Combinatorics, Probability and Computing (2026), **35**, pp. 554–567
doi:10.1017/S0963548326100443

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

Off-diagonal Ramsey numbers for linear hypergraphs

Xiaoyu He[1], Jiaxi Nie[1], Yuval Wigderson[2] and Hung-Hsun Yu[3]

[1]Georgia Institute of Technology, USA, [2]ETH Zürich, Switzerland, and [3]Princeton University, USA
Corresponding author: Jiaxi Nie; Email: jnie47@gatech.edu

(Received 12 September 2025; revised 9 March 2026; accepted 9 March 2026; first published online 14 April 2026)

Abstract

We study off-diagonal Ramsey numbers $r(H, K_n^{(k)})$ of k-uniform hypergraphs, where H is a fixed linear k-uniform hypergraph and $K_n^{(k)}$ is complete on n vertices. Recently, Conlon, Fox, Gunby, He, Mubayi, Suk, and Verstraëte disproved the folklore conjecture that $r(H, K_n^{(3)})$ always grows polynomially in n. In this paper, we show that much larger growth rates are possible in higher uniformity. In uniformity $k \geq 4$, we prove that for any constant $C > 0$, there exists a linear k-uniform hypergraph H for which

$$r(H, K_n^{(k)}) \geq \mathrm{twr}_{k-2}(2^{(\log n)^C}).$$

Keywords: Hypergraph; Ramsey number; stepping-up

2020 MSC Codes: Primary: 05D10; Secondary: 05D40, 05C65

1. Introduction

Let H_1, H_2 be two k-uniform hypergraphs (k-graphs for short). Their *Ramsey number* $r(H_1, H_2)$ is defined as the least integer N such that every red/blue edge-colouring of the complete k-graph $K_N^{(k)}$ contains a monochromatic red copy of H_1 or a monochromatic blue copy of H_2. In this paper, we will be concerned with the *off-diagonal* regime, where $H_1 = H$ is a fixed k-graph, and $H_2 = K_n^{(k)}$ is a clique, and we are interested in the asymptotics of $r(H, K_n^{(k)})$ as $n \to \infty$. Note that $r(H, K_n^{(k)})$ can also be viewed as the least integer N such that every N-vertex H-free k-graph contains an independent set of size n.

In the special case of $k = 2$, it follows from the classical Erdős–Szekeres bound [12] that $r(H, K_n^{(2)}) = n^{\Theta_H(1)}$, i.e. that $r(H, K_n^{(2)})$ grows polynomially in n for any fixed H. However, our knowledge of the correct order of polynomial growth is limited to extremely few graphs H. For example, while an early result of Erdős [8] implies that $r(K_3^{(2)}, K_n^{(2)}) = n^{2+o(1)}$ (with many further developments [1–3, 13, 15, 16, 23] obtaining much more precise asymptotic bounds), it was only extremely recently that Mattheus and Verstraëte [17] proved that $r(K_4^{(2)}, K_n^{(2)}) = n^{3+o(1)}$. It remains a major open problem to determine the order of polynomial growth of $r(H, K_n^{(2)})$ for most other graphs H, for example when H is a clique of order at least 5, or a cycle of length at least 4.

For hypergraphs of uniformity $k \geq 3$, our understanding is even more limited. For a general fixed H, the best known upper bound on $r(H, K_n^{(k)})$ follows from the stepping-down technique of Erdős and Rado [11], which implies that

$$r\big(H, K_n^{(k)}\big) \leq \mathrm{twr}_{k-1}\big(n^{O_H(1)}\big), \tag{1.1}$$

where the *tower function* is recursively defined by $\mathrm{twr}_1(x) = x$ and $\mathrm{twr}_i(x) = 2^{\mathrm{twr}_{i-1}(x)}$ for $i \geq 2$.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Perhaps surprisingly, the upper bound (1.1) turns out to be tight (up to the implicit constant) in many cases, including when H is a clique of order at least $k+2$ [19]. The proof of this result, and of many other lower bounds [4, 6, 7, 18] on hypergraph Ramsey numbers, uses a variant of the celebrated stepping-up construction of Erdős, Hajnal, and Rado [10], which we shortly discuss in more detail.

Although the bound in (1.1) is tight when H is a clique, it is natural to expect that it is far from tight when H is "far from being a clique". For example, a simple supersaturation argument, going back at least to the early work of Erdős and Hajnal [9], demonstrates that $r(H, K_n^{(k)}) = n^{O_H(1)}$ whenever H is k-partite, or, more generally, *iterated* k-partite.[1] It is conjectured in [5] that this is in fact a complete characterisation: H is iterated k-partite if and only if $r(H, K_n^{(k)}) = n^{O_H(1)}$.

Towards resolving this conjecture, or more generally towards understanding the behaviour of $r(H, K_n^{(k)})$ for general H, it is natural to ask what happens when H is still "far from clique-like", but is not necessarily iterated k-partite. One such class of hypergraphs is the class of *linear* hypergraphs, in which every pair of edges intersects in at most one vertex. A longstanding folklore conjecture in the field was that in uniformity $k=3$, every linear 3-graph H satisfies $r(H, K_n^{(3)}) = n^{O_H(1)}$. However, this conjecture was recently disproved in [4].

Theorem 1.1 ([4, Theorem 1.4]). *For every $C>1$, there exists a linear 3-graph H such that $r(H, K_n^{(3)}) \geq 2^{(\log n)^C}$ for all sufficiently large n.*

In particular, this lends some credence to the conjecture of [5]: since linear hypergraphs need not have polynomial off-diagonal Ramsey numbers, perhaps the only relevant structure is that of being iterated k-partite.

Our main result is a generalisation of Theorem 1.1 to arbitrary uniformities.

Theorem 1.2. *For every $C>1$ and every $k \geq 3$, there exists a linear k-graph H such that*

$$r\big(H, K_n^{(k)}\big) \geq twr_{k-2}\big(2^{(\log n)^C}\big)$$

for all sufficiently large n.

We note that the tower height in Theorem 1.2 is nearly best possible, by (1.1). We also remark that, if we let $K_{n,\dots,n}^{(k)}$ denote the complete k-partite k-graph with parts of size n, then a simple supersaturation argument (see [14, Proposition 7.2]) demonstrates that $r(H, K_{n,\dots,n}^{(k)}) \leq n^{O_H(1)}$ for any linear k-graph H. Thus, Theorem 1.2 implies that $r(H, K_n^{(k)})$ and $r(H, K_{n,\dots,n}^{(k)})$ can be extremely far apart, in that the latter is polynomial and the former grows as a tower of height at least $k-2$. Previously no such separation was known for any H; in particular, if H is a fixed clique, then $r(H, K_n^{(k)})$ and $r(H, K_{n,\dots,n}^{(k)})$ are expected to be of roughly the same order.

In order to prove Theorem 1.2, we need to both define the linear k-graph H and to define a colouring of $E(K_N^{(k)})$ containing no red copy of H and no large blue clique. As in many previous works, our colouring is based on the stepping-up construction of Erdős, Hajnal, and Rado [10]. In particular, our technique takes as input a linear $(k-1)$-graph F, and outputs a linear k-graph H whose off-diagonal Ramsey number is exponential in that of F, as stated in the following proposition.

Proposition 1.3. *Let $k \geq 4$ and let H be a linear $(k-1)$-graph. There exists a linear k-graph H' such that*

$$r\big(H', K_{2n+2k}^{(k)}\big) > 2^{r\left(H, K_n^{(k-1)}\right)-1}.$$

[1] The class of iterated k-partite hypergraphs is defined recursively as follows. First, a hypergraph with no edges is iterated k-partite. Next, H is iterated k-partite if $V(H)$ can be partitioned into k parts $V_1, \dots, V_k$ such that every edge either transverses the k parts or is fully contained in one part, and $H[V_i]$ is iterated k-partite for all i.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Note that Proposition 1.3 immediately implies Theorem 1.2 by induction on k, with Theorem 1.1 serving as the base case.

As discussed above, the colouring used in the proof of Proposition 1.3 is based on the stepping-up technique. Although variants of this technique appear frequently in the literature, most papers simply record certain ad hoc properties needed for the particular application. We therefore present in section 2 a slightly more general treatment of the properties of this construction, in the hopes that future researchers can use some of these in a black-box way. We then present the construction of the linear k-graph H in section 3, and complete the proof of Proposition 1.3, and thus of Theorem 1.2.

2. Stepping-up

As mentioned in the introduction, the problem of estimating $r(H, K_n^{(k)})$ can be equivalently viewed as estimating the minimum independence number of N-vertex H-free k-graph. In the remainder of this paper we use the latter perspective, as our arguments are slightly more natural in this language.

In classical stepping-up constructions, given a $(k-1)$-graph G on $\{0, \dots, N-1\}$, one constructs a k-graph $\widetilde{G}$ whose vertices are all vectors $\varepsilon = (\varepsilon_{N-1}, \varepsilon_{N-2}, \dots, \varepsilon_0) \in \{0,1\}^N$, ordered lexicographically. One also considers, for any two distinct strings $\varepsilon, \varepsilon'$, the largest index $\delta = \delta(\varepsilon, \varepsilon')$ where they differ. Observe that $\varepsilon < \varepsilon'$ if and only if $\varepsilon_\delta < \varepsilon'_\delta$ where $\delta = \delta(\varepsilon, \varepsilon')$. Finally, for any k strings $\varepsilon^{(1)} < \cdots < \varepsilon^{(k)}$, we decide whether they form an edge in $\widetilde{G}$ by some deterministic rule about the tuple $(\delta_1, \dots, \delta_{k-1})$ where $\delta_i = \delta(\varepsilon^{(i)}, \varepsilon^{(i+1)})$. The exact rule depends on the application, though typically if $\delta_1, \dots, \delta_{k-1}$ is monotone, then $\{\varepsilon^{(1)}, \dots, \varepsilon^{(k)}\}$ is an edge if and only if $\{\delta_1, \dots, \delta_{k-1}\}$ is an edge in G.

In this paper, we switch to a new set of notations that represents each vertex in the stepping-up construction with a single number instead of a string. We will also rephrase the conditions on $\delta_1, \dots, \delta_{k-1}$ in terms of binary trees, which should allow one to visualise the constructions better.

First of all, suppose that a, a' are two distinct numbers in $\{0, 1, \dots, 2^N - 1\}$. If their binary representations are $\varepsilon, \varepsilon'$, then it is clear that $a < a'$ if and only if $\varepsilon < \varepsilon'$. More importantly, $\delta(\varepsilon, \varepsilon')$ is simply the largest non-negative integer δ such that $\lfloor 2^{-\delta} a \rfloor \neq \lfloor 2^{-\delta} a' \rfloor$. We generalise this observation to the definition below.

Definition 2.1. *Let S be a set of at least two positive integers. The top splitting level $\ell(S)$ of S is the maximum non-negative integer ℓ such that $\{\lfloor 2^{-\ell} s \rfloor : s \in S\}$ has size at least two. The left subset S_{left} of S is the set of $s \in S$ with $\lfloor 2^{-\ell(S)} s \rfloor$ even, and the right subset S_{right} of S is the set of $s \in S$ with $\lfloor 2^{-\ell(S)} s \rfloor$ odd.*

For example, if $S = \{5, 6, 7, 8, 9\}$, then $\ell(S) = 3$, $S_{\text{left}} = \{5, 6, 7\}$, and $S_{\text{right}} = \{8, 9\}$. Note that, by the definition of $\ell(S)$, for any $s \in S$, there exists an integer q_s and an integer r_s such that $s = q_s 2^{\ell(S)+1} + r_s$ where $0 \leq r_s < 2^{\ell(S)+1}$. Thus $\lfloor 2^{-\ell(S)} s \rfloor$ is either $2q_s$ or $2q_s + 1$, depending on whether $r_s \leq 2^{\ell(S)}$ or not. Hence, S_{left}, S_{right} are well-defined and always non-empty.

To get the full list of $\delta_1, \dots, \delta_{k-1}$, we need to iteratively split the subsets into left subsets and right subsets until they all have size 1. If we record how the subsets split using a binary tree, we get the following definition.

Definition 2.2. *For any set of non-negative integers S, its binary structure $b(S)$ is a weighted rooted ordered binary tree defined iteratively as follows.*

(1) If $|S| = 1$, then $b(S)$ is a single root vertex with weight 1.

(2) If $|S| > 1$, let $b(S)$ be the binary tree so that the left subtree of the root is $b(S_{\text{left}})$, and the right subtree of the root is $b(S_{\text{right}})$. The weight of the root is $|S|$.

Note that the weight of any vertex v in $b(S)$ is equal to the number of leaves in the subtree rooted at v.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

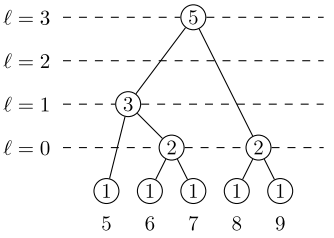


Figure 1. Binary structure of $S = \{5, 6, 7, 8, 9\}$.

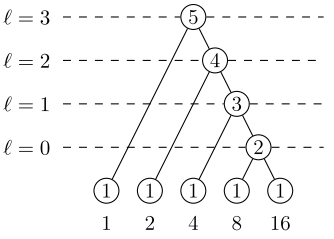


Figure 2. Binary structure of $S = \{1, 2, 4, 8, 16\}$.

For example, the following figure is the binary structure of $S = \{5, 6, 7, 8, 9\}$. To compute the δ's, we can simply read the internal nodes from left to right and see which level they split. For example, for $\varepsilon^{(1)} = (0, 1, 0, 1)$, $\varepsilon^{(2)} = (0, 1, 1, 0)$, . . ., $\varepsilon^{(5)} = (1, 0, 0, 1)$, we have $(\delta_1, \delta_2, \delta_3, \delta_4) = (1, 0, 3, 0)$.

As mentioned earlier, the case where $\delta_1, \ldots, \delta_{k-1}$ is monotone is usually special. Rephrasing this in terms of binary structures gives the following definition.

Definition 2.3. *The binary structure $b(S)$ is increasing if all right subtrees of internal nodes are singletons, and it is decreasing if all left subtrees of internal nodes are singletons. The binary structure is monotone if it is increasing or decreasing. If S is a set of non-negative integers whose binary structure is monotone, let $L(S)$ be iteratively defined as follows.*

(1) If $|S| = 1$, let $L(S)$ be the empty set.

(2) If $|S| > 1$, let S' be S_{left} if $b(S)$ is increasing, and let S' be S_{right} if $b(S)$ is decreasing. Then set $L(S) \stackrel{\text{def}}{=} L(S') \cup \{\ell(S)\}$.

For example, $\{5, 6, 7, 8, 9\}$ does not have a monotone binary structure (Figure 1), whereas $\{1, 2, 4, 8, 16\}$ is decreasing with $L(\{1, 2, 4, 8, 16\}) = \{0, 1, 2, 3\}$ (Figure 2). Note that if S is monotone, then $L(S)$ is simply the set of levels of the internal nodes in the binary structure of S.

Finally, instead of classifying k-tuples using the relative order of $\delta_1, \ldots, \delta_{k-1}$, we classify them based on types of their structures.

Definition 2.4. *A type of binary structure T is a weighted rooted ordered binary tree such that each leaf vertex has a positive integer weight, and each internal vertex has weight equal to the sum of the weights of its children. Its size is the weight of its root. A set of positive integers S is of type T if T can be obtained from $b(S)$ by iteratively removing leaves. The type T is monotone if there is some set of positive integer S of type T with $b(S)$ monotone.*

Before we introduce our reformulation of stepping-up, let us first review the classical construction given by Erdős, Hajnal and Rado [10]. For simplicity, assume that $k = 4$, and we are given a 3-graph G on N vertices with small independence number and clique number. The stepping-up

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

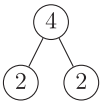


Figure 3. Type $T_{2,2}$.

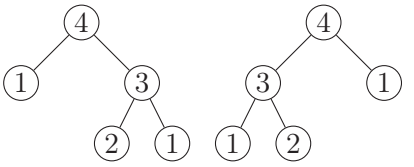


Figure 4. Type $T_{1,(2,1)}$ and type $T_{(1,2),1}$.

of G is a 4-graph on 2^N vertices defined as follows: for any $\{\varepsilon^{(1)}, \ldots, \varepsilon^{(4)}\}$, include it as an edge if one of the following holds:

- $\delta_1, \delta_2, \delta_3$ is monotone and forms an edge in G.
- $\delta_1 < \delta_2 > \delta_3$.

Using our terminology, the first condition is satisfied for some set S of size 4 precisely when $b(S)$ is monotone and $L(S)$ is an edge in G. With some thought, we can see that the second condition is precisely when $b(S)$ is of type $T_{2,2}$, where $T_{2,2}$ is the binary tree with weight 4 at the root and weight 2 at both of its children, both of which are leaves (Figure 3).

Our variant of stepping-up will also include two such conditions, which we now define in turn. We begin with the monotone edges.

Definition 2.5. *Let G be a $(k-1)$-graph on $\{0, 1, \ldots, N-1\}$. Its left stepping-up is the k-graph on $\{0, \ldots, 2^N - 1\}$ consisting of all edges e with $b(e)$ increasing and $L(e) \in E(G)$. Its right stepping-up is the k-graph on $\{0, \ldots, 2^N - 1\}$ consisting of all edges e with $b(e)$ decreasing and $L(e) \in E(G)$.*

Now we can define the stepping-up that we will use throughout the paper.

Definition 2.6. *For any tuple $(G_1, G_2, \mathcal{T})$ where G_1, G_2 are two $(k-1)$-graphs on $\{0, 1, \ldots, N-1\}$ and $\mathcal{T}$ is a family of types of binary structure of size k, its stepping-up is the k-graph G on vertex set $\{0, \ldots, 2^N - 1\}$ formed by taking the edge union*

$$G \stackrel{\text{def}}{=} \widetilde{G_1} \cup \widetilde{G_2} \cup G_{\mathcal{T}}$$

where $\widetilde{G_1}$ is the left stepping-up of G_1, $\widetilde{G_2}$ is the right stepping-up of G_2, and $G_{\mathcal{T}}$ is the k-graph with

$$E(G_{\mathcal{T}}) = \left\{ e \in \binom{\{0, \ldots, 2^N - 1\}}{k} : e \text{ is of type } T \text{ for some } T \in \mathcal{T} \right\}.$$

For example, the classical stepping-up construction discussed above for $k = 4$ is precisely the stepping-up of $(G, G, \{T_{2,2}\})$. Note that the complement can also be described as a stepping-up: it is the stepping-up of $(G^c, G^c, \{T_{1,(2,1)}, T_{(1,2),1}\})$ where $T_{1,(2,1)}$ and $T_{(1,2),1}$ are drawn in Figure 4 below.

Having rephrased the classical construction, we now turn to bounding its independence number, recalling that our ultimate goal is to construct a k-graph with no large independent sets. To bound the independence number of the stepped-up hypergraph, we define the following auxiliary functions.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Definition 2.7. *For any positive integers n_1, n_2 and any family $\mathcal{T}$ of types of binary structure, let $f(n_1, n_2, \mathcal{T})$ be the maximum possible size of set S of non-negative integers that does not contain any n_1-set with an increasing binary structure, any n_2-set with a decreasing binary structure, or any set of type T for any $T \in \mathcal{T}$.*

The following statement then follows easily from the definition.

Lemma 2.8. *Let G_1, G_2 be two $(k-1)$-graphs on $\{0, 1, \ldots, N-1\}$, and let $\mathcal{T}$ be a family of types of binary structure of size k. Let G be the stepping-up of $(G_1, G_2, \mathcal{T})$. Then*

$$\alpha(G) \leq f\left(\alpha(G_1)+2, \alpha(G_2)+2, \mathcal{T}\right).$$

Proof. Let $S \subseteq \{0, \ldots, 2^N - 1\}$ be an independent set in G. If S contains an $(\alpha(G_1)+2)$-set S' with an increasing binary structure, then $\left|L(S')\right| = \alpha(G_1)+1 > \alpha(G_1)$, hence $L(S')$ contains an edge e' in G_1. This shows that S contains a k-set e with $b(e)$ increasing and $L(e) = e'$, which is a contradiction. Therefore S does not contain any $(\alpha(G_1)+2)$-set S' with $b(S')$ increasing. Similarly, S does not contain any $(\alpha(G_2)+2)$-set S' with $b(S')$ decreasing. Lastly, from the way G is constructed, it is clear that S does not contain any sets of type T for any $T \in \mathcal{T}$. As a consequence, $|S| \leq f\left(\alpha(G_1)+2, \alpha(G_2)+2, \mathcal{T}\right)$. □

In order to apply Lemma 2.8, we also need to compute $f(n_1, n_2, \mathcal{T})$. This can be done by dynamic programming if $\mathcal{T}$ is explicitly known. Here we carry out the computation for two particularly useful cases.

For the first case, let $T_{a,b}$ be the type of binary structure where the children of the root are leaves with weight a on the left and b on the right. We show that $f(n_1, n_2, \mathcal{T})$ is linear in $n_1 + n_2$ if $\mathcal{T}$ contains some type of the form $T_{a,b}$.

Lemma 2.9. *Let n_1, n_2, k be positive integers with $n_1, n_2 \geq 2$. Let $\mathcal{T}$ be a family of types of binary structure of size k containing types of the form $T_{a,b}$ for some $a + b = k$. Furthermore, let $a_{\min}$ be the minimum possible value for a, and $b_{\min}$ be the minimum possible value for b. Then*

$$f\left(n_1, n_2, \mathcal{T}\right) \leq (a_{\min}-1)(n_2-2) + (b_{\min}-1)(n_1-2) + 2k.$$

Proof. We will proof by induction on $n_1 + n_2$. The case $n_1 + n_2 = 4$ holds trivially as $n_1 = n_2 = 2$ in this case, which forces $f(n_1, n_2, \mathcal{T}) = 1$ as any set of size 2 has an increasing binary structure. The inductive step is going to be established by the inequality

$$f(n_1, n_2, \mathcal{T}) \leq \max\left\{a_{\min} - 1 + f(n_1, n_2-1, \mathcal{T}), b_{\min} - 1 + f(n_1-1, n_2, \mathcal{T}), 2k\right\} \tag{2.1}$$

Indeed, assuming the claim holds for $f(n_1, n_2-1, \mathcal{T})$ and $f(n_1-1, n_2, \mathcal{T})$, then we have

$$a_{\min} - 1 + f(n_1, n_2-1, \mathcal{T}) \leq (a_{\min}-1)(n_2-2) + (b_{\min}-1)(n_1-2) + 2k,$$

$$b_{\min} - 1 + f(n_1-1, n_2, \mathcal{T}) \leq (a_{\min}-1)(n_2-2) + (b_{\min}-1)(n_1-2) + 2k,$$

and

$$2k \leq (a_{\min}-1)(n_2-2) + (b_{\min}-1)(n_1-2) + 2k.$$

Therefore (2.1) and the inductive hypothesis together imply

$$f(n_1, n_2, \mathcal{T}) \leq (a_{\min}-1)(n_2-2) + (b_{\min}-1)(n_1-2) + 2k,$$

and it remains to prove (2.1).

Suppose that S is a set with no n_1-sets with increasing binary structures, no n_2-sets with decreasing binary structures, and no sets of type T for any $T \in \mathcal{T}$. Note that S_{right} now cannot contain any (n_2-1)-set with decreasing binary structure: indeed, if $S' \subseteq S_{\text{right}}$ satisfies $\left|S'\right| = n_2 - 1$ and $b(S')$ is decreasing, then $\{s\} \cup S'$ also has decreasing structure for any $s \in S_{\text{left}}$, which is a contradiction.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Therefore

$$\left|S_{\text{right}}\right| \leq f(n_1, n_2 - 1, \mathcal{T}).$$

Similarly

$$|S_{\text{left}}| \leq f(n_1 - 1, n_2, \mathcal{T}).$$

If $|S_{\text{left}}| < a_{\min}$ or $\left|S_{\text{right}}\right| < b_{\min}$, then (2.1) now follows immediately. Therefore let us now assume $|S_{\text{left}}| \geq a_{\min}$ and $\left|S_{\text{right}}\right| \geq b_{\min}$. Since S does not contain any set of type $T_{a_{\min}, k-a_{\min}}$ or $T_{k-b_{\min}, b_{\min}}$, we see that $\left|S_{\text{right}}\right| < k - a_{\min}$ and $|S_{\text{left}}| < k - b_{\min}$. This shows that $|S| < 2k - a_{\min} - b_{\min} < 2k$, as desired. □

We will use Lemma 2.9 to get a better quantitative dependence when stepping-up. However, if one is only interested in understanding the tower heights of Ramsey numbers, then it suffices to use the much cheaper bound below.

The *depth* of a binary tree is the number of edges in the longest path from the root to a leaf.

Lemma 2.10. *Let k, d be two non-negative integers. If T is a type of binary structure of size k and depth d, then*

$$f(n_1, n_2, \{T\}) \leq k(n_1 + n_2)^d.$$

Proof. We induct on d. If $d = 0$, then any set of non-negative integers of size $|T| = k$ is of type T, which shows that $f(n_1, n_2, \{T\}) \leq k - 1 \leq k(n_1 + n_2)^d$, proving the base case of the induction.

Now suppose that the claim holds for all smaller depths, and we wish to prove it for d. We now do an internal induction on $n_1 + n_2$. The base case $n_1 = n_2 = 1$ is again trivial, since any 1-set has an increasing binary structure, hence $f(1, 1, \{T\}) = 0 \leq k(1 + 1)^d$. So we may now assume that the claim holds for any positive integers n_1', n_2' with $n_1' + n_2' < n_1 + n_2$.

Let S be a set of non-negative integers that avoids n_1-sets with increasing binary structures, n_2-sets with decreasing binary structures, and k-sets of type T. Let T_{left} and T_{right} be the left and right subtrees of the root of T, respectively. Note that both of these trees have order at most k and depth at most $d - 1$. Suppose that $|S| \geq 2$. As in the previous proof, we know that S_{left} does not contain any $(n_1 - 1)$-sets with increasing binary structures. Similarly, S_{right} does not contain any $(n_2 - 1)$-sets with decreasing binary structures.

Moreover, it cannot simultaneously hold that S_{left} contains a set of type T_{left}, and that S_{right} contains a set of type T_{right}, as these would combine to give a set of type T in S. If S_{left} does not contain a set of type T_{left}, then

$$\begin{aligned} |S| &= |S_{\text{left}}| + \left|S_{\text{right}}\right| \\ &\leq f(n_1 - 1, n_2, \{T_{\text{left}}\}) + f(n_1, n_2 - 1, \{T\}) \\ &\leq k(n_1 - 1 + n_2)^{d-1} + k(n_1 + n_2 - 1)^d \\ &\leq k(n_1 + n_2)^d, \end{aligned}$$

where we use the inductive hypothesis in the penultimate step and the inequality $(x - 1)^{d-1} + (x - 1)^d \leq x^d$ (valid for all integers $x, d \geq 1$) in the final step. Similarly, if S_{right} does not contain a set of type T_{right}, then

$$\begin{aligned} |S| &= |S_{\text{left}}| + \left|S_{\text{right}}\right| \\ &\leq f(n_1 - 1, n_2, \{T\}) + f(n_1, n_2 - 1, \{T_{\text{right}}\}) \\ &\leq k(n_1 - 1 + n_2)^d + k(n_1 + n_2 - 1)^{d-1} \\ &\leq k(n_1 + n_2)^d. \end{aligned}$$

In either case we get the desired bound, completing the induction. □

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

We briefly remark that from Lemma 2.9, it is easy to see that the stepping-up construction for $k = 4$ of Erdős, Hajnal and Rado only grows the independence number at most linearly. To bound the clique number, we can apply Lemma 2.10 to see that the independence number of the complement only grows at most quadratically as $T_{1,(2,1)}$ has depth 2. To show that the clique number grows at most linearly, one would need to compute $f(n_1, n_2, \{T_{1,(2,1)}, T_{(1,2),1}\})$, and it is not hard to see that once we take $T_{(1,2),1}$ into consideration, this function does indeed grow linearly in n_1 and n_2. We remark too that one can similarly rephrase stepping-up for general uniformities using our binary structures framework, but we omit the details as this does not improve on the classical construction.

3. Linear hypergraphs

We now turn to the proof of our technical result, Proposition 1.3. The main remaining difficulty is the construction of the linear k-graph; the following result gives the main property we need of this construction, namely that it is avoided by the stepping-up constructions discussed in the previous section.

Theorem 3.1. *Let $k \geq 4$, and let H be a linear $(k-1)$-graph with no isolated vertices. Then there exists a linear k-graph H' such that the following holds. Suppose that G is a $(k-1)$-graph that is H-free, and $\mathcal{T}$ is a family of types of binary structure of size k that are not monotone. Then the stepping-up $\widetilde{G}$ of $(G, G, \mathcal{T})$ is H'-free.*

We remark that we need $k \geq 4$ because any type of binary structure of size at most 3 is monotone. Given Theorem 3.1, we can readily complete the proof of Proposition 1.3 and hence of Theorem 1.2.

Proof of Theorem 1.3. For any linear $(k-1)$-graph H, let H' be the linear k-graph obtained from Theorem 3.1. By definition, there exists a $(k-1)$-graph G on $\{0, 1, \ldots, r(H, K_n^{(k-1)}) - 2\}$ that is H-free with $\alpha(G) \leq n - 1$. Let $\mathcal{T}$ be $\{T_{2,k-2}, T_{k-2,2}\}$. As $k \geq 4$, we know that both children have weight at least $k - 2 \geq 2$ and so $T_{2,k-2}$ and $T_{k-2,2}$ are not monotone.

Let $\widetilde{G}$ be the stepping-up of $(G, G, \mathcal{T})$. Then we know that $\widetilde{G}$ is H'-free and

$$\alpha(\widetilde{G}) \leq 2\alpha(G) + 2k < 2n + 2k$$

by Lemmas 2.8 and 2.9. This shows that

$$r\big(H', K_{2n+2k}^{(k)}\big) > |V(\widetilde{G})| = 2^{r\left(H, K_n^{(k-1)}\right)-1},$$

as desired. □

We now turn to the proof of Theorem 3.1, which occupies the remainder of this section. We first make the following two definitions, which arise naturally when analysing stepping-up constructions.

Definition 3.2. *A dyadic partition of a set A is a partition $A = A_1 \sqcup A_2 \sqcup \cdots \sqcup A_t$ with the property that $|A_i| \leq \frac{1}{2}(|A_i| + \cdots + |A_t|)$ for all $i < t$ and $|A_t| = 1$.*

The following lemma captures some properties we need of our construction of a linear k-graph.

Lemma 3.3. *Let $k \geq 3$, and fix an ordered linear k-graph $F^<$. There exists a linear k-graph H' with the following property. For every dyadic partition $V(H') = A_1 \sqcup \ldots \sqcup A_t$, and for every two-colouring $\beta : [t] \to \{L, R\}$, and for every ordering π of $V(H')$, there exists an ordered monochromatic transversal copy of $F^<$. That is, there is a copy of $F^<$ each of whose vertices lie in a different part A_i, such that all these A_i receive the same colour under β, and such that the order of these vertices under π agrees with the fixed ordering of $F^<$.*

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

We remark that Lemma 3.3 works for $k = 3$. However, our stepping-up construction only works for $k \geq 4$, as in the statement of Theorem 3.1. We will apply Lemma 3.3 with $F^<$ chosen to be a certain k-graph derived from a given $(k-1)$-graph, which we now define.

Definition 3.4. *Let H be a $(k-1)$-graph. Its expansion H^+ is the k-graph obtained from H by adding to each $e \in E(H)$ a new vertex v_e, which participates in no other edges.*

An ordered expansion of H is any ordering of $V(H^+)$ in which, for all $e \in E(H)$, the vertex v_e is the smallest vertex of $e \cup \{v_e\}$.

For example, we may obtain an ordered expansion by putting all the new vertices v_e first, followed by all the old vertices of H. Using Lemma 3.3 and the properties of the stepping-up construction discussed above, we are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. Let $F^<$ be an ordered expansion of H, and let H' be the k-graph given by Lemma 3.3. We will show that H' satisfies the requirement of Theorem 3.1. Suppose for the sake of contradiction that $\widetilde{G}$ contains a copy of H'. In this case, we will think of $V(H')$ as a subset of $V(\widetilde{G})$ and produce a dyadic partition $A_1 \sqcup \cdots \sqcup A_t$ with a colouring $\beta : [t] \to \{L, R\}$ of $V(H')$ as follows. Starting with $B_0 = V(H')$, suppose that we have defined B_i and $A_1, \ldots, A_i$ so that $V(H') = A_1 \sqcup \cdots \sqcup A_i \sqcup B_i$. If B_i is a singleton, then we set $t = i+1$, $A_{i+1} = B_i$ and then terminate. Therefore we will now assume that $|B_i| \geq 2$, and so its left and right subsets are well-defined and non-empty. Let (A_{i+1}, B_{i+1}) be a permutation of $((B_i)_{\text{left}}, (B_i)_{\text{right}})$ so that $|A_{i+1}| \leq |B_{i+1}|$. Moreover, we set $\beta(i+1) = L$ if A_{i+1} is the left subset of B_i, and $\beta(i+1) = R$ otherwise. Finally, we increase i by 1 and continue the process until it terminates. To check that $A_1 \sqcup \cdots \sqcup A_t$ is indeed a dyadic partition of $V(H')$, note that the procedure guarantees that it is a partition with $|A_t| = 1$. In addition, for all $1 \leq i < t$, we have that $|A_i| \leq \frac{1}{2}(|A_i| + |B_i|) = \frac{1}{2}(|A_i| + \cdots + |A_t|)$.

Having constructed the dyadic partition and the colouring, we now construct the ordering π of $V(H')$. For any $v \in V(H')$, let $i(v)$ be the unique i such that $v \in A_i$. Then for any $u, v \in V(H')$, we set $u >_\pi v$ if $i(u) < i(v)$. If $i(u) = i(v)$, then we arbitrarily order u and v with respect to π.

By Lemma 3.3, there is an ordered monochromatic transversal copy of $F^< = H^+$. We will use this copy to find a copy of H in G, which gives a contradiction. First assume that the ordered monochromatic transversal copy of H^+ is monochromatic of color R. We will again think of $V(H^+)$ as a subset of $V(H')$, which is a subset of $V(\widetilde{G})$. By the definition of H^+, there is an injection $\varphi : V(H) \to V(H^+)$ such that for each edge $e \in E(H)$, we have that $\varphi(e)$ is contained in an edge e' in $V(H^+)$ and $\varphi(e)$ does not contain the smallest element v_e (with respect to π) in e'. Now we define $\varphi^* : V(H) \to V(G) = \{0, 1, \ldots, N-1\}$ as $\varphi^*(v) = \ell\left(B_{i(\varphi(v))-1}\right)$, the top splitting level of $B_{i(\varphi(v))-1}$. We will show that this embeds H into G.

For each $e \in H$, we can order the vertices in e as $v_1, \ldots, v_{k-1}$ so that $i(\varphi(v_1)) > \cdots > i(\varphi(v_{k-1}))$. (Recall that they are distinct as the copy of H' is transversal). By the definition of π, we have that $v_e <_\pi \varphi(v_1) <_\pi \cdots <_\pi \varphi(v_{k-1})$, which shows that $i(v_e) > i(\varphi(v_1))$ again by the definition of π. Now for every $j \in [k-1]$, we will show that the top splitting level of $\{v_e, \varphi(v_1), \ldots, \varphi(v_j)\}$ is $\varphi^*(v_j)$ and its right subset is the singleton $\{\varphi(v_j)\}$. Note that by construction, it is clear that $B_{i(\varphi(v_j))-1} \supseteq \{v_e, \varphi(v_1), \ldots, \varphi(v_j)\}$. Moreover, as $\beta(i(\varphi(v_j))) = R$, the left subset of $B_{i(\varphi(v_j))-1}$ is $B_{i(\varphi(v_j))}$, which contains $v_e, \varphi(v_1) \ldots, \varphi(v_{j-1})$ as $i(v_e) > i(\varphi(v_1)) > \cdots > i(\varphi(v_{j-1})) > i(\varphi(v_j))$; whereas its right subset is $A_{i(\varphi(v_j))}$, which contains $\varphi(v_j)$ by definition. Therefore the top splitting level of $\{v_e, \varphi(v_1), \ldots, \varphi(v_j)\}$ is $\ell(B_{i(\varphi(v_j))-1}) = \varphi^*(v_j)$, and its right subset is $\{\varphi(v_j)\}$.

This immediately implies that $\{v_e\} \cup \varphi(e)$ has an increasing binary structure with $L(\{v_e\} \cup \varphi(e))$ is $\varphi^*(e)$. Note that for any $T \in \mathcal{T}$, as T is not monotone, we know that $\{v_e\} \cup \varphi(e)$ is not of type T. We also know that $\{v_e\} \cup \varphi(e)$ is not decreasing as $k \geq 3$. Therefore $\{v_e\} \cup \varphi(e)$ must be in the left stepping-up of G, and so $\varphi^*(e) = L(\{v_e\} \cup \varphi(e)) \in E(G)$ by definition. As a consequence, φ^* is an embedding from H to G, which is a contradiction with the assumption that G is H-free. Therefore we have reached a contradiction from assuming that there is a monochromatic ordered transversal copy of H^+ in H with color R.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

The other case where the ordered transversal copy of H^+ is monochromatic of color L can be dealt with analogously. We swap the roles of left and right, and the rest of the argument holds verbatim. As we reach a contradiction either way, we see that $\widetilde{G}$ must be H'-free. □

3.1 Properties of dyadic partitions

All that remains is to prove Lemma 3.3, which we now turn to. In this subsection, we record two useful facts we will need about the sizes of various parts in a dyadic partition, and which we use in the next subsection to complete the proof. The first demonstrates that it is not possible to put too many elements of A into a small collection of the sets A_i.

Lemma 3.5. *Let $A = A_1 \sqcup A_2 \sqcup \cdots \sqcup A_t$ be a dyadic partition. Then for every $I \subseteq [t]$,*

$$|A| - \sum_{i \in I} |A_i| \geq \left\lfloor \frac{|A|}{2^{|I|}} \right\rfloor.$$

Proof. Let $I = \{i_1, i_2, \ldots, i_s\}$ where $i_1 < i_2 < \ldots < i_s$, and let $b_j = |A| - \sum_{k=1}^{j} |A_{i_k}|$. Suppose first that $i_s < t$. By the definition of a dyadic partition, we have $b_1 \geq |A|/2$. For $j \geq 2$, suppose $b_{j-1} \geq |A|/2^{j-1}$. Then

$$\begin{aligned} b_j = |A| - \sum_{k=1}^{j} |A_{i_k}| &= \frac{1}{2}\left(|A| - \sum_{k=1}^{j-1} |A_{i_k}|\right) + \frac{1}{2}\left(|A| - \sum_{k=1}^{j-1} |A_{i_k}|\right) - |A_{i_j}| \\ &\geq \frac{b_{j-1}}{2} + \frac{1}{2}\left(|A| - \sum_{k=1}^{j-1} |A_{i_k}| - \sum_{i=i_j}^{t} |A_i|\right) \geq \frac{|A|}{2^j}, \end{aligned}$$

where the first inequality uses the fact that $|A_{i_j}| \leq \frac{1}{2}\sum_{i=i_j}^{t} |A_i|$, which follows from the definition of a dyadic partition and the fact that $i_j < t$, and the second inequality uses the inductive hypothesis and the fact that the term in parentheses is non-negative.

This completes the proof in case $i_s < t$. Additionally, the result is vacuous if $|A| < 2^{|I|}$. So we may assume that $i_s = t$ and $|A|/2^{|I|-1} > 2$. We then apply the argument above to $I \setminus \{t\}$ to conclude that

$$|A| - \sum_{i \in I} |A_i| = |A| - \sum_{i \in I \setminus \{t\}} |A_i| - |A_t| \geq \frac{|A|}{2^{|I|-1}} - 1 \geq \left\lfloor \frac{|A|}{2^{|I|}} \right\rfloor,$$

where we use the fact that $|A_t| = 1$ in a dyadic partition, as well as the inequality $x - 1 \geq \lfloor x/2 \rfloor$, valid for all real $x \geq 1$. □

Our next lemma, which will be crucial in the analysis of the stepping-up construction, is a simple corollary of Lemma 3.5.

Lemma 3.6. *Let $A = A_1 \sqcup A_2 \sqcup \cdots \sqcup A_t$ be a dyadic partition. For every $a \in A$, let $i(a)$ be the unique index i such that $a \in A_i$. For every integer $s < t/2$ and every 2-colouring $\beta : [t] \to \{L, R\}$, there exists $X \in \{L, R\}$ such that for every $I \subseteq [t]$ with $|I| = s - 1$,*

$$|\{a \in A \mid i(a) \notin I, \beta(i(a)) = X\}| \geq \left\lfloor \frac{|A|}{2^{2s-1}} \right\rfloor.$$

Proof. Suppose for contradiction that for each $X \in \{L, R\}$ there exists I_X with $|I_X| = s - 1$ such that

$$|\{a \in A \mid i(a) \notin I_X, \beta(i(a)) = X\}| < \left\lfloor \frac{|A|}{2^{2s-1}} \right\rfloor.$$

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Then we have

$$|A| - \sum_{i \in I_L \cup I_R} |A_i| \leq \sum_{X \in \{L,R\}} |\{a \in A \mid i(a) \notin I_X, \beta(i(a)) = X\}| < 2 \left\lfloor \frac{|A|}{2^{2s-1}} \right\rfloor \leq \left\lfloor \frac{|A|}{2^{2s-2}} \right\rfloor.$$

Note that $|I_L \cup I_R| \leq 2s-2$, hence by Lemma 3.5,

$$|A| - \sum_{i \in I_L \cup I_R} |A_i| \geq \frac{|A|}{2^{|I_L \cup I_R|}} \geq \left\lfloor \frac{|A|}{2^{2s-2}} \right\rfloor.$$

This gives a contradiction, and therefore completes the proof. □

3.2 Proof of Lemma 3.3

In this section we prove Lemma 3.3, thus completing the proof of Theorem 1.2.

Let us define an *oriented s-graph* to be an s-graph H_0, as well as a bijection $\psi_e : e \to [s]$ for every edge $e \in E(H_0)$. Note that in the case $s = 2$, this precisely recovers the definition of an oriented graph. We say an oriented s-graph is *linear* if its underlying unoriented s-graph is linear.

In the proof of Lemma 3.3, we will use the following statement, which is simply the same statement in case $F^<$ is an ordered s-uniform edge.

Lemma 3.7. *Let $s \geq 3$. There exists a linear oriented s-graph H_0 with the following property. For every dyadic partition $V(H_0) = A_1 \sqcup \ldots \sqcup A_t$, and for every two-colouring $\beta : [t] \to \{L, R\}$, and for every ordering π of $V(H_0)$, there exists an ordered monochromatic transversal edge. That is, there exists an edge $e \in E(H_0)$ which goes between s distinct parts A_i, all of which receive the same colour under β, such that the ordering given by ψ_e agrees with the ordering induced by π on e.*

Proof. Let n be a sufficiently large integer with respect to s, and let $c > 0$ be a sufficiently small constant. Let $p = cn^{2-s}$. Let $H_0 \sim G_s(n, p)$ be a random s-graph on k vertices in which each s-tuple is made an edge with probability p, independently over all choices. Additionally, we make H_0 into an oriented s-graph in a uniformly random way: each edge $e \in E(H_0)$ selects a uniformly random bijection $\psi_e : e \to [s]$, independently of all other choices. Let $\mathcal{E}$ be the event that there is an ordered monochromatic transversal edge in H_0 for every dyadic partition of $V(H_0)$, every two-colouring of $[t]$, and for every ordering of $V(H_0)$. We begin by estimating the probability of $\mathcal{E}$, and then show that with positive probability we can ensure both that $\mathcal{E}$ holds and that H_0 is linear.

Let us fix a dyadic partition $V(H_0) = A_1 \sqcup \ldots \sqcup A_t$, a colouring $\beta : [t] \to \{L, R\}$, and an ordering π of $V(H_0)$. Note that there are at most n^n dyadic partitions of a set of order n (because every such partition yields a unique function $[n] \to [t]$, and $t \leq n$), there are $2^t \leq 2^n$ choices for β, and there are $n! < n^n$ orderings of $V(H_0)$. So in total we need to union-bound over at most $(2n)^{2n}$ choices for these three data.

For any $a \in A_1 \sqcup \cdots \sqcup A_t$, let $i(a)$ be the unique i such that $a \in A_i$. We begin by estimating the number of potential monochromatic transversal edges. This number is exactly

$$T \overset{\text{def}}{=} \frac{1}{s!} \#\{a_1, \ldots, a_s \in V(H_0) : i(a_1), \ldots, i(a_s) \text{ distinct, and } \beta(i(a_1)) = \cdots = \beta(i(a_s))\}.$$

We first claim that $T = \Omega_s(n^s)$, which we prove as follows. Let $X \in \{L, R\}$ be such that the conclusion of Lemma 3.6 holds.[2] Then for any $i_1, \ldots, i_{s-1} \in [t]$, there exist at least $\Omega_s(n)$ choices for a_s with $i(a_s) \notin \{i_1, \ldots, i_{s-1}\}$ and $\beta(i(a_s)) = X$. This shows that for any $a_1, \ldots, a_m$ with $i(a_1), \ldots, i(a_m)$ distinct, $\beta(i(a_1)) = \cdots = \beta(i(a_m)) = X$ and $m < s$, there are at least $\Omega_s(n)$ ways to extend this sequence by a_{m+1} while maintaining the properties. This shows $T = \Omega_s(n^s)$, as claimed.

[2] Recall that n is sufficiently large, hence $t \geq \log n$ is also sufficiently large, so we may assume $t > 2s$ in order to apply Lemma 3.6

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

Now, let Z be the random variable counting the number of monochromatic transversal edges in H_0 which are correctly ordered, i.e. whose ordering ψ_e agrees with π. Recall that the random ordering is chosen independently of the other random choices, showing that Z is distributed as a binomial random variable with parameters T and $p/s!$. Therefore,

$$\mathbb{P}(Z=0) = (1-p/s!)^T \leq e^{-pT/s!} \leq e^{-\Omega_s(pn^s)} = e^{-\Omega_s(cn^2)}.$$

Applying the union bound over the at most $(2n)^{2n}$ choices for the dyadic partition as well as β and π, we see that with probability at least $1-e^{-\Omega_s(cn^2)}$, we have the claimed property for all possible choices of these three data. In other words, $\mathbb{P}(\mathcal{E}) \geq 1-e^{-\Omega_s(cn^2)}$.

It remains to ensure that H_0 is linear. For a pair of vertices u, v, let $\mathcal{E}_{u,v}$ be the event that u and v are contained in at most one edge of H_0. By the union bound, $\mathbb{P}(\overline{\mathcal{E}_{u,v}}) \leq 4^s n^{2s-4} p^2 = c^2 4^s$, as we have at most n^{2s-4} ways of picking the remaining vertices, at most 4^s ways of partitioning them into two edges, and probability p^2 that both edges appear. Now, we observe that each $\mathcal{E}_{u,v}$ is a down-event, so by Harris's inequality (or the FKG inequality), we have

$$\mathbb{P}\left(\bigwedge_{\{u,v\}\in\binom{V(G)}{2}} \mathcal{E}_{u,v}\right) \geq \prod_{u,v} (1-c^2 4^s) = (1-c^2 4^s)^{\binom{n}{2}} \geq e^{-O_s(c^2 n^2)}.$$

Recall that $\mathbb{P}(\mathcal{E}) \geq 1-e^{-\Omega_s(cn^2)}$. Therefore, for c sufficiently small in terms of s, and for n sufficiently large, we have that the sum of these two probabilities is larger than 1. Hence, with positive probability, H_0 is linear and contains an ordered monochromatic transversal edge for every dyadic partition, every colouring, and every ordering. □

Remark. The same proof shows that we can actually take H_0 to have (Berge) girth $> g$, for any fixed g. The linear case is simply this statement for $g=2$. Indeed, another application of Harris's inequality shows that the probability that H_0 does not contain a Berge cycle of length at most g is also at least $e^{-O_{s,g}(c^2 n^2)}$.

We are finally ready to prove Lemma 3.3.

Proof of Theorem 3.3. Let $s = |F^<|$. Let H_0 be the oriented s-graph given by Lemma 3.7. Obtain a k-graph H' by placing a copy of $F^<$ in each s-uniform edge of H_0, in an order-respecting way. That is, given an edge $e \in E(H_0)$, we insert $F^<$ into e by mapping the ith vertex in the ordering of $F^<$ to $\psi_e^{-1}(i)$, for all $1 \leq i \leq s$. Note that since $F^<$ and H_0 are both linear, so is H'. Additionally, by Lemma 3.7, we know that in any dyadic partition, colouring, and ordering of H_0, there is a monochromatic correctly-ordered transversal edge. This precisely yields the claimed monochromatic correctly-ordered transversal copy of $F^<$ in G. □

Remark. As above, if we assume that $F^<$ has girth $> g$, then we can ensure that G also has girth more than g, since any short Berge cycle in G must either stay entirely within one edge of H_0, or traverse at least g edges in H_0, and both of these are ruled out by the assumption that F and H_0 have girth $> g$.

4. Concluding remarks

One of the most fascinating special cases of a linear hypergraph is the Fano plane F_3; this is the seven-vertex 3-graph whose hyperedges are the lines in the projective geometry PG(2, 2) over $\mathbb{F}_2$. The growth rate of $r(F_3, K_n^{(3)})$ remains wide open: the best known lower bound is polynomial and the best known upper bound is exponential, and it would be very interesting to narrow the gap. In particular, F_3 is not iterated tripartite, so the conjecture of [5] predicts that $r(F_3, K_n^{(3)})$ grows super-polynomially; proving or disproving this is a natural step towards resolving the conjecture.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

There is a natural 4-uniform analogue of the Fano plane, namely the 13-vertex linear 4-graph F_4 whose hyperedges are the lines in the projective geometry PG(2, 3) over $\mathbb{F}_3$. F_4 is again not iterated 4-partite; curiously, however, it is quite straightforward to show that $r(F_4, K_n^{(4)})$ grows super-polynomially.

Proposition 4.1. *We have* $r(F_4, K_n^{(4)}) > 2^{n-2}$.

Proof. Let G_k be the 4-graph iteratively defined[3] as follows: G_1 is a single edge, and G_k is obtained from the disjoint union of two copies of G_{k-1} by adding in all edges intersecting each copy in exactly 2 vertices. Note that $\alpha(G_k) = \alpha(G_{k-1}) + 1$, since every independent set in G_k cannot contain two vertices from both copies of G_{k-1}. Together with the base case $\alpha(G_1) = 3$, we see that $\alpha(G_k) = k + 2$ for all k. Since G_k has 2^{k+1} vertices, the claimed result follows if we can show that G_k is F_4-free for all k.

If $F_4 \subseteq G_k$ for some k, then in particular there is a partition of $V(F_4)$ into two non-empty parts $A \cup B$ such that every edge is either fully contained in a part, or else intersects each part in exactly two vertices, and it can be checked by computer or minor casework that there is no such partition. For example, the following SageMath code does the job.

```
import sage.combinat.designs
F = designs.ProjectiveGeometryDesign(2,1,GF(3))
for A in Subsets(F.ground_set()):
    if all(len(A.intersection(Set(e))) in [0,2,4] for e in F.blocks()):
        print(A)
```

This code outputs only $A = \varnothing$ and $A = V(F_4)$, hence there is no such partition into non-empty parts. □

Note that Theorem 1.2 yields a linear 4-graph H with $r(H, K_n^{(4)}) \geq 2^{2^{(\log n)^C}}$, which is larger than the lower bound guaranteed by Proposition 4.1. However, the construction of H in Theorem 1.2 is probabilistic, whereas Proposition 4.1 gives a fully explicit linear 4-graph, namely F_4. We remark that by using this as our base case, one can similarly construct explicit linear k-graphs H with $r(H, K_n^{(k)}) \geq \mathrm{twr}_{k-2}(\Omega_k(n))$. Indeed, while the proof of Lemma 3.3 is probabilistic, it can be made explicit by using the Ramsey theorem for ordered linear hypergraphs of Nešetřil and Rödl [20]. One can use this powerful machinery to provide an alternative proof of Lemma 3.3; as the proof of the Nešetřil–Rödl theorem is fully deterministic, this would yield an explicit (but extraordinarily large and complicated) linear k-graph H with $r(H, K_n^{(k)}) \geq \mathrm{twr}_{k-2}(\Omega(n))$. To get the stronger result that H can have arbitrarily large girth, one could alternatively use the ordered version [22, Theorem 3.1] of the recent girth Ramsey theorem of Reiher and Rödl [21], whose proof is again explicit.

Finally, it would be very interesting to improve the lower bound in Theorem 1.1 to $r(H, K_n^{(3)}) \geq 2^{n^c}$ for some absolute constant $c > 0$, and some linear 3-graph H. By repeated applications of Proposition 1.3, this would immediately imply the existence of a linear k-graph H with $r(H, K_n^{(k)}) \geq \mathrm{twr}_{k-1}(n^c)$, matching the upper bound in (1.1) up to the constant.

Acknowledgements

We are grateful to the anonymous referees for several helpful comments which greatly improved the presentation of this paper.

[3] We remark that G_k can equivalently be defined as the stepping-up $(G_1, G_2, T_{2,2})$ where G_1 and G_2 are empty 3-graphs.

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press

References

[1] Ajtai, M., Komlós, J. and Szemerédi, E. (1980) A note on Ramsey numbers. *J. Combin. Theory Ser. A* **29** 354–360.
[2] Bohman, T. and Keevash, P. (2021) Dynamic concentration of the triangle-free process. *Random Struct. Algorithms* **58** 221–293.
[3] Campos, M., Jenssen, M., Michelen, M. and Sahasrabudhe, J. (2025) A new lower bound for the Ramsey numbers $R(3, k)$. Preprint available at arXiv: 2505.13371.
[4] Conlon, D., Fox, J., Gunby, B., et al. (2025) On off-diagonal hypergraph Ramsey numbers. *Int. Math. Res. Not. IMRN* (2025), Paper No. rnaf122, 22.
[5] Conlon, D., Fox, J., Gunby, B., et al. (2025) When are off-diagonal hypergraph Ramsey numbers polynomial? *Proc. Amer. Math. Soc* **153** 4605–4617.
[6] Conlon, D., Fox, J. and Sudakov, B. (2010) Hypergraph Ramsey numbers. *J. Amer. Math. Soc.* **23** 247–266.
[7] Conlon, D., Fox, J. and Sudakov, B. (2013) An improved bound for the stepping-up lemma. *Discrete Appl. Math.* **161** 1191–1196.
[8] Erdős, P. (1961) Graph Theory and Probability. II. *Canadian J. Math.* **13** 346–352.
[9] Erdős, P. and Hajnal, A. (1972) On Ramsey like theorems. Problems and results. In: *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972)*, Inst. Math. Appl., 123–140. Southend-on-Sea.
[10] Erdős, P., Hajnal, A. and Rado, R. (1965) Partition relations for cardinal numbers. *Acta Math. Acad. Sci. Hungar.* **16** 93–196.
[11] Erdős, P. and Rado, R. (1952) Combinatorial theorems on classifications of subsets of a given set. *Proc. London Math. Soc.* **s3-2** 417–439.
[12] Erdös, P. and Szekeres, G. (1935) A combinatorial problem in geometry. *Compositio Math.* **2** 463–470.
[13] Pontiveros, G. Fiz, Griffiths, S. and Morris, R. (2020) The triangle-free process and the Ramsey number $R(3, k)$. *Mem. Amer. Math. Soc.* **263** v+125.
[14] Fox, J. and He, X. (2021) Independent sets in hypergraphs with a forbidden link. *Proc. Lond. Math. Soc.* **123** 384–409.
[15] Hefty, Z., Horn, P., King, D. and Pfender, F. (2025) Improving $R(3, k)$ in just two bites. Preprint available at arXiv: 2510.19718.
[16] Kim, J. H. (1995) The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$. *Random Struct. Algorithms* **7** 173–207.
[17] Mattheus, S. and Verstraete, J. (2024) The asymptotics of $r(4, t)$. *Ann. Math.* **199** 919–941.
[18] Mubayi, D. and Suk, A. (2017) Off-diagonal hypergraph Ramsey numbers. *J. Combin. Theory Ser. B* **125** 168–177.
[19] Mubayi, D. and Suk, A. (2018) New lower bounds for hypergraph Ramsey numbers. *Bull. Lond. Math. Soc.* **50** 189–201.
[20] Nesetril, J. and Rödl, V. (1987) Strong Ramsey theorems for Steiner systems. *Trans. Amer. Math. Soc.* **303** 183–192.
[21] Reiher, C. and Rödl, V. (2023) The girth Ramsey theorem.Preprint available at arXiv: 2308.15589.
[22] Reiher, C., Rödl, V. and Schacht, M. (2025) Unavoidable subgraphs in Ramsey graphs. Preprint available at arXiv: 2502.09830.
[23] Shearer, J. B. (1983) A note on the independence number of triangle-free graphs. *Discrete Math.* **46** 83–87.

Cite this article: He X, Nie J, Wigderson Y, and Yu H (2026). Off-diagonal Ramsey numbers for linear hypergraphs. *Combinatorics, Probability and Computing* **35**, 554–567. https://doi.org/10.1017/S0963548326100443

https://doi.org/10.1017/S0963548326100443 Published online by Cambridge University Press