SAMPLING SPANNING TREES USING HDXS
(CSS.413.1: PSEUDORANDOMNESS - LECTURE 28)

Prahladh Harsha[*]

TIFR, Mumbai

In these notes, we give a self-contained exposition of the beautiful result of Nima Anari, Kuikui Liu, Shayan Oveis-Gharan and Cynthia Vinzant [ALOV19] that the Glauber Dynamics on spanning trees of a graph mixes in polynomial time.

CONTENTS

1 GLAUBER DYNAMICS ON SPANNING TREES

Let $G = (V, E)$ be an unweighted undirected connected graph. Let $\mathcal{T}_G$ be the set of spanning trees of G. Note that $\mathcal{T}_G$ can be exponentially large compared to the size of the graph (here let $n := |V|$). Consider the following random walk on $\mathcal{T}_G$, more commonly referred to as the *Glauber Dynamics* on spanning trees.

- On input $T \in \mathcal{T}_G$
  - Choose a uniformly random edge $e \in T$.

[*] prahladh@tifr.res.in

– Set $F \leftarrow T \setminus \{e\}$.
– Let $A, B \subset V$ be the two components of the forest F.
– Choose a uniformly random edge $e' \in E(A, B)$.
– Set $T' \leftarrow F \cup \{e'\}$.
– Output T'

We will refer to this random walk on the state space $\mathcal{T}_G$ as GD_G. It is easy to see that the stationary distribution for this walk is the uniform distribution $U_{\mathcal{T}_G}$ on $\mathcal{T}_G$, i.e., $U_{\mathcal{T}_G} \cdot \mathrm{GD}_G = U_{\mathcal{T}_G}$.

In a remarkable coming together of ideas from Markov chain sampling and high-dimensional expanders, Anari, Liu, Oveis-Gharan and Vinzant proved the following theorem bounding the spectral-gap of the random walk GD_G.

The spectral gap γ of a random walk is defined to be $1 - \max\{\lambda_2, |\lambda_n|\}$.

1.1 THEOREM ([ALOV19]). $\gamma(\mathrm{GD}_G) \geq \frac{1}{n-1}$.

This theorem proves that the spectral gap of the Glauber Dynamics is at least inverse polynomially large (in n) even though the state space of the random walk could be exponentially large in n. This immediately yields that the Glauber Dynamics on $\mathcal{T}_G$ mixes in polynomial time by the following well-known theorem on the mixing time of random walks in terms of their spectral gap.

1.2 THEOREM. *Let P be a random walk with stationary distribution π and spectral gap $\gamma \in (0,1)$. Then, the mixing time $t(\varepsilon)$ of the random walk P is upper bounded as follows:*

$$t(\varepsilon) \leq \frac{1}{\gamma}\left[\frac{1}{2}\log\left(\frac{1}{\pi_{\min}}\right) + \log\left(\frac{1}{2\varepsilon}\right)\right].$$

This gives a fast algorithm to *approximately* sample a uniformly random spanning tree in given undirected graph. The above Glauber dynamics has the advantage that it verbatim extends to random walks on bases of a matroid. Furthermore, the same approach can be used to give a bound on the modified log-Sobolev constant of GD_G, yielding even better bounds on the mixing time of the random walk [CGM19]. However, for the purpose of these notes, we will restrict our attention to sampling spanning trees of a given undirected graph.

We remark that there are other sampling algorithms for spanning trees (c.f., the extremely clever and cute sampling algorithm of Broder [Bro89] and Aldous [Ald90]).

We begin with some preliminaries on high-dimensional expanders (HDXs).

2 HIGH-DIMENSIONAL EXPANDERS

A *simplicial complex* X is a down-closed collection of sets. We will refer to the sets in X as faces. The *dimension* of a face $s \in X$ is $|s| - 1$. The dimension of X, denoted by $\dim(X)$, is the maximal dimension of any face $s \in X$. We let $X(i)$ denote the set of i-dimensional faces in X, also referred to as the set of i*-faces*. Note that if X is non-empty, then $X(-1) = \{\emptyset\}$. We will restrict our attention to

Note that $X(i)$ refers to the set of $(i+1)$-sized, and not i-sized, sets in X. In particular, $X(0)$ refers to the set of singletons.

pure simplicial complexes where all maximal faces have the same dimension, namely $\dim(X)$.

Let X be a k-dimensional simplicial complex. The simplicial complexes X we work with are typically accompanied with a probability distribution Π_k on the set of k-dimensional faces. If no distribution is explicitly specified, we assume that the distribution is the uniform distribution on $X(k)$. The distribution Π_k induces a joint distribution $\Pi = (\Pi_k, \Pi_{k-1}, \ldots, \Pi_0, \Pi_{-1})$ on $X(k) \times X(k-1) \times \cdots \times X(0) \times X(-1)$ as follows: pick a k-face $t_k \sim \Pi_k$, choose a random ordering $v_1, v_2, \ldots, v_{k+1}$ of the $k+1$ elements in t_k and set $t_{i-1} \leftarrow t_i \setminus \{v_{i+1}\}$ for $i \leftarrow k$ to 0. Then $(t_k, t_{k-1}, \ldots, t_0, t_{-1} = \emptyset) \sim \Pi$. We will refer to the pair (X, Π) as a *weighted simplicial complex*.

For each $-1 \leq i \leq k$, we define the function spaces $C(i)$ as follows:

$$C(i) := \{f \colon X(i) \to \mathbb{C}\}.$$

We equip these complex vector spaces $C(i)$'s with inner products as follows. Given functions $f, g \in C(i)$, the inner product $\langle \cdot, \cdot \rangle_{\Pi_i}$ is defined as

$$\langle f, g \rangle_{\Pi_i} := \mathop{\mathbb{E}}_{s \sim \Pi_i} \left[f(s) \cdot \overline{g(s)} \right].$$

We will drop the subscript Π_i if the domain of the functions f, g are clear from context.

2.1 DEFINITION (link). Let (X, Π) be a weighted simplical complex. Given any face $s \in X$, the *link of* s, denoted by $(X_s, \Pi^{(s)})$, is the following weighted simplicial complex.

$$X_s := \{t \setminus s \colon t \supset s, t \in X\}.$$

If s is an i-face then X_s is a $(k-i-1)$-dimensional simplical complex. The joint distribution $\Pi^{(s)}$ of the link X_s is the distribution Π conditioned on the facing containing s. More precisely, for any $-1 \leq j \leq k-i-1$, we have

$$\Pi_j^{(s)}(t') := \frac{\Pi_{j+i+1}(t' \cup s)}{\sum\limits_{t \in X(j+i+1) \colon t \supset s} \Pi_{j+i+1}(t)}.$$

Given a simplicial complex (X, Π) of dimension at least 1, the *underlying graph of* X, also referred to as the *1-skeleton* of X and denoted by $G(X)$, is the (weighted) graph given by $(X(0), X(1), \Pi_1)$. There is a natural random walk $P_{G(X)}$ on the vertices of this graph, induced by the edge distribution Π_1. For any $u, v \in X(0)$,

$$P_{G(X)}[u \to v] := \frac{\Pi_1(\{u, v\})}{\sum\limits_{e \in X(1) \colon e \ni u} \Pi_1(e)}.$$

Thus, in the language of links, we have $P_{G(X)}[u \to v] = \Pi_0^{(u)}(v)$. Hence, for any function $f \colon X(0) \to \mathbb{C}$, we have $P_{G(X)} f \colon X(0) \to \mathbb{C}$ given by the following expression.

$$(\mathrm{P}_{\mathrm{G(X)}}f)(u) = \mathop{\mathbb{E}}_{v\sim\Pi_0^{(u)}}[f(v)].$$

This gives the following nice expression for inner products of the form $\langle f, \mathrm{P}_{\mathrm{G(X)}}g\rangle$ where $f, g \in \mathrm{C}(0)$.

$$\begin{aligned}
\langle f, \mathrm{P}_{\mathrm{G(X)}}g\rangle &= \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[f(u)\cdot\overline{(\mathrm{P}_{\mathrm{G(X)}}g)(u)}\right] \\
&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[f(u)\cdot \mathop{\mathbb{E}}_{v\sim\Pi_0^{(u)}}\left[\overline{g(v)}\right]\right] \\
&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\mathop{\mathbb{E}}_{v\sim\Pi_0^{(u)}}\left[f(u)\cdot\overline{g(v)}\right] \\
&= \mathop{\mathbb{E}}_{\{u,v\}\sim\Pi_1}\left[f(u)\cdot\overline{g(v)}\right].
\end{aligned}$$

A similar calculation for $\langle \mathrm{P}_{\mathrm{G(X)}}f, g\rangle$ shows that $\langle f, \mathrm{P}_{\mathrm{G(X)}}g\rangle = \langle \mathrm{P}_{\mathrm{G(X)}}f, g\rangle$. In other words, $\mathrm{P}_{\mathrm{G(X)}}$ is self-adjoint (with respect to the inner product $\langle\cdot,\cdot\rangle_{\Pi_0}$) and hence has a complete eigen decomposition and real eigenvalues. We denote these eigenvalues by $1 = \lambda_1(\mathrm{G(X)}) \geq \lambda_2(\mathrm{G(X)}) \geq \cdots \geq \lambda_n(\mathrm{G(X)}) \geq -1$.

The following proposition is an easy consequence of the definition of inner product and link. Given any $f\colon \mathrm{X}(0)\to\mathbb{C}$ and $u\in\mathrm{X}(0)$, let $f_u\colon \mathrm{X}_u(0)\to\mathbb{C}$ be the restriction of the function f to $\mathrm{X}_u(0)$.

2.2 PROPOSITION. *For $f, g\colon \mathrm{X}(0)\to\mathbb{C}$, we have*

$$\langle f, g\rangle_{\Pi_0} = \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[\langle f_u, g_u\rangle_{\Pi_0^{(u)}}\right],$$

$$\langle \mathrm{P}_{\mathrm{G(X)}}f, g\rangle_{\Pi_0} = \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[\langle \mathrm{P}_{\mathrm{G(X}_u)}f_u, g_u\rangle_{\Pi_0^{(u)}}\right].$$

Proof.

$$\begin{aligned}
\langle f, g\rangle_{\Pi_0} &= \mathop{\mathbb{E}}_{v\sim\Pi_0}\left[f(v)\cdot\overline{g(v)}\right] &&= \mathop{\mathbb{E}}_{\{u,v\}\sim\Pi_1}\left[f(v)\cdot\overline{g(v)}\right] \\
&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\mathop{\mathbb{E}}_{v\sim\Pi_0^{(u)}}\left[f(v)\cdot\overline{g(v)}\right] &&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[\langle f_u, g_u\rangle_{\Pi_0^{(u)}}\right]. \\
\langle \mathrm{P}_{\mathrm{G(X)}}f, g\rangle_{\Pi_0} &= \mathop{\mathbb{E}}_{\{v,w\}\sim\Pi_1}\left[f(v)\cdot\overline{g(w)}\right] &&= \mathop{\mathbb{E}}_{\{u,v,w\}\sim\Pi_2}\left[f(v)\cdot\overline{g(w)}\right] \\
&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\mathop{\mathbb{E}}_{\{v,w\}\sim\Pi_1^{(u)}}\left[f(v)\cdot\overline{g(w)}\right] &&= \mathop{\mathbb{E}}_{u\sim\Pi_0}\left[\langle \mathrm{P}_{\mathrm{G(X}_u)}f_u, g_u\rangle_{\Pi_0^{(u)}}\right]. \quad\square
\end{aligned}$$

We now study various types of random walks on the faces of the simplicial complex.

2.1 *Up-down and Down-up walks*

There are two natural walks we can define on a the set $X(i)$ of i-faces.

- Up-Down walk $P_i^{\triangle}$:
  - On input $s \in X(i)$
    - * Choose a random $t \in X(i+1)$ from the distribution Π_{i+1} conditioned on $t \supset s$.
    - * Choose a random $v \in t$ and set $s' \leftarrow t \setminus \{v\}$.
    - * Output s'
- Down-Up walk $P_i^{\triangledown}$:
  - On input $s \in X(i)$
    - * Choose a random $v \in s$ and set $r \leftarrow s \setminus \{v\}$.
    - * Choose a random $s' \in X(i)$ from the distribution Π_i conditioned on $s' \supset r$.
    - * Output s'

The stationary distribution for both these walks is the distribution Π_i on layer $X(i)$. It is not hard to see that both these walks have a lazy component, i.e., for each i-face s there is a non-zero probability that the walk returns to the i-face s. We let $P_i^{\wedge}$ and $P_i^{\vee}$ be the corresponding non-lazy walks. The up-down walk $P_i^{\triangle}$ has a lazy $1/i+2$ lazy component. More precisely, rhe up-down walk $P_i^{\triangle}$ has the following nice decomposition into its lazy and non-lazy components.

The down-up walk $P_i^{\triangledown}$ does not necessarily have such a clean decomposition in terms of the corresponding non-lazy walk $P_i^{\vee}$. Why?

$$P_i^{\triangle} = \frac{1}{i+2} I_{X(i)} + \frac{i+1}{i+2} P_i^{\wedge}. \tag{1}$$

The (lazy) up-down and down-up walks can be further broken down in terms of a down and up walks as follows:

- Up walk $U_{i \to i+1}$:
  - On input $s \in X(i)$
    - * Choose a random $t \in X(i+1)$ from the distribution Π_{i+1} conditioned on $t \supset s$.
    - * Output $t \in X(i+1)$
- Down walk $D_{i \to i-1}$:
  - On input $s \in X(i)$
    - * Choose a random $v \in s$ and set $r \leftarrow s \setminus \{v\}$.
    - * Output $r \in X(i-1)$

It follows from the definitions that $P_i^{\triangle} = U_{i\to i+1} D_{i+1\to i}$ while $P_{i+1}^{\triangledown} = D_{i+1\to i} U_{i\to i+1}$. An immediate consequence of this decomposition of the up-down and down-up walks in terms of the up and down walks is the following.

If A and B are $r \times s$ and $s \times r$ matrices respectively, then AB and BA share all non-zero eigenvalues.

$$\lambda_2(P_i^{\triangle}) = \lambda_2(P_{i+1}^{\triangledown}). \tag{2}$$

This decomposition can be further used to show that the operators $P_i^{\triangle}$ and $P_{i+1}^{\triangledown}$ are positive semidefinite operators.

$$\begin{aligned}
\langle P_i^{\triangledown} f, f\rangle_{\Pi_i} &= \mathop{\mathbb{E}}_{s\sim\Pi_i}\left[(P_i^{\triangledown} f)(s)\cdot\overline{f(s)}\right] \\
&= \mathop{\mathbb{E}}_{s\sim\Pi_i}\left[(D_{i\to i-1} U_{i-1\to i} f)(s)\cdot\overline{f(s)}\right] \\
&= \mathop{\mathbb{E}}_{s\sim\Pi_i}\left[\mathop{\mathbb{E}}_{r\sim\Pi_{i-1}:\, r\subset s}\left[\mathop{\mathbb{E}}_{s'\sim\Pi_i:\, s'\supset r}\left[f(s')\right]\right]\cdot\overline{f(s)}\right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\mathop{\mathbb{E}}_{s'\sim\Pi_i:\, s'\supset r}\left[f(s')\right]\cdot\overline{\mathop{\mathbb{E}}_{s\sim\Pi_i:\, s\supset r}\left[f(s)\right]}\right] \qquad (3)\\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[(U_{i-1\to i} f)(r)\cdot\overline{(U_{i-1\to i} f)(r)}\right] \\
&= \langle U_{i-1\to i} f, U_{i-1\to i} f\rangle_{\Pi_{i-1}}.
\end{aligned}$$

A similar calculation shows $\langle P_i^{\triangle} f, f\rangle_{\Pi_i} = \langle D_{i+1\to i} f, D_{i+1\to i} f\rangle_{\Pi_{i+1}}$. Hence, both these operators are positive semidefinite.

2.2 *Link expansion*

2.3 DEFINITION (Link-HDX). A weighted simplicial complex (X, Π) is said to be a λ-onesided link-HDX if for every $-1 \le i < \dim(X)$ and $s \in X(i)$, we have that the underlying graph $G(X_s)$ of the link $(X_s, \Pi^{(s)})$ satisfies $\lambda_2(G(X_s)) \le \lambda$.

Similarly, (X, Π) is said to be a λ-twosided link-HDX if every face s satisfies $\max\{\lambda_2(G(X_s)), |\lambda_n(G(X_s))|\} \le \lambda$ (i.e, eigenvalue bounds on both sides). However, we won't need twosided link expansion for these notes.

The following theorem shows that if a simplical complex (X, Π) is a λ-onsided-link-HDX, then the non-lazy up-down walk can be λ-approximated by the down-up walk on the same layer (at least in one direction).

2.4 THEOREM ([KO20, DDFH18]). *If (X, Π) is a λ-onesided-link-HDX, then for every $0 \le i < \dim(X)$, we have*

$$P_i^{\wedge} - P_i^{\triangledown} \preccurlyeq \lambda I.$$

Proof. For any function $f\colon X(i) \to \mathbb{C}$ and $(i-1)$-face $r \in X(i-1)$, let $f_r\colon X_r(0) \to \mathbb{C}$ be the restriction of f to $X_r(0)$ defined as: $f_r(u) := f(r \cup \{u\})$.

To show that $P_i^{\wedge} - P_i^{\triangledown} \preccurlyeq \lambda I$, it suffices to show that for every $f \in C(i)$, we have $\langle (\mathrm{P}_i^{\wedge} - \mathrm{P}_i^{\triangledown}) f, f\rangle_{\Pi_i} \leq \lambda \langle f, f\rangle_{\Pi_i}$. To this end, we first express the inner products $\langle \mathrm{P}_i^{\wedge} f, f\rangle$ and $\langle \mathrm{P}_i^{\triangledown} f, f\rangle$ in terms of links $r \sim \Pi_{i-1}$.

We begin with the inner product $\langle \mathrm{P}_i^{\triangledown} f, f\rangle$. We know from (3) that

$$
\begin{aligned}
\langle \mathrm{P}_i^{\triangledown} f, f\rangle &= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \left[\mathop{\mathbb{E}}_{s'\sim\Pi_i :\, s'\supset r} [f(s)] \cdot \overline{\mathop{\mathbb{E}}_{s\sim\Pi_i :\, s\supset r} [f(s)]} \right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \left[\mathop{\mathbb{E}}_{u\sim\Pi_0^{(r)}} [f_r(u)] \cdot \overline{\mathop{\mathbb{E}}_{v\sim\Pi_0^{(r)}} [f_r(v)]} \right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \left[\mathop{\mathbb{E}}_{v\sim\Pi_0^{(r)}} \left[(\mathrm{J}_r f_r)(v) \cdot \overline{f_r(v)} \right] \right] \qquad \text{where } (\mathrm{J}_r f_r)(v) := \mathop{\mathbb{E}}_{u\sim\Pi_0^{(r)}} [f_r(u)] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \left[\langle \mathrm{J}_r f_r, f_r\rangle_{\Pi_0^{(r)}} \right]. \qquad (4)
\end{aligned}
$$

Observe that $(\mathrm{J}_r f_r)(v)$ is independent of v and hence $\mathrm{J}_r f_r = \mathrm{E}_{u\sim\Pi_0^{(r)}} [f_r(u)] \cdot \mathbb{1}_{\mathrm{X}_r(0)}$ where $\mathbb{1}_{\mathrm{X}_r(0)} \colon \mathrm{X}_r(0) \to \mathbb{C}$ is the constant one function on $\mathrm{X}_r(0)$.

We now move to the other inner product $\langle \mathrm{P}_i^{\wedge} f, f\rangle$. Let us first try to understand the non-lazy operator $\mathrm{P}_i^{\wedge}$. For any $s \in \mathrm{X}(i)$, we have

$$
(\mathrm{P}_i^{\wedge} f)(s) = \mathop{\mathbb{E}}_{u\sim\Pi_0^{(s)}} \mathop{\mathbb{E}}_{v\in s} f(s \cup \{u\} \setminus \{v\}) \qquad = \mathop{\mathbb{E}}_{r\sim\Pi_{i-1} :\, r\subset s} \mathop{\mathbb{E}}_{u\sim\Pi_0^{(s)}} f(r \cup \{u\}).
$$

Hence,

Recall that $\mathrm{G}(\mathrm{X}_r)$ refers to the underlying graph of the link $(\mathrm{X}_r, \Pi^{(r)})$ and $\mathrm{P}_{\mathrm{G}(\mathrm{X}_r)}$ the random walk on this graph.

$$
\begin{aligned}
\langle \mathrm{P}_i^{\wedge} f, f\rangle_{\Pi_i} &= \mathop{\mathbb{E}}_{s\sim\Pi_i} \left[\left(\mathop{\mathbb{E}}_{r\sim\Pi_{i-1} :\, r\subset s} \mathop{\mathbb{E}}_{u\sim\Pi_0^{(s)}} f(r\cup\{u\}) \right) \cdot \overline{f(s)} \right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \mathop{\mathbb{E}}_{\{u,v\}\sim\Pi_1^{(r)}} \left[f(r\cup\{u\}) \cdot \overline{f(r\cup\{v\})} \right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \mathop{\mathbb{E}}_{\{u,v\}\sim\Pi_1^{(r)}} \left[f_r(u) \cdot \overline{f_r(v)} \right] \\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}} \left[\langle \mathrm{P}_{\mathrm{G}(\mathrm{X}_r)} f_r, f_r\rangle_{\Pi_0^{(r)}} \right]. \qquad (5)
\end{aligned}
$$

For any $r \in \mathrm{X}(i-1)$, we can decompose the vector f_r as

$$
f_r = \mathrm{E}_{u\sim\Pi_0^{(r)}} [f_r(u)] \cdot \mathbb{1}_{\mathrm{X}_r(0)} + f_r^{\perp} = \mathrm{J}_r f_r + f_r^{\perp}
$$

where $\langle f_r^{\perp}, \mathbb{1}_{\mathrm{X}_r(0)}\rangle_{\Pi_0^{(r)}} = 0$. Applying the operator $\mathrm{P}_{\mathrm{G}(\mathrm{X}_r)}$ to f_r, we have

$$
\mathrm{P}_{\mathrm{G}(\mathrm{X}_r)} f_r = \mathrm{J}_r f_r + \mathrm{P}_{\mathrm{G}(\mathrm{X}_r)} f_r^{\perp}. \qquad (6)
$$

Since X is a λ-onesided-link-HDX, we have that $\langle \mathrm{P}_{\mathrm{G}(\mathrm{X}_r)} g, g\rangle_{\Pi_0^{(r)}} \leq \lambda \langle g, g\rangle_{\Pi_0^{(r)}}$ for any g satisfying $\langle g, \mathbb{1}_{X_r(0)}\rangle_{\Pi_0^{(r)}} = 0$. We are now ready to bound $\langle (P_i^{\wedge} - P_i^{\nabla})f, f\rangle$.

$$\begin{aligned}
\langle (P_i^{\wedge} - P_i^{\nabla})f, f\rangle_{\Pi_i} &= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\langle (P_{G(X_r)} - J_r)f_r, f_r\rangle_{\Pi_0^{(r)}}\right] && \text{[By (4) and (5)]}\\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\langle P_{G(X_r)} f_r^{\perp}, f_r\rangle_{\Pi_0^{(r)}}\right] && \text{[By (6)]}\\
&= \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\langle P_{G(X_r)} f_r^{\perp}, f_r^{\perp}\rangle_{\Pi_0^{(r)}}\right]\\
&\leq \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\lambda\cdot\langle f_r^{\perp}, f_r^{\perp}\rangle_{\Pi_0^{(r)}}\right]\\
&\leq \lambda \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\left[\langle f_r, f_r\rangle_{\Pi_0^{(r)}}\right]\\
&= \lambda\cdot \mathop{\mathbb{E}}_{r\sim\Pi_{i-1}}\mathop{\mathbb{E}}_{u\sim\Pi_0^r}\left[f(r\cup\{u\}\cdot\overline{f(r\cup\{u\})}\right]\\
&= \lambda\cdot \mathop{\mathbb{E}}_{s\sim\Pi_i}\left[f(s)\cdot\overline{f(s)}\right]\\
&= \lambda\cdot\langle f, f\rangle_{\Pi_i}.
\end{aligned}$$

Hence, $P_i^{\wedge} - P_i^{\nabla} \preccurlyeq \lambda I$. Thus, proved. □

2.3 *Oppenheim's Trickle-down Theorem*

Theorem 2.4 tells us that in order to show that the non-lazy up-down walk is close to the down-up walk, it suffices to show that X is a λ-onesided-link-HDX. The following theorem, due to Oppenheim [Opp18], says that it further suffices to show that the links corresponding to $X(k-2)$ are expanding.

2.5 THEOREM ([Opp18]). *Suppose* (X,Π) *is a k-dimensional weighted simplicial complex with the following properties.*

- *For all* $s \in X(k-2)$*, the link* $(X_s, \Pi^{(s)})$ *is a* λ*-onesided-link-HDX.*
- *The* 1*-skeleton of every link is connected.*

Then, (X,Π) *is a* $\left(\frac{\lambda}{1-(d-1)\lambda}\right)$*-onesided-link-HDX.*

This theorem is in turn proved by proving the following 2-dimensional version.

2.6 THEOREM. *Suppose* (X,Π) *is weighted* 2*-dimensional simplicial complex with the following two properties*

- *the 1-skeleton of* X *is connected and,*
- *for every vertex* $v \in X(0)$ *and for all* $f : X_v(0) \to \mathbb{C}$ *with* $f \perp \mathbb{1}_{X_v(0)}$*, we have*

$$\langle P_{G(X_v)} f, f\rangle_{\Pi_0^{(v)}} \leq \lambda\cdot\langle f, f\rangle_{\Pi_0^{(v)}}.$$

Then, for any $g\colon X(0) \to \mathbb{C}$ *with* $g \perp \mathbb{1}_{X(0)}$*, we have*

$$\langle P_{G(X)} g, g \rangle_{\Pi_0} \leq \frac{\lambda}{1-\lambda} \cdot \langle g, g \rangle_{\Pi_0}.$$

Let us first see how the 2-dimensional version implies the general trickle-down Theorem 2.5.

Proof of Theorem 2.5. For any $i \leq k-2$, let

$$\lambda_i := \min_{v \in X(i)} \max_{\substack{g: X_v(0) \to \mathbb{C} \\ g \perp \mathbb{1}_{X_v(0)}}} \frac{\langle P_{G(X_v)} g, g \rangle_{\Pi_0^{(v)}}}{\langle g, g \rangle_{\Pi_0^{(v)}}},$$

the smallest link expansion with respect to $X(i)$. From repeated applications of Theorem 2.6, we obtain

$$\lambda_{-1} \leq \frac{\lambda_0}{1-\lambda_0} \leq \frac{\lambda_1/(1-\lambda_1)}{1-(\lambda_1/(1-\lambda_1))} = \frac{\lambda_1}{1-2\lambda_1} \leq \cdots \leq \frac{\lambda_{d-2}}{1-(d-1)\lambda_{d-2}}$$

which eventually completes the proof of the trickle-down theorem. □

We now prove the 2-dimensional trickle-down Theorem 2.6

Proof of Theorem 2.6. Let $g\colon X(0) \to \mathbb{C}$ be an eigenvector that that maximises $\langle P_{G(X)} g, g \rangle_{\Pi_0}$ while satisfying $\langle g, g \rangle_{\Pi_0} = 1$ and $g \perp \mathbb{1}_{X(0)}$. Let $\eta := \langle P_{G(X)} g, g \rangle_{\Pi_0}$ be the maximal value attained. In particular, $P_{G(X)} g = \eta \cdot g$. From Theorem 2.2 we have $\eta = \langle P_{G(X)} g, g \rangle_{\Pi_0} = \mathbb{E}_{v \sim \Pi_0} \left[\langle P_{G(X_v)} g_v, g_v \rangle_{\Pi_0^{(v)}} \right]$.

Let $g_v\colon X_v(0) \to \mathbb{C}$ be the restriction of g to $X_v(0)$, i.e., $g_v(u) = g(u)$. Even though $g \perp \mathbb{1}_{X(0)}$, the *local* component g_v need not be perpendicular to $\mathbb{1}_{X_v(0)}$. Hence, let us write $g_v = \alpha_v \mathbb{1}_{X_v(0)} + g_v^{\perp}$ where $g_v^{\perp} \perp \mathbb{1}_{X_v(0)}$. Here $\alpha_v = \langle g_v, \mathbb{1}_{X_v(0)} \rangle_{\Pi_0^{(v)}} = \mathbb{E}_{u \sim \Pi_0^{(v)}} [g(u)] = (P_{G(X_v)} g)(v)$. We can now use this decomposition as follows.

$$\begin{aligned} \eta = \langle P_{G(X)} g, g \rangle_{\Pi_0} &= \mathop{\mathbb{E}}_{v \sim \Pi_0} \left[\langle P_{G(X_v)} g_v, g_v \rangle_{\Pi_0^{(v)}} \right] \\ &= \mathop{\mathbb{E}}_{v \sim \Pi_0} \left[\alpha_v^2 + \langle P_{G(X_v)} g_v^{\perp}, g_v^{\perp} \rangle_{\Pi_0^{(v)}} \right]. \end{aligned} \tag{7}$$

To further simplify the above expression, we make two observations.

- By the hypothesis, since $g_v^{\perp} \perp \mathbb{1}_{X_v(0)}$, we have

$$\langle P_{G(X_v)} g_v^{\perp}, g_v^{\perp} \langle_{\Pi_0^{(v)}} \leq \lambda \cdot \rangle g_v^{\perp}, g_v^{\perp} \rangle_{\Pi_0^{(v)}}. \tag{8}$$

- Since $\alpha_v = (P_{G(X_v)} g)(v)$, we have

$$\mathop{\mathbb{E}}_{v \sim \Pi_0} [\alpha_v^2] = \langle P_{G(X)} g, P_{G(X)} g \rangle_{\Pi_0} = \eta^2. \tag{9}$$

Continuing where we left off at (7), we have

$$\begin{aligned}
\eta &= \mathop{\mathbb{E}}_{v\sim\Pi_0}\left[\alpha_v^2 + \langle \mathrm{P}_{\mathrm{G}(\mathrm{X}_v)} g_v^\perp, g_v^\perp\rangle_{\Pi_0^{(v)}}\right] \\
&\le \mathop{\mathbb{E}}_{v\sim\Pi_0}\left[\alpha_v^2 + \lambda\langle g_v^\perp, g_v^\perp\rangle_{\Pi_0^{(v)}}\right] && \text{[By (8)]} \\
&= \mathop{\mathbb{E}}_{v\sim\Pi_0}\left[(1-\lambda)\alpha_v^2 + \langle g_v, g_v\rangle_{\Pi_0^{(v)}}\right] && \text{[Since } \langle g_v, g_v\rangle = \alpha_v^2 + \langle g_v^\perp, g_v^\perp\rangle\text{]} \\
&= (1-\lambda)\eta^2 + \lambda. && \text{[By (9)]}
\end{aligned}$$

This implies that

$$\begin{aligned}
&\eta(1-\eta) \le \lambda(1-\eta^2) \\
&\implies \eta \le \lambda(1+\eta) && \text{[Since X is connected, we have } \eta < 1\text{]} \\
&\implies \eta \le \frac{\lambda}{1-\lambda}.
\end{aligned}$$

Thus proved. □

This proof is from the exposition of Harsha and Saptharishi [HS22] on HDX constructions, which is in turn adapted from Yotam Dikstein's lectures notes [Dik19].

3 GLAUBER DYNAMICS AS A HDX RANDOM WALK

We now return to the question of analysing the Glauber Dynamics GD_G on the set $\mathcal{T}_\mathrm{G}$ of spanning trees. To this end, let $\mathcal{F}$ be the $(n-2)$-dimensional simplicial complex consisting of the forests of the graph G. Observe that the set of maximal dimensional faces in $\mathcal{F}$ is precisely the set of spanning trees of G, namely $\mathcal{T}_\mathrm{G}$ and furthermore that the Glauber Dynamics GD_G is the down-up walk on the top-most layer $\mathcal{F}(n-2)$. Thus, to prove Theorem 1.1, it suffices to understand the spectral gap of the down-up walk $\mathrm{P}_{n-2}^{\triangledown}$ of the weighted simplical complex $(\mathcal{F}, \Pi)$ on the $(n-2)$th layer (here Π is the joint distribution induced by the uniform distribution on $\mathcal{F}(n-2) = \mathcal{T}_\mathrm{G}$).

The key insight of Anari, Liu, Oveis-Gharan and Vinzant is the following lemma which shows that $\mathcal{F}$ is a 0-onesided-link-HDX.

3.1 LEMMA ([ALOV19]). *$\mathcal{F}$ is a 0-onesided-link-HDX.*

Let us first see how this lemma implies Theorem 1.1.

Proof of Theorem 1.1. The down-up walk $\mathrm{P}_{n-2}^{\triangledown}$ on $\mathcal{F}(n-2)$ is positive semidefinite. Hence, to bound its spectral gap $\gamma(\mathrm{P}_{n-2}^{\triangledown})$ it suffices to consider $1-\lambda_2(\mathrm{P}_{n-2}^{\triangledown})$. Since $\mathcal{F}$ is a 0-link-HDX, applying Theorem 2.4, we have $\mathrm{P}_i^{\wedge} \preccurlyeq \mathrm{P}_i^{\triangledown}$ for every $0 \le i < n-2$. This implies that $\lambda_2(\mathrm{P}_i^{\wedge}) \le \lambda_2(\mathrm{P}_i^{\triangledown})$ for every i.

$$\begin{aligned}
\gamma(P^{\triangledown}_{n-2}) &= 1 - \lambda_2(P^{\triangledown}_{n-2}) && \text{[Since } P^{\triangledown}_{n-2} \text{ is positive semidefinite]} \\
&= 1 - \lambda_2(P^{\vartriangle}_{n-3}) && \text{[By (2)]} \\
&= 1 - \lambda_2\left(\frac{1}{n-1} I_{\mathcal{F}(n-3)} + \frac{n-2}{n-1} P^{\wedge}_{n-3}\right) && \text{[By (1)]} \\
&= \frac{n-2}{n-1}\left(1 - \lambda_2(P^{\wedge}_{n-3})\right) \\
&\geq \frac{n-2}{n-1}\left(1 - \lambda_2(P^{\triangledown}_{n-3})\right) && \text{[Since } \lambda_2(P^{\wedge}_{n-3}) \leq \lambda_2(P^{\triangledown}_{n-3})\text{]} \\
&\geq \frac{n-2}{n-1} \cdot \frac{n-3}{n-2}\left(1 - \lambda_2(P^{\triangledown}_{n-4})\right) && \text{[Applying the same argument again]} \\
&\vdots \\
&\geq \frac{n-2}{n-1} \cdot \frac{n-3}{n-2} \cdots \frac{1}{2}\left(1 - \lambda_2(P^{\triangledown}_{0})\right) \\
&= \frac{1}{n-1}.
\end{aligned}$$

□

3.1 *Matroidal graphs are 0-onesided-link-HDXs*

In this section, we prove Theorem 3.1 by showing that for every $-1 \leq i \leq n-4$ and $F \in \mathcal{F}(i)$, we have that the 1-skeleton of the link of F is a 0-onesided-link-HDX. By Oppenheim's trickle down theorem, it suffices to show this for $i = n-4$. Let F be a forest in $\mathcal{F}(n-4)$ and $(\mathcal{F}_F, \Pi^{(F)})$ be the corresponding link.

This is the only place in the proof where we use the fact that the underlying state space is the set of spanning trees. The proof given here works verbatim if the set of spanning trees is replaced with the set of bases of a matroid.

Let us understand the 1-skeleton $G(F) = (\mathcal{F}_F(0), \mathcal{F}_F(1), \Pi_1^{(F)})$ of the link of F. Observe that $\Pi_1^{(F)}$ is the uniform distribution on the edges of G(F). F is a forest with $n-3$ edges. Let the 3 components of the forest be $V_1, V_2, V_3 \subseteq V$. The vertices of G(F) are precisely the edges in the graph G across these 3 components. In other words $\mathcal{F}_F(0) = E(V_1, V_2) \cup E(V_2, V_3) \cup E(V_3, V_1)$. What are the edges $\mathcal{F}_F(1)$ of G(F)? Two vertices in G(F) (equivalently two edges of $E(V_1, V_2) \cup E(V_2, V_3) \cup E(V_3, V_1)$) are connected iff they together with F combine to form a spanning tree of G. It immediately follows that the graph G(F) is the complete 3-partite graph. The following theorem shows that the second eigenvalue of any complete k-partite graph with the uniform distribution on the edges is at most 0.

3.2 THEOREM (eigenvalues of complete k-partite graph). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \pi_1)$ be a complete k-partite graph with parts $\mathcal{V} = \bigcup_{i=1}^{k} \mathcal{V}_i$ and $\mathcal{V}_i \cap \mathcal{V}_j = \emptyset$ if $i \neq j$ and π_1 the uniform distribution on the edges $\mathcal{E}$. Then, $\lambda_2(\mathcal{G}) \leq 0$.*

Proof. Let $n_i = |\mathcal{V}_i|$ be the size of the k parts and $n = \sum_{i=1}^{k} n_i$. The degree of any vertex in part $\mathcal{V}_i$ is $n - n_i$. Since π_1 is the uniform distribution on the edges, the induced distribution π_0 on the vertices is proportional to the degree of vertices. Hence, the vertex distribution π_0 is given as follows: If $v \in \mathcal{V}_i$, then

$$\pi_0(v) = \frac{n - n_i}{\sum\limits_j n_j(n - n_j)} = \frac{n - n_1}{n^2 - \sum\limits_j n_j^2}.$$

Let $f : \mathcal{V} \to \mathbb{C}$ be any vector orthogonal to the all one's vector $\mathbb{1}_{\mathcal{V}}$. In other words, $\langle f, \mathbb{1}_{\mathcal{V}} \rangle_{\pi_0} = 0$ or equivalently, $\sum\limits_i (n - n_i) \sum\limits_{v \in \mathcal{V}_i} f(v) = 0$. Let $F_i = \sum\limits_{v \in \mathcal{V}_i} f(v)$. Hence, we have

$$\sum_i (n - n_i) F_i = 0. \tag{10}$$

Now, let us consider the inner product $\langle P_{\mathcal{G}} f, f \rangle_{\pi_0}$.

$$\begin{aligned}
\langle P_{\mathcal{G}} f, f \rangle_{\pi_0} &= \mathop{\mathbb{E}}_{\{u,v\} \sim \pi_1} [f(u) f(v)] \\
&= \frac{1}{\sum\limits_{1 \le i < j \le k} n_i n_j} \cdot \sum_{1 \le i < j \le k} \left(\sum_{v \in \mathcal{V}_i} f(v) \right) \left(\sum_{v \in \mathcal{V}_j} f(v) \right) \\
&= \frac{1}{\sum\limits_{1 \le i < j \le k} n_i n_j} \cdot \sum_{1 \le i < j \le k} F_i \cdot F_j \\
&= \frac{1}{2\left(\sum\limits_{1 \le i < j \le k} n_i n_j \right)} \cdot \left[\left(\sum_i F_i \right)^2 - \sum_i F_i^2 \right] \\
&= \frac{1}{2\left(\sum\limits_{1 \le i < j \le k} n_i n_j \right)} \cdot \left[\left(\frac{\sum_i n_i F_i}{n} \right)^2 - \sum_i F_i^2 \right] && \text{[By (10)]} \\
&\le \frac{1}{2\left(\sum\limits_{1 \le i < j \le k} n_i n_j \right)} \cdot \left[\left(\frac{\sum_i n_i F_i^2}{n} \right) - \sum_i F_i^2 \right] && \text{[By Jensen's inequality]} \\
&= \frac{1}{2\left(\sum\limits_{1 \le i < j \le k} n_i n_j \right)} \left[- \frac{\sum_i (n - n_i) F_i^2}{n} \right] \\
&\le 0.
\end{aligned}$$

Hence, $\lambda_2(\mathcal{G}) \le 0$. $\square$

References

[Ald90] David J. Aldous. *The random walk construction of uniform spanning trees and uniform labelled trees*. SIAM J. Discrete Math., 3(4):450–465, 1990.

[ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis-Gharan, and Cynthia Vinzant. *Log-concave polynomials II: high-dimensional walks and an FPRAS for counting bases of a matroid*. In *Proc. 51st ACM Symp. on Theory of Computing (STOC)*, pages 1–12. 2019. arXiv:1811.01816.

[Bro89] Andrei Z. Broder. *Generating random spanning trees*. In *Proc. 30th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 442–447. 1989.

[CGM19] Mary Cryan, Heng Guo, and Giorgos Mousa. *Modified log-Sobolev inequalities for strongly log-concave distributions*. In *Proc. 60th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 1358–1370. 2019. arXiv:1903.06081.

[DDFH18] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha. *Boolean function analysis on high-dimensional expanders*. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, eds., *Proc. 22nd International Workshop on Randomization and Computation (RANDOM)*, volume 116 of *LIPIcs*, pages 38:1–38:20. Schloss Dagstuhl, 2018. arXiv:1804.08155, eccc:2018/TR18-075.

[Dik19] Yotam Dikstein. *Oppenheim's Trickling Down Theorem*, 2019. Lecture notes from the Error-Correcting Codes and High-Dimensional Expansion Boot Camp at Simons Institute for the Theory of Computing.

[HS22] Prahladh Harsha and Ramprasad Saptharishi. *A note on the elementary construction of high-dimensional expanders of Kaufman and Oppenheim*. Theory of Computing, Graduate Surveys, 2022. (To appear). arXiv:1912.11225.

[KO20] Tali Kaufman and Izhar Oppenheim. *High order random walks: Beyond spectral gap*. Combinatorica, 40(1):245–281, 2020. (Preliminary version in *20th RANDOM*, 2018). arXiv:1707.02799.

[Opp18] Izhar Oppenheim. *Local spectral expansion approach to high dimensional expanders part I: Descent of spectral gaps*. Discrete Comput. Geom., 59(2):293–330, 2018. arXiv:1709.04431.