CS 860 Spectral graph theory. Spring 2019, Waterloo.

Lecture 10: Higher order random walks

We study two types of random walks on simplicial complexes, called the up-walks and down-walks.

The main result is that they are fast mixing if the simplicial complex is a good link expander.

An immediate corollary is that the natural random walk on matroid bases is fast mixing, proving the matroid expansion conjecture.

Random walks on simplicial complexes

Kaufman and Mass define two natural random walks on faces of dimension k in a simplicial complexes, the up-walk that goes through faces of dimension $k+1$ and the down walk that goes through

To define these walks, consider the bipartite graph G_k with one side corresponding to faces in $X(k)$ and another side corresponding to faces in $X(k+1)$, where a face $\sigma \in X(k)$ has an edge to $\tau \in X(k+1)$ if and only if $\sigma \subset \tau$ and the weight of the edge is $w(\tau)$.

Now, consider the (weighted) random walk on G_k, where a vertex moves to a neighbor proportional to the weight of the edges connecting the two vertices.

As the graph is bipartite, the random walk matrix P_k of G_k is of the form $P_k = \begin{bmatrix} 0 & U_k \\ D_k & 0 \end{bmatrix}$, where U_k is $X(k)$ by $X(k+1)$ and D_k is $X(k+1)$ by $X(k)$.

When the weight function is balanced, the total weight incident on a face $\sigma \in X(k)$ is $\sum_{\tau \in X(k+1), \tau \supset \sigma} w(\tau) = w(\sigma)$, and so $U_k(\sigma, \tau) = \frac{w(\tau)}{w(\sigma)}$.

On the other side, all the edges incident on a face $\tau \in X(k+1)$ is of weight $w(\tau)$, and so $D_k(\tau, \sigma) = \frac{1}{k+2}$ for $\sigma \in X(k)$ and $\sigma \subset \tau$ (recall that $\tau \in X(k+1)$ is of size $k+2$).

Now, consider the two step random walk on G_k. the random walk matrix is $P_k^2 = \begin{bmatrix} U_k D_k & \\ & D_k U_k \end{bmatrix}$, with $U_k D_k$ is $X(k)$ by $X(k)$ and $D_k U_k$ is $X(k+1)$ by $X(k+1)$.

<u>Definition</u> (Up-walk matrix, down-walk matrix of a simplicial complex)

We call $P_k^\wedge := U_k D_k$ the up-walk matrix on $X(k)$ and $P_{k+1}^\vee := D_k U_k$ the down-walk matrix on $X(k+1)$.

We write $P_k^\wedge$ and $P_{k+1}^\vee$ explicitly as follows

For two faces $\sigma, \sigma' \in X(k)$,
$$P_k^\wedge(\sigma, \sigma') = \begin{cases} \frac{1}{k+2} & \text{if } \sigma = \sigma' \\ \frac{w(\sigma \cup \sigma')}{(k+2)\, w(\sigma)} & \text{if } \sigma \cup \sigma' \in X(k+1) \\ 0 & \text{otherwise.} \end{cases}$$

$$\begin{cases} \dfrac{w(\sigma\cup\sigma')}{(k+2)\,w(\sigma)} & \text{if } \sigma\cup\sigma' \in X(k+1) \\ 0 & \text{otherwise.} \end{cases}$$

For two faces $\tau, \tau' \in X(k+1)$,

$$P^{\vee}_{k+1}(\tau,\tau') = \begin{cases} \displaystyle\sum_{\sigma\in X(k):\,\sigma\subset\tau} \frac{w(\tau')}{(k+2)\,w(\sigma)} & \text{if } \tau=\tau' \\ \dfrac{w(\tau')}{(k+2)\,w(\tau\cap\tau')} & \text{if } \tau\cap\tau' \in X(k) \\ 0 & \text{otherwise} \end{cases}$$

Notice that $P^{\wedge}_0$ is just the standard lazy random walk on a graph (here the graph is the 1-skeleton).

<u>Definition</u> (Non-lazy up-walk matrix)

We write $\tilde{P}^{\wedge}_k := \frac{k+2}{k+1}\left(P_k - \frac{I}{k+2}\right)$ as the non-lazy up-walk matrix on $X(k)$.

Explicitly, for two faces $\sigma, \sigma' \in X(k)$,

$$\tilde{P}^{\wedge}_k(\sigma,\sigma') = \begin{cases} \dfrac{w(\sigma\cup\sigma')}{(k+1)\,w(\sigma)} & \text{if } \sigma\cup\sigma' \in X(k+1) \\ 0 & \text{otherwise} \end{cases}$$

The following is an important relation between the spectrum of the up-walk and the down-walk.

<u>Claim</u> $P^{\wedge}_k$ and $P^{\vee}_{k+1}$ have the same non-zero eigenvalues with multiplicity.

<u>proof</u> It follows from the fact that AB and BA have the same non-zero eigenvalues with multiplicity.

The proof is by showing that they have (essentially) the same characteristic polynomials (see e.g. wikipedia page "characteristic polynomial" for a proof). □

To compute the stationary distribution of the up-walk and the down-walk, we use the fact that

if $\pi_i P_{ij} = \pi_j P_{ji}$ for a vector $\pi \in \mathbb{R}^n$ and a row-stochastic matrix $P \in \mathbb{R}^{n\times n}$ (time-reversible)

then π is a Stationary distribution of P, i.e. $\pi P = \pi$. (Check this.)

<u>Claim</u> The stationary distribution of $P^{\wedge}_k$, $\tilde{P}^{\wedge}_k$ and $P^{\vee}_k$ is $w \in \mathbb{R}^{X(k)}$, the weight function on $X(k)$.

<u>proof</u> Check that the time reversible condition is satisfied with w and $P^{\wedge}_k$, $\tilde{P}^{\wedge}_k$, and $P^{\vee}_k$. □

A related fact is that all $P^{\wedge}_k$, $\tilde{P}^{\wedge}_k$ and $P^{\vee}_k$ can be written as $W^{-1}B$ for some symmetric matrix B,

where $W = \operatorname{diag}(w)$ and $w \in \mathbb{R}^{X(k)}$ is the weight function on $X(k)$.

We will later use this fact so that we can use the w-inner product $\langle u,v\rangle_w = \sum_i w(i)u(i)v(i)$ as in the previous lecture to bound the eigenvalues of $P^{\wedge}_k$, $\tilde{P}^{\wedge}_k$ and $P^{\vee}_k$.

Random walks on matroid bases To sample a uniform random basis of a matroid, we consider the

Random walks on matroid bases To sample a uniform random basis of a matroid, we consider the matroid complex with the weight of each basis to be one, and run the down-walk $P_d^{\vee}$.

Then, the stationary distribution is the uniform distribution.

We will prove that the down-walk is fast mixing when the complex is a good link expander.

Since we know that the matroid complex is a strong link expander, this gives a simple and efficient algorithm to sample a uniform matroid basis.

Garland's method

The idea of Garland's method is to decompose the global structure into local structures involving links.

For the random walk matrices, let's consider all the transitions involving a face $\eta \in X(k-1)$.

Let $P_\eta^{\vee}$ be the $X(k)$ by $X(k)$ matrix with $P_\eta^{\vee}(\sigma,\sigma') = \begin{cases} \frac{w(\sigma')}{(k+1)\,w(\eta)} & \text{if } \eta \subseteq \sigma\cap\sigma' \\ 0 & \text{otherwise.} \end{cases}$

Let $\tilde{P}_\eta^{\wedge}$ be the $X(k)$ by $X(k)$ matrix with $\tilde{P}_\eta^{\wedge}(\sigma,\sigma') = \begin{cases} \frac{w(\sigma\cup\sigma')}{(k+1)\,w(\sigma)} & \text{if } \eta = \sigma\cap\sigma' \\ 0 & \text{otherwise.} \end{cases}$

Then, we can write $P_k^{\vee} = \sum_{\eta\in X(k-1)} P_\eta^{\vee}$ and $\tilde{P}_k^{\wedge} = \sum_{\eta\in X(k-1)} \tilde{P}_\eta^{\wedge}$.

Observe that these matrices are (essentially) the same as the up-walk and down-walk matrices on the 0-faces of the link $\eta \in X(k-1)$.

Let $P_{\eta,0}^{\vee}$ be the down walk matrix on the 0-skeletons of the link $\eta\in X(k-1)$.

Then $P_{\eta,0}^{\vee}(i,j) = \frac{w_\eta(j)}{w_\eta(\phi)} = \frac{w(\eta\cup j)}{w(\eta)}$ for $i,j \in X_\eta(0)$.

Let $\tilde{P}_{\eta,0}^{\wedge}$ be the non-lazy up-walk matrix on the 0-skeletons of the link $\eta\in X(k-1)$.

Then $\tilde{P}_{\eta,0}^{\wedge}(i,j) = \begin{cases} \frac{w_\eta(i\cup j)}{w_\eta(i)} = \frac{w(\eta\cup i\cup j)}{w(\eta\cup i)} & \text{if } i\cup j \in X(1) \\ 0 & \text{otherwise.} \end{cases}$

Notice that $P_\eta^{\vee}$ is just $\frac{1}{k+1}$ times the "extended" matrix of $P_{\eta,0}^{\vee}$ (i.e. their dimensions are different, but the extra rows and columns in $P_\eta^{\vee}$ are all zero), with $\eta\cup j = \sigma'$.

Similarly, $\tilde{P}_\eta^{\wedge}$ is just $\frac{1}{k+1}$ times the "extended" matrix of $\tilde{P}_{\eta,0}^{\wedge}$, with $\sigma=\eta\cup i$, $\sigma'=\eta\cup j$ and $\sigma\cup\sigma'=\eta\cup i\cup j$.

Furthermore, $\tilde{P}_{\eta,0}^{\wedge}$ is exactly the random walk matrix of the link $\eta\in X(k-1)$ in the definition of link expander that we defined in L09.

So, if X is an α-link expander, then the second largest eigenvalue of $\tilde{P}^{\wedge}_{\eta,0}$ is at most α for every $\eta \in X(k-1)$.

Note also that both $\tilde{P}^{\wedge}_{\eta}$ and $P^{\vee}_{\eta}$ can be written as $W^{-1}B$ for some symmetric matrix B, and so we can use the w-inner product to reason about the eigenvalues of all these matrices.

<u>Definition</u> Given two matrices P, Q where $P = W^{-1}A$ and $Q = W^{-1}B$ for some symmetric matrices A, B and W a diagonal matrix, we say $P \preccurlyeq_W Q$ if $\langle x, Px \rangle_W \le \langle x, Qx \rangle_W$.

If $P \preccurlyeq_W Q$, then the Courant-Fischer theorem (with respect to this inner product, as discussed in L09) implies that if $p_1 \ge p_2 \ge \dots \ge p_n$ and $q_1 \ge q_2 \ge \dots \ge q_n$ are the eigenvalues of P and Q respectively then $q_i \ge p_i$ for all $1 \le i \le n$.

The following is an important lemma relating the spectrum of the non-lazy up-walk and the down-walk.

<u>Lemma</u> (Kaufman-Oppenheim) If X is an α-link expander, then $\tilde{P}^{\wedge}_{k} \preccurlyeq_W P^{\vee}_{k} + \alpha I$ for all k.

<u>Proof</u> We need to show that $\langle y, (\tilde{P}^{\wedge}_{k} - P^{\vee}_{k}) y \rangle_W \le \alpha \|y\|_W^2$ for all $y \in \mathbb{R}^{X(k)}$.

Recall that $\tilde{P}^{\wedge}_{k} = \sum_{\eta \in X(k-1)} \tilde{P}^{\wedge}_{\eta}$ and $P^{\vee}_{k} = \sum_{\eta \in X(k-1)} P^{\vee}_{\eta}$.

So, the LHS can be written as $\sum_{\eta \in X(k-1)} \langle y, (\tilde{P}^{\wedge}_{\eta} - P^{\vee}_{\eta}) y \rangle_W$.

Consider each term $\langle y, (\tilde{P}^{\wedge}_{\eta} - P^{\vee}_{\eta}) y \rangle_W$.

Recall that $\tilde{P}^{\wedge}_{\eta}$ and $P^{\vee}_{\eta}$ have many zero rows and columns, so we restrict our attention to the non-zero rows and columns of $\tilde{P}^{\wedge}_{\eta}$ and $P^{\vee}_{\eta}$, i.e. the faces $\sigma \in X(k)$ such that $\sigma \supset \eta$.

Call the restriction of the vector y and w to those faces y_η and w_η.

Then, $\langle y, (\tilde{P}^{\wedge}_{\eta} - P^{\vee}_{\eta}) y \rangle_W = \frac{1}{k+1} \langle y_\eta, (\tilde{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0}) y_\eta \rangle_{w_\eta}$.

Write $y_\eta = y_\eta^1 + y_\eta^\perp$ such that y_η^1 is the projection to the first eigenvector (which is parallel to $\vec{1}$) and $y_\eta^\perp$ so that $\langle y_\eta^\perp, y_\eta^1 \rangle_{w_\eta} = 0$.

Then $\langle y_\eta, (\hat{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0}) y_\eta \rangle_{w_\eta}$

$= \langle y_\eta^1 + y_\eta^\perp, (\tilde{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0})(y_\eta^1 + y_\eta^\perp) \rangle_{w_\eta}$

$= \langle y_\eta^1, (\tilde{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0}) y_\eta^1 \rangle_{w_\eta} + \langle y_\eta^\perp, (\tilde{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0}) y_\eta^\perp \rangle_{w_\eta}$ as the cross terms are equal to zero

$= \langle y_\eta^\perp, (\tilde{P}^{\wedge}_{\eta,0} - P^{\vee}_{\eta,0}) y_\eta^\perp \rangle_{w_\eta}$ as $\vec{1}$ is an eigenvector of eigenvalue 1 to both $\tilde{P}^{\wedge}_{\eta,0}$ and $P^{\vee}_{\eta,0}$

$= \langle y_\eta^\perp, \tilde{P}^{\wedge}_{\eta,0} y_\eta^\perp \rangle_{w_\eta}$ as $P^{\vee}_{\eta,0}$ is a rank one matrix

$$= \langle y_\eta^\perp, \tilde{P}^\wedge_{\eta,0} y_\eta^\perp \rangle_{w_\eta} \quad \text{as } P^\vee_{\eta,0} \text{ is a rank one matrix}$$

$$\leq \alpha \| y_\eta^\perp \|^2_{w_\eta} \quad \text{by the assumption that } X \text{ is an } \alpha\text{-link expander.}$$

$$(*) \quad \leq \alpha \| y_\eta \|^2_{w_\eta}$$

$$= \alpha \sum_{\substack{\sigma \in X(k) \\ \sigma \supset \eta}} w(\sigma) y(\sigma)^2.$$

Therefore, $\sum_{\eta \in X(k-1)} \langle y, (\tilde{P}^\wedge_\eta - P^\vee_\eta) y \rangle_w = \frac{1}{k+1} \sum_{\eta \in X(k-1)} \langle y_\eta, (\tilde{P}^\wedge_{\eta,0} - P^\vee_{\eta,0}) y_\eta \rangle_{w_\eta}$

$$\leq \frac{\alpha}{k+1} \sum_{\eta \in X(k-1)} \sum_{\substack{\sigma \in X(k) \\ \sigma \supset \eta}} w(\sigma) y(\sigma)^2$$

$$= \frac{\alpha}{k+1} \sum_{\sigma \in X(k)} \sum_{\substack{\eta \in X(k-1) \\ \eta \subset \sigma}} w(\sigma) y(\sigma)^2$$

$$= \alpha \sum_{\sigma \in X(k)} w(\sigma) y(\sigma)^2 \quad \text{as } |\sigma| = k+1$$

$$= \alpha \| y \|^2_w . \quad \square$$

Informally, the proof says that we replace each (weighted) complete graph in the down-walk of η by an expander graph in the up-walk of η, and so if $P_k^\vee$ corresponds to an expandor graph then $\tilde{P}_k^\wedge$ also corresponds to an expander graph.

Fast mixing

We are ready to bound the second eigenvalue of the up-walk and down walk, and hence the mixing time. The following theorem is by Kaufman and Oppenheim.

Theorem If X is an α-link expander, then the second eigenvalue of $P_k^\vee$ is at most $1 - \frac{1}{k+1} + k\alpha$.

Proof The proof goes by induction on k.

In the base case when $k=0$, the matrix $P_0^\vee$ is of rank one, and thus the second largest eigenvalue is at most zero, and the statement holds.

Now assume the statement holds for k, and we would like to prove the inductive step.

By the above theorem, $\tilde{P}_k^\wedge \preccurlyeq_w P_k^\vee + \alpha I$, and thus the second largest eigenvalue of $\tilde{P}_k^\wedge$ is at most $1 - \frac{1}{k+1} + (k+1)\alpha$.

Recall that $\tilde{P}_k^\wedge = \frac{k+2}{k+1}\left(P_k - \frac{I}{k+2}\right)$, and thus $P_k^\wedge = \frac{k+1}{k+2} \tilde{P}_k^\wedge + \frac{I}{k+2}$.

It follows that the second largest eigenvalue of $P_k^\wedge$ is at most

$$\left(\frac{k+1}{k+2}\right)\left(1 - \frac{1}{k+1} + (k+1)\alpha\right) + \frac{1}{k+2} = 1 - \frac{1}{k+2} + \frac{(k+1)^2\alpha}{k+2} \leq 1 - \frac{1}{k+2} + (k+1)\alpha .$$

Finally, recall that $P_{k+1}^{\vee}$ and $P_k^{\wedge}$ have the same non-zero eigenvalues, and this completes the proof. □

Sampling matroid bases

Recall from L09 that the matroid complex is a 0-link expander, and thus the second eigenvalue of $P_d^{\vee}$ is at most $1-\frac{1}{d+1}$.

By the results in L06, the mixing time is at most $O(d\log N) = O(dn\log d)$, where N is the number of bases and n is the number of elements in the ground set.

It also follows from Cheeger's inequality that the bases exchange graph is an expander graph, proving a long standing conjecture.

Improvements

We have observed some improvements of the Kaufman-Oppenheim theorem.

This is an ongoing joint work with Alev, Anari, Liu and Oveis Gharan.

The following lemma is a more careful version of the lemma of Kaufman-Oppenheim shown above.

Lemma If every link of dimension $k-1$ has second eigenvalue at most α, then $\tilde{P}_k^{\wedge} \preceq_w (1-\alpha) P_k^{\vee} + \alpha I$.

Proof We need to show that $\langle y, (\tilde{P}_k^{\wedge} - P_k^{\vee}) y\rangle_w \le \alpha \langle y, (I-P_k^{\vee}) y\rangle_w$, and this would imply that $\tilde{P}_k^{\wedge} - P_k^{\vee} \preceq \alpha(I-P_k^{\vee})$ as desired.

The proof follows the same lines as above, by writing the LHS as $\sum_{\eta\in X(k-1)} \langle y, (\tilde{P}_\eta^{\wedge} - P_\eta^{\vee}) y\rangle_w$.

Then, $\langle y, (\tilde{P}_\eta^{\wedge} - P_\eta^{\vee}) y\rangle_w = \frac{1}{k+1}\langle y_\eta, (\tilde{P}_{\eta,0}^{\wedge} - P_{\eta,0}^{\vee}) y_\eta\rangle_{w_\eta}$.

We bound each term as before until the equation (*) in the above proof.

Instead of bounding $\|y_\eta^{\perp}\|_{w_\eta}^2$ by $\|y_\eta\|_{w_\eta}^2$, we write the equality $\|y_\eta^{\perp}\|_{w_\eta}^2 = \|y_\eta\|_{w_\eta}^2 - \|y_\eta^{1}\|_{w_\eta}^2$.

So, $\langle y_\eta, (\tilde{P}_{\eta,0}^{\wedge} - P_{\eta,0}^{\vee}) y_\eta\rangle_{w_\eta} \le \alpha(\|y_\eta\|_{w_\eta}^2 - \|y_\eta\|_{w_\eta}^2)$.

As in the above proof, $\sum_{\eta\in X(k-1)} \langle y, (\tilde{P}_\eta^{\wedge} - P_\eta^{\vee}) y\rangle_w = \frac{1}{k+1}\sum_{\eta\in X(k-1)} \langle y_\eta, (\tilde{P}_{\eta,0}^{\wedge} - P_{\eta,0}^{\vee}) y_\eta\rangle_{w_\eta}$

$\le \frac{\alpha}{k+1}\sum_{\eta\in X(k-1)} (\|y_\eta\|_{w_\eta}^2 - \|y_\eta^{1}\|_{w_\eta}^2)$.

Using the same calculation as in the last paragraph of the above proof, $\sum_{\eta\in X(k-1)} \|y_\eta\|_{w_\eta}^2 = (k+1)\|y\|_w^2$.

So, to finish the proof, it remains to show that $\sum_{\eta\in X(k-1)} \|y_\eta^{1}\|_{w_\eta}^2 = (k+1)\langle y, P_k^{\vee} y\rangle$, and this is the only extra calculations done compared to the previous proof.

Recall that $y_\eta = y_\eta^1 + y_\eta^\perp$ where $y_\eta^1 = c\cdot 1$ for some constant c and $\langle y_\eta^1, y_\eta^\perp\rangle_{w_\eta} = 0$.

Therefore, $y_\eta^1 = \frac{\langle y_\eta, 1\rangle_{w_\eta}}{\langle 1,1\rangle_{w_\eta}} 1 = \frac{\sum_{\sigma\in X(k),\sigma\supset\eta} y(\sigma)w(\sigma)}{\sum_{\sigma\in X(k),\sigma\supset\eta} w(\sigma)} 1 = \frac{\sum_{\sigma\in X(k),\sigma\supset\eta} y(\sigma)w(\sigma)}{w(\eta)} 1$,

and hence $\|y_\eta^1\|_{w_\eta}^2 = \frac{\langle y_\eta, 1\rangle_{w_\eta}^2}{\langle 1,1\rangle_{w_\eta}} = \frac{\left(\sum_{\sigma\in X(k),\sigma\supset\eta} y(\sigma)w(\sigma)\right)^2}{w(\eta)}$.

So, $\frac{1}{k+1}\sum_{\eta\in X(k-1)} \|y_\eta^1\|_{w_\eta}^2 = \frac{1}{k+1}\sum_{\eta\in X(k-1)} \frac{\left(\sum_{\sigma\in X(k),\sigma\supset\eta} y(\sigma)w(\sigma)\right)^2}{w(\eta)}$

$$= \sum_{\sigma\in X(k)} y(\sigma)^2 w(\sigma)^2 \sum_{\eta\in X(k-1):\,\eta\subset\sigma} \frac{1}{(k+1)w(\eta)} + \sum_{\sigma\in X(k)} y(\sigma)w(\sigma) \sum_{\substack{\sigma'\in X(k)\\ \sigma'\cap\sigma\in X(k-1)}} y(\sigma')w(\sigma')\cdot\frac{1}{(k+1)\,w(\sigma\cap\sigma')}$$

Recall that $P_k^\vee(\sigma,\sigma) = \sum_{\eta\in X(k-1):\,\eta\subset\sigma} \frac{w(\sigma)}{(k+1)w(\eta)}$, and so the first summation is $\sum_{\sigma\in X(k)} w(\sigma)y(\sigma)P_k^\vee(\sigma,\sigma)y(\sigma)$.

Also, $P_k^\vee(\sigma,\sigma') = \frac{w(\sigma')}{(k+1)\,w(\sigma\cap\sigma')}$, and so the second summation is $\sum_{\sigma\in X(k)} w(\sigma)y(\sigma)\sum_{\substack{\sigma'\in X(k)\\ \sigma'\neq\sigma}} P_k^\vee(\sigma,\sigma')\,y(\sigma')$.

Combining these two sums in the above equality, we have

$$\frac{1}{k+1}\sum_{\eta\in X(k-1)}\|y_\eta^1\|_{w_\eta}^2 = \sum_{\sigma\in X(k)} w(\sigma)y(\sigma)\sum_{\sigma'\in X(k)} P_k^\vee(\sigma,\sigma')\,y(\sigma') = \sum_{\sigma\in X(k)} w(\sigma)y(\sigma)(P_k^\vee y)(\sigma) = \langle y, P_k^\vee y\rangle_w. \ \square$$

Somewhat surprisingly, the extra term that we saved is enough to prove a much sharper bound on the second eigenvalue of the down walk.

<u>Theorem</u> If X is an α-link expander, then the second eigenvalue of $P_k^\vee$ is at most $1-\frac{1}{k+1}\prod_{i=0}^{k-2}(1-\alpha_i)$, where α_i is the maximum second eigenvalue of the links of dimension i.

<u>Proof</u> The proof goes by induction on k.

In the base case when $k=0$, the matrix $P_0^\vee$ is of rank one, and thus the second largest eigenvalue is at most zero, and the statement holds.

Now assume the statement holds for $k\geq 0$, and we would like to prove the inductive step for $k\geq 1$.

By the above lemma, $\tilde{P}_k^\wedge \preccurlyeq_w (1-\alpha_{k-1})P_k^\vee + \alpha_{k-1} I$, and thus the second largest eigenvalue of $\tilde{P}_k^\wedge$ is at most $(1-\alpha_{k-1})\left(1-\frac{1}{k+1}\prod_{i=0}^{k-2}(1-\alpha_i)\right)+\alpha_{k-1} = 1-\frac{1}{k+1}\prod_{i=0}^{k-1}(1-\alpha_i)$.

Recall that $\tilde{P}_k^\wedge = \frac{k+2}{k+1}\left(P_k - \frac{I}{k+2}\right)$, and thus $P_k^\wedge = \frac{k+1}{k+2}\tilde{P}_k^\wedge + \frac{I}{k+2}$.

It follows that the second largest eigenvalue of $P_k^\wedge$ is at most

$$\left(\frac{k+1}{k+2}\right)\left(1-\frac{1}{k+1}\prod_{i=0}^{k-1}(1-\alpha_i)\right)+\frac{1}{k+2} = 1-\frac{1}{k+2}\prod_{i=0}^{k-1}(1-\alpha_i).$$

It follows that the second largest eigenvalue of P_k is at most

$$\left(\frac{k+1}{k+2}\right)\left(1-\frac{1}{k+1}\prod_{i=0}^{k-1}(1-\alpha_i)\right)+\frac{1}{k+2} = 1-\frac{1}{k+2}\prod_{i=0}^{k-1}(1-\alpha_i).$$

Finally, recall that $P_{k+1}^{\vee}$ and $P_k^{\wedge}$ have the same non-zero eigenvalues, and this completes the proof. $\square$

Recall that Kaufman and Oppenheim proved that the spectral gap of $P_k^{\vee}$ is at least $\frac{1}{k+1}-k\alpha$, and this is positive only if $\alpha = O(\frac{1}{k^2})$ where α is the maximum second eigenvalue of any link.

Using the above theorem, the spectral gap is at least $\frac{1}{k+1}\prod_{i=0}^{k-2}(1-\alpha_i)$, which is always positive as long as all $\alpha_i < 1$, and the spectral gap is $\Omega(\frac{1}{k})$ as long as $\alpha = O(\frac{1}{k})$.

Combining with Oppenheim's theorem, this is useful for the analysis of mixing time of Markov chains.

Very recently, Liu, Mohanty and Yang showed a very interesting construction of high dimensional expanders from expander graphs.

In their construction, every link has second eigenvalue at most $1/2$, so that Kaufman-Oppenheim's result cannot be directly applied.

Nevertheless, they proved that the spectral gap of $P_k^{\vee}$ is $\Omega(\frac{1}{k}\cdot\frac{1}{2^k})$ using some techniques in Markov chains, for their specific construction.

Our theorem recovers their spectral gap result using only the eigenvalue of the links.

References

- Higher order random walks: beyond spectral gap, by Kaufman and Oppenheim
- Log-concave polynomial II: High-dimensional walks and an FPRAS for counting bases of a matroid, by Anari, Liu, Oveis Gharan, Vinzant.
- High dimensional expanders from expanders, by Liu, Mohanty and Yang.