| | |
|---|---|
| Title: | Quantum Analogues of Markov Chains |
| Name: | Ashwin Nayak[1], Peter C. Richter[2], Mario Szegedy[2] |
| Affil./Addr. 1: | Department of Combinatorics and Optimization, and Institute for Quantum Computing, University of Waterloo, Waterloo, ON, Canada. Email: ashwin.nayak@uwaterloo.ca |
| Affil./Addr. 2: | Department of Computer Science, Rutgers, the State University of New Jersey, Piscataway, NJ, USA. Email: {richterp,szegedy}@cs.rutgers.edu |
| Keywords: | Markov chains; search problem; hitting time; quantum algorithms; quantum search; spatial search; quantum walks; element distinctness; triangle finding |
| SumOriWork: | 2004; Szegedy |

# Quantum Analogues of Markov Chains

Ashwin Nayak[1], Peter C. Richter[2], Mario Szegedy[2]

[1] Department of Combinatorics and Optimization, and Institute for Quantum Computing, University of Waterloo, Waterloo, ON, Canada.
Email: ashwin.nayak@uwaterloo.ca

[2] Department of Computer Science, Rutgers, the State University of New Jersey, Piscataway, NJ, USA.
Email: {richterp,szegedy}@cs.rutgers.edu

## Years and Authors of Summarized Original Work

2004; Szegedy

## Keywords

Markov chains; search problem; hitting time; quantum algorithms; quantum search; spatial search; quantum walks; element distinctness; triangle finding

## Problem Definition

***Spatial search and walk processes.*** *Spatial search* by *quantum walk* is database search with the additional constraint that one is required to move through the search space via a *quantum walk* that obeys some locality structure. For example, the data items may be stored at the vertices of a two-dimensional grid. The requirement of moves along the edges of the grid captures the cost of accessing different items starting from some fixed position in the database.

Quantum walks are analogues of classical random walks on graphs. The complexity of spatial search by quantum walk is strongly tied to the *quantum hitting time* [20] of the walk.

Let $S$, with $|S| = n$, be a finite set of *states*. Assume that a subset $M \subseteq S$ of states are *marked*. We are given a procedure $\mathcal{C}$ that, on input $x \in S$ and an associated data structure $d(x)$, checks whether the state $x$ is marked. The goal is either to find a marked state when promised that $M \neq \emptyset$ (*search version*), or to determine whether $M$ is nonempty (*decision version*).

The algorithm progresses in stages. In the *set-up stage* we access some state of $S$ (usually a random state). In the walk stage we move from state to state, performing a spatial walk as described below. The moves are called *updates*. In addition, in the walk stage we perform *checks* to see if the current state is marked at steps selected by the algorithm.

In the classical setting the *transition probabilities* of the spatial walk are described by a stochastic matrix $P = (p_{x,y})_{x,y \in S}$. This makes the walk a *Markov chain*. In every move the algorithm *must perform* a random transition according to $P$. The possible $x \to y$ moves, i.e., those with $p_{x,y} \neq 0$, form the edges of a (directed) graph $G$, and we say that the Markov chain $P$ has *locality structure* $G$.

We define the search problem in the classical setting, which carries over to the quantum case with little modification:

INPUT: Markov chain $P$ on set $S$, marked subset $M \subseteq S$ that is implicitly specified by a checking procedure $\mathcal{C}$, and the associated costs:

| Cost type | Setup | Update | Checking |
|---|---|---|---|
| Notation | $\mathsf{S}$ | $\mathsf{U}$ | $\mathsf{C}$ |

OUTPUT: a marked state if one exists (search version), or a Boolean return value that indicates whether $M$ is empty or not (decision version).

The algorithm is required to be correct with probability at least 2/3 in either case, the search or the decision problem. The significance of the set-up cost, which is incurred only once, will be clearer when we see some applications. Often we can choose between several competing walks, and we would like to design a one with minimum total cost.

In the *quantum* case, the random process $P$ is replaced by a *quantum walk* $W_P$ that has the same locality structure as $P$. The costs $\mathsf{S}, \mathsf{U}, \mathsf{C}$ reflect the costs of quantum operations.

***The quantum walk algorithm.*** Designing a quantum analogue of $P$ is not so straightforward, since stochastic matrices have no immediate unitary equivalents. One either needs to abandon the discrete-time nature of the walk [15] or define the walk operator on a space other than $\mathbb{C}^S$. Here we take the second route.

We say that a Markov chain $P$ is *irreducible* if its underlying digraph is strongly connected. Let $P$ be an irreducible Markov chain, let $\pi$ be its unique stationary distribution, and let $P^*$ (with $P^* = (p^*_{x,y})$) denote the *time-reversed Markov chain*, where $p^*_{x,y} := \pi_y p_{y,x}/\pi_x$. Define the following vectors in the vector space $\mathbb{C}^S$:

$$|p_x\rangle := \sum_{y \in X} \sqrt{p_{x,y}}\,|y\rangle \qquad \text{and} \qquad |p^*_y\rangle := \sum_{x \in X} \sqrt{p^*_{y,x}}\,|x\rangle \ .$$

Define the unitary operator $W_P := R_1 R_2$ on $\mathbb{C}^{S \times S}$ as the product of the two reflections $R_2 := \sum_{x \in S} |x\rangle\langle x| \otimes (2|p_x\rangle\langle p_x| - \mathrm{I})$ and $R_1 := \sum_{y \in S} (2|p^*_y\rangle\langle p^*_y| - \mathrm{I}) \otimes |y\rangle\langle y|$. The operator $W_P$ is called the *quantum analogue* of $P$, or the *discrete-time quantum walk operator* arising from $P$, and may be viewed as a walk on the *edges* of the underlying

graph $G$. We define a "checking" operator on $\mathbb{C}^S$, based on whether or not the current state is marked: $O_M := \sum_{x \notin M} |x\rangle\langle x| - \sum_{x \in M} |x\rangle\langle x|$.

In the above description, we have suppressed the data structure associated with a state in the Markov chain for the sake of simplicity. The precise description of the operators can be derived via the isometry $|x\rangle \mapsto |x\rangle|d(x)\rangle$ between the appropriate spaces (see, for example, Refs. [29, 30]). The data structure becomes especially significant in the context of the complexity of the operators.

A search algorithm by quantum walk is described by a quantum circuit that acts on "registers" or "wires" which are associated with the space $\mathbb{C}^S \otimes \mathbb{C}^S \otimes \mathbb{C}^k$, for some $k \geq 0$. We again suppress the registers carrying the data structure. The first two registers hold the current edge, and the last register holds auxiliary information, or work space, that drives the quantum walk. The quantum circuit implements the composition $X := X_t X_{t-1} \cdots X_1$, where each $X_i$ is either $W_P$ or $O_M$ acting on the edge registers, possibly controlled by the auxiliary register, or a unitary operator independent of $P$ and $M$ acting on any of the registers. The circuit $X$ is applied to a suitably constructed initial state $|\phi_0\rangle$.

We associate a cost with each operator as a measure of its complexity, with respect to a resource of interest. The resource could be circuit size, or in the query model (which is the more typical application) the number of queries. We denote the cost of implementing $W_P$ as a quantum circuit in the units of the resource of interest by $\mathsf{U}$ (*update cost*), the cost of constructing $O_M$ by $\mathsf{C}$ (*checking cost*), and the cost of preparing the initial state, $|\phi_0\rangle$, of the algorithm by $\mathsf{S}$ (*set-up cost*). Every time an operator is used, we incur the cost associated with it. This abstraction, implicit in Ref. [3] and made explicit in Ref. [30], allows $W_P$ and $O_M$ to be treated as black-box operators and provides a convenient way to capture *time complexity* or, in the quantum query model, *query complexity*. The cost of the sequence $X_t X_{t-1} \cdots X_1$ is the sum of the costs of the individual operators. The *observation probability* is the probability that we observe an element of $M$ on measuring the first register of the final state, $|\phi_t\rangle := X|\phi_0\rangle$, in the standard basis $(|x\rangle)_{x \in S}$. In the decision version of the problem, we measure a fixed single qubit of the auxiliary register in the standard basis to obtain the output of the algorithm.

## Key Results

***Walk definitions.*** Quantum walks were first introduced by David Meyer and John Watrous to study quantum cellular automata and quantum logspace, respectively. Discrete-time quantum walks were investigated for their own sake by Ambainis, Bach, Nayak, Vishwanath, and Watrous [4, 32] and Aharonov, Ambainis, Kempe, and Vazirani [2] on the infinite line and the $n$-cycle, respectively. The central issues in the early development of quantum walks included the definition of the walk operator, notions of mixing and hitting times, and the speed-up achievable compared to the classical setting.

***Hitting time.*** Exponential quantum speed-up of the hitting time between antipodes of the hypercube was shown by Kempe [20]. Childs, Cleve, Deotto, Farhi, Gutmann, and Spielman [12] presented the first oracle problem solvable exponentially faster by a quantum walk based algorithm than by any (not necessarily walk-based) classical algorithm.

The first systematic studies of quantum hitting time on the hypercube and the $d$-dimensional torus were conducted by Shenvi, Kempe, and Whaley [34] and Ambainis,

Kempe, and Rivosh [5]. Improving upon the Grover search based spatial search algorithm of Aaronson and Ambainis, Ambainis *et al.* [5] showed that the $d$-dimensional torus with $n$ nodes can be searched by quantum walk in $\sqrt{n}$ steps with observation probability $\Omega(1)$ for $d \geq 3$, and in $\sqrt{n \log n}$ in steps and observation probability $\Omega(1/\log n)$ for $d = 2$ (see also Ref. [11]). Combining the algorithm for $d = 2$ with amplitude amplification [9] we get an algorithm with observation probability $\Omega(1)$, at a cost that is a multiplicative factor of $\sqrt{\log n}$ larger.

In the results in Refs. [12, 20] the algorithm has implicit knowledge of the target state, as the walk starts from a state whose location is "related" to that of the target. It is not known if we can achieve an exponential speed-up when the walk starts in a state that is independent of the target.

***Element distinctness.*** The first result that used a quantum walk to solve a natural algorithmic problem, the so-called *element distinctness problem*, was due to Ambainis [3]. The problem is to find out if among the set of $s$ elements of a database two are identical. Ambainis constructed a walk on the *Johnson graph* $J(r, s)$ whose vertices are the $r$-size subsets of a universe of size $s$ (in his case the universe corresponds to the set of all database elements), with two subsets connected iff their symmetric difference has size two. A subset is marked, i.e., it is an element of $M$, if it captures two identical database elements. In the quantum (but also the classical) query model the set-up cost is $r$, which stands for the cost of downloading $r$ (random) database elements. Update incurs a constant cost, as it requires reading a new database element and forgetting and old one. Furthermore, since we are in the query model, the checking cost is zero, since whether a state is marked can be deduced from the currently held database elements without any further download. Ambainis ingeniously balanced the costs of $\mathsf{S}$ and $\mathsf{U}$ finding that in the quantum case the optimum choice for $r$ is $s^{2/3}$, leading to a query complexity of $s^{2/3}$ (this is a non-trivial balance: in the classical case the same walk gives no speed-up). In contrast, the Grover algorithm, the inspiration behind Ambainis' work, has no balancing option: its set-up and update costs are zero in the query model. (The Grover search may be viewed as a quantum walk on the complete graph.) It turns out that the above walk-based quantum query algorithm with complexity $\mathrm{O}(s^{2/3})$, matches the lower bound due to Aaronson and Shi [1].

***General Markov chains.*** Ambainis's result is based on the quantum hitting time of $J(r, s)$ for a marked set of relative size $\left(\frac{r}{s}\right)^2$. In Ref. [35], Szegedy investigates the hitting time of quantum walks arising from general Markov chains. His definitions (walk operator, hitting time) are abstracted directly from Ref. [3] and are consistent with prior literature, although slightly different in presentation.

For a Markov chain $P$, the (classical) *average hitting time* of $M$ can be expressed in terms of the *leaking walk matrix* $P_M$, which is obtained from $P$ by deleting all rows and columns indexed by states of $M$. Let $v_1, \ldots, v_{n-m}$, be the normalized eigenvectors of $P_M$, and let $\lambda_1, \ldots, \lambda_{n-m}$ be the associated eigenvalues, where $m = |M|$. Let $h(x, M)$ denote the expected time to reach $M$ from $x$. Let $\mu : S \to \mathbb{R}^+$ be the initial distribution from which we start, and $\mu'$ its restriction to $S \setminus M$. Denote the vector $\left(\sqrt{\mu'(x)}\right)_{x \in S \setminus M}$ by $u$. Then the average hitting time of $M$ is $h := \sum_{x \in S} \mu(x)\, h(x, M) = \sum_{k=1}^{n-m} \frac{|(v_k, u)|^2}{1-\lambda_k}$. Although the leaking walk matrix $P_M$ is not stochastic, one can consider the *absorbing walk* matrix $P' = \left[\begin{smallmatrix} P_M & P'' \\ 0 & \mathrm{I} \end{smallmatrix}\right]$, where $P''$ is the matrix obtained from $P$ by deleting the rows indexed by $M$ and the columns indexed by $S \setminus M$. The walk $P'$ behaves like $P$ but is absorbed by the first marked state it hits. Consider the quantum analogue $W_{P'}$ of $P'$ and $|\phi_0\rangle := \sum_{x \in S} \sqrt{\pi(x)}|x\rangle|p_x\rangle$, where $\pi$ is the stationary distribution of $P$. The state $|\phi_0\rangle$ is stationary for $W_P$, i.e., is an eigenvector with eigenvalue 1. Define the

*quantum hitting time*, $H$, of set $M$ to be the smallest $t$ for which $\|W_{P'}^t|\phi_0\rangle - |\phi_0\rangle\| \geq 0.1$. Note that the cost of $W_{P'}$ is proportional to $\mathsf{U} + \mathsf{C}$.

The motivation behind this definition of quantum hitting time is the following. The classical hitting time measures the number of iterations of the absorbing walk $P'$ required to noticeably skew the uniform starting distribution. Similarly, the quantum hitting time bounds the number of iterations of the following quantum algorithm for detecting whether $M$ is nonempty: At each step, apply operator $W_{P'}$. If $M$ is empty, then $P' = P$ and the starting state is left invariant. If $M$ is nonempty, then the angle between $W_{P'}^t|\phi_0\rangle$ and $W_P^t|\phi_0\rangle$ gradually increases (for $t$ not too large). Using an additional *control register* to apply either $W_{P'}$ or $W_P$ with quantum control, the divergence of these two states (should $M$ be nonempty) can be detected. The required number of iterations is characterized by $H$.

It remains to compute $H$. When $P$ is symmetric and *ergodic*, the expression for the classical hitting time has a quantum analogue [35] (we assume $m \leq n/2$ for technical reasons):

$$H \leq \sum_{k=1}^{n-m} \frac{\nu_k^2}{\sqrt{1-\lambda_k}}, \tag{1}$$

where $\nu_k = (v_k, u)$. Note that $u = \frac{1}{\sqrt{n}}(1, \ldots, 1)$, since $P$ is symmetric. So $\nu_k$ sum of the coordinates of $v_k$ divided by $1/\sqrt{n}$. From (1) and the expression for $h$ one can derive an amazing connection between the classical and quantum hitting times:

**Theorem 1 (Szegedy [35]).** *Let $P$ be symmetric and ergodic, and let $h$ be the classical hitting time for marked set $M$ and uniform starting distribution. Then the quantum hitting time of $M$ is at most $\sqrt{h}$. Therefore, the cost of solving the decision version of the problem is of order $\mathsf{S} + \sqrt{h}(\mathsf{U} + \mathsf{C})$.*

One can further show:

**Theorem 2 (Szegedy [35]).** *If $P$ is state-transitive and $|M| = 1$, then the marked state is observed with probability at least $n/h$ with cost $\mathrm{O}(\mathsf{S} + \sqrt{h}(\mathsf{U} + \mathsf{C}))$.*

The observation probability $n/h$ can be increased to $\Theta(1)$ with $\sqrt{h/n}$ iterations of the algorithm from Theorem 2, using amplitude amplification [9]. Theorems 1 and 2 imply most quantum hitting time results of the previous section *directly*, relying only on estimates of the corresponding classical hitting times. Expression (1) is based on a fundamental connection between the eigenvalues and eigenvectors of $P$ and $W_P$. Notice that $p_{y,x}^* = p_{y,x}$ for symmetric $P$, so $|p_y^*\rangle = |p_y\rangle$. So $R_1$ and $R_2$ are reflections through the subspaces generated by $\{|p_x\rangle \otimes |x\rangle |\ x \in S\}$ and $\{|x\rangle \otimes |p_x\rangle |\ x \in S\}$, respectively. The eigenvalues of $R_1 R_2$ can be expressed in terms of the eigenvalues of the mutual Gram matrix $D(P)$ of these systems. This matrix $D(P)$, the *discriminant matrix* of $P$, equals $P$ when $P$ is symmetric. The formula remains fairly simple even when $P$ is not symmetric. In particular, the absorbing walk $P'$ has discriminant matrix $\left[\begin{smallmatrix} P_M & 0 \\ 0 & \mathrm{I} \end{smallmatrix}\right]$. Finally, the relation between $D(P)$ and the spectral decomposition of $W_P$ is given by:

**Theorem 3 (Szegedy [35]).** *Let $P$ be an arbitrary Markov chain on a finite state space $S$ and let $\cos\theta_1 \geq \cdots \geq \cos\theta_l$ be those singular values of $D(P)$ lying in the open interval $(0, 1)$, with associated singular vector pairs $v_j, w_j$ for $1 \leq j \leq l$. Then the non-trivial eigenvalues of $W_P$ (namely those other than 1 and $-1$) and their corresponding eigenvectors are $(\mathrm{e}^{-2\mathrm{i}\theta_j}, R_1 w_j - \mathrm{e}^{-\mathrm{i}\theta_j} R_2 v_j)$ and $(\mathrm{e}^{2\mathrm{i}\theta_j}, R_1 w_j - \mathrm{e}^{\mathrm{i}\theta_j} R_2 v_j)$ for $1 \leq j \leq l$.*

***Subsequent developments.*** Magniez, Nayak, Roland, and Santha [29] used the Szegedy quantum analogue $W_P$ of an ergodic walk $P$, rather than that of its absorbing version $P'$, to develop a *search* algorithm in the style of Ambainis [3].

**Theorem 4 (Magniez, Nayak, Roland, Santha [29]).** *Let $P$ be reversible and ergodic with spectral gap $\delta > 0$. Let $M$ have probability either zero or $\varepsilon > 0$ under the stationary distribution of $P$. There is a quantum algorithm solving the search problem with cost $\mathsf{S} + \frac{1}{\sqrt{\varepsilon}}(\frac{1}{\sqrt{\delta}}\mathsf{U} + \mathsf{C})$.*

The main idea here is to apply quantum phase estimation [14, 22] to the quantum walk $W_P$ in order to implement an approximate reflection operator about the initial state. This operator is then used along with the checking operator $O_M$ in an amplitude amplification scheme to get the final algorithm.

The average classical hitting time $h$ may be bounded by $1/\delta\varepsilon$ (with $\delta, M, \varepsilon$ as in Theorem 4), and this bound is tight for most known applications. In these applications, the above algorithm *finds* marked elements with complexity at most that of the Szegedy algorithm. In other applications, for instance, Triangle Finding [30], where the checking cost $\mathsf{C}$ is much larger than the update cost $\mathsf{U}$, the complexity of the algorithm in Theorem 4 is asymptotically smaller.

In the case of the 2-D square grid with $n$ vertices, the average classical hitting time $h$ is $n \log n$. This is asymptotically lesser than $1/\delta\varepsilon$ when there is a single marked element. (In this case, $1/\delta\varepsilon = n^2$.) Algorithms due to Ambainis *et al.* [5] and Szegedy [35] find a unique marked state with $\mathrm{O}(\sqrt{n}\log n)$ steps of quantum walk, a $\sqrt{\log n}$ factor larger than $\sqrt{h}$. Tulsi [36] showed how we may find a unique marked element in $\mathrm{O}(\sqrt{h})$ steps. Magniez, Nayak, Richter, and Santha [28] extended this result to show that for any state-transitive Markov chain, a unique marked state can be found in $\mathrm{O}(\sqrt{h}\,)$ steps. They also devised a detection algorithm that solves the decision version of the problem for any reversible Markov chain and any number of marked elements, in $\mathrm{O}(\sqrt{h})$ steps (thus extending Theorem 1).

Krovi, Magniez, Ozols, and Roland [24] presented a different quantum algorithm for finding multiple marked elements in any *reversible* Markov chain. They introduced a notion of interpolation between any reversible chain $P$ and its absorbing counterpart $P'$, and used the quantum analogue of the interpolated walk. In the case of a unique marked element, the resulting algorithm solves the search version of the problem with cost $\mathsf{S} + \sqrt{h}(\mathsf{U} + \mathsf{C})$. The precise relationship between the number of steps of the quantum walk taken by the algorithm in the case of more than one marked element, and the corresponding classical hitting time remains open. It is known that for certain choices of $P$ and $M$, the former may be asymptotically larger than $\sqrt{h}$.

The schema due to Magniez *et al.* [29] described above has been extended in different ways. Jeffery, Kothari, and Magniez [18] use a *quantum state* as the data structure $d(x)$ associated with a state $x \in S$ in quantum algorithms with nested walks. In this manner, they avoid the repeated overhead of set-up cost in the inner quantum walks used for checking marked states. They solve several problems, including Triangle Finding, with query as well as time complexity matching, up to polylogarithmic factors, the performance of algorithms previously derived from *learning graphs* [7, 26]. Childs, Jeffery, Kothari, and Magniez [8] introduced the use of a data structure that depends on the state transition in the walk. Using this, they develop quantum algorithms with nested walks, where the recursion occurs in the update operation. The cost incurred is essentially what we would expect from Theorem 4. This extension leads to algorithms that are as efficient in *time* as in query complexity, for applications such as 3-Distinctness. Independently, Belovs designed a different quantum walk algorithm [8], which leads to a similar result for 3-Distinctness.

## Applications

We list some quantum walk based results for search problems that represent speed-ups over Grover search based solutions. All are inspired by Ambainis' algorithm for element distinctness.

***Triangle finding.*** Suppose we are given the adjacency matrix $A$ of a graph on $n$ vertices and are required to determine if the graph contains a triangle (i.e., a clique of size 3), using as few queries as possible to the entries of $A$. The classical query complexity of this problem is $\Theta(n^2)$. Magniez, Santha, and Szegedy [30] gave an $\tilde{\mathrm{O}}(n^{1.3})$ algorithm. This upper bound has been improved by a sequence of results [7, 16, 26, 29] (see also Ref. [18]) to $\tilde{\mathrm{O}}(n^{5/4})$. Several of these algorithms, including the current best algorithm due to Le Gall [16] are based on the quantum walk search framework.

***Matrix product verification and matrix multiplication.*** Suppose we are given three $n \times n$ matrices $A$, $B$, $C$ over a ring and are required to determine if $AB \neq C$, i.e., if there exist $i$, $j$ such that $\sum_k A_{ik}B_{kj} \neq C_{ij}$. We would like to make as few queries as possible to the entries of $A$, $B$, and $C$. This problem has classical query complexity $\Theta(n^2)$. Buhrman and Špalek [10] gave an $\mathrm{O}(n^{5/3})$ quantum query algorithm. They also observed that two Boolean matrices can be multiplied with query complexity $\mathrm{O}(n^{3/2}\sqrt{\ell}\,)$, where $\ell$ is the number of non-zero entries in the product. This has since been improved in a sequence of results [17, 25, 37] to $\mathrm{O}(n\sqrt{\ell}\,)$. The algorithm due to Le Gall [25] builds upon quantum walk algorithms. We refer the reader to Ref. [23] for further work on this topic.

***Group commutativity testing.*** Suppose we are presented with a black-box group specified by its $k$ generators and are required to determine if the group commutes using as few queries as possible to the group product operation (i.e., queries of the form "What is the product of elements $g$ and $h$?"). The classical query complexity is $\Theta(k)$ group operations. Magniez and Nayak [27] gave an (essentially optimal) $\tilde{\mathrm{O}}(k^{2/3})$ quantum query algorithm for this problem. The algorithm involves a quantum walk on the product of two graphs whose vertices are ordered $l$-tuples of distinct generators.

***Forbidden subgraph property.*** A property of graphs is called *minor closed* when the following condition holds: if a graph has the property, then all its minors also possess the property. A graph property (which need not be minor closed) is called a *forbidden subgraph property* (FSP) if it can be described by a finite set of forbidden subgraphs. Suppose we are given the adjacency matrix $A$ of a graph on $n$ vertices and are required to determine if the graph has a minor closed property $\Pi$, using as few queries as possible to the entries of $A$. Childs and Kothari [13] show that if $\Pi$ is non-trivial and is *not* FSP, then it has query complexity in $\Theta(n^{3/2})$. They complement this with a more efficient algorithm for any minor closed property $\Pi$ that *is* FSP. The algorithm has query complexity $\mathrm{O}(n^\alpha)$ for some $\alpha < 3/2$, and is based on the quantum walk search framework.

***3-Distinctness.*** This is a generalization of the element distinctness problem. Suppose we are given elements $x_1, \ldots, x_m \in \{1, \ldots, m\}$ and are asked if there exist *three* distinct indices $i, j, k$ such that $x_i = x_j = x_k$. The Ambainis quantum walk algorithm achieves query and time complexity $\mathrm{O}(m^{3/4})$. The *query* complexity was improved to $\mathrm{O}(m^{5/7})$ by Belovs [6] using a new technique—learning graphs, while the best time complexity remained unchanged. Childs *et al.* [8] later designed *time* efficient query algorithms with complexity $\tilde{\mathrm{O}}(m^{5/7})$, using extensions of the quantum walk search framework.

# Open Problems

Many issues regarding quantum analogues of Markov chains remain unresolved, both for the search problem and the closely related mixing problem.

***Search problem.*** Can the quadratic quantum speed-up of hitting time for the decision version of the problem be extended from all reversible Markov chains to all *ergodic* ones? Can quantum walks also *find* marked elements quadratically faster than classical walks, in the case of reversible Markov chains with *multiple* marked states? What other algorithmic applications of search by quantum walk can be found?

***Sampling problem.*** Another wide use of Markov chains in classical algorithms is in generating samples from certain probability distributions. In particular, *Markov chain Monte Carlo* algorithms work by running a carefully designed ergodic Markov chain. After a number of steps given by the *mixing time* of $P$, the distribution over states is guaranteed to be $\epsilon$-close to its stationary distribution $\pi$. Such algorithms form the basis of most randomized algorithms for approximating #P-complete problems (see, e.g., Ref. [19]). The sampling problem may be formalized as follows:

INPUT: Markov chain $P$, tolerance $\epsilon \in (0, 1)$.
OUTPUT: A sample from a distribution that is $\epsilon$-close to $\pi$ in total variation distance.

Notions of quantum mixing time were first proposed and analyzed on the line, the cycle, and the hypercube [2, 4, 31, 32]. Kendon and Tregenna [21] and Richter [33] have investigated the use of decoherence in improving mixing of quantum walks. Two fundamental questions about quantum mixing time remain open: What is the "most natural" definition? And, when is there a quantum speed-up over the classical mixing time?

# Recommended Reading


1. Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, July 2004.
2. Dorit Aharonov, Andris Ambainis, Julia Kempe, and Umesh Vazirani. Quantum walks on graphs. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 50–59, New York, NY, USA, 2001. ACM.
3. Andris Ambainis. Quantum walk algorithm for Element Distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007.
4. Andris Ambainis, Eric Bach, Ashwin Nayak, Ashvin Vishwanath, and John Watrous. One-dimensional quantum walks. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 37–49, New York, NY, USA, 2001. ACM.
5. Andris Ambainis, Julia Kempe, and Alexander Rivosh. Coins make quantum walks faster. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 1099–1108, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
6. Aleksandrs Belovs. Learning-graph-based quantum algorithm for $k$-Distinctness. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 207–216, Los Alamitos, CA, USA, Oct 2012. IEEE Computer Society.
7. Aleksandrs Belovs. Span programs for functions with constant-sized 1-certificates: Extended abstract. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 77–84, New York, NY, USA, 2012. ACM.
8. Aleksandrs Belovs, Andrew M. Childs, Stacey Jeffery, Robin Kothari, and Frédéric Magniez. Time-efficient quantum walks for 3-Distinctness. In Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming*, volume 7965 of *Lecture Notes in Computer Science*, pages 105–122. Springer Berlin Heidelberg, 2013.

9. Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum computation and information (Washington, DC, 2000)*, volume 305 of *Contemporary Mathematics*, pages 53–74. American Mathematical Society, Providence, RI, 2002.
10. Harry Buhrman and Robert Špalek. Quantum verification of matrix products. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '06, pages 880–889, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics.
11. Andrew Childs and Jeffrey Goldstone. Spatial search by quantum walk. *Physical Review A*, 70:022314, Aug 2004.
12. Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 59–68, New York, NY, USA, 2003. ACM.
13. Andrew M. Childs and Robin Kothari. Quantum query complexity of minor-closed graph properties. *SIAM Journal on Computing*, 41(6):1426–1450, 2012.
14. Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998.
15. Edward Farhi and Sam Gutmann. Quantum computation and decision trees. *Physical Review A*, 58:915–928, Aug 1998.
16. François Le Gall. Improved quantum algorithm for triangle finding via combinatorial arguments. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, pages 216–225, Los Alamitos, CA, USA, October 18–21 2014. IEEE Computer Society Press.
17. Stacey Jeffery, Robin Kothari, and Frédéric Magniez. Improving quantum query complexity of boolean matrix multiplication using graph collision. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming*, volume 7391 of *Lecture Notes in Computer Science*, pages 522–532. Springer Berlin Heidelberg, 2012.
18. Stacey Jeffery, Robin Kothari, and Frédéric Magniez. Nested quantum walks with quantum data structures. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1474–1485, Philadelphia, PA, USA, 2013. SIAM.
19. Mark Jerrum and Alistair Sinclair. The Markov Chain Monte Carlo method: An approach to approximate counting and integration. In Dorit S. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*, pages 482–520. PWS Publishing Co., Boston, MA, USA, 1997.
20. Julia Kempe. Discrete quantum walks hit exponentially faster. *Probability Theory and Related Fields*, 133(2):215–235, 2005.
21. Viv Kendon and Ben Tregenna. Decoherence can be useful in quantum walks. *Physical Review A*, 67:042315, Apr 2003.
22. Alexei Kitaev. Quantum measurements and the Abelian stabilizer problem. Technical Report quant-ph/9511026, arXiv.org, 1995.
23. Robin Kothari. *Efficient Algorithms in Quantum Query Complexity*. PhD thesis, University of Waterloo, Waterloo, ON, Canada, 2014.
24. Hari Krovi, Frédéric Magniez, Maris Ozols, and Jérémie Roland. Quantum walks can find a marked element on any graph. Technical Report arXiv:1002.2419v2, arXiv.org, 2014.
25. François Le Gall. Improved output-sensitive quantum algorithms for boolean matrix multiplication. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 1464–1476, Philadelphia, PA, USA, 2012. SIAM.
26. Troy Lee, Frédéric Magniez, and Miklos Santha. Improved quantum query algorithms for triangle finding and associativity testing. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1486–1502, Philadelphia, PA, USA, 2013. SIAM.
27. Frederic Magniez and Ashwin Nayak. Quantum complexity of testing group commutativity. *Algorithmica*, 48(3):221–232, July 2007.
28. Frédéric Magniez, Ashwin Nayak, Peter C. Richter, and Miklos Santha. On the hitting times of quantum versus random walks. *Algorithmica*, 63(1):91–116, 2012.
29. Frédéric Magniez, Ashwin Nayak, Jérémie Roland, and Miklos Santha. Search via quantum walk. *SIAM Journal on Computing*, 40:142–164, February 10 2011.
30. Frédéric Magniez, Miklos Santha, and Mario Szegedy. Quantum algorithms for the triangle problem. *SIAM Journal on Computing*, 37(2):413–424, 2007.
31. Cristopher Moore and Alexander Russell. Quantum walks on the hypercube. In José D.P. Rolim and Salil Vadhan, editors, *Randomization and Approximation Techniques in Computer Science*,

volume 2483 of *Lecture Notes in Computer Science*, pages 164–178. Springer Berlin Heidelberg, 2002.
32. Ashwin Nayak and Ashvin Vishwanath. Quantum walk on the line. Technical Report quant-ph/0010117, arXiv.org, 2000.
33. Peter Richter. Quantum speedup of classical mixing processes. *Physical Review A*, 76:042306, Oct 2007.
34. Neil Shenvi, Julia Kempe, and K. Birgitta Whaley. Quantum random-walk search algorithm. *Physical Review A*, 67:052307, May 2003.
35. Mario Szegedy. Quantum speed-up of Markov chain based algorithms. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 32–41, Los Alamitos, CA, USA, Oct 2004. IEEE Computer Society.
36. A. Tulsi. Faster quantum walk algorithm for the two dimensional spatial search. *Physical Review A*, 78:012310, 2008.
37. Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 645–654, Washington, DC, USA, 2010. IEEE Computer Society.