

Supersaturation for Ramsey-Turán problems

Dhruv Mubayi *

Vojtěch Rödl †

April 28, 2005

Abstract

For an l -graph \mathcal{G} , the Turán number $\text{ex}(n, \mathcal{G})$ is the maximum number of edges in an n -vertex l -graph \mathcal{H} containing no copy of \mathcal{G} . The limit $\pi(\mathcal{G}) = \lim_{n \rightarrow \infty} \text{ex}(n, \mathcal{G}) / \binom{n}{l}$ is known to exist [8]. The Ramsey-Turán density $\rho(\mathcal{G})$ is defined similarly to $\pi(\mathcal{G})$ except that we restrict to only those \mathcal{H} with independence number $o(n)$. A result of Erdős and Sós [3] states that $\pi(\mathcal{G}) = \rho(\mathcal{G})$ as long as for every edge E of \mathcal{G} there is another edge E' of \mathcal{G} for which $|E \cap E'| \geq 2$. Therefore a natural question is whether there exists \mathcal{G} for which $\rho(\mathcal{G}) < \pi(\mathcal{G})$.

Another variant $\tilde{\rho}(\mathcal{G})$ proposed in [3] requires the stronger condition that every set of vertices of \mathcal{H} of size at least ϵn ($0 < \epsilon < 1$) has density bounded below by some threshold. By definition, $\tilde{\rho}(\mathcal{G}) \leq \rho(\mathcal{G}) \leq \pi(\mathcal{G})$ for every \mathcal{G} . However, even $\tilde{\rho}(\mathcal{G}) < \pi(\mathcal{G})$ is not known for very many l -graphs \mathcal{G} when $l > 2$.

We prove the existence of a phenomenon similar to supersaturation for Turán problems for hypergraphs. As a consequence, we construct, for each $l \geq 3$, infinitely many l -graphs \mathcal{G} for which $0 < \tilde{\rho}(\mathcal{G}) < \pi(\mathcal{G})$.

We also prove that the 3-graph \mathcal{G} with triples 12a, 12b, 12c, 13a, 13b, 13c, 23a, 23b, 23c, abc, satisfies $0 < \rho(\mathcal{G}) < \pi(\mathcal{G})$. The existence of a hypergraph \mathcal{H} satisfying $0 < \rho(\mathcal{H}) < \pi(\mathcal{H})$ was conjectured by Erdős and Sós [3], proved by Frankl and Rödl [6], and later by Sidorenko [14].

Our short proof is based on different ideas and is simpler than these earlier proofs.

*Department of Mathematics, Statistics, and Computer Science, University of Illinois, 851 S. Morgan Street, Chicago, IL 60607-7045; research supported in part by the National Science Foundation under grants DMS-9970325 and DMS-0400812, and an Alfred P. Sloan Research Fellowship

†Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA; research supported in part by the National Science Foundation under grants DMS-0071261 and DMS-0300529

2000 Mathematics Subject Classification: 05C35, 05C65, 05D05

Keywords: *Ramsey Turán problems, extremal hypergraph theory, supersaturation*

1 Introduction

1.1 Background

An l -graph is a hypergraph $\mathcal{G} = (V(\mathcal{G}), E(\mathcal{G}))$, whose edge set $E(\mathcal{G})$ consists of l -element subsets of the vertex set $V(\mathcal{G})$. We write $n(\mathcal{G})$ and $e(\mathcal{G})$ for the number of vertices and edges, respectively, in a hypergraph \mathcal{G} . The Turán problem for l -graphs asks for the maximum number $t(n, l, r)$ of edges in an n -vertex l -graph that contains no copy of the complete l -graph on r vertices. For $l = 2$ Turán's theorem [18] provides an answer, while for each $r > l > 2$ this problem is unsolved, and is one of the most fundamental problems in extremal hypergraph theory. More generally, for an l -graph \mathcal{G} , the Turán number $\text{ex}(n, \mathcal{G})$ is the maximum number of edges in an n -vertex l -graph \mathcal{H} containing no copy of \mathcal{G} . The limit $\pi(\mathcal{G}) = \lim_{n \rightarrow \infty} \text{ex}(n, \mathcal{G}) / \binom{n}{l}$ is known to exist, is determined for every graph by the Erdős-Simonovits-Stone theorem [2, 4], but very difficult to compute even for specific l -graphs when $l > 2$.

One of the most important properties of these extremal problems is the supersaturation phenomenon [15]. This states that if an n -vertex l -graph has at least $(\pi(\mathcal{G}) + \epsilon) \binom{n}{l}$ edges, then it contains not just one, but at least $\delta \binom{n}{n(\mathcal{G})}$ copies of \mathcal{G} as a subgraph, where $\delta > 0$ depends only on ϵ . A consequence of this is that $\pi(\mathcal{G})$ remains unchanged if we blow up vertices of \mathcal{G} , i.e., replace a vertex v by several vertices v_1, \dots, v_t , such that $v \cup S \in E(\mathcal{G})$ if and only if $v_i \cup S \in E(\mathcal{G})$ for all $i = 1, \dots, t$. For graphs, the supersaturation phenomenon is essentially the content of the Erdős-Simonovits-Stone theorem.

Our motivation is to prove the supersaturation phenomenon in the context of Ramsey-Turán problems for hypergraphs. Throughout this paper we omit ceiling and floor symbols unless crucial.

Many of the conjectured extremal examples for hypergraph Turán problems have large independent sets. Motivated by this observation, Erdős and Sós imposed a restriction on the underlying l -graphs in this problem, namely that they should not have large independent sets. The resulting set of problems and results comprise Ramsey-Turán theory. More precisely, for $0 < \delta \leq 1$,

$$\text{ex}(n, \mathcal{G}, \delta) = \max\{e(\mathcal{F}) : \mathcal{F} \text{ is an } n\text{-vertex } l\text{-graph with } \mathcal{G} \not\subseteq \mathcal{F} \text{ and } \alpha(\mathcal{F}) < \delta n\},$$

or zero if no such hypergraph exists. The Ramsey-Turán number $\rho(\mathcal{G})$ is defined as

$$\sup_{\delta(n)} \left\{ \limsup_{n \rightarrow \infty} \frac{\text{ex}(n, \mathcal{G}, \delta(n))}{\binom{n}{l}} : \delta(n) \rightarrow 0 \text{ as } n \rightarrow \infty \right\}.$$

Obviously $\rho(\mathcal{G}) \leq \pi(\mathcal{G})$ for every \mathcal{G} , and there exist graphs G for which $\rho(G) < \pi(G)$ (the simplest example is $G = K_3$ which satisfies $\rho(K_3) = 0 < 1/2 = \pi(K_3)$).

Moving on to hypergraphs, a basic question is whether there exists an l -graph ($l > 2$) \mathcal{G} with $\rho(\mathcal{G}) < \pi(\mathcal{G})$. Erdős and Sós [3] proved that such \mathcal{G} do exist but with $\rho(\mathcal{G}) = 0$. They asked whether there exist l -graphs ($l > 2$) with

$$0 < \rho(\mathcal{G}) < \pi(\mathcal{G}). \quad (1)$$

This was answered positively by Frankl and Rödl [6] for every $l > 2$, who showed that there exist infinitely many l -graphs for which (1) holds; however, they were not able to obtain a single explicit \mathcal{G} . Subsequently, Sidorenko [14], using ideas from [6], proved that for the 3-graph \mathcal{G} whose edge set is $\{123, 145, 167, 245, 267, 345, 367, 467, 567\}$, inequality (1) holds.

We give another example for which (1) holds that has only six vertices. Moreover, our proof is much shorter than that in [14] and is entirely self-contained.

Definition 1.1. Let \mathcal{F}_1 be the 3-graph of one edge. For $r > 1$, \mathcal{F}_r is the 3-graph obtained from a copy Y of \mathcal{F}_{r-1} by adding a set X of $r + 1$ new vertices and all edges in $\{xx'y : x, x' \in X, y \in Y\}$.

Theorem 1.2. (Section 2)

$$2/7 \leq \rho(\mathcal{F}_2) \leq 2/3 < 3/4 \leq \pi(\mathcal{F}_2).$$

Another variant of the Ramsey-Turán problem, due to Erdős and Sós [3], is to require \mathcal{F} to have uniformly positive density. In fact, there is a hierarchy of uniformity conditions on hypergraphs ([3], [16] section 5), two of which we present below. We write $\mathcal{H}_n^{(l)}$ to denote an n -vertex l -graph, and write $a = b \pm c$ for $a \in [b - c, b + c]$.

It is more convenient to consider sequences $\{\mathcal{H}_i\}$ of l -graphs instead of individual l -graphs. Without loss of generality, we assume that $n(\mathcal{H}_i) \rightarrow \infty$ and $n(\mathcal{H}_{i+1}) \geq n(\mathcal{H}_i)$. A hypergraph \mathcal{H} contains \mathcal{G} if \mathcal{G} is a (not necessarily induced) subhypergraph of \mathcal{H} ; \mathcal{H} is \mathcal{G} -free if \mathcal{H} contains no copy of \mathcal{G} . The sequence $\{\mathcal{H}_i\}$ is \mathcal{G} -free if infinitely many \mathcal{H}_i are \mathcal{G} -free. The sequence $\{\mathcal{H}_i\}$ contains \mathcal{G} if all but finitely many of the \mathcal{H}_i 's contain \mathcal{G} .

An l -graph \mathcal{H} is (α, ξ) -uniform if every ξn vertices of \mathcal{H} span $(\alpha \pm \xi) \binom{\xi n}{l}$ edges; it is (α, ξ) -dense if every ξn vertices of \mathcal{H} span at least $\alpha \binom{\xi n}{l}$ edges. Note that the definitions of (α, ξ) -uniform and (α, ξ) -dense automatically imply their respective conclusions to sets of size at least ξn . To see this in the case of (α, ξ) -dense, suppose that $S \subset V(\mathcal{H})$ with $s = |S| \geq \xi n$. Then by averaging over all sets of size ξn in S ,

$$e(S) \geq \frac{\binom{s}{\xi n} \alpha \binom{\xi n}{l}}{\binom{s-l}{\xi n-l}} = \alpha \binom{s}{l}. \quad (2)$$

DENSE: $\{\mathcal{H}_i\}$ is α -dense if for all $\xi > 0$, there exists n_0 such that for all $n > n_0$, the hypergraph \mathcal{H}_n is (α, ξ) -dense.

UNIFORM: $\{\mathcal{H}_i\}$ is α -uniform if for all $\xi > 0$, there exists n_0 such that for all $n > n_0$, the hypergraph \mathcal{H}_n is (α, ξ) -uniform.

Remark 1.3. We use the term α -uniform when we speak about sequences of hypergraphs. This should not be confused with the common usage of l -uniform which denotes hypergraphs whose edges have size l .

Now we define Ramsey-Turán functions based on these definitions. Let

$$\begin{aligned} \rho_D(\mathcal{G}) &= \sup\{\alpha : \text{there exists an } \alpha\text{-dense sequence } \{\mathcal{H}_i\} \text{ that is } \mathcal{G}\text{-free} \} \\ &= \inf\{\alpha : \text{every } \alpha\text{-dense sequence } \{\mathcal{H}_i\} \text{ contains } \mathcal{G} \}, \end{aligned}$$

$$\begin{aligned} \rho_U(\mathcal{G}) &= \sup\{\alpha : \text{there exists an } \alpha\text{-uniform sequence } \{\mathcal{H}_i\} \text{ that is } \mathcal{G}\text{-free} \} \\ &= \inf\{\alpha : \text{every } \alpha\text{-uniform sequence } \{\mathcal{H}_i\} \text{ contains } \mathcal{G} \}. \end{aligned}$$

It is sometimes more convenient to work with l -graphs that are (α, ξ) -uniform than (α, ξ) -dense. Therefore, we consider both parameters ρ_D and ρ_U . Indeed, the following theorem also proves that these two concepts are the same for Turán problems. In fact, as a referee pointed out, it also shows this for every property that is closed under taking subgraphs. We say that $\{\mathcal{H}_i\}$ contains the subsequence $\{\mathcal{H}'_{j_i}\}$ if for every i , \mathcal{H}_{j_i} contains \mathcal{H}'_{j_i} .

Theorem 1.4. (Section 4) Fix $0 < \alpha < 1$. Every α -dense sequence $\{\mathcal{H}_i\}$ contains an α -uniform subsequence $\{\mathcal{H}'_{j_i}\}$.

From Theorem 1.4 together with the fact that for every $\epsilon > 0$, every $(\alpha + \epsilon)$ -uniform sequence is also α -dense, we can easily see that $\rho_D(\mathcal{G}) = \rho_U(\mathcal{G})$ for every \mathcal{G} . Consequently we can define

$$\tilde{\rho}(\mathcal{G}) = \rho_D(\mathcal{G}) = \rho_U(\mathcal{G}).$$

It is straightforward to verify that $\tilde{\rho}(G) = 0$ for every graph G , hence our focus is on hypergraphs. Using the reduction in Theorem 1.4 we describe infinitely many l -graphs ($l \geq 3$) with $0 < \tilde{\rho}(\mathcal{G}) < \pi(\mathcal{G})$ (in fact, the following result implies an even stronger statement).

Theorem 1.5. (Section 5) Let \mathcal{G} be an l -graph with $l \geq 3$. Then for any $\epsilon > 0$, there exists $\hat{\mathcal{G}}$ containing \mathcal{G} with $\rho(\hat{\mathcal{G}}) > 1 - \epsilon$ and $\tilde{\rho}(\hat{\mathcal{G}}) = \tilde{\rho}(\mathcal{G})$.

Let \mathcal{G} be such that $\tilde{\rho}(\mathcal{G}) > 0$. Choosing $\epsilon > 0$ so that $1 - \epsilon > \tilde{\rho}(\hat{\mathcal{G}})$, Theorem 1.5 yields the following.

Corollary 1.6. *If $\tilde{\rho}(\mathcal{G}) > 0$, then there exists $\hat{\mathcal{G}} \supset \mathcal{G}$ for which*

$$0 < \tilde{\rho}(\hat{\mathcal{G}}) < \rho(\hat{\mathcal{G}}) \leq \pi(\hat{\mathcal{G}}). \quad (3)$$

Remark 1.7. *Two simple examples of 3-graphs with positive $\tilde{\rho}$ are $H(4, 3)$, the 3-graph with four points and three triples, and K_4^3 , the complete 3-graph on four points. Füredi and the second author independently proved that $\tilde{\rho}(H(4, 3)) \geq 1/4$, and $\tilde{\rho}(K_4^3) \geq 1/2$.*

The hypergraphs in Definition 1.1 are another infinite family of 3-graphs with $0 < \tilde{\rho}(\mathcal{F}_r) < \pi(\mathcal{F}_r)$. These are the only known examples for which such a result has been proved using purely elementary arguments.

Theorem 1.8. *(Section 2) For $r > 2$,*

$$1 - \frac{3}{r} + \frac{2}{r^2} \leq \tilde{\rho}(\mathcal{F}_r) \leq 1 - \frac{2}{r(r+1)} < 1 - \frac{1}{r^2} \leq \pi(\mathcal{F}_r).$$

1.2 The main tool

In this subsection we describe the main tool that allows us to prove supersaturation for Ramsey-Turán problems. Informally, Theorem 1.9 says that in an l -graph whose edges are uniformly distributed, this property is inherited to almost all sets of constant size. Recall that $\mathcal{H}_n^{(l)}$ denotes an n -vertex l -graph.

Theorem 1.9. *(Mubayi-Rödl [11]) Let $0 < \alpha < 1, l \geq 2$ be fixed. For all $\tilde{\delta}$, there exists r_0 such that for all $r > r_0$, there exist δ, n_0 such that, if $n > n_0$ and $\mathcal{H}_n^{(l)}$ is (α, δ) -uniform, then all but $\exp\{-r^{1/l}/20\} \binom{n}{r}$ r -sets of vertices induce a subhypergraph that is $(\alpha, \tilde{\delta})$ -uniform.*

Remark 1.10. *Note that $r_0 = r_0(\tilde{\delta}) \rightarrow \infty$ as $\tilde{\delta} \rightarrow 0$.*

Remark 1.11. *Note also that if $\delta' < \delta$, and Theorem 1.9 holds for δ , then it also holds for δ' , since an (α, δ') -uniform hypergraph is also (α, δ) -uniform. We may therefore take δ to be as small as we need in relation to r , in particular, we assume that $\delta < 1/2r^{r+1}$. This will be needed in Section 5.*

Remark 1.12. *A closely related result to Theorem 1.9 was proved independently and simultaneously by Alon, de la Vega, Kannan, and Karpinski [1] (Theorem 9 in the extended version of [1]) using completely different techniques. One technical difficulty in applying that result to our situation is that it was proved only for l -partite l -graphs. Another difference is that the conclusion of Theorem 1.9 applies to all but an exponentially small (in r) proportion of r -sets, while the result of Alon et al. only gives the conclusion for 39/40 fraction of the r -sets. With the added steps of proving the result from [1] in a non l -partite framework, it could just as well be used for the main applications in this paper.*

2 Elementary proofs for $0 < \rho < \pi$ and $0 < \tilde{\rho} < \pi$

In this section we prove Theorems 1.2 and 1.8. We need the following well-known and easy fact.

Lemma 2.1. *Fix $l \geq 2$. If \mathcal{G} is an l -graph with $\chi(\mathcal{G}) \geq k$, then $\pi(\mathcal{G}) \geq 1 - (k - 1)^{1-l}$.*

Proof. Let \mathcal{H}_k be the n -vertex l -graph with vertex partition $V_1 \cup \dots \cup V_{k-1}$, where $\|V_i\| - \|V_j\| \leq 1$ for all $i \neq j$, and all edges except those that lie entirely within any one of the V_i 's. The density of \mathcal{H}_k is $1 - (k - 1)^{1-l} - o(1)$ (as $n \rightarrow \infty$) and $\chi(\mathcal{H}_k) \leq k - 1$. Since $\chi(\mathcal{G}) \geq k$, there is no copy of \mathcal{G} in \mathcal{H}_k . Therefore $\pi(\mathcal{G}) \geq 1 - (k - 1)^{1-l}$. \square

Lemma 2.2. $\chi(\mathcal{F}_r) \geq r + 1$ for every $r \geq 1$.

Proof. We proceed by induction on r , noting that $\chi(\mathcal{F}_1) \geq 2$. Suppose, for contradiction, that $r > 1$ and $\chi(\mathcal{F}_r) \leq r$. Recall that a copy of \mathcal{F}_r has vertex partition $X \cup Y$, where $|X| = r + 1$ and Y is a copy of \mathcal{F}_{r-1} . A proper r -coloring of \mathcal{F}_r assigns the same color, say 1, to at least two vertices in X . Consequently, color 1 is forbidden on Y , and therefore Y is properly $(r - 1)$ -colored. However, this contradicts the induction hypothesis $\chi(\mathcal{F}_{r-1}) \geq r$. \square

For vertices u, v in a hypergraph, we write $d_{u,v}$ for the number of edges containing both u and v . Recall that $H(4, 3)$ is the unique four vertex triple system with three triples.

Proof of Theorem 1.2: Recall that we must show $2/7 \leq \rho(\mathcal{F}_2) \leq 2/3 < 3/4 \leq \pi(\mathcal{F}_2)$.

Set $\mathcal{G} = \mathcal{F}_2$. Then \mathcal{G} contains $H(4, 3)$, and Frankl and Füredi [5] proved that $\rho(H(4, 3)) \geq 2/7$. Consequently, $\rho(\mathcal{G}) \geq \rho(H(4, 3)) \geq 2/7$.

Lemma 2.2 implies that $\chi(\mathcal{G}) \geq 3$, so by Lemma 2.1 we obtain $\pi(\mathcal{G}) \geq 1 - 2^{1-3} = 3/4$ (it is proved in [10] that $\pi(\mathcal{G}) = 3/4$).

We now prove that $\rho(\mathcal{G}) \leq 2/3$. Fix $\epsilon > 0$, and suppose that \mathcal{H} is an n -vertex 3-graph with $(2/3 + \epsilon)\binom{n}{3}$ edges and let n be large enough so that $\alpha(\mathcal{H}) < \epsilon n$. We will prove that \mathcal{H} contains a copy of \mathcal{G} thus completing the proof. Since

$$\sum_{\substack{S \subset V(\mathcal{H}), \\ |S|=3}} \sum_{u,v \in S} d_{u,v} = \sum_{u,v} d_{u,v}(n-2) \geq (2 + 3\epsilon)(n-2) \binom{n}{3},$$

there is an $S_0 \subset V(\mathcal{H})$ of size three with

$$\sum_{u,v \in S_0} d_{u,v} \geq (2 + 3\epsilon)(n-2).$$

If fewer than ϵn vertices from $V(\mathcal{H}) - S_0$ form edges with all pairs from S_0 , then

$$\sum_{u,v \in S_0} d_{u,v} < 3\epsilon n + 2(1 - \epsilon)n = \epsilon n + 2n < (2 + 3\epsilon)(n - 2) \leq \sum_{u,v \in S_0} d_{u,v}.$$

Consequently, at least ϵn vertices from $V(\mathcal{H}) - S_0$ form edges with all three pairs from S_0 . By hypothesis, there is an edge contained in these ϵn vertices, and together with S_0 , this forms a copy of \mathcal{G} . \square

The lower bound for $\tilde{\rho}(\mathcal{F}_r)$ in Theorem 1.8 follows from the following construction.

Construction 2.3. Fix $r \geq 2$. Color the edges of a complete graph on n vertices with r -colors, where each color appears on each edge independently with probability $1/r$. For each set S of three vertices, include S as an edge in a hypergraph if the colors on the three edges in S are distinct.

We first show that any resulting 3-graph $\mathcal{H}_{n,r}$ from Construction 2.3 contains no copy \mathcal{G} of \mathcal{F}_r . Suppose to the contrary that such a \mathcal{G} exists. Recall from Definition 1.1 that \mathcal{G} consists of a part Y that is a copy of \mathcal{F}_{r-1} , and a part X of size $r + 1$ with all triples $xx'y$ where $x, x' \in X, y \in Y$. Pick $y \in Y$. Then some two edges yx, yx' must have the same color, consequently $xx'y$ cannot be an edge in $\mathcal{H}_{n,r}$ and therefore in \mathcal{G} .

Next we show that for any $\epsilon > 0$, there is an n_0 such that for $n > n_0$ there exists an $\mathcal{H}_{n,r}$ that is (α, ϵ) -uniform, where $\alpha = (1 - 1/r)(1 - 2/r)$. Note that α is the probability that a set S of size three was selected as an edge.

The event that a specific edge is present in $\mathcal{H}_{n,r}$ is independent of the same event for some other edge. Hence we can apply the Chernoff inequality (see, e.g., [7] page 27) to conclude that the probability that there is a set of vertices of size ϵn with density outside $\alpha \pm \epsilon$ is at most

$$\binom{n}{\epsilon n} \exp \left\{ -c_\epsilon \alpha \binom{\epsilon n}{3} \right\} < 1,$$

since n is sufficiently large. Consequently, there is an $\mathcal{H} = \mathcal{H}_{n,r}$ that is (α, ϵ) -uniform, and \mathcal{H} contains no copy of \mathcal{F}_r . We conclude that

$$\tilde{\rho}(\mathcal{F}_r) \geq \alpha = 1 - 3/r + 2/r^2. \tag{4}$$

Proof of Theorem 1.8: We must prove that

$$1 - 3/r + 2/r^2 \leq \tilde{\rho}(\mathcal{F}_r) \leq 1 - \frac{1}{\binom{r+1}{2}} < 1 - \frac{1}{r^2} \leq \pi(\mathcal{F}_r).$$

The first inequality has been shown in (4). For the last inequality, we use Lemma 2.2 and then Lemma 2.1.

We now focus on proving that $\tilde{\rho}(\mathcal{F}_r) = \rho_D(\mathcal{F}_r) \leq 1 - 1/\binom{r+1}{2}$. Our approach is to use induction on r , and extends the argument in the proof of Theorem 1.2. Fix $\epsilon > 0$, and let $\delta_r = \epsilon^r$. Suppose that \mathcal{H} is an n -vertex 3-graph (n sufficiently large) that is (α, δ_r) -dense, where $\alpha = (1 - 1/\binom{r+1}{2}) + \epsilon$. We will show that \mathcal{H} contains a copy of \mathcal{F}_r . Consequently, every α -dense $\{\mathcal{H}_i\}$ contains \mathcal{F}_r . This proves that $\tilde{\rho}(\mathcal{F}_r) = \rho_D(\mathcal{F}_r) \leq \alpha = 1 - 1/\binom{r+1}{2} + \epsilon$, and since $\epsilon > 0$ is arbitrary we obtain $\tilde{\rho}(\mathcal{F}_r) \leq 1 - 1/\binom{r+1}{2}$.

The result is trivial for $r = 1$, so assume that $r > 1$. Writing V for $V(\mathcal{H})$, and counting codegrees in two ways, we have

$$\sum_{\substack{S \subset V, \\ |S|=r+1}} \sum_{u,v \in S} d_{u,v} = \sum_{x,y \in V} d_{x,y} \binom{n-2}{r-1} \geq \binom{n-2}{r-1} 3 \left(1 - \frac{1}{\binom{r+1}{2}} + \epsilon\right) \binom{n}{3}.$$

Consequently, there is an $(r+1)$ -set $S_0 \subset V$ for which

$$\sum_{u,v \in S_0} d_{u,v} \geq \frac{3 \binom{n-2}{r-1} \left(1 - \frac{1}{\binom{r+1}{2}} + \epsilon\right) \binom{n}{3}}{\binom{n}{r+1}} = \left[(1 + \epsilon) \binom{r+1}{2} - 1 \right] (n - r + 1). \quad (5)$$

Let $T \subset V - S_0$ be the set of vertices which do not form triples with all pairs of vertices from S_0 . Set $|T| = tn$ with $0 \leq t \leq 1$. Then

$$\begin{aligned} \sum_{u,v \in S_0} d_{u,v} &\leq tn \left[\binom{r+1}{2} - 1 \right] + (n - tn - (r+1)) \binom{r+1}{2} + 3 \binom{r+1}{3} \\ &= \left[t \left(\binom{r+1}{2} - 1 \right) + (1-t) \binom{r+1}{2} \right] n - (r+1) \binom{r+1}{2} + 3 \binom{r+1}{3} \\ &= \left[\binom{r+1}{2} - t \right] n - (r+1) \binom{r+1}{2} + 3 \binom{r+1}{3} \\ &\leq \left[\binom{r+1}{2} - t \right] n. \end{aligned}$$

Together with (5), this implies that $t \leq 1 - (\epsilon/2) \binom{r+1}{2}$. Therefore, every vertex in a set B of at least $(\epsilon/2) \binom{r+1}{2} n - r \geq \epsilon n$ vertices in $V - S_0$ forms edges with each pair from S_0 . Since \mathcal{H} is (α, δ_r) -dense, every set $B' \subset B$ of size at least $\epsilon^{r-1} |B| \geq \delta_r n$ has density at least α . Consequently, the hypergraph induced by B is (α, δ_{r-1}) -dense. We now apply induction to obtain a copy of \mathcal{F}_{r-1} in B . Together with S_0 this forms a copy of \mathcal{F}_r in \mathcal{H} as required. \square

3 The Regularity Lemma

In this section we describe our main tool needed to prove Theorem 1.4, the Szemerédi Regularity Lemma (see [17] and [9]). In an l -graph \mathcal{H} , let X_1, \dots, X_l be pairwise disjoint sets of vertices. Write

$e(X_1, \dots, X_l)$ for the number of edges with exactly one point in each X_i . The *density* of the l -tuple (X_1, \dots, X_l) is

$$d(X_1, \dots, X_l) = \frac{e(X_1, \dots, X_l)}{|X_1| \dots |X_l|}.$$

Given $\epsilon > 0$, the l -tuple (X_1, \dots, X_l) is ϵ -regular with density $\alpha = d(X_1, \dots, X_l)$ if for every choice of $Y_i \subseteq X_i$ with

$$|Y_i| \geq \epsilon |X_i| \text{ for each } i,$$

we have

$$d(Y_1, \dots, Y_l) = \alpha \pm \epsilon.$$

We also extend the definition of (α, ϵ) -uniform to an l -partite situation.

Definition 3.1. *The l -partite l -graph $\mathcal{H} = \mathcal{H}(V_1, \dots, V_l)$ is (α, ϵ) -uniform if for every choice of $V_i' \subset V_i$, with $|V_i'| = \epsilon |V_i|$, the density $d(V_1', \dots, V_l') = \alpha \pm \epsilon$.*

Remark 3.2. *For an l -tuple $L = (X_1, \dots, X_l)$ there is a subtle distinction between the following two statements*

- 1) L is ϵ -regular with density α
- 2) L is (α, ϵ) -uniform.

It is clear that 1) implies 2), but if 2) holds, then the best we can say regarding 1) is that L is 2ϵ -regular with density $\alpha' = \alpha \pm \epsilon$.

Consider a partition $\mathcal{P} = V_0 \cup V_1 \cup \dots \cup V_k$ of $V(\mathcal{H})$. We say that \mathcal{P} is an ϵ -regular partition if

- 1) $|V_0| < \epsilon |V|$,
- 2) $|V_1| = \dots = |V_k|$,
- 3) the l -tuple $(V_{i_1}, \dots, V_{i_l})$ is ϵ -regular for all but ϵk^l choices of l of the V_i 's.

With these notions we can state the celebrated Szemerédi Regularity Lemma. Below we state a version for l -graphs, $l \geq 2$. Its proof is essentially the same as for the case $l = 2$ (see [12] for an exposition).

Lemma 3.3. *(Regularity Lemma) Let $l \geq 2$ be fixed. For every $\epsilon > 0$ and every integer $k_0 \geq 1$ there exist integers K, n_0 such that every l -graph of order at least n_0 admits an ϵ -regular partition V_0, \dots, V_k with $k_0 \leq k \leq K$.*

4 Dense equals Uniform

In this section we prove Theorem 1.4. We think this result is of independent interest; however, its main application is to prove Theorem 1.5 by putting it in the context of Theorem 1.9. Our proof uses the Szemerédi Regularity Lemma described in Section 3.

Lemma 4.1. *For every $l \geq 2, \epsilon > 0, \beta \geq \alpha \geq 0$, there is an m_0 such that for $m > m_0$ the following holds: Let $\mathcal{H} = \mathcal{H}(V_1, \dots, V_l)$ be an ϵ -regular l -partite l -graph with density $\beta > 0$ and $|V_i| = m$ for all i . Then there is an l -partite $\mathcal{H}' \subset \mathcal{H}$ which is $(\alpha, 2\epsilon)$ -uniform.*

Proof. Form \mathcal{H}' by choosing each edge of \mathcal{H} independently with probability $p = \alpha/\beta$. Then the probability that there exist $V'_i \subset V_i$, $|V'_i| = 2\epsilon m$ with $d_{\mathcal{H}'}(V'_1, \dots, V'_l) \neq \alpha \pm 2\epsilon$ is by Chernoff's inequality at most

$$\binom{m}{2\epsilon m}^l \exp\{-c_\epsilon \alpha (2\epsilon m)^l\} < 1,$$

since $m > m_0$. Consequently there is an \mathcal{H}' for which $d_{\mathcal{H}'}(V'_1, \dots, V'_l) = \alpha \pm 2\epsilon$ for all choices of $V'_i \subset V_i$, $|V'_i| = 2\epsilon m$. \square

Proof of Theorem 1.4: Let $\{\mathcal{H}_i\}$ be an α -dense sequence of l -graphs. We will show that $\{\mathcal{H}_i\}$ contains an α -uniform subsequence $\{\mathcal{G}_t\}$ with $n(\mathcal{G}_t) \rightarrow \infty$. Since our definition of hypergraph sequence also requires the order of the individual hypergraphs to be nondecreasing, and $\{\mathcal{G}_t\}$ may not have this property, we can always take a subsequence $\{\mathcal{G}'_t\}$ of \mathcal{G}_t having this property (since $n(\mathcal{G}_t) \rightarrow \infty$). Our strategy is to show that

for every $\mu > 0$, there is an i_μ and an $\mathcal{H}(\mu) \subset \mathcal{H}_{i_\mu}$ which is $(\alpha, 4\mu)$ -uniform.

Our proof will also show that

$$n(\mathcal{H}(\mu)) \rightarrow \infty \text{ as } \mu \rightarrow 0. \tag{6}$$

This is clearly enough, since an easy diagonalization argument yields an α -uniform sequence. For example, $\{\mathcal{G}_t\} = \{\mathcal{H}(1/t)\}$ is an α -uniform sequence.

Fix $\mu > 0$ and

$$p = \max \left\{ \left\lceil \frac{2l^2}{\mu^{l+1}\alpha} \right\rceil, \left\lceil \frac{l^2}{\mu} \right\rceil \right\}. \tag{7}$$

Let $R^{(l)}(p)$ be the minimum number t such that every two-coloring of the edges of a complete l -graph on t vertices yields a monochromatic complete l -graph on p vertices. Ramsey's Theorem [13] states that $R^{(l)}(p)$ is finite for every fixed l, p . Next choose s so that

$$R^{(l)}(p) \leq s. \tag{8}$$

Now choose ϵ so small that

$$\epsilon < \frac{\alpha\mu^{l+1}}{10l} \quad (9)$$

and

$$s \leq \left\lfloor \frac{1}{4(l\epsilon)^{1/(l-1)}} \right\rfloor. \quad (10)$$

Also choose $k_0 > 1/[(l\epsilon)^{1/(l-1)}]$.

Given inputs ϵ, k_0 to the Regularity Lemma, assume that it outputs K, n_0 . Having α and l fixed, observe that p, s, ϵ, k_0 and consequently also K and n_0 can be viewed as functions of μ . Hence, for $\xi = \xi(\mu) = p/K$, there exists i_μ such that for all $i > i_\mu$, the hypergraph \mathcal{H}_i is (α, ξ) -dense. Take \mathcal{H}_i so that $i > i_\mu$ and \mathcal{H}_i has more than n_0 vertices. Because $n(\mathcal{H}_i) \rightarrow \infty$, this is possible. From now on we denote \mathcal{H}_i simply by \mathcal{H} .

Apply the Regularity Lemma to \mathcal{H} with inputs ϵ, k_0 . We obtain an ϵ -regular partition $V_0 \cup V_1 \cup \dots \cup V_k$ ($K > k > k_0$) of \mathcal{H} with all but ϵk^l of the l -tuples $(V_{j_1}, \dots, V_{j_l})$ being ϵ -regular. Then $V_1 \cup \dots \cup V_k$ is also an ϵ -regular partition with all but ϵk^l of the l -tuples $(V_{j_1}, \dots, V_{j_l})$ being ϵ -regular.

Turán's theorem states that a k -vertex l -graph with average degree $d^{l-1} > 1$ has an independent set of size at least $(1 - 1/l)k/d$ (the easy proof is to pick a random subset of vertices by choosing each vertex independently with probability $1/d$, compute the expected number of edges within this set, and delete a vertex from each of these edges). We apply this to the l -graph with vertex set $[k]$ and edge set consisting of the l -tuples (j_1, \dots, j_l) , where $(V_{j_1}, \dots, V_{j_l})$ is not ϵ -regular. For this l -graph, the average degree $d^{l-1} \leq l\epsilon k^l/k = l\epsilon k^{l-1}$, and we obtain the following lower bound for the independence number (for this computation to work, the probability $1/d$ must be less than one, and this holds since $k > k_0 > 1/[(l\epsilon)^{1/(l-1)}]$):

$$\left(1 - \frac{1}{l}\right) \frac{k}{d} = \frac{(1 - 1/l)k}{(l\epsilon)^{1/(l-1)}k} > \left\lfloor \frac{1}{4(l\epsilon)^{1/(l-1)}} \right\rfloor = s.$$

By relabelling if needed, we may therefore assume that all the l -tuples $(V_{j_1}, \dots, V_{j_l})$ are ϵ -regular for $1 \leq j_1 < \dots < j_l \leq s$. We now restrict only to those V_h for $1 \leq h \leq s$.

Divide these l -tuples $(V_{j_1}, \dots, V_{j_l})$, $1 \leq j_i \leq s$ into two groups,

- i) RED – those with density less than $\alpha - \mu$
- ii) BLUE – those with density at least $\alpha - \mu$.

Color these l -tuples by two colors depending on which of groups i) or ii) each one lies in. Now apply Ramsey's theorem to the l -graph whose edges are the l -tuples we have just colored. We obtain a monochromatic complete l -graph with p vertices by (8). Consequently, we may assume by relabelling, that all l -tuples $(V_{j_1}, \dots, V_{j_l})$, where $\{j_1, \dots, j_l\} \subseteq [p]$ have the same color.

Claim: Every l -tuple $(V_{j_1}, \dots, V_{j_l})$, where $\{j_1, \dots, j_l\} \subseteq [p]$, is blue.

Proof of Claim. Suppose for contradiction, that every l -tuple, $(V_{j_1}, \dots, V_{j_l})$, where $\{j_1, \dots, j_l\} \subseteq [p]$, is red. Let $\mathcal{H}[V_1, \dots, V_p]$ be the l -graph induced by $V_1 \cup \dots \cup V_p$. Set $m = n/p$ where $n = n(\mathcal{H}[V_1, \dots, V_p])$ (thus $m = |V_j|$ for each $1 \leq j \leq k$).

Since all l -tuples are colored red, the total number of edges in $\mathcal{H}[V_1, \dots, V_p]$ is at most

$$\begin{aligned} \binom{p}{l} m^l (\alpha - \mu) + p \binom{m}{2} \binom{pm}{l-2} &< \binom{pm}{l} (\alpha - \mu) + \frac{pm^2}{2} \frac{(pm)_{l-2}}{(l-2)!} \\ &< \binom{pm}{l} (\alpha - \mu) + \frac{pm^2 l^2}{2} \frac{(pm)_{l-2}}{l!} \\ &< \binom{pm}{l} (\alpha - \mu) + \frac{\mu(pm)^2}{2} \frac{(pm)_{l-2}}{l!}, \end{aligned} \quad (11)$$

where (11) follows from $p > l^2/\mu$ in (7). Since clearly $pm > l - 1$, this is at most

$$\binom{pm}{l} (\alpha - \mu) + \mu \frac{(pm)_l}{l!} = \alpha \binom{pm}{l} = \alpha \binom{n(\mathcal{H}[V_1, \dots, V_p])}{l}.$$

On the other hand, letting $\xi = p/K < p/k$, (in view of (2)), we obtain that \mathcal{H} is not (α, ξ) -dense. This contradiction completes the proof of the Claim. \square

The previous Claim implies that every l -tuple of the parts of $\mathcal{H}[V_1, \dots, V_p]$ has density at least $\alpha - \mu$. We next prove that $\mathcal{H}[V_1, \dots, V_p]$ contains an $(\alpha, 4\mu)$ -uniform hypergraph.

Let $\mathcal{H}[L]$ be the l -graph induced by the l -tuple $L = (V_{j_1}, \dots, V_{j_l})$. If the density of $\mathcal{H}[L]$ is greater than $\alpha + \mu$, then we apply Lemma 4.1 to $\mathcal{H}[L]$ to obtain $\mathcal{H}'[L]$ that is $(\alpha, 2\epsilon)$ -uniform (here we need $m > m_0$ from Lemma 4.1, which we may assume since n_0 is sufficiently large). If the density of $\mathcal{H}[L]$ is $\alpha \pm \mu$, then let $\mathcal{H}'[L] = \mathcal{H}[L]$. Consequently $\mathcal{H}'[L]$ is $(\alpha', 2\epsilon)$ -uniform, where $\alpha' = \alpha \pm \mu$.

Let $\mathcal{H}(\mu)$ be obtained from $\mathcal{H}[V_1, \dots, V_p]$ by replacing $\mathcal{H}[L]$ by $\mathcal{H}'[L]$ for each of the $\binom{p}{l}$ l -tuples L . Set $m = n(\mathcal{H}(\mu))/p = |V_j|$ for all $1 \leq j \leq p$. Note that $n(\mathcal{H}(\mu)) = mp$ can be made arbitrarily large, thus showing (6). Our next goal which will complete the proof is

Claim: $\mathcal{H}(\mu)$ is $(\alpha, 4\mu)$ -uniform.

Proof of Claim: Let $S \subset V(\mathcal{H}(\mu))$ be a set of size $4\mu n(\mathcal{H}(\mu)) = 4\mu mp$. We will show that S has density $\alpha \pm 4\mu$. Consider the partition $S = S_1 \cup S_2$, where S_1 consists of those elements of S which lie in a V_j that contains fewer than $2\epsilon|V_j| = 2\epsilon m$ elements of S . Then $|S_1| < 2\epsilon mp$.

Subclaim 1: The number of l -sets in S (not necessarily edges) that intersect S_1 is at most $(\mu\alpha/10) \binom{4\mu mp}{l}$.

Proof of Subclaim: We first choose a vertex in S_1 and then all possible $(l-1)$ -sets that contain it. This yields at most $2\epsilon mp \binom{mp}{l-1}$ l -sets. Now

$$2\epsilon mp \binom{mp}{l-1} = \frac{2\epsilon mp l}{l!} \prod_{i=0}^{l-2} (mp - i) < \frac{2\epsilon mp l}{l!} \prod_{i=0}^{l-2} (4mp - i/\mu) = \frac{2\epsilon mp l}{\mu^l (4mp - (l-1)/\mu)} \binom{4mp\mu}{l},$$

where the inequality holds since, for example, $p > (l-2)/\mu$. Because of (9) and $2mp > (l-1)/\mu$, this expression is upper bounded by

$$\frac{2mp\mu\alpha}{10(4mp - (l-1)/\mu)} \binom{4mp\mu}{l} < \frac{\mu\alpha}{10} \binom{4mp\mu}{l}. \quad \square$$

Subclaim 2: The number of l -sets in S intersecting some V_j in more than one vertex is at most $(\mu\alpha/10) \binom{4\mu mp}{l}$.

Proof of Subclaim: We first choose two vertices in some V_j and then all possible $(l-2)$ -sets that contain it. This yields at most $p \binom{m}{2} \binom{mp}{l-2}$ l -sets. This is at most

$$\frac{pm^2 \prod_{i=0}^{l-3} (mp - i)}{2(l-2)!} < \frac{pm^2 l^2}{2l!} \prod_{i=0}^{l-3} (4mp - i/\mu) = \frac{pm^2 l^2 \binom{4mp\mu}{l}}{2\mu^l (4mp - (l-2)/\mu) (4mp - (l-1)/\mu)}.$$

Because of $p > 2l^2/(\mu^{l+1}\alpha)$ from (7) and $2mp > (l-1)/\mu$, this expression is upper bounded by

$$\frac{pm^2 l^2}{2\mu^l (2mp)^2} \binom{4mp\mu}{l} < \frac{\mu\alpha}{10} \binom{4mp\mu}{l}. \quad \square$$

An upper bound for the number of l -sets from Subclaims 1 and 2 is $(2\mu\alpha/10) \binom{4\mu mp}{l}$. We have therefore proved that of the $\binom{4\mu mp}{l}$ l -sets $\{v_1, \dots, v_l\}$ of S , at least $(1 - 2\mu\alpha/10) \binom{4\mu mp}{l}$ are such that for some $1 \leq j_1 < \dots < j_l \leq p$

- 1) $v_i \in V_{j_i} \cap S$, and
- 2) $|V_{j_i} \cap S| \geq 2\epsilon m$ for all $i = 1, \dots, l$ (in other words $\{v_1, \dots, v_l\} \subseteq S_2$).

Since the l -tuple $(V_{j_1}, \dots, V_{j_l})$ is $(\alpha', 2\epsilon)$ -uniform, where $\alpha' = \alpha \pm \mu$, we infer that

$$e(S \cap V_{j_1}, \dots, S \cap V_{j_l}) = (\alpha \pm (2\epsilon + \mu)) \prod_{i=1}^l |S \cap V_{j_i}|$$

whenever 2) holds. Consequently, an upper bound for the number of edges in S is

$$\left(1 - \frac{2\mu}{10}\alpha\right) \binom{4\mu mp}{l} (\alpha + \mu + 2\epsilon) + \frac{2\mu}{10}\alpha \binom{4\mu mp}{l} < (\alpha + 4\mu) \binom{4\mu mp}{l}.$$

On the other hand, a lower bound is

$$\left(1 - \frac{2\mu}{10}\alpha\right) \binom{4\mu mp}{l} (\alpha - \mu - 2\epsilon) > (\alpha - 4\mu) \binom{4\mu mp}{l}.$$

Therefore S has density $(\alpha \pm 4\mu)$. This completes the proof of the claim and theorem. \square

5 Proof of Theorem 1.5

In this section we use Theorem 1.9 to prove Theorem 1.5, namely, show that there exist infinitely many l -graphs for which $0 < \tilde{\rho}(\mathcal{G}) < \rho(\mathcal{G}) \leq \pi(\mathcal{G})$. Note that this is indeed easier than the corresponding question of Erdős and Sós [3] for $\rho(\mathcal{G})$, who asked for l -graphs \mathcal{G} satisfying $0 < \rho(\mathcal{G}) < \pi(\mathcal{G})$ (see (1)), since $\tilde{\rho}(\mathcal{G}) \leq \rho(\mathcal{G})$ holds always. Nevertheless, we know of no simple method to generate such examples (other than the method in Section 2), and Theorem 1.5 implies that any \mathcal{G} with $\tilde{\rho}(\mathcal{G}) > 0$ can be suitably augmented to provide such an example. The proof of this implication is the content of this section.

A vertex v in an l -graph \mathcal{G} is said to be *multiplied* if we replace it by l new vertices v_1, \dots, v_l , replace each edge E containing v by the l edges $(E - \{v\}) \cup \{v_i\}$ for each i , and add the edge $\{v_1, \dots, v_l\}$. The following Lemma is a natural extension of the supersaturation proof for π applied to $\tilde{\rho}$.

Lemma 5.1. *Let \mathcal{G}' be an l -graph obtained from \mathcal{G} by performing a sequence of vertex multiplications. Then $\tilde{\rho}(\mathcal{G}') = \tilde{\rho}(\mathcal{G})$.*

Proof. We may assume that \mathcal{G}' is obtained by multiplying a single vertex v_1 of \mathcal{G} , since we can then repeat the argument. As obviously $\tilde{\rho}(\mathcal{G}') \geq \tilde{\rho}(\mathcal{G})$, we prove $\tilde{\rho}(\mathcal{G}') \leq \tilde{\rho}(\mathcal{G})$ by showing that for all $\tilde{\delta} > 0$, $\tilde{\rho}(\mathcal{G}') \leq \tilde{\rho}(\mathcal{G}) + 2\tilde{\delta}$. Now fix $\tilde{\delta} > 0$ and let $\alpha = \tilde{\rho}(\mathcal{G}) + 2\tilde{\delta}$. By definition of the function $\tilde{\rho}$, we must show that every α -uniform sequence of l -graphs contains \mathcal{G}' . Pick such a sequence $\{\mathcal{H}_i\}$.

Let $\alpha, l, \tilde{\delta}$ be the input of Theorem 1.9, and δ, r, n_0 be the output, with $\delta < 1/2r^{r+1}$ (see Remark 1.11). Choose t sufficiently large that

- 1) $n(\mathcal{H}_t) \geq n_0$, and
- 2) \mathcal{H}_t is (α, δ) -uniform.

Note that t exists since $\{\mathcal{H}_i\}$ is α -uniform. We will apply Theorem 1.9 to $\mathcal{H} = \mathcal{H}_t$ and prove that \mathcal{H} contains a copy of \mathcal{G}' . Our proof will apply to every element of $\{\mathcal{H}_i\}$ satisfying 1) and 2). Consequently, all but finitely many members of $\{\mathcal{H}_i\}$ contain \mathcal{G}' and therefore $\{\mathcal{H}_i\}$ contains \mathcal{G}' .

By Theorem 1.9, all but $e^{-r^{1/l}/20} \binom{n}{r}$ of the r -sets of vertices of \mathcal{H} induce a subsystem \mathcal{H}' that is $(\alpha, \tilde{\delta})$ -uniform. By the choice of α , the subsystem \mathcal{H}' has density at least $\tilde{\rho}(\mathcal{G}) + \tilde{\delta}$. Since r is sufficiently large, \mathcal{H}' contains a copy of \mathcal{G} . Letting $m = n(\mathcal{G})$, we deduce that the number of copies of \mathcal{G} in \mathcal{H} is at least

$$\frac{(1 - e^{-r^{1/l}/20}) \binom{n}{r}}{\binom{n-m}{r-m}} = \frac{(1 - e^{-r^{1/l}/20}) \binom{n}{m}}{\binom{r}{m}} > \frac{1}{r^m} \binom{n}{m}. \quad (12)$$

Now consider the m -graph \mathcal{F} whose edges consist of the copies of \mathcal{G} in \mathcal{H} , and associate to each edge of \mathcal{F} a special vertex (corresponding to the vertex that plays the role of v_1). For an $(m-1)$ -set M , let c_M be the number of edges of \mathcal{F} containing M whose special vertex is outside M . Then

$$\sum_{|M|=m-1} c_M = e(\mathcal{F}) > \frac{1}{r^m} \binom{n}{m},$$

so there is an M_0 with

$$c_{M_0} > \frac{1}{r^m} \frac{\binom{n}{m}}{\binom{n}{m-1}} > \frac{1}{2mr^m} n.$$

Now since $\delta < 1/2r^{r+1}$ and $r > m$ we conclude that $1/(2mr^m) > \delta$. Consequently, the set of all special vertices corresponding to M_0 contains an edge (in fact many). We have found a copy of \mathcal{G}' as required. This completes the proof. \square

Remark 5.2. *In the above proof we didn't need the full strength of Theorem 1.9 (see Remark 1.12). We only needed in (12) that some fixed proportion of r -sets inherit the property of being (α, δ) -uniform.*

Definition 5.3. *Fix $l \geq 3$. Let \mathcal{G}_0 be the l -graph of one vertex. For $i \geq 1$, let \mathcal{G}_i be the l -graph obtained by taking l disjoint copies of \mathcal{G}_{i-1} , and adding all edges with exactly one point in each copy.*

An l -graph is *simple* if every two edges have at most one point in common. For an l -graph \mathcal{G} , define $\chi_s(\mathcal{G})$ to be the minimum number of colors needed to partition $V(\mathcal{G})$ such that every color class induces a simple l -graph. Let $\alpha_s(\mathcal{G})$ be the maximum size of a subset of vertices that induces a simple l -graph.

Lemma 5.4. *Fix $l \geq 3$. Then $\chi_s(\mathcal{G}_i) \geq (1 + 1/(l-1))^{i-1}$. In particular, $\chi_s(\mathcal{G}_i) \rightarrow \infty$ as $i \rightarrow \infty$.*

Proof. Observe that $\chi_s(\mathcal{G}) \geq n/\alpha_s(\mathcal{G})$ for every l -graph with n vertices. It therefore suffices to prove that $\alpha_s(\mathcal{G}_i) \leq l(l-1)^{i-1}$, since \mathcal{G}_i has l^i vertices. We proceed by induction on i . The base case is $i = 1$, where the result is trivial. Now suppose the result holds for $i-1$, and consider \mathcal{G}_i comprising l copies of \mathcal{G}_{i-1} . The crucial observation is that taking two points from one copy of \mathcal{G}_{i-1} , and one point from each of the other $l-1$ copies of \mathcal{G}_{i-1} , yields a subset that does not induce a simple l -graph. Consequently,

$$\alpha_s(\mathcal{G}_i) \leq \max\{l, (l-1)\alpha_s(\mathcal{G}_{i-1})\} \leq (l-1)[l(l-1)^{i-2}] = l(l-1)^{i-1}.$$

Since l is fixed, the last statement is trivial. \square

Corollary 5.5. $\rho(\mathcal{G}_i) \rightarrow 1$ as $i \rightarrow \infty$.

Proof. Let $k + 1 = \chi_s(\mathcal{G}_i)$. Fix $\epsilon > 0$, and let \mathcal{A}_n be an n -vertex simple l -graph with independence number at most ϵn . Such l -graphs exist for every $n \geq n_0$ by an easy application of the probabilistic method (see, e.g., [3]). We will moreover assume that $n_0 \geq l/\epsilon$. Let \mathcal{H}_n be the nk -vertex l -graph with vertex set $\cup_{j=1}^k X_j$, where X_j is a copy of $V(\mathcal{A}_n)$ for each j . The edge set of \mathcal{H}_n consists of all l -sets intersecting at least two of the X_j 's, together with the l -sets within each X_j , where the hypergraph induced by each X_j is isomorphic to \mathcal{A}_n . Since $\epsilon n \geq l$, we have $\alpha(\mathcal{H}_n) \leq \epsilon n$. Also, $\chi_s(\mathcal{H}_n) \leq k$, which implies that \mathcal{H}_n contains no copy of \mathcal{G}_i . Furthermore, \mathcal{H}_n has density at least $1 - k^{1-l}$. By Lemma 5.4, $k \rightarrow \infty$ as $i \rightarrow \infty$, so $\rho(\mathcal{G}_i) \geq 1 - k^{1-l} \rightarrow 1$. \square

Lemma 5.6. Fix $l \geq 3$. For every $i \geq 0$, there exists an l -graph \mathcal{G}' obtained by a sequence of vertex multiplications (starting from a single vertex) such that $\mathcal{G}_i \subset \mathcal{G}'$.

Proof. We proceed by induction on i . When $i = 0$, the result is trivial, so let $i > 0$. Beginning with a vertex x_1 , form the edge $\{x_1, \dots, x_l\}$. Now by induction, for each j , start from x_j and perform a sequence of vertex multiplications to produce an l -graph containing \mathcal{G}_{i-1} . The resulting l -graph contains \mathcal{G}_i . \square

Proof of Theorem 1.5: Fix an l -graph \mathcal{G} , and $\epsilon > 0$. By Corollary 5.5, choose t so that $\rho(\mathcal{G}_t) > 1 - \epsilon$. Now pick a vertex x in \mathcal{G} , and by Lemma 5.6, perform a sequence of vertex multiplications to transform x to an l -graph \mathcal{G}' containing \mathcal{G}_t . Let $\hat{\mathcal{G}}$ be the l -graph with vertex set $V(\mathcal{G}) \cup V(\mathcal{G}')$ that results after these multiplications. Since $\hat{\mathcal{G}}$ is obtained from \mathcal{G} by a sequence of vertex multiplications, Lemma 5.1 implies that $\tilde{\rho}(\hat{\mathcal{G}}) = \tilde{\rho}(\mathcal{G})$. By the choice of t , and since $\mathcal{G}_t \subset \hat{\mathcal{G}}$, we also have $\rho(\hat{\mathcal{G}}) \geq \rho(\mathcal{G}_t) \geq 1 - \epsilon$. \square

6 Acknowledgments

Part of this work was done while the first author was a postdoctoral researcher in the Theory Group at Microsoft Research. The authors thank Mathias Schacht for useful comments, and Jozef Skokan for his careful reading of an earlier draft of this paper. The authors also thank the referees for their detailed comments.

References

- [1] N. Alon, W. F. de la Vega, R. Kannan and M. Karpinski, Random Sampling and Approximation of MAX-CSP Problems, Proc. of the 34 ACM STOC, ACM Press (2002) 232–239.
- [2] P. Erdős, M. Simonovits, A limit theorem in graph theory. *Studia Sci. Math. Hungar* 1 1966 51–57.
- [3] P. Erdős, V. T. Sós, On Ramsey-Turán type theorems for hypergraphs. *Combinatorica* 2 (1982), no. 3, 289–295.
- [4] P. Erdős, A. H. Stone, On the structure of linear graphs. *Bull. Amer. Math. Soc.* 52, (1946). 1087–1091.
- [5] P. Frankl, Z. Füredi, A new generalization of the Erdős-Ko-Rado theorem. *Combinatorica* 3 (1983), no. 3-4, 341–349.
- [6] P. Frankl, V. Rödl, Some Ramsey-Turán type results for hypergraphs. *Combinatorica* 8 (1988), no. 4, 323–332
- [7] S. Janson, T. Luczak, A. Ruciński, *Random Graphs*, Wiley (2000)
- [8] G. Katona, T. Nemetz, M. Simonovits, On a problem of Turán in the theory of graphs. (Hungarian. Russian, English summary) *Mat. Lapok* 15 1964 228–238.
- [9] J. Komlós, M. Simonovits, Szemerédi’s regularity lemma and its applications in graph theory. *Combinatorics, Paul Erdős is eighty, Vol. 2* (Keszthely, 1993), 295–352, *Bolyai Soc. Math. Stud.*, 2, János Bolyai Math. Soc., Budapest, 1996
- [10] D. Mubayi, V. Rödl, On the Turán number of triple systems, *J. Combin Theory, Ser. A*, 100 (2002), 136–152.
- [11] D. Mubayi, V. Rödl, Uniform edge distribution in hypergraphs is hereditary. *Electronic Journal of Combinatorics*, 11 (2004) Research paper R55 32 pp (electronic).
- [12] H. J. Prömel, A. Steger, Excluding induced subgraphs. III. A general asymptotic. *Random Structures and Algorithms* 3 (1992), no. 1, 19–31.
- [13] F. P. Ramsey, On a Problem of Formal Logic, *Proceedings of the London Mathematical Society*, 30 (1930) 264–286.

- [14] A. F. Sidorenko, On Ramsey-Turán numbers for 3-graphs. *J. Graph Theory* 16 (1992), no. 1, 73–78.
- [15] M. Simonovits, Extremal graph problems, degenerate extremal problems, and supersaturated graphs. *Progress in graph theory* (Waterloo, Ont., 1982), 419–437, Academic Press, Toronto, ON, 1984.
- [16] M. Simonovits, V. T. Sós, Ramsey-Turán theory, *Combinatorics, graph theory, algorithms and applications*. *Discrete Math.* 229 (2001), no. 1-3, 293–340.
- [17] E. Szemerédi, Regular partitions of graphs. (English. French summary) *Problèmes combinatoires et théorie des graphes* (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), pp. 399–401, Colloq. Internat. CNRS, 260, CNRS, Paris, 1978.
- [18] P. Turán, Eine Extremalaufgabe aus der Graphentheorie. (Hungarian) *Mat. Fiz. Lapok* 48, (1941). 436–452