Continuous-Time Quantum Walks on the Symmetric Group

Heath Gerhardt and John Watrous

Department of Computer Science
University of Calgary
Calgary, Alberta, Canada
{gerhardt, jwatrous}@cpsc.ucalgary.ca

Abstract. In this paper we study continuous-time quantum walks on Cayley graphs of the symmetric group, and prove various facts concerning such walks that demonstrate significant differences from their classical analogues. In particular, we show that for several natural choices for generating sets, these quantum walks do not have uniform limiting distributions, and are effectively blind to large areas of the graphs due to destructive interference.

1 Introduction

According to our current understanding of physics, quantum mechanics provides sources of true randomness, and mathematically speaking much of the underlying framework of quantum information and computation may be viewed as an extension of the study of random processes. The focus in quantum information and computation is often placed on finding information processing tasks that can be performed with the help of quantum information (such as factoring integers in polynomial time [20] or implementing unconditionally secure key distribution [7,21]) or on studying the distinctively non-classical aspects of quantum information (such as entanglement; see, for instance, [12]). However, it seems quite plausible that the study of quantum information and computation will also lead to new methods in the study of classical computation and random processes. Along these lines, Kerenidis and de Wolf [17] recently used quantum arguments to prove new results on (classical) locally decodable codes.

As a step toward understanding the possible implications of quantum methods for the study of random processes, it is natural to consider the differences between classical and quantum processes. One of the topics that has recently received attention in the quantum computing community that highlights these differences is the the study of quantum computational variants of random walks, or *quantum walks* [1,3,5,6,8,9,11,15,18,19,23]. (A recent survey on quantum walks by Kempe [16] is an ideal starting point for background on quantum walks.) In this paper we consider quantum walks on Cayley graphs of the symmetric group—a topic that has been suggested in at least two previous papers on quantum walks [16,3].

S. Arora et al. (Eds.): APPROX 2003+RANDOM 2003, LNCS 2764, pp. 290–301, 2003.

Two main variants of quantum walks have been considered: continuous-time quantum walks and discrete-time quantum walks. We restrict our attention to continuous-time quantum walks in this paper. Keeping in line with previous results on quantum walks, we find some significant differences between quantum and classical random walks on Cayley graphs of the symmetric group. In particular, we find that quantum walks on Cayley graphs of the symmetric group do not have uniform limiting distributions for several natural choices for the generators. This answers a question recently suggested by Ahmadi, Belk, Tamon, and Wendler [3] concerning non-uniform mixing of quantum walks.

One of the principle motivations for studying quantum walks has been that quantum walks may potentially be useful as algorithmic tools. This potential was recently demonstrated by Childs, Cleve, Deotto, Farhi, Gutmann and Spielman [8], who prove that there exists a black-box problem for which a quantum algorithm based on quantum walks gives an exponential speed-up over any classical randomized algorithm. The key to this algorithm is that a quantum walk is able to permeate a particular graph while any classical random walk (or any classical randomized algorithm, for that matter) cannot. One of the first problems that comes to mind as an obvious challenge for the quantum algorithms community is the graph isomorphism problem, and it is natural to ask whether quantum walks, and in particular quantum walks on Cayley graphs of the symmetric group, can be of any use for an algorithm for this problem. (While this was our primary motivation for studying quantum walks on the symmetric group, we have not found any way to apply our results to this problem.)

2 Definitions

2.1 Continuous-Time Quantum Walks on Graphs

A continuous-time quantum walk on an undirected graph $\Gamma = (V, E)$ can be defined in the following way. First, we let A be the $|V| \times |V|$ adjacency matrix of Γ, let D be the $|V| \times |V|$ diagonal matrix for which the diagonal entry corresponding to vertex v is $\deg(v)$, and let $L = D - A$. The matrix L is positive semidefinite and, under the assumption that Γ is connected, 0 is an eigenvalue with multiplicity 1; the uniform vector is a corresponding eigenvector. The quantum walk on Γ is then given by the unitary matrix $U(t) = e^{-itL}$ for $t \in \mathbb{R}$. If the quantum walk on Γ is run for time t starting at vertex u, then the amplitude associated with each vertex v is $U(t)[v, u]$, and thus measuring at this point (with respect to the standard basis) results in each vertex v with probability $|U(t)[v, u]|^2$. If instead of starting at a particular vertex u we have some quantum state described by $\psi : V \to \mathbb{C}$, and we run the quantum walk for time t, the new quantum state is described by $U(t)\psi$, and measuring results in each vertex v with probability $|(U(t)\psi)[v]|^2$. Other types of measurements can be considered, but we will focus just on this sort of measurement where the outcome is a vertex of the graph. To our knowledge, continuous-time quantum walks were first considered by Farhi and Gutmann [11].

Continuous-time quantum walks are analogous to continuous-time random walks on Γ, where the evolution is described by $M(t) = e^{-tL}$ rather than $U(t)$ as above. Specifically, if the continuous-time random walk is started at vertex u and run for time t, the probability of being at vertex v is given by $M(t)[v, u]$. Continuous-time random walks share many properties with their discrete-time variants [4].

This paper is concerned with quantum walks on Cayley graphs, which are regular graphs. In the case of regular graphs there is no difference between using the matrix L and the adjacency matrix for the definition of quantum walks, and we find it is more convenient to use the adjacency matrix for the graphs we are considering. (Of course one cannot replace L with the adjacency matrix when discussing the classical case, since this would not give rise to a stochastic process—the equivalence only holds for the quantum case.) The reasoning behind this equivalence is as follows. Because D and A commute for regular graphs, we see that $U(t) = e^{-itdI}e^{itA} = e^{-itd}e^{itA}$; the difference is a global phase factor, which has no significance when calculating the probabilities. So, from here after in this paper we will consider the unitary process given by $U(t) = e^{itA}$ rather than e^{-itL}.

In the case of classical random walks, there are various properties of random walks that are of interest. One of the most basic properties of a classical random walk is the limiting distribution (or stationary distribution). This distribution is the uniform distribution for random walks on connected, regular graphs, and in fact as a result of the way we have defined continuous-time random walks this distribution is uniform for any connected, undirected graph; this is apparent by considering the spectral decomposition of the matrix e^{-tL}.

As quantum walks are unitary (and therefore invertible) processes, they do not converge to any state, so one must be precise about what is meant by the limiting distribution. Suppose we have a quantum walk on some graph Γ and some vertex u has been designated as the starting vertex. The probability of measuring the walk at some vertex v after time t is, as described above, given by $P_t[v] = |U(t)[v, u]|^2$. If t is chosen uniformly from some range $[0, T]$ then the resulting distribution is

$$\bar{P}_T[v] = \frac{1}{T}\int_0^T P_t[v]\mathrm{d}t.$$

In the limit for large T these distributions converge to some distribution $\bar{P}$, which is the limiting distribution of the quantum walk. This notion of the limiting distribution for a quantum walk is discussed in [1].

2.2 Cayley Graphs and Representation Theory of the Symmetric Group

In this section we briefly discuss necessary background information on Cayley graphs of the symmetric group and on representation theory of the symmetric group, which is the main tool used in this paper to analyze quantum walks on Cayley graphs.

Let G be a finite group and let $R \subseteq G$ be a set of generators for G satisfying $g \in R \Leftrightarrow g^{-1} \in R$ for all $g \in G$. Then the Cayley graph of G with respect to R, which we denote by $\Gamma(G, R)$ in this paper, is an undirected graph defined as follows. The set of vertices of $\Gamma(G, R)$ coincides with G, and for any $g, h \in G$, $\{g, h\}$ is an edge in $\Gamma(G, R)$ if and only if $gh^{-1} \in R$. Equivalently, if $R = \{h_1, \ldots, h_d\}$ then each vertex g is adjacent to vertices $h_1 g, \ldots, h_d g$. Thus, $\Gamma(G, R)$ is a regular graph of degree $d = |R|$. We will restrict our attention to generating sets that form conjugacy classes. (The method we use for analyzing quantum walks on Cayley graphs is limited to such generating sets.) Recall that for some group G, elements g and h are conjugate if there exists some $a \in G$ such that $a^{-1}ga = h$. This is an equivalence relation that partitions G into conjugacy classes. A function $f : G \to \mathbb{C}$ is a *class function* if it is constant on conjugacy classes of G.

The conjugacy classes in S_n are determined by the cycle structures of elements when they are expressed in the usual cycle notation. Recall that a partition λ of n is a sequence $(\lambda_1, \ldots, \lambda_k)$ where $\lambda_1 \geq \cdots \geq \lambda_k \geq 1$ and $\lambda_1 + \cdots + \lambda_k = n$. The notation $\lambda \vdash n$ indicates that λ is a partition of n. There is one conjugacy class for each partition $\lambda \vdash n$ in S_n, which consists of those permutations having cycle structure described by λ. We denote by C_λ the conjugacy class of S_n consisting of all permutations having cycle structure described by λ.

A *representation* of a group G is a homomorphism from G to $\mathrm{GL}(d, \mathbb{C})$ for some positive integer d, where $\mathrm{GL}(d, \mathbb{C})$ denotes the general linear group of invertible $d \times d$ complex matrices. The *dimension* of such a representation is d, and we write $\dim(\rho)$ to denote the dimension of a given representation ρ. Two representations $\rho_1 : G \to \mathrm{GL}(d_1, \mathbb{C})$ and $\rho_2 : G \to \mathrm{GL}(d_2, \mathbb{C})$ are *equivalent* if there exists an invertible linear mapping $A : \mathbb{C}^{d_1} \to \mathbb{C}^{d_2}$ such that $A\rho_1(g) = \rho_2(g)A$ for all $g \in G$, otherwise they are *inequivalent*. A representation ρ of dimension d is *irreducible* if there are no non-trivial invariant subspaces of $\mathbb{C}^d$ under ρ. That is, if $W \subseteq \mathbb{C}^d$ is a subspace of $\mathbb{C}^d$ such that $\rho(g)W \subseteq W$ for all $g \in G$, then $W = \mathbb{C}^d$ or $W = \{0\}$. A collection of inequivalent, irreducible representations is said to be *complete* if every irreducible representation is equivalent to one of the representations in this set. It holds that any complete set of irreducible representations can be put into one-to-one correspondence with the conjugacy classes of the group in question.

The character corresponding to a representation ρ is a mapping $\chi_\rho : G \to \mathbb{C}$ obtained by taking the trace of the representation: $\chi_\rho(g) = \mathrm{tr}(\rho(g))$. Using the cyclic property of the trace it follows that the characters are constant on the conjugacy classes of a group. If we have a complete set of inequivalent, irreducible representations of a group, then the corresponding characters form an orthogonal basis for the space of all class functions.

The Fourier transform $\hat{f}$ of a complex-valued function f on G at a representation ρ is $\hat{f}(\rho) = \sum_{g \in G} f(g)\rho(g)$.

Fact 1 *Let f be a class function on a group G and ρ be an irreducible representation of G, then* $\hat{f}(\rho) = \frac{1}{\dim(\rho)} \left(\sum_{g \in G} f(g)\chi_\rho(g) \right) I$.

For the symmetric group on n elements there is a particular way of associating the partitions of n (which are in one-to-one correspondence with the conjugacy classes of S_n) with a complete set of inequivalent, irreducible representations of S_n. These particular representations are said to be in *Young normal form*. (Several text books describe the specific method for constructing these representations—see, for instance, James and Kerber [14]. It will not be important for this paper to discuss the actual construction of these representations.) These representations have the special property that all matrix entries in these representations are integers. Once we have these irreducible representations, it is possible to associate with each one an equivalent irreducible representation that has the property that $\rho(g)$ is a unitary matrix for every $g \in S_n$. The irreducible, unitary representation associated with a given partition $\lambda \vdash n$ will be denoted ρ_λ, and the corresponding character will be denoted χ_λ. The following fact will be a useful fact regarding these representations.

Fact 2 *Let λ and μ be partitions of n and let ρ_λ and ρ_μ be the associated unitary representations as described above. Then for all $1 \leq i, j \leq \dim(\rho_\lambda)$ and $1 \leq k, l \leq \dim(\rho_\mu)$,*

$$\sum_{g \in S_n} \rho_\lambda(g)[i,j]\overline{\rho_\mu(g)[k,l]} = \begin{cases} \frac{n!}{\dim(\rho_\lambda)} & \text{if } \lambda = \mu,\ i = k,\ \text{and } j = l \\ 0 & \text{otherwise} \end{cases}$$

When $\lambda, \nu \vdash n$, we write $\chi_\lambda(\nu)$ to denote the character χ_λ evaluated at an arbitrary $g \in C_\nu$, and more generally if f is a class function we write $f(\nu)$ to mean $f(g)$ for any $g \in C_\nu$.

Fact 3 *The sum of the squares of the characters of a conjugacy class over any complete, irreducible set of representations of a group G multiplied by the order of the class is the order of G. Thus, we have $|C_\lambda| \sum_{\nu \vdash n} \chi_\nu(\lambda)^2 = n!$ for every $\lambda \vdash n$.*

It will be necessary for us to be able to evaluate the characters associated with the irreducible representations of the symmetric group in certain instances. The Murnaghan-Nakayama rule provides a tool for doing this—information on the Murnaghan-Nakayama rule can be found in [22].

3 Continuous-Time Quantum Walks on $\Gamma(S_n, C_\lambda)$

In this section we analyze the quantum walk on $\Gamma(S_n, C_\lambda)$ for $\lambda \vdash n$. Our analysis implies that for some natural choices for λ the quantum walk on $\Gamma(S_n, C_\lambda)$ does not have a uniform limiting distribution with respect to the definition discussed in the previous section. In essence, the quantum walk has a significant "blind spot" consisting of all n-cycles (i.e., permutations having cycle-structure consisting of a single n-cycle).

This section is divided into three subsections. First we prove a general result concerning the spectral decomposition of quantum walks on S_n. We then consider

the case where the generating set consists of the set of all transposition, and finally the case where the generating set consists of all p-cycles for any choice of $p \in \{2, \ldots, n\}$.

3.1 Spectral Decomposition and Periodicity

Define $c_\lambda : S_n \to \mathbb{C}$ to be the unit vector that is uniform on the conjugacy class C_λ and zero everywhere else: $c_\lambda[g] = |C_\lambda|^{-1/2}$ if $g \in C_\lambda$, and $c_\lambda[g] = 0$ otherwise.

The analysis of quantum walks on $\Gamma(S_n, C_\lambda)$ is greatly simplified by the fact that these walks are constant on conjugacy classes, in the following sense.

Proposition 4 *Let $\alpha_t(g)$ denote the amplitude associated with vertex g after evolving the quantum walk on $\Gamma(S_n, C_\lambda)$ for time t, assuming the walk starts on a conjugacy class, i.e., $\alpha_t(g) = (U(t)c_\lambda)[g]$. Then for all t, α_t is a class function.*

The following theorem will be one of the main tools used in our analysis.

Theorem 5. *Assume $H[g,h] = f(g^{-1}h)$ for all $g, h \in S_n$, where f a class function on S_n, and let $U(t) = e^{itH}$ for all $t \in \mathbb{R}$. Then for any partitions $\lambda, \mu \vdash n$ we have*

$$c_\lambda^* U(t) c_\mu = \frac{\sqrt{|C_\lambda|}\sqrt{|C_\mu|}}{n!} \sum_{\nu \vdash n} \exp\left(\frac{it}{\dim(\rho_\nu)} \sum_{\gamma \vdash n} |C_\gamma| f(\gamma) \chi_\nu(\gamma)\right) \chi_\nu(\lambda)\chi_\nu(\mu).$$

In order to prove this theorem we will use the following lemma, by which a complete orthogonal set of eigenvectors and eigenvalues of $U(t)$ can be obtained.

Lemma 1. *Assume $H[g,h] = f(g^{-1}h)$ for all $g, h \in S_n$, where f is a class function on S_n. Define vectors $\psi_{\nu,i,j} : S_n \to \mathbb{C}$ for each $\nu \vdash n$, $1 \le i, j \le \dim(\rho_\nu)$ by $\psi_{\nu,i,j}[g] = \rho_\nu(g)[i,j]$ for all $g \in S_n$. Then each $\psi_{\nu,i,j}$ is an eigenvector of H with associated eigenvalue $\frac{1}{\dim(\rho_\nu)} \sum_{\gamma \vdash n} |C_\gamma| f(\gamma) \chi_\nu(\gamma)$. Moreover, these eigenvectors are pairwise orthogonal and span the space $\mathbb{C}^{S_n}$.*

Remark. The fact described in Lemma 1 is not new—for instance, it is discussed in Section 3E of [10] for general finite groups. A short proof of the lemma follows.

Proof of Lemma 1. For each $g \in S_n$ we have

$$(H\psi_{\nu,i,j})[g] = \sum_{h \in G} f(g^{-1}h)\rho_\nu(h)[i,j] = \sum_{h \in G} f(h)\rho_\nu(gh)[i,j].$$

Now, since ρ_ν is a homomorphism, we have $\rho_\nu(gh) = \rho_\nu(g)\rho_\nu(h)$, which implies

$$\begin{aligned}(H\psi_{\nu,i,j})[g] &= \sum_{k=1}^{\dim(\rho_\nu)} \rho_\nu(g)[i,k] \left(\sum_{h \in S_n} f(h)\rho_\nu(h)\right)[k,j] \\ &= \sum_{k=1}^{\dim(\rho_\nu)} \rho_\nu(g)[i,k]\hat{f}(\rho_\nu)[k,j].\end{aligned}$$

By Fact 1 we see that

$$\begin{aligned}(H\psi_{\nu,i,j})[g] &= \frac{1}{\dim(\rho_\nu)} \sum_{h \in S_n} f(h)\chi_\nu(h)\rho_\nu(g)[i,j] \\ &= \left(\frac{1}{\dim(\rho_\nu)} \sum_{\gamma \vdash n} |C_\gamma| f(\gamma)\chi_\nu(\gamma) \right) \psi_{\nu,i,j}[g].\end{aligned}$$

This establishes that the vectors $\psi_{\nu,i,j}$ are eigenvectors with associated eigenvalues as claimed. The fact that these eigenvectors are pairwise orthogonal follows from Fact 2 and the fact that they span the entire space $\mathbb{C}^{S_n}$ follows from this orthogonality along with Fact 3. ■

Proof of Theorem 5. By Lemma 1 we may write

$$H = \sum_{\nu,j,k} \left(\frac{1}{\dim(\rho_\nu)} \sum_{\gamma} |C_\gamma| f(\gamma)\chi_\nu(\gamma) \right) \frac{\psi_{\nu,j,k}\psi^*_{\nu,j,k}}{\|\psi_{\nu,j,k}\|^2}$$

and therefore

$$U(t) = \sum_{\nu,j,k} \exp\left(\frac{it}{\dim(\rho_\nu)} \sum_{\gamma} |C_\gamma| f(\gamma)\chi_\nu(\gamma) \right) \frac{\psi_{\nu,j,k}\psi^*_{\nu,j,k}}{\|\psi_{\nu,j,k}\|^2}$$

Let $X_\lambda : S_n \to \mathbb{C}$ denote the characteristic function of C_λ for $\lambda \vdash n$. Then we have that

$$c^*_\lambda \psi_{\nu,j,k} = \frac{1}{\sqrt{|C_\lambda|}} \hat{X}_\lambda(\rho_\nu)[j,k] = \begin{cases} \frac{\sqrt{|C_\lambda|}\chi_\nu(\lambda)}{\dim(\rho_\nu)} & \text{if } j = k \\ 0 & \text{otherwise} \end{cases}$$

by Fact 1. By Fact 2 we have $\|\psi_{\nu,j,k}\|^2 = \frac{n!}{\dim(\rho_\nu)}$. So,

$$\begin{aligned} & c^*_\lambda U(t) c_\mu \\ &= \frac{1}{n!} \sum_{\nu} \exp\left(\frac{it}{\dim(\rho_\nu)} \sum_{\gamma \vdash n} |C_\gamma| f(\gamma)\chi_\nu(\gamma) \right) \dim(\rho_\nu) \sum_{j,k=1}^{\dim(\rho_\nu)} c^*_\lambda \psi_{\nu,j,k}\, \psi^*_{\nu,j,k} c_\mu \\ &= \frac{\sqrt{|C_\lambda|}\sqrt{|C_\mu|}}{n!} \sum_{\nu \vdash n} \exp\left(\frac{it}{\dim(\rho_\nu)} \sum_{\gamma \vdash n} |C_\gamma| f(\gamma)\chi_\nu(\gamma) \right) \chi_\nu(\lambda)\chi_\nu(\mu), \end{aligned}$$

which is what we wanted to show. ■

Theorem 5 implies the following interesting fact.

Proposition 6 *Any continuous-time quantum walk on the Cayley graph of the symmetric group for which the generators form conjugacy classes is periodic, with period $2\pi/k$ for some $k \in \{1, 2, 3, \ldots\}$.*

Proof. Using Fact 1 we see that the quantity $|C_\gamma|\chi_\nu(\gamma)/\dim(\rho_\nu)$ is a sum of matrix elements of irreducible representations. This quantity is independent of the particular choice of the basis for the irreducible representations, so we may choose that basis that corresponds to Young's natural form, in which all of the matrix entries are integer valued, implying that the quantity itself is integer valued. Using Fact 3 and Theorem 5 therefore have that $U(2\pi) = U(0) = I$. Thus the period of the walk must divide 2π. ■

We have not discussed mixing times in this paper, but the previous proposition implies that quantum walks on Cayley graphs of S_n reach their limiting distribution quickly, and when calculating the limiting distribution it is only necessary to average over times in the range $[0, 2\pi]$. Note that in terms of implementation, this does not mean that the walk mixes in constant time; some number of operations that is polynomial in the degree of the graph and in some accuracy parameter is required to implement such a walk, assuming the ability to compute the neighbors of each vertex. See [2,8] for further details.

3.2 Cayley Graphs of S_n Generated by Transpositions

For the Cayley graph of S_n generated by the transpositions, Theorem 5 has various implications that we discuss in this section. We will require explicit values for various characters of the symmetric group, which we now mention. Using the Murnaghan-Nakayama rule it can be shown that

$$\chi_\nu\left((n)\right) = \begin{cases} (-1)^{n-k} & \text{for } \nu = (k, 1, \ldots, 1),\ k \in \{1, \ldots, n\} \\ 0 & \text{otherwise} \end{cases}$$

and $\chi_{(k,1,\ldots,1)}\left(\mathrm{id}\right) = \dim(\rho_{(k,1,\ldots,1)}) = \binom{n-1}{k-1}$. For the characters at the transpositions, it is known [13] that

$$\chi_\nu(\tau) = \frac{\dim(\rho_\nu)}{\binom{n}{2}} \sum_j \left(\binom{\nu_j}{2} - \binom{\nu'_j}{2} \right).$$

Here, τ is any transposition, ν' is the partition generated by transposing the Young diagram of ν, while ν_j and ν'_j are the j^{th} components of the partitions ν and ν'. Substituting these values into Theorem 5 gives

$$c_\lambda^* U(t) c_\mu = \frac{\sqrt{|C_\lambda|}\sqrt{|C_\mu|}}{n!} \sum_{\nu \vdash n} \exp\left(it \sum_j \left(\binom{\nu_j}{2} - \binom{\nu'_j}{2} \right) \right) \chi_\nu(\lambda)\chi_\nu(\mu)$$

for the quantum walk on $\Gamma(S_n, C_{(2,1,\ldots,1)})$, and specifically for the case where $\mu = (1, \ldots, 1)$ and $\lambda = (n)$ it follows that

$$\begin{aligned} & c_{(n)}^* U(t) c_{(1,\ldots,1)} \\ &= \frac{1}{\sqrt{n \cdot n!}} \sum_{k=1}^{n} \exp\left(it \left(\binom{k}{2} - \binom{n-k+1}{2} \right) \right) (-1)^{n-k} \binom{n-1}{k-1} \\ &= \frac{\left(2i \sin(tn/2)\right)^{n-1}}{\sqrt{n \cdot n!}}. \end{aligned}$$

In particular,

$$\max_t \left| c^*_{(n)} U(t) c_{(1,\ldots,1)} \right|^2 = \frac{2^{2n-2}}{n \cdot n!}, \tag{1}$$

where the maximum occurs for $t = (2k+1)\pi/n$, $k \in \mathbb{Z}$.

Eq. 1 has the following interpretation. If we start a quantum walk on $\Gamma(S_n, C_{(2,1,\ldots,1)})$ at the identity element and evolve for any amount of time and measure, the probability to measure some n-cycle is at most $\frac{2^{2n-2}}{n \cdot n!}$ as opposed to probability approaching $\frac{1}{n}$ for the classical case. The probability to measure any particular n-cycle is therefore at most $\frac{2^{2n-2}}{(n!)^2}$, as opposed to some number approaching $\frac{1}{n!}$ classically. The probabilities in the quantum case are smaller by a factor that is exponential in n.

As discussed in Section 2.1, we will denote by P_t the distribution on S_n obtained by performing the quantum walk on $\Gamma(S_n, C_{(2,1,\ldots,1)})$ for time t starting at the identity then measuring. The above analysis gives a lower bound for the total variation distance of P_t from the uniform distribution:

$$\|P_t - \text{uniform}\| \geq \frac{1}{n} - \frac{2^{2n-2}}{n \cdot n!}$$

for all values of t. This bound follows from considering only the n-cycles, and we believe the true bound to be much larger. Numerical simulations support this claim, but thus far we only have exact expressions for the n-cycles.

Given that we have an exact expression for the probability $P_t[g]$ for any n-cycle g, it is easy to determine the probability associated with any n-cycle in the limiting distribution. By the periodicity of our walks, we have

$$\bar{P}[g] = \frac{1}{2\pi} \int_0^{2\pi} P_t[g] \, dt$$

for each $g \in S_n$, and thus for any $g \in C_{(n)}$ we have $\bar{P}[g] = \frac{\binom{2n-2}{n-1}}{(n!)^2}$. Somewhat surprisingly, this average probability associated with reaching a given n-cycle is not unique to the particular choice of $C_{(2,1,\ldots,1)}$ as a generating set, as shown in the next subsection.

3.3 Other Generating Sets

We have not been able to obtain tractable expressions for the amplitudes associated with quantum walks for other generating sets besides $C_{(2,1,\ldots,1)}$. However, we can prove some facts concerning the limiting distributions for such walks in the case that the generating set consists of all p-cycles for any choice of p. (In case p is odd, we must keep in mind that only the alternating group is being generated.) Again we will focus on the probability of reaching n-cycles starting from the identity.

Consider the quantum walk on $\Gamma(S_n, C_\gamma)$, where γ is any partition. According to Theorem 5, the probability associated with a given conjugacy class C_λ when

starting from a uniform superposition on another class C_μ after time t is given by $|c_\lambda^* U(t) c_\mu|^2$, which may be written as

$$\frac{|C_\lambda||C_\mu|}{(n!)^2} \sum_{\nu,\eta} \exp\left(it\,|C_\gamma| \left(\frac{\chi_\nu(\gamma)}{\dim(\rho_\nu)} - \frac{\chi_\eta(\gamma)}{\dim(\rho_\eta)} \right) \right) \chi_\nu(\lambda)\chi_\nu(\mu)\chi_\eta(\lambda)\chi_\eta(\mu).$$

As before, we let $\bar{P}$ denote the limiting distribution of the walk when starting from the identity. Since our walks are periodic with period 2π, we therefore have

$$\bar{P}[g] = \frac{1}{(n!)^2} \sum_{\nu,\eta}^{*} \chi_\nu(g) \dim(\rho_\nu) \chi_\eta(g) \dim(\rho_\eta).$$

Here the asterisk denotes that the sum is over all partitions ν, η subject to the condition

$$\frac{\chi_\nu(\gamma)}{\dim(\rho_\nu)} = \frac{\chi_\eta(\gamma)}{\dim(\rho_\eta)}. \tag{2}$$

Observe that the choice of generators only affects the average distribution by determining what values other than $\nu = \eta$ are included in the sum. More generally, the average probability associated with obtaining some element in C_λ when starting the walk on the uniform superposition over C_μ is given by

$$\frac{|C_\lambda||C_\mu|}{(n!)^2} \sum_{\nu,\eta}^{*} \chi_\nu(\lambda)\chi_\nu(\mu)\chi_\eta(\lambda)\chi_\eta(\mu).$$

In the case that g is an n-cycle and $\gamma = (p, 1, \ldots, 1)$, the condition of Eq. 2 is relatively easy to characterize for those partitions ν and η for which $\chi_\nu(g)\chi_\eta(g) \neq 0$. Figure 1 summarizes the probability associated with each n-cycle g in the limiting distribution for the quantum walk on $\Gamma(S_n, C_\gamma)$. Due to space constraints, the derivation of these probabilities has been omitted. (See http://arxiv.org/abs/quant-ph/0305182 for a longer version of this paper containing these details.)

We have the following lower bounds on the total variation distance of the limiting distribution from the uniform distribution. As for the case of the quantum walk generated by the transpositions, this bound follows just from considering the n-cycles, and we believe the true distance from uniform to be much larger.

- Let $p \in \{2, \ldots, n\}$ be even, let $\gamma = (p, 1, \ldots, 1) \vdash n$ and let $\bar{P}$ denote the limiting distribution of the quantum walk on $\Gamma(S_n, C_\gamma)$. Then

$$\|\bar{P} - \text{uniform}(S_n)\| \geq \frac{1}{n} - \frac{1}{n \cdot n!} \binom{2n-2}{n-1}.$$

- Let n be odd, let $p \in \{2, \ldots, n\}$ be odd, let $\gamma = (p, 1, \ldots, 1) \vdash n$ and let $\bar{P}$ denote the limiting distribution of the quantum walk on $\Gamma(S_n, C_\gamma)$. Then

$$\|\bar{P} - \text{uniform}(A_n)\| \geq \frac{2}{n} - \frac{2}{n \cdot n!} \binom{2n-2}{n-1} + \frac{1}{n \cdot n!} \binom{n-1}{\frac{n-1}{2}}^2.$$

Parity of n	Parity of p	Range of p	Probability at each n-cycle
even or odd	even	$2 \leq p \leq \lceil \frac{n}{2} \rceil$	$\frac{1}{(n!)^2}\binom{2n-2}{n-1}$
even	even	$\frac{n}{2}+1 \leq p \leq n-1$	$\frac{2}{(n!)^2}\sum_{k=1}^{n-p}\binom{n-1}{k-1}^2$
odd	even	$\frac{n+1}{2}+1 \leq p \leq n-1$	$\frac{2}{(n!)^2}\sum_{k=1}^{n-p}\binom{n-1}{k-1}^2 + \frac{4}{(n!)^2}\binom{n-2}{p-1}^2$
even	even	$p = n$	$\frac{1}{(n!)^2}\binom{2n-2}{n-1}$
even	odd	—	0
odd	odd	$2 \leq p \leq \frac{n+1}{2}$	$\frac{2}{(n!)^2}\binom{2n-2}{n-1} - \frac{1}{(n!)^2}\binom{n-1}{\frac{n-1}{2}}^2$
odd	odd	$\frac{n+1}{2}+1 \leq p \leq n-1$	$\frac{4}{(n!)^2}\sum_{k=1}^{n-p}\binom{n-1}{k-1}^2 + \frac{4}{(n!)^2}\binom{n-2}{p-1}^2$
odd	odd	$p = n$	$\frac{2}{(n!)^2}\binom{2n-2}{n-1} - \frac{1}{(n!)^2}\binom{n-1}{\frac{n-1}{2}}^2$

Fig. 1. Probabilities associated with each n-cycle in the limiting distribution for $\Gamma(S_n, C_{(p,1,\ldots,1)})$.

4 Conclusion

In this paper we have studied some of the properties of continuous-time quantum walks on Cayley graphs of the symmetric group. Many questions concerning these walks remain unanswered. One obvious question that we have not attempted to address in this paper is whether quantum walks on the symmetric group can be applied in the context of quantum algorithms. In terms of specific properties of these walks, we have focused on the limiting distribution—is the limiting distribution bounded away from uniform by a constant? Many other properties of these walks may be of interest as well. For instance, the effect of decoherence on these walks is an interesting topic to consider.

References

1. D. Aharonov, A. Ambainis, J. Kempe, and U. Vazirani. Quantum walks on graphs. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pages 50–59, 2001.
2. D. Aharonov and A. Ta Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, 2003.
3. A. Ahmadi, R. Belk, C. Tamon, and C. Wendler. Mixing in continuous quantum walks on graphs. arXiv.org e-Print quant-ph/0209106, 2002.
4. D. Aldous and J. Fill. Reversible markov chains and random walks on graphs, 2002. See http://stat-www.berkeley.edu/users/aldous/RWG/book.html.
5. A. Ambainis, E. Bach, A. Nayak, A. Vishwanath, and J. Watrous. One-dimensional quantum walks. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pages 60–69, 2001.

6. E. Bach, S. Coppersmith, M. Goldschen, R. Joynt, and J. Watrous. One-dimensional quantum walks with absorbing boundaries. arXiv.org e-Print quant-ph/0207008, 2002.
7. C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
8. A. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. Spielman. Exponential algorithmic speedup by quantum walk. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, 2003.
9. A. Childs, E. Farhi, and S. Gutmann. An example of the difference between quantum and classical random walks. *Quantum Information Processing*, 1(35), 2002.
10. P. Diaconis. *Group Representations in Probability and Statistics*. Lecture Notes-Monograph Series. Institute of Mathematical Statistics, 1988.
11. E. Farhi and S. Gutmann. Quantum computation and decision trees. *Physical Review A*, 58:915–928, 1998.
12. M. Horodecki, P. Horodecki, and R. Horodecki. Mixed-state entanglement and quantum communication. In *Quantum Information: An Introduction to Basic Theoretical Concepts and Experiments*, volume 173 of *Springer Tracts in Modern Physics*. Springer-Verlag, 2001.
13. R. Ingram. Some characters of the symmetric group. In *Proc. Amer. Math. Soc.*, volume 1, pages 358–369, 1950.
14. G. James and A. Kerber. *The Representation Theory of the Symmetric Group*. Addison-Wesley, 1981.
15. J. Kempe. Quantum random walks hit exponentially faster. arXiv.org e-Print quant-ph/0205083, 2002.
16. J. Kempe. Quantum random walks – an introductory overview. *Contemporary Physics*, 2003. To appear. See also arXiv.org e-Print quant-ph/0303081.
17. I. Kerenidis and R. de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, 2003.
18. T. Mackay, S. Bartlett, L. Stephenson, and B. Sanders. Quantum walks in higher dimensions. arXiv.org e-Print quant-ph/0108004, 2001.
19. C. Moore and A. Russell. Quantum walks on the hypercube. In *Proceedings of the Sixth International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, 2002.
20. P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
21. P. Shor and J. Preskill. Simple proof of security of the BB84 quantum key distribution protocol. arXiv.org e-Print quant-ph/0003004, 2000.
22. S. Sternberg. *Group Theory and Physics*. Cambridge University Press, 1994.
23. T. Yamasaki, H. Kobayashi, and H. Imai. Analysis of absorbing times of quantum walks. arXiv.org e-Print quant-ph/0205045, 2002.