Moden Spectral Graph Theory **Fall 2024**

Lecture 10: Approximate Counting and Mixing Time of Markov Chains

Lecturer: Shayan Oveis Gharan

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

In this lecture we use the machinery we developed to study Glauber dynamics for the Ising model to study the hard-core model. We are given a graph $G = (V, E)$ with maximum degree Δ; we want to sample an independent set I with probability proportional to $\lambda^{|I|}$ for $0 < \lambda$. The main theorem we will prove is the following:

Theorem 11.1. *For any graph $G = (V, E)$ of maximum degree Δ, if $\lambda < (1-\delta)\frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}} =: (1-\delta)\lambda_c(\Delta)$ then π is $O_\delta(1)$ spectrally independent.*

Then, following the machinery we developed in the last lecture we have the Glauber-dynamics mixes in polynomial time.

Throughout this lecture we write

$$\mathcal{I}(u, v) := \Psi_\pi(u, v) := \mathbb{P}[v|u] - \mathbb{P}[v|\overline{u}].$$

In fact we prove the following theorem,

Theorem 11.2. *If $\lambda < (1-\delta)\lambda_c(\Delta)$ then for any vertex r,*

$$\sum_{v \neq r} |\mathcal{I}(r \to v)| \le O_\delta(1).$$

Having this Theorem 11.1 simply follows from the fact that

$$\lambda_{\max}(\Psi_\pi) \le \max_r \sum_v |\Psi_\pi(r, v)| \le O_\delta(1).$$

This in particular shows that π is $O_\delta(1)$ spectrally independent. Now, all pinnings of π also correspond to the hard-core model on graphs of maximum degree $\le \Delta$. So, the above theorem also implies all pinnings of π are $O_\delta(1)$-spectrally independent.

11.1 Self-avoiding Walk Tree

The main fundamental step in proving Theorem 11.2 is to reduce the theorem from arbitrary graphs G (with maximum degree Δ) to *trees* (with maximum degree Δ) in which we want to bound the maximum influence of the root to the rest of the vertices. This builds on Weitz's influential correlation decay technique [Wei06]

We start by defining the self-avoiding walk trees. Given a connected graph $G = (V, E)$ be a connected graph, and a specific vertex $r \in V$, and a total ordering of the vertex set V, the self-avoiding walk (SAW) tree rooted at r, $T_{SAW}(G, r)$ is defined as follows: It is a tree rooted at r of all paths starting at r in G except that whenever a path closes a cycle, say $r = v_0, v_1, \ldots, v_k, v_i$ where $0 \le i \le k-1$, the copy (in the tree) of of

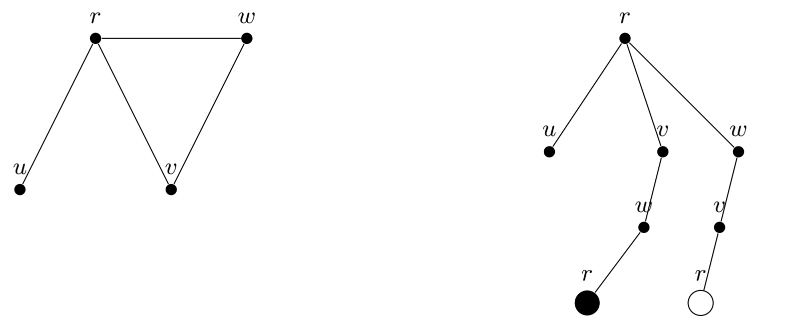

v_i (in G) is fixed to be occupied if $v_{i+1} < v_k$ in the total order and un-occupied otherwise. See the following picture for an example. So, observe that there are multiple copies of every vertex of G in the tree. For each $v \in V$ we denote the set of all unfixed copies of v in $T_{SAW}(G, r)$ by $\mathcal{C}_v$.

For the sake of the proof we assume that every vertex v has a distinct activity parameter, λ_v. In that case, all copies of v from $\mathcal{C}_v$ will have the *same* activity parameter λ_v in the SAW tree. As alluded to above, we will show that for any vertex $v \neq r$, $\mathcal{I}_G(r \to v) = \sum_{\hat{v} \in \mathcal{C}_v} \mathcal{I}_T(r \to \hat{v})$.

To establish that, the idea is to look the generating polynomial of the hardcore model as a multivariate polynomial in terms of vertex activities $\{\lambda_v\}_{v \in V}$ and relate the generating polynomial of G to the generating polynomial of T. Let $\lambda = \{\lambda_v\}_{v \in V}$ denote the vector of vertex activities. We define the partition function,

$$g_G(\lambda) = \sum_{I \text{ independent set}} \prod_{v \in I} \lambda_v$$

For a vertex v we write $g_{G,v}$ to denote the generating polynomial of the hard-core model when v is in, i.e., $\lambda_v \partial_{\lambda_v} g_G$ and similarly we let $g_{G,\bar{v}} = g_G(\lambda_v = 0)$.

Theorem 11.3. *Let $G = (V, E)$ be a connected graph, $r \in V$ be a vertex such that G is connected. Let $T = T_{SAW}(G, r)$ be the self-avoiding walk tree of G rooted at r. Then, $g_{G,r}(\lambda)$ divides $g_{T,r}(\lambda)$. More precisely, there exists a polynomial $p_{G(r)} = p_{G(r)}(\lambda_{-r})$ (that is independent of λ_r such that*

$$g_{T,r} = g_{G,r} \cdot p_{G(r)}, \qquad g_{T,\bar{r}} = g_{G,\bar{r}} \cdot p_{G(r)}.$$

For a vertex u we write $g_{G,u}$ to denote the generating polynomial of all independent sets that contain u and similarly we write $g_{G,\overline{u}}$ to denote the polynomial that u is out. First, we use the above theorem to prove the following lemma.

Lemma 11.4. *For any vertex $v \neq r$, $\mathcal{I}_G(r \to v) = \sum_{\hat{v} \in \mathcal{C}_v} \mathcal{I}_T(r \to \hat{v})$*

Proof. The main observation is that if $g(z_1, \ldots, z_n)$ is the generating polynomial of a probability distribution π over n items, then for any i, the marginal of i is exactly equal to $z_i \partial_{z_i} \log g$.

Having this we can write,

$$\begin{aligned} \lambda_v \partial_{\lambda_v} \log \frac{g_{G,r}(\lambda)}{g_{G,\bar{r}}(\lambda)} &= \lambda_v \partial_{\lambda_v} (\log g_{G,r}(\lambda) - \log g_{G,\bar{r}}) \\ &= \lambda_v \partial_{\lambda_v} \log g_{G,r}(\lambda) - \lambda_v \partial_{\lambda_v} \log g_{G,\bar{r}} \\ &= \mathbb{P}[v|r] - \mathbb{P}[v|\overline{r}] = \mathcal{I}_G(r \to v). \end{aligned} \tag{11.1}$$

In other words, the above calculations follows by a simple fact that if $g(z_1, \ldots, z_n)$ is a generating polynomial of a probability distribution over n items, then for any i, the marginal of i is exactly equal to $z_i \partial_{z_i} \log g$.

On the other hand, recall that for the SAW tree T , every free copy $\hat{v}$ of v has the same activity $\lambda_{\hat{v}} = \lambda_v$. So, by the above theorem,

$$\begin{aligned}
\lambda_v \partial_{\lambda_v} \log \frac{g_{G,r}(\lambda)}{g_{G,\bar{r}}(\lambda)} &= \lambda_v \partial_{\lambda_v} \log \frac{g_{T,r}(\lambda)}{g_{T,\bar{r}}(\lambda)} && \text{(Theorem 11.3)} \\
&= \sum_{\hat{v} \in \mathcal{C}_v} \lambda_{\hat{v}} \partial_{\lambda_{\hat{v}}} \log \frac{g_{T,r}(\lambda)}{g_{T,\bar{r}}(\lambda)} \cdot \frac{\partial \lambda_{\hat{v}}}{\partial \lambda_v} \lambda_{\hat{v}}(\lambda_v) && \text{(Chain Rule)} \\
&= \sum_{\hat{v} \in \mathcal{C}_v} \mathcal{I}_T(r \to \hat{v})
\end{aligned}$$

This completes the proof of the lemma. □

11.2 Reduction to Self Avoiding Walk Tree

In this section we prove Theorem 11.3. The proof is an inductive argument in which we condition on additional vertices of the graph G to be in/out. Therefore, we will need a stronger inductive hypothesis. For $\Lambda \subseteq V$ and a partial configuration $\sigma_\Lambda \in \{0,1\}^\Lambda$, we define the SAW tree with conditioning σ_Λ by assigning the configuring σ_v to every copy $\hat{v}$ of v from $\mathcal{C}_v$ and removing all descendants of $\hat{v}$ (from the tree), for each $v \in \Lambda$. Recall that in general, different copies of v from $\mathcal{C}_v$ can receive different in/out assignments. We define the generating polynomial $g^{\sigma_\Lambda}(.)$ to denote the generating polynomial of all independent sets consistent with the status of the set Λ of vertices.

We inductively prove that, there is a polynomial $p^{\sigma_\Lambda}_{G(r)}(\lambda)$ (that is independent of λ_r such that

$$g^{\sigma_\Lambda}_{T,r} = g^{\sigma_\Lambda}_{G,r} \cdot p^{\sigma_\Lambda}_{G(r)} \qquad \text{and} \qquad g^{\sigma_\Lambda}_{T,\bar{r}} = g^{\sigma_\Lambda}_{G,\bar{r}} \cdot p^{\sigma_\Lambda}_{G(r)}$$

We induct on the number of edges with (at least) one endpoint in the set $V \setminus \Lambda$.

Suppose that the root r has d neighbors $v_1, \dots, v_d$ in G. Define G' to be the graph obtained by replacing the vertex r with d vertices $r_1, \dots, r_d$ and then connecting $\{r_i, d_i\}$ for $1 \le i \le d$. For simplicity, we assume that $(G \setminus \{r\}) \setminus \Lambda$ is still connected. For each i, let $G_i = G' - r_i$. Consider the hardcore model on $G_i^{\sigma_\Lambda}$ together with an additional conditioning that the vertices $r_1, \dots, r_{i-1}$ are fixed to be **out** while $r_{i+1}, \dots, r_d$ are fixed to be **in**; we denote this conditioning by σ_{U_i} with $U_i := \{v_1, \dots, v_d\} \setminus \{v_i\}$. Then, $T = T_{SAW}(G, r)$ can be generated by the following recursive procedure.

Step 1) For each i, let $T_i = T_{SAW}(G_i, v_i)$ plus the conditioning σ_{U_i};

Step 2) Let $T = T_{SAW}(G, r)$ be the union of r and $T_1, \dots, T_d$ by connecting $\{r, v_i\}$ for $1 \le i \le d$; output T.

Observe that this algorithm exactly corresponds to the definition of the self-avoiding walk tree we gave in the previous section.

For the purpose of the proof we set $\lambda_{r_i} = 1$ for all $1 \le i \le d$ instead of λ_r (this is basically how we will avoid λ_r in as a parameter of $p^{\sigma_\Lambda}_{G,r}$). Observe that by definition

$$g^{\sigma_\Lambda}_{G,r} = \lambda_r g^{\sigma_\Lambda}_{G',r_1,\dots,r_d} \qquad g^{\sigma_\Lambda}_{G,\bar{r}} = g^{\sigma_\Lambda}_{G',\bar{r}_1,\dots,\bar{r}_d} \tag{11.2}$$

The main observation is that the graph G_i has one edge less than G, so by induction hypothesis, its generating polynomial divides the generating polynomial of a tree. Define $\Lambda_i = \Lambda \cup U_i$. We write

$$\begin{aligned} g^{\sigma_\Lambda}_{T,r} &= \lambda_r \prod_{i=1}^{d} g^{\sigma_{\Lambda_i}}_{T_i,\overline{v}_i} && \text{(recursion of a tree)} \\ &= \lambda_r \prod_{i=1}^{d} g^{\sigma_{\Lambda_i}}_{G_i,\overline{v}_i} \cdot p^{\sigma_{\Lambda_i}}_{G_i(v_i)} && \text{(Induction Hypothesis)} \\ &= \lambda_r \prod_{i=1}^{d} g^{\sigma_\Lambda}_{G',\overline{r}_1,\ldots,\overline{r}_{i-1},r_i,\ldots,r_d} \cdot \prod_{i=1}^{d} p^{\sigma_{\Lambda_i}}_{G_i(v_i)} \\ &= g^{\sigma_\Lambda}_{G,r} \prod_{i=2}^{d} g^{\sigma_\Lambda}_{G',\overline{r}_1,\ldots,\overline{r}_{i-1},r_i,\ldots,r_d} \cdot \prod_{i=1}^{d} p^{\sigma_{\Lambda_i}}_{G_i(v_i)} && \text{(by (11.2))} \end{aligned}$$

Similarly, we can write

$$\begin{aligned} g^{\sigma_\Lambda}_{T,\overline{r}} &= \prod_{i=1}^{d}(g^{\sigma_{\Lambda_i}}_{T_i,v_i} + g^{\sigma_{\Lambda_i}}_{T_i,\overline{v}_i}) = \prod_{i=1}^{d}(g^{\sigma_{\Lambda_i}}_{G_i,v_i} \cdot p^{\sigma_{\Lambda_i}}_{G_i(v_i)} + g^{\sigma_{\Lambda_i}}_{G_i,\overline{v}_i} \cdot p^{\sigma_{\Lambda_i}}_{G_i(v_i)}) \\ &= \prod_{i=1}^{d} g^{\sigma_\Lambda}_{G',\overline{r}_1,\ldots,\overline{r}_i,r_{i+1},\ldots,r_d} \cdot \prod_{i=1}^{d} p^{\sigma_{\Lambda_i}}_{G_i(v_i)} = g^{\sigma_\Lambda}_{G,\overline{r}} \prod_{i=1}^{d-1} g^{\sigma_\Lambda}_{G',\overline{r}_1,\ldots,\overline{r}_i,r_{i+1},\ldots,r_d} \cdot \prod_{i=1}^{d} p^{\sigma_{\Lambda_i}}_{G_i(v_i)} \end{aligned}$$

The inductive step simply follows by letting $g^{\sigma_\Lambda}_{G(r)} = \prod_{i=2}^{d} g^{\sigma_\Lambda}_{G',\overline{r}_1,\ldots,\overline{r}_i,r_{i+1},\ldots,r_d} \cdot \prod_{i=1}^{d} p^{\sigma_{\Lambda_i}}_{G_i(v_i)}$

This completes the proof of Theorem 11.3. Using Theorem 11.3 and Lemma 11.4 to prove Theorem 11.2 it is enough to prove the following theorem:

Theorem 11.5. *For any Δ-ary tree T rooted at a vertex r and any $\lambda \le (1-\delta)\lambda_c(\Delta)$, we have*

$$\sum_v \mathcal{I}(r \to v) \le O_\delta(1).$$

11.3 Bounding Influences on a Tree

Given a tree T (where every vertex has at most $\Delta - 1$ many children (note that root can really have Δ children but we ignore that for simplicity let $L_r(k)$ be the number of vertices at distance k of the root. [CLV20] proved that if the activity parameter $\lambda \le (1-\delta)\lambda_c(\Delta)$, then we have the following bound: For any $k \ge 1$,

$$\sum_{v \in L_r(k)} \mathcal{I}(r \to v) \le 4(1-\delta/2)^{k-1}$$

Summing this up for $k = 1 \to \infty$, even if T has infinitely many vertices, we get

$$\sum_v \mathcal{I}(r \to v) \le 8/\delta.$$

Next, we will explain the main ideas to prove the above bound. First, for a vertex $v \in T$, let T_v be the sub-tree of T rooted at v; thus $T_r = T$. Let $R_v := \frac{g_{T_v,v}(\lambda)}{g_{T_v,\overline{v}}(\lambda)} = \frac{\mathbb{P}[v \text{ in}]}{\mathbb{P}[v \text{ out}]}$. Say a vertex u has d children $v_1, \ldots, v_d$ in the tree; the tree recursion is a formula that computes R_u given $R_{v_1}, \ldots, R_{v_d}$ due to the independence of T_{v_i}'s. More specifically, there is a function $F_d : [0,\infty]^d \to [0,\infty]$ such that

$$R_u = F_d(R_{v_1}, \ldots, R_{v_d}) := \lambda \prod_{i=1}^{d} \frac{1}{R_{v_i} + 1}.$$

We leave it as an exercise to verify the above formula.

Recall that by equation (11.1), the influence of r to a vertex u is the derivative of $\log R_r$ with respect to the external field at u. So, it is natural to define an analogue of the F_d function for the log ratio quantity. More specifically, let $H_d := [-\infty, +\infty]^d \to [-\infty, +\infty]$ defined as follows:

$$\log R_u = H_d(\log R_{v_1}, \ldots, \log R_{v_d}) := \log \lambda + \sum_{i=1}^{d} \log \frac{1}{1 + e^{\log R_{v_i}}}$$

To put it differently, $H_d = \log \circ F_d \circ \exp$.

The following lemma follows from the fact that we are analyzing influences in a tree.

Lemma 11.6. *Suppose that $u, v, w \in T$ are three distinct vertices such that v is on the unique path from u to w. Then*

$$\mathcal{I}(u \to w) = \mathcal{I}(u \to v) \cdot \mathcal{I}(v \to w)$$

Proof.

$$\begin{aligned}
\mathcal{I}[u \to w] &= \mathbb{P}[w|u] - \mathbb{P}[w|\bar{u}] \\
&= \mathbb{P}[v|u]\,\mathbb{P}[w|u,v] + \mathbb{P}[\bar{v}|u]\,\mathbb{P}[w|u\bar{v}] - \mathbb{P}[v|\bar{u}]\,\mathbb{P}[w|\bar{u},v] - \mathbb{P}[\bar{v}|\bar{u}]\,\mathbb{P}[w|\bar{v},\bar{u}] \\
&= \mathbb{P}[v|u]\,\mathbb{P}[w|v] + \mathbb{P}[\bar{v}|u]\,\mathbb{P}[w|\bar{v}] - \mathbb{P}[v|\bar{u}]\,\mathbb{P}[w|v] - \mathbb{P}[\bar{v}|\bar{u}]\,\mathbb{P}[w|\bar{v}]
\end{aligned}$$

The last line uses that T is a tree and v is on the path from u to w. On the other hand,

$$\begin{aligned}
\mathcal{I}[u \to v]\mathcal{I}[v \to w] &= (\mathbb{P}[v|u] - \mathbb{P}[v|\bar{u}]) \cdot (\mathbb{P}[w|v] - \mathbb{P}[w|\bar{v}]) \\
&= \mathbb{P}[w|v] \cdot (\mathbb{P}[v|u] - \mathbb{P}[v|\bar{u}]) - \mathbb{P}[w|\bar{v}] \cdot (\mathbb{P}[v|u] - \mathbb{P}[v|\bar{u}]) \\
&= \mathbb{P}[w|v] \cdot (\mathbb{P}[v|u] - \mathbb{P}[v|\bar{u}]) - \mathbb{P}[w|\bar{v}] \cdot (\mathbb{P}[\bar{v}|\bar{u}] - \mathbb{P}[\bar{v}|u]) = \mathcal{I}[u \to w]
\end{aligned}$$

as desired. □

For the second lemma we need another notation: For $y \in [-\infty, \infty]$ define

$$h(y) := -\frac{e^y}{1 + e^y} = \frac{\partial}{\partial_y} H_d(y_1, \ldots, y_{i-1}, y, y_{i+1}, \ldots, y_d). \tag{11.3}$$

It follows by (11.1) that

Lemma 11.7. *For any vertex $v \in T$ and any child u of v we have*

$$\mathcal{I}(v \to u) = h(\log R_u).$$

Proof.

$$\begin{aligned}
\partial_y H_d &= \partial_y \left(\log \lambda + \log \frac{1}{1 + e^y} + \sum_{1 \le j \le d, j \ne i} \frac{1}{1 + e^{y_j}} \right) \\
&= \partial_y \log \frac{1}{1 + e^y} \\
&= -\frac{\partial_y (1 + e^y)}{1 + e^y} = -\frac{e^y}{1 + e^y} = -\frac{R_u}{1 + R_u}
\end{aligned}$$

□

Having the above two lemmas we can simply write the influence of r to vertices in $L_k(r)$ inductively. Now, the main issue is that the straightforward recursion gives us terms of the form $\prod_{i=0}^{k-1} h(\log R_{u_i})$ for any path $r = u_0, \ldots, u_{k-1}, u_k$. And, in principal we can have as many as $(\Delta - 1)^k$ many such paths. A direct upper bound on such a product does not give a tight bound on the influence (that is independent of Δ) as we have to multiply the upper-bound by $(\Delta - 1)^k$.

The trick is to use a method called the potential method: Instead of tracking log ratios in the tree recursion we apply a potential function Ψ and study how $\Psi(\log R_u)$ evolves in the tree. We also let $\psi := \Psi'$ be the derivative of the potential. More precisely define

$$H_d^{\Psi} := \Psi \circ H_d \circ \Psi^{-1}.$$

We prove inductively that for any vertex $u \in T$,

$$\sum_{v \in L_u(k)} \psi(\log R_u)|\mathcal{I}(u \to v)| \leq \max_{v \in L_u(k)} \{\psi(\log R_v)\} \cdot (1-\alpha)^k$$

where $L_u(k)$ is the set of vertices at distance k of u and α is a parameter that we choose later. The base case can be checked easily. Now, suppose the claim is checked for $k-1$. Say u has d children $w_1, \ldots, w_d$. We write,

$$\begin{aligned}
\sum_{v \in L_u(k)} \psi(\log R_u)|\mathcal{I}(u \to v)| &= \sum_{i=1}^{d} \psi(\log R_u)|\mathcal{I}(u \to w_i)| \sum_{v \in L_{w_i}(k-1)} |\mathcal{I}(w_i \to v)| && \text{(Lemma 11.6)}\\
&= \sum_{i=1}^{d} \frac{\psi(\log R_u)}{\psi(\log R_{w_i})}|h(\log R_{w_i})| \sum_{v \in L_{w_i}(k-1)} \psi(\log R_{w_i})\mathcal{I}(w_i \to v)| && \text{(Lemma 11.7)}\\
&\leq \sum_{i=1}^{d} \frac{\psi(\log R_u)}{\psi(\log R_{w_i})}|h(\log R_{w_i})| \max_{v \in L_{w_i}(k-1)} \psi(\log R_v) \cdot (1-\alpha)^{k-1} && \text{(IH)}\\
&\leq \max_{v \in L_u(k)} \psi(\log R_v)(1-\alpha)^{k-1} \cdot \sum_{i=1}^{d} \frac{\psi(\log R_u)}{\psi(\log R_{w_i})}|h(\log R_{w_i})|
\end{aligned}$$

Finally, the last observation is that the quantity in the sum is exactly $\left\|\nabla H_d^{\Psi}(\Psi(\log R_{w_1}), \ldots, \Psi(\log R_{w_d}))\right\|_1$. So, the main property of the potential function is that for any $y_1, \ldots, y_d$ in the range of Ψ we have

$$\left\|\nabla H_d^{\Psi}(y_1, \ldots, y_d)\right\|_1 \leq 1 - \alpha.$$

It turns out that this can be achieved for $\psi(y) = \sqrt{|h(y)|}$ and Ψ defined accordingly and for $\alpha \geq \delta/2$. In particular, we can write

$$\begin{aligned}
\sum_{i=1}^{d} \frac{\psi(\log R_u)}{\psi(\log R_{w_i})}|h(\log R_{w_i})| &= \sum_{i=1}^{d} \frac{\sqrt{|h(\log R_u)|}}{\sqrt{|h(\log R_{w_i})|}}|h(\log R_{w_i})|\\
&= \sum_{i=1}^{d} \sqrt{\frac{\lambda \prod_{j=1}^{d} \frac{1}{1+R_{w_j}}}{1 + \lambda \prod_{j=1}^{d} \frac{1}{1+R_{w_j}}}} \sqrt{\frac{R_{w_i}}{1+R_{w_i}}}
\end{aligned}$$

We leave it as an exercise to bound the RHS by $1 - \delta/2$ assuming $\lambda \leq (1-\delta)\lambda_c(\Delta)$. Note that $d \leq \Delta - 1$.

References

[CLV20] Z. Chen, K. Liu, and E. Vigoda. “Rapid Mixing of Glauber Dynamics up to Uniqueness via Contraction”. In: *FOCS*. 2020, pp. 1307–1318 (cit. on p. 11-4).

[Wei06] D. Weitz. "Counting independent sets up to the tree threshold". In: *STOC*. ACM, 2006, pp. 140–149 (cit. on p. 11-1).