10 Extremal numbers of general bipartite graphs

As such, it is a major research direction to study the extremal numbers of bipartite graphs. We have already discussed what we know about complete bipartite graphs, but we now turn our attention to $\mathrm{ex}(n, H)$ for general bipartite H.

As it turns out, we know a lot about this question, but there's much more that we don't know. We already saw a general-purpose upper bound, namely $\mathrm{ex}(n, H) \leqslant \mathrm{ex}(n, K_{s,t}) \leqslant O(n^{2-1/s})$ if H is a subgraph of $K_{s,t}$. But for many specific bipartite graphs, much better upper bounds are known, using arguments specific to the graph at hand. For example, the following is known about the extremal numbers of even cycles.

Theorem 10.1 (Erdős (unpublished), Bondy–Simonovits 1974)**.** *For every $\ell \geqslant 2$, we have*

$$\mathrm{ex}(n, C_{2\ell}) \leqslant O\left(n^{1+1/\ell}\right).$$

We won't quite prove this, but we'll prove a slightly weaker statement that is interesting in its own right; by being a bit more careful with essentially the same proof, one can prove Theorem 10.1. Recall from the homework that if $\mathcal{F}$ is a collection of graphs, then we say that G is $\mathcal{F}$-free if G contains no copy of any $H \in \mathcal{F}$, and we write

$$\mathrm{ex}(n, \mathcal{F}) = \max\{e(G) : G \text{ is an } n\text{-vertex } \mathcal{F}\text{-free graph}\}.$$

Theorem 10.2. *For every $\ell \geqslant 2$, if G is an n-vertex graph containing no cycle of length at most 2ℓ, then*

$$e(G) \leqslant O(n^{1+1/\ell}).$$

In other words,

$$\mathrm{ex}(n, \{C_3, C_4, \ldots, C_{2\ell}\}) \leqslant O(n^{1+1/\ell}).$$

In the course of the proof, we will need the following extremely useful lemma which you already encountered on the homework.

Lemma 10.3. *If G is an n-vertex graph, then it has a subgraph G' with minimum degree at least $e(G)/n$.*

Proof. We repeatedly delete from G any vertex of degree strictly less than $e(G)/n$, until we get stuck. Note that we delete at most n vertices, and each time we delete a vertex we delete strictly fewer than $e(G)/n$ edges, hence the total number of edges we delete is strictly less than $e(G)$. That is, when we stop the process, we necessarily have at least one edge left. In particular, when we stop the process, we have some non-empty subgraph G' of G. Since we stopped the process at G', every vertex in G' must have degree at least $e(G)/n$, as claimed. □

We are now ready to prove Theorem 10.2.

Proof of Theorem 10.2. Let G be an n-vertex graph with at least $Cn^{1+1/\ell}$ edges, for some constant C that we'll pick later. Our goal is to show that G must contain a cycle of length at most 2ℓ. By Lemma 10.3, there is a subgraph G' of G with minimum degree at least $e(G)/n \geqslant Cn^{1/\ell}$. Let $m \leqslant n$ be the number of vertices of G', and let $d \geqslant Cn^{1/\ell}$ be the minimum degree of G'.

We now fix some vertex $v \in V(G)$. Let $N^1(v)$ denote the set of neighbors of v, so that $|N^1(v)| \geqslant d$. If there were some edge inside the set $N^1(v)$, it would form a triangle together with v; we would thus find a cycle of length 3 in G', and hence also in G, and we'd be done. We may therefore assume that there are no edges in $N^1(v)$. Since every vertex in $N^1(v)$ has degree at least d, it must have at least $d-1$ neighbors outside of $\{v\} \cup N^1(v)$.

Moreover, suppose that we had $x, y \in N^1(v)$ which had some common neighbor $z \neq v$. Then the vertices v, x, z, y would form a cycle of length 4 in G', and we'd again be done. Therefore, we may assume that all neighbors (apart from v) of vertices in $N^1(v)$ are distinct. In other words, if we define $N^2(v)$ to be the set of vertices at distance 2 from v—that is, the set of neighbors of vertices in $N^1(v)$—then we have that

$$|N^2(v)| \geqslant d(d-1),$$

since there are at least d vertices in $N^1(v)$, each of which contributes at least $d-1$ vertices in $N^2(v)$, and we cannot have over-counted.

We now continue in this way, defining $N^3(v), \dots, N^\ell(v)$. At every step of the process, we cannot have any any edges inside any of these sets, nor any collisions: no two vertices in $N^i(v)$ can have a common neighbor in $N^{i+1}(v)$, as this would yield a cycle of length at most 2ℓ in G'. This implies that the size of these sets grows by at least a factor of $d-1$ at every step, that is,

$$|N^i(v)| \geqslant d(d-1)^{i-1}$$

for every $i \leqslant \ell$. In particular,

$$|N^\ell(v)| \geqslant d(d-1)^{\ell-1} \geqslant \Omega(d^\ell) \geqslant \Omega(C^\ell n).$$

However, we also trivially have the bound $|N^\ell(v)| \leqslant m \leqslant n$, since $N^\ell(v)$ is a subset of $V(G')$. This is a contradiction if we pick C sufficiently large. □

As with complete bipartite graphs, it is again widely conjectured that the bound in Theorems 10.1 and 10.2 are tight, but it is only known to be tight in case $\ell \in \{2, 3, 5\}$. This is pretty remarkable: we know that

$$\mathrm{ex}(n, C_4) = \Theta(n^{3/2}) \qquad \mathrm{ex}(n, C_6) = \Theta(n^{4/3}) \qquad \mathrm{ex}(n, C_{10}) = \Theta(n^{6/5})$$

but we have no idea what the value of $\mathrm{ex}(n, C_8)$ is! You will see the construction providing this lower bound on the homework.

To end this section, let me just mention two remarkable conjectures of Erdős and Simonovits, which roughly say that the behavior of $\mathrm{ex}(n, H)$ for general bipartite H is very complicated.

Conjecture 10.4 (Erdős–Simonovits rational exponents conjecture)**.** *For every bipartite H, there exists some rational number $\alpha \in [1, 2)$ so that*

$$\mathrm{ex}(n, H) = \Theta(n^\alpha).$$

Moreover, the converse holds: for every rational $\alpha \in [1, 2)$, there exists some bipartite graph H so that

$$\mathrm{ex}(n, H) = \Theta(n^\alpha)$$

The first part of this conjecture is doubted by some experts, though no one has any idea how to prove or disprove it. However, the second part of the conjecture—that there exists a graph for any rational α—is widely believed to be true, and we are in fact fairly close to proving it. Every few months, a new paper appears finding a new infinite set of rational numbers that are now known to be "achievable", i.e. to be the exponent of $\mathrm{ex}(n, H)$ for some bipartite H.

Moreover, a slight weakening of the second part of the conjecture was recently proved by Bukh and Conlon.

Theorem 10.5 (Bukh–Conlon 2018)**.** *For every rational $\alpha \in [1, 2)$, there exists some finite collection $\mathcal{F}$ of bipartite graphs for which*

$$\mathrm{ex}(n, \mathcal{F}) = \Theta(n^\alpha).$$

11 Stability

Many problems in extremal combinatorics exhibit a phenomenon called *stability*, first discovered by Erdős and Simonovits. Stability is a stronger statement than the determination of an extremal structure: it states that if some structure has *close to* the maximum size given some constraint, then it must be *close to*, in a structural sense, the extremal construction.

In extremal graph theory, the stability form of Turán's theorem says something like the following: if G is an n-vertex K_r-free graph such that $e(G)$ is not much less than $t_{r-1}(n)$, then one can add or delete a small number of edges to G to obtain the Turán graph $T_{r-1}(n)$. Sometimes stating such results formally is a bit cumbersome, but in a moment we'll see a very clean statement along these lines.

In addition to being interesting in their own right, stability results are extremely useful for many applications. Indeed, we will shortly see a proof of (a special case of) Theorem 9.3 using stability; this proof gives a flavor of how stability arguments are often used in extremal graph theory.

But first, here is a formal statement of the stability version of Turán's theorem.

Theorem 11.1 (Füredi 2015)**.** *Let G be an n-vertex K_r-free graph, and suppose that*

$$e(G) \geqslant t_{r-1}(n) - s,$$

for some integer $s \geqslant 0$. Then we may delete at most s edges from G to obtain an $(r-1)$-partite graph.

Note that the case $s = 0$ recovers Turán's theorem. In most applications, however, one usually takes s to be something like εn^2 for a small constant ε. We will only prove Theorem 11.1 in the special case $r = 3$, and you will prove the general case on the homework.

Proof of Theorem 11.1 for $r = 3$. Let v be a vertex of maximum degree in G. Let $B = N(v)$ be the neighborhood of v, and let $A = V(G) \setminus B$. Since v had maximum degree, we have that $\deg(w) \leqslant |B|$ for every vertex w. In particular, applying this to all vertices in A, we find that

$$\sum_{w \in A} \deg(w) \leqslant |A||B|.$$

Moreover, since $A \cup B$ is a partition of $V(G)$ into two parts, we have that $|A||B| \leqslant \lfloor n^2/4 \rfloor = t_2(n)$. On the other hand, the sum $\sum_{w \in A} \deg(w)$ counts every edge between A and B exactly once, and every edge inside A exactly twice, once for each endpoint. Therefore,

$$\sum_{w \in A} \deg(w) = e(A, B) + 2e(A).$$

Finally, we note that B is an independent set, since any edge inside B would form a triangle together with v, and we assumed that G is triangle-free. As a consequence, every edge in G lies either in A or between A and B. Putting this all together, we find that

$$\begin{aligned} e(G) &= e(A, B) + e(A) \\ &= (e(A, B) + 2e(A)) - e(A) \\ &= \left(\sum_{w \in A} \deg(w) \right) - e(A) \\ &\leqslant |A||B| - e(A) \\ &\leqslant t_2(n) - e(A). \end{aligned}$$

Rearranging, this says that

$$e(A) \leqslant t_2(n) - e(G) \leqslant s,$$

where the final inequality is our assumption on G. In other words, we have found a partition of $V(G)$ into two parts A, B, such that B is an independent set and A contains at most s edges. By deleting these s edges, we obtain a bipartite graph, completing the proof. □

While Theorem 11.1 only gives stability for Turán's theorem (i.e. gives an approximate structure for K_r-free graphs with many edges), one can deduce from it a stability version of the Erdős–Stone–Simonovits theorem, i.e. stability for H-free graphs for any non-bipartite H. You'll prove this on the homework. In the meantime, here is one such statement for $H = C_5$.

Proposition 11.2. *There is an absolute constant C such that the following holds. If G is an n-vertex C_5-free graph with $e(G) \geqslant \lfloor n^2/4 \rfloor$, then we may delete at most $Cn^{5/3}$ edges from G to obtain a bipartite graph.*

Proof. Let F be the subgraph of G consisting of all edges of G which lie on a triangle. We claim that F is $K_{3,3}$-free. Indeed, suppose that there were a $K_{3,3}$ in F, say with vertices $a_1, a_2, a_3, b_1, b_2, b_3$, where each a is adjacent to each b in F. As every edge in F lies on a triangle in G, there is some $c \in V(G)$ such that a_1, b_1, c form a triangle. If $c \notin \{a_2, b_2\}$, then we obtain a C_5 in G, namely $c \sim a_1 \sim b_2 \sim a_2 \sim b_1 \sim c$. Similarly, if $c \notin \{a_3, b_3\}$, we obtain a C_5 for the same reason, interchanging the roles of 2 and 3. As c cannot lie in both $\{a_2, b_2\}$ and $\{a_3, b_3\}$, we have proved that F is $K_{3,3}$-free. Consequently, $e(F) \leqslant \text{ex}(n, K_{3,3}) \leqslant cn^{5/3}$, for some absolute constant c, by Theorem 5.3.

That is, at most $cn^{5/3}$ edges of G lie on a triangle. By deleting these edges, we obtain a triangle-free subgraph G' of G, with $e(G') \geqslant e(G) - cn^{5/3} \geqslant t_2(n) - cn^{5/3}$. Now applying Theorem 11.1 to G' with $s = cn^{5/3}$, we conclude that we can delete at most $cn^{5/3}$ additional edges from G' to obtain a bipartite graph. Putting this together, we find that G can be made bipartite by deleting at most $2cn^{5/3}$ edges, proving the result by setting $C = 2c$. $\square$

We now turn to an application of the stability method, namely the proof of Theorem 9.3 in the case $H = C_5$. In order to keep things relatively simple, we will also assume a minimum degree condition, which can be removed with a fairly straightforward, but tedious, argument (see the homework). Here is the precise statement we will prove.

Proposition 11.3. *There is some integer n_0 such that for all $n \geqslant n_0$, the following holds. If G is an n-vertex C_5-free graph with minimum degree at least $n/3$, then $e(G) \leqslant \lfloor n^2/4 \rfloor$, with equality if and only if $G = T_2(n)$.*

Proof. We pick n_0 so that for all $n \geqslant n_0$, we have $Cn^{5/3} \leqslant n/1000$, where C is the constant from Proposition 11.2. By Proposition 11.2, we know that we can make G bipartite by deleting at most $Cn^{5/3}$ edges. That is, there exists a partition $V(G) = A \cup B$ such that $e(A) + e(B) \leqslant Cn^{5/3}$. We pick an optimal such partition, that is, a partition that minimizes $e(A) + e(B)$. We note that every vertex $a \in A$ has at least as many neighbors in B as in A. For if this were not the case, we could move a to B and strictly decrease $e(A) + e(B)$, contradicting our choice of the optimal partition. Similarly, every vertex $b \in B$ has at least as many neighbors in A as in B.

First, we note that we may assume $\frac{2}{5}n < |A|, |B| < \frac{3}{5}n$. For if this is not the case, say if $|A| < \frac{2}{5}n$, then $|A||B| \leqslant (\frac{2}{5}n)(\frac{3}{5}n) = \frac{6}{25}n^2$. But then

$$e(G) = e(A) + e(B) + e(A, B) \leqslant e(A) + e(B) + |A||B| \leqslant \left(\frac{1}{1000} + \frac{6}{25}\right) n^2 < \left\lfloor \frac{n^2}{4} \right\rfloor,$$

which is what we wanted to prove.

Next, suppose that some vertex $a \in A$ has at least $n/30$ neighbors in A. By the observation above, it also has at least $n/30$ neighbors in B. Let A', B' be its set of neighbors in A, B, respectively. Then the subgraph induced by A' and B' cannot contain a four-vertex path P_4, as such a path would yield a C_5 together with a. By an exercise on the homework,

$$e(A', B') \leqslant \text{ex}\,(n, P_4) \leqslant n.$$

Note that in a complete bipartite graph, A' and B' would have $|A'||B'| \geqslant n^2/900$ edges between them, so there are a ton of "missing edges". More precisely, we see that

$$e(A, B) \leqslant |A||B| - |A'||B'| + e(A', B') \leqslant \frac{n^2}{4} - \frac{n^2}{900} + n.$$

This again implies that

$$e(G) = e(A) + e(B) + e(A, B) < \left(\frac{1}{1000} + \frac{1}{4} - \frac{1}{900}\right) n^2 + n < \left\lfloor \frac{n^2}{4} \right\rfloor,$$

as claimed, where we use the fact that $n \geqslant n_0$ is sufficiently large in the final inequality.

So we may assume that every vertex in A has fewer than $n/30$ neighbors in A. By the minimum degree assumption, it has at least $n/3$ neighbors in total, hence at least $3n/10$ neighbors in B. This implies that any two vertices in A must have a common neighbor in B; indeed, if they did not, then their neighborhoods would be two disjoint subsets of B each of size at least $3n/10$, which is impossible since $|B| < \frac{3}{5}n = 2 \cdot (\frac{3n}{10})$.

We now claim that B is an independent set. Indeed, suppose that there were some edge bb', for $b, b' \in B$. Pick neighbors a, a' of b, b'. By the argument above, a and a' have a common neighbor in B, say b''. But then $b \sim a \sim b'' \sim a' \sim b' \sim b$ is a C_5, a contradiction. The exact same argument, but interchanging the roles of A and B, shows that A is an independent set.

Thus, we have boosted the stability result we started with to an exact structural description: G *is* bipartite, with parts A and B. But now, the only way for G to have the maximum possible number of edges is if $G = T_2(n)$, proving the claim. □