Stanford University — **Spring 2018**
Math 233: Non-constructive methods in combinatorics — **Instructor: Jan Vondrák**
Lecture date: January 23, 2018 — **Original scribe: Erik Bates**

Lecture 5. Symmetric Shearer's Lemma

Here we discuss a corollary of Shearer's Lemma that considers the symmetric case, in which all events are given the same probability bound.

Theorem 5.1 (Symmetric Shearer's Lemma) *Suppose there is a collection of events* $\{\mathcal{E}_i\}_{i=1}^n$ *such that each* $\mathcal{E}_i$ *is independent of all but* d *other events* ($d \geq 2$), *and*

$$\mathbb{P}(\mathcal{E}_i) \leq \frac{(d-1)^{d-1}}{d^d} =: p_{Shearer} \quad \forall\, i = 1, 2, \ldots, n. \tag{5.1}$$

Then

$$\mathbb{P}\left(\bigcap_{i=1}^n \overline{\mathcal{E}}_i\right) > 0.$$

Proof: Let G be a dependency graph for $\{\mathcal{E}_i\}$ with maximum degree d, and let $p = (d-1)^{d-1}/d^d$. We may assume that G is connected, since otherwise the problem reduces to a collection of independent problems. For G connected, we can find an ordering of vertices $(v_1, \ldots, v_n)$ such that each v_i, $i \geq 2$, has degree at most $d-1$ among $\{v_i, \ldots, v_n\}$. (However, this is not possible to arrange for v_1 if G is d-regular. Therefore, we need to handle this case separately later.)

By induction on $|S|$ we claim that

$$\frac{\breve{q}_S}{\breve{q}_{S-a}} > 1 - \frac{1}{d} \qquad \text{for } a \in S \text{ where } |S \cap \Gamma(a)| \leq d-1.$$

The base case of the induction is satisfied as

$$\frac{\breve{q}_{\{a\}}}{\breve{q}_\emptyset} = \frac{\breve{q}_{\{a\}}}{1} = 1 - p = 1 - \frac{(d-1)^{d-1}}{d^d} > 1 - \frac{d^{d-1}}{d^d} = 1 - \frac{1}{d}.$$

For the general case, we will use an identity established in the proof of the asymmetric case (see Lecture 5):

$$\frac{\breve{q}_S}{\breve{q}_{S-a}} = 1 - p \cdot \frac{\breve{q}_{S \setminus \Gamma^+(a)}}{\breve{q}_{S-a}}. \tag{5.2}$$

Assume that $a \in S$ is such that $|S \cap \Gamma(a)| \leq d-1$, and write $S \cap \Gamma^+(a) = \{a, a_1, a_2, \ldots, a_k\}$, $k \leq d-1$. Since each a_i has degree at most $d-1$ inside $S \setminus \{a\}$, the inductive hypothesis gives

$$\frac{\breve{q}_{S\setminus\Gamma^+(a)}}{\breve{q}_{S-a}} = \underbrace{\frac{\breve{q}_{S\setminus\Gamma^+(a)}}{\breve{q}_{(S\setminus\Gamma^+(a))+a_k}} \cdot \frac{\breve{q}_{(S\setminus\Gamma^+(a))+a_k}}{\breve{q}_{(S\setminus\Gamma^+(a))+a_k+a_{k-1}}} \cdots \frac{\breve{q}_{S-a-a_1}}{\breve{q}_{S-a}}}_{\text{at most } d-1 \text{ terms}} < \frac{1}{(1-1/d)^{d-1}}. \tag{5.3}$$

From (6.2), we get

$$\frac{\breve{q}_S}{\breve{q}_{S-a}} = 1 - p \cdot \frac{\breve{q}_{S \setminus \Gamma^+(a)}}{\breve{q}_{S-a}} > 1 - p \cdot \frac{d^{d-1}}{(d-1)^{d-1}} = 1 - \frac{1}{d}$$

which finishes the inductive claim.

Finally let us handle the case where $S = \{v_1, \ldots, v_n\}$ and v_1 has degree d. We can still use (6.2), but now the telescoping product in (6.3) may involve d terms, giving

$$\frac{\breve{q}_{[n]}}{\breve{q}_{[n]-v_1}} = 1 - p \cdot \frac{\breve{q}_{[n] \setminus \Gamma^+(v_1)}}{\breve{q}_{[n]-v_1}} > 1 - p \cdot \frac{d^d}{(d-1)^d} = 1 - \frac{1}{d-1}.$$

We note that $1 - \frac{1}{d-1}$ could be 0 (for $d = 2$) but the strict inequality ensures that the ratio is still positive. We conclude that

$$\mathbb{P}\left(\bigcap_{i=1}^{n} \overline{\mathcal{E}}_i\right) \geq \breve{q}_{[n]} = \frac{\breve{q}_{[n]}}{\breve{q}_{[n]-a_1}} \frac{\breve{q}_{[n]-v_1}}{\breve{q}_{[n]-v_1-v_2}} \cdots \frac{\breve{q}_{v_n}}{\breve{q}_{\emptyset}} > \left(1 - \frac{1}{d-1}\right)\left(1 - \frac{1}{d}\right)^{n-1} \geq 0,$$

completing the proof. □

Let us compare Symmetric Shearer's Lemma to the Lovász Local Lemma. In the LLL, assuming that all events get the same parameter x, it is required that

$$p \leq x(1-x)^d \tag{5.4}$$

for some $x \in (0,1)$. The optimal choice here can be shown to be $x = \frac{1}{d+1}$, which gives

$$p \leq \frac{d^d}{(d+1)^{d+1}} =: p_{\text{LLL}}. \tag{5.5}$$

Comparing (6.5) to (6.1), we see that the threshold probability in Shearer's lemma, $p_{\text{Shearer}} := \frac{(d-1)^{d-1}}{d^d}$, has the benefit of 1 additional dependency over the LLL. Further, the inequalities

$$\frac{(d+1)^d}{d^d} < e < \frac{d^d}{(d-1)^d}$$

show that

$$\frac{1}{e(d+1)} < p_{\text{LLL}} < \frac{1}{ed} < p_{\text{Shearer}} < \frac{1}{e(d-1)}.$$

Of course, as d grows large, p_{LLL} and p_{Shearer} are asymptotically the same.

5.1 Worst instance: d-regular trees

We would like to demonstrate that p_{Shearer} is optimal, in the sense that Theorem 6.1 fails if p_{Shearer} is taken any larger. The extreme case is when each $\mathcal{E}_i$ is dependent on exactly d other events and moreover the dependency graph is a (large) d-regular tree. Begin with a root vertex r, by itself called $\mathcal{T}_0$. A root with $d-1$ children is called $\mathcal{T}_1$. Constructed recursively, $\mathcal{T}_\ell$ is the tree obtained by taking a root with $d-1$ subtrees, each of which is $\mathcal{T}_{\ell-1}$. Note that all vertices in levels 1 through

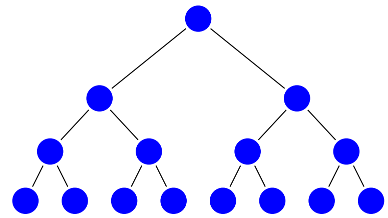

Figure 1: A binary tree ($d = 3$). Here, $\mathcal{T}_3$ is shown (levels 0 through 3).

$\ell-1$ have degree d; the root has degree $d-1$, and level ℓ consists of leaves. We call this a d-regular tree of depth ℓ (ignoring the slightly different degree at the root). For consistency, we also define $\mathcal{T}_{-1}$ to be the empty tree.

Suppose that the probability of each event is p. From (6.2), we have

$$\breve{q}_{\mathcal{T}_\ell} = \breve{q}_{\mathcal{T}_\ell \setminus r} - p \cdot \breve{q}_{\mathcal{T}_\ell \setminus \Gamma^+(r)}.$$

But $\mathcal{T}_\ell \setminus r$ is the union of $d-1$ disjoint copies of $\mathcal{T}_{\ell-1}$. Similarly, $\mathcal{T}_\ell \setminus \Gamma^+(r)$ is the union of $(d-1)^2$ disjoint copies of $\mathcal{T}_{\ell-2}$. Hence

$$\breve{q}_{\mathcal{T}_\ell} = \left(\breve{q}_{\mathcal{T}_{\ell-1}}\right)^{d-1} - p\left(\breve{q}_{\mathcal{T}_{\ell-2}}\right)^{(d-1)^2}.$$

Let us define

$$b_\ell := \frac{\breve{q}_{\mathcal{T}_\ell}}{(\breve{q}_{\mathcal{T}_{\ell-1}})^{d-1}} = 1 - p\left(\frac{(\breve{q}_{\mathcal{T}_{\ell-2}})^{d-1}}{\breve{q}_{\mathcal{T}_{\ell-1}}}\right)^{d-1}.$$

That is,

$$b_\ell = 1 - p\left(\frac{1}{b_{\ell-1}}\right)^{d-1}.$$

If Shearer's positivity conditions are satisfied for an arbitrarily large d-regular tree, then $b_\ell > 0$ for all $\ell \geq 0$, and also the sequence is decreasing by induction: $b_0 = 1-p$, $b_1 = 1 - \frac{p}{(1-p)^{d-1}} \leq b_0$, and if $b_\ell \leq b_{\ell-1}$, then $b_{\ell+1} = 1 - p/b_\ell^{d-1} \leq 1 - p/b_{\ell-1}^{d-1} = b_\ell$. Hence there is a limit,

$$\lambda := \lim_{\ell\to\infty} b_\ell,$$

which must satisfy $\lambda = 1 - \frac{p}{\lambda^{d-1}}$, and hence $p = \lambda^{d-1} - \lambda^d$. The maximum is attained at $\lambda = \frac{d-1}{d}$, which gives

$$p \leq \left(\frac{d-1}{d}\right)^{d-1} - \left(\frac{d-1}{d}\right)^{d} = \frac{1}{d}\left(\frac{d-1}{d}\right)^{d-1} = p_{\text{Shearer}}.$$

Indeed, p_{Shearer} is optimal.

5.2 Application of Shearer's Lemma: the multipartite Turán problem

Consider an r-partite graph G on $V_1 \cup V_2 \cup \cdots \cup V_r$. Suppose we have at least a certain density ρ between any two parts:

$$e(V_i, V_j) \geq \rho|V_i||V_j| \quad \forall\, i \neq j.$$

How large must ρ be to guarantee the existence of a clique K_r in G? More generally, given a graph H on r vertices, assume

$$\{i,j\} \in E(H) \quad \Rightarrow \quad e(V_i, V_j) \geq \rho |V_i||V_j|.$$

How large must ρ be to guarantee the existence of a copy of H in G? Following (Csikvári and Nagy, 2012), we show how to apply Shearer's Lemma.

Pick $x_i \in V_i$ independently and uniformly at random. For each $(i,j) \in E(H)$, define an event $\mathcal{E}_{ij} = \{\{x_i, x_j\} \notin E(G)\}$, so that if all $\mathcal{E}_{ij}$ are avoided, then a copy of H is present. Note that the probability of each event is at most $1-\rho$ by assumption. A dependency graph for the events $\mathcal{E}_{ij}$ is the *line graph* of H, which we call D: The vertices of D are the edges in H, and two of these vertices are adjacent if and only if the corresponding edges in H share a vertex. So independent sets in D are exactly matchings in H.

First, consider Symmetric Shearer's Lemma. The degrees in D are at most $2(\Delta(H)-1)$ where $\Delta(H)$ is the maximum degree in H. Hence, if the probability of each event is at most $\frac{1}{2e(\Delta(H)-1)}$, then by Theorem 6.1, $\mathbb{P}[\bigcap_{(i,j)\in E(H)} \mathcal{E}_{ij}] > 0$. Equivalently, if $\rho \geq 1 - \frac{1}{2e(\Delta(H)-1)}$ then G contains a copy of H.

This problem is actually a rare setting where we can apply Shearer's Lemma directly and obtain a stronger result. Consider the polynomial

$$q_\emptyset(p) = \sum_{I \in \mathrm{Ind}(D)} (-1)^{|I|} p^I = \sum_{\substack{M \subset H \\ M \text{matching}}} (-1)^{|M|} p^{|M|}.$$

This last sum is a variant of the *matching polynomial* of K_r. It is most commonly defined in the following form, which we refer to as the *matching defect polynomial*:

$$\mathcal{M}_H(x) = \sum_{\substack{M \subset H \\ M \text{matching}}} (-1)^{|M|} x^{r-2|M|}.$$

(Recall that $r = |V(H)|$.) A simple calculation gives

$$\mathcal{M}_H(x) = x^r q_\emptyset\left(\frac{1}{x^2}\right).$$

It is useful in this setting to consider Property 4 of Shearer's Lemma, stated in Lecture 5. In particular, we ask, for which p is it true that

$$q_\emptyset(\lambda p) > 0 \quad \forall\, \lambda \in [0,1]?$$

To answer this question, it suffices to locate the minimum positive root of $q_\emptyset$, or equivalently the maximum positive root of $\mathcal{M}_H$. Here we appeal to the following theorem (which we will prove later in this course).

Theorem 5.2 (Heilmann-Lieb) *For any graph H, the roots of the matching defect polynomial are all real and the maximum root is at most $2\sqrt{\Delta(H)-1}$.*

It follows that the minimum positive root of $q_\emptyset(p)$ for H is at least $\frac{1}{4(\Delta(H)-1)}$. Consequently, if $\rho \geq 1 - \frac{1}{4(\Delta(H)-1)}$ then G contains a copy of H, which improves the bound from above ($2e$ has been improved to 4).

For $H = K_r$, which is perhaps the most interesting special case here, we obtain that density $\rho \geq 1 - \frac{1}{4(r-2)}$ is sufficient to guarantee a copy of K_r. In fact, here we can go one step further a obtain a slightly tighter bound. $\mathcal{M}_{K_r}$ is known to be the Hermite polynomial of degree r. Recall that the Hermite polynomials are defined recursively, corresponding to the recursion in the context of matchings:

$$H_0(x) = 1$$
$$H_{r+1}(x) = xH_r(x) - rH_{r-1}(x).$$

For this special case, more accurate bounds are known. In particular, the maximum root of $\mathcal{M}_{K_r}$ is known to be $2\sqrt{r} - \Theta(r^{-1/6})$. Hence, the minimum positive root of $q_\emptyset$ is

$$\frac{1}{(2\sqrt{r} - \Theta(r^{-1/6}))^2} = \frac{1}{4r - \Theta(r^{1/3})}.$$

Consequently, if $\rho \geq 1 - \frac{1}{4r - \Theta(r^{1/3})}$ then G contains a copy of K_r, a slight improvement over the bound of $1 - \frac{1}{4(r-2)}$ from the Heilmann-Lieb theorem.

We conclude by mentioning that it is easy to construct an r-partitite graph of density $\rho = 1 - \frac{1}{r-1}$ which does not contain a K_r (an exercise). A better counterexample which can be found in [Csikvári-Nagy'12] implies that $\rho = 1 - \frac{1}{(2-\epsilon)r}$ is not sufficient to guarantee a copy of K_r for any $\epsilon > 0$. The gap between $1 - \frac{1}{2r}$ and $1 - \frac{1}{4r}$ remains open.

References

Péter Csikvári and Zoltán Lóránt Nagy. The density Turán problem. *Combinatorics, Probability and Computing*, 21(04):531–553, 2012.