Lecture 3 - Fox's proof of the removal lemma

The graph removal lemma states that for any graph H and any $\epsilon > 0$, there exists $\delta > 0$ such that any graph on n vertices with at most $\delta n^{v(H)}$ copies of H can be made H-free by removing at most ϵn^2 edges. The original proof of this uses the regularity lemma and gives a bound on δ^{-1} which is a tower of twos of height polynomial in ϵ^{-1}. In this lecture, we will discuss an alternative proof of this result, given by Fox, which improves the bound on δ^{-1} to a tower of twos of height logarithmic in ϵ^{-1}. While still huge, this result is interesting because it bypasses the lower bound for the regularity lemma.

The first ingredient we will need in the proof is a variant of the iterated weak regularity lemma which we proved in the previous lecture. Rather than fixing on the mean square density, it will be useful to define a mean-q density for any convex function q. If $\mathcal{A} = X_1 \cup \dots \cup X_k$ is a partition of the vertex set of a graph on n vertices, this is defined by

$$q(\mathcal{A}) = \sum_{1 \le i,j \le k} \frac{|X_i||X_j|}{n^2} q(d(X_i, X_j)).$$

We note that for any convex function q and any partition $\mathcal{A}$ of a graph with density d,

$$q(d) \le q(\mathcal{A}) \le dq(1) + (1-d)q(0).$$

The proof of the next lemma is essentially the same as that given in the previous lecture.

Lemma 1 *Let $q : [0,1] \to \mathbb{R}$ be a convex function, G be a graph with density d, $f : \mathbb{N} \to [0,1]$ be a decreasing function and $r = (dq(1) + (1-d)q(0) - q(d))/\gamma$. Then there are equitable partitions P and Q with Q a refinement of P satisfying $q(Q) \le q(P) + \gamma$, Q is weak $f(|P|)$-regular and $|Q| \le t_r$, where $t_0 = 1$, $t_i = t_{i-1}R(f(t_{i-1}))$ for $1 \le i \le r$ and $R(x) = 2^{cx^{-2}}$ as in the Frieze-Kannan weak regularity lemma.*

Rather than using the usual $q(x) = x^2$, we will use the convex function q on $[0,1]$ defined by $q(0) = 0$ and $q(x) = x \log x$ for $x \in (0,1]$. This entropy function is central to the proof since it captures the extra structural information coming from Lemma 3 below in a concise fashion. Note that $d \log d \le q(P) \le 0$ for every partition P.

The next lemma is a counting lemma that complements the Frieze-Kannan weak regularity lemma. As one might expect, this lemma gives a global count for the number of copies of H, whereas the counting lemma associated with the usual regularity lemma gives a means of counting copies of H between any $v(H)$ parts of the partition which are pairwise regular. Its proof, which we omit, is by a simple telescoping sum argument.

Lemma 2 *Let H be a graph on $\{1, \dots, h\}$ with m edges. Let $G = (V, E)$ be a graph on n vertices and $Q : V = V_1 \cup \dots \cup V_t$ be a vertex partition which is weak ϵ-regular. The number of homomorphisms from H to G is within $\epsilon m n^h$ of*

$$\sum_{1 \le i_1, \dots, i_h \le t} \prod_{(r,s) \in E(H)} d(V_{i_r}, V_{i_s}) \prod_{a=1}^{h} |V_{i_a}|.$$

Let P and Q be vertex partitions of a graph G with Q a refinement of P. A pair (V_i, V_j) of parts of P is (α, c)-*shattered* by Q if at least a c-fraction of the pairs $(u, v) \in V_i \times V_j$ go between pairs of parts of Q with edge density between them less than α.

The key component in the proof is the following lemma, which says that if P and Q are vertex partitions like those given by Lemma 1, then there are many pairs of vertex sets in P which are shattered by Q.

Lemma 3 *Let H be a graph on $\{1, \ldots, h\}$ with m edges and let $\alpha > 0$. Suppose G is a graph on n vertices for which there are fewer than δn^h homomorphisms of H into G, where $\delta = \frac{1}{4}\alpha^m(2k)^{-h}$. Suppose P and Q are equitable vertex partitions of G with $|P| = k \leq n$ and Q is a refinement of P which is weak $f(k)$-regular, where $f(k) = \frac{1}{4m}\alpha^m(2k)^{-h}$. For every h-tuple $V_1, \ldots, V_h$ of parts of P, there is an edge (i, j) of H for which the pair (V_i, V_j) is $(\alpha, \frac{1}{2m})$-shattered by Q.*

Proof As $|P| = k \leq n$, we have $|V_i| \geq \frac{n}{2k}$ for each i. Let Q_i denote the partition of V_i which consists of the parts of Q which are subsets of V_i. Consider an h-tuple $(v_1, \ldots, v_h) \in V_1 \times \cdots \times V_h$ picked uniformly at random. Also consider the event E that, for each edge (i, j) of H, the pair (v_i, v_j) goes between parts of Q_i and Q_j with density at least α. If E occurs with probability at least $1/2$, as Q is weak $f(k)$-regular, Lemma 2 implies that the number of homomorphisms of H into G where the copy of vertex i is in V_i for $1 \leq i \leq h$ is at least

$$\frac{1}{2}\alpha^m \prod_{i=1}^{h} |V_i| - mf(k)n^h \geq \left(\frac{1}{2}\alpha^m(2k)^{-h} - mf(k)\right) n^h = \delta n^h,$$

contradicting that there are fewer than δn^h homomorphisms of H into G. So E occurs with probability less than $1/2$. Hence, for at least $1/2$ of the h-tuples $(v_1, \ldots, v_h) \in V_1 \times \cdots \times V_h$, there is an edge (i, j) of H such that the pair (v_i, v_j) goes between parts of Q_i and Q_j with density less than α. This implies that for at least one edge (i, j) of H, the pair (V_i, V_j) is $(\alpha, \frac{1}{2m})$-shattered by Q. □

We will need the following lemma which tells us that if a pair of parts from P is shattered by Q then there is an increment in the mean-entropy density. Its proof, which we again omit, is a simple application of Jensen's inequality.

Lemma 4 *Let $q : [0, 1] \to \mathbb{R}$ be the convex function given by $q(0) = 0$ and $q(x) = x \log x$ for $x > 0$. Let $\epsilon_1, \ldots, \epsilon_r$ and $d_1, \ldots, d_r$ be nonnegative real numbers with $\sum_{i=1}^{r} \epsilon_i = 1$ and $d = \sum_{i=1}^{r} \epsilon_i d_i$. Suppose $\beta < 1$ and $I \subset [r]$ is such that $d_i \leq \beta d$ for $i \in I$ and let $s = \sum_{i \in I} \epsilon_i$. Then*

$$\sum_{i=1}^{r} \epsilon_i q(d_i) \geq q(d) + (1 - \beta + q(\beta))sd.$$

We are now ready to prove the theorem in the following precise form.

Theorem 1 *Let H be a graph on $\{1, \ldots, h\}$ with m edges. Let $\epsilon > 0$ and δ^{-1} be a tower of twos of height $8h^4 \log \epsilon^{-1}$. If G is a graph on n vertices in which at least ϵn^2 edges need to be removed to make it H-free, then G contains at least δn^h copies of H.*

Proof Suppose for contradiction that there is a graph G on n vertices in which at least ϵn^2 edges need to be removed from G to delete all copies of H, but G contains fewer than δn^h copies of H. If

$n \leq \delta^{-1/h}$, then the number of copies of H in G is less than $\delta n^h \leq 1$, so G is H-free, contradicting that at least ϵn^2 edges need to be removed to make the graph H-free. Hence, $n > \delta^{-1/h}$. Note that the number of mappings from $V(H)$ to $V(G)$ which are not one-to-one is $n^h - h!\binom{n}{h} \leq h^2 n^{h-1} < h^2 \delta^{1/h} n^h$. Let $\delta' = 2h^2\delta^{1/h}$, so the number of homomorphisms from H to G is at most $\delta' n^h$.

The graph G contains at least $\epsilon n^2/m$ edge-disjoint copies of H. Let G' be the graph on the same vertex set which consists entirely of the at least $\epsilon n^2/m$ edge-disjoint copies of H. Then $d(G') \geq m \cdot \epsilon/m = \epsilon$ and G' consists of $\frac{d(G')}{m}n^2$ edge-disjoint copies of H. We will show that there are at least $\delta' n^h$ homomorphisms from H to G' (and hence to G as well). For the rest of the argument, we will assume the underlying graph is G'.

Let $\alpha = \frac{\epsilon}{8m}$. Apply Lemma 1 to G' with $f(k) = \frac{1}{4m}\alpha^m(2k)^{-h}$ and $\gamma = \frac{d(G')}{2h^4}$. Note that r as in Lemma 1 is

$$r = d(G')\log(1/d(G'))/\gamma = 2h^4\log(1/d(G')) \leq 2h^4 \log \epsilon^{-1}.$$

Hence, we get a pair of equitable vertex partitions P and Q, with Q a refinement of P, $q(Q) \leq q(P)+\gamma$, Q is weak $f(|P|)$-regular and $|Q|$ is at most a tower of twos of height $3r \leq 6h^4\log\epsilon^{-1}$. Let $V_1, \ldots, V_k$ denote the parts of P and Q_i denote the partition of V_i consisting of the parts of Q which are subsets of V_i.

Suppose that (V_a, V_b) is a pair of parts of P with edge density $d = d(V_a, V_b) \geq \epsilon/m$ which is $(\alpha, \frac{1}{2m})$-shattered by Q. Note that $\alpha \leq d/8$. Arbitrarily order the pairs $U_i \times W_i \in Q_a \times Q_b$, letting $d_i = d(U_i, W_i)$ and $\epsilon_i = \frac{|U_i||W_i|}{|V_a||V_b|}$, so that the conditions of Lemma 4 with $\beta = 1/8$ are satisfied. Applying Lemma 4, we get, since $q(\beta) = -\frac{1}{8}\log 8 = -\frac{3}{8}$, that

$$q(Q_a, Q_b) - q(V_a, V_b) \geq (1 - \beta + q(\beta))\frac{1}{2m}d(V_a, V_b)|V_a||V_b|/n^2 \geq \frac{1}{4m}e(V_a, V_b)/n^2.$$

Note that

$$q(Q) - q(P) = \sum_{1 \leq a,b \leq k} (q(Q_a, Q_b) - q(V_a, V_b)),$$

which shows that $q(Q) - q(P)$ is the sum of nonnegative summands.

There are at most $\frac{\epsilon}{m}n^2/2$ edges of G' going between pairs of parts of P with density at most $\frac{\epsilon}{m}$. Hence, at least $1/2$ of the edge-disjoint copies of H making up G' have all its edges going between pairs of parts of P of density at least $\frac{\epsilon}{m}$. By Lemma 3, for each copy of H, at least one of its edges goes between a pair of parts of P which is $(\alpha, \frac{1}{2m})$-shattered by Q. Thus,

$$q(Q) - q(P) \geq \sum \frac{1}{4m}e(V_a, V_b)/n^2 \geq \frac{1}{4m} \cdot \frac{d(G')}{2m} = \frac{d(G')}{8m^2} > \gamma,$$

where the sum is over all ordered pairs (V_a, V_b) of parts of P which are $(\alpha, \frac{1}{2m})$-shattered by Q and with $d(V_a, V_b) \geq \frac{\epsilon}{m}$. This contradicts $q(Q) \leq q(P) + \gamma$ and completes the proof. $\square$