Pair constructions for hypergraph Ramsey numbers

Talk by Xiaoyu He

Notes by Sanjana Das

September 22, 2023

This is based on joint work with David Conlon, Ben Gunby, Jacob Fox, Dhruv Mubayi, Andrew Suk, and Jacques Verstraete.

§1 Background

Today, we'll talk about lower bound constructions for the Ramsey numbers of 3-uniform hypergraphs. First, we'll give some background — we probably already think Ramsey numbers are interesting, but we'll now see why 3 is interesting, and what we know about lower bounds.

§1.1 Hypergraphs

Definition 1.1. A k-graph (or k-uniform hypergraph) is a generalization of a graph where edges are (unordered) k-tuples of vertices — in other words, a k-graph is an object $\mathcal{H} = (V, E)$ with $E \subseteq \binom{V}{k}$.

Here are some examples of hypergraphs.

Example 1.2

The complete k-uniform hypergraph on n vertices, denoted $K_n^{(k)}$, is the hypergraph where we have n vertices and an edge for *every* $\binom{n}{k}$ k-tuple.

Example 1.3

The complete k-partite k-graph $K_{n,\dots,n}^{(k)}$ is the blowup of a single edge — we have k vertex sets, each with n vertices, and all the edges consisting of one vertex from each set.

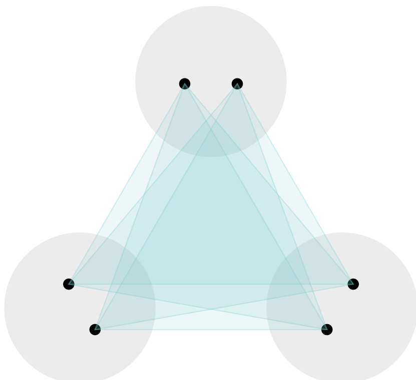

Example 1.4

The k-uniform tight cycle on s vertices, denoted $C_s^{(k)}$, is obtained by taking s points on a circle as our vertices and drawing an edge for every consecutive k points.

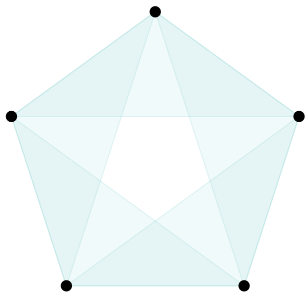

Example 1.5

The star $S_s^{(k)}$ consists of $s+1$ vertices and $\binom{s}{k-1}$ edges; we have a single distinguished vertex, and we draw an edge for all k-tuples containing that vertex.

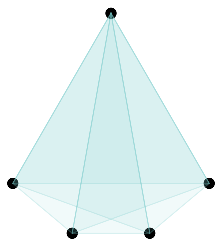

We'll look at the Ramsey numbers of these hypergraphs — we're most interested in Ramsey numbers of complete k-graphs, but the others are interesting as well (and we understand some of them better).

§1.2 Ramsey numbers

Definition 1.6. Given k-graphs $\mathcal{H}_1, \ldots, \mathcal{H}_r$, their Ramsey number $r(\mathcal{H}_1, \ldots, \mathcal{H}_r)$ is the smallest n such that any r-coloring of the edges of $K_n^{(k)}$ contains a monochromatic $\mathcal{H}_i$ in color i for some i.

Theorem 1.7 (Ramsey's theorem)

For any collection of hypergraphs $\mathcal{H}_1, \ldots, \mathcal{H}_r$, their Ramsey number $r(\mathcal{H}_1, \ldots, \mathcal{H}_r)$ is finite.

In other words, given any $\mathcal{H}_1, \ldots, \mathcal{H}_r$, there's some large number n that guarantees that no matter how we color our n-vertex complete k-graph, we'll find at least one of these monochromatic structures (in the correct color).

Question 1.8. How big are these Ramsey numbers quantitatively?

We're particularly interested in $r(K_s^{(k)}, K_t^{(k)})$ is; we'll denote this by $r_k(s,t)$ for convenience.

§1.3 Known bounds for graph Ramsey numbers

We'll now talk about the previously known bounds for hypergraph Ramsey numbers; we'll first start with *graph* Ramsey numbers (i.e., the case $k = 2$).

In the *diagonal* case — where we're interested in $r_2(t,t)$ — we now know that

$$\sqrt{2}^t \leq r_2(t,t) \leq 3.999^t.$$

(The upper bound is a spectacular recent breakthrough of Campos, Griffiths, Morris, and Sahasrabudhe.) So broadly speaking, we know that $r_2(t,t)$ is exponential in t, but we don't know the correct base of the exponent.

In the *off-diagonal* case, we fix s and take $t \to \infty$. For a long time, we were stuck at

$$t^{(s+1)/2} \lesssim r_2(s,t) \lesssim t^{s-1}$$

(we're dropping log factors in these bounds). But recently this year, we essentially solved $r_2(4,t)$ — Mattheus and Verstraete (in 2023) showed that

$$r_2(4,t) \asymp t^3.$$

§1.4 Known bounds for hypergraph Ramsey numbers

In the case of graph Ramsey numbers, we have a somewhat big gap between the upper and lower bounds — the two bounds have different bases of their exponents. But the gap for *hypergraphs* is actually much bigger. For $k = 3$, the best bounds we have are

$$2^{\Omega(t^2)} \leq r_3(t,t) \leq 2^{2^{O(t)}}.$$

(The lower bound, similarly to the $k = 2$ case, is probabilistic — we take a random coloring.)

So starting from $k = 3$, we don't even know the correct *tower height* — the two bounds have a gap of 1 in tower height.

But the good news is that this gap doesn't get any bigger as we increase the uniformity (i.e., we have a gap of 1 in tower height of 1 for *all* $k \geq 4$) — for all $k \geq 4$, we know

$$2^{r_{k-1}(\varepsilon t, \varepsilon t)} \leq r_k(t,t) \leq 2^{r_{k-1}(t,t)^{k-1}}$$

(for some ε). So we can bound $r_k(t,t)$ both above and below by something exponential in Ramsey numbers of the previous uniformity, which means the tower height increases by 1 when we increase uniformity. (The lower bound is due to Erdős–Hajnal — called the *stepping up lemma* — and the upper bound is due to Erdős–Rado.)

This means $k = 3$ is in some sense the critical case — if we can understand the tower height for $k = 3$, then we can step up to any higher uniformity. We'll focus on lower bounds because for a long time, people have believed that the upper bound for $r_3(t,t)$ should be the truth (i.e., $r_3(t,t)$ should be double-exponential). One reason to think this is that when we have *four* colors, then we *do* get something double-exponential — we know that

$$r_3(t,t,t,t) = 2^{2^{\Theta(t)}}$$

(this is due to Hajnal). So we don't know what happens for 2 colors, but we do know that we get double-exponential behavior for 4, and we might expect that the number of colors shouldn't affect the tower height. In fact, proving that $r_3(t,t)$ is double-exponential was one of Erdős's problems.

> **Question 1.9** (Erdős \$500). Prove that $r_3(t,t) = 2^{2^{\Theta(t)}}$.

Today, we'll talk about some ideas for improving lower bounds for 3-uniform Ramsey numbers.

§2 Pair constructions

The bound $r_3(t,t) \geq 2^{\Omega(t^2)}$ is proven using a random coloring. But there's another style of coloring that's often quite useful; we'll first see this construction for the off-diagonal case, but it turns out to actually also be useful for the diagonal case.

§2.1 The off-diagonal case

> **Theorem 2.1** (Conlon–Fox–Sudakov)
> We have $r_3(4,t) \geq 2^{\Omega(t \log t)}$.

This improves previous work of Erdős–Hajnal that showed $r_3(4,t) \geq 2^{\Omega(t)}$. We'll start by proving this result, to get some idea of how these proofs work.

> **Theorem 2.2** (Erdős–Hajnal)
> We have $r_3(4,t) \geq 2^{\Omega(t)}$.

Proof. We start with a random *tournament* T on $n = 2^{ct}$ vertices (for some c) — i.e., a random orientation of the complete graph. We then use this graph structure to define a hypergraph coloring χ of $\binom{[n]}{3}$. For each triple uvw, we look at uvw in our tournament, and there's two possibilities — it forms either a cyclic triangle or a transitive triangle. If it's a cyclic triangle then we color uvw red, and if it's a transitive triangle then we color it blue.

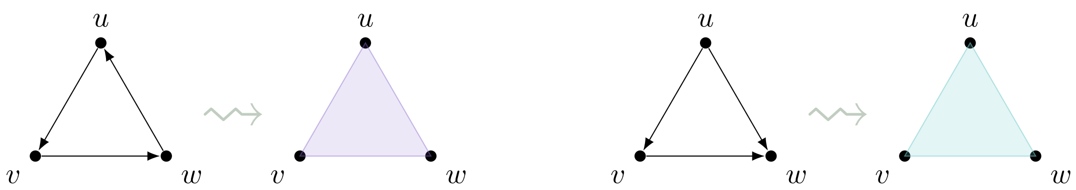

The beautiful property of this hypergraph coloring is that it can't have a red 4-clique, because we can't have too many cyclic triangles among four vertices — every vertex has either in-degree or out-degree 2, and if u has out-edges to both v and w, then uvw is not a cyclic triangle. This means we've *deterministically* guaranteed that there is no red 4-clique.

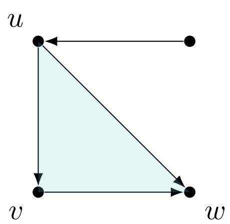

The randomness comes in to show that there's no big blue clique. A blue clique on t vertices in χ would correspond to a set of t vertices in our tournament for which all triples are transitive, meaning that the *entire* set of t vertices has to be transitive. This is very unlikely — for a given set of t vertices, the probability it is transitive is

$$\mathbb{P}[v_1, \ldots, v_t \text{ transitive}] = \frac{t!}{2^{\binom{t}{2}}}$$

(we first choose an order of the vertices, and then we must orient all edges according to it). The term of $t!$ is tiny (compared to the denominator), so we end up with the same computation as in the probabilistic proof for the *graph* lower bound. So with positive probability there's no blue t-clique in χ, and therefore we get $r_3(t,t) > n$. □

§2.2 Pair constructions and pair complexity

The idea of Conlon–Fox–Sudakov to get $2^{\Omega(t \log t)}$ is that there's actually a whole *family* of constructions you can build that look like this. We won't describe exactly what their construction is, but we'll give a general definition of the family we consider.

Definition 2.3. A 3-uniform pair construction is a coloring $\chi\colon \binom{[n]}{3} \to \{\text{red}, \text{blue}\}$ such that χ factors through two functions $f\colon \binom{[n]}{2} \to [p]$ and $g\colon [p]^3 \to \{\text{red}, \text{blue}\}$ — i.e., for all $u < v < w$ we have

$$\chi(u, v, w) = g(f(uv), f(vw), f(wu)).$$

In other words, a pair construction is a coloring that we can induce from a 2-uniform coloring — we color a *graph* with finitely many colors (corresponding to $f\colon \binom{[n]}{2} \to [p]$ — here p is the number of colors), and then use some deterministic rule to lift this to a hypergraph coloring (corresponding to $g\colon [p]^3 \to \{\text{red}, \text{blue}\}$).

We'll use $\chi_{f,g}$ to denote the pair construction corresponding to f and g.

Remark 2.4. This definition generalizes naturally to higher uniformities as well.

Example 2.5

The Erdős–Hajnal construction is a pair construction with $p = 2$ (we can represent a tournament as an edge-coloring of K_n, where we color an edge based on whether it's directed from the smaller to bigger or bigger to smaller vertex).

The construction by Conlon–Fox–Sudakov is also a pair construction, where p is logarithmic in n rather than constant.

Definition 2.6. Given a coloring $\chi\colon \binom{[n]}{3} \to \{\text{red}, \text{blue}\}$, the pair complexity of χ is the smallest p such that χ can be written as $\chi_{f,g}$ for some $f\colon \binom{[n]}{2} \to [p]$ and $g\colon [p]^3 \to \{\text{red}, \text{blue}\}$.

Note that *every* coloring χ is a pair construction for sufficiently large p — if p is large, then we can choose f such that from $(f(uv), f(vw), f(wu))$ we can read off the original triple (u, v, w), which allows us to define g. So we are really interested in what we can do with *small* pair complexity.

It turns out that most 3-uniform Ramsey constructions that we know of have constant or logarithmic pair complexity; in contrast, a random construction would have linear pair complexity.

Question 2.7. Can *all* 3-uniform Ramsey bounds be proven with a construction with 'small' pair complexity, or do we ever need a construction with e.g. linear pair complexity?

If we take 'small' to mean *constant*, then there's an easy answer — we can't do very much with constant pair complexity. The reason for this is that if we have a constant number of colors in the f-layer (i.e., our graph coloring), then we can use multicolor Ramsey to find monochromatic cliques there, which will produce monochromatic cliques in χ as well; this means the best constructions we can get this way will be exponential in t. (The Erdős–Hajnal construction used constant pair complexity and got an exponential bound; and here we've seen this is the best we can do.)

But with *logarithmic* pair complexity (where by 'logarithmic' we mean logarithmic in the number of vertices — i.e., $p \asymp \log n$), this doesn't happen; we might be able to get any bound we could want.

Remark 2.8. For comparison (to see what sorts of bounds we would *like* to prove), the best-known *upper* bound for $r_3(4,t)$ is

$$r_3(4,t) \leq 2^{O(t^2 \log t)}.$$

So we still have a gap between the upper and lower bounds; in particular, the possible range we have for $r_3(4,t)$ overlaps with the one for $r_3(t,t)$.

Remark 2.9. The construction by Conlon–Fox–Sudakov nicely interpolates between the off-diagonal and diagonal case as well — it shows that for *any* s and t, we have

$$r_3(s,t) \geq 2^{\Omega(st \log(t/s))}.$$

In some sense, this means the lower bound of $2^{\Omega(t^2)}$ in the diagonal case should be on the same level of 'difficulty' as the bound of $2^{\Omega(t \log t)}$ in the off-diagonal case (in the sense that both come from the same construction). So if we improve the bound in the off-diagonal case, we might be able to improve it in the diagonal case as well.

§2.3 Stepping up constructions

Another example of a construction with logarithmic complexity is the *stepping up* construction.

Suppose that χ is any stepping-up coloring, as defined in the morning — this means we have a $(k-1)$-uniform coloring of $[m]$ and want to define a k-uniform coloring of $\{0,1\}^m$. To do so, for $u, v \in \{0,1\}^m$ we define $\delta(u,v)$ as the first bit at which binary strings u and v differ. Then for each $v_1 < \cdots < v_k$, we define $\chi(v_1, \ldots, v_k)$ based on the pattern formed by $\delta(v_1, v_2)$, $\delta(v_2, v_3)$, …, $\delta(v_{k-1}, v_k)$ and the color of this $(k-1)$-tuple in the original coloring. This can be used to prove, for example, the lower bound

$$r_3(t,t,t,t) \geq 2^{2^{\Omega(t)}}$$

(the four-color double-exponential lower bound mentioned earlier).

Such constructions can be viewed as pair constructions, with $f(uv) = \delta(u,v)$; this means we have logarithmic pair complexity (since if there's $n = 2^m$ vertices, there's $m = \log n$ possible labels).

§2.4 Some results

One result in this direction (previous work with Jacob Fox) is a proof of the Conlon–Fox–Sudakov bound for sparser hypergraphs.

Theorem 2.10 (Fox–He)
For any $s \geq 3$, we have

$$r(S_s^{(3)}, K_t^{(3)}) \geq 2^{\Omega(st \log(t/s))}.$$

Recall that $S_s^{(3)}$ is obtained by taking a single vertex and only putting in edges that contain this vertex, as opposed to putting in *all* edges — this means it has only quadratically many edges, rather than cubically many. Still, we get the same lower bound as for $K_s^{(3)}$ in the Conlon–Fox–Sudakov result, even though $S_s^{(3)}$ is much sparser. We might intuitively expect that the Ramsey number for $K_s^{(3)}$ should be bigger in order than the one for $S_s^{(3)}$ (since $K_s^{(3)}$ is much denser, so it should be easier to avoid); so this is another reason to think that the Ramsey numbers for complete hypergraphs should be bigger than the current best lower bound.

This construction again has logarithmic pair complexity — there's some delicate way of coloring pairs that guarantees you never have a red star.

The authors have been trying to generalize this kind of construction as much as possible and see what we can do with it — what's the best possible bound we can prove? The following result is in some sense the most general lower bound we can expect to prove with this method:

Theorem 2.11

If H satisfies the property that every pair homomorphic image of H contains a Berge cycle, then

$$r(H, K_t^{(3)}) \geq 2^{\Omega(t \log t)}.$$

(The authors believe that this is exactly the property of H that characterizes when you can prove lower bounds of this type using a *random* pair construction — one where we choose f uniformly at random, and p is logarithmic.)

First, let's talk about what this property means.

Definition 2.12. A pair homomorphism $f\colon H \to G$ (for two 3-uniform hypergraphs H and G) is a function $f\colon \binom{v(H)}{2} \to v(G)$ such that if uvw is an edge of H, then $f(uv)f(vw)f(wu)$ is an edge of G.

(This is a new concept that's tailored towards understanding pair constructions. We're lying a bit in this definition — we actually need to order H and orient G. We'll sweep this under the rug; but some things will look trivial without it, and they don't look trivial when you take it into account.)

Example 2.13

There is a pair homomorphism from the star on 4 vertices to the 'loose triangle' on six vertices.

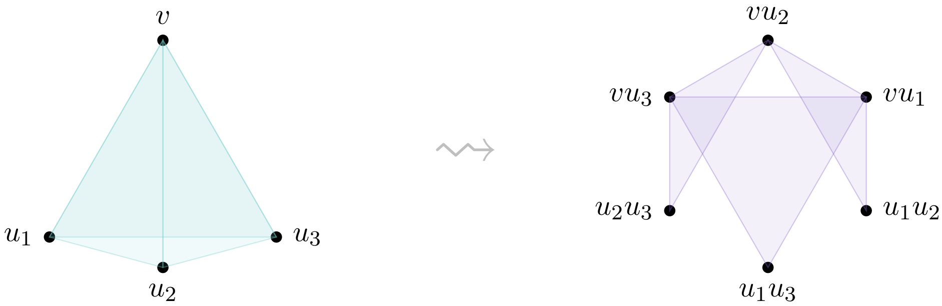

Pair homomorphisms generally take something tight and turn it into something loose.

Definition 2.14. A Berge cycle is a cycle which allows both tight and loose moves — in other words, you put points on a circle and you get to choose edges by taking certain consecutive triples, as long as every two adjacent edges overlap.

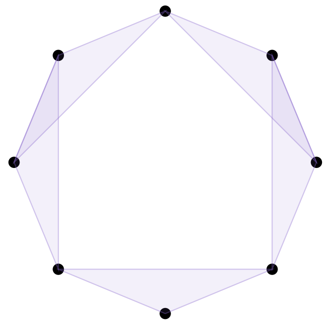

The idea behind Theorem 2.11 is that if every pair-homomorphic image of H contains a cycle, we can find a pair construction χ where g avoids all cycles of length at most $v(H)$ in red, and this will mean we won't have a red H (deterministically). Then showing that we don't have a blue $K_t^{(3)}$ is a probabilistic argument.

§3 Linear hypergraphs

Question 3.1. If H doesn't satisfy the property in Theorem 2.11, do we expect that the bound is false?

The authors thought that this might be the case, and then immediately disproved it — even though Theorem 2.11 is the limit of *one* kind of construction, we can still use other methods to get bounds for hypergraphs that don't satisfy this condition at all.

Definition 3.2. A hypergraph is linear if every two edges intersect in at most one vertex.

Theorem 3.3
For all $c > 1$, there exists a hypergraph H which is linear and such that $r(H, K_t^{(3)}) > 2^{(\log n)^c}$.

This doesn't quite answer Question 3.1. But the point is that we have a system for proving lower bounds using random pair constructions, and this system does nothing for linear hypergraphs — you can map a linear hypergraph to whatever you want under a pair homomorphism, even a single edge. So you can't hope to avoid a linear hypergraph using this kind of machinery; and yet we can still prove lower bounds better than the easy polynomial ones.

Remark 3.4. The construction for Theorem 3.3 is also a pair construction with logarithmic pair complexity (based on stepping up), but here f is very deterministic (unlike the random pair constructions from earlier).

§3.1 Some further questions

Question 3.5. Is Theorem 3.3 true for almost all linear hypergraphs H?

(More precisely, by 'almost all' we mean that you fix c and a number of vertices $n \gg c$, and sample from all linear hypergraphs with this number of vertices.)

The proof of Theorem 3.3 does involve sampling H randomly, but there we sample from the Erdős–Rényi graph $\mathcal{G}^{(3)}(n, \frac{1}{n})$ and then delete some edges. The authors believe that if actually sampling H uniformly at random would also work; but proving this is a bit tricky.

Here's another question, which is probably harder:

Question 3.6. Is Theorem 3.3 true for the Fano plane?

We still don't know superpolynomial lower bounds for the Fano plane — any lower bound $r(F, K_t^{(3)}) > t^{\omega(1)}$ would be interesting.

Finally, here's one final question; if we knew the answer, it'd probably tell us whether we should keep going with pair constructions or try something else altogether.

Question 3.7. We know that $t^{\Omega(1)} \leq r(C_5^{(3)} \setminus e, K_t^{(3)}) \leq 2^{O(t \log t)}$. Which is correct?

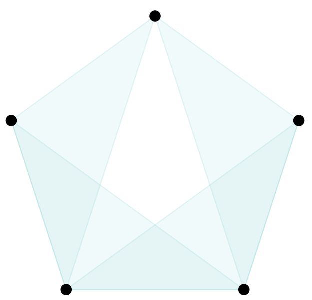

This is the only 5-vertex hypergraph for which we don't know whether the answer is polynomial or exponential. It looks exactly like a tight path, except that you identify the last two vertices; and pair homomorphisms don't se that. So somehow we need a different way of distinguishing the two.