}\rangle$ for even/odd \bar{x} . Single qubit rotations $R'_p = |0\rangle\langle 0| + \exp[i(N^2\delta\tau/2)2^p]|1\rangle\langle 1|$ give the desired diagonal phase shifts for $\exp(i\Lambda^{(N,1)}\tau/2)$. Single qubit rotations $X'_p = \exp[i(-N^2\delta\tau/2)2^p\sigma_x]$, where $\sigma_x = |0\rangle\langle 1| + |1\rangle\langle 0|$, to implement the desired off diagonal couplings of $\exp(i\Lambda^{(N,2)}\tau/2)$ and $\exp(i\Lambda^{(N,3)}\tau/2)$. The values of τ and δ can be inferred from the restriction that the operator splitting error is comparable to truncation error, described in Sec. IV. The parity shift operators, defined $D^\pm|\bar{x}\rangle = |\bar{x} \pm 1\rangle$, are required to shift the block structure of $\Lambda^{(N,3)}$ so that only operations on the least significant qubit \bar{x}_0 need be performed. The D^\pm can be implemented using quantum Fourier transforms [at cost of $O(\ln^2 N)$ operations] and single qubit rotations. Final disentanglement of ancillae is achieved by a second application of “ $d_{\bar{x}}$ ” or “ $a_{\bar{x}+1}$.”

introduced by using U_Π the splitting error, which has value $O(\|\Lambda^{(N)}\|_2^3\tau^2)$ provided $\|\Lambda^{(N)}\|_2\tau < 1$. The splitting error will be shown in the next section to limit the computational efficiency of estimating eigenvalues.

Applying U_Π rather than U , Eq. (34) takes the form

$$|\Psi'_{f,\Pi}\rangle = \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} \exp(ij\lambda_{f,\Pi}^{(N)}\tau) |\phi_{f,\Pi}^{(N)}\rangle |j\rangle. \quad (39)$$

The eigenvalue is encoded in the phase periodicity of $|\phi_{f,\Pi}^{(N)}\rangle |j\rangle$, and can be determined to at most the $\pm\pi/M$ precision allowed by a $\ln_2 M$ bit representation of a full 2π radians. We briefly review the procedure for retrieving the phase [12], beginning with the application of the quantum Fourier transform,

$$\text{QFT} = \frac{1}{\sqrt{M}} \sum_{l=0}^{M-1} \sum_{m=0}^{M-1} \exp(-2\pi ilm/M) |l\rangle\langle m|, \quad (40)$$

to the index qubits. The resulting state is

$$\text{QFT}|\Psi'_{f,\Pi}\rangle = \sum_{l=0}^{M-1} b_{l,f} |\phi_{f,\Pi}^{(N)}\rangle |l\rangle \quad (41)$$

with the coefficients

$$b_{l,f} = \frac{1}{M} \sum_{j=0}^{M-1} \exp[ij(\lambda_{f,\Pi}^{(N)}\tau - 2\pi l/M)], \quad (42)$$

which have square modulus

$$|b_{l,f}|^2 = \frac{\sin^2[M(2\pi l/M - \lambda_{f,\Pi}^{(N)}\tau)/2]}{M^2 \sin^2[(2\pi l/M - \lambda_{f,\Pi}^{(N)}\tau)/2]}. \quad (43)$$

A projective measurement of the index produces $|l'\rangle$ where $|\lambda_{f,\Pi}^{(N)}\tau/2\pi - l'/M| < 1/2M$ with a probability $|b_{l',f}|^2 \geq [M^2 \sin^2(\pi/2M)]^{-1}$. All eigenvalues will satisfy $|\lambda_{f,\Pi}^{(N)}| < \pi/\tau$ since we will impose $\|\Lambda^{(N)}\|_2\tau \ll 1$ (to be made precise in the next section), so identification of l' will determine an eigenvalue $\lambda_{f,\Pi}^{(N)}$ uniquely to a precision $\pm\pi/M\tau$.

Since we began not with the desired state alone, but with a superposition $|\Psi\rangle = \sum_{k=0}^{N-1} \alpha_k |\Psi_f\rangle$, measurement of the index will determine a particular $\lambda_{f,\Pi}^{(N)}$ with relative probability $|\alpha_f|^2$. It is the initial trial wave function $|\psi^{(N_0)}\rangle$ that determines the probability $|\alpha_f|^2$ of the eigenvalue/eigenvector pair being selected by a projective measurement.

Upon completion of the eigenvalue readout (via index bits l'), the accumulator is left in the eigenstate $|\phi_f^{(N)}\rangle$. This is useful since it allows further information to be extracted. For instance, one can efficiently test whether $|\phi_f^{(N)}\rangle$ has a particular symmetry, such as inversion symmetry about a particular point \bar{x} in the domain. This can serve as a partial check as to whether the desired $|\phi_f^{(N)}\rangle$ was indeed selected by the projective measurement.

IV. COMPUTATIONAL COST—ONE-DIMENSION

We now analyze the computational cost for implementing the Abrams-Lloyd algorithm for the one-dimensional Hermitian problem described in the preceding sections. As pointed out, there are three sources of error that must be considered to determine the required number of operations for a given accuracy in eigenvalue estimation.

First, uniform discretization of the continuous problem to $N=2^n$ points on the spatial domain introduces a truncation error,

$$\left| \frac{\lambda_f^{(N)} - \lambda_f}{\lambda_f} \right| = \Theta\left(\frac{1}{N^2}\right). \quad (44)$$

The truncation error quantifies the accuracy with which the discrete problem represents the continuous problem for low order [i.e., $f=O(1)$] eigenfunctions ϕ_f . To compare algorithms, classical or quantum, we may ask how many operations are required to achieve the $\Theta(1/N^2)$ accuracy in the solution of the discrete eigenvalue problem.

Second, splitting $\Lambda^{(N)}$ into parts so as to approximate U with a product U_{Π} of local operators results in what we have termed splitting error. From Eq. (38) the eigenvalue $\lambda_{f,\Pi}^{(N)}$ of U_{Π} is

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O(\|\Lambda^{(N)}\|_2^3 \tau^2), \quad (45)$$

where we choose τ such that $\|\Lambda^{(N)}\|_2 \tau = \|\mathcal{L}^{(N)}\|_2 \tau < 1$. However, from the finite difference formula Eq. (14) and the form of $\mathcal{L}^{(N)}$ in Eq. (15), the spectral radius $\|\mathcal{L}^{(N)}\|_2 = \Theta(N^{2S})$. Hence, the splitting error is

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O(N^{6S} \tau^2) \quad (46)$$

which, unlike truncation error, increases polynomially with an increase in the number of discretization points N . The splitting error results from the fact that the product U_{Π} creates deviations from the true advancement in phase at high spatial frequencies. For example, in the system described in Eq. (35), it is the noncommuting nature of advancing even pairings of points and odd pairings of points that generates an error with spatial frequency $N/2$.

Third, the measurement of phase $\lambda_{f,\Pi}^{(N)} \tau$ via the quantum Fourier transform is limited by the uniform discretization of 2π radians into $M=2^m$ intervals. The limited phase resolution allows us to specify $\lambda_{f,\Pi}^{(N)}$ upon completion of the algorithm to a precision $2\pi/M\tau$.

The three sources of error allow us to determine the optimal number of index bits $m=\ln_2 M$, the value of the constant τ , and thus the complexity of the algorithm. Obviously, there is nothing gained in solving the discretized problem to an accuracy greater than the truncation error $\Theta(1/N^2)$ if the goal is to study the continuous problem. We can thus allow the splitting error $O(N^{6S} \tau^2)$ to be of the same order as the truncation error,

$$\Theta\left(\frac{1}{N^2}\right) \geq O(N^{6S} \tau^2) \rightarrow \tau \leq \Omega\left(\frac{1}{N^{3S+1}}\right). \quad (47)$$

Since $\lambda_{f,\Pi}^{(N)} = O(1)$ for our low order eigenvalue with $f=O(1)$, the phase advancement $\lambda_{f,\Pi}^{(N)} \tau \leq \Omega(1/N^{3S+1})$ for the low order eigenfunction becomes exceedingly small. In order to resolve this phase so that our final eigenvalue uncertainty does not exceed the truncation error, we require

$$\frac{2\pi}{M\tau} \leq \Theta\left(\frac{1}{N^2}\right) \rightarrow M \geq O(N^{3(S+1)}) \quad (48)$$

thus prescribing the number $m=\ln_2 M$ of index register qubits.

The complexity of the eigenvalue estimation can now be stated. The determination of a suitable initial guess eigenstate $\phi_f^{(N_0)}$ requires the determination of an eigenvector of an $N_0 \times N_0$ problem. This can be done classically in $\Omega(N_0)$ steps, since each of N_0 points in the spatial domain description of $\mathcal{L}^{(N_0)}$ must contribute to the eigenvalue. Near optimal classical methods are in fact known. In the case of a tridiagonal $\mathcal{L}^{(N_0)}$, bisection gives an eigenvalue to $\Theta(1/N_0^2)$ precision with $\Theta(N_0 \ln N_0)$ operations [13]. Low order eigenvalues of wider bandwidth $\mathcal{L}^{(N_0)}$ matrices can be determined to the same precision with the same order of operations using more complex classical techniques (1). Only a modest N_0 is required for the probability of a successful iteration of the quantum algorithm, $1 - O(1/N_0^2)$, to be comparable to unity. Following the construction of an initial eigenstate estimate, this estimate must be loaded into the accumulator register, which can be done in $\Theta(N_0)$ steps. We suppose that N will exceed N_0 by a substantial factor, so that the initial state preparation is a negligible cost compared to the remainder of the algorithm. The majority of the computational steps in the quantum algorithm are accounted for by the $M \geq O(N^{3(S+1)})$ applications of U_{Π} , each of which requires $O(\ln^c N)$ gate operations for some constant $c=O(1)$. The final quantum Fourier transform requires $\Theta(\ln^2 M)$ gate operations, a negligible $\ln(N)$ contribution compared to the M applications of U_{Π} .

Thus, to achieve $\Theta(1/N^2)$ accuracy in the final eigenvalue, at least $O(N^{3(S+1)} \ln^c N)$ operations and $\Theta(D \ln N)$ qubits are required. In contrast, an eigenvalue can be found using classical techniques to $\Theta(1/N^2)$ accuracy using $\Theta(N \ln N)$ operations. The quantum algorithm requires significantly more work than classical algorithms for the one-dimensional problem. Nonetheless, we show in the next section that the quantum algorithm is easily extended to higher dimensional problems where increased efficiency over classical techniques is indeed possible.

V. HIGHER DIMENSIONAL PROBLEMS

Here we will generalize the results of the one-dimensional problem to the multidimensional problem. Many of the arguments presented in the earlier sections are not specific to the single dimension domain, and in many cases we can simply replace scalars with vectors. The continuous problem we wish to solve involves an operator \mathcal{D} mapping functions $\psi(\mathbf{x})$, defined over a D -dimensional cubic domain $\mathbf{x} \in \mathcal{S} = [0, 1]^{\otimes D}$, to functions $\mathcal{D}\psi(\mathbf{x})$. Rather than explicitly writing out the general form of a multidimensional Hermitian operator \mathcal{D} analogous to the single dimensional operator of Eq. (1), we simply state that \mathcal{D} must satisfy

$$\int_0^1 dx_1 \cdots \int_0^1 dx_D (\phi_f^* \mathcal{D} \phi_{f'} - \phi_{f'}^* \mathcal{D} \phi_f) = 0 \quad (49)$$

for any eigenfunctions ϕ_f satisfying $\mathcal{D} \phi_f = \lambda_f \phi_f$. We can then define an ‘‘equivalent’’ bilinear operator \mathcal{L} that maps any two

vector functions $\psi^*(\mathbf{x})$ and $\varphi(\mathbf{x})$ to a scalar function $\psi^* \mathcal{L} \varphi$. This can be done by using the higher dimensional forms of integration by parts, which in one dimension allowed us to relate \mathcal{D} to \mathcal{L} . We exclude “trivial” problems that are readily expressed as a tensor product of single dimensional problems, $\mathcal{L} = \mathcal{L}_1 \otimes \mathcal{L}_2 \otimes \dots \otimes \mathcal{L}_D$. We are therefore considering problems whose structure is instead a sum of tensor product terms

$$\mathcal{L} = \sum_{\beta=1}^B \mathcal{L}_{\beta,1} \otimes \mathcal{L}_{\beta,2} \otimes \dots \otimes \mathcal{L}_{\beta,D} \quad (50)$$

for some constant $B > 1$, and where the differential order of each one-dimensional $\mathcal{L}_{\beta,\alpha}$ is $2S_{\beta,\alpha}$. The differential order of \mathcal{L} is then $2S = \max_{\beta} \{\sum_{\alpha} 2S_{\beta,\alpha}\}$. Of course, we retain the Hermitian property

$$\int_0^1 dx_1 \dots \int_0^1 dx_D (\phi_f^* \mathcal{L} \phi_{f'} - \phi_{f'}^* \mathcal{L} \phi_f) = 0 \quad (51)$$

and the associated eigenvalue/eigenvector properties. Normalizing the eigenfunctions allows us to write

$$\lambda_f = \int_0^1 dx_1 \dots \int_0^1 dx_D \phi_f^* \mathcal{L} \phi_f \quad (52)$$

which is simply the Rayleigh quotient.

Discretization proceeds as in the single dimensional case, with each domain coordinate $x_i \in [0, 1]$ discretized to N points. Functions $\psi(\mathbf{x})$ are represented by rank D tensors $\psi_{\mathbf{x}}^{(N)}$, i.e., for $D=2$ dimensions $\psi_{\mathbf{x}}^{(N)}$ is a matrix of numbers, for $D=3$ dimensions $\psi_{\mathbf{x}}^{(N)}$ is a “cube” of numbers, and so forth. Partial derivatives are converted to finite differences as in the one-dimensional case. The operator \mathcal{L} can thus be discretized to a tensor $\mathcal{L}_{\mathbf{x},\mathbf{x}'}^{(N)}$. The truncation error in the multidimensional problem is

$$\left| \frac{\lambda_f^{(N)} - \lambda_f}{\lambda_f} \right| = \Theta \left(\frac{1}{N^2} \right) \quad (53)$$

which is identical to the one-dimensional case because the relative finite difference errors on each coordinate are $\Theta(1/N^2)$.

The implementation of the Abrams-Lloyd algorithm for the multidimensional problem proceeds in a completely analogous fashion to the one-dimensional case, with the number of accumulator qubits $D \ln_2 N = Dn$ so as to represent a volume $V = N^D$. As before, an initial estimate of the desired eigenvector $\phi_f^{(N)}$ is required. A coarse classical simulation can produce an eigenvector $\phi_f^{(N_0)}$ with $N_0 < N$. Since truncation error scales as $\Theta(1/N^2)$, the required value of N_0 is such that the probability of a successful iteration of the algorithm, $1 - O(1/N_0^2)$, approaches unity. The computational cost is $\Theta(N_0^D \ln N_0)$ classical gate operations for generating the initial eigenstate and $\Theta(N_0^D)$ gate operations to load the state into the accumulator.

The heart of the algorithm is the controlled application of the unitary $U = \exp(i\Lambda^{(N)} \tau)$ where $\Lambda^{(N)} = \sum |\bar{\mathbf{x}}\rangle \mathcal{L}_{\mathbf{x},\mathbf{x}'}^{(N)} \langle \bar{\mathbf{x}}'|$. As before, U acts within a large Hilbert space, so an approximat-

ing operator U_{Π} is applied instead. The operator U_{Π} is a sequence of operations acting conditionally upon a much smaller Hilbert space than the full $D \ln_2 N$ qubits. We quantify the size of this Hilbert space now. The multidimensional $\Lambda^{(N)}$ is no longer represented by a band diagonal matrix, but has the structure of a sum of tensor products as in Eq. (50),

$$\Lambda^{(N)} = \sum_{\beta=1}^B \Lambda_{\beta,1}^{(N)} \otimes \Lambda_{\beta,2}^{(N)} \otimes \dots \otimes \Lambda_{\beta,D}^{(N)}. \quad (54)$$

The local nature of $\Lambda^{(N)}$ is quantified by the maximum number of states $|\mathbf{x}'\rangle$ for which $\langle \bar{\mathbf{x}}'| \Lambda^{(N)} |\bar{\mathbf{x}}\rangle$ is not zero (maximizing over all possible $|\bar{\mathbf{x}}\rangle$). This volume, v , is the maximum product of matrix bandwidths,

$$v = \max_{\beta} \{(2S_{\beta,1} + 1)(2S_{\beta,2} + 1) \dots (2S_{\beta,D} + 1)\} \leq \left(1 + \frac{2S}{D}\right)^D, \quad (55)$$

where we have used the restriction $2S = \max_{\beta} \{\sum_{\alpha} 2S_{\beta,\alpha}\}$ to arrive at the bound on v . It follows that we can split $\Lambda^{(N)} = \sum_{p=1}^R \Lambda^{(N,p)}$ where the $\Lambda^{(N,p)}$ act conditionally upon a Hilbert space of $r = \lceil \ln_2 v \rceil$ qubits. The size of this reduced Hilbert space is independent of domain size N^D , so that $\exp(i\Lambda^{(N,p)} \tau)$ can be applied to the requisite accuracy (polynomial in $1/N$) with only $\Theta(\ln^c N)$ universal gates for some constant c . As before, we assume that the function evaluations required for conditional action upon the r -qubit subspace entails at most $O(\ln N)$ universal gates. The total number of the split up operators $\Lambda^{(N,p)}$ is bounded $R \leq v = (1 + 2S/D)^D$ independently of the domain size N^D . Thus, U_{Π} can be applied with $O(\ln^c N)$ work for some $c = O(1)$.

The approximation U_{Π} can be composed by using a symmetric product as in Eq. (37) so that the splitting error is

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O(\|\Lambda^{(N)}\|_2^3 \tau^2). \quad (56)$$

More generally, an approximation of U correct to higher order in $\|\Lambda^{(N)}\|_2 \tau$ can be implemented [14–16]. For the sake of generality, we assume we have a product operator U_{Π} correct to order $\|\Lambda^{(N)}\|_2^{\nu} \tau^{\nu}$, and set $\nu=2$ to recover the simple symmetric product results. In practice, one cannot take ν arbitrarily large since the number of terms in U_{Π} grows exponentially in ν . The optimal choice of ν is that which minimizes the overall computational cost.

Using the fact that $2S$ is the differential order of \mathcal{L} , the splitting error becomes

$$\lambda_{f,\Pi}^{(N)} = \lambda_f^{(N)} + O(N^{2S(\nu+1)} \tau^{\nu}). \quad (57)$$

The final phase measurement through a quantum Fourier transform proceeds as in the one-dimensional case, with the same precision of $\pm \pi/(M\tau)$ in determining $\lambda_{f,\Pi}^{(N)}$, where $m = \ln_2 M$ is the number of index qubits. Requiring that the final eigenvalue be determined to the truncation error limit as in the one-dimensional case, the same line of reasoning as in the preceding section leads to

$$\tau \leq \Omega \left(\frac{1}{N^{2(S(1+1/\nu)+1/\nu)}} \right),$$

$$M \geq O(N^{2(S+1)(1+1/\nu)}). \quad (58)$$

The computational cost of the algorithm is dominated by the M applications of U_{Π} , each application of U_{Π} requiring $O(\ln^c N)$ number of operations. The computational cost for the quantum algorithm is

$$\mathfrak{N}_Q = O(M \ln^c N) = O(N^{2(S+1)(1+1/\nu)} \ln^c N) \quad (59)$$

in addition to the cost for finding and loading an eigenstate with coarse discretization N_0 along each axis. We assume $O(N_0)$. The number of qubits required by the quantum algorithm is $\Theta(\ln N)$.

We now consider classical costs associated with the multidimensional eigenvalue equation. Discretization and reduction of the continuous problem to a matrix equation results in a sparse $N^D \times N^D$ matrix with a number of bands depending on the spatial derivatives and dimensions in the continuous problem. The most efficient and near optimal classical method requires

$$\mathfrak{N}_C = O(N^D \ln N) \quad (60)$$

operations in order to attain a low order eigenvalue with truncation error accuracy $\Theta(1/N^2)$. The method is near optimal in the classical case since the computational cost per each of $V=N^D$ points in the domain is merely $O(\ln N)$. Any classical method must “visit” each point in the simulation domain in order for that point to influence the outcome of the classical calculation, hence the classical computation cost is $\Omega(N^D)$. Of course, the number of bits required is $\Theta(N^D)$.

The maximum improvement in computational efficiency provided by the quantum algorithm presented is

$$\max\left(\frac{\mathfrak{N}_C}{\mathfrak{N}_Q}\right) = O\left(\frac{N^{D-2(S+1)(1+1/\nu)}}{\ln^{c-1} N}\right) \quad (61)$$

with respect to the best known (near optimal) classical algorithm. From the above, we see that the domain dimension must satisfy $D > 2(S+1)(1+1/\nu)$ in order to see any improvement using the Abrams-Lloyd algorithm. In particular, we have $S=1$ for Schrödinger’s equation and we can identify $D/3$ as the number of particles in space (3 degrees of freedom per particle, neglecting spin). A many-body eigenvalue calculation is more efficient than classical simulation for particle number $D/3 > (4/3)(1+1/\nu)$. For the case where U_{Π} is a simple symmetric product, $\nu=2$ and we require $D/3 > 2$ in order to see improved computational efficiency. Higher order approximations, $\nu > 2$ will result in two (spinless) particle calculations already being done more efficiently using the Abrams-Lloyd algorithm.

We now discuss the generality of the results for domains other than the simple hypercube $\mathcal{S}=[0,1]^{\otimes D}$ discretized to $V=N^D$ points. A more complex domain \mathcal{S}' can be had by deleting regions from \mathcal{S} along planes defined by the uniform discretization scheme. The computational cost incurred is that required to ensure the probability amplitudes in $|\psi\rangle$ do not “spill” into the deleted regions $\mathcal{S}-\mathcal{S}'$. This is easily done by circuits such as those in Fig. 1, wherein quantum gates can be used to determine the conditional application of a few-qubit operator throughout the simulation domain. The

computational cost is therefore proportional to the classical cost of determining whether a point \bar{x} is in or out of the specified domain \mathcal{S}' subtended by the hypercube \mathcal{S} . As an explicit example, the subcircuit $a_{\bar{x}+1}$ of Fig. 1 for applying $\exp(i\Lambda^{(N,2)}\tau)$ of Eq. (35) can be made to compute $|\bar{x}\rangle|\text{anc}\rangle \rightarrow |\bar{x}\rangle|\text{anc}\rangle$ for $\bar{x} \in \mathcal{S}-\mathcal{S}'$ and $|\bar{x}\rangle|\text{anc}\rangle \rightarrow |\bar{x}\rangle|a_{\bar{x}+1} \oplus \text{anc}\rangle$ for $\bar{x} \in \mathcal{S}'$. The effect of this operation is to conditionally apply $\exp(i\Lambda^{(N,2)}\tau)$ to those points $\bar{x} \in \mathcal{S}'$. Clearly, $\mathcal{S}'=\mathcal{S}$ is the simplest domain to consider as there is no added computational cost, but more complex domains are accessible at only the modest cost of describing the domain with a Boolean function.

VI. CONCLUSION

Our analysis of the Abrams-Lloyd algorithm raises several questions. First, it is natural to ask what sort of qubit phase rotation accuracy is required during the application of U_{Π} to the initial guess eigenstate. The phase that is applied to qubits by the operator U_{Π} during the computation is of the same order as the phase applied to the highest order eigenvector, $\lambda_{N,\Pi}^{(N)}\tau$ where the eigenvalue $\lambda_{N,\Pi}^{(N)}=\Theta(N^{2S})$ for a differential operator of order $2S$ and $\tau=\Omega(1/N^{3S+1})$ for a second order splitting formula. The magnitude of the phase rotations applied to qubits is therefore $\Omega(1/N^{S+1})$. The relative accuracy with which the phase must be applied is $\Theta(1/N^2)$ if the final eigenvalue estimation is to be accurate to the truncation error limit of $\Theta(1/N^2)$. Thus, the absolute accuracy required from single qubit rotations is $\Omega(1/N^{S+3})$, independent of the number of dimensions D . The absolute accuracy is a small quantity for very modest values of $N=100$ (representing a relative eigenvalue accuracy of 10^{-4}) with a second order operator ($2S=2$). Angular resolution of 10^{-8} in the control of qubits represents a technical feat, but thankfully the principles of fault tolerant quantum computation [17,18] can be applied here to lessen the accuracy requirements for physical qubit operations.

Second, it is tempting to compare the quantum and classical algorithms for the simulation of dynamical evolution. The Abrams-Lloyd algorithm simulates the dynamics of the Schrödinger equation $\partial\psi/\partial t=\mathcal{D}\psi$ for some (possibly fictitious) Hamiltonian represented by \mathcal{D} , but only limited detail of the dynamics in a quantum simulation are accessible. The probability amplitudes characterizing a register of $D \ln N + \ln M$ qubits can result in at most $D \ln N + \ln M$ classical bits of information being extracted by measurement (by the Holevo bound). For instance, in order to obtain the eigenvector coefficients $\phi_f^{(N)}$, at least $\Theta(N^D/\ln N)$ iterations of the algorithm would be required. This is in contrast to a classical simulation of dynamical evolution where $\Theta(N)$ bits would be required to store a state at a single dynamical step, and $\Theta(NM)$ bits are required to store the entire evolution of an initial state over M dynamical steps. We emphasize that the strength of the Abrams-Lloyd algorithm is not in its ability to provide great detail into dynamical evolution but rather in extracting useful classical information (such as eigenvalues) from a very compact representation of that dynamical evolution.

Finally, the analysis of the Abrams-Lloyd algorithm raises the question as to why the eigenvalue convergence for low dimensional problems (i.e., small D) is less than that of optimal classical approaches. Part of the answer lies in the classical theory of the matrix eigenvalue calculation. An important tool for numerical estimation of eigenvalues is the Krylov subspace, which is defined to be the span of the set $\{\psi, A\psi, A^2\psi, \dots, A^{M'-1}\psi\}$ for some initial guess vector ψ , some hopefully small constant $M' < N^D$, and some $N^D \times N^D$ matrix A of which we seek several low order eigenvalues. The Krylov subspace is spanned by at most M' vectors, rather than the full N^D vector space of A , and so projecting onto the Krylov subspace gives an efficient means of estimating eigenvalues/eigenvectors of A . If the matrix whose lowest eigenvalue is sought is $\Lambda^{(N)}$, then we might choose $A = (\Lambda^{(N)} - \mu I)^{-1}$ where μ is an initial estimate of the eigenvalue sought (the eigenvalues of $\Lambda^{(N)}$ being simply related to those of A). With $A = (\Lambda^{(N)} - \mu I)^{-1}$, the vector $A^i \psi$ converges exponentially towards the eigenvector $\phi_f^{(N)}$ whose eigenvalue minimizes $|\lambda_f^{(N)} - \mu|$. In contrast, if $A = \exp(i\Lambda^{(N)})$ as in the

Abrams-Lloyd algorithm, there is no such convergence towards a target eigenvector since the eigenvalues of A are of unit norm. The unitarity of quantum gates restricts eigenvalues to lie on the unit circle in the complex plane, which is a poor eigenvalue distribution from the perspective of estimating a target eigenvalue [19]. This leads to the question of whether controlled decoherence can be used to produce non-unitarity evolution to accelerate the selection of a target eigenvector with a net reduction in gate operations and/or delay.

ACKNOWLEDGMENTS

The authors thank Chris Anderson, Oscar P. Boykin, Salman Habib, Colin Williams, Seth Lloyd, Joseph F. Traub, and Isaac Chuang for stimulating discussions and useful suggestions. This work was supported by the Defense Advanced Research Projects Agency and the Defense MicroElectronics Activity.

-
- [1] R. P. Feynman, *The Feynman Lectures on Computation* (Perseus, Massachusetts, 2000).
- [2] S. Lloyd, *Science* **273**, 1073 (1996).
- [3] D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **79**, 2586 (1997).
- [4] D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **83**, 5162 (1999).
- [5] B. Boghosian and W. Taylor, quant-ph/9701019 (unpublished).
- [6] C. Zalka, *Proc. R. Soc. London, Ser. A* **454**, 313 (1998).
- [7] D. A. Meyer, quant-ph/0111069 (unpublished).
- [8] D. A. Lidar and O. Biham, *Phys. Rev. E* **56**, 3661 (1997).
- [9] A. Papageorgiou and H. Woźniakowski, *Quantum Inf. Process.* **4**, 87 (2005).
- [10] P. Jaksch and A. Papageorgiou, *Phys. Rev. Lett.* **91**, 257902 (2003).
- [11] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000), Appendix 3.
- [12] R. Cleve, A. Ekert, C. Machiavello, and M. Mosca, *Proc. R. Soc. London, Ser. A* **454**, 339 (1998).
- [13] J. W. Demmel, *Applied Numerical Linear Algebra* (SIAM, New York, 1997).
- [14] H. Yoshida, *Phys. Lett. A* **150**, 262 (1990).
- [15] N. Hatano and M. Suzuki, quant-ph/0506007 (unpublished).
- [16] A. T. Sornborger and E. D. Stewart, *Phys. Rev. A* **60**, 1956 (1999).
- [17] J. Preskill, *Proc. R. Soc. London, Ser. A* **454**, 385 (1998).
- [18] D. Gottesman, *Phys. Rev. A* **57**, 127 (1998).
- [19] Y. Saad, *Numerical Methods for Large Eigenvalue Problems* (Halsted, New York, 1992).
- [20] A. Brandt, S. McCormick, and J. Ruge, *SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput.* **4**, 244 (1983).