Controlled Quantum Amplification*

Cătălin Dohotaru[1] and Peter Høyer[2]

1 **Department of Computer Science, University of Calgary, Calgary, Canada**
cdohotaru@gmail.com
2 **Department of Computer Science, University of Calgary, Calgary, Canada**
hoyer@ucalgary.ca

Abstract

We propose a new framework for turning quantum search algorithms that decide into quantum algorithms for finding a solution. Consider we are given an abstract quantum search algorithm A that can determine whether a target g exists or not. We give a general construction of another operator U that both determines and finds the target, whenever one exists. Our amplification method at most doubles the cost over using A, has little overhead, and works by controlling the evolution of A. This is the first known general framework to the open question of turning abstract quantum search algorithms into quantum algorithms for finding a solution.

We next apply the framework to random walks. We develop a new classical algorithm and a new quantum algorithm for finding a unique marked element. Our new random walk finds a unique marked element using H update operations and $1/\epsilon$ checking operations. Here H is the hitting time, and ϵ is the probability that the stationary distribution of the walk is in the marked state. Our classical walk is derived via quantum arguments. Our new quantum algorithm finds a unique marked element using $\sqrt{H}$ update operations and $\sqrt{1/\epsilon}$ checking operations, up to logarithmic factors. This is the first known quantum algorithm being simultaneously quadratically faster in both parameters. We also show that the framework can simulate Grover's quantum search algorithm, amplitude amplification, Szegedy's quantum walks, and quantum interpolated walks.

1998 ACM Subject Classification F.1.2 Modes of Computation, F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases Quantum algorithms, quantum walks, random walks, quantum search

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.18

1 Introduction

Grover's search algorithm [15], amplitude amplification [8], and quantum walks [3, 30] are highly successful methodologies for searching in quantum algorithms. They are used in search problems in which we are given some unitary operator W as a black-box. We start in some initial state $|\text{init}\rangle$ which is a $(+1)$-eigenvector of the unitary W. Our goal is to produce some unknown target state $|g\rangle$. We are given a reflection[1] operator $\mathsf{G} = 1 - 2|g\rangle\langle g|$ that permits us to distinguish the target state $|g\rangle$ from any other orthogonal state. Our task is to construct an algorithm that evolves the initial state $|\text{init}\rangle$ into a state that has constant overlap with

* This work has been supported in part by the Canadian Institute for Advanced Research (CIFAR) and Canada's Natural Sciences and Engineering Research Council (NSERC).

[1] To simplify later calculations, we define the operator G as the reflection about the subspace orthogonal to $|g\rangle$.

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).
Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;
Article No. 18; pp. 18:1–18:13

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the target state $|g\rangle$, using the operator W. The standard approach in quantum algorithmics for doing so is to use operator $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$ [5].

Much work on search problems on quantum computers has been specifically developed within quantum walks. The operator W is often derived from a random walk, in which case it is referred to as the *walk* or *update* operator. The cost of the quantum search algorithm is then typically phrased in terms of the spectral gap δ or the hitting time of the associated random walk.

The study of quantum walks has been a highly successful and active line of research, with recent results such as a quantum algorithm for triangle finding [13] using only $O(n^{5/4})$ queries. Other applications of quantum walks include verification of matrix products [9], testing group commutativity [27], formula evaluation [4], subgraph finding [10], triangle finding [26, 13, 14], and 3-distinctness [7].

These and other applications were found after the seminal works of Ambainis [3] and Szegedy [30]. Ambainis [3] gave a quantum walk for element distinctness, and Szegedy [30] gave a general method for obtaining an quantum search algorithm from any symmetric random walk [30]. Magniez et al. [25] gave a quantum algorithm that finds a unique marked element for any reversible random walk[2] of cost in the order of $\mathsf{S} + \sqrt{1/(\epsilon\delta)}\mathsf{U} + \sqrt{1/\epsilon}\mathsf{C}$. Here S, U, and C are the setup, update and checking costs of the quantum walk, δ the spectral gap of the walk, and ϵ the probability that the stationary state is in the marked state [29, 28].

Magniez et al. [23] gave a quantum algorithm that finds a unique marked element for any state-transitive random walk of cost in the order of $\mathsf{S} + \sqrt{H}\mathsf{U} + \sqrt{H}\mathsf{C}$. Here H is the hitting time of the random walk. Krovi et al. [22] introduced the novel idea of interpolating walks. Krovi et al. show in [20, 21] that interpolated walks can find a marked element for any reversible random walk. Their algorithm works for multiple marked elements and has cost in the order of $\mathsf{S} + \sqrt{H^+}\mathsf{U} + \sqrt{H^+}\mathsf{C}$, where H^+ is a quantity introduced in [21] and referred to as the extended hitting time. When there is a unique marked element, the extended hitting time and hitting time coincide, $H^+ = H$. When there are multiple marked elements, the extended hitting time is bounded by $H \leq H^+ \leq \frac{1}{\epsilon\delta}$. Excellent surveys on quantum walks, their history and applications, include [2, 18, 29, 32, 28].

In this work, we propose a new framework for the general setting of search problems. Our framework does *not* require that the operator W is derived from a random walk and it makes no explicit use of properties of quantum walks or random walks. The most obvious application of our framework is naturally to random walks, but is not limited to such cases. We will assume that W has only real entries, and that our target $|g\rangle$ has real coordinates in a canonical orthogonal basis for the space acted upon by W. This assumption is fulfilled in all the applications we consider here, including quantum walks.

Our framework consists of several parts. We give a new generic quantum algorithm for solving search problems. Our algorithm is based on a circuit U which controls the search, and we refer to this process as *controlled quantum amplification*. We prove that whenever the operator $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$ determines whether a target $|g\rangle$ exists or not, in some number of iterations T, with constant success probability, then our circuit U both determines and finds the target using at most $2T$ iterations, with constant success probability.

We apply our analysis to quantum walks and prove that we can obtain a quadratic speedup for reversible random walks. We show that our framework can simulate quantum interpolated

[2] A Markov chain P with a unique stationary distribution $\pi = (\pi_x)$ is said to be *reversible* if $\pi_x P_{y,x} = \pi_y P_{x,y}$ for all states x, y. All Markov chains obtained from a random walk on undirected graphs are reversible. We will below use the wording "random walk" and "Markov chain" interchangeably.

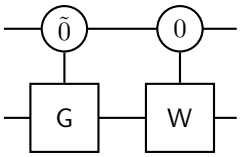

Figure 1 Our circuit U for controlled quantum amplification, expressed in terms of the reflection operator G and the unitary W, here given in its simplest form. We apply circuit U successively to the initial state $|0\rangle|\overline{\text{init}}\rangle$, say T times, thus producing the final state $\mathsf{U}^T|0\rangle|\overline{\text{init}}\rangle$, which we measure. If the outcome of the measurement is $|\tilde{1}\rangle|g\rangle$, the circuit has successfully found the marked state g.

walks [21], thus eliminating the need for running an interpolation of a random walk and its absorbing analog. We prove a relation between the operator W and the operator A, and we use this relation to construct a new quantum algorithm that, up to a logarithmic factor, finds a unique marked element in cost $\mathsf{S}+\sqrt{H}\mathsf{U}+\sqrt{1/\epsilon}\mathsf{C}$. This is superior to the existing algorithms. We also use this relation to construct a new *classical* algorithm that finds a unique marked element in cost $\mathsf{S}+H\mathsf{U}+1/\epsilon\mathsf{C}$. Our classical walk is derived using quantum arguments.

2 Controlled quantum amplification

In a quantum search problem, we are given an arbitrary unitary operator W, through which we would like to extract some unknown target state $|g\rangle$. The operator W can be given as a black box, which we can apply to any state $|\psi\rangle$, producing the state $\mathsf{W}|\psi\rangle$. An application of W has some cost, which is called the *update cost*.

We are also given a reflection operator $\mathsf{G}=1-2\Pi_G$, where $\Pi_G=|g\rangle\langle g|$ is a projection on the target state $|g\rangle$. Operator G permits us to distinguish the target state $|g\rangle$ from any other orthogonal state. The operator G can be given as a black box as well. An application of G has some cost, which is called the *checking cost*.

Our first step is to prepare the initial state $|\text{init}\rangle$. Preparing this state has a cost, which is called the *setup cost*. We then measure whether the initial state contains the target state or not by applying the measurement $\{\Pi_G, 1-\Pi_G\}$. The probability by which we measure $|g\rangle$ is some $\epsilon=\sin^2(\theta)$ where $\sin(\theta)=\langle g|\text{init}\rangle$. Let $|g\rangle$ and this initial angle θ be so that $0\le\theta\le\pi/2$. If $\theta=0$, our initial success probability is zero and the quantum search algorithm $\mathsf{A}=\mathsf{W}\cdot\mathsf{G}$ will not amplify this. If $\theta=\pi/2$, our initial state *is* the target state and there is nothing to be amplified. We shall therefore assume that $0<\theta<\pi/2$. Typically θ is very close to zero, corresponding to that the non-amplified success probability ϵ is small.

If our initial measurement yields the outcome $|g\rangle$, we terminate the algorithm as we have produced our target state $|g\rangle$. Otherwise, our initial state becomes $|\overline{\text{init}}\rangle=\frac{1}{\cos(\theta)}(|\text{init}\rangle-\sin(\theta)|g\rangle)$, which is orthogonal to $|g\rangle$ by construction. Thus starting with $|\overline{\text{init}}\rangle$, we then want to produce a state with large overlap with $|g\rangle$. To achieve this, we propose the following circuit U.

The circuit U acts on two registers. The first register contains a qubit which we use to control the evolution in the second register. The circuit is comprised of two operators, both of which are controlled on the state in the first register. We fix an angle $0<\tilde{\theta}<\pi/2$ and use the rotated orthogonal basis $|\tilde{0}\rangle=\cos(\tilde{\theta})|0\rangle+\sin(\tilde{\theta})|1\rangle$ and $|\tilde{1}\rangle=-\sin(\tilde{\theta})|0\rangle+\cos(\tilde{\theta})|1\rangle$. The first operator $|\tilde{0}\rangle\langle\tilde{0}|\otimes\mathsf{G}+|\tilde{1}\rangle\langle\tilde{1}|\otimes 1$ in our circuit reflects about the target state $|g\rangle$ conditional on that the control qubit is in state $|\tilde{0}\rangle$. The second operator $|0\rangle\langle 0|\otimes\mathsf{W}+|1\rangle\langle 1|\otimes 1$ applies the update W conditional on that the control qubit is in state $|0\rangle$.

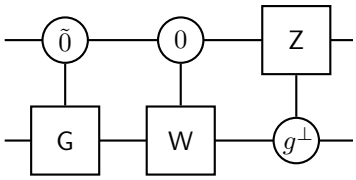

Figure 2 W with reflection.

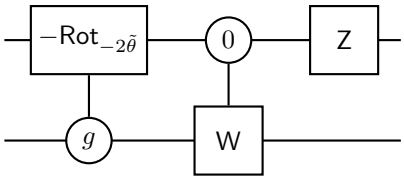

Figure 3 W with rotation.

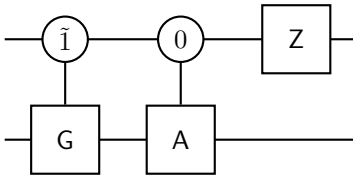

Figure 4 $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$ with reflection.

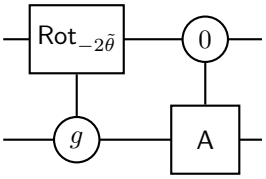

Figure 5 A with rotation.

The circuit U is parameterized by the angle $\tilde{\theta}$. The choice of $\tilde{\theta}$ has some algorithmic consequences, which we may handle by exponential searching similar to past work [8, 23]. The framework applies with no increase in asymptotic cost if we are given a multiplicatively approximate value for the initial success amplitude $\sin(\theta)$. We apply circuit U successively to the initial state $|0\rangle|\overline{\text{init}}\rangle$, say T times, thus producing the final state $\mathsf{U}^T|0\rangle|\overline{\text{init}}\rangle$, which we measure. We prove here that for an appropriately chosen value of T, the final measurement yields the target $|g\rangle$ with probability at least a constant, and we thus refer to this process as *controlled quantum amplification*.

The circuit has multiple interpretations and forms which both provide us with flexibility in terms of implementations as well as a foundation for proving properties on quantum amplification processes. We give here four circuits, all of which act equivalently to U. In Figure 2, the third gate $\mathsf{Z} \otimes (1 - |g\rangle\langle g|) + 1 \otimes |g\rangle\langle g|)$ applies the phase gate $\mathsf{Z} = |0\rangle\langle 0| - |1\rangle\langle 1|$ on the control qubit conditional on that the search space does not contain the target state $|g\rangle$. In Figures 3 and 5, the first gate rotates the ancilla qubit by an angle of $\pi - 2\tilde{\theta}$ and $-2\tilde{\theta}$, respectively, conditional on the search space contains the target state $|g\rangle$.

One of our aims is to compare our new controlled amplification circuit U with the operator $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$. The operator A has been used extensively in quantum walks and, when applied to random walks, it corresponds to an absorbing random walk [30, 3, 5]. The main limitation of the operator A is that it does not necessarily produce the target state, even when one exists. Significant research has been put into understanding when the operator A does and does not produce the target state. Our proposed controlled circuit circumvents this barrier in all generality. We prove that whenever the operator A determines whether a unique target exists or not, in some number of iterations, then our circuit U both determines and finds the target in the same asymptotic number of iterations.

Controlled quantum computations have been successfully applied in quantum computing dating back to at least the notion of phase kick-back [11]. Tulsi used a quantum walk with controlled operators for the problem of finding a unique target $|g\rangle$ on a grid [31]. His algorithm finds a unique target in cost $O(\sqrt{n \log n})$, which is quadratically smaller than the classical hitting time of $\Theta(n \log n)$ [1]. The grid graph is a two-dimensional torus of size $\sqrt{n} \times \sqrt{n}$ and has been a notoriously hard case for quantum searching because all of its edges are local. Magniez, Nayak, Richter, and Santha [23] extended this and gave a controlled operator that finds a unique target $|g\rangle$ for any state-transitive graph.

Our controlled amplifier does not rely on any graph-theoretic properties. We prove the general statement that whenever A determines whether a unique target state exists or not, our controlled amplifier finds the target state in asymptotically the same cost.

3 Quantum hitting times

The hitting time is a notion used in the analysis of stochastic processes such as random walks. It is the expected number of steps some stochastic process U uses to reach the target state g, starting from some appropriately defined initial distribution π. The choice of an appropriate corresponding definition of quantum hitting time is non-trivial.

Let U be any real unitary, and let $|w\rangle$ be any normalized target state with real coordinates in a canonical orthogonal basis for the space acted upon by U. The possible eigenvalues for U are $+1$, -1, and conjugated pairs $(e^{i\alpha}, e^{-i\alpha})$ for some eigenphase $0 < \alpha < \pi$. Each such pair of eigenvalues corresponds to a distinct two-dimensional subspace acted upon by U by a rotation of angle α. We order these non-trivial eigenphases of U as $0 < \alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_m < \pi$, for some $m \geq 0$. Let $|U_j^+\rangle$ and $|U_j^-\rangle$ be the conjugated eigenvectors of U corresponding to the eigenvalues $e^{i\alpha_j}$ and $e^{-i\alpha_j}$, for each $1 \leq j \leq m$. We decompose $|w\rangle$ into this ordered eigenbasis of U, as $w_0|U_0\rangle + \sum_{j=1}^{m} \left(w_j^+|U_j^+\rangle + w_j^-|U_j^-\rangle\right) + w_{-1}|U_{-1}\rangle$. Here we group all $(+1)$-eigenvectors of U into $|U_0\rangle$, and all the (-1)-eigenvectors into $|U_{-1}\rangle$. Since U and $|w\rangle$ have real components, we can choose the scalars of the eigenvectors of U such that $w_j^+ = w_j^- = w_j \in \mathbb{R}$. Given this basis, we define the quantum hitting time as follows.

▶ **Definition 1.** The *quantum hitting time* of U on $|w\rangle = w_0|U_0\rangle + \sum_{j=1}^{m} w_j\left(|U_j^+\rangle + |U_j^-\rangle\right) + w_{-1}|U_{-1}\rangle$ is

$$\mathsf{QHT}_\alpha(\mathsf{U}, |w\rangle) = \sqrt{2\sum_{j=1}^{m} |w_j|^2 \frac{1}{\alpha_j^2}}. \tag{1}$$

Since our proofs use the specifics of this definition, we mention that our definition differs from the more commonly used quantity $\mathsf{QHT}_1(\mathsf{U}, |w\rangle) = 2\sum_{j=1}^{m} |w_j|^2 \frac{1}{\alpha_j} + |w_{-1}|$. Our notion of quantum hitting time QHT_α is quadratically smaller than the classical hitting time for reversible random walks, by Szegedy's correspondence [30]. Thus, if a quantum algorithm has cost in the order of QHT_α, then it has cost quadratically smaller than the classical hitting time. For technical reasons, we need to introduce a second notion of quantum hitting time, which we refer to as the *cotangent quantum hitting time* and define as $\mathsf{QHT}_{\cot}(\mathsf{U}, |w\rangle) = \sqrt{2\sum_{j=1}^{m} |w_j|^2 \cot^2(\frac{\alpha_j}{2})}$. Our two notions of quantum hitting times QHT_α and $\mathsf{QHT}_{\cot}$ are asymptotically of the same order, as shown in Lemma 9 in the appendix. We use QHT as a shorthand for QHT_α.

4 Finding in the quantum hitting time

Our goal is to prove that the circuit U finds a target state in the quantum hitting time, stated as Corollary 6 below. We first identify the principal eigenvector of the circuit U.

▶ **Lemma 2.** *The unnormalized state* $|v_0\rangle = \sin(\tilde{\theta})\,|0, \overline{\text{init}}\rangle - \frac{\sin(\theta)}{\cos(\theta)}\,|\tilde{1}, g\rangle$ *is a* $(+1)$*-eigenvector of circuit* U.

Proof. The proof follows be considering the action of the circuit U given in Figure 1.

The first operator in the circuit reflects the state $|\tilde{0}, g\rangle$. The first term in the state $|v_0\rangle$ is orthogonal to this reflection state since $|\overline{\text{init}}\rangle$ is orthogonal to $|g\rangle$ by definition. The second term in $|v_0\rangle$ is also orthogonal to the reflection state since $|\tilde{0}\rangle$ and $|\tilde{1}\rangle$ constitute an orthonormal basis. The reflection operator thus acts trivially on $|v_0\rangle$.

The second operator in the circuit applies the operator W on the search register conditional on that the control qubit is in state $|0\rangle$. Rewrite $|v_0\rangle$ on the form

$$\frac{\sin(\tilde{\theta})}{\cos(\theta)}|0, \text{init}\rangle - \frac{\sin(\theta)}{\cos(\theta)}\cos(\tilde{\theta})\,|1, g\rangle.$$

Then each of the two terms is again invariant, and we conclude that the second operator similarly acts trivially on $|v_0\rangle$. ◀

We want this principal eigenvector to be an equally weighted superposition of our initial state $|0, \overline{\text{init}}\rangle$ and our target state $|\tilde{1}, g\rangle$. We therefore choose and fix the angle $\tilde{\theta}$ such that $0 < \tilde{\theta} < \pi/2$ and $\sin(\tilde{\theta}) = \frac{\sin(\theta)}{\cos(\theta)}$, and write the principal eigenvector on the normalized form $|U_0\rangle = \frac{1}{\sqrt{2}}\big(|0, \overline{\text{init}}\rangle - |\tilde{1}, g\rangle\big)$.

When considering quantum search problems, it is commonly assumed that the walk operator W has a unique $(+1)$-eigenvector (up to scalars), which we adapt here for convenience. The purpose of a quantum search problem is to produce a state that has large overlap with the target state $|g\rangle$. Our proposed algorithm for doing so, is to apply our circuit U successively a number of T times on the initial state $|0, \overline{\text{init}}\rangle$, producing the final state $\mathsf{U}^T|0, \overline{\text{init}}\rangle$. We show that it suffices to pick the number of iterations T to be in the order of the quantum hitting time $\mathsf{QHT}(\mathsf{U}, |\tilde{1}, g\rangle)$ of U. We divide the proof up into two cases. When W is a reflection, our circuit emulates amplitude amplification. For general operators W, we use that the principal $(+1)$-eigenvector of U has large overlap with both the initial state and the target state $|\tilde{1}, g\rangle$.

Consider first that $\mathsf{W} = 2|\text{init}\rangle\langle\text{init}| - 1$ is a reflection about the initial state. In amplitude amplification [8], we apply the operator $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$ a number of T times, thus constructing the state $|\text{final}\rangle = (\mathsf{W} \cdot \mathsf{G})^T|\text{init}\rangle$. By picking $T = \lceil \frac{\pi}{4\theta} \rceil \in \Theta(\frac{1}{\sqrt{\epsilon}})$, a measurement of the final state $|\text{final}\rangle$ yields the target state $|g\rangle$ with probability at least $1 - \epsilon$, where $\epsilon = \sin^2(\theta) = \langle g|\text{init}\rangle^2$ is the initial success probability. Amplitude amplification thus amplifies quadratically faster than classical repetition.

Now consider circuit U given in Figure 1 when $\mathsf{W} = 2|\text{init}\rangle\langle\text{init}| - 1$ is a reflection about the initial state. Then U is the product of two reflections and effectively implements a rotation in the two-dimensional subspace spanned by $|\tilde{0}, g\rangle$ and $|0, \overline{g}\rangle$, where $|\overline{g}\rangle = \frac{1}{\cos(\theta)}(|g\rangle - \sin(\theta)|\text{init}\rangle)$ is defined analogously to $|\overline{\text{init}}\rangle$. The rotational angle is 2φ, where φ is given by the inner product $\cos(\varphi) = \langle 0, \overline{g}|\tilde{0}, g\rangle = \cos(\tilde{\theta})\cos(\theta) = \sqrt{\cos(2\theta)}$. This gives us that $2\varphi \approx 2\sqrt{2}\theta$. The initial state is $|0, \overline{\text{init}}\rangle = \frac{1}{\sqrt{2}}(|U_0\rangle + |U_{\text{rot}}\rangle)$, where $|U_{\text{rot}}\rangle$ belongs to the two-dimensional rotational subspace. We pick $T = \lceil \frac{\pi}{2\sqrt{2}\theta} \rceil$, and produce the final state $\mathsf{U}^T|0, \overline{\text{init}}\rangle$, a measurement of which yields the target $|\tilde{1}, g\rangle = \frac{1}{\sqrt{2}}(|U_0\rangle - |U_{\text{rot}}\rangle)$ with probability at least $1 - O(\epsilon)$.

4.1 Finding in the quantum hitting time for general A

When W is a reflection about the initial state $|\text{init}\rangle$, the previous subsection implies that it suffices to choose T to be in the order of $\frac{1}{\theta}$, just as in amplitude amplification. For arbitrary operators W, it suffices to choose T to be in the order of the quantum hitting time $\mathsf{QHT}(\mathsf{U}, |0, \overline{\text{init}}\rangle)$ of U on the initial state $|0, \overline{\text{init}}\rangle$.

▶ **Theorem 3.** *There is an algorithm that applies* U *an expected number of order* $\mathsf{QHT}(\mathsf{U}, |0,\overline{\text{init}}\rangle)$ *times to the initial state* $|0,\overline{\text{init}}\rangle$*, and produces a final state with constant overlap with* $|\tilde{1},g\rangle$*.*

The proof is as follows. By Lemma 2, the initial state $|0,\overline{\text{init}}\rangle$ can be written as an equal superposition of the principal eigenvector $|U_0\rangle$ of U and the state $|U_{\text{rot}}\rangle = \frac{1}{\sqrt{2}}\big(|0,\overline{\text{init}}\rangle + |\tilde{1},g\rangle\big)$. The state $|U_0\rangle$ has constant overlap with the target state $|\tilde{1},g\rangle$. The state $|U_0\rangle$ is an eigenvector of U with eigenphase 0, whereas the state $|U_{\text{rot}}\rangle$ is a superposition of states with non-zero eigenphases. Following a standard argument via phase estimation [19, 11, 8, 24, 23], we can determine which is the case by successive (controlled) U applications. The primary observations are that the success probability in phase estimation can be expressed naturally in terms of our quantum hitting time QHT_α, and the controls on U can be dropped.

The theorem remains true if we replace the quantum hitting time with the effective quantum hitting time, at the expense of a drop in the overlap by at most a small constant. Here the effective quantum hitting time is the smallest number of applications of (controlled) U required to produce a final state with constant overlap with $|\tilde{1},g\rangle$. By Markov's inequality, the effective quantum hitting time is at most in the order of the quantum hitting time.

Theorem 3 provides us with an expression of the cost of U in terms of the quantum hitting time of U itself. We next relate the quantum hitting time of U to the quantum hitting times of quantum search operators A and W. Let $\epsilon = \sin^2(\theta)$ be the initial success probability, where angle $0 < \theta < \pi/2$ is so that $\sin(\theta) = |\langle g|\text{init}\rangle|$, and set angle $0 < \tilde{\theta} < \pi/2$ so that $\sin(\tilde{\theta}) = \frac{\sin(\theta)}{\cos(\theta)}$.

▶ **Theorem 4.**

$$\mathsf{QHT}(\mathsf{U}, |0,\overline{\text{init}}\rangle) = \mathsf{QHT}(\mathsf{U}, |\tilde{1},g\rangle) = \Theta\left(\frac{1}{\sqrt{\epsilon}}\mathsf{QHT}(\mathsf{W}, |\bar{g}\rangle)\right) = \Theta\big(\mathsf{QHT}(\mathsf{A}, |\overline{\text{init}}\rangle)\big).$$

Consider Theorem 4. The first equality follows since the principal $(+1)$-eigenvector $|U_0\rangle = \frac{1}{\sqrt{2}}(|0,\overline{\text{init}}\rangle - |\tilde{1},g\rangle)$ of the controlled amplifier U is an equal superposition of the starting state and the target state. To prove the remaining two equalities, we relate the computational quantity QHT to a structural quantity, an inner product.

▶ **Lemma 5.** *Let* V *be any real unitary, and let* $|w\rangle$ *be a real state that does not overlap the* $(+1)$*-eigenspace* V_+ *of* V*. Then the operator* $\mathsf{S} = \mathsf{V}\cdot(1-2|w\rangle\langle w|)$ *has a unique* $(+1)$*-eigenvector* $|+S\rangle$ *orthogonal to* V_+*, and it satisfies that* $\mathsf{QHT}(\mathsf{V}, |w\rangle) = \Theta\big(\frac{1}{|\langle w|+S\rangle|}\big)$*.*

Proof. We first decompose $|w\rangle$ into the eigenbasis of V as $|w\rangle = \sum_j w_j(|V_j^+\rangle + |V_j^-\rangle) + w_{-1}|V_{-1}\rangle$, where $w_j \in \mathbb{R}$ for all j. The states $|V_j^+\rangle$ and $|V_j^-\rangle$ are the two conjugated eigenstates for the j^{th} rotational subspace of V with eigenphases $\pm\varphi_j$. The state $|V_{-1}\rangle$ is the normalized projection of $|w\rangle$ onto the (-1)-eigenspace of V (when it exists). Then S has a $(+1)$-eigenvector on the form $|+w\rangle = |w\rangle + i|w^\perp\rangle = |w\rangle + i\sum_j w_j \cot\left(\frac{\varphi_j}{2}\right)\big(|V_j^-\rangle - |V_j^+\rangle\big)$, and it is the unique $(+1)$-eigenvector orthogonal to the trivial $(+1)$-eigenspace V_+. The norm $\|+w\|$ of $|+w\rangle$ is

$$\sqrt{1+\||w^\perp\rangle\|^2} = \sqrt{1+2\sum_j w_j^2\cot^2\left(\frac{\varphi_j}{2}\right)} = \Theta(\mathsf{QHT}_{\cot}(\mathsf{V}, |w\rangle)) = \Theta(\mathsf{QHT}(\mathsf{V}, |w\rangle)),$$

where the last equality follows by Lemma 9 given in the appendix. The normalized eigenvector is then $|+S\rangle = \frac{1}{\Theta(\mathsf{QHT}(\mathsf{V},|w\rangle))}\left(|w\rangle + i|w^\perp\rangle\right)$. ◀

Note that in Theorem 4, the last two vectors are orthogonal to the $(+1)$-eigenspaces of the respective operators, and hence Lemma 5 applies. Since the corresponding inner products

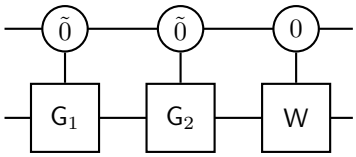

Figure 6 The circuit U rewritten to permit an analysis of its action when there are multiple targets.

are of the same order and also of the same order as the quantity $\mathsf{QHT}(\mathsf{U}, |\tilde{1}, g\rangle)$, we deduce the four quantum hitting times are of the same order, implying Theorem 4.

Theorem 4 implies that the quantum hitting time of U on $|0, \overline{\text{init}}\rangle$ is asymptotically the same as the quantum hitting time of A on $|\overline{\text{init}}\rangle$.

▶ **Corollary 6.** *There is an algorithm that applies circuit* U *an expected number in the order of* $\mathsf{QHT}(\mathsf{A}, |\overline{\text{init}}\rangle)$ *times to the initial state* $|0, \overline{\text{init}}\rangle$ *and produces a final state with constant overlap with* $|\tilde{1}, g\rangle$.

When A is a quantum walk derived from a reversible random walk P using Szegedy's construction [30], then A can detect the presence of a unique target in cost in the order of $\mathsf{QHT}(\mathsf{A}, |\overline{\text{init}}\rangle) \in O(\sqrt{\mathsf{HT}(\mathsf{P}, \{m\})})$, which is quadratically less than the hitting time of P [23, 28]. The corollary thus implies that U finds a unique target quadratically faster than classically.

5 Finding with multiple targets

We have analyzed our circuit U given in Figure 1 for a single target state $|g\rangle$. Consider now we have t targets. Let $|g_1\rangle, \ldots, |g_t\rangle$ be a set of t orthogonal states spanning this target subspace $\mathcal{G}$, and let Π_G be the projection onto $\mathcal{G}$. In our circuit U in Figure 1, consider we now have $\mathsf{G} = 1 - 2\Pi_G$. Let $|g_\pi\rangle$ be the normalized projection of $|\text{init}\rangle$ onto the target subspace. Then $\epsilon_\pi = \sin^2(\theta) = |\langle g_\pi|\text{init}\rangle|^2$ is the total probability that a measurement of the initial state would successfully produce any of the target states.

To analyze the circuit U when there are multiple targets, we again rewrite the circuit. Set $\mathsf{G}_1 = 1 - 2|g_\pi\rangle\langle g_\pi|$ and let $\mathsf{G}_2 = 1 - 2(\Pi_G - |g_\pi\rangle\langle g_\pi|)$. We can then write our circuit U on the form given in Figure 6. The rewritten circuit differs in form from our original circuit by the second gate $|\tilde{0}\rangle\langle\tilde{0}| \otimes \mathsf{G}_2 + |\tilde{1}\rangle\langle\tilde{1}| \otimes 1$. In general, this second gate changes the quantum hitting time of the circuit. For some classes of operators W, though, the second gate has no impact. In particular, for quantum walks on reversible graphs when the operator W is comprised of a reflection and a swap operator [30], we can compute an explicit expression of the complexity.

▶ **Theorem 7.** *Let* W *be a reversible quantum walk with multiple marked elements. There is an algorithm that applies* U *an expected number in the order of* $\frac{1}{\sqrt{\epsilon_\pi}} \cdot \mathsf{QHT}(\mathsf{W}, |\overline{g_\pi}\rangle)$ *times on the initial state* $|0, \overline{\text{init}}\rangle$ *and produces a final state with constant overlap with* $|\tilde{1}, g_\pi\rangle$.

Note that in this theorem, $|g_\pi\rangle$ is the normalized projection of $|\text{init}\rangle$ onto the marked subspace and ϵ_π is the total probability a measurement of $|\text{init}\rangle$ yields a marked element. The state $|\overline{g_\pi}\rangle$ is the normalized projection of $|g_\pi\rangle$ onto the subspace orthogonal to $|\text{init}\rangle$. The theorem follows by observing that the quantum hitting times on the input $|0, \overline{\text{init}}\rangle$ with or without the second gate are the same.

6 Faster algorithms for a unique marked element

Consider there is a unique marked element g. By Theorem 4, we can write $\mathsf{QHT}^2(\mathsf{A}, |\overline{\text{init}}\rangle)$ as a product of the two factors $\frac{1}{\epsilon}$ and $\mathsf{QHT}^2(\mathsf{W}, |\overline{g}\rangle)$. This decomposition permits us to device a new quantum algorithm for finding a unique marked element with constant success probability of cost in the order of

$$\mathsf{S} + \sqrt{H}\mathsf{U} + \frac{1}{\sqrt{\epsilon}}\mathsf{C}. \tag{2}$$

Here $H \in \widetilde{O}(\mathsf{QHT}^2(\mathsf{A}, |\overline{\text{init}}\rangle))$ is in the order of the square of the quantum hitting time, up to logarithmic factors. When A is the quantum walk derived from a reversible random walk P using Szegedy's construction, then $\mathsf{QHT}^2(\mathsf{A}, |\overline{\text{init}}\rangle) \in O(\mathsf{HT}(\mathsf{P}, \{g\}))$ is in the order of the hitting time of P with unique marked element g.

Since the hitting time $\mathsf{HT}(\mathsf{P}, \{g\})$ is upper bounded by $\frac{1}{\epsilon\delta}$, the cost of our algorithm is, up to logarithmic factors, upper bounded by the cost of the existing algorithms in [25, 23, 21]. Our algorithm is derived by using Theorem 4 and is based on recursive amplitude amplification [17, 16, 25]. In terms of the checking cost C, our algorithm has the same cost as amplitude amplification would have, which offers a quadratic speed-up over any classical algorithm.

We also obtain a new *classical* algorithm that has cost of order

$$\mathsf{S} + H\mathsf{U} + \frac{1}{\epsilon}\mathsf{C}. \tag{3}$$

Here $H \in O(\mathsf{HT}(\mathsf{P}, \{g\}))$ is the hitting time of the reversible walk P with unique marked element g, without logarithmic factors. This combines the best of two natural choices of random walks having cost in the order of $\mathsf{S}+H(\mathsf{U}+\mathsf{C})$ and $\mathsf{S}+\frac{1}{\epsilon}(\frac{1}{\delta}\mathsf{U}+\mathsf{C})$ (see e.g. Santha [29] for a discussion on classical search algorithms). The main structure in our classical algorithm is as follows.

1. Sample an initial vertex x according to the stationary distribution π.
2. Repeat the following of order H/E times
   a. Let x denote the current state.
   b. Check if x is marked. If so, halt and output x.
   c. Else apply the random walk P a number of E times, starting from x.
3. If the current state x is marked, output x. Otherwise output "no marked element found."

The proof is by showing that $\mathsf{QHT}^2(\mathsf{W}(\mathsf{P}^E), \{g\})$ is upper bounded by a constant when E is of order $\mathsf{QHT}^2(\mathsf{W}(\mathsf{P}), \{g\})$. This statement is effectively an example of a classical theorem derived via quantum arguments. We have not been able to find this classical algorithm discussed in the literature before, and it is to the best of our knowledge new. We neither know of a way to prove that this classical algorithm finds a marked element in cost in the order of the expression in Eq. 3 without explicitly or implicitly applying arguments resembling the arguments introduced in this paper. It is, as far as we know, the first known random walk derived through the notion of quantum walks. We refer the reader to the excellent survey by Drucker and de Wolf [12] for further examples on quantum proofs for classical theorems.

7 Simulation of quantum interpolated walks

Our controlled amplifier can be applied to arbitrary real operators W, and we prove a general bound on the cost of the amplifier given in terms of the quantum hitting time QHT_α. We

now consider operators W derived from reversible random walks. This is the broadest class of random walks for which quantum algorithms of costs in the order of the quantum hitting time have been derived. The general problem is given by a state space X on which we defined a reversible random walk P. A subset $\mathcal{M}$ of the states of X are marked (the elements of $\mathcal{M}$ correspond to the solutions to some computational problem). Our goal is to find an element of $\mathcal{M}$.

Szegedy gives a general method for constructing a quantum walk $\mathsf{W}(\mathsf{P})$ from a reversible random walk P [30, 23]. Krovi et al. give in [21] a notion of interpolation between a reversible random walk P and its corresponding absorbing walk P', which is obtained from P by replacing all the transitions from marked vertices with self-loops. The interpolation $\mathsf{P}(s)$ is between two classical walks, where we transition according to P' with some fixed probability s, and transition according to P with complementary probability $1-s$. The resulting walk $\mathsf{P}(s)$ then yields a quantum walk $\mathsf{W}(\mathsf{P}(s))$ by Szegedy's construction. We use W and $\mathsf{W}(s)$ as shorthands for $\mathsf{W}(\mathsf{P})$ and $\mathsf{W}(\mathsf{P}(s))$, respectively. Let $\mathsf{HT}(\mathsf{P}(s),\mathcal{M})$ denote the classical hitting time of $\mathsf{P}(s)$. The quantum interpolated walk introduced in [21] finds a marked element. It takes a number of steps that is in the order of $\mathsf{HT}^+(\mathsf{P},\mathcal{M})$, where $\mathsf{HT}^+(\mathsf{P},\mathcal{M})$ is defined as $\mathsf{HT}^+(\mathsf{P},\mathcal{M}) = \lim_{s\to 1}\mathsf{HT}(\mathsf{P}(s),\mathcal{M})$. This limit is well-defined and referred to as the extended hitting time [21]. We use HT and HT^+ as shorthands for $\mathsf{HT}(\mathsf{P},\mathcal{M})$ and $\mathsf{HT}^+(\mathsf{P},\mathcal{M})$ when both P and $\mathcal{M}$ are fixed.

We show that controlled quantum amplifiers can simulate quantum interpolated walks. We do so by giving a constructive embedding E_s of $\mathsf{W}(s)$ into our framework. For a given parameter s, we choose the angle $\tilde{\theta}$ so that it satisfies that $0 \le \tilde{\theta} \le \pi/2$ and that

$$\sin\tilde{\theta} = \sqrt{1-s}, \tag{4}$$

obtaining the circuit $\mathsf{U} = \mathsf{U}(\tilde{\theta})$ (see Figure 1).

Let $\mathcal{H}_{\mathcal{M}}$ be the subspace of $\mathcal{H}_{\mathsf{W}}$ with marked items in the first register. Denote by $\epsilon = \sin^2(\theta)$ the initial success probability, namely the probability that we obtain a marked state in the first register by measuring $|\mathrm{init}\rangle$ according to $\{\Pi(\mathcal{H}_{\mathcal{M}}), 1-\Pi(\mathcal{H}_{\mathcal{M}})\}$. The optimal value for $\tilde{\theta}$, which is $\tilde{\theta} = \arcsin\left(\frac{\sin(\theta)}{\cos(\theta)}\right)$, corresponds to the optimal value of s, which is given by $s = 1 - \frac{\epsilon}{1-\epsilon}$.

Denote by $\mathcal{H}_{\mathsf{W}(s)}$ the space on which the quantum walk $\mathsf{W}(s)$ acts non-trivially (which is the same as the space on which W acts non-trivially), and let $\mathcal{H}_{\mathsf{U}}$ denote the space on which $\mathsf{U}(\tilde{\theta})$ acts. We define an embedding E_s from $\mathcal{H}_{\mathsf{W}(s)}$ to a subspace of $\mathcal{H}_{\mathsf{U}}$. We define E_s on a spanning set of $\mathcal{H}_{\mathsf{W}(s)}$ and then extend it by linearity,

$$\mathsf{E}_s \begin{cases} |u, p(s)_u\rangle & \mapsto & |0\rangle|u, p_u\rangle \\ |p(s)_u, u\rangle & \mapsto & |0\rangle|p_u, u\rangle \\ |m, p(s)_m\rangle & \mapsto & -|\tilde{1}\rangle|m, p_m\rangle \\ |p(s)_m, m\rangle & \mapsto & \sin\tilde{\theta}|0\rangle|p_m, m\rangle - \cos\tilde{\theta}|1\rangle|m, p_m\rangle. \end{cases} \tag{5}$$

The state $|p_x\rangle$ is a superposition of the neighbors of x and is defined by $\langle y|p_x\rangle$ being the square-root of the probability of transition from state x to state y in the random walk P. The states $|p(s)_x\rangle$ are defined similarly the random walk $\mathsf{P}(s)$.

By direct inspection, the embedding preserves inner products and is thus well-defined. Consider the following two maps from $\mathcal{H}_{\mathsf{W}(s)}$ to $\mathcal{H}_{\mathsf{U}}$ given by $\mathsf{E}_s\mathsf{W}(s)$ and $\mathsf{U}(\tilde{\theta})\mathsf{E}_s$. The first map first applies the quantum interpolated walk and then the embedding. The second map first applies the embedding and then our controlled quantum walk $\mathsf{U}(\tilde{\theta})$. These two operators act identically on each of the states given on the left hand sides in Eq. 5, and we thus conclude that $\mathsf{E}_s\mathsf{W}(s) = \mathsf{U}(\tilde{\theta})\mathsf{E}_s$, implying Theorem 8.

▶ **Theorem 8.** *Fix any* $0 \leq s < 1$. *There is an inner-product preserving map* E_s *from* $\mathcal{H}_\mathsf{W}$ *to a subspace of* $\mathcal{H}_\mathsf{U}$ *such that* $\mathsf{E}_s\mathsf{W}(s) = \mathsf{U}(\tilde{\theta})\mathsf{E}_s$.

Theorem 8 readily implies that controlled quantum walks can simulate quantum interpolated walks. Instead of running a quantum interpolated walk on its initial state $|\overline{\text{init}}\rangle$, we run a controlled quantum walk on the initial state $|0\rangle|\overline{\text{init}}\rangle = \mathsf{E}_s|\overline{\text{init}}\rangle$. Instead of measuring whether we have produced the state $|m, p(s)_m\rangle$ in a quantum interpolated walk, we run a controlled quantum walk and measure whether we have produced the state $|\tilde{1}\rangle|m, p_m\rangle = -\mathsf{E}_s|m, p(s)_m\rangle$.

By combining Theorems 7 and 8, we obtain that the extended hitting time $\mathsf{HT}^+(\mathsf{P}, \mathcal{M})$ of a reversible walk P on a marked subset $\mathcal{M}$ is in the order of $\frac{1}{\epsilon_\pi}\mathsf{QHT}^2(\mathsf{W}(\mathsf{P}), |\overline{g_\pi}\rangle)$. Further, we also re-derive the following result already shown by Ambainis and Kokainis [6]. Since $\mathsf{QHT}(\mathsf{W}, |\overline{g_\pi}\rangle) \leq \frac{1}{\sqrt{\delta}}$, we get that the extended hitting time is never more than a factor of $1/\delta$ larger than the hitting time, and that the extended hitting time $\mathsf{HT}^+(\mathsf{P}, \mathcal{M})$ is in the order of $\frac{1}{\epsilon_\pi \delta}$ for all marked subsets $\mathcal{M}$. Here δ is the spectral gap of the random walk P.

8 Concluding remarks

An quantum search algorithm takes two ingredients: an operator W and a reflection operator G. The standard method in quantum search algorithms is to apply the composed operator $\mathsf{A} = \mathsf{W} \cdot \mathsf{G}$. This method permits one to distinguish between the case when there is a marked state and the case when there are none. It does not in general produce a marked state, even when one exists.

We have here proposed a general method for amplifying the success probability of quantum search algorithms. The method applies readily to arbitrary (real) operators W, including operators derived from random walks. Our circuit U can be based equally well on either of the two operators W or A, depending on the application in mind. We prove that the controlled amplifier U finds a unique marked element in the same asymptotic cost as the standard circuit A determines whether one exists or not. We then prove and use properties of the controlled amplifier U to simulate amplitude amplification and interpolated walks, and to derive a new quantum and classical algorithm. Up to logarithmic factors, the costs are in the order of $\mathsf{S} + \sqrt{H}\mathsf{U} + 1/\sqrt{\epsilon}\mathsf{C}$ and $\mathsf{S} + H\mathsf{U} + 1/\epsilon\mathsf{C}$, respectively. Both algorithms improve upon the best known quantum and classical algorithms.

References

1 D. Aldous and J. Fill. Reversible Markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book`.

2 A. Ambainis. Quantum walks and their algorithmic applications. *International Journal of Quantum Information*, 1:507–518, 2003. `arXiv:quant-ph/0403120`.

3 A. Ambainis. Quantum walk algorithm for element distinctness. In *45th IEEE Symposium on Foundations of Computer Science*, FOCS'04, pages 22–31, 2004. `doi:10.1109/FOCS.2004.54`.

4 A. Ambainis, A. M. Childs, B. Reichardt, R. Špalek, and S. Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM Journal on Computing*, 39:2513–2530, 2010. `doi:10.1109/FOCS.2007.57`.

5 A. Ambainis, J. Kempe, and A. Rivosh. Coins make quantum walks faster. In *16th ACM Symposium on Discrete Algorithms*, SODA'05, pages 1099–1108, 2005. `arXiv:quant-ph/0402107`.

6 A. Ambainis and M. Kokainis. Analysis of the extended hitting time and its properties. *Poster presented at QIP 2015*, 2015.

7 A. Belovs, A. M. Childs, S. Jeffery, R. Kothari, and F. Magniez. Time-efficient quantum walks for 3-distinctness. In *40th International Colloquium on Automata, Languages, and Programming*, ICALP'13, pages 105–122, 2013. doi:10.1007/978-3-642-39206-1_10.

8 G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002. arXiv:quant-ph/0005055.

9 H. Buhrman and R. Špalek. Quantum verification of matrix products. In *17th ACM-SIAM Symposium on Discrete Algorithms*, SODA'06, pages 880–889, 2006. arXiv:quant-ph/0409035.

10 A. M. Childs and R. Kothari. Quantum query complexity of minor-closed graph properties. In *28th Symposium on Theoretical Aspects of Computer Science*, STACS'11, pages 661–672, 2011. doi:10.4230/LIPIcs.STACS.2011.661.

11 R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London, Series A*, 454:339–354, 1998. arXiv:quant-ph/9708016.

12 A. Drucker and R. de Wolf. *Quantum Proofs for Classical Theorems*. Number 2 in Graduate Surveys. Theory of Computing Library, 2011. doi:10.4086/toc.gs.2011.002.

13 F. Le Gall. Improved quantum algorithm for triangle finding via combinatorial arguments. In *55th IEEE Symposium on Foundations of Computer Science*, FOCS'14, pages 216–225, 2014. doi:10.1109/FOCS.2014.31.

14 F. Le Gall and S. Nakajima. Quantum algorithm for triangle finding in sparse graphs. In *15th Asian Quantum Information Science Conference*, AQIS'15, 2015. arXiv:1507.06878.

15 L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79:325–328, 1997. doi:10.1103/PhysRevLett.79.325.

16 P. Høyer and R. de Wolf. Improved quantum communication complexity bounds for disjointness and equality. In *19th Symp. on Theoretical Aspects of Computer Science*, STACS'02, pages 299–310, 2002. doi:10.1007/3-540-45841-7_24.

17 P. Høyer, M. Mosca, and R. de Wolf. Quantum search on bounded-error inputs. In *30th International Colloquium on Automata, Languages and Programming*, ICALP'03, pages 291–299, 2003. doi:10.1007/3-540-45061-0_25.

18 J. Kempe. Quantum random walks: An introductory overview. *Contemporary Physics*, 44(4):307–327, 2003. doi:10.1080/0010715103100011077 6.

19 A. Kitaev. Quantum measurements and the abelian stabilizer problem, 1995. arXiv:quant-ph/9511026.

20 H. Krovi, F. Magniez, M. Ozols, and J. Roland. Finding is as easy as detecting for quantum walks. In *37st International Colloquium on Automata, Languages and Programming*, ICALP'10, pages 540—-551, 2010. arXiv:1002.2419v1.

21 H. Krovi, F. Magniez, M. Ozols, and J. Roland. Quantum walks can find a marked element on any graph. *Algorithmica*, 74:851–907, February 2016. doi:10.1007/s00453-015-9979-8.

22 H. Krovi, M. Ozols, and J. Roland. Adiabatic condition and the quantum hitting time of Markov chains. *Physical Review A*, 82:022333, 2010. doi:10.1103/PhysRevA.82.022333.

23 F. Magniez, A. Nayak, P. Richter, and M. Santha. On the hitting times of quantum versus random walks. *Algorithmica*, 63(1):91–116, 2012. doi:10.1007/s00453-011-9521-6.

24 F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *39th ACM Symposium on Theory of Computing*, pages 575–584, 2007. doi:10.1137/090745854.

25 F. Magniez, A. Nayak, J. Roland, and M. Santha. Search via quantum walk. *SIAM Journal on Computing*, 40(1):142–164, Jan 2011. doi:10.1137/090745854.

26 F. Magniez, M. Santha, and M. Szegedy. Quantum algorithms for the triangle problem. *SIAM Journal on Computing*, 27:413–424, 2007. doi:10.1137/050643684.

27 A. Nayak and F. Magniez. Quantum complexity of testing group commutativity. *Algorithmica*, 48:221–232, 2007. doi:10.1007/s00453-007-0057-8.

28 A. Nayak, P. C. Richter, and M. Szegedy. Quantum analogs of markov chains. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*, pages 1–10. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. doi:10.1007/978-3-642-27848-8_302-2.

29 M. Santha. Quantum walk based search algorithms. In *International Conference on Theory and Applications of Models of Computation*, pages 31–46, 2008. doi:10.1007/978-3-540-79228-4_3.

30 M. Szegedy. Quantum speed-up of Markov chain based algorithms. In *45th IEEE Symposium on Foundations of Computer Science*, FOCS'04, pages 32–41, 2004. doi:10.1109/FOCS.2004.53.

31 A. Tulsi. Faster quantum walk algorithm for the two dimensional spatial search. *Physical Review A*, 78:012310, 2008. doi:10.1103/PhysRevA.78.012310.

32 S. E. Venegas-Andraca. Quantum walks: A comprehensive review. *Quantum Information Processing*, 11(5):1015–1106, 2012. arXiv:1201.4780.

A A lemma

The following lemma, referenced in Section 3, shows that the two quantum hitting times, $\mathsf{QHT}_{\cot}$ and QHT_{α} are of the same asymptotic order.

▶ **Lemma 9.** *For any real unitary* U *and any real state* $|w\rangle$,

$$\mathsf{QHT}_{\cot}(\mathsf{U},|w\rangle) \le 2\,\mathsf{QHT}_{\alpha}(\mathsf{U},|w\rangle) \le \sqrt{2} + \mathsf{QHT}_{\cot}(\mathsf{U},|w\rangle).$$

Proof. Since $\sin x < x < \tan x$ for any $x \in (0,\pi/2)$,

$$\cot^2 x < \frac{1}{x^2} < 1 + \cot^2 x. \tag{6}$$

By the first inequality in Eq. 6,

$$\mathsf{QHT}_{\cot}(\mathsf{U},|w\rangle) = \sqrt{2\sum_{j=1}^{m} |w_j|^2 \cot^2\left(\frac{\alpha_j}{2}\right)} \le \sqrt{2\sum_{j=1}^{m} |w_j|^2 \left(\frac{4}{\alpha_j^2}\right)} = 2\,\mathsf{QHT}_{\alpha}(\mathsf{U},|w\rangle),$$

proving the first inequality. Since the eigenphases α_j of U belong to $(0,\pi)$, the half angles $\alpha_j/2$ belong to the interval $(0,\pi/2)$. By the second inequality in Eq. 6, then

$$\begin{aligned}\mathsf{QHT}_{\alpha}(\mathsf{U},|w\rangle) &= \frac{1}{\sqrt{2}}\sqrt{\sum_{j=1}^{m} |w_j|^2 \left(\frac{4}{\alpha_j^2}\right)} \le \frac{1}{\sqrt{2}}\sqrt{\sum_{j=1}^{m} |w_j|^2 \left(1+\cot^2\left(\frac{\alpha_j}{2}\right)\right)} \\ &= \frac{1}{\sqrt{2}}\sqrt{1+\frac{1}{2}\mathsf{QHT}^2_{\cot}(\mathsf{U},|w\rangle)}.\end{aligned}$$

We obtain the second inequality in the lemma by applying the inequality $\sqrt{a+b} \le \sqrt{a}+\sqrt{b}$, for positive reals a and b. ◀