The Method of Interlacing Polynomials

Kuikui Liu

June 2017

Abstract

We survey the method of interlacing polynomials, the heart of the recent solution to the Kadison-Singer problem as well as breakthrough results in spectral graph theory. We will focus on four specific problems that build on each other, and discuss their connections with each other. Algorithms for constructing promised solutions are also discussed. Finally, we will mention some applications of these breakthroughs to spectral graph theory, and approximation algorithms, and discuss current open questions and directions for future work.

Contents

1 Introduction and Organization

The method of interlacing polynomials is a new technique for proving the existence of (and constructing) objects whose naturally associated polynomial has bounded roots; for matrices, these are the characteristic polynomials. Polynomials with interlacing roots provided some of the intuition behind improved results on problems such as graph sparsification [27]; they were also a key component in the recent positive resolution to the Kadison-Singer problem [35], as well as breakthroughs in constructing Ramanujan graphs [34], [36], [42] which are very important objects in spectral graph theory with numerous applications.

The goal of this survey is to:

1. discuss the method of interlacing polynomials, intuition why it is so powerful, and how it was applied to the Kadison-Singer problem and other related problems
2. discuss other root-bounding methods that work in conjunction with interlacing polynomials to produce improved guarantees
3. give some applications of the problems that were solved using these methods

We aim for the majority of this survey to be self-contained, beyond assumed knowledge of linear algebra and calculus. We will attempt to skip as few steps as possible in proofs for the first three applications we'll study.

2 Preliminaries

2.1 Some Notation

1. We will let $\mathbb{R}[x_1,\ldots,x_n]$ be the set of polynomials in the variables $x_1,\ldots,x_n$ with real coefficients. Similarly, we will let $\mathbb{C}[z_1,\ldots,z_n]$ be the set of polynomials in the variables $z_1,\ldots,z_n$ with complex coefficients.
2. For a matrix $A \in \mathbb{R}^{n\times m}$, one can view it as a linear operator from $\mathbb{R}^m$ to $\mathbb{R}^n$. For this, we write $\|A\|_2$ to be the **operator norm** of A defined by $\sup_{x\in\mathbb{R}^m} \|Ax\|_2$. We also write $\|A\|_F$ for its **Frobenius norm**, defined by $\sqrt{\sum_{i=1}^n \sum_{j=1}^m A_{ij}^2}$.
3. For a matrix $A \in \mathbb{C}^{n\times n}$, A^* will denote its conjugate transpose (i.e. its adjoint when viewed as an operator on $\mathbb{C}^n$). When A has all real entries, we'll use $A^\top$ for its transpose (again, its adjoint when viewed as an operator on $\mathbb{R}^n$).
4. For a matrix $A \in \mathbb{C}^{n\times n}$, we'll use both the notation $\chi_A(x)$ and $\chi[A](x)$ for its **characteristic polynomial** $\det(xI - A)$. Oftentimes, we'll use the first notation when the expression for A is small and easy to write down; we'll usually reserve the second for more complicated expressions for A.
5. For a Hermitian matrix $A \in \mathbb{C}^{n\times n}$ (or symmetric matrix $A \in \mathbb{R}^{n\times n}$), we will write $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ for its largest and smallest eigenvalues, respectively. For a real-rooted polynomial, we will use $\mathrm{maxroot}(p)$ to denote its maximum root.
6. For a vector $v \in \mathbb{C}^n$, we let $\mathrm{diag}(v) \in \mathbb{C}^{n\times n}$ is a diagonal matrix D with $D_{ii} = v_i$ for $1 \le i \le n$.
7. We will use $\mathcal{O}$ for "Big-O" notation, and $O(n)$ for the orthogonal group of $n \times n$ matrices.
8. We will use $\mathcal{S}_n$ for the set of permutations on n letters. We will also abuse notation and have this also be the set of $n \times n$ permutation matrices.
9. For a matrix $A \in \mathbb{C}^{n\times m}$, $A_{i,j}$ and $A(i,j)$ will be used interchangeably, both to denote the entry with row index i and column index j.

2.2 Linear Algebra

2.2.1 Eigenvalues and Eigenvectors

Theorem 1 (Spectral Theorem)**.** *If $A \in \mathbb{C}^{n\times n}$ is Hermitian, then A has all real eigenvalues $\lambda_1,\dots,\lambda_n$ with a corresponding orthonormal basis $u_1,\dots,u_n$ of $\mathbb{C}^n$ such that $A = \sum_{i=1}^n \lambda_i u_i u_i^*$.*

Theorem 2 (Courant-Fischer)**.** *For a Hermitian matrix $A \in \mathbb{C}^{n\times n}$ with eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$, we have the following identities.*

$$\lambda_k = \min_{x_1 \perp \cdots \perp x_{n-k+1}} \min_{x\neq 0, \|x\|_2=1} \{x^* A x : x \in \operatorname{span}\{x_1,\dots,x_{n-k+1}\}\} \tag{1}$$

$$= \max_{x_1 \perp \perp, x_k} \min_{x\neq 0, \|x\|_2=1} \{x^* A x : x \in \operatorname{span}\{x_1,\dots,x_k\}\} \tag{2}$$

Lemma 1. *If $A \in \mathbb{C}^{n\times n}$ is Hermitian, with eigendecomposition $\sum_{i=1}^n \lambda_i u_i u_i^*$, then $A^k = \sum_{i=1}^n \lambda_i^k u_i u_i^*$ for all $k \in \mathbb{N}$.*

Definition 1 (Moore-Penrose Pseudoinverse)**.** If $A \in \mathbb{C}^{n\times n}$ is Hermitian, we can diagonalize it as $A = \sum_{i=1}^k \lambda_i u_i u_i^*$, where $\lambda_1,\dots,\lambda_k$ are the nonzero real eigenvalues of A, and $u_1,\dots,u_k$ is a set of orthonormal eigenvectors. Define the **Moore-Penrose Pseudoinverse of** A by

$$A^\dagger = \sum_{i=1}^k \frac{1}{\lambda_i} u_i u_i^*$$

Then $\ker(A) = \ker(A^\dagger)$ clearly and

$$AA^\dagger = A^\dagger A = \sum_{i=1}^k u_i u_i^*$$

is an orthonormal projection onto $\operatorname{im}(A) = \ker(A)^\perp = \operatorname{span}\{u_1,\dots,u_k\}$ (and is the identity on this subspace). Note that if $k = n$, then $A^\dagger = A^{-1}$.

Definition 2 (Mutually Diagonalizable)**.** If $\mathcal{H} \subset \mathbb{C}^{n\times n}$ is a set of Hermitian matrices, then they are **mutually diagonalizable** if there exists an invertible square matrix $V \in \mathbb{C}^{n\times n}$ such that for every $A \in \mathcal{H}$, there exists a real diagonal matrix $D \in \mathbb{R}^{n\times n}$ with $A = VDV^*$.

Lemma 2. *If $A, B \in \mathbb{C}^{n\times n}$ Hermitian are mutually diagonalizable, then $AB = BA$, i.e. they commute.*

Lemma 3. *Let $A \in \mathbb{C}^{n\times n}$ be a Hermitian matrix and let $b \in \mathbb{R}$ be such that $bI - A$ is invertible (i.e. b is not an eigenvalue of A). Let $u_1,\dots,u_n$ be the orthonormal eigenvectors of A. Then, $\{(bI-A)^k : k \in \mathbb{Z}\} \cup \{\sum_{i\in S} u_i u_i^* : S \subset [n]\}$ is mutually diagonalizable.*

2.2.2 Trace and Determinant Properties

Lemma 4 (Trace Properties)**.** *1. If $A, B \in \mathbb{C}^{n\times n}$ and $a, b \in \mathbb{C}$, then $\operatorname{tr}(aA + bB) = a\operatorname{tr}(A) + b\operatorname{tr}(B)$ (linearity)*

2. If $A \in \mathbb{C}^{n\times k}$ and $B \in \mathbb{C}^{k\times n}$, then

$$\operatorname{tr}(AB) = \operatorname{tr}(BA) \tag{3}$$

and

$$A \cdot B = \sum_{i=1}^n \sum_{j=1}^k A_{ij} B_{ji} = \operatorname{tr}(AB)$$

3. If $A \in \mathbb{C}^{n\times n}$ and $\lambda_1,\dots,\lambda_n$ are its eigenvalues, then

$$\operatorname{tr}(A) = \sum_{i=1}^n \lambda_i$$

Corollary 1. *If $A \in \mathbb{C}^{d\times d}$ is positive semidefinite, then* $\operatorname{tr}(A) \geq 0$.

Corollary 2. *Let $v_1, \ldots, v_m \in \mathbb{C}^n$ satisfy $\sum_{i=1}^m v_i v_i^* = I$. Then* $\operatorname{tr}(M) = \sum_{i=1}^m v^* M v$.

Corollary 3. *Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix and P, Q be the orthogonal projections on* $\operatorname{im}(A), \ker(A)$, *respectively. Let $B \in \mathbb{R}^{n\times n}$ be arbitrary. Then,*

$$\operatorname{tr}\left[B^\top (bI - A)^k B\right] = \operatorname{tr}\left[B^\top P (bI - A)^k P B\right] + \operatorname{tr}\left[B^\top Q (bI - A)^k Q B\right]$$

for every $k \in \mathbb{Z}$, since $A, P, Q, (bI - A)^k$ are mutually diagonalizable.

Lemma 5 (Determinant Properties)**.** *For $A, B \in \mathbb{C}^{n\times n}$, and eigenvalues $\lambda_1, \ldots, \lambda_n$ of A*

1. $\det(AB) = \det(A)\det(B)$
2. $\det(A) = \det(A^*)$
3. $\det(A) = \prod_{i=1}^n \lambda_i$

2.2.3 Other Useful Formulas

Theorem 3 (Jacobi Formula)**.** *If $f : \mathbb{C} \to \mathbb{C}^{d\times d}$, where $f(t)$ is invertible in some open interval I, then*

$$\partial_t \det(f(t)) = \det(f(t)) \cdot \operatorname{tr}\left(f(t)^{-1} \cdot \partial_t f(t)\right)$$

in I.

Corollary 4. *Let $A, B \in \mathbb{C}^{d\times d}$ where A is invertible. Then $\partial_t \det(A + tB) = \det(A + tB) \cdot \operatorname{tr}((A + tB)^{-1} B)$. In particular, $[\partial_t \det(A + tB)]_{t=0} = \det(A)\operatorname{tr}(A^{-1}B)$*

Lemma 6 (Matrix Determinant Lemma)**.** *Let $A \in \mathbb{C}^{n\times n}$ be invertible, and $u, v \in \mathbb{C}^n$. Then*

$$\det(A + uv^*) = (1 + v^* A^{-1} u)\det(A)$$

Lemma 7 (Sherman-Morrison Formula)**.** *Let $A \in \mathbb{R}^{n\times n}$ be invertible, and $u, v \in \mathbb{R}^n$. Then $A + uv^*$ is invertible if and only if $1 + v^* A^{-1} u \neq 0$. Furthermore, if $A + uv^*$ is invertible, then it has inverse*

$$(A + uv^*)^{-1} = A^{-1} - \frac{A^{-1} u v^* A^{-1}}{1 + v^* A^{-1} u}$$

2.3 A Very Quick Primer on Spectral Graph Theory

We will only work with undirected graphs in this survey.

Definition 3 (Adjacency Matrix)**.** Let $G = (V = [n], E)$ be a graph with $|E| = m$. The **adjacency matrix** A_G of G is the symmetric $n \times n$ matrix with entries $A_G(i, j) = 1$ if and only if $(i, j) \in E$. If G is a multigraph, then $A_G(i, j)$ is the number of edges between i and j. If G is a weighted graph with weights $w : E \to \mathbb{R}$, then $A_G(i, j) = w(i, j)$.

Definition 4 (Laplacian Matrix)**.** Let $G = (V = [n], E)$ be a graph with $|E| = m$. Let d_i be the degree of $i \in V$, defined by $d_i = |\{(i, j) : j \in V\}|$ (or, if G is weighted with $w : E \to \mathbb{R}$, $d_i = \sum_{j\in V} w(i, j)$). Let $D_G = \operatorname{diag}(\{d_i\}_{i=1}^n)$ be the **degree matrix** of G. We define the **Laplacian matrix** L_G by $L_G = D_G - A_G$.

One can alternatively construct L_G as follows: First, orient the edges of G arbitrarily (as in for every edge $e = (u, v) \in E$, arbitrarily set u or v to be the "head" of e and the other vertex to be the "tail" of e). Then, define the **signed edge-vertex incidence matrix of** G $B \in \mathbb{R}^{m\times n}$ as the matrix with entries

$$B(e, v) = \begin{cases} 1 & \text{if } v \text{ is the head of } e \\ -1 & \text{if } v \text{ is the tail of } e \\ 0 & \text{otherwise} \end{cases}$$

Then, set $L_G = B^\top W B$, where $W \in \mathbb{R}^{m\times m}$ has $W(e, e) = w(e)$ for all $e \in E$.

Remark 1. If w assigns nonnegative weight to every edge, associated with the Laplacian is a nice quadratic form, which can easily be derived using the definition (identity) $L_G = B^\top WB$: for a weighted graph $G = (V, E, w)$, for all $x \in \mathbb{R}^n$,

$$x^T L_G x = x^\top B^\top WBx = \left\| W^{1/2} Bx \right\|_2^2 = \sum_{(i,j)\in E} w(i,j)(x_i - x_j)^2 \geq 0$$

In particular, L_G is positive semidefinite when the edge weights are nonnegative.

The idea behind spectral graph theory is to analyze the properties of graphs through linear algebraic techniques. For example, through the eigenvalues of A_G and L_G, we can learn about connectivity, bipartiteness, etc. Here are some example results.

Theorem 4. *Let $\lambda_1 \leq \cdots \leq \lambda_n$ be the eigenvalues of L_G. Then, $\lambda_k = 0$ if and only if G has at least k connected components.*

Theorem 5. *Let $\lambda_1 \leq \cdots \leq \lambda_n$ be the eigenvalues of A_G. Then G is bipartite if and only if $\{\lambda_1, \ldots, \lambda_n\}$ is symmetric about 0.*

3 Small-Scale Interlacing

We begin with interlacing in its most basic form. We will motivate why this is a nice property for a set of real-rooted polynomials to have.

Definition 5 (Interlacing Polynomials)**.** A degree-$(d-1)$ polynomial $q \in \mathbb{R}[x]$ with all real roots $\alpha_1 \le \cdots \le \alpha_{d-1}$ **interlaces** a degree-d polynomial $p \in \mathbb{R}[x]$ with all real roots $\beta_1 \le \cdots \le \beta_d$ if $\beta_1 \le \alpha_1 \le \beta_2 \le \cdots \le \alpha_{d-1} \le \beta_d$; if $q \in \mathbb{R}[x]$ is a degree-d polynomial with all real roots $\alpha_1 \le \cdots \le \alpha_d$, then $q(x)$ interlaces $p(x)$ if $\alpha_1 \le \beta_1 \le \cdots \le \alpha_d \le \beta_d$.

We say a finite set $\{p_i\}_1^n \subset \mathbb{R}[x]$ of degree-d real-rooted polynomials have a **common interlacing** if there is a degree-$(d-1)$ (or, equivalently, degree-d) real-rooted polynomial $q \in \mathbb{R}[x]$ that interlaces each p_i.

Remark 2. It's not hard to see that if $p(x) = \prod_{i=1}^n (x - \lambda_i)$ and $q(x) = \prod_{i=1}^n (x - \mu_i)$ have a common interlacing, where $\lambda_1 \le \cdots \le \lambda_n$ and $\mu_1 \le \cdots \le \mu_n$, then it follows that for any nonempty $S \subset [n]$, $\prod_{i \in S}(x - \lambda_i)$ and $\prod_{i \in S}(x - \mu_i)$ also have a common interlacing.

We now state a crucial theorem that is the motivation for using interlacing.

Theorem 6. *Let $p_1, \ldots, p_n \in \mathbb{R}[z]$ be a set of real-rooted degree-d polynomials with positive leading coefficient and let $p_\emptyset = \sum_{i=1}^n p_i$. If $\{p_i\}_1^n$ have a common interlacing, then*

1. *$p_\emptyset$ is real-rooted*
2. *for every $1 \le k \le d$, there exists $1 \le i \le n$ such that $\lambda_k(p_i) \le \lambda_k(p_\emptyset)$*

where $\lambda_k(p)$ denotes the kth smallest root of p.

Example 1. Interlacing is a particularly special because in general, adding (or averaging) real-rooted polynomials does not produce anything useful; the sum (or average) of real-rooted polynomials might not even be real-rooted, as demonstrated by the $p_1(x) = (x-1)^2$ and $p_2(x) = (x+1)^2$, for which $(p_1 + p_2)(x) = 2x^2 + 2$.

However, if we had, say, $p_1(x) = x(x-2)$ and $p_2(x) = (x-1)(x-4)$, then p_1, p_2 have a common interlacing, and $(p_1 + p_2)(x) = 2x^2 - 6x + 4 = 2(x-2)(x-1)$ and $\mathrm{maxroot}(p_2) \le \mathrm{maxroot}(p_1 + p_2)$.

An interesting observation is that if $\{p_i\}_1^n$ has a common interlacing, then so does $\{\mu_i p_i\}_1^n$, for any choice of constants $\mu_1, \ldots, \mu_n \in \mathbb{C}$. Hence, if we form a probability distribution $\mathcal{D}$ over $\{p_i\}_1^n$, assigning probability $\mu_i \ge 0$ to p_i, then this theorem states there is a $\mu_i p_i$ (and hence, a p_i) whose maximum root is upper bounded by $\sum_{i=1}^n \mu_i p_i = \mathbb{E}_{p \sim \mathcal{D}}[p]$. Thus, to prove that there is a p_i whose maximum root is upper bounded by some quantity α, it suffices to bound the maximum root of $\mathbb{E}_{p \sim \mathcal{D}}[p]$ by α, which can be much easier, as we will see in the applications below.

Proof Sketch of Theorem 6. For real-rootedness, observe that the roots of $p_\emptyset$ are the same as those of $\mathbb{E}_{i \sim \mathrm{unif}([n])}[p_i] = p_\emptyset / n$. Since $p_1, \ldots, p_n$ have a common interlacing, by Theorem 7 (see below), $p_\emptyset / n$ is real-rooted. Hence, so is $p_\emptyset$. This proves the first guarantee.

For the second guarantee, by Remark 2, it suffices to show that there exists $1 \le i \le n$ such that $\mathrm{maxroot}(p_i) \le \mathrm{maxroot}(p_\emptyset)$, and then go by induction on the degree d. Let g be a degree-$(d-1)$ common interlacing for $p_1, \ldots, p_n$.

Since the p_i have positive leading coefficient, there exists x_i such that $p_i(x) > 0$ for all $x > x_i$, for all $1 \le i \le n$; this also shows that $p_\emptyset(x) > 0$ for all sufficiently large x. Furthermore, since every p_i has exactly one root that is at least $\lambda_{d-1}(g)$, namely $\lambda_d(p_i)$, we have $p_i(\lambda_{d-1}(g)) \le 0$ for all $1 \le i \le n$. Thus, $p_\emptyset(\lambda_{d-1}(g)) \le 0$. Since $p_\emptyset(x) > 0$ for all sufficiently large x, the largest root $\lambda_d(p_\emptyset)$ of $p_\emptyset$ is at least $\lambda_{d-1}(g)$.

Now, since $p_\emptyset(\lambda_d(p_\emptyset)) \ge 0$, there exists $1 \le i \le n$ such that $p_i(\lambda_d(p_\emptyset)) \ge 0$. Since p_i has at most one root that is at least $\lambda_{d-1}(g)$, $\lambda_d(p_i) \in [\lambda_{d-1}(g), \lambda_d(p_\emptyset)]$. In particular, $\lambda_d(p_i) \le \lambda_d(p_\emptyset)$. □

So, how does one show that a set of polynomials has a common interlacing? The following result is a particularly useful characterization for when a common interlacing exists. It reduces the problem of showing the existence of common interlacings into the problem of showing certain polynomials are real-rooted. Throughout, whenever we need to prove the existence of common interlacings, we will apply this lemma. This greatly simplifies our work, because then, we can apply techniques such as those from the theory of real stable polynomials. We will discuss these techniques in greater depth later. We will not need this for our first two applications.

Theorem 7 (Theorem 2.1 of [9]). *Let $p_1, \dots, p_n \in \mathbb{R}[z]$ be degree-d polynomials with positive leading coefficient. Then $\{p_i\}_1^n$ have a common interlacing if and only if $\sum_{i=1}^n \lambda_i p_i$ is real-rooted for all λ_i satisfying $\sum_{i=1}^n \lambda_i = 1$, that is, all convex combinations of $\{p_i\}_1^n$ are real-rooted.*

We now turn to an important example of when we know interlacing occurs. This will be particularly useful for our first two applications.

Theorem 8 (Cauchy Interlacing Theorem). *Let $A \in \mathbb{R}^{n\times n}$ be symmetric and $v \in \mathbb{R}^n$. Then $\chi_A(x)$ interlaces $\chi_{A+vv^\top}(x)$.*

Proof. By the Matrix Determinant Lemma (Lemma 6),

$$\det(xI - A - vv^\top) = (1 - v^\top (xI - A)^{-1} v)\det(xI - A)$$

The claim is trivial if $v = 0$ so assume otherwise. Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of A, and $u_1, \dots, u_n$ be orthonormal eigenvectors corresponding to $\lambda_1, \dots, \lambda_n$. Observe that $A = \sum_{i=1}^n \lambda_i u_i u_i^\top$, and by orthonormality of $u_1, \dots, u_n$, $I = \sum_{i=1}^n u_i u_i^\top$. Thus, $xI - A = \sum_{i=1}^n (x - \lambda_i) u_i u_i^\top$ so we may write

$$(xI - A)^{-1} = \sum_{i=1}^n \frac{1}{x - \lambda_i} u_i u_i^\top$$

Hence,

$$1 - v^\top (xI - A)^{-1} v = 1 - \sum_{i=1}^n \frac{1}{x - \lambda_i} v^\top u_i u_i^\top v = 1 - \sum_{i=1}^n \frac{|\langle u_i, v\rangle|^2}{x - \lambda_i}$$

Putting these together, we have

$$\chi_{A+vv^\top}(x) = \chi_A(x) \cdot \left(1 - \sum_{i=1}^n \frac{|\langle u_i, v\rangle|^2}{x - \lambda_i}\right)$$

Let μ be an eigenvalue of $A + vv^\top$. Assume without loss of generality that the set of indices i for which $\langle u_i, v\rangle \neq 0$ is precisely $[k]$, for some k (note, we can guarantee that $k \geq 1$ since if $\langle u_i, v\rangle = 0$ for all $1 \leq i \leq n$, by linear independence of $u_1, \dots, u_n$, $v = 0$). We have two cases:

1. $\chi_A(\mu) = 0$, i.e. μ is also an eigenvalue of A. This occurs when μ is an eigenvalue of A corresponding to an eigenvector that is orthogonal to v.

2. $f(\mu) \stackrel{\text{def}}{=} \sum_{i=1}^k |\langle u_i, v\rangle|^2/(\mu - \lambda_i) = 1$. This is a real rational function that blows up when μ is approaches λ_i, for any $1 \leq i \leq k$. More concretely, for every $1 \leq i \leq k$,

$$\lim_{x\to\lambda_i^-} f(x) = -\infty \qquad \lim_{x\to\lambda_i^+} f(x) = +\infty$$

   where $\lim_{x\to\lambda_i^-}$ is the limit as x approaches λ_i from the left (and similarly, for $\lim_{x\to\lambda_i^+}$). We also have $\lambda_{x\to\infty} f(x) = \lim_{x\to-\infty} f(x) = 0$. Observe that $f(x)$ is continuous and injective on each of the open intervals $(\lambda_i, \lambda_{i+1})$ for all $0 \leq i \leq k$ (where for convenience, we define $\lambda_0 = -\infty$ and $\lambda_{k+1} = \infty$); furthermore, f is a bijection from $(\lambda_i, \lambda_{i+1})$ to $\mathbb{R}$ for $1 \leq i < k$, and a bijection from $(\lambda_k, \lambda_{k+1})$ to $\mathbb{R}_{>0}$. By the Intermediate Value Theorem, for each $1 \leq i \leq k$, there exists precisely one value of μ in $(\lambda_i, \lambda_{i+1})$ for which $f(\mu) = 1$.

Combining these two cases, we have interlacing. □

4 Warm-Up I: The Column Subset Selection Problem

Problem 1 (Column Subset Selection)**.** *For a given matrix $A \in \mathbb{R}^{n\times m}$, is there is a subset of columns $S \subset [m]$ such that the submatrix $B \in \mathbb{R}^{n\times|S|}$ formed by those columns satisfies "certain linear algebraic properties", and if so, can it be found efficiently?*

This is a general problem, where "certain linear algebraic properties" can be anything, but are usually ones useful in computation and analysis, such as condition number, rank, etc., with many applications in numerical linear algebra and functional analysis. In this section, we will be considering "well-invertibility", which means whether or not the matrix $A \in \mathbb{R}^{n\times m}$ has a column submatrix $A_{[n],S} \in \mathbb{R}^{n\times|S|}$, for some $S \subset [m]$, such that

1. $A_{[n],S}$ is injective, so that it is invertible on a subspace of $\mathbb{R}^m$.
2. $A^{-1}_{[n],S}$ has singular values bounded away from 0.

Well-invertible matrices are desirable because they are robust to additive noise, and algorithms that are sensitive to floating point precision remain "numerically stable" when run on these types of matrices. The main question we will answer is: How large can $|S|$ be for there to exist a well-invertible submatrix with $|S|$ columns?

4.1 Bourgain-Tzafriri Restricted Invertibility

We begin with a celebrated result of Bourgain-Tzafriri [6].

Theorem 9 (Theorem 1.2 from [6])**.** *There exist universal constants $0 < c, C < 1$ such that for all linear operators $T : \mathbb{R}^n \to \mathbb{R}^n$ satisfying $\|Te_i\|_2 = 1$, for all $1 \le i \le n$, there exists a subset $S \subset [n]$ with cardinality $|S| \ge \lfloor cn/\|T\|_2^2 \rfloor$ such that*

$$\|Tv\|_2^2 \ge C\|v\|_2^2$$

for all $v \in \operatorname{span}\{e_i : i \in S\}$. That is, for every linear operator $T : \mathbb{R}^n \to \mathbb{R}^n$ satisfying $\|Te_i\|_2 = 1$ for all $1 \le i \le n$, there exists an $S \subset [n]$ of size $\Omega(n/\|T\|_2^2)$ such that the operator norm of T as a linear operator from $\operatorname{span}\{e_i : i \in S\}$ is $\Omega(1)$.

This fact has been known since the 1990s. Here, we will prove a more general and more powerful version. This next result will also make more clear the connection between the stable rank and the existence of well-invertible column submatrices.

Theorem 10 (Theorem 2 from [30])**.** *Let $v_1, \dots, v_m \in \mathbb{R}^n$ satisfies $\sum_{i=1}^m v_i v_i^\top = I$. Let $0 < \epsilon < 1$ be arbitrary and let $T : \mathbb{R}^n \to \mathbb{R}^n$ be a linear operator. Then, there exists $S \subset [m]$ such that all of the following hold:*

1. *$\{Tv_i\}_{i\in S}$ is a linearly independent set, that is, T is invertible on the subspace $\operatorname{span}\{v_i : i \in S\}$ of $\mathbb{R}^m$.*
2. *$|S| = \lfloor \epsilon^2 \|T\|_F^2 / \|T\|_2^2 \rfloor$*
3. *The $|S|$th largest eigenvalue of $\sum_{i\in S}(Tv_i)(Tv_i)^\top$ is greater than $(1-\epsilon)^2 \cdot \|T\|_F^2/m$.*

$$\lambda_{|S|}\left(\sum_{i\in S}(Tv_i)(Tv_i)^\top\right) > \frac{(1-\epsilon)^2\|T\|_F^2}{m}$$

The original restricted invertibility of Bourgain-Tzafriri follows as an immediately corollary. Let's see how.

Proof of Theorem 9. Fix an arbitrary $0 < \epsilon < 1$, and let $c = c(\epsilon) = \epsilon^2$ and $C = C(\epsilon) = (1-\epsilon)^2$. Let $\{v_i\}_1^m = \{e_i\}_1^n$ (so $m = n$). We must show that there exists a subset $S \subset [n]$ with cardinality $|S| \ge cn/\|T\|_2^2$ such that $\|Tv\|_2^2 \ge C\|v\|_2^2$ for all $v \in \operatorname{span}\{e_i : i \in S\}$.

Observe that since $\|Te_i\|_2^2 = 1$ for all $1 \le i \le n$,

$$\|T\|_F^2 = \sum_{i=1}^{n} \|Te_i\|_2^2 = n$$

so that

$$\lfloor \epsilon^2 \|T\|_F^2 / \|T\|_2^2 \rfloor = \lfloor cn/\|T\|_2^2 \rfloor \qquad \text{and} \qquad \frac{(1-\epsilon)^2 \|T\|_F^2}{m} = (1-\epsilon)^2 = C(\epsilon)$$

Hence, applying Theorem 10 immediately yields a subset $S \subset [n]$ with cardinality $|S| \ge \lfloor cn/\|T\|_2^2 \rfloor$ such that $\{Te_i\}_{i \in S}$ is linearly independent. Finally, observe that for all $v \in \operatorname{span}\{e_i : i \in S\}$, we have

$$\|Tv\|_2^2 \ge \lambda_{|S|}\left(\sum_{i \in S} (Tv_i)(Tv_i)^\top\right) > C(\epsilon)$$

where P_U is the orthogonal projection onto the $|S|$-dimensional subspace $U = \operatorname{span}\{e_i : i \in S\}$. This completes the proof. □

Note that Theorem 10 also implies another simpler form of restricted invertibility proven in [33] using slightly more advanced machinery. We will not discuss that proof here; instead, we will observe that it follows immediately from Theorem 10.

Theorem 11 (Theorem 3.1 from [33]). *Suppose $v_1, \dots, v_m \in \mathbb{R}^n$ are vectors with $\sum_{i=1}^m v_i v_i^\top = I$. Then for every $k < n$, there is a subset $S \subset [m]$ with $|S| = k$ and satisfying*

$$\lambda_k\left(\sum_{i \in S} v_i v_i^\top\right) \ge \left(1 - \sqrt{\frac{k}{n}}\right)^2 \cdot \frac{n}{m}$$

Proof. Let $T = I$ and set $\epsilon = \sqrt{k/n}$. Note that $\|T\|_F^2 = n$ since T has n nonzero entries, all equal to 1, and $\|T\|^2 = 1$ since $\|Tv\|_2 = \|Iv\|_2 = \|v\|_2$ for all $v \in \mathbb{R}^n$. Theorem 10 shows there is a set $S \subset [m]$ such that $|S| = \lfloor \epsilon^2 \|T\|_F^2 / \|T\|_2^2 \rfloor = \lfloor \sqrt{k/n}^2 \cdot n \rfloor = k$, $\{v_i\}_{i \in S}$ is a linearly independent set, and

$$\left\langle v, \left(\sum_{i \in S} v_i v_i^\top\right) v \right\rangle > \left(1 - \sqrt{\frac{k}{n}}\right)^2 \cdot \frac{n}{m} \cdot \|v\|_2^2$$

for all $v \in \operatorname{span}\{v_i : i \in S\}$. Since $\{v_i\}_{i \in S}$ is a linearly independent set, $\dim \operatorname{span}\{v_i : i \in S\} = |S| = k$ so that by Theorem 10

$$\lambda_k\left(\sum_{i \in S} v_i v_i^\top\right) \ge \left(1 - \sqrt{\frac{k}{n}}\right)^2 \cdot \frac{n}{m}$$

as desired. □

4.2 A Single Univariate Barrier with Interlacing

We will build up S iteratively. At the start of step j:

1. S will contain j indices.
2. We will also have an operator $A^{(j)} = \sum_{i \in S_j} (Tv_i)(Tv_i)^\top$, where S_j is the state of S at step j

We will begin with $S = \emptyset$ and $A^{(0)} = 0$. We must show that at the end of every step $0 \le j < t$, for $t = \lfloor \epsilon^2 \|T\|_F^2 / \|T\|^2 \rfloor$ total iterations, there is a "good" $i \in [m] \setminus S_j$ to add to S_j, i.e. a "good" vector v_i for which we can perform a rank-1 update to $A^{(j)}$ with v_i. We will measure this "goodness" by a potential function; this potential will guide this process throughout. For a $b \in \mathbb{R}$, define the potential function w.r.t. b by

$$\Phi_b(A) = \sum_{i=1}^{m} (Tv_i)^\top (A - bI)^{-1} (Tv_i) = \sum_{i=1}^{m} v_i^\top (T^\top (A - bI)^{-1} T) v_i$$

for all $A \in \mathbb{R}^{n \times n}$. Note that since $\sum_{i=1}^m v_i v_i^\top = I$, by Corollary 2, we may also compactly express $\Phi_b(A)$ as

$$\Phi_b(A) = \operatorname{tr}\left[T^\top (A - bI)^{-1} T \right]$$

We will show that if the $\Phi_b(A^{(j)})$ is sufficiently small, then there will be a "good" vector v_i to add. Since each rank-1 update can shift some eigenvalues of $A^{(j)}$, in order to keep adding vectors, we will also need to shift the barrier appropriately at each step so as to keep the value of the potential sufficiently small. Specifically, at iteration j, we will have a value b_j for which $\Phi_{b_j}(A^{(j)})$ is sufficiently small. This will be made precise in Lemma 9. The invariant we will maintain is: at the end of step j (after j vectors have been added), $A^{(j)}$ has exactly j nonzero eigenvalues, which are all strictly greater than b_j.

Thus, to satisfy the third guarantee of Theorem 10, it suffices to lower bound b_t for $t = \lfloor \epsilon^2 \|T\|_F^2 / \|T\|^2 \rfloor$. The challenge is showing this process can be done without shifting the barriers too much; if we shift the barrier too much, the lower bound on $\lambda_{|S|}\left(\sum_{i \in S} v_i v_i^\top\right)$ at the end will not be interesting.

An interpretation of the parameter b in the potential function definition is that it serves as a "barrier" for the eigenvalues of the $A^{(j)}$; for the purposes of our proof, this "barrier" will separate the nonzero (in particular, positive) eigenvalues of $A^{(j)}$ from the zero eigenvalues.

We will now make this discussion formal. Specifically, we prove the two lemmas. The first is a supporting claim. The second is the core of the argument; it will tell us when we can update $A^{(j)}$ and how much we are allowed to move our barrier.

Lemma 8 (Lemma 3 from [30]). *Suppose $A \succeq 0$ has k nonzero eigenvalues, which are all strictly greater than $b > 0$. If $w \neq 0$ and $w^\top (A - bI)^{-1} w < -1$, then $A + ww^\top$ has $k+1$ nonzero eigenvalues, which are all strictly greater than b.*

Lemma 9 (Lemma 4 from [30]). *Suppose $A \succeq 0$ and $b_{\text{old}} \in \mathbb{R}$; let Q be the orthonormal projection onto $\ker A$. If all of the following hold*

(1) has k nonzero eigenvalues, which are all strictly greater than b_{old}

(2) $\Phi_{b_{\text{old}}}(A) \le -m - (\|T\|_2^2 / \delta)$

(3) $0 < \delta < b \le \delta \|QT\|_F^2 / \|T\|_2^2$

w.r.t. b, then there exists an $i \in [m]$ such that:

1. *$A + (Tv_i)(Tv_i)^\top$ has $k+1$ nonzero eigenvalues, which are all strictly greater than $b_{\text{new}} = b_{\text{old}} - \delta$*
2. *$\Phi_{b_{\text{new}}}(A + (Tv_i)(Tv_i)^\top) \le \Phi_{b_{\text{old}}}(A)$.*

First, let us see how these two lemmas combined allow us to prove Theorem 10. Notice that in the previous lemma, guarantee 1 tells us how to shift the barrier, and guarantee 2 ensures that if $\Phi_{b_j}(A^{(j)})$ is sufficiently small, then so is $\Phi_{b_{j+1}}(A^{(j+1)})$. This naturally leads to an inductive argument; we just need to choose the initial barrier value b_0 and the "step size" $\delta > 0$ of the barrier at each iteration.

Proof of Theorem 10. Initialize this iterative process with $A^{(0)} = 0$ and

$$b_0 = \frac{(1-\epsilon)\|T\|_F^2}{m}$$

and step size

$$\delta = \frac{(1-\epsilon)\|T\|_2^2}{\epsilon m}$$

Let $0 < \epsilon < 1$. Note that we may assume that $\epsilon^2 \|T\|_F^2 / \|T\|_2^2 \geq 1$ since the theorem is trivially true otherwise.

Our goal is to inductively prove that $A^{(j)}$ satisfies all conditions (1), (2) and (3) of Lemma 9 w.r.t. b_j, for all steps $0 \leq j \leq t-1$, where $t = \lfloor \epsilon^2 \|T\|_F^2 / \|T\|_2^2 \rfloor$, so that at least t vectors are added through this entire process.

For the base case $j = 0$, note that $A^{(0)} = 0$ trivially has 0 nonzero eigenvalues, satisfying (1); as a result, we also have $Q^{(0)} = P_{\ker A^{(0)}} = I$. Furthermore, since $A^{(0)} = 0$, we have by definition of the potential that

$$\Phi_{b_0}(A^{(0)}) = \operatorname{tr}\left[T^\top (A^{(0)} - b_0 I)^{-1} T\right] = -\frac{1}{b_0}\operatorname{tr}\left[T^\top T\right] = -\frac{\|T\|_F^2}{b_0} = -\frac{m}{1-\epsilon} = -m - \frac{\epsilon m}{1-\epsilon} = -m - \frac{\|T\|_2^2}{\delta}$$

so that (2) is satisfied as well. Finally, since $\epsilon < 1$ and $\epsilon^2 \|T\|_F^2 / \|T\|_2^2 \geq 1$, we have $\|T\|_F^2 / \|T\|_2^2 \geq 1/\epsilon^2 > 1/\epsilon$. Thus,

$$\delta = \frac{(1-\epsilon)\|T\|_2^2}{m} \cdot \frac{1}{\epsilon} < \frac{(1-\epsilon)\|T\|_2^2}{m} \cdot \frac{\|T\|_F^2}{\|T\|_2^2} = b_0$$

and

$$b_0 = \frac{(1-\epsilon)\|T\|_F^2}{m} = \frac{(1-\epsilon)\|T\|_2^2}{m} \cdot \frac{\|T\|_F^2}{\|T\|_2^2} < \frac{(1-\epsilon)\|T\|_2^2}{\epsilon m} \cdot \frac{\|T\|_F^2}{\|T\|_2^2} = \delta \cdot \frac{\|T\|_F^2}{\|T\|_2^2} = \delta \cdot \frac{\left\|Q^{(0)}T\right\|_F^2}{\|T\|_2^2}$$

using the fact that $Q^{(0)} = I$. This verifies (3) and we have the base case for the induction.

Suppose (1), (2) and (3) are all satisfied up to some $0 \leq j < t-1$. Since $A^{(j)}$ satisfies (1), (2) and (3) w.r.t. b_j, by Lemma 9, we may choose $1 \leq i \leq m$ such that $A^{(j+1)} = A^{(j)} + (Tv_i)(Tv_i)^\top$ has exactly $k+1$ nonzero eigenvalues, all of which are strictly greater than $b_{j+1} = b_j - \delta$, and $\Phi_{b_{j+1}}(A^{(j+1)}) \leq \Phi_{b_j}(A^{(j)})$. Note Lemma 9 immediately tells us that $A^{(j+1)}$ satisfies (1); (2) also comes immediately since $\Phi_{b_{j+1}}(A^{(j+1)}) \leq \Phi_{b_j}(A^{(j)}) \leq -m - (\|T\|_2^2/\delta)$, where in the last inequality, we use the induction hypothesis. For (3), observe that

$$\delta < b_k = b_0 - k\delta \iff \delta(k+1) < b_0 \iff k < \frac{b_0}{\delta} - 1 = \frac{\epsilon\|T\|_F^2}{\|T\|_2^2} - 1 < \Longleftarrow\ k \leq \left\lfloor \frac{\epsilon^2\|T\|_F^2}{\|T\|_2^2} \right\rfloor - 1 = t - 1$$

Since $j < t-1 \implies j+1 \leq t-1$, we have $\delta < b_{j+1}$, giving the left half of (3). Finally, observe that

$$\left\|Q^{(j+1)}T\right\|_F^2 \geq \left\|Q^{(j)}T\right\|_F^2 - \|T\|_2^2$$

so that

$$b_j \leq \delta\frac{\left\|Q^{(j)}T\right\|_F^2}{\|T\|_2^2} \iff b_{j+1} \leq \delta\frac{\left\|Q^{(j)}T\right\|_F^2}{\|T\|_2^2} - \delta \leq \delta\frac{\left\|Q^{(j+1)}T\right\|_F^2 + \|T\|_2^2}{\|T\|_2^2} - \delta = \delta\frac{\left\|Q^{(j+1)}T\right\|_F^2}{\|T\|_2^2}$$

This verifies that the right half of (3) holds for $A^{(j+1)}$ w.r.t. b_{j+1}. This completes the induction.

Finally, after adding a total of t vectors, our barrier now sits at

$$b_0 - t\delta \geq \frac{(1-\epsilon)\|T\|_F^2}{m} - \epsilon^2 \frac{\|T\|_F^2}{\|T\|_2^2} \cdot \frac{(1-\epsilon)\|T\|_2^2}{\epsilon m} = \frac{(1-\epsilon)\|T\|_F^2}{m} - \frac{\epsilon(1-\epsilon)\|T\|_F^2}{m} = \frac{(1-\epsilon)^2\|T\|_F^2}{m}$$

Since all t nonzero eigenvalues of $A^{(t)}$ are strictly greater than $b_t = b_0 - t\delta$, we have the desired bound. Note that we also get the linear independence of $\{Tv_i\}_{i\in S}$ for free, since $A^{(t)}$ has exactly $t = |S|$ nonzero eigenvalues. This proves the theorem. □

Now, it remains to prove the two lemmas.

Proof of Lemma 8. Let $0 < b < \lambda_1 \leq \cdots \leq \lambda_k$ be the nonzero eigenvalues of A. Let $\mu_1 \leq \cdots \leq \mu_{k+1}$ be the $k+1$ largest eigenvalues of $A + ww^\top$. By the Cauchy Interlacing Theorem (Theorem 8), we have that these eigenvalues interlace.

$$\mu_1 \leq \lambda_1 \leq \cdots \leq \lambda_k \leq \mu_{k+1}$$

Now, the Sherman-Morrison formula (Lemma 7) shows that

$$\begin{aligned}
\operatorname{tr}[(A + ww^\top - bI)^{-1}] - \operatorname{tr}[(A-bI)^{-1}] &= \operatorname{tr}\left[(A-bI)^{-1} - \frac{(A-bI)^{-1}ww^\top(A-bI)^{-1}}{1 + w^\top(A-bI)^{-1}w}\right] - \operatorname{tr}[(A-bI)^{-1}] \\
&= -\frac{1}{1 + w^\top(A-bI)^{-1}w} \operatorname{tr}[(A-bI)^{-1}ww^\top(A-bI)^{-1}] \\
&= \frac{w^\top(A-bI)^{-2}w}{1 + w^\top(A-bI)^{-1}w}
\end{aligned}$$

where we use linearity of trace, the equality $\operatorname{tr}(AB) = \operatorname{tr}(BA)$ (Lemma 4), as well as invertibility of $A - bI$ (which follows from the fact that all eigenvalues of A are either zero or strictly greater than b); note this last quantity is positive, since $(A-bI)^{-2} = ((A-bI)^{-1})^2$ being positive semidefinite and $w \neq 0$ implies the numerator is positive, and the assumption that $w^\top(A-bI)^{-1}w < -1$ implies the denominator is negative.

Now,

$$\operatorname{tr}[(A-bI)^{-1}] = \sum_{i=1}^{k} \frac{1}{\lambda_i - b} + \sum_{i=k+1}^{n} \frac{1}{0-b} = -\frac{n-k}{b} + \sum_{i=1}^{k} \frac{1}{\lambda_i - b}$$

and

$$\operatorname{tr}[(A + ww^\top - bI)^{-1}] = \sum_{i=1}^{k+1} \frac{1}{\mu_i - b} + \sum_{i=k+2}^{n} \frac{1}{0-b} = -\frac{n-k-1}{b} + \sum_{i=1}^{k+1} \frac{1}{\mu_i - b}$$

Combining these, we have

$$0 < \operatorname{tr}[(A + ww^\top - bI)^{-1}] - \operatorname{tr}[(A-bI)^{-1}] = \frac{1}{\mu_{k+1} - b} + \frac{1}{b} + \sum_{i=1}^{k}\left(\frac{1}{\mu_i - b} - \frac{1}{\lambda_i - b}\right) \leq \frac{1}{\mu_{k+1} - b} + \frac{1}{b}$$

where we use the fact that $\frac{1}{\mu_i - b} \leq \frac{1}{\lambda_i - b}$, which follows from interlacing. Finally, since

$$0 < \frac{1}{\mu_{k+1} - b} + \frac{1}{b} = \frac{b}{b(\mu_{k+1} - b)} + \frac{\mu_{k+1} - b}{b(\mu_{k+1} - b)} = \frac{\mu_{k+1}}{b(\mu_{k+1} - b)}$$

and $\mu_{k+1} > 0$, we must have $\mu_{k+1} > b$. □

Proof of Lemma 9. Our goal is to prove that by moving the barrier by an amount $\delta > 0$ that is not too large and not too small, we can add a vector without increasing the value of the potential function. First, let us write out the potential function of the matrix after the update w.r.t. b_{new} in terms of the potential function of the matrix prior to the update w.r.t. b_{new}. Again, we use the Sherman-Morrison formula (Lemma 7).

$$\begin{aligned}\Phi_{b_{\text{new}}}(A + ww^\top) &= \operatorname{tr}\left[T^\top (A + ww^\top - b_{\text{new}} I)^{-1} T\right] \\ &= \operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} T\right] - \frac{\operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} ww^\top (A - b_{\text{new}})^{-1} T\right]}{1 + w^\top (A - b_{\text{new}} I)^{-1} w} \\ &= \Phi_{b_{\text{new}}}(A) - \frac{w^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} w}{1 + w^\top (A - b_{\text{new}} I)^{-1} w}\end{aligned}$$

Again, our goal to ensure there exists a nonzero $w \in \{Tv_i\}_{i=1}^m$ such that the above is upper bounded by $\Phi_b(A)$, the potential function evaluated at the previous matrix w.r.t. the previous barrier value.

Now, the set of nonzero vectors w satisfying both

$$\Phi_{b_{\text{new}}}(A) - \frac{w^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} w}{1 + w^\top (A - b_{\text{new}} I)^{-1} w} \le \Phi_b(A) \qquad \text{and} \qquad w^\top (A - b_{\text{new}} I)^{-1} w < -1$$

are precisely those for which

$$w^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} w \le (\Phi_b(A) - \Phi_{b-\delta}(A)) \cdot (-1) \cdot (1 + w^\top (A - b_{\text{new}} I)^{-1} w)$$

by applying the second inequality to the first.

Now, to show that there exists a $w \in \{Tv_i\}_{i=1}^m$ satisfying this, it suffices to show that the sum over $\{Tv_i\}_{i=1}^m$ of the left-hand side is less than or equal to the sum over $\{Tv_i\}_{i=1}^m$ of the right-hand side.

$$\sum_{i=1}^m v_i^\top T^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} Tv_i \le (\Phi_b(A) - \Phi_{b-\delta}(A)) \cdot \sum_{i=1}^m (-1) \cdot (1 + v_i^\top T^\top (A - b_{\text{new}} I)^{-1} Tv_i)$$

By Corollary 2, we can rewrite both sides of the above inequality as

$$\begin{aligned}\operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} T\right] &\le (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \cdot (-1) \cdot \left(m + \operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} T\right]\right) \\ &= (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \cdot (-m - \Phi_{b_{\text{new}}}(A))\end{aligned}$$

For this, it suffices to prove the following tighter bound.

$$\operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} T\right] \le (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \cdot \left(\frac{\|T\|_2^2}{\delta} + (\Phi_b(A) - \Phi_{b_{\text{new}}}(A))\right)$$

since $\Phi_b(A) \le -m - (\|T\|_2^2/\delta)$ implies

$$\Phi_{b-\delta}(A) = \Phi_b(A) - (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \le -m - \frac{\|T\|_2^2}{\delta} - (\Phi_b(A) - \Phi_{b-\delta}(A))$$

Towards this, we first make the observation that since $TT^\top \preceq \|T\|_2^2 I$,

$$T^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} T \preceq \|T\|_2^2 \cdot T^\top (A - b_{\text{new}} I)^{-2} T$$

and so by Corollary 1 we can bound the trace as

$$\operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-1} TT^\top (A - b_{\text{new}} I)^{-1} T\right] \le \|T\|_2^2 \operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-2} T\right]$$

Thus, to prove the lemma, it suffices to show that

$$\|T\|_2^2 \operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-2} T\right] \leq (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \cdot \left(\frac{\|T\|_2^2}{\delta} + (\Phi_b(A) - \Phi_{b_{\text{new}}}(A))\right)$$

Now, decompose I into $P + Q$, where P, Q are the orthogonal projections (i.e. satisfying $P^2 = P = P^\top$ and $Q^2 = Q = Q^\top$) onto $\operatorname{im}(A)$ and $\ker(A)$, respectively. Since $P, Q, A, (A - b_{\text{new}})^{-1}$ and $(A - b_{\text{new}} I)^{-2}$ are mutually diagonalizable by Lemma 3, we may decompose our potential function into

$$\Phi_b(A) = \Phi_b^P(A) + \Phi_b^Q(A), \quad \text{where} \quad \Phi_b^P(A) \stackrel{\text{def}}{=} \operatorname{tr}\left[T^\top P (A - b_{\text{old}} I)^{-1} P T\right], \Phi_b^Q(A) \stackrel{\text{def}}{=} \operatorname{tr}\left[T^\top Q (A - b_{\text{old}} I)^{-1} Q T\right]$$

and the trace above into

$$\operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-2} T\right] = \operatorname{tr}\left[T^\top P (A - b_{\text{new}} I)^{-2} P T\right] + \operatorname{tr}\left[T^\top Q (A - b_{\text{new}} I)^{-2} Q T\right]$$

using Corollary 3. Will then proceed by bounding each of the two terms on the right-hand side separately.

Let's consider the term corresponding to P first. Now, $x^\top (A - b_{\text{old}} I)^{-1} x \geq 0$ for all $x \in \operatorname{im}(A)$ since the eigenvectors of $(A - b_{\text{old}} I)^{-1}$ corresponding to the k nonnegative eigenvalues of $(A - b_{\text{old}} I)^{-1}$ (which exist since A has exactly k nonzero eigenvalues, all larger than b) span $\operatorname{im}(A)$. Since P is the projection onto $\operatorname{im}(A)$ and $P = P^\top$, $Px = P^\top x \in \operatorname{im}(A)$ so that $x^\top P (A - b_{\text{old}} I)^{-1} P x = (Px)^\top (A - b_{\text{old}} I)^{-1} (Px) \geq 0$ for all $x \in \mathbb{R}^n$. Thus, $P(A - b_{\text{old}} I)^{-1} P \geq 0$. By the same reasoning, since $0 < b_{\text{new}} < b$, $P(A - b_{\text{new}} I)^{-1} P \geq 0$. As a result, we have

$$\delta P (A - b_{\text{new}} I)^{-1} P \leq P (A - b_{\text{old}} I)^{-1} P - P (A - b_{\text{new}} I)^{-1} P$$

which immediately implies

$$\begin{aligned}
\operatorname{tr}\left[T^\top P (A - b_{\text{new}} I)^{-2} P T\right] &\leq \frac{1}{\delta} \operatorname{tr}\left[T^\top \left[P (A - b_{\text{old}} I)^{-1} P - P (A - b_{\text{new}} I)^{-1} P\right] T\right] \\
&= \frac{1}{\delta} \operatorname{tr}\left[T^\top P (A - b_{\text{old}} I)^{-1} P T\right] - \operatorname{tr}\left[T^\top P (A - b_{\text{new}} I)^{-1} P T\right] = \frac{1}{\delta}\left(\Phi_b^P(A) - \Phi_{b_{\text{new}}}^P(A)\right)
\end{aligned}$$

Now, let's turn to the term corresponding to Q. First, we make the following trace computations:

$$\operatorname{tr}\left[T^\top Q (A - b_{\text{new}} I)^{-2} Q T\right] = \frac{\|QT\|_F^2}{b_{\text{new}}^2}$$

and

$$\Phi_b^Q(A) - \Phi_{b_{\text{new}}}^Q(A) = \operatorname{tr}\left[T^\top Q \left((A - b_{\text{old}} I)^{-1} - (A - b_{\text{new}} I)^{-1}\right) Q T\right] = \delta \frac{\|QT\|_F^2}{b b_{\text{new}}}$$

This is where we will use the assumption that $b \leq \delta \|QT\|_F^2 / \|T\|_2^2 \iff \|T\|_2^2 \leq \delta \|QT\|_F^2 / b$. Combining this with our trace computations, we have

$$\|T\|_2^2 \operatorname{tr}\left[T^\top Q (A - b_{\text{new}} I)^{-2} Q T\right] \leq \left(\Phi_b^Q(A) - \Phi_{b_{\text{new}}}^Q(A)\right)\left(\frac{\|T\|_2^2}{\delta} + \left(\Phi_b^Q(A) - \Phi_{b_{\text{new}}}^Q(A)\right)\right)$$

This immediately implies

$$\|T\|_2^2 \operatorname{tr}\left[T^\top Q (A - b_{\text{new}} I)^{-2} Q T\right] \leq (\Phi_b(A) - \Phi_{b-\delta}(A))\left(\frac{\|T\|_2^2}{\delta} + (\Phi_b(A) - \Phi_{b-\delta}(A))\right) - (\Phi_b(A) - \Phi_{b-\delta}(A)) \cdot \frac{\|T\|_2^2}{\delta}$$

Combining with the bound

$$\operatorname{tr}\left[T^\top P (A - b_{\text{new}} I)^{-2} P T\right] \leq \frac{1}{\delta}\left(\Phi_b^P(A) - \Phi_{b_{\text{new}}}^P(A)\right)$$

we obtained earlier for the term corresponding to P, we have

$$\|T\|_2^2 \operatorname{tr}\left[T^\top (A - b_{\text{new}} I)^{-2} T\right] \leq (\Phi_b(A) - \Phi_{b_{\text{new}}}(A)) \cdot \left(\frac{\|T\|_2^2}{\delta} + (\Phi_b(A) - \Phi_{b_{\text{new}}}(A))\right)$$

proving the lemma. □

4.3 An Algorithm

The original proof of the restricted invertibility was nonconstructive. It did not naturally give an algorithm to find such a well-invertible column submatrix. Tropp [23] gave the first randomized polynomial time algorithm to compute such a submatrix.

Theorem 12 (Theorem 5.7 from [23]). *There exists an algorithm such that given a matrix $A \in \mathbb{R}^{m\times n}$ satisfying $\|A_i\|_2 = 1$, for all columns A_i of A, produces a $S \subset [m]$ such that $|S| \geq c\|A\|_F^2/\|A\|_2^2$ and $\kappa(A_S) \leq \sqrt{3}$, where $c > 0$ is an absolute constant. Furthermore, this algorithm runs in $\tilde{\mathcal{O}}(|S|^2 m + |S|^{3.5})$ time.*

The core of the algorithm given in [23] is based on repeatedly sampling a uniformly random subset σ of $[m]$ and then solving a convex program over the probability simplex (to compute something known as a Grothendieck factorization) to obtain a subset of σ as a candidate solution. We refer the reader to [23] for details.

The proof of Theorem 10 suggests a completely different algorithm that is much simpler to understand and analyze. It also has the added benefit of being completely deterministic, and runs in $\mathcal{O}(mn^3)$-time.

Algorithm 1

Input: vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ satisfying $\sum_{i=1}^m v_i v_i^\top = I$; a linear operator $T : \mathbb{R}^n \to \mathbb{R}^n$ provided as a square matrix in $\mathbb{R}^{n\times n}$; $0 < \epsilon < 1$

Output: a set S of indices satisfying the guarantees of Theorem 10

compute and cache $\|T\|_F^2$ and $\|T\|_2^2$
compute and cache Tv_i for all $1 \leq i \leq m$
compute and cache outer product $(Tv_i)(Tv_i)^\top$ for all $1 \leq i \leq m$
initialize $S = \emptyset$, $b_0 = (1-\epsilon)\|T\|_F^2/m$, $\delta = (1-\epsilon)\|T\|_2^2/\epsilon m$, $A^{(0)} = 0$
while True **do**
    $b_{j+1} \leftarrow b_j - \delta$
    if $b_{j+1} > (1-\epsilon)^2 \cdot \|T\|_F^2/m$ **then**
        return S
    end if
    compute $(A - b_{j+1}I)^{-1}$
    for $i = 1, \ldots, m$ **do**
        compute eigenvalues of $A^{(j)} + (Tv_i)(Tv_i)^\top$
        compute $\Phi_{b_{j+1}}(A^{(j)} + (Tv_i)(Tv_i)^\top)$
        if $\lambda_{j+1}(A^{(j)} + (Tv_i)(Tv_i)^\top) > b_{j+1}$ && $\Phi_{b_{j+1}}(A^{(j)} + (Tv_i)(Tv_i)^\top) \leq \Phi_{b_j}(A^{(j)})$ **then**
            $A^{(j+1)} \leftarrow A^{(j)} + (Tv_i)(Tv_i)^\top$
            $S \leftarrow S \cup \{i\}$
            break out of for-loop
        end if
    end for
    if $A^{(j+1)}$ is not created (no good update remaining) **then**
        return S
    end if
end while
return S

5 Warm-Up II: Spectral Graph Sparsification

5.1 Prior Work

Problem 2 (Graph Sparsification)**.** *Suppose we are given an undirected weighted graph $G = (V, E, w)$, where $w : E \to \mathbb{R}_{\geq 0}$. Does G have a (re)weighted subgraph $H = (V, \tilde{E}, \tilde{w})$ such that H is sparse, having ,say, $\tilde{\mathcal{O}}(n)$ edges, but still approximates G in some sense that we are interested in?*

Graph sparsification naturally finds applications in big data. Large graphs are pervasive in machine learning, social network analysis, etc. in today's world. In most cases, these graphs cannot be stored in the memory of a single or a small number of computers and so we need new techniques to handling these graphs efficiently. Without preprocessing, such as sparsification, scaling up our algorithms will be difficult.

One metric for "approximation" considered in the late 90s in a work of Benczur-Karger was cut size. They proved the following result, which immediately led to

1. an $\tilde{\mathcal{O}}(n^2/\epsilon^2)$-time algorithm for computing a $(1+\epsilon)$-approximation to the value of the minimum s-t cut of an undirected graph (Corollary 1.3 from [12])
2. an $\tilde{\mathcal{O}}(n^2/\epsilon^2)$-time algorithm for computing an $\mathcal{O}(\log n)$-approximation to the sparsest cut (Corollary 1.5 from [12])

Theorem 13 (Theorem 1.2 from [12])**.** *Let $G = (V, E, w)$ be an undirected weighted graph, and $\epsilon > 0$. There exists an $\mathcal{O}(m \log^3 n)$ algorithm that, with high probability, produces a subgraph $H = (V, \tilde{E}, \tilde{w})$ such that:*

1. *$|\tilde{E}| \leq \mathcal{O}(n \log n/\epsilon^2)$, i.e. H is sparse.*
2. *$(1-\epsilon) \cdot |E_G(S, \overline{S})| \leq |E_H(S, \overline{S})| \leq (1+\epsilon) \cdot |E_G(S, \overline{S})|$ for all cuts $(S, \overline{S})$.*

We will concern ourselves with a stronger form of approximation, namely that of *spectral sparsification*. We say that $H = (V, \tilde{E}, \tilde{w})$ is a $(1 \pm \epsilon)$-spectral sparsifier of G if

$$(1-\epsilon) \cdot L_G \preceq L_H \preceq (1+\epsilon) \cdot L_G$$

that is,

$$(1-\epsilon) \cdot x^\top L_G x \leq x^\top L_H x \leq (1+\epsilon) \cdot x^\top L_G x$$

for all $x \in \mathbb{R}^n$. This is a stronger form of approximation than cut sparsification in the sense that a $(1 \pm \epsilon)$-spectral sparsifier is also a $(1 \pm \epsilon)$-cut sparsifier, since H is a cut sparsifier if

$$(1-\epsilon) \cdot x^\top L_G x \leq x^\top L_H x \leq (1+\epsilon) \cdot x^\top L_G x$$

holds only for all $x \in \{0,1\}^n = \{\mathbf{1}_S : S \subset V\}$. This notion of sparsification was first used in a work of Spielman-Teng [18], who used it to produce nearly linear-time solvers for certain types of linear systems. They were formally introduced by Spielman-Teng in 2008 [22], where they gave an algorithm that produces a spectral sparsifier with $\mathcal{O}(n \log^c n)$ edges, for some (large) absolute constant c.

Soon after, Spielman-Srivastava made a significant improvement to previous works based off of sampling edges with probability proportional to their effective resistances, and then applying dimensionality reduction and matrix concentration.

Theorem 14 (Theorem 1 from [21])**.** *There exists a $\tilde{\mathcal{O}}(m/\epsilon^2)$-time algorithm such that, given an undirected weighted graph $G = (V, E, w)$, with probability at least $1/2$, returns a $(1 \pm \epsilon)$-spectral sparsifier of G with $\mathcal{O}(n \log n/\epsilon^2)$ edges.*

Below, we will give a fully deterministic algorithm that produces a $(1+\epsilon)$-spectral sparsifier in $\mathcal{O}(n/\epsilon^2)$ edges in time $\mathcal{O}(mn^3/\epsilon^2)$. While this is significantly slower, the algorithm is fully deterministic, and the techniques used are of theoretical interest in themselves.

5.2 "Twice Ramanujan" Sparsifiers

The main result of this section is as follows. The constant $(d+1+2\sqrt{d})/(d+1-2\sqrt{d})$ is what motivates the name "twice Ramanujan"; in the special case when the graph is an unweighted complete graph, it says we can spectrally approximate at least this well with a weighted graph that has at most twice as many edges as a Ramanujan graph (see Definition 11 below).

Theorem 15 (Theorem 1.1 from [27]). *For every $d > 1$ and every undirected weighted graph $G = (V, E, w)$ with $n = |V|, m = |E|$, there is a weighted subgraph $H = (V, \tilde{E}, \tilde{w})$ with $|\tilde{E}| = \lceil d \cdot (n-1) \rceil$ that satisfies*

$$L_G \preceq L_H \preceq \left(\frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}} \right) \cdot L_G$$

An immeidate corollary is the fact that we can obtain $(1+\epsilon)$-spectral sparsifiers.

Corollary 5. *For every $\epsilon > 0$ and every undirected weighted graph $G = (V, E, w)$ with $n = |V|, m = |E|$, there is a weighted subgraph $H = (V, \tilde{E}, \tilde{w})$ with $|\tilde{E}| \leq \mathcal{O}(n/\epsilon^2)$ that satisfies*

$$L_G \preceq L_H \preceq (1+\epsilon) \cdot L_G$$

that is, there is a $(1+\epsilon)$-approximate spectral sparsifier of G with $\mathcal{O}(n/\epsilon^2)$ edges.

Proof. Observe that

$$\frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}} = \frac{(d+1-2\sqrt{d})+4\sqrt{d}}{d+1-2\sqrt{d}} = 1 + \frac{4\sqrt{d}}{d+1-2\sqrt{d}} \leq 1 + \frac{c}{\sqrt{d}}$$

for some constant $c > 0$. Let $\epsilon > 0$ and $G = (V, E, w)$ be an undirected weighted graph. Choose $d = \lceil c^2/\epsilon^2 \rceil$ so that $\sqrt{d} \geq \sqrt{c^2/\epsilon^2} = c/\epsilon \implies c/\sqrt{d} \leq \epsilon$, which implies

$$1 \leq \frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}} \leq 1 + \frac{c}{\sqrt{d}} \leq 1 + \epsilon$$

Theorem 15 shows that there is a weighted subgraph $H = (V, \tilde{E}, \tilde{w})$ with $|\tilde{E}| = \lceil d \cdot (n-1) \rceil \leq \mathcal{O}(n/\epsilon^2)$ such that

$$L_G \preceq L_H \preceq \left(\frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}} \right) \cdot L_G \leq (1+\epsilon) \cdot L_G$$

□

The core of the proof lies in a purely linear algebraic statement very similar to Theorem 10 on restricted invertibility.

Theorem 16 (Theorem 3.1 from [27]). *Let $d > 1$, and $v_1, \ldots, v_m \in \mathbb{R}^n$ satisfy $\sum_{i=1}^m v_i v_i^\top = I$. Then, there is a subset $S \subset [m]$ with $|S| \leq dn$ and nonzero scalars $s_i > 0$, for all $i \in S$, such that*

$$I \preceq \sum_{i \in S} s_i v_i v_i^\top \preceq \left(\frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}} \right) \cdot I$$

Let's first see how Theorem 15 follows from this fact.

Proof of Theorem 15. Let $G = (V, E, w)$ be an undirected weighted graph on n vertices with m edges; order the edges of G and label them by indices $i \in [m]$. We will assume G is connected: if G is disconnected, we can apply the same argument to the connected components of G. Let $L_G = B^\top W B$ be the Laplacian of G. We will restrict our attention to $\mathrm{im}(L_G) = \mathrm{span}\{\mathbf{1}\}^\perp \cong \mathbb{R}^{n-1}$ and let $d > 1$ be arbitrary. Consider the columns $v_1, \ldots, v_m$ of $V = (L_G^\dagger)^{1/2} B^\top W^{1/2} \in \mathbb{R}^{n \times m}$. Note that

$$\sum_{i=1}^m v_i v_i^\top = V V^\top = (L_G^\dagger)^{1/2} B^\top W^{1/2} (W^{1/2})^\top B \left((L_G^\dagger)^{1/2} \right)^\top = (L_G^\dagger)^{1/2} B^\top W B (L_G^\dagger)^{1/2} = (L_G^+)^{1/2} L_G^{1/2} L_G^{1/2} (L_G^\dagger)^{1/2} = I_{\mathrm{im}(L_G)}$$

Theorem 16 furnishes a subset $S \subset [m]$ of size at most $d(n-1)$ and scalars $s_i \geq 0$ with $s_i > 0$ for all $i \in S$ and $s_i = 0$ for all $i \notin S$, such that

$$I_{\mathrm{im}(L_G)} \preceq \sum_{i=1}^{m} s_i v_i v_i^\top = V S V^\top \preceq \kappa \cdot I_{\mathrm{im}(L_G)}$$

where for convenience, we set $S = \mathrm{diag}(s_1, \ldots, s_m) \in \mathbb{R}^{m \times m}$, and $\kappa = (d+1+2\sqrt{d})/(d+1-2\sqrt{d})$. Let $L_H = B^\top W^{1/2} S W^{1/2} B$ be the Laplacian of a reweighted subgraph H with edge weights $\tilde{w}_i = w_i s_i$ for all $i \in E$. Since at most $d(n-1)$ indices i have $s_i > 0$, H has edge set $\tilde{E} = \{i \in E : s_i > 0\}$, which is of size at most $d(n-1)$. It remains to verify that H is a κ-approximate spectral sparsifier of G. For this, observe that

$$\begin{aligned}
I_{\mathrm{im}(L_G)} \preceq V S V^\top \preceq \kappa \cdot I_{\mathrm{im}(L_G)} &\iff I_{\mathrm{im}(L_G)} \preceq (L_G^\dagger)^{1/2} B^\top W^{1/2} S W^{1/2} B (L_G^\dagger)^{1/2} \preceq \kappa \cdot I_{\mathrm{im}(L_G)} \\
&\iff I_{\mathrm{im}(L_G)} \preceq (L_G^\dagger)^{1/2} L_H (L_G^\dagger)^{1/2} \preceq \kappa \cdot I_{\mathrm{im}(L_G)} \\
&\iff y^\top y \leq y^\top (L_G^\dagger)^{1/2} L_H (L_G^\dagger)^{1/2} y \leq \kappa \cdot y^\top y, \quad \forall y \in \mathrm{im}(L_G) = \mathrm{im}(L_G^{1/2}) \\
&\iff x^\top (L_G^{1/2})^\top (L_G^{1/2}) x \leq x^\top (L_G^{1/2})(L_G^\dagger)^{1/2} L_H (L_G^\dagger)^{1/2} (L_G^{1/2}) x \leq \kappa \cdot x^\top (L_G^{1/2})^\top (L_G^{1/2}) x \\
&\qquad \forall x \in \ker(L_G)^\perp = \mathrm{span}\{\mathbf{1}\}^\perp \\
&\iff x^\top L_G x \leq x^\top L_H x \leq x^\top L_G x, \quad \forall x \perp \mathbf{1} \\
&\iff L_G \preceq L_H \preceq \kappa \cdot L_G
\end{aligned}$$

□

5.3 Intuition with Interlacing

Before we dive into the details of the proof, here is a heuristic calculation that tells us why we might be able to prove a result like Theorem 16.

Suppose we have some symmetric matrix $A \in \mathbb{R}^{n \times n}$ with real eigenvalues $\lambda_1, \ldots, \lambda_n$ and orthonormal eigenvectors $u_1, \ldots, u_n$. choose a uniformly random vector v from $\{v_i\}_1^m$. Then, in expectation,

$$\mathbb{E}[|\langle v, u_j \rangle|^2] = \frac{1}{m} \sum_{i=1}^{m} u_j^\top v_i v_i^\top u_j = \frac{1}{m} u_j^\top \left(\sum_{i=1}^{m} v_i v_i^\top \right) u_j = \frac{\|u_j\|^2}{m} = \frac{1}{m}$$

In particular, making a uniformly random rank-one update, in expectation, shifts all eigenvalues upward around the same amount of $1/m$. Thus, after updating A over many iterations, we'd expect the condition number $\lambda_{\max}(A)/\lambda_{\min}(A)$ to remain bounded; the largest and smallest eigenvalues won't drift too much.

It turns out that this averaging intuition does work in the following sense: if we explicitly compute the expected characteristic polynomial of $A + vv^\top$, using the same argument as in the Cauchy Interlacing Theorem (Theorem 8), we have

$$\mathbb{E}[\chi_{A+vv^\top}(x)] = \mathbb{E}\left[\chi_A(x) \left(1 - \sum_{j=1}^{n} \frac{|\langle v, u_j \rangle|^2}{x - \lambda_j} \right) \right] = \chi_A(x) \left(1 - \sum_{j=1}^{n} \frac{\mathbb{E}[|\langle v, u_j \rangle|^2]}{x - \lambda_j} \right) = \chi_A(x) \left(1 - \frac{1}{m} \sum_{j=1}^{n} \frac{1}{x - \lambda_j} \right)$$

Expanding out this last expression and applying the differentiation product rule,

$$\prod_{i=1}^{n} (x - \lambda_i) - \frac{1}{m} \sum_{j=1}^{n} \prod_{k \neq j} (x - \lambda_k) = \chi_A(x) - \frac{1}{m} \chi_A'(x)$$

Thus, starting with $A = 0$ and adding averages iteratively, after k rounds, we get

$$\chi_{A^{(k)}}(x) = \left(1 - \frac{1}{m} \cdot \frac{d}{dx} \right)^k \chi_{A^{(0)}}(x) = \left(1 - \frac{1}{m} \cdot \frac{d}{dx} \right)^k x^n$$

This last expression is a well-known object: the set of such polynomials of that form, for $k \geq 0$ forms a family of orthogonal polynomials known as the **Laguerre polynomials**. It's roots are very-well studied and it is known that after $k = dn$ rounds (see Theorem 4.4 from [10]), we will have

$$\frac{\lambda_{\max}(A^{(k)})}{\lambda_{\min}(A^{(k)})} = \frac{d+1+2\sqrt{d}}{d+1-2\sqrt{d}}$$

This is precisely the bound we hope for. This would be ideal also because it would lead to *unweighted spectral sparsifiers*, since at every step, we add the same amount of every vector v_i.

The main problem with this argument of course is that we are not allowed to add averages of vectors. We actually need to choose a single vector from $\{v_i\}_{i=1}^m$ at each step. This is why we need the weights s_i and also why the argument is much more complicated.

5.4 Two Univariate Barriers

Our goal now is to prove Theorem 16. As in restricted invertibility, build up S and the constants s_i iteratively. The argument will inductive, except now, we will employ both a "lower" barrier as well as an "upper" barrier. We need both because now, we are trying to guarantee both a lower bound as well as an upper bound on the eigenvalues of $\sum_{i\in S} s_i v_i v_i^\top$. One can visualize the lower barrier as repelling all of the eigenvalues of $\sum_{i\in S} s_i v_i v_i^\top$ upwards (which naturally happens as we make more rank-one updates), while the upper barrier keeps the top eigenvalue of $\sum_{i\in S} s_i v_i v_i^\top$ close the smallest eigenvalue.

For a symmetric matrix $A \in \mathbb{R}^{n\times n}$ with eigenvalues $\lambda_1 \leq \cdots \leq \lambda_n$, and "barrier" values $\ell, u \in \mathbb{R}$ (with $\ell \leq u$), we define the lower barrier w.r.t. ℓ of A by

$$\Phi_\ell(A) \stackrel{\text{def}}{=} \operatorname{tr}[(A-\ell I)^{-1}] = \sum_{i=1}^n \frac{1}{\lambda_i - \ell}$$

and the upper barrier w.r.t. u of A by

$$\Phi^u(A) \stackrel{\text{def}}{=} \operatorname{tr}[(uI-A)^{-1}] = \sum_{i=1}^n \frac{1}{u-\lambda_i}$$

Let's first analyze each barrier separately. We prove the following two lemmas on their movement.

Lemma 10 (Shifting Upper Barrier; Lemma 3.3 from [27])**.** *Let $A \in \mathbb{R}^{n\times n}$ be symmetric and let $v \in \mathbb{R}^n$ be any vector. Let $u_{\text{old}} > 0$ satisfy $\lambda_{\max}(A) < u_{\text{old}}$, and $u_{\text{new}} = u_{\text{old}} + \delta_u$ for some $\delta_u > 0$. Let $s > 0$. If*

$$U_A(v) \stackrel{\text{def}}{=} \frac{v^\top (u_{\text{new}} I - A)^{-2} v}{\Phi^{u_{\text{old}}}(A) - \Phi^{u_{\text{new}}}(A)} + v^\top (u_{\text{new}} I - A)^{-1} v \leq \frac{1}{s}$$

then

$$\Phi^{u_{\text{new}}}(A + svv^\top) \leq \Phi^u(A) \qquad \text{and} \qquad \lambda_{\max}(A + svv^\top) < u_{\text{new}}$$

Proof. First, we apply the Sherman-Morrison Formula (Lemma 7) to express the potential of the updated matrix w.r.t. the new barrier in terms of the potential of the old matrix w.r.t. the old and new barrier values.

$$\begin{aligned}
\Phi^{u_{\text{new}}}(A+svv^\top) &= \operatorname{tr}\left[\left(u_{\text{new}} I - A - svv^\top\right)^{-1}\right] = \operatorname{tr}\left[(u_{\text{new}} I - A)^{-1} + \frac{s(u_{\text{new}} I - A)^{-1} vv^\top (u_{\text{new}} I - A)^{-1}}{1 - sv^\top (u_{\text{new}} I - A)^{-1} v}\right] \\
&= \operatorname{tr}\left[(u_{\text{new}} I - A)^{-1}\right] + \frac{s \operatorname{tr}\left[(u_{\text{new}} I - A)^{-1} vv^\top (u_{\text{new}} I - A)^{-1}\right]}{1 - sv^\top (u_{\text{new}} I - A)^{-1} v} \\
&= \Phi^{u_{\text{new}}}(A) + \frac{v^\top (u_{\text{new}} I - A)^{-2} v}{(1/s) - v^\top (u_{\text{new}} I - A)^{-1} v} \\
&= \Phi^{u_{\text{old}}}(A) - (\Phi^{u_{\text{old}}}(A) - \Phi^{u_{\text{new}}}(A)) + \frac{v^\top (u_{\text{new}} I - A)^{-2} v}{(1/t) - v^\top (u_{\text{new}} I - A)^{-1} v}
\end{aligned}$$

Note since $1/s \geq U_A(v) > v^\top (u_{\text{new}} I - A)^{-1} v$ by assumption, the denominator of the last term is finite. Thus, to prove that $\Phi^{u_{\text{new}}}(A + svv^\top) \leq \Phi^u(A)$, it suffices to prove

$$-(\Phi^{u_{\text{old}}}(A) - \Phi^{u_{\text{new}}}(A)) + \frac{v^\top (u_{\text{new}} I - A)^{-2} v}{(1/t) - v^\top (u_{\text{new}} I - A)^{-1} v} \leq 0$$

Rearranging, this is precisely $U_A(v) \leq 1/s$, which is what we assumed. This gives the first guarantee.

For the second guarantee, observe that $\Phi^{u_{\text{new}}}(A + svv^\top)$ is finite for every $s > 0$. Now, if $\lambda_{\max}(A + svv^\top) \geq u_{\text{new}}$, then there exists some $0 < \tilde{s} \leq s$ such that $\lambda_{\max}(A + \tilde{s}vv^\top) = u_{\text{new}}$, which would lead to $\Phi^{u_{\text{new}}}(A + \tilde{s}vv^\top)$, a contradiction. Thus, we have $\lambda_{\max}(A + svv^\top) < u_{\text{new}}$ as well. □

Lemma 11 (Shifting Lower Barrier; Lemma 3.4 from [27])**.** *Let $A \in \mathbb{R}^{n \times n}$ be symmetric and let $v \in \mathbb{R}^n$ be any vector. Let $\ell_{\text{old}} > 0$ satisfy $\lambda_{\min}(A) > \ell_{\text{old}}$, $\delta_\ell > 0$ satisfy $\Phi_\ell(A) \leq 1/\delta_\ell$, and $\ell_{\text{new}} = \ell_{\text{old}} + \delta_\ell$. Let $s > 0$. If*

$$L_A(v) \stackrel{\text{def}}{=} \frac{v^\top (A - \ell_{\text{new}} I)^{-2} v}{\Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)} - v^\top (A - \ell_{\text{new}} I)^{-1} v \geq \frac{1}{s} > 0$$

then

$$\Phi_{\ell_{\text{new}}}(A + svv^\top) \leq \Phi_{\ell_{\text{old}}}(A) \qquad \text{and} \qquad \lambda_{\min}(A + svv^\top) > \ell_{\text{new}}$$

Proof. Observe that since $\lambda_{\min}(A) > \ell_{\text{old}}$, all eigenvalues of A are strictly larger than ℓ_{old}. Hence, $\frac{1}{\lambda_i - \ell_{\text{old}}} > 0$ for all $1 \leq i \leq n$. Combining this with $\Phi_{\ell_{\text{old}}}(A) = \sum_{i=1}^n \frac{1}{\lambda_i - \ell_{\text{old}}} \leq 1/\delta_\ell$ shows that $\frac{1}{\lambda_{\min}(A) - \ell_{\text{old}}} < 1/\delta_\ell$. Rearranging shows $\lambda_{\min}(A) > \ell_{\text{old}} + \delta_\ell = \ell_{\text{new}}$. Hence, for every $s > 0$, $\lambda_{\min}(A + svv^\top) > \lambda_{\min}(A) > \ell_{\text{new}}$. The intuition here is that $\Phi_{\ell_{\text{old}}}(A)$ being bounded implies the eigenvalues of A are a "safe" distance away from ℓ_{old}, the previous barrier value. This proves the second guarantee.

As above, for the first guarantee, we begin by applying the Sherman-Morrison Formula (Lemma 7) to express the potential of the updated matrix w.r.t. the new barrier in terms of the potential of the old matrix w.r.t. the old and new barrier values.

$$\begin{aligned}
\Phi_{\ell_{\text{new}}}(A + svv^\top) &= \operatorname{tr}\left[\left(A + svv^\top - \ell_{\text{new}} I\right)^{-1}\right] = \operatorname{tr}\left[(A - \ell_{\text{new}} I)^{-1} - \frac{s(A - \ell_{\text{new}} I)^{-1} vv^\top (A - \ell_{\text{new}} I)^{-1}}{1 + sv^\top (A - \ell_{\text{new}} I)^{-1} v}\right] \\
&= \operatorname{tr}\left[(A - \ell_{\text{new}} I)^{-1}\right] - \frac{s \operatorname{tr}\left[(A - \ell_{\text{new}} I)^{-1} vv^\top (A - \ell_{\text{new}} I)^{-1}\right]}{1 + sv^\top (A - \ell_{\text{new}} I)^{-1} v} \\
&= \Phi_{\ell_{\text{new}}}(A) - \frac{v^\top (A - \ell_{\text{new}} I)^{-2} v}{(1/s) + v^\top (A - \ell_{\text{new}} I)^{-1} v} \\
&= \Phi_{\ell_{\text{old}}}(A) + \left(\Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)\right) - \frac{v^\top (A - \ell_{\text{new}} I)^{-2} v}{(1/s) + v^\top (A - \ell_{\text{new}} I)^{-1} v}
\end{aligned}$$

Thus, to prove $\Phi_{\ell_{\text{new}}}(A + svv^\top) \leq \Phi_{\ell_{\text{old}}}(A)$, it suffices to show

$$\left(\Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)\right) - \frac{v^\top (A - \ell_{\text{new}} I)^{-2} v}{(1/s) + v^\top (A - \ell_{\text{new}} I)^{-1} v} \leq 0$$

Rearranging, we have $L_A(v) \geq 1/s$, which is what we assumed. □

Now, we are almost ready for movement of both barriers. Before we dive in, we need one final technical lemma.

Lemma 12 (Claim 3.6 from [27])**.** *Suppose $\lambda_{\min}(A) > \ell_{\text{old}}$, $0 \leq \Phi_{\ell_{\text{old}}}(A) \leq \epsilon_\ell$ and $\frac{1}{\delta_\ell} - \epsilon_\ell \geq 0$. Then*

$$\frac{\sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}}{\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1} (\lambda_i - \ell_{\text{old}})^{-1}} - \Phi_{\ell_{\text{new}}}(A) \geq \frac{1}{\delta_\ell} - \Phi_{\ell_{\text{old}}}(A)$$

Proof. Observe that since $\frac{1}{\delta_\ell} - \epsilon_\ell \geq 0$, $\delta_\ell \leq 1/\epsilon_\ell$. Furthermore, since $\frac{1}{\lambda_{\min}(A)-\ell_{\text{old}}} < \sum_{i=1}^n \frac{1}{\lambda_i - \ell_{\text{old}}} = \Phi_{\ell_{\text{old}}}(A) \leq \epsilon_\ell$, we have $\lambda_{\min}(A) - \ell_{\text{old}} > 1/\epsilon_\ell$. Putting these together, we see that $\lambda_{\min}(A) - \ell_{\text{old}} > \delta_\ell \implies \lambda_{\min}(A) > \ell_{\text{new}}$. Thus, every term of $\sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}$ is strictly positive.

Now, rearranging the desired result gives

$$\begin{aligned}
\sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2} &\geq \left(\frac{1}{\delta_\ell} + \Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)\right)\left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}\right) \\
&= \left(\frac{1}{\delta_\ell} + \delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}\right)\left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}\right) \\
&= \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1} + \left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}\right)^2
\end{aligned}$$

Moving the first term of the right-hand side to the left-hand side, we equivalently have

$$\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}(\lambda_i - \ell_{\text{old}})^{-1} \geq \left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}\right)^2$$

Thus, it suffices to prove this identity. Now, let $x, y \in \mathbb{R}^n$ be the vectors with entries $x_i = (\lambda_i - \ell_{\text{old}})^{-1/2}$ and $y_i = (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1/2}$ respectively. Then, the right hand side becomes $\delta_\ell^2 |\langle x, y\rangle|^2$. Applying the Cauchy-Schwarz Inequality and the fact that $\frac{1}{\delta_\ell} - \epsilon_\ell \geq 0 \implies 1 \geq \delta_\ell \epsilon_\ell$, we have

$$\begin{aligned}
\delta_\ell^2 |\langle x, y\rangle|^2 \leq \delta_\ell^2 \|x\|_2^2 \left\|y\right\|_2^2 &= \delta_\ell \left(\sum_{i=1}^n (\lambda_i - \ell_{\text{old}})^{-1}\right)\left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}(\lambda_i - \ell_{\text{old}})^{-1}\right) \\
&= \delta_\ell \Phi_{\ell_{\text{old}}}(A)\left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}(\lambda_i - \ell_{\text{old}})^{-1}\right) \\
&\leq (\delta_\ell \epsilon_\ell)\left(\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}(\lambda_i - \ell_{\text{old}})^{-1}\right) \\
&\leq \delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}(\lambda_i - \ell_{\text{old}})^{-1}
\end{aligned}$$

This is the desired bound and the lemma is proved. □

Lemma 13 (Shifting Both Barriers; Lemma 3.5 from [27])**.** *Let $A \in \mathbb{R}^{n\times n}$ be symmetric. Let $u_{\text{old}}, \ell_{\text{old}} > 0$ satisfy $\lambda_{\max}(A) < u_{\text{old}}$ and $\lambda_{\min}(A) > \ell_{\text{old}}$. Let $\epsilon_u, \epsilon_\ell, \delta_u, \delta_\ell > 0$ satisfy $\Phi^{u_{\text{old}}}(A) \leq \epsilon_u$, $\Phi_{\ell_{\text{old}}}(A) \leq \epsilon_\ell$, and*

$$0 \leq \frac{1}{\delta_u} + \epsilon_u \leq \frac{1}{\delta_\ell} - \epsilon_\ell$$

Set $u_{\text{new}} = u_{\text{old}} + \delta_u$ and $\ell_{\text{new}} = \ell_{\text{old}} + \delta_\ell$. Then, there exists an $1 \leq i \leq m$ and a $s > 0$ such that

1. $U_A(v_i) \leq 1/s \leq L_A(v_i)$
2. $\lambda_{\max}(A + sv_iv_i^\top) < u_{\text{new}}$
3. $\lambda_{\min}(A + sv_iv_i^\top) > \ell_{\text{new}}$
4. $\Phi^{u_{\text{new}}}(A + sv_iv_i^\top) \leq \epsilon_u$
5. $\Phi_{\ell_{\text{new}}}(A + sv_iv_i^\top) \leq \epsilon_\ell$

where $U_A(v), L_A(v)$ are defined as in Lemma 10 and Lemma 11.

Proof. We will first show that $\sum_{i=1}^m L_A(v_i) \geq \sum_{i=1}^m U_A(v_i)$. First, we write out the right-hand side and apply Corollary 2 with the fact that $\sum_{i=1}^m v_i v_i^\top = I$.

$$\begin{aligned}
\sum_{i=1}^m U_A(v_i) &= \frac{\sum_{i=1}^m v_i^\top (u_{\text{new}} I - A)^{-2} v_i}{\Phi^{u_{\text{old}}}(A) - \Phi^{u_{\text{new}}}(A)} + \sum_{i=1}^m v_i^\top (u_{\text{new}} I - A)^{-1} v_i \\
&= \frac{\text{tr}\left[(u_{\text{new}} I - A)^{-2}\right]}{\Phi^{u_{\text{old}}}(A) - \Phi^{u_{\text{new}}}(A)} + \text{tr}\left[(u_{\text{new}} I - A)^{-1}\right] \\
&= \frac{\sum_{i=1}^n (u_{\text{new}} - \lambda_i)^2}{\sum_{i=1}^n \left((u_{\text{old}} - \lambda_i)^{-1} - (u_{\text{new}} - \lambda_i)^{-1}\right)} + \Phi^{u_{\text{new}}}(A) \\
&= \frac{\sum_{i=1}^n (u_{\text{new}} - \lambda_i)^2}{\delta_u \sum_{i=1}^n (u_{\text{old}} - \lambda_i)^{-1}(u_{\text{new}} - \lambda_i)^{-1}} + \Phi^{u_{\text{new}}}(A)
\end{aligned}$$

Since $u_{\text{new}} > u_{\text{old}}$, we have the denominator of the first term is larger than the numerator of the first term. Hence, we can upper bound the above as

$$\frac{1}{\delta_u} + \Phi^{u_{\text{new}}}(A) \leq \frac{1}{\delta_u} + \Phi^{u_{\text{old}}}(A) \leq \frac{1}{\delta_u} + \epsilon_u$$

Now, we write out $\sum_{i=1}^m L_A(v_i)$. Using almost exactly the same manipulations, we have

$$\begin{aligned}
\sum_{i=1}^m L_A(v_i) &= \frac{\sum_{i=1}^m v_i^\top (A - \ell_{\text{new}} I)^{-2} v_i}{\Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)} - \sum_{i=1}^m v_i^\top (A - \ell_{\text{new}} I)^{-1} v_i \\
&= \frac{\text{tr}\left[(A - \ell_{\text{new}} I)^{-2}\right]}{\Phi_{\ell_{\text{new}}}(A) - \Phi_{\ell_{\text{old}}}(A)} - \text{tr}\left[(A - \ell_{\text{new}} I)^{-1}\right] \\
&= \frac{\sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}}{\sum_{i=1}^n \left((\lambda_i - \ell_{\text{new}})^{-1} - (\lambda_i - \ell_{\text{old}})^{-1}\right)} - \Phi_{\ell_{\text{new}}}(A) \\
&= \frac{\sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-2}}{\delta_\ell \sum_{i=1}^n (\lambda_i - \ell_{\text{new}})^{-1}(\lambda_i - \ell_{\text{old}})^{-1}} - \Phi_{\ell_{\text{new}}}(A)
\end{aligned}$$

Note, we cannot use a similar argument as when we proved an upper bound on $\sum_{i=1}^m U_A(v_i)$ to prove a lower bound for $\sum_{i=1}^m L_A(v_i)$, since now, both the lower barrier as well as the eigenvalues of A have shifted upwards. The denominator of the first term above, after removing δ_ℓ, is still larger than the numerator, so we cannot lower bound the above expression by $\frac{1}{\delta_\ell} - \Phi_{\ell_{\text{new}}}(A)$. Fortunately, by Lemma 12, we can still lower bound it by

$$\frac{1}{\delta_\ell} - \Phi_{\ell_{\text{old}}}(A) \geq \frac{1}{\delta_\ell} - \epsilon_\ell$$

Thus, we have the following desired chain of inequalities:

$$\sum_{i=1}^m U_A(v_i) \leq \frac{1}{\delta_u} + \epsilon_u \leq \frac{1}{\delta_\ell} - \epsilon_\ell \leq \sum_{i=1}^m L_A(v_i)$$

This shows that there exists $1 \leq i \leq m$ such that $U_A(v_i) \leq L_A(v_i)$, since otherwise, we'd have $\sum_{i=1}^m U_A(v_i) > \sum_{i=1}^m L_A(v_i)$, a contradiction. Since $L_A(v_i) \geq U_A(v_i) > 0$, we also have that there exists $s > 0$ such that $U_A(v_i) \leq 1/s \leq L_A(v_i)$ (for example, one can use continuity of $1/x$ with the Intermediate Value Theorem). Furthermore, for this choice of $1 \leq i \leq m$ and $s > 0$, we have:

1. $\lambda_{\max}(A + s v_i v_i^\top) < u_{\text{new}}$ and $\Phi^{u_{\text{new}}}(A + s v_i v_i^\top) \leq \Phi^{u_{\text{old}}}(A) \leq \epsilon_u$ by Lemma 10, since $U_A(v_i) \leq 1/s$ and $\lambda_{\max}(A) < u_{\text{old}}$

2. $\lambda_{\min}(A + sv_iv_i^\top) > \ell_{\text{new}}$ and $\Phi_{\ell_{\text{new}}}(A + sv_iv_i^\top) \le \Phi_{\ell_{\text{old}}}(A) \le \epsilon_\ell$ by Lemma 11, since $L_A(v_i) \ge 1/s$ and $\lambda_{\min}(A) > \ell_{\text{old}}$

This completes the proof. □

With the ability to shift both barriers and reason about the resulting potential function values, eigenvalues, and barrier values, we are ready to prove the main theorem.

Proof of Theorem 16. We will inductively build a sequence of matrices $A^{(t)}$ given by $\sum_{i\in S_t} s_i v_i v_i^\top$, where S_t is the state of S after round t; we will also keep track of upper and lower barrier values u_t, ℓ_t. Specifically, we will prove that for each of $0 \le t < dn$ steps:

(1) $\lambda_{\max}(A^{(t)}) < u_t$

(2) $\lambda_{\min}(A^{(t)}) > \ell_t$

(3) $\Phi^{u_t}(A^{(t)}) \le \epsilon_u$

(4) $\Phi_{\ell_t}(A^{(t)}) \le \epsilon_\ell$

We still need to choose our step size constants $\delta_u, \delta_\ell > 0$, our potential upper bounds $\epsilon_u, \epsilon_\ell > 0$ and our initial barrier values u_0, ℓ_0 so that the requirements of Lemma 13 are satisfied. For this, we will take

$$\delta_\ell = 1 \qquad \delta_u = \frac{\sqrt{d}+1}{\sqrt{d}-1} \qquad \epsilon_\ell = \frac{1}{\sqrt{d}} \qquad \epsilon_u = \frac{\sqrt{d}-1}{d+\sqrt{d}} \qquad \ell_0 = \frac{-n}{\epsilon_\ell} \qquad u_0 = \frac{n}{\epsilon_u}$$

It is easy to check that

$$\frac{1}{\delta_u} + \epsilon_u = \frac{1}{\delta_\ell} - \epsilon_\ell$$

Initially, we have $\lambda_{\max}(A^{(0)}) = 0 < u_0$, $\lambda_{\min}(A^{(0)}) = 0 > \ell_0$, and

$$\Phi^{u_0}(A^{(0)}) = \operatorname{tr}\left[(u_0 I)^{-1}\right] = \frac{\operatorname{tr}[I]}{u_0} = \epsilon_u$$
$$\Phi_{\ell_0}(A^{(0)}) = \operatorname{tr}\left[(-\ell_0 I)^{-1}\right] = \frac{\operatorname{tr}[I]}{-\ell_0} = \epsilon_\ell$$

If we assume (1), (2), (3), and (4) all hold up to some $t \ge 0$, then by Lemma 13, there exists some $1 \le i \le m$ and $s_i > 0$ such that $A^{(t+1)} = A^{(t)} + s_i v_i v_i^\top$ satisfy (1), (2), (3), and (4). This completes the induction.

Since we run at most dn iterations of this process, we obtain a set $S \subset [m]$ of size at most dn. It remains to reason about the eigenvalues of $A^{(dn)} = \sum_{i\in S} s_i v_i v_i^\top$. For this, using (1) and (2), we have

$$d + 1 - 2\sqrt{d} = \ell_0 + dn\delta_\ell = \ell_{dn} < \lambda_{\min}(A^{(dn)}) \le \lambda_{\max}(A^{(dn)}) < u_{dn} = u_0 + dn\delta_u = d + 1 + 2\sqrt{d}$$

and so

$$(d + 1 - 2\sqrt{d}) \cdot I \le \sum_{i\in S} s_i v_i v_i^\top \le (d + 1 + 2\sqrt{d}) \cdot I$$

Dividing by $d + 1 - 2\sqrt{d}$ and replacing each nonzero s_i with $s_i/(d + 1 - 2\sqrt{d})$ proves the theorem. □

5.5 An Algorithm

Just as in restricted invertibility, the barrier argument naturally lends itself to an algorithm for constructing spectral sparsifiers. The algorithm has the advantage of being fully deterministic, and also improves upon the sparsity of the subgraph produced. The only down-side is that it is slow, taking $\mathcal{O}(dn^3m)$-time to compute a $(d + 1 + 2\sqrt{d})/(d + 1 - 2\sqrt{d})$-spectral sparsifier. To produce a $(1 \pm \epsilon)$-spectral sparsifier, one needs to set $d \approx 1/\epsilon^2$. Hence, this algorithm requires $\mathcal{O}(n^3m/\epsilon^2)$ time to compute a $(1 \pm \epsilon)$-spectral sparsifier.

Algorithm 2

Input: weighted graph $G = (V, E, w)$; $d > 1$

Output: a $(d+1+2\sqrt{d})/(d+1-2\sqrt{d})$-spectral sparsifier $H = (V, \tilde{E}, \tilde{w})$ with $\tilde{E} \subset E$

set $\delta_\ell = 1$, $\delta_u = (\sqrt{d}+1)/(\sqrt{d}-1)$, $\epsilon_\ell = 1/\sqrt{d}$, $\epsilon_u = (\sqrt{d}-1)/(d+\sqrt{d})$, $\ell_0 = -n/\epsilon_\ell$ and $u_0 = n/\epsilon_u$
initialize $A = 0$
initialize sparsifier H with vertex set V, $\tilde{E} = \emptyset$, $\tilde{w}(e) = 0$ for all $e \in E$
compute $B \in \mathbb{R}^{m \times n}$, $W \in \mathbb{R}^{m \times m}$ and $L_G = B^\top W B$ from G
compute eigendecomposition of L_G and then compute $(L_G^\dagger)^{1/2}$
compute $V = (L_G^\dagger)^{1/2} B^\top W^{1/2}$; let $v_1, \ldots, v_m$ be the columns of V
for round $t = 0, \ldots, dn-1$ **do**
    set $u_{t+1} = u_t + \delta_u$, $\ell_{t+1} = \ell_t + \delta_\ell$
    compute $(u_t I - A)^{-1}$, $(A - \ell_t I)^{-1}$, $(u_{t+1} I - A)^{-1}$, $(A - \ell_{t+1} I)^{-1}$, $(u_{t+1} I - A)^{-2}$, $(A - \ell_{t+1} I)^{-2}$
    take traces to obtain $\Phi^{u_{t+1}}(A)$, $\Phi^{u_t}(A)$, $\Phi_{\ell_{t+1}}(A)$, $\Phi_{\ell_t}(A)$
    for $i = 1, \ldots, m$ **do**
        compute $U_A(v_i)$, $L_A(v_i)$
        if $U_A(v_i) \le L_A(v_i)$ **then**
            choose $s_i > 0$ such that $U_A(v_i) \le 1/s_i \le L_A(v_i)$ (for example, one can take $s_i = 1/L_A(v_i)$)
            update $A \leftarrow A + s_i v_i v_i^\top$
            $\tilde{E} \leftarrow \tilde{E} \cup \{i\}$
            set $\tilde{w}(i) = w(i) s_i$
            break out of inner for-loop
        end if
    end for
end for
downweight all weights $\tilde{w}$ by a multiplicative factor of $1/(d+1-2\sqrt{d})$
return $H = (V, \tilde{E}, \tilde{w})$

6 Interlacing Families

Now, let's return to interlacing polynomials. In the previous two sections, we mainly just needed the Cauchy Interlacing Theorem so that we could reason about the eigenvalues of matrices after applying rank-one perturbations. In the following applications, we'll need these earlier results and more.

Suppose we have a "large" set of polynomials. It may be too restrictive to demand that this set has a common interlacing. However, we still want to obtain the same guarantees as in Theorem 6, i.e. ensure there is a polynomial in the set that has its maximum root upper bounded by the maximum root of the average.

If we know that certain sets of sums of these polynomials have common interlacings, it turns out we can still make this guarantee; the trick is to inductively apply Theorem 6 on the "layers" of sums of these polynomials. The intuition is that each "layer" progressively reveals more and more information about the largest root of the total sum. We formalize this as follows.

Definition 6 (Interlacing Families)**.** Let $S_1, \ldots, S_m$ be finite sets. For each $(s_1, \ldots, s_m) \in \prod_{i=1}^m S_i$, associate a real-rooted degree-n polynomial $p_{s_1,\ldots,s_m}$ with positive leading coefficient. For every $1 \le k < m$ and partial assignment in the first k coordinates $(s_1, \ldots, s_k) \in \prod_{i\in[k]} S_i$, define

$$p_{s_1,\ldots,s_k} = \sum_{(s_{k+1},\ldots,s_m)\in S_{k+1}\times\cdots\times S_m} p_{s_1,\ldots,s_m}$$

For $k = 0$, define

$$p_\emptyset = \sum_{(s_1,\ldots,s_m)\in S_1\times\cdots\times S_m} p_{s_1,\ldots,s_m}$$

The polynomials $\{p_{s_1,\ldots,s_m}\}$ form an **interlacing family** if for all $k = 0, \ldots, m-1$ and all $(s_1, \ldots, s_k) \in \prod_{i\in[k]} S_i$, the polynomials $\{p_{s_1,\ldots,s_k,t}\}_{t\in S_{k+1}}$ have a common interlacing.

Here, "layer k" is $\{p_{s_1,\ldots,s_k}\}$. By increasing k incrementally and fixing more and more elements in the partial assignment $(s_1, \ldots, s_k)$, we get closer and closer to a/the desired polynomial in the original family $\{p_{s_1,\ldots,s_m}\}$. With this in mind, we see the definition lends itself naturally to using induction.

Theorem 17. *Let $S_1, \ldots, S_m$ be finite sets and let $\{p_{s_1,\ldots,s_m}\}_{(s_1,\ldots,s_m)\in\prod_{i\in[m]} S_i}$ be an interlacing family of polynomials. Then, for every $k \in [m]$, there exists $(s_1, \ldots, s_k) \in \prod_{i\in[k]} S_i$ such that the largest root of $p_{s_1,\ldots,s_k}$ is upper bounded by the largest root of $p_\emptyset$.*

Proof. For the induction basis, observe that

$$p_\emptyset = \sum_{(s_1,\ldots,s_m)\in S_1\times\cdots\times S_m} p_{s_1,\ldots,s_m} = \sum_{s_1\in S_1} \sum_{(s_2,\ldots,s_m)\in S_2\times\cdots\times S_m} p_{s_1,\ldots,s_m} = \sum_{t\in S_1} p_t$$

Since $\{p_t\}_{t\in S_1}$ have a common interlacing, there exists s_1 such that the largest root of p_{s_1} is at most the largest root of $p_\emptyset$ by Theorem 6.

Now, assume that there exists $(s_1, \ldots, s_k) \in S_1 \times \cdots \times S_k$ such that the largest root of $p_{s_1,\ldots,s_k}$ is at most the largest root of $p_\emptyset$, for some $k \ge 1$. The following is a similar observation to the one made in the base case.

$$p_{s_1,\ldots,s_k} = \sum_{(s_{k+1},\ldots,s_m)\in S_{k+1}\times\cdots\times S_m} p_{s_1,\ldots,s_m} = \sum_{s_{k+1}\in S_{k+1}} \sum_{(s_{k+2},\ldots,s_m)\in S_{k+2}\times\cdots\times S_m} p_{s_1,\ldots,s_m} = \sum_{t\in S_{k+1}} p_{s_1,\ldots,s_k,t}$$

Since $\{p_{s_1,\ldots,s_k,t}\}_{t\in S_{k+1}}$ has a common interlacing, there exists s_{k+1} such that the largest root of $p_{s_1,\ldots,s_{k+1}}$ is at most the largest root of $p_{s_1,\ldots,s_k}$ by Lemma 4, which is at most the largest root of $p_\emptyset$ by the inductive hypothesis. □

As a special case, we have the following corollary, which extends the applicability of interlacing.

Corollary 6. *For an interlacing family $\mathcal{P} = \{p_{s_1,\dots,s_m}\}_{(s_1,\dots,s_m)\in\prod_{i\in[m]} S_i}$, there exists $p \in \mathcal{P}$ such that the largest root of p is at most the largest root of $p_\emptyset = \sum_{q\in\mathcal{P}} q$.*

The real power of interlacing comes when we incorporate a probabilistic perspective.

Corollary 7. *Let $\mathcal{D}_1,\dots,\mathcal{D}_m$ be independent probability distributions with $\mathcal{D}_i$ supported on S_i, for all $1 \le i \le m$. Let $\mathcal{P} = \{p_{s_1,\dots,s_m}\}$ be an interlacing family. Then, there exists $(s_1,\dots,s_m) \in S_1 \times \cdots \times S_m$ such that*

$$\operatorname{maxroot}(p_{s_1,\dots,s_m}) \le \operatorname{maxroot}\left(\mathbb{E}_{(s_1,\dots,s_m)\sim\mathcal{D}_1\times\cdots\times\mathcal{D}_m}[p_{s_1,\dots,s_m}]\right)$$

where $\mathcal{D}_1 \times \cdots \times \mathcal{D}_m$ is the product distribution on $S_1 \times \cdots \times S_m$.

Proof. Let $p_{i,s_i} = \mathbb{P}_{s\sim\mathcal{D}_i}[s = s_i]$ be the probability of seeing s_i under $\mathcal{D}_i$. Since the $\mathcal{D}_1,\dots,\mathcal{D}_m$ are independent,

$$\mathbb{P}_{s\sim\mathcal{D}_1\times\cdots\times\mathcal{D}_m}[s = (s_1,\dots,s_m)] = \prod_{i=1}^m p_{i,s_i}$$

Now, since for any nonzero scalar $c \in \mathbb{R}$ and univariate polynomial $p \in \mathbb{R}[x]$, the roots of $c \cdot p$ are precisely the roots of p. In particular,

$$\left\{\left(\prod_{i=1}^m p_{i,s_i}\right) p_{s_1,\dots,s_m}\right\}$$

forms an interlacing family. By Corollary 6, there exists $(s_1,\dots,s_m) \in S_1 \times \cdots \times S_m$ such that

$$\begin{aligned}\operatorname{maxroot}(p_{s_1,\dots,s_m}) &= \operatorname{maxroot}\left(\left(\prod_{i=1}^m p_{i,s_i}\right) p_{s_1,\dots,s_m}\right) \le \operatorname{maxroot}\left(\sum_{\tilde{s_1}\in S_1,\dots,\tilde{s_m}\in S_m}\left(\prod_{i=1}^m p_{i,\tilde{s_i}}\right) p_{\tilde{s_1},\dots,\tilde{s_m}}\right)\\ &= \operatorname{maxroot}\left(\mathbb{E}_{(s_1,\dots,s_m)\sim\mathcal{D}_1\times\cdots\times\mathcal{D}_m}[p_{s_1,\dots,s_m}]\right)\end{aligned}$$

□

The interlacing families proof technique generally follows the following structure:

1. show that the collection of polynomials of interest forms an interlacing family
2. prove a bound on the maximum root of their average (over some probability distribution of interest)

In view of Corollary 7, one can see that this proof technique has the flavor of a probabilistic method. Recall that if f is a function (random variable) from a probability space $\Omega = (X, \mu)$ to, say, $\mathbb{R}$, then $\mathbb{P}_{x\sim\mu}[f(x) \le b]$ shows that there exists $x \in X$ with $f(x) \le b$. One can view that the given distributions $\mathcal{D}_1,\dots,\mathcal{D}_m$ form a probability space over $X = \{p_{s_1,\dots,s_m}\}$, and f is the maxroot function.

6.1 An Algorithmic Perspective

One can visualize this process with a tree, with "layer k" i.e. $\{p_{s_1,\dots,s_k}\}$ is represented in the kth level of the tree. If S_i has values $\{s_i^1,\dots,s_i^{\ell_i}\}$ for all $i \in [m]$, then

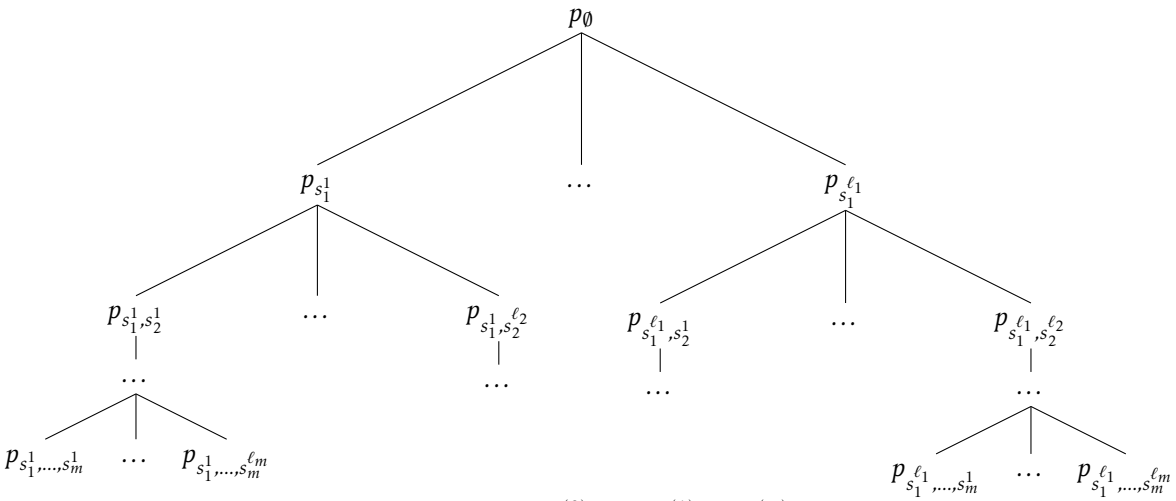

The Theorem simply allows us to find a path $p^{(0)} = p_\emptyset, p^{(1)}, \ldots, p^{(m)} = p_{s_1,\ldots,s_m}$ from the root $p_\emptyset$ to some leaf polynomial $p_{s_1,\ldots,s_m}$ in our interlacing family such that the largest root of $p^{(i)}$ is at most the largest root of $p^{(i-1)}$ for all $i \in [m]$. $p^{(m)} = p_{s_1,\ldots,s_m}$ is then the polynomial we are looking for.

This visualization also suggests an algorithm, namely, "walk down the tree" starting at the root, moving to a child with smaller maximum root, and terminating when a leaf node is hit.

Algorithm 3 Interlacing-Families-Algorithm-Outline

Input: an interlacing family $\{p_{s_1,\ldots,s_m}\}$

compute (maximum root of) $p_\emptyset$
for $i = 1, \ldots, m$ **do**
  for $s = s_i^1, \ldots, s_i^{\ell_i}$ **do**
    compute (maximum root of) $p_{s_1,\ldots,s_{i-1},s}$
    if $\text{maxroot}(p_{s_1,\ldots,s_{i-1},s}) \leq \text{maxroot}(p_{s_1,\ldots,s_{i-1}})$ **then**
      set $p_{s_1,\ldots,s_i} \leftarrow p_{s_1,\ldots,s_{i-1},s}$
    end if
  end for
end for
return $p_{s_1,\ldots,s_m}$

The challenge here is efficiency. In general, one needs:

1. polynomial depth of the interlacing family tree
2. polynomially many children for each node in the interlacing family tree
3. to be able to compute (the maximum root of) $p_\emptyset$ efficiently
4. to be able to compute (the maximum root of) each child polynomial $p_{s_1,\ldots,s_{i+1}}$ of $p_{s_1,\ldots,s_i}$ efficiently

For many interlacing families of interest (including all considered in this survey, with a single exception), we do not expect to be able to run this "natural" algorithm in polynomial time. Just computing $p_\emptyset$ can be #**P**-Hard, as we will see.

7 Stability

To prove that a family of polynomials $\{p_{s_1,\dots,s_m}\}$ is an interlacing family, one needs to prove the existence of a common interlacing for $\{p_{s_1,\dots,s_k,t}\}_{t\in S_{k+1}}$ for all k and all $(s_1,\dots,s_k)\in S_1\times\cdots\times S_k$. This is where Theorem 7 will be particularly useful. To prove that existence of a common interlacing for $\{p_{s_1,\dots,s_k,t}\}_{t\in S_{k+1}}$, it suffices to show that all convex combinations

$$\sum_{t\in S_{k+1}}\lambda_t p_{s_1,\dots,s_k,t}$$

are real-rooted. For this, we can then employ the power of real stability, a generalization of real-rootedness; real stable polynomials are very well-studied.

Let $\mathbb{H}=\{z\in\mathbb{C}:\mathrm{Im}(z)>0\}$.

Definition 7. A nonzero polynomial $p\in\mathbb{C}[z_1,\dots,z_n]$ is **stable** if $\mathbf{z}\in\mathbb{H}^n \implies p(\mathbf{z})\neq 0$. A stable polynomial is **real stable** if it has real coefficients.

Observe that real stability is indeed a generalization of real-rootedness. Indeed, if p is a real univariate polynomial, then, since complex roots come in conjugate pairs, p is real-rooted if and only if p is real stable. Now, let's see an equivalent characterization of (real) stability that, while we will not use directly, can be useful for intuition.

Lemma 14. *A polynomial $p(\mathbf{z})$ is stable if and only if the univariate restriction $p(t\mathbf{e}+\mathbf{x})$ is stable for all $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$.*

Proof. Suppose p is stable. Suppose for contradiction that there exists $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$ such that $p(t\mathbf{e}+\mathbf{x})$ is not stable; choose $t_0\in\mathbb{H}$ such that $p(t_0\mathbf{e}+\mathbf{x})=0$. Then, since $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$ and $\mathrm{Im}\,t_0>0$, all entries of $t_0\mathbf{e}+\mathbf{x}$ have positive imaginary part, i.e. $t\mathbf{e}+\mathbf{x}\in\mathbb{H}^n$, contradicting p being stable.

Now suppose p is not stable. Choose $\mathbf{z}_0\in\mathbb{H}^n$ such that $p(\mathbf{z}_0)\neq 0$. Decompose $\mathbf{z}_0$ into $t_0\mathbf{e}_0+\mathbf{x}_0$, where $\mathbf{x}_0=\mathrm{Re}\,\mathbf{z}_0=(\mathrm{Re}(z_0)_1,\dots,\mathrm{Re}(z_0)_n)\in\mathbb{R}^n$, $\mathbf{e}=\mathrm{Im}\,\mathbf{z}_0=(\mathrm{Im}(z_0)_1,\dots,\mathrm{Im}(z_0)_n)$ and $t_0=i$. Since $\mathbf{z}_0$ is a root of p, $t_0=i$ is a root of $p(t\mathbf{e}+\mathbf{x})$. Since $t_0\in\mathbb{H}$, there exists $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$ such that $p(t\mathbf{e}+\mathbf{x})$ is not stable. □

Corollary 8. *A polynomial $p(\mathbf{z})$ is real stable if and only if the univariate restriction $p(t\mathbf{e}+\mathbf{x})$ is real-rooted for all $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$.*

Proof. Observe that p has real coefficients if and only if $p(t\mathbf{e}+\mathbf{x})$ has real coefficients for all $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$ (one direction is obvious; the other direction follows from if $c\prod_{i\in S}z_i^{d_i}$ is some monomial in p with $c\notin\mathbb{R}$ and $S\subset[n]$, then $c\prod_{i\in S}(te_i+z_i)^{d_i}$ is a monomial in $p(t\mathbf{e}+\mathbf{x})$ so that for $\mathbf{e}=\mathbf{1}_S+\epsilon\mathbf{1}_{[n]\setminus S},\mathbf{x}=\mathbf{0}$, we get $p(t\mathbf{e}+\mathbf{x})=c\prod_{i\in S\subset[n]}(te_i+z_i)^{d_i}+\epsilon s(t)$, which has a nonreal coefficient of $\approx c$ for sufficiently small $\epsilon>0$). Hence, p is real stable if and only if $p(t\mathbf{e}+\mathbf{x})$ is real stable for all $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$ if and only if $p(t\mathbf{e}+\mathbf{x})$ is real-rooted for all $\mathbf{e}\in\mathbb{R}^n_{>0},\mathbf{x}\in\mathbb{R}^n$. □

The main reason stability is such a nice property for polynomials to have is that there is a large class of useful operators under which this property is preserved.

Theorem 18 (Stability-Preserving Operators)**.** *Let $p\in\mathbb{C}[z_1,\dots,z_n]$ be a stable polynomial. Then the following are also all stable.*

1. *$p\cdot q$, where $m\in\mathbb{N}$ and $q\in\mathbb{C}[z_1,\dots,z_m]$ is also stable (**Closure Under Products**)*
2. *$p(z_{\sigma(1)},\dots,z_{\sigma(n)})$ for any permutation $\sigma:[n]\to[n]$ (**Coordinate Permutation**)*
3. *$p(\dots,az_i,\dots)$ for any $a>0$ and $i\in[n]$ (**Positive Coordinate Scaling**)*
4. *$p(\dots,z_i,\dots,z_i,\dots)\in\mathbb{C}[\dots,z_i,\dots,z_i,\dots]$ (**Diagonalization: restricting some z_j to be equal to some other z_i for any $i,j\in[n]$**)*

5. $z_i^d p(\dots, -1/z_i, \dots)$ *where* $d = \deg_i(p)$ *is the degree of* z_i *in* p *viewing other variables as constants, for any* $i \in [n]$ (***Coordinate Inversion***)
6. $p(\dots, a, \dots)$ *for any* $a \in \overline{\mathbb{H}} = \mathbb{H} \cup \mathbb{R}$ *and* $i \in [n]$ (***Specialization: fixing a coordinate***)
7. $\partial_{z_i} p$, *for any* $1 \le i \le n$ (***Differentiation***)
8. $(1 - \partial_{z_i})p$, *for any* $1 \le i \le n$

7.1 Examples

The closure properties of stable polynomials given in Theorem 18 already give us many tools. Let's start with an example of how to prove a univariate polynomial is real-rooted (recall, this was the original motivation for bringing in this theory). The idea is to start with a (usually multivariate) polynomial that we already know to be real stable, and apply a sequence of stability-preserving operators that also reduce this polynomial to the original polynomial of interest.

Example 2. *Claim:* For any $k \in [n]$ and any $\lambda_1, \dots, \lambda_n \in \mathbb{R}$, the polynomial $p(t) = \sum_{S \in \binom{[n]}{k}} \prod_{i \in S} (1 + t\lambda_i)$ is real-rooted.

Proof. Consider the following bivariate polynomial.

$$p(y,t) = \prod_{i=1}^{n} (y + (t - \lambda_i)) = \sum_{j=0}^{n} y^{n-j} \left(\sum_{S \in \binom{[n]}{j}} \prod_{i \in S} (t - \lambda_i) \right)$$

$p(y,t)$ clearly has real coefficients and is stable, for if $(y,t) \in \mathbb{H}^2$, then $\mathrm{Im}(y + t - \lambda_i) > 0$ (since λ_i is real and $\mathrm{Im}(y), \mathrm{Im}(t) > 0$) for all $i \in [n]$, which shows $p(y,t) \neq 0$. Differentiating $n - k$ times with respect to y, we obtain

$$\frac{\partial^{n-k}}{\partial y^{n-k}} p(y,t) = \sum_{j=0}^{k} \frac{(n-j)!}{(k-j)!} y^{k-j} \left(\sum_{S \in \binom{[n]}{j}} \prod_{i \in S} (t - \lambda_i) \right)$$

Specializing to $y = 0$, the $j = k$ term in the sum is left behind

$$\left[\frac{\partial^{n-k}}{\partial y^{n-k}} p(y,t) \right]_{y=0} = (n-k)! \cdot \sum_{S \in \binom{[n]}{k}} \prod_{i \in S} (t - \lambda_i)$$

By Theorem 18, we've only applied operators that preserve real stability, so $q(t) = \sum_{S \in \binom{[n]}{k}} \prod_{i \in S} (t - \lambda_i)$ is real stable; since it is a polynomial in a single variable, it is real-rooted. Finally, $p(t) = (-t)^k q(-1/t)$ proves p is real-rooted. □

The following is an important example. It is a large class of polynomials that are real stable.

Theorem 19 (Proposition 2.4 from [24])**.** *Let* $A_1, \dots, A_m$ *be Hermitian positive semidefinite matrices. Then the polynomial*

$$\det \left(\sum_{i=1}^{m} z_i A_i \right)$$

is real stable.

Finally, in the special case of two variables, it turns out we can completely characterize all real stable polynomials.

Theorem 20 (Corollary 6.7 from [29])**.** *Let* $p(z_1, z_2)$ *be a real stable polynomial of degree* d*. Then, there exist positive semidefinite matrices* $A, B \in \mathbb{R}^{d \times d}$*, with* $A + B$ *positive definite, and a symmetric matrix* $C \in \mathbb{R}^{d \times d}$ *such that either* $p(z_1, z_2) = \det(z_1 A + z_2 B + C)$ *or* $p(z_1, z_2) = -\det(z_1 A + z_2 B + C)$.

8 The Kadison-Singer Problem

Problem 3. *Does every pure state on the (abelian) von Neumann algebra $\mathbb{D}$ of bounded diagonal operators on ℓ_2 have a unique extension to a pure state on $B(\ell_2)$, the von Neumann algebra of all bounded operators on ℓ_2?*

Setting up the technical machinery required to understand what exactly the problem statement is is far beyond the scope of this survey. Fortunately, through a series of many works between 1959, the time the problem was first posed by Kadison-Singer [1], and the early 2000s, it was shown that a positive answer to the problem is equivalent to several other conjectures. We focus on a particular equivalent conjecture, stated by Weaver in 2004 [17], and resolve it.

Conjecture 1 (KS_2; [17])**.** *There exist universal constants $\alpha \geq 2$ and $\beta > 0$ such that if both*

1. *$w_1, \ldots, w_m \in \mathbb{C}^d$ satisfy $\|w_i\| \leq 1$ for all i*
2. *$\sum_{i=1}^m |\langle u, w_i \rangle|^2 = \alpha$ for every unit vector $u \in \mathbb{C}^d$*

then there exists a partition S_1, S_2 of $[m]$ such that

$$\sum_{i \in S_j} |\langle u, w_i \rangle|^2 \leq \alpha - \beta$$

for every unit vector $u \in \mathbb{C}^d$, for every $j \in \{1, 2\}$.

The main technical theorem is a purely probabilistic result that we will use the method of interlacing polynomials to prove.

Theorem 21 (Theorem 1.4 from [35])**.** *Let $\epsilon > 0$ and $\mathcal{D}_1, \ldots, \mathcal{D}_m$ be independent distributions each supported on finitely many vectors in $\mathbb{C}^d$. Let $v_1 \sim \mathcal{D}_1, \ldots, v_m \sim \mathcal{D}_m$ be random vectors drawn independently such that*

$$\sum_{i=1}^m \mathbb{E}[v_i v_i^*] = I$$

and $\mathbb{E}[\|v_i\|^2] \leq \epsilon$ for all $1 \leq i \leq m$. Then

$$\mathbb{P}\left[\left\|\sum_{i=1}^m v_i v_i^*\right\| \leq (1 + \sqrt{\epsilon})^2\right] > 0$$

This theorem implies a kind of generalization of KS_2, where instead of considering a partition of size 2, we consider a partition of size $r \geq 1$.

Corollary 9. *Let $r \in \mathbb{N}$ and $u_1, \ldots, u_m \in \mathbb{C}^d$ satisfy $\sum_{i=1}^m u_i u_i^* = I$ and $\|u_i\|^2 \leq \delta$ for all $1 \leq i \leq m$. Then, there exists a partition $S_1, \ldots, S_r$ of m such that*

$$\left\|\sum_{i \in S_j} u_i u_i^*\right\| \leq \left(\frac{1}{\sqrt{r}} + \sqrt{\delta}\right)^2$$

for all $1 \leq j \leq r$.

Proof. For every $i \in [m]$ and $k \in [r]$, define $w_{i,k} = \mathbf{0} \oplus \cdots \oplus u_i \oplus \cdots \oplus \mathbf{0} \in \mathbb{C}^{rd}$, where u_i appears in the kth position in the direct sum; $w_{i,k}$ will have $r(k-1)$ zero entries, followed by the entries of u_i, and then followed by zeros. For intuition,

$$w_{i,1} = \begin{pmatrix} u_i \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{pmatrix} \qquad w_{i,2} = \begin{pmatrix} \mathbf{0} \\ u_i \\ \vdots \\ \mathbf{0} \end{pmatrix} \qquad \cdots \qquad w_{i,r-1} = \begin{pmatrix} \mathbf{0} \\ \vdots \\ u_i \\ \mathbf{0} \end{pmatrix} \qquad w_{i,r} = \begin{pmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ u_i \end{pmatrix}$$

Let $\mathcal{D}_i$ be a uniform distribution supported on $\{\sqrt{r}w_{i,k}\}_{k=1}^r$, for all $1 \le i \le m$; let $v_i \sim \mathcal{D}_i$ for all $1 \le i \le m$ be independent. Then,

$$\mathbb{E}_{\mathcal{D}_i}[v_i v_i^*] = \sum_{k=1}^r \frac{1}{r}\left(\sqrt{r}w_{i,k}\right)\left(\sqrt{r}w_{i,k}\right)^* = \sum_{k=1}^r w_{i,k}w_{i,k}^* = \begin{pmatrix} u_i u_i^* & 0 & \dots & 0 \\ 0 & u_i u_i^* & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & u_i u_i^* \end{pmatrix} \in \mathbb{C}^{(rd)\times(rd)}$$

so that by the assumption that $\sum_{i=1}^m u_i u_i^* = I$, we have $\sum_{i=1}^m \mathbb{E}_{\mathcal{D}_i}[v_i v_i^*] = I_{rd}$. Furthermore, $\|v_i\|^2 = \left\|\sqrt{r}u_i\right\|^2 = r\|u_i\|^2 \le r\delta$.

Setting $\epsilon = r\delta$, Theorem 21 shows there exists an assignment to the $v_1, \dots, v_m$ such that

$$\left\|\sum_{i=1}^m v_i v_i^*\right\| \le \left(1 + \sqrt{r\delta}\right)^2$$

Define $S_k = \{i \in [m] : v_i = w_{i,k}\}$. Then, for all $k \in [r]$,

$$\begin{aligned}\left\|\sum_{i\in S_k} u_i u_i^*\right\| = \left\|\sum_{i\in S_k} w_{i,k}w_{i,k}^*\right\| = \left\|\sum_{i=1}^m v i v_i^*\right\| \le \left\|\sum_{k=1}^r \sum_{i\in S_k} w_{i,k}w_{i,k}^*\right\| \\ = \frac{1}{r}\left\|\sum_{k=1}^r \sum_{i\in S_k}\left(\sqrt{r}w_{i,k}\right)\left(\sqrt{r}w_{i,k}\right)^*\right\| = \frac{1}{r}\left\|\sum_{i=1}^m v_i v_i^*\right\| \le \frac{1}{r}\left(1+\sqrt{r\delta}\right)^2 = \left(\frac{1}{\sqrt{r}} + \sqrt{\delta}\right)^2\end{aligned}$$

□

Weaver's KS_2 conjecture now immediately follows as a special case.

Proof of KS_2 (Conjecture 1). Set $r = 2$, $\delta = 1/18$, $\alpha = 18$, $\beta = 2$ and $u_i = w_i/\sqrt{\alpha}$ for all $1 \le i \le m$. Then, since $\sum_{i=1}^m |\langle u, w_i\rangle|^2 = \alpha$ for all unit vectors $u \in \mathbb{C}^d$, $\sum_{i=1}^m |\langle u, u_i\rangle|^2 = u^*\left(\sum_{i=1}^m u_i u_i^*\right)u = 1$ for all unit vectors $u \in \mathbb{C}^d$ so $\sum_{i=1}^m u_i u_i^* = I$. Furthermore, $\|u_i\|^2 \le \|w_i\|^2/\alpha = \delta\|w_i\|^2 \le \delta$ for all $1 \le i \le m$. Corollary 9 then shows there exists a partition S_1, S_2 of m such that

$$\sum_{i\in S_j} |\langle u, u_i\rangle|^2 = u^*\left(\sum_{i\in S_j} u_i u_i^*\right)u \le \left\|\sum_{i\in S_j} u_i u_i^*\right\| \le \left(\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{18}}\right)^2$$

for every unit vector $u \in \mathbb{C}^d$ and every $j \in \{1, 2\}$. In particular,

$$\sum_{i\in S_j} |\langle u, w_i\rangle|^2 = \alpha \cdot \sum_{i=1}^m |\langle u, u_i\rangle|^2 \le 18 \cdot \left(\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{18}}\right)^2 = 16 = \alpha - \beta$$

for every unit vector $u \in \mathbb{C}^d$ and every $j \in \{1, 2\}$, as desired. Thus KS_2 holds with constants $\alpha = 18$ and $\beta = 2$. □

8.1 A Bound with High Probability

Before diving into the proof, let's contrast Theorem 21 with a "with high probability" bound, which is a direct application of matrix concentration bounds. Essentially, in Theorem 21, we are trading probability guarantees for good upper bounds on the maximum norm: instead of getting a weaker upper bound, that is dependent on the dimension, with high probability, we are getting a very strong dimension-independent bound but with just nonzero probability.

Theorem 22 (Matrix Chernoff; Theorem 1.1 from [31]). *Consider a finite collection of independent Hermitian positive semidefinite $d \times d$ matrix-valued random variables $X_1, \dots, X_m$. Suppose $\lambda_{\max}(X_k) \le R$ with probability 1, for all $1 \le k \le m$. Define $\mu_{\min} = \lambda_{\min}\left(\sum_{k=1}^m \mathbb{E}[X_k]\right)$ and $\mu_{\max} = \lambda_{\max}\left(\sum_{k=1}^m \mathbb{E}[X_k]\right)$. Then for $\delta \in [0,1]$,*

$$\mathbb{P}\left\{\lambda_{\min}\left(\sum_{k=1}^m X_k\right) \le (1-\delta)\mu_{\min}\right\} \le d \cdot \left[\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right]^{\mu_{\min}/R}$$

and for $\delta \ge 0$,

$$\mathbb{P}\left\{\lambda_{\max}\left(\sum_{k=1}^m X_k\right) \ge (1+\delta)\mu_{\max}\right\} \le d \cdot \left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{\mu_{\max}/R}$$

Proposition 1. *Let $v_1, \dots, v_m \in \mathbb{C}^d$ be random vectors independently drawn from distributions $\mathcal{D}_1, \dots, \mathcal{D}_m$, where $\ell_i = |\mathrm{supp}(\mathcal{D}_i)| < \infty$. Assume that $\mathbb{E}\sum_1^m v_i v_i^* = I$ and $\mathbb{E}\|v_i\|^2 \le \epsilon$ for all $i \in [m]$. Furthermore, assume that $\|v_i\|^2 \le c\epsilon$ for all $v_i \in \mathcal{D}_i$ and all $1 \le i \le m$, for some constant $c > 0$. Then*

$$\left\|\sum_{i=1}^m v_i v_i^*\right\| < \mathcal{O}(\epsilon \log d)$$

with high probability, that is, there exists a constant $C > 0$ such that

$$\mathbb{P}\left[\left\|\sum_{i=1}^m v_i v_i^*\right\| < C\epsilon \log d\right] \to 1 \qquad as \qquad d \to \infty$$

Proof. We wish to apply Matrix Chernoff. Let $X_i = v_i v_i^*$ be a random Hermitian positive semidefinite $d \times d$ matrix. By assumption, $\lambda_{\max}(X_i) \le \mathrm{tr}(X_i) = \|v_i\|^2 \le c\epsilon$ with probability 1, for all $1 \le i \le m$. Furthermore,

$$\mu_{\min} = \lambda_{\min}\left(\sum_1^m \mathbb{E} v_i v_i^*\right) = \lambda_{\min}(I_d) = 1$$
$$\mu_{\max} = \lambda_{\max}\left(\sum_1^m \mathbb{E} v_i v_i^*\right) = \lambda_{\min}(I_d) = 1$$

Then, by the Matrix Chernoff bound (Theorem 22), for every $\delta \ge 0$,

$$\mathbb{P}\left[\left\|\sum_{i=1}^m v_i v_i^*\right\| \ge 1+\delta\right] = \mathbb{P}\left[\lambda_{\max}\left(\sum_{i=1}^m X_i\right) \ge (1+\delta)\mu_{\max}\right] \le d \cdot \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{1/c\epsilon}$$

Since $c\epsilon$ is fixed, to prove the claim, it suffices to choose δ such that

$$\frac{e^{\delta}}{(1+\delta)^{1+\delta}} \le o(d^{-c\epsilon})$$

For this, $\delta \le \mathcal{O}(\epsilon \log d)$ suffices. □

It turns out, in order to obtain any kind of "with high probability" guarantees on the operator norm of $\sum_{i=1}^m v_i v_i^*$, one cannot avoid a $\log d$ dependence in the upper bound. Furthermore, one needs a bound on the maximum norm of any vector in each distribution.

8.2 Interlacing Families and the Mixed Characteristic Polynomial

Let $\mathcal{D}_1, \dots, \mathcal{D}_m$ be the given independent distributions, and let $v_{i,1}, \dots, v_{i,\ell_i}$ be the vectors $\mathcal{D}_i$ is supported on, for all $1 \le i \le m$. For convenience, we also view $\mathcal{D}_i$ as a probability distribution over the indices $[\ell_i]$. Let p_{i,j_i} be the probability of v_{i,j_i} (or $j_i \in [\ell_i]$) under $\mathcal{D}_i$. Our strategy is to:

1. show the set of polynomials

$$p_{j_1,\dots,j_m}(x) = \left(\prod_{i=1}^{m} p_{i,j_i}\right)\chi\left[\sum_{i=1}^{m} v_{i,j_i}v_{i,j_i}^*\right](x)$$

   form an interlacing family

2. prove an upper bound on the maximum root of

$$\mathbb{E}_{j_1\sim\mathcal{D}_1,\dots,j_m\sim\mathcal{D}_m}\left[p_{j_1,\dots,j_m}(x)\right]$$

We begin with step 1. By Theorem 7, it suffices to show that averages of the polynomials $p_{j_1,\dots,j_m}(x)$ are real-rooted. This is where we'll employ the real stability results of the previous section. The main theorem is as follows; step 1 is an immediate consequence.

Theorem 23 (Theorem 4.1 from [35])**.** *Let $\mathcal{D}_1,\dots,\mathcal{D}_m$ be independent distributions each supported on finitely many vectors in $\mathbb{C}^d$. Set $A_i = \mathbb{E}[v_i v_i^*]$. Then, for every matrix $M \in \mathbb{C}^{d\times d}$,*

$$\mathbb{E}\det\left(M - \sum_{i=1}^{m} v_i v_i^*\right) = \left[\left(\prod_{i=1}^{m}(1-\partial_{z_i})\right)\det\left(M + \sum_{i=1}^{m} z_i A_i\right)\right]_{z_1=\cdots=z_m=0}$$

In particular, setting $M = xI$, we have the following for the mixed characteristic polynomial.

$$\mu[A_1,\dots,A_m](x) \stackrel{\text{def}}{=} \mathbb{E}\chi\left[\sum_{i=1}^{m} v_i v_i^*\right](x) = \left[\left(\prod_{i=1}^{m}(1-\partial_{z_i})\right)\det\left(xI + \sum_{i=1}^{m} z_i A_i\right)\right]_{z_1=\cdots=z_m=0}$$

First, we need a technical lemma.

Lemma 15. *Let $A \in \mathbb{C}^{n\times n}$ be arbitrary. If $v \in \mathbb{C}^n$ is a random vector, then*

$$\mathbb{E}_v[\det(A - vv^*)] = (1-\partial_t)\det\left(A + t\mathbb{E}_v[vv^*]\right)|_{t=0}$$

Proof. First, assume A is invertible. By the Matrix Determinant Lemma (Lemma 6), we may write

$$\begin{aligned}
\mathbb{E}_v[\det(A - vv^*)] &= \mathbb{E}_v[\det(A)(1 - v^*A^{-1}v)] = \det(A) - \det(A)\mathbb{E}_v[v^*A^{-1}v] \\
&= \det(A) - \det(A)\mathbb{E}_v[\operatorname{tr}(A^{-1}vv^*)] = \det(A) - \det(A)\operatorname{tr}\left(A^{-1}\mathbb{E}_v[vv^*]\right) \\
&= \left[\det\left(A + t\mathbb{E}_v[vv^*]\right) - \det(A)\operatorname{tr}\left(A^{-1}\mathbb{E}_v[vv^*]\right)\right]_{t=0} \\
&= (1-\partial_t)\det\left(A + t\mathbb{E}_v[vv^*]\right)|_{t=0}
\end{aligned}$$

where in the last step, we use Corollary 4.

Now, assume A isn't invertible. Consider a sequence of invertible matrices $\{A_k\}_{k=1}^{\infty}$ such that $A_k \to A$ entry-wise. Since $\mathbb{E}_v[\det(A - vv^*)]$ and $(1-\partial_t)\det(A + t\mathbb{E}_v[vv^*])|_{t=0}$ are continuous functions in A and

$$\mathbb{E}_v[\det(A_k - vv^*)] = (1-\partial_t)\det\left(A_k + t\mathbb{E}_v[vv^*]\right)|_{t=0}$$

for every $k \geq 1$, we have

$$\mathbb{E}_v[\det(A - vv^*)] = \lim_{k\to\infty}\mathbb{E}_v[\det(A_k - vv^*)] = \lim_{k\to\infty}(1-\partial_t)\det\left(A_k + t\mathbb{E}_v[vv^*]\right)|_{t=0} = (1-\partial_t)\det\left(A + t\mathbb{E}_v[vv^*]\right)|_{t=0}$$

□

Proof of Theorem 23. We will prove the claim by induction on $0 \le k \le m$. When $k = 0$, there is nothing to prove. Assume that for every matrix M,

$$\mathbb{E}\det\left(M - \sum_{i=1}^{k} v_i v_i^*\right) = \left[\left(\prod_{i=1}^{k}(1-\partial_{z_i})\right)\det\left(M + \sum_{i=1}^{k} z_i A_i\right)\right]_{z_1=\cdots=z_k=0}$$

for some $0 \le k < m$. We will prove equality for $k+1$. Applying independence of the v_i, we may break the expectation on the left-hand side into

$$\mathbb{E}_{v_{k+1}\sim\mathcal{D}_{k+1}}\mathbb{E}_{v_1\sim\mathcal{D}_1,\ldots,v_k\sim\mathcal{D}_k}\det\left(M - \sum_{i=1}^{k} v_i v_i^* - v_{k+1}v_{k+1}^*\right)$$

By the induction hypothesis, the inner expectation is

$$\mathbb{E}_{v_{k+1}\sim\mathcal{D}_{k+1}}\left[\left(\prod_{i=1}^{k}(1-\partial_{z_i})\right)\det\left(M + \sum_{i=1}^{k} z_k A_k - v_{k+1}v_{k+1}^*\right)\right]_{z_1=\cdots=z_k=0}$$

which, after we use linearity of expectation to move out the differential operators, gives

$$\left[\left(\prod_{i=1}^{k}(1-\partial_{z_i})\right)\mathbb{E}_{v_{k+1}\sim\mathcal{D}_{k+1}}\det\left(M + \sum_{i=1}^{k} z_k A_k - v_{k+1}v_{k+1}^*\right)\right]_{z_1=\cdots=z_k=0}$$

Finally, applying Lemma 15, this inner expectation reduces to

$$\left[\left(\prod_{i=1}^{k}(1-\partial_{z_i})\right)(1-\partial_{z_{k+1}})\det\left(M + \sum_{i=1}^{k} z_k A_k + z_{k+1}A_{k+1}\right)\right]_{z_1=\cdots=z_{k+1}=0}$$

which proves the inductive step. □

Corollary 10. *Let $\mathcal{D}_1,\ldots,\mathcal{D}_m$ be independent distributions each supported on finitely many vectors in $\mathbb{C}^d$. Set $A_i = \mathbb{E}_{v_i\sim\mathcal{D}}[v_i v_i^*]$. Then the mixed characteristic polynomial*

$$\mu[A_1,\ldots,A_m](x) = \mathbb{E}\chi\left[\sum_{i=1}^{m} v_i v_i^*\right](x)$$

is real-rooted.

Proof. Since $A_1,\ldots,A_m,I$ are all positive semidefinite, Theorem 19 shows

$$\det\left(xI + \sum_{i=1}^{m} z_i A_i\right)$$

is real stable in $x, z_1,\ldots,z_m$. Theorem 23 shows that

$$\mu[A_1,\ldots,A_m](x) = \left[\left(\prod_{i=1}^{m}(1-\partial_{z_i})\right)\det\left(xI + \sum_{i=1}^{m} z_i A_i\right)\right]_{z_1=\cdots=z_m=0}$$

i.e. $\mu[A_1,\ldots,A_m](x)$ is obtained from $\det\left(xI + \sum_{i=1}^{m} z_i A_i\right)$ by applying a sequence of stability-preserving operators (Theorem 18). Hence, $\mu[A_1,\ldots,A_m](x)$ is real stable; since it is univariate, it is real-rooted. □

Remark 3. Note that Corollary holds even when $\mathcal{D}_i$ is supported on a single vector, for any $1 \le i \le m$. In the extreme case when $\mathcal{D}_i$ is supported on a single vector v_i for every $1 \le i \le m$, we have that $\mu[A_1,\ldots,A_m](x)$ is just the characteristic polynomial of $\sum_{i=1}^{m} v_i v_i^*$, which we know is real-rooted since it $\sum_{i=1}^{m} v_i v_i^*$ is Hermitian positive semidefinite.

Corollary 11 (Theorem 4.5 from [35]). *The set of polynomials $p_{j_1,\dots,j_m}$ form an interlacing family.*

Proof. Let $j_1,\dots,j_k$ be a partial assignment. We must show that $\{p_{j_1,\dots,j_k,t}(x)\}_{t=1}^{j_{k+1}}$ has a common interlacing. By Theorem 7, it suffices to prove that the convex combination

$$\sum_{t=1}^{\ell_{k+1}} \lambda_t p_{j_1,\dots,j_k,t}(x)$$

is real-rooted for every choice of $\lambda_t \geq 0$ satisfying $\sum_{t=1}^{\ell_{k+1}} \lambda_t = 1$. For this, let w_{k+1} be a random vector that is equal to $v_{k+1,t}$ with probability λ_t. Then

$$\sum_{t=1}^{\ell_{k+1}} \lambda_t p_{j_1,\dots,j_k,t}(x) = \left(\prod_{i=1}^{k} p_{i,j_i}\right) \mathbb{E}_{w_{k+1},v_{k+1},\dots,v_m} \chi\left[\sum_{i=1}^{k} v_{i,j_i} v_{i,j_i}^* + w_{k+1} w_{k+1}^* + \sum_{i=k+2}^{m} v_i v_i^*\right](x)$$

Observe that the right-hand side is a multiple of a mixed characteristic polynomial. Hence, Corollary 10 shows it is real-rooted and we're done. □

8.3 Multivariate Barriers

In this subsection, we prove a general upper bound on the largest root of the mixed characteristic polynomial of positive semidefinite matrices $A_1,\dots,A_m$ in the case when $\sum_{i=1}^m A_i = I$.

Theorem 24 (Theorem 5.1 from [35]). *Suppose $A_1,\dots,A_m$ are Hermitian positive semidefinite matrices satisfying $\sum_{i=1}^m A_i = I$ and $\mathrm{tr}(A_i) \leq \epsilon$ for all $1 \leq i \leq m$. Then $\mathrm{maxroot}(\mu[A_1,\dots,A_m](x)) \leq (1+\sqrt{\epsilon})^2$.*

This will be useful for the main theorem we want to prove for Kadison-Singer, since $A_i = \mathbb{E}[v_i v_i^*]$ so that $\mathrm{tr}(A_i) = \mathrm{tr}(\mathbb{E}[v_i v_i^*]) = \mathbb{E}[\mathrm{tr}(v_i v_i^*)] = \mathbb{E}[\|v_i\|^2]$. In particular, our assumption that $\mathbb{E}[\|v_i\|^2] \leq \epsilon$ directly translates to $\mathrm{tr}(A_i) \leq \epsilon$.

The idea here is similar the argument we used for building well-invertible submatrices and Ramanujan sparsifiers. Theorem 23 shows us the mixed characteristic polynomial $\mu[A_1,\dots,A_m](x)$ is the result of applying a sequence of differential and specialization operators to $\det\left(\sum_{i=1}^m z_i A_i\right)$. In particular, we can think of $\mu[A_1,\dots,A_m](x)$ as being the result of a process that slowly moves the roots of $\det\left(\sum_{i=1}^m z_i A_i\right)$ to the roots of $\left(\prod_{i=1}^m (1-\partial_{z_i})\right)\det\left(\sum_{i=1}^m z_i A_i\right)$. We will devise multivariate barrier functions that will "guide" this process. Specifically, we will use the barriers given by

$$\Phi_p^i(z) = \frac{\partial_{z_i} p(z)}{p(z)} = \partial_{z_i} \log p(z)$$

For convenience, we say a vector $y \in \mathbb{R}^d$ is *above the roots of* $p \in \mathbb{R}[z_1,\dots,z_d]$ if $p(y+t) > 0$ for all $t \in \mathbb{R}^d_{\geq 0}$; the set of points that are above the roots of p is denoted by $\mathrm{ABOVE}(p)$. When p is a real-rooted univariate polynomial, $\mathrm{ABOVE}(p)$ is precisely the set of real numbers larger than the largest root of p. To prove the desired bound, it then suffices to prove that

$$(1+\sqrt{\epsilon})^2 \cdot \mathbf{1} \in \mathrm{ABOVE}\left(\left(\prod_{i=1}^m (1-\partial_{z_i})\right)\det\left(\sum_{i=1}^m z_i A_i\right)\right)$$

since then, $(1+\sqrt{\epsilon})^2 \in \mathrm{ABOVE}(\mu[A_1,\dots,A_m](x)) \implies (1+\sqrt{\epsilon})^2 \geq \mathrm{maxroot}(\mu[A_1,\dots,A_m](x))$.

We begin by showing a nice identity for the barrier in a special case that we will need later.

Lemma 16. *Suppose* $p(z_1,\dots,z_m) = \det\left(\sum_{i=1}^m z_i A_i\right)$*, where* $A_1,\dots,A_m$ *are Hermitian positive semidefinite matrices. Then*

$$\Phi_p^j(z_1,\dots,z_m) = \operatorname{tr}\left[\left(\sum_{i=1}^m z_i A_i\right)^{-1} A_j\right]$$

for every $z = (z_1,\dots,z_m) \in \mathrm{ABOVE}(p)$.

Proof. For every $z = (z_1,\dots,z_m) \in \mathrm{ABOVE}(p)$, by definition, $p(z_1,\dots,z_m) > 0$. In particular, $\sum_{i=1}^m z_i A_i$ is invertible for every $z \in \mathrm{ABOVE}(p)$. Thus, $\sum_{i=1}^m z_i A_i$ is invertible for every $z \in \mathrm{ABOVE}(p)$. By Corollary 4,

$$\partial_{z_j} p(z_1,\dots,z_m) = \det\left(\sum_{i=1}^m z_i A_i\right)\operatorname{tr}\left[\left(\sum_{i=1}^m z_i A_i\right)^{-1} A_j\right] = p(z_1,\dots,z_m)\cdot\operatorname{tr}\left[\left(\sum_{i=1}^m z_i A_i\right)^{-1} A_j\right]$$

Hence, substituting this into the definition of Φ_p^j, we have the claim.

$$\Phi_p^j(z_1,\dots,z_m) = \frac{\partial_{z_j} p(z_1,\dots,z_m)}{p(z_1,\dots,zm)} = \operatorname{tr}\left[\left(\sum_{i=1}^m z_i A_i\right)^{-1} A_j\right]$$

□

Now, let's turn to some nice properties of the barriers we've defined and their relation to the set $\mathrm{ABOVE}(p)$. Specifically, they are nonincreasing and convex in every coordinate, holding other variables fixed.

Lemma 17 (Lemma 5.7 from [35])**.** *Suppose* p *is real stable and* $z \in \mathrm{ABOVE}(p)$*. Then for all* $1 \le i,j \le m$ *and* $\delta \ge 0$*:*

1. $\Phi_p^i(z+\delta e_j) \le \Phi_p^i(z)$ *(monotonicity)*
2. $\Phi_p^i(z+\delta e_j) \le \Phi_p^i(z) + \delta\cdot\partial_{z_j}\Phi_p^i(z+\delta e_j)$ *(convexity)*

Proof. We consider the two cases $i = j$ and $i \ne j$ separately. To avoid confusion between the variable arguments of Φ_p^i, p, etc. and $z \in \mathrm{ABOVE}(p)$, we will use $x = (x_1,\dots,x_m)$ for the variable arguments.

Let's start with the easier one: $i = j$. Let $q_{z,i}(x_i) = \prod_{k=1}^d (x_i - \lambda_k)$, where $\lambda_1,\dots,\lambda_d$ are the roots of $q_{z,i}(x_i)$, the restriction of p to coordinate i while keeping all other coordinates fixed as constants; note $\lambda_1,\dots,\lambda_d$ are real by Theorem 18 since p is real stable. Then

$$\Phi_p^i(x) = \frac{q'_{z,i}(x_i)}{q_{z,i}(x_i)} = \sum_{k=1}^d \frac{1}{x_i - \lambda_k}$$

Differentiating, we have

$$\partial_{x_i}\Phi_p^i(x) = \sum_{k=1}^d \frac{\partial}{\partial x_i}\frac{1}{x_i-\lambda_k} = -\sum_{k=1}^d \frac{1}{(x_i-\lambda_k)^2} \le 0$$

proving monotonicity. Convexity follows by differentiating once more (recall that a twice differentiable univariate function is convex if and only if its second derivative is nonnegative).

$$\partial_{x_i}^2\Phi_p^i(x)\,|_{x=z} = \left[\sum_{k=1}^d \frac{\partial^2}{\partial x_i^2}\frac{1}{x_i-\lambda_k}\right]_{x=z} = \sum_{k=1}^d \frac{2}{(z_i-\lambda_k)^3} > 0$$

where in the last inequality, we use the fact that $z \in \text{ABOVE}(p) \implies z_i > \lambda_k$.

Now, let's turn to the case when $i \neq j$. Similarly, define $q_{z,i,j}(x_i, x_j)$ to be the bivariate restriction of p to the coordinates x_i, x_j, fixing all other coordinates as constants. Again, note that $q_{z,i,j}$ is real stable and $(z_i, z_j) \in \text{ABOVE}(q_{z,i,j})$. By Theorem 20, there are Hermitian positive semidefinite matrices B_i, B_j and a Hermitian matrix C such that either $q_{z,i,j}(x_i, x_j) = \det(x_i B_i + x_j B_j + C)$ or $q_{z,i,j}(x_i, x_j) = -\det(x_i B_i + x_j B_j + C)$; since $q_{z,i,j}(z_i + t_i, z_j + t_j) > 0$ for every $(t_i, t_j) \in \mathbb{R}^2_{\geq 0}$, $q_{z,i,j}(x_i, x_j) = \det(x_i B_i + x_j B_j + C)$.

For convenience, set $M = z_i B_i + z_j B_j + C$; note M is positive definite (in particular, invertible) since $z \in \text{ABOVE}(p)$ and $B_i + B_j$ is positive definite by Theorem 20. By Lemma 4 and the definition of M, we may write $\Phi^i_p(z) = \operatorname{tr}\left(M^{-1} B_i\right) = \operatorname{tr}\left(I B_i M^{-1}\right)$ and

$$\begin{aligned}\Phi^i_p(z + \delta e_j) &= \operatorname{tr}\left((M + \delta B_j)^{-1} B_i\right) = \operatorname{tr}\left(((I + \delta B_j M^{-1})M)^{-1} B_i\right) \\ &= \operatorname{tr}\left(M^{-1}(I + \delta B_j M^{-1})^{-1} B_i\right) = \operatorname{tr}\left((I + \delta B_j M^{-1})^{-1} B_i M^{-1}\right)\end{aligned}$$

For $\delta \geq 0$ sufficiently small, we can write out $(I + \delta B_j M^{-1})^{-1}$ as the following power series.

$$I - \delta B_j M^{-1} + \delta^2 (B_j M^{-1})^2 + \sum_{k=3}^{\infty} (-\delta B_j M^{-1})^k$$

Hence,

$$\begin{aligned}\Phi^i_p(z + \delta e_j) &= \operatorname{tr}\left(\left(I - \delta B_j M^{-1} + \delta^2 (B_j M^{-1})^2 + \sum_{k=3}^{\infty} (-\delta B_j M^{-1})^k\right) B_i M^{-1}\right) \\ &= \operatorname{tr}(I B_i M^{-1}) + \operatorname{tr}\left(\left(-\delta B_j M^{-1} + \delta^2 (B_j M^{-1})^2 + \sum_{k=3}^{\infty} (-\delta B_j M^{-1})^k\right) B_i M^{-1}\right) \\ &= \Phi^i_p(z) + \operatorname{tr}\left(\left(-\delta B_j M^{-1} + \delta^2 (B_j M^{-1})^2 + \sum_{k=3}^{\infty} (-\delta B_j M^{-1})^k\right) B_i M^{-1}\right)\end{aligned}$$

Writing out the definition of the partial derivative, we have

$$\begin{aligned}\partial_{z_j} \Phi^i_p(z) &= \lim_{\delta \to 0} \frac{\Phi^i_p(z + \delta e_j) - \Phi^i_p(z)}{\delta} = \lim_{\delta \to 0} \operatorname{tr}\left(\left(-B_j M^{-1} + \delta (B_j M^{-1})^2 + \sum_{k=3}^{\infty} (-1)^k \delta^{k-1} (B_j M^{-1})^k\right) B_i M^{-1}\right) \\ &= \operatorname{tr}\left(-B_j M^{-1} B_i M^{-1}\right)\end{aligned}$$

To prove monotonicity, observe that $\partial_{z_j} \Phi^i_p(z) = -\operatorname{tr}\left(B_j M^{-1} B_i M^{-1}\right) \leq 0$ since B_j and $M^{-1} B_i M^{-1}$ are positive semidefinite.

It remains to prove convexity in the direction e_j. Following a similar computation, if we differentiate again, we obtain

$$\partial^2_{z_j} \Phi^i_p(z) = \operatorname{tr}\left((B_j M^{-1} B_j)(M^{-1} B_i M^{-1})\right)$$

which is nonnegative since $B_j M^{-1} B_j$ and $M^{-1} B_i M^{-1}$ are positive semidefinite (since they are conjugate to M^{-1} and B_i respectively, which are both positive semidefinite). This proves convexity. □

Now, we derive update rules for our barriers that will allow us to reason about the evolution of the maximum root.

Lemma 18 (Lemmas 5.9 and 5.10 from [35]). *Suppose p is real stable and $z \in \text{ABOVE}(p)$.*

1. *If* $\Phi_p^i(z) < 1$, *then* $z \in \text{ABOVE}((1-\partial_{z_i})p)$.

2. *If* $\delta > 0$ *satisfies* $\Phi_p^j(z) \le 1-\delta^{-1}$, *then* $\Phi^i_{(1-\partial_{z_j})p}(z+\delta e_j) \le \Phi_p^i(z)$ *for all* i.

Proof. 1. Let $t \in \mathbb{R}^m_{\ge 0}$. By Lemma 17 part 1, Φ_p is nonincreasing in each coordinate. Since $\Phi_p^i(z) < 1$, we have $\Phi_p^i(z+t) \le \Phi_p^i(z) < 1$. Writing out the definition of Φ_p^i, we see this is equivalent to $\partial_{z_i} p(z+t) < p(z+t)$, which implies $(1-\partial_{z_i})p(x+t) > 0$, for all $1 \le i \le m$. Since $t \in \mathbb{R}^m_{\ge 0}$ was arbitrary, $z \in \text{ABOVE}((1-\partial_{z_i})p)$.

2. For convenience, we will take ∂_i to mean ∂_{z_i}. Let's first write $\Phi^i_{(1-\partial_{z_j})p}$ in terms of Φ_p^j, Φ_p^i and $\partial_j \Phi_p^i$. Observe that

$$p - \partial_j p = \left(1 - \frac{\partial_j p}{p}\right) p = (1-\Phi_p^j)p$$

Hence,

$$\Phi^i_{(1-\partial_{z_j})p} = \frac{\partial_i(p-\partial_j p)}{p-\partial_j p} = \frac{\partial_i\left((1-\Phi_p^j)p\right)}{(1-\Phi_p^j)p} = \frac{(1-\Phi_p^j)(\partial_i p)}{(1-\Phi_p^j)p} + \frac{p(\partial_i(1-\Phi_p^i))}{(1-\Phi_p^j)p} = \Phi_p^i - \frac{\partial_i \Phi_p^j}{1-\Phi_p^j}$$
$$= \Phi_p^i - \frac{\partial_j \Phi_p^i}{1-\Phi_p^j}$$

where in the last step we use the fact that we can change the order of partial differentiation:

$$\Phi_p^i - \frac{\partial_i \Phi_p^j}{1-\Phi_p^j} = \Phi_p^i - \frac{\partial_i \partial_j \log p(z)}{1-\Phi_p^j} = \Phi_p^i - \frac{\partial_j \partial_i \log p(z)}{1-\Phi_p^j} = \Phi_p^i - \frac{\partial_j \Phi_p^i}{1-\Phi_p^j}$$

Thus, $\Phi^i_{(1-\partial_{z_j})p}(z+\delta e_j) \le \Phi_p^i(z)$ is equivalent to

$$-\frac{\partial_j \Phi_p^i(z+\delta e_j)}{1-\Phi_p^j(z+\delta e_j)} \le \Phi_p^i(z) - \Phi_p^i(z+\delta e_j)$$

By Lemma 17 part 2 (convexity of Φ_p^i in coordinate j, for all i, j), it suffices to show that the left-hand side obeys the following tighter upper bound.

$$-\frac{\partial_j \Phi_p^i(z+\delta e_j)}{1-\Phi_p^j(z+\delta e_j)} \le -\delta \cdot \partial_j \Phi_p^i(z+\delta e_j)$$

Again, since $\Phi_p^j(z)$ is nonincreasing for every coordinate by Lemma 17 part 1, we have $\partial_j \Phi_p^i(z+\delta e_j) \le 0$. Thus, dividing by $-\partial_j \Phi_p^i(z+\delta e_j)$ from both sides, we have the equivalent inequality

$$\frac{1}{1-\Phi_p^j(z+\delta e_j)} \le \delta$$

This is immediately implied by the fact that $\Phi_p^j(z+\delta e_j) \le \Phi_p^j(z) \le 1-\delta^{-1} < 1$.

□

Proof of Theorem 21. Consider the polynomial

$$p_k(z_1, \ldots, z_m) = \left(\prod_{i=1}^k (1-\partial_{z_i})\right) \det\left(\sum_{i=1}^m z_i A_i\right)$$

defined for all $0 \le k \le m$; note p_k is real stable by Theorem 19 and Theorem 18 for all $0 \le k \le m$, and $p_{k+1} = (1 - \partial_{z_{k+1}})p_k$ for all $0 \le k < m$. For convenience, define the variables $t = \sqrt{\epsilon} + \epsilon$, $\phi = 1 - \frac{1}{1+\sqrt{\epsilon}} < 1$, and $\delta = 1/(1-\phi) = 1 + \sqrt{\epsilon} > 0$; note $1 - \delta^{-1} = \phi$. Finally, for $0 \le k \le m$, define $x^k = t\mathbf{1} + \sum_{i=1}^{k} \delta e_i$ to be the vector with $t + \delta$ in the first k coordinates and t in the remaining coordinates. We will prove by induction on $0 \le k \le m$ that

1. $x^k \in \text{ABOVE}(p_k)$
2. $\Phi^i_{p_k}(x^k) \le \phi$ for all $1 \le i \le m$

It will then follow that with $k = m$, we have $x^m = (t+\delta)\mathbf{1} = (1+\sqrt{\epsilon})^2 \cdot \mathbf{1} \in \text{ABOVE}(p_m)$. Recalling that by Theorem 23, p_m satisfies $p_m(x,\dots,x) = \mu[A_1,\dots,A_m](x)$, this then proves the $(1+\sqrt{\epsilon})^2$ upper bound on $\text{maxroot}(\mu[A_1,\dots,A_m](x))$.

Since $A_1,\dots,A_m$ are positive semidefinite and $\sum_{i=1}^m A_i = I \implies \sum_{i=1}^m x^0(i)A_i = t\sum_{i=1}^m A_i = tI \implies p(x^0) = \det(tI) > 0$, $x^0 \in \text{ABOVE}(p)$. Furthermore, by Corollary 4,

$$\Phi^i_{p_0}(z_1,\dots,z_m) = \frac{\partial_i p_0(z_1,\dots,z_m)}{p_0(z_1,\dots,z_m)} = \text{tr}\left(\left(\sum_{i=1}^m z_i A_i\right)^{-1} A_i\right)$$

so that

$$\Phi^i_{p_0}(x^0) = \text{tr}\left((tI)^{-1} \cdot A_i\right) = t^{-1}\,\text{tr}(A_i) \le \frac{\epsilon}{t} = \frac{\epsilon}{\epsilon + \sqrt{\epsilon}} = \phi$$

This verifies that the base case holds.

Suppose the claim holds for some $k \ge 0$. Since p_k is real stable, $x^k \in \text{ABOVE}(p_k)$ and $\Phi^{k+1}_{p_k}(x^k) \le \phi < 1$, we have $x^k \in \text{ABOVE}((1-\partial_{z_{k+1}})p_k) = \text{ABOVE}(p_{k+1})$ by Lemma 18 part 1; since the coordinates of $x^{k+1} = x^k + \delta e_{k+1}$ are at least those of x^k, $x^{k+1} \in \text{ABOVE}(p_{k+1})$. Furthermore, since $\delta > 0$ satisfies $\Phi^{k+1}_{p_k}(x^k) \le \phi = 1 - \delta^{-1}$, we have by Lemma 18 part 2 that

$$\Phi^i_{p_{k+1}}(x^{k+1}) = \Phi^i_{(1-\partial_{z_{k+1}})p_k}(x^k + \delta e_{k+1}) \le \Phi^i_{p_k}(x^k) \le \phi$$

for all $1 \le i \le m$. Thus, we also have $\Phi^i_{p_{k+1}}(x^{k+1}) \le \phi$. This completes the induction and proves the theorem. □

With the results of this subsection and the previous subsection, we can now prove the main probabilistic result that implies a positive answer to the Kadison-Singer problem.

Proof of Theorem 21. Let $\mathcal{D}_1,\dots,\mathcal{D}_m$ be the given independent distributions such that $\sum_{i=1}^m \mathbb{E}_{\mathcal{D}_i}[v_i v_i^*] = I$ and $\mathbb{E}_{\mathcal{D}_i}[\|v_i\|^2] \le \epsilon$. Define $A_i = \mathbb{E}_{\mathcal{D}_i}[v_i v_i^*]$. Then, by linearity of trace and expectation, for all $1 \le i \le m$,

$$\text{tr}[A_i] = \text{tr}[\mathbb{E}_{\mathcal{D}_i}[v_i v_i^*]] = \mathbb{E}_{\mathcal{D}_i}[\text{tr}[v_i v_i^*]] = \mathbb{E}_{\mathcal{D}_i}[v_i^* v_i] = \mathbb{E}_{\mathcal{D}_i}[\|v_i\|^2] \le \epsilon$$

By Theorem 24, $\text{maxroot}(\mu[A_1,\dots,A_m](x)) \le (1+\sqrt{\epsilon})^2$. Now, Corollary 11 shows the set of polynomials $p_{j_1,\dots,j_m}$ form an interlacing family so by Corollary 7, there exists an assignment $j_1,\dots,j_m$, i.e. vectors $v_{1,j_1} \in \mathcal{D}_1,\dots,v_{m,j_m} \in \mathcal{D}_m$ such that

$$\text{maxroot}\left(\chi\left[\sum_{i=1}^m v_{i,j_i} v_{i,j_i}^*\right](x)\right) \le \text{maxroot}\left(\mu[A_1,\dots,A_m](x)\right) \le (1+\sqrt{\epsilon})^2$$

that is,

$$\left\|\sum_{i=1}^m v_{i,j_i} v_{i,j_i}^*\right\| \le (1+\sqrt{\epsilon})^2$$

as desired. □

8.4 Applications

8.4.1 Revisiting Restricted Invertibility

Recall that Restricted Invertibility (Theorem 10) guarantees the existence of a well-invertible column submatrix A_S of A. One can ask if we can make stronger guarantees. For example, can we "split up" A into several well-invertible submatrices? Under what conditions can this occur?

Problem 4 (Strong Bourgain-Tzafriri)**.** *There exists a universal constant $c > 0$ such that for every $B > 0$, there exists $m \in \mathbb{N}$ such that for every linear operator $T : \mathbb{R}^n \to \mathbb{R}^n$ with $\|Te_i\| = 1$ for all $1 \le i \le n$ and $\|T\| \le B$, then there is a partition $\{S_j\}_1^m$ of $[n]$ such that*

$$\|Tv\|_2^2 \ge c\|v\|_2^2$$

for every $v \in \operatorname{span}\{e_i : i \in S_j\}$, for every $1 \le j \le m$.

Casazza-Vershynin proved that this is in fact equivalent to the Kadison-Singer problem.

Theorem 25 (Theorem 2.4 from [26])**.** *The Kadison-Singer Problem has a positive solution if and only if the Strong Bourgain-Tzafriri Problem has a positive solution.*

Thus, the positive resolution to the Kadison-Singer Problem immediately implies a much stronger claim over Theorem 10.

8.4.2 The Asymmetric Traveling Salesman Problem

Theorem 21 made the assumption that the $\mathcal{D}_1, \dots, \mathcal{D}_m$ are independent. One can ask under what looser assumptions on the relationship between the $\mathcal{D}_1, \dots, \mathcal{D}_m$ gives similar results. One can indeed relax the independence assumptions and for this, we need a couple definitions.

Definition 8. Let μ be a probability distribution on the collection of subsets of $[m]$. Define the **generating polynomial of** μ in the variables $x_1, \dots, x_m$ by

$$g_\mu(x_1, \dots, x_m) = \sum_{S \subset [m]} \mu(S) \cdot x^S$$

where $x^S = \prod_{i \in S} x^S$ for every $S \subset [m]$ (with $x^\emptyset = 1$). μ is a **homogeneous** distribution if g_μ is a homogeneous polynomial. μ is a **strongly Rayleigh** distribution if g_μ is real stable.

Theorem 26 (Theorem 1.2 from [37])**.** *Let μ be a homogeneous strongly Rayleigh probability distribution on $[m]$ such that $\mathbb{P}_{S \sim \mu}[i \in S] \le \epsilon_1$ for all $1 \le i \le m$. Let $v_1, \dots, v_m \in \mathbb{R}^d$ be vectors satisfying*

$$\sum_{i=1}^m v_i v_i^* = I$$

and $\|v_i\|^2 \le \epsilon_2$ for all $1 \le i \le m$. Then,

$$\mathbb{P}_{S \sim \mu}\left[\left\|\sum_{i \in S} v_i v_i^*\right\| \le 4(\epsilon_1 + \epsilon_2) + 2(\epsilon_1 + \epsilon_2)^2\right] > 0$$

This result is actually incomparable to Theorem 21 but, just as Theorem 21, does imply a positive answer to the Kadison-Singer Problem. The main benefit here is that it allows one to make similar existence claims but when the objects (vectors) being analyzed are negatively dependent.

[37] proved this result and used it to show the existence of a thin basis in a set of vectors, and the existence of a spectrally thin tree in a graph. These results together with the results of [39] imply that the Held-Karp linear program relaxation for the Asymmetric Traveling Salesman Problem is bounded asymptotically by a polynomial in $\log \log n$, where n is the number of vertices in the graph. We refer interested readers to [37] and [39] for details.

9 Building Bipartite Ramanujan Graphs I: Ramanujan Covers

The goal of this section is to present a construction of infinite families of simple d-regular bipartite Ramanujan graphs, for arbitrary $d > 1$. The idea is based on taking a simple d-regular bipartite Ramanujan graph and "lifting" it to obtain a larger simple d-regular Ramanujan graph. This construction is primarily existential; we do not know of a way to make the natural interlacing families algorithm efficient on the interlacing family we will analyze here. Later, we'll see a completely different construction that actually has a polynomial time algorithm for construction.

9.1 Expansion

We begin with a definition.

Definition 9 (Edge Expansion)**.** Let $G = (V, E)$ be a d-regular graph. We define the **edge expansion of** $S \subset V$ by

$$\phi(S) = \frac{|E(S, \overline{S})|}{d \cdot |S|}$$

where $\overline{S} = V \setminus S$ and $E(S, \overline{S}) = \{(u, v) \in E : u \in S, v \in \overline{S}\}$ is the set of edges crossing the cut $(S, \overline{S})$. The **edge expansion of** G is defined to be $\phi(G) = \min_{S \subset V : |S| \leq |V|/2} \phi(S)$.

Note that in a d-regular graph, $d \cdot |S|$ is the total number of edges incident to some vertex in S, while $|E(S, \overline{S})|$ is the total number of edges incident to some vertex in S **and** leaving S. Hence, $\phi(S)$ in some sense measures the connectivity of S to the rest of the graph, reweighted by its "importance", namely its size.

The intuition here is as follows: we want to measure the connectivity of a graph. Of course, if this is all we cared about, we could just use the complete graph, which has maximum connectivity (each vertex is connected to as many other vertices as possible). However, we also want to be cost efficient (the complete graph has $\Omega(n^2)$ edges). Hence, we need to trade some connectivity for sparsity.

How should this tradeoff be made? Well, if only a single vertex has very low connectivity with the rest of the graph, then maybe we don't care that much. Low connectivity only really becomes a problem when removal of only a few edges disconnects are large portion (say, a constant fraction of the vertices) of the graph from another large portion.

Example 3. Here is an example. Although the following graph is not d-regular, it illustrates why we not only care about connectivity but also a quantity like expansion, where connectivity is weighted by "importance" (quantified by the size of the vertex sets on each side of the cut).

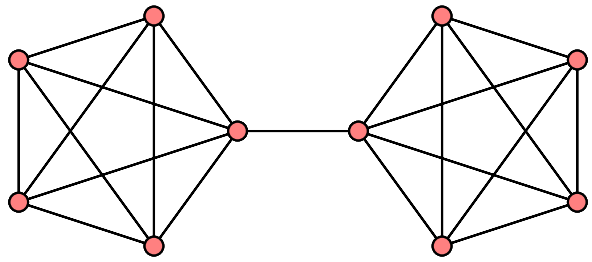

Figure 1: This is a dumbbell graph consisting of two 5-cliques connected by a single edge.

The main problem with this graph is that removal of just a single edge, namely, the middle edge, disconnects the two halves, both with $n/2 \geq \Omega(n)$ vertices, from each other. This would make a poor network indeed!

Here is a simple result making the preceding intuition more rigorous.

Lemma 19 (Lemma 1.2 from [32]). *Let $G = (V, E)$ be a d-regular graph. If any ϵ-fraction of the edges of G are removed, for some $0 < \epsilon < \phi(G)$, then G will still have a connected component that contains at least a $\left(1 - \frac{\epsilon}{2\phi(G)}\right)$-fraction of the vertices.*

With this intuition, it becomes clear that large edge expansion is desirable. Expanders are graphs "with $\Omega(1)$ expansion", i.e. those for which every cut S has $|E(S, \overline{S})|$ within a constant (times d) factor of $|S|$ (so that large sets subsets of vertices have many edges leaving it). More formally:

Definition 10 (Expanders). A sequence $\{G_n = (V_n, E_n)\}_{n\in\mathbb{N}}$ of d_n-regular graphs with $|V_n| = n$ are a family of **expanders** if there exists an absolute constant $c > 0$ independent of n such that $\phi(G_n) \geq c$ for all $n \in \mathbb{N}$.

Now, let's look at a key tool to approximating the edge expansion of a graph.

Theorem 27 (Cheeger's Inequalities). *Let $\mu_1 \leq \cdots \leq \mu_n$ be the eigenvalues of L_G, where $G = (V, E)$ is a d-regular graph. Then*

$$\frac{\mu_2}{2d} \leq \phi(G) \leq \sqrt{\frac{2\mu_2}{d}}$$

Translated into the eigenvalues $\lambda_1 \leq \cdots \leq \lambda_n$ of the adjacency matrix A_G of G, we have

$$\frac{1}{2} - \frac{\lambda_2}{2d} \leq \phi(G) \leq \sqrt{2 - \frac{2\lambda_2}{d}}$$

Thus, to prove that G has large expansion, it suffices to show that λ_2 is small w.r.t. d.

Note: When analyzing spectral properties of d-regular graphs, we generally do not care about the largest eigenvalue, which is d (consider the vector $\mathbf{1}$; d-regularity of G shows that sum of each row is d so that $A_G\mathbf{1} = d\mathbf{1}$). This only tells us about d-regularity of G and nothing else. Similarly, if G is also bipartite, we do not care about the smallest eigenvalue, since symmetry of the spectrum of A_G (Theorem 5) shows that the smallest eigenvalue is $-d$; again, this only tells us d-regularity and bipartiteness, nothing else. We refer to these eigenvalues as the "trivial" eigenvalues.

We conclude this section with the definition of Ramanujan graphs. We will see soon why these graphs are considered "optimal" expanders.

Definition 11 (Ramanujan Graph). A d-regular graph G is **Ramanujan** if the largest nontrivial eigenvalue in absolute value of its adjacency matrix (not $\pm d$) is upper bounded by $2\sqrt{d-1}$

Remark 4. One can actually define Ramanujan graphs much generally via the **unverisal cover** of G. We will not discuss this notion in this survey.

9.2 Applications of Expanders

9.2.1 Rapid Random Walk Mixing

Definition 12 (Markov Chain). A **Markov chain** is a set of states Ω with a **transition (probability) function** $P : \Omega \times \Omega \to [0, 1]$ with $P(a, b) = \mathbb{P}[a \to b]$, i.e. the probability of "transitioning" from a to b. If Ω is finite, we often write P as a **transition matrix**. Furthermore, we often represent the chain as a weighted directed graph $G = (V = \Omega, E = \Omega \times \Omega)$, where each edge $(a, b) \in E$ has weight $P(a, b)$.

Remark 5. Given a (possibly directed and/or nonnegatively weighted) graph $G = (V, E)$, one can define a Markov chain with $\Omega = V$ and transition probabilities given by $P(a, b) = w_{a,b}/\sum_{v\sim a} w_{a,v} = w_{a,b}/d_w(a)$; if $(u, v) \notin E$), then $P(u, v) = 0$. More compactly, we can write the transition probabilities as $P = D^{-1}A$, where $D = \text{diag}(\{d_w(u)\}_{u\in V})$.

Definition 13 (Random Walk). Let (Ω, P) be a Markov chain. A **walk** on Ω is a sequence of states $\{s_i\}_1^m$ (we allow $m = \infty$). A **random walk** on Ω with initial distribution Q is a random sequence $\{s_i\}_1^m$ (we allow $m = \infty$) such that $s_1 \sim Q$ initially, and we generate s_{i+1} with probability $P[s_{i+1} \mid s_i] = P(s_i, s_{i+1})$, for all i.

Markov chains are very useful in approximation algorithms. For example, Markov Chain Monte Carlo (MCMC) is used to sample from complex distributions that are intractable to learn exactly. An application from statistical machine learning is approximate learning of probabilistic graphical models (Hidden Markov Models, Bayesian networks, etc.). It has also found applications in web search and page rank, modeling users' website visits as a Markov chain. In physics, Markov chains are used to model thermodynamic systems. Finally, in economics, Markov chains are used to model the evolution of stock prices.

A key property useful in modeling with Markov chains is that they converge to a "stationary distribution" on the vertices; this is usually the intractable distribution we are trying to sample from (by the design of the chain). To generate samples, one then runs a random walk for some time on the Markov chain so that it converges to an approximate stationary distribution. More specifically, at each time step t of the walk, the probability that s_t is at a certain vertex is given by $(P^\top)^t \pi$.

Definition 14 (Stationary Distribution)**.** A distribution π on the vertices of $G = (V, E)$ is a **stationary distribution** if $P^\top \pi = \pi$.

Theorem 28. *For a Markov chain given as a weighted undirected graph G, if G is connected and not bipartite, then G has a unique stationary distribution. Furthermore, $\lim_{t\to\infty}(P^\top)^t \pi_0 = \pi$ for any initial distribution π_0 on V.*

Remark 6. For general weighted graphs, the bipartiteness condition can be fixed by making the chain "lazy"; we do this by adding "self-edges" (or "loops") to each vertex to make G not bipartite. Furthermore, we make the weight of the self-loop on vertex v equal to $d_w(v)$. It turns out this preserves the stationary distribution of the chain.

In order to efficiently generate these samples from an approximate stationary distribution, one needs to show that the Markov chain "converges" quickly to the stationary distribution.

Definition 15 (Mixing Time)**.** The **mixing time** of a Markov chain corresponding to a weighted graph G, with stationary distribution π, is the smallest time t such that for any initial distribution π_0 on V, $\left\|(P^\top)^t \pi_0 - \pi\right\|_1 \leq 1/2e$, where $\|\cdot\|_1$ is the total variation distance (i.e. ℓ_1-distance).

Thus, when one designs a Markov chain for a problem (such as learning a Bayesian network), one needs that the chain not only converges to the correct stationary distribution, but also that it doesn't take too long to converge; only then will the chain be useful.

Now, here is the connection to expander graphs. For simplicity, we'll work with d-regular unweighted undirected graphs.

Theorem 29. *Let $G = (V, E)$ be a d-regular, unweightd, undirected graph. Then the mixing time of the lazy Markov chain associated with G mixes in*

$$\mathcal{O}\left(\frac{\log n}{\phi(G)^2}\right)$$

steps.

Recall that expanders are graphs for which $\phi(G)$ is $\Omega(1)$. Hence, this result states that running a lazy Markov chain on an expander will converge quickly (become $1/2e$-close in total variation distance in $\mathcal{O}(\log n)$ steps) to the correct stationary distribution that we want.

9.2.2 Error-Correcting Codes

Suppose we would like to transmit a m-bit message over a (possibly insecure) channel. We want the bits that the sender ends up sending to be "robust" to errors, where some bits may be flipped along the way, in the sense that the receiver can still recover the original message being sent. Of course, if the sender just sends his original m-bit message without modification, any error in the received message is unrecoverable. Some redundancy/additional bits will be needed.

For example, one can send three copies of the same m-bit message concatenated to each other. The receiver can then take a majority vote over all three copies for each bit. If no pair of bits corresponding to the same location in the message are corrupted, then the receiver will recover the message correctly. The problem with this is that it is communication efficient (we need to transmit three times as many bits) with weak guarantees.

First, let's formalize a model. We will view $\{0,1\}^m$ as the vector space $\mathbb{F}_2^m$, where addition and scalar multiplication are down element-wise modulo 2.

Definition 16 (Code)**.** Fix m and consider the set of all possible m-bit messages $\{0,1\}^m$. Fix $n > m$. An **error-correcting code** is an injective map $C : \{0,1\}^m \to \{0,1\}^n$. Each n-bit string in the image of C is called a **codeword**. A code is **linear** if every linear combination of codewords is a codeword, that is, there exists a **parity-check matrix** $M \in \{0,1\}^{(n-m)\times n}$ such that $y \in \text{im}(C)$ if and only if $My = 0$. Note every linear code has a matrix $A \in \{0,1\}^{n\times m}$ such that $C(x) = Ax$.

The **rate** of a code $C : \{0,1\}^m \to \{0,1\}^n$ is equal to $r(C) = m/n$. The **distance** of a code $C : \{0,1\}^m \to \{0,1\}^n$ is $\min_{c_1,c_2\in\text{im}(C)} \|c_1 - c_2\|_1$.

Remark 7. Of course, we also want encoding and decoding to be done efficiently. This is why we will focus on linear codes in this section.

Coding theory studies codes and the tradeoffs between the rate of the code, a measure of its bit efficiency, and the distance of the code, a measure of its robustness to noise (the distance is the minimum number of bit flips, i.e. errors, required to transform a codeword into another).

Somewhat surprisingly, expander graphs have found numerous applications here; they have been used to design efficient and robust error-correcting codes. The idea was first introduced in [2] and is as follows. We will construct a sparse bipartite graph $G = (V = L \cup R, E)$ with n vertices on one side of the bipartition L, and $n - m$ vertices on the other side R. The parity-check matrix $M \in \{0,1\}^{(n-m)\times n}$ is then defined by $M_{ij} = 1$ if and only if vertex $i \in L$ and vertex $j \in R$ are connected by an edge (this is also known as the Tutte matrix of G). It turns out that the rate and distance of the code associated with this parity-check matrix M is intimately related to the expansion properties of the graph G. This construction also, for example, gives a code with $\Omega(1)$ rate and $\Omega(n)$ distance; note the stark contrast between the distance achieved here and the naive "copy-the-message-several-times" approach, which achieved $\Omega(1)$ rate but $\mathcal{O}(1)$ distance. For more details, we refer the reader to [2], [13], and [20].

9.2.3 Additional Applications

Other applications of expander graphs include the theory of pseudorandomness and embeddings of finite-point metric spaces into Euclidean space with small dimension. For more applications, see [20].

9.3 Prior Work on Ramanujan Graphs

The seminal work of Alon-Boppana showed that for any infinite sequence of simple undirected d-regular graphs of size tending to infinity, the largest nontrivial eigenvalue in absolute value tends to at least $2\sqrt{d-1}$; that is, asymptotically, the $2\sqrt{d-1}$ bound is optimal. This is why Ramanujan graphs are particularly interesting: they are in this sense, "optimal" expanders.

Theorem 30 (Alon-Boppana Bound; [4] and [8])**.** *For a fixed d, if $\{X_n\}$ is a sequence of simple undirected graphs with X_n having n vertices and maximum degree d, then $\liminf_{n\to\infty} |\lambda(X_n)| \geq 2\sqrt{d-1}$, where $\lambda(G) = \max_{1\leq i\leq n:\lambda_i(G)\neq\pm d} |\lambda_i(G)|$ is the largest nontrivial eigenvalue in absolute value.*

Simultaneously, we also believe that random d-regular graphs (in a model defined in the following theorem) are Ramanujan with "reasonable" probability, for any $d > 1$. The first step towards this was the work of Friedman, who showed that random d-regular (multi-)graphs are "almost" Ramanujan with high probability.

Theorem 31 (Theorem 1.1 from [16]). *Fix $\epsilon > 0$ and an even positive integer d. Then there are a constants $c, c' > 0$ such that for a random multigraph G, formed by choosing $d/2$ uniformly random permutations $\pi_1, \ldots, \pi_{d/2}$ on $[n]$ independently and taking the edge set to be $\{(i, \pi_j(i)), (i, \pi_j^{-1}(i))\}_{i \in [n], j \in [d/2]}$, we have*

$$\mathbb{P}[\forall i > 1, |\lambda_i(G)| \leq 2\sqrt{d-1} + \epsilon] \geq 1 - \frac{c}{n^{\lceil (\sqrt{d-1}+1)/2 \rceil - 1}}$$

and

$$\mathbb{P}[\lambda_2(G) > 2\sqrt{d-1}] \geq \frac{c'}{n^{\lfloor (\sqrt{d-1}+1)/2 \rfloor}}$$

Empirical work calculations has also shown that Ramanujan graphs actually exist in great abundance. Specifically, the work of Miller-Novikoff-Sabelli [25] empirically showed that the largest nontrivial positive eigenvalue and the smallest nontrivial negative eigenvalue of a d-regular graph can be well-modeled by a Tracy-Widom distribution. As a result, it is conjectured that if the mean and standard deviation of the corresponding Tracy-Widom distribution grow asymptotically in a desirable way as the number of vertices n grows, then:

1. approximately 52% of d-regular graphs from families of bipartite d-regular graphs should be be Ramanujan
2. approximately 27% of d-regular graphs from families of nonbipartite d-regular graphs should be be Ramanujan

Tracy-Widom distributions commonly seen in random matrix theory, specifically when one examines the spectral distributions of random matrix ensembles such as the Gaussian Unitary Ensemble. Such distributions are beyond the scope of this survey and we will not discuss further their results here. We refer the reader to [25].

Finally, there are has also been work on explicitly constructing infinite families of Ramanujan graphs for a fixed degree. Prior works [5], [7], [11] were group and number theoretical constructions, and hence, only worked for d-regular graphs of specific (but infinitely many) degrees d. However, they had the benefit of being "explicit", in the sense that:

1. they were Cayley graphs and can be described by a set of group generators
2. the graphs were efficiently constructable: one can query the set of neighbors of any vertex in polynomial time
3. they did not rely on bipartiteness

Theorem 32 ([5]). *Let p, q be unequal prime numbers congruent to 1 modulo 4. Then there exists a $(p+1)$-regular (not necessarily) Ramanujan graph with $n = q(q^2-1)$ or $n = q(q^2-1)/2$ vertices. In particular, for every prime p congruent to 1 modulo 4, there exists an infinite family of $(p+1)$-regular Ramanujan graphs.*

Theorem 33 (Informal; [11]). *For every prime p and positive integer m, there exists an infinite family of (p^m+1)-regular Ramanujan graph.*

Again, we will not discuss these constructions.

9.4 2-Covers

As stated earlier, the argument is based on a "lifting" idea proposed by Bilu-Linial [19], who used it to give a deterministic polynomial time algorithm to construct arbitrarily large expanders. These expanders produced have nontrivial eigenvalues bounded in absolute value by $\mathcal{O}\left(\sqrt{d \log^3 d}\right)$. Specifically, we will show that every d-regular Ramanujan graph can be "lifted" to a d-regular Ramanujan graph that has twice as

many vertices; applying a sequence of lifts then enables us to build arbitrarily large d-regular Ramanujan graphs starting with a small Ramanujan graph.

Let's first formalize what a "lift" means. These are also known as "covers". While "lift" has a nice intuitive meaning for the process being described, "cover" is, in a sense, a more appropriate name, stemming from "covering spaces" in topology. From now on, we will stick with "cover".

Definition 17. Let $G = (V, E)$ be a graph, with $m = |E|$, and $s : E \to \{\pm 1\}$ (or $s \in \{\pm 1\}^m$ more commonly), be a signing of the edges of G. Then the **2-cover** G_s **of** G **associated with** s is the graph $G_s = (V_s, E_s)$ whose vertices and edges can be labeled as $V_s = \{u_L, u_R : u \in V\}$ and $E_s = \{(u_L, v_L), (u_R, v_R) : (u, v) \in E, s(u, v) = 1\} \cup \{(u_L, v_R), (u_R, v_L) : (u, v) \in E, s(u, v) = -1\}$.

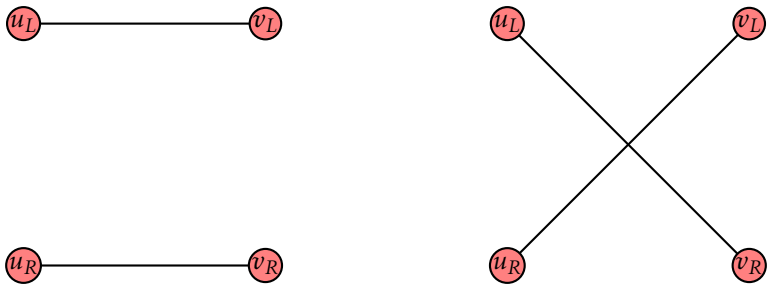

Figure 2: These are the possible 2-covers of a graph with two vertices u, v and an edge connecting them. The 2-cover on the left corresponds to the sign $+1$, while the 2-cover on the right corresponds to the sign -1. This idea extends more generally: every edge (with endpoints) of a 2-cover G_s of G corresponding to the signing $s \in \{\pm 1\}^m$ is replaced with either a copy of the left (corresponding to $+1$ on that edge) or a copy of the right (corresponding to -1 on that edge).

Lemma 20. *If $G = (V, E)$ is a bipartite graph, then so is G_s for every 2-cover G_s of G.*

It was shown that the eigenvalues of a 2-cover can be related to the eigenvalues of the base graph.

Lemma 21 (Lemma 3.1 from [19])**.** *Let G be a graph and A be its adjacency matrix, and A_s be the signed adjacency matrix of A w.r.t. a signing $s \in \{\pm 1\}^m$ of the edges. Let G_s be the 2-cover of G associated with the signing s and B be its adjacency matrix. Then every eigenvalue of A and every eigenvalue of A_s is an eigenvalue of B. Furthermore, the multiplicity of each eigenvalue of B is the sum of its multiplicities in A and A_s.*

We will refer to the eigenvalues of A_s (with multiplicity) as the "new eigenvalues", and the eigenvalues (with multiplicity) of A as the "old eigenvalues". The following conjecture formalizes the idea behind the construction.

Conjecture 2 (Conjecture 3.1 from [19])**.** *For every d-regular graph G, there exists a 2-cover $\tilde{G}$ such that all new eigenvalues, i.e. the eigenvalues of A_s, where s is the signing associated with the 2-lift, are in the range $[-2\sqrt{d-1}, 2\sqrt{d-1}]$.*

We will present an affirmative resolution for the bipartite case using the technique of interlacing families.

Theorem 34. *Let $G = (V, E)$ be a d-regular bipartite graph. Then there exists a signing $s \in \{\pm 1\}^m$ such that the corresponding 2-cover $G_s = (V_s, E_s)$ has all new eigenvalues bounded in the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. In particular, if G is a d-regular bipartite Ramanujan graph, then so is G_s, for some signing $s \in \{\pm 1\}^m$.*

9.5 Interlacing Families and the Matching Polynomial

Let $G = (V, E)$, with $|V| = n$ and $|E| = m$, be a d-regular graph. Our goal is to show that in the set of edge signings $\{\pm 1\}^m$, there exists a signing s such that the eigenvalues of A_s are bounded in $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. To do this, we will:

1. show that $\{\det(xI - A_s)\}_{s\in\{\pm 1\}^m}$ forms an interlacing family
2. bound the maximum root of $\sum_{s\in\{\pm 1\}^m} \det(xI - A_s)$ by $2\sqrt{d-1}$

Note that these combined do **not** imply the desired result. This is because step 2 only upper bounds the maximum root; it does not lower bound the minimum root. Both are required since we are bounding the largest root in absolute value. This is where we crucially use the bipartiteness of G, since all 2-covers of a bipartite graph are bipartite (Lemma 20), and the spectrum of a bipartite graph is symmetric about zero (Theorem 5); an upper bound on the maximum nontrivial eigenvalue immediately gives a lower bound on the minimum nontrivial eigenvalue.

We will do step 2 first. It turns out much of the work for this step had already been done. We will relate $\sum_{s\in\{\pm 1\}^m} \det(xI - A_s)$ to the matching polynomial $\mathcal{M}_G(x)$ of G, which is defined by

$$\mathcal{M}_G(x) \stackrel{\text{def}}{=} \sum_{k=0}^{\lfloor n/2\rfloor} (-1)^k m_k x^{n-2k}$$

where m_k is the number of matchings with k edges in G (with the convention that $m_0 = 1$). From this, we will apply a well-known bound due to Heilmann-Lieb [3] on the maximum root of $\mathcal{M}_G(x)$.

Theorem 35 (Theorem 3.6 from [34])**.**

$$\mathbb{E}_{s\sim\text{unif}(\{\pm 1\}^m)}[\det(xI - A_s)] = \mathcal{M}_G(x)$$

Proof. For convenience, we drop the subscript for the expectation, as all expectations will be over $\text{unif}(\{\pm 1\}^m)$. Recall that we may expand the determinant out as:

$$\begin{aligned}\det(xI - A_s) &= \sum_{\sigma\in\mathcal{S}_n} \text{sign}(\sigma) \prod_{i=1}^n (xI - A_s)_{i,\sigma(i)} = \sum_{\sigma\in\mathcal{S}_n} \text{sign}(\sigma) \prod_{i=1}^n (x - A_s(i,\sigma(i))) \\ &= \sum_{\sigma\in\mathcal{S}_n} \text{sign}(\sigma) \sum_{k=0}^n x^{n-k} \sum_{S\subset[n]:|S|=k} \left[\prod_{i\in S} A_s(i,\sigma(i))\right]\end{aligned}$$

Hence, applying linearity of expectation,

$$\begin{aligned}\mathbb{E}_s[\det(xI - A_s)] &= \sum_{\sigma\in\mathcal{S}_n} \text{sign}(\sigma) \sum_{k=0}^n x^{n-k} \sum_{S\subset[n]:|S|=k} \mathbb{E}_s\left[\prod_{i\in S} A_s(i,\sigma(i))\right] \\ &= \sum_{k=0}^n x^{n-k} \sum_{S\subset[n]:|S|=k} \sum_{\sigma\in\mathcal{S}_n} \text{sign}(\sigma)\mathbb{E}_s\left[\prod_{i\in S} A_s(i,\sigma(i))\right] \\ &= \sum_{k=0}^n x^{n-k} \sum_{S\subset[n]:|S|=k} \sum_{\sigma\in\mathcal{S}_n:(i,\sigma(i))\in E,\forall i\in S} \text{sign}(\sigma)\mathbb{E}_s\left[\prod_{i\in S} s(i,\sigma(i))\right]\end{aligned}$$

Note, we cannot take a step further and simplify this last expression into

$$\sum_{k=0}^n x^{n-k} \sum_{S\subset[n]:|S|=k} \sum_{\sigma\in\mathcal{S}_n:(i,\sigma(i))\in E,\forall i\in S} \left[\text{sign}(\sigma) \prod_{i\in S} \mathbb{E}_s[s(i,\sigma(i))]\right]$$

because, while we do know that $s(i,\sigma(i))$ is independent of $\sigma(j,\sigma(j))$ for $(i,\sigma(i)) \neq (j,\sigma(j))$, the product may include $(i,\sigma(i))$ and $(\sigma(i),i)$, which are perfectly correlated.

Now, let's look at each expectation $\mathbb{E}_s[\prod_{i\in S} s(i,\sigma(i))]$ for any fixed $S \subset [n]$ and $\sigma \in \mathcal{S}_n$ with $(i,\sigma(i)) \in E$ for all $i \in S$. Note the number of times the sign for any fixed $(i,\sigma(i))$ appears in the product is in $\{0,1,2\}$

(since at most, the product accounts for both $(i,\sigma(i))$ and $(\sigma(i),i)$, when both $i,\sigma(i)\in S$, where we have $s(i,\sigma(i)) = s(\sigma(i),i)$). Furthermore, recall that the signs assigned to each edge are independent and uniformly random from $\{\pm 1\}$. Hence, the expectation is zero if the sign for any edge $(i,\sigma(i))$ appears exactly once in the product.

In particular, the expectation is nonzero only when $\sigma \in \mathcal{S}_n$ satisfies $(i,\sigma(i)) \in E$ for all $i \in S$, as well as both $i,\sigma(i)\in S$, in which case the expectation is 1 (since $s(i,\sigma(i))^0 = s(i,\sigma(i))^2 = 1$ for all i). These are precisely the perfect matchings on S, i.e. matchings on $[n]$ such that every vertex of S is matched to another vertex of S, and no vertex outside of S is matched. Note such permutations also have $\operatorname{sign}(\sigma) = (-1)^{|S|/2}$. Thus, we can write the entire sum as

$$\begin{aligned}\sum_{k=0}^{n} x^{n-k} \sum_{S\subset[n]:|S|=k,k \text{ even}} \sum_{\text{perfect matching } \sigma\in\mathcal{S}_n \text{ on } S} (-1)^{k/2} &= \sum_{k=0}^{\lfloor n/2\rfloor} (-1)^k x^{n-2k} \sum_{S\subset[n]:|S|=k} \sum_{\text{perfect matching } \sigma\in\mathcal{S}_n \text{ on } S} 1 \\ &= \sum_{k=0}^{\lfloor n/2\rfloor} (-1)^k m_k x^{n-2k} = \mathcal{M}_G(x)\end{aligned}$$

□

Theorem 36 (Theorems 4.2 and 4.3 from [3]). *Let $G=(V,E)$ be a graph with maximum degree d. Then $\mathcal{M}_G(x)$ is real-rooted and* $\operatorname{maxroot}(\mathcal{M}_G(x)) \le 2\sqrt{d-1}$.

Now, it remains to prove that $\{\det(xI - A_s)\}_{s\in\{\pm1\}^m}$ forms an interlacing family. The main technical work lies in the following theorem.

Theorem 37 (Theorem 5.1 from [34]). *Let $p_1,\dots,p_m \in [0,1]$. Then*

$$\sum_{s\in\{\pm1\}^m} \left(\prod_{i:s_i=1} p_i\right)\left(\prod_{i:s_i=-1}(1-p_i)\right)\det(xI-A_s)$$

is real-rooted.

Corollary 12 (Theorem 5.2 from [34]). *The polynomials $\{\det(xI-A_s)\}_{s\in\{\pm1\}^m}$ form an interlacing family.*

Proof. For convenience, for a signing $s=(s_1,\dots,s_m)\in\{\pm1\}^m$, set $f_s(x) = f_{s_1,\dots,s_m}(x) = \det(xI-A_s)$. By the definition of an interlacing family, we must show that for every partial assignment $(s_1,\dots,s_k)\in\{\pm1\}^k$, and every $\lambda\in[0,1]$, the convex combination

$$\begin{aligned}q_{s_1,\dots,s_k}(x) = \lambda f_{s_1,\dots,s_k,1}(x) + (1-\lambda) f_{s_1,\dots,s_k,-1}(x) &= \lambda \sum_{s'_{k+2}\in\{\pm1\},\dots,s'_m\in\{\pm1\}} f_{s_1,\dots,s_k,1,s'_{k+2},\dots,s'_m}(x) \\ &\quad + (1-\lambda) \sum_{s'_{k+2}\in\{\pm1\},\dots,s'_m\in\{\pm1\}} f_{s_1,\dots,s_k,-1,s'_{k+2},\dots,s'_m}(x)\end{aligned}$$

is real-rooted. It suffices to show that $q_{s_1,\dots,s_k}(x)/2^{m-k-1}$ is real-rooted. For this, observe that we may write

$$\frac{1}{2^{m-k-1}} q_{s_1,\dots,s_k}(x) = \sum_{s\in\{\pm1\}^m} \left(\prod_{i:s_i=1} p_i\right)\left(\prod_{i:s_i=-1}(1-p_i)\right) f_s(x)$$

with

$$p_i = \begin{cases} (1+s_i)/2 & \text{if } 1\le i\le k \\ \lambda & \text{if } i = k+1 \\ 1/2 & \text{otherwise}\end{cases}$$

Case 1 corresponds to the partial assignment that we've already fixed $s_1,\dots,s_k$. We have $p_i = 1$ if $s_i = 1$ and $1 - p_i = 1$ if $s_i = -1$. Case 2 corresponds to the current convex combination weights λ on $f_{s_1,\dots,s_k,1}$ and $f_{s_1,\dots,s_k,-1}$. Case 3 corresponds to the remaining signs; this is where we distribute the $1/2^{m-k-1}$.

Applying Theorem 37 gives the result. □

It remains to prove Theorem 37. For this, we will need the real-rootedness of more general class of polynomials.

Theorem 38 (Theorem 6.6 from [34]). *Let $a_1,\dots,a_m,b_1,\dots,b_m \in \mathbb{R}^n$, $p_1,\dots,p_m \in [0,1]$ and D be a Hermitian positive semidefinite matrix. Then*

$$\sum_{S\subset[m]}\left(\prod_{i\in S} p_i\right)\left(\prod_{i\notin S}(1-p_i)\right)\det\left(xI + D + \sum_{i\in S} a_i a_i^* + \sum_{i\notin S} b_i b_i^*\right)$$

is real-rooted.

Proof. An easy way is to use the machinery we already developed in resolving the Kadison-Singer problem. Specifically, we observe that the polynomial of interest is essentially a mixed characteristic polynomial.

Consider distributions $\mathcal{D}_i$ over $\{a_i, b_i\}$, with $\mathbb{P}_{v\sim\mathcal{D}_i}[v = a_i] = p_i$ and $\mathbb{P}_{v\sim\mathcal{D}_i}[v = b_i] = 1 - p_i$. Let $A_i = p_i a_i a_i^* + (1-p_i) b_i b_i^*$. We will draw v_i from $\mathcal{D}_i$ for all $1 \le i \le m$ independently. Corollary 10 shows

$$\begin{aligned}\mu[A_1,\dots,A_m](x) &= \mathbb{E}_{v_1\sim\mathcal{D}_1,\dots,v_m\sim\mathcal{D}_m}\det\left(xI - D - \sum_{i=1}^m v_i v_i^*\right)\\ &= \sum_{S\subset[m]}\left(\prod_{i\in S} p_i\right)\left(\prod_{i\notin S}(1-p_i)\right)\det\left(xI - D - \sum_{i\in S} a_i a_i^* - \sum_{i\notin S} b_i b_i^*\right)\end{aligned}$$

is real-rooted. Finally, using the fact that $\det(-A) = (-1)^n\det(A)$ for $A \in \mathbb{C}^{n\times n}$, we see that the original polynomial of interest may be written as

$$(-1)^n\sum_{S\subset[m]}\left(\prod_{i\in S} p_i\right)\left(\prod_{i\notin S}(1-p_i)\right)\det\left(-xI - D - \sum_{i\in S} a_i a_i^* - \sum_{i\notin S} b_i b_i^*\right) = (-1)^n\mu[A_1,\dots,A_m](-x)$$

we have the result. □

We are now ready to prove Theorem 37.

Proof of Theorem 37. Let d_v be the degree of v in G, and $d = \max_v d_v$ be the maximum degree of G. To show that

$$\sum_{s\in\{\pm1\}^m}\left(\prod_{i:s_i=1} p_i\right)\left(\prod_{i:s_i=-1}(1-p_i)\right)\det(xI - A_s)$$

is real-rooted, it suffices to show that

$$\sum_{s\in\{\pm1\}^m}\left(\prod_{i:s_i=1} p_i\right)\left(\prod_{i:s_i=-1}(1-p_i)\right)\det(xI + dI - A_s)$$

is real-rooted, as the roots of these two polynomials differ by d.

Next, we note that $dI - A_s$ may be written as a signed Laplacian matrix of G plus a diagonal matrix with nonnegative entries, so that it is positive semidefinite. More precisely, if for each edge $(u,v) \in E$, we set

$$L^1_{u,v} = (e_u - e_v)(e_u - e_v)^*$$

if (u,v) has sign $+1$ in s, and

$$L_{u,v}^{-1} = (e_u + e_v)(e_u + e_v)^*$$

if (u,v) has sign -1 in s, then

$$dI - A_s = \sum_{(u,v)\in E} L_{u,v}^{s_{u,v}} + D$$

where $D = \mathrm{diag}(\{d - d_v\}_{v\in V})$. Thus, if we take $a_{u,v} = e_u - e_v$ and $b_{u,v} = e_u + e_v$, then our polynomial of interest becomes

$$\sum_{s\in\{\pm1\}^m} \left(\prod_{i:s_i=1} p_i\right)\left(\prod_{i:s_i=-1}(1-p_i)\right)\det\left(xI + D + \sum_{(u,v)\in E:s_{u,v}=1} a_{u,v}a_{u,v}^* + \sum_{(u,v)\in E:s_{u,v}=-1} b_{u,v}b_{u,v}^*\right)$$
$$= \sum_{S\subset E}\left(\prod_{i\in S} p_i\right)\left(\prod_{i\notin S}(1-p_i)\right)\det\left(xI + D + \sum_{(u,v)\in S} a_{u,v}a_{u,v}^* + \sum_{(u,v)\in E\setminus S} b_{u,v}b_{u,v}^*\right)$$

where we use S to denote the set of edges signed with $+1$. It follows immediately from Theorem 38 that this polynomial is real-rooted. □

Proof of Theorem 34. Corollary 12 shows that $\{\det(xI - A_s)\}_{s\in\{\pm1\}^m}$ forms an interlacing family and so there exists a signing $\tilde{s}\in\{\pm1\}^m$ such that $\mathrm{maxroot}(\det(xI - A_{\tilde{s}})) \le \mathrm{maxroot}(\mathbb{E}_{s\sim\mathrm{unif}(\{\pm1\}^m)}[\det(xI - A_s)])$. Since the expectation on the right-hand side is $\mathcal{M}_G(x)$ by Theorem 35, which has maximum root upper bounded by $2\sqrt{d-1}$ by Theorem 36, the adjacency matrix of the 2-cover $G_{\tilde{s}}$ has all eigenvalues bounded in $[-2\sqrt{d-1}, 2\sqrt{d-1}]$ as desired. □

9.6 r-Covers

The proof above shows how to construct simple d-regular bipartite Ramanujan graphs for every degree $d > 1$, but only of specific sizes, namely powers of 2 multiplied by the size of the base graph. Ideally, we'd like to be able to strengthen the argument to show the existence of d-regular bipartite Ramanujan graphs of arbitrary size. Now, 2-covers correspond to signings of edges of the base graph. By viewing the set of signs $\{\pm1\}$ as the set of permutations on 2 item S_2, one can see how to generalize the notion of a 2-covers to that of an r-covers, for $r \ge 1$.

Definition 18. Let $G = (V,E)$ be a graph, with $m = |E|$, and $\sigma : E \to S_r$. Then the **r-cover G_σ of G associated with** σ is the graph $G_\sigma = (V_\sigma, E_\sigma)$ whose vertices and edges can be labeled as $V_\sigma = \{v_1,\dots,v_r : v\in V\}$ and $E_\sigma = \{(u_1, v_{\sigma(1)}),\dots,(u_r, v_{\sigma(r)}) : (u,v)\in E\}$. The set of r-covers of G is denoted $\mathcal{C}_{r,G}$.

Lemma 22. *If $G = (V,E)$ is a bipartite graph, then so is G_σ for every r-cover G_σ of G, for every $r > 1$.*

Then, it is natural to ask if every simple d-regular graph has a good r-lift, for any $r \ge 1$. If this were true, then we would indeed be able to show the existence of d-regular bipartite Ramanujan graphs of all sizes. In this section, we'll give a very high level sketch of how to resolve this question affirmatively. For more detail, we refer to [41].

Before we continue, we need an analogue of the signed adjacency matrix we used for 2-covers in order to talk about "new eigenvalues". It doesn't seem immediately obvious (at least if one doesn't know representation theory) how to generalize the notion of a signed adjacency matrix, since in this special case, all we had to do was change the sign of the entries of A appropriately.

However, what we really are doing is encoding the information about the edge signs in a matrix. Now, we just have to encode the information about the permutations described by $\sigma : E \to S_r$ in a matrix for more general $r \ge 1$. As it turns out, the right tool here to use is that of group representations. While representation theory well lies outside the scope, we'll just give the definitions needed to state the main results of [41].

9.6.1 A Quick Detour Into Group Representations

Recall that a group G is a set with a binary operation $\cdot : G \times G \to G$ (the group law of G) such that

1. $a \cdot b \in G$ for all $a, b \in G$ (closure).
2. There exists an element $e \in G$ such that $a \cdot e = e \cdot a = a$ for all $a \in G$ (identity).
3. For every $a \in G$, there exists an element $a^{-1} \in G$ such that $a \cdot a^{-1} = a^{-1} \cdot a = e$ (inverses).

A group is abelian if $a \cdot b = b \cdot a$ for all $a, b \in G$. A map $\varphi : G \to H$ where G, H are group is a group homomorphism if $\varphi(a \cdot b) = \varphi(a) \cdot \varphi(b)$ for all $a, b \in G$. Finally, if G is a group and S is a set, then a left (resp. right) group action is a function $\varphi : G \times X \to X$ (resp. $\varphi : X \times G \to X$) satisfying

1. $\varphi(e, x) = x$ (resp. $\varphi(x, e) = x$) for all $x \in X$
2. $\varphi(ab, x) = \varphi(a, \varphi(b, x))$ (resp. $\varphi(x, ab) = \varphi(\varphi(x, a), b)$) for all $a, b \in G, x \in X$

Definition 19. A **d-dimensional representation** of a finite group Γ is a group homomorphism $\pi : \Gamma \to \mathrm{GL}(V)$, where V is a d-dimensional vector space and $\mathrm{GL}(V)$ is the group of invertible linear transformations from V to itself (with function composition as the group operation). For $d \in \mathbb{N}$, a **d-dimensional complex representation** of Γ is a representation where $V = \mathbb{C}^d$ and $\mathrm{GL}(V) = \mathrm{GL}_d(\mathbb{C})$ is the set of invertible $d \times d$ complex matrices. If Γ additionally has a topological structure (i.e. it is a topological group), then we require π to also be continuous.

Two representations $\pi_1, \pi_2 : \Gamma \to \mathrm{GL}_d(\mathbb{C})$ are **isomorphic** if there exists $B \in \mathrm{GL}_d(\mathbb{C})$ such that $\pi_1(g) = B^{-1}\pi_2(g)B$ for all $g \in \Gamma$. For two representations $\pi_1 : \Gamma \to \mathrm{GL}_{d_1}(\mathbb{C})$ and $\pi_2 : \Gamma \to \mathrm{GL}_{d_2}(\mathbb{C})$, we define their **direct sum** $\pi_1 \oplus \pi_2 : \Gamma \to \mathrm{GL}_{d_1+d_2}(\mathbb{C})$ by

$$(\pi_1 \oplus \pi_2)(g) = \begin{pmatrix} \pi_1(g) & 0 \\ 0 & \pi_2(g) \end{pmatrix}$$

A representation is **irreducible** if it is not isomorphic to a sum of representations. A representation is **unitary** if it is isomorphic to a representation whose image is a subgroup of the group of unitary matrices. A representation is **faithful** if it is injective.

Example 4. The **trivial representation** $\pi_{\mathrm{triv}} : \Gamma \to \mathrm{GL}_1(\mathbb{C}) = \mathbb{C}$ is given by $g \mapsto 1$ for all $g \in \Gamma$.

Example 5. Consider the action of Γ on itself by right multiplication, namely $g(h) = hg$. If we view $\mathbb{C}^{|\Gamma|}$ as the set of complex-valued functions on Γ, then the action gives a representation $\pi_{\mathrm{reg}} : \Gamma \to \mathrm{GL}_{|\Gamma|}(\mathbb{C})$ such that $\pi_{\mathrm{reg}}(g)$ is a (permutation) matrix satisfying $\pi_{\mathrm{reg}}(g)v(h) = v(hg)$ for all complex-valued functions $v \in \mathbb{C}^{|\Gamma|}$, for all $g, h \in \Gamma$. π_{reg} is known as the **(right) regular** representation of Γ. One can similarly define a **(left) regular** representation of Γ by considering the group action of Γ on itself by left multiplication. For $\Gamma = S_r$, these two representations coincide.

Example 6. Consider the regular representation π_{reg} of S_r, which maps $\sigma \in S_r$ to the permutation matrix corresponding to it. Since the subspace $\mathrm{span}\{\mathbf{1}\}$ of $\mathbb{C}^r$ is invariant under all permutation matrices, π_{reg} is isomorphic to the trivial representation from Γ to $\mathrm{span}\{\mathbf{1}\}$. Hence, the action of π_{reg} on $\mathrm{span}\{\mathbf{1}\}^{\perp}$ is an $(r-1)$-dimensional representation called the **standard representation** of Γ and is denoted π_{std}.

For example, the standard representation of $S_2 = \{\mathrm{id}, \mathrm{swap}\}$ is simply $\pi_{\mathrm{std}}(\mathrm{id}) = 1$ and $\pi_{\mathrm{std}}(\mathrm{swap}) = -1$. The standard representation of S_3 is given by

$$\mathrm{id} \mapsto I_2 \qquad (123) \mapsto \begin{pmatrix} e^{2\pi i/3} & 0 \\ 0 & e^{-2\pi i/3} \end{pmatrix} \qquad (132) \mapsto \begin{pmatrix} e^{-2\pi i/3} & 0 \\ 0 & e^{2\pi i/3} \end{pmatrix}$$

$$(12) \mapsto \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad (23) \mapsto \begin{pmatrix} 0 & e^{-2\pi i/3} \\ e^{2\pi i/3} & 0 \end{pmatrix} \qquad (13) \mapsto \begin{pmatrix} 0 & e^{2\pi i/3} \\ e^{-2\pi i/3} & 0 \end{pmatrix}$$

With this final example, we can now define the analogue of the signed adjacency matrix for an r-cover $\sigma : E \to S_r$. Specifically, let G be a graph, $r \geq 1$, and fix an ordering on the vertices $V = [n]$ of G. For a $\sigma : E \to S_r$, we define

$$A_\sigma = \begin{pmatrix} \pi_{\text{std}}(\sigma(1,1)) & \dots & \pi_{\text{std}}(\sigma(n,1)) \\ \vdots & \ddots & \vdots \\ \pi_{\text{std}}(\sigma(1,n)) & \dots & \pi_{\text{std}}(\sigma(n,n)) \end{pmatrix} \in \mathbb{C}^{n(r-1)\times n(r-1)}$$

where we make the abuse of notation by still writing $\pi_{\text{std}}(\sigma(u,v))$ even if $(u,v) \notin E$; in this case, $\pi_{\text{std}}(\sigma(u,v)) = 0$. Observe that when $r = 2$ and $\sigma : E \to S_2$ corresponds to a signing, we recover the original notion of a signed adjacency matrix.

9.6.2 Interlacing Families and the d-Matching Polynomial

The key consequence of [41] as we alluded to earlier is the following.

Theorem 39. *Let $G = (V,E)$ be a d-regular bipartite graph. Let $r \geq 1$ be arbitrary. Then there exists a $\sigma : E \to S_r$ such that the corresponding r-cover $G_\sigma = (V_\sigma, E_\sigma)$ has all new eigenvalues bounded in the interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$. In particular, if G is a d-regular bipartite Ramanujan graph, then so is G_σ, for some $\sigma : E \to S_r$.*

To prove this, the same strategy works as in the special case of $r = 2$. Specifically, we will:

1. show that $\{\det(xI - A_\sigma)\}_{\sigma \in S_r^E}$ forms an interlacing family
2. bound the maximum root of $\sum_{\sigma \in S_r^E} \det(xI - A_\sigma)$ by $2\sqrt{d-1}$

For step 2, we will relate

$$\sum_{\sigma \in S_r^E} \det(xI - A_\sigma)$$

to a special polynomial whose roots we know how to analyze. When $r = 2$, we use the matching polynomial of a graph. For $r > 2$, we need a generalization (yet again). [41] proposed to use what they call the k-matching polynomial, which for a general graph G is defined by

$$\mathcal{M}_{k,G}(x) = \mathbb{E}_{H \sim \text{unif}(\mathcal{C}_{k,G})}[\mathcal{M}_H(x)]$$

for $k = r - 1$. In the case when $r = 2$, we recover the regular matching polynomial of G: $\mathcal{M}_{1,G}(x) = \mathcal{M}_G(x)$.

Theorem 40 (Theorem 1.8 from [41])**.** *Let G be a simple, connected d-regular graph. Let $r \in \mathbb{N}$. Then*

$$\mathbb{E}_{\sigma \sim \text{unif}(S_r^E)}[\det(xI - A_\sigma)] = \mathcal{M}_{r-1,G}(x)$$

Theorem 41 (Theorem 2.7 from [41])**.** *If G is a simple, connected d-regular graph, and $k \in \mathbb{N}$, then $\mathcal{M}_{k,G}(x)$ is real-rooted. Furthermore, all root of $\mathcal{M}_{k,G}(x)$ lie in $[-2\sqrt{d-1}, 2\sqrt{d-1}]$.*

We refer the details of the proofs of these statements to [41]. One key thing to note is that we have reduced and rephrased the results of [41] to the special case when:

1. G is d-regular and simple; many results of [41] allow G to have multiedges and do not assume regularity (they phrase their results in the language of the universal cover of G)
2. the lifts are given by labeling the edges of G with elements from the symmetric group S_r and representing them in A_σ by block matrices given by the standard representation of S_r; the results of [41] allow labelings from a finite group Γ other than S_r along with a complex representation $\pi : \Gamma \to \text{GL}_{r-1}(\mathbb{C})$, subject to some nice conditions that we will not present

9.7 Kadison-Singer and Ramanujan Coverings

Now, let's look at an interesting connection between the Kadison-Singer Problem and Ramanujan graphs. Let $G = (V, E)$ be a d-regular graph. Recall the proof of Theorem 34. Recall we defined $a_{u,v} = e_u - e_v$ and $b_{u,v} = e_u + e_v$, and showed that

$$dI - A_s = \sum_{(u,v)\in E: s_{u,v}=1} a_{u,v}a_{u,v}^* + \sum_{(u,v)\in E: s_{u,v}=-1} b_{u,v}b_{u,v}^*$$

Then

$$\mathbb{E}_{s\sim\text{unif}(\{\pm 1\}^m)}\left[\sum_{(u,v)\in E: s_{u,v}=1} a_{u,v}a_{u,v}^* + \sum_{(u,v)\in E: s_{u,v}=-1} b_{u,v}b_{u,v}^*\right] = \mathbb{E}_{s\sim\text{unif}(\{\pm 1\}^m)}[dI - A_s] = dI - \mathbb{E}_{s\sim\text{unif}(\{\pm 1\}^m)}[A_s] = dI$$

More compactly, if $r_{u,v} = a_{u,v}/\sqrt{d}$ with probability 1/2 and $r_{u,v} = b_{u,v}/\sqrt{d}$ with probability 1/2, then

$$\mathbb{E}\left[\sum_{(u,v)\in E} r_{u,v}r_{u,v}^*\right] = I$$

Now, $\left\|a_{u,v}/\sqrt{d}\right\|^2 = \left\|b_{u,v}/\sqrt{d}\right\|^2 = 2/d$. Hence, if we formulate this entire problem of demonstrating there exists a signing $s \in \{\pm 1\}^m$ as an instance of Theorem 21 that implied Kadison-Singer, specifically with $\epsilon = 2/d$ and $\mathcal{D}_{u,v} = \text{unif}(\{a_{u,v}/\sqrt{d}, b_{u,v}/\sqrt{d}\})$ for all $(u,v) \in E$, then we have

$$\mathbb{P}_s\left[\|dI - A_s\| \le \left(1+\sqrt{\frac{2}{d}}\right)^2\right] = \mathbb{P}_s\left[\left\|\sum_{(u,v)\in E} r_{u,v}r_{u,v}^*\right\| \le \left(1+\sqrt{\frac{2}{d}}\right)^2\right] > 0$$

that is, there exists a signing such that $\lambda_{\max}(dI - A_s) \le d(1+\sqrt{2/d})^2 = d + 2 + 2\sqrt{2d} \implies \lambda_{\max}(A_s) \le 2 + 2\sqrt{2d}$. Note that asymptotically, this is optimal (but suboptimal w.r.t. the $2\sqrt{d-1}$ bound we obtained using the matching polynomial; see Theorem 35 and Theorem 36).

What is particularly interesting about this reduction is that this shows the dependence on ϵ in the bound

$$\mathbb{P}\left[\left\|\sum_{i=1}^m v_i v_i^*\right\| \le (1+\sqrt{\epsilon})^2\right] > 0$$

is asymptotically optimal; if we could do better, then we'd be able to show the existence of the existence of a signing $s \in \{\pm 1\}^m$ such that $\lambda_{\max}(A_s) \le o(\sqrt{d})$. This would immediately contradict the Alon-Boppana bound (Theorem 30).

9.8 Shortcomings of this Proof

9.8.1 The Natural Algorithm is Inefficient

Theorem 35 shows that

$$\mathbb{E}_{s\sim\text{unif}(\{\pm 1\}^m)}[\det(xI - A_s)] = \mathcal{M}_G(x)$$

In particular, the root of the interlacing family is the matching polynomial. Now, it is known that computing even the just the constant coefficient of $\mathcal{M}_G(x)$ (which is the number of perfect matchings in G) is #**P**-Hard to compute. Hence, we cannot hope to exactly and efficiently compute the coefficients of $\mathcal{M}_G(x)$. This makes computing the maximum root of $\mathcal{M}_G(x)$ much more difficult.

Now, one can say that by Theorem 36, $\text{maxroot}(\mathcal{M}_G(x)) \le 2\sqrt{d-1}$, so it doesn't really matter what $\text{maxroot}(\mathcal{M}_G(x))$

is precisely (because we only really care about the $2\sqrt{d-1}$ upper bound). However, how do we proceed to the next step? That is, how do we compute the maximum roots of p_1, p_{-1}, the children of the root $\mathcal{M}_G(x)$? Well, if could compute their coefficients efficiently, then we can also compute the coefficients of $\mathcal{M}_G(x)$ by just writing it as $p_1 + p_{-1}$. In particular, it is #**P**-Hard to compute the coefficients of p_1, p_{-1}. Now, we really are stuck, because here, we actually need to know which polynomial, p_1 or p_{-1}, has its maximum root smaller than $2\sqrt{d-1}$.

This is actually a problem for the natural interlacing family algorithm for the Kadison-Singer problem as well. Specifically, the constant term of

$$\mu[A_1, \ldots, A_m](x)$$

is given by

$$\begin{aligned}\left[\left(\prod_{i=1}^{m}(1-\partial_{z_i})\right)\det\left(xI+\sum_{i=1}^{m}z_iA_i\right)\right]_{z_1=\cdots=z_m=x=0} &= \left[\left(\prod_{i=1}^{m}(1-\partial_{z_i})\right)\det\left(\sum_{i=1}^{m}z_iA_i\right)\right]_{z_1=\cdots=z_m=0} \\ &= (-1)^m\left[\frac{\partial^m}{\partial_{z_1}\cdots\partial_{z_m}}\det\left(\sum_{i=1}^{m}z_iA_i\right)\right]_{z_1=\cdots=z_m=0}\end{aligned}$$

which is (up to a sign) the **mixed discriminant** $D(A_1, \ldots, A_m)$ **of** $A_1, \ldots, A_m$. Since $A_i = \mathbb{E}_{\mathcal{D}_i}[v_i v_i^*]$ and we make no assumptions on $\mathcal{D}_i$ other than finite support, computing $\mu[A_1, \ldots, A_m](x)$ in all cases requires being able to compute the mixed discriminant $D(A_1, \ldots, A_m)$ efficiently in all cases. Again, we cannot hope to do this as mixed discriminants can encode the permanent, which is #**P**-Hard to compute in general (for a matrix $B \in \mathbb{R}^{m\times m}$, $\operatorname{per} B = D(A_1, \ldots, A_m)$, where $A_i = \operatorname{diag}(b_{1i}, \ldots, b_{mi})$ for all columns $1 \le i \le m$). The best known polynomial-time algorithms for approximating mixed discriminants in general achieve exponential multiplicative approximation ratios; for example, [15] achieve, through a deterministic algorithm, an approximation ratio of $m^m/m! \approx e^m$.

With this in mind, we can only hope to approximate the maximum roots of the desired polynomials. This question was addressed in [43], where they gave algorithms that, given access to the top k coefficients of a degree-n polynomial, approximates the maximum root of p to within a multiplicative factor of $n^{1/k}$ when $k \le \log n$, and $1 + \mathcal{O}(k^{-2}\log^2 n)$ when $k > \log n$, respectively. These algorithms run in time polynomial in k; in the context of running the natural interlacing families algorithm, they lead to $2^{\tilde{\mathcal{O}}(n^{1/3})}$-time algorithms for finding the desired polynomials. Unfortunately, [43] also proves nearly matching lower bounds.

Another approach has been to restrict attention to special kinds of covers that are easy to work with and for which there are only polynomially many such covers. [38] considered covers known as "shift covers" (in their terminology, "shift lifts"), which are covers for which σ is a power of the "shift" permutation $(1 \cdots r) \in S_r$ (it sends 1 to 2, 2 to 3, etc. and r to 1).

9.8.2 Bipartiteness

To prove that there is a lift satisfying the Ramanujan bound, we needed to use the fact that the eigenvalues are symmetric about 0, since we are only able to directly bound the largest eigenvalue. However, we'd ideally like to be able to show that infinite families of nonbipartite Ramanujan graphs exist as well. Here, we present an attempt to remove this dependency that does **not** work, but is interesting nonetheless.

The idea here is to consider a kind of "symmetrized" characteristic polynomial Ξ_A, where the roots of this polynomial are symmetric about 0 and each root of the original characteristic polynomial χ_A is a root of Ξ_A. Then, one can attempt to use the method of interlacing polynomials in a similar fashion to prove a bound on the largest root of Ξ_A; then, by symmetry, we can conclude a bound on both the largest and

smallest roots of χ_A. Let A be a symmetric matrix with eigenvalues $\lambda_1 \leq \cdots \leq \lambda_n$. Define

$$\Xi_A(x) = \prod_{i=1}^{n}(x-\lambda_i)(x+\lambda_i) = \det(xI-A)\det(xI+A) = \det(x^2I-A^2) = \det\left(xI - \begin{bmatrix} 0 & A \\ A & 0 \end{bmatrix}\right)$$

This is the lowest degree polynomial with $\lambda_1, \ldots, \lambda_n$ as roots (with multiplicity), and whose roots are symmetric about 0. Now, if we consider the adjacency matrix A of a graph G, and a signing $s \in \{\pm 1\}^m$, then

$$\Xi_{A_s}(x) = \det\left(xI - \begin{bmatrix} 0 & A_s \\ A_s & 0 \end{bmatrix}\right)$$

This has a nice interpretation: A_s corresponds to a 2-cover G_s of G, and so $\begin{bmatrix} 0 & A_s \\ A_s & 0 \end{bmatrix}$ corresponds to the double cover of G_s.

Ideally, the next step in this idea is to show that $\{\Xi_{A_s}(x)\}_{s\in\{\pm1\}^m}$ forms an interlacing family. Unfortunately, this is actually not true. It turns that

$$\mathbb{E}_{s\sim\text{unif}(\{\pm1\}^m)}[\Xi_{A_s}(x)]$$

fails to be real-rooted, even for small graphs (a simple experiment with $K_{2,2}$ shows this average is not real-rooted).

10 Building Bipartite Ramanujan Graphs II: Unions of Matchings

In this section, we'll present a completely different construction of d-regular bipartite Ramanujan graphs. The key differences in results are that

1. we will be able to give a deterministic polynomial time algorithm for building such graphs
2. the graphs produced may have multiedges

Here, instead of taking a base d-regular bipartite Ramanujan graph and lifting it to a larger d-regular bipartite Ramanujan graph, we will just build the desired multigraph in one shot. Specifically, if given a degree d and an even size $2n$, we will be able to return, in polynomial time a d-regular bipartite Ramanujan multigraph.

The idea is similar to the model Friedman was working with to prove Theorem 31. Specifically, we'll carefully choose d permutations on $[2n]$ and then take G to be the union of the matchings corresponding to the permutations. This requires analyzing a completely different interlacing family.

For convenience, we will overload the notation of $\mathcal{S}_n$ to mean both the set (group) of permutations on n letters as well as the group of $n \times n$ permutation matrices.

The main result is as follows.

Theorem 42. *Let $P_1, \ldots, P_d \sim \mathrm{unif}(\mathcal{S}_n)$ be independent, where $d \geq 3$ and $n > d/2$. Then there exists a choice of $n \times n$ permutations $Q_1, \ldots, Q_d$ such that all nontrivial eigenvalues of the adjacency matrix*

$$A = \sum_{i=1}^{d} \begin{bmatrix} 0 & Q_i \\ Q_i^\top & 0 \end{bmatrix}$$

are upper bounded by $2\sqrt{d-1}$ in absolute value. In particular, there exists a choice of d perfect matchings on $2n$ vertices $[2n] = V$ such that their union is a d-regular bipartite Ramanujan multigraph.

The proof strategy is exactly the same as in the previous applications we've seen. We'll begin with the first step of proving that the polynomials we are interested in form an interlacing family. This is the simpler part. Following this step, we'll spend a few subsections developing the machinery needed to bound the largest nontrivial eigenvalue of the average of the adjacency matrices of unions of d perfect matchings.

Just as in the previous method of building Ramanujan graphs, this construction relies crucially on bipartiteness, since we need the eigenvalues of the adjacency matrix to be symmetric about 0 in order to bound the smallest nontrivial eigenvalue.

Finally, unlike in previous sections, here, we will not go through all details of proofs. We will sketch their main ideas and refer the reader to [36] for long calculations, etc.

10.1 Interlacing Families for Permutations

The main goal of this subsection is to prove the following theorem.

Theorem 43. *The set of polynomials*

$$\chi\left[\sum_{i=1}^{d} \begin{bmatrix} 0 & P_i S_i^\top \\ (P_i S_i^\top)^\top & 0 \end{bmatrix}\right](x)$$

where the P_i, S_i range over $\mathcal{S}_n$, form an interlacing family.

Remark 8. This theorem is stated in a way that is nicer to work with when we try to bound the roots of the average of these polynomials later on. Of course, the product of uniformly random, independent permutations is a uniformly random permutation so what one should keep in mind is that we are really saying that the set of polynomials

$$\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & P_i \\ P_i^\top & 0 \end{bmatrix}\right](x)$$

form an interlacing family. Note there are $(n!)^d$ many of these polynomials.

We will actually prove a more general fact. But first, we need a nice condition on the distributions allowed on $\mathcal{S}_n$ which will allow us to decompose a random variable taking values in $\mathcal{S}_n$.

Definition 20 (Realizability by Swaps)**.** A **random swap** S is a matrix-valued random variable taking values in $\{I, \sigma\}$, where σ is some transposition of two indices $i \neq j$ in $[n]$. A random variable P taking values in $\mathcal{S}_n$ is **realizable by swaps** if there exist random swaps $S_1, \ldots, S_N$ such that the distribution of P is the same as the distribution of $S_N \cdot S_{N-1} \cdots S_2 \cdot S_1$.

Lemma 23. *Let $P, Q \sim \mathrm{unif}(\mathcal{S}_n)$. Then, P, Q and $P \oplus Q = \begin{pmatrix} P & 0 \\ 0 & Q \end{pmatrix}$ are realizable by swaps.*

Proof Sketch. For P, Q, see the Fisher-Yates algorithm for generating a uniformly random permutation via n random swaps. This also proves that $P \oplus I$ and $I \oplus Q$ are realizable by swaps. Finally, observe that $P \oplus Q = (P \oplus I) \cdot (I \oplus Q)$ so that $P \oplus Q$ is realizable by swaps. □

With these notions in hand, we first state a simpler "version" of our goal.

Theorem 44. *Let $k \in \mathbb{N}$ be arbitrary. Suppose $A_1, \ldots, A_d$ are symmetric $k \times k$ matrices, and $\mathcal{D}_1, \ldots, \mathcal{D}_d$ be independent distributions on the set of $k \times k$ permutation matrices such that $P_i \sim \mathcal{D}_i$ is realizable by swaps, for all $1 \leq i \leq d$. Then, the set of polynomials*

$$\left(\prod_{i=1}^{d} \mathbb{P}_{\mathcal{D}_i}[P_i]\right) \chi\left[\sum_{i=1}^{d} P_i A_i P_i^\top\right](x)$$

form an interlacing family.

Again, the key is to decompose each P_i into a product of random swap matrices. This theorem will follow immediately using Theorem 7 and the following, which is the main technical work of this subsection.

Theorem 45 (Theorem 3.3 from [36])**.** *Let $A_1, \ldots, A_d$ be symmetric matrices in $\mathbb{R}^{n \times n}$ and let $\{S_{ij}\}_{i \in [d], j \in [N]}$ be independent (not necessarily identical) random swaps. Then*

$$\mathbb{E}_{\{S_{ij}\}} \chi\left[\sum_{i=1}^{d}\left(\prod_{j=N}^{1} S_{ij}\right) A_i \left(\prod_{j=1}^{N} S_{ij}^\top\right)\right](x)$$

is real-rooted.

Proof Sketch. The first key observation is that the polynomial $p(X_1, \ldots, X_d) = \det(X_1 + \cdots + X_d)$ satisfies two important properties:

1. The univariate restriction $p(tI - A_1, \ldots, tI - A_d)$ is real-rooted for symmetric matrices $A_1, \ldots, A_d$ (referred to as **hyperbolicity** in [36]). This follows from the Spectral Theorem; for symmetric matrices $A_1, \ldots, A_d$, $p(tI - A_1, \ldots, tI - A_d)$ is the characteristic polynomial of a real symmetric matrix.

2. $$p(X_1, \ldots, X_{i-1}, X_i + svv^\top, X_{i+1}, \ldots, X_d) = p(X_1, \ldots, X_d) + s \cdot D_{i, vv^\top}[p(X_1, \ldots, X_d)]$$

for every vector $v \in \mathbb{R}^n$, every $1 \le i \le d$, and every $s \in \mathbb{R}$, where

$$D_{i,vv^\top}[p(X_1,\dots,X_d)] \stackrel{\text{def}}{=} \left[\partial_t p(X_1,\dots,X_{i-1},X_i+tvv^\top,X_{i+1},\dots,X_d)\right]_{t=0}$$

is a linear operator (referred to as **rank-1 linearity** in [36]). This can be proved using the Matrix Determinant Lemma (Lemma 6) and Corollary 4:

$$\det\left(\sum_{i=1}^d X_i + svv^\top\right) = \left(1 + sv^\top\left(\sum_{i=1}^d X_i\right)^{-1} v\right)\det\left(\sum_{i=1}^d X_i\right)$$

and

$$\left(\partial_t p(X_1,\dots,X_{i-1},X_i+tvv^\top,X_{i+1},\dots,X_m)\right)|_{t=0} = \det\left(\sum_{i=1}^d X_i\right)\mathrm{tr}\left(\left(\sum_{i=1}^d X_i\right)^{-1} vv^\top\right) = \left(v^\top\left(\sum_{i=1}^d X_i\right)^{-1} v\right)\det\left(\sum_{i=1}^d X_i\right)$$

The other key observation is that random swaps preserve these nice properties in expectation:

Lemma 24 (Lemma 3.7 from [36]). *For any random swap S and any $1 \le i \le d$, $\mathbb{E}_S[p(X_1,\dots,X_{i-1},SX_iS^\top,X_{i+1},\dots,X_d)]$ is hyperbolic and rank-1 linear.*

Once we have this, the theorem follows immediately, since applying this lemma inductively dN times, we have

$$\mathbb{E}_{\{S_{ij}\}}\det\left(\sum_{i=1}^d\left(\prod_{j=N}^1 S_{ij}\right)X_i\left(\prod_{j=1}^N S_{ij}^\top\right)\right)$$

is rank-1 linear and hyperbolic, and then restricting this polynomial to $(t/m)I - A_1,\dots,(t/m)I - A_d$ yields real-rootedness of

$$\mathbb{E}_{\{S_{ij}\}}\chi\left[\sum_{i=1}^d\left(\prod_{j=N}^1 S_{ij}\right)A_i\left(\prod_{j=1}^N S_{ij}^\top\right)\right]$$

by hyperbolicity. Hence, it remains to prove Lemma 24.

Proof Sketch of Lemma 24. Let S be any random swap equal to σ with probability $\lambda \ge 0$ and identity with probability $1-\lambda$. Let $1 \le i \le d$ be arbitrary. We must show that

$$q(X_1,\dots,X_m) = (1-\lambda)p(X_1,\dots,X_m) + \lambda p(X_1,\dots,\sigma X_i\sigma^\top,\dots,X_m)$$

is hyperbolic and rank-1 linear. For rank-1 linearity, observe that since $D_{i,vv^\top}$ is a linear operator, the sum of rank-1 linear polynomials is rank-1 linear.

Hyperbolicity is a bit more involved. The steps of the proof proceeds as follows:

1. Show that if $P(X_1,\dots,X_m)$ is rank-1 linear and hyperbolic, then the roots of $P(tI-A_1,\dots,tI-A_d)$ interlace the roots of $P(tI-A_1,\dots,tI-A_{i-1},tI-A_i-vv^\top,tI-A_{i+1},\dots,tI-A_d)$. In the special case of $P(X_1,\dots,X_d) = \det\left(\sum_{i=1}^d X_i\right)$, this is just the Cauchy Interlacing Theorem (Theorem 8). For the general case, see Lemma 3.8 from [36].

2. Prove that if σ is a transposition and A is symmetric, then $A - \sigma A\sigma^\top$ has rank 2 and trace 0 so that we can write $A - \sigma A\sigma^\top = aa^\top - bb^\top$ for some vectors $a, b \in \mathbb{R}^n$. This is a simple calculation for which we refer the reader to Lemma 3.10 from [36]. Then, hyperbolicity reduces to proving that

$$Q(tI-A_1,\dots,tI-A_d) = (1-\lambda)P(tI-A_1,\dots,tI-A_d) + \alpha P(tI-A_1,\dots,tI-A_i-aa^\top+bb^\top,\dots,tI-A_d)$$

is real-rooted. Using step 1, we know that

$$P(tI - A_1, \ldots, tI - A_i + bb^\top, \ldots, tI - A_d)$$

interlaces both $P(tI - A_1, \ldots, tI - A_d)$ and

$$P(tI - A_1, \ldots, tI - A_i - aa^\top + bb^\top, \ldots, tI - A_d)$$

so that the two polynomials, whose $(1-\lambda, \lambda)$ convex combination is $Q(tI - A_1, \ldots, tI - A_d)$, have a common interlacing. Since $\lambda \in [0,1]$ was arbitrary, applying Theorem 7 completes the proof.

□

□

Proof of Theorem 43. Set $k = 2n$ and

$$M = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \in \mathbb{R}^{(2n)\times(2n)}$$

and set $A_1 = \cdots = A_d = M$. Let $P_1, \ldots, P_d, S_1, \ldots, S_d$ be independent, uniformly random $n \times n$ permutation matrices. Lemma 23 shows that the random matrices $P_i \oplus S_i$ are realizable by swaps. Writing

$$\begin{bmatrix} P_i & 0 \\ 0 & S_i \end{bmatrix} M \begin{bmatrix} P_i & 0 \\ 0 & S_i \end{bmatrix}^\top = \begin{bmatrix} 0 & (P_i S_i^\top) \\ (P_i S_i^\top)^\top & 0 \end{bmatrix}$$

and applying Theorem 44, we have the result. □

Corollary 13. *There exist $n \times n$ permutation matrices $Q_1, \ldots, Q_d$ such that*

$$\lambda_2\left(\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & Q_i \\ Q_i^\top & 0 \end{bmatrix}\right](x)\right) \leq \lambda_2\left(\mathbb{E}_{P_1,\ldots,P_d,S_1,\ldots,S_d}\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & P_i S_i^\top \\ (P_i S_i^\top)^\top & 0 \end{bmatrix}(x)\right]\right)$$

where $\lambda_2(p)$ denotes the second largest root of the polynomial p (in this case, this is the largest nontrivial eigenvalue).

10.2 Finite Free Convolutions

As in previous applications of the interlacing families proof technique, one needs to analyze the maximum root of the average of the polynomials in the interlacing family. The polynomials we discuss in this subsection are somewhat different from the ones in the interlacing family we ultimately want to analyze. We will relate these soon. For now, we will be considering the average of characteristic polynomials of matrices of the form $A + QBQ^\top$ and $(A + QBR^\top)(A + QBR^\top)^\top$, where Q, R are orthonormal matrices.

The inspiration for studying these polynomials stems from free probability, which studies "noncommutative random variables" and as an application, studies limiting distributions of the spectra of random matrices. We refer interested readers to [28].

Definition 21 (Symmetric Additive Convolution)**.** Let $p(x) = \chi[A](x)$ and $q(x) = \chi[B](x)$ be two real-rooted polynomials, where $A, B \in \mathbb{R}^{n\times n}$ are symmetric. We define the **symmetric additive convolution of p and q** by

$$(p \boxplus_n q)(x) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))}\chi[A + QBQ^\top](x)$$

Definition 22 (Asymmetric Additive Convolution)**.** Let $p(x) = \chi[AA^\top](x)$ and $q(x) = \chi[BB^\top](x)$ be two real-rooted polynomials with nonnegative roots, for some $A, B \in \mathbb{R}^{n\times n}$. We define the **asymmetric additive convolution of** p **and** q by

$$(p \boxplus\!\!\boxplus_n q)(x) = \mathbb{E}_{Q,R\sim \mathrm{unif}(O(n))}\chi[(A+QBR^\top)(A+QBR^\top)^\top](x)$$

Alternatively, we can define it as the polynomial satisfying the relation

$$\mathbb{S}((p \boxplus\!\!\boxplus_n q)(x)) = \mathbb{E}_{Q,R\sim \mathrm{unif}(O(n))}\chi\left[\begin{bmatrix} 0 & A \\ A^\top & 0 \end{bmatrix} + \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix}\begin{bmatrix} 0 & B \\ B^\top & 0 \end{bmatrix}\begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix}^\top\right](x)$$

where $\mathbb{S}$ is the linear operator on $\mathbb{C}[x]$ defined by $(\mathbb{S}p)(x) = p(x^2)$. This follows from the fact that

$$\mathbb{S}(\chi[M^\top M](x)) = \chi\begin{bmatrix} 0 & M \\ M^\top & 0 \end{bmatrix}(x)$$

for every $M \in \mathbb{R}^{n\times n}$.

From now on, whenever we write $p \boxplus_n q$ for two polynomials p, q we will implicitly assume p, q are characteristic polynomials of Hermitian $n\times n$ matrices A, B respectively. Similarly, whenever we write $p \boxplus\!\!\boxplus_n q$ for two polynomials p, q, we will implicity assume p, q are characteristic polynomials of positive semidefinite Hermitian $n\times n$ matrices.

The following are some nice properties of these convolutions.

Lemma 25 (Properties of $\boxplus_n$ and $\boxplus\!\!\boxplus_n$; Lemma 2.5 from [36])**.**

1. $(p \boxplus_n q)(x)$ *is real-rooted.*
2. $(p \boxplus\!\!\boxplus_n q)(x)$ *is real-rooted, with all roots being nonnegative.*
3. $\boxplus_n$ *and* $\boxplus\!\!\boxplus_n$ *are bilinear operators (in the coefficients of the polynomials on which they operate) and associative.*

As alluded to mentioned earlier, in this section, we will be able to give a polynomial time algorithm for building the elusive Ramanujan graphs. That the natural interlacing families algorithm runs in polynomial time crucially depends on the following explicit analytic formulas for the convolutions of two polynomials; specifically, given polynomials p, q through their coefficients, we can compute $p \boxplus_n q$ and $p \boxplus\!\!\boxplus_n q$ in polynomial time using the following formulas.

Theorem 46 (Theorem 1.1 from [40])**.** *Suppose*

$$p(x) = \sum_{k=0}^{n}(-1)^k a_k x^{n-k} \qquad q(x) = \sum_{k=0}^{n}(-1)^k b_k x^{n-k}$$

are real-rooted polynomials. Then,

$$(p \boxplus_n q)(x) = \sum_{k=0}^{n}(-1)^k x^{n-k}\cdot\left(\sum_{i+j=k}\frac{(n-i)!(n-j)!}{n!(n-i-j)!}a_i b_j\right)$$

Theorem 47 (Theorem 1.3 from [40])**.** *Suppose*

$$p(x) = \sum_{k=0}^{n}(-1)^k a_k x^{n-k} \qquad q(x) = \sum_{k=0}^{n}(-1)^k b_k x^{n-k}$$

are polynomials with nonnegative real roots. Then,

$$(p \boxplus\!\!\boxplus_n q)(x) = \sum_{k=0}^{n}(-1)^k x^{n-k}\cdot\left(\sum_{i+j=k}\frac{(n-i)!(n-j)!}{n!(n-i-j)!}a_i b_j\right)^2$$

10.3 The Cauchy Transform

Definition 23. Let $p(x) = \prod_{i=1}^{d}(x - \lambda_i)$ be a degree-d real-rooted polynomial. We define the **Cauchy Transform of** p to be the function

$$\mathcal{G}_p(x) = \frac{1}{d} \cdot \frac{p'(x)}{p(x)} = \frac{1}{d} \sum_{i=1}^{d} \frac{1}{x - \lambda_i}$$

Similarly, we define the **Inverse Cauchy Transform of** p to be

$$\mathcal{K}_p(w) = \max\{x \in \mathbb{R} : \mathcal{G}_p(x) = w\}$$

Remark 9. Observe that $\mathcal{G}_p$ blows up at the roots of p and is decreasing for $x > \text{maxroot}(p)$. Thus, $\mathcal{K}_p(w)$ is the unique value of x such that $x > \text{maxroot}(p)$ and $\mathcal{G}_p(x) = w$. Note two key facts about $\mathcal{K}_p(w)$:

1. $\mathcal{K}_p(w)$ is an upper bound on $\text{maxroot}(p)$
2. $\mathcal{K}_p(w) \to \text{maxroot}(p)$ as $w \to \infty$

Note that the Cauchy Transform is essentially the barrier from Kadison-Singer but for a single variable and the barrier from Ramanujan Sparsifiers, and is also very similar to the barrier from Restricted Invertibility. [36] note that using the barrier argument from Ramanujan Sparsifiers [27] is suboptimal. The following results (proven in [40] and but also stated in [36]) are what will allow us to achieve the optimal Ramanujan bound.

Theorem 48 (Theorem 2.7 from [36])**.** *Let p, q be real-rooted polynomials, and let $w > 0$. Then $\mathcal{K}_{p \boxplus_n q}(w) \leq \mathcal{K}_p(w) + \mathcal{K}_q(w) - w^{-1}$.*

Theorem 49 (Theorem 2.8 from [36])**.** *Let p, q be real-rooted polynomials. Assume that all roots of $p(x)$ and $q(x)$ are nonnegative. Then $\mathcal{K}_{\mathbb{S}(p \boxplus \boxplus_n q)}(w) \leq \mathcal{K}_{\mathbb{S}p}(w) + \mathcal{K}_{\mathbb{S}q}(w) - w^{-1}$.*

10.4 Quadrature

In this subsection, we will focus on relating finite free convolutions to averaging over $\mathcal{S}_n$. The name "quadrature" (also known as "Gaussian quadrature") comes from numerical integration. It is a method of computing the integral of a function f approximately or, in some cases, exactly, by reducing the integral to a finite sum of f evaluated at a small number of special points known as "quadrature points". The approximation guarantees are usually made through Taylor series. This theory of integral approximation is also a beautiful example of a practical application of the theory of orthogonal polynomials. We will not discuss those applications here.

We refer to the following result as a "quadrature" result because we are relating an average over the orthogonal group, namely, an integral, to an average over $\mathcal{S}_n$, namely, a finite sum. However, our goal is actually "to go in the opposite direction", in the sense that we are trying to use an integral to compute a sum that has superexponentially many terms. The reason we want to do this is precisely because this integral is (essentially) a finite free convolution, for which we have bounds on maximum roots via Theorem 48 and Theorem 49, as well as polynomial-time computable formulas via Theorem 46 and Theorem 47.

Theorem 50 (Quadrature Theorem; Theorem 4.2 from [36])**.** *Let $A, B \in \mathbb{R}^{n \times n}$ be symmetric. Then*

$$\mathbb{E}_{P \sim \text{unif}(\mathcal{S}_n)} \det(A + PBP^\top) = \mathbb{E}_{Q \sim \text{unif}(O(n))} \det(A + QBQ^\top)$$

In order to prove this, we will show that $M \mapsto \det(A + MBM^\top)$ satisfies invariance under right multiplication by orthonormal matrices. That this invariance implies the quadrature result, Theorem 50, is a result of the following supporting lemma.

Lemma 26 (Lemma 4.3 from [36]). *Let $f : O(n) \to \mathbb{R}$ be any function, and let H be a finite subgroup of $O(n)$. If f is "invariant" in the sense that*

$$\mathbb{E}_{P\sim\mathrm{unif}(H)} f(P) = \mathbb{E}_{P\sim\mathrm{unif}(H)} f(PQ_0)$$

for all $Q_0 \in O(n)$, then

$$\mathbb{E}_{P\sim\mathrm{unif}(H)} f(P) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(Q)$$

Proof. It suffices to prove that $\mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(Q) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(PQ)$ for every $P \in H$, since then,

$$\begin{aligned}\mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(Q) &= \mathbb{E}_{P\sim\mathrm{unif}(H)}\mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(PQ) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))}\mathbb{E}_{P\sim\mathrm{unif}(H)} f(PQ)\\ &= \mathbb{E}_{Q\sim\mathrm{unif}(O(n))}\mathbb{E}_{P\sim\mathrm{unif}(H)} f(P) = \mathbb{E}_{P\sim\mathrm{unif}(H)} f(P)\end{aligned}$$

where in the last step, we use invariance. To prove $\mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(Q) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))} f(PQ)$, it suffices to show that the map $Q \mapsto PQ$ is bijective, for all $P \in H$. For this, we'll use the fact that H is a finite subgroup of $O(n)$. If $M \in O(n)$ is arbitrary, then $P^{-1}M \in O(n)$ so that $P^{-1}M \mapsto M$. This shows $Q \mapsto PQ$ is surjective. Similarly, if $PA = PB$, then $A = P^{-1}PA = P^{-1}PB = B$ so that $Q \mapsto PQ$ is injective. This proves bijectivity and hence, the lemma. □

10.4.1 Proving Invariance

Here, we will outline how to show that $M \mapsto \det(A + MBM^\top)$ satisfies the invariance property of Lemma 26 with $H = \mathcal{S}_n$. We restate the goal in the following lemma.

Lemma 27 (Invariance). *Define $f_{A,B} : O(n) \to \mathbb{R}$ by $f_{A,B}(M) = \det(A + MBM^\top)$, where A, B are fixed matrices in $\mathbb{R}^{n\times n}$. Let $H = \mathcal{S}_n$. Then $\mathbb{E}_{P\sim\mathrm{unif}(H)} f_{A,B}(P) = \mathbb{E}_{P\sim\mathrm{unif}(H)} f_{A,B}(PQ_0)$ for every $Q_0 \in O(n)$.*

The plan is as follows: first, we view $\mathcal{S}_n$ as the set of permutations on $e_1, \ldots, e_n$, the vertices of the probability simplex. Then, we will decompose $\mathcal{S}_n$ into subgroups $\mathcal{A}_{i,j,k}$, where $\mathcal{A}_{i,j,k} \subset \mathcal{S}_n$ is the set of all permutations that fix e_ℓ for all $\ell \notin \{i, j, k\}$. Similarly, we will decompose $O(n)$ into subgroups $O_{i,j,k}$, where $O_{i,j,k} \subset O(n)$ is the set of all orthogonal transformations that fix $V^\perp$, where V is the 2-dimensional subspace parallel to the 2-dimensional affine subspace containing e_i, e_j, e_k. With the following lemma, this will essentially allow us to just consider invariance with respect to $\mathcal{A}_{i,j,k}$ and $O_{i,j,k}$.

Lemma 28. *$\bigcup_{i,j,k} \mathcal{A}_{i,j,k}$ generate $\mathcal{S}_n$. Similarly, $\bigcup_{i,j,k} O_{i,j,k}$ generate $O(n)$.*

Proof. $\bigcup_{i,j,k} \mathcal{A}_{i,j,k}$ contains the set of all transpositions on n elements. The fact that the set of transpositions on n elements generate $\mathcal{S}_n$ is a standard result from group theory. For the second fact, see [36] Lemma 4.7. □

Lemma 29 (Invariance w.r.t. $\mathcal{A}_{i,j,k}$ and $O_{i,j,k}$). *For every i, j, k not all equal, and every $Q_0 \in O_{i,j,k}$,*

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{A,B}(P) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{A,B}(PQ_0)$$

Proof Sketch. Since $\mathcal{A}_{i,j,k}$ is isomorphic to $\mathcal{A}_2 \cong S_3$, and $O_{i,j,k}$ is isomorphic to $O(2)$, the claim is equivalent to proving

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_2)} \det(A + (P \oplus I_{n-2})B(P \oplus I_{n-2})^\top) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_2)} \det(A + (PQ_0 \oplus I_{n-2})B(PQ_0 \oplus I_{n-2})^\top)$$

for every $Q_0 \in O(2)$. For this, first decompose $O(2)$ into $SO(2)$, the subgroup of rotations

$$R_\theta = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}$$

and the subgroup of reflections

$$F = \left\{ I_2, \begin{bmatrix} 1 & 0 \\ 1 & -1 \end{bmatrix} \right\}$$

Similarly, decompose $\mathcal{A}_2$ into $\mathcal{Z}_3 = \{R_0, R_{2\pi/3}, R_{4\pi/3}\}$, the three cyclic permutations of e_1, e_2, e_3, and F. Then, first prove that

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{Z}_3)} \det(A + (P \oplus I_{n-2})B(P \oplus I_{n-2})^\top) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{Z}_3)} \det(A + (PR_\theta \oplus I_{n-2})B(PR_\theta \oplus I_{n-2})^\top)$$

which involves proving

$$\det(A + (R_\theta \oplus I_{n-2})B(R_\theta \oplus I_{n-2})^\top) = \sum_{k=-2}^{2} c_k e^{ik\theta}$$

for some constants $c_{-2}, c_{-1}, c_0, c_1, c_2$. We refer the details to Lemma 4.5 from [36]. Once we have this, combining with the fact that

$$\begin{aligned}\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_2)} \det(A + (P \oplus I_{n-2})B(P \oplus I_{n-2})^\top) &= \mathbb{E}_{D\sim\mathrm{unif}(F)}\mathbb{E}_{R\sim\mathrm{unif}(\mathcal{Z}_3)} \det(A + (RD \oplus I_{n-2})B(RD \oplus I_{n-2})^\top) \\ &= \mathbb{E}_{D\sim\mathrm{unif}(F)}\mathbb{E}_{R\sim\mathrm{unif}(\mathcal{Z}_3)} \det(A + (R \oplus I_{n-2})(D \oplus I_{n-2})B(D \oplus I_{n-2})^\top(R \oplus I_{n-2})^\top)\end{aligned}$$

we have the claim, since every $Q_0 \in O(2)$ may be written as $R_\theta D$ for some $R_\theta \in SO(2)$ and $D \in F$. □

With these, we can prove invariance with respect to $\mathcal{S}_n$ and $O(n)$.

Proof of Lemma 27. First, observe that

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(P) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)}\mathbb{E}_{P_1\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{A,B}(PP_1) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)}\mathbb{E}_{P_1\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{P^\top AP,B}(P_1)$$

By Lemma 29, for every i, j, k not all equal and every $Q_1 \in O_{i,j,k}$,

$$\begin{aligned}\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)}\mathbb{E}_{P_1\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{P^\top AP,B}(P_1) &= \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)}\mathbb{E}_{P_1\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{P^\top AP,B}(P_1Q_1) \\ &= \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)}\mathbb{E}_{P_1\sim\mathrm{unif}(\mathcal{A}_{i,j,k})} f_{A,B}(PP_1Q_1) \\ &= \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(PQ_1)\end{aligned}$$

With these, we conclude that for every $Q_1 \in O_{i,j,k}$,

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(P) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(PQ_1)$$

Now, let $Q_0 \in O(n)$ be arbitrary. By Lemma 28, we may write $Q_0 = \prod_{\ell=1}^m Q_\ell$, where for ever ℓ, there exists i, j, k not all equal such that $Q_\ell \in \mathcal{A}_{i,j,k}$. Then inductively applying the above equality m times gives $\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(P) = \mathbb{E}_{P\sim\mathrm{unif}(\mathcal{A}_n)} f_{A,B}(PQ_0)$ as desired. □

10.4.2 Bipartite Quadrature

Finally, here, we apply our main Quadrature Theorem (Theorem 50) to adjacency matrices of bipartite graphs. For convenience of notation, for $A, B \in \mathbb{R}^{n\times n}$ (not necessarily symmetric), we define

$$M_{A,B}(P,Q) = \begin{bmatrix} 0 & A \\ A^\top & 0 \end{bmatrix} + \begin{bmatrix} P & 0 \\ 0 & Q \end{bmatrix}\begin{bmatrix} 0 & B \\ B^\top & 0 \end{bmatrix}\begin{bmatrix} P & 0 \\ 0 & Q \end{bmatrix}^\top \in \mathbb{R}^{(2n)\times(2n)}$$

Since Theorem 50 does not deal with matrices of the form $P \oplus I$ and $Q \oplus I$ nor $I \oplus P$ and $I \oplus Q$, where $P \sim \mathrm{unif}(\mathcal{S}_n)$ and $Q \sim \mathrm{unif}(O(n))$, we cannot apply it directly. We need a "bipartite" analogue. However, it turns out the proof is nearly identical; for $A, B \in \mathbb{R}^{n\times n}$ not necessarily symmetric, one first needs to replace $f_{A,B}(M) = \det(A + MBM^\top)$ with $f_{A,B}(C) = \det(M_{A,B}(C, I))$ (resp. $f_{A,B}(C) = \det(M_{A,B}(I, C))$), all occurrences of P in $f_{A,B}$ by $P \oplus I$ (resp. $I \oplus P$), and all occurrences of Q in $f_{A,B}$ by $Q \oplus I$ (resp. $I \oplus Q$). This gives:

Theorem 51 (Corollary 4.12 from [36]). *For $A, B \in \mathbb{R}^{n\times n}$ not necessarily symmetric,*

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{S}_n)} \det\left(M_{A,B}(P, I)\right) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))} \det\left(M_{A,B}(Q, I)\right)$$

and

$$\mathbb{E}_{P\sim\mathrm{unif}(\mathcal{S}_n)} \det\left(M_{A,B}(I, P)\right) = \mathbb{E}_{Q\sim\mathrm{unif}(O(n))} \det\left(M_{A,B}(I, Q)\right)$$

Then, as an immediate consequence, through two applications of the preceding theorem, we have:

Corollary 14 ("Bipartite Analogue" of Theorem 50; Theorem 4.11 from [36]). *For $A, B \in \mathbb{R}^{n\times n}$ not necessarily symmetric,*

$$\mathbb{E}_{P,S\sim\text{unif}(\mathcal{S}_n)} \det(M_{A,B}(P,S)) = \mathbb{E}_{Q,R\sim\text{unif}(O(n))} \det(M_{A,B}(Q,R))$$

Now, we have the main result.

Theorem 52 (Theorem 4.10 from [36]). *Suppose $A, B \in \mathbb{R}^{n\times n}$ are (not necessarily symmetric) matrices with $A\mathbf{1} = A^\top\mathbf{1} = a\mathbf{1}$ and $B\mathbf{1} = B^\top\mathbf{1} = b\mathbf{1}$. Let p, q be degree-$(n-1)$ real-rooted polynomials such that $\chi[AA^\top](x) = (x-a^2)p(x)$ and $\chi[BB^\top](x) = (x-b^2)q(x)$. Then*

$$\mathbb{E}_{P,S\sim\text{unif}(\mathcal{S}_n)}\left[\chi_{M_{A,B}(P,Q)}(x)\right] = (x^2-(a+b)^2)\mathbb{S}((p \boxplus\!\!\boxplus_{n-1} q)(x))$$

Proof Sketch. First, apply a change of basis V that simultaneously block diagonalizes A, B, P, S as

$$VAV^\top = \hat{A}\oplus a \qquad VBV^\top = \hat{B}\oplus b \qquad VPV^\top = \hat{P}\oplus 1 \qquad VSV^\top = \hat{S}\oplus 1$$

This isolates the eigenvector $\mathbf{1}$ common to all of those these matrices and allows us to restrict our attention to $\hat{A}, \hat{B}, \hat{P}, \hat{S} \in \mathbb{R}^{(n-1)\times(n-1)}$.

Conjugating $M_{A,B}(P,S)$ by $V\oplus V$, we have

$$\mathbb{E}_{P,Q\sim\text{unif}(\mathcal{S}_n)}\left[\chi_{M_{A,B}(P,S)}(x)\right] = (x^2-(a+b)^2)\mathbb{E}_{P,S\sim\text{unif}(\mathcal{S}_n)}\left[\chi_{M_{\hat{A},\hat{B}}(\hat{P},\hat{S})}(x)\right]$$

(For the full calculation, see [36].)

Applying Corollary 14, the definition of the asymmetric additive convolution, and recalling that $p(x) = \chi[\hat{A}\hat{A}^\top](x)$ and $q(x) = \chi[\hat{A}\hat{A}^\top](x)$, we have

$$\mathbb{E}_{P,Q}\left[\chi_{M_{\hat{A},\hat{B}}(\hat{P},\hat{Q})}(x)\right] = \mathbb{E}_{P,Q\sim\text{unif}(O(n))}\left[\chi_{M_{\hat{A},\hat{B}}(\hat{P},\hat{Q})}(x)\right] = \mathbb{S}((p \boxplus\!\!\boxplus_{n-1} q)(x))$$

Combining these we have the result. □

Corollary 15 (Corollary 4.13 from [36]). *Suppose $A_1, \ldots, A_d \in \mathbb{R}^{n\times n}$ are (not necessarily symmetric) matrices, with $A_i\mathbf{1} = A_i^\top\mathbf{1} = a_i$ for all $1 \le i \le d$. Suppose $p_1, \ldots, p_d$ are degree-$(n-1)$ real-rooted polynomials satisfying $\chi[A_iA_i^\top](x) = (x-a_i^2)p_i(x)$ for every $1 \le i \le d$. Then*

$$\mathbb{E}_{P_1,\ldots,P_d,S_1,\ldots,S_d\sim\text{unif}(\mathcal{S}_n)}\chi\left[\sum_{i=1}^d (P_i\oplus S_i)\begin{bmatrix} 0 & A_i \\ A_i^\top & 0\end{bmatrix}(P_i\oplus S_i)^\top\right](x) = \left(x^2-\left(\sum_{i=1}^d a_i\right)^2\right)\mathbb{S}[(p_1 \boxplus\!\!\boxplus_{n-1} \ldots \boxplus\!\!\boxplus_{n-1} p_d)(x)]$$

Proof Sketch. Apply Theorem 52 inductively. □

10.5 Combining All of these Techniques

We now know that the set of polynomials

$$\chi\left[\sum_{i=1}^d \begin{bmatrix} 0 & P_i \\ P_i^\top & 0\end{bmatrix}\right](x)$$

where the P_i range over $\mathcal{S}_n$, form an interlacing family. It remains to bound the largest nontrivial root of

$$\mathbb{E}_{P_1,\ldots,P_d\sim\text{unif}(\mathcal{S}_n)}\left[\chi\left[\sum_{i=1}^d \begin{bmatrix} 0 & P_i \\ P_i^\top & 0\end{bmatrix}\right](x)\right]$$

This is where we will apply the finite free convolution results. The preceding four subsections were devoted to setting up the machinery required to perform this analysis. We are now ready to prove Theorem 42.

Proof of Theorem 42. Again, let $M = \begin{pmatrix} 0 & I \\ I & 0 \end{pmatrix}$ be a fixed perfect matching on n vertices. By Theorem 43, all that is left to do is to upper bound

$$\lambda_2\left(\mathbb{E}_{P_1,\ldots,P_d,S_1,\ldots,S_d}\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & P_iS_i^\top \\ (P_iS_i^\top)^\top & 0 \end{bmatrix}(x)\right]\right)$$

Since $I\mathbf{1} = 1\cdot\mathbf{1}$, Corollary 15 shows that

$$\mathbb{E}_{P_1,\ldots,P_d,S_1,\ldots,S_d}\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & P_iS_i^\top \\ (P_iS_i^\top)^\top & 0 \end{bmatrix}(x)\right] = (x^2-d^2)\mathbb{S}\left[\underbrace{(p\boxplus\boxplus_{n-1}\ldots\boxplus\boxplus_{n-1}p)(x)}_{d\text{ times}}\right]$$

where $p(x) = (x-1)^{n-1}$. Hence, it suffices to upper bound $\mathrm{maxroot}(\mathbb{S}[(p\boxplus\boxplus_{n-1}\ldots\boxplus\boxplus_{n-1}p)(x)])$.

Now, observe that

$$\mathcal{G}_p(x) = \frac{1}{2(n-1)}\frac{2(n-1)x(x^2-1)^{n-2}}{(x^2-1)^{n-1}} = \frac{x}{x^2-1}$$

is decreasing for all $x>1$, and, applying Theorem 49 with the inverse Cauchy transform $d-1$ times inductively, we also have

$$\mathcal{K}_{\mathbb{S}((p\boxplus\boxplus_{n-1}\ldots\boxplus\boxplus_{n-1}p)(x))}(w) \le d\cdot\mathcal{K}_{\mathbb{S}(p(x))} - \frac{d-1}{w} = d\cdot\mathcal{K}_{(x^2-1)^{n-1}}(w) - \frac{d-1}{w}$$

Thus, $\mathcal{K}_{(x^2-1)^{n-1}}(w) = x$ if and only if $w = \mathcal{G}_p(x) = \frac{x}{x^2-1}$. Plugging this in to the right-hand side of the above bound, we have an upper bound of

$$d\cdot\mathcal{K}_{(x^2-1)^{n-1}}(w) - \frac{d-1}{w} = dx - \frac{(d-1)(x^2-1)}{x} = \frac{x^2+(d-1)}{x} = x + \frac{d-1}{x}$$

for any $x>1$. This quantity is clearly minimized at $x=\sqrt{d-1}$, for which the bound becomes $2\sqrt{d-1}$. This is exactly what we wanted to show. □

10.6 An Algorithm

Now, let's discuss the natural interlacing families algorithm. Specifically, in this subsection, we'll explain the problem with the natural interlacing families algorithm, and sketch the main ideas in [42] behind obtaining a polynomial time algorithm for produce the desired Ramanujan multigraphs.

What is nice in this method of computing the Ramanujan multigraphs is that Theorem 47 provides us with a polynomial-time method for computing

$$\mathbb{E}_{P_1,\ldots,P_d,S_1,\ldots,S_d}\chi\left[\sum_{i=1}^{d}\begin{bmatrix} 0 & P_iS_i^\top \\ (P_iS_i^\top)^\top & 0 \end{bmatrix}(x)\right] = (x^2-d^2)\mathbb{S}\left[\underbrace{(p\boxplus\boxplus_{n-1}\ldots\boxplus\boxplus_{n-1}p)(x)}_{d\text{ times}}\right]$$

which is the root of the interlacing family tree. The tree also has depth d. The algorithm can then be written as:
There is one major problem with this algorithm: Each internal node of the tree has $n!$ many children!

Algorithm 4

Input: a degree $d \geq 3$; a size $n > d/2$

Output: the adjacency matrix of a d-regular bipartite Ramanujan multigraph $G = (V, E)$ with $|V| = 2n$

set $M = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ # adjacency matrix of a fixed perfect matching on $2n$ vertices
set $p(x) = p_0(x) = (x-1)^{n-1}$
for $1 \leq k \leq d$ **do**
    # compute and cache k-wise convolution of p with itself:
    $p_k(x) = (p \boxplus_{n-1} \ldots \boxplus_{n-1} p)(x) = (p_{k-1} \boxplus_{n-1} p)(x)$ # use convolution formula in Theorem 47
end for
compute $q_0(x) = \mathbb{S}(p_d(x))$ # this is the root of the interlacing family tree (after dividing out $x^2 - d^2$)
compute $r_0 = \text{maxroot}(q_0)$, the largest nontrivial root
keep track of the adjacency matrix of current multigraph $A \leftarrow 0 \in \mathbb{R}^{(2n)\times(2n)}$
for $1 \leq k \leq d$ **do**
    for permutation $P \in \mathcal{S}_n$ **do**
        # compute current characteristic polynomial and divide out trivial roots $\pm k$ (current multigraph is k-regular)
        $t(x) \leftarrow \chi\left[A + \begin{bmatrix} 0 & P \\ P^\top & 0 \end{bmatrix}\right](x)/(x^2 - k^2)$
        # compute the expected characteristic polynomial over remaining perfect matchings having fixed those in A and the new one specified by P
        $\tilde{q}_k(x) \leftarrow (t \boxplus_{n-1} p_{d-k})(x)$ # use convolution formula in Theorem 47
        if $\text{maxroot}(q) \leq r_{k-1} = \text{maxroot}(q_{k-1}(x))$ **then**
            $q_k(x) \leftarrow \tilde{q}_k(x)$
            $r_k \leftarrow \text{maxroot}(q_k)$
            $A \leftarrow A + \begin{bmatrix} 0 & P \\ P^\top & 0 \end{bmatrix}$
            break out of inner for-loop
        end if
    end for
end for
return A

10.6.1 A Modified Interlacing Family

The key observation of [42] was that by Theorem 45, there is no need fix an entire matching at once. In particular, [42] proposed the following variant of the interlacing family we considered earlier. For convenience, we index the $2n$ vertices by $[2n]$, with $[n]$ indexing the vertices on one side of the bipartition specified by M.

Definition 24 (Partially Specified Matching)**.** A **partially specified matching** is a bipartite matching such that vertices $1,\dots,t$ are matched and no vertex $t+1,\dots,n$ are matched. The **Matching Interlacing Family** is the interlacing family obtained by taking the nodes to be partially specified or perfect matchings, with the child of an internal node that has matched vertices $1,\dots,t$ either matching vertex $t+1$ if $t<n$, or starting a new partially specified matching (matching vertex 1) if $t=n$.

Note that the Matching Interlacing Family still has the same leaves as that of the original interlacing family; thus, the roots of the Matching Interlacing Family and the original interlacing family are also equal. Hence, the upper bound on the largest nontrivial root remains valid. The only thing one has to prove is that the Matching Interlacing Family is in fact, an interlacing family.

Theorem 53 (Lemma 3.1 from [42])**.** *The Matching Interlacing Family is an interlacing family.*

Proof Sketch. Fix permutations $P_1,\dots,P_{k-1}\in\mathcal{S}_n$ specifying perfect matchings, and a permutation $P_{k,t}$ specifying a partially specified matching up to and including node $1\le t\le n$. The main detail here is to instead consider the random adjacency matrix

$$\begin{aligned}A_{k,t}&=A_{k,t}(S_{k,t+1},S_{k+1},\dots,S_d;P_1,\dots,P_{k-1},P_{k,t})\\&\stackrel{\text{def}}{=}\sum_{i=1}^{k-1}\begin{bmatrix}0&P_i\\P_i^\top&0\end{bmatrix}+\begin{bmatrix}0&S_{k,t+1}P_{k,t}S_{k,t+1}^\top\\(S_{k,t+1}P_{k,t}S_{k,t+1}^\top)^\top&0\end{bmatrix}+\sum_{i=k+1}^{d}\begin{bmatrix}0&S_i\\S_i^\top&0\end{bmatrix}\end{aligned}$$

and its expected characteristic polynomial

$$\begin{aligned}E_{k,t}(x)&=E_{k,t}(x;P_1,\dots,P_{k-1},P_{k,t})\\&\stackrel{\text{def}}{=}\mathop{\mathbb{E}}_{S_{k+1},\dots,S_d\sim\text{unif}(\mathcal{S}_n),S_{k,t+1}\sim\text{unif}(\mathcal{S}_n^{(t+1)})}\left[\chi\left[A_{k,t}(S_{k,t+1},S_{k+1},\dots,S_d;P_1,\dots,P_{k-1},P_{k,t})\right](x)\right]\end{aligned}$$

for every $1\le k\le d$ and every $1\le t\le n$, where $S_{k+1},\dots,S_d$ are independent uniformly random permutations for yet-to-be-determined perfect matchings, $S_{k,t+1}$ is a random permutation such that $S_{k,t}P_{k,t}$ gives a perfect matching, and $\mathcal{S}_n^{(t+1)}$ is the subgroup of $\mathcal{S}_n$ of permutations fixing $1,\dots,t$ (with convention $\mathcal{S}_n^{(t+1)}=\mathcal{S}_n$ when $t=0$).

Even though we have not explicitly analyzed these random adjacency matrix, we've already set up all of the machinery required to prove this result. For details, we refer to [42]. □

The main benefit of the Matching Interlacing Family is that the number of children of an internal node drops to n; while the depth of the tree increases to dn, the total number of nodes in this interlacing family tree is still superexponentially smaller. The new algorithm is now as follows:
As stated in the algorithm, the one thing preventing us from obtaining an efficient algorithm is the lack of nice convolutional formulas for computing the expected characteristic polynomial $E_{k,t}(x)$. The majority of the work of [42] was to show how to compute the coefficients of $E_{k,t}(x)$. [42] first proves a new quadrature result:

Theorem 54 (Theorem 4.1 from [42])**.** *Let $A\in\mathbb{R}^{n\times n}$ be arbitrary (not necessarily symmetric) and let $P\sim\text{unif}(\mathcal{S}_n^{(t+1)})$, which we view as a random permutation matrix that randomly permutes $e_{t+1},\dots,e_n$ and fixes $e_1,\dots,e_t$. Then, the expected characteristic polynomial of*

$$\begin{bmatrix}0&A\\A^\top&0\end{bmatrix}+\begin{bmatrix}0&P\\P^\top&0\end{bmatrix}$$

Algorithm 5

Input: a degree $d \geq 3$; a size $n > d/2$

Output: the adjacency matrix of a d-regular bipartite Ramanujan multigraph $G = (V, E)$ with $|V| = 2n$

set $M = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ # adjacency matrix of a fixed perfect matching on $2n$ vertices
set $p(x) = p_0(x) = (x-1)^{n-1}$
for $1 \leq k \leq d$ **do**
    # compute and cache k-wise convolution of p with itself:
    $p_k(x) = (p \boxplus_{n-1} \ldots \boxplus_{n-1} p)(x) = (p_{k-1} \boxplus_{n-1} p)(x)$ # use convolution formula in Theorem 47
end for
compute $q_{0,0}(x) = \mathbb{S}(p_d(x))$; this is the root of the interlacing family tree (after dividing out $x^2 - d^2$)
compute $r_{0,0} = \text{maxroot}(q_{0,0})$, the largest nontrivial root
keep track of the adjacency matrix of current multigraph $A \leftarrow 0$
for $1 \leq k \leq d$ **do**
    keep track of current matched vertices with permutation P_t; initialize $P_0 \leftarrow I$
    for vertex $1 \leq t \leq n$ **do**
        for unmatched vertex $n+1 \leq v \leq 2n$ **do** # potential neighbor of t
            set $S_{t,v}$ to be the a fixed (deterministic) swap that matches t to v
            compute $E_{k,t}(x)$ # NOTE: need new routine (that we don't already know how to compute)
            if $\text{maxroot}(E_{k,t}) \leq r_{k,t} = \text{maxroot}(q_{k,t})$ **then**
                $q_{k,t}(x) \leftarrow E_{k,t}(x)$
                $r_{k,t} \leftarrow \text{maxroot}(q_{k,t})$
                $P_t \leftarrow S_{t,v} P_{t-1}$
                break out of inner for-loop
            end if
        end for
    end for
    update $A \leftarrow A + \begin{bmatrix} 0 & P_n \\ P_n^\top & 0 \end{bmatrix}$
end for
return A

is equal to the expected characteristic polynomial of

$$\begin{bmatrix} 0 & A \\ A^\top & 0 \end{bmatrix} + \begin{bmatrix} 0 & Q \\ Q^\top & 0 \end{bmatrix}$$

where Q is a uniformly random orthogonal matrix that fixes the subspace spanned by $e_1, \ldots, e_t$.

Then, [42] reduces computing $E_{k,t}(x)$ to computing the coefficients of the expected characteristic polynomial of $(\hat{A} + \hat{Q})^\top(\hat{A} + \hat{Q})$, where again, the expectation is taken over orthogonal matrices that fix the subspace spanned by $e_1, \ldots, e_t$. Finally, to compute these coefficients, [42] introduces a trivariate polynomial, written as a determinant of a $(2n) \times (2n)$ matrix. It turns out the kth coefficient of the characteristic polynomial we are interested in is a linear combination, with polynomially many terms, of the coefficients of this trivariate polynomial. We refer the reader to [42] for details of the analysis. These ingredients combined make Algorithm 5 a polynomial time algorithm for producing the adjacency matrix of a d-regular, bipartite Ramanujan multigraph.

11 Open Questions

1. Is there an efficient algorithm to find the vectors promised in Theorem 21?
2. Is there an efficient algorithm to find the good graph covers promised in Theorem 34 and Theorem 39?
3. How can we reduce the computational complexity of the algorithm from [42]?
4. How can we remove the bipartiteness condition in the Ramanujan graph constructions?
5. What are other interesting applications of the method of interlacing polynomials?

12 Acknowledgements

I am very grateful to both Professor Shayan Oveis Gharan, as well as Professor Rekha Thomas, for all of their patience and encouragement, for very helpful discussions and for giving useful feedback on this thesis. I also would like to thank Alireza Rezaei and Yuan Gao for helping me proofread this thesis and also providing helpful feedback. Last but not least, I want to thank my friends and family for the incredible love and support they've given me.

References

[1] R. V. Kadison, I. M. Singer. Extensions of pure states. *American Journal of Mathematics 81, pp.383-400* (1959)

[2] R. G. Gallager. Low Density Parity Check Codes. *MIT Press, Cambridge, MA* (1963)

[3] O. J. Heilmann, E. H. Lieb. Theory of monomer-dimer systems. *Communications in Mathematical Physics, 25(3), pp.190-232* (1972)

[4] N. Alon. Eigenvalues and Expanders. *Combinatorica 6 (2), pp.83-96* (1986)

[5] A. Lubotzky, R. Phillips, P. Sarnak. Ramanujan Graphs. *Combinatorica 8 (3), pp.261-277* (1986)

[6] J. Bourgain, L. Tzafriri. Invertibility of "Large" Submatrices with Applications to the Geometry of Banach Spaces and Harmonic Analysis. *Israel Journal of Mathematics, Volume 57, pp.137-224* (1987)

[7] G. A. Margulis. Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problems of Information Transmission, 24(1), pp.39-46* (1988)

[8] A. Nilli. On the Second Eigenvalue of a Graph. *Discrete Mathematics 91, pp.207-210* (1991)

[9] J. P. Dedieu. Obreschkoff's theorem revisited: what convex sets are contained in the set of hyperbolic polynomials? *Journal of Pure and Applied Algebra, Vol. 81, Issue 3, pp.269-278* (1992)

[10] H. Dette, W. J. Studden. Some New Asymptotic Properties for the Zeros of Jacobi, Laguerre and Hermite Polynomials. *arXiv:math/9406224v1* (1994)

[11] M. Morgenstern. Existence and Explicit Constructions of $q+1$ Regular Ramanujan Graphs for Every Prime Power q. *Journal of Combinatorial Theory, Series B 62, pp.44-62* (1994)

[12] A. Benczur, D. Karger, Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time. *Proceedings of the twenty-eighth annual ACM symposium on the Theory of Computing (STOC '96), pp.47-55* (1996)

[13] M. Sipser, D. A. Spielman. Expander Codes. *IEEE Transactions on Information Theory, Vol. 42, No. 6, pp. 1710-1722* (1996)

[14] R. Vershynin. John's Decompositions: Selecting a Large Part. *Israel Journal of Mathematics, Volume 122, pp.253-277* (2001)

[15] L. Gurvits, A. Samorodnitsky. A Deterministic Algorithm for Approximating the Mixed Discriminant and Mixed Volume, and a Combinatorial Corollary. *Discrete & Computational Geometry, Vol. 27, Issue 4, pp.531-550* (2002)

[16] J. Friedman. A Proof of Alon's Second Eigenvalue Conjecture and Related Problems. *arXiv:cs/0405020* (2004)

[17] N. Weaver. The Kadison-Singer problem in discrepancy theory. *Discrete Mathematics, 278(1-3), pp.227-239* (2004)

[18] D. Spielman, S. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. *STOC, pp.81-90* (2004). Full version available at http://arxiv.org/abs/cs.DS/0310051.

[19] Y. Bilu, N. Linial. Lifts, Discrepancy and Nearly Optimal Spectral Gap. *Combinatorica, Volume 26, pp.495-519* (2006)

[20] S. Hoory, N. Linial, A. Wigderson. Expander Graphs and their Applications. *American Mathematical Society, Vol. 43, No. 4, pp. 439-561* (2006)

[21] D. Spielman, N. Srivastasa. Graph Sparsification by Effective Resistances. *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC '08), pp. 563–568* (2008)

[22] D. Spielman, S. Teng. Spectral Sparsification of Graphs. *arXiv:0808.4134* (2008)

[23] J. A. Tropp. Column Subset Selection, Matrix Factorization, and Eigenvalue Optimization. *arXiv:0806.4404* (2008)

[24] J. Borcea, P. Brändén. Applications of Stable Polynomials to Mixed Determinants: Johnson's Conjectures, Unimodality and Symmetrized Fischer Products. *arXiv:math/0607755v5* (2008)

[25] S. J. Miller, T. Novikoff, A. Sabelli. The Distribution of the Largest Nontrivial Eigenvalues in Families of Random Regular Graphs. *Experimental Mathematics Vol. 17, No. 2, pp.231-244* (2008)

[26] P. G. Casazza, R. Vershynin. Kadison-Singer meets Bourgain-Tzafriri. Preprint.

[27] J. Batson, D. A. Spielman, N. Srivastava. Twice-Ramanujan Sparsifiers. *arXiv:0808.0163* (2009)

[28] R. Speicher. Free Probability Theory. *arXiv:0911.0087v1* (2009)

[29] J. Borcea, P. Brändén. Multivariate Pólya-Schur Classification Problems in the Weyl Algebra. *arXiv:math/0606360v6* (2009)

[30] D. A. Spielman, N. Srivastava. An Elementary Proof of the Restricted Invertibility Theorem. *arXiv:0911.1114* (2010)

[31] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *arXiv:1004.4389* (2010)

[32] L. Trevisan. Lecture Notes on Expansion, Sparsest Cut, and Spectral Graph Theory.

[33] A. Marcus, D. A. Spielman, N. Srivastava. Ramanujan Graphs and the Solution of the Kadison-Singer Problem. *arXiv:1408.4421* (2014)

[34] A. Marcus, D. A. Spielman, N. Srivastava. Interlacing Families I: Bipartite Ramanujan Graphs of All Degrees. *arXiv:1304.4132* (2014)

[35] A. Marcus, D. A. Spielman, N. Srivastava. Interlacing Families II: Mixed Characteristic Polynomials and the Kadison-Singer Problem. *arXiv:1306.3969* (2014)

[36] A. Marcus, D. A. Spielman, N. Srivastava. Interlacing Families IV: Bipartite Ramanujan Graphs of All Sizes. *arXiv:1505:08010* (2015)

[37] N. Anari, S. O. Gharan. The Kadison-Singer Problem for Strongly Rayleigh Measures and Applications to Asymmetric TSP. *arXiv:1412.1143* (2015)

[38] K. Chandrasekaran, A. Velingker. Towards Constructing Ramanujan Graphs Using Shift Lifts. *arXiv:1502.07410v3* (2015)

[39] N. Anari, S. O. Gharan. Effective-Resistance-Reducing Flows, Spectrally Thin Trees, and Asymmetric TSP. *arXiv:1411.4613* (2015)

[40] A. Marcus, D. A. Spielman, N. Srivastava. Finite free convolutions of polynomials. *arXiv:1504.00350* (2015)

[41] C. Hall, D. Puder, W. F. Sawin. Ramanujan Coverings of Graphs. *arXiv:1506.02335* (2016)

[42] M. B. Cohen. Ramanujan Graphs in Polynomial Time. *arXiv:1604.03544* (2016)

[43] N. Anari, S. O. Gharan, A. Saberi, N. Srivastava. Approximating the Largest Root and Applications to Interlacing Families. *arXiv:1704.03892* (2017)

[44] Lecture 11: Clustering and the Spectral Partitioning Algorithm. Lecturer: Shayan Oveis Gharan. Course: CSE 521 Design and Analysis of Algorithms I (2016)

[45] Lecture 12: Introduction to Spectral Graph Theory, Cheeger's Inequality. Lecturer: Shayan Oveis Gharan. Course: CSE 521 Design and Analysis of Algorithms I (2016)

[46] Lecture 13: Random walks on graphs and electrical networks. Lecturer: James R. Lee. Course: CSE 525 Randomized Algorithms (2016)

[47] Lecture 18: Markov chains and mixing times. Lecturer: James R. Lee. Course: CSE 525 Randomized Algorithms (2016)