On Cayley Graphs on the Symmetric Group Generated by Tranpositions

Joel Friedman*
Department of Mathematics
University of British Columbia
Vancouver, BC V6T 1Z2
CANADA
jf@math.ubc.ca

May 10, 2000

Abstract

Given a connected graph, X, we denote by $\lambda_2 = \lambda_2(X)$ its smallest non-zero Laplacian eigenvalue.

In this paper we show that among all sets of $n-1$ transpositions which generate the symmetric group, S_n, the set whose associated Cayley graph has the highest λ_2 is the set $\{(1,n),(2,n),\ldots,(n-1,n)\}$ (or the same with n and i exchanged for any $i<n$). For this set we have $\lambda_2 = 1$.

This result follows easily from the following result. For any set of transpositions, T, we can form a graph on n vertices, G_P, by forming an edge $\{i,j\}$ in G_P for each transposition $(i,j)\in T$. We prove that if G_P is bipartite, then λ_2 of the Cayley graph associated to T is at most $\lambda_2(G_P)$; left open is the compelling conjecture that the two λ_2's are always equal. We discuss this and other generalizations of Bacher's work, which dealt with the case $T=\{(1,2),(2,3),\ldots,(n-1,n)\}$.

1 Introduction

Let $X=(V,E)$ be a finite graph on $n=|V|$ vertices, possibly with self-loops and multiple edges. The Laplacian associated with X is the $n\times n$ matrix $\Delta = D - A$, where A is the usual adjacency matrix and D is the diagonal matrix whose i,i entry is d_i, the degree of i. We have that Δ is symmetric and positive-semidefinite,

*The author also wishes to thank the NSERC for supporting this research in part.

and has real eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$; $\lambda_2 = \lambda_2(X)$ is non-zero iff X is connected. λ_2 upper and lower bounds certain expansion coefficients in X (see [Tan84, AM85, Dod84, Alo86, Lub94]), and so λ_2 can be viewed as a measure of expansion.

By $\mathrm{Spec}(X)$ we mean the multiset of eigenvalues of $\Delta = \Delta(X)$. We will sometimes denote λ_n by $\lambda_{\max}$. By a *function* on X or V we mean a real-valued function on V; by an eigenfunction we mean one with respect to Δ. If X is a graph, $V(X)$ denotes X's vertex set.

If G is a group and T a subset of G, the *Cayley graph of G with respect to T* is the directed graph $X(G,T)$ with vertex set G and an edge (g, gt) for each $g \in G$ and $t \in T$. If T is closed under taking inverses (i.e. for all T we have $t \in T$ iff $t^{-1} \in T$), then $X(G,T)$ can be viewed as a graph (i.e. an undirected graph).

If X is d-regular, i.e. $d_v = d$ for all v, then A's eigenvalues are precisely d minus Δ's eigenvalues. Thus the study of the spectrum of the adjacency matrix is essentially the same as that of the Laplacian for regular graphs. This holds for Cayley graphs, since they are d-regular for $d = |T|$.

Let S_n be the symmetric group on n elements, which we take to be $\{1, \ldots, n\}$. In this article we study λ_2 of certain Cayley graphs on S_n, especially those generated by a minimal number of transpositions. In [Bac94], Bacher shows that the Cayley graph on S_n with respect to the generators $T_1 = \{(1,2), (2,3), \ldots, (n-1, n)\}$ has $\lambda_2 = 2 - 2\cos(\pi/n)$. In [FOW85, FH98, FH99], it was observed that with respect to the generators $T_2 = \{(1,n), (2,n), \ldots, (n-1,n)\}$ we have $\lambda_2 = 1$. So in a sense the Cayley graph with respect to T_2 is a much better expander than that with respect to T_1. One of the main goals of this paper is to prove that with respect to a set of $n-1$ transpositions, the highest λ_2 which can be achieved is 1.

Theorem 1.1 *Let T be a set of $n-1$ transpositions from S_n. Then the Cayley graph on S_n with respect to T has $\lambda_2 \leq 1$, with equality iff for some i we have $T = \{(i,j) | j \neq i\}$.*

Theorem 1.1 is an easy consequence of the following interesting fact. To any set, T, of transpositions on S_n, we associate a *primitive graph*, G_P, with vertex set $\{1, \ldots, n\}$ and with an edge $\{i, j\}$ for each $(i,j) \in T$.

Theorem 1.2 *Let T be a set of transpositions in S_n whose primitive graph, G_P, is bipartite. Then $\lambda_2(G_P)$ occurs as an eigenvalue in $X(S_n, T)$ with multiplicity at least $n-1$ times its multiplicity in G_P.*

The above proposition is a more or less straightforward consequence of the methods of Bacher in [Bac94], where he proved the above for $T = (1,2), (2,3), \ldots, (n-1, n)$.

The multiplicity statement in theorem 1.2 follows from standard representation theory (see [Dia88]). Namely, the eigenvalues of a Cayley graph $X(G,T)$ (for any G and T) are those of the natural matrix associated with T acting on each irreducible;

each eigenvalue of an irreducible gives rise to f eigenvalues in X, where f is the dimension of the irreducible. One irreducible representation on S_n is the standard representation (see [Ser82]), which can be viewed as the functions on $\{1,\ldots,n\}$ whose sum of values is zero (with S_n acting in the obvious way). It follows that the eigenvalues of $X = X(S_n,T)$ corresponding to this representation are just $\lambda_2,\ldots,\lambda_n$ of G_P. Hence each $\lambda_i(G_P)$, $i \geq 2$, contributes an $n-1$ multiplicity of the same eigenvalue in $X(S_n,T)$. (It also follows, by taking the k-th exterior power of the standard representation, that each sum of k distinct eigenvalues, $\lambda_i(G_P)$, with $i \geq 2$, contributes an $\binom{n-1}{k}$ multiplicity of the same eigenvalue in X.) The same remark holds for any Cayley graph on S_n, with G_P replaced by $X[1]$, where $X[1]$ is defined below in section 2.

Theorem 1.1 follows from theorem 1.2 from the well known fact that $\lambda_2 \leq 1$ for any tree, with equality iff the tree has $n-1$ leaves and one interior vertex of degree $n-1$ (see [Moh91]).

Theorem 1.2 leads to the following fascinating conjecture:

Conjecture 1.3 $\lambda_2(X(S_n,T)) = \lambda_2(G_P)$ *whenever* T *is a set of transpositions with* G_P *bipartite.*

As mentioned before, we know the above conjecture holds for T being T_1 or T_2 as above. It is also valid when G_P is the complete graph (see [DS81]), and a natural generalization of this conclusion is known to hold also when T is any conjugacy class in S_n (see [Roi96]).

The question of the exact multiplicity of λ_2 in $X(S_n,T)$ is interesting. For $T = T_2$ as before, its multiplicity is exactly $(n-1)(n-2)$ (see [FOW85, FH98, FH99]), which is $n-1$ times its multiplicity in G_P. In theorem 1.2 we do not generally know when equality holds. When it doesn't, it means there are some extra representations with the same eigenvalues; it would be interesting to know if this can happen.

The rest of this paper is organized as follows. In section 2 we recall the notion of a Schreier graph, and discuss a natural sequence of Schreier graphs needed here. In section 3 we outline the proof of theorem 1.2. In section 4 we discuss the pseudo-bipartiteness of the sequence of Schreier graphs; this shows a limitation of Bacher's method to the case where G_P is bipartite. We also discuss some facts about eigenvalues in section 4. The rest of the details of the proof of theorem 1.2 are in section 5. In section 6 we close with some remarks on future directions, and prove that the representation theory eigenfunctions are the same as those constructed in sections 3–5. In appendix A we make some technical and bibliographical remarks on the notion of $X[k]$ of section 2.

The author wishes to thank Jean-Pierre Tillich, Yuval Roichman, and Nick Pippenger for helpful discussions.

2 Some Schreier Graphs

Recall that for a subgroup H of a group G with T a subset of G, the *Schreier graph* on $H\backslash G$ with respect to T is the graph $X = X(G, H, T)$ with vertex set being the right cosets $H\backslash G = \{Hg | g \in G\}$ and edges (Hg, Hgt) for each coset Hg and each $t \in T$. If T is closed under taking inverses then X can be viewed as a graph. Note that if 1 is the identity in G, we have $X(G, \{1\}, T) = X(G, T)$, so the notion of a Schreier graph generalizes that of a Cayley graph. If $H \subset K$ are subgroups of G, then the natural map $H\backslash G \to K\backslash G$ gives rise to a map $X(G, H, T) \to X(G, K, T)$ which is a covering map (see [Fri93]); in particular, the eigenvalues of $X(G, K, T)$ are a subset of those of $X(G, H, T)$.

Notice that Schreier graphs are regular, and so studying their Laplacian spectrum is equivalent to studying thier adjacency matrix spectrum.

Let T be a subset of S_n. The inclusion of $\{1, \ldots, n-k\}$ into $\{1, \ldots, n\}$ for $k > 0$ gives rise to an inclusion $S_{n-k} \subset S_n$, giving rise to a Schreier graph $X[k] = X(S_n, S_{n-k}, T)$. The inclusion $S_{n-m-1} \subset S_{n-m}$ for $m = 0, \ldots, n-1$ yields a sequence of covering maps:

$$X[n] \to X[n-1] \to \cdots \to X[1] \to X[0]. \tag{2.1}$$

(Here we take S_0 to be the one element group, so $X[n] = X[n-1] = X(S_n, T)$; also $X[0]$ has one vertex.) Note (see [Bac94]) that $X[k]$ can be viewed as the graph $(V[k], E[k])$, whose vertices are sequences $I = (i_1, \ldots, i_k)$ of length k of distinct integers of $\{1, \ldots, n\}$, with an edge $(I, t(I)) \in E$ for each $t \in T$ and each $I \in V$, where $t(I) = (t(i_1), \ldots, t(i_k))$; actually S_{n-k} embeds naturally into S_n in $\binom{n}{k}$ different ways, by acting on a subset of size $n-k$ of $\{1, \ldots, n\}$, and each embedding gives rise to an isomorphic $X[k]$; in fact, the $I \in V[k]$ simply represent the image with respect to a (an $S_{n-k}\backslash S_n$) coset element of the k elements of $\{1, \ldots, n\}$ which are not acted upon by S_{n-k}. It follows that the mapping $X[m] \to X[m-1]$ on the vertices is just deletion of the first element.

For later use we mention that S_k acts on $V[k]$ in the obvious way, and this action gives rise to one on $V[k]$ functions which commutes with Δ.

3 The Proof of Theorem 1.2

In this section we outline the proof of theorem 1.2 and describe how the technique might generalize further. The details and proofs of certain lemmas are given in sections 4 and 5.

For theorem 1.2 we follow Bacher's approach. Let T be, for the time being, an arbitrary collection of transpositions in S_n and let $X[k]$ be as before. The first lemma is interesting but not hard:

Lemma 3.1 $\lambda_2(X[n]) = 2|T| - \lambda_{\max}(X[n-2])$.

It would be interesting to know what is the appropriate generalization for different T. We discuss possible generalizations and a proof of this lemma in section 4.

We will describe a natural candidate $\lambda_{\max}(X[n-2])$, or more generally $\lambda_{\max}(X[k])$ for arbitrary k. Recall that S_k acts on $V[k]$, and its resulting action on $V[k]$ functions commutes with Δ.

Definition 3.2 *A function u on $V[k]$ is* antisymmetric *if for any $s \in S_k$ and $I \in V[k]$ we have $u(s(I)) = \operatorname{sign}(s)u(I)$. We denote the space of such functions by $\mathcal{A}[k]$.*

Since Δ commutes with the S_k action, Δ leaves $\mathcal{A}[k]$ invariant. The dimension of $\mathcal{A}[k]$ is clearly $\binom{n}{k}$, and so there are $\binom{n}{k}$ independent eigenfuntions in $\mathcal{A}[k]$. The remarkable fact is that we can exhibit all the $\binom{n}{k}$ antisymmetric eigenfunctions and eigenvalues.

Theorem 3.3 *The eigenvalues of Δ in $\mathcal{A}[k]$ are precisely*

$$\left\{ \sum_{i\in I} \lambda_i \;\middle|\; I \subset \{1,\ldots,n\},\ |I| = k \right\},$$

where $\{\lambda_1, \cdots, \lambda_n\} = Spec(G_P)$ and G_P is the primitive graph of T.

We shall outline the proof of this theorem at the end of this section.

So the largest $\mathcal{A}[n-2]$ eigenvalue is $\nu = \lambda_3 + \cdots + \lambda_n$, and taking the trace of Δ shows that $\nu = 2|T| - \lambda_2$. If ν is the largest $X[n-2]$ eigenvalue, then

$$\lambda_2(X[n]) = 2|T| - \lambda_{\max}(X[n-2]) = 2|T| - \nu = \lambda_2(G_P).$$

Conversely if ν is not the largest $X[n-2]$ eigenvalue, then $\lambda_2(X[n]) < \lambda_2(G_P)$.

We remark that Bacher essentially showed that ν is the largest $X[n-2]$ eigenvalue when $X[n-2]$ is "pseudo-bipartite," in the sense below.

Definition 3.4 *We say that a graph, X, is* pseudo-bipartite *if X is bipartite when all self-loops are removed. A* bipartition *of pseudo-bipartite graph, $X = (V, E)$, is a map $\operatorname{sign}\colon V \to \{\pm 1\}$ such that $\operatorname{sign}(u)$, $\operatorname{sign}(v)$ differ whenever $u \neq v$ and $\{u, v\} \in E$.*

Unfortunately $X[n-2]$ is pseudo-bipartite only when T is T_1 as in the introduction. Still, as remarked above, $\nu = \lambda_3 + \cdots + \lambda_n$ is an eigenvalue of $X[n-2]$, and so theorem 1.2 follows. It remains to discuss the proof of theorem 3.3 in more detail.

There is a standard notion of the (Cartesian) product of graphs (see section 5), giving rise to a notion of a power, X^r, for any graph $X = (V, E)$ and integer $r \geq 1$. We define another notion of antisymmetric, for functions on X^r. A $\sigma \in S_r$ acts on $V(X^r) = V^r$ by permuting indices, i.e.

$$\sigma(v_1, \ldots, v_r) = (v_{\sigma(1)}, \ldots, v_{\sigma(r)}).$$

Definition 3.5 *A function u on X^r is* antisymmetric *if $u(\sigma(v)) =$ sign$(\sigma)u(v)$ for all $v \in V(X^r)$ and $\sigma \in S_r$.*

Again, it is easy to find the dimension of the antisymmetric functions; it is $\binom{n}{r}$ where $n = |V(X)|$. And again, Δ of X^r commutes with the S_r action on functions, and hence leaves the antisymmetric functions invariant. It is easy to compute the antisymmetric part of the spectrum:

Lemma 3.6 *The eigenvalues corresponding to antisymmetric eigenvalues of Δ on X^r are precisely:*

$$\left\{ \sum_{i \in I} \lambda_i \;\middle|\; I \subset \{1, \ldots, n\},\ |I| = r \right\}$$

where $\{\lambda_1, \cdots, \lambda_n\} =$ Spec(X).

Proof See section 5.

The crux of Bacher's method is the following remarkable observation, appropriately generalized to our setting:

Theorem 3.7 *The antisymmetric spectrum of $X[k]$ equals that of G_P^k.*

Proof There is an obvious map of vertex sets, $B\colon V[k] \to V(G_P^k)$, namely inclusion. While this map is not a morphism of graphs (i.e. does not preserve edge incidences), remarkably enough B^*, the pullback of G_P^k functions to $X[k]$ functions, does preserve Δ when restricted to antisymmetric functions (and pulls back antisymmetric functions to antisymmetric functions). See section 5 for the details.

$\square$

Theorem 3.7 and lemma 3.6 yield theorem 3.3. This completes our proof of theorem 1.2 (modulo the details in sections 4 and 5).

4 Eigenvalues and Quotients

In this section we prove lemma 3.1. We prove it from two observations. The first one is clear:

Proposition 4.1 *If X is d-regular, bipartite graph, and if* sign *is a bipartition of X, then the map $f \mapsto f$* sign *maps eigenfunctions to eigenfunctions, with $\lambda \mapsto 2d - \lambda$ the corresponding map on eigenvalues.*

Since $X[n]$ is $|T|$-regular and bipartite, we have $\lambda_{N-1} = 2|T| - \lambda_2$ where $N = n!$. This holds whenever T is a collection of sign -1 elements of S_n. The following proposition requires T to consist entirely of elements of order two:

Proposition 4.2 *Let S be a group of order N, and T a set of generators with $t^2 = 1$ for all $t \in T$. Then*

$$\lambda_{N-1}(X(S,T)) = \max_{t \in T} \lambda_{\max}(X(S,\{1,t\},T)).$$

Proof We begin with two remarks. First, if f is an $X(S,T)$ eigenfunction and $s \in S$, then the function g given by $g(u) = f(su)$ is also an eigenfunction of the same eigenvalue; more generally, Δ commutes with the "left multiplication" action of S on functions. Secondly, if H is a subgroup of S and f is an $X(S,T)$ eigenfunction which is constant on all $H\backslash S$ cosets, then f can be viewed as a function $\tilde{f}$ on $X(S,H,T)$, and $\tilde{f}$ is an eigenfunction with the same eigenvalue as f; more generally, $\pi \circ \Delta_{X(S,T)} = \Delta_{X(S,H,T)} \circ \pi$, where π is the natural map from functions on S constant on right H cosets to functions on $H\backslash S$.

Let f be an eigenfunction corresponding to $\lambda_{N-1} = \lambda_{N-1}(X(S,T))$. Let v be such that $f(v) \neq 0$. Then $(\lambda_{N-1} - |T|)f(v)$ is $-\sum_{t \in T} f(vt)$, and $\lambda_{N-1} = 2|T| - \lambda_2 < 2|T|$, so we have that $f(v) + f(vt) \neq 0$ for some t. Then the function g_t given by $g_t(u) = f(vu) + f(vtu)$ gives another eigenfunction with the same eigenvalue as that of f. Also g_t does not vanish identically, and is invariant on right $\{1,t\}$ cosets. Hence g_t is an $X(S,\{1,t\},T)$ eigenfunction, and thus $\lambda_{N-1} \le \lambda_{\max}(X(S,\{1,t\},T))$, and so

$$\lambda_{N-1} \le \max_{t \in T} \lambda_{\max}(X(S,\{1,t\},T)). \tag{4.1}$$

We now prove the reverse inequality to equation 4.1. We begin by claiming that for any t we have $\lambda_{\max}(X(S,\{1,t\},T)) < 2|T|$; indeed, $X(S,\{1,t\},T)$ is connected (since $X(S,T)$ is), is $|T|$-regular, and has a self-loop. So if f is an eigenfunction with eigenvalue $2|T|$, then it is easy to see that the set where $|f|$ takes its maximum is: (1) nonempty, (2) closed under taking neighbors, and (3) cannot include vertices where Δ has a self-loop. This contradiction shows that $\lambda_{\max}(X(S,\{1,t\},T)) < 2|T|$.

Now notice that $\mathrm{Spec}(X(S,\{1,t\},T)) \subset \mathrm{Spec}(X(S,T))$, and $\lambda_N(X(S,T)) = 2|T|$. So $\lambda_{\max}(X(S,\{1,t\},T)) < 2|T|$ implies that

$$\lambda_{\max}(X(S,\{1,t\},T)) \le \lambda_{N-1}.$$

Since this holds for all t we conclude, using equation 4.1, the proposition.

$\square$

We finish the proof of lemma 3.1 by noticing that when T is a subset of transpositions of S_n, $X[n-2]$ is isomorphic to $X(S_n,\{1,t\},T)$ for any $t \in T$.

$\square$

5 More Details of the Proof

We begin by proving lemma 3.6, starting by describing the Cartesian product of two graphs $X = (V, E)$ and $Y = (W, F)$. It is the graph $X \times Y = (V \times W, \mathcal{E})$, where $\mathcal{E}$ consist of two types of edges: for each $\{v_1, v_2\} \in E$ and each $w \in W$ we form an edge $\{v_1 \times w, v_2 \times w\}$, and similarly for each F edge and V vertex. We have $\Delta(X \times Y) = \Delta(X) \otimes I_W + I_V \otimes \Delta(Y)$, where I_V, I_W are identity matrices. It follows that if $\{f_i\}, \{g_j\}$ are bases of X, Y eigenfunctions with eigenvalues $\{\lambda_i\}, \{\nu_j\}$, then $\{f_i \otimes g_j\}$ is a $X \times Y$ basis of eigenfunctions with eigenvalues $\{\lambda_i + \nu_j\}$ (here $f_i \otimes g_j(v, w) = f_i(v) g_j(w)$).

The Cartesian product gives a notion of a power of a graph, X^r. If $\mathcal{F}(Y)$ denotes the functions on a graph, Y, then $\mathcal{F}(X^r)$ is isomorphic to $(\mathcal{F}(X))^{\otimes r}$. Under the S_r action described in section 3, we have that the space of antisymmetric functions on X is just $\bigwedge^r(\mathcal{F}(X))$ viewed as a subspace of $(\mathcal{F}(X))^{\otimes r}$. It follows that if $\{f_i\}$ is a basis for $\mathcal{F}(X)$, then

$$\{f_{i_1} \wedge \cdots \wedge f_{i_r} \mid 1 \le i_1 < \cdots < i_r \le n\}$$

is a basis for the antisymmetric functions, where

$$f_{i_1} \wedge \cdots \wedge f_{i_r} = \frac{1}{r!} \sum_{\sigma \in S_r} \operatorname{sign}(\sigma) f_{i_{\sigma(1)}} \otimes \cdots \otimes f_{i_{\sigma(r)}}.$$

If the $\{f_i\}$ are also eigenfunction of eigenvalues $\{\lambda_i\}$, then clearly each $f_{i_1} \wedge \cdots \wedge f_{i_r}$ is an eigenfunction of eigenvalue $\lambda_{i_1} + \cdots + \lambda_{i_r}$. This therefore produces a basis of eigenfunctions, and so proves lemma 3.6.

□

We finish with the heart of Bacher's method, which is the proof of theorem 3.7. So let $G_P = (V, E)$, note $V(G_P^k) = V^k$, note that $V[k] \subset V^k$ (since $V = \{1, \ldots, n\}$), and let $B\colon V[k] \to V^k$ be the inclusion. Consider an antisymmetric eigenfunction, f, of eigenvalue λ on G_P^k. Fix $I = (i_1, \ldots, i_k) \in V[k]$; we wish to calculate $(\Delta(B^* f))(I)$ and to compare it to $(\Delta(f))(I)$. First consider the $J \in V[k]$ adjacent to I in $X[k]$; we will divide them into two types. Let $\mathcal{S}$ denote the set of $J = (j_1, \ldots, j_k) \in V[k]$ such that for some $a \in \{1, \ldots, k\}$ we have $j_m = i_m$ for $m \ne a$ and $(i_a, j_a) \in T$ and j_a is distinct from $i_1, \ldots, i_k$. Next let $\mathcal{T}$ denote the set of $J = (j_1, \ldots, j_k) \in V[k]$ such that for some $a, b \in \{1, \ldots, k\}$ with $a < b$ we have $(i_a, i_b) \in T$, $j_m = i_m$ for $m \ne a, b$, $j_a = i_b$, and $j_b = i_a$. Clearly $\mathcal{S}$ and $\mathcal{T}$ are disjoint, and clearly their union is precisely the set of J which are $\ne I$ and adjacent to I in $X[k]$. Notice also that for each such J, the edge $\{I, J\}$ in $X[k]$ has multiplicity one. Now for $J \in \mathcal{T}$ we have that J is obtained from I by switching two coordinates, and so $f(J) = -f(I)$. It follows that

$$(\Delta(B^* f))(I) = |\mathcal{S} \cup \mathcal{T}| f(I) - \sum_{J \in \mathcal{S} \cup \mathcal{T}} f(J) = |\mathcal{S} \cup \mathcal{T}| f(I) - \sum_{J \in \mathcal{S}} f(J) - \sum_{J \in \mathcal{T}} f(J)$$

$$= |\mathcal{S} \cup \mathcal{T}| f(I) - \sum_{J \in \mathcal{S}} f(J) \; - |\mathcal{T}|(-f(I)) = \big(|\mathcal{S}| + 2|\mathcal{T}|\big) f(I) - \sum_{J \in \mathcal{S}} f(J).$$

Now consider the $J \in V^k$ adjacent to I in G_P^k. Clearly if $J \in \mathcal{S}$, then J is adjacent to I in G_P^k. However, $J \in \mathcal{T}$ are not adjacent to I in G_P^k; instead, for every $a < b$ such that $(i_a, i_b) \in T$, we have *two* G_P^k edges to I; namely the J such that $j_a = j_b = i_a$ and $j_m = i_m$ for all $m \neq a, b$, and J' with $j'_a = j'_b = i_b$ and $j'_m = i_m$ for all $m \neq a, b$. Notice that such J, J' have the same a-th and b-th coordinates, and so $f(J) = f(J') = 0$. So if the set of such J, J' is denoted $\mathcal{T}'$, we have $|\mathcal{T}'| = 2|\mathcal{T}|$. Clearly $\mathcal{S}, \mathcal{T}'$ give all J which are $\neq I$ and adjacent to I in G_P^k, and for each such J the edge $\{I, J\}$ has multiplicity one in G_P^k, and so

$$(\Delta(f))(I) = |\mathcal{S} \cup \mathcal{T}'| f(I) - \sum_{J \in \mathcal{S} \cup \mathcal{T}'} f(J) = \Big(|\mathcal{S}| + 2|\mathcal{T}|\Big) f(I) - \sum_{J \in \mathcal{S}} f(J).$$

It follows that $(\Delta(B^* f))(I) = (\Delta(f))(I)$, and so $(\Delta(B^* f))(I) = \lambda f(I)$ for all $I \in V[k]$. Hence $B^* f$ is a $X[k]$ eigenfunction with eigenvalue λ. Since all antisymmetric functions on G_P^k vanish on the vertices $V^k - V[k]$, we have that B^* is an injection on the space of antisymmetric V^k functions. Hence the basis of $\binom{n}{k}$ antisymmetric eigenfunctions for V^k maps to a family of linearly independent $V[k]$ eigenfunctions under B^*. Furthermore, B^* clearly maps antisymmetric functions to antisymmetric functions. Since the dimension of antisymmetric function on $V[k]$ is also $\binom{n}{k}$, the antisymmetric eigenvalues of V^k are precisely those of $V[k]$.

6 Concluding Remarks

As for future work, especially regarding the question of Lubotzky, perhaps the following approach is the most interesting.

An *involution*, σ, is an element of S_n such that $\sigma^2 = 1$. Such a σ is a product of disjoint transpositions, and we further call σ an *odd involution* if it has an odd number of transpositions. If T consists entirely of odd involutions of S_n then from the propositions at the beginning of section 4 we have:

$$\lambda_2(X(S_n, T)) = 2|T| - \max_{t \in T} \lambda_{\max}(X(S, \{1, t\}, T)).$$

While the class of odd involution generated Cayley graphs on S_n is a very restricted class, and may not yield the best expanders among all possible Cayley graphs on S_n (of a fixed degree), it seems that there could be enough "randomness" among such graphs to get a pretty good expander (compared with the best possible). For example, if $n \equiv 2 \pmod 4$, then any perfect matching of $\{1, \ldots, n\}$ gives an odd involution, and a "random" collection of a small number of such involutions might behave almost as "randomly" (i.e. give expansion as good) as a random Cayley graph of the same degree. Perhaps the above formula could be of use.

Notice that in the above we stick to involutions. It would be nice to have a formula for $\lambda_{N-1}(X(S, T))$ in terms of $\lambda_{\max}$ of certain quotients, for (more) general groups S

and generators T. Notice that the expression

$$\max_{t \in T} \lambda_{\max}(X(S, \langle t \rangle, T))$$

is not generally equal to $\lambda_{N-1}(X(S,T))$ (here $\langle t \rangle$ is the subgroup generated by t). (Take S to be a cyclic group of order > 2, and $T = \{t, t^{-1}\}$ where t is a generator of S.)

Finally we mention the interesting question of the multiplicity of λ_2 in general Cayley graphs $X(S_n, T)$. If the multiplicity of λ_2 in $X(S_n, T)$ is greater than $n-1$ times that in $X[1]$, it would mean there are some other representations whose corresponding eigenfunctions achieve an eigenvalue of λ_2. It would be interesting to know what these representations and eigenfunctions are. Perhaps such T represent borderline cases between those where $\lambda_2(X(S_n, T)) < \lambda_2(X[1])$ and those where they are equal.

As for eigenfunctions, it is interesting to compare the representation theory eigenfunctions with those constructed in Bacher's approach. We finish this section by showing that the two sets of eigenfunctions constructed span the same space.

The representation theory eigenfunctions arise from how the standard representation lies in the regular representation. To understand this, let $G_P = (V, E)$ as before, so that $V = \{1, \dots, n\}$, let $\mathcal{C}(A)$ for a set A denote the real-valued functions on A, and let $\mathcal{S}$ be those elements of $\mathcal{C}(V)$ with sum over all vertices equal zero. Then S_n acts on V in the obvious way, which gives the S_n action on $\mathcal{S}$, $\sigma(f)(v) = f(\sigma(v))$ for $v \in V$, $f \in \mathcal{S}$, and $\sigma \in S_n$. Clearly $\mathcal{S}$ is the standard representation.

Recall that $V[n]$ is the set of n-tuples $(v_1, \dots, v_n)$ with distinct v_i in $\{1, \dots, n\}$, and that $V[n]$ is identified with S_n with a $\sigma \in S_n$ corresponding to $(\sigma(1), \dots, \sigma(n))$. Projection onto the j-th coordinate gives a map $\pi_j \colon V[n] \to V$; this gives rise to an embedding π_j^* of $\mathcal{S}$ into $\mathcal{C}(V[n]) \simeq \mathcal{C}(S_n)$, and hence to a copy of $\mathcal{S}$ in $\mathcal{C}(S[n])$ which we denote $\mathcal{S}_j = \pi_j^*(\mathcal{S})$. We claim that $\mathcal{S}_1, \dots, \mathcal{S}_{n-1}$ are independent subspaces of $\mathcal{C}(S_n)$. This follows immediately from the lemma below:

Lemma 6.1 *Let $f_i \in \mathcal{C}(V)$ for $i = 1, \dots, n-1$ satisfy $\pi_1^* f_1 + \cdots + \pi_{n-1}^* f_{n-1} = 0$. Then each f_i is constant on V.*

Proof For any distinct $v_1, \dots, v_n$ in V we have

$$f_1(v_1) + f_2(v_2) + \cdots + f_{n-1}(v_{n-1}) = 0 = f_1(v_n) + f_2(v_2) + \cdots + f_{n-1}(v_{n-1}),$$

and so $f_1(v_1) = f_1(v_n)$. It follows that for any distinct $v_1, v_n \in V$ (we can find $v_2, \dots$ so that the v_i are distinct and hence) we have $f_1(v_1) = f_1(v_n)$. Hence f_1 is constant on V. Similarly so are the other f_i.

□

The above gives us $n-1$ independent embeddings of $\mathcal{S}$ into $\mathcal{C}(S_n)$ (and we know this is the maximum possible). So if $f_1, \dots, f_n$ are a basis of eigenfunctions for G_P corresponding to eigenvalues $0 = \lambda_1 \le \cdots \le \lambda_n$, then $f_{ij} = \pi_j^* f_i$, i.e.

$$f_{ij}(v_1, \dots, v_n) = f_i(v_j),$$

gives for $i \geq 2$ and $j = 1, \ldots, n-1$ independent $X(S_n, T)$ eigenfunctions of eigenvalue λ_i. These are "the" representation theory eigenfunctions; of course, the particular functions depend on our $n-1$ embeddings of $\mathcal{S}$ in $\mathcal{C}(S_n)$; yet their span does not depend on the embeddings.

Of course, $\pi_1^* + \cdots + \pi_n^*$ is zero on $\mathcal{S}$, so we cannot add $j = n$ to the above collection of functions. Similarly f_{ij} for $i \geq 2$ and j ranging over any $n-1$ values of $\{1, \ldots, n\}$ gives independent eigenfunctions with the same span as above.

Next we show that the eigenfunctions implicitly described in sections 3–5 give a different but equivalent set of eigenfunctions. So for each $t \in T$, $t = (j, k)$, let $\pi_t \colon V[n] \to V[n-2]$ be the map dropping the j-th and k-th coordinates. For each $i \in \{1, \ldots, n\}$ let

$$F_i = f_2 \wedge f_3 \wedge \cdots \wedge f_{i-1} \wedge f_{i+1} \wedge \cdots \wedge f_n.$$

Then $\tilde{f}_{it} = \operatorname{sign} \cdot \pi_t^* F_i$ is an eigenfunction of eigenvalue λ_i.

Theorem 6.2 *The $\tilde{f}_{it}$, for $i \geq 2$ and $t \in T$ span the same space as the f_{ij}, $i \geq 2$ and $j \in \{1, \ldots, n\}$.*

Proof This theorem follows quite easily from the theorem below. For any set A there is a natural inner product on $\mathcal{C}(A)$, $(f, g) = \sum_{a \in A} f(a) g(a)$.

Theorem 6.3 *Assume that $f_1, \ldots, f_n$ are orthogonal. Then for each $i \geq 2$ and $t = (j, k) \in T$ there is a $c \neq 0$ such that $\tilde{f}_{it} = c(f_{ik} - f_{jk})$.*

Proof Without loss of generality we may assume $t = (1, 2)$. Clearly $\tilde{f}_{it}(v_1, \ldots, v_n)$ is invariant upon exchanging any two of $v_3, \ldots, v_n$ (since the sign and the $\pi_t^* F_i$ each contribute a -1 to such an exchange). Hence $\tilde{f}_{it}(v_1, \ldots, v_n) = \tilde{f}_{it}(v_1, v_2)$ is a function of v_1 and v_2. Also $\tilde{f}_{it}(v_2, v_1) = -\tilde{f}_{it}(v_1, v_2)$, since $\pi_t^* F_i$ is invariant upon exchanging the first two variables, and sign gives a -1 upon this exchange. It follows that $\tilde{f}_{it} \in \bigwedge^2(\mathcal{C}(V))$, and so

$$\tilde{f}_{it}(v_1, v_2) = \sum_{a<b} c_{ab} f_a \wedge f_b(v_1, v_2)$$

for some constants c_{ab}. A standard (and easy) calculation show that since the f_a are orthogonal, so are the $f_a \wedge f_b$ (in the $\mathcal{C}(V[2])$ inner product). So $c_{ab} = 0$ iff $(\tilde{f}_{it}, f_a \wedge f_b) = 0$; we now check the latter. Now

$$(\tilde{f}_{it}, f_a \wedge f_b)_{V[2]} = (n-2)!(\operatorname{sign} \cdot \pi_t^* F_i, f_a \wedge f_b \otimes 1 \otimes \cdots \otimes 1)_{V[n]}$$

$$= (n-2)!(\operatorname{sign}, f_a \wedge f_b \otimes F_i)_{V[n]},$$

and the right-hand-side clearly vanishes unless $a = 1$ and $b = i$. Hence $\tilde{f}_{it} = c_{1i} f_1 \wedge f_i(v_1, v_2) = c(f_i(v_2) - f_i(v_1))$ since f_1 is constant. Thus $\tilde{f}_{it} = c(f_{ik} - f_{ij})$, and clearly $c \neq 0$ since $\tilde{f}_{it}$ is not identically zero (since F_i isn't).

□

We remark that the above theorem needs the orthogonality condition. For f_{ij}, f_{ik} depend only on f_i, while $\tilde{f}_{it}$ depends on all f_m except $m = i, 1$. So if f_i corresponds to an eigenvalue of multiplicity $s > 1$, holding f_m fixed for $m \neq i, 1$ and varying f_i keeps $\tilde{f}_{it}$ fixed while lets $f_1 \wedge f_i$ vary over an s-dimensional space minus some 1-dimensional subspaces.

Now we finish the proof of theorem 6.2. First assume that the f_i are orthogonal. We are reduced to proving the following: let W be a vector space with $w_1, \ldots, w_n \in W$ satisfying $w_1 + \cdots + w_n = 0$ but are otherwise independent (i.e. any $n-1$ of them are independent). Then the span of the w_i equals the span of $w_k - w_j$ over all $(j,k) \in T$. Clearly the span of the latter lies in that of the former. So let $H = (V, F)$ be a spanning tree of G_P, let $r \in V$ be an arbitrary vertex (thought of as the root of H), and orient the edges of F "away from r" (so H is now a directed tree). It suffices to show that $\widetilde{w}_t = w_k - w_j$ over $t = (j,k) \in F$ are linearly independent.

First note that $\sum d_i w_i = 0$ and $\sum d_i = 0$ implies that all the d_i are zero. Now assume that $\sum_{t \in F} c_t \widetilde{w}_t = 0$ with $c_t \neq 0$ for some t. Then $\sum c_t \widetilde{w}_t = \sum d_i w_i$ with

$$d_i = \sum_{t=(j,i) \in F} c_t - \sum_{t=(i,j) \in F} c_t.$$

So clearly $\sum d_i = 0$. But if $t = (i,j)$ is an edge of maximal distance from r with $c_t \neq 0$, the $d_j = c_t \neq 0$, a contradiction

Finally, we consider general $f_2, \ldots, f_n$ (i.e. not necessarily orthogonal). We have that the span of $\{f_i\}_{i \geq 2}$ is the orthogonal complement of f_1, regardless of the choice of $\{f_i\}_{i \geq 2}$. Hence the span of the $f_{ij} = \pi_j^* f_i$ is independent of the choice of $\{f_i\}_{i \geq 2}$. Similarly the span of F_i (with $i \geq 2$) is the orthogonal complement of

$$\{f_1 \wedge w \mid w \in \textstyle\bigwedge^{n-3}(\mathcal{F}(V))\}.$$

Hence the span of $\tilde{f}_{it} = \pi_t^* F_i$ is independent of the choice of $\{f_i\}_{i \geq 2}$.

$\square$

A Remarks on $X[k]$

After one paragraph of minor technical remarks on $X(G, H, T)$ we spend the rest of this section making bibliographical remarks about $X[k]$.

Notice that there is some ambiguity in how to make $X(G, H, T)$ into an undirected graph (in the presence of multiple edges or self-loops). But these ambiguities don't affect Δ (or the adjacency matrix), and are easily resolved by insisting that T come with an involution σ such that $t(\sigma(t)) = 1$ for all t. Also, we can weaken the condition that T be closed under taking inverses to T coming with an involution σ such that $t(\sigma(t))$ lies in the core[1] of H in G for all t.

[1] The core of H in G is $\bigcap_{g \in G} gHg^{-1}$, i.e. the largest normal subgroup of G lying in H.

We now briefly remark on where λ_2 of the graphs $X[k]$ has been studied elsewhere. First, in the contexts of random graphs, if we take d random elements of S_n, take their inverses, we get a set, T, of $2d$ elements of S_n closed under taking inverses. The resulting $X[1]$ is the usual model of a random regular graph studied in [BS87, FKS89, Fri91]; in [Fri91] it is shown that $\lambda_2 = 2d - 2\sqrt{2d} + O(\log d)$ with high probability as $n \to \infty$; unresolved is the conjecture of Alon (see [Alo86]) that $\lambda_2 \geq 2d - 2\sqrt{2d-1} - \epsilon$ for any $\epsilon > 0$ with high probability, even though numerical evidence indicates that $\lambda_2 \geq 2d - 2\sqrt{2d-1}$ probably holds in most cases (see [Fri93]). The study of λ_2 of $X[k]$ becomes important to the "quick k-transitivity" of the graph $X[1]$ (see [FJR+96]), where a question in cryptography can be resolved using the fact that λ_2 of $X[k]$ is $\geq 2d - 2d(1+\epsilon)d^{-1/(k+1)}(2d-1)^{1/(2k+2)}$ with high probability (for any fixed $\epsilon > 0$, with k, d fixed and $n \to \infty$).

Second, Lubotzky (see [Lub95]) asks whether there are Cayley graphs on S_n which are expanders, in the sense that there should exist a d and $\epsilon > 0$ and sequence X_n of d-regular Cayley graphs on S_n with $\lambda_2(X_n) \geq \epsilon$. In other words, is there a fixed d and $\epsilon > 0$ such that for any n there is a subset $T \subset S_n$ of size d with $\lambda_2(X(S_n, T)) \geq \epsilon$. In the language of the previous paragraph, $X(S_n, T) = X[n]$, and so the previous question is one about $\lambda_2(X[n])$ for the optimal T of a certain size. Equation 2.1 gives rise to an inclusion of spectra:

$$\mathrm{Spec}(X[0]) \subset \mathrm{Spec}(X[1]) \subset \cdots \subset \mathrm{Spec}(X[n]). \tag{A.1}$$

From the previous paragraph we know that for any k there are d and $\epsilon > 0$ such that for any n we have $\lambda_2(X[k]) \geq \epsilon$ for some T of size d. So even if the answer to Lubotzky's question is negative, we might ask for which functions $f(n)$ is there a d and $\epsilon > 0$ such that for all n we have the same holding with $k = f(n)$; in other words, at which point in equation A.1 do small Laplacian eigenvalues enter in? We can also ask a similar question where d varies with n; this has been studied for T being a conjugacy class by Roichman in [Roi97].

Thirdly, in the context of the main theorems in this paper, we are saying that in equation A.1, between Spec($X[1]$) and Spec($X[n]$), all new eigenvalues arising are at least as great as $\lambda_2(X[1])$, in a certain special case. Indeed, if T consists entirely of transpositions, then if we take $X[1]$ and delete all its self-loops (of which there are very very many), we get T's primitive graph, G_P. So $\mathrm{Spec}(G_P) = \mathrm{Spec}(X[1])$; and the relationship between theorem 1.2 and the above is clear. It would be quite remarkable if such relations between Spec($X[1]$) and Spec($X[n]$) held for a large class of T (if it held for most T of fixed size d as $n \to \infty$ then Lubotzky's question would be true); it might be interesting to ask for which T this holds.

References

[Alo86] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[AM85] N. Alon and V.D. Milman. λ_1, isoperimetric inequalities for graphs and superconcentrators. *J. Comb. Theory Ser. B*, 38:73–88, 1985.

[Bac94] Roland Bacher. Valeur propre minimale du laplacien de Coxeter pour le group symétrique. *Journal of Algebra*, 167:460–472, 1994.

[BS87] Andrei Broder and Eli Shamir. On the second eigenvalue of random regular graphs. In *28th Annual Symposium on Foundations of Computer Science*, pages 286–294, 1987.

[Dia88] Persi Diaconis. *Group Representations in Probability and Statistics.* Institute of Mathematical Statistics, Hayward, California, 1988.

[Dod84] Jozef Dodziuk. Difference equations, isoperimetric inequality and transcience of certain random walks. *Trans. AMS*, 284:787–794, 1984.

[DS81] P. Diaconis and M. Shahshahani. Generating a random permutation with random transpositions. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 57:159–179, 1981.

[FH98] Joel Friedman and Phil Hanlon. On the Betti numbers of chessboard complexes. *J. Algebraic Combin.*, 8(2):193–203, 1998.

[FH99] Joel Friedman and Phil Hanlon. Erratum: "On the Betti numbers of chessboard complexes" [J. Algebraic Combin. **8** (1998), no. 2, 193–203; MR 2000c:05155] by J. Friedman and P. Hanlon. *J. Algebraic Combin.*, 9(1):95+11, 1999.

[FJR+96] J. Friedman, A. Joux, Y. Roichman, J. Stern, and J-P. Tillich. The action of a few random permutations on r-tuples and an application to cryptography. In *STACS*, pages 375–386, 1996.

[FKS89] J. Friedman, J. Kahn, and E. Szemerédi. On the second eigenvalue of random regular graphs. In *21st Annual ACM Symposium on Theory of Computing*, pages 587–598, 1989.

[FOW85] L. Flatto, A.M. Odlyzko, and D.B. Wales. Random shuffles and group representations. *Annals of Probability*, 13:154–178, 1985.

[Fri91] J. Friedman. On the second eigenvalue and random walks in random d-regular graphs. *Combinatorica*, 11:331–362, 1991.

[Fri93] Joel Friedman. Some geometric aspects of graphs and their eigenfunctions. *Duke J. Math*, 69:487–525, 1993.

[Lub94] Alexander Lubotzky. *Discrete Groups, Expanding Graphs and Invariant Measures*. Birkhauser, 1994.

[Lub95] A. Lubotzky. Cayley graphs: Eigenvalues, expanders and random walks. 1994/95. preprint.

[Moh91] B. Mohar. The Laplacian spectrum of graphs. In O.R. Ollermann Y. Alavi, G. Chartrand and A.K. Schwenk, editors, *Graph theory, combinatorics and applications*, pages 871–898, New York, 1991. Wiley.

[Roi96] Yuval Roichman. Upper bound on the characters of the symmetric groups. *Invent. Math.*, 125(3):451–485, 1996.

[Roi97] Yuval Roichman. Expansion properties of Cayley graphs of the alternating groups. *J. Combin. Theory Ser. A*, 79(2):281–297, 1997.

[Ser82] Jean-Pierre Serre. *Linear representations of finite groups*. Springer-Verlag, 1982.

[Tan84] R.M. Tanner. Explicit concentrators from generalized n-gons. *SIAM J. Alg. Disc. Methods*, 5:287–293, 1984.