CPSC 536F: NOTES ON EIGENVALUES, EXPANSION, AND MIXING TIMES

JOEL FRIEDMAN

Copyright: Copyright Joel Friedman 2022. Not to be copied, used, or revised without explicit written permission from the copyright owner.

Disclaimer: The material may sketchy and/or contain errors, which I will elaborate upon and/or correct in class. For those not in CPSC 536F: use this material at your own risk...

The point of these notes is to state some results regarding the eigenvalues of graphs and Markov chains, as they apply to expansion and mixing times. We many not have time to prove them all.

These notes rely on definitions given in class.

1. Examples of Eigenvalues of Regular Graphs

If A_G is a d-regular graph, then A_G/d is a Markov matrix. This is important in thinking of examples.

We will discuss the eigenvalues of the following graph:

(1) the cycle of length n;
(2) the cartesian product of k cycles (possibly of different lengths), which is a $2k$-regular "k-dimensional grid graph," and only a weak expander for k fixed and number of vertices large;
(3) the Boolean hypercube;
(4) more generally, the cartesian product $G_1 \times G_2$ of any two graphs in terms of the eigenvalues/vectors of G_1 and G_2;
(5) other products;
(6) Cayley graphs of abelian groups and (some remarks) about non-abelian groups.

For fixed d, it is known that if A_G is any d-regular graph on n vertices, then

$$\lambda_2 \geq 2\sqrt{d-1}\bigl(1 - 1/(\log_{d-1} n)^2\bigr);$$

the bound of Alon-Boppana bound gives the same with $\log_{d-1} n$ instead of $(\log_{d-1} n)^2$; the tighter bound was proven independently by Nabil Kahale and me. The Broder-Shamir bound says that most d-regular graphs on n vertices (constructed from $d/2$ permutations on n vertices) has $\rho = \max_{i>1} |\lambda_i|$ bounded by $2d^{3/4} + \epsilon$ for any $\epsilon > 0$, and one can ultimately improve this to $2\sqrt{d-1} + \epsilon$ (this was proven by me, and is a rather long story). The $2\sqrt{d-1}$ is also the L^2 norm of the adjacency operator on the d-regular (infinite) tree, and this is not a coincidence.

The Boolean cube is a good model for various "configuration models," where the second eigenvalue is very close to the first, but far enough separated to have

Research supported in part by an NSERC grant.

algorithmic consequences. For example, if G is a bipartite graph on $2N$ with each vertex connected to only one other vertex, then the set of all matchings (not perfect matching, of which there is only one...), where the adjacency matrix is to delete or add an edge, is simply a walk on the Boolean hypercube of dimension N. Hence $n = 2^N$ is the number of configurations (vertices), this is a $d = N$ regular graph, and the second eigenvalue is of size $d - 2 = N - 2$.

2. Expansion and Mixing via Eigenvalues Regular Graphs

Many theorems in algebraic graph theory are simpler for regular graphs; hence we often make this assumption when needed.

Theorem 2.1. *Let G be a d-regular graph, and A_G its adjacency matrix, and let*

$$d = \lambda_1(G) \geq \cdots \geq \lambda_n(G).$$

Let $\rho = \max_{i>1} |\lambda_i(G)|$. Then for any subsets $A, B \subset V_G$ of vertices, the number of edges from A to B, denoted $e(A, B)$, satisfies

$$\bigl|e(A,B) - d\,|A|\,|B|/n\bigr| \leq \rho \sqrt{\frac{|A|(n-|A|)}{n}} \sqrt{\frac{|B|(n-|B|)}{n}}.$$

The proof follows easily from the more conceptual formula where we show that

$$A_G = \frac{d}{n} E + \mathcal{E},$$

where E is the all 1's matrix, and

$$\mathcal{E} = \sum_{i>1} \lambda_i v_i v_i^{\mathrm{T}},$$

where $v_1, \ldots, v_n$ are an orthonormal eigenbasis corresponding to the eigenvalues $\lambda_1, \ldots, \lambda_n$. Hence $\mathcal{E}$ takes the orthogonal complement of the constant vector $\mathbf{1}$ to itself, and has operator norm ρ on this subspace. The projection of the characteristic vector of A onto this subspace has norm

$$\sqrt{\frac{|A|(n-|A|)}{n}},$$

and similarly for B.

Notice that A_G^k is the matrix counting the number of walks on a graph to itself of length k, and the above theorem implies (since A_G^k has the same eigenvectors, with eigenvalues λ_i^k)

$$A_G^k = \frac{d^k}{n} E + \mathcal{E}^k,$$

where the L^2-norm of $\mathcal{E}^k$ is at most ρ^k.

We remark that

$$\sqrt{\frac{|A|(n-|A|)}{n}} \leq \sqrt{|A|},$$

and the simpler right-hand-side gives up a factor of $\sqrt{2}$ or less when $|A| \leq n/2$. This weaker bound is often called the "Expander Mixing Lemma," i.e.,

$$\bigl|e(A,B) - d\,|A|\,|B|/n\bigr| \leq \rho \sqrt{|A|\,|B|}.$$

2.1. **Why a "Mixing Lemma"?** To see why this is called a "mixing lemma," we use the above theorem to get a (generally weak) bound on the mixing time of A_G as a Markov chain. The associated Markov matrix to a d-regular graph is $P = A_G/d$, i.e., the Markov chain that traverses each edge out of a vertex with probability $1/d$. The *mixing time* of an $n \times n$ Markov matrix, P, is usually defined as the smallest value of $t = t_{\mathrm{mix}}$ such that if $e_1, \ldots, e_n$ denotes the standard basis of $\mathbb{R}^n$ (interpreted as a stochasitc vector), then

$$\|P^t e_i - \mathbf{1}/n\|_{L^1} \le 1/2.$$

This implies that[1] that for any $k \in \mathbb{N}$ we have

$$\mathrm{TotalVar}(P^{kt} e_i, \mathbf{1}/n) = (1/2)\|P^{kt} e_i - \mathbf{1}/n\|_{L^1} \le 1/2^{k+1}.$$

The reason why eigenvalue bounds tend to be weak is that eigenvalue bounds are better suited to estimate the L^2-norm.

If $P = A_G/d$ and ρ is as above, we easily see that $e_i - \mathbf{1}/n$ has zero projection to $\not\Bbbk$, and $\mathbf{1}/n$ is fixed by P. Hence for any $t \in \mathbb{N}$,

$$\|P^t(e_i - \mathbf{1}/n)\|_{L^2} = \|P^t e_i - \mathbf{1}/n)\|_{L^2} \le (\rho/d)^t \|e_i - \mathbf{1}/n\|_{L_2} \le (\rho/d)^t \sqrt{(n-1)/n}$$

(since e_i is the characteristic vector of a set, A, of vertices of size 1); this is likely a pretty good bound. The problem is that we have $\|v\|_{L_1} \le \sqrt{n}\|v\|_{L_2}$ (which is optimal when all components of v are of the same absolute value), and so to bound the mixing time we have to use the the bound

$$\|P^t(e_i - \mathbf{1}/n)\|_{L^1} \le \sqrt{n}(\rho/d)^t \sqrt{(n-1)/n} = (\rho/d)^t \sqrt{n-1},$$

which is at most $1/2$ when $t \ge (1/2)\log(4(n-1))/\log(\rho/d)$, which tends to be an overestimate of the true mixing time.

Similarly, for any real $\epsilon > 0$ one defines the ϵ-mixing time to be the smallest $t = t_{\mathrm{mix}}(\epsilon)$ such that for all $i \in [n]$

$$\text{Total Variation}(P^t e_i, \mathbf{1}/n) = (1/2)\|P^t e_i - \mathbf{1}/n\|_{L^1} \le \epsilon,$$

and we can get a similar estimate of $t_{\mathrm{mix}}(\epsilon)$ for any $\epsilon > 0$.

3. Alon-Boppana Bound and Improvements

We will prove that if G is a d-regular graph of diameter k (largest distance between two vertices)

$$\lambda_2(A_G) \ge 2\sqrt{d-1}\,\bigl(1 - f(k)\bigr),$$

where $f(k) \to 0$ as $k \to \infty$; specifically one can prove that $f(k) \le C/k^2$ where C is an absolute constant (independently discovered by Nabil Kahale and me), although

---

[1] Here are some details: for a general Markov matrix, P, of an irreducible Markov chain with stationary distribution π, then set $d(t) = \sup_i \|e_i P^t - \pi\|_{\mathrm{TV}}$, where TV is "total variation, which equals exactly 1/2 the L^1 distance, and set $\overline{d}(t) = \sup_{i,j} \|e_i P^t - e_j P^t\|_{\mathrm{TV}}$. Then a coupling argument shows that $\overline{d}(s+t) \le \overline{d}(s)\overline{d}(t)$; see *Markov Chains and Mixing Times*, by Levin, Peres, and Wilmer, Section 4.4. (Our definition of d and $\overline{d}$ considers only the distributions e_i, e_j rather than arbitrary distributions, but this does not alter the definitions given in the above textbook.) One also sees that $d(t) \le \overline{d}(t) \le 2d(t)$. Formally one defines the ϵ-mixing time, for any real $\epsilon > 0$, denoted $t_{\mathrm{mix}}(\epsilon)$, as the smallest t such that $d(t) \le \epsilon$; one typically defines t_{mix} to be $t_{\mathrm{mix}}(1/4)$. In applications one is typically interested in $t_{\mathrm{mix}}(\epsilon)$ for some ϵ; however, the above inequalities imply that $d(s+t) \le 2\overline{d}(s)\overline{d}(t) \le 2d(s)d(t)$, and hence $d(kt) \le 2^{k-1}d(t)$, and hence $\|e_i P^{kt} - \mathbf{1}/n\|_{L^1} \le 1/2^{k+1}$.

it is easier to prove $f(k) \le C/k$. Since the diameter of G is at least $\log_{d-1} n - c$, this result implies the usual Alon-Boppana bound

$$\lambda_2 \ge 2\sqrt{d-1}\bigl(1 - C/\log_{d-1} n\bigr).$$

We develop a number of important ideas regarding the spectral analysis of graphs.

3.1. **Max-Min Theorem.** We will use the following standard theorem in linear algebra. Recall that if A is a real $n \times n$ matrix, then we define the *Rayleigh quotient* of a vector $v \in \mathbb{R}^n$ at A to be

$$\mathcal{R}_A(v) = \frac{(Av)\cdot v}{v \cdot v}.$$

Assuming that A is symmetric, there is an orthonormal eigenbasis $v_1, \ldots, v_n$ for A, i.e., $Av_i = \lambda_i v_i$ for real λ_i, and we easily check that if $v = c_1 v_1 + \cdots + c_n v_n$, then

$$\begin{aligned}\mathcal{R}_A(v) &= \frac{c_1^2\lambda_1 + \cdots + c_n^2\lambda_n}{c_1^2 + \cdots + c_n^2} \\ &= \lambda_1 \frac{c_1^2}{c_1^2 + \cdots + c_n^2} + \cdots + \lambda_n \frac{c_n^2}{c_1^2 + \cdots + c_n^2},\end{aligned}$$

which is a “convex linear combination” of $\lambda_1, \ldots, \lambda_n$. Hence, for example, if λ_1 is the largest eigenvalue of A, then

$$\mathcal{R}_A(v) \le \lambda_1 \frac{c_1^2}{c_1^2 + \cdots + c_n^2} + \cdots + \lambda_1 \frac{c_n^2}{c_1^2 + \cdots + c_n^2} = \lambda_1,$$

and similarly $\mathcal{R}_A(v) \ge \lambda_n$ if λ_n is the smallest eigenvalue of A.

Theorem 3.1. *Let A be a real $n \times n$ symmetric matrix, whose eigenvalues are be ordered*

$$\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_n.$$

Let $W \subset \mathbb{R}^n$ be a subspace of dimension r. Then there exists a nonzero $w \in W$ such that

$$\mathcal{R}_A(w) \le \lambda_r.$$

Proof. Let $v_1, \ldots, v_n$ be an orthonormal eigenbasis for A, i.e., with $Av_i = \lambda v_i$. We claim that there is a nonzero $w \in W$ that is orthogonal to $v_1, \ldots, v_{r-1}$: indeed, if $w_1, \ldots, w_r$ is a basis for W then $w = \alpha_1 w_1 + \cdots + \alpha_r w_r$ is orthogonal to some v_i iff

$$\alpha_1(w_1 \cdot v_i) + \cdots + \alpha_r(w_r \cdot v_i) = 0.$$

Hence if we impose this condition for $i = 1, \ldots, r-1$, we get $r-1$ (homogeneous) linear equations for the r variables α_i, which therefore has a nontrivial solution, and for any such nontrivial solution $w = \alpha_1 w_1 + \cdots + \alpha_r w_r$ is a nonzero vector in W orthogonal to $v_1, \ldots, v_{i-1}$. It follows that such a w can be written as

$$w = c_1 v_1 + \cdots + c_n v_n,$$

where $c_1 = \ldots = c_r = 0$. Since $c_1 = \ldots = c_{r-1} = 0$ in the equation for w above, we have

$$\begin{aligned}\mathcal{R}_A(w) &= \lambda_r \frac{c_r^2}{c_1^2 + \cdots + c_n^2} + \cdots + \lambda_n \frac{c_n^2}{c_1^2 + \cdots + c_n^2} \\ &\le \lambda_r \frac{c_r^2}{c_1^2 + \cdots + c_n^2} + \cdots + \lambda_r \frac{c_n^2}{c_1^2 + \cdots + c_n^2} = \lambda_r.\end{aligned}$$

$\square$

The above theorem is called the *max-min theorem* because it can be written

$$\min_{w\in W\setminus\{0\}} \mathcal{R}_A(w) \le \lambda_r$$

for any W of dimension r, and we have equality when $W = \mathrm{Span}(v_1, \dots, v_r)$, since if $w = c_1v_1 + \cdots + c_rv_r$, then

$$\mathcal{R}_A(w) = \lambda_1 \frac{c_1^2}{c_1^2 + \cdots + c_r^2} + \cdots + \lambda_r \frac{c_r^2}{c_1^2 + \cdots + c_r^2}$$

$$\ge \lambda_r \frac{c_1^2}{c_1^2 + \cdots + c_r^2} + \cdots + \lambda_r \frac{c_r^2}{c_1^2 + \cdots + c_r^2} = \lambda_r.$$

Hence we have

$$\max_{\dim(W)=r} \quad \min_{w\in W\setminus\{0\}} \mathcal{R}_A(w) = \lambda_r.$$

Similarly, there is a min-max principle by applying the max-min principle to $-A$, i.e.,

$$\min_{\dim(W)=r} \quad \max_{w\in W\setminus\{0\}} \mathcal{R}_A(w) = \lambda_{n-r}.$$

One often uses the following corollary of Theorem 3.1.

Corollary 3.2. *Let A be a real $n \times n$ symmetric matrix, whose eigenvalues are be ordered*

$$\lambda_1 \ge \lambda_2 \ge \cdots \ge \lambda_n.$$

Let $u_1, \dots, u_r$ be non-zero vectors such that for all $i \ne j$ we have $u_i \cdot u_j = 0$ and $(Au_i) \cdot u_j = 0$. Then

$$\lambda_r \ge \min\bigl(\mathcal{R}_A(u_1), \dots, \mathcal{R}_A(u_r)\bigr).$$

Proof. We may scale each u_i so that $u_i \cdot u_i = 1$. Let W be the span of $u_1, \dots, u_r$. Then for any $w \in W$ with $w \ne 0$ we have $w = c_1u_1 + \cdots + c_ru_r$ with real $c_1, \dots, c_r$ that are not all zero; hence

$$(Aw) \cdot w = \sum_{i,j} c_ic_j(Au_i) \cdot u_j = \sum_i c_i^2(Au_i) \cdot u_i = \sum_i \mathcal{R}_A(u_i)$$

(since $u_i \cdot u_i = 1$), and

$$w \cdot w = \sum_i c_i^2 u_i \cdot u_i = \sum_i c_i^2.$$

Hence

$$\mathcal{R}_A(w) = \sum_i \frac{c_i^2}{c_1^2 + \cdots + c_r^2} \mathcal{R}_A(u_i) \ge \min\bigl(\mathcal{R}_A(u_1), \dots, \mathcal{R}_A(u_r)\bigr).$$

$\square$

3.2. **Covering and Etale Maps of Digraphs.**

Definition 3.3. Let $G = (V_G, E_G^{\mathrm{dir}}, h_G, t_G)$ and $H = (V_H, E_H^{\mathrm{dir}}, h_H, t_H)$ be directed graphs. By a *morphism* (often called a *graph homomorphism* in the literature) *from* H *to* G we mean a pair $\pi = (\pi_V, \pi_E)$ with $\pi_V\colon V_H \to V_G$, and $\pi_E\colon E_H^{\mathrm{dir}} \to E_G^{\mathrm{dir}}$ that respect the heads and tails maps in the sense that $\pi_E h_H = h_G \pi_E$ and $\pi_E t_H = t_G \pi_E$; in this case we write $\pi\colon H \to G$. We say that π is a *covering map* (respectively, an *étale map*) if for each $v \in V_H$, both (1) the map π_E restricted to

$$t_H^{-1}(v) \to t_G^{-1}(\pi_V v)$$

(so $t_H^{-1}(v) = \{e \in E_H^{\mathrm{dir}} \mid t_H e = v\}$) is a bijection (respectively, an injection), and similarly (2) so is the restriction of π_E as a map $h_H^{-1}(v) \to h_G^{-1}(\pi_V v)$.

In the literature, an étale map is sometimes alternatively called a *closed immersion*.

Example 3.4. If G' is a subgraph of G, then the inclusion $G' \to G$ is étale.

Example 3.5. If $m, n \in \mathbb{N}$, then ther is a covering map from the cycle of length n to that of length m iff n is divisible by m.

The following proposition is easy to prove and gives some intuition for covering maps. (See examples given in class.)

Theorem 3.6. *Let G be strongly connected directed graph, and $\pi\colon H \to G$ a covering map. Then if for some $v \in V_G$ we have that $\pi_V^{-1}(v)$ is of size $k \in \mathbb{N}$, then for any $v \in V_G$, $\pi_V^{-1}(v)$ is of size k, and for any $e \in E_G^{\mathrm{dir}}$, $\pi_E^{-1}(e)$ is of size k. In this case we say that π is* a k-to-1 (covering) map.

Homework: the composition of two covering maps is a covering map; the composition of two étale maps is an étale map.

Homework: if $\pi\colon H \to G$ is an étale map, then there is a graph H' such that H is a subgraph of H', there is a covering map $\pi'\colon H' \to G$, such that the inclusion of H to H' followed by the covering map $H \to G$ equals the map $\pi\colon H \to G$.

3.3. **The Comparison Lemma for Digraphs.** Recall that a *walk of length k* in a digraph, G, is an alternating sequences of vertices and directed edges

$$(v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k)$$

such that for all $i = 1, \ldots, k$, $t_G e_i = v_{i-1}$ and $h_G e_i = v_i$; such a walk *is from v_0 to v_k* or *begins in v_0 and ends in v_k*. For any $v, v' \in V_G$ we let $\mathrm{walks}_{\le k}(v, v')$ denote the number of walks from v to v' of length at most k.

We recall that the Perron-Frobenius theorem implies that if G is a strongly connected digraph (i.e., for any $v, v' \in V_G$, there is a walk from v to v'), then its adjacency matrix, A_G has a positive, real, eigenvalue λ_1 that is as large in absolute value as any other eigenvalue; we also call λ_1 the *Perron-Frobenius eigenvalue* of G and denote it by λ_{PF}; furthermore, for any $v, v' \in V_G$,

$$\lim_{k\to\infty} \bigl(\mathrm{walks}_{\le k}(v, v')\bigr)^{1/k} = \lambda_{\mathrm{PF}}(G).$$

To prove this theorem it is helpful to know the *path (or walk) lifting lemma.* Consider any graph morphism $\pi\colon \to H \to G$ and a walk in H

$$w = (v_0', e_1', \ldots, e_k', v_k').$$

Then we define $\pi(w)$ to be the sequence

$$\pi(w) = (\pi_V v_0', \pi_E e_1', \dots, \pi_E e_k', \pi_V v_k'),$$

which we easily verify is a walk in G.

Theorem 3.7. *Let $\pi\colon H \to G$ be a covering map (respectively, an étale map). Let*

$$w = (v_0, e_1, v_1, \dots, v_{k-1}, e_k, v_k)$$

be a walk in G, and let $v_0' \in V_H$ with $\pi(v_0') = v_0$. Then there exists a unique (respectively at most one) walk

$$w' = (v_0', e_1', v_1', \dots, v_{k-1}', e_k', v_k')$$

such that $\pi(w') = w$.

Homework: Prove this theorem, using induction on k.

Theorem 3.8. *Let H, G be graphs such that there is an étale morphism $H \to G$. Then $\lambda_{\mathrm{PF}}(H) \le \lambda_{\mathrm{PF}}(G)$. Furthermore equality holds when H is a covering graph.*

Proof. Fix any two vertices $v, v' \in V_H$. For every walk of length k from v to v'

$$w = (v_0, e_1, v_1, \dots, v_{k-1}, e^k, v_k),$$

we easily check that there is a corresponding walk

$$\pi(w) = \big(\pi_V(v_0), \pi_E(e_1), \dots, \pi_E(e_{k-1}), \pi_V(v_k)\big).$$

By the path lifting lemma, the map π from walks of length at most k from v to v' is an injection of walks from $\pi(v)$ to $\pi(v')$. Hence

$$\mathrm{walks}^H_{\le k}(v, v') \le \mathrm{walks}^G_{\le k}\big(\pi(v), \pi(v')\big).$$

Now we take the $1/k$-th power of both sides and take limits. □

3.4. **Covering and Etale Maps of Gaphs.** Most all of the above notions for digraphs such as covering maps, étale maps, the path lifting lemma, etc., also hold for graphs in the following sense: namely, if $G = (V_G, E_G^{\mathrm{dir}}, h_G, t_G, \iota_G)$ and $H = (V_H, E_H^{\mathrm{dir}}, h_H, t_H, \iota_H)$ are graphs, then a *morphism of graphs* (or *graph homomorphism* $\pi\colon H \to G$ is any morphism $\pi = (\pi_V, \pi_E)$ from the underlying directed graph of H, $(V_H, E_H^{\mathrm{dir}}, h_H, t_H)$, to that of G, $(V_G, E_G^{\mathrm{dir}}, h_G, t_G)$ which, in addition, respects the edge pairings ι_H and ι_H in the sense that $\pi\iota_H = \iota_G\pi$. We then say that π is a covering map or étale map if the underlying map of directed graphs is. A *walk* in a graph, G, is just a walk in the underlying directed graph.

Hence any theorem above such as the path lifting lemma or the Perron-Frobenius eigenvalue comparison theorem, therefore holds for graphs just as it holds for digraphs.

3.5. **The Universal Cover of a Graph.** If we fix a graph, G, there is a "largest" connected graph H such that $\pi\colon H \to G$ is a covering map, in the sense that for any covering map $\pi'\colon H' \to G$ and any vertices $v \in V_H$ and $v' \in V_{H'}$ with $\pi(v) = \pi'(v')$, there is a unique covering map $\eta\colon H \to H'$ such that $\eta(v) = v'$. In this sense, $\pi\colon H \to G$ is a *universal cover*.

In fact, H is a tree (therefore, generally infinite). If G is d-regular, then this tree is the (unique, infinite) d-regular tree, T_d.

Example 3.9. The universal cover of a cycle of length n is the "infinite path," whose vertex set is the integers, $\mathbb{Z}$, and where $x, y \in \mathbb{Z}$ are adjacent (with one edge from x to y) iff $y = x \pm 1$.

Here is one way to build the universal cover $\pi\colon T \to G$ of a graph, G. By a *non-backtracking walk* in G, we mean a walk

$$w = (v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k)$$

such that for all $i = 1, \ldots, k-1$, $\iota e_{i+1} \neq e_i$. We let the vertices of T be the set of non-backtracking walks; we declare the walk of length k above to be adjacent to the walk of length $k-1$

$$(v_0, e_1, v_1, \ldots, v_{k-2}, e_{k-1}, vk-1),$$

provided that $k \geq 1$, and—for any k—to any non-backtracking walk that extends w by one vertex:

$$(v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k, e', v')$$

with $\iota e' \neq e_k$. We define the map $V_T \to V_G$ to be the map taking the non-backtracking walk

$$w = (v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k)$$

to the vertex v_k. We extend this to a covering map in the natural way; for example we define the map $E_T^{\rm dir} \to E_G^{\rm dir}$ by mapping the edge from

$$(v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k)$$

to the walk

$$(v_0, e_1, v_1, \ldots, v_{k-1}, e_k, v_k, e_{k+1}, v_{k+1})$$

to be the edge e_{k+1}.

3.6. The Alon-Boppana Theorem and Stronger Results.

Lemma 3.10. *Let $T_{d,\ell}$ be the d-regular infinite tree, where we pick a vertex, v, in this tree and take the induced subgraph on all vertices of distance ℓ to v; we call v the* root *of $T_{d,\ell}$. Then for fixed d there exist $C_1, C_2 > 0$ such that for all ℓ,*

$$2\sqrt{d-1}\bigl(1 - C_1/\ell^2\bigr) \leq \lambda_{\rm PF}(T_{d,\ell}) \leq 2\sqrt{d-1}\bigl(1 - C_1/\ell^2\bigr).$$

To prove this lemma we make a detailed calculation of the Perron-Frobenius eigenvalue by building the appropriate eigenfunction. Namely, we look for a function $u\colon T_d \to \mathbb{R}$ that is everywhere positive, and such that $u(w)$ depends only on the distance from w to the root of $T_{d,\ell}$. If we content ourselves with the weaker bound, namely

$$2\sqrt{d-1}\bigl(1 - C_1/\ell\bigr) \leq \lambda_{\rm PF}(T_{d,\ell})$$

as in the original Alon-Boppana theorem, we take $u(w)$ to be $(d-1)^{-k/2}$, where k is the distance from w to the root. In this case we easily see that

$$\mathcal{R}(u) = 2\sqrt{d-1}\bigl(1 - O(1)/\ell\bigr).$$

It follows that

$$\mathcal{R}(u) \leq \lambda_{\rm PF}(T_{d,\ell}).$$

Department of Computer Science, University of British Columbia, Vancouver, BC V6T 1Z4, CANADA.

E-mail address: `jf@cs.ubc.ca`

URL: `http://www.cs.ubc.ca/~jf`