EMBEDDINGS AND RAMSEY NUMBERS OF SPARSE k-UNIFORM HYPERGRAPHS

OLIVER COOLEY, NIKOLAOS FOUNTOULAKIS, DANIELA KÜHN AND DERYK OSTHUS

ABSTRACT. Chvátal, Rödl, Szemerédi and Trotter [3] proved that the Ramsey numbers of graphs of bounded maximum degree are linear in their order. In [5, 19] the same result was proved for 3-uniform hypergraphs. Here we extend this result to k-uniform hypergraphs for any integer $k \geq 3$. As in the 3-uniform case, the main new tool which we prove and use is an embedding lemma for k-uniform hypergraphs of bounded maximum degree into suitable k-uniform 'quasi-random' hypergraphs.

keywords: hypergraphs; regularity lemma; Ramsey numbers; embedding problems

1. Introduction

1.1. **Ramsey numbers.** The *Ramsey number* $R(\mathcal{H})$ of a k-uniform hypergraph $\mathcal{H}$ is the smallest $N \in \mathbb{N}$ such that for every 2-colouring of the hyperedges of the complete k-uniform hypergraph on N vertices one can find a monochromatic copy of $\mathcal{H}$. For general $\mathcal{H}$, the best upper bound is due to Erdős and Rado [7]. It implies that $R(\mathcal{H})$ is at most $|\mathcal{H}|$ raised to a tower of height k (where a tower of height 2 is simply an exponential function). For the case of graphs (i.e. when $k = 2$) it is known that there are many families of graphs H for which the Ramsey number is much smaller than exponential. In particular, Burr and Erdős [2] asked for which graphs H the Ramsey number $R(H)$ is linear in the order $|H|$ of H and conjectured this to be true for graphs of bounded maximum degree. This was proved by Chvátal, Rödl, Szemerédi and Trotter [3]. Here we show that their result extends to k-uniform hypergraphs $\mathcal{H}$ of bounded maximum degree, where the *degree of a vertex* x in $\mathcal{H}$ is defined to be the number of hyperedges which contain x.

Theorem 1. *For all $\Delta, k \in \mathbb{N}$ there exists a constant $C = C(\Delta, k)$ such that all k-uniform hypergraphs $\mathcal{H}$ of maximum degree at most Δ satisfy $R(\mathcal{H}) \leq C|\mathcal{H}|$.*

The case of 3-uniform hypergraphs was proved earlier in [5] and independently by Nagle, Olsen, Rödl and Schacht [19]. Also, Kostochka and Rödl [17] recently proved an approximate version of Theorem 1: for all $\varepsilon, \Delta, k > 0$ there is a constant C such that $R(\mathcal{H}) \leq C|\mathcal{H}|^{1+\varepsilon}$ if $\mathcal{H}$ has maximum degree at most Δ. Apart from this, the only other result on the Ramsey number of sparse hypergraphs is due to Haxell et al. [12, 13], who asymptotically determined the Ramsey numbers of 3-uniform tight and loose cycles.

The overall strategy of our proof of Theorem 1 is related to that of Chvatál et al. [3], which is based on the regularity lemma for graphs. We apply a version of the regularity lemma for k-uniform hypergraphs due to Rödl and Schacht [23]. Roughly speaking, it guarantees a partition of an arbitrary dense k-uniform hypergraph into 'quasi-random' subhypergraphs. Our main contribution is an embedding lemma which guarantees the existence of a copy of a hypergraph $\mathcal{H}$ of bounded maximum degree inside a suitable 'quasi-random' hypergraph $\mathcal{G}$ even if the order of $\mathcal{H}$ is linear in that of $\mathcal{G}$.

1.2. **Open questions.** Several open problems immediately suggest themselves. Firstly, it would be desirable to extend Theorem 1 to a larger class of hypergraphs. For instance the graph analogue of Theorem 1 is known for so-called p-arrangeable graphs [1], which include the class of all planar graphs. However, Rödl and Kostochka [17] showed that a natural hypergraph analogue of the famous Burr-Erdős conjecture on Ramsey numbers of d-degenerate graphs fails for k-uniform hypergraphs if $k \geq 3$. (A graph is d-degenerate if the maximum average degree over all its subgraphs is at most d. If a graph is p-arrangeable, then it is also d-degenerate for some d.) But it may still be possible to generalize the Burr-Erdős conjecture to hypergraphs in a different way.

Secondly, it may be possible to strengthen the embedding lemma to allow for embeddings of a k-uniform hypergraph $\mathcal{H}$ of bounded degree into a suitable k-unifom 'quasi-random' hypergraph $\mathcal{G}$, where the order of $\mathcal{H}$ is allowed to be almost the same as that of $\mathcal{G}$. But it seems that this would require rather different methods. However, one case where we can prove such a result is when $\mathcal{H}$ has bounded bandwidth. This will be done in Section 9. Ideally, one would of course like to have a hypergraph analogue of the blow-up lemma for graphs, which would allow both hypergraphs to have exactly the same order.

Thirdly, Graham, Rödl and Ruciński [11] gave a proof of the result of Chvátal et al. [3] which does not rely on the regularity lemma. It would be interesting to know whether this can also be done in the case of Theorem 1.

1.3. **Organization of the paper.** In Section 2 we state the embedding lemma mentioned in the introduction and give an overview of its proof. Our proof relies on a more general version (Lemma 4) of the well-known counting lemma for hypergraphs as well as on an 'extension lemma' (Lemma 5). We introduce these lemmas, along with further tools, in Section 3. In Section 4 we derive our version of the counting lemma from that in [23]. In Section 5 we use it to deduce the extension lemma. We then prove the embedding lemma in Section 6. The regularity lemma for k-uniform hypergraphs is introduced in Section 7. In Section 8 we deduce Theorem 1 from the regularity lemma and the embedding lemma. In the final section, we prove and discuss an embedding lemma for k-uniform hypergraphs which guarantees the existence of almost spanning subhypergraphs of bounded bandwidth.

2. Overview of the proof of Theorem 1 and statement of the embedding lemma

2.1. **Overview of the proof of Theorem 1.** The proof in [3] that graphs of bounded degree have linear Ramsey numbers proceeds roughly as follows: Let H be a graph of maximum degree Δ. Take a complete graph K_n, where n is a sufficiently large integer. Colour the edges of K_n with red and blue, and apply the graph regularity lemma to the denser of the two monochromatic graphs, G_{red} say, to obtain a partition of the vertex set into a bounded number of clusters. Since almost all pairs of clusters are regular or 'quasi-random', by Turán's theorem there will be a set of r clusters, where $r := R(K_{\Delta+1})$, in which each pair of clusters is regular. A pair of clusters will be coloured red if its density in G_{red} is at least $1/2$, and blue otherwise. By the definition of r, there must be a set of $\Delta + 1$ clusters such that all the pairs have the same colour. If this colour is red, then one can apply the so-called embedding or key lemma for graphs to find a (red) copy of H in the subgraph of G_{red} spanned by these $\Delta + 1$ clusters. This is possible since $\chi(H) \leq \Delta + 1$. If all the pairs of clusters are coloured blue we apply the embedding lemma in the blue subgraph G_{blue} of K_n to find a blue copy of H. It turns out that in this proof we only needed $n \geq C|H|$, where C is a constant dependent only on Δ. Thus $R(H) \leq C|H|$.

We will generalize this approach to k-uniform hypergraphs. As mentioned in Section 1, the main obstacle is the proof of an embedding lemma for k-uniform hypergraphs (Lemma 2 below), which allows us to embed a k-uniform hypergraph $\mathcal{H}$ within a suitable 'quasi-random'

k-uniform hypergraph $\mathcal{G}$, where the order of $\mathcal{H}$ might be linear in the order of $\mathcal{G}$. Our proof uses ideas from [5].

2.2. **Notation and statement of the embedding lemma.** Before we can state the embedding lemma, we first have to say what we mean by regular or 'quasi-random' hypergraphs. In the setup below, this will involve the relationship between certain i-uniform hypergraphs and $(i-1)$-uniform hypergraphs on the same vertex set. Given a hypergraph $\mathcal{G}$, we write $E(\mathcal{G})$ for the set of its hyperedges and define $e(\mathcal{G}) := |E(\mathcal{G})|$. We write $K_i^{(j)}$ for the complete j-uniform hypergraph on i vertices. Given an i-partite j-uniform hypergraph $\mathcal{G}$, we write $\mathcal{K}_i(\mathcal{G})$ for the set of i-sets of vertices of $\mathcal{G}$ which form a copy of $K_i^{(j)}$ in $\mathcal{G}$. We often write $|K_i^{(j)}|_{\mathcal{G}}$ instead of $|\mathcal{K}_i(\mathcal{G})|$. Given an i-partite i-uniform hypergraph $\mathcal{G}_i$, and an i-partite $(i-1)$-uniform hypergraph $\mathcal{G}_{i-1}$ on the same vertex set, we define the *density of $\mathcal{G}_i$ with respect to $\mathcal{G}_{i-1}$* to be

$$d(\mathcal{G}_i|\mathcal{G}_{i-1}) := \frac{|\mathcal{K}_i(\mathcal{G}_{i-1}) \cap E(\mathcal{G}_i)|}{|\mathcal{K}_i(\mathcal{G}_{i-1})|}$$

if $|\mathcal{K}_i(\mathcal{G}_{i-1})| > 0$, and $d(\mathcal{G}_i|\mathcal{G}_{i-1}) := 0$ otherwise. More generally, if $\mathbf{Q} := (Q(1), Q(2), \dots, Q(r))$ is a collection of r subhypergraphs of $\mathcal{G}_{i-1}$, we define $\mathcal{K}_i(\mathbf{Q}) := \bigcup_{j=1}^r \mathcal{K}_i(Q(j))$ and

$$d(\mathcal{G}_i|\mathbf{Q}) := \frac{|\mathcal{K}_i(\mathbf{Q}) \cap E(\mathcal{G}_i)|}{|\mathcal{K}_i(\mathbf{Q})|}$$

if $|\mathcal{K}_i(\mathbf{Q})| > 0$, and $d(\mathcal{G}_i|\mathbf{Q}) := 0$ otherwise. Again, we often write $|K_i^{(j)}|_{\mathbf{Q}}$ instead of $|\mathcal{K}_i(\mathbf{Q})|$.

We say that $\mathcal{G}_i$ is *(d_i, δ, r)-regular with respect to $\mathcal{G}_{i-1}$* if every r-tuple $\mathbf{Q}$ with $|\mathcal{K}_i(\mathbf{Q})| > \delta|\mathcal{K}_i(\mathcal{G}_{i-1})|$ satisfies

$$d(\mathcal{G}_i|\mathbf{Q}) = d_i \pm \delta.$$

Given $\ell \geq i \geq 3$, an ℓ-partite i-uniform hypergraph $\mathcal{G}_i$ and an ℓ-partite $(i-1)$-uniform hypergraph $\mathcal{G}_{i-1}$ on the same vertex set, we say that $\mathcal{G}_i$ is (d_i, δ, r)-regular with respect to $\mathcal{G}_{i-1}$ if for every i-tuple K of vertex classes, either $\mathcal{G}_i[K]$ is (d_i, δ, r)-regular with respect to $\mathcal{G}_{i-1}[K]$ or $d(\mathcal{G}_i[K]|\mathcal{G}_{i-1}[K]) = 0$ (but the latter should not hold for all K). Instead of $(d_i, \delta, 1)$-regularity we sometimes refer to *(d_i, δ)-regularity.*

The density of a bipartite graph G with vertex classes A and B is defined by $d(A, B) := e(A, B)/|A||B|$ and G is (d, δ)-regular if for all sets $X \subseteq A$ and $Y \subseteq B$ with $|X| \geq \delta|A|$ and $|Y| \geq \delta|B|$ we have $d(X, Y) = d \pm \delta$. We say that an ℓ-partite graph $\mathcal{G}_2$ is *(d_2, δ)-regular* if each of the $\binom{\ell}{2}$ bipartite subgraphs forming it is either (d_2, δ)-regular or has density 0 (and if for at least one of them the former holds).

Suppose that we have $\ell \geq k$ vertex classes $V_1, \dots, V_\ell$, and that for each $i = 2, \dots, k$ we are given an ℓ-partite i-uniform hypergraph $\mathcal{G}_i$ with these vertex classes. Suppose also that $\mathcal{H}$ is an ℓ-partite k-uniform hypergraph with vertex classes $X_1, \dots, X_\ell$. We will aim to embed $\mathcal{H}$ into $\mathcal{G}_k$, and in particular to embed X_j into V_j for each $j = 1, \dots, \ell$. So we make the following definition: We say that $(\mathcal{G}_k, \dots, \mathcal{G}_2)$ *respects the partition of $\mathcal{H}$* if whenever $\mathcal{H}$ contains a hyperedge with vertices in $X_{j_1}, \dots, X_{j_k}$, then there is a hyperedge of $\mathcal{G}_k$ with vertices in $V_{j_1}, \dots, V_{j_k}$ which also forms a copy of $K_k^{(i)}$ in $\mathcal{G}_i$ for each $i = 2, \dots, k-1$.

Lemma 2 (Embedding lemma for hypergraphs)**.** *Let Δ, k, ℓ, r, n_0 be positive integers with $k \leq \ell$ and let $c, d_2, d_3, \dots, d_k, \delta, \delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$,*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2, \dots, d_{k-1}\} \leq \delta_k \ll d_k, 1/\Delta, 1/\ell$$

and

$$c \ll d_2, \dots, d_k, 1/\Delta, 1/\ell.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}$ is an ℓ-partite k-uniform hypergraph of maximum degree at most Δ with vertex classes $X_1, \dots, X_\ell$ such that $|X_i| \leq cn$ for all $i = 1, \dots, \ell$. Suppose that for each $i = 2, \dots, k$, $\mathcal{G}_i$ is an ℓ-partite i-uniform hypergraph with vertex classes $V_1, \dots, V_\ell$, which all have size n. Suppose also that $\mathcal{G}_k$ is (d_k, δ_k, r)-regular with respect to $\mathcal{G}_{k-1}$, that for each $i = 3, \dots, k-1$, $\mathcal{G}_i$ is (d_i, δ)-regular with respect to $\mathcal{G}_{i-1}$, that $\mathcal{G}_2$ is (d_2, δ)-regular, and that $(\mathcal{G}_k, \dots, \mathcal{G}_2)$ respects the partition of $\mathcal{H}$. Then $\mathcal{G}_k$ contains a copy of $\mathcal{H}$.

In the statement of Lemma 2 we used the following notation (which will be used frequently later on as well). Given constants a_1, a_2, a_3, we write $a_1 \ll a_2 \ll a_3$ to mean that we choose these constants from right to left, and there are increasing functions f and g such that the lemma holds provided that $a_2 \leq f(a_3)$ and $a_1 \leq g(a_2)$. The functions f and g are determined by the calculations in the proof of Lemma 2, but for clarity of the exposition we will not determine them explicitly.

3. Further notation and tools

3.1. **Embedding lemma for complexes.** Instead of Lemma 2, we will prove a considerably stronger version which appears as Lemma 3 below. It allows the embedding of hypergraphs which are not necessarily uniform and gives a lower bound on the number of such embeddings. This enables us to prove the lemma by induction on $|\mathcal{H}|$. Before we can state Lemma 3, we need to make the following definitions.

A *complex* $\mathcal{H}$ on a vertex set V is a collection of subsets of V, each of size at least 2, such that if $B \in \mathcal{H}$, and if $A \subseteq B$ has size at least 2, then $A \in \mathcal{H}$. A *k-complex* is a complex in which no member has size greater than k. The members of size $i \geq 2$ are called the *i-edges* of $\mathcal{H}$ and the elements of V are called the *vertices* of $\mathcal{H}$. We write $E_i(\mathcal{H})$ for the set of all i-edges of $\mathcal{H}$ and set $e_i(\mathcal{H}) := |E_i(\mathcal{H})|$. We also write $|\mathcal{H}| := |V|$ for the *order* of $\mathcal{H}$. Note that a k-uniform hypergraph can be turned into a k-complex by making every hyperedge into a complete i-uniform hypergraph $K_k^{(i)}$, for each $2 \leq i \leq k$. (In a more general k-complex we may have i-edges which do not lie within an $(i+1)$-edge.) Given $k \leq \ell$, a *(k, ℓ)-complex* is an ℓ-partite k-complex. Given a k-complex $\mathcal{H}$, for each $i = 2, \dots, k$ we denote by $\mathcal{H}_i$ the *underlying i-uniform hypergraph of* $\mathcal{H}$. So the vertices of $\mathcal{H}_i$ are those of $\mathcal{H}$ and the hyperedges of $\mathcal{H}_i$ are the i-edges of $\mathcal{H}$.

Two vertices x and y in a k-complex are *neighbours* if they are joined by a 2-edge. (Note that if x and y lie in a common i-edge for some $2 \leq i \leq k$, then they are joined by a 2-edge.) The *degree* $d(x)$ of a vertex x is the maximum (over $2 \leq i \leq k$) of the number of i-edges containing x. Thus x has at most $d(x)$ neighbours. The *maximum degree* of the complex $\mathcal{H}$ is the greatest degree of any vertex. Note that if $\mathcal{H}$ is a k-uniform hypergraph of maximum degree Δ, the maximum degree of the corresponding k-complex is still bounded as a function of Δ and k. The *distance* between two vertices x and y in a k-complex $\mathcal{H}$ is the length of the shortest path between x and y in the underlying 2-graph $\mathcal{H}_2$ of $\mathcal{H}$. The *components* of $\mathcal{H}$ are the subcomplexes induced by the components of $\mathcal{H}_2$.

We say that a k-complex $\mathcal{G}$ is $(d_k, \dots, d_2, \delta_k, \delta, r)$-regular if $\mathcal{G}_k$ is (d_k, δ_k, r)-regular with respect to $\mathcal{G}_{k-1}$, if $\mathcal{G}_i$ is (d_i, δ)-regular with respect to $\mathcal{G}_{i-1}$ for each $i = 3, \dots, k-1$, and if $\mathcal{G}_2$ is (d_2, δ_2)-regular. We denote $(d_k, \dots, d_2)$ by $\mathbf{d}$ and refer to $(\mathbf{d}, \delta_k, \delta, r)$-regularity.

Suppose that $\mathcal{G}$ is a (k, ℓ)-complex with vertex classes $V_1, \dots, V_\ell$, which all have size n. Suppose also that $\mathcal{H}$ is a (k, ℓ)-complex with vertex classes $X_1, \dots, X_\ell$ of size at most n. Similarly as for hypergraphs we say that $\mathcal{G}$ *respects the partition of* $\mathcal{H}$ if whenever $\mathcal{H}$ contains an i-edge with vertices in $X_{j_1}, \dots, X_{j_i}$, then there is an i-edge of $\mathcal{G}$ with vertices in $V_{j_1}, \dots, V_{j_i}$. On the other hand, we say that a labelled copy of $\mathcal{H}$ in $\mathcal{G}$ is *partition-respecting* if for each $i = 1, \dots, \ell$ the vertices corresponding to those in X_i lie within V_i. We denote by $|\mathcal{H}|_{\mathcal{G}}$ the number of labelled, partition-respecting copies of $\mathcal{H}$ in $\mathcal{G}$.

Lemma 3 (Embedding lemma for complexes)**.** *Let Δ, k, ℓ, r, n_0 be positive integers and let $c, \alpha, d_2, \dots, d_k, \delta, \delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$,*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2, \dots, d_{k-1}\} \leq \delta_k \ll \alpha \ll d_k, 1/\Delta, 1/\ell$$

and

$$c \ll \alpha, d_2, \dots, d_k.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}$ is a (k, ℓ)-complex of maximum degree at most Δ with vertex classes $X_1, \dots, X_\ell$ such that $|X_i| \leq cn$ for all $i = 1, \dots, \ell$. Suppose also that $\mathcal{G}$ is a $(\mathbf{d}, \delta_k, \delta, r)$-regular (k, ℓ)-complex with vertex classes $V_1, \dots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}$. Then for every vertex $h \in \mathcal{H}$ we have that

$$|\mathcal{H}|_{\mathcal{G}} \geq (1-\alpha)n \left(\prod_{i=2}^{k} d_i^{e_i(\mathcal{H}) - e_i(\mathcal{H}_h)} \right) |\mathcal{H}_h|_{\mathcal{G}},$$

where $\mathcal{H}_h$ denotes the induced subcomplex of $\mathcal{H}$ obtained by removing h. In particular, $\mathcal{G}$ contains at least $((1-\alpha)n)^{|\mathcal{H}|} \prod_{i=2}^{k} d_i^{e_i(\mathcal{H})}$ labelled partition-respecting copies of $\mathcal{H}$.

Note that the bound relating $|\mathcal{H}|_{\mathcal{G}}$ to $|\mathcal{H}_h|_{\mathcal{G}}$ in Lemma 3 is close to what one would get with high probability if $\mathcal{G}$ were a random complex. This also shows that the bound is close to best possible. Lemma 3 will be proved in Section 6. In the proof we will need two lemmas on embeddings of complexes of bounded order, which are stated in the next subsection.

Recall that if the maximum degree of a k-uniform hypergraph $\mathcal{H}$ is at most Δ then the maximum degree of the corresponding k-complex is bounded by a function of Δ and k. So it is easy to see that Lemma 3 does indeed imply Lemma 2.

3.2. **Counting lemma and extension lemma.** We will need a variant (Lemma 4) of the counting lemma for k-unifom hypergraphs due to Rödl and Schacht [23, Thm 21]. (A similar result was proved earlier by Gowers [9] as well as Nagle, Rödl and Schacht [21].) It states that if $|\mathcal{H}|$ is bounded and $\mathcal{G}$ is suitably regular, then the number of copies of $\mathcal{H}$ in $\mathcal{G}$ is as large as one would expect if $\mathcal{G}$ were random. The main difference to the result in [23] is that Lemma 4 counts copies of k-complexes $\mathcal{H}$ instead of copies of k-uniform hypergraphs $\mathcal{H}$ and also includes an upper bound on the number of these copies. We will derive Lemma 4 from the result in [23] in Section 4.

Lemma 4 (Counting lemma)**.** *Let k, ℓ, r, t, n_0 be positive integers and let $\varepsilon, d_2, \dots, d_k, \delta, \delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$ and*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2, \dots, d_{k-1}\} \leq \delta_k \ll \varepsilon, d_k, 1/\ell, 1/t.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}$ is a (k, ℓ)-complex on t vertices with vertex classes $X_1, \dots, X_\ell$. Suppose also that $\mathcal{G}$ is a $(\mathbf{d}, \delta_k, \delta, r)$-regular (k, ℓ)-complex with vertex classes $V_1, \dots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}$. Then

$$|\mathcal{H}|_{\mathcal{G}} = (1 \pm \varepsilon) n^t \prod_{i=2}^{k} d_i^{e_i(\mathcal{H})}.$$

The main difference between the counting lemma and the embedding lemma is that the counting lemma only allows for complexes $\mathcal{H}$ of bounded order. We will apply the counting lemma to embed complexes of order $\leq f(\Delta, k)$ for some appropriate function f. Note that the upper and lower bounds of the counting lemma imply Lemma 3 for the case when $|\mathcal{H}|$ is bounded. A formal proof of this (which settles the base case for the induction in the proof of Lemma 3) can be found at the beginning of Section 6.

In the induction step of the proof of Lemma 3 we will also need the following extension lemma, which states that if $\mathcal{H}'$ is a complex of bounded order, $\mathcal{H} \subseteq \mathcal{H}'$ is an induced subcomplex and $\mathcal{G}$ is suitably regular, then almost all copies of $\mathcal{H}$ in $\mathcal{G}$ can be extended to the 'right' number of copies of $\mathcal{H}'$, where the 'right' number is the number one would expect if $\mathcal{G}$ were random.

Lemma 5 (Extension lemma). *Let k, ℓ, r, t, t', n_0 be positive integers, where $t < t'$, and let $\beta, d_2, \ldots, d_k, \delta, \delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$ and*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2, \ldots, d_{k-1}\} \leq \delta_k \ll \beta, d_k, 1/\ell, 1/t'.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}'$ is a (k,ℓ)-complex on t' vertices with vertex classes $X_1, \ldots, X_\ell$ and let $\mathcal{H}$ be an induced subcomplex of $\mathcal{H}'$ on t vertices. Suppose also that $\mathcal{G}$ is a $(\mathbf{d}, \delta_k, \delta, r)$-regular (k,ℓ)-complex with vertex classes $V_1, \ldots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}'$. Then all but at most $\beta|\mathcal{H}|_{\mathcal{G}}$ labelled partition-respecting copies of $\mathcal{H}$ in $\mathcal{G}$ are extendible into

$$(1 \pm \beta) n^{t'-t} \prod_{i=2}^{k} d_i^{e_i(\mathcal{H}') - e_i(\mathcal{H})}$$

labelled partition-respecting copies of $\mathcal{H}'$ in $\mathcal{G}$.

As well as these versions of the counting lemma and extension lemma, we will need to be able to apply versions of these lemmas to underlying $(k-1)$-complexes. In this case, we have that the regularity constant δ is much smaller than all the densities $d_2, \ldots, d_{k-1}$, but on the other hand we have no r in the highest level and thus we cannot apply Lemmas 4 and 5. So instead of Lemma 4 we will use the following variant of a result of Kohayakawa, Rödl and Skokan [14, Cor. 6.11].

Lemma 6 (Dense counting lemma). *Let k, ℓ, t, n_0 be positive integers and let $\varepsilon, d_2, \ldots, d_{k-1}, \delta$ be positive constants such that*

$$1/n_0 \ll \delta \ll \varepsilon \ll d_2, \ldots, d_{k-1}, 1/\ell, 1/t.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}$ is a $(k-1,\ell)$-complex on t vertices with vertex classes $X_1, \ldots, X_\ell$. Suppose also that $\mathcal{G}$ is a $(d_{k-1}, \ldots, d_2, \delta, \delta, 1)$-regular $(k-1,\ell)$-complex with vertex classes $V_1, \ldots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}$. Then

$$|\mathcal{H}|_{\mathcal{G}} = (1 \pm \varepsilon) n^t \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H})}.$$

In Section 4 we will show how Lemma 6 can be deduced from the result in [14]. The following dense version of the extension lemma can be deduced from the dense counting lemma in exactly the same way that Lemma 5 is derived from Lemma 4 in Section 5.

Lemma 7 (Dense extension lemma). *Let k, ℓ, t, t', n_0 be positive integers and let $\beta, d_2, \ldots, d_{k-1}, \delta$ be positive constants such that*

$$1/n_0 \ll \delta \ll \beta \ll d_2, \ldots, d_{k-1}, 1/\ell, 1/t'.$$

Then the following holds for all integers $n \geq n_0$. Suppose that $\mathcal{H}'$ is a $(k-1,\ell)$-complex on t' vertices with vertex classes $X_1, \ldots, X_\ell$ and let $\mathcal{H}$ be an induced subcomplex of $\mathcal{H}'$ on t vertices. Suppose also that $\mathcal{G}$ is a $(d_{k-1}, \ldots, d_2, \delta, \delta, 1)$-regular $(k-1,\ell)$-complex with vertex *classes $V_1, \dots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}'$. Then all but at most $\beta|\mathcal{H}|_\mathcal{G}$ labelled partition-respecting copies of $\mathcal{H}$ in $\mathcal{G}$ can be extended into*

$$(1 \pm \beta) n^{|\mathcal{H}'|-|\mathcal{H}|} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}')-e_i(\mathcal{H})}$$

labelled partition-respecting copies of $\mathcal{H}'$ in $\mathcal{G}$.

4. Deriving Lemmas 4 and 6

Throughout this section, and in the remainder of the paper, whenever we talk about a copy of a complex $\mathcal{H}$ in $\mathcal{G}$ we mean that this copy is labelled and partition-respecting, without mentioning this explicitly. First, we deduce Lemma 6 from [14, Cor. 6.11]. The difference between the two is that the latter result only counts complete hypergraphs but on the other hand it allows for different densities within each level. We need a few definitions that make this notion precise. Let $\mathcal{G}$ be a (k,t)-complex. Recall that $\mathcal{G}_i$ denotes the underlying i-uniform hypergraph of $\mathcal{G}$. For each $3 \le i < k$, we say that $\mathcal{G}_i$ *is $(\ge d_i, \delta_i)$-regular with respect to $\mathcal{G}_{i-1}$*, if for every i-tuple Λ_i of vertex classes of $\mathcal{G}$ the induced hypergraph $\mathcal{G}_i[\Lambda_i]$ is $(d_{\Lambda_i}, \delta_i)$-regular with respect to $\mathcal{G}_{i-1}[\Lambda_i]$, for some $d_{\Lambda_i} \ge d_i$. Similarly we define when $\mathcal{G}_k$ *is $(\ge d_k, \delta_k, r)$-regular with respect to $\mathcal{G}_{k-1}$* and when $\mathcal{G}_2$ *is $(\ge d_2, \delta_2)$-regular*. Let $\mathbf{d} := (d_k, \dots, d_2)$. We say that a (k,t)-complex $\mathcal{G}$ is *$(\ge \mathbf{d}, \delta_k, \delta, r)$-regular* if

- $\mathcal{G}_k$ is $(\ge d_k, \delta_k, r)$-regular with respect to $\mathcal{G}_{k-1}$;
- $\mathcal{G}_i$ is $(\ge d_i, \delta)$-regular with respect to $\mathcal{G}_{i-1}$ for each $3 \le i < k$;
- $\mathcal{G}_2$ is $(\ge d_2, \delta)$-regular.

Lemma 8 ([14])**.** *Let k, t, n_0 be positive integers and let $\varepsilon, d_2, \dots, d_{k-1}, \delta$ be positive constants such that*

$$1/n_0 \ll \delta \ll \varepsilon \ll d_2, \dots, d_{k-1}, 1/t.$$

Then the following holds for all integers $n \ge n_0$. Suppose that $\mathcal{G}$ is a $(\ge (d_{k-1}, \dots, d_2), \delta, \delta, 1)$-regular $(k-1,t)$-complex with vertex classes $V_1, \dots, V_t$, all of size n. Then

$$|K_t^{(k-1)}|_\mathcal{G} = (1 \pm \varepsilon) n^t \prod_{i=2}^{k-1} \prod_{\Lambda_i} d_{\Lambda_i},$$

where the second product is taken over all i-tuples Λ_i of vertex classes of $\mathcal{G}$.

We now show how to deduce Lemma 6 from this. Full details can be found in [4].

Proof of Lemma 6. First we prove the lemma for the case when $\ell = t$, i.e. when each of the vertex classes $X_1, \dots, X_t$ of $\mathcal{H}$ consists of exactly one vertex, say $X_i := \{h_i\}$. Given such an $\mathcal{H}$ and a complex $\mathcal{G}$ as in Lemma 6, we construct a complex $\mathcal{G}'$ from $\mathcal{G}$ as follows: Starting with $i = 2$, for all i with $2 \le i \le k$ in turn, we successively consider each i-tuple $\Lambda_i = (V_{j_1}, \dots, V_{j_i})$ of vertex classes of $\mathcal{G}$. If $h_{j_1}, \dots, h_{j_i}$ forms an i-edge of $\mathcal{H}$ we let $\mathcal{G}'_i[\Lambda_i] = \mathcal{G}_i[\Lambda_i]$. If $h_{j_1}, \dots, h_{j_i}$ does not from an i-edge we make each copy of $K_i^{(i-1)}$ in $\mathcal{G}'_{i-1}[\Lambda_i]$ into an i-edge of $\mathcal{G}'_i$. Thus in the latter case the density of $\mathcal{G}'_i[\Lambda_i]$ with respect to $\mathcal{G}'_{i-1}[\Lambda_i]$ will be 1. (If $i = 2$, this means that we let $\mathcal{G}'_i[\Lambda_i]$ be the complete bipartite graph with vertex classes V_{j_1} and V_{j_2}.) Using that $\mathcal{H}$ is a complex, it is easy to see that $\mathcal{G}'$ is also $(\ge (d_{k-1}, \dots, d_2), \delta, \delta, 1)$-regular. Clearly, there is a bijection between the copies of $\mathcal{H}$ in $\mathcal{G}$ and the copies of $K_t^{(k)}$ in $\mathcal{G}'$. So $|\mathcal{H}|_\mathcal{G} = |K_t^{(k)}|_{\mathcal{G}'}$. The result now follows if we apply Lemma 8 to $\mathcal{G}'$.

It now remains to deduce Lemma 6 for arbitrary ℓ-partite complexes $\mathcal{H}$ from the result for the above case. For this, we use a simple argument that was also used in [5] to obtain Lemma 4 in the case $k = 3$. We define a complex $\mathcal{G}^*$ from $\mathcal{G}$ by making $|X_i|$ copies

$V_i^1,\dots,V_i^{|X_i|}$ of each vertex class V_i in such a way that for any selection of indices $i_1,\dots,i_t$ the complex $\mathcal{G}^*[V_1^{i_1},\dots,V_t^{i_t}]$ is isomorphic to $\mathcal{G}$. Note that $\mathcal{G}^*$ is $|\mathcal{H}|$-partite. Also, we can turn $\mathcal{H}$ into an $|\mathcal{H}|$-partite complex $\mathcal{H}^*$ by viewing each vertex as a single vertex class. Note that different copies of $\mathcal{H}$ in $\mathcal{G}$ give rise to different copies of $\mathcal{H}^*$ in $\mathcal{G}^*$. Thus $|\mathcal{H}|_{\mathcal{G}} \le |\mathcal{H}^*|_{\mathcal{G}^*}$. Conversely, the only case where a copy of $\mathcal{H}^*$ in $\mathcal{G}^*$ does not correspond to a copy of $\mathcal{H}$ in $\mathcal{G}$ is when there is some i and indices $j_1 \neq j_2$ such that the vertices that are used by $\mathcal{H}^*$ in $V_i^{j_1}$ and $V_i^{j_2}$ correspond to the same vertex of V_i. It is easy to see that the number of such copies is comparatively small. Thus the desired bounds on $|\mathcal{H}|_{\mathcal{G}}$ immediately follow from the bounds on $|\mathcal{H}^*|_{\mathcal{G}^*}$ which we obtained in the previous paragraph. □

We now prove Lemma 4. Its proof is based on the following version of the counting lemma that accompanies the hypergraph regularity Lemma from [23]. We will state and apply this regularity lemma in Sections 7 and 8. Lemma 9 gives a lower bound on the number of complete complexes $K_t^{(k)}$ in a regular (k,t)-complex $\mathcal{G}$, under less restrictive assumptions on the regularity constants than those in Lemma 8.

Lemma 9 (Counting lemma for complete complexes [23])**.** *Let k,r,t,n_0 be positive integers and let $\varepsilon,d_2,\dots,d_k,\delta,\delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$ for $i=2,\dots,k-1$ and*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2,\dots,d_{k-1}\} \le \delta_k \ll \varepsilon, d_k, 1/t.$$

Then the following holds for all integers $n \ge n_0$. Suppose that $\mathcal{G}$ is a $(\mathbf{d},\delta_k,\delta,r)$-regular (k,t)-complex with vertex classes $V_1,\dots,V_t$, all of size n, which respects the partition of $K_t^{(k)}$. Then

$$|K_t^{(k)}|_{\mathcal{G}} \ge (1-\varepsilon)n^t \prod_{i=2}^{k} d_i^{\binom{k}{i}}.$$

Lemma 4 is more general in the sense that it counts copies of complexes that may not be complete, and also gives an upper bound on their number. We will deduce Lemma 4 from Lemma 9 in several steps. The first (and main) step is to deduce a counting lemma which still gives just a lower bound on the number of copies of complete complexes, but now in a (k,t)-complex $\mathcal{G}$ where the density of $\mathcal{G}_i[\Lambda_i]$ with respect to $\mathcal{G}_{i-1}[\Lambda_i]$ might be different for different i-tuples Λ_i of vertex classes of $\mathcal{G}$.

Lemma 10 (Counting lemma for complete complexes – different densities)**.** *Let k,r,t,n_0 be positive integers and let $\varepsilon,d_2,\dots,d_k,\delta,\delta_k$ be positive constants such that*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2,\dots,d_{k-1}\} \le \delta_k \ll \varepsilon, d_k, 1/t.$$

Then the following holds for all integers $n \ge n_0$. Suppose $\mathcal{G}$ is a $(\ge \mathbf{d},\delta_k,\delta,r)$-regular (k,t)-complex with vertex classes $V_1,\dots,V_t$, all of size n, such that for all $2 < i < k$ and all i-tuples Λ_i of vertex classes of $\mathcal{G}$ the hypergraph $\mathcal{G}_i[\Lambda_i]$ is (d_{Λ_i},δ)-regular with respect to $\mathcal{G}_{i-1}[\Lambda_i]$ where d_{Λ_i} can be written as $d_{\Lambda_i} = p_{\Lambda_i}/q_{\Lambda_i}$ such that $p_{\Lambda_i}, q_{\Lambda_i} \in \mathbb{N}$ and $1/q_{\Lambda_i} \ge d_i$. Suppose that the analogue holds for all the d_{Λ_2} and all the d_{Λ_k}. Then

$$|K_t^{(k)}|_{\mathcal{G}} \ge (1-\varepsilon)n^t \prod_{i=2}^{k} \prod_{\Lambda_i} d_{\Lambda_i},$$

where the second product is taken over all i-tuples Λ_i of vertex classes of $\mathcal{G}$.

Proof. We will prove this lemma by an inductive argument, in which we allow for different densities in the top levels but not in the lower levels, and show that we can always move down another level, until we allow different densities in all levels. This leads to the following definition. For any $2 < j \le k$, we say that a complex $\mathcal{G}$ is $(\ge d_k,\dots,\ge d_j, d_{j-1},\dots,d_2,\delta_k,\delta,r)$-*regular* if

- $\mathcal{G}_k$ is $(\geq d_k, \delta_k, r)$-regular with respect to $\mathcal{G}_{k-1}$;
- $\mathcal{G}_i$ is $(\geq d_i, \delta)$-regular with respect to $\mathcal{G}_{i-1}$ for each $j \leq i \leq k-1$;
- $\mathcal{G}_i$ is (d_i, δ)-regular with respect to $\mathcal{G}_{i-1}$ for each $3 \leq i \leq j-1$;
- $\mathcal{G}_2$ is (d_2, δ)-regular.

Choose new constants $\eta_i, \xi_i, \varepsilon_i$ and integers r_i satisfying

$$1/n_0 \ll \delta = \xi_2 \ll \cdots \ll \xi_k \ll \min\{\delta_k, d_2, \dots, d_{k-1}\} \leq \delta_k = \eta_2 \ll \cdots \ll \eta_{k+1} \\ \ll \varepsilon_k \ll \cdots \ll \varepsilon_2 = \varepsilon, d_k, 1/t$$

and $1/n_0 \ll 1/r = 1/r_2 \ll \cdots \ll 1/r_k \ll \min\{\delta_k, d_2, \dots, d_{k-1}\}$. Then the following claim immediately implies the lemma:

Claim. *Let $2 \leq j \leq k$. Suppose that $\mathcal{G}$ satisfies the conditions of Lemma 10 but is $(\geq d_k, \dots, \geq d_j, d_{j-1}, \dots, d_2, \eta_j, \xi_j, r_j)$-regular instead of $(\geq \mathbf{d}, \delta_k, \delta, r)$-regular if $j > 2$, where $1/d_i \in \mathbb{N}$ for all $i = 2, \dots, j-1$. Then*

$$|K_t^{(k)}|_{\mathcal{G}} \geq (1-\varepsilon_j) n^t \left(\prod_{i=2}^{j-1} d_i^{\binom{t}{i}} \right) \prod_{i=j}^{k} \prod_{\Lambda_i} d_{\Lambda_i}.$$

We prove this claim by backward induction on j as follows: given a t-partite complex $\mathcal{G}$ which is $(\geq d_k, \dots, \geq d_j, d_{j-1}, \dots, d_2, \eta_j, \xi_j, r_j)$-regular, we will partition the hyperedges of $\mathcal{G}_j$ to obtain several $(\geq d_k, \dots, \geq d_{j+1}, d'_j, d_{j-1}, \dots, d_2, \eta_{j+1}, \xi_{j+1}, r_{j+1})$-regular complexes for some d'_j. We will then apply the lower bound from the induction hypothesis to each of these complexes. Summing over all of them will give the lower bound in the claim.

We first consider the case $j = k$. We will apply the slicing lemma [23, Prop. 33] to split the kth level $\mathcal{G}_k$ of the complex $\mathcal{G}$ to obtain regular complexes whose densities within the kth level are the same. (The slicing lemma itself can be proved using a simple application of a Chernoff bound.) Set $d'_k := 1/\prod_{\Lambda_k} q_{\Lambda_k}$. The slicing lemma implies that for all Λ_k there is a partition $P(\Lambda_k)$ of the set $E(\mathcal{G}_k[\Lambda_k])$ of k-edges induced on Λ_k such that each part is (d'_k, η_{k+1}, r_k)-regular with respect to $\mathcal{G}_{k-1}[\Lambda_k]$. So for each Λ_k, $P(\Lambda_k)$ has d_{Λ_k}/d'_k parts. Now for each Λ_k, choose one part from $P(\Lambda_k)$ and let $\mathcal{C}_k$ denote the resulting k-uniform t-partite hypergraph. Let $\mathcal{G}^{\mathcal{C}_k}$ denote the k-complex obtained from $\mathcal{G}$ by replacing $\mathcal{G}_k$ with $\mathcal{C}_k$. Then

$$|K_t^{(k)}|_{\mathcal{G}} = \sum_{\mathcal{C}_k} |K_t^{(k)}|_{\mathcal{G}^{\mathcal{C}_k}}.$$

Here the summation is over all possible choices of parts from each of the $\binom{t}{k}$ partitions $P(\Lambda_k)$. So the number of summands is $\prod_{\Lambda_k} d_{\Lambda_k}/d'_k = {d'}_k^{-\binom{t}{k}} \prod_{\Lambda_k} d_{\Lambda_k}$. Moreover, by Lemma 9 each summand in the above sum can be bounded below:

$$|K_t^{(k)}|_{\mathcal{G}^{\mathcal{C}_k}} \geq (1-\varepsilon_k) n^t \left(\prod_{i=2}^{k-1} d_i^{\binom{t}{i}} \right) {d'}_k^{\binom{t}{k}}.$$

Altogether, this implies the claim for $j = k$.

Now suppose that $j < k$ and that the claim holds for $j+1$. To apply the induction hypothesis, we now need to get equal densities in the jth level. We will achieve this by applying the slicing lemma to this level. Set $d'_j := 1/\prod_{\Lambda_j} q_{\Lambda_j}$. So $1/d'_j \in \mathbb{N}$. The slicing lemma implies that for every j-tuple Λ_j of vertex classes of $\mathcal{G}$ there is a partition $P(\Lambda_j)$ of the set $E(\mathcal{G}_j[\Lambda_j])$ of j-edges induced on Λ_j such that each part is (d'_j, ξ_{j+1})-regular with respect to $\mathcal{G}_{j-1}[\Lambda_j]$. For each Λ_j, the corresponding partition $P(\Lambda_j)$ will have $a_{\Lambda_j} := d_{\Lambda_j}/d'_j$ parts. Now for each Λ_j, choose one part from $P(\Lambda_j)$ and let $\mathcal{C}_j$ denote the resulting j-uniform t-partite hypergraph. We let $\mathcal{G}^{\mathcal{C}_j}$ denote the (k,t)-complex obtained from $\mathcal{G}$ as follows: we replace $\mathcal{G}_j$ by $\mathcal{C}_j$ and for each $j < i \le k$ we replace $\mathcal{G}_i$ with the subhypergraph whose i-edges span a $K_i^{(j)}$ in $\mathcal{C}_j$. Thus $\mathcal{G}_j^{\mathcal{C}_j}$ is (d'_j, ξ_{j+1})-regular with respect to $\mathcal{G}_{j-1} = \mathcal{G}_{j-1}^{\mathcal{C}_j}$. However, to apply the induction hypothesis this is not enough. We also need to prove the following more general assertion.

For all $i = j, \dots, k$ and any Λ_i the following holds. If $i = j$ then $\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i]$ is (d'_j, ξ_{j+1})-regular with respect to $\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]$. If $j < i < k$ then $\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i]$ is $(d_{\Lambda_i}, \xi_{j+1})$-regular with respect to $\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]$. If $i = k$ then $\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i]$ is $(d_{\Lambda_i}, \eta_{j+1}, r_{j+1})$-regular with respect to $\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]$ for all but at most $\sqrt{\eta_{j+1}}\prod_{\Lambda_j} a_{\Lambda_j}$ hypergraphs $\mathcal{C}_j$. $(*)$

We will prove $(*)$ by induction on i. If $i = j$ then we already know that the assertion is true. So suppose that $i > j$ and that the claim holds for $i-1$. We will first consider the case when $i < k$. The induction hypothesis together with the Dense Counting Lemma for complete complexes (Lemma 8) implies that

$$|K_i^{(i-1)}|_{\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]} \ge \frac{1}{2} n^i \left(\prod_{\ell=2}^{j-1} d_\ell^{\binom{i}{\ell}} \right) {d'_j}^{\binom{i}{j}} \prod_{s=j+1}^{i-1} \prod_{\Lambda_s \subseteq \Lambda_i} d_{\Lambda_s}. \tag{1}$$

Similarly, the assumptions on $\mathcal{G}$ in the claim together with Lemma 8 imply

$$|K_i^{(i-1)}|_{\mathcal{G}_{i-1}[\Lambda_i]} \le 2n^i \left(\prod_{\ell=2}^{j-1} d_\ell^{\binom{i}{\ell}} \right) \prod_{s=j}^{i-1} \prod_{\Lambda_s \subseteq \Lambda_i} d_{\Lambda_s}. \tag{2}$$

If we combine these inequalities and use the fact that $\xi_j \ll \xi_{j+1} \ll d_j, 1/k$, we obtain

$$|K_i^{(i-1)}|_{\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]} \ge \sqrt{\xi_{j+1}} |K_i^{(i-1)}|_{\mathcal{G}_{i-1}[\Lambda_i]} \ge \frac{\xi_j}{\xi_{j+1}} |K_i^{(i-1)}|_{\mathcal{G}_{i-1}[\Lambda_i]}. \tag{3}$$

In other words, a ξ_{j+1}-proportion of copies of $K_i^{(i-1)}$ in $\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]$ gives rise to a ξ_j-proportion of copies in $\mathcal{G}_{i-1}[\Lambda_i]$. Moreover, $\mathcal{K}_i(\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]) \cap E(\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i]) = \mathcal{K}_i(\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]) \cap E(\mathcal{G}_i[\Lambda_i])$ by the definition of $\mathcal{G}^{\mathcal{C}_j}$ and so $d(\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i] | \mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]) = d(\mathcal{G}_i[\Lambda_i] | \mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]) = d_{\Lambda_i} \pm \xi_j$ by (3) and the (d_{Λ_i}, ξ_j)-regularity of $\mathcal{G}_i[\Lambda_i]$ with respect to $\mathcal{G}_{i-1}[\Lambda_i]$. Thus the $(d_{\Lambda_i}, \xi_{j+1})$-regularity of $\mathcal{G}_i^{\mathcal{C}_j}[\Lambda_i]$ with respect to $\mathcal{G}_{i-1}^{\mathcal{C}_j}[\Lambda_i]$ follows from the (d_{Λ_i}, ξ_j)-regularity of $\mathcal{G}_i[\Lambda_i]$ with respect to $\mathcal{G}_{i-1}[\Lambda_i]$.

But if $i = k$, this might not be true, as η_{j+1} may not be small compared to d_j. However, given a k-tuple Λ_k of vertex classes of $\mathcal{G}$, it is true for most complexes $\mathcal{G}^{\mathcal{C}_j}[\Lambda_k]$. To see this, given Λ_k, let $\mathcal{B}$ be a (k,k)-complex obtained as follows: For each $\Lambda_j \subset \Lambda_k$, choose one part from $P(\Lambda_j)$ and let $\mathcal{B}_j$ denote the resulting j-uniform k-partite hypergraph. To obtain $\mathcal{B}$ from $\mathcal{G}[\Lambda_k]$, we replace $\mathcal{G}_j[\Lambda_k]$ by $\mathcal{B}_j$ and for each $j < i \le k$ we replace $\mathcal{G}_i[\Lambda_k]$ with the subhypergraph whose i-edges span a $K_i^{(j)}$ in $\mathcal{B}_j$. Thus there are $\prod_{\Lambda_j \subset \Lambda_k} a_{\Lambda_j} =: A_{\Lambda_k}$ such complexes $\mathcal{B}$. (Recall that $a_{\Lambda_j} = d_{\Lambda_j}/d'_j$ was the number of parts of the partition $P(\Lambda_j)$.) Using that $(*)$ holds for all $i < k$, similarly as in (1)–(3) one can show that

$$|K_k^{(k-1)}|_{\mathcal{B}_{k-1}} \ge \frac{{d'_j}^{\binom{k}{j}}}{4\prod_{\Lambda_j \subset \Lambda_k} d_{\Lambda_j}} |K_k^{(k-1)}|_{\mathcal{G}_{k-1}[\Lambda_k]} = \frac{|K_k^{(k-1)}|_{\mathcal{G}_{k-1}[\Lambda_k]}}{4A_{\Lambda_k}}. \tag{4}$$

We will now prove the following:

The underlying k-uniform hypergraph $\mathcal{B}_k$ is not $(d_{\Lambda_k}, \eta_{j+1}, r_{j+1})$-regular with respect to $\mathcal{B}_{k-1}$ for less than $\eta_{j+1} A_{\Lambda_k}$ of the complexes $\mathcal{B}$. $(**)$

If $(**)$ is false then we can find $T := \eta_{j+1}A_{\Lambda_k}/2$ such complexes $\mathcal{B}^1, \dots, \mathcal{B}^T$, such that each $\mathcal{B}^\ell$ has a $\mathbf{Q}^\ell = (Q_1^\ell, \dots, Q_{r_{j+1}}^\ell)$ satisfying $Q_s^\ell \subseteq \mathcal{B}_{k-1}^\ell$ for all $s = 1, \dots, r_{j+1}$ and $|K_k^{(k-1)}|_{\mathbf{Q}^\ell} \geq \eta_{j+1}|K_k^{(k-1)}|_{\mathcal{B}_{k-1}^\ell}$, but either $d(\mathcal{B}_k^\ell|\mathbf{Q}^\ell) > d_{\Lambda_k} + \eta_{j+1}$ for each ℓ or $d(\mathcal{B}_k^\ell|\mathbf{Q}^\ell) < d_{\Lambda_k} - \eta_{j+1}$ for each ℓ. We will assume the latter – the proof in the former case is similar. But then let $\mathbf{Q} = (\mathbf{Q}^1, \mathbf{Q}^2, \dots, \mathbf{Q}^T)$. Thus $\mathbf{Q}$ is a Tr_{j+1}-tuple and

$$|K_k^{(k-1)}|_{\mathbf{Q}} \geq \sum_{\ell=1}^{T} \eta_{j+1}|K_k^{(k-1)}|_{\mathcal{B}_{k-1}^\ell} \overset{(4)}{\geq} \eta_j |K_k^{(k-1)}|_{\mathcal{G}_{k-1}[\Lambda_k]}.$$

Since we may assume that $Tr_{j+1} \leq r_j$ our assumption on the regularity of $\mathcal{G}_k[\Lambda_k]$ with respect to $\mathcal{G}_{k-1}[\Lambda_k]$ implies that $d(\mathcal{G}_k[\Lambda_k]|\mathbf{Q}) \geq d_{\Lambda_k} - \eta_j$. On the other hand, the definition of $\mathcal{B}$ implies that $d(\mathcal{B}_k^\ell|\mathbf{Q}^\ell) = d(\mathcal{G}_k[\Lambda_k]|\mathbf{Q}^\ell)$. Thus $d(\mathcal{G}_k[\Lambda_k]|\mathbf{Q}) \leq \max_{1\leq \ell \leq T} d(\mathcal{G}_k[\Lambda_k]|\mathbf{Q}^\ell) = \max_{1\leq\ell\leq T} d(\mathcal{B}_k^\ell|\mathbf{Q}^\ell) < d_{\Lambda_k} - \eta_{j+1}$. This is a contradiction, and so $(**)$ holds.

Note that $(**)$ implies that for all but at most $\binom{t}{k}\eta_{j+1}\prod_{\Lambda_j} a_{\Lambda_j}$ hypergraphs $\mathcal{C}_j$ the hypergraph $\mathcal{G}_k^{\mathcal{C}_j}$ is $(d_{\Lambda_k}, \eta_{j+1}, r_{j+1})$-regular with respect to $\mathcal{G}_{k-1}^{\mathcal{C}_j}$ – we call these $\mathcal{C}_j$ *nice*. Since $\eta_{j+1} \ll 1/t$, this completes the proof of $(*)$.

We are now ready to finish the proof of the induction step of the claim. The induction

$$|K_t^{(k)}|_{\mathcal{G}} \quad \geq \quad \sum_{\text{nice } \mathcal{C}_j} |K_t^{(k)}|_{\mathcal{G}^{\mathcal{C}_j}} \geq (1-\varepsilon_{j+1}) \sum_{\text{nice } \mathcal{C}_j} n^t \left(\prod_{i=2}^{j-1} d_i^{\binom{t}{i}}\right) {d'}_j^{\binom{t}{j}} \prod_{i=j+1}^{k} \prod_{\Lambda_i} d_{\Lambda_i}.$$

The summation is over all possible choices of nice $\mathcal{C}_j$. So the number of summands is at least $(1-\sqrt{\eta_{j+1}})\prod_{\Lambda_j} a_{\Lambda_j}$ and for each Λ_j we have $a_{\Lambda_j}d'_j = d_{\Lambda_j}$. Since $\eta_{j+1}, \varepsilon_{j+1} \ll \varepsilon_j$, the claim follows and hence Lemma 10 as well. □

Instead of making use of the parameter r in the proof of $(*)$ we could have also used the fact that the partitions guaranteed by the slicing lemma are obtained by considering random partitions.

It is straightforward to obtain a corresponding upper bound from the lower bound in Lemma 10.

Lemma 11 (Counting lemma for complete complexes – upper bound)**.** *Under the conditions of Lemma 10,*

$$|K_t^{(k)}|_{\mathcal{G}} = (1 \pm \varepsilon) n^t \prod_{i=2}^{k} \prod_{\Lambda_i} d_{\Lambda_i}.$$

Proof. Clearly, all we have to prove is the upper bound. The proof is based on an argument that was used in [20] and later in [5] to derive a similar upper bound in the case of 3-complexes and thus we only give a sketch of it. A detailed proof can be found in [4]. Let $[t]^k$ denote the set of all k-subsets of $[t] = \{1, \dots, t\}$. Given $S \subseteq [t]^k$, we let $\mathcal{G}^S$ denote the (k,t)-complex obtained from $\mathcal{G}$ as follows: for each $\{i_1, \dots, i_k\} \in S$ we replace the set $E_k(\mathcal{G}[\Lambda_k])$ of all k-edges of $\mathcal{G}$ induced on $\Lambda_k := \{V_{i_1}, \dots, V_{i_k}\}$ by $\mathcal{K}_k(\mathcal{G}_{k-1}[\Lambda_k]) \setminus E_k(\mathcal{G}[\Lambda_k])$. Thus the density of $\mathcal{G}_k^S[\Lambda_k]$ with respect to $\mathcal{G}_{k-1}^S[\Lambda_k]$ is now $1 - d_{\Lambda_k}$. Moreover,

$$|K_t^{(k-1)}|_{\mathcal{G}_{k-1}} = \sum_{S \subseteq [t]^k} |K_t^{(k)}|_{\mathcal{G}^S}.$$

Observe that $|K_t^{(k)}|_{\mathcal{G}} = |K_t^{(k)}|_{\mathcal{G}^\emptyset}$ and hence

$$|K_t^{(k)}|_{\mathcal{G}} = |K_t^{(k-1)}|_{\mathcal{G}_{k-1}} - \sum_{S\subseteq [t]^k, S \neq \emptyset} |K_t^{(k)}|_{\mathcal{G}^S}.$$

Thus, to obtain an upper bound on $|K_t^{(k)}|_{\mathcal{G}}$ all we have to do now is to obtain an upper bound on $|K_t^{(k-1)}|_{\mathcal{G}_{k-1}}$ and a lower bound on $|K_t^{(k)}|_{\mathcal{G}^S}$, for every non-empty S. But the former follows from the dense counting lemma (Lemma 6) and the latter follows from Lemma 10 above. (This is why in Lemma 10 we need to allow more general densities than just $1/a$, for $a \in \mathbb{N}$.) □

Lemma 4 now follows from Lemma 11 in exactly the same way as Lemma 6 followed from Lemma 8.

5. Proof of the Extension Lemma

The proof idea is similar to that of [23, Cor. 26], [10, Lemma 6.6] and [5, Lemma 5]. We first introduce some more notation. Given k-complexes $\mathcal{H}' \subseteq \mathcal{H}''$ such that $\mathcal{H}'$ is induced, and a copy H' of $\mathcal{H}'$ in $\mathcal{G}$, we define $|H' \to \mathcal{H}''|$ to be the number of ways in which H' can be extended to a copy of $\mathcal{H}''$ in $\mathcal{G}$. We also define

$$\overline{|\mathcal{H}' \to \mathcal{H}''|} := n^{|\mathcal{H}''|-|\mathcal{H}'|} \prod_{i=2}^{k} d_i^{e_i(\mathcal{H}'')-e_i(\mathcal{H}')}.$$

Thus $\overline{|\mathcal{H}' \to \mathcal{H}''|}$ is roughly the expected number of ways H' could be extended to a copy of $\mathcal{H}''$ if $\mathcal{G}$ were a random complex.

To prove the extension lemma, pick a copy H of $\mathcal{H}$ in $\mathcal{G}$ uniformly at random, and define $X := |H \to \mathcal{H}'|$. Then X is a random variable. We have $\mathbb{E}(X) = \frac{1}{|\mathcal{H}|_{\mathcal{G}}} \sum_{H \in \mathcal{G}} |H \to \mathcal{H}'| = |\mathcal{H}'|_{\mathcal{G}}/|\mathcal{H}|_{\mathcal{G}}$. (Here the sum $\sum_{H \in \mathcal{G}}$ is over all copies of $\mathcal{H}$ in $\mathcal{G}$.) We pick some constant ε satisfying $\delta_k \ll \varepsilon \ll \beta$. By applying the upper bound of the counting lemma to $\mathcal{H}$ and the lower bound to $\mathcal{H}'$ we obtain a lower bound for $\mathbb{E}(X)$. Similarly we obtain an upper bound. In this way we can easily deduce that

$$\mathbb{E}(X) = (1 \pm \sqrt{\varepsilon})\overline{|\mathcal{H} \to \mathcal{H}'|}. \tag{5}$$

Now consider $\mathbb{E}(X^2)$. We aim to show that its value is approximately $\overline{|\mathcal{H} \to \mathcal{H}'|}^2$, and so X has a low variance. Using Chebyshev's inequality, this will then imply that X is concentrated around its mean. In other words, only a few copies of $\mathcal{H}$ do not extend to the correct number of copies of $\mathcal{H}'$ in $\mathcal{G}$.

Observe that $\mathbb{E}(X^2) = \frac{1}{|\mathcal{H}|_{\mathcal{G}}} \sum_{H \in \mathcal{G}} |H \to \mathcal{H}'|^2$. We view $|H \to \mathcal{H}'|^2$ as the number of pairs H_1', H_2' of copies of $\mathcal{H}'$ which extend H. Here the pairs are allowed to overlap, but we first obtain a rough estimate by insisting that they intersect precisely in H. So let $\mathcal{H}^*$ be the (k,ℓ)-complex obtained from two disjoint copies of $\mathcal{H}'$ by identifying them on $\mathcal{H}$. Thus any copy of $\mathcal{H}^*$ in $\mathcal{G}$ extending H corresponds to a pair H_1', H_2'. However, we will later need to take account of those pairs H_1', H_2' which do not arise from a copy of $\mathcal{H}^*$. These pairs are exactly those whose intersection is strictly larger than H.

By applying the counting lemma to $\mathcal{H}^*$ and to $\mathcal{H}$, as before we obtain

$$\frac{1}{|\mathcal{H}|_{\mathcal{G}}} \sum_{H \in \mathcal{G}} |H \to \mathcal{H}^*| = (1 \pm \sqrt{\varepsilon})\overline{|\mathcal{H} \to \mathcal{H}^*|} = (1 \pm \sqrt{\varepsilon})\overline{|\mathcal{H} \to \mathcal{H}'|}^2.$$

On the other hand, the number of pairs H_1', H_2' which do not arise from a copy of $\mathcal{H}^*$ is at most $(t'-t)^2 n^{2(t'-t)-1} < \varepsilon((\prod_{i=2}^k d_i^{e_i(\mathcal{H}')-e_i(\mathcal{H})})n^{t'-t})^2 = \varepsilon\overline{|\mathcal{H} \to \mathcal{H}'|}^2$. Thus

$$\frac{1}{|\mathcal{H}|_{\mathcal{G}}} \sum_{H \in \mathcal{G}} |H \to \mathcal{H}'|^2 = (1 \pm 2\sqrt{\varepsilon})\overline{|\mathcal{H} \to \mathcal{H}'|}^2. \tag{6}$$

Putting (5) and (6) together, we obtain

$$var(X) = \mathbb{E}(X^2) - (\mathbb{E}(X))^2 < 5\sqrt{\varepsilon}\,\overline{|\mathcal{H} \to \mathcal{H}'|}^2.$$

Now recall Chebyshev's inequality: $\mathbb{P}(|X - \mathbb{E}(X)| \geq t) \leq var(X)/t^2$. We apply this inequality with $t := \beta\overline{|\mathcal{H} \to \mathcal{H}'|}$. This implies that the probability that a randomly chosen copy of $\mathcal{H}$ in $\mathcal{G}$ does not satisfy the conclusion of the extension lemma is at most $var(X)/\beta^2\overline{|\mathcal{H} \to \mathcal{H}'|}^2 < 5\sqrt{\varepsilon}/\beta^2 < \beta$, and so at most $\beta|\mathcal{H}|_{\mathcal{G}}$ copies of $\mathcal{H}$ do not satisfy the conclusion, as required.

6. Proof of the embedding lemma for complexes

We prove Lemma 3 by induction on $|\mathcal{H}|$. We first suppose that the connected component of $\mathcal{H}$ which contains the vertex h has order less than Δ^4. In this case we will use the counting lemma to prove the embedding lemma. So let $\mathcal{C}$ be the component of $\mathcal{H}$ containing h, and let $\mathcal{D} := \mathcal{H} - \mathcal{C}$. Also, let $\mathcal{C}_h := \mathcal{C} - h$. Note that a copy of $\mathcal{H}$ consists of disjoint copies of $\mathcal{C}$ and $\mathcal{D}$, while $\mathcal{H}_h$ consists of disjoint copies of $\mathcal{C}_h$ and $\mathcal{D}$. Copies of these complexes in $\mathcal{G}$ will be denoted by C, D and C_h.

Choose a new constant β such that $c, \delta_k \ll \beta \ll \alpha$. Now note that $|\mathcal{H}|_{\mathcal{G}} = \sum_{D \in \mathcal{G}} |\mathcal{C}|_{\mathcal{G}-D}$, and by applying the upper and lower bounds of the counting lemma to copies of $\mathcal{C}$ in $\mathcal{G}$ and $\mathcal{G} - D$ respectively, we obtain $|\mathcal{C}|_{\mathcal{G}-D} \geq \frac{(1-c)^{\Delta^4}(1-\beta)}{(1+\beta)}|\mathcal{C}|_{\mathcal{G}} \geq (1-3\beta)|\mathcal{C}|_{\mathcal{G}}$. So

$$|\mathcal{H}|_{\mathcal{G}} \geq \sum_{D \in \mathcal{G}} (1-3\beta)|\mathcal{C}|_{\mathcal{G}} = (1-3\beta)|\mathcal{C}|_{\mathcal{G}}|\mathcal{D}|_{\mathcal{G}}. \tag{7}$$

On the other hand, by a similar argument using the upper and lower bounds from the counting lemma in $\mathcal{G}$ for $\mathcal{C}_h$ and $\mathcal{C}$ respectively,

$$|\mathcal{H}_h|_{\mathcal{G}} \leq |\mathcal{C}_h|_{\mathcal{G}}|\mathcal{D}|_{\mathcal{G}} \leq \frac{1+\beta}{1-\beta} \frac{|\mathcal{C}|_{\mathcal{G}}|\mathcal{D}|_{\mathcal{G}}}{n\prod_{i=2}^{k} d_i^{e_i(\mathcal{C})-e_i(\mathcal{C}_h)}}. \tag{8}$$

Combining (7) and (8) gives the desired result.

Thus we may assume that the component of $\mathcal{H}$ containing h has order at least Δ^4. This deals with the base case of the inductive argument, and it also means that the third neighbourhood of h in $\mathcal{H}$ will be non-empty, which will be convenient later on in the proof.

We pick new constants ε_k and ε_{k-1} satisfying the following hierarchies:

$$\delta \ll \varepsilon_{k-1} \ll d_2, d_3, \ldots, d_k, 1/\Delta,$$
$$c, \delta_k, \varepsilon_{k-1} \ll \varepsilon_k \ll \alpha.$$

Let $\mathcal{N}_h$ be the subcomplex of $\mathcal{H}$ induced by the neighbours of h, and let $\mathcal{B}$ be the subcomplex of $\mathcal{H}$ induced by h and the neighbours of h. Then any copy of $\mathcal{H}$ in $\mathcal{G}$ extending a copy N_h of $\mathcal{N}_h$ can be obtained by first extending N_h into a copy of $\mathcal{H}_h$ and then extending N_h into a copy of $\mathcal{B}$, where the vertex chosen for h has to be distinct from all the vertices chosen for $\mathcal{H}_h$.

We now define a copy N_h of $\mathcal{N}_h$ to be *typical* if it has about the correct number of extensions into $\mathcal{B}$, i.e. if $|N_h \to \mathcal{B}| = (1 \pm \varepsilon_k)\overline{|\mathcal{N}_h \to \mathcal{B}|}$. An application of the extension lemma shows that at most $\varepsilon_k|\mathcal{N}_h|_{\mathcal{G}}$ copies of $\mathcal{N}_h$ in $\mathcal{G}$ are not typical. We denote the set of typical copies of $\mathcal{N}_h$ by `typ`, and the set of all atypical copies by `atyp`.

Now observe that if all of the copies of $\mathcal{N}_h$ were typical, the proof would be complete, since then

$$\begin{aligned}|\mathcal{H}|_{\mathcal{G}} &\geq \sum_{N_h \in \mathcal{G}} |N_h \to \mathcal{H}_h|(|N_h \to \mathcal{B}| - cn) \geq \Big((1-\varepsilon_k)\overline{|\mathcal{N}_h \to \mathcal{B}|} - cn\Big) \sum_{N_h} |N_h \to \mathcal{H}_h| \\ &\geq (1-\alpha)\overline{|\mathcal{N}_h \to \mathcal{B}|}|\mathcal{H}_h|_{\mathcal{G}} = (1-\alpha)\overline{|\mathcal{H}_h \to \mathcal{H}|}|\mathcal{H}_h|_{\mathcal{G}}.\end{aligned}$$

The third inequality follows since $c \ll \alpha, d_2, \dots, d_k$, and $\varepsilon_k \ll \alpha$.

However, we also need to take account of the atypical copies of $\mathcal{N}_h$. The proportion of these is about ε_k, which may be larger than some d_i. It will turn out that this is too large for our purposes, and so we will need to consider the atypical copies more carefully.

Before we do this, we define, instead of $|H' \to \mathcal{H}''|$, the expression $|H' \stackrel{k-1}{\to} \mathcal{H}''|$, where $\mathcal{H}' \subseteq \mathcal{H}''$ are induced subcomplexes of $\mathcal{H}$ and H' is a copy of $\mathcal{H}'$ in $\mathcal{G}$. We consider the underlying $(k-1)$-complexes in each case, and define $|H' \stackrel{k-1}{\to} \mathcal{H}''|$ to be the number of ways in which the underlying $(k-1)$-complex of H' can be extended to the underlying $(k-1)$-complex of $\mathcal{H}''$ within (the underlying $(k-1)$-complex of) $\mathcal{G}$. Clearly $|H' \stackrel{k-1}{\to} \mathcal{H}''| \geq |H' \to \mathcal{H}''|$. We also define

$$\overline{|\mathcal{H}' \stackrel{k-1}{\to} \mathcal{H}''|} := n^{|\mathcal{H}''|-|\mathcal{H}'|} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}'')-e_i(\mathcal{H}')}.$$

Thus $\overline{|\mathcal{H}' \stackrel{k-1}{\to} \mathcal{H}''|}$ is roughly the expected value of $|H' \stackrel{k-1}{\to} \mathcal{H}''|$ if $\mathcal{G}$ were a random complex. Also,

$$\overline{|\mathcal{H}' \stackrel{k-1}{\to} \mathcal{H}''|} = \overline{|\mathcal{H}' \to \mathcal{H}''|} / d_k^{e_k(\mathcal{H}'')-e_k(\mathcal{H}')} \geq \overline{|\mathcal{H}' \to \mathcal{H}''|}.$$

Now define a copy N_h of $\mathcal{N}_h$ in $\mathcal{G}$ to be *useful* if it has about the expected number of extensions to $\mathcal{B}$ as a $(k-1)$-complex. More precisely, N_h is useful if $|N_h \stackrel{k-1}{\to} \mathcal{B}| = (1 \pm \varepsilon_{k-1})\overline{|\mathcal{N}_h \stackrel{k-1}{\to} \mathcal{B}|}$. We denote the set of useful copies of $\mathcal{N}_h$ by `Usef`. Our next aim is to show that all but at most $\sqrt{\varepsilon_{k-1}}|\mathcal{N}_h|_{\mathcal{G}}$ copies of $\mathcal{N}_h$ are useful. This proportion will turn out to be small enough that it does not affect our bounds on $|\mathcal{H}|_{\mathcal{G}}$ significantly.

So let $|\mathcal{N}_h|_{\mathcal{G}}^{k-1}$ denote the number of copies of the underlying $(k-1)$-complex of $\mathcal{N}_h$ in $\mathcal{G}$. Then Lemmas 4 and 6 together imply that $|\mathcal{N}_h|_{\mathcal{G}}^{k-1} \leq (1+2\varepsilon_k)|\mathcal{N}_h|_{\mathcal{G}}/d_k^{e_k(\mathcal{N}_h)}$. Moreover, the dense extension lemma (Lemma 7) shows that all but at most $\varepsilon_{k-1}|\mathcal{N}_h|_{\mathcal{G}}^{k-1}$ copies of the underlying $(k-1)$-complex of $\mathcal{N}_h$ in $\mathcal{G}$ are useful (i.e. have about the expected number of extensions to the underlying $(k-1)$-complex of $\mathcal{B}$). Altogether this shows that all but at most

$$\varepsilon_{k-1}(1+2\varepsilon_k)|\mathcal{N}_h|_{\mathcal{G}}/d_k^{e_k(\mathcal{N}_h)} \leq \sqrt{\varepsilon_{k-1}}|\mathcal{N}_h|_{\mathcal{G}} \tag{9}$$

copies of $\mathcal{N}_h$ are useful.

So we will be able to obtain the expected lower bound on $|N_h \to \mathcal{B}|$ for the typical copies N_h of $\mathcal{N}_h$ in $\mathcal{G}$, and it will turn out that the non-useful copies of $\mathcal{N}_h$ will not affect the calculations significantly, but we do still need to deal with those copies of $\mathcal{N}_h$ which are useful but atypical. To do this, we define $\mathcal{N}_h^*$ to be the subcomplex of $\mathcal{H}$ induced by the vertices at distance 3 from h. We also define $\mathcal{F}$ to be the subcomplex of $\mathcal{H}$ induced by the vertices at distance 1, 2 or 3 from h, i.e. the subcomplex induced by $\mathcal{N}_h$, $\mathcal{N}_h^*$ and the vertices in between (see Figure 1).

Given copies N_h of $\mathcal{N}_h$ and N_h^* of $\mathcal{N}_h^*$, we say that the pair N_h, N_h^* is *useful* if N_h and N_h^* are disjoint and if the pair has about the expected number of extensions into copies of $\mathcal{F}$ as $(k-1)$-complexes, i.e. if $|N_h \cup N_h^* \stackrel{k-1}{\to} \mathcal{F}| = (1 \pm \varepsilon_{k-1})\overline{|\mathcal{N}_h \cup \mathcal{N}_h^* \stackrel{k-1}{\to} \mathcal{F}|}$. Once again, one can use Lemmas 4, 6 and 7 applied to $\mathcal{N}_h \cup \mathcal{N}_h^*$ to show that at most $\sqrt{\varepsilon_{k-1}}|\mathcal{N}_h|_{\mathcal{G}}|\mathcal{N}_h^*|_{\mathcal{G}}$ disjoint pairs N_h, N_h^* are not useful. Together with the fact that only a few of the pairs N_h, N_h^* will intersect, this shows that at most $2\sqrt{\varepsilon_{k-1}}|\mathcal{N}_h|_{\mathcal{G}}|\mathcal{N}_h^*|_{\mathcal{G}}$ pairs N_h, N_h^* are not useful. Hence at most $\varepsilon_{k-1}^{1/4}|\mathcal{N}_h|_{\mathcal{G}}$ copies of $\mathcal{N}_h$ form a non-useful pair together with more than $2\varepsilon_{k-1}^{1/4}|\mathcal{N}_h^*|_{\mathcal{G}}$ copies of $\mathcal{N}_h^*$. We will remove such copies of $\mathcal{N}_h$ from the set of useful copies of $\mathcal{N}_h$. Then

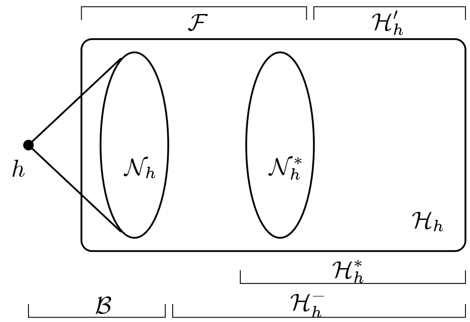

FIGURE 1. The complex $\mathcal{H}$

together with (9), this implies

$$|\mathcal{N}_h|_{\mathcal{G}} - |\mathtt{Usef}| \le 2\varepsilon_{k-1}^{1/4}|\mathcal{N}_h|_{\mathcal{G}}. \tag{10}$$

We denote by $\mathtt{Usef}^*(N_h)$ the set of all N_h^* which form a useful pair together with N_h.

Claim. *Any useful copy N_h of $\mathcal{N}_h$ satisfies*

$$|N_h \to \mathcal{H}_h| \le \frac{10}{d_k^{\Delta^3}} \frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}}.$$

Note that $\sum_{N_h} |N_h \to \mathcal{H}_h| = |\mathcal{H}_h|_{\mathcal{G}}$, so $|\mathcal{H}_h|_{\mathcal{G}}/|\mathcal{N}_h|_{\mathcal{G}}$ is the average value of $|N_h \to \mathcal{H}_h|$ over all copies N_h of $\mathcal{N}_h$. Later on, we will apply the claim to show that only a small fraction of copies of $\mathcal{H}$ contain a useful but atypical copy of $\mathcal{N}_h$. To prove the claim, fix a useful copy N_h of $\mathcal{N}_h$. Put $\mathcal{H}_h^* := \mathcal{H}_h - (\mathcal{F} - \mathcal{N}_h^*)$. We aim to extend N_h to a copy of $\mathcal{H}_h$ by first picking a copy N_h^* of $\mathcal{N}_h^*$, then extending this to a copy of $\mathcal{H}_h^*$ and also extending $N_h \cup N_h^*$ to a copy of $\mathcal{F}$. We must also make sure that no vertices are used more than once. However, since we are only looking for an upper bound on $|N_h \to \mathcal{H}_h|$, and ignoring this restriction can only increase the number of extensions we find, we may ignore this difficulty. Thus

$$|N_h \to \mathcal{H}_h| \le \sum_{N_h^* \in \mathtt{Usef}^*(N_h)} |N_h \cup N_h^* \to \mathcal{F}||N_h^* \to \mathcal{H}_h^*| + \sum_{N_h^* \notin \mathtt{Usef}^*(N_h)} |N_h \cup N_h^* \to \mathcal{F}||N_h^* \to \mathcal{H}_h^*|. \tag{11}$$

We bound the two sums separately. To bound the first sum, we need to bound $|N_h \cup N_h^* \to \mathcal{F}|$ in the case when the pair N_h, N_h^* is useful. But clearly $|N_h \cup N_h^* \to \mathcal{F}| \le |N_h \cup N_h^* \overset{k-1}{\to} \mathcal{F}|$, and

$$|N_h \cup N_h^* \overset{k-1}{\to} \mathcal{F}| \le (1+\varepsilon_{k-1})\overline{|\mathcal{N}_h \cup \mathcal{N}_h^* \overset{k-1}{\to} \mathcal{F}|} = \frac{(1+\varepsilon_{k-1})\overline{|\mathcal{N}_h \cup \mathcal{N}_h^* \to \mathcal{F}|}}{d_k^{e_k(\mathcal{F})-e_k(\mathcal{N}_h)-e_k(\mathcal{N}_h^*)}}$$

whenever $N_h^* \in \mathtt{Usef}^*(N_h)$. So the first sum in (11) is bounded by

$$\frac{1+\varepsilon_{k-1}}{d_k^{e_k(\mathcal{F})-e_k(\mathcal{N}_h)-e_k(\mathcal{N}_h^*)}}\overline{|\mathcal{N}_h \cup \mathcal{N}_h^* \to \mathcal{F}|}|\mathcal{H}_h^*|_{\mathcal{G}} \le \frac{2}{d_k^{\Delta^3}}\overline{|\mathcal{N}_h \cup \mathcal{N}_h^* \to \mathcal{F}|}|\mathcal{H}_h^*|_{\mathcal{G}}. \tag{12}$$

To see the bound of Δ^3 on the number of k-edges which we used in the final inequality, note that $|\mathcal{F} - \mathcal{N}_h - \mathcal{N}_h^*| \le \Delta^2$ and that the number of k-edges each of these vertices lies in is at most Δ. We now want to express the bound in (12) in terms of $|\mathcal{H}_h^-|_{\mathcal{G}}$, where $\mathcal{H}_h^- := \mathcal{H}_h - \mathcal{N}_h$.

By the induction hypothesis applied several times,

$$\begin{aligned}|\mathcal{H}_h^*|_{\mathcal{G}} &\le ((1-\alpha)n)^{-(|\mathcal{H}_h^-|-|\mathcal{H}_h^*|)}\left(\prod_{i=2}^k d_i^{-(e_i(\mathcal{H}_h^-)-e_i(\mathcal{H}_h^*))}\right)|\mathcal{H}_h^-|_{\mathcal{G}}\\ &\le 2\frac{\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}}{\overline{|\mathcal{N}_h\cup\mathcal{N}_h^*\to\mathcal{F}|}}|\mathcal{H}_h^-|_{\mathcal{G}}.\end{aligned}$$

In the last line we used that $e_i(\mathcal{H}_h)=e_i(\mathcal{H}_h^*)+e_i(\mathcal{F})-e_i(\mathcal{N}_h^*)$ and $|\mathcal{F}|-|\mathcal{N}_h|-|\mathcal{N}_h^*|=|\mathcal{H}_h^-|-|\mathcal{H}_h^*|$ (see Figure 1). We also used that $(1-\alpha)^{-(|\mathcal{H}_h^-|-|\mathcal{H}_h^*|)}\le 2$. So we obtain

$$\sum_{N_h^*\in\texttt{Usef}^*(N_h)}|N_h\cup N_h^*\to\mathcal{F}||N_h^*\to\mathcal{H}_h^*|\le\frac{4\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}}{d_k^{\Delta^3}}|\mathcal{H}_h^-|_{\mathcal{G}}.\tag{13}$$

To bound the second sum in (11), we define $\mathcal{H}_h':=\mathcal{H}_h^*-\mathcal{N}_h^*$, and observe that trivially any copy N_h^* of $\mathcal{N}_h^*$ satisfies $|N_h^*\to\mathcal{H}_h^*|\le|\mathcal{H}_h'|_{\mathcal{G}}$. On the other hand, by the induction hypothesis applied several times,

$$|\mathcal{H}_h'|_{\mathcal{G}}\le((1-\alpha)n)^{|\mathcal{H}_h'|-|\mathcal{H}_h^-|}\left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h')-e_i(\mathcal{H}_h^-)}\right)|\mathcal{H}_h^-|_{\mathcal{G}}\le\frac{2|\mathcal{H}_h^-|_{\mathcal{G}}}{\left(\prod_{i=2}^k d_i\right)^{2\Delta^4}n^{|\mathcal{H}_h^-|-|\mathcal{H}_h'|}}.$$

Since at most $2\varepsilon_{k-1}^{1/4}|\mathcal{N}_h^*|_{\mathcal{G}}\le 2\varepsilon_{k-1}^{1/4}n^{|\mathcal{N}_h^*|}$ copies of $\mathcal{N}_h^*$ do not lie in $\texttt{Usef}(N_h)$, the second sum in (11) is bounded by

$$\begin{aligned}\sum_{N_h^*\notin\texttt{Usef}(N_h)}|N_h\cup N_h^*\to\mathcal{F}||N_h^*\to\mathcal{H}_h^*| &\le 2\varepsilon_{k-1}^{1/4}n^{|\mathcal{N}_h^*|}n^{|\mathcal{F}|-|\mathcal{N}_h|-|\mathcal{N}_h^*|}\frac{2|\mathcal{H}_h^-|_{\mathcal{G}}}{\left(\prod_{i=2}^k d_i\right)^{2\Delta^4}n^{|\mathcal{H}_h^-|-|\mathcal{H}_h'|}}\\ &= 2\varepsilon_{k-1}^{1/4}\frac{2|\mathcal{H}_h^-|_{\mathcal{G}}}{\left(\prod_{i=2}^k d_i\right)^{2\Delta^4}}\\ &\le \left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}\right)|\mathcal{H}_h^-|_{\mathcal{G}}.\end{aligned}\tag{14}$$

The last inequality follows since $\varepsilon_{k-1}\ll d_2,d_3,\dots,d_k,1/\Delta$. Substituting (13) and (14) into (11) we obtain

$$\begin{aligned}|N_h\to\mathcal{H}_h| &\le \left(1+\frac{4}{d_k^{\Delta^3}}\right)\left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}\right)|\mathcal{H}_h^-|_{\mathcal{G}}\\ &\le \frac{5\left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}\right)}{d_k^{\Delta^3}}|\mathcal{H}_h^-|_{\mathcal{G}}.\end{aligned}\tag{15}$$

It now remains only to relate $|\mathcal{H}_h^-|_{\mathcal{G}}$ to $|\mathcal{H}_h|_{\mathcal{G}}/|\mathcal{N}_h|_{\mathcal{G}}$. Once again we apply the induction hypothesis several times to obtain

$$|\mathcal{H}_h|_{\mathcal{G}}\ge((1-\alpha)n)^{|\mathcal{H}_h|-|\mathcal{H}_h^-|}\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)}|\mathcal{H}_h^-|_{\mathcal{G}}.$$

On the other hand, the counting lemma implies that $|\mathcal{N}_h|_{\mathcal{G}} \le (1+\alpha)\left(\prod_{i=2}^k d_i^{e_i(\mathcal{N}_h)}\right) n^{|\mathcal{N}_h|}$. Putting these two bounds together, we obtain

$$\begin{aligned} \frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}} &\ge \frac{((1-\alpha)n)^{|\mathcal{H}_h|-|\mathcal{H}_h^-|}\left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)}\right)|\mathcal{H}_h^-|}{(1+\alpha)\left(\prod_{i=2}^k d_i^{e_i(\mathcal{N}_h)}\right)n^{|\mathcal{N}_h|}} \\ &\ge \frac{1}{2}\left(\prod_{i=2}^k d_i^{e_i(\mathcal{H}_h)-e_i(\mathcal{H}_h^-)-e_i(\mathcal{N}_h)}\right)|\mathcal{H}_h^-|_{\mathcal{G}}. \end{aligned} \tag{16}$$

Together with (15), this shows that

$$|N_h \to \mathcal{H}_h| \le \frac{5\cdot 2}{d_k^{\Delta^3}}\frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}},$$

which completes the proof of the claim.

Using the claim we now go on to prove the induction step. Given a copy H_h of $\mathcal{H}_h$, we denote by $N_h(H_h)$ the induced copy of $\mathcal{N}_h$. We have

$$\begin{aligned} |\mathcal{H}|_{\mathcal{G}} &= \sum_{H_h\in\mathcal{G}} |H_h \to \mathcal{H}| \ge \sum_{H_h\in\mathcal{G}} (|N_h(H_h)\to\mathcal{B}| - cn) \\ &= \sum_{N_h\in\mathcal{G}} |N_h\to\mathcal{H}_h||N_h\to\mathcal{B}| - cn|\mathcal{H}_h|_{\mathcal{G}} \\ &\ge (1-\varepsilon_k)\overline{|\mathcal{N}_h\to\mathcal{B}|}\left(\sum_{N_h\in\mathcal{G}}|N_h\to\mathcal{H}_h| - \sum_{N_h\notin\mathtt{typ}}|N_h\to\mathcal{H}_h|\right) - cn|\mathcal{H}_h|_{\mathcal{G}}. \end{aligned} \tag{17}$$

We want to show that the term in this expression which comes from the atypical copies of $\mathcal{N}_h$ does not affect the calculations too much, and so we aim to bound the contribution from atypical copies of $\mathcal{N}_h$. We have

$$\sum_{N_h\notin\mathtt{typ}}|N_h\to\mathcal{H}_h| = \sum_{N_h\notin\mathtt{typ},N_h\in\mathtt{Usef}}|N_h\to\mathcal{H}_h| + \sum_{N_h\notin\mathtt{typ},N_h\notin\mathtt{Usef}}|N_h\to\mathcal{H}_h|. \tag{18}$$

Now the claim implies that we can bound the first sum in (18) by

$$\sum_{N_h\notin\mathtt{typ},N_h\in\mathtt{Usef}}|N_h\to\mathcal{H}_h| \le \sum_{N_h\notin\mathtt{typ},N_h\in\mathtt{Usef}}\frac{10}{d_k^{\Delta^3}}\frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}} \le |\mathtt{atyp}|\frac{10}{d_k^{\Delta^3}}\frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}} \le \sqrt{\varepsilon_k}|\mathcal{H}_h|_{\mathcal{G}}. \tag{19}$$

Meanwhile we can also bound the second sum by

$$\begin{aligned} \sum_{N_h\notin\mathtt{typ},N_h\notin\mathtt{Usef}}|N_h\to\mathcal{H}_h| &\le \sum_{N_h\notin\mathtt{typ},N_h\notin\mathtt{Usef}}|\mathcal{H}_h^-|_{\mathcal{G}} \\ &\overset{(16)}{\le} (|\mathcal{N}_h|_{\mathcal{G}} - |\mathtt{Usef}|)\frac{2}{\prod_{i=1}^k d_i^{\Delta^2}}\frac{|\mathcal{H}_h|_{\mathcal{G}}}{|\mathcal{N}_h|_{\mathcal{G}}} \\ &\overset{(10)}{\le} \varepsilon_{k-1}^{1/5}|\mathcal{H}_h|_{\mathcal{G}}. \end{aligned} \tag{20}$$

Combining (18), (19) and (20), we have

$$\sum_{N_h\notin\mathtt{typ}}|N_h\to\mathcal{H}_h| \le 2\sqrt{\varepsilon_k}|\mathcal{H}_h|_{\mathcal{G}}$$

and combining this with (17), we obtain

$$\begin{aligned}
|\mathcal{H}|_{\mathcal{G}} &\geq (1-\varepsilon_k)\overline{|\mathcal{N}_h \to \mathcal{B}|}\,(|\mathcal{H}_h|_{\mathcal{G}} - 2\sqrt{\varepsilon_k}|\mathcal{H}_h|_{\mathcal{G}}) - cn|\mathcal{H}_h|_{\mathcal{G}} \\
&= (1-\varepsilon_k)n\left(\prod_{i=2}^{k} d_i^{e_i(\mathcal{B})-e_i(\mathcal{N}_h)}\right)(1-2\sqrt{\varepsilon_k})|\mathcal{H}_h|_{\mathcal{G}} - cn|\mathcal{H}_h|_{\mathcal{G}} \\
&\geq (1-\alpha)n\left(\prod_{i=2}^{k} d_i^{e_i(\mathcal{H})-e_i(\mathcal{H}_h)}\right)|\mathcal{H}_h|_{\mathcal{G}},
\end{aligned}$$

as required. This completes the proof of Lemma 3.

7. The regularity lemma for k-uniform hypergraphs

7.1. **Preliminary definitions and statement.** In this section we state the version of the regularity lemma for k-uniform hypergraphs due to Rödl and Schacht [23], which we use in the proof of Theorem 1 in the next section. To prepare for this we will first need some notation. We follow [23]. Given a finite set V of vertices, we will define a family $\mathcal{P} = \{\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k-1)}\}$ where each $\mathcal{P}^{(j)}$ is a partition of certain j-subsets of V. These partitions will satisfy properties which we will describe below. We denote by $[V]^j$ the set of all j-subsets of V. Suppose that we are given a partition $\mathcal{P}^{(1)} = \{V_1, \dots, V_{|\mathcal{P}^{(1)}|}\}$ of $[V]^1 = V$. We will call the V_i *clusters*. We denote by $\mathrm{Cross}_j = \mathrm{Cross}_j(\mathcal{P}^{(1)})$ the set of all those j-subsets of V that meet each part of $\mathcal{P}^{(1)}$ in at most 1 element. Each $\mathcal{P}^{(j)}$ will be a partition of Cross_j. Moreover, any two j-sets that belong to the same part of $\mathcal{P}^{(j)}$ will meet the same j clusters. This means that each part of $\mathcal{P}^{(j)}$ can be viewed as a j-partite j-uniform hypergraph whose vertex classes are these clusters. In particular, the parts of $\mathcal{P}^{(2)}$ can be thought of as bipartite subgraphs between two of the clusters. Moreover, for each part A of $\mathcal{P}^{(3)}$ there will be 3 clusters and 3 bipartite graphs belonging to $\mathcal{P}^{(2)}$ between these clusters such that all the 3-sets in A form triangles in the union of these 3 bipartite graphs.

More generally, suppose that we have already defined partitions $\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(j-1)}$ and are about to define $\mathcal{P}^{(j)}$. Given $i < j$ and $I \in \mathrm{Cross}_i$, we let $P^{(i)}(I)$ denote the part of $\mathcal{P}^{(i)}$ the set I belongs to. Given $J \in \mathrm{Cross}_j$, the *polyad* $\hat{P}^{(j-1)}(J)$ *of* J is defined by

$$\hat{P}^{(j-1)}(J) := \bigcup\{P^{(j-1)}(I) :\ I \in [J]^{j-1}\}.$$

Thus $\hat{P}^{(j-1)}(J)$ is the unique collection of j parts of $\mathcal{P}^{(j-1)}$ in which J spans a copy of the complete $(j-1)$-uniform hypergraph $K_j^{(j-1)}$ on j vertices. Moreover, note that $\hat{P}^{(j-1)}(J)$ can be viewed as a j-partite $(j-1)$-uniform hypergraph whose vertex classes are the j clusters containing the vertices of J. We set

$$\hat{\mathcal{P}}^{(j-1)} := \{\hat{P}^{(j-1)}(J) :\ J \in \mathrm{Cross}_j\}.$$

Note that the polyads $\hat{P}^{(j-1)}(J)$ and $\hat{P}^{(j-1)}(J')$ need not be distinct for different $J, J' \in [V]^j$. However, if these polyads are distinct then $\mathcal{K}_j(\hat{P}^{(j-1)}(J)) \cap \mathcal{K}_j(\hat{P}^{(j-1)}(J')) = \emptyset$. (Recall that $\mathcal{K}_j(\hat{P}^{(j-1)}(J))$ is the set of all j-sets of vertices which form a $K_j^{(j-1)}$ in $\hat{P}^{(j-1)}(J)$. So in particular, $\mathcal{K}_j(\hat{P}^{(j-1)}(J))$ contains J.) This implies that $\{\mathcal{K}_j(\hat{P}^{(j-1)}) : \hat{P}^{(j-1)} \in \hat{\mathcal{P}}^{(j-1)}\}$ is a partition of Cross_j. The property of $\mathcal{P}^{(j)}$ which we require is that it refines $\{\mathcal{K}_j(\hat{P}^{(j-1)}) : \hat{P}^{(j-1)} \in \hat{\mathcal{P}}^{(j-1)}\}$, i.e. each part of $\mathcal{P}^{(j)}$ has to be contained in some $\mathcal{K}_j(\hat{P}^{(j-1)})$.

We also need a notion which generalizes that of a polyad: given $J \in \mathrm{Cross}_j$ and $i < j$ we set

$$\hat{P}^{(i)}(J) := \bigcup\{P^{(i)}(I) :\ I \in [J]^i\}.$$

Then the properties of our partitions imply that $\bigcup_{i=1}^{j-1} \hat{P}^{(i)}(J)$ is a $(j-1,j)$-complex.

Altogether, given $\boldsymbol{a} = (a_1, \dots, a_{k-1}) \in \mathbb{N}^{k-1}$ we say that $\mathcal{P}(k-1, \boldsymbol{a}) = \{\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k-1)}\}$ is a *family of partitions on V* if

(1) $\mathcal{P}^{(1)}$ is a partition of V into a_1 clusters.
(2) For all $j = 2, \dots, k-1$, $\mathcal{P}^{(j)}$ is a partition of Cross_j such that for each part there is a polyad $\hat{P}^{(j-1)} \in \hat{\mathcal{P}}^{(j-1)}$ so that the part is contained in $\mathcal{K}_j(\hat{P}^{(j-1)})$. Moreover, for each polyad $\hat{P}^{(j-1)} \in \hat{\mathcal{P}}^{(j-1)}$, the set $\mathcal{K}_j(\hat{P}^{(j-1)})$ is the union of a_j parts of $\mathcal{P}^{(j)}$.

We say that $\mathcal{P} = \mathcal{P}(k-1, \boldsymbol{a})$ is *t-bounded* if $a_1, \dots, a_{k-1} \le t$. Suppose that a_1 divides $|V|$. Then $\mathcal{P} = \mathcal{P}(k-1, \boldsymbol{a})$ is called *$(\eta, \delta, \boldsymbol{a})$-equitable* if

(1) $\mathcal{P}^{(1)}$ is a partition of V into a_1 clusters of equal size;
(2) $|[V]^k \setminus \mathrm{Cross}_k| \le \eta\binom{|V|}{k}$;
(3) for every $K \in \mathrm{Cross}_k$, the $(k-1,k)$-complex $\bigcup_{i=1}^{k-1} \hat{P}^{(i)}(K)$ is $(\mathbf{d}, \delta, \delta, 1)$-regular, where $\mathbf{d} = (1/a_{k-1}, \dots, 1/a_2)$.

In particular, the second condition implies that $1/a_1$ is small compared to η.

Let $\delta_k > 0$ and $r \in \mathbb{N}$. Suppose that $\mathcal{G}$ is a k-uniform hypergraph on V and $\mathcal{P} = \mathcal{P}(k-1, \boldsymbol{a})$ is a family of partitions on V. Recall that we can view each polyad $\hat{P}^{(k-1)} \in \hat{\mathcal{P}}^{(k-1)}$ as a $(k-1)$-uniform k-partite hypergraph. $\mathcal{G}$ is called *(δ_k, r)-regular with respect to $\hat{P}^{(k-1)}$* if $\mathcal{G}$ is (δ_k, d, r)-regular with respect to $\hat{P}^{(k-1)}$ for some d. We say that $\mathcal{G}$ is *(δ_k, r)-regular with respect to $\mathcal{P}$* if

$$\left|\bigcup\{\mathcal{K}_k(\hat{P}^{(k-1)}) : \ \mathcal{G} \text{ is not } (\delta_k, r)\text{-regular with respect to } \hat{P}^{(k-1)} \in \hat{\mathcal{P}}^{(k-1)}\}\right| \le \delta_k |V|^k.$$

This means that not much more than a δ_k-fraction of the k-subsets of V form a $K_k^{(k-1)}$ that lies within a polyad with respect to which $\mathcal{G}$ is not regular.

Now, we are ready to state the regularity lemma, which we are going to use in the proof of Theorem 1.

Theorem 12 (Rödl and Schacht [23])**.** *Let $k \ge 2$ be a fixed integer. For all positive constants η and δ_k and all functions $r : \mathbb{N}^{k-1} \to \mathbb{N}$ and $\delta : \mathbb{N}^{k-1} \to (0,1]$, there are integers t and m_0 such that the following holds for all $m \ge m_0$ which are divisible by $t!$. Suppose that $\mathcal{G}$ is a k-uniform hypergraph of order m. Then there exists an $\boldsymbol{a} \in \mathbb{N}^{k-1}$ and a family of partitions $\mathcal{P} = \mathcal{P}(k-1, \boldsymbol{a})$ of the vertex set V of $\mathcal{G}$ such that*

(1) *$\mathcal{P}$ is $(\eta, \delta(\boldsymbol{a}), \boldsymbol{a})$-equitable and t-bounded and*
(2) *$\mathcal{G}$ is $(\delta_k, r(\boldsymbol{a}))$-regular with respect to $\mathcal{P}$.*

The advantage of this regularity lemma compared to the one proved earlier by Rödl and Skokan [24] is that it uses only two regularity constants δ and δ_k instead of $k-1$ different ones. The regularity constants $\delta_2, \dots, \delta_k$ produced by the regularity lemma in [24] might satisfy $\delta_2 \ll 1/a_2 \ll \delta_3 \ll 1/a_3 \ll \dots \ll 1/a_{k-1} \ll \delta_k$, which would make the proof of the corresponding embedding lemma more complicated.

Note that the constants in Theorem 12 can be chosen such that they satisfy the following hierarchy:

$$\frac{1}{m_0} \ll \frac{1}{r} = \frac{1}{r(\boldsymbol{a})}, \delta = \delta(\boldsymbol{a}) \ll \min\{\delta_k, \eta, 1/a_1, 1/a_2, \dots, 1/a_{k-1}\}. \tag{21}$$

7.2. **The reduced hypergraph.** In the proof of Theorem 1 that follows in the next section, we will use the so-called reduced hypergraph. If $\mathcal{P} = \{\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(k-1)}\}$ is the partition of the vertex set of $\mathcal{G}$ given by the regularity lemma, the *reduced hypergraph* $\mathcal{R} = \mathcal{R}(\mathcal{G}, \mathcal{P})$ is a k-uniform hypergraph whose vertices are the clusters, i.e. the parts of $\mathcal{P}^{(1)}$. To define the set of hyperedges we need the following notion. We say that a k-tuple of clusters is *useful* if $\mathcal{G}$

is (δ_k, r)-regular with respect to all but at most a $\sqrt{\delta_k}$-fraction of all those polyads $\hat{P}^{(k-1)}$ which are induced on these k clusters. The set of hyperedges of $\mathcal{R}$ consists of precisely those k-tuples that are useful. In the proof of Theorem 1, we shall need an estimate on the number of these hyperedges. In particular, we need to show that $\mathcal{R}$ is very dense. This is conveyed in the following proposition.

Proposition 13. *All but at most $2\sqrt{\delta_k}a_1^k$ of the k-tuples of clusters are useful.*

Proof. By the dense counting lemma (Lemma 6) each polyad in $\hat{\mathcal{P}}^{(k-1)}$ contains at least

$$f(m, \boldsymbol{a}) := \frac{1}{2}\left(\frac{m}{a_1}\right)^k \prod_{i=2}^{k-1}\left(\frac{1}{a_i}\right)^{\binom{k}{i}}$$

copies of $K_k^{(k-1)}$. Since $\mathcal{G}$ is (δ_k, r)-regular with respect to $\mathcal{P}$, the number of polyads in $\hat{\mathcal{P}}^{(k-1)}$ with respect to which $\mathcal{G}$ is not (δ_k, r)-regular is at most

$$\frac{\delta_k m^k}{f(m, \boldsymbol{a})} = \frac{2\prod_{i=1}^{k-1} a_i^{\binom{k}{i}}}{m^k}\delta_k m^k = 2\delta_k \prod_{i=1}^{k-1} a_i^{\binom{k}{i}}. \tag{22}$$

We call these polyads *bad*. Now, each k-tuple of clusters induces $\prod_{i=2}^{k-1} a_i^{\binom{k}{i}}$ polyads in $\hat{\mathcal{P}}^{(k-1)}$. Thus if there were more than $2\sqrt{\delta_k}a_1^k$ k-tuples of clusters each inducing more than $\sqrt{\delta_k}\prod_{i=2}^{k-1} a_i^{\binom{k}{i}}$ bad polyads, the total number of bad polyads would exceed the bound given in (22), yielding a contradiction. □

8. Proof of Theorem 1

We now give a brief outline of the proof of Theorem 1: consider any red/blue colouring of the hyperedges of $K_m^{(k)}$, where $m = C|\mathcal{H}|$ and C is a large constant depending only on k and the maximum degree of $\mathcal{H}$. We apply the hypergraph regularity lemma to the red subhypergraph $\mathcal{G}_{red}$ to obtain a reduced hypergraph $\mathcal{R}$ which is very dense. Thus the following lemma will show that $\mathcal{R}$ contains a copy of $K_\ell^{(k)}$ with $\ell := R(K_{k\Delta}^{(k)})$. The lemma itself is a weak hypergraph analogue of Turán's theorem and can be found in [6].

Lemma 14. *For all $\ell, k \in \mathbb{N}$ with $\ell \geq k$, there exists a constant $c_0 = c_0(k, \ell) < 1$ such that every k-uniform hypergraph $\mathcal{R}$ on $t \geq \ell$ vertices with $e(\mathcal{R}) \geq c_0\binom{t}{k}$ contains a copy of $K_\ell^{(k)}$.*

The copy of $K_\ell^{(k)}$ in $\mathcal{R}$ involves ℓ clusters and for each k-tuple of them the red hypergraph $\mathcal{G}_{red}$ is regular with respect to almost all of the polyads induced on it. We will then show that we can find a $(k-1, \ell)$-complex $\mathcal{S}$ on these clusters such that for each $j = 2, \ldots, k-1$ the restriction of its underlying j-uniform hypergraph $\mathcal{S}_j$ to any $(j+1)$-tuple of clusters is a polyad. Moreover, $\mathcal{G}_{red}$ will be regular with respect to $\mathcal{S}_{k-1}$. By combining $E(\mathcal{G}_{red}) \cap \mathcal{K}_k(\mathcal{S}_{k-1})$ with $\mathcal{S}$, we will obtain a regular k-complex $\mathcal{S}_{red}$. Similarly we obtain a k-complex $\mathcal{S}_{blue}$ which also turns out to be regular. We then consider the following red/blue colouring of $K_\ell^{(k)}$. We colour a hyperedge red if $\mathcal{G}_{red}$ has density at least $1/2$ with respect to the corresponding polyad in $\mathcal{S}_{k-1}$ and blue otherwise. By the definition of ℓ, we can find a monochromatic $K_{k\Delta}^{(k)}$. If it is red, then we can apply the embedding lemma to $\mathcal{S}_{red}$ to find a red copy of $\mathcal{H}$. This can be done since $\Delta(\mathcal{H}) \leq \Delta$ implies that the chromatic number of $\mathcal{H}$ is at most $(k-1)\Delta + 1 \leq k\Delta$. If our monochromatic copy of $K_{k\Delta}^{(k)}$ is blue, then we can apply the embedding lemma to $\mathcal{S}_{blue}$ and obtain a blue copy of $\mathcal{H}$.

Proof of Theorem 1. Given Δ and k, we choose C to be a sufficiently large constant. We will describe the bounds that C has to satisfy at the end of the proof. Let $m := C|\mathcal{H}|$ and consider any red/blue colouring of the hyperedges of $K_m^{(k)}$. Let $\mathcal{G}_{red}$ be the red and $\mathcal{G}_{blue}$ be the blue subhypergraph on $V = V(K_m^{(k)})$. We may assume without loss of generality that $e(\mathcal{G}_{red}) \geq e(\mathcal{G}_{blue})$. We apply the hypergraph regularity lemma to $\mathcal{G}_{red}$ with constants $\eta, \delta_k \ll 1/\Delta, 1/k$ as well as functions r and δ satisfying the hierarchy in (21). This gives us clusters $V_1, \dots, V_{a_1}$, each of size n say, together with a t-bounded $(\eta, \delta, \boldsymbol{a})$-equitable family of partitions $\mathcal{P} = \mathcal{P}(k-1, \boldsymbol{a})$ on V where $\boldsymbol{a} = (a_1, \dots, a_{k-1})$. (Note that by deleting some vertices of $\mathcal{G}_{red}$ if necessary we may assume that $m = |\mathcal{G}_{red}|$ is divisible by $t!$.) Since $\eta \ll 1/\Delta, 1/k$, condition (2) in the definition of an $(\eta, \delta, \boldsymbol{a})$-equitable family of partitions implies that the a_1 which we obtain from the regularity lemma satisfies

$$a_1 \geq R(K_{k\Delta}^{(k)}) =: \ell.$$

Note that the definition of ℓ involves a hypergraph Ramsey number whose value is unknown. However, for the argument below all we need is that this number exists.

Let $\mathcal{R}$ denote the reduced hypergraph, defined in the previous section. Proposition 13 implies that $\mathcal{R}$ has at least $(1-\varepsilon)\binom{a_1}{k}$ hyperedges, where $\varepsilon := 4\sqrt{\delta_k}k!$. Since $\delta_k \ll 1/\Delta, 1/k$, we may assume that $e(\mathcal{R}) \geq (1-\varepsilon)\binom{|\mathcal{R}|}{k} > c_0\binom{|\mathcal{R}|}{k}$, where c_0 is as defined in Lemma 14. Since $|\mathcal{R}| = a_1 \geq \ell$, this means that we can apply Lemma 14 to $\mathcal{R}$ to obtain a copy of $K_\ell^{(k)}$ in $\mathcal{R}$. Without loss of generality we may assume that the vertices of this copy are the clusters $V_1, \dots, V_\ell$.

As mentioned above, we now want to find a $(k-1, \ell)$-complex $\mathcal{S}$ on these clusters such that for each $j = 2, \dots, k-1$ its underlying j-uniform hypergraph $\mathcal{S}_j$ is a union of parts of $\mathcal{P}^{(j)}$ and $\mathcal{G}_{red}$ is regular with respect to $\mathcal{S}_{k-1}$. We construct $\mathcal{S}$ inductively starting from the lower levels. To begin with, for each pair V_i, V_j $(1 \leq i < j \leq \ell)$ independently, we choose with probability $1/a_2$ one of the parts of $\mathcal{P}^{(2)}$ induced on V_i, V_j. $\mathcal{S}_2$ will be the union of these parts. Now suppose that we have chosen $\mathcal{S}_{j-1}$ such that its restriction to any j-tuple of clusters forms a polyad (clearly this is the case for $\mathcal{S}_2$). Now, if $\hat{P}^{(j-1)}$ is such a polyad, we choose a part of $\mathcal{P}^{(j)}$ uniformly at random among the a_j parts of $\mathcal{P}^{(j)}$ that form $\mathcal{K}_j(\hat{P}^{(j-1)})$, independently for each j-tuple of clusters. We let $\mathcal{S}$ be the $(k-1, \ell)$-complex thus obtained.

We will show that there is some choice of $\mathcal{S}$ such that for every k-tuple among the clusters $V_1, \dots, V_\ell$ the hypergraph $\mathcal{G}_{red}$ is (δ_k, r)-regular with respect to the restriction of $\mathcal{S}_{k-1}$ to this k-tuple. Note that $\mathcal{S}_{k-1}$ restricted to any particular k-tuple of clusters is in fact a polyad selected uniformly at random among all polyads $\hat{P}^{(k-1)}$ induced by these k clusters. Therefore, since all the k-tuples of clusters are useful, the definition of a useful k-tuple implies that the probability that $\mathcal{G}_{red}$ has the necessary regularity is at least

$$1 - \sqrt{\delta_k}\binom{\ell}{k} > \frac{1}{2}.$$

The final inequality holds since we may assume that δ_k is sufficiently small compared to $1/\ell$. This shows the existence of a $(k-1, \ell)$-complex $\mathcal{S}$ with the required properties. In what follows, $P_{\mathcal{S}}$ will always denote a $(k-1)$-uniform subhypergraph of $\mathcal{S}$ induced by k of the clusters $V_1, \dots, V_\ell$. So each such $P_{\mathcal{S}}$ is a polyad and to each hyperedge of the subhypergraph of $\mathcal{R}$ induced by the clusters $V_1, \dots, V_\ell$ there corresponds such a polyad $P_{\mathcal{S}}$.

We now use the densities of $\mathcal{G}_{red}$ with respect to $\mathcal{S}_{k-1}$ to define a red/blue colouring of the $K_\ell^{(k)}$ which we found in $\mathcal{R}$: we colour a hyperedge of this $K_\ell^{(k)}$ red if the polyad $P_{\mathcal{S}}$ corresponding to this hyperedge satisfies $d(\mathcal{G}_{red}|P_{\mathcal{S}}) \geq 1/2$, otherwise we colour it blue. Since $\ell = R(K_{k\Delta}^{(k)})$, we find a monochromatic copy K of $K_{k\Delta}^{(k)}$ in our $K_\ell^{(k)}$. We now greedily assign the vertices of $\mathcal{H}$ to the clusters that form the vertex set of K in such a way that if k vertices of $\mathcal{H}$ form a hyperedge, then they are assigned to k different clusters. (We may think of this as a $(k\Delta)$-vertex-colouring of $\mathcal{H}$.) We now need to show that with this assignment we can apply the embedding lemma to find a monochromatic copy of $\mathcal{H}$ in either the subhypergraph of $\mathcal{G}_{red}$ induced by the $k\Delta$ clusters in K or the subhypergraph of $\mathcal{G}_{blue}$ induced by these clusters.

First suppose that K is red, so we want to apply the embedding lemma to the k-complex formed by $\mathcal{G}_{red}$ and $\mathcal{S}$ (induced on the $k\Delta$ clusters in K). However, the embedding lemma requires all the densities involved to be equal and of the from $1/a$ for $a \in \mathbb{N}$, whereas all we know is that for every polyad $P_{\mathcal{S}}$ corresponding to a hyperedge of K, we have $d(\mathcal{G}_{red}|P_{\mathcal{S}}) \geq 1/2$. This minor obstacle can be overcome by choosing a hypergraph $\mathcal{G}'_{red} \subseteq \mathcal{G}_{red}$ such that $\mathcal{G}'_{red}$ is $(1/2, 3\delta_k, r)$-regular with respect to each polyad $P_{\mathcal{S}}$. The existence of such a $\mathcal{G}'_{red}$ follows immediately from the slicing lemma [23, Prop. 33]. We then add $E(\mathcal{G}'_{red}) \cap \mathcal{K}_k(\mathcal{S}_{k-1})$ to the subcomplex of $\mathcal{S}$ induced by the clusters in K to obtain a regular $(k, k\Delta)$-complex $\mathcal{S}_{red}$ and we apply the embedding lemma (Lemma 2) there to find a copy of $\mathcal{H}$ in $\mathcal{G}'_{red}$, and therefore also in $\mathcal{G}_{red}$.

On the other hand, if K is blue, we need to prove that $\mathcal{G}_{blue}$ is regular with respect to all chosen polyads $P_{\mathcal{S}}$. So suppose $\mathbf{Q} = (Q(1), \dots, Q(r))$ is an r-tuple of subhypergraphs of one of these polyads $P_{\mathcal{S}}$, satisfying $|\mathcal{K}_k(\mathbf{Q})| > \delta_k |\mathcal{K}_k(P_{\mathcal{S}})|$. Let d be such that $\mathcal{G}_{red}$ is (d, δ_k, r)-regular with respect to $P_{\mathcal{S}}$. Then

$$|(1-d) - d(\mathcal{G}_{blue}|\mathbf{Q})| = |d - (1 - d(\mathcal{G}_{blue}|\mathbf{Q}))| = |d - d(\mathcal{G}_{red}|\mathbf{Q})| < \delta_k.$$

Thus $\mathcal{G}_{blue}$ is $(1-d, \delta_k, r)$-regular with respect to $P_{\mathcal{S}}$ (note that $\delta_k \ll 1/2 \leq 1-d$). Following the same argument as in the previous case, we add $E(\mathcal{G}'_{blue}) \cap \mathcal{K}_k(\mathcal{S}_{k-1})$ to the subcomplex of $\mathcal{S}$ induced by the clusters in K to derive the regular $(k, k\Delta)$-complex $\mathcal{S}_{blue}$ to which we can apply the embedding lemma to obtain a copy of $\mathcal{H}$ in $\mathcal{G}_{blue}$.

It remains to check that we can choose C to be a constant depending only on Δ and k. Note that the constants and functions η, δ_k, r and δ we defined at the beginning of the proof all depend only on Δ and k. So this is also true for the integers m_0 and t and the vector $\boldsymbol{a} = (a_1, \dots, a_{k-1})$ which we then obtained from the regularity lemma. Note that in order to be able to apply the regularity lemma to $\mathcal{G}_{red}$ we needed $m \geq m_0$, where $m = C|\mathcal{H}|$. This is certainly true if we set $C \geq m_0$. The embedding lemma allows us to embed subcomplexes of size at most cn, where n is the cluster size and where c satisfies $c \ll 1/a_2, , \dots, 1/a_{k-1}, d_k, 1/(k\Delta)$ (recall that $d_k = 1/2$ and $d_i = 1/a_i$ for all $i = 2, \dots, k-1$). Thus c too depends only on Δ and k. In order to apply the embedding lemma we needed that $n \geq n_0$, where n_0 as defined in the embedding lemma depends only on Δ and k. Since the number of clusters is at most t, this is satisfied if $m \geq tn_0$, which in turn is certainly true if $C \geq tn_0$. When we applied the embedding lemma to $\mathcal{H}$, we needed that $|\mathcal{H}| \leq cn$. Since $n = m/a_1 = C|\mathcal{H}|/a_1 \geq C|\mathcal{H}|/t$, it suffices to choose $C \geq t/c$ for this. Altogether, this shows that we can define the constant C in Theorem 1 by $C := \max\{m_0, tn_0, t/c\}$. □

9. Embedding almost spanning hypergraphs of bounded bandwidth

In Lemma 3 the complex $\mathcal{H}$ to be embedded had vertex classes of size at most cn, where n was the size of the vertex classes of the regular complex $\mathcal{G}$ and c was assumed to be relatively small, more precisely $c \ll \alpha, d_2, \dots, d_k$. Of course one would like to have an embedding lemma for complexes $\mathcal{H}$ whose vertex classes have size almost n (or even exactly n). This has been achieved in the case of graphs H of bounded maximum degree by means of the Blow-up lemma [15] as well as in some very special examples for 3-uniform hypergraphs $\mathcal{H}$ [18].

The aim of this section is to prove such an embedding lemma for ℓ-partite k-complexes $\mathcal{H}$ of bounded bandwidth. (So the order of $\mathcal{H}$ is allowed to be almost the same as that of $\mathcal{G}$.) A graph G has *bandwidth* at most B if there exists a linear ordering of its vertices such that the endvertices of each edge of G have distance at most B in this ordering. Of course, if G has bandwidth at most B then its maximum degree is at most $2B$. We say that a complex $\mathcal{H}$ has *bandwidth* at most B if this holds for its underlying graph $\mathcal{H}_2$. Note that if a k-complex $\mathcal{H}$ has bandwidth at most B then $k \le B + 1$. Moreover, for any i-edge of $\mathcal{H}$ the maximum distance between any two of its vertices in the linear ordering of the vertices of $\mathcal{H}$ is at most B.

Lemma 15 (Embedding lemma for complexes of bounded bandwidth)**.** *Let B, k, ℓ, r, n_0 be positive integers and let $c, d_2, \dots, d_k, \delta, \delta_k$ be positive constants such that $1/d_i \in \mathbb{N}$ and*

$$1/n_0 \ll 1/r, \delta \ll \min\{\delta_k, d_2, \dots, d_{k-1}\} \le \delta_k \ll d_k, 1-c, 1/B, 1/\ell.$$

Then the following holds for all integers $n \ge n_0$ and $\mathbf{d} := (d_k, \dots, d_2)$. Suppose that $\mathcal{H}$ is a (k,ℓ)-complex of bandwidth at most B with vertex classes $X_1, \dots, X_\ell$ such that $|X_i| \le cn$ for all $i = 1, \dots, \ell$. Suppose also that $\mathcal{G}$ is a $(\mathbf{d}, \delta_k, \delta, r)$-regular (k,ℓ)-complex with vertex classes $V_1, \dots, V_\ell$, all of size n, which respects the partition of $\mathcal{H}$. Then $\mathcal{G}$ contains a copy of $\mathcal{H}$.

Note that similarly as with Lemmas 2 and 3, the hierachy of the constants in Lemma 15 fits with the one provided by the regularity lemma (Theorem 12).

The proof of Lemma 15 differs from that of Lemma 3 – it is similar to the proof of the embedding lemma for graphs (see for example the proof of the Key Lemma in [16]). In the latter proof, the graph H is embedded vertex by vertex and at every step we maintain a set of possible candidates for each vertex of H that we still have to embed. In our proof of Lemma 15 though, we do not embed $\mathcal{H}$ vertex by vertex. Using the fact that $\mathcal{H}$ has bounded bandwidth we split its vertex set into a chain of small segments, which are embedded one by one according to their ordering in this chain.

One example of a hypergraph of bounded bandwidth is a hypergraph cycle (there are several possible definitions, but this applies to the usual ones). In [13] a result similar to Lemma 15 was used to determine the asymptotics of Ramsey numbers of 3-uniform tight cycles (see also [22] for related results and a sketch of the argument). So Lemma 15 may be a useful tool e.g. when extending this to k-uniform hypergraphs.

Proof of Lemma 15. Consider a linear ordering of the vertices of $\mathcal{H}$ in which any two vertices that lie in a common hyperedge have distance at most B. One can naturally split the vertices into consecutive (overlapping) segments $\mathcal{S}_1, \dots, \mathcal{S}_s$ as seen in Figure 2. Each segment

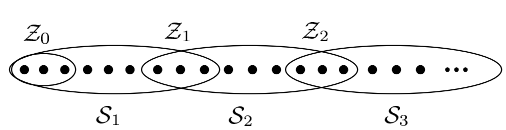

FIGURE 2. Splitting the vertices of $\mathcal{H}$ when $B = 3$.

is a subcomplex of $\mathcal{H}$ of order $3B$ and any two consecutive segments $\mathcal{S}_i$ and $\mathcal{S}_{i+1}$ overlap on exactly B vertices (by adding a bounded number of isolated vertices to $\mathcal{H}$ if necessary, we may assume that the final segment $\mathcal{S}_s$ has $3B$ vertices too). We let $\mathcal{Z}_i := \mathcal{S}_i \cap \mathcal{S}_{i+1}$ be the subcomplex induced by this overlap. Note that if $x \in \mathcal{Z}_i$ and $y \in \mathcal{Z}_{i-1}$ then no hyperedge of $\mathcal{H}$ contains both x and y. So $\mathcal{Z}_i \cup \mathcal{Z}_{i-1}$ will be an induced subcomplex of $\mathcal{S}_i$. This will be useful in the course of our proof. We denote by $\mathcal{Z}_s$ the subcomplex of $\mathcal{H}$ induced by the last B vertices and by $\mathcal{Z}_0$ the subcomplex of $\mathcal{H}$ induced by the first B vertices. For all $i = 1, \dots, s$ we let $\mathcal{H}_i := \bigcup_{j=1}^{i} \mathcal{S}_j$. We also set $\mathcal{H}_0 := \mathcal{S}_0 := \mathcal{Z}_0$ and $\mathcal{H}_{-1} := \emptyset$.

As before, when referring to a copy of a subcomplex of $\mathcal{H}$ in $\mathcal{G}$ we mean that this copy is labelled and partition-respecting without mentioning this explicitly. The proof of Lemma 15 proceeds as follows: Suppose that we have already embedded $\mathcal{H}_{i-1}$, i.e. we have fixed a copy H_{i-1} of $\mathcal{H}_{i-1}$ in $\mathcal{G}$. Suppose also that we have a collection of candidate copies of $\mathcal{Z}_i$ such that together with H_{i-1} each of these candidate copies extends into a copy of $\mathcal{H}_i$. We will prove that there is at least one amongst these candidate copies for $\mathcal{Z}_i$ for which one can find many candidate copies of $\mathcal{Z}_{i+1}$. We will then choose such a candidate copy Z_i of $\mathcal{Z}_i$ and an extension of $Z_i \cup H_{i-1}$ into a copy of $\mathcal{H}_i$ and continue. More formally, we will prove the following assertion by induction on i.

For every $i = 0, \dots, s$ there exists a copy H_{i-1} of $\mathcal{H}_{i-1}$ in $\mathcal{G}$ and a set $\overline{\mathcal{Z}}_i$ of at least $(1-c)^B |\mathcal{Z}_i|_{\mathcal{G}}/16$ copies of $\mathcal{Z}_i$ in $\mathcal{G} - V(H_{i-1})$ such that for every copy $Z \in \overline{\mathcal{Z}}_i$ the subcomplex $Z \cup H_{i-1}$ of $\mathcal{G}$ is extendible into a copy of $\mathcal{H}_i$ in $\mathcal{G}$. $(*)$

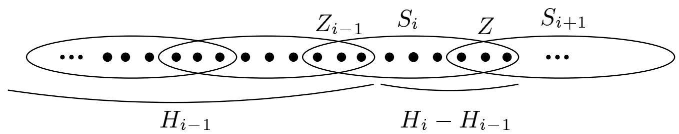

FIGURE 3. Extending the copy H_{i-1} of $\mathcal{H}_{i-1}$ into a copy of $\mathcal{H}_i$.

As $|\mathcal{Z}_s|_{\mathcal{G}} > 0$ by the counting lemma, Lemma 15 is an immediate consequence of $(*)$ for $i = s$. So it remains to prove $(*)$. If $i = 0$ then $\mathcal{H}_{i-1}$ is empty and $\mathcal{H}_i = \mathcal{Z}_0$. So we can take $\overline{\mathcal{Z}}_0$ to be the set of all copies of $\mathcal{Z}_0$ in $\mathcal{G}$. So suppose that $i \geq 0$ and $(*)$ holds for i. We have to prove that $(*)$ holds for $i+1$. If $i \geq 1$ let Z_{i-1} denote the copy of $\mathcal{Z}_{i-1}$ inside H_{i-1}. Given $Z \in \overline{\mathcal{Z}}_i$, note that an extension of $Z \cup H_{i-1}$ into a copy of $\mathcal{H}_i$ can be obtained by extending $Z \cup Z_{i-1}$ into a copy of $\mathcal{S}_i$ which meets H_{i-1} precisely in Z_{i-1}. (Indeed, this is true since no hyperedge of $\mathcal{H}$ meets both $\mathcal{H}_{i-1} - \mathcal{Z}_{i-1}$ and $\mathcal{S}_i - \mathcal{Z}_{i-1} = \mathcal{H}_i - \mathcal{H}_{i-1}$, see Figure 3.) Note that $(*)$ guarantees that for every $Z \in \overline{\mathcal{Z}}_i$ there is such a copy S_Z of $\mathcal{S}_i$. (If $i = 0$ then we take $S_Z := Z$.) Our aim is to find a $Z \in \overline{\mathcal{Z}}_i$ such that for the copy $H_{i-1} \cup S_Z$ of $\mathcal{H}_i$ there is a set $\overline{\mathcal{Z}}_{i+1}$ of copies of $\mathcal{Z}_{i+1}$ as in $(*)$.

Recall that H_{i-1} contains at most cn vertices from each vertex class of $\mathcal{G}$, and therefore there are at least $(1-c)n$ vertices available to embed the remainder of $\mathcal{H}$. We choose an arbitrary set of $(1-c)n$ such vertices in each vertex class of $\mathcal{G}$ and let $\mathcal{G}'$ denote the subcomplex of $\mathcal{G}$ that is induced by these vertices. Since $1 - c \gg \delta_k, \delta$, an easy inductive argument (starting with the edges and ending with the k-edges) shows that $\mathcal{G}'$ is still $(\mathbf{d}, \sqrt{\delta_k}, \sqrt{\delta}, r)$-regular.

Pick a new constant β such that

$$\delta_k \ll \beta \ll 1 - c, 1/B.$$

Apply the extension lemma (Lemma 5) to $\mathcal{G}'$ with $\mathcal{Z}_i \cup \mathcal{Z}_{i+1}$ playing the role of $\mathcal{H}$ and $\mathcal{S}_{i+1}$ playing the role of $\mathcal{H}'$ to see that all but at most $\beta |\mathcal{Z}_i \cup \mathcal{Z}_{i+1}|_{\mathcal{G}'} \leq \beta |\mathcal{Z}_i|_{\mathcal{G}'} |\mathcal{Z}_{i+1}|_{\mathcal{G}'}$ pairs Z_i, Z_{i+1} of disjoint copies of $\mathcal{Z}_i, \mathcal{Z}_{i+1}$ are extendible into at least

$$\mathtt{Ext} := \frac{1}{2}(1-c)^B n^B \prod_{j=2}^{k} d_j^{e_j(\mathcal{S}_{i+1}) - e_j(\mathcal{Z}_i) - e_j(\mathcal{Z}_{i+1})}$$

copies of $\mathcal{S}_{i+1}$ in $\mathcal{G}'$. Let $\hat{\mathcal{Z}}_i$ denote the set of copies Z_i of $\mathcal{Z}_i$ for which there exist at least $\sqrt{\beta}|\mathcal{Z}_{i+1}|_{\mathcal{G}'}$ copies Z_{i+1} of $\mathcal{Z}_{i+1}$ such that the pair Z_i, Z_{i+1} extends to less than `Ext` copies of $\mathcal{S}_{i+1}$ in $\mathcal{G}'$. Therefore $|\hat{\mathcal{Z}}_i| < \sqrt{\beta}|\mathcal{Z}_i|_{\mathcal{G}'}$. Fix any $Z_i \in \overline{\mathcal{Z}}_i \setminus \hat{\mathcal{Z}}_i$. (Such a Z_i exists since $|\overline{\mathcal{Z}}_i \setminus \hat{\mathcal{Z}}_i| \geq |\overline{\mathcal{Z}}_i| - \sqrt{\beta}|\mathcal{Z}_i|_{\mathcal{G}'} \geq ((1-c)^B/16 - \sqrt{\beta})|\mathcal{Z}_i|_{\mathcal{G}'} > 0$ since $|\mathcal{Z}_i|_{\mathcal{G}'} \leq |\mathcal{Z}_i|_{\mathcal{G}}$.) As $Z_i \notin \hat{\mathcal{Z}}_i$ there is a set $\overline{\mathcal{Z}}'_{i+1}$ of copies of $\mathcal{Z}_{i+1}$ in $\mathcal{G}'$ such that $|\overline{\mathcal{Z}}'_{i+1}| \geq (1-\sqrt{\beta})|\mathcal{Z}_{i+1}|_{\mathcal{G}'}$ and such that for every $Z_{i+1} \in \overline{\mathcal{Z}}'_{i+1}$ there are at least `Ext` extensions of $Z_i \cup Z_{i+1}$ into a copy of $\mathcal{S}_{i+1}$ in $\mathcal{G}'$. We take $H_i := H_{i-1} \cup S_{Z_i}$ to be the copy of $\mathcal{H}_i$ which is obtained by extending $H_{i-1} \cup Z_i$. $\overline{\mathcal{Z}}_{i+1}$ will be a subset of $\overline{\mathcal{Z}}'_{i+1}$.

Our next aim is to give a lower bound on $|\overline{\mathcal{Z}}'_{i+1}|$ in terms of $|\mathcal{Z}_{i+1}|_{\mathcal{G}}$. Thus, we first need to obtain a lower bound on $|\mathcal{Z}_{i+1}|_{\mathcal{G}'}$ in terms of $|\mathcal{Z}_{i+1}|_{\mathcal{G}}$. Applying the counting lemma (Lemma 4) to $\mathcal{G}'$, we obtain

$$|\mathcal{Z}_{i+1}|_{\mathcal{G}'} \geq \frac{1}{2}(1-c)^B n^B \prod_{j=2}^{k} d_j^{e_j(\mathcal{Z}_{i+1})}.$$

Also, the counting lemma applied to $\mathcal{G}$ itself yields

$$|\mathcal{Z}_{i+1}|_{\mathcal{G}} \leq 2n^B \prod_{j=2}^{k} d_j^{e_j(\mathcal{Z}_{i+1})},$$

which along with the previous inequality implies that $|\mathcal{Z}_{i+1}|_{\mathcal{G}'} \geq (1-c)^B |\mathcal{Z}_{i+1}|_{\mathcal{G}}/4$. Thus

$$|\overline{\mathcal{Z}}'_{i+1}| \geq \frac{(1-\sqrt{\beta})(1-c)^B}{4}|\mathcal{Z}_{i+1}|_{\mathcal{G}}.$$

Recall that we have fixed a copy S_{Z_i} of $\mathcal{S}_i$ which together with H_{i-1} forms the copy H_i of $\mathcal{H}_i$ in $\mathcal{G}$. As $|S_{Z_i}| = 3B$ if $i \geq 1$ and $|S_{Z_0}| = B$, there are at most $3B^2 n^{B-1}$ copies of $\mathcal{Z}_{i+1}$ that are not disjoint from S_{Z_i}. But by the counting lemma

$$|\mathcal{Z}_{i+1}|_{\mathcal{G}} \geq \frac{1}{2} n^B \prod_{j=2}^{k} d_j^{e_j(\mathcal{Z}_{i+1})}.$$

Therefore, the set $\overline{\mathcal{Z}}_{i+1}$ consisting of all those elements of $\overline{\mathcal{Z}}'_{i+1}$ that avoid S_{Z_i} satisfies

$$|\overline{\mathcal{Z}}_{i+1}| \geq |\overline{\mathcal{Z}}'_{i+1}| - 3B^2 n^{B-1} \geq \frac{(1-\sqrt{\beta})(1-c)^B}{8}|\mathcal{Z}_{i+1}|_{\mathcal{G}} \geq \frac{(1-c)^B}{16}|\mathcal{Z}_{i+1}|_{\mathcal{G}}.$$

Arguing similarly, we also deduce that for each copy Z_{i+1} of $\mathcal{Z}_{i+1}$ in $\overline{\mathcal{Z}}_{i+1}$ there is at least one extension of $Z_i \cup Z_{i+1}$ into a copy of $\mathcal{S}_{i+1}$ that apart from avoiding H_{i-1} also avoids $S_{Z_i} - Z_i$. (Recall that for every such Z_{i+1} there are at least `Ext` extensions that avoid H_{i-1}.) Thus the copy H_i of $\mathcal{H}_i$ and the set $\overline{\mathcal{Z}}_{i+1}$ satisfy $(*)$ for $i+1$. □

References

[1] G. Chen and R. Schelp, Graphs with linearly bounded Ramsey numbers, *J. Combinatorial Theory B* **57** (1993), 138–149.

[2] S.A. Burr and P. Erdős, On the magnitude of generalized Ramsey numbers for graphs. In *Infinite and Finite Sets, Colloquia Mathematica Societatis János Bolyai vol. 10* **1** (1975), 214-240.

[3] V. Chvátal, V. Rödl, E. Szemerédi and W.T. Trotter, Jr., The Ramsey number of a graph with a bounded maximum degree, *J. Combinatorial Theory B* **34** (1983), 239–243.

[4] O. Cooley, Ph.D. thesis, University of Birmingham, in preparation.

[5] O. Cooley, N. Fountoulakis, D. Kühn and D. Osthus, 3-uniform hypergraphs of bounded degree have linear Ramsey numbers, submitted.

[6] D. de Caen, Extension of the theorem of Moon and Moser on complete subgraphs, *Ars Combin.* **16** (1983), 5-10.
[7] P. Erdős and R. Rado, Combinatorial theorems on classifications of subsets of a given set, *Proc. London Mathematical Society* **3** (1952), 417–439.
[8] P. Frankl and V. Rödl, Extremal problems on set systems, *Random Structures & Algorithms* **20** (2002), 131–164.
[9] W.T. Gowers, Hypergraph regularity and the multidimensional Szemerédi theorem, preprint.
[10] W.T. Gowers, Quasirandomness, counting and regularity for 3-uniform hypergraphs, *Combinatorics, Probability & Computing* **15** (2006), 143–184.
[11] R.L. Graham, V. Rödl and A. Ruciński, On graphs with linear Ramsey numbers, *J. Graph Theory* **35** (2000), 176-192.
[12] P.E. Haxell, T. Łuczak, Y. Peng, V. Rödl, A. Ruciński, M. Simonovits and J. Skokan, The Ramsey number for hypergraph cycles I, *J. Combinatorial Theory A* **113** (2006), 67–83.
[13] P.E. Haxell, T. Łuczak, Y. Peng, V. Rödl, A. Ruciński, M. Simonovits and J. Skokan, The Ramsey number for hypergraph cycles II, preprint.
[14] Y. Kohayakawa, V. Rödl and J. Skokan, Hypergraphs, quasi-randomness, and conditions for regularity, *J. Combinatorial Theory A* **97** (2002), 307–352.
[15] J. Komlós, G. Sarkőzy and E. Szemerédi, The blow-up lemma, *Combinatorica* **17** (1997), 109–123.
[16] J. Komlós and M. Simonovits, Szemerédi's Regularity Lemma and its applications in graph theory, *Bolyai Society Mathematical Studies 2, Combinatorics, Paul Erdős is Eighty (Vol. 2)* (D. Miklós, V. T. Sós and T. Szőnyi eds.), Budapest (1996), 295–352.
[17] A. Kostochka and V. Rödl, On Ramsey numbers of uniform hypergraphs with given maximum degree, *J. Combinatorial Theory A* **113** (2006), 1555–1564.
[18] D. Kühn and D. Osthus, Loose Hamilton cycles in 3-uniform hypergraphs of large minimum degree, *J. Combinatorial Theory B* **96** (2006), 767–821.
[19] B. Nagle, S. Olsen, V. Rödl and M. Schacht, On the Ramsey number of sparse 3-graphs, preprint.
[20] B. Nagle and V. Rödl, Regularity properties for triple systems, *Random Structures & Algorithms* **23** (2003), 264–332.
[21] B. Nagle, V. Rödl and M. Schacht, The counting lemma for k-uniform hypergraphs, *Random Structures & Algorithms* **28** (2006), 113–179.
[22] J. Polcyn, V. Rödl, A. Ruciński and E. Szemerédi, Short paths in quasi-random triple systems with sparse underlying graphs, *J. Combinatorial Theory B* **96** (2006), 584–607.
[23] V. Rödl and M. Schacht, Regular partitions of hypergraphs, *Combinatorics, Probability & Computing*, to appear.
[24] V. Rödl and J. Skokan, Regularity lemma for k-uniform hypergraphs, *Random Structures & Algorithms* **25** (2004), 1–42.

Oliver Cooley, Nikolaos Fountoulakis, Daniela Kühn & Deryk Osthus,

School of Mathematics, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

E-mail addresses: {cooleyo,nikolaos,kuehn,osthus}@maths.bham.ac.uk