

Approximating fractional time quantum evolution

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2009 J. Phys. A: Math. Theor. 42 185302

(<http://iopscience.iop.org/1751-8121/42/18/185302>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.153

The article was downloaded on 03/06/2010 at 07:38

Please note that [terms and conditions apply](#).

Approximating fractional time quantum evolution

L Sheridan¹, D Maslov^{1,2} and M Mosca^{1,3,4,5}

¹ Institute for Quantum Computing, University of Waterloo, Waterloo, ON N2L 3G1, Canada

² National Science Foundation, 4201 Wilson Blvd, Arlington, VA 22230, USA

³ Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada

⁴ St Jerome's University, Waterloo, ON N2L 3G3, Canada

⁵ Department of Combinatorics & Optimization, University of Waterloo, ON N2L 3G1

E-mail: lssherid@iqc.ca

Received 15 January 2009, in final form 17 March 2009

Published 17 April 2009

Online at stacks.iop.org/JPhysA/42/185302

Abstract

An algorithm is presented for approximating the arbitrary powers of a black box unitary operation, \mathcal{U}^t , where t is a real number and \mathcal{U} is a black box implementing an unknown unitary. The complexity of this algorithm is calculated in terms of the number of calls to the black box, the errors in the approximation and a certain 'gap' parameter. For general \mathcal{U} and large t , one should apply \mathcal{U} a total of $\lfloor t \rfloor$ times followed by our procedure for approximating the fractional power $\mathcal{U}^{t-\lfloor t \rfloor}$. An example is also given where for large integers t , this method is more efficient than direct application of t copies of \mathcal{U} . Further applications and related algorithms are also discussed.

PACS numbers: 03.67.Ac, 03.67.Lx

1. Introduction

Suppose an n -qubit unitary \mathcal{U} is implemented by evolving (or simulating the evolution of) a *time-independent* Hamiltonian H for a period of time $\tau = 1$, that is, $U = e^{-iH}$. Then for any $t \in \mathbb{R}_+$, one can implement \mathcal{U}^t by simply evolving the Hamiltonian for a period of time t . For example, if $t = \frac{1}{2}$, then a square root of \mathcal{U} , $e^{-i\frac{1}{2}H}$, could be implemented in this way, and in such a model of computation the cost would be half of the cost of implementing \mathcal{U} .

In our work, we question what can be done if \mathcal{U} is realized in some other way, such as a non-trivial sequence of time-dependent Hamiltonians or a quantum circuit. In other words, consider the situation when \mathcal{U} is given in the form of a black box. The goal is to implement real-valued powers t of this unitary operation by making use of the multiple copies of the black box implementing \mathcal{U} . The complexity of such a procedure is measured in terms of the total number of calls to the unitary.

It is possible to find the t th power of an unknown unitary by first performing a sufficiently precise complete quantum process tomography of a $(2^n \times 2^n)$ -dimensional unitary \mathcal{U} , which

uses $O(4^n)$ calls to the unitary with various input states [1, 2] and measurements to achieve \mathcal{U} with constant precision. The exponential scaling with n is necessary. In particular, a lower bound on the number of calls to the unitary can be derived from constructing an ϵ -net over unitaries⁶. If the space of unitary operations is divided up into balls of radius ϵ , then the total number of unitaries that can be specified up to this resolution is $(\frac{c}{\epsilon})^{4^n}$, where c is a constant. To specify each of these requires $4^n \log(\frac{c}{\epsilon})$ bits. To discover information about \mathcal{U} , one might supply states to the unitary tensored with the identity operation of the same dimension (a larger dimension does not help) and perform measurements on the output states. The states will have dimension 4^n , and by Holevo's bound one could hope to obtain no more than $2n$ bits of information for each of these calls to \mathcal{U} (in general, we will get less than this). Therefore, it is not possible by tomography to use fewer than $\Omega(\frac{4^n}{2n} \log(\frac{c}{\epsilon}))$ calls to \mathcal{U} . Even if we allow for error in some fraction δ of the basis states, this still cannot reduce the required number of calls to any function polynomial in n [3].

In this paper, we report on a more efficient approach to implementing powers. In particular, we present an algorithm for approximating with error ϵ (using the trace norm) any constant power of an unknown unitary using only $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ calls to the unitary itself, assuming that a certain 'gap' parameter g we explain later satisfies $g \in \Omega(\epsilon)$ (otherwise, the complexity is $O(\frac{1}{g} \log \frac{1}{\epsilon})$ calls; see section 2.2 for more detail). Note that this complexity is independent of the number of qubits n .

There has been some relevant prior work concentrating on the relation between discrete and continuous oracles for various problems. Farhi and Gutmann [4] introduced the concept of (continuous) Hamiltonian black box oracles for quantum computing. Ioannou [5] and Roland and Cerf [6] consider the problem of simulating a Hamiltonian for Grover search [7] on a discrete computer. Mochon [8] extends this to a more general setting where he is concerned with finding lower bounds for discrete oracle problems by considering them in the Hamiltonian setting and mapping them to the problem of finding geodesics in manifolds. In particular, he considers one-item Grover search and oracle interrogation, highlighting the case of computing the XOR function on a hidden bit string. In these cases, he is able to exploit symmetries in the oracle problems to solve what is otherwise a very difficult problem. Our focus is different. First, we do not allow for an oracle to be applied a fractional amount of time (instead, we are concerned with simulation of fractional oracles in the case when an oracle is only allowed to be applied an integer number of times). Second, we present an algorithm, with our focus on (constructive) upper bounds. Third, we do not make symmetry assumptions on our oracles. Finally, we do not restrict the eigenvalues of these unitaries to ± 1 as in the Grover search and XOR case, or the work in [9] which shows how to simulate a general continuous *Boolean* oracle evolving for a total of time T with $O(T \log T / \log \log T)$ discrete oracle queries.

The remainder of the paper is organized as follows. In section 2, we describe the basic construction for efficient computation of the powers of \mathcal{U} , discuss our assumptions and constraints, give the complexity of the computation and calculate the precision of our approximate solution. The full calculation of these results is included in appendix B. An example of a special case in which t is a large positive integer and our algorithm is more efficient than direct application of the t copies of the unitary is given in section 3. Related applications of our algorithm to computing the fractional Fourier transform and noise filtering are explored in section 4. Finally, we discuss the implications of this work in the concluding section.

⁶ We thank P Hayden for this observation.




Figure 1. The quantum circuit diagram outlining the implementation of the algorithm that raises an unknown unitary \mathcal{U} to some power t . Stage I performs eigenvalue estimation, where value ℓ in the first register corresponds to the eigenvalue parameter estimate $\lambda = \ell/2^m$. Stage II then uses this eigenvalue estimate to induce a corresponding phase shift $e^{2\pi i\ell t/2^m}$. Stage III uncomputes the eigenvalue estimation step. In the case that the eigenvalue parameters λ_k are of the form $\ell_k/2^m$, the eigenvalue estimation is exact and the uncomputation step returns the control register back to the initial all-zeros state. In general, when we do not have this assumption on the form of the eigenvalues, we perform a slightly more complicated eigenvalue estimation algorithm that gives an estimate with error at most $1/2^m$ with probability in $1 - O(1/2^m)$.

2. The algorithm

In this section, we describe our basic construction. For any unitary \mathcal{U} on a finite-dimensional state space, consider its spectral decomposition $\mathcal{U} = P\Lambda P^\dagger$, where P is a unitary matrix composed of the eigenvectors of \mathcal{U} and Λ is a diagonal matrix containing the eigenvalues of \mathcal{U} . For such a decomposition, powers of \mathcal{U} may be computed as follows:

$$\mathcal{U}^t = P\Lambda^t P^\dagger. \tag{1}$$

A key observation is that in the quantum case, one does not need to actually implement the basis change operators P and P^\dagger in order to exploit the above feature of the spectral decomposition.

For now, let us assume that t is a real number between 0 and 1.

Our algorithm is described in three stages. In the first stage, approximations to the eigenvalues of \mathcal{U} are calculated using an eigenvalue estimation algorithm. Let us initially assume that the eigenvalues are of the form $e^{2\pi i\lambda_k}$ with $\lambda_k = \frac{\ell_k}{2^m}$ for some integer $\ell_k \in \{0, 1, \dots, 2^m - 1\}$ (so that we can initially ignore the effect of precision errors). In the second stage, phase shifts are applied to the eigenstates of \mathcal{U} . The third stage uncomputes the eigenvalue estimation. Refer to figure 1 for details.

Suppose that $|\Psi\rangle$ is a quantum state on which we wish to perform the transformation \mathcal{U}^t . Note that any state $|\Psi\rangle$ is a superposition of eigenvectors, $|\psi_k\rangle$, of \mathcal{U} , $|\Psi\rangle = \sum_k \alpha_k |\psi_k\rangle$. Prepare m ancilla qubits in the ‘all-zeros’ state, where m is a user chosen parameter that characterizes the complexity of and errors in our algorithm; its role is discussed in more detail in appendix B. These ancilla are input to the quantum Fourier transform (QFT) circuit [10] (or equivalently, since applied to the all-zeros state, a circuit consisting of only Hadamard operations on each qubit) and later they are used as the controls for controlled- \mathcal{U} operators, while the state $|\Psi\rangle$ is the target. The notation $c - \mathcal{U}^j$ denotes the operation $|j\rangle|\phi\rangle \mapsto |j\rangle\mathcal{U}^j|\phi\rangle$. Using the eigenstate expansion of $|\Psi\rangle$ in the target, it can be shown that such a controlled operation causes a phase kick-back to the control register, so that the total state becomes $\sum_k \sum_j \alpha_k e^{2\pi i\lambda_k j} |j\rangle|\psi_k\rangle$. Next, an inverse QFT operation is performed on the ancilla register. The result left on the ancilla register is estimates of the eigenvalue parameter, λ_k , specifically $\ell_k = 2^m \lambda_k$, in superposition (as dictated by the input vector $|\Psi\rangle$). For a more detailed description of the eigenvalue estimation algorithm, see [10].

For stage II, we apply the controlled operation $c - \mathcal{R}(2\pi\ell_k t/2^m)$ to induce the phase shift $e^{2\pi i\ell_k t/2^m}$ when the eigenvalue parameter estimate ℓ_k is in the control register. This

computation corresponds to the exponentiation of the diagonal matrix in the spectral decomposition formula. There are a few natural ways to implement this step, as shown in [11] or [12].

Finally, the state of the control register (containing the ℓ_k values) is uncomputed back to the all-zeros (input ancilla) state. This completes the third and final stage of the computation.

This leaves the registers in the final state $|\mathbf{0}\rangle \otimes \sum_k \alpha_k e^{2\pi i \lambda_k t} |\psi_k\rangle$, which is simply the result of the application of \mathcal{U}^t to the state $|\Psi\rangle$. When we drop the assumption that the eigenvalue parameters λ_k are exact ratios of some integers ℓ_k divided by 2^m , we will in general get precision errors in the eigenvalue estimation. This imprecision in the resolution of λ_k means that the phases applied to the state are not exact, which further implies that the uncomputation step does not return the ancilla register precisely to all zeros. Still, by a proper choice of parameters, and a ‘gap’ assumption discussed below, the errors can be managed—in particular, we prove that they are exponentially small in the parameter m . This error reduction is done by choosing an eigenvalue estimation algorithm that outputs an estimate with error at most $\frac{1}{2^m}$ with probability in $1 - O(\frac{1}{2^m})$. This can be done with $O(m)$ repetitions of the standard QFT-based eigenvalue estimation algorithm [10] and applying the Chernoff bound. This uses $O(m2^m)$ calls to the black box. Relevant calculations may be found in appendix B.

2.1. Underlying assumptions

Our algorithm, as formulated in the previous section, asserts that in addition to the black box implementation of \mathcal{U} itself we also have access to a black box implementing controlled- \mathcal{U} and a black box implementing \mathcal{U}^{-1} . Furthermore, we also must assume, as discussed in more detail in the following section, that the spectrum of matrix \mathcal{U} has a small gap, in particular that \mathcal{U} has no eigenvalue $\lambda = e^{i\phi}$, where $\phi \in (2\pi(1-g), 2\pi)$ for some value g . For simplicity, we assume that $g \geq \frac{1}{2^m}$; otherwise, we would need to replace some of the $O(2^m)$ terms with $O(\max\{2^m, \frac{1}{g}\})$ (the reason is that we need to do eigenvalue estimation with precision less than the gap size with high probability). In this case, we achieve an approximation with error in $O(\frac{1}{2^m})$. Note that there are no other assumptions about gaps elsewhere in the spectrum. Furthermore, an assumption of this form is essentially necessary in the case of worst-case complexity (as we explain in section 2.2), which is what we are interested in. Average-case performance is discussed more in appendix D.

Not all of the above restrictions are necessary, but they facilitate a clear and transparent description of the algorithm and its analysis. In some cases of interest, the individual restrictions may be dropped. In particular, in cases where only the black box implementation of \mathcal{U} is given (and not the controlled- \mathcal{U}), one can exploit the technique for implementing a controlled- $e^{-i\phi}U$, where $e^{i\phi}$ is a random eigenvalue of \mathcal{U} , and use it in the algorithm to find arbitrary powers, as outlined in [13] (see appendix A for details). In some special cases, we can remove the assumption of having a black box for \mathcal{U}^{-1} . For example, for non-integer values of $t \geq 2^m$, we can approximate \mathcal{U}^t with error in $O(\frac{1}{t})$ using $O(t)$ calls to the controlled- \mathcal{U} , without any use of a black box for \mathcal{U}^{-1} (see appendix C). We note that the above restriction on the form of the spectrum is naturally satisfied in many practical unitaries, including the quantum Fourier transform, standard oracles $|j\rangle \mapsto (-1)^{f(j)}|j\rangle$ for computing Boolean functions $f(\cdot)$ and oracles such as those implicitly used in the adiabatic algorithms of Farhi *et al* [14] that reversibly compute a function with a finite discrete spectrum.

It is also worth noting that, for non-integer values of the power t , there are many operators that one might reasonably call the t th power of \mathcal{U} . In other words, for any given \mathcal{U} , there are many Hamiltonians H such that $\mathcal{U} = e^{-iH}$, and thus one could naturally define \mathcal{U}^t as e^{-iHt} for any one of these Hamiltonians. For instance, consider $t = \frac{1}{2}$ and the identity matrix $\mathbb{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$.

Let us find all unitary matrices B such that $BB = \mathbb{I}$. In the most natural understanding, each such B is a possible square root of \mathbb{I} . First, matrices \mathbb{I} and $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$ are both square roots of \mathbb{I} . This happens because there are two possible square roots of the complex number 1, which is the eigenvalue of \mathbb{I} . In addition, each eigenspace of a unitary matrix may be broken into subspaces such that different square roots of the corresponding eigenvalue may be used to construct square roots of a given matrix. In our case, this helps to construct two more square roots of the form $\pm \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. Furthermore, every matrix of the form $\pm \begin{pmatrix} \cos a & e^b \sin a \\ e^{-b} \sin a & -\cos a \end{pmatrix}$, and its transpose, where a and b are real-valued parameters, is a square root of the operator \mathbb{I} . For the choice of parameters $a = \frac{\pi}{2}$ and $b = 0$ this allows us to construct root $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, also known as the Pauli-X matrix. Roots of the identity matrix in higher dimensions get more complicated. Thus, the problem of finding fractional powers is not completely straightforward to define. In our solution, we construct fractional powers based on the spectral decomposition, choosing the primitive complex fractional power of every eigenvalue (though, in principle, our method allows us to break the eigenspaces into subspaces and take different roots of the eigenvalue for each subspace). Thus, in a sense, the fractional power we construct is the most natural one, and also it is the fractional power of a matrix most commonly defined in linear algebra textbooks. With the above example, our algorithm implements the \mathbb{I} root of \mathbb{I} , though modifications are possible that would allow us to explore a wider range of square roots of \mathbb{I} .

2.2. Optimality

For such a general black box and assuming a spectral gap of size $g = \frac{1}{2^m}$, we cannot compute $\mathcal{U}^{\frac{1}{2}}$ with constant error with fewer than $\Omega(2^m)$ calls to \mathcal{U} . To see why this is the case, suppose we are given either $\mathcal{U} = \mathbb{I}$ or $\mathcal{U} = \mathcal{U}_{1-1/2^m} = |0\rangle\langle 0| + e^{-2\pi i/2^m} |1\rangle\langle 1|$. It follows from the following simple lemma that $\Omega(2^m)$ calls to \mathcal{U} are necessary to correctly guess which \mathcal{U} we were given with probability greater than $\frac{2}{3}$.

Lemma 1. *Suppose we are given a black box \mathcal{U} that implements either \mathbb{I} or $\mathcal{U}_{1-\delta} = |0\rangle\langle 0| + e^{-i\delta} |1\rangle\langle 1|$, for some small $\delta > 0$. Any algorithm that correctly guesses the identity of \mathcal{U} with probability at least $\frac{2}{3}$, for any prior distribution of the two possible values of \mathcal{U} , must make $\Omega(\frac{1}{\delta})$ evaluations of \mathcal{U} .*

This can be easily proved, e.g., by the ‘hybrid’ method [15].

Note that the root $\mathcal{U}_{1-1/2^m}^{\frac{1}{2}}(|0\rangle + |1\rangle) = |0\rangle - e^{-\pi i/2^m} |1\rangle$ is almost perfectly distinguishable from $\mathbb{I}^{\frac{1}{2}}(|0\rangle + |1\rangle) = |0\rangle + |1\rangle$. Thus, the power finding algorithm must take $\Omega(2^m)$ calls to \mathcal{U} in order to compute $\mathcal{U}^{\frac{1}{2}}$ with (small) constant error. The same holds for any power $t \in (0, 1)$ that is bounded away from both 0 and 1.

Thus, while more efficient ways to compute fractional powers of \mathcal{U} are possible in special cases, in general one cannot do much better than what we have described. In the case when $\frac{1}{2^m} < g < \pi$, we can still obtain a lower bound of $\Omega(2^m)$ queries by letting \mathcal{U} be either Z or $|0\rangle\langle 0| - e^{-2\pi i/2^m} |1\rangle\langle 1|$, and computing square roots of \mathcal{U}^2 as a means of distinguishing the two cases.

This same argument also shows why a gap assumption is necessary. Without the gap assumption, there is no upper bound on the worst-case complexity of computing the square root of a black box unitary, since such a black box could be used to distinguish \mathbb{I} from $\mathcal{U}_{1-\delta} = |0\rangle\langle 0| + e^{-i\delta} |1\rangle\langle 1|$ for arbitrarily small values of δ . As pointed out earlier, if we have a known gap g without a promise that $g \in \Omega(\epsilon) = \Omega(1/2^m)$, then a complexity of

$O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$ suffices, and the example above shows that $\Omega\left(\frac{1}{g}\right)$ queries are necessary. If g is unknown, then there are several cases one might consider. If we are promised a lower bound $0 < g_{\min} \leq g$, then in the worst case, we must use $O\left(\frac{1}{g_{\min}} \log \frac{1}{\epsilon}\right)$ calls to the unitary. If instead we have a means of reliably recognizing when the algorithm has succeeded, then trying a sequence of lower bounds that decreases exponentially until the algorithm succeeds will also give a complexity in $O\left(\frac{1}{\min\{g, \epsilon\}} \log \frac{1}{\epsilon}\right)$, which is the same complexity as when g is known. If, however, we do not know g and have no means of recognizing when the algorithm has succeeded, then no upper bound is possible.

3. Large powers and an example of exponential improvement

When it comes to computing large powers of \mathcal{U} by this method, precision in the eigenvalue estimates becomes a limiting factor. If instead of the exact eigenvalue $e^{2\pi i \lambda_k}$ we have its estimate $e^{2\pi i \tilde{\lambda}_k}$, where $|\tilde{\lambda}_k - \lambda_k| = \epsilon$, then mapping $|\psi_k\rangle \mapsto e^{2\pi i \tilde{\lambda}_k t} |\psi_k\rangle$ will lead to an error in the phase parameter of size $t\epsilon$. For arbitrary \mathcal{U} , it takes $O\left(\frac{1}{\epsilon}\right)$ calls to \mathcal{U} in order to get an estimate with error at most ϵ with high probability. Thus to achieve an error of $\delta = t\epsilon$ in the application of \mathcal{U}^t , we need $\epsilon = \delta/t$ precision in the estimate, amounting to $O\left(\frac{t}{\delta}\right)$ calls to \mathcal{U} . For $t > 1$, a better thing to do in this general case is to first apply \mathcal{U} a total of $\lceil t \rceil$ times, and then use the method from the previous section in order to compute the remaining fractional power of \mathcal{U} .

However, if for special cases of \mathcal{U} we are able to get smaller errors in the eigenvalue estimates, we can do much better. For example, for black boxes that map $|j\rangle \mapsto (-1)^{X_j} |j\rangle$ for some $j \in \{0, 1\}^n$ and $X_1 X_2 \dots X_j \dots X_N \in \{0, 1\}^N$, it is easy to exactly determine the eigenvalue parameter X_j with only one call to the black box. As discussed later, it is easy to exactly determine the eigenvalue of an eigenvector of the quantum Fourier transform, since the QFT has order 4. In both of these examples, our technique does not help with computing higher powers of \mathcal{U} since \mathcal{U} has very low order, say r , and thus computing \mathcal{U}^t is equivalent to computing $\mathcal{U}^{t \bmod r}$, where $t \bmod r$ is the unique element of $t - r\mathbb{Z}$ that lies in the interval $[0, r)$. We next construct an example where \mathcal{U} has very high order, but one can still compute precise estimates of its eigenvalues with relatively few calls to \mathcal{U} , and thus exponentiate \mathcal{U} with much fewer than t calls, even for t less than the order of \mathcal{U} .

Suppose the eigenvalues of the desired unitary \mathcal{U} are all of the form $e^{2\pi i \frac{\ell_k}{p_k}}$, for $k \in \{1, 2, \dots, b\}$, where ℓ_k is an integer, $p_k \in \{p_1 = 2, p_2 = 3, p_3 = 5, \dots, p_b | p_i = i\text{th prime}\}$ and b is a natural number. Then $\mathcal{U}^B = \mathbb{I}$, where $B = \prod_{k=1}^b p_k$, and (assuming that each p_k occurs at least once with a non-zero ℓ_k) no lower power of \mathcal{U} is equal to the identity. If t is on the order of B , then the number of calls to the unitary required for the straightforward implementation of \mathcal{U}^t is in $O(B) = O\left(\prod_{k=1}^b k \log k\right)$. This product is called a *primorial* [16], and it is in $O(e^{(1+o(1))b \log b}) = O(b^{(1+o(1))b})$.

We next demonstrate how our basic algorithm introduced in section 2 may be updated to approximate \mathcal{U}^t with exponentially fewer calls to \mathcal{U} than the number of calls it takes to obtain \mathcal{U}^t via a direct application of t unitaries.

The method used for this problem employs the continued fraction algorithm, which was also used in the order-finding part of Shor's algorithm [17]. The eigenvalue estimation algorithm will return an estimate $\frac{h}{2^m}$ of the eigenvalue parameter $\frac{\ell_k}{p_k}$ to within an error at most $\frac{1}{2^m}$ with probability at least $\frac{8}{\pi^2}$. The error probability can be amplified down to $\frac{1}{2^m}$ by repeating the estimation $O(m)$ times and taking the majority vote. If m is chosen such that $m \in O(\log b)$, so that $2^m > 2p_b p_{b-1}$, then the fraction $\frac{\ell_k}{p_k}$ will be the only fraction, or *convergent*, in the continued fraction expansion of $\frac{h}{2^m}$ that has distance at most $\frac{1}{2^m}$ from $\frac{h}{2^m}$ and has a denominator at most p_b .




Figure 2. A diagram illustrating the procedure. The estimate is $\frac{h}{2^m}$. The smallest gap between any two eigenvalues is $\frac{1}{p_b p_{b-1}}$ and we require this to be larger than $\frac{2}{2^m}$ to ensure that only one rational fraction with the denominator no greater than p_b will be found in the region $\frac{h-1}{2^m}$ to $\frac{h+1}{2^m}$.

This convergent $\frac{\ell_k}{p_k}$ can be found with $O(m)$ arithmetic operations over integer numbers with at most $\log m$ bits each (e.g., no more than $O(m^3)$ binary operations using trivial algorithms for division and multiplication) using the continued fraction algorithm. Once we have the exact value of the eigenvalue parameter $\frac{\ell_k}{p_k}$, with probability $\frac{1}{2^m}$, we can correctly compute any positive power t of \mathcal{U} with error in $O(\frac{1}{2^m})$ using $\text{poly}(b, \log t)$ calls to \mathcal{U} and other elementary operations. This is done by mapping, in superposition, $|\ell_k\rangle|p_k\rangle|\psi_k\rangle \mapsto e^{2\pi i t \frac{\ell_k}{p_k} + O(\frac{1}{2^m})} |\ell_k\rangle|p_k\rangle|\psi_k\rangle$, which can be done by standard reversible computing methods and m -bit arithmetic operations (the $O(\frac{1}{2^m})$ error is due to precision in the arithmetic operations). Then, as usual, one uncomputes the computation of $|\ell_k\rangle|p_k\rangle$. As alluded to before, this gives an overall complexity of $O(m2^m) = O(b \log b) (\subseteq O(2^m))$ to perform the operation up to error $O(\frac{1}{2^m})$. This follows from (B.11) in appendix B.

This complexity is polynomial in $\log t$ (versus $\text{poly}(t)$) because we only need to compute an m -bit approximation to $t\ell_k/p_k$. This also allows us to uncompute the eigenvalue estimates found in the first stage of the algorithm without ever needing to use the operation \mathcal{U}^{-1} . Note that t could be very large, for example $t \sim B/2 = 2^{b-1}$. In this case, we can apply \mathcal{U}^t to any state of appropriate dimension with the output state containing an error term with probability amplitude $O(\frac{1}{b})$ with only $O(b \log b) \ll t \in O(2^b)$ calls. The error probability can efficiently be made exponentially small by repeating the phase estimation part.

4. Fractional quantum Fourier transform

In classical computation, roots of the Fourier transform are used for efficient filtering of the frequency noise that is some non-constant function of a conjugate variable, such as time [18]. Given the role the QFT plays in the quantum computation, and the importance of roots of Fourier transforms in classical computation, there may as well be other interesting techniques based on the fractional powers of the quantum Fourier transform.

Note that $\text{QFT}^4 = \mathbb{I}$. Let us illustrate the computation of a fractional power $t, t < 1$ of the QFT using the algorithm introduced in section 2. First, note that the QFT has four distinct eigenvalues: $+1, i, -1$ and $-i$. This means that we need only two ancilla bits for the eigenvalue estimation part of the exponentiation algorithm, and such estimation results in the exact values. During the eigenvalue estimation algorithm, the controlled-QFT is applied three times ($2^0 = 1$ with the first ancilla qubit as a control and $2^1 = 2$ more with the second ancilla as a control). We need three more applications of the controlled-QFT in the third stage of the algorithm to uncompute ancilla. This means that the query complexity of the computation of a fractional power of the QFT is exactly six queries to the controlled-QFT. Moreover, during such computation we introduced no errors due to the approximations of eigenvalues, and thus, given all gates are perfect, the computed fractional power is exact.

This specific example of fractional QFT was alternatively derived as a special case of the work of Klappenecker and Rötteler [19], where they address the different problem of

how to simulate a known unitary matrix A given other operators that generate an algebra containing A .

5. Conclusions and discussion

In this work we have further developed the connection between continuous time black boxes (i.e., ‘Hamiltonian oracles’) and their discrete counterparts, unitary black boxes. Previous work has mostly focused on important special cases of this relationship. Our work pushes the boundary in a new direction, where the unknown \mathcal{U} is a unitary acting on an arbitrarily large Hilbert space. It also demonstrates constructively that an algorithm with a continuous oracle can also be implemented with a discrete oracle that implements the continuous oracle for some fixed time (with some assumptions on the Hamiltonian encoded by the oracle; e.g., if the energies, the eigenvalues of the Hamiltonian, are between 0 and E_{\max} , then the fixed time interval should be bounded below $2\pi/E_{\max}$ so that our gap assumption is satisfied).

For the case that t is a positive non-integer, the complexity of this algorithm is substantially higher than that in the continuous setting, where we have access to a Hamiltonian H satisfying $\mathcal{U} = e^{-iH}$. If we have direct access to the Hamiltonian, H , then in principle we can easily perform any unitary of the form e^{-iHt} with a cost t in terms of calls to the Hamiltonian. For an algorithm with total query complexity $t < 1$ in the Hamiltonian black box model, it is hard to imagine a reasonable way to effect e^{-iHt} with fewer than one evaluation of \mathcal{U} and indeed our algorithm makes $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ calls to \mathcal{U} to effect e^{-iHt} with precision ϵ (we show a corresponding lower bound in section 2.2), which is notably lower than the number of calls required in the obvious approach of doing process tomography on \mathcal{U} , approximating \mathcal{U}^t and then synthesizing a circuit to implement such an approximation. In particular, our approach does not depend on the dimension of the Hilbert space that \mathcal{U} acts on.

We have presented an algorithm for finding the powers of unknown unitary operations raised to any real power $t > 0$. We found the complexity of the algorithm to be in $O(\lfloor t \rfloor + 2^m m)$ where the error is in $O(\frac{1}{2^m})$. For large integers t , we presented a non-trivial example where our method is exponentially more efficient than the direct repeated application of \mathcal{U} a total of t times. Note that the same exponential speed-up is also possible in the continuous case if the Hamiltonian has some structure that enables eigenvalue estimation much better than the standard worst-case limits; in such cases, our procedure, using a modest amount of ancilla workspace, can simulate the evolution of H for a time τ in a time polynomial in $\log(\tau)$.

Our algorithm could be of practical use in a situation where we have very precise clocks and control, say with error less than a very small positive ϵ , but for some technical reason we must run the unknown Hamiltonian for a minimum time $T_{\min} \gg \epsilon$ (e.g., because of a minimum reset time on some critical piece of apparatus, such as a detector, used to simulate or execute the Hamiltonian), and T_{\min} is still smaller than $2\pi/E_{\max}$ (assuming the energies are non-negative). In this case, there is a smallest amount of time for which a given Hamiltonian can naturally be run. If we wish to run the Hamiltonian for less than that amount of time, then our method could be used.

Further, if the Hamiltonian that solves some problem is time dependent, as is often the case in quantum circuits or in models based on the permutation or braiding of particles on a surface, it may be of interest to apply some fraction of this total transformation to turn the solution to the problem into a subroutine of a larger algorithm. The method we propose provides a general way of obtaining this fractional application.

Our technique may also be applied to compute other functions of \mathcal{U} , by replacing each eigenvalue $e^{i\lambda_k}$ with $e^{if(\lambda_k)}$ for any function $f(\cdot)$. This may be useful in contexts such as the Childs method [20] for simulating continuous time quantum walk Hamiltonians with discrete

time walks. This requires applying phases that are the sines or arcsines of the eigenvalues of the original Hamiltonian. Our method could achieve this for an unknown Hamiltonian to be simulated. The properties of the function $f(\cdot)$ will determine the efficiency as a function of the precision. In the case when $f(\cdot)$ is continuous with a bounded first derivative, and 2π -periodic, then we can also drop the assumption that there must be a gap in the spectrum.


In addition, there are a number of classical cryptography schemes that rely on the computational difficulty of finding roots modulo some number, such as RSA. One might naturally attempt to construct quantum cryptography schemes where knowledge of the ‘secret key’ corresponds to the ability to compute a fractional power (i.e., a root) of some unitary operator. This algorithm provides a way of implementing such roots, implying serious limitations for such an approach in quantum cryptographic schemes.

Acknowledgments

We thank Scott Aaronson, Patrick Hayden, Rolando Somma and John Watrous for discussions about the relevance of this problem. Dr Aaronson, Dr Hayden and Dr Somma were interested in black box square root question, and Dr Watrous sketched a solution similar to the one that we had. We thank Dr Donny Cheung at the University of Calgary for helpful insights. This research has been supported by DTO-ARO, NSERC, CFI, CIFAR, CRC, ORF, Ontario-MRI, QuantumWorks and the Government of Canada. L Sheridan gratefully acknowledges the support of the Mike and Ophelia Lazaridis Fellowship. This paper was based on work partially supported by the National Science Foundation, during D Maslov’s assignment at the Foundation. Any opinion, finding and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

Appendix A. Controlled unitaries

Note that the eigenvalue estimation algorithm requires the use of a controlled unitary $c - \mathcal{U}$. If we are given an eigenvector of \mathcal{U} with eigenvalue +1, then this transformation can be implemented using a controlled-SWAP (equivalently, a series of Fredkin gates [21]), as was shown by Kitaev [13] and is illustrated in the following circuit diagram:



It is easy to verify that the above circuit implements

$$|0\rangle_{\text{control}}|q_2\rangle_{\text{target}}|\psi_k\rangle \mapsto e^{2\pi i\lambda_k}|0\rangle_{\text{control}}|q_2\rangle_{\text{target}}|\psi_k\rangle$$

$$|1\rangle_{\text{control}}|q_2\rangle_{\text{target}}|\psi_k\rangle \mapsto |1\rangle_{\text{control}}(\mathcal{U}|q_2\rangle_{\text{target}})|\psi_k\rangle.$$

In other words, up to a global phase, this implements the controlled- $(e^{-2\pi i\lambda_k}\mathcal{U})$. If $\lambda_k = 0$, i.e., $|\psi_k\rangle$ has eigenvalue +1, then we have implemented $c - \mathcal{U}$. Otherwise, we have a very similar operation, which may or may not be sufficient, depending on the application. If, for example, any $c - (e^{i\phi}\mathcal{U})$ will suffice, as long as we use the same one consistently, then this is easy to achieve.

We do not need to assume that we know the eigenvectors $\{|\psi_k\rangle\}$ of \mathcal{U} or are given a specific eigenvector. Instead, we can use any state in the target register. For simplicity, we can consider the completely mixed state, which can be decomposed as an equal mixture over the eigenstates. Provided we keep the state that is in the target register, and reuse the resulting state every time we construct a controlled- \mathcal{U} between these two registers, the phase (though random) will be the same for the entire computation. That is, this is equivalent to implementing $c - (e^{i\phi}\mathcal{U})$ consistently throughout the computation, where ϕ is an unknown random value (whose distribution depends on the weight of the eigenspace of $e^{-i\phi}$ in the initial mixed state). In essence, the target system serves as a phase reference.

Appendix B. Complexity and error analysis

We will describe upper bounds on the precision of our algorithm, which is constrained by the lack of perfect precision in computing the eigenvalue estimates. We give the analysis for $t = \frac{1}{2}$; however, the same final bound applies for any $t \in (0, 1)$.

Let us first describe an ‘ideal’ algorithm that exactly computes the square root of \mathcal{U} . For the ideal algorithm, some of the steps will be unnecessary, but we include them to make the comparison with the actual algorithm simple.

Stage I computes a near-optimal eigenvalue estimation (in the sense of nearly optimizing the chance of obtaining an estimate with error at most $\frac{1}{2^m}$). We do this by repeating the ‘standard’ eigenvalue estimation algorithm [11] a total of $r \in O(m)$ times and take the majority answer, with the constant being chosen so that each eigenvalue estimate $e^{2\pi i \tilde{\lambda}_k}$ of $e^{2\pi i \lambda_k}$ satisfies $|e^{2\pi i \tilde{\lambda}_k} - e^{2\pi i \lambda_k}| \leq \frac{1}{2^m}$ with probability in $1 - O(\frac{1}{2^m})$. One can fine-tune the optimal eigenvalue estimation procedure further (e.g., [22]); however, this simple procedure is sufficient.

This eigenvalue estimation procedure is described by the transformation

$$|0\rangle \sum_k \alpha_k |\psi_k\rangle \mapsto \sum_k \alpha_k \left(\sum_{\mathbf{y}} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle \right) |\psi_k\rangle, \quad (\text{B.1})$$

where the values $\mathbf{y} \in \mathbb{Z}_{2^m}^r$ (where $\mathbb{Z}_{2^m} = \{0, 1, \dots, 2^m - 1\}$) are r -tuples of integers and $\lambda_{\mathbf{y}}$ is the value obtained by taking the most commonly occurring integer y_{mode} in the r -tuple of integers \mathbf{y} (with some convention for breaking ties) and letting $\lambda_{\mathbf{y}} = y_{\text{mode}}/2^m$.

In the ideal stage II, we exactly implement the map

$$\sum_k \alpha_k \left(\sum_{\mathbf{y}} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle \right) |\psi_k\rangle \mapsto \sum_k \alpha_k \left(\sum_{\mathbf{y}} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle \right) e^{2\pi i \lambda_k/2} |\psi_k\rangle. \quad (\text{B.2})$$

Then stage III uncomputes the eigenvalue estimation exactly (since with the ideal stage II, there is no coupling of the control registers with the target registers), leaving us with the ideal output

$$\sum_k \alpha_k |00 \dots 0\rangle e^{2\pi i \lambda_k/2} |\psi_k\rangle. \quad (\text{B.3})$$

The only difference between this ideal algorithm and our actual algorithm is that stage II actually effects the unitary

$$\sum_k \alpha_k \left(\sum_{\mathbf{y}} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle \right) |\psi_k\rangle \mapsto \sum_k \alpha_k \sum_{\mathbf{y}} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle e^{2\pi i \lambda_{\mathbf{y}}/2} |\psi_k\rangle. \quad (\text{B.4})$$

Since the probability that $|e^{2\pi i\lambda_y/2} - e^{2\pi i\lambda_k/2}| > \frac{1}{2^m}$ is in $O\left(\frac{1}{2^m}\right)$ (note that here we use the ‘gap’ assumption), we can write the state in stage II as

$$|\Phi\rangle = \sum_k \alpha_k \sum_{\mathbf{y} \in S} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle e^{2\pi i\lambda_{\mathbf{y}}/2} |\psi_k\rangle + \sum_k \alpha_k \sum_{\mathbf{y} \in S'} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle e^{2\pi i\lambda_{\mathbf{y}}/2} |\psi_k\rangle, \quad (\text{B.5})$$

where the values of $\mathbf{y} \in S$ include the ‘good’ values of \mathbf{y} that produce estimates $\lambda_{\mathbf{y}}$ satisfying $|e^{2\pi i\lambda_{\mathbf{y}}} - e^{2\pi i\lambda_k}| \leq \frac{1}{2^m}$ and the ‘bad’ values, $\mathbf{y} \in S'$, are the rest. The norm of the bad part of the state is in $O\left(\frac{1}{2^{2m}}\right)$ and the norm of the ‘good’ part is $1 - O\left(\frac{1}{2^{2m}}\right)$.

Let Λ denote the ideal phase shift operator and $\tilde{\Lambda}$ denote the actual operator.

We thus have

$$\begin{aligned} \Lambda|\Phi\rangle - \tilde{\Lambda}|\Phi\rangle &= \sum_k \alpha_k \sum_{\mathbf{y} \in S} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle (e^{2\pi i\lambda_k/2} - e^{2\pi i\lambda_{\mathbf{y}}/2}) |\psi_k\rangle \\ &\quad + \sum_k \alpha_k \sum_{\mathbf{y} \in S'} \beta_{\mathbf{y},k} |\mathbf{y}\rangle |\lambda_{\mathbf{y}}\rangle (e^{2\pi i\lambda_k/2} - e^{2\pi i\lambda_{\mathbf{y}}/2}) |\psi_k\rangle. \end{aligned}$$

Thus we can bound the norm of this difference by

$$\|\Lambda|\Phi\rangle - \tilde{\Lambda}|\Phi\rangle\|^2 \leq \sum_k \sum_{\mathbf{y} \in S} |\alpha_k \beta_{\mathbf{y},k}|^2 |\delta_{k,\mathbf{y}}|^2 + \sum_k \sum_{\mathbf{y} \in S'} |\alpha_k \beta_{\mathbf{y},k}|^2 |\delta_{k,\mathbf{y}}|^2, \quad (\text{B.6})$$

where $|\delta_{k,\mathbf{y}}|^2 = |e^{i\lambda_k/2} - e^{i\lambda_{\mathbf{y}}/2}|^2$.

We know that $\sum_{\mathbf{y} \in S} \sum_k |\alpha_k \beta_{\mathbf{y},k}|^2 \in 1 - O\left(\frac{1}{2^{2m}}\right)$, $\sum_{\mathbf{y} \in S'} \sum_k |\alpha_k \beta_{\mathbf{y},k}|^2 \in O\left(\frac{1}{2^{2m}}\right)$, and that $|\delta_{k,\mathbf{y}}|^2 \in O\left(\frac{1}{2^{2m}}\right)$ for the good values, and at most 1 for the bad values.

Thus

$$\|\Lambda|\Phi\rangle - \tilde{\Lambda}|\Phi\rangle\|^2 \in O\left(\frac{1}{2^{2m}}\right). \quad (\text{B.7})$$

Noting that

$$\| |u\rangle - |v\rangle \|_2^2 = (\langle u| - \langle v|)(|u\rangle - |v\rangle) \geq 2(1 - |\langle u|v\rangle|) \quad (\text{B.8})$$

and using the equality

$$\| |u\rangle\langle u| - |v\rangle\langle v| \|_{\text{Tr}} = 2\sqrt{1 - |\langle u|v\rangle|}, \quad (\text{B.9})$$

which implies that

$$\| |u\rangle\langle u| - |v\rangle\langle v| \|_{\text{Tr}} \leq 2\| |u\rangle - |v\rangle \|_2, \quad (\text{B.10})$$

$\|\Lambda(|\Phi\rangle\langle\Phi|) - \tilde{\Lambda}(|\Phi\rangle\langle\Phi|)\|_{\text{Tr}} \in O\left(\frac{1}{2^m}\right)$. The trace norm of the difference between the two unitary superoperators is simply given by taking the maximization over states $|\Phi\rangle$. Since this holds for any state $|\Phi\rangle$, we have

$$\|\Lambda - \tilde{\Lambda}\|_{\text{Tr}} \in O\left(\frac{1}{2^m}\right). \quad (\text{B.11})$$

This is the case after stage II of the algorithm. Note that stage III is a unitary operation that is the same for both the ideal and the actual cases. The trace norm is invariant under unitary transformations, so the trace norm of the unitary operators for the full ideal and actual algorithms have the same bound. In the case of unitary operators, this is equivalent to the diamond norm.

Appendix C. Inverses

Our algorithm for generating \mathcal{U}^t uses \mathcal{U}^{-1} when the eigenvalue estimates need to be uncomputed. More specifically, we start with $|0\rangle|\psi_k\rangle$, and we apply $(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})c - \mathcal{U}^j(\text{QFT}_{2^m} \otimes \mathbb{I})$ to compute

$$\sum_j \left(\text{QFT}_{2^m} \frac{e^{2\pi i j \lambda_k}}{\sqrt{2^m}} |j\rangle \right) |\psi_k\rangle. \quad (\text{C.1})$$

We then use the value in the first register to control the phase shift corresponding to $|\psi_k\rangle$. For simplicity, let us assume that we do not bother to induce a phase shift, but we are still left with the task of uncomputing the ‘damage’ we have done. It is trivial if we have access to $c - \mathcal{U}^{-1}$ (which can be implemented with a $c - \mathcal{U}$ and \mathcal{U}^{-1}), since we just apply $(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})c - \mathcal{U}^{-j}(\text{QFT}_{2^m} \otimes \mathbb{I})$ to get back $|0\rangle|\psi_k\rangle$.

However, in some cases, we do not need the \mathcal{U}^{-1} operations. We could attempt to circumvent the need for \mathcal{U}^{-1} operations by applying $(\text{QFT}_{2^m} \otimes \mathbb{I})c - \mathcal{U}^j(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})$ to yield

$$\begin{aligned} & (\text{QFT}_{2^m} \otimes \mathbb{I})c - \mathcal{U}^j(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})c - \mathcal{U}^j(\text{QFT}_{2^m} \otimes \mathbb{I})|0\rangle|\psi_k\rangle \\ &= (\text{QFT}_{2^m} \otimes \mathbb{I})c - \mathcal{U}^j(M \otimes \mathbb{I})c - \mathcal{U}^j(\text{QFT}_{2^m} \otimes \mathbb{I})|0\rangle|\psi_k\rangle, \end{aligned} \quad (\text{C.2})$$

where $M = \text{QFT}_{2^m}^{-2}$, which implies $M|x\rangle = |2^m - x \bmod 2^m\rangle$.

We can easily verify that $(M \otimes \mathbb{I})c - \mathcal{U}^j(M \otimes \mathbb{I}) = c - \mathcal{U}^{2^m-j} \bmod 2^m$, and thus the above state equals

$$\begin{aligned} &= (\text{QFT}_{2^m} \otimes \mathbb{I})(M \otimes \mathbb{I})c - \mathcal{U}^{2^m-j} \bmod 2^m (c - \mathcal{U}^j(\text{QFT}_{2^m} \otimes \mathbb{I})|0\rangle|\psi_k\rangle) \\ &= (\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})c - \mathcal{U}^{2^m}(\text{QFT}_{2^m} \otimes \mathbb{I})|0\rangle|\psi_k\rangle. \end{aligned} \quad (\text{C.3})$$

Note that since $c - \mathcal{U}^{2^m}$ no longer depends on the value of the first register, we can commute it through the QFTs to get

$$c - \mathcal{U}^{2^m}(\text{QFT}_{2^m}^{-1} \otimes \mathbb{I})(\text{QFT}_{2^m} \otimes \mathbb{I})|0\rangle|\psi_k\rangle = c - \mathcal{U}^{2^m}|0\rangle|\psi_k\rangle = e^{2\pi i 2^m \lambda_k}|0\rangle|\psi_k\rangle. \quad (\text{C.4})$$

This phase factor of $e^{2\pi i 2^m \lambda_k}$ is in general a problem when the target register is in a superposition of eigenstates, since different eigenstates will pick up a different phase factor.

In some cases, this is not a problem. One observation is that such a transformation is equivalent to applying \mathcal{U}^{2^m} . Thus, if our goal is to apply \mathcal{U}^t for $t = \lfloor t \rfloor + \delta$, $0 \leq \delta < 1$, where $t \geq 2^m$, we can apply our algorithm for implementing \mathcal{U}^δ , but instead with the above modification. This yields an approximation to $\mathcal{U}^{2^m+\delta}$ with error in $O(\frac{1}{2^m})$. We can then apply \mathcal{U}^{t-2^m} to complete the approximation of \mathcal{U}^t with error in $O(\frac{1}{t}) \subseteq O(\frac{1}{2^m})$.

Also note that if $\lambda_k = \frac{\ell_k}{2^m}$ for an integer ℓ_k , then the phase factors $e^{2\pi i 2^m \lambda_k} = 1$ all equal 1, and thus pose no problem.

This problem can also be remedied in any other cases where λ_k can be determined exactly (or more significantly precisely than error $\frac{1}{2^m}$), since before uncomputing λ_k , one can add an additional phase shift of $e^{-2\pi i 2^m \lambda_k}$ conditional on the value of λ_k . This will eliminate the final unwanted phase factor of $e^{2\pi i 2^m \lambda_k}$ associated with the eigenvector $|\psi_k\rangle$.

Appendix D. Coping with no gap assumption

If one is only interested in the average-case performance (e.g., averaging over all possible input states to $\mathcal{U}^{\frac{1}{2}}$ according to the Haar measure), one is still faced with the potential problem that most or all of the spectrum of \mathcal{U} is in the difficult region (see section 2.2) of $e^{i\phi} \approx 1$, $\pi < \phi \leq 2\pi$.

If one is content with obtaining a square root of $e^{i\theta}\mathcal{U}$ for a random $\theta \in [0, 2\pi)$, then one can show that the average-case error (for any input distribution) is still in $O\left(\frac{1}{2^m}\right)$ for an algorithm that only uses the square root of $e^{i\theta}\mathcal{U}$ once.

For an algorithm that uses $(e^{i\theta}\mathcal{U})^{\frac{1}{2}}$ a total of k times, the error could get magnified to $O\left(\frac{k^2}{2^m}\right)$. Thus, one could select $m' = m + \lceil 2 \log k \rceil$ for each simulation of $(e^{i\theta}\mathcal{U})^{\frac{1}{2}}$.

In order to see why we have a k^2 factor instead of k , we present an algorithm where substituting an exact square root with our approximate square root yields an error in $\Omega\left(\frac{k^2}{2^m}\right)$. The value θ is of the form $2\pi\ell/2^m - \epsilon$, for some integer ℓ and non-negative value $\epsilon \ll 1/2^m$. Consider an operator \mathcal{U} that has every 2^m th root of unity $e^{2\pi i\ell/2^m}$ as an eigenvalue with equal multiplicities. Consider an algorithm that does eigenvalue estimation on the identity operator with precision in $\Theta(1/2^m)$, and uses quantum searching to find an eigenvector with eigenvalue -1 . Clearly, this algorithm should not find such an eigenvalue since all the eigenvalues are $+1$. Now let $U_1 = (e^{i\theta}\mathcal{U})^{-\frac{1}{2}}$ and let U_2 (for simplicity, let us assume it is unitary, even though in practice it will be a map that is almost unitary) be our approximation to $(e^{i\theta}\mathcal{U})^{\frac{1}{2}}$ for the same θ . Note that $U_1(e^{i\theta}\mathcal{U})^{\frac{1}{2}} = I$, but U_1U_2 will have most eigenvalues near $+1$ and proportion $1/2^m$ of the eigenvalues near -1 . Thus, a generalized version of quantum searching (see appendix E) will generate these (close to) -1 eigenvectors with an amplitude in $\Theta\left(\frac{k}{\sqrt{2^m}}\right)$ using only k evaluations of U_1U_2 . Recall that with the ideal $(e^{i\theta}\mathcal{U})^{\frac{1}{2}}$ (instead of U_2), there would be no such eigenvectors for the search algorithm to find; thus the part of the wave function that found these -1 eigenvalues is an error. This implies an error with probability in $\Theta\left(\frac{k^2}{2^m}\right)$.

Appendix E. A generalization of quantum search

Suppose we are given a black box O_ϕ that flags an unknown quantum state $|\phi\rangle \in H_N$. In other words, $O_\phi|\phi\rangle|b\rangle = |\phi\rangle|b \oplus 1\rangle$ and $O_\phi|\phi'\rangle|b\rangle = |\phi'\rangle|b\rangle$ for any $|\phi'\rangle$ orthogonal to $|\phi\rangle$. For brevity, let us directly use the phase flip version of this black box $O_\phi|\phi\rangle = -|\phi\rangle$ and $O_\phi|\phi'\rangle = |\phi'\rangle$ for $\langle\phi|\phi'\rangle = 0$ (see section 8.1 of [10] for a more detailed discussion of the relationship between these black boxes).

Note that if we have a unitary operation A that maps $|00\dots 0\rangle \mapsto \sin(\theta)|\phi\rangle + \cos(\theta)|\phi'\rangle$, where $\langle\phi|\phi'\rangle = 0$, then we can define the slight generalization of a quantum search iterate, $Q = -AU_{00\dots 0}A^{-1}O_\phi$, where $U_{00\dots 0} = I - 2|00\dots 0\rangle\langle 00\dots 0|$ and note that $Q^k|00\dots 0\rangle = \sin((2k+1)\theta)|\phi\rangle + \cos((2k+1)\theta)|\phi'\rangle$. Thus, if we have an algorithm A that ‘guesses’ $|\phi\rangle$ with probability $p = \sin^2(\theta)$, then after $O\left(\frac{1}{\sqrt{p}}\right)$ calls to Q , we obtain $|\phi\rangle$ with probability very close to 1.

In traditional quantum searching, the state $|\phi\rangle$ is known to be a computational basis state, and thus it is very easy to produce a unitary A such that $|\langle\phi|A|00\dots 0\rangle| = \frac{1}{\sqrt{N}}$; in particular, any operator that maps $|00\dots 0\rangle \mapsto \sum_x \frac{1}{\sqrt{N}}|x\rangle$ will do.

However, for an arbitrary quantum state $|\phi\rangle$, it is non-trivial to produce such an initial state. One cannot simply use a maximally mixed state, since amplitude amplification requires the inverse of the operation A that generates the ‘guess’. We could resort to techniques for generating pure states that for practical purposes are ‘random’ (e.g., [23]); however, for this particular application, it is sufficient to use the following idea.

Let A act on an N^2 -dimensional Hilbert space and map $|00\dots 0\rangle|00\dots 0\rangle \mapsto \sum_x \frac{1}{\sqrt{N}}|x\rangle|x\rangle$. Note that for any basis $\{|\phi_j\rangle\}$, this maximally mixed state can be rewritten as $\sum_j \frac{1}{\sqrt{N}}|\phi_j\rangle|\phi_j^*\rangle$ (where the coefficients of $|\phi_j^*\rangle$ are the complex conjugates of those of $|\phi_j\rangle$). Note that the maximally entangled state has equal support for every eigenvector in the

basis of eigenvalues of the black box unitary, which we require for the procedure outlined in this section to succeed. Let us assume, without loss of generality, that $|\phi_0\rangle = |\phi\rangle$. If we apply amplitude amplification, using $Q = -AU_{00\dots 0,0\dots 0}A^{-1}(O_\phi \otimes I)$, then it is easy to verify that

$$Q^k|00\dots 0\rangle|00\dots 0\rangle = \sin((2k+1)\theta)|\phi_0\rangle|\phi_0^*\rangle + \cos((2k+1)\theta) \left(\sum_{j=1}^{N-1} \frac{1}{\sqrt{N-1}} |\phi_j\rangle|\phi_j^*\rangle \right), \quad (\text{E.1})$$

where $\sin(\theta) = \frac{1}{\sqrt{N}}$. Thus, by choosing $k \approx \frac{\pi}{4\theta} \in O\left(\frac{1}{\sqrt{p}}\right)$, we obtain $|\phi_0\rangle|\phi_0^*\rangle$ with high fidelity. For applications where the extra $|\phi^*\rangle$ system will pose a problem, this approach will not work, and techniques for generating pure states that behave like random quantum states could be used instead.

This technique also works if the black box flags a subspace of dimension d . Again, assume without loss of generality that this subspace is spanned by $|\phi_0\rangle, |\phi_1\rangle, \dots, |\phi_{d-1}\rangle$. Then the maximally entangled state can be rewritten as

$$\sin(\theta) \left(\sum_{j=0}^{d-1} \frac{1}{\sqrt{d}} |\phi_j\rangle|\phi_j^*\rangle \right) + \cos(\theta) \left(\sum_{j=d}^{N-1} \frac{1}{\sqrt{N-d}} |\phi_j\rangle|\phi_j^*\rangle \right), \quad (\text{E.2})$$

where $\sin^2(\theta) = \frac{d}{N}$.

Thus $Q^k|00\dots 0\rangle|00\dots 0\rangle \approx \sum_{j=0}^{d-1} \frac{1}{\sqrt{d}} |\phi_j\rangle|\phi_j^*\rangle$ for $k \approx \frac{\pi}{4} \sqrt{\frac{N}{d}}$, assuming $\sqrt{\frac{d}{N}}$ is much less than 1.

Furthermore, if we applied amplitude estimation [24] in this case, we would be approximating the value of $\theta = \arcsin\left(\sqrt{\frac{d}{N}}\right)$, which would allow us to approximate the dimension of the subspace, d .

References

- [1] Chuang I and Nielsen M 1997 *J. Mod. Opt.* **44** 2455–67
- [2] Nielsen M and Chuang I 2000 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [3] Aaronson S 2006 *New J. Phys.* **8** 58
- [4] Farhi E and Gutmann S 1998 *Phys. Rev. A* **57** 2403
- [5] Ioannou L 2002 Continuous-time quantum algorithms: searching and adiabatic computation *Master's Thesis* University of Waterloo
- [6] Roland J and Cerf N J 2003 *Phys. Rev. A* **68** 062311
- [7] Grover L K 1996 *Proc. 28th Annual ACM Symp. on the Theory of Computing* pp 212–9
- [8] Mochon C 2007 *Phys. Rev. A* **75** 042313
- [9] Cleve R, Gottesman D, Mosca M, Somma R and Yonge-Mallo D 2009 Efficient discrete-time simulations of continuous-time quantum query algorithms *Proc. 41st Annual ACM Symp. on Theory of Computing (STOC '09) (Bethesda, USA)* to be published
- [10] Kaye P, Laflamme R and Mosca M 2007 *An Introduction to Quantum Computing* (Oxford: Oxford University Press) pp 110–30
- [11] Cleve R, Ekert A, Macchiavello C and Mosca M 1998 *Proc. R. Soc. A* **454** 339–54
- [12] Zalka C 1998 *Proc. R. Soc. A* **454** 313–22
- [13] Kitaev A 1995 arXiv:quant-ph/9511026
- [14] Farhi E, Goldstone J and Gutmann S 2000 arXiv:quant-ph/0007071
- [15] Bennett C, Bernstein E, Brassard G and Vazirani U 1997 *SIAM J. Comput.* **26** 1510–23
- [16] Fletcher A, Miller J C P, Rosenhead L and Comrie L J 1962 *An Index of Mathematical Tables* vols 1,–2 2nd edn (Reading, MA: Blackwell)
- [17] Shor P 1997 *SIAM J. Comput.* **26** 1484–509
- [18] Almeida L B 1994 *IEEE Trans. Signal Process.* **42** 3084–91

- [19] Klappenecker A and Rötteler M 2003 *Int. J. Found. Comput. Sci.* **14** 777–96
- [20] Childs A M 2008 On the relationship between continuous- and discrete-time quantum walk arXiv:0810.0312v1
- [21] Fredkin E and Toffoli T 1982 *Int. J. Theor. Phys.* **21** 219–253
- [22] van Dam W, D’Ariano G M, Ekert A, Macchiavello C and Mosca M 2007 *Phys. Rev. Lett.* **98** 090501
- [23] Ambainis A and Emerson J 2007 *22nd Annual IEEE Conf. on Computational Complexity (CCC’07)* pp 129–40
- [24] Brassard G, Høyer P, Mosca M and Tapp A 2002 Quantum amplitude amplification and estimation *Quantum Computation and Quantum Information Science (AMS Contemporary Math Series Vol 305)* pp 53–74