Course notes for 6.S981: Approximate Counting and Sampling

Fall 2023

These are lecture notes I (Noah Golowich) have taken for the course "6.S891: Algorithm Counting and Sampling", taught by Kuikui Liu at MIT in the Fall of 2023. Please note that they are very rough and have not been subjected to any sort of scrutiny or editing (and thus likely to contain errors). Any errors are my own.

1 September 7, 2023

Suppose we have a huge space Ω (e.g., $\mathbb{R}^n, \{-1,1\}^n$), as well as a function $\omega : \Omega \to \mathbb{R}_{\geq 0}$. We are interested with two problems:

1. Counting: compute $Z := \sum_{x \in \Omega} \omega(x)$.
2. Sammpling: sample random $x \in \Omega$ according to $\mu(x) \sim \omega(x)$, i.e., $\mu(x) = \omega(x)/Z$.

The above has many applications, including detecting gerrymandering (detect if a map is anomalous), privacy, statistical inference (namely, if we have samples from a generative model $\mathbb{P}(X|\theta)$, we want to sample $\mathbb{P}(\theta|X) \propto \mathbb{P}(X|\theta) \cdot \mathbb{P}(\theta)$, and we know $\mathbb{P}(X|\theta)$ and the prior $\mathbb{P}(\theta)$), statistical mechanics.

1.1 Complexity Theory

For counting problems, we have $\sharp\mathsf{P}$. Given an NP problem L, we map it to the $\sharp\mathsf{P}$ problem $\sharp L$, which is the problem of counting the number of solutions (i.e., which make the verifier accept). Given SAT, we would map it to $\sharp SAT$, which asks for the number of satisfying solutions. A problem L is $\sharp\mathsf{P}$-hard if any problem in $\sharp\mathsf{P}$ reduces to $\sharp L$. A problem is $\sharp\mathsf{P}$-complete if it is $\sharp\mathsf{P}$-hard and it is in $\sharp\mathsf{P}$.

It turns out that $\sharp SAT$ is $\sharp\mathsf{P}$-complete (same Cook-Levin reduction, trace it through, and you get that counting the number of solutions is $\sharp P$-hard). Surprisingly, not every $\sharp\mathsf{P}$-complete problem arises from a NP-hard problem.

Theorem 1.1 (Valiant)**.** *The problem of counting the number of bipartite perfect matchings is $\sharp\mathsf{P}$-complete.*

Edmonds (1970s) showed that the decision analogue, namely determining if a graph has a perfect matching, is in P. But the counting analogue is hard.

As a corollary, we have:

Corollary 1.2. *Computing the permanent of a matrix is $\sharp\mathsf{P}$-hard.*

Recall that the permanent is defined by $\mathrm{per}(A) = \sum_{\sigma \in S_n} \prod_{i \in [n]} A_{i,\sigma(i)}$. (Recall that $\det(A)$ can be computed in polynomial time.)

There are various counting problems which we can solve efficientlt:

Theorem 1.3 (Matrix tree theorem)**.** *Counting the number of spanning trees of a graph is efficient. In particular, the number of spanning trees is* $\det(L)$*, where* L *is the Laplacian.*

The course will focus on approximate versions of counting.

1.2 Approximate counting/sampling

Definition 1.1 (FPRAS)**.** A *fully poly-time randomized approximation scheme (FPRAS)*, on error $\epsilon \in (0,1)$, failure probability $\delta \in (0,1)$, outputs $\hat{Z}$ such that $\Pr((1-\epsilon)Z \le \hat{Z} \le (1+\epsilon)Z) \ge 1-\delta$, in time $\mathrm{poly}(n, 1/\epsilon, \log 1/\delta)$. Here n is a parameter which quantifies the input size.

As for sampling, we measure closeness by total variation distance: $\|\mu - \nu\|_{\mathsf{TV}} = \frac{1}{2}\sum_{x \in \Omega} |\mu(x) - \nu(x)|$.

Definition 1.2. An efficient sampling algorithm on input δ outputs a random variable $X \in \Omega$ such that $\|\mathrm{Law}(X) - \mu\|_{\mathsf{TV}} \le \delta$ in time $\mathrm{poly}(n, \log 1/\delta)$, where again n is the input size.

For a random variable X, we write $\mathrm{Law}(X)$ to denote the distribution of X.

Total variation distance is nice since

$$\|\mu - \nu\|_{\mathsf{TV}} = \sup_{f: X \to [0,1]} |\mathbb{E}_\nu[f] - \mathbb{E}_\mu[f]|.$$

1.3 Applications of sampling

Sampling seems more powerful than counting since, a priori, it seems to give more. Using Monte Carlo method, you can approximate several statistics of the distribution μ.

In particular, given $f : \Omega \to [0,1]$, if we want to compute $\mathbb{E}_\mu[f]$, the algorithm, given a sampler, draws samples $X_1, \ldots, X_t$ i.i.d. and outputs the empirical mean $\frac{1}{t}\sum_{i=1}^t f(X_i)$. We use Hoeffding bound:

Theorem 1.4 (Hoeffding/Chernoff/Bernstein)**.**

$$\Pr\left(\left|\frac{1}{T}\sum_{i=1}^{T} f(X_i) - \mathbb{E}_\mu[f]\right| > \epsilon \cdot \mathbb{E}_\mu[f]\right) \le 2\exp\left(-\frac{\epsilon^2 T \cdot \mathbb{E}_\mu[f]}{3}\right).$$

In particular, if we take $T \approx \frac{1}{\epsilon^2 \cdot \mathbb{E}_\mu(f)} \cdot \log 1/\delta$, then we can approximate $\mathbb{E}_\mu[f]$ up to an ϵ-fraction. Note that we need $\mathbb{E}_\mu[f]$ to not be too small. If f is unbounded, but $\mathrm{Var}_\mu(f) < \infty$, thenwe should use the "median of means" estimator. Namely, partition the sample into a bunch of subsets, and use the median of means of the subsets. If we only have an approximate sampler (say have TV error of ϵ'), then we incur an additional error of ϵ'.

1.4 Equivalence of counting and sampling

Theorem 1.5 (Jerrum-Valiant-Vazirani)**.** *For "self-reducible" problems, there exists an FPRAS if and only if there exists and efficient sampler.*

Intuitively, *self-reducibility* means that we can decompose a counting/sampling problem into many sub-problems of the same type.

We prove the above theorem for a special case, namely the problem of counting/sampling from matchings. Let $G = (V, E)$ be a graph, and let Ω_G be the set of matchings $M \subset E$ of G. Recall that a matching is a subset of edges so that no vertex is incident to more than one vertex of G. We take the weight function to be uniform over matchings.

In what sense is this problem self-reducible? For any $e \in E$, we can partition matchings based on the membership of e:

1. Matchings containing e are in 1-1 correspondence with matchings of $G\backslash\{u, v\}$. (Just add e to a matching of $G\backslash\{u, v\}$.)
2. Matchings not containing e are in 1-1 correspondence with $G\backslash\{e\}$ (remove e but keep its endpoints).

Why is self-reducibility useful? For any edge e, we have $\Pr_M(e \in M) = \frac{|\Omega_{G-u-v}|}{|\Omega_G|}$. Thus, rearranging, we have $|\Omega_G| = \frac{1}{p(e)} \cdot |\Omega_{G-u-v}|$. We can estimate $p(e) = \mathbb{E}_M[\mathbb{1}e \in M]$ using Monte-Carlo methods. In particular, if we can estimate $|\Omega_{G-u-v}|$, then we can estimate $|\Omega_G|$. Similarly, we have $\Pr_M(e \notin M) = |\Omega_{G-e}|/|\Omega_G|$, and so $|\Omega_G| = 1/(1-p(e)) \cdot |\Omega_{G-e}|$, which is more convenient since typically $p(e)$ is small.

We first prove the following direction of the JVV theorem:

Proof that sampler implies FPRAS. We do this inductively. We order the edges $e_1, \ldots, e_m$ arbitrarily. Define $G_i = G_{i-1} - e_i$, where $G_0 = G$ and G_m is empty, so that $|\Omega_{G_0}| = 1$. We have, by the argument above,

$$|\Omega_G| = \prod_{i=1}^{m} \frac{|\Omega_{G_{i-1}}|}{|\Omega_{G_i}|} = \prod_{i=1}^{m} \frac{1}{\Pr_{M \sim G_{i-1}}(e_i \notin M)}.$$

If we get a $(1 \pm \epsilon/m)$-approximation to all of the probabilities $\Pr_{G_{i-1}}(e_i \notin M)$, then we get a $(1 \pm \epsilon)$-approximation to $|\Omega_G$. If we have an approximate sampler, we can use the monte carlo method to get a $(1 \pm \epsilon/m)$-approximation to each of these probabilities: we need to make sure that these probabilities are not too small. It is easy to see that $\Pr_{M \sim G_{i-1}}(e_i \notin M) \geq 1/2$: for any matching with edge e, there is a matching without edge e obtained by removing e. □

Note that in general, we may not be able to show such a strong bound on the probability, but in general we can estimate both and use whichever is bigger.

Next, we prove the other direction, which roughly speaking runs the above argument in reverse.

Proof that FPRAS implies sampler. Order the edges $e_1, \ldots, e_m$ arbitrarily:

1. We can estimate $\Pr(e_1 \in M) = p_1$.
2. With probability p_1 add e_1 to the matching $\hat{M}$. Remove u_1, v_1, namely the endpoints of e_1.
3. Otherwise, keep the matching $\hat{M}$.

4. Continue with the remaining edges.

□

Let's assume that the estimate of the edge probabilities are all good, up to error $1 \pm \delta/m$. Let M be any fixed matching. We want to analze the TVD between the random matching we generated and the true matching. Define E_i as the event that e_i has the "correct" status (i.e., the same status as the edge e_i in M). Then

$$\Pr(\hat{M} = M) = \prod_{i=1}^{m} \Pr(E_i | E_{1:i-1}) \leq \prod_{i=1}^{m} (1 + \delta/m) \cdot p_i \leq (1 + O(\delta)) \cdot \mu(M),$$

and similarly we have $\Pr(\hat{M} = M) \geq (1 - O(\delta)) \cdot \mu(M)$. Summing over all matchings M, we end up with $\left\| \mathrm{Law}(\hat{M}) - \mu \right\|_{\mathsf{TV}} \leq O(\delta)$.

Note that there's one slight trick here: if we have a $1 \pm \epsilon$-approximation to some probability p, we don't in general get a $1 \pm \epsilon$-approximation to $1 - p$. But this is true if $p \leq 1/2$. Note that, depending on whether each edge is in M, to compute $\Pr(E_i | E_{i-1})$ we may need either $\Pr_{M \sim G_i}(e_i)$ or $1 - \Pr_{M \sim G_i}(e_i)$, where $G_i = G_{i-1} - e_i$ with probability p_i, and $G_i = G_{i-1} - u_i - v_i$ with probability $1 - p_i$.

(In the general, non-matching case, we should estimate both of these probabilities.)

There's one issue with the above proof: the dependence on δ is $\mathrm{poly}(1/\delta)$ since to estimate the edge probabilities we need $O(1/\delta^2)$ samples. To get $\log(1/\delta)$ dependence, we add a rejection sampling step:

Proof for $\log 1/\delta$ *sample complexity of sampling.* Let $\eta = c/m$ for some constant η. We generate $\hat{M}$ as above. If $\hat{M}$ is the output, then accept with probability $c' \cdot \frac{\mu(\hat{M})}{\nu(\hat{M})}$, and otherwise trye again, where $\nu = \mathrm{Law}(\hat{M})$. Note that $\nu = \mathrm{Law}(\hat{M})$ can be computed exactly since it is given by $\prod_{i=1}^{m}$ estimate of p_i, which we computed exactly. Moreover, we can estimate $\mu(\hat{M})$ since we have an FPRAS. Moreover, ignoring the error in this FPRAS, we get that, if we accept, $\Pr(\hat{M}) = \nu(\hat{M}) \cdot (\mu(\hat{M})/\nu(\hat{M})) = \mu(\hat{M})$, so the correct matching probability.

The final algorithm tries the above for $O(\log 1/\delta)$ iterations: if we fail on all of them, then we fail, and we eat this probability-δ of failure. What we need is that we accept on each iteration with constant probability. To see this, as long as the failure probability is η, we have $\|\mu - \nu\|_{\mathsf{TV}} \leq O(\eta)$. Moreover, by inspecting the sampling argument above, we can ensure that, with, say, constant probability over the estimates from sampling, we get a uniform bound on $\|\mu/\nu\|_\infty \leq 2$, which ensures that c' can be taken to be, say, $1/2$, and we will accept with constant probability.

The proof has to condition on the event that all our approximation errors are small enough: the failure probabilities here account for the TV error:

- Chance of rejecting all propsed matchings (roughly $\log 1/\delta$ of those).
- Chance of bad estimates (take care of this by making δ small).

□

Note that the probabilities p_i are random, but we can just let ν be the distribution with whatever probabilities p_i we compute, and under the high-probability event we will have ν is sufficiently close to μ.

2 September 12, 2023

Remember that if μ, ν are probability measures on a set Ω, then we have the definition $\|\mu - \nu\|_{\mathsf{TV}} = \frac{1}{2}\sum_{x\in\Omega} |\mu(x) - \nu(x)|$.

Lemma 2.1. *It holds that* $\|\mu - \nu\|_{\mathsf{TV}} = \sup_{f:\Omega\to[0,1]} |\mathbb{E}_\mu f - \mathbb{E}_\nu f|$.

We omit the proof.

Examples of models we'll discuss in this course.

1. For $A \in \mathbb{R}^{n\times n}, \beta \in \mathbb{R}$, $\Omega = \{\pm 1\}^n$, we have the Ising model
$$\mu(\sigma) \propto \exp\left(\frac{\beta}{2}\sigma^\top A\sigma\right).$$

2. More generally, we have graphical models: $\Omega = [q]^V$, $G = (V, E)$, so that Ω is a space of configurations. For each edge we have a mapping from the configurations of its endpoints to $\mathbb{R}$, namely $\{\psi_e : [q] \times [q] \to \mathbb{R}\}$. Then the Hamiltonian is, for $\sigma \in [q]^V$ (i.e., an assignment of vertices),
$$H(\sigma) = \sum_{e=(u,v)} \psi_e(\sigma(u), \sigma(v)),$$
and we have $\mu(\sigma) \propto \exp(\beta \cdot H(\sigma))$. We can recover the Ising model by taking $\Omega = \{\pm 1\}^n$ and $\psi_{uv} = A_{uv} \cdot \sigma(u) \cdot \sigma(v)$.

   Another example is to take μ uniform over proper colorings. Here we map illegal configurations (i.e., adjacent vertices of the same color) to $-\infty$. Other examples is uniform over independent sets, or uniform over matchings.

   To model matchings, we consider a graph G and consider the line graph $L(G)$, whose vertices are edges of G. To sample matchings of G we sample an independent set of $L(G)$.

3. Spanning trees in graphs (matroids).

2.1 Rejection sampling argument for counting impiles sampling

Remember that the sampling argument was as follows: we order the edges $e_1, \dots, e_m \in E$. We use the approximate counting argument to estimate $q_1 := \Pr_M[e_1 \in M] = |\Omega_{G-u_1-v_1}|/|\Omega_G|$. Then decide to add e_1 to the matching with probability q_1, and delete u_1, v_1; otherwise, we delete e_1.

The naive argument leads to $\operatorname{poly} 1/\delta$ dependence; to get $\operatorname{poly}\log 1/\delta$, we need to do rejection sampling. Here's a cleaner argument for that step: consider a big tree where at each step we branch into the cases where we add e_i or don't add e_i, and end up in the appropriate graph. Note that the probabilities $q_2, \dots$ depend on choices we've made in the previous step (as they affect the current graph), but we omit this dependence.

Imagine that someone has precomputed all probabilities in the tree, and that, along our (random) path in the tree, we just query the probabilities q_i along that path. If everything is deterministic, this defines some well-defined distribution over matchings, and we do rejection sampling over that distribution.

In particular, for all estimates $\vec{q}$ of the probabilities, we get a distribution $\nu_{\vec{q}}$ over matchings. If we have sampled M, then we have a sample from $\nu_{\vec{q}}(M)$ exactly. If instead the probabilities $\vec{q}$ are random, then all of this is random over the probabilities $\vec{q}$, but nothing else changes. The rejection sampling step is based off of the probabilities $v_{\vec{q}}(M)$, though note that to show L_∞ boundedness of the ratio μ/ν, we need to condition on the (high) probability event that the estimates $v_{\vec{q}}(M)$ are all sufficiently accurate, and then use the simple telescoping argument presented last time. (In fact, it seems that we only need this to be ae constant probability event for the argument to work.)

2.2 Markov chains

At a high level, we have some complicated μ we want to sample from: we start with an arbitrary $X \in \Omega$, and (via a MC) we want to add enoug hrandomness in the right ways.

Definition 2.1 (Markov chain)**.** A matrix $P \in \mathbb{R}_{\geq 0}^{\Omega \times \Omega}$ sampling $\sum_{y \in \Omega} P(x \to y) = 1$ is a Markov chain (MC).

A MC generates a stochastic process $(X_t)_{t=0}^{\infty}$ so that

$$\Pr(X_{t+1} = y | X_t = x) = P(x \to y).$$

The MCMC Paradigm is as follows: design some P so that $\text{Law}(P_t) \to \mu$ as $t \to \infty$. Note that $\text{Law}(P_t) = \mu_0 \cdot P^t$, where $\mu_0 = \text{Law}(X_0)$.

Definition 2.2 (Stationary distribution)**.** A distribution μ is stationary wrt P is $\mu P = \mu$.

Examples:

- For $G = (V, E)$, we start out at some $v \in V$, we transform to a uniformly random neighbor of u. In particular, $P_G = D_G^{-1} \cdot A_G$, where D_G is the diagonal matrix of vertex degrees. It is straightforward to see that $\mu(v) \propto \deg(v)$ is the stationary distribution.

  Note that if G is disconnected, we cannot possibly have a unique stationary distribution. (If we start from one component, we will never end up at the other one.) Also note that if G is bipartite, we will not converge to a stationary distribution (we will oscillate between the two sides at each step indefinitely). In some sense, these are the only two obstacles to convergence to a unique stationary distribution (formalized below).

Definition 2.3 (Ergodicity)**.** P is *ergodic* if:

1. It is *irreducible*: for all $x, y \in \Omega$, there exists t so that $P^t(x, y) > 0$ (this is just connectivity in the above example).
2. If is *aperiodic*: for all $x \in \Omega$, $\gcd\{t : P^t(x, x) > 0\} = 1$. (This rules out, e.g., bipartite graphs.)

Note that to ensure aperiodicty, it suffices to ensure that $P(x, x) > 0$ for all x. In particular, given P, we replace it with $(I + P)/2$, which inherits all useful properties of P and is aperiodic. Note also that ergodicity can be more simply defined as: for all $x, y \in \Omega$, there exists N so that for all $t > N$, $P^t(x, y) > 0$. Or more generally, there exists t so that for all $x, y \in \Omega$, $P^t(x \to y) > 0$.

Theorem 2.2 (Fundamental theorem of MCs)**.** *If P is ergodic, then there is a unique stationary distribution μ. Moreover, for all $x \in \Omega$, $\left\|\delta_x \cdot P^t - \mu\right\|_{\mathsf{TV}} \to 0$ as $t \to \infty$.*

Often designing Markov chains with goal staitonary distributions is challenging. But there is one easy condition under which we can do this:

Definition 2.4 (Reversibility)**.** We say that μ is *reversible* wrt P if for all $x, y \in \Omega$,

$$\mu(x) \cdot P(x \to y) = \mu(y) \cdot P(y \to x).$$

Example: simple random walk on undirected graph. This is a useful condition since it has that up to some rescaling, the matrix P is symmetric, so its eigenvalues are real, etc. Note that if μ is reversible wrt P, then it is the stationary distribution.

Here's an easy recipe for constructing Markov chains corresponding to distributions μ. Suppose $\mu(x) \propto w(x)$ on Ω, where we can compute $w(x)$. (E.g., $w(x)$ is the exponential of the Hamiltonian, for graphical models.) We start with a symmetric Markov chain Q on Ω (i.e., $Q(x \to y) = Q(y \to x)$). We then ahve the following *filter step*:

- We accept a proposal x (i.e., $X_{t+1} = x$) with probability $\min\{1, \frac{w(x)}{w(X_t)}\}$.
- Otherwise, stay with $X_{t+1} = X_t$.

In particular, we have defined

$$P(x \to y) = Q(x \to y) \cdot \min\left\{1, \frac{w(y)}{w(x)}\right\},$$

for $y \neq x$, and we stay on x with the remaining probability.

Lemma 2.3. *For all Q, the Metropolis filter wrt w leads to a P that is reversible wrt μ.*

Proof. Calculation. □

Example: Glauber dynamics Consider μ on $\{\pm 1\}^n$. For $\sigma \in \{\pm 1\}^n$ and $i \in [n]$, write $\sigma^{\oplus i}$ to denote σ with the ith bit flipped. Then the Glauber dynamics is defined as follows:

1. Pick $i \sim [n]$ uniformly.
2. Flip i w.p. $\frac{\mu(\sigma^{\oplus i})}{\mu(\sigma)+\mu(\sigma^{\oplus i})}$, and otherwise stay at σ. Note that $\mu(\sigma), \mu(\sigma^{\oplus i})$ above can be replaced by $w(\sigma), w(\sigma^{\oplus i})$ since Z cancels.

It is straightforward to check that the above μ is reversible wrt the Glauber dynamics MC.

Mixing time.

Definition 2.5 (Mixing time)**.** For $\epsilon \in (0, 1)$, define the mixing time $\tau(\epsilon)$ to be the maximum over all $x \in \Omega$ of the minimum t so that $\left\|\delta_x \cdot P^t - \mu\right\|_{\mathsf{TV}} \leq \epsilon$. In other words,

$$\tau(\epsilon) = \sup_{x \in \Omega} \min\{t : \left\|\delta_x P^t - \mu\right\|_{\mathsf{TV}} \leq \epsilon\}.$$

Note that total variation distance between μ, ν always decreases at each step of the Markov chain.

3 September 14, 2023

3.1 Coupling and spectral methods

Definition 3.1. Given distributions μ, ν on Ω, a distribution ξ on $\Omega \times \Omega$ is a *coupling* if $\sum_y \xi(x,y) = \mu(x)$ and $\sum_x \xi(x,y) = \nu(y)$ for all y.

There is always a coupling since we can take $\xi(x,y) = \mu(x)\nu(y)$.

Lemma 3.1 (Coupling)**.** $\|\mu - \nu\|_{\mathsf{TV}} = \inf_\xi \Pr_\xi(x \neq y)$.

Proof. For the upper bound

$$\|\mu-\nu\|_{\mathsf{TV}} = \sum_{x:\mu(x)>\nu(x)} (\mu(x)-\nu(x)) = \sum_x (\mu(x)-\min\{\mu(x),\nu(x)\}) = 1-\sum_x \min\{\mu(x),\nu(x)\} \leq 1-\sum_x \xi(x,x),$$

where the inequality follows since $\xi(x,x) \leq \min\{\mu(x),\nu(x)\}$ by the coupling constraints.

for the reverse direction, we want to show that there exists a coupling ξ so that $\xi(x,x) = \min\{\mu(x),\nu(x)\}$. Now define $A = \{x : \mu(x) > \nu(x)\}$, $B = \{x : \mu(x) < \nu(x)\}$, and let C be everything else. By the design of the coupling, all sub-matrices except the $A \times B$ one have to have all zeros (except the diagonal, which is populated with $\min\{\mu(x),\nu(x)\}$). In particular, we need to design a $A \times B$ matrix M so that

$$\sum_{y \in B} M(x,y) = \mu(x) - \nu(x)$$

for all $x \in A$ and

$$\sum_{x \in A} M(x,y) = \nu(y) - \mu(y)$$

for all $y \in B$. It turns out that we can find a distribution M as required but we don't do the details. (Or, see Proposition 4.7 of the Markov Chain and Mixing Time book for a different perpsective.) □

Definition 3.2 (Coupling of MCs)**.** A coupling for a Markov chain P is a stochastic process $(X_t, Y_t)_{t=0}^\infty$ on $\Omega \times \Omega$ so that $\Pr(X_{t+1} = b | X_t = a) = P(a \to b)$ and $\Pr(Y_{t+1} = b | Y_t = a) = P(a \to b)$.

A trivial coupling is the identity coupling: $X_t = Y_t$ for all t.

We say that a coupling is *Markovian* if $\Pr(X_{t+1} = c | X_t = a, Y_t = b) = \Pr(a \to b)$, and $\Pr(Y_{t+1} = c | X_t = a, Y_t = b) = \Pr(b \to c)$. We can think of this in the following alternative way: for all $(x,y) \in \Omega^2$, we couple $P(x \to \cdot), P(y \to \cdot)$.

Lemma 3.2. *Suppose P is a Markov chain, $X_0 \sim \mu_0, Y_0 \sim \mu$, then*

$$\left\|\mu_0 P^t - \mu\right\|_{\mathsf{TV}} \leq \Pr(X_t \neq Y_t)$$

for any coupling of P.

Example: simple random walk on $\{\pm 1\}^n$. Define a coupling as follows: for $X_t, Y_t \in \{\pm 1\}^n$, we (a) choose the same uniformly random coordinate $i \in [n]$; (b) Toss the *same* coin. Note that if $X_t(i) = Y_t(i)$ for some t, then for all $t' \geq t$, $X_{t'}(i) = Y_{t'}(i)$. By the coupon collector problem, X_t, Y_t coalesce in time $\frac{1}{2}n \log n + cn$ with probability $1 - e^{-c}$.

Recall the fundamental theorem:

Theorem 3.3. *If P is ergodic, then P has a unique stationary distribution μ, and for all $x \in \Omega$, $\left\|\delta_x P^t - \mu\right\|_{\mathsf{TV}} \to 0$ as $t \to \infty$.*

Proof. First we show there exists a stationary distribution, which is an eigenvalue problem. Since $P\mathbf{1} = \mathbf{1}$ and a matrix has the same eigenvalues as its transpose, we can find some $v \in \mathbb{R}^\Omega$ so that $vP = v$. Using the Perron-Frobenius theorem, we get that in fact v has to have non-negative entries, which implies that by normalizing it, we get a stationary distribution.

Note that if the theorem is true, then convergence tells us we must have uniqueness of the stationary distribution.

So the main nontrivial thing is to show convergence. The idea is to use: there exists $t^\star$ so that $P^{t^\star}(x \to y) > 0$ for all $x, y \in \Omega$. In particular, $\varepsilon := \min_{x,y} P^{t^\star}(x \to y) > 0$.

Thus, if we couple two copies of $P^{k \cdot t^\star}$ for any $k \in \mathbb{N}$, there is always an ε probability of coalescence. So, with probability 1, the two chains eventually hit, and once the two things in the coupling hit each other, they will be equal at all timesteps. □

3.2 Constructing good couplings

At a high level: we want to define some graph on Ω and couple $P(x \to \cdot), P(y \to \cdot)$ for neighbors x, y. Then we extend this by a "composition" technique to a full coupling.

Lemma 3.4 (Composition)**.** *If $\mu_1, \mu_2, \mu_3 \in \Delta(\Omega)$ and we have couplings ξ_{12}, ξ_{23} between $(\mu_1, \mu_2), (\mu_2, \mu_3)$, respectively, then $\xi_{13}(x, z) := \sum_y \frac{\xi_{12}(x,y)\xi_{23}(y,z)}{\mu_2(y)}$ is a valid coupling of (μ_1, μ_3).*

Proof. Computation. □

Theorem 3.5 (Bubley-Dyer)**.** *Let $E \subset \binom{\Omega}{2}$ be so that (Ω, E) is a graph. There exists $\alpha > 0$ so that the following holds. If for all neighbors x, y, there exists a coupling of $P(x \to \cdot), P(y \to \cdot)$ so that*

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) | X_t \sim Y_t] \leq 1 - \alpha,$$

then we have $T_{\text{mix}}(\epsilon) \leq O\left(\frac{1}{\alpha} \log(D/\epsilon)\right)$, where D denotes the diameter of the graph and $d(\cdot, \cdot)$ denotes distance on the graph.

Proof. Let $X_t, Y_t \in \Omega$ (not necessarily neighbors). Let $Z_t^{(i)}$ be defined as follows: $X_t = Z_t^{(0)}, \ldots, Y_t = Z_t^{(d)}$ denotes a shortest path between X_t, Y_t. We want a coupling so that

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) | X_t, Y_t] \leq (1 - \alpha) \cdot d(X_t, Y_t).$$

Using the composition of couplings lemma and triangle inequality, we get that there is a coupling between $P(X_t \to \cdot)$ and $P(Y_t \to \cdot)$ so that

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) | X_t, Y_t] \leq \sum_{j=0}^{d-1} \mathbb{E}[d(Z_{t+1}^{(j+1)}, Z_{t+1}^{(j)}) | Z_t^{(j)}, Z_t^{(j_1)}] \leq \sum_{j=0}^{d-1} d(Z_t^{(j)}, Z_t^{(j+1)}) \leq (1 - \alpha) \cdot d(X_t, Y_t),$$

where we write $d = d(X_t, Y_t)$.

The above argument tells us that $d(\mu, \nu) := \inf_\xi \mathbb{E}_{(x,y)\sim\xi} d(x, y) \geq \|\mu - \nu\|_{\mathsf{TV}}$ is contracting at rate $1 - \alpha$. Thus, we get the mixing result. (Here $d(\mu, \nu)$ is the *Wasserstein distance*, studied in optimal transport.) □

3.3 Application: graph colorings

Recall that a q-coloring of a graph is an assignment $\chi : V \to [q]$ so that adjacent vertices have different colors. We will always assume $q \geq \Delta + 1$ where Δ denotes maximal degree (otherwise it becomes hard to find a coloring).

We consider the Glauber dynamics here: If we are at a coloring χ_t, then:

1. Pick $v \sim V$ uniformly at random.
2. Out of the colors available to v (namely all colors not given to neighbors of v), choose one such c uniformly.
3. Set $\chi_{t+1}(v) = c$ and $\chi_{t+1}(u) = \chi_t(u)$ for all $u \neq v$.

Theorem 3.6 (Jerrum '95)**.** *If $q \geq 2\Delta + 1$, then $T_{\text{mix}}(\epsilon) \leq O(qn \log(n/\epsilon))$.*

The upper bound on mixing time is essentially as good as you can hope for. Note that you need $q \geq \Delta + 2$ for Glauber to be ergodic. (For instance, consider the complete graph $K_{\Delta+1}$ – each coloring has no neighborrs.) Some remarks (there has been much recent work on this question):

- If $q \leq \Delta$, then there is no FPRAS unless $\mathsf{RP} = \mathsf{NP}$.
- A major conjecture is that if $q \geq \Delta + 1$ there is a FPRAS; and if $q \geq \Delta + 2$, then you have $O(n \log n)$ mixing.

Now we prove the above theorem:

Proof. We have to construct a coupling when χ_t, χ'_t differ on a unique vertex $v \in V$. We argue as follows: We can ensure that we select the same update vertex u: We need to couple $\mathsf{unif}([q]\backslash\{\chi_t(w) : w \sim u\})$ and $\mathsf{unif}([q]\backslash\{\chi'_t(w) : w \sim u\})$.

We consider various cases for the relation between u and vertex v (the unique vertex at which χ_t, χ'_t disagree).

- If $u \notin N(v) \cup \{v\}$, then $c = c'$, as the two distributions we want to couple are identical, so we can ensure that the new color is identical between the coupled processes.
- Another good case is $u = v$: here we destroy all disagreement since again the two distributions we want to couple are identical, so that the two coloring at the next time step are in fact equal!
- The bad case is that we choose a vertex u that is a neighbor of v: there might be a disagreement in the set of available colors: in particular, looking at the set of available colors availble to u, there will be at most 2 differences, since $\chi'_t(v) \neq \chi_t(v)$. We can couple the two uniform distributions so that they only differ in at most 1 slot of the set of available colors (namely, the slots corresponding to $\chi'_t(v), \chi_t(v)$). (You also have to deal with the fact that the sets of available colors differ in size: you can handle it in the same manner.)

We now analyze the contraction as follows;

$$\mathbb{E}[d(\chi_{t+1}, \chi'_{t+1})|\chi_t, \chi'_t] \leq \frac{n-\Delta-1}{n} + 2 \cdot \frac{\Delta}{n} \cdot \frac{1}{q-\Delta} + \frac{\Delta}{n} \cdot \frac{q-\Delta-1}{q-\Delta} = 1 - \frac{q-2\Delta}{q-\Delta} \cdot \frac{1}{n} < 1,$$

where the final inequality holds if $q \geq 2\Delta + 1$. Above, the first term deals with the $n - \Delta - 1$ vertices in $V \backslash N(v) \cup \{v\}$, for which the distance remains 1; the second term deals with the Δ vertices in $N(v)$ for which the distance increases to 2 with probability $1/(q-\Delta)$, as the number of available colors is at least $q - \Delta$ and the two lists of available colors differ in at most a single slot; and the third term dea lwit hthe Δ vertices in $N(v)$ for which the distance stays at 1 with probability $1 - 1/(q-\Delta)$, again due to our coupling as constructed above. Note that technically we have assumed that the degree of v is Δ, though more generally we're fine by monotonicity.

If $q \geq 2\Delta + 1$, then we can take $\alpha \geq 1/((q-\Delta)n) = \Omega(1/(qn))$ in Theorem 3.5, and, using that the diameter of the colorings graph is at most n, we get the desired bound $O(qn \log(n/\epsilon))$ on mixing time.

□

4 September 19, 2023

Today: we talk about spectral/conductance methods. Consider the Ferromagnetic Ising Model: given a graph $G = (V, E)$ and $\beta > 0$, we want to sample from

$$\mu(\sigma) \propto \exp\left(\frac{\beta}{2}\sigma^\top A\sigma + h \cdot \langle \sigma, \mathbb{1}\rangle\right) \ \forall \sigma \in \{\pm 1\}^V.$$

Theorem 4.1 (Jerrum-Sinclair)**.** *For all G, β, h, there exists as FPRAS for computing $Z_G(\beta, h) = \sum_\sigma \exp\left(\frac{\beta}{2}\sigma^\top A\sigma + h \cdot \langle \sigma, \mathbb{1}\rangle\right)$. (Thus you can efficiently do sampling as well.)*

4.1 Spectral methods

Suppose we have a distribution μ on Ω. This induces a weighted inner product on $\mathbb{R}^\Omega$, given by $\langle f, g\rangle_\mu = \mathbb{E}_{X\sim\mu}[f(x)g(x)]$.

Lemma 4.2. *P is reversible wrt μ iff P is self-adjoint wrt $\langle \cdot, \cdot \rangle_\mu$.*

Proof. We have $\langle \delta_x, P\delta_y\rangle_\mu = P(y \to x) \cdot \mu(y)$ and $\langle \delta_y, P\delta_x\rangle_\mu = P(x \to y) \cdot \mu(x)$. Then extend to inner products on all vectors by linearity. □

Since P is self-adjoint with respect to this inner product, we have that all eigenvlaues of P are real. We can order the eigenvalues by $-1 \leq \lambda_{|\Omega|} \leq \cdots \leq \lambda_2 \leq \lambda_1 = 1$. Irreducibility corresponds to $\lambda_2 < 1$, and aperiodicity corresponds to $\lambda_n > -1$.

Define $\lambda_* := \max_{i>1} |\lambda_i| = \max\{\lambda_2, |\lambda_n|\}$.

We have the following quantitative bound on mixing time:

Theorem 4.3. *If P is ergodic and reversible, then for all $\epsilon > 0$,*

$$T_{\text{mix}}(\epsilon) \leq \frac{1}{1-\lambda_*} \cdot \log \frac{1}{\epsilon \mu_{\min}},$$

where $\mu_{\min} = \min_{x:\mu(x)>0} \mu(x)$.

We remark that there is also a lower bound showing that you must have:

$$T_{\text{mix}}(\epsilon) \geq \frac{1}{1-\lambda_*} \log \frac{1}{2\epsilon}.$$

The idea of proof is to bounds things in terms of χ^2 divergence and thne relate it to TVD. Now,

$$\chi^2(\nu||\mu) = \mathbb{V}_\mu(d\nu/d\mu) = \mathbb{E}_{X\sim\mu}\left[\left|\frac{\nu(x)}{\mu(x)} - 1\right|^2\right].$$

We have the following standard fact:

Lemma 4.4. $\|\mu - \nu\|_{\mathsf{TV}}^2 \leq \frac{1}{4}\chi^2(\nu||\mu)$.

Understanding the spectral gap via the Poincaré inequality. Given $f : \Omega \to \mathbb{R}$, we can define the following quadratic form:

$$\mathcal{E}_P(f,f) = \langle f, (I-P)f\rangle_\mu = \frac{1}{2}\sum_{x,y\in\Omega} \mu(x)P(x\to y)\cdot(f(x)-f(y))^2.$$

We say that P satisfies the Poincare inequality with constant α if for all f, $\alpha \cdot \mathbb{V}_\mu(f) \leq \mathcal{E}_P(f,f)$. We have:

Lemma 4.5. $1-\lambda_* = \inf_f \frac{\mathcal{E}_P(f,f)}{\mathbb{V}_\mu(f)}$. *Thus,* $1-\lambda_*$ *is the best constant you can put in the Poincare inequality.*

Recall that

$$T_{\text{mix}}(\epsilon) = \max_{x\in\Omega} \min\{t \geq 0 : \ \left\|\delta_x P^t - \mu\right\|_{\mathsf{TV}} \leq \epsilon\}.$$

Proof of Theorem 4.3. Let $f = d\delta_x/d\mu$ for some $x \in \Omega$. Note that $\|f\|_\infty \leq 1/\mu_{\min}$. In light of the definition of χ^2 divergence above, it suffices to prove that

$$\mathbb{V}_\mu\left(\frac{d(\nu P^t)}{d\mu}\right) = \mathbb{V}_\mu(P^t f) \leq \lambda_*^{2t}\cdot\mathbb{V}_\mu(f) \leq \lambda_*^{2t}\cdot\frac{1}{\mu_{\min}}\ \forall f.$$

(Note that $P^t\frac{d\nu}{d\mu} = \frac{d(\nu P^t)}{d\mu}$ holds by reversibility, since, for $\nu = \delta_x$, it is equivalent to $\frac{P^t(y\to x)}{\mu(x)} = \frac{P^t(x\to y)}{\mu(y)}$ for all y.) We have that $\mathbb{V}_\mu(f) = \langle f,f\rangle - \langle f,\mathbb{1}\rangle_\mu^2$. Now we can write

$$\mathbb{V}_\mu(f) - \mathbb{V}_\mu(Pf) = \langle f,f\rangle_\mu - \langle f,\mathbb{1}\rangle_\mu^2 - \langle Pf,Pf\rangle_\mu + \langle Pf,\mathbb{1}\rangle_\mu^2 = \langle f,(I-P^2)f\rangle_\mu.$$

The eigenvalues of P^2 are the squares of the eigenvlaues of P, so we have from the Poincare inequality

$$\langle f,(I-P^2)f\rangle_\mu \geq (1-\lambda_*^2)\cdot\mathbb{V}_\mu(f).$$

Thus $\mathbb{V}_\mu(Pf) \leq \lambda_*^2\cdot\mathbb{V}_\mu(f)$. Now use induction on t. □

4.2 Combinatorial methods to bound λ_*.

Definition 4.1 (Conductance)**.** For $S \subset \Omega$, define

$$\Phi(S) = \sum_{x \in S, y \notin S} \mu(x) P(x \to y),$$

where $\mu(S) = \sum_{x \in S} \mu(x)$.

We now define $\Phi(P) = \inf_{S:\mu(S) \leq 1/2} \Phi(S)$. If conductance is lower, there is a set which it is hard to "escape from".

Theorem 4.6 (Cheeger's inequality)**.** $\Phi(P)^2/2 \leq 1 - \lambda_* \leq 2\Phi(P)$.

The conductance is useful for proving lower bounds on mixing: we just need to find some S for which $\Phi(S)$ is small, which shows that $1 - \lambda_*$ is small.

4.3 Flows/paths

For all $x, y \in \Omega$, suppose that to it we have associated a path from x to y in Ω, denoted $P_{x \to y}$. The idea is to route $\mu(x)\mu(y)$ "goods" from $x \to y$ using the transitions in P. Each transition $a \to b$ has cost $\mu(a) \cdot P(a \to b)$. We quantify the efficiency by:

- $\max_{x,y} |P_{x \to y}|$.
- $\max_{a \to b} C_P(a \to b)$, where $C_P(a \to b) = \frac{1}{\mu(a) \cdot P(a \to b)} \sum_{x,y:(a \to b) \in P_{x \to y}} \mu(x)\mu(y)$.

Theorem 4.7. *For all $\{P_{x \to y}\}_{x,y}$, it holds that*

$$\frac{1}{1 - \lambda_2} \leq \max_{x,y} |P_{x \to y}| \cdot \max_{a \to b} C_P(a, b).$$

Proof. Note the following two expressions that appear in the Poincare inequality:

$$\mathbb{V}_\mu(f) = \frac{1}{2} \sum_{x,y} \mu(x)\mu(y)(f(x) - f(y))^2$$

$$\mathcal{E}_P(f, f) = \frac{1}{2} \sum_{x,y} \mu(x) P(x \to y)(f(x) - f(y))^2$$

Now we have $f(x) - f(y) = \sum_{(a \to b) \in P_{x \to y}} f(a) - f(b)$, so we have (from C-S)

$$\begin{aligned}\mathbb{V}_\mu(f) \leq & \max_{x,y} |P_{x \to y}| \cdot \frac{1}{2} \sum_{x,y} \mu(x)\mu(y) \sum_{(a \to b) \in P_{x \to y}} (f(a) - f(b))^2 \\ = & \frac{1}{2} \max_{x,y} |P_{x \to y}| \sum_{a \to b} (f(a) - f(b))^2 \sum_{x,y:(a \to b) \in P_{x \to y}} \mu(x)\mu(y).\end{aligned}$$

The inner summation is $C_P(a \to b) \cdot \mu(a) P(a \to b)$, and this gives us that the above is upper bounded by

$$\max |P_{x \to y}| \cdot \max_{a \to b} C_P(a \to b) \cdot \frac{1}{2} \sum_{a,b} \mu(a) P(a \to b)(f(a) - f(b))^2.$$

The desired lemma statmeent follows by the Poincare inequality. □

Usually, bounding the length of the path $|P_{x\to y}|$ is easy: the challenging part is bounding the congestion $C_P(a,b)$.

Example: μ is uniform on $\{\pm 1\}^n$. Let's order the coordinates $1,\dots,n$. Given $x,y\in\{\pm1\}^n$, we define a path $P_{x\to y}$ by flipping the coordinates in order (if they're not already equal). This leads to $|P_{x\to y}|\le n$. We just need to understand the congestion. Fix $a\to b$ which corresponds to flipping coordinate k. Note that $C_P(a\to b)\le C$ if and only if $|\{(x,y):(a\to b)\in P_{x\to y}\}|\le C\cdot n$. Given a path of the form $x,\dots,a\to b,\cdots y$, note that x must agree with b on coordinates $k+1,\dots,n$ and y must agree with a on coordinates $1,\dots,k$. The number of choices for y is 2^{n-k} and the number of choices for x is 2^k, so the number of possibly $\{x,y\}$ is at most 2^n. So, we can take $C=O(1)$. Summarizing $C_P(a\to b)\le n$ and $|P_{x\to y}|\le n$, meaning that $1-\lambda_*\ge 1/n^2$. This doesn't quite give us the right answer (mixing time should be $O(n)$, not $O(n^2)$). If we plug this into Theorem 4.3 then using that $\log 1/\mu_{\min}=n$, we get mixing time of $O(n^3)$, which is very far off. Recall from last class we can get $O(n\log n)$ using coupling.

4.4 Ferromagnetic Ising model

One natural attempt is to run the Glauber dynamics directly for μ. This may work when β is small (since the quadratic form looks roughly linear). But, it will fail for large β, where $T_{\text{mix}}\ge\exp(\Omega(n))$. This is because you, e.g., won't transition between a component where most vertices are $+1$ and a component where most vertices are -1.

The fix is to transform μ: let $0\le\rho,\lambda\le 1$ be two parameters. Define $\hat\mu$ on 2^E by $\hat\mu(F)\propto\rho^{|odd(F)|}\cdot\lambda^{|F|}$, where $odd(F)$ is the set of odd-degree vertices in (V,F). We are penalizing subsets F that induces lots of odd-degree vertices.

Now we define $\hat Z_G(\rho,\lambda)=\sum_{F\subset E}\rho^{|odd(F)|}\lambda^{|F|}$.

Proposition 4.8. *For all β,h, let $\rho=\tanh(\beta)$, $\lambda=\tanh(h)$. Then $Z_G(\beta,h)=C(\beta,h)\cdot\hat Z_G(\rho,\lambda)$, where*

$$C(\beta,h)=2^{|V|}\cosh(h)^{|V|}\cosh(\beta)^{|E|}.$$

The idea is that roughly speaking, by flipping all the signs from some configuration that has high probability under μ, you can get another configuration that has high probability. One we have this, we can sample from the modified distribution, compute its partition function, and then use the above proposition to transform back to the ferromagnetic Ising model. To prove the above proposition:

Proof. Assume $h=0$ for convenience. The key identity is that $e^x=\cosh(x)\cdot(1+\tanh(x))$. Now we can write

$$\exp(\beta/2\cdot\sigma^\top A\sigma)=\prod_{uv\in E}\exp(\beta\sigma_u\sigma_v)=\cosh(\beta)^{|E|}\prod_{uv\in E}(1+\sigma_u\sigma_v\tanh(\beta)),$$

where the second inequality uses that tanh is an odd function. Now we have

$$\prod_{uv\in E}(1+\sigma_u\sigma_v\tanh(\beta))=\sum_{F\subset E}\tanh(\beta)^{|F|}\prod_{v\in odd(F)}\sigma_u,$$

since vertices with even degree have the corresponding product of σ_u cancel out to 1. Now we sum over all configurations:

$$Z_G(\beta, 0) = \sum_{\sigma} \exp(\beta/2 \cdot \sigma^\top A \sigma) = 2^n \cosh(\beta)^{|E|} \sum_{F \subset E} \tanh(\beta)^{|F|} \cdot \mathbb{E}_{\sigma \sim \{\pm 1\}^n} \left[\prod_{v \in odd(F)} \sigma_v \right].$$

The product on the RHS is 1 if $odd(F) = \emptyset$, and 0 otherwise. Now

$$C(\beta, 0) \cdot \hat{Z}_G(\rho, \lambda) = 2^n \cdot \cosh(\beta)^{|E|} \cdot \sum_{F \subset E} \tanh(\beta)^{|odd(F)|} \cdot \mathbb{1}\{|odd(F)| = 0\},$$

since $\lambda = 0$, which by inspection we see is equal to $Z_G(\beta, 0)$. The case of general $h \in \mathbb{R}$ is similar (see the class notes). □

5 September 21, 2023

Recall Proposition 4.8, which allows us to reduce sampling from the Ferromagnetic Ising model to computing the partition function for the distribution $\hat{\mu}(F) \propto \rho^{|odd(F)|} \cdot \lambda^{|F|}$. Today we do this latter task:

Theorem 5.1. *For the Glauber dynamics for $\hat{\mu}$, $T_{\text{mix}}(\epsilon) \leq O\left(\frac{m^2}{\rho^6} \cdot (m + \log 1/\epsilon)\right)$.*

Note that we need some dependence on ρ: if we consider the cycle, then as $\rho \to 0$, most of the mass is concentrated on the empty set and the set of all vertices, which are far apart. So, there must be some $1/\rho$ dependence.

Remark. Note that it's not clear that $\hat{\mu}$ is self-reducible, so whether the counting-sampling equivalence works is not clear. We handel this later.

5.1 Flow encodings

Let's fix a collection of paths $P_{x \to y}$ for each pair (x, y). Define $CP(a \to b) = \{(x, y) : (a \to b) \in P_{x \to y}\}$. The main challenge is bounding $\sum_{(x,y) \in CP(a \to b)} \mu(x)\mu(y)$ by $\text{poly}(n) \cdot \mu(a) \cdot P(a \to b)$.

Definition 5.1 (Flow encoding)**.** A flow encoding is a collection of mappings $\{\eta_{a \to b} : CP(a \to b) \to \Omega\}_{a \to b}$, where $\eta_{a \to b}$ are all injective.

Lemma 5.2. *If we have a flow encoding $\{\eta_{a \to b} : CP(a \to b) \to \Omega\}$ so that for all transitions $a \to b$, for all $(x, y) \in CP(a \to b)$,*

$$\mu(x)\mu(y) \leq \alpha \cdot \mu(a)P(a \to b) \cdot \mu(\eta_{a \to b}(x, y)),$$

then the congestion $C_P(a \to b) \leq \alpha$.

Proof. We have

$$\sum_{(x,y) \in CP(a \to b)} \mu(x)\mu(y) \leq \alpha \mu(a)P(a \to b) \sum_{(x,y) \in CP(a \to b)} \mu(\eta_{a \to b}(x, y)) \leq \alpha \mu(a)P(a \to b),$$

where the final inequality uses that $\sum_{(x,y) \in CP(a \to b)} \mu(\eta_{a \to b}(x, y)) \leq 1$ by injectivity of $\eta_{a \to b}$. □

Example: hypercube. Remember that, given $a \to b$, to get from x to y, we first flip coordinates $1 \dots k-1$, then we flip $a \to b$ at coordinate k, and then we flip $k+1 \dots n$. A way to get from y to x is to flip $1 \dots k-1$, then are at a', flip k to get to b', then flip $k+1 \dots n$. There is a sort of duality between (a,b) and (a',b') here. So we could define $\eta_{a\to b}(x,y) = a'$. This turns out to be a good choice.

Now we prove the main theorem:

Proof of Theorem 5.1. Consider subsets $I, F \subset E$, and denote their symmetric difference by $I \oplus F \subset E$. We can break $I \oplus F$ into a collection of paths and cycles which are all edge-disjoint, so that the number of paths in this collection is equal to half the number of odd degree vertices: i.e., $2|\{paths\}| = |odd(I \oplus F)|$. (In particular, if there are any odd-degree vertices, take a path from one odd-degree vertex to another. Repeat until no odd-degree vertices left, and then we can find a cycle.)

Defining $P_{x\to y}$. The idea in constructing $P_{x\to y}$ is to "flip in order of traversal of the paths and cycles". We illustrate with an example. Suppose G is a graph on 5 vertices with edges $12, 23, 34, 41, 15, 52$. Suppose I has edges $23, 34, 41$ and F has edges $41, 12, 23, 15, 52$. then $I \oplus F$ is $12, 25, 51, 34$. The decomposition of $I \oplus F$ is a path 34 and a cycle $12, 25, 51$. Then the path from I to F is (where we write the edges that we are flipping)

$$34,\ 15,\ 12,\ 52.$$

The probabilities of the respective subsets are $\hat{\mu}(I) \propto \rho^2\lambda^3$, $\rho^4\lambda^2$, $\rho^4\lambda^3$, $\rho^4\lambda^4$, $\rho^4\lambda^5 \propto \hat{\mu}(F)$.

Note that we can also flip going from F to I using the same order of edges we flip (i.e., $34, 15, 12, 52$).

Defining the flow encoding. Given $I \to \dots \to A \to_{e_k} B \to \dots \to F$, we define the flow encoding to be

$$\eta_{A\to B}(I,F) = I \oplus F \oplus (A \cup B).$$

We could have replaced $A \cup B$ with the larger of A, B, though the above gives us a better bound. Note that this is exactly as we did in the hypercube example above.

To prove injectivity, suppose htat $M \in image(\eta_{A\to B})$. We can definitely recover $I \oplus F = M \oplus (A \cup B)$, since A, B are fixed. Note that for each pair (I, F) we have chosen a canonical order on the edges we're flipping (i.e., for each $I \oplus F$, which we know, there is a canonical ordering). We just need to know which point on this canonical path we're at: but this point is defined by the edge $e_k = A \oplus B$ that we're flipping at this step. More precisely, if the sequence of edges we flip to get from I to F is $e_1, \dots, e_n$, then we know that B agrees with I on $e_{k+1}, \dots, e_n$ and with F on $e_1, \dots, e_k$. So, once we know e_k, we can flip $e_k, e_{k+1}, \dots, e_n$ to get to F and flip $e_1, \dots, e_{k-1}$ to get to I.

Congestion bound. We claim that the hypothesis of Lemma 5.2 holds with $\alpha = m/\rho^6$. The first claim is that

$$\mu(A) \cdot P(A \to B) \geq \frac{\rho^2}{m} \cdot \mu(A \cup B).$$

To prove this, since A, B differ on one edge, we have by reversibility that $\mu(A) \cdot P(A \to B) = \mu(A \cup B) \cdot P(A \cup B \to A \cap B)$. Since $A \cup B$ is bigger than $A \cap B$, we have that $P(A \cup B \to A \cap B) \geq \rho^2/M$

(with the $1/m$ factor coming from the fact that we choose a random edge to flip and thoe ρ^2 comes from the fact that we may increase the number of odd vertices by 2 – but we don't pick up a factor of λ since $A \cup B$ is bigger). This proves our first claim.

So, it remains to show that

$$\mu(I)\mu(F) \leq \rho^{-4} \cdot \mu(T) \cdot \mu(I \oplus F \oplus T),$$

where $T = A \cup B$. Here the "dual paths" have intermediate verticese T and its dual $I \oplus F \oplus T$. Equivalently (by cancellations of the partition function and cancellation of λ factors) we want to show that

$$\rho^{|odd(F)|}\rho^{|odd(I)|} \leq \rho^{-4}\rho^{|odd(T)|}\rho^{|odd(I \oplus F \oplus T)|}.$$

Since $0 \leq \rho \leq 1$, the above is equivalent to

$$|odd(T)| + |odd(I \oplus F \oplus T)| \leq 4 + |odd(I)| + |odd(F)|.$$

To prove the above, the key insight is that each time we flip an edge, we change the numer of odd degree vertices by at most 2. Moreover, if we flip adjacent edges, then the number of odd degree vertices also changes by at most 2 (since the intermediate vertex has degree changing by 2). But what about disjoint paths? To deal with that, we first make the following claim:

Claim 5.3. *If we have finished a path/cycle to get to T, then*

$$|odd(T)| + |odd(I \oplus F \oplus T)| = |odd(I)| + |odd(F)|.$$

Proof. If T is a cycle the equality is immediate. So suppose T is gotten from I by following a path from u to v. By construction we know that u and v must both be odd-degree vertices in $I \oplus F$: so, u is in exactly one of $odd(F), odd(I)$, and same for v. Suppose that $u \in odd(I), u \notin odd(F)$ and $v \in odd(F), v \notin odd(I)$ (we treat the general case below). Suppose for simplicity that this path is the first in the sequence: in particular, that T is gotten by traversing this path starting from I, and that $I \oplus F \oplus T$ is gotten by tranversing a path starting from F.

To compute $odd(T)$: we have removed $u \in odd(I)$ from $odd(I)$, and added v to $odd(T)$. So, after this path, we have that $odd(T) = odd(I) - u + v$. Similarly, $odd(I \oplus F \oplus T) = odd(F) + u - v$.

(If we add u to get to T, we must remove u to get to $I \oplus F \oplus T$, and vice versa.)

Another way of seeing it is as follows: in general, we have $odd(A \oplus B) = odd(A) \oplus odd(B)$. If the path that we have traversed is denoted by P (so that $T = I \oplus P$), then it follows that $odd(I \oplus F \oplus T) = odd(F) \oplus odd(P)$ and $odd(T) = odd(I) \oplus odd(P)$. Now clearly $|odd(I)| + |odd(F)| = |odd(F) \oplus odd(P)| + |odd(I) \oplus odd(P)|$, since the only odd-degree vertices of P are its endpoints. □

Given the above claim, we then show that if we're in the middle of a path or a cycle, then the sum of odd degree vertices of the two sets can't change by more than 4. In particular, T is in the middle of a path or cycle: we want to compare $odd(T)$ to what we would have if we have finished traversing the rest of the path or cycle. In particular, suppose this path/cycle is defined by $e_1, \ldots, e_k$, with T corresponding to $e_1, \ldots, e_\ell$. Let T_{init} be what we get starting from T and untraversing $e_1, \ldots, e_\ell$, and T_{final} be what we get starting from T and traversing $e_{\ell+1}, \ldots, e_k$. The second claim is:

Claim 5.4. $|odd(T)| - \frac{|odd(T_{init})| + |odd(T_{final})|}{2} \leq 2.$

Proof. Each time we traverse two consecutive edges along a traversal, the shared vertex has the parity of its degree remain the same. So, $|odd(T)|$ is within 2 of both $|odd(T_{init})|$ and $|odd(T_{final})|$. □

Given the above claim, we apply it to both T as well as the analogous case $I \oplus F \oplus T$. □

6 September 26, 2023

Last time, we showed how to sample from the ferromagnetic Ising model. When does Glauber dynamics mix? For most classes of graphs, there exists β_c so that for $\beta < \beta_c$, Glauber dynamics mixes in $O(n \log n)$ steps and for $\beta > \beta_c$, Glauber needs $\exp(n)$ steps (rougly).

How do we understand this phase transition? Use something called *correlation decay* [Weitz '06], [Bandyopadhyay-Gamarnik]. But first we introduce some useful definitions/concepts.

Definition 6.1 (Spin system)**.** Given $q \in \mathbb{N}$, a q-spin system consists of:

- An interaction matrix $A \in \mathbb{R}_{\geq 0}^{q \times q}$.
- External fields $\lambda \in \mathbb{R}_{\geq 0}^{q}$.
- A graph $G = (V, E)$.

The Gibbs distribution is defined as follows: for $\sigma : V \to [q]$,

$$\mu(\sigma) \propto \prod_{uv \in E} A_{\sigma(u),\sigma(v)} \prod_{v \in V} \lambda_{\sigma(v)} \tag{1}$$

Some examples:

- If $A = \begin{pmatrix} e^{2\beta} & 1 \\ 1 & e^{2\beta} \end{pmatrix}$, we recover the ferromagnetic Ising model.
- If $A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, we recover sampling from independent sets.
- If A is the all-1s matrix with 0s on the diagonal, we recover sampling from proper q-colorings.

We define a *pinning* to be a map $\tau : S \to [q]$, for some $S \subset V$. This induces a conditional distribution

$$\mu^{\tau}(\sigma) \propto \begin{cases} \mu(\sigma) & : \sigma|_S = \tau \\ 0 & : \text{otherwise.} \end{cases}$$

Lemma 6.1 (Global Markov property)**.** *For all partitions $A \sqcup S \sqcup B$ so that all paths $A \ni u \rightsquigarrow v \in B$ must go through S and all pinnings $\tau : S \to [q]$, then $\mu^{\tau} = \mu_A^{\tau} \otimes \mu_B^{\tau}$, where μ_A^{τ} denotes the marginal on A (and similarly for μ_B^{τ}).*

We don't give a formal proof but note that it is intuitively clear (all correlation between A, B is separated by S).

Correlation decay.

Definition 6.2 (Weak spatial mixing)**.** For all $r \in V$ and $S \subset V - \{r\}$ (think of S as a "sphere" around r), we have *weak spatial mixing* if for all $\tau, \sigma : S \to [q]$ we have

$$\|\mu_r^\tau - \mu_r^\sigma\|_{\mathsf{TV}} \leq C \cdot (1-\delta)^{\mathrm{dist}(r,S)},$$

where $0 < \delta < 1$ is a constant and $C > 0$ is a constant.

Intuitively, if r is far from S, then the spins of vertices of S shouldn't influence the distribution of the spin at r.

Definition 6.3 (Strong spatial mixing)**.** We have *strong spatial mixing* if, in the context of the above definition,

$$\|\mu_r^\tau - \mu_r^\sigma\|_{\mathsf{TV}} \leq C \cdot (1-\delta)^{\mathrm{dist}(r,S_{\tau,\sigma})},$$

where $S_{\tau,\sigma} = \{v \in S : \ \tau(v) \neq \sigma(v)\}$.

You can think of strong spatial mixing as asking for weak spatial mixing on all possible conditional distributions, where you condition on only the nodes where τ, σ disagree.

What is a phase transition? For this lecture, think of it as follows: there exists a threshold for some parameter, λ_c, so that for $\lambda > \lambda_c$ WSM/SSM fails but for $\lambda < \lambda_c$, WSM/SSM holds.

6.1 Hardcore model

This is the special case of a spin system regarding sampling from independent sets. To spell it out, given $G = (V, E)$ and $\lambda \geq 0$, we have $\mu(I) \propto \lambda^{|I|}$, for $I \subset V$ an independent set. This is a discretization of a famous statistical physics model called the *hard gas* model. Mathematically, the hardcore model is universal in some sense. Also, if you pin a vertex, it corresponds to deleting that vertex and its entire neighborhood (so convenient).

We denote the partition function by $Z_G(\lambda) = \sum_{I \subset V} \lambda^{|I|}$ (called the independence polynomial of the graph). We let Δ denote the max-degree of G. If λ is large, this distribution concentrates on maximum independent sets, which are hard to find. If λ is small, concentrates on small independent sets (should be easy). We define:

Definition 6.4. We define the critical threshold $\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta} \approx e/(\Delta-1)$.

Theorem 6.2 (Weitz '06)**.** *there exists an FPTAS for computing $Z_G(\lambda)$ for all $\lambda < \lambda_c(\Delta)$ on all graphs of max-degree Δ. SSM holds on all such graphs.*

We also have a converse:

Theorem 6.3 (Sly '10)**.** *For all $\lambda > \lambda_c(\Delta)$, if there exists an FPRAS for $Z_G(\lambda)$ on all max-degree Δ graphs G, then $\mathsf{NP} = \mathsf{RP}$.*

We also remark that it was previously well-known that WSM (and thus SSM) fails for $\lambda > \lambda_c(\Delta)$.

The focus of the next two lectures is on the Weitz proof. As way of intuition, we remark that the worst-case graph G for the Weitz theorem is the infinite Δ-regular tree.

We remark that it is an open problem to design an efficient algorithm for $\lambda = \lambda_c(\Delta)$. We know that correlations still decay here, but not exponentially fast.

6.2 Correlation decay on trees

We begin by proving correlation decay on trees.

Proposition 6.4. *If $\lambda < \lambda_c(\Delta)$, then for all max-degree Δ trees T rooted at r, SSM holds for r.*

If we have a tree rooted at r, then given a vertex u, the subtree rooted at u, T_u, consists of all nodes further away from r than u.

Definition 6.5. Define $P_u = \Pr_{I \sim \mu_{T_u,\lambda}}[u \in I]$.

Lemma 6.5 (Tree recursion)**.** *If $u_1, \dots, u_d$ are neighbors of r, then*

$$P_r = F_d(p_1, \dots, p_d) := \frac{\lambda \prod_{i=1}^d (1-p_i)}{1 + \lambda \prod_{i=1}^d (1-p_i)}.$$

Here we write $p_i = p_{u_i}$.

This lemma gives us a way to compute the marginals for each vertex in a tree: first compute the marginals for the leaves (which are just $p_u = \lambda/(1+\lambda)$), and then we work upwards using tree recursion. Then, we can use the marginals to compute the partition function. We now prove the tree recursion lemma:

Proof. We write $Z_T(\lambda) = Z_{T-r}(\lambda) + \lambda \cdot Z_{T-N(r)}(\lambda)$. Moreover, $Z_{T-r}(\lambda) = \prod_{i=1}^e Z_{T_{u_i}}(\lambda)$. Similarly, $Z_{T-N(r)}(\lambda)$ decomposes. We don't write out the full details. □

We can now write, using the tree recursion lemma,

$$\|\mu_r^\tau - \mu_r^\sigma\|_{\mathsf{TV}} = |p_r^\tau - p_r^\sigma| = |F_d(p_{u_1}^\tau, \dots, p_{u_d}^\tau) - F_d(p_{u_1}^\sigma, \dots, p_{u_d}^\sigma)|,$$

where τ, σ are restricted to the appropriate subtrees. We want to show htat F_d is a contraction, to bound the above by $(1-\delta) \cdot \max_i |p_{u_i}^\tau - p_{u_i}^\sigma|$. We then repeatedly apply this until we get to some level where there is a vertex in $S_{\tau,\sigma}$.

The extremal case is as follows: the tree $\hat{\mathbb{T}}_{\Delta,L}$ which is a depth-L, $\Lambda - 1$-ary tree there τ is identically 0 on layer L and σ is identiclaly 1 on layer L. We can then specialize the tree recursion formula, which gives (by symmetry)

$$f_d(p) = \frac{\lambda(1-p)^d}{1 + \lambda(1-p)^d}.$$

The corresponding special case of what we want to show is (by monotonicty of f_d)

$$|f_d^{\circ L}(0) - f_d^{\circ L}(1)| \le C \cdot (1-\delta)^L.$$

If this is the case, then for all $p \in [0,1]$, we must have $\{f_d^{\circ L}(p)\}_{L=0}^\infty \to \hat{p}$, for some fixed $\hat{p}$ as $L \to \infty$. Moreover, $\hat{p}$ must be a fixed point of f_d. (Note that $\hat{p}$ depends on λ, p.) So, it all comes down to whether a univariate function has a unique fixed point (which it does, since it's monotone decreasing) and whether it converges to this fixed point (which it may not, if λ is too large).

Proposition 6.6. *If $\lambda \le \lambda_c(\Delta)$ then for all p, $\{f_d^{\circ L}(p)\}_L \to \hat{p}$. Furthermore, if $\lambda \le (1-\delta)\lambda_c(\Delta)$, then $|f_d'(\hat{p})| \le 1 - O(\delta)$ and $|f_d^{\circ L}(p) - \hat{p}| \le C \cdot (1 - O(\delta))^L$.*

If $\lambda > \lambda_c(d)$, then there exist $\hat{p}_{odd} < \hat{p} < \hat{p}_{even}$ so that $f_d^{\circ 2L}(p) \to \begin{cases} \hat{p}_{odd} & : p < \hat{p} \\ \hat{p}_{even} & : p > \hat{p} \\ \hat{p} & : p = \hat{p} \end{cases}$.

We remark that $\hat{p}_{odd}, \hat{p}_{even}$ are fixed points of $f_d^{\circ 2}$. So, as L gets large, we actually don't approach $\hat{p}$ in the λ large regime.

We don't prove the above proposition (it's just direct calculation).

6.3 The multivariate case

To solve the multivariate case, we will reduce to the single-variate case. We let $\varphi : [0,1] \to \mathbb{R}_{\ge 0}$ be an increasing "potential function". We measure $|\varphi(p_r^\tau) - \varphi(p_r^\sigma)|$. In particular, we have a new tree recursion, defined by:

$$G_d(m_1, \dots, m_d) = \varphi(F_d(\varphi^{-1}(m_1), \dots, \varphi^{-1}(m_d))),$$

where $m_i = \varphi(p_i)$. We claim that

$$\|\nabla G_d(m)\|_1 \le 1 - O(\delta),$$

which will imply that we get the contraction we want (essentially by mean value theorem).

Lemma 6.7. *Suppose $\lambda \le (1-\delta)\lambda_c(\Delta)$. Let $\varphi(p)$ be so that $\Phi(p) = \varphi'(p) = \frac{1}{\sqrt{p}\cdot(1-p)}$. Then $\|\nabla G_d(m)\|_1 \le 1 - O(\delta)$ for all m and for all $1 \le d \le \Delta - 1$.*

Proof. We have

$$|\partial_{m_i} G_d| = \left| \frac{\Phi(F_d(p))}{\Phi(p_i)} \cdot \partial_{p_i}(F_d(p)) \right| = \sqrt{F_d(p)} \cdot \sqrt{p_i}$$

by the Chain rule. Why is this a good thing? If p_i is large, then by the repulsive nature of the hardcore model it will force $F_d(p)$ to be smaller. The two terms above, $F_d(p)$ and p_i essentially control each other, and thus their produce is small. Hence the partial derivative will be small. □

We complete the proof next time.

7 September 28, 2023

Remember the setting: we study the hardcore model, $\mu(I) = \mu_{G,\lambda}(I) \propto \lambda^{|I|}$ for independent sets I. We have the tree recursion formula, which gives that $p_r = F_d(p_1, \dots, p_d) = \frac{\lambda \prod_{i=1}^d (1-p_i)}{1+\lambda \prod_{i=1}^d (1-p_i)}$. Moreover, we write $f_d(p) = F_d(p \cdot \mathbf{1})$ and $g_d(p) = G_d(p \cdot \mathbf{1}) = \varphi \circ f_d \circ \varphi^{-1}$.

Today the goal is to prove Weitz Theorem In particular, we want to show that if $\lambda \le (1 - \delta) \cdot \lambda_c(\Delta)$, then SSM holds on all graphs of max-deg Δ. Furthermore, there exists an FPTAS for computing $Z_G(\lambda)$.

The subgoal is to establish contraction for the single-variable case, which we did last time. Now we extend this to the multivariate case:

Lemma 7.1. *For all $m \in \mathbb{R}^d_{\geq 0}$, there exists $\bar{m}$ so that $\|\nabla G_d(m)\|_1 \leq |g_d'(\bar{m})|$.*

Proof. Writing $p = \varphi^{-1}(m)$ (i.e., so that φ is applied coordinate-wise), we have

$$\partial_i G_d(m) = \frac{\Phi(F_d(p))}{\Phi(p_i)} \cdot \partial_i F_d(p) = -\sqrt{F_d(p)} \cdot \sqrt{p_i},$$

where we use thatt $\partial_i F_d(p) = \frac{1-F_d(p)}{1-p_i} \cdot F_d(p)$ by definition of F_d. Recall that $\Phi = \varphi' = \frac{1}{\sqrt{p}(1-p)}$, and $G_d = \varphi \circ F_d \circ \varphi^{-1}$.

It follows that

$$\|\nabla G_d(m)\|_1^2 = F_d(p) \left(\sum_{i=1}^d \sqrt{p_i} \right)^2 \leq d \cdot F_d(p) \sum_{i=1}^d p_i,$$

where the inequality uses Cauchy-Schwarz. Using the definition of F_d, the above is equal to

$$d^2 \cdot \frac{\lambda \prod_{i=1}^d (1-p_i)}{1 + \lambda \prod_i (1-p_i)} \cdot \left(1 - \frac{1}{d} \sum_{i=1}^d (1-p_i) \right) \leq d^2 \frac{\lambda \prod_i (1-p_i)}{1 + \lambda \prod_i (1-p_i)} \cdot \left(1 - \prod_{i=1}^d (1-p_i)^{1/d} \right),$$

where we use AM-GM. Now we set $\bar{p} = 1 - \prod_{i=1}^d (1-p_i)^{1/d}$ and $\bar{m} = \varphi(\bar{p})$. Since $g_d'(m) = d \cdot \partial_1 G_d(p \cdot \mathbf{1}) = -d\sqrt{f_d(p)} \cdot \sqrt{p}$ (where again $p = \varphi^{-1}(m)$), the above is equal to

$$d^2 \cdot \frac{\lambda(1-\bar{p})^d}{1 + \lambda(1-\bar{p})^d} \cdot \bar{p} = (g_d'(m))^2,$$

as desired. □

Lemma 7.2. *For all $m \in \mathbb{R}_{\geq 0}$, $|g_d'(m)| \leq \sqrt{|f_d'(\hat{p})|}$, where $\hat{p}$ is defined to satisfy $\hat{p} = f_d(\hat{p})$.*

The main point is that $\sqrt{|f_d'(\hat{p})|}$ is less than $1 - O(\delta)$ if $\lambda \leq (1-\delta)\lambda_c(\Delta)$ (see the notes for a proof of this UB on fixed point). We skip the proof of the above lemma. Summarizing, we have shown that $\|\nabla G_d(m)\|_1 \leq 1 - O(\delta)$ for all $m \in \mathbb{R}^d_{\geq 0}$, which completes the proof of Lemma 6.7 from last time.

7.1 Reduction from general graphs to trees

Tree recursion no longer works since if we remove a vertex, we no longer break the graph into disjoint subgraphs. We begin with an approach to compute marginal probabilities inefficiently.

Theorem 7.3. *Lete G be a graph, $r \in V$, with neighbors $u_1, \ldots, u_d$. For $k \in [d]$, let G_k be octained from G by deleting $r, u_1, \ldots, u_{k-1}$. Write $p_{G_r} := \Pr_{I \sim \mu_{G,\lambda}}[r \in I]$. Then*

$$p_{G,r} = F_d(p_{G_1,u_1}, \ldots, p_{G_d,u_d}).$$

Note that the above gives an inefficient but exact way of computing marginal probabilities: we recursively call the algorithm to query each of p_{G_k,u_k}, using that G_k is smaller than G for each k. The recursion must always terminate. We can actually imagine an exponentially large tree corresponding to the recursive execution of this algorithm: the root is (G, r), and internal nodes are nodes of the form (H, u) where H is a subgraph of G and $u \in H$, and the leaves are graphs with isolated vertices (together with a chosen vertex from that set of isolated vertices). This is

sometimes called a *computation tree* or *self-avoiding walk tree*: its root-to-leaf paths correspond to paths in G which never backtrack along already-visited nodes. We denote this tree by $T_{SAW}(G, r)$.

We notes the following properties: the max-degree of T_{SAW} is bounded above by the max-degree of G. Second, it preserves the shortest path distance to the root r.

Moreover, suppose that $\Lambda \subset V - \{r\}$ and $\tau : \Lambda \to \{in, out\}$ is a pinning. We can lift τ to a pinning $\tau_{SAW} : \Lambda_{SAW} \to \{in, out\}$, where $\Lambda_{SAW} = \{(H, u) : u \in \Lambda\}$.

It turns out that if we have correlation decay in T_{SAW}, then we get correlation decay in the original graph G: correlation decay is only harder to achieve in G.

Now we prove the reduction lemma:

Proof. We have

$$p_{G,r} = \frac{Z_g(r \leftarrow \text{"}in\text{"})}{Z_G(r \leftarrow \text{"}in\text{"}) + Z_G(r \leftarrow \text{"}out\text{"})},$$

where $Z_G(r \leftarrow \text{"}in\text{"})$ denotes the measure of all independent sets including r, and similarly for $Z_G(r \leftarrow \text{"}out\text{"})$. Given the graph G with root r and neighbors $u_1, \dots, u_d$, we create a new graph $\widetilde{G}$ with d copies of r, denoted $r_1, \dots, r_d$. Each r_i has a single neighbor which is u_i. Moreover, we downweight the λ's for each r_i to $\lambda^{1/d}$ (to make sure weighting works out).

Then we now have:

$$p_{G,r} = \frac{Z_{\widetilde{G}}(r_1, \dots, r_d \leftarrow \text{"}in\text{"})}{Z_{\widetilde{G}}(r_1, \dots, r_d \leftarrow \text{"}in\text{"}) + Z_{\widetilde{G}}(r_1, \dots, r_d \leftarrow \text{"}out\text{"})}.$$

Now we can peel out each vertex, one at a time. We will shoot for a recursion of the form

$$Z_{\widetilde{G}}(r_1, \dots, r_d \leftarrow \text{"}in\text{"}) \approx \prod_{k=1}^{d} \frac{Z_{\widetilde{G}}(r_1, \dots, r_k \leftarrow \text{"}in\text{"})}{Z_{\widetilde{G}}(r_1, \dots, r_{k-1} \leftarrow \text{"}in\text{"})}$$

To formally do this, we let $\widetilde{G}_k$ denote the graph G with edges between r_i, u_i removed for $i \geq k+1$. Then by telescoping we have

$$Z_{\widetilde{G}_0}(r_1, \dots, r_d \leftarrow \text{"}in\text{"}) = \widetilde{Z}_{G_0} \cdot \prod_{k=1}^{d} \frac{Z_{\widetilde{G}_k}(r_1, \dots, r_k \leftarrow \text{"}in\text{"})}{Z_{\widetilde{G}_{k-1}}(r_1, \dots, r_{k-1} \leftarrow \text{"}in\text{"})}$$

Note that in the hardcore model, we have $Z_{\widetilde{G}}(r_1, \dots, r_d \leftarrow \text{"}out\text{"}) = \widetilde{Z}_{G_0}$.

We now have

$$Z_{\widetilde{G}}(r_1, \dots, r_{k-1} \leftarrow \text{"}in\text{"}, r_k \leftarrow \text{"}in\text{"}) = \lambda^{1/2} \cdot Z_{\widetilde{G}_{k-1}}(r_1, \dots, r_{k-1} \leftarrow \text{"}in\text{"}, u_k \leftarrow \text{"}out\text{"}).$$

Now, each ratio in the product above then is equal to $\lambda^{1/d} \cdot (1 - \Pr_{\widetilde{G}_k}(u_k \leftarrow \text{"}in\text{"}))$, where G_k is $\widetilde{G}$ but with $r, r_1, \dots, r_{k-1}, u_1, \dots, u_{k-1}$ deleted.

By definition of F_d, this gives the desired form of $p_{G,r}$. □

One can try generalize the above to multiple spins: naively, you blow up the number of instances (i.e., degree) by a factor of q, since you have to look at the probability that the root is labeled by each element of $[q]$. One open problem is to find a generalization of the above argument where you don't blow up the degree like this.

7.2 Proof of strong spatial mixing

Proof. We are given two pinnings $\tau, \sigma : \Lambda \to \{in, out\}$, where $\Lambda \subset V - \{r\}$. We want to upper bound $\|\mu_r^\tau - \mu_r^\sigma\|_{\mathsf{TV}} = |p_{G,r}^\tau - p_{G,r}^\sigma|$.

The above is upper bounded by

$$|\varphi(p_{G,r}^\tau) - \varphi(p_{G,r}^\sigma)| = |G_d(\varphi(p_{G_1,u_1}^\tau), \ldots, \varphi(p_{G_d,u_d}^\tau)) - G_d(\varphi(p_{G_1,u_1}^\sigma), \ldots, \varphi(p_{G_d,u_d}^\sigma))|$$

We previously showed that the above is upper bounded by

$$(1 - O(\delta)) \cdot \max_{1 \le i \le d} |\varphi(p_{G_i,u_i}^\tau) - \varphi(p_{G_i,u_i}^\sigma)|.$$

(Here we have used also that $|G_d(m) - G_d(m')| \le \sup_{\bar{m}} \|\nabla G_d(\bar{m})\|_1 \cdot \|m - m'\|_\infty$, as well as our upper bound on $\sup_{\bar{m}} \|\nabla G_d(\bar{m})\|_1 \le 1 - O(\delta)$.)

By induction, we get that it is upper bounded by $(1 - O(\delta))^{dist(r, \Lambda_{\tau,\sigma})}$. In the course of this recursion, we might come across vertices in $\Lambda - \Lambda_{\tau,\sigma}$: but any such vertex (H, u) (i.e., with H a subgraph of G and $u \in \Lambda - \Lambda_{\tau,\sigma}$) must satisfy $p_{H,u}^\tau = p_{H,u}^\sigma$, and hence $\varphi(p_{H,u}^\tau) = \varphi(p_{H,u}^\sigma)$ since such vertices u must satisfy $\tau(u) = \sigma(u)$. In particular, the recursion can stop at such vertices. Moreover, to bound $|p_{G,r}^\tau - p_{G,r}^\sigma| \le C \cdot |\varphi(p_{G,r}^\tau) - \varphi(p_{G,r}^\sigma)|$, we need $\inf_{q \in [0,1]} |\varphi'(q)| \ge c > 0$, since $\varphi'(q) = \frac{1}{\sqrt{q} \cdot (1-q)}$ (e.g., we may take $c = 1/10$). Finally, at the last step of recursion, we need a bound on $\max_{p,q \in [0,1]} |\varphi(p) - \varphi(q)|$, which is similarly upper bounded by an absolute constant.

Now here is the algorithm to estimate $Z_G(\lambda)$:

1. First, note that: to estimate $Z_G(\lambda)$ to $(1 \pm \epsilon)$ multiplicative error, it suffices to estimate marginals to $\pm O(\epsilon/n)$ additive error. This is by a similar telescoping reduction to in the first lecture. Namely, for any specified vertex r, we have $|Z_G| = \frac{|Z_{G-\{r\}}|}{|Z_G|} \cdot |Z_{G-\{r\}}|$. Then the ratio $\frac{|Z_{G-\{r\}}|}{|Z_G|}$ is the probability that r is not in an independent set (which is its marginal), and the second term $|Z_{G-\{r\}}|$ can be estimated by recursion.
2. To estimate $p_{G,r}$ we run a recursion to depth $L = O(1/\delta \cdot \log(n/\epsilon))$.
3. If we hit an instances (H, u) at this depth, then we return an arbitrary estimate $\widetilde{p}_{H,u}$.

That this works follows immediately from strong spatial mixing (we get contraction at a rate of $(1 - O(\delta))$, and the choice of L ensures we get additive approximation to $\pm\epsilon/n$).

The running time is $\Delta^{\log(n/\epsilon)/\delta} = (n/\epsilon)^{\log(\Delta)/\delta}$, since we branch by a factor of Δ at each step. □

Note that this contrasts to belief propagation (run in practice by physicists), which has no guarantee of correctness in general.

7.3 Finishing contraction

Finally, we prove Lemma 7.2.

Proof of Lemma 7.2. Remembr that $f_d(p) = \frac{\lambda(1-p)^d}{1+\lambda(1-p)^d}$. We have $g_d = \varphi \circ f_d \circ \varphi^{-1}$. We want to show that for all $m = \varphi(p)$, $|g_d'(m)| \le \sqrt{f_d'(\hat{p})}$, where $\hat{p} = f_d(\hat{p})$ is the fixed point of f.

We have that $|g_d'(m)| = d\sqrt{f_d(p)} \cdot \sqrt{p}$, where p is defined by $m = \varphi(p)$.

Moreover, by direct computation we have $f_d'(\hat{p}) = -d \cdot \frac{1 - f_d(\hat{p})}{1 - \hat{p}} \cdot f_d(\hat{p}) = -d \cdot \hat{p}$, using that $\hat{p}$ is a fixed point of f_d. So, what we want to prove is equivalent to: for all $p \in [0, 1]$, $d \cdot f_d(p) \cdot p \le \hat{p}$. Then we can complete the lemma by doing a bit more analysis of the fixed point here; see notes. □

8 October 3, 2023

Recall that the permanent of a matrix is defined by $\operatorname{per}(A) = \sum_{\sigma \in S_n} \prod_{i \in [n]} A_{i,\sigma(i)}$, for a complex-valued matrix A. Our main goal is to prove the following theorem:

Theorem 8.1 (Barvinok)**.** *Assume that $\|A - \mathbf{1}\mathbf{1}^\top\|_\infty \le \eta$ for $\eta < 1/2$.*[1] *Then there is a polynomial T_m of degree $m \le O_\eta(\log(n/\epsilon))$ so that $|T_m(A) - \log \operatorname{per}(A)| \le \epsilon$. Furthermore, T_m can be computed in $2^{\operatorname{poly}\log n}$ time (i.e., quasipolynomial time).*

Note that the dependence of degree on η ends up being $O(1/(1-\eta))$.

8.1 Taylor series approach

In general, we expect that for a degree-m Taylor series, we have error $\le C_{\eta,m}$ if the function is relatively smooth (i.e., its derivatives are not too large).

Proposition 8.2. *Suppose $g : \mathbb{C} \to \mathbb{C}$ is a univariate polynomial of degree d. Assume that $g(z) \neq 0$ forall $z \in \mathbb{C}$ so that $|z| \le R$ for some $R > 1$. Write $f = \log g$ and define*

$$T_m(z) = \sum_{k=0}^{m} \frac{f^{(k)}(0)}{k!} \cdot z^k,$$

where $f^{(k)}$ is the kth derivative of f. Then for all $z \in D$ (where D denotes the unit disc in $\mathbb{C}$),

$$|f(z) - T_m(z)| \le \frac{d}{(m+1)R^m(R-1)}.$$

Note that ultimately we will take $g(z) = \operatorname{per}(zA + (1-z)\cdot \mathbf{1}\mathbf{1}^\top)$. Note also that the derivatives of g can be readily computed, so that the derivatives of $\log g$ can also be computed.

We now prove the above proposition:

Proof. Let $\zeta_1, \dots, \zeta_d$ be the roots of g. By assumption $|\zeta_i| > R$ for all i. We can write

$$g(z) = g(0) \cdot \prod_{i=1}^{d} (1 - \frac{z}{\zeta_i}).$$

Taking log, we get

$$f(z) = f(0) + \sum_{i=1}^{d} \log(1 - z/\zeta_i).$$

Since $\log(1-z)$ is holomorphic inside the unit disc, the Taylor series for $\log(1-z)$ converges inside D. In particular,

$$\left|\log(1 - z/\zeta_i) - \left(\sum_{k=1}^{m} \frac{1}{k} \cdot (z/\zeta_i)^k\right)\right| \le \left|\sum_{k=m+1}^{\infty} \frac{1}{k} \cdot (z/\zeta_i)^k\right| \le \frac{1}{m+1} \sum_{k=m+1}^{\infty} \frac{|z|^k}{|\zeta_i|^k} \le \frac{1}{m+1} \sum_{k=m+1}^{\infty} R^{-k} = \frac{1}{(m+1)R^m(R-1)}.$$

We have used that $|z| \le 1$ and $|\zeta_i| \ge R$. If we sum over these approximations for all i, then for any fixed m, the error is at most d times the quantity above, yielding the proof of the proposition. □

[1] This is ℓ_∞ norm as a vector

Lemma 8.3. *Given $\{g^{(k)}(0)\}_{k=0}^m$, we can compute $\{f^{(k)}(0)\}_{k=0}^m$ in time $O(m^2)$.*

Proof. We have $f = \log g$, so $f^{(1)} = g^{(1)}/g^{(0)}$, i.e., $g^{(1)} = g^{(0)} \cdot f^{(1)}$. Now we can differentiate this identity k times, and we get

$$g^{(k)} = \sum_{j=0}^{k-1} \binom{k-1}{j} f^{(k-j)} g^{(j)}.$$

In particular, given the derivatives of g and $f^{(0)}, \dots, f^{(k-1)}$, the above gives a formula for $f^{(k)}$ which can be computed in time $O(k)$. Then just use induction on k, which leads to $O(m^2)$ computation claim. □

To prove correctness of our algorithm, we need to control the roots of the polynomial that we're analyzing. In particular, Barvinok's theorem is:

Theorem 8.4 (Barvinok)**.** *There exists $\eta_0 > 1/2$ so that $\|A - \mathbf{1}\mathbf{1}^\top\|_\infty \le \eta_0$ implies that $\operatorname{per}(A) \neq 0$.*

First, we prove Theorem 8.1.

Proof of Theorem 8.1. We take $\eta_0 = 1/2$. Let $g(z) = \operatorname{per}(\mathbf{1}\mathbf{1}^\top + z(A - \mathbf{1}\mathbf{1}^\top))$. By Theorem 8.4, $g(z/2) \neq 0$ for all z so that $|z| \le 1/\eta_0 = 2$. We need to compute $\{g^{(k)}(0)\}_{k=0}^m$.[2] Let's write down what this derivative is:

$$\frac{d^k}{dz^k} \sum_{\sigma \in S_n} \prod_{i=1}^n (1 + z(A_{i,\sigma(i)} - 1)).$$

The inner product is

$$\sum_{S \subset [n]} \prod_{i \in S} (A_{i,\sigma(i)} - 1) \cdot z^{|S|}.$$

To compute the derivative at 0, we kill off all subsets $|S| < k$, and also those with $|S| > k$ (since there will be a term $z^{|S|-k}$). Thus,

$$g^{(k)}(0) = k! \sum_{\sigma} \sum_{|S|=k} \prod_{i \in S} (A_{i,\sigma(i)} - 1).$$

We claim that the above can be computed in $n^{O(k)}$ time. Let's reorder the summation: we claim that it is equal to

$$g^{(k)}(0) = k!(n-k)! \sum_{|S|=k} \sum_{\sigma : S \to [n] \text{ injective}} \prod_{i \in S} (A_{i,\sigma(i)} - 1).$$

In particular, we have used that for all $(n-k)!$ ways to extend σ outside of its values on S, the value of the product will remain the same. So we get a blowup of a factor $(n-k)!$. The above summation is now over $n^{O(k)}$ terms, so we can just brute-force compute it exactly.

Then by the previous lemma, we may compute the polynomial $T_m(z)$ in time $n^{O(k)}$, and by the first proposition, with $R = 2$, get get ϵ-approximation with degree $\log(n/\epsilon)$, as desired. □

[2] Technically we want the derivatives of the mapping $z \to g(z/2)$, but it suffices to show we can compute the derivatives of g by the chain rule.

8.2 Proof of Barvinok's zero-free theorem

We prove Theorem 8.4 for the rest of lecture. Note that permanent is a multilinear polynomial. If we let A_j denote the submatrix where we delete the 1st row and jth column, then we can write

$$\mathrm{per}(A) = \sum_{j=1}^{n} A_{1,j} \cdot \mathrm{per}(A_j).$$

Note that A_j submatrices are also close to the all-1s matrix. We want to say that all permanents in the above sum are nonzero complex numbers, and since $A_{1,j}$ are close to 1, when we sum them up we get something that is nonzero. The problem is that each of $\mathrm{per}(A_j)$ can be pointing in different directions and they all cancel each other out. Thus, we will actually show that the complex numbers $\mathrm{per}(A_j)$ all point in a similar direction.

In particular, the inductive hypothesis is the following statement $IH(n)$:

1. For all $A \in \mathbb{C}^{n\times n}$ so that $\|A - \mathbf{1}\mathbf{1}^\top\|_\infty \le 1/2$, $\mathrm{per}(A) \neq 0$.
2. For all pairs $A, B \in \mathbb{C}^{n\times n}$ which are $1/2$-close to the all-1s matrix and A, B differ on at most one row/column, then $\angle(\mathrm{per}(A), \mathrm{per}(B)) \le \alpha$. (We will ultimately take $\alpha = \pi/2$.)

The following is the main geometric lemma:

Lemma 8.5. *Suppose $\alpha \le 2\pi/3$. Let $u_1, \dots, u_n \in \mathbb{C}$ be nonzero and such that for all i, j, $\angle(u_i, u_j) \le \alpha$. Let $a_1, \dots, a_n \in \mathbb{C}$ satisfy $\|a - \mathbf{1}\|_\infty \le 1/2$. Then*

1. *$\sum_{i=1}^n a_i u_i \neq 0$.*
2. *$\angle\left(\sum_{i=1}^n a_i u_i, \sum_{i=1}^n u_i\right) \le \sin^{-1}\left(\frac{1/2}{\cos(\alpha/2)}\right)$. This expression is $\pi/4$ if we plug in $\alpha = \pi/2$.*

Think of $a_1, \dots, a_n$ as being the entries of the first row of A.

We first show how to use Lemma 8.5 to finish the proof of the zero-free theorem.

Proof of Theorem 8.4. We leave $IH(1)$ as an exercise. Now assume $IH(n-1)$ holds. We write $\mathrm{per}(A) = \sum_{i=1}^n A_{1,j}\mathrm{per}(A_j)$. We will apply Lemma 8.5 with $\{a_j\} = \{A_{1,j}\}$. Note that the matrices A_j differ, pairwise, by at most one column. Thus, by $IH(n-1)$, we have that, with $\{u_j\} = \{\mathrm{per}(A_j)\}$, $\angle(u_i, u_j) = \angle(\mathrm{per}(A_i), \mathrm{per}(A_j)) \le \pi/2$. The first item of the geometric lemma implies that $\mathrm{per}(A) = \sum_{i=1}^n a_i u_i \neq 0$, establishing the first part of $IH(n)$.

Next we need to prove the second part of $IH(n)$. So consider A, B which differ in a single row or column (say wlog they differ in the first row). The main point is to use the triangle inequality (note that the angle between complex vectors satisfies the triangle inequality). Then

$$\begin{aligned}
&\angle(\mathrm{per}(A), \mathrm{per}(B)) \\
=&\angle\left(\sum_j A_{1,j}\mathrm{per}(A_j), \sum_{j=1}^n B_{1,j}\mathrm{per}(B_j)\right) \\
\le&\angle\left(\sum_{j=1}^n A_{1,j}\mathrm{per}(A_j), \sum_{j=1}^n \mathrm{per}(A_j)\right) + \angle\left(\sum_{j=1}^n \mathrm{per}(A_j), \sum_{j=1}^n \mathrm{per}(B_j)\right) + \angle\left(\sum_{j=1}^n \mathrm{per}(B_j), \sum_{j=1}^n B_{1,j}\mathrm{per}(B_j)\right).
\end{aligned}$$

The second term in the final expression is 0 since $A_j = B_j$ for all j. The first and third terms are at most $\pi/4$ by the second statement of the geometric lemma with $\alpha = \pi/2$, applied to each of the sequences $\{u_j\} = \{\text{per}(A_j)\}$ and $\{u_j\} = \{\text{per}(B_j)\}$. This completes the induction. □

Proof of Lemma 8.5. Write $u = \sum_{i=1}^n u_i$, $v = \sum_{i=1}^n a_i u_i$.

First, we claim that $|u| \geq \cos(\alpha/2) \sum_{i=1}^n |u_i|$. (This is a sort of reverse triangle inequality.) The idea behind the proof of this claim is that there is some cone with internal angle α so that all u_i are in this cone. Then we project all u_i onto the bisector of this cone. This argument crucially relies on the fact that $\alpha \leq 2\pi/3$, as otherwise can take $1, e^{2\pi i/3}, e^{4\pi i/3}$, which do not all lie in a convex cone of size $2\pi/3$.

Given the above claim, $|u - v| = |\sum_{i=1}^n (a_i - 1) \cdot u_i| \leq \|a - \mathbf{1}\|_\infty \cdot \sum_{i=1}^n |u_i| \leq 1/2 \sum_{i=1}^n |u_i| < \cos(\alpha/2) \sum_i |u_i| \leq |u|$, where we have taken $\alpha = \pi/2$. Thus, $|v| \geq -|u - v| + u| > 0$, and hence $v \neq 0$, proving the first claim.

To prove the second claim, consider the vectors $u, v, u - v$. We know that $|u - v| < |u|$. We want to bound $\angle(u, v)$, which we denote by θ. Note that $\sin\theta \leq \frac{|u-v|}{|u|} \leq \frac{1/2}{\cos \pi/4}$ (can see the first inequality by drawing a picture, and the second inequality since $|u-v| \leq 1/2 \sum_i |u_i| \leq \frac{1/2}{\cos(\pi/4)} \cdot |u|$). Thus, taking $\sin^{-1}$, we get $\theta \leq \sin^{-1}\left(\frac{1/2}{\cos(\pi/4)}\right)$. □

9 October 5, 2023

Given a graph $G = (V, E)$, a *matching* is a subset $M \subset E$ so that no vertex is incident to more than one edge of the graph (not necessarily a perfect matching). Today we discuss:

Theorem 9.1 (Heilman-Lieb,Patel-Rogts)**.** *There is a deterministic algorithm to count all matchings up to $(1 \pm \epsilon)$-multiplicative error in time n^C, where $C = C(\Delta)$ is a function of the max degree.*

We use similar analytic techniques to last time. In particular, define the following matching polynomial:

$$\mathcal{M}_G(\lambda) = \sum_{M \subset E, \text{ matching}} \lambda^{|M|} = \sum_{k=0}^{\lfloor n/2 \rfloor} m_k \lambda^k,$$

where m_k is the number of k-matchings.

Theorem 9.2 (HL)**.** *$\mathcal{M}_G$ has negative real roots, and if G has max-degree Δ, then each root r satisfies $r \leq -\frac{1}{4(\Delta-1)}$.*

In particular, the region to the right of the line $\Re(z) = -1/(4(\Delta - 1))$ is zero-free. We can use the technique of last time to estimate the matching polynomial in this region in quasipolynomial time, but today we will in fact get a poly-time algorithm.

Our strategy is similar to the lecture on correlation decay: we will reduce to the tree case. Recall that for $r \in G$, we defined $T = T_{SAW}(G, r)$:

- The vertices of T correspond to self-avoiding walks $r = u_0 \to \cdots \to u_k$.
- Twowalks are adjacent if one extends the other.

Theorem 9.3 (Godsil-Gutman)**.** *$\mathcal{M}_G$ divides $\mathcal{M}_T$, where $T = T_{SAW}(G, r)$, for any choice of the root vertex r.*

In particular, it suffices to get control on the roots of $\mathcal{M}_T$, since the roots of $\mathcal{M}_G$ are a subset.

Proof of Theorem 9.3. We will prove the following more precise statement:

$$\frac{\mathcal{M}_{G-r}(\lambda)}{\mathcal{M}_G(\lambda)} = \frac{\mathcal{M}_{T-r}(\lambda)}{\mathcal{M}_T(\lambda)}. \tag{2}$$

By rearranging, this shows that $\mathcal{M}_G$ divides $\mathcal{M}_T$. In particular, we get that

$$\mathcal{M}_T = \mathcal{M}_G \cdot \frac{\mathcal{M}_{T-r}}{\mathcal{M}_{G-r}}.$$

By induction, we have that $\mathcal{M}_{G-r} \mid \mathcal{M}_{T-r}$ (in particular, note the T_{SAW} of $G - r$ with root r is equal to $\mathcal{M}_{T-r}$ since the walks have to be self-avoiding).

To prove (2), the intuition is that $\mathcal{M}_{G-r}/\mathcal{M}_G$ is the probability over a random matching (where matching M is given weight $\lambda^{|M|}$) that r is not included, i.e., $\Pr_M(r \text{ not included})$. The proof will be similar to in the lecture on correlation decay (use a tree recursion).[3] In particular, we have that

$$\mathcal{M}_G(\lambda) = \mathcal{M}_{G-r}(\lambda) + \lambda \sum_{v \sim r} \mathcal{M}_{G-r-v}(\lambda). \tag{3}$$

Rearranging, we get

$$\frac{\mathcal{M}_G(\lambda)}{\mathcal{M}_{G-r}(\lambda)} = 1 + \lambda \sum_{v \sim r} \frac{\mathcal{M}_{G-r-v}(\lambda)}{\mathcal{M}_{G-r}(\lambda)} = 1 + \lambda \sum_{v \sim R} \frac{\mathcal{M}_{T_{SAW}(G-r,v)-v}(\lambda)}{\mathcal{M}_{T_{SAW}(G-r,v)}(\lambda)},$$

where the second equality uses induction with v being the root node. Now the key claim is that the subtree of $T_{SAW}(G, r)$ rooted at the node corresponding to v is exactly $T_{SAW}(G - r, v)$. Thus, the above expression is equal to

$$1 + \lambda \sum_{v \sim r} \frac{\mathcal{M}_{T_{SAW}(G,r)-r-v}(\lambda)}{\mathcal{M}_{T_{SAW}(G,r)-r}(\lambda)} = \frac{\mathcal{M}_T(\lambda)}{\mathcal{M}_{T-r}(\lambda)}, \tag{4}$$

Note that we have used the fact that since we removed r, the subtrees given by its children are all disjoint, so the matching polynomial is a product, and thus for all subtrees which don't contain v, we get cancellations on the LHS of (4). Note that to show the equality in (4) we used (3).

This completes the inductive step. □

So, it suffices to prove the HL result for trees. We define

$$\widetilde{\mathcal{M}}_G(z) = \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k m_k z^{n-2k} = z^n \mathcal{M}_G(-1/z^2).$$

Lemma 9.4. *If G is a tree, then $\widetilde{\mathcal{M}}_G(z) = \det(zI - A_G)$.*

[3] We can't use exactly the same recursion from a few lectures ago (tree recursion for independent sets) since that loses the special structure of matchings.

This will givve what we want: since A_G is symmetric, its eigenvalues are real and moreover can be bounded by trace moment method.

Proof. We have

$$\det(zI - A) = \sum_{\sigma: V \to V} (-1)^{sign(\sigma)} \prod_{v \in V} (zI - A)_{v,\sigma(v)}.$$

Now let's pick out vertices v so that σ maps them to itself. That gives us powers of z. In particular, the above is equal to

$$\sum_{k=0}^{n} z^{n-k} (-1)^k \sum_{S \in \binom{V}{k}} \det(A_{S,S}),$$

where $A_{S,S}$ denotes the $S \times S$ submatrix corresponding to $S \subset V$. We claim that $(-1)^{k/2} \det(A_{S \times S})$ corresponds to the number of perfect matchings in the induced subgraph $G[S]$. This will imply that $\sum_{S \in \binom{V}{k}} \det(A_{S \times S}) = (-1)^{k/2} m_k$.

Since $G[S]$ is a tree, it suffices to prove WLOG that $\det(A) = (-1)^{n/2} m_n$. So write $\det(A) = \sum_{\sigma: V \to V} (-1)^{sign(\sigma)} \prod_v A_{v,\sigma(v)}$. We want to show that only permutations corresponding to matchings have a nonzero contribution to this sum, i.e., only if $(v, \sigma(v)) \in E$ for all $v \in V$.

We also know that all permutations σ have a cycle decomposition. We want that each of the successive pair of edges in such a cycle corresponds to an edge of G. But G is a tree and so has no cycles, and so we have a nonzero contribution if and only if the cycle decomposition of σ is a decomposition into pairs of vertices which correspond to edges of G. In particular, we must have $\sigma(\sigma(v)) = v$ for all v. Such a decomposition clearly corresponds to a matching. □

The above lemma implies that $\widetilde{\mathcal{M}}_G$ has real roots. Then clearly by definition, so does $\mathcal{M}_G$. So we just need to prove a bound on the roots. We use the following:

Lemma 9.5. *If T is a tree of max-degree Δ, then all eigenvalues of A_T are $\leq 2\sqrt{\Delta - 1}$.*

Proof. We look at $\lim_{k \to \infty} \mathrm{Tr}(A_T^k)^{1/k} = \lambda_{\max}(A_t)$. Now, $A_T^k(u, u)$ is the number of walks in the tree that start at u and return to u. Roughly speaking, the number of walks of length k scales as $\Delta^{k/2}$, since after walking downwards for $k/2$ steps we need to come back up. We also use the fact that trees are bipartite, so all eigenvalues are symmetric about 0, and thus we get a bound on $\lambda_{\min}$ as well. □

By the above lemma, all roots of $\widetilde{\mathcal{M}}_G(z)$ are at most $2\sqrt{\Delta - 1}$. Thus using the definition of $\widetilde{\mathcal{M}}_G$ in terms of $\mathcal{M}_G$, we get that all roots of $\mathcal{M}_G$ are bounded above by $(2\sqrt{\Delta - 1})^{-2} = 1/(4(\Delta - 1))$.

9.1 Using the Heilman-Lieb theorem

Recall from last time that, for $f(\lambda) = \log \mathcal{M}_G(\lambda)$, to estimate f in a zero-free region containing 0 we need to compute $\{f^{(k)}(0)\}_{k=0}^m$ for m sufficiently large. Last time, we showed that we can do this if we know the low-order Taylor coefficients of $\mathcal{M}_G$, i.e., $\{\mathcal{M}_G^{(k)}\}_{k=0}^m = \{m_k\}_{k=0}^m$. Last time, we brute-forced this, which took $n^{O(m)}$ time – this gave quasipolynomial time, by taking $m = O(\log n)$. Today, we do this more intelligently, using a combinatorial interpretation of $f^{(k)}(0)$.

Let $r_1, \ldots, r_{n/2}$ denote the roots of $\mathcal{M}_G$. Then we can write $\mathcal{M}_G(z) = \prod_{i=1}^{n/2}(1 - z/r_i)$. There is no coefficient in front since $\mathcal{M}_G(0) = 1$. As we saw last time,

$$f^{(k)}(0)/k! = \frac{1}{k}\sum_{i=1}^{\lfloor n/2 \rfloor} r_i^{-k}. \tag{5}$$

You get this by taking log of both sides and looking at the Taylor series of $\log(1 - z/r_i)$.

Lemma 9.6. *Let $T(r) = T_{SAW}(G, r)$ for each r. Then for all k,*

$$\sum_{i=1}^{\lfloor n/2 \rfloor} |r_i|^{-k} = \sum_{r \in V} (A_{T(r)}^{2k})(r, r).$$

The above lemma is quite remarkable: not even clear that the LHS is an integer! This actually implies a $\Delta^{O(k)}$ algorithm for exactly computing $f^{(k)}(0)$:

Proof of Theorem 9.1. We just walk along all length-$2k$ paths starting from r in $T(r)$ to compute $\sum_r (A_{T(r)}^{2k})(r, r)$, which takes $\Delta^{O(k)}$ time. Taking $k = \log(n/\epsilon)$, this gives a $\Delta^{O(\log n/\epsilon)} = (n/\epsilon)^{C_\Delta}$-time algorithm to implement the Barvniok tecnique from last time: in particular, by Lemma 9.6 and (5), we can compute $f^{(k)}(0)$ for all $k \leq O(\log n/\epsilon)$, in the stated time, which then gives us the desired polynomial approximation by the standard Taylor series bounds discussed last time. □

Finally we prove Lemma 9.6.

Proof of Lemma 9.6. The proof strategy is as follows: we will use $\widetilde{\mathcal{M}}_G(z) = z^n \mathcal{M}_G(-1/z^2) = z^{n-2\lfloor n/2 \rfloor} \prod_{i=1}^{\lfloor n/2 \rfloor}(z - \sqrt{|r_i|})(z + \sqrt{|r_i|})$. Note that the $-k$th moments of $\mathcal{M}_G$ (i.e., $\sum_{i=1}^{\lfloor n/2 \rfloor} |r_i|^{-k}$) are equal to the $2k$th moments of $\widetilde{\mathcal{M}}_G$.

Look at $(\log \widetilde{\mathcal{M}}_G(z))' = \frac{\widetilde{\mathcal{M}}_G'(z)}{\widetilde{\mathcal{M}}_G(z)}$. Expanding this out, we have:

$$\frac{\widetilde{\mathcal{M}}_G'(z)}{\widetilde{\mathcal{M}}_G(z)} = z^{-1} \sum_{i=1}^{\lfloor n/2 \rfloor} \frac{1}{1 - 1/(|r_i| z^2)} = z^{-1} \sum_{k=0}^{\infty} z^{-2k} \sum_{i=1}^{\lfloor n/2 \rfloor} |r_i|^{-k}. \tag{6}$$

We now need two lemmas.

Lemma 9.7. $\widetilde{\mathcal{M}}_G'(z) = \sum_{r \in V} \widetilde{\mathcal{M}}_{G-r}(z)$.

Proof. Double-counting argument; left as exercise. □

The above lemma implies that

$$\frac{\widetilde{\mathcal{M}}_G'(z)}{\widetilde{\mathcal{M}}_G(z)} = \sum_r \frac{\widetilde{\mathcal{M}}_{G-r}(z)}{\widetilde{\mathcal{M}}_G(z)} = z \cdot \sum_{r \in V} \frac{\widetilde{\mathcal{M}}_{T(r)-r}(z)}{\widetilde{\mathcal{M}}_{T(r)}(z)}, \tag{7}$$

where the second equality uses (2), together with the fact that the extra factors of $z^{|T(r)|-1}$ cancel out on the RHS.

Next, we need:

Lemma 9.8. *For all trees T, for all $r \in T$,*

$$z^{-1}\frac{\widetilde{\mathcal{M}}_{T-r}(z^{-1})}{\widetilde{\mathcal{M}}_T(z^{-1})} = \sum_{k=0}^{\infty} z^{2k}(A_T)^{2k}(r,r).$$

Proof. We saw before that the LHS of the lemma statement is equal to

$$z^{-1} \cdot \frac{\det(z^{-1}I - A_{T-r})}{\det(z^{-1}I - A_T)}.$$

Since T is a tree, there are no odd-length walks that go back to their starting point, and so RHS is equal to

$$\sum_{k=0}^{\infty} z^k(A_T^k)(r,r) = (I - zA_T)^{-1}(r,r).$$

Next, for a matrix B, we can define its adjugate, $adj(B)$, to satisfy $B \cdot adj(B) = \det(B) \cdot I$. By Cramer's rule, we have that

$$adj(B)_{ij} = \det(B^{-i,-j}),$$

where $B^{-i,-j}$ is the submatrix of B after deleting row i and column j. Now we take $B = z^{-1}I \cdot A_T$. Then

$$(I - zA)^{-1} = z^{-1}(z^{-1}I - A)^{-1} = \frac{z^{-1}}{\det(z^{-1}I - A)} \cdot adj(z^{-1}I - A).$$

Now we take the (r,r)th entry of both sides. The (r,r) entry of $adj(z^{-1}I - A)$ is equal to $\det(z^{-1}I - A_{T-r})$ by Cramer's rule, which must be equal to the (r,r) entry of $(I - zA_T)^{-1}$, and thus is equal to the RHS Of the lemma statement. □

The proof of Lemma 9.6 is concluded by combining (6) with (7) and Lemma 9.8. □

10 October 12, 2023

Definition 10.1 (Stability)**.** Let $p(z_1, \ldots, z_n)$ be a polynomial. Let $T_1, \ldots, T_n \subset \mathbb{C}$. We say that p is $\prod_i T_i$-stable if $p(z) \neq 0$ whenever $z_i \in T_i$ for all i. If $T_i = T$ for all i, then we say that p is T-stable.

Recall the Ising model:

$$Z_{G,\beta}(\lambda) = \sum_{\sigma: V \to \{\pm 1\}} \exp\left(\frac{\beta}{2}\sigma^\top A_G \sigma\right) \cdot \prod_{v:\sigma(v)=1} \lambda_v,$$

where λ denotes the external field. Let $S = \{v : \sigma(v) = 1\}$, so that $\sigma^\top A_G \sigma = |E(S,S)| + |E(V - S, V - S)| - 2|E(S, V - S)|$. Then

$$Z_{G,\beta}(\lambda) \propto \sum_{S \subseteq V} \exp\left(-2\beta|E(S, V - S)|\right) \cdot \lambda^S,$$

where $\lambda^S = \prod_{v \in S} \lambda_v$.

We have the following theorem:

Theorem 10.1 (Lee-Yang, '52)**.** *For all G and $\beta \geq 0$, $Z_{G,\beta}(\lambda)$ is D-stable and $\bar{D}^c$-stable, where D denotes the unit disk. Moreover, all zeros of $Z_{G,\beta}(t \cdot \mathbf{1})$ lie on the unit circle.*

By Barvinok's algorithm, this gives us a quasipolynomial-time algorithm for estimating the partition function, as long as λ is bounded slightly away from 1. What was the original motivation? Lee-Yang were looking for phase transitions. In particular, if you let $F_\beta(\lambda) = \lim_{n\to\infty} \frac{1}{n} \log Z_{G_n,\beta}$, then discontinuities of $F_\beta(\lambda)$ correspond to a phase transition. In particular, the Lee-Yang theorem implies that if $|\lambda - 1| > 0$, then there is no phase transition.

To prove the theorem, we use induction: the base case is an edge, i.e., a graph with two vertices and an edge between them.

Lemma 10.2. *let $a \in \mathbb{C}$ with $|a| \leq 1$. Then the bivariate polynomial $1+az_1+\bar{a}z_2+z_1z_2$ is D-stable.*

Note that the Tutte polynomial $Z_{G,\beta}$ is a special case of the above, letting $a = \exp(-2\beta)$, for $\beta \geq 0$ (so that $a \leq 1$).

Proof. Suppose that $|a| = 1$. Then $p(z) = (1 + az_1)(1 + \bar{z}z_2)$. Now clearly each of the factors is D-stable, since $|a| = 1$.

Now suppose $|a| < 1$. For all z_2, there is a unique z_1 so that $p(z_1, z_2) = 0$. In particular,

$$z_1 = f(z_2) = -\frac{1 + \bar{a}z_2}{a + z_2}.$$

It is now easy to check that $f(D) \subset \bar{D}^c$. f is a Mobius transformation, very well studied; here's one way to prove this: easy to check that $f^{-1}(z_1) = -\frac{1+az_1}{\bar{a}+z_1}$. We claim that $f(\partial D) = \partial D$. Since f maps 0 outside of the circle, then by continuity, f maps the whole interior of the unit circle to the exterior. To check that f maps ∂D to ∂D, we simply need to check that when $|z| = 1$, $|1 + \bar{a}z| = |\bar{a} + z|$: $|1 + \bar{a}z| = |\bar{z} + \bar{a}| = |a + z|$, where the first equality is multiplying by $|\bar{z}|$. □

To prove the inductive step, we need the following theorem:

Theorem 10.3. *Let $p(z) = \sum_{S\subset[n]} a_S z^S$ and $q(z) = \prod_{S\subset[n]} b_S z^S$ be D-stable polynomials. Then $(p * q)(z) = \sum_{S\subset[n]} a_S b_S z^S$ is also D-stable.*

We first show that Theorem 10.3 implies the Lee-Yang theorem.

Proof of Theorem 10.1. For all $e = (u, v) \in E$, let

$$p_e(z) = (1 + \exp(-2\beta)z_u + \exp(-2\beta)z_v + z_uz_v) \cdot \prod_{w\neq u,v} (1 + z_w).$$

Observe:

$$Z_{\beta,G}(\lambda) = *_{e\in E} p_e.$$

If we know that each p_e is D-stable, then by induction, the whole thing is D-stable. Moreover, each p_e is D-stable since the part of it not involving the product over w is D-stable. □

To prove Theorem 10.3, we use the Assano-Raelle lemma:

Lemma 10.4 (Asano-Raelle)**.** *Let $p(z_1, z_2) = a + bz_1 + cz_2 + dz_1z_2$ be D-stable. Then so is $q(z) = a + dz$.*

The idea is we're contracting out (removing) terms that only involve z_1 or z_2, but not both.

Proof. We can assume that a is nonzero (or else p is not D-stable). Note that q has a unique root $-a/d$. We need $|a| \geq |d|$. WLOG, we can assume $|b| \geq |c|$.

Suppose $|a| < |d|$ for contradiction. If $p(z) = 0$, then $z_1 = -\frac{a+cz_2}{b+dz_2}$. We want there to exist $z_2 \in D$ so that $z_1 \in D$ aswell, i.e., $|a + cz_2| < |b + dz_2|$. We know $|b| \geq |c|, |d| > |a|$. Hence $|b| + |d| > |a| + |c| \geq |a + cz_2|$ for all $z_2 \in D$. We just need $z_2 \in D$ so that $|b + dz_2| \approx |b| + |d|$. We can make them arbitrarily close: in particular, first choose z_2 so that dz_2 has the same direction as b, and then take $|z_2| \to 1$. □

Now we prove Theorem 10.3.

Proof of Theorem 10.3. Let $r(x_1, \ldots, x_n, y_1, \ldots, y_n) = p(x) \cdot q(y)$. We apply the Asano-Ruelle contractions to the tuples of variables $\{(x_i, y_i)\}_{i=1}^n$. In particular, for each i, we drop monomials for which only x_i or y_i (but not both) occur in the monomial. By an inductive application of Asano-Ruelle, at each step along these contractions, we are always preserving D-stability. At the end of this, we get the polynomial $\sum_{S \subset [n]} a_S b_S z^S$. (Technically, the polynomial is $\sum_S a_S b_S (x_i y_i)^n$, whose D-stability implies the D-stability of $\sum_S a_S b_S z^S$.)

Note that at each step of the contraction, we fix the values of all other variables (to something arbitrary in the unit circle), that "restricted" polynomial is still D-stable. Then we apply Asano-Ruelle, and since the fixing is arbitrary, we get from Asano-Ruelle that the contracted polynomial is D-stable. □

10.1 Independence polynomials

Now we study $Z_G(z) = \sum_{I \subset V \text{indep}} z^I$, where the sum is over independent sets I.

Theorem 10.5 (Scott-Sokal '05)**.** *Let $p \in \mathbb{R}^V_{\geq 0}$. TFAE:*

- *Z_G is stable wrt $\prod_{v \in V} \overline{D(0, p_v)}$. (Think of p_v as something that depends on the degree of v.)*
- *For all $S \subset V$, $Z_{G[S]}(-p) > 0$.*

Note that Shearer's condition is the sharpest conition under which we get zero-freeness (since the above is an equivalence). Recall that $G[S]$ is the induced subgraph on S and $D(0, p_v)$ is the disk with center 0 and radius p_v. Note also that Sherer's condition is the sharpest condition under which the Lovasz Local Lemma holds. Note that the above theorem tells us that we only have to worry about the negative real axis wrt avoiding zeros.

Proof of Theorem 10.5. Clearly, Z_G is stable with respect to $\prod_{v \in V} \overline{D(0, p_v)}$ implies that for all $S \subset V$, $Z_{G[S]}(z_S) = Z_G(y)$, where $y_v = z_v$ if $v \in S$, and 0 otherwise satisfies that $Z_{G[S]}(-p) \neq 0$. It is in fact positive since it is positive at the point 0, and by continuity it can never be 0 when you move between 0 and $-p_v$.

Next we prove that Shearer's condition implies stability. It suffices to prove that, for all $v \in S$ and all $S \subset V$ and all $z \in \prod_{v \in V} \overline{D(0, p_v)}$,

$$\frac{|Z_{G[S]}(z)|}{|Z_{G[S-v]}(z)|} \geq \frac{Z_{G[S]}(-p)}{Z_{G[S-v]}(-p)}. \tag{8}$$

We now decompose

$$Z_G(z) = Z_{G-v}(z) + z_v \cdot Z_{G-N[v]}(z),$$

as we have done before. For simplicity we assume that $S = V$. By triangle inequality, we have

$$\frac{|Z_G(z)|}{|Z_{G-v}(z)|} \geq 1 - |z_v| \cdot \frac{|Z_{G-N[v]}(z)|}{|Z_{G-v}(z)|}.$$

We have $|z_v| \leq p_v$. Moreover, for $z = -p$, the above is actually an equality!

So, it suffices to prove that

$$\frac{Z_{G-N[v]}(-p)}{Z_{G-v}(-p)} \geq \frac{|Z_{G-N[v]}(z)|}{|Z_{G-v}(z)|}.$$

To prove this, we use telescoping (same trick as for correlation decay). We order $N(v) = u_1, \ldots, u_k$. In particular, the left-hand side above is

$$\prod_{i=1}^{k} \frac{Z_{G-v-u_1-\cdots-u_i}(-p)}{Z_{G-v-u_1-\cdots-u_{i-1}(-p)}}.$$

Each of these ratios is lower bounded by the corresponding ratio for the corresponding smaller graph $G - v - u_1 - \cdots - u_{i-1}$ (using induction, i.e, the reciprocal of (8)).

Note that we also need to do the base case: $Z_{\{v\}}(z) = 1 + z \geq 1 - p_v$. Note that $Z_\emptyset(z) = 1$, and this is immediate. □

When is Shearer's condition satisfied?

Theorem 10.6 (Dobrushin)**.** *If there is $y \in \mathbb{R}^V_{\geq 0}$ so that $p_v \leq \frac{y_v}{\prod_{u \in N[v]}(1+y_u)}$ for all $v \in V$, then Shearer's condition is satisfied.*

We don't prove this, but give an example: suppose that G has max-degree Δ. Take $y = \mathbf{1}/\Delta$. Then the univariate independent polynomial $Z_G(\lambda)$ is zero-free in $D(0, \Delta^\Delta \cdot \Delta^{-\Delta-1})$.

Now we have that $\Delta^\Delta/(\Delta+1)^{\Delta+1} \approx 1/(e(\Delta+1))$. the optimal threshold is $(\Delta-1)^{\Delta-1}/\Delta^\Delta$: this is attained in the limit for $(\Delta-1)$-ary trees.

Recall that the critical threshold for correlation decay is $\lambda_c(\Delta) = (\Delta-2)^\Delta/(\Delta-1)^{\Delta-1} \approx e/(\Delta-1)$. This is larger than what we have above by a factor of e^2 roughly. The zero-free region of the independencee polynomial is asymmetric about the imginary axis. It is still an open problem to characterize its full shape!

Here's another (easier) condition:

Theorem 10.7 (Kotecky-Preiss)**.** *Fix $p \in \mathbb{R}^V_{\geq 0}$. Suppose there is $a \in \mathbb{R}^V_{\geq 0}$ so that $\sum_{u \in N[v]} p_u \cdot e^{a(u)} \leq a(v)$ for all $v \in V$. Then Shearer's condition holds.*

The above requires only checking something on the neighborhood of each vertex. To prove the above theorem:

Proof. Plug in $y_v = p_v \cdot e^{a(v)}$ into Dobrushin. □