

On off-diagonal hypergraph Ramsey numbers

David Conlon* Jacob Fox† Benjamin Gunby‡ Xiaoyu He§ Dhruv Mubayi¶
Andrew Suk^{||} Jacques Verstraëte**

Abstract

A fundamental problem in Ramsey theory is to determine the growth rate in terms of n of the Ramsey number $r(H, K_n^{(3)})$ of a fixed 3-uniform hypergraph H versus the complete 3-uniform hypergraph with n vertices. We study this problem, proving two main results. First, we show that for a broad class of H , including links of odd cycles and tight cycles of length not divisible by three, $r(H, K_n^{(3)}) \geq 2^{\Omega_H(n \log n)}$. This significantly generalizes and simplifies an earlier construction of Fox and He which handled the case of links of odd cycles and is sharp both in this case and for all but finitely many tight cycles of length not divisible by three. Second, disproving a folklore conjecture in the area, we show that there exists a linear hypergraph H for which $r(H, K_n^{(3)})$ is superpolynomial in n . This provides the first example of a separation between $r(H, K_n^{(3)})$ and $r(H, K_{n,n,n}^{(3)})$, since the latter is known to be polynomial in n when H is linear.

1 Introduction

Given two k -uniform hypergraphs (or, henceforth, k -graphs) G and H , their *Ramsey number* $r(G, H)$ is the smallest N such that whenever the edges of the complete k -graph $K_N^{(k)}$ are colored in red or blue, there must be either a red copy of G or a blue copy of H . That these numbers exist is the statement of Ramsey's original theorem [16], but many questions remain about their quantitative behavior.

The graph case $k = 2$ has received particular attention. A result of Erdős and Szekeres [8] from 1935 says that

$$r(K_s, K_n) \leq \binom{n+s-2}{s-1}$$

and one of the main driving forces in the area has been to decide if and when this bound is close to being tight. When $s = n$, their bound implies that $r(K_n, K_n) \leq 4^n$ and the problem of improving this bound by an exponential factor, which was recently resolved in a breakthrough paper of Campos, Griffiths, Morris and Sahasrabudhe [3], has occupied a central place in extremal combinatorics.

*Department of Mathematics, California Institute of Technology, Pasadena, CA 91125. Email: dconlon@caltech.edu. Research supported by NSF Award DMS-2054452.

†Department of Mathematics, Stanford University, Stanford, CA 94305. Email: jacobfox@stanford.edu. Research supported by NSF Award DMS-2154129.

‡Department of Mathematics, Rutgers University, Piscataway, NJ 08854. Email: bg570@rutgers.edu.

§Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: xiaoyuh@princeton.edu. Research supported by NSF Award DMS-2103154.

¶Department of Mathematics, Statistics and Computer Science, University of Illinois, Chicago, IL 60607. Email: mubayi@uic.edu. Research partially supported by NSF Awards DMS-1763317, DMS-1952767 and DMS-2153576, by a Humboldt Research Award and by a Simons Fellowship.

^{||}Department of Mathematics, University of California at San Diego, La Jolla, CA 92093. Email: asuk@ucsd.edu. Research supported by an NSF CAREER Award and by NSF Awards DMS-1952786 and DMS-2246847.

**Department of Mathematics, University of California at San Diego, La Jolla, CA 92093. Email: jacques@ucsd.edu. Research supported by NSF Award DMS-1800332.

The problem that we will be concerned with in this paper is more closely related to the estimation of the off-diagonal Ramsey numbers $r(K_s, K_n)$ where s is fixed and n tends to infinity. In this case, the Erdős–Szekeres bound $r(K_s, K_n) \leq n^{s-1}$ can be close to tight. For instance, by combining important results of Ajtai, Komlós and Szemerédi [1] and Kim [11], we have

$$r(K_3, K_n) = \Theta\left(\frac{n^2}{\log n}\right).$$

Until lately, no similar result was known for any other s , but another recent breakthrough result by Mattheus and Verstraëte [13] shows that $r(K_4, K_n) = \tilde{\Theta}(n^3)$ (where the tilde in the big-O notation means that the estimate is tight up to logarithmic factors). At present, no similar result is known for $s \geq 5$.

Despite the problems that remain in the graph case, even less is known about Ramsey numbers of hypergraphs. It has been known since the 1970s that there are positive constants c and C depending only on k such that

$$t_{k-1}(cn^2) \leq r(K_n^{(k)}, K_n^{(k)}) \leq t_k(Cn),$$

where the tower function is given by $t_1(x) = x$ and $t_{i+1}(x) = 2^{t_i(x)}$. However, the correct tower height in these bounds remains unknown for all $k \geq 3$. The case $k = 3$ is of particular importance, since the ingenious stepping-up lemma of Erdős and Hajnal (see, for instance, [10]) allows us, starting from $k = 3$, to construct lower bound colorings for uniformity $k + 1$ from colorings for uniformity k , gaining an extra exponential each time. In particular, if we could show that $r(K_n^{(3)}, K_n^{(3)})$ grows as a double exponential in n , this would essentially resolve the problem of estimating $r(K_n^{(k)}, K_n^{(k)})$ for all k .

Our concern in this paper will be with studying off-diagonal Ramsey numbers for 3-graphs. As above, variants of the stepping-up lemma (see, for example, [6, 15]) can be used to lift bounds on off-diagonal Ramsey numbers for 3-graphs to higher uniformities, so we will focus on the critical 3-uniform case. For $r(K_s^{(3)}, K_n^{(3)})$ with s fixed and n growing, the best known bounds, due to Conlon, Fox and Sudakov [4], are that

$$2^{cn \log n} \leq r(K_s^{(3)}, K_n^{(3)}) \leq 2^{Cn^{s-2} \log n}$$

for some positive constants c and C depending only on s .

In general, there are very few H for which the growth rate of $\log(r(H, K_n^{(3)}))$ in terms of n is well understood. The only exceptions are tripartite 3-graphs and their iterated blowups, for which Erdős and Hajnal [7] proved that $r(H, K_n^{(3)}) \leq n^{O_H(1)}$, and links of odd cycles, for which Fox and He [9] showed that

$$r(H, K_n^{(3)}) = 2^{\Theta_H(n \log n)},$$

proving that an old upper bound of Erdős and Hajnal [7] is tight. Note that for us the *link* of a (2-)graph G will mean the 3-graph L_G with vertex set $V(G) \cup \{u\}$, where u is a new vertex, and edge set $e \cup \{u\}$ for every $e \in E(G)$. More generally, the result of Fox and He showed that $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n \log n)}$ whenever H is the link of an odd cycle.

Our first theorem extends this latter result by showing that $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n \log n)}$ for a large class of 3-graphs H , including links of odd cycles and all tight cycles of length n with $3 \nmid n$ and, as we shall see below, is sharp in both these cases. We state our result in terms of a function of H that we denote by $m_{\text{pair}}(H)$, though we will hold off on formally defining this parameter until Section 2.

Theorem 1.1. *If $m_{\text{pair}}(H) \geq \frac{1}{2}$, then $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n \log n)}$.*

Since $K_{n,n,n}^{(3)} \subseteq K_{3n}^{(3)}$, this theorem also holds with $K_{n,n,n}^{(3)}$ replaced by $K_n^{(3)}$. Though we will only define $m_{\text{pair}}(H)$ later on, some sense of it may be gained by noting that an early result of Erdős and Hajnal [7] on off-diagonal hypergraph Ramsey numbers may also be couched in this language.

Theorem 1.2. *[Rephrasing of a result in [7]] If $m_{\text{pair}}(H) > \frac{1}{3}$, then $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n)}$.*

As well as being much more general than the result of Fox and He [9], Theorem 1.1 also has a simpler proof, in particular avoiding the entropy techniques of [9] and perhaps suggesting that $m_{\text{pair}}(H)$ is in fact the correct parameter to work with. It is also often tight, as shown by the following result. The *shadow* (or *1-skeleton*) of a 3-graph H is the graph G with $V(G) = V(H)$ whose edge set consists of all pairs of vertices uv such that uvw is an edge of H for some vertex w . We use the standard notation $\partial H = E(G)$.

Theorem 1.3. *Let H be a 3-graph that can be created by starting with the empty graph and iteratively adding edges that strictly increase the number of edges in the shadow of H at each step. Then*

$$r(H, K_{n,n,n}^{(3)}) \leq 2^{O_H(n \log n)}.$$

Note that Theorem 1.3 immediately implies that Theorem 1.1 is tight for links of odd cycles and for tight cycles, as these may clearly be constructed by repeatedly adding 3-edges that add at least one 2-edge to the shadow. However, we also note that, apart from finitely many exceptions in the case of tight cycles, the stronger bound $r(H, K_n^{(3)}) \leq 2^{O_H(n \log n)}$ was already known in both of these cases [7, 14].

Our second main result addresses a longstanding problem in the area, asking whether $r(F, K_n^{(3)})$ is polynomial in n if F is a *linear* 3-graph (also known as a linear triple system, partial triple system or partial Steiner system), that is, a 3-graph where any two edges share at most one vertex. That this should be the case was something of a folklore conjecture, but had been resistant to attack. The following result shows why.

Theorem 1.4. *For all sufficiently large k , there exists a linear 3-graph F on k vertices such that*

$$r(F, K_n^{(3)}) > 2^{c(\log n)^{\sqrt{k}/(128 \log^5 k)}},$$

where $c > 0$ depends only on k .

In contrast, a result of Fox and He [9] shows that if F is a linear 3-graph, then the Ramsey number $r(F, K_{n,n,n}^{(3)})$ grows polynomially in n . Combined, these results give the first examples of 3-graphs F for which the Ramsey numbers $r_3(F, K_{n,n,n}^{(3)})$ and $r_3(F, K_n^{(3)})$ are known to have very different growth rates.

Theorem 1.4 also has a further corollary. If, for a triple system F , we define $m_k(F)$ to be the minimum number of edges in an F -free triple system with chromatic number at least k , Bohman, Frieze and Mubayi [2] conjectured that there exists a linear triple system F for which $m_k(F) = k^{3+o(1)}$. It is easy to see that $m_k(F) \geq k^{3+o(1)}$, since any triple system with f edges has chromatic number $O(f^{1/3})$, so this conjecture would be asymptotically tight. Theorem 1.4 settles the conjecture, since it produces N -vertex F -free 3-graphs with independence number $n = N^{o(1)}$ and, hence, chromatic number at least $k = N^{1-o(1)}$. Moreover, the number of edges is trivially at most $\binom{N}{3} = k^{3+o(1)}$.

We now proceed to our proofs, beginning with Theorem 1.1. We will return to Theorem 1.4 in Section 3 and conclude in Section 4 with some further remarks and conjectures. We note that, unless otherwise indicated, all of our logarithms are natural logarithms.

2 Proof of Theorem 1.1

In this section, we prove Theorem 1.1, saying that if $m_{\text{pair}}(H) \geq \frac{1}{2}$, then $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n \log n)}$. Our first task is to formally define $m_{\text{pair}}(H)$, which we do in Section 2.1. Our construction is described in detail in Section 2.2 and then, in Section 2.3, we show that it has the required properties. In Section 2.4, we show that $m_{\text{pair}}(H) \geq \frac{1}{2}$ for all links of odd cycles and all tight cycles whose length is not a multiple of 3, so that our results do indeed apply in these cases. Finally, in Section 2.5, we give the short proof of Theorem 1.3, which says that Theorem 1.1 is tight for a large class of 3-graphs.

2.1 Avoidability

In this section, we formally define the “pair density” statistic $m_{\text{pair}}(H)$ that was mentioned in the introduction and give a characterization of all 3-graphs H with $m_{\text{pair}}(H) \geq \frac{1}{2}$. We will call such 3-graphs “avoidable”. However, before getting to that, we need several other definitions.

Definition. A hypergraph is *ordered* if its vertex set is a linearly ordered set and *nonempty* if it has at least one edge.

Definition. An *oriented 3-graph* G is a pair of sets (V, E) such that $E \subseteq V^3$ contains at most one of the six permutations of any given ordered triple $(v_1, v_2, v_3) \in V^3$.

Definition. Given an ordered 3-graph H and an oriented 3-graph G , a *pair homomorphism* from H to G is a function $f : \partial H \rightarrow V(G)$ such that, for every edge $\{u, v, w\} \in E(H)$ with $u < v < w$, $f(uv)$, $f(vw)$ and $f(wu)$ are distinct and $(f(uv), f(vw), f(wu)) \in E(G)$. The map sending uvw to $(f(uv), f(vw), f(wu))$ induces another map, which we also call f , from $E(H)$ to $E(G)$. Write $f(H)$ for the subgraph of G whose edge set is the image of this map and whose vertex set is $f(\partial H)$. Note that $v(f(H)) = |f(\partial H)|$ and $e(f(H)) = |f(E(H))|$.

Definition. If H is a nonempty ordered 3-graph, define

$$m_{\text{pair}}(H) := \min_{f: \partial H \rightarrow V(G)} \max_{H' \subseteq H} \frac{e(f(H'))}{v(f(H'))},$$

where f ranges over all pair homomorphisms from H to any oriented G and H' ranges over all nonempty subhypergraphs of H . We define $m_{\text{pair}}(H)$ for an unordered 3-graph H to be the minimum of $m_{\text{pair}}(H)$ over all orderings of H .

Since $f(H)$ has no isolated vertices by definition, $v(f(H)) \leq \sum \deg_{f(H)}(v) = 3e(f(H))$ and consequently $m_{\text{pair}}(H) \geq 1/3$ for all nonempty H . On the other hand, $m_{\text{pair}}(K_s^{(3)}) = \Theta(s)$. Indeed, the upper bound follows immediately from the Kruskal–Katona theorem and for the lower bound we may argue as follows. Suppose that f is a surjective pair homomorphism from $K_s^{(3)}$ to an oriented graph G . Then, for any incident pair e and e' in $\partial K_s^{(3)}$, we know that $f(e) \neq f(e')$ and e and e' lie in an edge of G . As we vary e' over all elements of $\partial K_s^{(3)}$ incident to e , we see that $f(e)$ lies in at least $\Omega(s)$ edges of G . Hence, the minimum degree of G is at least $\Omega(s)$, so $e(G)/v(G) = \Omega(s)$ as well.

In order to say what it means for a 3-graph to be avoidable, we must first recall that a Berge cycle is a hypergraph with (distinct) vertices v_1, \dots, v_t and distinct edges e_1, \dots, e_t , where $\{v_i, v_{i+1}\} \subseteq e_i$ and indices are taken modulo t .

Definition. We say that a 3-graph H is *avoidable* if every 3-graph G for which there is a pair homomorphism f from H to G contains a Berge cycle.

The following lemma shows that being avoidable is equivalent to having pair density at least $1/2$.

Lemma 2.1. *A 3-graph H is avoidable if and only if $m_{\text{pair}}(H) \geq \frac{1}{2}$.*

Proof. This boils down to the observation that a 3-graph G contains a subhypergraph G' with $e(G')/v(G') \geq \frac{1}{2}$ if and only if it contains a Berge cycle. Indeed, any Berge cycle satisfies this property and, conversely, any Berge acyclic hypergraph satisfies $v(G) > 2e(G)$ by applying Lemma 2.7 below to each connected component. \square

The following definition is also of interest.

Definition. A 3-graph H is *123-inducible* if there is an ordering of $V(H)$ and a labelling of ∂H by 1, 2 and 3 in such a way that, for every $uvw \in E(H)$ with $u < v < w$, uv is labelled 1, vw is labelled 2 and uw is labelled 3.

Lemma 2.2. *H is 123-inducible if and only if $m_{\text{pair}}(H) = \frac{1}{3}$.*

Proof. As a 123-coloring is a pair homomorphism yielding one edge and three vertices (and, clearly, $3e \geq v$ for any 3-graph without isolated vertices), any 123-inducible graph has m_{pair} equal to $1/3$. Conversely, if $m_{\text{pair}}(H) = \frac{1}{3}$, there is a pair homomorphism f such that $3e(f(H)) = v(f(H))$. But then $f(H)$ must consist entirely of disjoint triples, whose vertices we may color with the colors 1, 2, 3 in accordance with the ordering on $v(H)$. That is, H is 123-inducible. \square

We note that Theorem 1.2, the result of Erdős and Hajnal [7] quoted in the introduction, was originally proved in terms of 123-inducibility and then states that if a 3-graph H is not 123-inducible, then $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n)}$. The proof is rather simple: take a random coloring of the edges of $K_N^{(2)}$ in three colors 1, 2 and 3 and then color an edge uvw of $K_N^{(3)}$ with $u < v < w$ in red if and only if uv is colored 1, vw is colored 2 and wu is colored 3. It is easy to verify that there exists a positive constant C such that, with positive probability, this coloring contains no red copy of any 3-graph H which is not 123-inducible and no blue $K_{n,n,n}^{(3)}$ with $n = C \log N$.

2.2 The construction

In this section, we describe our construction for Theorem 1.1. In light of Lemma 2.1, it will suffice to show the following.

Theorem 2.3. *If H is avoidable, then $r(H, K_{n,n,n}^{(3)}) \geq 2^{\Omega_H(n \log n)}$.*

The construction for Theorem 2.3 is as follows. Take ε sufficiently small in terms of H , c sufficiently small in terms of ε and n sufficiently large in terms of c . Let $N = 2^{cn \log n}$ and $t = n^\varepsilon$ and define $f : \binom{[N]}{2} \rightarrow [t]$ uniformly at random. Pick $g : [t]^3 \rightarrow \{\text{red}, \text{blue}\}$ to satisfy:

1. The red color in g is an oriented 3-graph, i.e., for all $e \in [t]^3$ at most one of the six permutations $\sigma(e)$ of e satisfies $g(\sigma(e)) = \text{red}$. In particular, if e has at least two equal coordinates, then $g(e) = \text{blue}$.
2. Among any $k \leq \binom{v(H)}{2}$ elements of $[t]$ there are fewer than $k/2$ red edges. In particular, the red edges in g form a linear 3-graph.
3. In blue, g does not have complete balanced oriented tripartite subgraphs of order $t^{1-\varepsilon}$. That is, for any $X, Y, Z \subseteq [t]$ of size $t^{1-\varepsilon}$, there must be at least one triple $(x, y, z) \in X \times Y \times Z$ for which $g(x, y, z) = \text{red}$.

Thus, the red color in g may be seen as an oriented 3-uniform analog of a graph with high girth and small independence number. Our coloring $\chi : \binom{[N]}{3} \rightarrow \{\text{red}, \text{blue}\}$ is now defined by $\chi(u, v, w) = g(f(uv), f(vw), f(wu))$ for all $u < v < w$ in $[N]$. Note that condition 2 already implies that χ contains no red copy of H provided H is avoidable.

We now prove the existence of a function g satisfying conditions 1-3 before moving on to proving Theorem 2.3 in the next section.

Lemma 2.4. *For any 3-graph H , there exists a function g satisfying conditions 1-3 above for any ε sufficiently small and t sufficiently large in terms of H .*

Proof. We would like to construct a function g with the required properties. First, note that we must have $g(i, i, i) = g(i, i, j) = g(i, j, i) = g(j, i, i) = \text{blue}$ for all i, j . To color the remaining edges, we use the alteration method. Let $g : [2t]^3 \rightarrow \{\text{red}, \text{blue}\}$ be random, where $g(i, j, k) = \text{red}$ with probability $p = t^{-2+\delta}$ independently for each triple (i, j, k) . The expected number of pairs of permutations of the same triple which are both colored red is at most

$$\binom{2t}{3} \binom{6}{2} p^2 = O(t^3 p^2) < t/3$$

and the expected number of red $(k, k/2)$ -configurations is at most

$$\binom{2t}{k} \binom{k^3}{k/2} p^{k/2} = O_k(t^k p^{k/2}) < t/3$$

if δ is sufficiently small and t is sufficiently large in terms of $|V(H)|$. The expected number of blue tripartite $X, Y, Z \subset [2t]$ of size $|X| = |Y| = |Z| = t^{1-\varepsilon}$ is at most

$$(2t)^{3t^{1-\varepsilon}} (1-p)^{t^{3-3\varepsilon}-O(t^2)} \leq e^{O(t^{1-\varepsilon} \log t) - t^{1-3\varepsilon+\delta}} < t/3$$

if ε is sufficiently small compared to δ . Here the $O(t^2)$ subtracted in the exponent accounts for the possibility that X, Y, Z overlap, in which case up to $O(t^2)$ triples (x, y, z) with a repeated coordinate were colored blue in advance.

Thus, in total, there exists a sample of g for which the number of bad red and blue configurations is together less than t , so we may delete t vertices from $[2t]$ to obtain the desired g . \square

2.3 Proof of Theorem 2.3

We note that g is fixed and all probabilistic statements are with respect to the random choice of f . By construction, for all choices of f , the coloring χ is guaranteed to contain no red copy of H . It therefore remains to show that with positive probability χ contains no blue copy of $K_{n,n,n}^{(3)}$. We remark that the only property of g that we will use in the proof of this statement is property 3, which will be used in the proof of Claim 2 below.

Observe that every $K_{n,n,n}^{(3)}$ contains a complete tripartite subgraph with parts I, J, K of order $n/3$ such that $I < J < K$ (that is, for all $i \in I, j \in J, k \in K, i < j < k$). Thus, up to this factor of three that we will ignore, it suffices to show that χ contains no blue copy of $K_{n,n,n}^{(3)}$ whose parts I, J, K satisfy $I < J < K$. By taking a union bound over all N^{3n} possible choices of I, J, K , we see that it suffices to show the following.

Lemma 2.5. *If $I < J < K$ are three n -element subsets of $[N]$, then the probability that $\chi(i, j, k) = \text{blue}$ for all $(i, j, k) \in I \times J \times K$ is at most $t^{-\varepsilon n^2/50}$.*

Proof. We would like to count *blue functions* $f : (I \times J) \cup (J \times K) \cup (I \times K) \rightarrow [t]$, i.e., those for which $g(f(i, j), f(j, k), f(k, i)) = \text{blue}$ for all $(i, j, k) \in I \times J \times K$. Let \mathcal{F} be the family of all such blue f . If we set $\delta = \varepsilon/50$, Lemma 2.5 is equivalent to the inequality

$$|\mathcal{F}| \leq t^{(3-\delta)n^2}, \tag{1}$$

which is now our goal. We will split the argument into three steps. For simplicity, we often abuse notation by writing $\sum_X f(x)$ and $\prod_X f(x)$ instead of $\sum_{x \in X} f(x)$ and $\prod_{x \in X} f(x)$.

1. Coloring $(I \times J) \cup (J \times K)$ first. Let \mathcal{F}_J be the family of functions $f_J : (I \times J) \cup (J \times K) \rightarrow [t]$ for which there are at least $t^{(1-2\delta)n^2}$ blue functions $f \in \mathcal{F}$ extending f_J , i.e., for which $f|_{(I \times J) \cup (J \times K)} = f_J$. It suffices to show that

$$|\mathcal{F}_J| \leq t^{(2-2\delta)n^2}, \tag{2}$$

since that would imply that

$$|\mathcal{F}| \leq |\mathcal{F}_J| \cdot t^{n^2} + t^{2n^2} \cdot t^{(1-2\delta)n^2} \leq 2t^{(3-2\delta)n^2} \leq t^{(3-\delta)n^2},$$

by conditioning on whether or not $f|_{(I \times J) \cup (J \times K)} \in \mathcal{F}_J$. That is, (1) would hold.

It will be convenient to set up some notation for counting blue extensions of f_J . For $f_J : (I \times J) \cup (J \times K) \rightarrow [t]$, let $f_J^*(i, k)$ be the number of choices of $c = f(i, k)$ for which $g(f_J(i, j), f_J(j, k), c) = \text{blue}$ for all j . Observe that the number of extensions $P(f_J)$ of f_J to blue functions $f \in \mathcal{F}$ is exactly the product

$P(f_J) = \prod_{I \times K} f_J^*(i, k)$. Thus, \mathcal{F}_J is exactly the family of f_J satisfying $P(f_J) \geq t^{(1-2\delta)n^2}$, so, by the AM-GM inequality, every $f_J \in \mathcal{F}_J$ satisfies

$$S(f_J) := \sum_{I \times K} f_J^*(i, k) \geq n^2 \left(\prod_{I \times K} f_J^*(i, k) \right)^{1/n^2} = n^2 P(f_J)^{1/n^2} \geq n^2 t^{1-2\delta}.$$

2. Coloring $(I \times \{j\}) \cup (\{j\} \times K)$ one j at a time. To prove (2), we reveal the values of f_J one j at a time and track the evolution of $P(f_J)$ and $S(f_J)$. To do so, we extend the definitions of f_J^* , $P(f_J)$ and $S(f_J)$ to subsets $J' \subseteq J$. Given $J' \subseteq J$ and a function $f_{J'} : (I \times J') \cup (J' \times K) \rightarrow [t]$ defined on the pairs incident to J' , we take $f_{J'}^*(i, k)$ to be the number of choices of $c = f(i, k)$ for which $g(f_{J'}(i, j), f_{J'}(j, k), c) = \text{blue}$ for all $j \in J'$, $P(f_{J'}) := \prod_{I \times K} f_{J'}^*(i, k)$ and $S(f_{J'}) = \sum_{I \times K} f_{J'}^*(i, k)$.

If $J = \{j_1, \dots, j_n\}$, let $J_r = \{j_1, \dots, j_r\}$, where $J_0 = \emptyset$. Observe that if $f_J : (I \times J) \cup (J \times K) \rightarrow [t]$ is restricted to $f_{J_r} := f_J|_{(I \times J_r) \cup (J_r \times K)}$, then the values $S(f_{J_r})$ satisfy

$$n^2 t = S(f_{J_0}) \geq S(f_{J_1}) \geq S(f_{J_2}) \geq \dots \geq S(f_{J_n}) = S(f_J) \geq 0.$$

Thus, for any f_J we have that at least $n/2$ values of r satisfy $S(f_{J_{r-1}}) - S(f_{J_r}) \leq 2nt$. In words, this means that revealing the colors on the pairs incident to j_r decreases the total number of blue choices between I and K by at most $2nt$. Call r *slow* (with respect to f_J) if $S(f_{J_{r-1}}) - S(f_{J_r}) \leq 2nt$. It will suffice to show the following.

Claim 1. If $P(f_{J_{r-1}}) \geq t^{(1-2\delta)n^2}$, then the number of ways to color $f_{j_r} := f_J|_{(I \times \{j_r\}) \cup (\{j_r\} \times K)}$ to make r slow is at most $t^{(2-5\delta)n}$.

Indeed, if the claim is true, we have

$$|\mathcal{F}_J| \leq 2^n \cdot (t^{2n})^{n/2} \cdot (t^{(2-5\delta)n})^{n/2} \leq t^{(2-2\delta)n^2},$$

since, to construct $f_J \in \mathcal{F}_J$, there are at most 2^n ways to pick which of the steps are slow and, on each of the $\geq n/2$ slow steps, there are $\leq t^{(2-5\delta)n}$ ways to color f_{j_r} because, throughout the process, $P(f_{J_{r-1}}) \geq P(f_J) \geq t^{(1-2\delta)n^2}$. This would prove (2), as desired.

3. Counting slow steps. It remains to prove Claim 1. Fix r satisfying $P(f_{J_{r-1}}) \geq t^{(1-2\delta)n^2}$. We are interested in counting the number of “slow colorings” $f_{j_r} : (I \times \{j_r\}) \cup (\{j_r\} \times K) \rightarrow [t]$ for which $S(f_{J_{r-1}}) - S(f_{J_r}) \leq 2nt$. That is, we wish to count the number of choices of f_{j_r} for which

$$\sum_I \sum_K (f_{J_{r-1}}^*(i, k) - f_{J_r}^*(i, k)) \leq 2nt.$$

Since $\sum_K (f_{J_{r-1}}^*(i, k) - f_{J_r}^*(i, k))$ is nonnegative, at least $3n/4$ values of i satisfy

$$\sum_K (f_{J_{r-1}}^*(i, k) - f_{J_r}^*(i, k)) \leq 8t. \tag{3}$$

Expanding out the assumption $P(f_{J_{r-1}}) \geq t^{(1-2\delta)n^2}$, we have

$$\prod_I \prod_K f_{J_{r-1}}^*(i, k) \geq t^{(1-2\delta)n^2}.$$

Since $f_{J_{r-1}}^*(i, k) \leq t$ for all i, k , at least $3n/4$ values of i satisfy

$$\prod_K f_{J_{r-1}}^*(i, k) \geq t^{(1-8\delta)n}. \tag{4}$$

Thus, there must exist a set $I' \subseteq I$ of exactly $n/2$ values of i for which (3) and (4) hold simultaneously. Let us now fix such an I' of size $n/2$ and count the number of f_{j_r} satisfying (3) for all $i \in I'$. Call such colorings I' -slow.

Our goal is to prove that the number of I' -slow colorings of f_{j_r} is at most $t^{(2-\epsilon/8)n}$. This will finish the proof of Claim 1 as, summing over the at most 2^n choices of I' , the number of ways to make a slow coloring f_{j_r} is at most $t^{(2-\epsilon/10)n} = t^{(2-5\delta)n}$, as desired.

Choosing f_{j_r} is the same as choosing $f_{\{j_r\} \times K}$ and $f_{I \times \{j_r\}}$. Suppose we have already picked $f_{\{j_r\} \times K}$, the colors incident to K . Then the remaining pairs from j_r to I can be colored independently: those outside of I' in t ways each and those incident to I' so as to satisfy (3). For a particular $i \in I'$, let $C_i(f_{\{j_r\} \times K})$ be the set of colors consisting of those $c = f_{j_r}(i, j_r)$ satisfying (3). Such a color c has to be chosen so that for at most $8t$ of the n values of k , $f_{j_r}^*(i, k)$ decreases.

Claim 2. If i satisfies (4), then the number of choices of $f_{\{j_r\} \times K}$ for which $|C_i(f_{\{j_r\} \times K})| \geq t^{1-\epsilon}$ is at most $2^t \cdot \binom{n}{8t} \cdot t^{(1-\epsilon/2)n} \leq t^{(1-\epsilon/4)n}$.

Proof. Fix a choice of $X = C_i(f_{\{j_r\} \times K})$ in at most 2^t ways and a choice of the at most $8t$ values of k for which $f_{j_r}^*(i, k)$ decreases in at most $\binom{n}{8t}$ ways. Once these choices are fixed, the color choices for $j_r \times k$ can be made independently for each k , so long as, for those k outside of the special $8t$, we do not eliminate any possible colors across (i, k) . Let Y_k denote the set of valid color choices for $j_r \times k$, so that $|Y_k| = t$ for those in the special set of $8t$.

By (4), we see that, for at least $3n/4$ values of k , we have a set $Z_k \subseteq [t]$ of size $f_{j_r}^*(i, k) \geq t^{1-\epsilon}$ of possible colors for ik . For all except $8t$ of these, none of these colors can be eliminated by choosing any of the possible colors ij from X and so, for each of at least $3n/4 - 8t \geq n/2$ choices of k , we have that there is no triple $(x, y, z) \in X \times Y_k \times Z_k$ for which $g(x, y, z) = \text{red}$. Since X and Z_k are both at least $t^{1-\epsilon}$, we find that Y_k is less than $t^{1-\epsilon}$. Thus, the number of choices of $f_{\{j_r\} \times K}$ for this particular choice of X and the exceptional $8t$ values of k is at most $(t^{1-\epsilon})^{n/2} \cdot t^{n/2}$, so the claim follows. \square

This claim essentially finishes the proof. We have

$$|\{I'\text{-slow colorings } f_{j_r}\}| = \sum_{f_{\{j_r\} \times K}} \prod_i |C_i(f_{\{j_r\} \times K})|,$$

a sum that has t^n total terms, where each summand is a product of factors which lie in $\{0, 1, \dots, t\}$. We can partition the products above based on whether $i \in I'$, so that, using the trivial bound t for $|C_i(f_{\{j_r\} \times K})|$ when $i \notin I'$, we obtain

$$\sum_{f_{\{j_r\} \times K}} \prod_i |C_i(f_{\{j_r\} \times K})| \leq \sum_{f_{\{j_r\} \times K}} \left(\prod_{i \in I'} |C_i(f_{\{j_r\} \times K})| \right) t^{n-|I'|}.$$

By Claim 2 and the definition of I' , we have that if $i \in I'$, then the corresponding factor $|C_i(f_{\{j_r\} \times K})|$ is at least $t^{1-\epsilon}$ for at most $t^{(1-\epsilon/4)n}$ choices of $f_{\{j_r\} \times K}$. This allows us to partition the sum as

$$\sum_{f_{\{j_r\} \times K}} \prod_{i \in I'} |C_i(f_{\{j_r\} \times K})| = \sum_{|C_i(f_{\{j_r\} \times K})| \geq t^{1-\epsilon}} \prod_{i \in I'} |C_i(f_{\{j_r\} \times K})| + \sum_{|C_i(f_{\{j_r\} \times K})| < t^{1-\epsilon}} \prod_{i \in I'} |C_i(f_{\{j_r\} \times K})|.$$

The first sum can be bounded by $t^{(1-\epsilon/4)n} \cdot t^{|I'|}$ and the second sum can be bounded by $t^n \cdot (t^{(1-\epsilon)})^{|I'|}$. Multiplying by $t^{n-|I'|}$ and recalling that $|I'| = n/2$, this yields

$$|\{I'\text{-slow colorings } f_{j_r}\}| \leq t^{n-|I'|} (t^{(1-\epsilon/4)n} \cdot t^{|I'|} + t^n \cdot (t^{(1-\epsilon)})^{|I'|}) = t^{(2-\epsilon/4)n} + t^{(2-\epsilon/2)n} < t^{(2-\epsilon/8)n},$$

so the proof is complete. \square

2.4 Avoidability proofs

In this section, we calculate the pair densities of tight cycles and links of cycles, as a result showing that non-tripartite tight cycles and links of non-bipartite graphs are avoidable. Write $C_n^{(3)}$ for the 3-uniform tight cycle of length n . Concretely, we can take $V(C_n^{(3)}) = \mathbb{Z}_n$ and $E(C_n^{(3)}) = \{\{k, k+1, k+2\} : k \in \mathbb{Z}_n\}$.

Proposition 2.6.

$$m_{\text{pair}}(C_n^{(3)}) = \begin{cases} 1/2 & 3 \nmid n, \\ 1/3 & 3|n. \end{cases}$$

Proof. As any subgraph H of $C_n^{(3)}$ satisfies $e(H) \leq |\partial H|/2$, we can see that $m_{\text{pair}}(C_n^{(3)}) \leq 1/2$ by mapping each element of $\partial C_n^{(3)}$ to a distinct vertex.

If $3|n$, it is not difficult to see that $C_n^{(3)}$ is 123-inducible and thus has $m_{\text{pair}} = 1/3$. It therefore suffices to show that if $m_{\text{pair}}(C_n^{(3)}) < 1/2$, then $3|n$.

Take n with $m_{\text{pair}}(C_n^{(3)}) < 1/2$. By definition, we can take an ordering of $V(C_n^{(3)})$, a nonnegative integer k and a surjection $f : \partial C_n^{(3)} \rightarrow [k]$ such that $|f(E(C_n^{(3)}))| < k/2$. Consider the 3-graph G on $[k]$ with edges given by $f(e)$, $e \in E(C_n^{(3)})$. Then $v(G) = k$ and $e(G) < k/2$. Note that G is connected, as in a tight cycle we can go from any edge to any other by taking a sequence of edges any adjacent pair of which share an element of the shadow.

Lemma 2.7. *For any connected 3-graph G ,*

$$v(G) \leq 2e(G) + 1,$$

with equality if and only if G is Berge-acyclic. Moreover, in the case of equality, the shadow of G has no cycle of length at least four.

Proof. Note that if $c(G)$ is the number of connected components in G , then $2e(G) + c(G) - v(G)$ cannot decrease upon adding an edge e to G , as easily seen by a short case analysis on the number of new vertices in e . For the empty 3-graph, this quantity is 0, so, for a connected 3-graph, $2e(G) + 1 - v(G) = 2e(G) + c(G) - v(G) \geq 0$, as desired. Moreover, if G' is a Berge cycle within G , then $2e(G') + c(G') - v(G') \geq 1$, so, again adding edges one at a time, we may conclude that $2e(G) + 1 - v(G) = 2e(G) + c(G) - v(G) \geq 1$.

Conversely, suppose that G is Berge-acyclic. If every edge of G has at least 2 vertices each contained in another edge, then we can create a walk of the form $e_1 v_1 e_2 v_2 \dots$, where $v_i \in e_i, e_{i+1}$ and $e_i \neq e_{i+1}, v_i \neq v_{i+1}$. If we get to a vertex or an edge that we have seen before, we obtain a Berge cycle, a contradiction. Thus, G must contain at least one edge e that contains at most one vertex that is contained in any other edge. Removing e decreases $e(G)$ by one and $v(G)$ by 2, but preserves Berge-acyclicity and connectedness, so a simple induction shows that $v(G) = 2e(G) + 1$.

Finally, suppose that G is Berge-acyclic and assume, for the sake of contradiction, that ∂G has a cycle of length at least 4. Let e_1, \dots, e_n be the edges in a shortest such cycle of ∂G . If $n \geq 5$, e_i and e_j cannot be in the same edge $e \in E(G)$ for any i, j , as if they were then you could replace them with the third edge of ∂G within e and make a cycle of length $n - 1$. Thus, e_1, \dots, e_n are all in different edges of G and form a Berge cycle, a contradiction. If $n = 4$, either e_1, \dots, e_4 are in different edges of G and form a Berge 4-cycle or two of them are in the same edge $e' \in E(G)$, say e_1 and e_2 . Replacing e_1 and e_2 with the third element of ∂G within e' , we either obtain a Berge 3-cycle or the remaining three edges of ∂G are all contained in one edge of G . In this latter case, however, we have two edges of G spanning only 4 vertices, yielding a Berge 2-cycle, a contradiction. \square

Applying the lemma, since $v(G) > 2e(G)$, we must have that G is Berge-acyclic and that the shadow of G only has cycles of length 3.

Consider now the closed walk W of length n on the complete graph with $2k$ vertices $[k] \times \{+, -\}$ defined as follows. Label the vertices of $C_n^{(3)}$ as v_1, \dots, v_n cyclically. If the edge $v_i v_{i+1}$ has color j (i.e., it is mapped

to j under f), then we map it to $(j, +)$ if $v_i < v_{i+1}$ in our ordering on $V(C_n^{(3)})$ and $(j, -)$ otherwise. The moves in the walk W go from the image of $v_i v_{i+1}$ to the image of $v_{i+1} v_{i+2}$ (with indices considered cyclically).

Lemma 2.8. *The walk W has no cycles of length more than 3 and no edges that are traversed in both directions.*

Proof. Consider the map p from $[k] \times \{+, -\} \rightarrow [k]$ that ignores the sign. Edges in W are sent to edges in the shadow of G , so all cycles in the image of W must have length 3.

Suppose now, for the sake of contradiction, that W has a cycle C of length at least 4. The image $p(C)$ is a closed walk in the shadow of G . The minimum closed subwalk of $p(C)$ coming from a subpath of C must have length 2 or 3, as the shadow of G only has cycles of length 3. Since C has length greater than 3, this must come from a subpath of C of the form

$$(j_1, s_1), (j_2, s_2), (j_1, -s_1)$$

or

$$(j_1, s_1), (j_2, s_2), (j_3, s_3), (j_1, -s_1)$$

for some $j_i \in [k]$, $s_i \in \{+, -\}$. We now show that both of these cases give contradictory conclusions about the vertex ordering on v_1, \dots, v_n .

Case 1: In the first case, there are two edges $v_i v_{i+1} v_{i+2}$ and $v_j v_{j+1} v_{j+2}$ of $C_n^{(3)}$ such that $f(v_i v_{i+1}) = f(v_{j+1} v_{j+2}) = j_1$ and $f(v_{i+1} v_{i+2}) = f(v_j v_{j+1}) = j_2$. Thus, in G , the images of $v_i v_{i+1} v_{i+2}$ and $v_j v_{j+1} v_{j+2}$ share two vertices and so must be the same edge by Lemma 2.7. This implies that $f(v_i v_{i+2}) = f(v_j v_{j+2}) = j_3$ for some color j_3 . Since f cannot send two edges to two different permutations of the same triple, we see that the relative orderings of $v_i v_{i+1} v_{i+2}$ and $v_{j+2} v_{j+1} v_j$ must be the same. However, since the sign s_2 of $v_{i+1} v_{i+2}$ is the same as the sign of $v_j v_{j+1}$, this cannot be the case.

Case 2: In the second case, consider the edges $v_i v_{i+1} v_{i+2}$, $v_j v_{j+1} v_{j+2}$ and $v_k v_{k+1} v_{k+2}$ that form the three edges of our cycle in $p(C)$. Then we have that

$$f(v_i v_{i+1}) = f(v_{k+1} v_{k+2}) = j_1,$$

$$f(v_j v_{j+1}) = f(v_{i+1} v_{i+2}) = j_2$$

and

$$f(v_k v_{k+1}) = f(v_{j+1} v_{j+2}) = j_3.$$

Thus, each pair of the edges $f(v_i v_{i+1} v_{i+2})$, $f(v_i v_{i+1} v_{i+2})$ and $f(v_i v_{i+1} v_{i+2})$ in G must share a vertex. No two different edges can share two vertices in G , so, since this cannot form a loose 3-cycle, it must be that these three edges of G are the same. Thus, $f(v_j v_{j+2}) = j_1$, $f(v_k v_{k+2}) = j_2$ and $f(v_i v_{i+2}) = j_3$. Once again, since f cannot send different edges to different permutations of the same triple, this gives that the relative orderings of $v_i v_{i+1} v_{i+2}$, $v_{j+2} v_{j+1} v_j$ and $v_{k+1} v_{k+2} v_k$ are identical. However, since $v_i v_{i+1}$ maps to (j_1, s_1) but $v_{k+1} v_{k+2}$ maps to $(j_1, -s_1)$, we see that $v_i v_{i+1}$ and $v_{k+1} v_{k+2}$ have different orderings, contradicting that $v_i v_{i+1} v_{i+2}$ and $v_{k+1} v_{k+2} v_k$ have the same relative ordering. This concludes the proof that W has no cycles of length more than 3.

Finally, we show that no edge of W is traversed in both directions. Suppose, for the sake of contradiction, that there is an edge $(j_1, s_1), (j_2, s_2)$ that is traversed in both directions. Suppose that in the forward direction this edge comes from $v_i v_{i+1} v_{i+2}$ and in the reverse direction it comes from $v_j v_{j+1} v_{j+2}$. Then $f(v_i v_{i+1}) = f(v_{j+1} v_{j+2}) = j_1$ and $f(v_j v_{j+1}) = f(v_{i+1} v_{i+2}) = j_2$. As in the arguments above, we must have that $f(v_i v_{i+2}) = f(v_j v_{j+2})$ and $v_i v_{i+1} v_{i+2}$ has the same relative ordering as $v_{j+2} v_{j+1} v_j$. However, this contradicts the fact that $v_i v_{i+1}$ and $v_{j+1} v_{j+2}$ have the same relative ordering (given by s_1). Thus, no edge of W is traversed in both directions, finishing the proof of Lemma 2.8. \square

It is an easy induction to show that any closed walk satisfying the conditions of Lemma 2.8 has length divisible by 3 (it must have at least one triangle and removing it gives a closed walk of length $n - 3$ also satisfying the conditions). Since W has length n , we must have $3|n$, finishing the proof of Proposition 2.6. \square

We now prove a similar statement for links of cycles. Write $L_n^{(3)}$ for the 3-graph that is the link of the cycle C_n . Concretely, $V(L_n^{(3)}) = \mathbb{Z}_n \cup \{x\}$ and $E(L_n^{(3)}) = \{\{x, k, k + 1\} : k \in \mathbb{Z}_n\}$.

Proposition 2.9.

$$m_{\text{pair}}(L_n^{(3)}) = \begin{cases} 1/2 & 2 \nmid n, \\ 1/3 & 2|n. \end{cases}$$

Proof. The proof at first proceeds along similar lines to the previous one, but there are differences in the details. First, since any subgraph H of $L_n^{(3)}$ satisfies $e(H) \leq |\partial H|/2$, we must have $m_{\text{pair}}(L_n^{(3)}) \leq 1/2$ for all n . If $2|n$, it is easy to see that $L_n^{(3)}$ is tripartite and thus $m_{\text{pair}}(L_n^{(3)}) = 1/3$. It therefore remains to show that $m_{\text{pair}}(L_n^{(3)}) \geq 1/2$ for all odd n .

Suppose, for the sake of contradiction, that we have n with $m_{\text{pair}}(L_n^{(3)}) < 1/2$. We must show that n is even. We may take an ordering of $V(L_n^{(3)})$, a nonnegative integer k and a surjection $h : \partial L_n^{(3)} \rightarrow [k]$ such that $|h(E(L_n^{(3)}))| < k/2$. As before, let G be the 3-graph with edges given by $h(e), e \in E(L_n^{(3)})$. Since G is a connected 3-graph with $v(G) = k$ and $e(G) < k/2$, Lemma 2.7 implies that G is Berge-acyclic.

Consider now the closed walk W of length n on the complete graph with $2k$ vertices $[k] \times \{+, -\}$ defined as follows. Label the center of $L_n^{(3)}$ as v and the remaining vertices cyclically as v_1, \dots, v_n . If $vv_i \in \partial L_n^{(3)}$ has color j , map it to $(j, +)$ if $v < v_i$ in our ordering on $V(L_n^{(3)})$ and $(j, -)$ otherwise. (Note that this is different from how our mapping uses the ordering in the previous proof!) The moves in the walk W go from the image of vv_i to the image of vv_{i+1} (with indices considered cyclically).

The equivalent of Lemma 2.8 in this case is the following.

Lemma 2.10. *The walk W is bipartite.*

Proof. Suppose, for the sake of contradiction, that W has an odd cycle C . Consider the map $p : [k] \times \{+, -\} \rightarrow [k]$ that ignores the sign. Edges in C are sent to edges in the shadow of G , so all cycles in $p(C)$ must have length 3. Thus, in the graph $p(C)$, all minimal closed subwalks have length 2 or 3. So we may partition the edges of $p(C)$ (which possibly occur with repetition) into closed subwalks of lengths 2 or 3.

Call an edge of C (and the corresponding edge of $p(C)$) *switching* if it goes from $+$ to $-$ or $-$ to $+$. Since C is a cycle, it contains an even number of switching edges. As $p(C)$ has odd length, one of the closed subwalks in our partition must have a number of switching edges of different parity to its length. Thus, $p(C)$ contains one of the following:

1. A closed walk of length 2 using exactly 1 switching edge.
2. A closed walk of length 3 using 0 or 2 switching edges.

We now show that both of these are impossible.

Case 1: Suppose we have a closed walk of length 2 using exactly 1 switching edge. We may assume the walk goes $a \rightarrow b \rightarrow a$ for some $a, b \in [k]$. This implies that there are indices i, j such that vv_i and vv_{j+1} are colored a , vv_j and vv_{i+1} are colored b . Now $v_i v_{i+1}$ and $v_j v_{j+1}$ must be colored the same color, call it c , or else G would have two distinct 3-edges sharing two vertices. Thus, $vv_i v_{i+1}$ and $vv_j v_{j+1}$ must correspond to the same edge in G . Comparing which 2-edges are colored a, b, c , we see that this implies that $vv_i v_{i+1}$ and $vv_{j+1} v_j$ have the same relative ordering. But now note that the edge in our walk corresponding to $vv_i v_{i+1}$ is switching if and only if v lies between v_i and v_{i+1} . Thus, the edge corresponding to $vv_i v_{i+1}$ is switching if and only if the edge corresponding to $vv_j v_{j+1}$ is, meaning it is impossible for exactly one such edge to be switching.

Case 2: Suppose we have a closed walk of length 3 using 0 or 2 switching edges. We may assume the walk goes $a \rightarrow b \rightarrow c \rightarrow a$ for some $a, b, c \in [k]$. This implies that there are indices h, i, j such that

- vv_h and vv_{j+1} are colored a ,
- vv_i and vv_{h+1} are colored b ,
- vv_j and vv_{i+1} are colored c .

As this cannot yield a Berge 3-cycle in G or two 3-edges sharing two vertices, we must have that $v_h v_{h+1}$ is colored c , $v_i v_{i+1}$ is colored a and $v_j v_{j+1}$ is colored b . As in the previous case, this implies that $vv_h v_{h+1}$, $v_i v_{i+1} v$ and $v_j v_{j+1} v v_j$ have the same relative ordering. Once again, the edge of our walk corresponding to $vv_i v_{i+1}$ is switching if and only if v lies between v_i and v_{i+1} in the ordering (and similarly for j and k), so exactly one of our three edges must be switching (depending on which is the middle element in the ordering). Thus, it is impossible for 0 or 2 switching edges to appear, finishing the proof. \square

Now Proposition 2.9 immediately follows, as Lemma 2.10 implies that W is a closed walk in a bipartite graph and thus its length n must be even. \square

2.5 Proof of Theorem 1.3

In this section, we prove Theorem 1.3, which says that Theorem 1.1 is tight for a large class of 3-graphs, namely, any 3-graph that can be created by starting with the empty graph and iteratively adding edges that strictly increase the number of edges in the shadow of H at each step. This theorem will follow immediately from the following two lemmas, both of which are proved using supersaturation techniques.

Lemma 2.11. *For all 3 graphs H , $r(H, K_{n,n,n}^{(3)}) \leq r(H, K_{1,n,n}^{(3)})^{3n}$.*

Lemma 2.12. *If H is a 3-graph such that $r(H, K_{1,n,n}^{(3)}) = n^{O(1)}$ and x, y, z are vertices such that at least one of xy, yz and zx is not in the shadow of H , then*

$$r(H \cup \{xyz\}, K_{1,n,n}^{(3)}) = n^{O(1)}.$$

To derive Theorem 1.3, observe that if H satisfies the conditions of that theorem, we can apply Lemma 2.12 repeatedly to get $r(H, K_{1,n,n}^{(3)}) = n^{O(1)}$ and Lemma 2.11 then implies that $r(H, K_{n,n,n}^{(3)}) \leq n^{O(n)}$. It will therefore suffice to prove the two lemmas.

Proof of Lemma 2.11. Fix n and let $A = r(H, K_{1,n,n}^{(3)})$. Let χ be a 2-coloring of the complete 3-graph on A^{3n} vertices and suppose that χ has no red copy of H .

Taking a random induced subgraph on A vertices, as there is no red copy of H , there must be a blue copy of $K_{1,n,n}^{(3)}$. Since every copy of $K_{1,n,n}^{(3)}$ has probability at most

$$\frac{\binom{A^{3n} - (2n+1)}{A - (2n+1)}}{\binom{A^{3n}}{A}} = \frac{\binom{A}{2n+1}}{\binom{A^{3n}}{2n+1}} \leq A^{-(2n+1)(3n-1)} < A^{-6n^2-1}$$

of being chosen, there must be at least A^{6n^2+1} blue copies of $K_{1,n,n}^{(3)}$. Thus, at least

$$\frac{A^{6n^2+1}}{\binom{A^{3n}}{2n} \binom{2n}{n}} > A \geq n$$

copies of $K_{1,n,n}^{(3)}$ must use the same underlying $K_{n,n}$, yielding a copy of $K_{n,n,n}^{(3)}$ in blue. \square

Proof of Lemma 2.12. By adding isolated vertices if necessary, we can and will assume that $\{x, y, z\} \subseteq V(H)$. Suppose that $r(H, K_{1,n,n}^{(3)}) = A$. Assume, without loss of generality, that xy is not in the shadow of H . Let χ be a 2-coloring of the complete graph on $A^{2v(H)}$ vertices and suppose that χ has no blue copy of $K_{1,n,n}^{(3)}$.

Taking a random induced subgraph on A vertices, as there is no blue copy of $K_{1,n,n}^{(3)}$, there must be a red copy of H . Since every copy of H has probability

$$\frac{\binom{A^{2v(H)} - v(H)}{A - v(H)}}{\binom{A^{2v(H)}}{A}} = \frac{\binom{A}{v(H)}}{\binom{A^{2v(H)}}{v(H)}} \leq A^{v(H) - 2v(H)^2}$$

of being chosen, χ must have at least $A^{2v(H)^2 - v(H)}$ red copies of H . Thus, there must be at least

$$A^{2v(H)^2 - v(H) - 2v(H)(v(H) - 2)} = A^{3v(H)}$$

copies of H in red that have all vertices except x and y in a fixed location. Let F be the set consisting of these copies. Let z_0 be the location where z appears in these copies and let S and T be the possible locations for x and y among these copies. Then $|S||T| \geq A^{3v(H)}$, so $|S|, |T| \geq A^{v(H)} > 2n$. Since χ contains no blue $K_{1,n,n}$, there must be a red edge with one vertex at z_0 and one vertex each in S and T , say at s and t .

Since xy is not in the shadow of H , the positions of x and y in S and T are independent after the positions of the remaining vertices are fixed, so there must be a copy of H in F such that x and y are sent to s and t . Thus, adding the red edge z_0st to this copy, we obtain a copy of $H \cup \{xyz\}$. Therefore,

$$r(H \cup \{xyz\}, K_{1,n,n}^{(3)}) \leq r(H, K_{1,n,n}^{(3)})^{2v(H)},$$

as required. □

3 Stepping up for linear hypergraphs versus cliques

In this section, we prove Theorem 1.4, which states that for all sufficiently large k there exists a linear 3-graph F on k vertices such that

$$r(F, K_n^{(3)}) > 2^{c(\log n)^{\sqrt{k}/(128 \log^5 k)}},$$

where $c > 0$ depends only on k , though we note that we have not made a serious attempt at optimizing the power on the log. The proof divides into two parts: first constructing a “quasirandom” F which has many edges between certain appropriately chosen vertex subsets and then using the stepping-up technique to give a coloring with no red copy of F and no blue copy of $K_n^{(3)}$. These objectives are accomplished separately in the next two subsections.

3.1 Constructing F

In this section, we make use of a probabilistic technique pioneered by Krivelevich [12] to show that there are linear hypergraphs with certain ambient properties that will be needed for our stepping-up argument.

Lemma 3.1. *For all sufficiently large k , there is a linear triple system F on k vertices with the following three properties:*

1. *Every vertex subset of order $k^{3/5}$ has at least one edge.*
2. *For any two disjoint subsets $S, T \subset V(F)$ with $|S \cup T| > k/4$, $|S| > \log^2 k$ and $|T| > k^{1/2} \log k$, there is at least one edge with one vertex in S and two vertices in T .*

3. For any disjoint vertex subsets $A_1, \dots, A_t, S \in V(F)$ with $|S| > k/4$, $|A_i| \leq k^{1/2} \log k$, $|A_1 \cup \dots \cup A_t| \geq k/(16 \log k)$, there are at least

$$\frac{|S|}{500k} \sum_{i,j} |A_i||A_j|$$

edges in F with the property that one vertex is in S and the other two are in distinct parts A_i and A_j .

Proof. Consider the binomial random 3-graph $H \sim H^{(3)}(k, p)$ with k vertices where each edge appears independently with probability $p = 1/200k$. Let B denote the unique 3-graph with 4 vertices and 2 edges and let \mathcal{B} denote a maximal collection of edge-disjoint copies of B in H . Form F by starting with H and deleting both edges from every copy of B in \mathcal{B} . Then, by the maximality of \mathcal{B} , the remaining 3-graph F has no copies of B and, hence, is a linear triple system. We will now show that with high probability F has the three required properties. We start with the second and third properties and then observe that the first property can be proved similarly.

Pick $S, T \subset V(F)$ as in the second property of the lemma. Call an edge in H with one vertex in S and two vertices in T an STT edge. Let $X = X_{S,T}$ be the number of STT edges, let $Y = Y_{S,T}$ be the number of copies of B that contain at least one STT edge and let $Z = Z_{S,T}$ be the maximal number of pairwise edge-disjoint copies of B , each containing at least one STT edge. Obviously, $Z \leq Y$. Define the event

$$A = A_{S,T} = \{X > 10Z\}.$$

We note that if $A_{S,T}$ holds for every appropriate S, T , then F satisfies the second property. Indeed, if we delete all edges in all copies of B in \mathcal{B} , then the number of STT edges that are deleted is at most $2Z$, leaving at least one STT edge.

Set $s = |S|$ and $t = |T|$. Since $\mathbb{E}X = ps \binom{t}{2}$ and $p = 1/200k$,

$$\frac{1}{200} \mathbb{E}X = p^2 ks \binom{t}{2} \leq \mathbb{E}Y \leq 3p^2 ks \binom{t}{2} = \frac{3}{200} \mathbb{E}X.$$

Consequently, $a \mathbb{E}Y \leq \mathbb{E}X \leq b \mathbb{E}Y$ for $a = 200/3 > 60$ and $b = 200$. Now

$$\Pr(\bar{A}) = \Pr(X \leq 10Z) \leq \Pr\left(X \leq \frac{\mathbb{E}X}{2}\right) + \Pr\left(Z \geq \frac{\mathbb{E}X}{20}\right) \leq \Pr\left(X \leq \frac{\mathbb{E}X}{2}\right) + \Pr\left(Z \geq \frac{a \mathbb{E}Y}{20}\right).$$

Krivelevich [12, Claim 1] proved that, for any constant $c > 0$,

$$\Pr(Z \geq c \mathbb{E}Y) < e^{-c(\log c - 1) \mathbb{E}Y}$$

and, therefore,

$$\Pr\left(Z \geq \frac{a \mathbb{E}Y}{20}\right) \leq e^{-\frac{a}{20}(\log(a/20) - 1) \mathbb{E}Y} \leq e^{-\frac{a}{20t}(\log(a/20) - 1) \mathbb{E}X}.$$

Note that $\log(a/20) > 1$, since $a > 60$. Moreover, the standard Chernoff bound gives that

$$\Pr\left(X \leq \frac{\mathbb{E}X}{2}\right) < e^{-\mathbb{E}X/8}.$$

Recall that $\mathbb{E}X = ps \binom{t}{2} = s \binom{t}{2} / 200k$. The number of choices for S, T of orders s, t is at most $\binom{k}{s} \binom{k}{t} < k^{s+t}$. Thus, in order to show that $\sum_{S,T} \Pr(\bar{A}_{S,T}) = o(1)$, the inequalities above imply that we just need $st^2/k \gg (s+t) \log k$. Note that $st^2/k \gg s \log k$, since $t > \sqrt{k} \log k$. Moreover, $st^2/k \gg t \log k$, since $s+t > k/4$, $s > \log^2 k$ and $t > \sqrt{k}$. Therefore, F does indeed satisfy the second property with high probability.

We now turn to the third property. Writing $X' = X'_{S,T}$ for the number of STT edges in F , we first observe that for all S, T with $|S| > k/4$ and $|T| \geq k/(16 \log k)$, w.h.p. $X' > \mathbb{E}X'/2$ in F . Indeed, the previous argument shows that w.h.p. the number of STT edges in F is at least $4X'/5 > \mathbb{E}X'/2$, where, in the inequality, we used that, by the Chernoff bound, $X' > 3\mathbb{E}X'/4$ w.h.p.

Now suppose that we are given S and A_1, \dots, A_t as in the third property. Set $T = \cup_i A_i$. Since F is linear and $|A_i| \leq k^{1/2} \log k$ for all i , the number of SA_iA_i edges in F over all i is at most

$$\sum_i |A_i|^2 \leq \frac{k}{k^{1/2} \log k} (k^{1/2} \log k)^2 = k^{3/2} \log k. \quad (5)$$

The STT edges are either SA_iA_j edges for $i \neq j$ or SA_iA_i edges. Note that $\binom{t}{2} > \sum_{i \neq j} |A_i||A_j|$ and $\mathbb{E}X_{S,T} = ps \binom{t}{2}$. In view of (5), and k sufficiently large, we have that $\sum_{i \neq j} |A_i||A_j| > k^2/2 \log^2 k$. Writing $X_{SA_iA_i}$ for the number of SA_iA_i edges, we conclude that w.h.p. the number of SA_iA_j edges over all $i \neq j$ in F is at least

$$X'_{S,T} - \sum_i X_{SA_iA_i} \geq \frac{\mathbb{E}X_{S,T}}{2} - k^{3/2} \log k > \frac{s}{400k} \sum_{i \neq j} |A_i||A_j| - k^{3/2} \log k > \frac{s}{500k} \sum_{i \neq j} |A_i||A_j|.$$

Finally, we consider the first property. To this end, suppose $S \subset V(F)$ with $|S| = k^{3/5}$. Define $X = X_S$, $Y = Y_S$, $Z = Z_S$ in the obvious way as above and let $A = A_S$ be the event that $X > 10Z$. Proceeding as we did for the second property and using that $\mathbb{E}X = p \binom{s}{3}$, $\Pr(\bar{A}) = o(1)$ if the condition $ps^3 \gg s \log k$ holds. But this is equivalent to $s^2 \gg k \log k$, which holds with plenty of room to spare.

We have therefore shown that F satisfies each property w.h.p., so there is indeed a linear F that satisfies all three properties simultaneously. \square

In our analysis, we will also need the following result, which comes from a standard application of the probabilistic method.

Lemma 3.2. *For every $k \geq 3$, there is a positive constant $c = c(k)$ such that, for every $n > 2^k$, there is a graph G on $m = c(n/\log n)^{\frac{\sqrt{k}}{64 \log^5 k}}$ vertices with independence number less than n where every subset with r vertices, for every $k^{1/2}/(16 \log^2 k) \leq r \leq k$, induces fewer than $r^2/\log^3 k$ edges.*

Proof. Consider $G(m, p)$ with $p = m^{-\frac{64 \log^5 k}{\sqrt{k}}}$. Then, summing over all r with $k^{1/2}/(16 \log^2 k) \leq r \leq k$, the expected number of subsets of order r that induce at least $r^2/\log^3 k$ edges is at most

$$\sum_{r \geq \frac{k^{1/2}}{16 \log^2 k}} \binom{m}{r} 2^{r^2} p^{r^2/\log^3 k} < \sum_{r \geq \frac{k^{1/2}}{16 \log^2 k}} m^{2r} m^{-\frac{64r^2 \log^2 k}{\sqrt{k}}} < 1/3.$$

Moreover, the expected number of independent sets of size n is at most

$$\binom{m}{n} (1-p)^{\binom{n}{2}} \leq m^n e^{-pn^2/2} = e^{n \log m - \frac{n^2}{2m \cdot 64 \log^5 k / \sqrt{k}}} < 1/3,$$

where the last inequality follows from the fact that $m = c(n/\log n)^{\frac{\sqrt{k}}{64 \log^5 k}}$ and that c may be taken sufficiently small in terms of k . By Markov's inequality and the union bound, the statement follows. \square

3.2 Stepping up

Proof of Theorem 1.4. Let k be a sufficiently large constant that will be determined later. Let F be the linear triple system on k vertices given by Lemma 3.1 and let G be a graph on $m = c(n/\log n)^{\frac{\sqrt{k}}{64 \log^5 k}}$ vertices, where $V(G) = \{0, 1, \dots, m-1\}$, with the properties described in Lemma 3.2. Set $V = \{0, 1, \dots, 2^m - 1\}$. Given subsets $U_1, U_2 \subset V$, we write $U_1 < U_2$ if $u_1 < u_2$ for all $u_1 \in U_1$ and $u_2 \in U_2$. In what follows, we will use G to define a coloring $\chi : \binom{V}{3} \rightarrow \{\text{red}, \text{blue}\}$ of the triples of V with no red copy of F and no blue clique on $\binom{m+n-1}{n-1}$ vertices.

For each $v \in V$, write $v = \sum_{i=0}^{m-1} v(i)2^i$ with $v(i) \in \{0, 1\}$ for each i . For $u \neq v$, let $\delta(u, v) \in V(G)$ denote the largest i for which $u(i) \neq v(i)$. It is easy to verify the following properties (see, for example, [10]).

Property I: For every triple $u < v < w$, $\delta(u, v) \neq \delta(v, w)$.

Property II: For $v_1 < \dots < v_r$, $\delta(v_1, v_r) = \max_{1 \leq j \leq r-1} \delta(v_j, v_{j+1})$.

We define $\chi : \binom{V}{3} \rightarrow \{\text{red, blue}\}$ as follows. For vertices $v_1 < v_2 < v_3$ in V , let $\delta_i = \delta(v_i, v_{i+1})$ and set $\chi(v_1, v_2, v_3) = \text{red}$ if and only if $\delta_1 > \delta_2$ and $\delta_1 \delta_2 \in E(G)$. We claim that this coloring does not contain a blue clique of order $\binom{m+n-1}{n-1}$.

Lemma 3.3. *The coloring χ does not contain a blue clique on $\binom{m+n-1}{n-1}$ vertices.*

Lemma 3.3 will follow from the claim below. To state it, we need some definitions. Given a subset $A \subset V$, we say that A contains a δ -increasing set of order s if there are vertices $v_1, \dots, v_s \in A$ such that $v_1 < v_2 < \dots < v_s$ and

$$\delta(v_1, v_2) < \delta(v_2, v_3) < \dots < \delta(v_{s-1}, v_s).$$

Likewise, we say that A contains a δ -decreasing set of order t if there are vertices $v_1, \dots, v_t \in A$ such that $v_1 < v_2 < \dots < v_t$ and

$$\delta(v_1, v_2) > \delta(v_2, v_3) > \dots > \delta(v_{t-1}, v_t).$$

Claim 3. Let $A \subset V$ be such that A does not contain a δ -increasing set of order s or a δ -decreasing set of order t . Then $|A| \leq \binom{s+t-4}{t-2}$.

Proof. We proceed by induction on s and t . Let $A = \{v_1, \dots, v_{|A|}\}$, where $v_1 < v_2 < \dots < v_{|A|}$. For the base case $s = 3$ and $t \geq 3$, if A does not contain a δ -increasing set of order 3, then we must have

$$\delta(v_1, v_2) > \delta(v_2, v_3) > \dots > \delta(v_{|A|-1}, v_{|A|}).$$

Since A does not contain a δ -decreasing set of order t , we must have $|A| \leq t - 1 = \binom{s+t-4}{t-2}$. A symmetric argument works for the other base case $s \geq 3$ and $t = 3$.

For the inductive step assume that the statement holds for $s' < s$ or $t' < t$. Suppose A does not contain a δ -increasing set of order s or a δ -decreasing set of order t . Let $j \in \{1, \dots, m-1\}$ be such that

$$\delta_j = \delta(v_j, v_{j+1}) = \max_i \delta(v_i, v_{i+1})$$

and let $A_1 = \{v_1, \dots, v_j\}$ and $A_2 = \{v_{j+1}, \dots, v_{|A|}\}$. Then, by Properties I and II above, A_1 does not contain a δ -increasing set of size $s-1$, since otherwise $A_1 \cup v_{j+1}$ contains a δ -increasing set of size s , a contradiction. Likewise, A_2 does not contain a δ -decreasing set of size $t-1$, since otherwise Properties I and II would imply that $A_2 \cup v_j$ contains a δ -decreasing set of size t , again a contradiction. Therefore, by the induction hypothesis, we have

$$|A| = |A_1| + |A_2| \leq \binom{s+t-5}{t-2} + \binom{s+t-5}{t-3} = \binom{s+t-4}{t-2},$$

as required. □

Proof of Lemma 3.3. Let $A \subset V$ be such that χ colors every triple in A blue. Clearly, A cannot contain a δ -increasing set of order $m+2$. Moreover, by Properties I and II above, A does not contain a δ -decreasing set on $n+1$ vertices, as this would correspond to an independent set of order n in G . Therefore, by Claim 3, we have $|A| \leq \binom{m+n-1}{n-1}$. □

Next, we show that χ does not contain a red copy of F . For the sake of contradiction, suppose χ contains a copy of F with vertex set $V' = \{v_1, \dots, v_k\} \subset V$, where $v_1 < \dots < v_k$, and let $\delta_i = \delta(v_i, v_{i+1})$. In what follows, we will define a vertex partition

$$\mathcal{P}_t : V' = A_1 \cup A_2 \cup \dots \cup A_r \cup S_t \cup B_s \cup B_{s-1} \cup \dots \cup B_1$$

such that the following conditions hold:

1. $t = r + s + 1$.
2. $A_1 < \dots < A_r < S_t < B_s < \dots < B_1$.
3. $|A_i| \leq k^{1/2} \log k$, $|B_j| \leq \log^2 k$ and $|S_t| > k/4$.
4. There are $\delta_1^* > \dots > \delta_r^*$ such that for any $u_1 \in A_i$, $u_2 \in A_j$, $u_3 \in S_t$, where $i < j$, we have $u_1 < u_2 < u_3$ and

$$\delta(u_1, u_2) = \delta_i^* > \delta_j^* = \delta(u_2, u_3).$$
5. There are $\tilde{\delta}_1 > \dots > \tilde{\delta}_s$ such that for any $u_1 \in B_i$, $u_2 \in B_j$, $u_3 \in S_t$, where $i < j$, we have $u_3 < u_2 < u_1$ and

$$\delta(u_3, u_2) = \tilde{\delta}_j < \tilde{\delta}_i = \delta(u_2, u_1).$$
6. $|A_1 \cup \dots \cup A_r| < k/4$ and $|B_1 \cup \dots \cup B_s| < k^{3/4}$.

We start with $\mathcal{P}_1 : V' = S_1$. Suppose we have the partition

$$\mathcal{P}_{t-1} : V' = A_1 \cup \dots \cup A_r \cup S_{t-1} \cup B_s \cup \dots \cup B_1,$$

with the properties described above. If $|A_1 \cup \dots \cup A_r| < \frac{k}{4} - k^{1/2} \log k$ and $|B_1 \cup \dots \cup B_s| < k^{3/4} - \log^2 k$, we define \mathcal{P}_t by partitioning S_{t-1} as follows.

Let $S_{t-1} = \{v_w, v_{w+1}, \dots, v_{w'}\}$, where $w < w'$. We define $w \leq z < w'$ such that

$$\delta(v_z, v_{z+1}) = \max_{w \leq i < w'} \delta(v_i, v_{i+1}).$$

Let $T = \{v_w, \dots, v_z\}$ and $S = \{v_{z+1}, \dots, v_{w'}\}$. Note that $|S \cup T| = |S_{t-1}| > k/4$. Suppose $|T| > k^{1/2} \log k$ and $|S| > \log^2 k$. By Properties I and II, χ colors every triple with two vertices in T and one in S blue. However, by the second property of F , one such triple must be red, which is a contradiction. Hence, we must have either $|T| \leq k^{1/2} \log k$ or $|S| \leq \log^2 k$.

If $|T| \leq k^{1/2} \log k$, then we set $A_{r+1} = T = \{v_w, v_{w+1}, \dots, v_z\}$ and $S_t = S = \{v_{z+1}, v_{z+2}, \dots, v_{w'}\}$ and we have the partition

$$\mathcal{P}_t = A_1 \cup \dots \cup A_{r+1} \cup S_t \cup B_s \cup \dots \cup B_1.$$

Moreover, for $\delta_{r+1}^* = \delta(v_z, v_{z+1})$, this partition satisfies the required properties.

If $|S| \leq \log^2 k$, then we set $S_t = \{v_w, v_{w+1}, \dots, v_z\}$ and $B_{s+1} = \{v_{z+1}, \dots, v_{w'}\}$ and we have the partition

$$\mathcal{P}_t = A_1 \cup \dots \cup A_r \cup S_t \cup B_{s+1} \cup B_s \cup \dots \cup B_1.$$

Moreover, for $\tilde{\delta}_{s+1} = \delta(v_z, v_{z+1})$, this partition again satisfies the required properties.

Let t be the maximum integer such that the partition

$$\mathcal{P}_t = A_1 \cup \dots \cup A_r \cup S_t \cup B_s \cup \dots \cup B_1$$

satisfies all six properties described above. Then $|A_1 \cup \dots \cup A_r| > k/8$ or $|B_1 \cup \dots \cup B_s| > k^{3/4}/2$. The proof now falls into two cases.

Case 1: $|A_1 \cup \dots \cup A_r| > k/8$.

By partitioning dyadically and averaging, we can conclude that there is an integer K and r' indices $j_1 < j_2 < \dots < j_{r'}$ such that

$$2^K \leq |A_{j_i}| < 2^{K+1}$$

for all $i \in \{1, \dots, r'\}$ and

$$\frac{k}{16 \log k} \leq \frac{k}{8 \log_2 k} \leq \sum_{i=1}^{r'} |A_{j_i}| \leq k.$$

Hence, each part satisfies

$$\frac{k}{32r' \log k} \leq |A_{j_i}| \leq \frac{2k}{r'}.$$

In what follows, we will show that $R = \{\delta_{j_1}^*, \dots, \delta_{j_{r'}}^*\}$ induces (nearly) quadratically many edges, contradicting the properties of G .

Since $|S_t| > k/4$, $|A_{j_i}| \leq k^{1/2} \log k$ and $|A_{j_1} \cup \dots \cup A_{j_{r'}}| \geq \frac{k}{16 \log k}$, the third property of F implies that there is an absolute constant c' such that χ colors

$$f \geq \sum_{i,\ell} |A_{j_i}| |A_{j_\ell}| \frac{|S|}{500k} \geq \frac{k^2}{c' \log^2 k}$$

triples, with one vertex in S and the other two in distinct parts A_{j_i}, A_{j_ℓ} , red. On the other hand, by property 4 of the partition, each such red triple is obtained from an edge $\delta_{j_i}^* \delta_{j_\ell}^* \in E(G)$. Since F is linear, the number of such triples going across parts (A_{j_i}, A_{j_ℓ}, S) , with one vertex in each part, is at most $|A_{j_i}| |A_{j_\ell}|$. Writing e_R for the number of edges in the induced subgraph $G[R]$, we obtain

$$f \leq \sum_{\delta_{j_i}^* \delta_{j_\ell}^* \in E(G)} |A_{j_i}| |A_{j_\ell}| \leq e_R \frac{4k^2}{(r')^2}.$$

This yields

$$e_R \geq \frac{(r')^2}{4c' \log^2 k}.$$

From

$$\frac{k}{16 \log k} \leq \sum_{i=1}^{r'} |A_{j_i}| \leq r' (k^{1/2} \log k),$$

we obtain $r' \geq k^{1/2}/(16 \log^2 k)$. On the other hand, since k may be assumed to be sufficiently large, the properties of G imply that $e_R \leq (r')^2/(\log^3 k) < (r')^2/(4c' \log^2 k)$, a contradiction.

Case 2: $|B_1 \cup \dots \cup B_s| > k^{3/4}/2$.

Since $1 \leq |B_i| \leq \log^2 k$, we have $s > k^{3/5}$. Let $u_i \in B_i$ and consider the subset $\{u_1, \dots, u_s\} \subset V'$. Then, by property 5 of the partition, for $u_i > u_j > u_k$, we have $\delta(u_i, u_j) > \delta(u_j, u_k)$. Hence, by Properties I and II, every triple in $\{u_1, \dots, u_s\}$ is blue. However, by the first property of F , there must be at least one red triple, a contradiction.

Since $|V| = 2^m = 2^{c(n/\log n) \frac{\sqrt{k}}{64 \log^5 k}}$, for $n' = \binom{m+n}{n} \leq (m+n)^n \leq m^{2n}$, we have $\log(n') = \Theta_k(n \log n)$ and $\log \log(n') = \Theta_k(\log n)$. Hence, there is a constant c_0 depending only on k such that

$$r_3(F, n') > |V| = 2^m = 2^{c(n/\log n) \frac{\sqrt{k}}{64 \log^5 k}} \geq 2^{c_0 \left(\frac{\log n'}{(\log \log n')^2} \right) \frac{\sqrt{k}}{64 \log^5 k}} > 2^{c_0 (\log n') \frac{\sqrt{k}}{128 \log^5 k}},$$

as required. \square

4 Concluding remarks

4.1 Pair constructions

All of the best known lower-bound constructions for 3-uniform hypergraph Ramsey numbers are “pair constructions”, which can be formally defined as follows.

Definition. A 3-uniform pair construction (or, henceforth, pair construction) on vertex set $[N]$ is a coloring $\chi_{f,g}$ of the complete 3-graph $K_N^{(3)}$ obtained from a pair of functions $f : \binom{[N]}{2} \rightarrow [p]$ and $g : [p]^3 \rightarrow \{\text{red, blue}\}$ by the rule $\chi_{f,g}(\{i, j, k\}) = g(f(i, j), f(j, k), f(k, i))$ for any edge $\{i, j, k\} \in \binom{[N]}{3}$ satisfying $i < j < k$. We call the parameter p the *complexity* of the pair construction.

To see how this definition relates to existing constructions, we note that stepping-up constructions are exactly pair constructions $\chi_{f,g}$ where $f(i, j) = \delta(i, j)$ is the index of the highest significant binary digit where i and j differ and typically have complexity $O(\log N)$. Similarly, inducibility constructions, like those based on random tournaments or 3-colorings, are pair constructions of constant complexity where $f : \binom{[N]}{2} \rightarrow [p]$ is a uniform random function. The two constructions described in this paper also fit firmly into this paradigm.

These observations suggest the following natural question: do all Ramsey colorings of 3-graphs admit low-complexity encodings? More formally, if, for any 3-graphs G and H , we define $r^p(G, H)$ to be the smallest n such that every 2-edge-coloring of $K_n^{(3)}$ of complexity at most p contains either a red copy of G or a blue copy of H , then we have the following broad conjecture.

Conjecture 4.1. *There exists an absolute constant $a > 0$ such that, for all 3-graphs G and H , $r^p(G, H) \geq r(G, H)^a$ whenever $p > r(G, H)^{o(1)}$.*

In other words, we believe that if there is a Ramsey coloring of order N , then there is also a Ramsey coloring of order $N^{\Omega(1)}$ with much smaller complexity. This would at least explain the prevalence of such constructions in Ramsey theory, but might also help shed some light on the possible growth rates for $r(H, K_n^{(3)})$, a topic which we now discuss.

While almost all known Ramsey colorings can be encoded with complexity which is polylogarithmic in $r(G, H)$, we now show that such a strong dependence is not possible in general.

Proposition 4.2. *For $n, p \geq 1$, $r^p(K_{n,n,n}^{(3)}) \leq 2^{O(n \log p)}$. In particular, if $p = 2^{o(n)}$, then $r^p(K_{n,n,n}^{(3)}) = r(K_{n,n,n}^{(3)})^{o(1)}$.*

Proof. The Kővári–Sós–Turán theorem implies that any bipartite graph with parts of order $s \leq t$ and density q with $t \geq q^{-10}$ has a complete subgraph with $\Omega(\log s / \log(2/q))$ vertices in the part of order s and at least $t^{0.9}$ vertices in the part of order t .

Let $\chi_{f,g}$ be a pair construction for $K_N^{(3)}$ with complexity p and $N = 2^{Cn \log p}$ for some large constant $C > 0$. We would like to show that χ contains a monochromatic $K_{n,n,n}^{(3)}$. Divide the vertex set into three intervals I, J, K of order $N/3$. We will apply the Kővári–Sós–Turán theorem to shrink I, J, K until the edges between $I \times J, J \times K$ and $K \times I$ are each monochromatic in f .

First, we build an auxiliary complete bipartite graph $B = (I, J \times K; E)$ and define an edge-coloring of B by $\phi(i, (j, k)) = (f(i, j), f(i, k))$. Observing that ϕ takes at most p^2 values, there must exist c_1 and c_2 for which $\phi(i, (j, k)) = (c_1, c_2)$ for at least $|I||J||K|/p^2$ choices of $(i, j, k) \in I \times J \times K$. By the Kővári–Sós–Turán theorem, we obtain subsets $I' \subseteq I, S \subseteq J \times K$ for which $|I'| = \Omega(\log N / \log(2p^2))$, $|S| \geq (|J||K|)^{0.9}$ and ϕ is constant on $I' \times S$. Letting J', K' be the set of all $j \in J, k \in K$ that appear in S , this implies that $|J'| \geq |S|/|K| \geq |J|^{0.8}$, $|K'| \geq |S|/|J| \geq |K|^{0.8}$ and $f(i, j) = c_1, f(i, k) = c_2$ for all $(i, j, k) \in I' \times J' \times K'$.

Let c_3 be the most common value of $f(j, k)$ between J' and K' . By the Kővári–Sós–Turán theorem, we can find subsets $J'' \subseteq J', K'' \subseteq K'$ where $|J''|, |K''| = \Omega(\log N / \log(2p))$ and $f(j, k) = c_3$ for all pairs $(j, k) \in J'' \times K''$. Observe that, by our choice of N , $|I'|, |J''|, |K''| = \Omega(\log N / \log(2p^2)) \geq n$. Since every triple $(i, j, k) \in I' \times J'' \times K''$ is assigned the same color $\chi_{f,g}(i, j, k) = g(c_1, c_2, c_3)$, we are done. \square

Therefore, if $N = r(K_{n,n,n}^{(3)}) = 2^{\Theta(n^2)}$, the minimum complexity needed to obtain $r^p(K_{n,n,n}^{(3)}) \geq N^{\Omega(1)}$ is $p = 2^{\Omega(n)} \geq 2^{\Omega(\sqrt{\log N})}$. One can show that this is tight by picking uniformly random f and g .

4.2 What growth rates are possible?

Adding to the earlier results of Fox and He [9], this paper shows that there are many fixed 3-graphs H for which $r(H, K_n^{(3)}) = 2^{\Theta_H(n \log n)}$. Strangely, we do not have any examples of H for which $r(H, K_n^{(3)}) = 2^{\Theta_H(n)}$, though it is natural to conjecture their existence.

Conjecture 4.3. *There exists a 3-graph H for which $r(H, K_n^{(3)}) = 2^{\Theta_H(n)}$.*

We now explicitly construct a 3-graph H with $\frac{1}{3} < m_{\text{pair}}(H) < \frac{1}{2}$. Since Theorem 1.1 does not apply to this H and yet $r(H, K_n^{(3)}) \geq 2^{\Omega_H(n)}$ by Theorem 1.2, it may be a promising candidate for verifying Conjecture 4.3.

Proposition 4.4. *There exists a 3-graph H with $\frac{1}{3} < m_{\text{pair}}(H) < \frac{1}{2}$.*

Proof. First we show that if H can be constructed by repeatedly adding edges $\{u, v, w\}$ such that at most one of the three pairs $\{u, v\}$, $\{v, w\}$ and $\{w, u\}$ was previously in an edge together, then $m_{\text{pair}}(H) < 1/2$. Indeed, suppose $f : \partial H \rightarrow V(G)$ is a pair homomorphism with $e(G)/v(G) < \frac{1}{2}$, $H' = H \cup \{u, v, w\}$ and, without loss of generality, among the three pairs only $\{u, v\}$ lies in the shadow of H . Then there is a pair homomorphism $f' : \partial H' \rightarrow V(G')$ where G' is obtained from G by adding two new vertices $x = f'(v, w)$ and $y = f'(w, u)$ and one new edge $(f(u, v), x, y)$. But then

$$\frac{e(G')}{v(G')} = \frac{e(G) + 1}{v(G) + 2} < \frac{1}{2}.$$

A link path is the link 3-graph of a path. We construct H by linearly gluing two link paths, by which we mean that we identify some of the vertices in the two paths while insisting that they do not share edges. Such a gluing satisfies the property above, since each link path can be built iteratively, and so $m_{\text{pair}}(H) < \frac{1}{2}$. It remains to check that $m_{\text{pair}}(H) > \frac{1}{3}$, i.e., that H is not 123-inducible.

Our construction is as follows. Let U with vertices u, u_1, \dots, u_{12} and V with vertices v, v_1, \dots, v_{12} be the links of the paths $u_1 u_2 \dots u_{12}$ and $v_1 v_2 \dots v_{12}$, respectively. We identify certain u_i to certain v_j as follows:

- Glue v_1 to u_2 , v_2 to u_5 , v_5 to u_6 and v_6 to u_1 .
- Glue v_7 to u_{11} , v_8 to u_8 , v_{11} to u_7 and v_{12} to u_{12} .

It is easy to check that this is a linear gluing. We call the resulting graph H and claim that it is not 123-inducible.

Lemma 4.5. *If W is the w -link of the path $w_1 w_2 \dots w_k$, then, in any total order on w, w_1, \dots, w_k that leads to a valid 123-coloring of W , either $w_1 < w_2 > w_3 < w_4 > \dots$ or $w_1 > w_2 < w_3 > w_4 < \dots$.*

Proof. For any valid 123-coloring and any edge abc , we may determine the relative order of a , b and c if we know the label of the 2-edge ab and the relative order of a and b . Thus, since the edges $ww_i w_{i+1}$ and $ww_{i+1} w_{i+2}$ share the 2-edge ww_{i+1} , $ww_{i+1} w_i$ and $ww_{i+1} w_{i+2}$ must have the same relative order. In particular, $w_i < w_{i+1}$ if and only if $w_{i+1} > w_{i+2}$, completing the proof of the lemma. \square

Using this lemma, it is not difficult to show that H is not 123-inducible. Suppose, for the sake of contradiction, that we have a valid total ordering on the vertices of H leading to a 123-coloring. We have four cases:

Case 1: $u_1 < u_2$ and $v_1 < v_2$. In this case, the lemma implies $u_5 < u_6$ and $v_5 < v_6$. However, $v_1 = u_2$ and $v_2 = u_5$, so $u_1 < u_2 = v_1 < v_2 = u_5 < u_6$. However, $v_5 = u_6$ and $v_6 = u_1$, contradicting $v_5 < v_6$.

Case 2: $u_1 > u_2$ and $v_1 > v_2$. This is exactly the same as Case 1, but with the inequality signs flipped.

Case 3: $u_1 < u_2$ and $v_1 > v_2$. The lemma implies that $u_7 < u_8$ and $u_{11} < u_{12}$, while $v_7 > v_8$ and $v_{11} > v_{12}$. Since $v_7 = u_{11}$ and $v_8 = u_8$, we have that $u_7 < u_8 = v_8 < v_7 = u_{11} < u_{12}$. But $v_{11} = u_7$ and $v_{12} = u_{12}$, contradicting $v_{11} > v_{12}$.

Case 4: $u_1 > u_2$ and $v_1 < v_2$. This is exactly the same as Case 3, but with the inequality signs flipped.

Since we have a contradiction in each case, H is not 123-inducible, as required. \square

If we allow families \mathcal{H} instead of single 3-graphs H , then the growth rate $2^{\Theta_{\mathcal{H}}(n)}$ can occur. Indeed, letting \mathcal{H} be the family of links of nonbipartite graphs, we quickly see that the 3-graph whose edges are the cyclic triangles in a random tournament provides the lower bound. On the other hand, the upper bound $2^n - 1$ follows by applying induction to the larger part of the bipartition determined by the link of a vertex in an \mathcal{H} -free 3-graph.

More generally, one might ask what growth rates are possible. For instance, are there fixed 3-graphs H for which $r(H, K_n^{(3)})$ equals, say, $2^{\Theta_H(\sqrt{n})}$, $2^{\Theta_H(n \log \log n)}$ or even $2^{\Theta_H(n \log^* n)}$?

However, the outstanding question on off-diagonal hypergraph Ramsey numbers is to decide whether there are fixed 3-graphs H for which $r(H, K_n^{(3)}) \geq 2^{\omega(n \log n)}$. It seems likely that this should already be the case for $H = K_4^{(3)}$, but our methods seem insufficient for proving this. We even suspect the following should be true.

Conjecture 4.6. *For every $c > 0$, there exists s such that $r(K_s^{(3)}, K_n^{(3)}) \geq 2^{\Omega_s(n^c)}$.*

We note that this is not true if we replace $K_n^{(3)}$ by $K_{n,n,n}^{(3)}$, since it is easy to see that $r(H, K_{n,n,n}^{(3)}) \leq 2^{O_H(n^2)}$ for all H . In fact, using the methods of [5], one can show that for every 3-graph H there exists $\epsilon > 0$ such that $r(H, K_{n,n,n}^{(3)}) \leq 2^{O_H(n^{2-\epsilon})}$. We suspect that it may even be the case that $r(H, K_{n,n,n}^{(3)}) \leq 2^{O_H(n \log n)}$ for all H , which would go some way towards explaining the difficulties in bypassing this bound.

4.3 Linear hypergraphs

The conjecture that $r(F, K_n^{(3)})$ is at most polynomial in n for any linear F was a special case of another unpublished conjecture, saying that if F is 123-inducible, then $r(F, K_n^{(3)})$ is at most polynomial in n . If that had been true, then, together with Theorem 1.2, we would have had a clean characterization saying that $r(F, K_n^{(3)})$ is at most polynomial in n if F is 123-inducible and at least exponential in n if F is not 123-inducible. Instead, Theorem 1.4 leaves us in a messier, though arguably more interesting, situation. If F is linear with s vertices, then $r(F, K_n^{(3)}) \leq r(K_s^{(3)}, K_n^{(3)})$ and the best known upper bound for $r(K_s^{(3)}, K_n^{(3)})$ is $2^{O(n^{s-2} \log n)}$ from [4]. However, for linear F , the arguments in [4] can be extended to prove the stronger upper bound $r(F, K_n^{(3)}) \leq 2^{n^{s-\Omega(s^{1/2})}}$. We even believe that the following may be true.

Conjecture 4.7. *There is an absolute constant $C > 0$ such that $r(F, K_n^{(3)}) \leq 2^{O_F(n^C)}$ for every fixed linear 3-graph F .*

Perhaps this conjecture could also hold with n^C replaced by $n^{o(1)}$ or with the family of linear 3-graphs replaced by the family of 123-inducible 3-graphs, though we are rather less certain about these possibilities.

Finally, we would also like to know whether Theorem 1.4 holds not just for some linear F but for most linear F . To show this, one would just need to show that the pseudorandomness conditions proved in Lemma 3.1 hold with high probability for random (partial) Steiner systems. We may return to this problem, whose study would require methods rather different to those used here, in future work.

Acknowledgements. This research was initiated during a visit to the American Institute of Mathematics under their SQuaREs program.

References

- [1] M. Ajtai, J. Komlós and E. Szemerédi, A note on Ramsey numbers, *J. Combin. Theory Ser. A* **29** (1980), 354–360.
- [2] T. Bohman, A. Frieze and D. Mubayi, Coloring H -free hypergraphs, *Random Structures Algorithms* **36** (2010), 11–25.

- [3] M. Campos, S. Griffiths, R. Morris and J. Sahasrabudhe, An exponential improvement for diagonal Ramsey, preprint available at arXiv:2303.09521 [math.CO].
- [4] D. Conlon, J. Fox and B. Sudakov, Hypergraph Ramsey numbers, *J. Amer. Math. Soc.* **23** (2010), 247–266.
- [5] D. Conlon, J. Fox and B. Sudakov, Erdős–Hajnal-type theorems in hypergraphs, *J. Combin. Theory Ser. B* **102** (2012), 1142–1154.
- [6] D. Conlon, J. Fox and B. Sudakov, An improved bound for the stepping-up lemma, *Discrete Appl. Math.* **161** (2013), 1191–1196.
- [7] P. Erdős and A. Hajnal, On Ramsey like theorems. Problems and results, in Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972), 123–140, Inst. Math. Appl., Southend-on-Sea, 1972.
- [8] P. Erdős and G. Szekeres, A combinatorial problem in geometry, *Compos. Math.* **2** (1935), 463–470.
- [9] J. Fox and X. He, Independent sets in hypergraphs with a forbidden link, *Proc. London Math. Soc.* **123** (2021), 384–409.
- [10] R. L. Graham, B. L. Rothschild and J. H. Spencer, **Ramsey Theory**, 2nd Edition, Wiley-Intersci. Ser. Discrete Math. Optim., John Wiley & Sons, Inc., New York, 1990.
- [11] J. H. Kim, The Ramsey number $R(3, t)$ has order of magnitude $t^2 / \log t$, *Random Structures Algorithms* **7** (1995), 173–207.
- [12] M. Krivelevich, Bounding Ramsey numbers through large deviation inequalities, *Random Structures Algorithms* **7** (1995), 145–155.
- [13] S. Mattheus and J. Verstraëte, The asymptotics of $r(4, t)$, *Annals of Math.* **199** (2024), 919–941.
- [14] D. Mubayi, Improved bounds for the Ramsey number of tight cycles versus cliques, *Combin. Probab. Comput.* **25** (2016), 791–796.
- [15] D. Mubayi and A. Suk, New lower bounds for hypergraph Ramsey numbers, *Bull. London Math. Soc.* **50** (2018), 189–201.
- [16] F. P. Ramsey, On a problem of formal logic, *Proc. London Math. Soc.* **30** (1930), 264–286.