EXTREMAL GRAPH THEORY 2 - THE RANDOM ALGEBRAIC METHOD

DAVID CONLON

In this lecture, we describe a method, first proposed by Boris Bukh and then developed further by Bukh and the lecturer, for proving lower bounds for extremal numbers. We will demonstrate this in the simplest case by proving the following result.

Theorem 1. *For any integer $s \geq 1$, there exists an integer $t \geq s$ and a constant $c > 0$ such that*

$$ex(n, K_{s,t}) \geq cn^{2-1/s}.$$

The set-up requires some definitions. Let q be a prime power and let $\mathbb{F}_q$ be the finite field of order q. We will consider polynomials in t variables over $\mathbb{F}_q$, writing any such polynomial as $f(X)$, where $X = (X_1, \dots, X_t)$. We let $\mathcal{P}_d$ be the set of polynomials in X of degree at most d, that is, the set of linear combinations over $\mathbb{F}_q$ of monomials of the form $X_1^{a_1} \cdots X_t^{a_t}$ with $\sum_{i=1}^t a_i \leq d$. By a random polynomial, we just mean a polynomial chosen uniformly from the set $\mathcal{P}_d$. One may produce such a random polynomial by choosing the coefficients of the monomials above to be random elements of $\mathbb{F}_q$.

We will now determine the probability that a randomly chosen polynomial from $\mathcal{P}_d$ passes through a given set of points. For one point, the probability is $1/q$, as shown by the following simple result.

Lemma 1. *If f is a random polynomial from $\mathcal{P}_d$, then, for any fixed $x \in \mathbb{F}_q^t$,*

$$\mathbb{P}[f(x) = 0] = 1/q.$$

Proof. Let $\mathcal{Q}_d = \{f \in \mathcal{P}_d : f(0) = 0\}$, that is, the collection of polynomials in $\mathcal{P}_d$ with zero constant term. Since every $f \in \mathcal{P}_d$ can be written as $g + h$ where $g \in \mathcal{Q}_d$ and h is a constant, we may sample a random element of $\mathcal{P}_d$ by adding a random element g of $\mathcal{Q}_d$ and a random element h of $\mathbb{F}_q$. Since, for any fixed choice of g, there is only one choice out of q for h such that $f(x) = 0$, the result follows. □

The following result shows that once q and d are sufficiently large, the probability that a randomly chosen polynomial from $\mathcal{P}_d$ contains each of m distinct points is exactly $1/q^m$. That is, the events are independent.

Lemma 2. *Suppose that $q > \binom{m}{2}$ and $d \geq m - 1$. Then, if f is a random polynomial from $\mathcal{P}_d$ and $x_1, \dots, x_m$ are m distinct points in $\mathbb{F}_q^t$,*

$$\mathbb{P}[f(x_i) = 0 \text{ for all } i = 1, \dots, m] = 1/q^m.$$

Proof. Let $x_i = (x_{i,1}, \dots, x_{i,t})$ for each $i = 1, \dots, m$. We will choose elements $a_2, \dots, a_t \in \mathbb{F}_q$ such that $x_{i,1} + \sum_{j=2}^t a_j x_{i,j}$ is distinct for all $i = 1, \dots, m$. To see that this is possible, note that there are exactly $\binom{m}{2}$ equations

$$x_{i,1} + \sum_{j=2}^t a_j x_{i,j} = x_{i',1} + \sum_{j=2}^t a_j x_{i',j},$$

each with at most q^{t-2} solutions $(a_2, \dots, a_t)$. Therefore, since the total number of choices for $(a_2, \dots, a_t)$ is q^{t-1} and $q^{t-1} > q^{t-2}\binom{m}{2}$, we can make an appropriate choice.

We now consider $\mathcal{P}'_d$, the set of polynomials of degree at most d in Z, where $Z_1 = X_1 + \sum_{j=2}^t a_j X_j$ and $Z_j = X_j$ for all $2 \leq j \leq t$. Since this change of variables is an invertible linear map, $\mathcal{P}'_d$ is identical to $\mathcal{P}_d$, so it will suffice to show that a randomly chosen polynomial from $\mathcal{P}'_d$ passes through all of the points $z_1, \ldots, z_m$ corresponding to $x_1, \ldots, x_m$ with probability q^{-m}. Note that, by our choice above, $z_{i,1} \neq z_{i',1}$ for any $1 \leq i < i' \leq m$.

For any f in $\mathcal{P}'_d$, we may write $f = g + h$, where h contains all monomials of the form Z_1^j for $j = 0, 1, \ldots, m-1$ and g contains all other monomials. For any fixed choice of g, there is exactly one choice of h such that $f(z_i) = 0$ for all $i = 1, \ldots, m$, namely, the unique polynomial of degree at most $m-1$ which takes the value $-g(z_i)$ at $z_{i,1}$ for all $i = 1, 2, \ldots, m$, where uniqueness follows from the fact that the $z_{i,1}$ are distinct. Since there were q^m possible choices for h, the result follows. □

We also need to note some basic facts about affine varieties over finite fields. If we write $\overline{\mathbb{F}}_q$ for the algebraic closure of $\mathbb{F}_q$, a variety over $\overline{\mathbb{F}}_q$ is a set of the form

$$W = \{x \in \overline{\mathbb{F}}_q^t : f_1(x) = \cdots = f_s(x) = 0\}$$

for some collection of polynomials $f_1, \ldots, f_s : \overline{\mathbb{F}}_q^t \to \overline{\mathbb{F}}_q$. We say that W is defined over $\mathbb{F}_q$ if the coefficients of these polynomials are in $\mathbb{F}_q$ and write $W(\mathbb{F}_q) = W \cap \mathbb{F}_q^t$. We say that W has complexity at most M if s, t and the degrees of the f_i are all bounded by M. Finally, we say that a variety is absolutely irreducible if it is irreducible over $\overline{\mathbb{F}}_q$, reserving the term irreducibility for irreducibility over $\mathbb{F}_q$ of varieties defined over $\mathbb{F}_q$.

The first result we will need is the Lang–Weil bound relating the dimension of a variety W to the number of points in $W(\mathbb{F}_q)$.

Lemma 3. *Suppose that W is a variety over $\overline{\mathbb{F}}_q$ of complexity at most M. Then*

$$|W(\mathbb{F}_q)| = O_M(q^{\dim W}).$$

Moreover, if W is defined over $\mathbb{F}_q$ and absolutely irreducible, then

$$|W(\mathbb{F}_q)| = q^{\dim W}(1 + O_M(q^{-1/2})).$$

Secondly, we need the following result.

Lemma 4. *Suppose that W is a variety over $\overline{\mathbb{F}}_q$ of complexity at most M which is defined over $\mathbb{F}_q$. Then there are $O_M(1)$ absolutely irreducible varieties $Y_1, \ldots, Y_s$, each of which is defined over $\mathbb{F}_q$ and has complexity $O_M(1)$, such that $\cup_{i=1}^s Y_i(\mathbb{F}_q) = W(\mathbb{F}_q)$.*

In reality, what we need is the following combination of these results.

Lemma 5. *Suppose W is a variety over $\overline{\mathbb{F}}_q$ of complexity at most M which is defined over $\mathbb{F}_q$. Then, for all q sufficiently large in terms of M, either $|W(\mathbb{F}_q)| \geq q/2$ or $|W(\mathbb{F}_q)| \leq C$ for some $C = C_M$ depending only on M.*

Proof. By Lemma 4, there is a decomposition $W(\mathbb{F}_q) = \bigcup_{i=1}^s Y_i(\mathbb{F}_q)$ for some bounded-complexity absolutely irreducible varieties Y_i defined over $\mathbb{F}_q$. If $\dim Y_i \geq 1$ for some i, Lemma 3 tells us that $|W(\mathbb{F}_q)| \geq |Y_i(\mathbb{F}_q)| = q^{\dim Y_i}(1 + O_M(q^{-1/2})) \geq q/2$. On the other hand, if $\dim Y_i = 0$ for every Y_i, Lemma 3 tells us that $|W(\mathbb{F}_q)| \leq \sum |Y_i(\mathbb{F}_q)| = O_M(1)$. □

We are now ready to prove our main theorem.

Proof of Theorem 1. Let $t = 2s$, $d = s^3$, $N = q^s$ and suppose that q is sufficiently large. Let $f : \mathbb{F}_q^s \times \mathbb{F}_q^s \to \mathbb{F}_q$ be a random polynomial in $\mathcal{P}_d$. We consider the bipartite graph G between two

copies U and V of $\mathbb{F}_q^s$, each of order $N = q^s$, where (u, v) is an edge of G if and only if $f(u, v) = 0$. Lemma 1 tells us that the probability a given edge (u, v) is in G is $1/q$. Therefore, the expected number of edges in G is $N^2/q = N^{2-1/s}$.

Suppose that $u_1, \dots, u_s$ are distinct points, all in U (a similar argument applies if they are all in V), and consider the variety

$$W = \{v \in \mathbb{F}_q^s : f(u_1, v) = f(u_2, v) = \cdots = f(u_s, v) = 0\}.$$

Then, by Lemma 5, $|W| \leq C$ or $|W| \geq q/2$. Moreover,

$$\mathbb{E}[|W|^r] = \sum_{v_1,\dots,v_r \in V} \mathbb{P}[f(u_i, v_j) = 0 \text{ for all } i = 1, \dots, s, j = 1, \dots, r].$$

By Lemma 2, if ℓ of the v_j are distinct and $d \geq s\ell$, this probability is $q^{-s\ell}$. If we let N_ℓ be the number of surjective functions from an r-element set to an ℓ-element set, we therefore see that

$$\mathbb{E}[|W|^r] = \sum_{\ell \leq r} \binom{q^s}{\ell} N_\ell q^{-s\ell} \leq \sum_{\ell \leq r} N_\ell = O_r(1).$$

Hence, by Markov's inequality,

$$\mathbb{P}[|W| > C] = \mathbb{P}[|W| \geq q/2] = \mathbb{P}[|W|^r \geq (q/2)^r] \leq \frac{O_r(1)}{(q/2)^r}.$$

We call an s-tuple $(u_1, \dots, u_s)$ bad if the vertices in the s-tuple have more than C common neighbours. If we let B be the random variable counting the number of bad s-tuples, we have, by taking a union bound over the $\binom{q^s}{s} \leq q^{s^2}$ possible s-tuples in U, that, for $r = s^2$,

$$\mathbb{E}[B] \leq q^{s^2} \cdot \frac{O_r(1)}{(q/2)^r} = O_r(q^{s^2-r}) = O_s(1).$$

We now remove a vertex from each bad s-tuple to form a new graph G'. Since each vertex has degree at most N, the total number of edges removed is at most BN. Hence, the expected number of edges is

$$N^{2-1/s} - \mathbb{E}[B]N = \Omega_s(N^{2-1/s}).$$

Therefore, there is a graph with at most $2N$ vertices and $\Omega_s(N^{2-1/s})$ edges such that no s-tuple of vertices have more than C common neighbours. As stated, this result only holds when q is a prime power and $N = q^k$. However, it is a simple matter to use Bertrand's postulate to show that the same conclusion holds for all N. □

This method has since seen several applications. The first was an adaptation of Bukh's technique to prove an analogue of the Kollár–Rónyai–Szabó result for the even cycle problem.

Theorem 2 (Conlon)**.** *For every natural number $k \geq 2$, there exists a natural number ℓ such that, for every n, there is a graph with n vertices and $\Omega(n^{1+1/k})$ edges with at most ℓ paths of length k between any two vertices.*

This was then further generalised by Bukh and Conlon to make progress on the rational exponents conjecture of Erdős and Simonovits. For the statement, note that, given a finite family of graphs $\mathcal{H}$, the extremal number $\mathrm{ex}(n, \mathcal{H})$ is the largest number of edges in an n-vertex graph containing no graph from $\mathcal{H}$ as a subgraph.

Theorem 3 (Bukh–Conlon)**.** *For every rational number $r \in [1, 2]$, there exists a finite family of graphs $\mathcal{H}_r$ such that $ex(n, \mathcal{H}_r) = \Theta(n^r)$.*

The original conjecture asked for single graphs rather than families. This remains open, though, building on the work of Bukh and Conlon, it has been verified for a broad range of exponents, including all rationals of the form $1 + a/b$ and $2 - a/b$ with $b > a^2$.

Two other basic conjectures of Erdős and Simonovits remain wide open. The first is a sort of converse of the conjecture above.

Conjecture 1. *For every graph H, there exists a rational $r \in [1, 2]$ such that $ex(n, \mathcal{H}_r) = \Theta(n^r)$.*

The second, which would show that the result of Bukh and Conlon implies the original conjecture for singletons, is the compactness conjecture, which reads as follows.

Conjecture 2. *For any finite family of graphs $\mathcal{H}$, there exists $H \in \mathcal{H}$ such that $ex(n, \mathcal{H}) = \Theta(ex(n, H))$.*

This lecturer is disinclined from believing either conjecture, which are known to be false for hypergraphs. For instance, if one takes the 3-uniform hypergraphs H_1 and H_2 with edge sets $E(H_1) = \{abc, abd\}$ and $E(H_2) = \{abd, bce, caf\}$, then $\mathrm{ex}(n, H_1) = \Theta(n^2)$ and $\mathrm{ex}(n, H_2) = \Theta(n^2)$, while $\mathrm{ex}(n, \{H_1, H_2\}) = o(n^2)$, disproving the compactness conjecture in this case. We note that the bound $\mathrm{ex}(n, \{H_1, H_2\}) = o(n^2)$ is a consequence of the celebrated triangle removal lemma, which, unfortunately, we will not have time to explore further here.