Math 301: The Erdős-Stone-Simonovitz Theorem and Extremal Numbers for Bipartite Graphs

Mary Radcliffe

1 The Erdős-Stone-Simonovitz Theorem

Recall, in class we proved Turán's Graph Theorem, namely

Theorem 1 (Turán's Theorem). *Let $k \geq 3$. Then $ex(n, K_k) = \left\lfloor \frac{(k-2)n^2}{2(k-1)} \right\rfloor$.*

That is to say, if G is a graph having at least $\left\lfloor \frac{(k-2)n^2}{2(k-1)} \right\rfloor + 1$ edges, then G contains K_k as a subgraph. On the other hand, there exist graphs having $\left\lfloor \frac{(k-2)n^2}{2(k-1)} \right\rfloor$ edges that do not contain K_k as a subgraph. Moreover, we showed that the extremal graphs (that is, the graphs achieving the bound) are the complete balanced $(k-1)$-partite graphs (shown in Figure 1). As these will come up again, we shall denote this graph by $T_{k-1}(n)$, called a Turán graph.

Our goal for the next few theorems is to extend this result as best we can to other classes of graphs. First, let's think about what was special about the Turán graphs. These graphs work for our situation because we can group the vertices up into $k-1$ independent sets, and the Pigeonhole Principle allows us to say that if a K_k were present, it would have to have two vertices in one of these independent sets (which is, of course, impossible).

So perhaps, then, what is special about Turán graphs is the independent sets. That is to say, it is certainly clear that any graph whose vertex set can be written as a union of $k-1$ independent sets can contain no copy of K_k. Perhaps this is the key to unlocking extensions of this extremal number.

So when can a graph be written as a union of $k-1$ independent sets?

Lemma 1. *Let G be a graph on n vertices. The vertex set of G can be partitioned into $k-1$ independent sets if and only if G can be properly colored with at most $k-1$ colors.*

The proof of this lemma is essentially trivial. Note that if we properly color the vertices of G, then any two vertices having the same color must be independent. Taking the partition as exactly the color classes will yield the result, and vice versa: given a partition as above, we may color each independent set with one color to obtain a $(k-1)$-coloring.

Fundamentally, this is the problem with finding a K_k inside of $T_{k-1}(n)$: the Turán graph is $(k-1)$-colorable, so any subgraph of $T_{k-1}(n)$ is also $(k-1)$-colorable. But K_k is not. This leads to the following immediate generalization:

Theorem 2 (Erdős-Stone-Simonovitz Theorem). *Let H be a graph with chromatic number $\chi(H) = k \geq 3$. Then $ex(n, H) \leq \frac{(k-1)n^2}{2(k-1)} + o(n^2)$; i.e., the extremal number for H is asymptotically equal to the number of edges in $T_{k-1}(n)$.*

Furthermore, it should be clear that the extremal graphs will be the Turán graphs. (If that is not clear, prove it as an exercise. Use induction.)

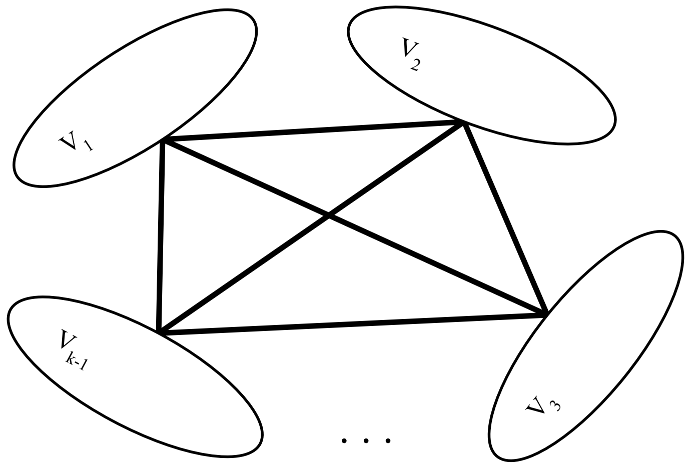

Figure 1: The Turán Graph. Here, we assume that $|V_i|$ and $|V_j|$ are sets of vertices, whose sizes differ by at most 1 for any i, j, and the solid lines between sets of vertices indicate that all edges are present. This graph is a complete $(k-1)$-partite graph, in that we have all edges between partition sets, and is balanced in the sense that the sizes of the vertex partition are balanced.

The proof of this theorem is messy, so before we get into the details, let's sketch it out. Since we are looking at an asymptotic result, we actually need to show that for any $\epsilon > 0$, there exists some N sufficiently large that if $n > N$, then $ex(n, H) \le \frac{(k-1)n^2}{2(k-1)} + \epsilon n^2$. Since ϵ is arbitrarily small, we essentially have written that $ex(n, H) \le \frac{(k-1)n^2}{2(k-1)} + o\big(n^2\big)$, just as we wanted.

So here will be our strategy:

1. Choose some small ϵ.
2. Look at a graph having n vertices and at least $\frac{(k-1)n^2}{2(k-1)} + \epsilon n^2$ edges.
3. Show that if n is large enough, we can find, as a subgraph here, the complete k-partite graph having partite sets of size at least t (for any t. We will specify later). Actually, we will show it contains a copy of $T_k(kt)$ as a subgraph (which is the same thing).
4. Notice that since H is k-colorable, as long as t is larger then the size of each color class, we can embed h inside $T_k(kt)$, by embedding each color class in one of the partite sets, and then using whichever edges we need.
5. Realize that these steps are enough to prove the statement: we have shown that in any graph G having at least $\frac{(k-1)n^2}{2(k-1)} + \epsilon n^2$ (with n large enough), we can find a copy of H by lazily picking $t = |V(H)|$.
6. WIN THE GAME!!!!

Certainly, the difficulty of this strategy lies entirely in Step 3. This will be the bulk of the work in proving the theorem. Since this part of the proof is hairy, we shall split it up into some lemmas.

Lemma 2. *Fix $\epsilon > 0$, and let G be a graph on n vertices having at least $\frac{(k-1)n^2}{2k-2} + \epsilon n^2$ edges. Then G contains a subgraph D having at least $m = \sqrt{\epsilon(k-1)}n$ vertices, such that the degree of every vertex in D is at least $\left(\frac{k-2}{k-1} + \epsilon\right) m + 1$.*

Proof. We shall produce D using the following algorithm:

1. Initialize: $D_0 = G$, $n_0 = n$.
2. Given D_i, n_i: let $v \in V(D_i)$ be a vertex in D_i having degree $\leq \left(\frac{k-2}{k-1} + \epsilon\right) n_i$
3. Define $D_{i+1} = D_i \backslash \{v\}$; that is, D_{i+1} is obtained from D_i by removing the vertex v. Define $n_{i+1} = |V(D_{i+1})| = n_i - 1 = n - (i+1)$.
4. Repeat steps 2 and 3 until no vertex v satisfies 2.
5. Output D_i.

Clearly, this algorithm will terminate eventually, since eventually we will either have a graph whose vertices all have a large degree, or a graph with no vertices. We'd really like it to be the first thing, so let's consider how long the algorithm will run for. To do so, let's look carefully at the i^{th} step.

Notice that in the $(i+1)^{\text{st}}$ step, we remove at most $\left(\frac{k-2}{k-1} + \epsilon\right) n_i$ edges from D_i. Hence, the subgraph D_i is obtained by removing at most

$$\begin{aligned}
\sum_{j=0}^{i-1} \left(\frac{k-2}{k-1} + \epsilon\right) n_j &= \sum_{j=0}^{i-1} \left(\frac{k-2}{k-1} + \epsilon\right)(n-j) \\
&= \left(\frac{k-2}{k-1} + \epsilon\right) \sum_{j=0}^{i-1} (n-j) \\
&= \left(\frac{k-2}{k-1} + \epsilon\right) \left(ni - \frac{i(i-1)}{2}\right) \\
&= \left(\frac{k-2}{k-1} + \epsilon\right) \frac{2n(n-n_i) - (n-n_i)(n-n_i-1)}{2} && \text{since } i = n - n_i \\
&= \left(\frac{k-2}{k-1} + \epsilon\right) \frac{n^2 - n_i^2 + n - n_i}{2}
\end{aligned}$$

edges from G. Also, the total number of edges in D_i is at most $\binom{n_i}{2} \leq \frac{n_i^2}{2}$. Therefore, the total number of edges in G is at most

$$\frac{n_i^2}{2} + \left(\frac{k-2}{k-1} + \epsilon\right) \frac{n^2 - n_i^2 + n - n_i}{2}.$$

On the other hand, the hypothesis was that G contains at least $\frac{(k-1)n^2}{2(k-1)} + \epsilon n^2$ edges. Therefore, it must be that

$$\frac{(k-1)n^2}{2(k-1)} + \epsilon n^2 \leq \frac{n_i^2}{2} + \left(\frac{k-2}{k-1} + \epsilon\right) \frac{n^2 - n_i^2 + n - n_i}{2}.$$

Therefore, this process must stop whenever the above inequality fails to hold. That is, the process terminates when

$$\begin{aligned}
\frac{(k-1)n^2}{2k-2} + \epsilon n^2 &> \frac{n_i^2}{2} + \left(\frac{k-2}{k-1} + \epsilon\right) \frac{n^2 - n_i^2 + n - n_i}{2} \\
n^2 \left(\frac{k-2}{2(k-1)} + \epsilon - \frac{k-2}{2(k-1)} - \frac{\epsilon}{2}\right) - \frac{n}{2}\left(\frac{k-2}{k-1} + \epsilon\right) &> n_i^2 \left(\frac{1}{2} - \frac{k-2}{2(k-1)} - \frac{\epsilon}{2}\right) - \frac{n_i}{2}\left(\frac{k-2}{k-1} + \epsilon\right) \\
n^2\epsilon - n\left(\frac{k-2}{k-1} + \epsilon\right) &> n_i^2 \left(\frac{1}{k-1} - \epsilon\right) - n_i \left(\frac{k-2}{k-1} + \epsilon\right).
\end{aligned}$$

Notice, the left hand side of this inequality is just a constant (when G is fixed, which it is), and the left hand side is a parabola that decreases when n_i decreases. If you use the quadratic formula, you will find that this inequality is true so long as $n_i < \sqrt{\epsilon(k-1)}n$. That is to say, after at most $n\left(1 - \sqrt{\epsilon(k-1)}\right)$ steps, the algorithm must terminate, and thus the result holds. $\square$

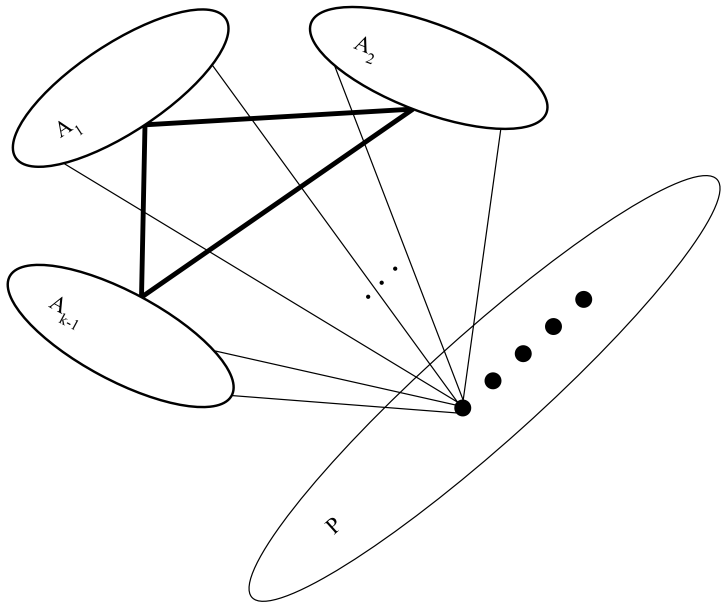

Figure 2: The set P defined in the proof of Lemma 3. Note that if we consider a vertex in P, it has a lot of edges (namely, t edges) into each set A_i. However, the edges need not always be to the same vertices.

Lemma 3. *Let $\epsilon > 0$, $k \geq 1$, and let G be a graph on n such that every vertex of G has degree at least $\left(\frac{k-2}{k-1} + \epsilon\right) n$. Fix $t > 0$. Then if n is sufficiently large, $T_k(kt)$ is a subgraph of G.*

Proof. We shall work by induction on k. Note that the case of $k = 1$ is trivial, since $T_1(t)$ is the empty graph on t vertices.

Suppose the result is known for $k - 1$ for any choice of t. Let $s = \left\lceil \frac{t}{\epsilon} \right\rceil$. Notice that $\frac{k-2}{k-1} = 1 - \frac{1}{k-1} \leq 1 - \frac{1}{k-2}$, and hence G has sufficiently many edges to apply the induction hypothesis. Choose n large enough that we can find a copy of $T_{k-1}((k-1)s)$ in G. Label the partite sets as $A_1, A_2, \ldots, A_{k-1}$, and let $A = A_1 \cup A_2 \cup \cdots \cup A_{k-1}$. Let W be the set of vertices in $V(G)$ that do not appear in A.

Define $P = \{v \in W \mid N(v) \cap A_i \geq t \text{ for every } i\}$; that is, P is the set of vertices that have at least t edges into each A_i. This is illustrated in Figure 2. Notice that our goal now is to construct a set $B_k \subset P$, such that for each i, there is a subset $B_i \subset A_i$ with every possible edge between B_k and B_i, and $|B_i| = t$. This will be exactly a $T_k(kt)$. We have restricted to P, since these are the only possible vertices to be included in such a set B_k. See Figure 3.

Most of the rest of the proof comes down to counting. First, let us count the number of nonedges between W and A.

First, if a vertex $v \in W$ is not in P, then there exists some i such that the number of edges from v to A_i is at most $t - 1$. Hence, the number of nonedges between W and A is at least $(|W| - |P|)(s - t + 1) \geq (|W| - |P|)(s - t)$.

On the other hand, the degree of every vertex in G is at least $\left(\frac{k-2}{k-1} + \epsilon\right) n$, and hence given any vertex v, the number of nonedges involving v is at most $n - \left(\frac{k-2}{k-1} + \epsilon\right) n = n\left(\frac{1}{k-1} - \epsilon\right)$. Summing over all vertices in A, we thus obtain that the number of nonedges between W and A is at most $s(k-1)n\left(\frac{1}{k-1} - \epsilon\right)$.

Taking these two bounds together yields

$$s(k-1)n\left(\frac{1}{k-1} - \epsilon\right) \geq (|W| - |P|)(s - t).$$

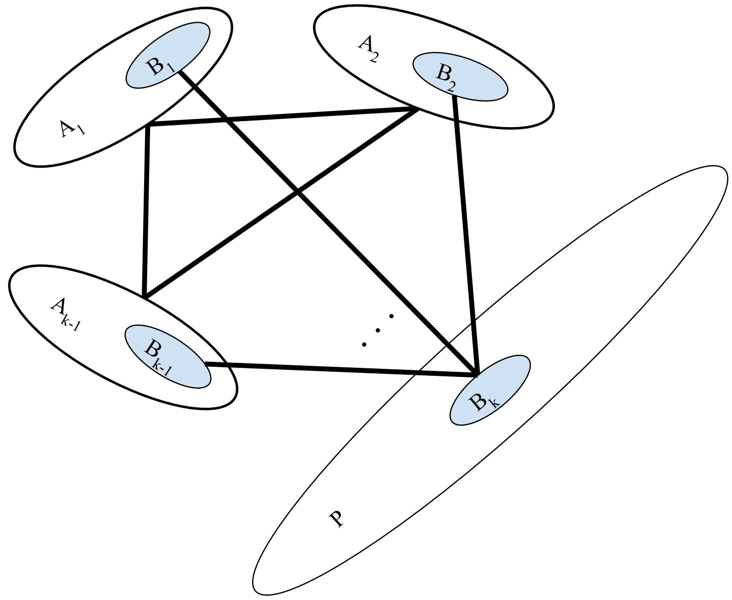

Figure 3: The goal of the proof: to find a (size t) subset B_i in each A_i, and a set B_k in P, with all the edges between B_k and each B_i. Notice that we already have all the edges between B_i and B_j when $i \neq j$ and both $i, j < k$, since they are subsets of the A_i sets.

Recalling that $s = \frac{t}{\epsilon}$ and $|W| = n - s(k-1)$, this yields

$$t(k-1)n\left(\frac{1}{k-1} - \epsilon\right) \geq (n - s(k-1) - |P|)(t - \epsilon t).$$

Solving this inequality for $|P|$ and using lots of algebra, we obtain

$$|P| \geq n\left(\frac{\epsilon(k-2)}{1-\epsilon}\right) - s(k-1).$$

The important thing here is that $|P|$ is monotonically increasing with n, even as $|A|$ stays the same. Hence, we may choose $|P|$ as large as we like. Moreover, note that by the pigeonhole principle, as long as $|P|$ is large enough, we will be able to find a set B_k in P satisfying the desired properties. Hence, by taking n sufficiently large, we can produce such B_i, and the result is demonstrated.

□

Note that these two lemmas immediately prove the Erdős-Stone-Simonovitz theorem, as we can first find a good subgraph D in G by Lemma 2, and then find a $T_k(kt)$ inside D (and hence inside G) using Lemma 3. Once we have a $T_k(kt)$ as a subgraph of G, applying the observation made in the proof sketch then yields the result.

2 Extremal numbers for bipartite graphs

Note that the Erdős-Stone-Simonovits Theorem works for any graph H whose chromatic number is at least 3. This leaves out the entire class of graphs having chromatic number exactly two. Recall the following:

Lemma 4. *A graph G has chromatic number 2 if and only if G is bipartite.*

Hence, we are left asking, how do we compute extremal numbers for bipartite graphs?

The answer, unfortunately, is a lot of *we don't know.* We do have the following bound for any bipartite graphs:

Theorem 3. *Let H be a bipartite graph having v_H vertices. Then there exists a constant c (depending on H) such that $ex(n, H) \leq cn^{2-1/v_H}$.*

To prove this bound, we shall make use of the following binomial identity. While we will not prove this here, you should be able to prove this yourself if you know how to use Janson's inequality (you can prove that $\binom{n}{r}$ is convex in n.

Lemma 5. *Let $a_1, a_2, \ldots, a_N$ be positive integers, with average $A = \frac{1}{N}\sum a_i$. Then $\sum \binom{a_i}{r} \geq N\binom{A}{r}$.*

That is to say, replacing each $\binom{a_i}{r}$ with a copy of $\binom{A}{r}$ only increases the sum.

Proof. For simplicity let's write $r = v_H$. Notice that if $K_{r,r}$ is a subgraph of G, then so is H, since we may embed H just as we did in the Turá n graphs for the proof of the Erdős-Stone-Simonovits Theorem.

Now, suppose that for each $v \in V(G)$, we take S_v to be a set of r neighbors of v. Note that if there exists a set of r vertices, say A, such that for all $u, v \in A$, $S_u = S_v$, then we can form a copy of $K_{r,r}$ inside of v having one partite set as A, and the other partite set as S_v for any $v \in A$. Hence, it must be that this construction is impossible.

Therefore, $\sum_{v\in V(G)} \binom{d_v}{r} \leq (r-1)\binom{n}{r}$, since no set in $\binom{[n]}{r}$ can be represented in the sum r or more times. Applying the bound from the previous Lemma, and setting $e(G)$ equal to the number of edges in G, we thus obtain

$$
\begin{aligned}
\sum_{v\in V(G)} \binom{d_v}{r} &\leq (r-1)\binom{n}{r} \\
n\binom{\frac{1}{n}\sum d_v}{r} &\leq (r-1)\binom{n}{r} \\
n\binom{2e(G)/n}{r} &\leq (r-1)\binom{n}{r}
\end{aligned}
$$

Note that $\binom{n}{r} \leq n^r$, and $\binom{2e(G)/n}{r} \geq \left(\frac{2e(G)}{nr}\right)^r$ (use Stirling's formula, for example). Therefore, we have

$$n\left(\frac{2e(G)}{nr}\right)^r \leq (r-1)n^r.$$

Solving for $e(G)$, we obtain $e(G) \leq \frac{r(r-1)^{1/r}}{2} n^{2-1/r}$. Note that the initial fraction is independent of n, but does depend upon H; it is this fraction that we take to be c in the statement of the theorem. $\square$

At this point, you might think "Awesome, we did it! We got a bound! DONE." And you're right to celebrate, this is useful, sort of. But imagine if H has only a few vertices, say 3. The exponent on n in this case is $n^{2-1/3} = n^{5/3}$. But the only bipartite graph on 3 vertices is the path of length 2, and we know that the extremal number for this path is $\lfloor n \rfloor 2$, which is a far, far cry from $n^{5/3}$.

In fact, **there is NO bipartite graph containing a cycle for which $ex(n, H)$ is known.** We know a few of them asymptotically, but we don't know the constants.

For an example, let's consider the first nontrivial case: $H = C_4$. By applying our theorem carefully (actually, by redoing the analysis in this somewhat special case), we obtain

$$ex(n, C_4) \leq \frac{1}{2}n^{3/2} + \frac{1}{2}n.$$

This upper bound is asymptotically tight in SOME cases. Specifically, it is known (analysis complicated) that

$$ex(n, C_4) = \frac{1}{2}n^{3/2} + o(n^{3/2}),$$

in the case that $n = p^2 + p + 1$ for some prime p. (I know, this n looks ridiculous. It comes from analyzing certain graphs built out of finite fields, which are always of order p^k for a prime p.) This construction is due to Füredi, and I'd be happy to show it to you in office hours. And in fact, in a construction due to Klein, we have that $ex(n, C_4) = \Theta(n^{3/2})$, that is, we are asymptotically on the right order here. And while it is conjectured that the constant $1/2$ is correct, we cannot prove it for any n other than some fancy forms involving primes (like the above).