Polynomial Paradigm in Algorithms **Winter 2020**

Lecture 12-15: Simplicial Complexes

Lecturer: Shayan Oveis Gharan *Feb 24th*

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications.*

Parts of this lecture and the next are based on talks by Irit Dinur and Yotam Dikstein at Simon's Program on High dimensional expanders.

The basic object that we work with is a *simplicial complex* that we call X. Unless specified otherwise, we assume X is defined on a ground set of elements $[n]$. More concretely, $X \subseteq 2^{[n]}$ is a simplicial complex if it is a downward closed family of sets. We call sets $\sigma \in X$, *faces* of X. For a face σ, the dimension of σ is one less than the number of elements in σ,

$$\dim(\sigma) := |\sigma| - 1.$$

The dimension of X is defined to be the largest dimension of any face of X:

$$\dim(X) := \max_{\sigma \in X} \dim(\sigma).$$

For X of dimension d, we write $X(-1), X(0), \ldots, X(d)$ to denote faces of dimension $-1, \ldots, d$ in X. Note that $X(-1) = \{\emptyset\}$ as the $\emptyset$ is the only face of dimension -1 in X.

We say a simplicial X is *pure* if for any face $\sigma \in X$, there is a face $\tau \in X$ such that $\sigma \subseteq \tau$ and $\dim(\tau) = \dim(X)$. All complexes that we will work with are pure. We call a face $\tau \in X$ such that $\dim(\tau) = \dim(X)$ a top face or a maximal face of X.

One of the main goals of this lecture is to define what it means for X to be a high dimensional expander. Normally, one may study expander graphs as a family of regular graphs. But, for the reason that will become clear shortly, we need to also study expansion when the graph is not regular.

12.1 Assigning Weights to Top Faces

We start by assigning weights to top faces of X. Consider the natural 2-dimensional complex defined by the following graph where we have a face for every vertex, edge and triangle. Say, we assign a uniform

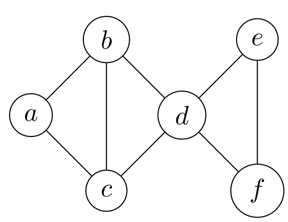

distribution to triangles. This induces a distribution on edges where the probability of an edge $\{a, b\}$ is the probability that by choosing a random triangle and knocking down a vertex we get $\{a, b\}$. This implies for example

$$\mathbb{P}[\{a, b\}] = \mathbb{P}[\{a, b, c\}]\,\mathbb{P}[\{a, b\}|\{a, b, c\}] = \frac{1}{3} \cdot \frac{1}{3} = \frac{1}{9}.$$

Similarly, we have $\mathbb{P}[\{b,c\}] = 2/9$. We also get a distribution over vertices: For example, $\mathbb{P}[a] = \mathbb{P}[a|\{a,b\}] + \mathbb{P}[a|\{a,c\}] = 1/9$ whereas $\mathbb{P}[d] = 2/9$.

More generally, if we start with a distribution Π_d over faces of dimension d, it induces a distribution Π_{d-1} over faces of dimension $d-1$ and so on where each time we knock out a vertex from the given face uniformly at random. Note that we can also construct Π_d from Π_{d-1} as follows: First we sample $\sigma \sim \Pi_{d-1}$, then among all faces $\tau \supset \sigma$ of dimension d we choose one proportional to its probability in Π_d.

For $0 \leq k \leq d$, we abuse notation and write $\sigma \sim X(k)$ to denote that σ is sampled from the distribution Π_k.

Definition 12.1 (1-skeleton of a Complex)**.** *Given a d dimensional complex X defined on ground elements $[n]$, the 1-skeleton (or the graph of X) is the undirected weighted graph with vertex set $X(0) = [n]$, edge set $X(1)$ and the weight of every edge $e = \{i,j\}$ is $\Pi_1(\{i,j\})$.*

In reality, we do not know how to construct high dimensional expanders such that all these distributions are uniform and that is a reason that we have to study spectral expansion for non-regular graphs.

Let me immediately point a simple connection between these complexes and generating polynomials. Given a complex X on elements $[n]$ and dimension d with distribution Π_d on top faces; we can define the $d+1$-homogeneous generating polynomial of X as follows:

$$g_X(z_1, \dots, z_n) = \sum_{\sigma \in X : \dim(\sigma) = d} \Pi_d(\sigma) z^\sigma.$$

So, Π_{d-1}, the underling distribution on faces of dimension $d-1$, corresponds to the coefficient of

$$\left(\sum_{i=1}^{n} \frac{1}{d+1} \partial_{z_i} \right) g_X.$$

12.2 Links

Given a vertex $i \in X$, we define the *link* of i as

$$X_i = \{\sigma : \sigma \cup \{i\} \in X, i \notin \sigma\}.$$

Note that the link of i is also a simplicial complex.

In other words, we look at every face of X that has i and we knock out i to obtain a face of X_i. You should think of a link as a generalization of the *neighborhood* of a vertex in a graph.

In the above example, $X_a = \{\{b,c\}, \{b\}, \{c\}, \emptyset\}$. At higher dimensions we can also look at the links of high dimensional faces. More generally, we can look at the link of any face τ as

$$X_\tau = \{\sigma : \sigma \cup \tau \in X, \sigma \cap \tau = \emptyset\}.$$

One of the main property of these high dimensional expanders is that if we know "expansion" properties of these local neighborhoods, we can deduce global "expansion" properties of the whole complex.

12.3 Spectral Link Expansion

What does it mean for a link to be spectrally expanding? For start, we want every link to be connected. First of all, since we start from a weighted complex, every link is also a weighted complex.

Given a vertex i, the distribution Π_d on d-faces of X imposes a distribution on top faces of X_i that corresponds to the coefficients of the following polynomial:

$$\partial_{z_1} g_X$$

So, we write $\sigma \sim X_i(k)$ to denote that $\sigma \cup i$ is sampled from Π_{k+1} conditioned on i.

We will say a complex is spectral expander if every link is an spectral expander. In other words, if all of these "local" structures are expanding. One amazing property of the links is that we can write the whole complex as a convex combination of the links.

This fact essentially corresponds to the following Euler identity for homogeneous polynomials:

$$g_X = \sum_i z_i \partial_{z_i} g_X$$

12.3.1 Spectral Expanders

Let G be the underlying graph of X, i.e., G is the 1-skeleton of X. In this section we see what it means for G to be a spectral expander.

For two functions $f, g : X(0) \to \mathbb{R}$, define

$$\langle f, g \rangle := \mathbb{E}_{i \sim X(0)} \left[f(i) g(i) \right],$$

where recall that $i \sim X(0)$ means that i is sampled from the distribution Π_0 on the vertices of X. We can look at the random walk operator on G that says if I am at vertex i, I choose an edge incident to i with probability proportional to its weight:

$$\mathbb{P}\left[i \to j\right] = \mathbb{P}\left[\{i,j\} | i\}\right] = \frac{\mathbb{P}\left[\{i,j\}\right]}{\sum_k \mathbb{P}\left[\{i,k\}\right]}.$$

Having this, we can define the transition probability operator $P : \mathbb{R}^{X(0)} \to \mathbb{R}^{X(0)}$ as follows: For a function $f : X(0) \to \mathbb{R}$, define

$$Pf(i) := \mathbb{E}_{\{i,j\}|i} \left[f(j) \right] = \sum_j \mathbb{P}\left[i \to j\right] f(j).$$

It follows that P is self-adjoint with respect to the above inner product, i.e., for any two $f, g : X(0) \to \mathbb{R}$,

$$\langle f, Pg \rangle = \mathbb{E}_i \left[f(i) Pg(i) \right] = \mathbb{E}_i \left[f(i) \mathbb{E}_{\{i,j\}|i} \left[g \right](j) \right] = \mathbb{E}_{\{i,j\}} \left[f(i) g(j) \right] = \langle Pf, g \rangle. \tag{12.1}$$

i.e., we are using that the following two operations are equivalent:

- First choose an edge then a vertex in it
- First choose a vertex and then an edge that contains this vertex.

As a corollary the matrix P has n eigenfunctions. The top eigenfunction is the all-ones function and all eigenvalues of P are between $-1, +1$.

Definition 12.2. *A graph G is a two sided λ-expander if the eigenvalues*

$$-1 \le \lambda_n \le \dots \le \lambda = 1$$

if $-\lambda \le \lambda_n, \lambda_2 \le \lambda$. We say G is a one-sided λ-expander if $\lambda_2 \le \lambda$.

12.3.2 Back to Simplicial Complexes

The following definition is the key in this area:

Definition 12.3 (High Dimensional Spectral Expander)**.** *A d-dimensional complex X is a two-sided λ-spectral expander if the graph underlying link of every face $\sigma \in X(i)$ for $i < d-1$ is itself a two-side λ-spectral expander.*

Where does this definition come from? First, it does not seem such a thing is realizable. First of all, it is not hard to see that the complete complex is a 0-spectral expander. This is because the underlying graph of the link of any face is also a complete graph, the second eigenvalue of the random walk of a complete graph is $\frac{-1}{n-1}$. More generally, if we choose a random complex with a high enough density we get a high dimensional expander. One non-trivial question is if there is sparse high dimensional expanders? Such objects were first constructed by Lubotzky Samuels and Vishne. More recently, Kaufman and Oppenheim provided a "simpler" group theoretic construction:

Theorem 12.4. *Given $d \geq 2$ and a prime power q where $\sqrt{q} \gg d-1$ there exists in finite sequence of simplicial complexes $X^1, X^2, \ldots$ such that the number of vertices goes to infinity, each X^i is a $\frac{1}{\sqrt{q}-(d-1)}$-one spectral expander and that every vertex of each complex is in at most $f(q,d)$ many d-dimensional faces (independent of the number of underlying vertices in the complex).*

Fact 12.5. *Given a complex X, let A be the (weighted) adjacency matrix of the 1-skeleton of X. There is a normalizing constant $C > 0$ such that $A = \binom{d+1}{2} \nabla^2 g_X|_{z=\mathbf{1}}$.*

Proof. First note that for any edge $\{i,j\}$,

$$\begin{aligned} \partial_{z_i}\partial_{z_j} g_X|_{z=\mathbf{1}} &= \sum_{\sigma:\{i,j\}\in\sigma} \Pi_d(\sigma) \\ &= \sum_{\sigma:\{i,j\}\in\sigma} \frac{\mathbb{P}\left[\{i,j\}|\sigma\right]}{\binom{d+1}{2}} \Pi_d(\sigma) = \frac{\mathbb{P}\left[\{i,j\}\right]}{\binom{d+1}{2}}. \end{aligned}$$

as desired. □

Fact 12.6. *Given a d-dimensional complex X on elements $[n]$, let P be the random walk operator on 1-skeleton of X (as defined above). Then,*

$$P = \frac{1}{d}\,\mathrm{diag}(\nabla g_X)^{-1} \nabla^2 g_X.$$

where all differential operators are evaluated at $z = \mathbf{1}$, and diag *maps a vector $v \in \mathbb{R}^n$ to a diagonal $n \times n$ matrix with v on the diagonal.*

Proof. Fix $1 \leq i,j \leq n$. Then,

$$P(i,j) = \frac{\mathbb{P}\left[\{i,j\}\right]}{\sum_k \mathbb{P}\left[\{i,k\}\right]} = \frac{\nabla^2 g_X(i,j)}{\sum_k \nabla^2 g_X(i,k)} = \frac{\nabla^2 g_X(i,j)}{d\nabla g_X(i)},$$

as desired. The first equality uses the previous fact, and the last equality uses Euler's identity:

$$\sum_j z_j \partial_{z_j}(\partial_{z_i} g_X) = d\partial_{z_i} g_X$$

as g_x is $d+1$-homogeneous. □

Lemma 12.7. *Given a d-dimensional complex X on $[n]$ and random walk operator P,*

$$\lambda_2(P) = \lambda_1\left(\frac{1}{d}\,\mathrm{diag}(\nabla g_X)^{-1}\nabla^2 \log g_X\right)$$

Proof. Without loss of generality assume that $g_X|_{z=\mathbf{1}} = 1$. It follows that

$$\nabla^2 \log g_X = \nabla^2 g_X - (\nabla g_X)(\nabla g_X)^T.$$

Therefore,

$$\frac{1}{d}\,\mathrm{diag}(\nabla g_X)^{-1}\nabla^2 \log g_X = P - \frac{1}{d}\,\mathrm{diag}(\nabla g_X)^{-1}(\nabla g_X)(\nabla g_X)^T = P - \frac{1}{d}\mathbf{1}(\nabla g_X)^T.$$

Recall that $\mathbf{1}$ is the largest eigenfunction of P. But,

$$\frac{1}{d}\mathbf{1}(\nabla g_X)^T\mathbf{1} = \frac{1}{d}\mathbf{1}(\sum_i \partial_{z_i} g_X) = \frac{d+1}{d}\mathbf{1}.$$

Therefore, subtracting $\frac{1}{d}\mathbf{1}(\nabla g_X)^T$ from P reduces the eigenvalue 1 of the all-ones eigenfunction to $-1/d$ and the rest of the eigenvalues remain invariant. The lemma's statement follows. □

As an immediate consequence of the above lemma we can write g_X is log-concave at $\mathbf{1}$ iff $\lambda_2(P) \le 0$. Let me state another consequence:

Lemma 12.8. *A d-dimensional complex X is a one-sided 0-spectral expander iff g_X is log-concave (over $\mathbb{R}^n_{\ge 0}$).*

Proof. Suppose g_X is log-concave. Fix a face $\tau \in X(k)$ for $k \le d-2$. Then (up to a normalizing constant), the generating polynomial of $g_{X_\tau} = \partial^\tau g_X$. Since g_X is log-concave, by HW2 Problem 5, it is completely log-concave; so g_{X_τ} is log-concave. Therefore, $\nabla^2 \log g_{X_\tau}|_{z=\mathbf{1}} \preceq 0$. But by the previous lemma, this implies that $\lambda_2(P_\tau) \le 0$ as desired.

Conversely, suppose X is a one-sided 0-spectral expander. By Theorem 8.15, to show g_X is log-concave we need to verify

i) For any τ of dimension at most $d-2$, that $\partial^\tau g_X$ is in-decomposable.

ii) For any τ of dimension $d-2$, $\partial^\tau g_X$ is log-concave.

We start by checking the first condition: For τ of co-dimension at least 2. If $\tau \notin X$, then $\partial^\tau g_X = 0$ so there is nothing to prove. Otherwise, we know underlying graph of X_τ is a 0-spectral expander. So $\lambda_2(P_\tau) \le 0$. But this implies that the underlying graph of X_τ is connected. So, $g_{X_\tau} = \partial^\tau g_X$ is in-decomposable. Now, for τ of co-dimension 2, X_τ being 0-spectral expander implies $\lambda_2(P_\tau) \le 0$. But by the previous lemma this implies $\nabla^2 \log g_{X_\tau} \preceq 0$ as desired. □

Since the support of any homogeneous multilinear log-concave polynomial correspond to bases of a matroid (see corollary 10.6), we obtain that the top faces of any one-sided 0-spectral expander correspond to bases of a matroid.

12.4 Oppenheim's Trickling Down Theorem

Next, we prove the following theorem due to Oppenheim:

Theorem 12.9. *Let X be a d-dimensional complex. Suppose for all vertices i,*

i) X_i is a one-sided λ-spectral expander.

ii) The 1-skeleton of X is connected.

Then, the 1-skeleton of X is a $\frac{\lambda}{1-\lambda}$-spectral expander.

Roughly speaking, in the case $d = 2$, the above theorem says the following: Say we are given a graph G such that the local neighborhood of every vertex is a very good expander. If G is also connected, then G is a very good expander.

Note that the second assumption is also necessary as we can just let G be the union of two very good expander. In such a case even though the local neighborhood of every vertex is a very good expander the whole graph is far from an expander.

By a repeated application of the above theorem we obtain the following theorem:

Theorem 12.10. *Let X be a d-dimensional complex such that*

1. *For any $\tau \in X(d-2)$, X_τ is one-sided λ-spectral expander.*
2. *For any $\tau \in X(k)$ where $k \leq d-2$, the underlying graph of X_τ is connected.*

Then, X is a one-sided $\frac{\lambda}{1-(d-1)\lambda}$-spectral expander.

Note that for the case $\lambda = 0$ the above theorem is special case of theorem 8.15.

In the rest of this section we prove Theorem 12.9 using a well-known technique called Garland's method.

First, we state the localization lemmas that correspond to the Garlands method:

Lemma 12.11. *Suppose $f, g : X(0) \to \mathbb{R}$. Then,*

$$\langle f, g\rangle = \mathbb{E}_{i\sim X(0)}\left[\langle f_i, g_i\rangle_{X_i}\right],$$

where f_i is the restriction of f to the vertices link of i, i.e., for any $j \in X_i(0)$, $f_i(j) = f(j)$.

Proof. Observe that

$$\begin{aligned}
\mathbb{E}_{j\sim X(0)}\left[f(j)g(j)\right] &= \mathbb{E}_{\{i,j\}\sim X(1)}\mathbb{E}_{j|\{i,j\}}f(j)g(j)\\
&= \mathbb{E}_{i\sim X(0)}\mathbb{E}_{\{i,j\}|i}f(j)g(j)\\
&= \mathbb{E}_{i\sim X(0)}\left[\langle f_i, g_i\rangle_{X_i}\right].
\end{aligned}$$

where in the second inequality we used that the following two operations are the same: (1) First sample an edge $\{i, j\}$, then choose j uniformly at random, (2) First sample i, then an edge $\{i, j\}$ incident to i and then pick the "other" vertex j. □

Lemma 12.12. *For $f, g : X(0) \to \mathbb{R}$,*

$$\langle Pf, g\rangle = \mathbb{E}_{k\sim X(0)}\left[\langle P_k f_k, g_k\rangle_{X_k}\right],$$

where P_k is the transition probability matrix corresponding to the 1-skeleton of the link of k.

Proof. First, by (12.1) we can write

$$\begin{aligned}\langle Pf, g\rangle &= \mathbb{E}_{\{i,j\}\sim X(1)}\left[f(i)g(j)\right]\\ &= \mathbb{E}_{\{i,j,k\}\sim X(2)}\mathbb{E}_{\{i,j\}|\{i,j,k\}}f(i)g(j)\\ &= \mathbb{E}_{k\sim X(0)}\mathbb{E}_{\{i,j,k\}|k}f(i)g(j)\\ &= \mathbb{E}_{k\sim X(0)}\mathbb{E}_{\{i,j\}\sim X_k(1)}f(i)g(j)\\ &= \mathbb{E}_{k\sim X(0)}\langle P_k f_k, g_k\rangle_{X_k}.\end{aligned}$$

where we used that the following two operations are the same: (1) Choose a random triangle $\{i, j, k\} \sim X(2)$ and drop one of them uniformly at random, (2) Choose a random vertex $k \sim X(0)$, choose a random triangle $\{i, j, k\} \sim X(2)$ among all those that have k, drop k to get $\{i, j\}$. □

Now, we are ready to prove Theorem 12.9. Let G be the 1-skeleton of X. Since G is connected, the second eigenvalue of P is smaller than 1. Let γ be an eigenvalue of P other than than the top eigenvalue, i.e., $\gamma < 1$. Let $f : X(0) \to \mathbb{R}$ be the corresponding eigenfunction, i.e., $Pf = \gamma f$. Since eigenfunctions of P are orthogonal, f is orthogonal to the top eigenfunction, the all ones function, i.e., $\langle f, \mathbf{1}\rangle = 0$. . Furthermore, assume that f is normalized such that $\langle f, f\rangle = 1$.

It follows that

$$\gamma = \langle Pf, f\rangle = \mathbb{E}_{i\sim X(0)}\left[\langle P_i f_i, f_i\rangle\right], \tag{12.2}$$

where we used the localization lemma 12.12. Now, for a second, suppose for all i, f_i is orthogonal to $\mathbf{1}_i$, then by the assumption of the theorem,

$$\langle P_i f_i, f_i\rangle \le \lambda\|f_i\|^2.$$

Therefore,

$$\mathbb{E}_{i\sim X(0)}\left[\langle P_i f_i, f_i\rangle\right] \le \mathbb{E}_{i\sim X(0)}\left[\lambda\|f_i\|^2\right] = \lambda\|f\|^2.$$

where the equality follows by another application of the localization lemma 12.11. So, we get $\gamma \le \lambda$. But this this assumption is too good to be true, the fact that f is orthogonal to all-ones does not imply that even a part of the support of f is orthogonal to all-ones in that part.

Instead, we decompose f_i as follows:

$$f_i = \alpha_i\mathbf{1}_i + f_i^\perp,$$

where $\langle f_i^\perp, \mathbf{1}_i\rangle = 0$. Since $\|\mathbf{1}_i\|_{X_i} = 1$, we can simply let $\alpha_i = \langle f_i, \mathbf{1}_i\rangle$. We can then write,

$$\begin{aligned}\mathbb{E}_{i\sim X(0)}\langle P_i f_i, f_i\rangle &= \mathbb{E}_{i\sim X(0)}\langle P_i f_i^\perp, f_i^\perp\rangle + \langle\alpha_i\mathbf{1}_i, \alpha_i\mathbf{1}_i\rangle\\ &\le \mathbb{E}_{i\sim X(0)}\lambda\|f_i^\perp\|^2 + \mathbb{E}_{i\sim X(0)}\langle f_i, \mathbf{1}_i\rangle^2\\ &= \mathbb{E}_{i\sim X(0)}\lambda\|f_i^2\| + (1-\lambda)\langle f_i, \mathbf{1}_i\rangle^2\\ &= \lambda + (1-\lambda)\mathbb{E}_{i\sim X(0)}\langle f_i, \mathbf{1}_i\rangle^2\end{aligned} \tag{12.3}$$

where in the first equality, we used that $\langle\mathbf{1}_i, f_i^\perp\rangle = 0$, and in the second inequality we used that $\|f_i\|^2 = \alpha_i^2 + \|f_i^\perp\|^2$.

But,

$$\langle f_i, \mathbf{1}_i\rangle = \mathbb{E}_{j\sim X_i(0)}f(j) = Pf(i)$$

Therefore,

$$\mathbb{E}_{i\sim X(0)}\langle f_i, \mathbf{1}_i\rangle^2 = \mathbb{E}_{i\sim X(0)} Pf(i)\cdot Pf(i) = \langle Pf, Pf\rangle = \gamma^2.$$

Putting the above equation together with (12.2), (12.3) we obtain

$$\gamma \le \lambda + (1-\lambda)\gamma^2$$

Solving the above for γ we obtain either $\gamma \ge 1$ or $\gamma \le \frac{\lambda}{1-\lambda}$ as desired. This completes the proof of Theorem 12.9.

12.5 High Dimensional Walks

Given a d-dimensional simplicial complex X defined on $[n]$ equipped with a distribution distribution Π_d on top faces, we can define a random walk on faces in $X(k)$. Fix $0 \le k \le d-1$. Given a face $\sigma \in X(k)$, the up-down-walk is defined as follows: First we choose $\tau \in X(k+1)$ condition on σ. Then, we remove a uniformly random element of τ to get $\sigma' \in X(k)$. We write $P^{k\wedge}$ to denote the transition probability matrix of this walk.

So, this operator acts as follows: Given a function $f : X(k) \to \mathbb{R}$, for any $\sigma \in X(k)$,

$$P^{k\wedge} f(\sigma) = \mathbb{E}_{i\sim X_\sigma(0)}\mathbb{E}_{\sigma'|\sigma\cup\{i\}} f(\sigma').$$

Observe that this walk is *lazy*, i.e., starting from $\sigma \in X(k)$, with probability $\frac{1}{k+2}$, we return back to σ, and we get $\sigma' = \sigma$.

Similarly, we can define the down-up walk, $P^{k\vee}$ where for a given $\sigma \in X(k)$, we first knock down a uniformly random element of σ, i.e., we choose $\tau \in X(k-1)|\sigma$, and then we go up and choose $\sigma' \in X(k)|\tau$.

Example 12.1 ($k = 0$ Case)**.** *In the above definition consider $k = 0$ case. Then, $P^{0\wedge}$ is exactly the half-lazy random walk on the 1-skeleton of X. Namely, given a vertex i, first we choose an edge $\{i, j\}|i$, then we drop one of i, j uniformly at random. So, we return back to i with probability* $1/2$. *So, we can write*

$$P^{0\wedge} = \frac{1}{2}P + I/2. \tag{12.4}$$

On the other hand, $P^{0\vee}$ is just the constant operator. Starting from a vertex i, first we remove i and get to $\emptyset$, then we choose $j \sim X(0)$. In particular, for $f : X(0) \to \mathbb{R}$,

$$P^{0\vee} f(i) = \mathbb{E}_{j\sim X(0)} f(j).$$

Observe that

$$\lambda_2(P) = \lambda_1(P^{0\wedge} - I/2 - P^{0\vee}). \tag{12.5}$$

The proof is similar to the proof of Lemma 12.7.

Example 12.2 (Matroid Case)**.** *For a well-defined example, we explain these walks on the* independence complex *of a matroid. Let $M = ([n], \mathcal{I})$ be a matroid of rank $r + 1$. Define a simplicial complex X_M on ground set $[n]$ where $\sigma \subseteq [n]$ is a face of X_M iff $\sigma \in \mathcal{I}$. Observe that X_M is a simplicial complex simply because the downward closed property of independence sets. Furthermore, X_M is pure because of the exchange property: Namely, given any independent set $\sigma \in cI$ we can add elements to σ to turn it into a base. So, the top faces of X_M are bases and X_M has rank r. We let Π_r be the uniformly distribution over bases of M.*

Since $g_M = g_{X_M}$ is a log-concave polynomial, by Lemma 12.8, X_M is a one-sided 0-spectral expander. Now, consider the down-up walk $P^{r\vee}$: Given a face $\sigma \in X_M(r)$, i.e., a base, first we drop an element, say i, of σ

uniformly at random, i.e., we go to an independent set $\sigma - \{i\}$ of rank r. Then, among all bases $\sigma' \supset \sigma - \{i\}$ we choose one uniformly at random. This walk is called the bases exchange walk. Observe that we can run this walk for any given matroid having access to an independence oracle. In this section, we will study second eigenvalue of $P^{r\vee}$ and we will use that in the following section to generate a uniformly random base of a matroid.

Lemma 12.13. *For $0 \leq k \leq d-1$ define M^k such that $P^{k\wedge} = \frac{I}{k+2} + \frac{k+1}{k+2} M^k$. Then, Then,*

$$\lambda_1(M^k - P^{k\vee}) \leq \max_{\tau \in X(k-1)} \lambda_2(P_\tau),$$

where P_τ is the random walk operator on the 1-skeleton of link of τ. In particular, if X is a λ-spectral expander, then for all $0 \leq k \leq d-1$, we have

$$\lambda_1(M^k - P^{k\vee}) \leq \lambda.$$

Proof. Let $f : X(k) \to \mathbb{R}$. Then,

$$\begin{aligned} \langle M^k f, f \rangle &= \mathbb{E}_{\tau \sim X(k-1)} \langle M_\tau f_\tau, f_\tau \rangle \\ &= \mathbb{E}_{\tau \sim X(k-1)} \mathbb{E}_{i \sim X_\tau(0)} \left(\mathbb{E}_{j \sim X_0(\tau \cup \{i\})} f(j) \right) \cdot f(i) \\ &= \mathbb{E}_{\tau \sim X(k-1)} \mathbb{E}_{(i,j) \sim X_\tau(1)} f(j) f(i) \\ &= \mathbb{E}_{\tau \sim X(k-1)} \langle P_\tau f, f \rangle \end{aligned}$$

where in the first equality we simply used the localization lemma, and as usual f_τ is the restriction of f to the link of τ. Similarly, we can write

$$\begin{aligned} \langle P^{k\vee} f, f \rangle &= \mathbb{E}_{\tau \sim X(k-1)} \langle P_\tau^{k\vee} f_\tau, f_\tau \rangle \\ &= \mathbb{E}_{\tau \sim X(k-1)} \mathbb{E}_{i \sim X_\tau(0)} \left(\mathbb{E}_{j \sim X_0(\tau)} f(j) \right) f(i) \\ &= \mathbb{E}_{\tau \sim X(k-1)} \langle P_\tau^{0\vee} f_\tau, f_\tau \rangle \end{aligned}$$

It then follows that

$$\begin{aligned} \langle (M^k - P^{k\vee}) f_\tau, f_\tau \rangle &= \mathbb{E}_{\tau \sim X(k-1)} \langle (P_\tau - P_\tau^{0\vee}) f_\tau, f_\tau \rangle \\ &\leq \mathbb{E}_{\tau \sim X(k-1)} \lambda_2(P_\tau) \|f_\tau\|^2 \\ &\leq \max_{\tau \in X(k-1)} \lambda_2(P_\tau) \cdot \|f\| \end{aligned}$$

where the inequality follows by (12.5). This completes the proof of the lemma. □

A stronger version of the above lemma was recently proved by Alev Levi and Lau:

Theorem 12.14. *For $0 \leq k \leq d-1$ and $M^k = P^{k\wedge} - \frac{I}{k+2}$ be the non-lazy up-down walk we have*

$$M^k - P^{k\vee} \preceq_{\Pi_k} \max_{\tau \in X(k-1)} \lambda_2(P_\tau) \cdot (I - P^{k\vee}),$$

where we write $A \preceq_{\Pi_k} B$ when for any $f \in X(k) \to \mathbb{R}$, $\langle Af, f \rangle \leq \langle Bf, f \rangle$.

Lemma 12.15. *For any $0 < k \leq d$, spectrum$(P^{k-1\wedge})$ = spectrum$(P^{k\vee})$, where by spectrum I mean the set of non-zero eigenvalues of the given operator.*

Having the above two lemmas, we can bound the second eigenvalue of $P^{k\wedge}$ inductively.

Corollary 12.16. *If X is a 0-spectral expander, then $\lambda_2(P^{k\wedge}) \leq 1 - \frac{1}{k+2}$.*

Proof. We prove by induction. Firstly, since $\lambda_2(P) \leq 0$, by (12.4) we have $\lambda_2(P^{0\wedge}) = 1/2$. Now, suppose

$$\lambda_2(P^{k-1\wedge}) \leq 1 - \frac{1}{k+1}.$$

By Lemma 12.15, we get that

$$\lambda_2(P^{k\vee}) \leq 1 - \frac{1}{k+1}.$$

By definition of M^k,

$$P^{k\wedge} \quad = \quad \frac{I}{k+2} + \frac{k+1}{k+2}(P^{k\vee} + (M^k - P^{k\vee})).$$

Therefore,

$$\begin{aligned} \lambda_2(P^{k\wedge}) &\leq \frac{1}{k+2} + \frac{k+1}{k+2}(\lambda_2(P^{k\vee}) + \lambda_1(M^k - P^{k\vee})) \\ &\leq \frac{1}{k+2} + \frac{k+1}{k+2}(1 - \frac{1}{k+1} + 0) = 1 - \frac{1}{k+2}, \end{aligned}$$

where in the second inequality we used that, since X is a 0-spectral expander, by Lemma 12.13, so $\lambda_1(M^k - P^{k\vee}) \leq 0$. □

Having the above fact it follows that for any d-homogeneous multilinear log-concave polynomial $p \in \mathbb{R}_{\geq 0}[z_1, \ldots, z_n]$ the corresponding $d-1$-dimensional simplicial complex is a 0-spectral expander and the the $\lambda_2(P^{d-1\vee}) \leq 1 - 1/(k+1)$. In the next section we will see that by simulating the down-up chain we can evaluate p at any point in $\mathbb{R}^n_{\geq 0}$.

If X is a λ-spectral expander it follows from the above proof that

$$\lambda_2(P^{k\wedge}) \leq 1 - \frac{1}{k+2} + \frac{k+1}{2}\lambda.$$

This implies that if $\lambda = O(1/k^2)$ the second eigenvalue is bounded away from 1 and the walk mixes rapidly as we will discuss in the following section.

Following the recent work of Levi and Lau the following stronger theorem holds:

Theorem 12.17. *Let X be a d-dimensional simplicial complex and for any $0 \leq k \leq d-2$ let*

$$\lambda_k = \max_{\sigma \in X(k)} \lambda_2(P_\tau).$$

Then,

$$\lambda_2(P^{k\wedge}) \leq 1 - \frac{1}{k+2}\prod_{i=0}^{k-1}(1 - \lambda_k).$$

In particular, if for all k, $\lambda_{d-k} \leq O(1)/k$ then the second eigenvalue of $P^{k\wedge}$ is bounded away from 1 and the walk mixes rapidly.

12.6 Mixing Time and Approximate Counting

Given a d-dimensional λ-spectral expander X Suppose we are given a function $f : X(k) \to \mathbb{R}$ and we want to estimate

$$\mathbb{E}_{\sigma\sim X(k)} f(\sigma) = \langle f, \mathbf{1}\rangle. \tag{12.6}$$

One way to estimate this quantity is to run the up-down (or down-up) Markov chain started from a fixed state $\sigma \in X(k)$. The question is how long we should run the Markov chain so we can approximate the above quantity?

Theorem 12.18. *Let $Q = P^{\vee k}$. For any $\sigma \in X(k)$, $\epsilon > 0$ and any function $f \in X(k) \to \mathbb{R}$ we have*

$$|Q^t f(\sigma) - \mathbb{E}f| \le \epsilon \|f\|^2,$$

if

$$t \ge \frac{\log\left(\frac{1}{\epsilon} \cdot \max_{\sigma\in X(k)} \frac{1}{\Pi_k(\sigma)}\right)}{1 - \lambda_2(P^{\vee k})}$$

Say $f : X(k) \to [0, 1]$. Then, $\langle f, f\rangle \le 1$. Having this, we can simply estimate $\mathbb{E}f$ within an additive ϵ error (with probability $1 - \delta$) by averaging $f(\tau)$, where τ is the t-th state that the chain lands on, for $O(\frac{1}{\epsilon^2} \log \delta)$ many independent samples.

Proof. Let $F = \mathbb{E}_{\sigma\sim X(k)} f(\sigma)$. We write

$$\begin{aligned}
\max_{\sigma\in X(k)} |Q^t f(\sigma) - F| &\le \max_{\sigma\in X(k)} \frac{1}{\Pi_k(\sigma)} \cdot \mathbb{E}_{\sigma\sim X(k)} |Q^t f(\sigma) - F| \\
&\le \max_{\sigma\in X(k)} \frac{1}{\Pi_k(\sigma)} \cdot \mathbb{E}_{\sigma\sim X(k)} |Q^t f(\sigma) - F|^2 \\
&= \max_{\sigma\in X(k)} \frac{1}{\Pi_k(\sigma)} \cdot \langle Q^t f - F\mathbf{1}, Q^t f - F\mathbf{1}\rangle.
\end{aligned}$$

where the second inequality follows by the Cauchy-Schwarz inequality and Now, let $g_1, \dots, g_{|X(k)|}$ be the eigenfunctions of Q with corresponding eigenvalues $\lambda_1 \ge \dots \ge \lambda_{|X(k)|}$. Recall $\lambda_1 = 1$ and $g_1 = \mathbf{1}$. Also that Q^t have exactly the same eigenfunctions with eigenvalues $\lambda_1^t, \dots, \lambda_{|X(k)|}^t$. Therefore, we can write $f = \sum_i \langle f, g_i\rangle g_i$.

$$\begin{aligned}
\langle Q^t f - F\mathbf{1}, Q^t f - F\mathbf{1}\rangle &= \langle Q^{2t} f, f\rangle - 2F\langle f, \mathbf{1}\rangle + F^2\langle \mathbf{1}, \mathbf{1}\rangle \\
&= \langle Q^{2t} f, f\rangle - F^2 \\
&= \sum_{i=1}^{|X(k)|} \lambda_i^{2t} \langle f, g_i\rangle^2 - F^2 \\
&= \sum_{i=2}^{|X(k)|} \lambda_i^{2t} \langle f, g_i\rangle^2 \\
&\le \lambda_2^{2t} \|f\|^2
\end{aligned}$$

where in the second equality we used (12.6) and that $\|\mathbf{1}\| = 1$. We will use without proof that all eigenvalues of $Q = P^{k\vee}$ are non-negative, i.e., the matrix is PSD. Therefore, for $t \ge \frac{\log\left(\frac{1}{\epsilon} \cdot \max_{\sigma\in X(k)} \frac{1}{\Pi_k(\sigma)}\right)}{1 - \lambda_2(P^{\vee k})}$ we have

$$\max_{\sigma\in X(k)} |Q^t f(\sigma) - F| \le \lambda_2^{2t} \|f\|^2 \max_{\sigma} \frac{1}{\Pi_k(\sigma)} \le \epsilon \cdot \|f\|^2$$

as desired. ∎

Using the above tool let me explain how to use it to count the number of bases of a matroid.

Theorem 12.19. *There is a randomized algorithm that for any given matroid $M = ([n], \mathcal{I})$ of rank $r+1$ counts the number of bases of M within $1 \pm \eta$ multiplicative factor in time polynomial in $n, r, 1/\eta$.*

Proof. First, construct an r dimensional simplicial complex as explained in Example 12.2 with uniform weights on the top faces, i.e., bases. Fix an element n. Observe that either $\mathbb{P}[\sigma \sim X(r)]\, n \in \sigma \geq 1/2$ or $\mathbb{P}[\sigma \sim X(r)]\, n \notin \sigma \geq 1/2$. Without loss of generality assume the former. Consider the following function $f : X(r) \to \{0, 1\}$, where for any $\sigma \in X(r)$,

$$f(\sigma) = \mathbb{I}[n \in \sigma].$$

Since $\|f\| \leq 1$ by the previous theorem we can run a Markov chain starting from an arbitrary state $\sigma \in X(r)$ to find an estimate Z of $\mathbb{E}[f] = \mathbb{P}_\sigma[n \in \sigma]$ within ϵ additive error for $\epsilon < \eta/10n$, i.e.,

$$Z - \epsilon \leq \mathbb{E}[f] \leq Z + \epsilon.$$

All we need to do is to run the chain for $\frac{1}{\epsilon^2 \log \delta}$ many times for $t \geq r + 1 \log(\frac{1}{\epsilon}|X(r)|)$ which is a polynomial in $n, r, 1/\eta$ and calculate the fraction of the states at time t have n. Since $\mathbb{P}_\sigma[n \notin \sigma] \geq 1/2$ this gives a $1 \pm 2\epsilon = 1 \pm \eta/5n$ multiplicative error of the same quantity, i.e.,

$$(1 - \eta/5n)\mathbb{E}[f] \leq Z \leq (1 + \eta/5n)\mathbb{E}[f].$$

Now, we recursively find an estimate N that is a $1 \pm \eta(1 - 1/n)$ multiplicative approximation of the number of bases of $M/\{n\}$, i.e.,

$$(1 - \eta(1 - 1/n))\#\text{Base of } M/\{n\} \leq N \leq (1 + \eta(1 - 1/n))\#\text{Base of } M/\{n\}$$

Since the number of bases of M is equal to number of bases of $M/\{n\}$ divided by $\mathbb{P}[n]$, we can just output N/Z as an estimate. It follows from the above two inequalities that

$$(1 - \eta)|X(r)| \leq N/Z \leq (1 + \eta)|X(r)|$$

as desired. □