The Geometry of Polynomials in Algorithms

Lecture 2: The Matching Polynomial and Heilmann-Lieb

Lecturer: Nathan Klein

1 Recap

Last time we showed several key statements:

1. **A Classification of Real Rooted Generating Polynomials.** A $p \in \mathbb{R}_{\geq 0}[t]$ with $p(1) = 1$ is the generating polynomial of a Poisson binomial distribution μ if and only if p is real rooted. A student made a nice comment in class, which is that this implies that the generating polynomial of any distribution which is *not* Poisson Binomial must have a nonreal root.

   Also, it's good to notice that any non-zero $p \in \mathbb{R}_{\geq 0}[t]$ with $p(1) \neq 1$ still naturally maps to a generating polynomial, namely to $p(t)/p(1)$. In the same way, any real-rooted polynomial with non-negative coefficients corresponds to a Poisson Binomial after scaling by $1/p(1)$.

2. **Closure Properties of Real Rooted Polynomials.** We showed that real rooted polynomials are closed under certain operations like differentiation: if p is real rooted, so is p' (so long as p, p' are not the 0 polynomial).

3. **Newton's Inequalities.** We used the closure properties to show that the sequence of coefficients $a_0, \ldots, a_d$ of any real rooted polynomial $p \in \mathbb{R}_{\geq 0}[t]$ is ultra-log concave, i.e., for all $1 \leq i \leq d-1$,

$$\left(\frac{a_i}{\binom{d}{i}}\right)^2 \geq \frac{a_{i-1}}{\binom{d}{i-1}} \frac{a_{i+1}}{\binom{d}{i+1}}$$

   we also showed that this implied that the sequence $a_0, \ldots, a_d$ was unimodal.

1.1 Concentration

A nice consequence of this is that for any distribution μ with a real-rooted generating polynomial, the output of μ is concentrated around its expectation. In particular, we can apply Chernoff bounds, since Poisson Binomial distributions are sums of independent Bernoullis:

Theorem 1.1 (Chernoff Lower Tail). *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables. Then if $X = \sum_{i=1}^n X_i$ and $0 \leq \delta \leq 1$, we have:*

$$\mathbb{P}\left[X \leq (1-\delta)\mathbb{E}\left[X\right]\right] \leq \exp(-\mathbb{E}\left[X\right] \delta^2/2)$$

Theorem 1.2 (Chernoff Upper Tails). *Let $X_1, \ldots, X_n$ be independent Bernoulli random variables. Let $X = \sum_{i=1}^n X_i$. For all $\delta > 0$, we have*

$$\mathbb{P}\left[X \geq (1+\delta)\mathbb{E}\left[X\right]\right] \leq \exp(-\mathbb{E}\left[X\right] \delta^2/(2+\delta)).$$

So we should think of the coefficients of a real rooted polynomial is tightly concentrated around its center of mass.

2 Counting Matchings

In this lecture, we'll examine a distribution μ which outputs the number of edges in a uniformly random matching of a graph. Recall a matching is a subset of the edges of a graph in which every vertex is adjacent to at most one edge. It will not be at all clear why μ should be a Poisson Binomial, but we will be able to prove that it is by going through its generating polynomial.

So suppose we have a graph $G = (V, E)$ with $|V| = n$. Let m_i denote the number of matchings that contain i edges. We now want to understand the sequence $m_0, \ldots, m_{\lfloor n/2 \rfloor}$, and the following polynomial:

$$p(t) = \sum_{k=0}^{\lfloor n/2 \rfloor} m_k t^k$$

Where for simplicity we state the unscaled version: the distribution corresponds to $p(t)/p(1)$. We will prove the following:

Theorem 2.1 (Heilmann-Lieb Theorem)**.** *p is real rooted.*

Using the Chernoff bounds above, this will demonstrate that if you take a random matching in a graph, the number of edges it contains will be tightly concentrated around its expected value. This is an interesting fact on its own, but also has applications in physics.

2.1 A Problem in Physics

Imagine there is some system where atoms (or molecules) can either pair up with one other or remain unpaired. Certain atoms can pair with one another, and between these we have edges. This is called a monomer-dimer configuration in physics.

As we increase the temperature of a system, the atoms start moving around more, and bonds become broken more easily. So the higher the temperature the fewer the bonds we will see. The question we want to ask is: *do monomer-dimer systems exhibit a phase transition?* In other words, is there a temperature threshold which, once reached, changes the behavior of the system drastically? For example, the behavior of collections of water molecules has phase transitions at freezing and at boiling temperatures.

We will use an activity parameter λ which combines various properties of the system such as temperature and pressure. We can then model the probability of a particular configuration (matching) of size k as proportional to λ^k. So the probability of a matching M is:

$$\mathbb{P}\left[M\right] = \frac{\lambda^{|M|}}{\sum_{M' \in \mathcal{M}} \lambda^{|M'|}}$$

where $\mathcal{M}$ is the set of all matchings. Notice that when $\lambda \to \infty$, only the largest matchings will be sampled. As $\lambda \to 0$, the matching of size 0 will be sampled. If you're curious, a good approximation of *why* this is the model is as follows: suppose every bond occurs independently with probability $\frac{\lambda}{1+\lambda}$. Then, if we condition on the bonds forming a matching, we get the above expression. In particular:

$$\mathbb{P}\left[M\right] = \frac{(\frac{\lambda}{1+\lambda})^{|M|}(\frac{1}{1+\lambda})^{|E|-|M|}}{\sum_{M' \in \mathcal{M}}(\frac{\lambda}{1+\lambda})^{|M'|}(\frac{1}{1+\lambda})^{|E|-|M'|}} = \frac{(\frac{1}{1+\lambda})^{|E|}\lambda^{|M|}}{(\frac{1}{1+\lambda})^{|E|}\sum_{M' \in \mathcal{M}} \lambda^{|M'|}} = \frac{\lambda^{|M|}}{\sum_{M' \in \mathcal{M}} \lambda^{|M'|}}$$

What we are interested in is the number of bonds, which will be indicative of our phase transition. Where X is the random variable counting our number of bonds, this can be modeled as:

$$\mathbb{P}[X = k] = \frac{m_k \lambda^k}{\sum_{i=0}^{\lfloor n/2 \rfloor} m_i \lambda^i}$$

If this distribution is Poisson Binomial, then by Chernoff bounds, there is a small range of sizes k that dominate and have most of the mass for any choice of λ, as a λ scaling is just the polynomial $p(\lambda t)$ (which must also be real rooted by our closure properties). To move towards systems with matchings much larger than the expectation, we must increase λ by a lot so that those become appropriately highly weighted. Similarly to get small matchings, we must decrease λ by a lot. Intuitively, this means we will *not* see a sharp phase transition, although more work is needed to formalize this.

To understand the intuition better, it may help to see a non-example. Suppose that at $\lambda = 1$, with probability 99%, we see a matching of size $0.01n$, and otherwise we see a matching of size $0.05n$. (If our polynomial is real rooted, this is not possible, since we should get concentration around $0.01n$ and thus an exponentially small probability of seeing a matching 5 times larger than usual.) But now if I increase λ to $1 + \epsilon$, a sampled matching will suddenly have size $0.05n$ with probability close to 1 for even $\epsilon = O(1/n)$.[1] So this is a very sharp phase transition: an activity change which goes to 0 as n goes to infinity has the effect of changing the size of the matching by a factor of 5.

There is a lot more to learn about this, and diving deeper could be a nice project. For example, one thing physicists care about is that $\lim_{n\to\infty} \frac{1}{|V(G_n)|} \log p$ is analytic for a sequence $G_1, \ldots,$ of graphs, which is related to the fact that the zeros of p are all negative reals.

2.2 The Heilmann-Lieb Theorem: Proof of Real Rootedness

Given this motivation, let's prove that

$$p(t) = \sum_{k=0}^{\lfloor n/2 \rfloor} m_k t^k$$

is real rooted. We will actually prove something stronger: that the above works even if the graph is weighted, so that each edge has a weight $w_e \geq 0$ and:

$$m_k = \sum_{M \in \mathcal{M}: |M| = k} \prod_{e \in M} w_e$$

in this way, WLOG we may assume we have a complete graph, with some weights possibly equal to 0. We can then recover our claim on the original polynomial by setting $w_e = 1$ for edges in G and 0 otherwise.

The proof will work by induction. Notice that for any vertex v, we have:

$$p(t) = p_{G \setminus \{v\}}(t) + t \sum_{u \sim v} w_{\{u,v\}} p_{G \setminus \{u,v\}}(t)$$

[1]Formally: $\mathbb{P}[X = 0.05n \mid \lambda = 1 + \epsilon] = \frac{(1+\epsilon)^{0.05n} 0.01}{(1+\epsilon)^{0.01n} 0.99 + (1+\epsilon)^{0.05n} 0.01} = \frac{(1+\epsilon)^{0.04n} 0.01}{0.99 + (1+\epsilon)^{0.04n} 0.01}$. For example, choosing $1 + \epsilon = 99^{50/n}$, this probability is $\frac{99^2 \cdot 0.01}{0.99 + 99^2 \cdot 0.01} = 0.99$. And note $\epsilon = O(1/n)$.

The reason is simple: every matching either does not include v, in which case the first term applies, or it does, in which case it also includes some neighbor of v. In that case we increment the size of the matching by 1 and remove v and its neighbor.

Our proof strategy will be to show that if $p_{G\setminus\{v\}}$ and $p_{G\setminus\{u,v\}}$ are real rooted, so is the above polynomial. By induction, this will show $p_G(t)$ is real rooted for any graph G.

So, we need some strategy that says: if p, q are real rooted, so is some combination of p and q. In general, this is not true: for example, 1 is real rooted, and so is x^2, but x^2+1 is not. This motivates the concept of *interlacing*.

Definition 2.2 (Interlacing). *Let p be a degree d real rooted polynomial with roots $\lambda_1 \geq \lambda_2 \geq \dots, \geq \lambda_d$ listed with multiplicity.*

1. *If q has degree d with roots $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_d$ respectively, then q interlaces p if:*
$$\beta_d \leq \lambda_d \leq \beta_{d-1} \leq \cdots \leq \beta_1 \leq \lambda_1$$
2. *If q has degree $d-1$ with roots $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_{d-1}$, then q interlaces p if:*
$$\lambda_d \leq \beta_{d-1} \leq \cdots \leq \beta_1 \leq \lambda_1$$

If the inequalities are all strict in the above definitions, then we say q strictly interlaces p.

You might immediately notice that p' interlaces p by Rolle's theorem. We will come back to interlacing later when we study the Kadison-Singer problem, but for now let's see it in action in the proof that p is real rooted.

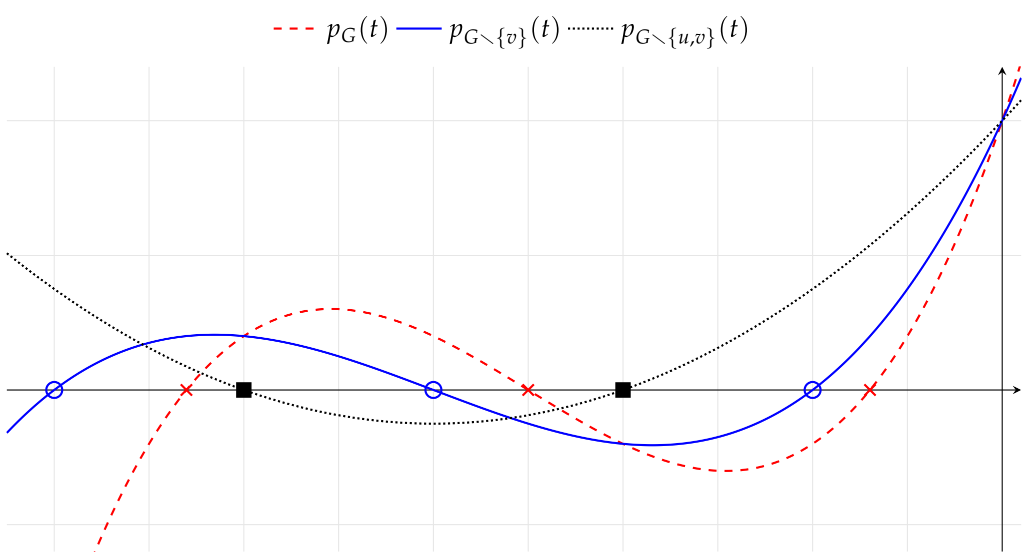

Figure 1: Illustration of the roots and signs when the degree of $p_{G\setminus\{v\}}$ is equal to the degree of p_G.

Proof. Let $G = (V, E)$ be our input graph. First we will assume all edges have non-zero weight, setting $w_e = \epsilon$ for all edges initially set to 0. We can then take the limit as $\epsilon \to 0$, and apply a theorem about limits of real rooted polynomials to finish the proof.

Our inductive hypothesis will be that for every complete graph G with positive edge weights:

1. All roots of p_G are negative and appear with multiplicity 1,

2. For each $v \in V$, $p_{G \setminus \{v\}}$ strictly interlaces p_G.

For the base case, we have $p(t) = 1$ for a single vertex and $p(t) = 1 + w_e t$ for two vertices with an edge e. These interlace, and the graph with a single edge has a negative root $-1/w_e$.

Now, given G with $|V| = n$, let v be an arbitrary vertex. By induction, for each $u \in V$ with $v \neq u$, $p_{G \setminus \{u,v\}}$ strictly interlaces $p_{G \setminus \{v\}}$. Recall our inductive statement:

$$p(t) = p_{G \setminus \{v\}}(t) + t \sum_{u \sim v} w_{\{u,v\}} p_{G \setminus \{u,v\}}(t)$$

Consider any $p_{G \setminus \{u,v\}}$ with roots $\beta_d < \beta_{d-1} < \dots \beta_1 < 0$. The roots index where the polynomial changes sign. Since $p_{G \setminus \{u,v\}}$ strictly interlaces $p_{G \setminus \{v\}}$, if $\lambda_{d'} < \lambda_{d'-1} < \dots \lambda_1$ are the roots of $p_{G \setminus \{v\}}$, it must be that $p_{G \setminus \{u,v\}}(\lambda_i)$ has sign $(-1)^{i+1}$: this is because $p_{G \setminus \{u,v\}}(0) = 1$, so $p_{G \setminus \{u,v\}}(\lambda_1)$ has positive sign.

However this is true for all u. Therefore, p itself has sign $(-1)^i$ at each root λ_i, since the contribution from $p_{G \setminus \{v\}}$ is 0, and we multiply the other polynomials by λ_i, which is negative. (Notice we are using that we have a complete graph here, as otherwise it is possible that p could be 0 at λ_i.) This gives us d' sign changes, since $p(0) > 0$. Note p has degree d' or $d' + 1$; if it has degree d' we are done. Otherwise, notice that the sign of p at $\lambda_{d'}$ is $(-1)^{d'}$. However, $\lim_{t \to -\infty} p(t)$ has sign $(-1)^{d'+1}$ because for p to have degree $d' + 1$, there must be a matching of size $d' + 1$.

This gives us that p is real rooted, and we can see that all roots are negative and have multiplicity 1. It remains to show that $p_{G \setminus \{v\}}$ strictly interlaces p for each v, but one can read this off of our construction of the roots of p (the same works for any choice of v).

As discussed, to finish the proof, we take the limit of the polynomial as we take $\epsilon \to 0$. This implies real rootedness of the polynomial attained in the limit due to the following lemma, as p is not the 0 polynomial. $\square$

Lemma 2.3. *Let $p_1, \dots \in \mathbb{R}[t]$ be a sequence of real rooted polynomials of bounded degree. If p_n converges to p coeficient-wise, then p is either the 0 polynomial or is real rooted.*

Proof. By coefficient-wise convergence and bounded degree, $\lim_{n \to \infty} p_n(z) = p(z)$ for every $z \in \mathbb{C}$. Suppose that $p(a + ib) = 0$ for $b \neq 0$. We will use this to show that $p(a + \alpha i b) = 0$ for every $\alpha \in [0, 1]$, which will imply that p has infinitely many roots and therefore must be the 0 polynomial.

First notice that for any $r \in \mathbb{R}$ and any $\alpha \in [0, 1]$, $|a - r + i\alpha b|^2 \leq |a - r + ib|^2$, as we are just making the complex part smaller. But this means that for each n,

$$|p_n(a + i\alpha b)| = |a_d \prod_{i=1}^{d} (a - \lambda_i + i\alpha b)| \leq |a_d \prod_{i=1}^{d} (a - \lambda_i + ib)| = |p_n(a + ib)|$$

where $\lambda_1, \dots, \lambda_d \in \mathbb{R}$ are the roots of p_n. Therefore:

$$\lim_{n \to \infty} |p_n(a + i\alpha b)| \leq \lim_{n \to \infty} |p_n(a + ib)| = |p(a + ib)| = 0$$

as desired. $\square$