

TIGHT CYCLES AND REGULAR SLICES IN DENSE HYPERGRAPHS

PETER ALLEN, JULIA BÖTTCHER, OLIVER COOLEY, AND RICHARD MYCROFT

ABSTRACT. We study properties of random subcomplexes of partitions returned by (a suitable form of) the Strong Hypergraph Regularity Lemma, which we call regular slices. We argue that these subcomplexes capture many important structural properties of the original hypergraph. Accordingly we advocate their use in extremal hypergraph theory, and explain how they can lead to considerable simplifications in existing proofs in this field. We also use them for establishing the following two new results.

Firstly, we prove a hypergraph extension of the Erdős-Gallai Theorem: for every $\delta > 0$ every sufficiently large k -uniform hypergraph with at least $(\alpha + \delta)\binom{n}{k}$ edges contains a tight cycle of length αn for each $\alpha \in [0, 1]$.

Secondly, we find (asymptotically) the minimum codegree requirement for a k -uniform k -partite hypergraph, each of whose parts has n vertices, to contain a tight cycle of length αkn , for each $0 < \alpha < 1$.

1. INTRODUCTION

The Szemerédi Regularity Lemma [35] is a powerful tool in extremal graph theory. It is probably fair to say that the majority of the advances in the last decade in extremal graph theory either rely on, or at least were inspired by, the Regularity Lemma. Finding the right extension of this result for hypergraphs turned out to be a challenging endeavour, which culminated in the proof of the Strong Hypergraph Regularity Lemma together with a corresponding Counting Lemma (see [11, 25, 31, 32, 33]), which provide an analogous machinery for extremal problems in hypergraphs. The difficulty with these tools is their technical intricacy, which leads to significant additional complexity in applications of the regularity method in extremal hypergraph theory.

In this paper we argue that often much of this complexity can be avoided by using instead of the complicated structure returned by the Strong Hypergraph Regularity Lemma a more accessible structure, which we call a *regular slice*. We provide a lemma (which is a consequence of the Strong Hypergraph Regularity Lemma) that asserts the existence of such regular slices which inherit enough structure from the original hypergraph to be useful for embedding problems. In addition we provide two applications of this lemma concerning the existence of tight cycles in dense hypergraphs.

In the remainder of this introduction we will first describe these applications, and then provide more details on our lemma concerning regular slices.

1.1. Cycles in dense graphs and hypergraphs. Dirac's Theorem [9] asserts that any n -vertex graph G with minimum degree $\delta(G) \geq n/2$ contains a Hamilton cycle. Bondy [4] extended this result and showed that Hamiltonian n -vertex graphs with at least $n^2/4$ edges are in fact either complete bipartite or pancyclic, that is, they contain

Date: August 20, 2018.

PA was partially supported by FAPESP (Proc. 2010/09555-7); JB by FAPESP (Proc. 2009/17831-7); PA and JB by CNPq (Proc. 484154/2010-9); OC by the DFG (TA 309/2-2); The cooperation of the authors was supported by a joint CAPES-DAAD project (415/ppp-probral/po/D08/11629, Proj. no. 333/09). The authors are grateful to NUMEC/USP, Núcleo de Modelagem Estocástica e Complexidade of the University of São Paulo, and Project MaCLinC/USP, for supporting this research.

cycles of all lengths $\ell \leq n$. Cycles of shorter lengths were also considered by Erdős and Gallai [10], whose celebrated theorem states that for any integer $d \geq 3$, any n -vertex graph with at least $(d-1)(n-1)/2 + 1$ edges contains a cycle of length at least d . This is the best possible bound which is linear in d . The analogous problem for minimum degree conditions was settled by Alon [2], who proved that large graphs with minimum degree at least n/k contain a cycle of length at least $n/(k-1)$. Nikiforov and Schelp [26] considered exact cycle lengths. They showed that for all $\alpha > 0$ there is n_0 such that for $n \geq n_0$ all n -vertex graphs G with $\delta(G) \geq \alpha n$ contain all even ℓ -cycles with $4 \leq \ell \leq \delta(G) + 1$ (they also considered odd cycles for non-bipartite G ; and even sharper results are obtained in [1]).

For hypergraphs much less is known. A *tight cycle*¹ in a k -uniform hypergraph G is a cyclically ordered list of vertices such that every k consecutive vertices form an edge in G . Only recently Rödl, Ruciński and Szemerédi [29] established an approximate extension of Dirac's Theorem for k -uniform hypergraphs (the 3-uniform case was resolved exactly by the same authors in [30]): for all $\gamma > 0$ every sufficiently large k -uniform hypergraph such that every $(k-1)$ -set of vertices lies in at least $(\frac{1}{2} + \gamma)n$ edges has a tight Hamilton cycle. A hypergraph analogue of the Erdős-Gallai Theorem is not yet known. Györi, Katona and Lemons [12] made a first step in this direction. They showed that k -uniform hypergraphs G with more than $(\alpha n - k) \binom{n}{k-1}$ edges contain a tight path on αn vertices. Note that this statement is vacuous for $\alpha > 1/k$.

Our first result is an approximate Erdős-Gallai type result that establishes (up to the error term δ) the best possible linear density bound for the containment of tight cycles of a given length.

Theorem 1. *For every positive δ and every integer $k \geq 3$, there is an integer n_* such that the following holds for all $\alpha \in [0, 1]$. If G is a k -uniform hypergraph on $n \geq n_*$ vertices with $e(G) \geq (\alpha + \delta) \binom{n}{k}$, then G contains a tight cycle of length ℓ for every $\ell \leq \alpha n$ that is divisible by k .*

Partite versions of the graph embedding results mentioned above also have been studied extensively. Moon and Moser [24] showed that the minimum degree condition in Dirac's Theorem can (almost) be halved when G is balanced bipartite. They proved that a bipartite graph G with n vertices in each partition class contains a Hamilton cycle if $\delta(G) \geq (n/2) + 1$. A corresponding (bi)pancyclicity result was established by Schmeichel and Mitchem [34] and balanced k -partite graphs were considered by Chen, Faudree, Gould, Jacobson and Lesniak in [5]. For hypergraphs, Rödl and Ruciński [28] proved that for all $\gamma > 0$ every sufficiently large k -uniform hypergraph which is k -partite with n vertices in each part, such that each $(k-1)$ -set of vertices, one from each partition class, is contained in at least $(\frac{1}{2} + \gamma)n$ edges, contains a tight Hamilton cycle.

Our second result establishes asymptotically best possible codegree bounds for the containment of tight cycles of a given length in the partite setting.

Theorem 2. *For every positive δ and every integer $k \geq 3$ there is an integer n_* such that the following holds for each $\alpha \in [0, 1]$. If G is a k -uniform k -partite hypergraph with parts of size $n \geq n_*$, such that any collection of $k-1$ vertices, one in each of $k-1$ parts of G , lies in at least $(\alpha + \delta)n$ edges of G , then*

(a) *G contains a tight cycle of length ℓ for every $\ell \leq \alpha kn$ that is divisible by k , and*

¹There are several other notions of cycles in hypergraphs. Since our results are about tight cycles we concentrate on these here.

(b) if $\alpha \geq \frac{1}{2}$ then G contains a tight cycle of length ℓ for every $\ell \leq (1 - \delta)kn$ that is divisible by k .

We believe that with substantial additional work part (b) of this Theorem would also follow from the proof in [28]. However, our new tools allow for a much simpler proof.

1.2. Regular slices. Briefly, given an n -vertex graph G the Regularity Lemma returns an equipartition of the vertices of G into a constant number of parts such that almost all of the bipartite subgraphs induced by pairs of parts approximate random graphs. This so-called *regular partition* can be represented by a constant size weighted graph, the *weighted reduced graph*, in which the (weighted) density of copies of any small graph approximates the density of copies in the original graph. In fact, we can use a regular partition and the reduced graph also to find copies of large bounded degree subgraphs in G , among other things. In this way we can reduce difficult extremal graph problems to relatively simple problems on the weighted reduced graph. This approach to extremal graph theory is called the regularity method.

It is then natural to ask for an analogue of regular partitions and reduced graphs for hypergraphs. However, the partition returned by the Strong Hypergraph Regularity Lemma is no longer given by a vertex partition and cannot be represented by a weighted hypergraph. Instead, the partition has the structure of a weighted multi-complex with edges of sizes up to and including k , which makes applications of the Strong Hypergraph Regularity Lemma significantly harder. Whilst simpler forms of hypergraph regularity do exist (see [7, 19]), they are substantially less powerful for hypergraph embedding.

The *Regular Slice Lemma* (Lemma 6) that we propose in this paper for approaching problems in extremal hypergraph theory bypasses these difficulties by taking a random subcomplex of the weighted multi-complex returned by the Strong Hypergraph Regularity Lemma (after a considerable number of suitable modifications). Related ideas were applied already by Haxell, Łuczak, Peng, Rödl, Ruciński and Skokan [13]. The advantage of this random subcomplex, which we call a *regular slice*, is that it does correspond to a vertex partition and a weighted reduced hypergraph. The disadvantage of this approach is that a regular slice discards most of the hyperedges of the original hypergraph. However, as our lemma asserts, a lot of information about the original hypergraph is still captured: edge densities, and more generally small subgraph densities, minimum degree and codegree conditions, and edge densities in link hypergraphs are approximated in regular slices. Since these are exactly the type of conditions which typically appear in extremal hypergraph theory we believe that regular slices will be useful for solving many problems in this area.

1.3. Structure of the paper. In the next section we give the definitions needed to work with hypergraphs, then in Section 3 we discuss lower bounds for Theorems 1 and 2. Section 4 introduces the definitions required for and the statement of the Regular Slice Lemma. It also includes the statement of a Cycle Embedding Lemma, which enables us to apply the Regular Slice Lemma to prove our two main theorems. In Sections 5 and 6 we assume these results to prove Theorems 1 and 2 respectively. Section 7 deals with the definitions and machinery of hypergraph regularity, while Section 8 contains the proof of the Regular Slice Lemma, and in Section 9, we prove the Cycle Embedding Lemma. We conclude with a discussion of open problems and possible further applications in Section 10. Finally, Appendix A contains the derivation of the version of the Hypergraph Counting Lemma we use, and Appendix B contains the proof of a variant of the Strong Hypergraph Regularity Lemma which is used in the proof of the Regular Slice Lemma.

2. PRELIMINARIES

A *hypergraph* consists of a vertex set V and an edge set E , where each edge $e \in E$ is a subset of V . We frequently identify a hypergraph \mathcal{H} with its edge set, writing $e \in \mathcal{H}$ to mean that e is an edge of \mathcal{H} , and $|\mathcal{H}|$ for $e(\mathcal{H})$, the number of edges of \mathcal{H} . A *k-uniform hypergraph*, or *k-graph*, is a hypergraph in which every edge has size k ; it is *l-partite* for $l \geq k$ if there exists a partition of the vertex set into l parts so that every edge has at most one vertex in each part. Given a hypergraph \mathcal{H} , we denote by $\mathcal{H}^{(i)}$ the *i-graph* on $V(\mathcal{H})$ formed by the *i-edges* of \mathcal{H} , i.e. the edges of cardinality i , and write $e_i(\mathcal{H}) := |\mathcal{H}^{(i)}|$ for the number of such edges. The *order* of a hypergraph is the number of vertices $v(\mathcal{H}) := |V(\mathcal{H})|$. The *degree* of a set of i vertices in a k -graph, where $1 \leq i \leq k - 1$, is the number of edges containing all i vertices. In the case when $i = k - 1$ we talk of the *codegree*. If \mathcal{H} is a hypergraph on V , and $V' \subseteq V$, then the induced subgraph $\mathcal{H}[V']$ is the hypergraph on V' whose edges are all $e \in \mathcal{H}$ with $e \subseteq V'$. A *k-complex* \mathcal{H} is a hypergraph in which all edges have size at most k , and which is down-closed: that is, if $e \in \mathcal{H}$ is an edge, and $e' \subseteq e$, then $e' \in \mathcal{H}$. Note that for any non-empty complex \mathcal{H} we have $\emptyset \in \mathcal{H}$. We prefer the terms *subgraph* and *subcomplex* to subhypergraph, sub-*i*-graph, or sub-*k*-complex.

Throughout this paper we will maintain the convention of using normal letters for uniform hypergraphs, and calligraphic letters for not necessarily uniform hypergraphs (which will usually be complexes). The exceptions to this rule are the definition of $\mathcal{H}^{(i)}$ above and of \mathcal{H}_X in Section 4: both define a specific uniform subgraph of a not necessarily uniform hypergraph \mathcal{H} .

Let G be a k -graph. A *tight path* in G is an ordered list of distinct vertices of G , such that each set of k consecutive vertices induces an edge of G ; a *tight cycle* is a cyclically ordered list with the same property. The *length* of a tight path or cycle is the number of edges in the path or cycle; the number of vertices in a tight cycle is therefore equal to its length, whilst the number of vertices in a tight path is $k - 1$ greater than its length. We will denote the tight cycle or path on ℓ vertices by C_ℓ or P_ℓ respectively.

Any k -graph G naturally corresponds to a k -complex \mathcal{G} , whose edges are all subsets $e' \subseteq e$ of edges $e \in G$. We refer to \mathcal{G} as the complex *generated by the down-closure of G* . We can then work with the k -complex \mathcal{G} to obtain useful information about the k -graph G . We say that a set S of $k + 1$ vertices of G is *supported on G* if every subset $S' \subseteq S$ of size k is an edge of G (i.e. S forms a clique in G). Similarly, we say that a $(k + 1)$ -graph G' on $V(G)$ is *supported on G* if every edge of G' is supported on G . Note in particular that if \mathcal{H} is a k -complex, then $\mathcal{H}^{(i)}$ is supported on $\mathcal{H}^{(i-1)}$ for any $1 \leq i \leq k$.

We write $[r]$ to denote the set $\{1, 2, \dots, r\}$. For a set A , we often write $\binom{A}{r}$ to denote the collection of subsets of A of size r . We use the notation $x = a \pm b$ to denote $a - b \leq x \leq a + b$. If $S(x, y)$ is a statement, we say ' $S(x, y)$ holds for $x \ll y$ ' if for any $y > 0$ there exists $x_0 > 0$ such that for any $0 < x \leq x_0$ the statement $S(x, y)$ holds; similar expressions with more constants are defined analogously. Throughout this paper floors and ceilings are often omitted where they do not affect the argument.

3. LOWER BOUNDS

3.1. Lower bounds for Theorem 1. Given a k -graph H , we define $\text{ex}(n, H)$ to be the maximum number of edges an n -vertex k -graph can have without containing H as a subgraph. Then Theorem 1 says that $\text{ex}(n, P_{\alpha n}), \text{ex}(n, C_{\alpha n}) \leq (\alpha + o(1)) \binom{n}{k}$. We will now establish several lower bounds on $\text{ex}(n, P_{\alpha n})$ and $\text{ex}(n, C_{\alpha n})$.

Györi, Katona and Lemons [12] (who were interested in the case $\alpha n = \ell$ where ℓ is constant) observed that given any set system \mathcal{S} on n vertices with sets of size at most $\alpha n - 1$, no pair of which intersect in more than $k - 2$ vertices, one can construct a k -graph G without any tight path on $\alpha n = \ell$ vertices by taking every k -set contained in a set of \mathcal{S} . In the event that all members of \mathcal{S} have size exactly $\alpha n - 1$ and every $(k - 1)$ -set in $[n]$ is in some member of \mathcal{S} , \mathcal{S} is a combinatorial design and the corresponding G has $(\alpha n - k) \binom{n}{k-1} / k \approx \alpha \binom{n}{k}$ edges. A celebrated recent result of Keevash [16] states that for each fixed $\alpha n = \ell$ and k , these combinatorial designs exist for all sufficiently large n satisfying a necessary divisibility condition. Györi, Katona and Lemons (writing before [16] appeared) used Rödl's solution [27] of the Erdős-Hanani Conjecture to prove, for each $\delta > 0$, the existence of hypergraphs with $(1 - \delta) \alpha \binom{n}{k}$ edges and no ℓ -vertex tight path, provided that n is sufficiently large compared to α . This provides the lower bound $\text{ex}(n, P_{\alpha n}) \geq (1 - o(1)) \alpha \binom{n}{k}$. Unfortunately both Keevash's and Rödl's proofs only work when n is superexponentially large in ℓ , and it is easy to see that in fact no designs (or even approximations to designs) can exist if $\ell \gg n^{1/2}$, so that this construction does not work in the constant α range where Theorem 1 gives a non-trivial result.

A very similar construction gives G without tight cycles on ℓ or more vertices. Namely, if there exists a set system \mathcal{S} on $n - 1$ vertices with all sets of size $\ell - 2$, covering each $(k - 1)$ -set exactly once, then we can construct a graph G whose edges are all k -sets contained in members of \mathcal{S} together with all k -sets containing the n th vertex. We have $e(G) = \frac{\ell-1}{k} \binom{n-1}{k-1}$, and it is easy to check that G contains no tight cycle on ℓ or more vertices. However, if we only want to exclude tight cycles on exactly ℓ vertices, we obtain

$$\text{ex}(n, C_\ell) \geq \frac{1}{2} \left(\frac{1}{k!}\right)^{\frac{\ell}{\ell-1}} n^{k-1 + \frac{k-1}{\ell-1}}$$

by deleting an edge from each copy of C_ℓ in the random k -graph with edge probability

$$p = \left(\frac{2}{k!} n^{k-\ell}\right)^{\frac{1}{\ell-1}}.$$

A simple lower bound on both $\text{ex}(n, P_{\alpha n})$ and $\text{ex}(n, C_{\alpha n})$ is provided by the following construction. Let G be the k -graph on vertex set $V = A \dot{\cup} B$, where $|A| = \alpha n / k - 1$ and $|B| = (1 - \alpha/k)n + 1$, and whose edge set is $E(G) := \{e \in \binom{V}{k} : e \cap A \neq \emptyset\}$. It is easy to see that the longest tight path in G has at most $k|A| + k - 1 < \alpha n$ vertices (the longest tight cycle is even slightly shorter). For large n , the number of edges in G is

$$\left(1 - \left(1 - \frac{\alpha}{k}\right)^k\right) \binom{n}{k} + o(n^k) = \left(\sum_{i=1}^k (-1)^{i-1} \binom{k}{i} \left(\frac{\alpha}{k}\right)^i\right) \binom{n}{k} + o(n^k).$$

Hence $\text{ex}(n, C_{\alpha n}), \text{ex}(n, P_{\alpha n}) \geq (\alpha + O(\alpha^2)) \binom{n}{k}$, matching the first order asymptotics in Theorem 1 for small α .

Moreover, for each $2 \leq r \leq k$, if $\alpha n r / k > (1 - \alpha)n + 1$, then the k -graph G on vertex set $V = A \dot{\cup} B$, where $|A| = \alpha n - 1$ and $|B| = (1 - \alpha)n + 1$, with edge set consisting of all k -sets either contained in A or with at least r vertices in B , contains no αn -vertex tight cycle. The densest of these is the construction with $r = 2$, which is permissible when $1 - \alpha$ is small, giving $\text{ex}(n, C_{\alpha n}) \geq (1 - k\alpha^{k-1}(1 - \alpha) + O(1/n)) \binom{n}{k}$. (Minor modifications of the construction give the same asymptotic bound for $\text{ex}(n, P_{\alpha n})$.) This tends to 1 as $1 - \alpha$ tends to zero, but does not match the first order asymptotics of Theorem 1.

Finally, we note that the divisibility condition on ℓ in Theorem 1 is necessary, at least for $\alpha \leq k!/k^k$, as can be seen by considering the complete k -partite k -graph with vertex classes of size n/k .

3.2. Lower bounds for Theorem 2. Clearly Theorem 2 part (b) is asymptotically best possible. For part (a), consider the following construction. Given the k -partite vertex set $V_1 \cup \dots \cup V_k$ with $|V_i| = n$ for each i and $\alpha \leq \frac{1}{2}$, we partition each V_i into sets V_i^0 and V_i^1 of sizes respectively $(1 - \alpha)n$ and αn . We let a set of k vertices e , with one vertex in each of $V_1^{j_1}, \dots, V_k^{j_k}$, where $j_i \in \{0, 1\}$ for each $i \in [k]$, be an edge of the k -graph G if and only if $j_1 + \dots + j_k$ is odd. It is easy to check that the partite codegree condition of Theorem 2 is almost satisfied, and similarly easy to check that a tight cycle with one edge e whose vertices are in $V_1^{j_1}, \dots, V_k^{j_k}$ must have all its vertices in $V_1^{j_1}, \dots, V_k^{j_k}$, from which it follows that we cannot have tight cycles on more than $k \min(|V_1^{j_1}|, \dots, |V_k^{j_k}|) = k\alpha n$ vertices in G . This shows that Theorem 2 part (a) is asymptotically best possible.

4. REGULAR SLICES

In this section we state our Regular Slice Lemma. We first need to introduce the required notation. Much of this notation is standard, and will also be needed in Section 7 for the Strong Hypergraph Regularity Lemma. We will then state a Cycle Embedding Lemma, which allows us to use the nice properties of regular slices to find tight cycles in a hypergraph satisfying the appropriate conditions.

4.1. Multipartite hypergraphs and regular complexes. Let \mathcal{P} partition a vertex set V into parts V_1, \dots, V_s . Then we say that a subset $S \subseteq V$ is \mathcal{P} -partite if $|S \cap V_i| \leq 1$ for every $i \in [s]$. Similarly, we say that a hypergraph \mathcal{H} is \mathcal{P} -partite if all of its edges are \mathcal{P} -partite. In this case we refer to the parts of \mathcal{P} as the *vertex classes* of \mathcal{H} . As defined previously for k -graphs, we say that a hypergraph \mathcal{H} is s -partite if there is some partition \mathcal{P} of $V(\mathcal{H})$ into s parts for which \mathcal{H} is \mathcal{P} -partite.

Let \mathcal{H} be a \mathcal{P} -partite hypergraph. Then for any $A \subseteq [s]$ we write V_A for $\bigcup_{i \in A} V_i$. The *index* of a \mathcal{P} -partite set $S \subseteq V$ is $i(S) := \{i \in [s] : |S \cap V_i| = 1\}$. We write \mathcal{H}_A to denote the collection of edges in \mathcal{H} with index A . So \mathcal{H}_A can be regarded as an $|A|$ -partite $|A|$ -graph on vertex set V_A , with vertex classes V_i for $i \in A$. It is often convenient to refer to the subgraph induced by a set of vertex classes rather than with a given index; if X is a k -set of vertex classes of \mathcal{H} we write \mathcal{H}_X for the k -partite k -uniform subgraph of $\mathcal{H}^{(k)}$ induced by $\bigcup X$, whose vertex classes are the members of X . Note that $\mathcal{H}_X = \mathcal{H}_{\{i: V_i \in X\}}$. In a similar manner we write $\mathcal{H}_{X<}$ for the k -partite hypergraph on vertex set $\bigcup X$ whose edge set is $\bigcup_{X' \subsetneq X} \mathcal{H}_{X'}$. Note that if \mathcal{H} is a complex, then $\mathcal{H}_{X<}$ is a $(k - 1)$ -complex because X is a k -set.

Let $i \geq 2$, let H_i be any i -partite i -graph, and let H_{i-1} be any i -partite $(i - 1)$ -graph, on a common vertex set V partitioned into i common vertex classes. We denote by $K_i(H_{i-1})$ the i -partite i -graph on V whose edges are all i -sets in V which are supported on H_{i-1} (i.e. induce a copy of the complete $(i - 1)$ -graph K_{i-1}^{i-1} on i vertices in H_{i-1}). The *density of H_i with respect to H_{i-1}* is then defined to be

$$d(H_i|H_{i-1}) := \frac{|K_i(H_{i-1}) \cap H_i|}{|K_i(H_{i-1})|}$$

if $|K_i(H_{i-1})| > 0$. For convenience we take $d(H_i|H_{i-1}) := 0$ if $|K_i(H_{i-1})| = 0$. So $d(H_i|H_{i-1})$ is the proportion of copies of K_i^{i-1} in H_{i-1} which are also edges of H_i . When H_{i-1} is clear from the context, we simply refer to $d(H_i|H_{i-1})$ as the *relative density of H_i* . More generally, if $\mathbf{Q} := (Q_1, Q_2, \dots, Q_r)$ is a collection of r not necessarily disjoint subgraphs of H_{i-1} , we define $K_i(\mathbf{Q}) := \bigcup_{j=1}^r K_i(Q_j)$ and

$$d(H_i|\mathbf{Q}) := \frac{|K_i(\mathbf{Q}) \cap H_i|}{|K_i(\mathbf{Q})|}$$

if $|K_i(\mathbf{Q})| > 0$. Similarly as before we take $d(H_i|\mathbf{Q}) := 0$ if $|K_i(\mathbf{Q})| = 0$. We say that H_i is (d_i, ε, r) -regular with respect to H_{i-1} if we have $d(H_i|\mathbf{Q}) = d_i \pm \varepsilon$ for every r -set \mathbf{Q} of subgraphs of H_{i-1} such that $|K_i(\mathbf{Q})| > \varepsilon|K_i(H_{i-1})|$. We often refer to $(d_i, \varepsilon, 1)$ -regularity simply as (d_i, ε) -regularity; also, we say simply that H_i is (ε, r) -regular with respect to H_{i-1} to mean that there exists some d_i for which H_i is (d_i, ε, r) -regular with respect to H_{i-1} . Finally, given an i -graph G whose vertex set contains that of H_{i-1} , we say that G is (d_i, ε, r) -regular with respect to H_{i-1} if the i -partite subgraph of G induced by the vertex classes of H_{i-1} is (d_i, ε, r) -regular with respect to H_{i-1} . Similarly as before, when H_{i-1} is clear from the context, we refer to the relative density of this i -partite subgraph of G with respect to H_{i-1} as the *relative density of G* .

Now let \mathcal{H} be an s -partite k -complex on vertex classes V_1, \dots, V_s , where $s \geq k \geq 3$. Recall that this means that if $e \in \mathcal{H}$ and $e' \subseteq e$ then $e' \in \mathcal{H}$. So if $e \in \mathcal{H}^{(i)}$ for some $2 \leq i \leq k$, then the vertices of e induce a copy of K_i^{i-1} in $\mathcal{H}^{(i-1)}$. This means that for any index $A \in \binom{[s]}{i}$ the density $d(\mathcal{H}^{(i)}[V_A]|\mathcal{H}^{(i-1)}[V_A])$ can be regarded as the proportion of ‘possible edges’ of $\mathcal{H}^{(i)}[V_A]$ which are indeed edges. (Here a ‘possible edge’ is a subset of $V(\mathcal{H})$ of index A all of whose proper subsets are edges of \mathcal{H}). We therefore say that \mathcal{H} is $(d_k, \dots, d_2, \varepsilon_k, \varepsilon, r)$ -regular if

- (a) for any $2 \leq i \leq k-1$ and any $A \in \binom{[s]}{i}$, the induced subgraph $\mathcal{H}^{(i)}[V_A]$ is (d_i, ε) -regular with respect to $\mathcal{H}^{(i-1)}[V_A]$, and
- (b) for any $A \in \binom{[s]}{k}$, the induced subgraph $\mathcal{H}^{(k)}[V_A]$ is (d_k, ε_k, r) -regular with respect to $\mathcal{H}^{(k-1)}[V_A]$.

So each constant d_i approximates the relative density of each subgraph $\mathcal{H}^{(i)}[V_A]$ for $A \in \binom{[s]}{i}$ for which $\mathcal{H}^{(i)}[V_A]$ is non-empty. For a $(k-1)$ -tuple $\mathbf{d} = (d_k, \dots, d_2)$ we write $(\mathbf{d}, \varepsilon_k, \varepsilon, r)$ -regular to mean $(d_k, \dots, d_2, \varepsilon_k, \varepsilon, r)$ -regular. A regular complex is the correct notion of ‘approximately random’ for hypergraph regularity.

4.2. Regular slices and the reduced k -graph. The Regular Slice Lemma says that any k -graph G admits a ‘regular slice’. This is a multipartite $(k-1)$ -complex \mathcal{J} whose vertex classes have equal size, which is regular, and which moreover has the property that G is regular with respect to \mathcal{J} . The first two of these conditions are formalised in the following definition: we say that a $(k-1)$ -complex \mathcal{J} is (t_0, t_1, ε) -equitable if it has the following properties.

- (a) \mathcal{J} is \mathcal{P} -partite for some \mathcal{P} which partitions $V(\mathcal{J})$ into t parts, where $t_0 \leq t \leq t_1$, of equal size. We refer to \mathcal{P} as the *ground partition* of \mathcal{J} , and to the parts of \mathcal{P} as the *clusters* of \mathcal{J} .
- (b) There exists a *density vector* $\mathbf{d} = (d_{k-1}, \dots, d_2)$ such that for each $2 \leq i \leq k-1$ we have $d_i \geq 1/t_1$ and $1/d_i \in \mathbb{N}$, and the $(k-1)$ -complex \mathcal{J} is $(\mathbf{d}, \varepsilon, \varepsilon, 1)$ -regular.

For any k -set X of clusters of \mathcal{J} , we write $\hat{\mathcal{J}}_X$ for the k -partite $(k-1)$ -graph $\mathcal{J}_{X^{(k-1)}}$. For reasons that will become apparent later, we sometimes refer to $\hat{\mathcal{J}}_X$ as a *polyad*. The following well-known fact, a special case of the Dense Counting Lemma (see [20, Theorem 6.5]), tells us approximately the number of edges in each layer of \mathcal{J} , and also how many k -sets are supported on any polyad. Note that the definition of an equitable complex implies that $\varepsilon \ll 1/t_1 \leq d_i$ in this fact.

Fact 3. *Suppose that $1/m_0 \ll \varepsilon \ll 1/t_1, 1/t_0, \beta, 1/k$, and that \mathcal{J} is a (t_0, t_1, ε) -equitable $(k-1)$ -complex with density vector (d_{k-1}, \dots, d_2) whose clusters each have size $m \geq m_0$.*

Let X be a set of k clusters of \mathcal{J} . Then

$$|K_k(\hat{\mathcal{J}}_X)| = (1 \pm \beta)m^k \prod_{i=2}^{k-1} d_i^{\binom{k}{i}},$$

and for any proper subset $X' \subsetneq X$ we have

$$|\mathcal{J}_{X'}| = (1 \pm \beta)m^{|X'|} \prod_{i=2}^{|X'|} d_i^{\binom{|X'|}{i}}.$$

□

Given a (t_0, t_1, ε) -equitable $(k-1)$ -complex \mathcal{J} and a k -graph G on $V(\mathcal{J})$, we say that G is (ε_k, r) -regular with respect to a k -set X of clusters of \mathcal{J} if there exists some d such that G is (d, ε_k, r) -regular with respect to the polyad $\hat{\mathcal{J}}_X$. We also write $d_{\mathcal{J}}^*(X)$ for the relative density of G with respect to $\hat{\mathcal{J}}_X$, or simply $d^*(X)$ if \mathcal{J} is clear from the context, which will usually be the case in applications. Note in particular that Fact 3 then implies that the number of edges of G which are supported on $\hat{\mathcal{J}}_X$ is approximately $d^*(X)m^k \prod_{i=2}^{k-1} d_i^{\binom{k}{i}}$. We now give the key definition of the Regular Slice Lemma.

Definition 4 (Regular slice). *Given $\varepsilon, \varepsilon_k > 0$, $r, t_0, t_1 \in \mathbb{N}$, a k -graph G and a $(k-1)$ -complex \mathcal{J} on $V(G)$, we call \mathcal{J} a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G if \mathcal{J} is (t_0, t_1, ε) -equitable and G is (ε_k, r) -regular with respect to all but at most $\varepsilon_k \binom{t}{k}$ of the k -sets of clusters of \mathcal{J} , where t is the number of clusters of \mathcal{J} .*

It will sometimes be convenient not to specify all of the parameters: we may write that \mathcal{J} is $(\cdot, \cdot, \varepsilon)$ -equitable, or is a $(\cdot, \cdot, \varepsilon, \varepsilon_k, r)$ -slice for G , if we do not wish to specify t_0 or t_1 (as will be the case if we specify instead the density vector \mathbf{d} and the number of clusters t).

Given a regular slice \mathcal{J} for a k -graph G , it will be important to know the relative densities $d^*(X)$ for k -sets X of clusters of \mathcal{J} . To keep track of these we make the following definition.

Definition 5 (Weighted reduced k -graph). *Given a k -graph G and a (t_0, t_1, ε) -equitable $(k-1)$ -complex \mathcal{J} on $V(G)$, we let $R_{\mathcal{J}}(G)$ be the complete weighted k -graph whose vertices are the clusters of \mathcal{J} , and where each edge X is given weight $d^*(X)$ (in particular, the weight is in $[0, 1]$). When \mathcal{J} is clear from the context we often simply write $R(G)$ instead of $R_{\mathcal{J}}(G)$.*

In applications we will usually take \mathcal{J} to be a regular slice for G . In this case we will want to ensure that G is regular with respect to all k -sets X of clusters of \mathcal{J} with $d^*(X) > 0$, which can be achieved by the simple expedient of deleting all edges of G lying in k -sets with respect to which G is not regular; the definition of a regular slice implies that there are few such edges. The reason that we do not specify this in the definition of a regular slice, or specify that only k -sets with respect to which G is regular are given positive weight in the weighted reduced graph, is that we will also make use of this definition in the process of proving that a certain equitable $(k-1)$ -complex is in fact a regular slice. Note that although different choices of \mathcal{J} may well produce different weighted reduced k -graphs (and only some of these will have ‘good’ properties), since \mathcal{J} will always be clear from the context we will speak of ‘the’ reduced k -graph $R(G)$ of G .

In general, it is not very helpful to know that \mathcal{J} is a regular slice for a k -graph G ; the reduced graph of G with respect to \mathcal{J} does not necessarily resemble G in the way that the reduced 2-graph of a 2-graph H with respect to a Szemerédi partition resembles H . The

Regular Slice Lemma states that there is a regular slice \mathcal{J} with respect to which $R(G)$ does resemble G , in the sense that densities of small subgraphs (part (a) of the Regular Slice Lemma) and degree conditions (part (b)) are preserved. Furthermore, all vertices of G are in some sense ‘represented’ on \mathcal{J} (part (c))—this is useful for embedding spanning subgraphs. In order to make this precise, we need the following further definitions.

Given hypergraphs G and H , we write $n_H(G)$ for the number of labelled copies of H in G . In order to extend this definition to weighted k -graphs G , we let

$$n_H(G) := \sum_{\phi: V(H) \rightarrow V(G)} \prod_{e \in E(H)} d^*(\phi(e))$$

where ϕ ranges over all injective maps and d^* is the weight function on $E(G)$. We then define the H -density of G as

$$d_H(G) := \frac{n_H(G)}{\binom{v(G)}{v(H)} \cdot v(H)!}.$$

In other words, $d_H(G)$ is the expectation of $\prod_{e \in E(H)} d^*(\phi(e))$ for an injective map $\phi : V(H) \rightarrow V(G)$ chosen uniformly at random. In the special case that H consists of a single edge, we will write simply $d(G)$, and speak of the *density of G* or the *edge-density of G* .

Now let G be a k -graph on n vertices. Given a set $S \subseteq V(G)$ of size j for some $1 \leq j \leq k - 1$, and a subset $X \subseteq V(G)$, the *relative degree* $\overline{\text{deg}}(S; G, X)$ of S in X with respect to G is defined to be

$$\overline{\text{deg}}(S; G, X) := \frac{|\{e \in G[S \cup X] : S \subseteq e\}|}{\binom{|X \setminus S|}{k-j}}.$$

In other words, $\overline{\text{deg}}(S; G, X)$ is the proportion of k -sets of vertices of $G[X \cup S]$ extending S which are in fact edges of G . To extend this definition to weighted k -graphs G , we replace the number of edges in $G[X \cup S]$ including S with the sum of the weights of edges of $G[X \cup S]$ including S . Finally, if \mathcal{S} is a collection of j -sets in $V(G)$, then $\overline{\text{deg}}(\mathcal{S}; G, X)$ is defined to be the mean of $\overline{\text{deg}}(S; G, X)$ over all sets $S \in \mathcal{S}$.

Given a k -graph G and distinct ‘root’ vertices v_1, \dots, v_ℓ of G , and a k -graph H equipped with a set of distinct ‘root’ vertices x_1, \dots, x_ℓ , we define the *number of labelled rooted copies of H in G* , written $n_H(G; v_1, \dots, v_\ell)$, to be the number of injective maps from $V(H)$ to $V(G)$ which embed H in G and take x_j to v_j for each $1 \leq j \leq \ell$. Then the *density of rooted copies of H in G* is defined to be

$$d_H(G; v_1, \dots, v_\ell) := \frac{n_H(G; v_1, \dots, v_\ell)}{\binom{v(G) - \ell}{v(H) - \ell} \cdot (v(H) - \ell)!}.$$

This density has a natural probabilistic interpretation: choose uniformly at random an injective map $\psi : V(H) \rightarrow V(G)$ such that $\psi(x_j) = v_j$ for every $j \in [\ell]$. Then $d_H(G; v_1, \dots, v_\ell)$ is the probability that ψ embeds H in G . Next, we define \mathcal{H}^{skel} to be the $(k - 1)$ -complex on $V(H) - \ell$ vertices which is obtained from the complex \mathcal{H} generated by the down-closure of H by deleting the vertices x_1, \dots, x_ℓ (and all edges containing them) and deleting all edges of size k . Given a (t_0, t_1, ε) -equitable $(k - 1)$ -complex \mathcal{J} on $V(G)$, the *number of rooted copies of H supported by \mathcal{J}* , written $n_H(G; v_1, \dots, v_\ell, \mathcal{J})$, is defined to be the number of labelled rooted copies of H in G such that each vertex of \mathcal{H}^{skel} lies in a distinct cluster of \mathcal{J} and the image of \mathcal{H}^{skel} is in \mathcal{J} (but note that we do *not* require the edges involving v_1, \dots, v_ℓ to be contained in or supported by \mathcal{J} , and indeed typically they will not be). We also define $n'_{\mathcal{H}^{skel}}(\mathcal{J})$ to be the number of labelled

copies of \mathcal{H}^{skel} in \mathcal{J} with each vertex of \mathcal{H}^{skel} embedded in a distinct cluster of \mathcal{J} . Then the *density* $d_H(G; v_1, \dots, v_\ell, \mathcal{J})$ of rooted copies of H in G supported by \mathcal{J} is then defined by

$$d_H(G; v_1, \dots, v_\ell, \mathcal{J}) := \frac{n_H(G; v_1, \dots, v_\ell, \mathcal{J})}{n'_{\mathcal{H}^{skel}}(\mathcal{J})}.$$

Again we have a natural probabilistic interpretation: let $\psi : V(\mathcal{H}^{skel}) \rightarrow V(G)$ be an injective map chosen uniformly at random, and extend ψ to a map $\psi' : V(H) \rightarrow V(G)$ by taking $\psi'(x_i) = v_i$ for each $i \in [\ell]$. Then $d_H(G; v_1, \dots, v_\ell, \mathcal{J})$ is the conditional probability that ψ' embeds H in G given that ψ embeds \mathcal{H}^{skel} in \mathcal{J} with each vertex of \mathcal{H}^{skel} embedded in a different cluster of \mathcal{J} .

We can now state the Regular Slice Lemma. We remark that for many applications it suffices to take $q = s = 1$, with \mathcal{Q} being the trivial partition of V with one part.

Lemma 6 (Regular Slice Lemma). *Let $k \geq 3$ be a fixed integer. For all positive integers q, t_0 and s , positive ε_k and all functions $r : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$, there are integers t_1 and n_0 such that the following holds for all $n \geq n_0$ which are divisible by $t_1!$. Let V be a set of n vertices, and suppose that G_1, \dots, G_s are edge-disjoint k -graphs on V , and that \mathcal{Q} is a partition of V into at most q parts of equal size. Then there exists a $(k-1)$ -complex \mathcal{J} on V which is a $(t_0, t_1, \varepsilon(t_1), \varepsilon_k, r(t_1))$ -regular slice for each G_i , such that the ground partition \mathcal{P} of \mathcal{J} refines \mathcal{Q} , and such that \mathcal{J} has the following additional properties.*

- (a) *For each $1 \leq i \leq s$, any k -graph H with $v(H) \leq 1/\varepsilon_k$ and each set X of at least $\varepsilon_k t$ clusters of \mathcal{J} (where t is the total number of clusters of \mathcal{J}), we have*

$$\left| d_H(R(G_i)[X]) - d_H\left(G_i\left[\bigcup X\right]\right) \right| < \varepsilon_k.$$

- (b) *For each $1 \leq i \leq s$, each $1 \leq j \leq k-1$, each set Y of j clusters of \mathcal{J} , and each set X of clusters of \mathcal{J} for which $\bigcup X$ is the union of some parts of \mathcal{Q} , we have*

$$\left| \overline{\deg}(Y; R(G_i), X) - \overline{\deg}(\mathcal{J}_Y; G_i, \bigcup X) \right| < \varepsilon_k.$$

- (c) *For each $1 \leq i \leq s$, each $1 \leq \ell \leq 1/\varepsilon_k$, each k -graph H equipped with a set of distinct root vertices x_1, \dots, x_ℓ such that $v(H) \leq 1/\varepsilon_k$, and any distinct vertices v_1, \dots, v_ℓ in V , we have*

$$\left| d_H(G_i; v_1, \dots, v_\ell, \mathcal{J}) - d_H(G_i; v_1, \dots, v_\ell) \right| < \varepsilon_k.$$

To understand this lemma consider first the special case of regularising only one graph $G = G_1$, i.e. $s = 1$, with \mathcal{Q} being the trivial partition with one part, i.e. $q = 1$, and where (for properties (a) and (c)) H consists of a single edge (in the latter case, rooted at one vertex v). Then the result is that \mathcal{J} is a regular slice for G , with not too few clusters (bounded below by t_0) but also not too many (bounded above by t_1), with very strong regularity properties (typically ε is much smaller than any density d_i) and G is (ε_k, r) -regular with respect to \mathcal{J} . This looks so far very much like the ‘usual’ hypergraph regularity (see Lemma 20 later for a statement) except that the statement that G is regular with respect to \mathcal{J} says something only about a very small part of G . Then property (a) states that $R(G)$ ‘looks like’ G in that edge densities on large sets agree. Property (b) tells us that degree conditions on G transfer to $R(G)$. Finally property (c) says: for every vertex v of G , the fraction of $(k-1)$ -sets of $V(G)$ which make edges with v is close to the fraction of $(k-1)$ -edges of \mathcal{J} which make edges with v . That is, vertex degrees in G are inherited when we consider only extensions supported on \mathcal{J} .

We note that it is often necessary to estimate (for example) the number of edges of G which have one vertex in each of some pairwise-disjoint sets X_1, \dots, X_k rather than just

the number of edges in a set X : provided that each of the X_i is large, it follows by use of the inclusion-exclusion principle and the ability conferred by (a) to estimate the number of edges in a large set X that we can also make such estimates.

For embedding subgraphs, it is important to distinguish dense regular k -sets (i.e. k -sets on which G is dense and regular with respect to \mathcal{J}). In order to state our Cycle Embedding Lemma we therefore give the following variation of the definition of the reduced graph.

Definition 7 (The d -reduced k -graph). *Let G be a k -graph and let \mathcal{J} be a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G . Then for $d > 0$ we define the d -reduced k -graph $R_d(G)$ to be the k -graph whose vertices are the clusters of \mathcal{J} and whose edges are all k -sets X of clusters of \mathcal{J} such that G is (ε_k, r) -regular with respect to X and $d^*(X) \geq d$.*

The next lemma states that for regular slices \mathcal{J} from Lemma 6, H -densities and degrees are also preserved by $R_d(G)$, allowing us to work with this structure also.

Lemma 8. *Let G be a k -graph and let \mathcal{J} be a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G with t clusters. Also let X be a set of clusters of \mathcal{J} . Then for any k -graph H we have*

$$d_H(R_d(G)[X]) \geq d_H(R(G)[X]) - d - \frac{\varepsilon_k e(H) \binom{t}{k}}{\binom{|X|}{k}},$$

and for any set Y of at most $k - 1$ clusters of \mathcal{J} we have

$$\overline{\deg}(Y; R_d(G), X) \geq \overline{\deg}(Y; R(G), X) - d - \zeta(Y),$$

where $\zeta(Y)$ is defined to be the proportion of k -sets of clusters Z with $Y \subseteq Z \subseteq Y \cup X$ which are not (ε_k, r) -regular with respect to G .

Proof. Observe that we can transform $R(G)$ to $R_d(G)$ by editing the edge-weights of $R(G)$ in three stages. First, for any edge S of $R(G)$ with $d^*(S) \geq d$, we increase the weight of S from $d^*(S)$ to 1. Second, for any edge S of $R(G)$ with $d^*(S) < d$, we decrease the weight of S from $d^*(S)$ to zero. Finally, for any edge S of $R(G)$ such that G is not (ε_k, r) -regular with respect to S , we reduce the weight of S to zero. Note that the number of k -sets S of the latter type is at most $\varepsilon_k \binom{t}{k}$ since \mathcal{J} is a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G .

To prove the first equation, we consider the effect of each of these changes on the quantity $d_H(R(G)[X])$. Recall that this was defined to be the average of $\prod_{e \in H} d^*(\phi(e))$ taken over all the $\binom{|X|}{h} h!$ injections $\phi : V(H) \rightarrow X$, where $h := |V(H)|$. If the weight of any edge of $R(G)$ is increased, this average cannot decrease. Likewise, if the weights of any subset of edges of $R(G)$ are decreased from at most d to zero, then this average will decrease by at most d also. Finally, if the weights of some set of m chosen edges of $R(G)$ are reduced to zero, then this reduces $\prod_{e \in H} d^*(\phi(e))$ by at most one for each of the at most $e(H)m \binom{|X|-k}{h-k} (h-k)!k!$ injections ϕ for which $\phi(e)$ is a chosen edge for some $e \in H$. So $d_H(R(G)[X])$ decreases by at most

$$\frac{e(H)m \binom{|X|-k}{h-k} (h-k)!k!}{\binom{|X|}{h} h!} = \frac{e(H)m}{\binom{|X|}{k}}.$$

Combining these changes for the three steps described above, we obtain the first equation.

Similarly, for the second equation we consider the effect of each change on the quantity $\overline{\deg}(Y; R(G), X)$, which was defined to be the average of $d^*(Z)$ over all k -sets Z with $Y \subseteq Z \subseteq Y \cup X$. As before, increasing the weight of any edge of $R(G)$ cannot cause this average to decrease, and decreasing the weights of any subset of edges of $R(G)$ by at most d will cause this average to decrease by at most d . Finally, since $d^*(Z)$ must be

between zero and one for any Z , reducing the weight of a ζ -proportion of the k -sets Z to zero will reduce this average by at most ζ ; combining these changes for the three steps gives the second equation. \square

4.3. Cycle Embedding Lemma. A standard and very useful result (originally proved by Łuczak [22]) in extremal graph theory is that for all $d > 0$ and sufficiently small $\varepsilon > 0$, the following holds. Given a graph G and a partition of $V(G)$ into clusters of equal size, let $R_d(G)$ be the graph whose vertices correspond to clusters and where an edge indicates that the bipartite graph induced by the corresponding clusters is ε -regular with density at least d . If there is a connected component of $R_d(G)$ which contains a matching covering at least $(\alpha + d)v(R_d(G))$ of the vertices of $R_d(G)$, then there are paths and even cycles in G of each length up to $\alpha v(G)$. In fact, a stronger result is true, as observed by Hladký, Král' and Piguet (see [6], where this idea was also used): we need only a fractional matching with weight $(\alpha + d)v(R_d(G))$. In this subsection we state the corresponding result for tight paths and cycles in k -graphs.

Let G be an n -vertex k -graph. Then a *matching* in G is a set $M \subseteq E(G)$ of vertex-disjoint edges of G , and the *matching number* $\nu(G)$ denotes the maximum size of a matching in G . A *fractional matching* M in G assigns a weight $w_e \in [0, 1]$ to each edge $e \in G$ so that for every vertex $v \in V(G)$ we have $\sum_{e: v \in e} w_e \leq 1$. The *weight* of M is the sum of all the edge weights, which must lie between zero and n/k . We say that M is *perfect* if it has weight n/k .

Next, we define a *tight walk* W in G to be a sequence of vertices of G such that each set of k consecutive vertices induces an edge of G . For edges $e, f \in G$ we say that W is a walk from e to f if W begins with the vertices of e (in some order) and concludes with the vertices of f (in some order). If such a walk exists then we say that e and f are *tightly connected*. This gives an equivalence relation on the edges of G . To see this, observe that if e and f are edges of G with $|e \cap f| = k - 1$, then given any $(k - 1)$ -tuple in e and $(k - 1)$ -set in f , there is a tight walk in G whose first $k - 1$ vertices are the chosen $(k - 1)$ -tuple in e (in order) and whose last $k - 1$ vertices are the chosen $(k - 1)$ -set in f (in some order which we cannot choose). Applying this observation repeatedly establishes the transitivity of the ‘tightly connected’ relation, and moreover shows that the observation still holds if we replace the assumption $|e \cap f| = k - 1$ by the weaker assumption that e and f are tightly connected in G . A *tight component* of G is an equivalence class of this relation, that is, an edge maximal set $C \subseteq G$ such that each pair e, f of edges in C are tightly connected (recall we identify this edge set with the subgraph of G with vertex set $\bigcup C$ and edge set C). A *tightly connected matching* in G is a matching in which all edges are tightly connected (that is, they all lie in the same tight component of G). Finally, a *tightly connected fractional matching* is a fractional matching in which the same is true of all edges of non-zero weight.

We can now state our main embedding result, which we prove in Section 9.

Lemma 9 (Cycle Embedding Lemma). *Let k, r, n_0, t be positive integers, and $\psi, d_2, \dots, d_k, \varepsilon, \varepsilon_k$ be positive constants such that $1/d_i \in \mathbb{N}$ for each $2 \leq i \leq k - 1$, and such that $1/n_0 \ll 1/t$,*

$$\frac{1}{n_0} \ll \frac{1}{r}, \varepsilon \ll \varepsilon_k, d_2, \dots, d_{k-1} \quad \text{and} \quad \varepsilon_k \ll \psi, d_k, \frac{1}{k}.$$

Then the following holds for all integers $n \geq n_0$. Let G be a k -graph on n vertices, and \mathcal{J} be a $(\cdot, \cdot, \varepsilon, \varepsilon_k, r)$ -regular slice for G with t clusters and density vector (d_{k-1}, \dots, d_2) . Suppose that $R_{d_k}(G)$ contains a tightly connected fractional matching with weight μ . Then G contains a tight cycle of length ℓ for every $\ell \leq (1 - \psi)k\mu n/t$ that is divisible by k .

5. ERDŐS-GALLAI FOR HYPERGRAPHS

In this section we aim to demonstrate the value of the tools presented in the previous section by proving Theorem 1.

Our strategy is simple: by Lemmas 6 and 9 it is enough to show that any k -graph G with edge density α contains a tightly connected matching with at least $\alpha v(G)/k$ edges. To obtain this, we show (Proposition 10) that any G with edge density α contains a tight component G^* with $e_k(\mathcal{G}^*) \geq \alpha e_{k-1}(\mathcal{G}^*)$, and then that any k -graph which satisfies this inequality contains the desired matching (Lemma 11).

We first justify our assertion that there is a dense tight component.

Proposition 10. *Let G be a k -graph on n vertices. Then there is a tight component G^* of G such that*

$$e_k(\mathcal{G}^*) \geq \frac{e_{k-1}(\mathcal{G}^*)}{\binom{n}{k-1}} e(G),$$

where \mathcal{G}^* denotes the k -complex generated by the down-closure of G^* .

Proof. Let G_1, \dots, G_s be the tight components of G , and let \mathcal{G}_i denote the k -complex generated by the down-closure of G_i . Fix $\ell \in [s]$ which maximises $e_k(\mathcal{G}_\ell)/e_{k-1}(\mathcal{G}_\ell)$. By the definition of tight components we have $\sum_{i \in [s]} e_{k-1}(\mathcal{G}_i) \leq \binom{n}{k-1}$. Hence

$$e(G) = \sum_{i \in [s]} e_k(\mathcal{G}_i) \leq \frac{e_k(\mathcal{G}_\ell)}{e_{k-1}(\mathcal{G}_\ell)} \sum_{i \in [s]} e_{k-1}(\mathcal{G}_i) \leq \frac{e_k(\mathcal{G}_\ell)}{e_{k-1}(\mathcal{G}_\ell)} \binom{n}{k-1}.$$

□

We next state the lemma guaranteeing that a dense k -graph has a large matching.

Lemma 11. *Let k, r be any natural numbers and let \mathcal{G} be a k -complex in which*

$$(1) \quad e_k(\mathcal{G}) \geq (r-1)e_{k-1}(\mathcal{G}) + 1.$$

Then $\nu(\mathcal{G}^{(k)}) \geq r$.

The proof of this lemma proceeds inductively and uses a classical concept from extremal set theory called compression. Let \mathcal{H} be a hypergraph on vertex set $[n]$, and choose $i, j \in [n]$ with $i < j$. Then the ij -compression S_{ij} performs the following operation on \mathcal{H} . For every edge $e \in E(\mathcal{H})$ such that $i \notin e$, $j \in e$ and $\{i\} \cup e \setminus \{j\} \notin E(\mathcal{H})$, delete e from $E(\mathcal{H})$ and replace it by $\{i\} \cup e \setminus \{j\}$. We denote the resulting hypergraph by $S_{ij}(\mathcal{H})$. If $S_{ij}(\mathcal{H}) = \mathcal{H}$ for every $i < j$, then we say that \mathcal{H} is *fully-compressed*.

The next proposition sets out various properties of compressions of complexes which we shall use.

Proposition 12. *Let \mathcal{H} be a k -complex on vertex set $[n]$. Then for any $1 \leq i < j \leq n$ we have that*

- (a) $S_{ij}(\mathcal{H})$ is a k -complex with $e_\ell(S_{ij}(\mathcal{H})) = e_\ell(\mathcal{H})$ for any ℓ ,
- (b) $\nu(S_{ij}(\mathcal{H}^{(k)})) \leq \nu(\mathcal{H}^{(k)})$, and
- (c) if \mathcal{H} is fully-compressed, then for any edge $e \in \mathcal{H}$ such that $j \in e$ and $i \notin e$ we have that $\{i\} \cup e \setminus \{j\} \in \mathcal{H}$.

Proof. (a) is immediate from the definition, whilst (b) follows from Lemma 2.1 in [14], or is easy to prove directly. Finally (c) follows since in this case $S_{ij}(\mathcal{H}) = \mathcal{H}$. □

The second proposition needed to prove Lemma 11 shows that if G satisfies (1) and some edge of \mathcal{G} is contained in few higher-level edges of \mathcal{G} , then removing this edge and all edges containing it from \mathcal{G} gives a subcomplex which also satisfies (1).

Proposition 13. *Let k, r be any natural numbers and let \mathcal{G} be a k -complex in which $e_k(\mathcal{G}) \geq (r-1)e_{k-1}(\mathcal{G}) + 1$. Fix $0 \leq j \leq k-1$, and suppose that $e \in E(\mathcal{G}^{(j)})$ lies in fewer than $(k-j)r$ edges of $\mathcal{G}^{(j+1)}$. Let \mathcal{G}' be the k -complex obtained by deleting e and any edge of \mathcal{G} which contains e from \mathcal{G} . Then $e_k(\mathcal{G}') \geq (r-1)e_{k-1}(\mathcal{G}') + 1$.*

To prove this we use the local LYM inequality, which can be found in e.g. [3].

Theorem 14 (Local LYM inequality). *Let \mathcal{G} be an i -complex on n vertices. Then $e_{i-1}(\mathcal{G}) \geq \frac{i}{n-i+1}e_i(\mathcal{G})$. \square*

Proof of Proposition 13. Let $A = \{x \in V(\mathcal{G}) : e \cup \{x\} \in \mathcal{G}^{(j+1)}\}$. By assumption we have $|A| \leq (k-j)r - 1$. Let \mathcal{H} be the $(k-j)$ -complex on vertex set A with edge set $\{e' \setminus e : e \subseteq e' \in \mathcal{G}\}$. Note that when we delete e and all edges containing it from \mathcal{G} , we delete exactly $d := e_{k-j}(\mathcal{H})$ edges from $\mathcal{G}^{(k)}$ and exactly $d' := e_{k-j-1}(\mathcal{H})$ edges from $\mathcal{G}^{(k-1)}$. By Theorem 14,

$$\begin{aligned} \frac{d'}{d} &= \frac{e_{k-j-1}(\mathcal{H})}{e_{k-j}(\mathcal{H})} \geq \frac{k-j}{|A| - k + j + 1} \geq \frac{k-j}{(k-j)r - 1 - k + j + 1} \\ &= \frac{k-j}{(k-j)(r-1)} = \frac{1}{r-1}. \end{aligned}$$

Thus the complex \mathcal{G}' obtained from our deletions satisfies

$$\begin{aligned} e_k(\mathcal{G}') &= e_k(\mathcal{G}) - d \geq e_k(\mathcal{G}) - (r-1)d' \\ &\geq (r-1)e_{k-1}(\mathcal{G}) + 1 - (r-1)d' = (r-1)e_{k-1}(\mathcal{G}') + 1 \end{aligned}$$

where we used the assumption that $e_k(\mathcal{G}) \geq (r-1)e_{k-1}(\mathcal{G}) + 1$ for the second inequality. \square

The case $j = 0$ of Proposition 13 immediately gives the following corollary, which we use in the proof of Lemma 11.

Corollary 15. *Let k, r be any natural numbers and let \mathcal{G} be a k -complex in which $e_k(\mathcal{G}) \geq (r-1)e_{k-1}(\mathcal{G}) + 1$. Then $|V(\mathcal{G})| \geq kr$. \square*

We now give the proof of Lemma 11.

Proof of Lemma 11. Let \mathcal{G} be a k -complex on vertex set $[n]$ in which $e_k(\mathcal{G}) \geq (r-1)e_{k-1}(\mathcal{G}) + 1$. First, we perform repeatedly the following two operations. If for some $\ell \in [k-1]$ there is an edge $e \in \mathcal{G}^{(\ell)}$ which is contained in fewer than $(k-\ell)r$ edges of $\mathcal{G}^{(\ell+1)}$, we delete it and all its supersets from \mathcal{G} . If there are $1 \leq i < j \leq n$ such that $S_{ij}(\mathcal{G}) \neq \mathcal{G}$ then we replace \mathcal{G} with $S_{ij}(\mathcal{G})$. Eventually we reach a complex \mathcal{H} where neither operation is possible.

Observe that Proposition 12 part (a) and Proposition 13 together tell us that, since we started with a complex \mathcal{G} satisfying (1), \mathcal{H} satisfies (1) also. Moreover, Proposition 12 part (b) together with the trivial fact that deleting edges from a complex does not increase the matching number of any level of the complex implies that $\nu(\mathcal{H}^{(k)}) \leq \nu(\mathcal{G}^{(k)})$. By definition, \mathcal{H} is fully-compressed. Now given $1 \leq \ell \leq k-1$, let e be an edge of $\mathcal{H}^{(\ell)}$. Because e is in \mathcal{H} , there are at least $(k-\ell)r$ edges of $\mathcal{H}^{(\ell+1)}$ containing e , and in particular there is some $j \geq (k-\ell)r$ such that $e \cup \{j\}$ is in $\mathcal{H}^{(\ell+1)}$. Now Proposition 12 part (c) states that for any $i < j$ with $i \notin e$ we have $\{i\} \cup e \cup \{j\} \setminus \{j\} \in \mathcal{H}^{(\ell+1)}$, and thus we have

$$(2) \quad \{i\} \cup e \in \mathcal{H}^{(\ell+1)} \quad \text{for each } i \in [(k-\ell)r] \quad \text{such that } i \notin e.$$

We will now show by induction on ℓ that the ℓ -edges

$$e_{\ell,m} := \{ (k-\ell)r + m, (k-\ell+1)r + m, \dots, (k-1)r + m \}$$

for $m = 1, \dots, r$ form a matching in $\mathcal{H}^{(\ell)}$. Note that for any $\ell \in [k]$ the sets $e_{\ell, m}$ for $m \in [r]$ are by definition pairwise-disjoint; it remains to show that all $e_{\ell, m}$ are edges of \mathcal{H} .

The base case $\ell = 1$ is trivial, since we are looking only for distinct vertices, and the singletons $(k-1)r+1, (k-1)r+2, \dots, kr$ are in \mathcal{H} by Corollary 15.

Now for some $1 \leq \ell \leq k-1$ and $m \in [r]$ suppose that $e_{\ell, m}$ is an edge of $\mathcal{H}^{(\ell)}$. Let $i = (k-\ell-1)r+m \leq (k-\ell)r$. By definition we have $i \notin e_{\ell, m}$, and so by (2) we have $e_{\ell+1, m} = \{i\} \cup e_{\ell, m} \in \mathcal{H}^{(\ell+1)}$, as desired.

We conclude that the sets $e_{k, m}$ for $m \in [r]$ form a matching in $\mathcal{H}^{(k)}$, so $r \leq \nu(\mathcal{H}^{(k)}) \leq \nu(\mathcal{G}^{(k)})$ as desired. \square

We note that Lemma 11 is tight. Indeed, for any k and r , let \mathcal{K} be the k -complex generated by the down-closure of the complete k -graph $K_{kr-1}^{(k)}$ on $kr-1$ vertices. Then

$$e_k(\mathcal{K}) = \binom{kr-1}{k} = (r-1) \binom{kr-1}{k-1} = (r-1)e_{k-1}(\mathcal{K}),$$

and $\nu(\mathcal{K}^{(k)}) = r-1$.

We can now complete our proof of Theorem 1, which we restate for the convenience of the reader.

Theorem 1. *For every positive δ and every integer $k \geq 3$, there is an integer n_* such that the following holds for all $\alpha \in [0, 1]$. If G is a k -uniform hypergraph on $n \geq n_*$ vertices with $e(G) \geq (\alpha + \delta) \binom{n}{k}$, then G contains a tight cycle of length ℓ for every $\ell \leq \alpha n$ that is divisible by k .*

Proof. Given $k \geq 3$ and $\delta > 0$, we choose $d_k = \delta/3$, $t_0 = 24k/\delta$ and $\psi = \delta/8$. We let $\varepsilon_k \leq \delta/12$ be sufficiently small for Lemma 9. We choose functions $\varepsilon(\cdot)$ tending to zero sufficiently rapidly, and $r(\cdot)$ growing sufficiently rapidly, so that for any $t \in \mathbb{N}$ and $d_2, \dots, d_{k-1} \geq 1/t$ we may apply Lemma 9 with $r(t)$, $\varepsilon(t)$. Obtain t_1 and n_0 by applying Lemma 6 with inputs $t_0, \varepsilon_k, r(\cdot)$ and $\varepsilon(\cdot)$ (taking $q = s = 1$); for the rest of this proof we write r and ε for $r(t_1)$ and $\varepsilon(t_1)$ respectively. Then by Lemma 6, for any k -graph G on n vertices, where $n \geq n_0$ is divisible by $t_1!$, there is a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G such that $d(R(G)) \geq d(G) - \varepsilon_k$. Finally, we choose $n_1 \geq n_0$ sufficiently large for us to apply Lemma 9 with t_1 in place of t , n_1 in place of n_0 and all other constants as above, and also such that $t_1! \binom{n_1}{k-1} < \delta \binom{n_1}{k} / 12$.

Set $n_* := n_1 + t_1!$, and let G' be a k -graph on $n' \geq n_*$ vertices with $e(G') \geq (\alpha + \delta) \binom{n'}{k}$. Delete at most $t_1! - 1$ vertices from G' to obtain a k -graph G on n vertices, where n is divisible by $t_1!$. By choice of n_1 we have $e(G) \geq (\alpha + 11\delta/12) \binom{n}{k}$. So Lemma 6 gives us a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice \mathcal{J} for G with t clusters (where $t_0 \leq t \leq t_1$) such that $d(R(G)) \geq d(G) - \varepsilon_k \geq \alpha + 10\delta/12$. Finally by Lemma 8 we have $d(R_{d_k}(G)) \geq d(R(G)) - d_k - \varepsilon_k \geq \alpha + 5\delta/12$.

The total number of edges in $R_{d_k}(G)$ is therefore at least $(\alpha + 5\delta/12) \binom{t}{k}$. It follows by Proposition 10 that there is a tight component R of $R_{d_k}(G)$ such that the k -complex \mathcal{R} generated by the down-closure of R satisfies

$$\begin{aligned} e_k(\mathcal{R}) &\geq \frac{e_{k-1}(\mathcal{R})}{\binom{t}{k-1}} \left(\alpha + \frac{5\delta}{12} \right) \binom{t}{k} = \frac{t-k+1}{k} \left(\alpha + \frac{5\delta}{12} \right) e_{k-1}(\mathcal{R}) \\ &\geq \left(\alpha + \frac{4\delta}{12} \right) t \frac{e_{k-1}(\mathcal{R})}{k} \end{aligned}$$

where the final inequality is guaranteed by $t \geq t_0 = 24k/\delta$.

By Lemma 11 it follows that R contains a matching of at least $(\alpha + \delta/4)t/k$ edges. Since this matching is contained in R it is tightly connected. Lemma 9 then implies that G contains a tight cycle of length ℓ for every ℓ at most $(1 - \psi)k((\alpha + \delta/4)t/k)n/t \geq \alpha n$ which is divisible by k (where the inequality follows from $\psi = \delta/8$). \square

6. CYCLES IN k -PARTITE HYPERGRAPHS

In this section we will prove Theorem 2. As in the previous section, our strategy is to show that the partite codegree condition implies the existence of a sufficiently large (this time fractional) matching in a connected component of $R_d(G)$, after which applying Lemma 9 will give the existence of the desired tight paths and cycles in G . Unfortunately, while the minimum codegree of G transfers to $R(G)$ by Lemma 6, as we saw in Lemma 8 it does not transfer perfectly to $R_d(G)$: some $(k - 1)$ -sets of clusters may lie in many irregular k -sets of clusters, and so have small relative degree in $R_d(G)$. To handle this, we follow the approach of Keevash and Mycroft [18], moving to non-uniform hypergraphs \mathcal{R} in which an edge of $\mathcal{R}^{(k-1)}$ corresponds to a $(k - 1)$ -set of $R_d(G)$ which *does* acquire the desired codegree from G , and the presence of an edge of \mathcal{R} of size $j < k - 1$ implies that most of the $(j + 1)$ -supersets of this edge are also edges. The hypergraph-theoretic core of our proof of Theorem 2 is Lemma 17; to prove this we use the following special case of Lemma 7.2 of [18], which was proved by a straightforward application of Farkas's Lemma.

Lemma 16. *Let \mathcal{H} be a k -partite hypergraph whose parts X_1, \dots, X_k each have size t . Suppose that $\emptyset \in \mathcal{H}$ and that for any $0 \leq i \leq k - 1$ and $j \in [k]$, any edge of $\mathcal{H}^{(i)}$ which does not intersect X_j is contained in at least $t - it/k$ edges of $\mathcal{H}^{(i+1)}$ which do intersect X_j . Then $\mathcal{H}^{(k)}$ admits a perfect fractional matching.* \square

Lemma 17. *Given $0 \leq \alpha < 1$ and $\beta > 0$, let \mathcal{R} be a k -partite hypergraph with vertex classes X_1, \dots, X_k of size t with the following properties.*

- (i) $\emptyset \in \mathcal{R}$, and $\{v\} \in \mathcal{R}$ for any $v \in V(\mathcal{R}) = \bigcup_{i \in [k]} X_i$.
- (ii) For any $1 \leq i \leq k - 2$ and $j \in [k]$, any edge of $\mathcal{R}^{(i)}$ which has no vertex in X_j is contained in at least $(1 - \beta)t$ edges of $\mathcal{R}^{(i+1)}$ which intersect X_j .
- (iii) Each edge of $\mathcal{R}^{(k-1)}$ is contained in at least $(\alpha + (2^k + 1)\beta)t$ edges of $\mathcal{R}^{(k)}$.

Then

- (a) $\mathcal{R}^{(k)}$ contains a tightly connected matching of αt edges, and
- (b) if $\alpha \geq \frac{1}{2}$ then $\mathcal{R}^{(k)}$ contains a tightly connected perfect fractional matching.

Proof. First note that for any $1 \leq i \leq k$, if some edge $e \in \mathcal{R}^{(i)}$ does not include any edge of $\mathcal{R}^{(i-1)}$ as a subset, then deleting e from \mathcal{R} yields a subgraph which also meets the conditions of the lemma. So we may assume that every edge in $\mathcal{R}^{(i)}$ includes at least one edge of $\mathcal{R}^{(i-1)}$ as a subset.

The edges $f \in \mathcal{R}^{(k)}$ for which every subset of f is an edge of \mathcal{R} are the most useful in the sense that for each subset of f we can apply the degree conditions (ii) and (iii). We say that such edges are *excellent*. We will show that each tightly connected component of $\mathcal{R}^{(k)}$ contains a matching of αt excellent edges, giving (a).

Claim 18. *Each edge of $\mathcal{R}^{(k)}$ is tightly connected to an excellent edge of $\mathcal{R}^{(k)}$.*

Proof. Let $e = \{v_1, \dots, v_k\}$ be an edge of $\mathcal{R}^{(k)}$ with $v_i \in X_i$ for each i . Since every edge in $\mathcal{R}^{(i)}$ includes at least one edge of $\mathcal{R}^{(i-1)}$ as a subset, we may assume without loss of generality that the subsets $\{v_1, \dots, v_j\}$ are in \mathcal{R} for each $1 \leq j \leq k - 1$. Choose $w_k \in X_k$

such that $\{v_1, \dots, v_j, w_k\}$ is in \mathcal{R} for each $0 \leq j \leq k-1$. This is possible because $\{v_1, \dots, v_j\}$ is in \mathcal{R} for each j , and so at most $(k-2)\beta t + (1-\alpha - (2^k+1)\beta)t < t$ vertices in X_k do not satisfy the given condition. Now for each $i = k-1, \dots, 1$ in that order, we choose a vertex $w_i \in X_i$ such that $\{v_1, \dots, v_j, S\}$ is an edge of \mathcal{R} for each $0 \leq j \leq i-1$ and each subset S of $\{w_i, \dots, w_k\}$. Again, since $2^{k-1}\beta t + (1-\alpha - (2^k+1)\beta)t < t$ this is always possible. Then the sequence of edges of the form $\{v_1, \dots, v_i, w_{i+1}, \dots, w_k\}$ for $0 \leq i \leq k$ is a tight walk between e and $\{w_1, \dots, w_k\}$ in \mathcal{R} (since by construction each of these k -sets is in \mathcal{R}) and also by construction $\{w_1, \dots, w_k\}$ is an excellent edge. \square

We next show that for any small vertex set S and any excellent edge e , there is a tight walk from e to another excellent edge outside S .

Claim 19. *Given any set S of vertices of \mathcal{R} such that $|S \cap X_i| < \alpha t$ for each $1 \leq i \leq k$, and an excellent edge e , there is a tight walk from e to an excellent edge e' of $\mathcal{R}^{(k)}$ with $e' \cap S = \emptyset$.*

Proof. Let $e = \{u_1, \dots, u_k\}$, where $u_i \in X_i$ for each i . For each $i = 1, \dots, k$ in that order, we wish to choose v_i to be a vertex of $X_i \setminus S$ such that $\{v_1, \dots, v_i, u_{i+1}, \dots, u_k\}$ is an excellent edge. This means that for each subset T of $\{v_1, \dots, v_{i-1}, u_{i+1}, \dots, u_k\}$, we need to guarantee that $T \cup \{v_i\}$ is an edge of \mathcal{R} . In total, at most

$$2^{k-1}\beta t + (1-\alpha - (2^k+1)\beta)t < t - \alpha t - \beta t$$

vertices of X_i are not suitable. Thus at least $(\alpha + \beta)t > |S \cap X_i|$ vertices are suitable, and it is always possible to choose $v_i \in X_i \setminus S$ as desired. Now the sequence of edges of the form $\{v_1, \dots, v_i, u_{i+1}, \dots, u_k\}$ for $0 \leq i \leq k$ is a tight walk from e to an excellent edge $e' = \{v_1, \dots, v_k\}$ which does not intersect S as desired. \square

Next we show that any tight component of $\mathcal{R}^{(k)}$ has a matching with at least αt edges. So let C be a tight component of $\mathcal{R}^{(k)}$; then C contains at least one edge. Since by Claim 18 any edge of $\mathcal{R}^{(k)}$ is tightly connected to an excellent edge of $\mathcal{R}^{(k)}$, C contains an excellent edge e_1 .

We construct a matching in C as follows. We let $M_1 = \{e_1\}$. Now for each $2 \leq i \leq \alpha t$, let e_i be the edge returned by applying Claim 19 with $e = e_{i-1}$ and S consisting of all vertices covered by M_{i-1} , and let $M_i = M_{i-1} \cup \{e_i\}$. Then $M_{\alpha t}$ is the desired matching in C . (In fact, this process actually returns a tight path in C on αt vertices. The matching $M_{\alpha t}$ consists of every k th edge of this path.)

To complete the proof we show that in the case $\alpha \geq \frac{1}{2}$ we have more: $\mathcal{R}^{(k)}$ is tightly connected and admits a perfect fractional matching. For the latter we use Lemma 16. Indeed, any edge of $\mathcal{R}^{(k-1)}$ is contained in at least $\alpha t \geq t/k = t - \frac{k-1}{k}t$ edges of $\mathcal{R}^{(k)}$, the edge of $\mathcal{R}^{(0)}$ (i.e. \emptyset) is contained in the edges $\{v\} \in \mathcal{R}^{(1)}$ for $v \in V(\mathcal{R})$, and for any $1 \leq i \leq k-2$, any edge of $\mathcal{R}^{(i)}$ which does not intersect some part X_j is contained in at least $(1-\beta)t \geq (k-1)t/k \geq t - \frac{i}{k}t$ edges of $\mathcal{R}^{(i+1)}$ which do intersect X_j (the first inequality follows from $\alpha + (2^{k-1}+1)\beta \leq 1$). So \mathcal{R} meets the conditions of Lemma 16, from which we deduce that $\mathcal{R}^{(k)}$ admits a perfect fractional matching.

To show that $\mathcal{R}^{(k)}$ is tightly connected, it is by Claim 18 enough to show that any two excellent edges are tightly connected. Given two excellent edges $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$, where $u_i, v_i \in X_i$ for each i , we have that both $e := \{u_1, \dots, u_{k-1}\}$ and $f := \{v_1, \dots, v_{k-1}\}$ are edges of \mathcal{R} : in particular there are at least $2(\alpha + (2^k+1)\beta)t - t \geq (2^k+1)\beta t$ vertices in X_k which form edges with both e and f . Of these, all but at most $2 \cdot 2^{k-1}\beta t$ form excellent edges with both e and f , so we can choose a vertex $w_k \in X_k$ so that $e \cup \{w_k\}$ and $f \cup \{w_k\}$ are both excellent edges of $\mathcal{R}^{(k)}$. We repeat the same

process for each $i = k - 1, \dots, 1$ in that order to find $w_i \in X_i$ which forms excellent edges with each of $\{u_1, \dots, u_{i-1}, w_{i+1}, \dots, w_k\}$ and $\{v_1, \dots, v_{i-1}, w_{i+1}, \dots, w_k\}$. We have thus constructed tight walks from each of $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$ to $\{w_1, \dots, w_k\}$, and thus the two edges are tightly connected, as desired. \square

We now prove Theorem 2, which we restate for the reader's convenience. We note that an ordering of constants and functions:

$$1/n_1 \ll \varepsilon(\cdot), 1/r(\cdot) \ll 1/t_1 \ll \varepsilon_k \ll \beta, d_k, \psi \ll \delta, 1/k$$

would be enough in the proof, but for convenience we will give values explicitly.

Theorem 2. *For every positive δ and every integer $k \geq 3$ there is an integer n_* such that the following holds for each $\alpha \in [0, 1]$. If G is a k -uniform k -partite hypergraph with parts of size $n \geq n_*$, such that any collection of $k - 1$ vertices, one in each of $k - 1$ parts of G , lies in at least $(\alpha + \delta)n$ edges of G , then*

- (a) G contains a tight cycle of length ℓ for every $\ell \leq \alpha kn$ that is divisible by k , and
- (b) if $\alpha \geq \frac{1}{2}$ then G contains a tight cycle of length ℓ for every $\ell \leq (1 - \delta)kn$ that is divisible by k .

Proof. Given $0 < \delta < 1$ and $k \geq 3$, we choose

$$\psi = \frac{\delta}{28}, \quad d_k = \frac{\delta}{7}, \quad \beta = \frac{\delta}{7(2^k + 1)}, \quad t_0 = 1.$$

We let

$$\varepsilon_k \leq \frac{\beta^{k-1}\delta}{7 \cdot 2^{k-1}}$$

be sufficiently small for applying Lemma 9 with constants ψ, ε_k and d_k . Next we choose functions $r : \mathbb{N} \rightarrow \mathbb{N}$ growing sufficiently rapidly and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$ tending to zero sufficiently rapidly so that for any $t \in \mathbb{N}$ we can apply Lemma 9 with ε_k as chosen above, $r(t)$, $\varepsilon(t)$ and any $d_2, \dots, d_k \geq 1/t$. Now Lemma 6 returns t_1 and n_0 such that for any $n \geq n_0$ divisible by $t_1!$, any kn -vertex k -graph G and any partition \mathcal{Q} of $V(G)$ into k parts of size n , there is a $(t_0, t_1, \varepsilon(t_1), \varepsilon_k, r(t_1))$ -regular slice for G which satisfies conditions (a) and (b) of Lemma 6 and whose ground partition refines \mathcal{Q} . Fix this t_1 ; for the rest of the proof we will write ε and r for $\varepsilon(t_1)$ and $r(t_1)$ respectively. Finally we choose $n_1 \geq \max(n_0, 7t_1!/\delta)$ to be sufficiently large to apply Lemma 9 with t_1 in place of t , n_1 in place of n_0 and all other constants as above.

Let G' be a k -partite k -graph with parts U'_1, \dots, U'_k of size $n' \geq n_* := n_1 + t_1!$ such that any $(k - 1)$ -set of vertices, one in each of $k - 1$ parts of G' , lies in at least $(\alpha + \delta)n'$ edges of G' . Choose n divisible by $t_1!$ with $n' - t_1! \leq n \leq n'$ and a subset $U_i \subseteq U'_i$ of size n for each $i \in [k]$, and let G be the subgraph of G' induced by $\bigcup_{i \in [k]} U_i$. Let \mathcal{Q} denote the partition of $V(G)$ into the classes U_i . Since at most $t_1!$ vertices were deleted from each vertex class of G' , and $n' \geq n_1 \geq 7t_1!/\delta$, any \mathcal{Q} -partite $(k - 1)$ -set S of vertices of G lies in at least $(\alpha + 6\delta/7)n'$ edges of G , so $\overline{\deg}(S; G, U_i) \geq \alpha + 6\delta/7$, where U_i is the part of \mathcal{Q} which S does not intersect. By Lemma 6 we may choose a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice \mathcal{J} for G which has the properties (a) and (b) of Lemma 6 and whose ground partition refines \mathcal{Q} . Since each cluster of \mathcal{J} has the same size, and the same is true of the parts of \mathcal{Q} , the number of clusters of \mathcal{J} must be divisible by k . So let t be such that kt is the number of clusters. Then each cluster has size $m := n/t$, and since \mathcal{J} is (t_0, t_1, ε) -equitable we have $t_0 \leq kt \leq t_1$.

Let \mathcal{Q}_R denote the natural partition of the clusters of \mathcal{J} ; so the parts of \mathcal{Q}_R are V_1, \dots, V_k , where V_i consists of all clusters which are subsets of U_i . Observe that we

have $|V_i| = t$, and that the reduced k -graph $R_{d_k}(G)$ is \mathcal{Q}_R -partite. We would like to immediately apply Lemma 17 to the k -complex generated by the down-closure of $R_{d_k}(G)$, but unfortunately this complex may not satisfy the required degree conditions. Instead we obtain a subcomplex which does satisfy the necessary conditions.

To do this, consider a \mathcal{Q}_R -partite set \mathcal{X} of $k - 1$ clusters. There is exactly one vertex class V_i which \mathcal{X} does not intersect, and we say that \mathcal{X} is *good* if there are fewer than $\delta t/7$ clusters $X \in V_i$ such that G is not (ε_k, r) -regular with respect to the k -set $\mathcal{X} \cup \{X\}$ (for convenience, we henceforth refer to such k -sets simply as *irregular k -sets*). Next, for $\ell = k - 2, \dots, 0$ (in that order) we say that a \mathcal{Q}_R -partite set \mathcal{X} of ℓ clusters is *good* if it is contained in at most $\beta t/2$ \mathcal{Q}_R -partite sets of $\ell + 1$ clusters which are *bad*, i.e. not good.

At the end of this process we have labelled each individual cluster as good or bad. Suppose that at least $\beta t/2$ clusters are bad. Then we can construct an irregular k -set in $R(G)$ by choosing any one of these bad clusters, any of the at least $\beta t/2$ clusters with which it forms a bad pair, and so on up to any of the at least $\delta t/7$ irregular k -sets which contain the bad $(k - 1)$ -set so constructed. Since we could construct a given irregular k -set in at most $k!$ ways, we conclude that $R(G)$ contains at least $(\beta t/2)^{k-1}(\delta t/7)/k!$ irregular k -sets. Since $\varepsilon_k \leq \beta^{k-1}\delta/(7 \cdot 2^{k-1})$, this is greater than $\varepsilon_k \binom{t}{k}$, contradicting the fact that \mathcal{J} is an $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for G .

We conclude that fewer than $\beta t/2$ clusters are bad. So for each $1 \leq i \leq k$ we may choose a set V'_i of $t' := (1 - \beta/2)t$ clusters from V_i , all of which are good. We now define a k -partite hypergraph \mathcal{R} on $V' := V'_1 \cup \dots \cup V'_k$ as follows. For each $0 \leq j \leq k - 1$ we take $\mathcal{R}^{(j)}$ to consist of all good j -sets of clusters (so in particular, $\emptyset \in \mathcal{R}$ and $\{v\} \in \mathcal{R}$ for any $v \in V'$). For $\mathcal{R}^{(k)}$ we instead take all edges of $R_{d_k}(G)$ which are contained in V' .

We want to apply Lemma 17 with t' and $\alpha + \delta/7$ in place of t and α respectively. Observe that condition (i) is satisfied. For condition (ii) note that by definition of \mathcal{R} , for any $1 \leq i \leq k - 2$ and $j \in [k]$, any edge $\mathcal{X} \in \mathcal{R}^{(i)}$ which does not intersect V'_j is contained in at least $|V'_j| - \beta t/2 \geq (1 - \beta)t'$ edges of $\mathcal{R}^{(i+1)}$ which do intersect V'_j .

We now check condition (iii). Consider any edge \mathcal{X} of $\mathcal{R}^{(k-1)}$, so \mathcal{X} is a good set of $k - 1$ clusters, and let j be such that \mathcal{X} does not intersect V_j . Recall that $\overline{\deg}(S; G, U_j) \geq \alpha + 6\delta/7$ for any \mathcal{Q} -partite $(k - 1)$ -set of vertices of G which does not intersect U_j ; it follows that $\overline{\deg}(\mathcal{J}_{\mathcal{X}}; G, U_j) \geq \alpha + 6\delta/7$. By property (b) of Lemma 6 we have $|\overline{\deg}(\mathcal{X}; R(G), V_j) - \overline{\deg}(\mathcal{J}_{\mathcal{X}}; G, U_j)| < \varepsilon_k$, so $\overline{\deg}(\mathcal{X}; R(G), V_j) \geq \alpha + 5\delta/7$. Since \mathcal{X} is good, the proportion of clusters X in V_j for which $\mathcal{X} \cup \{X\}$ is an irregular k -set is at most $\delta/7$, and so Lemma 8 implies that $\overline{\deg}(\mathcal{X}; R_{d_k}(G), V_j) \geq \overline{\deg}(\mathcal{X}; R(G), V_j) - d_k - \delta/7 \geq \alpha + 3\delta/7$, and therefore $\overline{\deg}(\mathcal{X}; R_{d_k}(G), V'_j) \geq \alpha + 2\delta/7$. As $(2^k + 1)\beta = \delta/7$, it follows that any edge \mathcal{X} of $\mathcal{R}^{(k-1)}$ lies in at least $(\alpha + \delta/7 + (2^k + 1)\beta)t'$ edges of $\mathcal{R}^{(k)}$.

So \mathcal{R} also satisfies condition (iii) of Lemma 17. It follows that $\mathcal{R}^{(k)}$ has a tight component with a matching containing $(\alpha + \delta/7)t' = (\alpha + \delta/7)(1 - \beta/2)t \geq (\alpha + \delta/14)t$ edges. Since $\mathcal{R}^{(k)}$ is a subgraph of $R_{d_k}(G)$, by Lemma 9 G contains a tight cycle of length ℓ for any $k < \ell \leq (1 - \psi)(\alpha + \delta/14)kn$ which is divisible by k . Since $(1 - \psi)(\alpha + \delta/14)kn \geq \alpha kn$ this completes the proof of (a). If $\alpha \geq \frac{1}{2}$ then Lemma 17 also shows that $\mathcal{R}^{(k)}$ admits a tightly connected perfect fractional matching, and (b) follows analogously. \square

7. STRONG HYPERGRAPH REGULARITY

In this section we introduce the Strong Hypergraph Regularity Lemma and some related tools.

7.1. Families of partitions and the Strong Hypergraph Regularity Lemma. In order to state the Strong Hypergraph Regularity Lemma we need some more notation. Fix $k \geq 3$, and let \mathcal{P} partition a vertex set V into parts V_1, \dots, V_t . For any $A \subseteq [t]$, we denote by Cross_A the collection of \mathcal{P} -partite subsets $S \subseteq V$ of index $i(S) = A$. Likewise, we denote by Cross_j the union of Cross_A for each $A \in \binom{[t]}{j}$, so Cross_j contains all \mathcal{P} -partite subsets $S \subseteq V$ of size j . Note that Cross_A and Cross_j are dependent on the choice of partition \mathcal{P} , but this will always be clear from the context. For each $2 \leq j \leq k-1$ and $A \in \binom{[t]}{j}$ let \mathcal{P}_A be a partition of Cross_A . For consistency of notation we also define the trivial partitions $\mathcal{P}_{\{s\}} := \{V_s\}$ for $s \in [t]$ and $\mathcal{P}_\emptyset := \{\emptyset\}$. Let \mathcal{P}^* consist of the partitions \mathcal{P}_A for each $A \in \binom{[t]}{j}$ and each $0 \leq j \leq k-1$. We say that \mathcal{P}^* is a $(k-1)$ -family of partitions on V if whenever $S, T \in \text{Cross}_A$ lie in the same part of \mathcal{P}_A and $B \subseteq A$, then $S \cap \bigcup_{j \in B} V_j$ and $T \cap \bigcup_{j \in B} V_j$ lie in the same part of \mathcal{P}_B . In other words, given $A \in \binom{[t]}{j}$, if we specify one part of each \mathcal{P}_B with $B \in \binom{A}{j-1}$, then we obtain a subset of Cross_A consisting of all $S \in \text{Cross}_A$ whose $(j-1)$ -subsets are in the specified parts. Thus the partitions \mathcal{P}_B give a natural partition of Cross_A , and we are saying that \mathcal{P}_A must refine it.

We refer to the parts of each member of \mathcal{P}^* as *cells*. Also, we refer to \mathcal{P} as the *ground partition* of \mathcal{P}^* , and the parts of \mathcal{P} (i.e. the vertex classes V_i) as the *clusters* of \mathcal{P}^* . For each $0 \leq j \leq k-1$ let $\mathcal{P}^{(j)}$ denote the partition of Cross_j formed by the parts (which we call *j-cells*) of each of the partitions \mathcal{P}_A with $A \in \binom{[t]}{j}$ (so in particular $\mathcal{P}^{(1)} = \mathcal{P}$).

The cells of \mathcal{P}^* naturally form $(k-1)$ -complexes on V . Indeed, for any $0 \leq j \leq k-1$, any $A \in \binom{[t]}{j}$ and any $Q' \in \text{Cross}_A$, let $C_{Q'}$ denote the cell of \mathcal{P}_A which contains Q' . Then the fact that \mathcal{P}^* is a family of partitions implies that for any $Q \in \text{Cross}_k$ the union $\mathcal{J}(Q) := \bigcup_{Q' \subseteq Q} C_{Q'}$ of cells containing subsets of Q is a k -partite $(k-1)$ -complex. We say that the $(k-1)$ -family of partitions \mathcal{P}^* is (t_0, t_1, ε) -equitable if

- (a) \mathcal{P} partitions V into t clusters of equal size, where $t_0 \leq t \leq t_1$,
- (b) for each $2 \leq j \leq k-1$, $\mathcal{P}^{(j)}$ partitions Cross_j into at most t_1 cells,
- (c) there exists $\mathbf{d} = (d_{k-1}, \dots, d_2)$ such that for each $2 \leq j \leq k-1$ we have $d_j \geq 1/t_1$ and $1/d_j \in \mathbb{N}$, and for every $Q \in \text{Cross}_k$ the k -partite $(k-1)$ -complex $\mathcal{J}(Q)$ is $(\mathbf{d}, \varepsilon, \varepsilon, 1)$ -regular.

Note that conditions (a) and (c) imply that $\mathcal{J}(Q)$ is a $(1, t_1, \varepsilon)$ -equitable $(k-1)$ -complex (with the same density vector \mathbf{d}) for any $Q \in \text{Cross}_k$.

Next, for any \mathcal{P} -partite set Q with $2 \leq |Q| \leq k$, define $\hat{P}(Q; \mathcal{P}^*)$ to be the $|Q|$ -partite $(|Q|-1)$ -graph on $V_{i(Q)}$ with edge set $\bigcup_{Q' \in \binom{Q}{|Q|-1}} C_{Q'}$. We refer to $\hat{P}(Q; \mathcal{P}^*)$ as a *polyad*; when the family of partitions \mathcal{P}^* is clear from the context, we write simply $\hat{P}(Q)$ rather than $\hat{P}(Q; \mathcal{P}^*)$. Note that the condition for \mathcal{P}^* to be a $(k-1)$ -family of partitions can then be rephrased as saying that if $2 \leq |Q| \leq k-1$ then the cell C_Q is supported on $\hat{P}(Q)$. Moreover, we will show in the proof of Lemma 6 (Claim 28) that if \mathcal{P}^* is (t_0, t_1, ε) -equitable for sufficiently small ε , then for any $2 \leq j \leq k-1$ and any $Q \in \text{Cross}_j$ the number of j -cells of \mathcal{P}^* supported on $\hat{P}(Q)$ is precisely equal to $1/d_j$.

Now let G be a k -graph on V , and let \mathcal{P}^* be a $(k-1)$ -family of partitions on V . Let $Q \in \text{Cross}_k$, so the polyad $\hat{P}(Q)$ is a k -partite $(k-1)$ -graph. Recall (see Section 4.1) that G is (ε_k, r) -regular with respect to $\hat{P}(Q)$ if there is some d such that G is (d, ε_k, r) -regular with respect to $\hat{P}(Q)$. We say that G is (ε_k, r) -regular with respect to \mathcal{P}^* if there are at most $\varepsilon_k \binom{|V|}{k}$ sets $Q \in \text{Cross}_k$ for which G is not (ε_k, r) -regular with respect to the

polyad $\hat{P}(Q)$. That is, at most an ε_k -proportion of subsets of V of size k yield polyads with respect to which G is not regular (though some subsets of V of size k do not yield any polyad due to not being members of Cross_k). Similarly, we say that G is *perfectly* (ε_k, r) -regular with respect to \mathcal{P}^* if for every set $Q \in \text{Cross}_k$, the graph G is (ε_k, r) -regular with respect to the polyad $\hat{P}(Q)$, i.e. there are no polyads with respect to which G is not regular.

Finally, we define the notion of a slice through a family of partitions. Indeed, if \mathcal{P}^* is a $(k-1)$ -family of partitions, then a *slice through* \mathcal{P}^* is a $(k-1)$ -complex \mathcal{J} on V such that for each $0 \leq i \leq k-1$ we may write $\mathcal{J}^{(i)} = \bigcup_{A \in \binom{[t]}{i}} C(A)$, where each $C(A)$ is a cell in \mathcal{P}_A . That is, a slice consists of a single cell from each \mathcal{P}_A , but the requirement that \mathcal{J} should be a $(k-1)$ -complex requires that the choices of cells are ‘consistent’, meaning that each chosen 3-cell is supported on the chosen 2-cells, and so forth. Observe in particular that any slice through \mathcal{P}^* must be \mathcal{P} -partite. Moreover, if \mathcal{P}^* is (t_0, t_1, ε) -equitable then any slice through \mathcal{P}^* must be (t_0, t_1, ε) -equitable also, and the polyads of any slice through \mathcal{P}^* are also polyads of \mathcal{P}^* . The reason for the term ‘regular slice’ should now be apparent: we take a slice through a $(k-1)$ -family of partitions \mathcal{P}^* so that the chosen slice has desirable regularity properties.

We are now ready to state the Strong Hypergraph Regularity Lemma. The form we consider, due to Rödl and Schacht, is Lemma 23 in [32]².

Lemma 20 (Strong Hypergraph Regularity Lemma). *Let $k \geq 3$ be a fixed integer. For all positive integers q, t_0 and s , positive ε_k and functions $r : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$, there exist integers t_1 and n_0 such that the following holds for all $n \geq n_0$ which are divisible by $t_1!$. Let V be a vertex set of size n , and suppose that G_1, \dots, G_s are edge-disjoint k -graphs on V , and that \mathcal{Q} is a partition of V into at most q parts of equal size. Then there exists a $(k-1)$ -family of partitions \mathcal{P}^* on V such that*

- (a) *the ground partition of \mathcal{P}^* refines \mathcal{Q} ,*
- (b) *\mathcal{P}^* is $(t_0, t_1, \varepsilon(t_1))$ -equitable, and*
- (c) *for each $1 \leq i \leq s$, G_i is $(\varepsilon_k, r(t_1))$ -regular with respect to \mathcal{P}^* .* □

Similar results were proved previously by Rödl and Skokan [33] and Gowers [11]. In applications of Lemma 20, the regularity parameter ε_k of the graphs G_i is typically much larger than the entries of the density vector of \mathcal{P}^* , which may cause substantial technical difficulties. The next result, also due to Rödl and Schacht (the form we state is Theorem 25 in [32]), is a Regular Approximation Lemma, in which the regularity parameter may be taken to be much smaller than the densities of \mathcal{P}^* . However, the price for this is that it is not our k -graph G which is regularised, but instead a k -graph G' which differs from G only in a small number of edges.

Lemma 21 (Regular Approximation Lemma). *Let $k \geq 3$ be a fixed integer. For all positive integers q, t_0 and s , positive ν and functions $\varepsilon : \mathbb{N} \rightarrow (0, 1]$, there exist integers t_1 and n_0 such that the following holds for all $n \geq n_0$ which are divisible by $t_1!$. Let V be a vertex set of size n , and suppose that G_1, \dots, G_s are edge-disjoint k -graphs on V , and that \mathcal{Q} is a partition of V into at most q parts of equal size. Then there exist k -graphs G'_1, \dots, G'_s on V and a $(k-1)$ -family of partitions \mathcal{P}^* on V such that*

- (a) *the ground partition of \mathcal{P}^* refines \mathcal{Q} ,*
- (b) *\mathcal{P}^* is $(t_0, t_1, \varepsilon(t_1))$ -equitable,*

²In fact, their lemma allows for an initial family of partitions rather than just a partition \mathcal{Q} of the vertex set V .

- (c) for each $1 \leq i \leq s$ the graph G'_i is perfectly $(\varepsilon(t_1), 1)$ -regular with respect to \mathcal{P}^* , and
(d) for each $1 \leq i \leq s$, we have $|G'_i \Delta G_i| \leq \nu n^k$. \square

In fact, in order to prove Lemma 6 we need a slightly strengthened version of the Strong Hypergraph Regularity Lemma, which we deduce from Lemma 21.

Given families of partitions \mathcal{P}^* and $\tilde{\mathcal{P}}^*$, such that the ground partition \mathcal{P} of \mathcal{P}^* is a refinement of the ground partition $\tilde{\mathcal{P}}$ of $\tilde{\mathcal{P}}^*$, we say that \mathcal{P}^* is generated from $\tilde{\mathcal{P}}^*$ by \mathcal{P} if every $\tilde{\mathcal{P}}$ -partite j -cell of \mathcal{P}^* is an induced subgraph of a j -cell of $\tilde{\mathcal{P}}^*$. Note that this condition does not determine a unique \mathcal{P}^* , since the j -cells of \mathcal{P}^* which are not $\tilde{\mathcal{P}}$ -partite are not determined. However, if $\tilde{\mathcal{P}}^*$ has many clusters, it does determine most j -cells of \mathcal{P}^* . If $\tilde{\mathcal{J}}$ is a slice through $\tilde{\mathcal{P}}^*$ and \tilde{X} is a set of clusters of $\tilde{\mathcal{P}}$, then $\tilde{\mathcal{J}}[\tilde{X}]$ is the subset of elements of $\tilde{\mathcal{J}}$ supported on \tilde{X} . Next, given a set \tilde{X} of ℓ clusters of $\tilde{\mathcal{P}}^*$, we say that a set X of clusters of \mathcal{P}^* is \tilde{X} -consistent if X has precisely ℓ members, each of which is a subset of precisely one member of \tilde{X} . Finally, given such X and \tilde{X} , for any slice \mathcal{J} through \mathcal{P}^* we say that $\mathcal{J}[X]$ is contained in $\tilde{\mathcal{J}}[\tilde{X}]$, and write $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}[\tilde{X}]$, if every j -cell of $\mathcal{J}[X]$ is an induced subgraph of a cell of $\tilde{\mathcal{J}}[\tilde{X}]$.

The Strengthened Regularity Lemma we require follows. The difference between this lemma and Lemma 20 is the appearance of a ‘coarse’ family of partitions $\tilde{\mathcal{P}}^*$ containing the ‘fine’ family of partitions \mathcal{P}^* on which we guarantee regularity properties. This lemma guarantees that the neighbourhood of some root vertices looks about the same on a part of the coarse family of partitions as on any of the corresponding parts in the fine family of partitions.

Lemma 22 (Strengthened Regularity Lemma). *Let $k \geq 3$ be a fixed integer. For all positive integers q, t_0 and s , positive ε_k , functions $r : \mathbb{N} \rightarrow \mathbb{N}$, $\varepsilon : \mathbb{N} \rightarrow (0, 1]$ and monotone increasing functions $p : \mathbb{N} \rightarrow \mathbb{N}$, there exist integers t_1^*, t_1 and n_0 with $t_1 = p(t_1^*)t_1^*$ such that the following holds for all $n \geq n_0$ which are divisible by $t_1!$. Let V be a vertex set of size n , and suppose that G_1, \dots, G_s are edge-disjoint k -graphs on V , and that \mathcal{Q} is a partition of V into at most q parts of equal size. Then there exist $(k-1)$ -families of partitions \mathcal{P}^* and $\tilde{\mathcal{P}}^*$ on V with ground partitions \mathcal{P} and $\tilde{\mathcal{P}}$ respectively such that*

- (a) \mathcal{P} refines \mathcal{Q} ,
- (b) \mathcal{P}^* is generated from $\tilde{\mathcal{P}}^*$ by \mathcal{P} ,
- (c) \mathcal{P} has $t = p(t_1^*)\tilde{t}$ parts, where \tilde{t} is the number of parts of $\tilde{\mathcal{P}}$,
- (d) $\tilde{\mathcal{P}}^*$ and \mathcal{P}^* are $(t_0, t_1, \varepsilon(t_1))$ -equitable with equal density vectors, and all densities are at least $1/t_1^*$,
- (e) for each $1 \leq i \leq s$, G_i is $(\varepsilon_k, r(t_1))$ -regular with respect to \mathcal{P}^* ,
- (f) for each $1 \leq i \leq s$, each $1 \leq \ell \leq 1/\varepsilon_k$, each k -graph H equipped with a set of distinct root vertices x_1, \dots, x_ℓ such that $v(H) \leq 1/\varepsilon_k$, any distinct vertices v_1, \dots, v_ℓ in V , any slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ and \mathcal{J} through \mathcal{P}^* , any $(v(H) - \ell)$ -set of clusters \tilde{X} of $\tilde{\mathcal{P}}^*$ and any \tilde{X} -consistent $(v(H) - \ell)$ -set X of clusters of \mathcal{P}^* such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}[\tilde{X}]$, we have

$$d_H(G_i; v_1, \dots, v_\ell, \tilde{\mathcal{J}}[\tilde{X}]) = d_H(G_i; v_1, \dots, v_\ell, \mathcal{J}[X]) \pm \varepsilon_k.$$

We will prove Lemma 22 in Appendix B. However, it is worth noting that we need it only for proving property (c) of Lemma 6. The only use of this property in this paper is in the sketch proof of Theorem 39 given in Section 10, so the reader whose interest lies only in verifying Theorems 1 and 2 can safely forget Lemma 22 and pretend that Lemma 20 provides the $(k-1)$ -family of partitions \mathcal{P}^* in the proof of Lemma 6.

7.2. Tools for working with regularity. We will need various standard results in the proofs of Lemmas 6 and 9, which we present here.

We start with the Hypergraph Counting Lemma. The version we present here is slightly modified from a result in [8], which in turn was derived from [31, Theorem 9]. We sketch in Appendix A how this modified version can be obtained. Similar results were also proved previously by Gowers [11] and by Nagle, Rödl and Schacht [25]. We remark that Fact 3 is a special case of this lemma. Note though that in contrast to Fact 3, Lemma 23 allows for a k -graph G on top of the equitable complex \mathcal{J} , whose regularity ε_k with respect to \mathcal{J} can be bigger than the entries of the density vector of \mathcal{J} .

Lemma 23 (Counting Lemma, [8, Lemma 4]). *Let k, s, r, m_0 be positive integers, and let $d_2, \dots, d_{k-1}, \varepsilon, \varepsilon_k, \beta$ be positive constants such that $1/d_i \in \mathbb{N}$ for any $2 \leq i \leq k-1$ and*

$$\frac{1}{m_0} \ll \frac{1}{r}, \varepsilon \ll \varepsilon_k, d_2, \dots, d_{k-1} \quad \text{and} \quad \varepsilon_k \ll \beta, \frac{1}{s}.$$

Then the following holds for all integers $m \geq m_0$. Let H be a k -graph on s vertices $1, \dots, s$, and let \mathcal{H} be the k -complex generated by the down-closure of H . Also let \mathcal{J} be a $(\cdot, \cdot, \varepsilon)$ -equitable $(k-1)$ -complex with s clusters V_1, \dots, V_s each of size m and with density vector \mathbf{d} . Finally, let G be a k -graph on $\bigcup_{i \in [s]} V_i$ which is supported on $\mathcal{J}^{(k-1)}$ such that for any edge $e \in H$ the graph G is (ε_k, r) -regular with respect to the k -set of clusters $\{V_j : j \in e\}$. Then the number of copies of \mathcal{H} in G such that i is in V_i for each i is

$$\left(\prod_{e \in H} d^*(\{V_j : j \in e\}) \pm \beta \right) \left(\prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H})} \right) m^s.$$

□

A key property of regular complexes is that the restriction of such a complex to a large subset of its vertex set is also a regular complex, with the same relative densities at each level of the complex, albeit with somewhat degraded regularity properties. The next lemma states this property formally. Its proof is as sketched for Lemma 4.1 in [21] (the quantification there is slightly different but this does not affect the proof).

Lemma 24 (Regular Restriction Lemma). *Suppose integers k, m and real $\alpha, \varepsilon, \varepsilon_k, d_2, \dots, d_k > 0$ are such that*

$$\frac{1}{m} \ll \varepsilon \ll \varepsilon_k, d_2, \dots, d_{k-1} \quad \text{and} \quad \varepsilon_k \ll \alpha, \frac{1}{k}.$$

For any $r, s \in \mathbb{N}$ and $d_k > 0$, set $\mathbf{d} = (d_k, \dots, d_2)$, and let \mathcal{G} be an s -partite k -complex whose vertex classes V_1, \dots, V_s each have size m and which is $(\mathbf{d}, \varepsilon_k, \varepsilon, r)$ -regular. Choose any $V'_i \subseteq V_i$ with $|V'_i| \geq \alpha m$ for each $i \in [s]$. Then the induced subcomplex $\mathcal{G}[V'_1 \cup \dots \cup V'_s]$ is $(\mathbf{d}, \sqrt{\varepsilon_k}, \sqrt{\varepsilon}, r)$ -regular. □

Given a copy of some subgraph $H' \subseteq H$ in $\mathcal{G}^{(k)}$, how many ways are there to extend H' to a copy of H in $\mathcal{G}^{(k)}$? The next lemma gives a lower bound on this number for almost all copies of H' in $\mathcal{G}^{(k)}$. To state this precisely we make the following definitions.

Let \mathcal{G} be an s -partite k -complex whose vertex classes V_1, \dots, V_s are each of size m , and let \mathcal{H} be an s -partite k -complex whose vertex classes X_1, \dots, X_s each have size at most m . We say that an embedding of \mathcal{H} in \mathcal{G} is *partition-respecting* if for any $i \in [s]$ the vertices of X_i are embedded within V_i . We denote the set of labelled partition-respecting copies of \mathcal{H} in \mathcal{G} by $\mathcal{H}_{\mathcal{G}}$. The Extension Lemma [8, Lemma 5] states that if \mathcal{H}' is an induced subcomplex of \mathcal{H} , and \mathcal{G} is regular with $\mathcal{G}^{(k)}$ reasonably dense, then almost all partition-respecting copies of \mathcal{H}' in \mathcal{G} can be extended to a large number of copies of \mathcal{H}

in \mathcal{G} . We present a slightly modified form of this lemma, implied by [21, Lemma 4.6]. Again, we explain in Appendix A how the modification can be obtained.

Lemma 25 (Extension Lemma, [21, Lemma 4.6]). *Let k, s, r, b, b', m be positive integers, where $b' < b$, and let $c, \beta, d_2, \dots, d_k, \varepsilon, \varepsilon_k$ be positive constants such that $1/d_i \in \mathbb{N}$ for any $2 \leq i \leq k-1$ and*

$$\frac{1}{m} \ll \frac{1}{r}, \varepsilon \ll c \ll \varepsilon_k, d_2, \dots, d_{k-1} \quad \text{and} \quad \varepsilon_k \ll \beta, d_k, \frac{1}{s}, \frac{1}{b}.$$

Suppose that \mathcal{H} is an s -partite k -complex on b vertices with vertex classes X_1, \dots, X_s and let \mathcal{H}' be an induced subcomplex of \mathcal{H} on b' vertices. Suppose that \mathcal{G} is an s -partite k -complex with vertex classes V_1, \dots, V_s , all of size m , such that $\bigcup_{i=0}^{k-1} \mathcal{G}^{(i)}$ is $(\cdot, \cdot, \varepsilon)$ -equitable with density vector (d_{k-1}, \dots, d_2) . Suppose further that for each $e \in \mathcal{H}^{(k)}$ with index $A \in \binom{[s]}{k}$, the k -graph $\mathcal{G}^{(k)}[V_A]$ is (d, ε_k, r) -regular with respect to $\mathcal{G}^{(k-1)}[V_A]$ for some $d \geq d_k$. Then all but at most $\beta|\mathcal{H}'_G|$ labelled partition-respecting copies of \mathcal{H}' in \mathcal{G} extend to at least $cm^{b-b'}$ labelled partition-respecting copies of \mathcal{H} in \mathcal{G} . \square

8. PROOF OF THE REGULAR SLICE LEMMA

In this section we prove Lemma 6. We begin with an outline of the proof, considering the case of regularising only one k -graph $G = G_1$. Let \mathcal{P}^* be an equitable $(k-1)$ -family of partitions obtained from Lemma 22. As previously noted, we will take a ‘slice’ through \mathcal{P}^* so that the i -cells of this ‘slice’ are consistent. This can be done by the following procedure: for each pair of clusters we choose one of the 2-cells of \mathcal{P}^* on these clusters. We then throw out all other 2-cells and everything in higher levels of \mathcal{P}^* which is not supported on our chosen 2-cells. Now for each triple of clusters we choose one of the 3-cells of \mathcal{P}^* on these clusters, and so on.

Since the family of partitions \mathcal{P}^* is (t_0, t_1, ε) -equitable, the procedure described will always output a (t_0, t_1, ε) -equitable $(k-1)$ -complex \mathcal{J} . Because G is also regular with respect to \mathcal{P}^* , we could hope that \mathcal{J} will be a regular slice for G . But in order for this to be true, it is necessary that we do not accidentally choose the $(k-1)$ -cells of disproportionately many polyads $\hat{P}(Q)$ such that G is not regular with respect to $\hat{P}(Q)$. This already suggests our proof method: we will follow the above procedure, and when we are required to pick an i -cell for some i -set of clusters, we will choose uniformly at random from the $1/d_i$ possibilities. It is easy to see that the expected fraction of irregular k -sets of clusters of \mathcal{J} is then equal to the fraction of irregular polyads in \mathcal{P}^* . So by linearity of expectation there exists a regular slice for G , an idea which was previously observed in the $k=3$ case by Haxell, Łuczak, Peng, Rödl, Ruciński and Skokan [13]. However, a straightforward application of McDiarmid’s inequality shows that the fraction of irregular sets is actually sharply concentrated, so that the $(k-1)$ -complex \mathcal{J} obtained from the random procedure is very likely to be a regular slice. A similar argument together with the union bound guarantees that \mathcal{J} is very likely to satisfy both properties (a) and (b) of Lemma 6.

So far the properties of \mathcal{P}^* we used would also be guaranteed by Lemma 20. However, this will not be true for the argument proving property (c). The obstacle to obtaining property (c) is as follows: for any H we are asking for roughly n^{1/ε_k} distinct inequalities to be satisfied, one for each choice of root vertices. Although each individual one of these inequalities is very likely to be true, their success probability only depends on the regularity ε_k and the number of clusters t , but not on n . Hence a union bound fails miserably. In order to circumvent this, we make use of the coarse $(k-1)$ -family of partitions $\tilde{\mathcal{P}}^*$ given by Lemma 22. We will illustrate our strategy in the case that H is

a k -edge, rooted at a single vertex. Each vertex v of G defines a $(k-1)$ -graph, called the *link of v* , whose edges are those $(k-1)$ -sets which together with v form edges of G . Then in this case, property (c) states that for each $v \in V(G)$, the density of the link of v is close to the density of the link of v supported on \mathcal{J} .

Observe that most $(k-1)$ -sets in $V(G)$ are $\tilde{\mathcal{P}}$ -partite, and the $\tilde{\mathcal{P}}$ -partite $(k-1)$ -sets are partitioned by $\tilde{\mathcal{P}}^*$ into its $(k-1)$ -cells. These $(k-1)$ -cells are of approximately equal size by Fact 3, and it follows that for each v the density of the link of v is close to the average, over $(k-1)$ -cells C of $\tilde{\mathcal{P}}^*$, of the density of the link of v on C . Now each $(k-1)$ -cell C of $\tilde{\mathcal{P}}^*$ is partitioned by \mathcal{P} into a large number of $(k-1)$ -cells of \mathcal{P}^* , and Lemma 22(f) guarantees that for each v the link density of v on each of these parts is close to the link density of v on C . It follows that if \mathcal{J} is a slice through $\tilde{\mathcal{P}}^*$ whose $(k-1)$ -cells are chosen with about equal frequency from the $(k-1)$ -cells of $\tilde{\mathcal{P}}^*$ then \mathcal{J} satisfies this case of property (c). We can again prove that this is the likely outcome using McDiarmid's inequality.

The form of McDiarmid's inequality we use is the following.

Theorem 26 (McDiarmid's inequality, [23]). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be such that there exists a vector $(c_1, \dots, c_n) \in \mathbb{R}^n$ with the following property. For each \mathbf{x} , if \mathbf{x}' differs from \mathbf{x} only in coordinate i then $|f(\mathbf{x}) - f(\mathbf{x}')| \leq c_i$. Now if (X_1, \dots, X_n) is a vector of independent random variables, then for each $a > 0$ we have*

$$\Pr \left(|f(X_1, \dots, X_n) - \mathbb{E} f(X_1, \dots, X_n)| \geq a \right) \leq 2 \exp \left(\frac{-a^2}{2 \sum_{i=1}^n c_i^2} \right).$$

□

We will also make several uses of the following probabilistic statement which shows that $\mathbb{P}(E | C \cap S)$ and $\mathbb{P}(E | S)$ are nearly the same when S 'almost implies' C .

Proposition 27. *Let E, C and S be events in some probability space such that $\mathbb{P}(S) > 0$ and $\mathbb{P}(C^c | S) \leq \varepsilon$ for some constant $\varepsilon \leq \frac{1}{2}$. Then $\mathbb{P}(E | C \cap S) = \mathbb{P}(E | S) \pm 2\varepsilon$.*

Proof. Observe that $\mathbb{P}(C \cap S) = \mathbb{P}(S)(1 - \mathbb{P}(C^c | S)) > 0$, so $\mathbb{P}(E | C \cap S)$ is well-defined, and

$$\begin{aligned} (1 + 2\varepsilon)\mathbb{P}(E | S) &\geq \frac{\mathbb{P}(E \cap S)}{(1 - \varepsilon)\mathbb{P}(S)} \geq \frac{\mathbb{P}(E \cap S)}{\mathbb{P}(S) - \mathbb{P}(C^c \cap S)} \\ &\geq \mathbb{P}(E | C \cap S) = \frac{\mathbb{P}(E \cap C \cap S)}{\mathbb{P}(C \cap S)} \geq \frac{\mathbb{P}(E \cap S) - \mathbb{P}(E \cap C^c \cap S)}{\mathbb{P}(S)} \\ &\geq \frac{\mathbb{P}(E \cap S)}{\mathbb{P}(S)} - \frac{\mathbb{P}(C^c \cap S)}{\mathbb{P}(S)} \geq \mathbb{P}(E | S) - \varepsilon. \end{aligned}$$

□

Proof of Lemma 6. Given integers q, t_0 and s , a constant ε_k , and functions $r : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$, we define further constants as follows. Without loss of generality we assume that $\varepsilon_k \leq 1$, and define

$$(3) \quad t_0^* := \max \left(t_0, \frac{128k^2}{\varepsilon_k^{2+2/\varepsilon_k+k}}, \frac{4096k!^4 k^{2k} q^k s^2}{\varepsilon_k^4} \right).$$

Next we choose a constant

$$(4) \quad \varepsilon_k^* \leq \frac{\varepsilon_k^{2k+1}}{40s \cdot 2^k}.$$

so that ε_k^* is sufficiently small to apply Lemma 23 with ε_k^* , $\varepsilon_k/100$ and $1/\varepsilon_k$ in place of ε_k , β and s respectively. We choose $p : \mathbb{N} \rightarrow \mathbb{N}$ to be any monotone increasing function such that for all $x \in \mathbb{N}$ we have

$$(5) \quad 2 \exp \left(- 2^{-x-9} \varepsilon_k^2 p(x)^2 \prod_{j=2}^{k-1} \left(\frac{1}{x} \right)^{2 \binom{x}{j}} \right) < 2^{-xp(x)} \prod_{j=2}^{k-1} \left(\frac{1}{x} \right)^{\binom{x}{j}}.$$

It may not be immediately obvious that this is possible, but observe that if $p(x) > \max(2x/c_1(x), 2/c_2(x))$ then we have

$$2 \exp \left(- c_1(x)p(x)^2 \right) < 2 \exp \left(- 2xp(x) \right) < c_2(x)2^{-xp(x)},$$

from which (5) follows by appropriate choice of c_1, c_2 . We choose strictly monotone functions $r^* : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon^* : \mathbb{N} \rightarrow (0, 1]$ such that

$$(6) \quad \varepsilon^*(t_1) \leq \min \left(\varepsilon(t_1), \frac{\varepsilon_k}{10 \cdot 2^{2^{1/\varepsilon_k}}}, \frac{1}{2t_1^2} \right) \text{ and } r^*(t_1) \geq r(t_1)$$

for any $t_1 \in \mathbb{N}$. Moreover, we make these choices so that for any $t_1 \in \mathbb{N}$ we can apply Lemma 23 to count graphs on up to $1/\varepsilon_k$ vertices, with $r^*(t_1)$, $\varepsilon^*(t_1)$, ε_k^* and $\varepsilon_k/100$ in place of r , ε , ε_k and β respectively, with $1/t_1$ in place of each d_i , and so that we can apply Fact 3 with $\varepsilon_k/10$ in place of β to any $(t_0^*, t_1, \varepsilon^*(t_1))$ -equitable $(k-1)$ -complex with sufficiently large clusters. Let n_0^* , t_1^* and t_1 be obtained by applying Lemma 22 with inputs $q, t_0^*, s, \varepsilon_k^*$ and functions $r^*(\cdot)$, $\varepsilon^*(\cdot)$ and $p(\cdot)$. Note that we have $t_1 = p(t_1^*)t_1^*$ by Lemma 22. Finally, let $n_0 \geq n_0^*$ be sufficiently large to apply Lemma 23 and Fact 3 with $m_0 = n_0/t_1$ and all other constants as before. For the remainder of the proof we write ε , ε^* , r and r^* to denote $\varepsilon(t_1)$, $\varepsilon^*(t_1)$, $r(t_1)$ and $r^*(t_1)$ respectively.

Let V be a set of $n \geq n_0$ vertices, where n is divisible by $t_1!$, and let \mathcal{Q} partition V into at most q parts of equal size. Let G_1, \dots, G_s be edge-disjoint k -graphs on the vertex set V . Then we must show that there exists a $(k-1)$ -complex \mathcal{J} on V which is a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for each G_i , whose ground partition \mathcal{P} refines \mathcal{Q} , and which satisfies properties (a), (b) and (c) of the lemma. We start by applying Lemma 22 (with the inputs stated above), which yields partitions \mathcal{P}^* and $\tilde{\mathcal{P}}^*$ with the properties stated in that lemma. In particular, \mathcal{P}^* and $\tilde{\mathcal{P}}^*$ are both $(t_0^*, t_1, \varepsilon^*)$ -equitable with the same density vector $\mathbf{d} = (d_{k-1}, \dots, d_2)$, the ground partition \mathcal{P} of \mathcal{P}^* refines \mathcal{Q} , and each G_i is (ε_k^*, r^*) -regular with respect to \mathcal{P}^* .

Claim 28. *For each $2 \leq i \leq k-1$ and each \mathcal{P} -partite set $Q \in \binom{V}{i}$, the number of i -cells of \mathcal{P}^* supported on the polyad $\hat{P}(Q, \mathcal{P}^*)$ is precisely equal to $1/d_i$. Moreover, the same statement holds with $\tilde{\mathcal{P}}$ and $\tilde{\mathcal{P}}^*$ in place of \mathcal{P} and \mathcal{P}^* respectively.*

Proof. We prove the claim for \mathcal{P}^* ; the proof for $\tilde{\mathcal{P}}^*$ is identical. Recall that part (c) of the definition of an equitable family of partitions implies that $1/d_i$ must be an integer, and that each i -cell of \mathcal{P}^* supported on $\hat{P}(Q)$ is $(d_i, \varepsilon^*, 1)$ -regular with respect to $\hat{P}(Q)$. In particular, the number of i -sets in each of these i -cells is $(d_i \pm \varepsilon^*) |K_i(\hat{P}(Q))|$. Suppose for a contradiction that for some polyad $\hat{P}(Q)$, the number of i -cells of \mathcal{P}^* supported on $\hat{P}(Q)$ is at least $\frac{1}{d_i} + 1$. Since these i -cells are pairwise-disjoint and cover $K_i(\hat{P}(Q))$, we conclude that

$$|K_i(\hat{P}(Q))| \geq \left(\frac{1}{d_i} + 1 \right) (d_i \pm \varepsilon^*) |K_i(\hat{P}(Q))|.$$

By Fact 3 this number is non-zero, hence $1 \geq \left(\frac{1}{d_i} + 1 \right) (d_i \pm \varepsilon^*)$, which is a contradiction, since we have $d_i \geq 1/t_1$ and $\varepsilon^* \leq 1/(2t_1^2)$. A similar argument shows that the number of i -cells of \mathcal{P}^* supported on $\hat{P}(Q)$ is bigger than $\frac{1}{d_i} - 1$, giving the desired result. \square

We now construct our regular slice \mathcal{J} , whose clusters are the clusters of \mathcal{P}^* (that is, the parts of \mathcal{P}), by applying Algorithm 1 to \mathcal{P}^* . We claim that with positive probability \mathcal{J} satisfies the conclusions of Lemma 6.

Algorithm 1: Taking a slice

```

let  $\mathcal{J} = \{\emptyset\} \cup \{\{v\} : v \in V\}$ ;
foreach  $2 \leq i \leq k - 1$  do
    foreach  $X \in \binom{\mathcal{P}}{i}$  do
        let  $\mathcal{C} := \{C : C \text{ is an } i\text{-cell of } \mathcal{P}^* \text{ and is supported on } \mathcal{J}_{X^{(i-1)}}\}$ ;
        choose  $C \in \mathcal{C}$  uniformly at random;
        let  $\mathcal{J} := \mathcal{J} \cup C$ ;
    end
end
return  $\mathcal{J}$ ;

```

It is immediate from the algorithm that the output \mathcal{J} is a $(k - 1)$ -complex which is a slice through \mathcal{P}^* . Since \mathcal{P}^* is $(t_0^*, t_1, \varepsilon^*)$ -equitable with density vector $\mathbf{d} = (d_{k-1}, \dots, d_2)$, \mathcal{J} is also $(t_0^*, t_1, \varepsilon^*)$ -equitable with the same density vector. Let t denote the number of clusters of \mathcal{J} , so $t_0^* \leq t \leq t_1$, and let $m := n/t$, so m is the common size of each cluster. Note that by Claim 28 every slice through \mathcal{P}^* has probability precisely $\prod_{i=2}^{k-1} d_i^{\binom{t}{i}}$ of being selected as \mathcal{J} .

To show that \mathcal{J} is likely to be a regular slice for each G_i , we must bound the number of *irregular polyads*, that is, polyads $\hat{\mathcal{J}}_X$ for which some G_i with $1 \leq i \leq s$ is not (ε_k^*, r^*) -regular with respect to $\hat{\mathcal{J}}_X$. Since each G_i is (ε_k^*, r^*) -regular with respect to \mathcal{P}^* , the number of \mathcal{P} -partite sets $Q \in \binom{V}{k}$ for which G_i is not (ε_k^*, r^*) -regular with respect to $\hat{P}(Q)$ is at most $\varepsilon_k^* \binom{n}{k}$. On the other hand, since ε^* is chosen small enough so that we can apply Fact 3 with $\varepsilon_k/10$ in place of β , each polyad $\hat{P}(Q)$ of \mathcal{P}^* supports

$$(1 \pm \frac{\varepsilon_k}{10}) m^k \prod_{i=2}^{k-1} d_i^{\binom{k}{i}}$$

\mathcal{P} -partite members of $\binom{V}{k}$. It follows that the number of irregular polyads in \mathcal{P}^* is at most

$$(7) \quad \frac{s \varepsilon_k^* \binom{n}{k}}{(1 - \frac{\varepsilon_k}{10}) m^k \prod_{i=2}^{k-1} d_i^{\binom{k}{i}}} \leq 2s \varepsilon_k^* \binom{t}{k} \prod_{i=2}^{k-1} d_i^{-\binom{k}{i}}.$$

By Claim 28 the probability that a given polyad of \mathcal{P}^* is chosen for \mathcal{J} is precisely $\prod_{i=2}^{k-1} d_i^{\binom{k}{i}}$. So by linearity of expectation, the expected number of k -sets of clusters X such that \mathcal{J}_X is an irregular polyad is at most $2s \varepsilon_k^* \binom{t}{k}$. By Markov's inequality, we conclude that with probability at least $1/2$ this number is at most $4s \varepsilon_k^* \binom{t}{k}$, which by (4) is at most $\varepsilon_k \binom{t}{k}$. In particular, \mathcal{J} is a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for each G_i with probability at least $1/2$. We stress however that later in this proof we will make use of the stronger property that \mathcal{J} contains at most $4s \varepsilon_k^* \binom{t}{k}$ polyads with respect to which some G_i is not (ε_k^*, r^*) -regular.

It remains to show that properties (a), (b) and (c) each hold with high probability. In each case, this amounts to verifying that some random variables are likely to all be

close to certain values, and our proof in each case follows the pattern of first showing that the expectations of these random variables are indeed close to the desired values, then showing that the random variables are sufficiently concentrated to apply a union bound. We begin with property (a).

The following Claim 29 shows that for any i the reduced k -graph $R(G_i)$ of G_i (with respect to \mathcal{J}) has approximately the same H -density as G_i for each small H in expectation, and that this remains true when considering the subgraph induced on a large subset of the clusters. For this, we say that a copy of H in G_i is \mathcal{P} -crossing if it has at most one vertex in any part of \mathcal{P} .

If it happens to be the case that the density of H -copies in G equals the density of \mathcal{P} -crossing H -copies, and ε and ε_k are zero, this claim is a triviality. The number of \mathcal{P} -crossing copies of H in $G_i[\bigcup X]$ then equals the number obtained by applying Lemma 23 to \mathcal{P}^* , which in this case means that $d_H(R(G_i)[X])$ equals $d_H(G[\bigcup X])$ precisely.

Of course none of these assumptions are true, and consequently errors are introduced: the proof of the following claim amounts to showing that since the assumptions are ‘almost’ true, the errors are small. It is convenient in the proof to use probabilistic language: the probability that a random injective map from $V(H)$ to $V(G)$ is a graph embedding is exactly the density of H -copies in G .

Claim 29. *Let H be a labelled k -graph on $h \leq 1/\varepsilon_k$ vertices, X be a set of at least $\varepsilon_k t$ clusters of \mathcal{J} , and $G = G_i$ for some $i \in [s]$. Then*

$$\left| \mathbb{E}_{\mathcal{J}} d_H(R(G)[X]) - d_H\left(G\left[\bigcup X\right]\right) \right| \leq \varepsilon_k/2.$$

Proof. Let the vertices of H be labelled with the integers $1, \dots, h$. Choose an injective map $\phi : V(H) \rightarrow \bigcup X$ uniformly at random. Say that ϕ is \mathcal{P} -crossing if no two vertices of H are mapped to the same part of \mathcal{P} , and let CROSS be the event that ϕ is \mathcal{P} -crossing. Also let EMB denote the event that ϕ is an embedding of H into G . We prove the claim by calculating the conditional probability $\mathbb{P}(\text{EMB} \mid \text{CROSS})$ (that is, the proportion of \mathcal{P} -crossing h -tuples which form \mathcal{P} -crossing copies of H in $G[\bigcup X]$) in two different ways.

First, note that we have $\mathbb{P}(\text{EMB}) = d_H(G[\bigcup X])$ by definition. Recall that $m = n/t$ is the size of each cluster of \mathcal{J} . Then there are $h! \binom{|X|m}{h}$ h -tuples in $\bigcup X$. At most $h! \binom{|X|m}{h-1} hm$ of these h -tuples are not \mathcal{P} -crossing. Since $t \geq t_0^* \geq 20/\varepsilon_k^4$ by (3) and $|X| \geq \varepsilon_k t$ we have $h \leq 1/\varepsilon_k \leq \varepsilon_k t \leq |X|m/2$. So the probability of the complement CROSS^c is at most

$$\mathbb{P}(\text{CROSS}^c) \leq \frac{h! \binom{|X|m}{h-1} hm}{h! \binom{|X|m}{h}} = \frac{h^2 m}{|X|m - h + 1} \leq \frac{2h^2}{|X|} \leq \frac{2}{\varepsilon_k^3 t} \leq \frac{\varepsilon_k}{10}.$$

We may therefore apply Proposition 27 with EMB, CROSS and the certain event in place of E, C and S respectively, which gives

$$(8) \quad \mathbb{P}(\text{EMB} \mid \text{CROSS}) = \mathbb{P}(\text{EMB}) \pm \frac{\varepsilon_k}{5} = d_H\left(G\left[\bigcup X\right]\right) \pm \frac{\varepsilon_k}{5}.$$

We now turn to the second way of evaluating $\mathbb{P}(\text{EMB} \mid \text{CROSS})$, this time in terms of $\mathbb{E}_{\mathcal{J}} d_H(R(G)[X])$. For any h -tuple $\mathcal{C} := (C_1, \dots, C_h)$ of distinct clusters of \mathcal{J} , we say that a copy of H in G is \mathcal{C} -crossing if vertex j of H lies in the cluster C_j for each $j \in [h]$. Similarly, we say that ϕ is \mathcal{C} -crossing if $\phi(j) \in C_j$ for each $j \in [h]$, an event which we denote by $\text{CROSS}_{\mathcal{C}}$.

Fix some h -tuple $\mathcal{C} := (C_1, \dots, C_h)$, and let \mathcal{H} be the k -complex generated by the down-closure of H . For any slice \mathcal{J} through \mathcal{P}^* and any $e \in \mathcal{H}$ let \mathcal{J}_e denote the cell of \mathcal{J} on the clusters $\{C_j : j \in e\}$, and define $\mathcal{J}(H, \mathcal{C}) := \bigcup_{e \in \mathcal{H}} \mathcal{J}_e$ (so if ϕ is \mathcal{C} -crossing, then

$\mathcal{J}(H, \mathcal{C})$ is the subcomplex of \mathcal{J} consisting of all cells of \mathcal{J} on sets of clusters to which ϕ maps an edge of \mathcal{H} . We can then extend $\mathcal{J}(H, \mathcal{C})$ to a slice through \mathcal{P}^* by appropriate choices of cells on each remaining set of clusters, and by Claim 28, the number of such extensions does not depend on \mathcal{J} . We therefore have

$$(9) \quad \mathbb{E}_{\mathcal{J}} \prod_{e \in H} d^*(\{C_j : j \in e\}) = \mathbb{E}_{\mathcal{J}(H, \mathcal{C})} \prod_{e \in H} d^*(\{C_j : j \in e\}).$$

Let $q_{\mathcal{J}(H, \mathcal{C})}$ denote the number of \mathcal{C} -crossing copies of H in G which are supported on $\mathcal{J}(H, \mathcal{C})$. Then for each fixed $\mathcal{J}(H, \mathcal{C})$ we have

$$\mathbb{P}(\text{EMB and } \phi(H) \text{ is supported on } \mathcal{J}(H, \mathcal{C}) \mid \text{CROSS}_{\mathcal{C}}) = \frac{q_{\mathcal{J}(H, \mathcal{C})}}{m^h}.$$

Observe that for each \mathcal{C} -crossing copy of H in G there is precisely one possibility for $\mathcal{J}(H, \mathcal{C})$ on which this copy is supported. Moreover, by Claim 28 there are precisely $\prod_{j=2}^{k-1} d_j^{-e_j(\mathcal{H})}$ different choices for $\mathcal{J}(H, \mathcal{C})$, each of which is equally likely to occur. Hence we have

$$(10) \quad \begin{aligned} \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\mathcal{C}}) &= \frac{1}{m^h} \sum_{\mathcal{J}(H, \mathcal{C})} q_{\mathcal{J}(H, \mathcal{C})} \\ &= \frac{1}{m^h} \left(\prod_{j=2}^{k-1} d_j^{-e_j(\mathcal{H})} \right) \mathbb{E}_{\mathcal{J}(H, \mathcal{C})} q_{\mathcal{J}(H, \mathcal{C})}. \end{aligned}$$

If \mathcal{J} has the property that G is (ε_k^*, r^*) -regular with respect to $\hat{\mathcal{J}}_Z$ for every k -set $Z \subseteq \mathcal{C}$, then, because we chose ε_k^* , ε^* and r^* such that we can apply Lemma 23 with $\beta = \varepsilon_k/100$, the number of \mathcal{C} -crossing copies of H in G which are supported on $\mathcal{J}(H, \mathcal{C})$ is

$$q_{\mathcal{J}(H, \mathcal{C})} = \left(\prod_{e \in H} d^*(\{C_j : j \in e\}) \pm \frac{\varepsilon_k}{100} \right) \left(\prod_{j=2}^{k-1} d_j^{e_j(\mathcal{H})} \right) m^h.$$

If on the other hand \mathcal{J} does not satisfy this property, we also need some bound on $q_{\mathcal{J}(H, \mathcal{C})}$. This number is obviously at least zero, and is at most the number of copies of H in the k -graph whose edges are all k -tuples supported on $\bigcup_{Z \in \binom{\mathcal{C}}{k}} \hat{\mathcal{J}}_Z$. Since this k -graph is automatically (ε_k^*, r^*) -regular with respect to \mathcal{J} , we can apply Lemma 23 to find that the number of \mathcal{C} -crossing copies of H in G supported on $\mathcal{J}(H, \mathcal{C})$ satisfies

$$q_{\mathcal{J}(H, \mathcal{C})} \leq \left(1 + \frac{\varepsilon_k}{100} \right) \left(\prod_{j=2}^{k-1} d_j^{e_j(\mathcal{H})} \right) m^h \leq 2 \left(\prod_{j=2}^{k-1} d_j^{e_j(\mathcal{H})} \right) m^h,$$

and hence

$$q_{\mathcal{J}(H, \mathcal{C})} = \left(\prod_{e \in H} d^*(\{C_j : j \in e\}) \pm 2 \right) \left(\prod_{j=2}^{k-1} d_j^{e_j(\mathcal{H})} \right) m^h.$$

Now let $\zeta(\mathcal{C})$ be the fraction of slices \mathcal{J} through \mathcal{P}^* for which G is not (ε_k^*, r^*) -regular with respect to $\hat{\mathcal{J}}_Z$ for some k -set $Z \subseteq \mathcal{C}$. Putting the above two estimates for $q_{\mathcal{J}(H, \mathcal{C})}$ together with (10), we have

$$\begin{aligned} \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\mathcal{C}}) &= \left(\mathbb{E}_{\mathcal{J}(H, \mathcal{C})} \prod_{e \in H} d^*(\{C_j : j \in e\}) \right) \pm \left(2\zeta(\mathcal{C}) + \frac{\varepsilon_k}{100} \right) \\ &\stackrel{(9)}{=} \left(\mathbb{E}_{\mathcal{J}} \prod_{e \in H} d^*(\{C_j : j \in e\}) \right) \pm \left(2\zeta(\mathcal{C}) + \frac{\varepsilon_k}{100} \right). \end{aligned}$$

We now condition on the event CROSS. Observe that there is then precisely one h -tuple $\mathcal{C} = (C_1, \dots, C_h)$ for which the event CROSS $_{\mathcal{C}}$ occurs (i.e. ϕ is \mathcal{C} -crossing). We take \mathcal{C} to be this k -tuple, so \mathcal{C} is now a random variable. Moreover, since ϕ was chosen uniformly at random and each cluster has equal size, the k -tuple \mathcal{C} is chosen uniformly at random among all possibilities, which implies that

$$(11) \quad \mathbb{P}(\text{EMB} \mid \text{CROSS}) = \left(\mathbb{E}_{\mathcal{C}} \mathbb{E}_{\mathcal{J}} \prod_{e \in H} d^*(\{C_j : j \in e\}) \right) \pm \left(2 \mathbb{E}_{\mathcal{C}} \zeta(\mathcal{C}) + \frac{\varepsilon_k}{100} \right).$$

However, $\mathbb{E}_{\mathcal{C}} \zeta(\mathcal{C})$ is simply the probability that our uniformly random \mathcal{C} has the property that G is not (ε_k^*, r^*) -regular with respect to $\hat{\mathcal{J}}_Z$ for some k -set $Z \subseteq \mathcal{C}$. Taking a union bound, this is at most $\binom{h}{k}$ multiplied by the probability that G is not (ε_k^*, r^*) -regular with respect to $\hat{\mathcal{J}}_Z$ for a k -set Z of clusters of X chosen uniformly at random. Since there are at least $\binom{\varepsilon_k t}{k}$ k -sets Z of clusters in X , and by (7) there are at most $2\varepsilon_k^* s \binom{t}{k} \prod_{j=2}^{k-1} d_j^{-\binom{k}{j}}$ irregular polyads of \mathcal{P}^* , each of which is chosen for \mathcal{J} with probability $\prod_{j=2}^{k-1} d_j^{\binom{k}{j}}$, we find that

$$(12) \quad \mathbb{E}_{\mathcal{C}} \zeta(\mathcal{C}) \leq \binom{h}{k} \frac{2\varepsilon_k^* s \binom{t}{k}}{\binom{\varepsilon_k t}{k}} \stackrel{(3)}{\leq} \frac{2\varepsilon_k^* s t^k h^k}{(\varepsilon_k t/2)^k} \stackrel{(4)}{\leq} \frac{\varepsilon_k}{20}.$$

Finally, since \mathcal{C} is chosen uniformly at random, by definition we have $d_H(R(G)[X]) = \mathbb{E}_{\mathcal{C}} \prod_{e \in H} d^*(\{C_j : j \in e\})$, and so

$$(13) \quad \mathbb{E}_{\mathcal{J}} d_H(R(G)[X]) = \mathbb{E}_{\mathcal{C}} \mathbb{E}_{\mathcal{J}} \prod_{e \in H} d^*(\{C_j : j \in e\}).$$

Combining equations (11), (12) and (13) we conclude that

$$\mathbb{E}_{\mathcal{J}} d_H(R(G)[X]) = \mathbb{P}(\text{EMB} \mid \text{CROSS}) \pm \frac{\varepsilon_k}{5}$$

and so by (8) we have

$$\mathbb{E}_{\mathcal{J}} d_H(R(G)[X]) = d_H \left(G \left[\bigcup X \right] \right) \pm \frac{\varepsilon_k}{2}.$$

□

We next show that the random variable $d_H(R(G_i)[X])$ is concentrated about its mean using McDiarmid's inequality. We will require enough concentration of $d_H(R(G_i)[X])$ to make use of a union bound over all possible H and X .

Claim 30. *Let H be a k -graph on $h \leq 1/\varepsilon_k$ vertices, X a set of at least $\varepsilon_k t$ clusters of \mathcal{J} , and $G = G_i$ for some $i \in [s]$. Then the probability that $d_H(R(G)[X])$ deviates from $\mathbb{E} d_H(R(G)[X])$ by more than $\varepsilon_k/2$ is at most*

$$2^{-\binom{1/\varepsilon_k}{k}} 2^{-2t}.$$

Proof. We want to apply McDiarmid's inequality, which holds for a random variable which can be written as $f(\mathbf{x})$ for a vector \mathbf{x} whose entries are independent random variables. Hence we first argue that $d_H(R(G)[X])$ is of this form. We put an arbitrary order on the j -cells of \mathcal{P}^* . Let Y be a j -set of clusters. Then observe that when Algorithm 1 chooses the j -cell of \mathcal{J} for Y , we have already chosen the cells of $\mathcal{J}^{(j-1)}$, so, by Claim 28, \mathcal{J}_Y is chosen uniformly at random from the exactly d_j^{-1} j -cells supported on $\mathcal{J}^{(j-1)}[Y]$.

Equivalently, we choose a number p from 1 to d_j^{-1} uniformly at random and take the p th j -cell supported on $\mathcal{J}^{(j-1)}[Y]$ in \mathcal{P}^* . It follows that we can write

$$d_H(R(G)[X]) = f(\mathbf{x})$$

where \mathbf{x} is a vector of integers whose first $\binom{|X|}{2}$ coordinates are chosen independently uniformly at random in $[d_2^{-1}]$, whose next $\binom{|X|}{3}$ coordinates are chosen independently and uniformly at random in $[d_3^{-1}]$, and so on.

We now want to bound $|f(\mathbf{x}) - f(\mathbf{x}')|$ for two vectors \mathbf{x}, \mathbf{x}' differing in only one coordinate. In other words, we want to bound the change c_Y of $d_H(R(G)[X])$ when we change the choice of the j -cell on the set $Y \in \binom{X}{j}$ in Algorithm 1. Observe that this change can only affect weighted copies of H which use the j -cell Y . Moreover, $t \geq t_0^* \geq 2/\varepsilon_k^3$ by definition of t_0^* , and therefore $\varepsilon_k t \geq 2/\varepsilon_k^2 \geq 2h^2$. Consequently we have

$$(14) \quad c_Y \leq \frac{|X|^{h-j}}{\binom{|X|}{h} \cdot h!} \leq \frac{t^{h-j}}{\binom{\varepsilon_k t}{h} \cdot h!} \leq \frac{t^{h-j}}{(\varepsilon_k t - h)^h} \leq 2\varepsilon_k^{-h} t^{-j}$$

We therefore have

$$(15) \quad \sum_{j=2}^{k-1} \sum_{Y \in \binom{X}{j}} c_Y^2 \stackrel{(14)}{\leq} \sum_{j=2}^{k-1} \binom{t}{j} (2\varepsilon_k^{-h} t^{-j})^2 \leq \sum_{j=2}^{k-1} 4\varepsilon_k^{-2h} t^{-j} \leq 4k\varepsilon_k^{-2h} t^{-2}$$

Hence using Theorem 26, the probability that $f(\mathbf{x})$ differs from $\mathbb{E} f(\mathbf{x})$ by more than $\varepsilon_k/2$ is at most

$$\begin{aligned} 2 \exp\left(\frac{-\varepsilon_k^2/4}{8k\varepsilon_k^{-2h} t^{-2}}\right) &= 2 \exp\left(\frac{-\varepsilon_k^{2+2h} t^2}{32k}\right) \leq 2 \exp\left(\frac{-4t}{\varepsilon_k^k}\right) \\ &< \exp\left(\frac{-1}{\varepsilon_k^k}\right) \exp(-2t) < 2^{-(1/\varepsilon_k^k)} 2^{-2t}, \end{aligned}$$

where the first inequality holds since $h \leq 1/\varepsilon - k$ and $t \geq t_0^* \geq 128k^2/\varepsilon_k^{2+2/\varepsilon_k+k}$ by definition of t_0^* . \square

We can now prove that property (a) holds with high probability by taking a union bound over the s different k -graphs G_i , the at most 2^t choices of X and the $2^{\binom{1/\varepsilon_k}{k}}$ choices of H . Applying Claims 29 and 30 we deduce that with probability at least $1 - s2^{-t}$ we have that $d_H(G_i[\bigcup X]) = d_H(R(G_i)[X]) \pm \varepsilon_k$ for all $1 \leq i \leq s$, all k -graphs H with at most $1/\varepsilon_k$ vertices and all sets X of at least $\varepsilon_k t$ clusters.

Next we show that (b) holds with high probability also. This is achieved through a similar but simpler argument as for (a), so we will be brief. Note first that since each part of \mathcal{Q} has equal size, any set X of clusters for which $\bigcup X$ is the union of some parts of \mathcal{Q} has size at least t/q . So fix some $1 \leq j \leq k-1$, a set X of at least t/q clusters of \mathcal{J} , a j -set Y of clusters of \mathcal{J} , and some k -graph G on V . We first prove an analogue of Claim 29, that the expected value of $\overline{\deg}(Y; R_{\mathcal{J}}(G), X)$ conditioned upon the choices of all cells on subsets of Y , is within $\varepsilon_k/2$ of $\overline{\deg}(\mathcal{J}_Y; G, \bigcup X)$. To prove this we choose uniformly at random an edge $\{v_1, \dots, v_j\}$ of \mathcal{J}_Y (which is fixed since we are conditioning on the choice of cell for Y), and also choose uniformly at random vertices v_{j+1}, \dots, v_k from $\bigcup X$ so that the vertices v_i with $1 \leq i \leq k$ are all distinct. Let $e = \{v_1, \dots, v_k\}$. By definition $\overline{\deg}(\mathcal{J}_Y; G, \bigcup X) = \mathbb{P}(e \in G)$. Let CROSS be the event that e is \mathcal{P} -partite.

Using the fact that the set $\{v_1, \dots, v_j\}$ is automatically \mathcal{P} -partite, we have, similarly as in the proof of Claim 29, that

$$\mathbb{P}(\text{CROSS}) \geq 1 - \frac{\binom{|X|^{m-j}}{k-j-1} km}{\binom{|X|^{m-j}}{k-j}} \geq 1 - \frac{k^2 m}{|X|^{m-k+1}} \geq 1 - \frac{\varepsilon_k}{10},$$

and therefore

$$(16) \quad \overline{\text{deg}}(\mathcal{J}_Y; G, \bigcup X) = \mathbb{P}(e \in G \mid \text{CROSS}) \pm \frac{\varepsilon_k}{4}.$$

For each $i \in [j]$ let C_i be the cluster containing v_i . Now fix any $(k-j)$ distinct clusters C_{j+1}, \dots, C_k , and define \mathcal{C} to be the k -tuple (C_1, \dots, C_k) . Recall that $\mathcal{J}[\mathcal{C}]$ was defined to be the subcomplex of \mathcal{J} consisting of all cells of \mathcal{J} which are supported on the clusters of \mathcal{C} . By definition of relative density and Fact 3, the number of edges in $E_{\mathcal{C}}$ corresponding to any given choice of $\mathcal{J}[\mathcal{C}]$ is

$$d_{\mathcal{J}[\mathcal{C}]}^*(\mathcal{C} \mid K_k(\hat{\mathcal{J}}_{\mathcal{C}})) = (1 \pm \frac{\varepsilon_k}{10}) d_{\mathcal{J}[\mathcal{C}]}^*(\mathcal{C}) \prod_{i=2}^{k-1} d_i^{\binom{k}{i}} m^k,$$

Let $E_{\mathcal{C}}$ consist of the edges $f \in G$ which have precisely one vertex in each cluster of \mathcal{C} and which satisfy $f \cap Y \in \mathcal{J}_Y$. Then for each fixed \mathcal{C} , summing over all possibilities for $\mathcal{J}[\mathcal{C}]$ we have

$$\begin{aligned} |E_{\mathcal{C}}| &= \sum_{\mathcal{J}[\mathcal{C}]} (1 \pm \frac{\varepsilon_k}{10}) d_{\mathcal{J}[\mathcal{C}]}^*(\mathcal{C}) \prod_{i=2}^{k-1} d_i^{\binom{k}{i}} m^k \\ &= (1 \pm \frac{\varepsilon_k}{10}) \mathbb{E}_{\mathcal{J}[\mathcal{C}]} d_{\mathcal{J}[\mathcal{C}]}^*(\mathcal{C}) \prod_{i=2}^j d_i^{\binom{j}{i}} m^k \\ &= (1 \pm \frac{\varepsilon_k}{10}) \mathbb{E}_{\mathcal{J}} d_{\mathcal{J}}^*(\mathcal{C}) \prod_{i=2}^j d_i^{\binom{j}{i}} m^k. \end{aligned}$$

The second line comes from observing that by Claim 28, $E_{\mathcal{C}}$ is partitioned into $\prod_{i=2}^{k-1} d_i^{-\binom{k}{i} + \binom{j}{i}}$ sets according to the choice of $\mathcal{J}[\mathcal{C}]$, and the distribution over choices of $\mathcal{J}[\mathcal{C}]$ given by \mathcal{J} is uniform, which also gives the third line. The total number of possibilities for e which give \mathcal{C} is $|\mathcal{J}_Y| m^{k-j}$, and by Fact 3 we have

$$|\mathcal{J}_Y| m^{k-j} = (1 \pm \frac{\varepsilon_k}{10}) \prod_{i=2}^j d_i^{\binom{j}{i}} m^k.$$

We conclude that, writing $\text{CROSS}_{\mathcal{C}}$ for the event that $\phi(v_i) \in C_i$ for each $i \in [k]$, we have

$$\mathbb{P}(e \in G \mid \text{CROSS}_{\mathcal{C}}) = \frac{|E_{\mathcal{C}}|}{|\mathcal{J}_Y| m^{k-j}} = (1 \pm \frac{\varepsilon_k}{4}) \mathbb{E}_{\mathcal{J}}[d_{\mathcal{J}}^*(\mathcal{C})],$$

We now condition on the event CROSS . This implies that there is precisely one choice of distinct clusters C_{j+1}, \dots, C_k of $X \setminus Y$ for which the event $\text{CROSS}_{\mathcal{C}}$ occurs; we now take \mathcal{C} to be given by this choice. Since v_{j+1}, \dots, v_k were chosen uniformly at random, and each cluster has equal size, it follows that the $(k-j)$ -set $\{C_{j+1}, \dots, C_k\}$ is chosen uniformly at random from all $(k-j)$ -sets in $X \setminus Y$. This implies that

$$\mathbb{P}(e \in G \mid \text{CROSS}) = \mathbb{E}_{\mathcal{C}} \mathbb{E}_{\mathcal{J}}[d_{\mathcal{J}}^*(\mathcal{C})] \pm \varepsilon_k/4.$$

Furthermore, by definition we have $\mathbb{E}_{\mathcal{C}}(d_{\mathcal{J}}^*(\mathcal{C})) = \overline{\deg}(Y; R_{\mathcal{J}}(G), X)$, so

$$\mathbb{E}_{\mathcal{J}}[\overline{\deg}(Y; R_{\mathcal{J}}(G), X)] = \mathbb{P}(e \in G \mid \text{CROSS}) \pm \frac{\varepsilon_k}{4} \stackrel{(16)}{=} \overline{\deg}(\mathcal{J}_Y; G, \bigcup X) \pm \frac{\varepsilon_k}{2},$$

proving the analogue of Claim 29.

Next, we prove an analogue of Claim 30, namely that for any k -graph G on V , any $1 \leq j \leq k-1$, any j -set Y of clusters of \mathcal{J} and any set X of at least t/q clusters of \mathcal{J} , with high probability the random variable $\overline{\deg}(Y; R_{\mathcal{J}}(G), X)$ conditioned on the choice of all cells on subsets of Y is within $\varepsilon_k/2$ of its expectation conditioned on the choice of all cells on subsets of Y . This argument is very similar to the proof of Claim 30. As there, we can write $\overline{\deg}(Y; R_{\mathcal{J}}(G), X)$ conditioned on the choice of all cells on subsets of Y as a function $f'(\mathbf{x})$ where the entries of \mathbf{x} are integers corresponding to the choice of i -cell on each i -set in $X \cup Y$ not contained in Y , for each $2 \leq i \leq k-1$. Again, for each $Z \subseteq X \cup Y$ of size between 2 and $k-1$ not contained in Y , we let c_Z bound the difference $|f'(\mathbf{x}) - f'(\mathbf{x}')|$ for pairs of vectors differing only on the entry corresponding to Z . The value of $f'(\mathbf{x})$ is the average value of $d_{\mathcal{J}}^*(\mathcal{C})$ for sets \mathcal{C} of clusters in X containing Y , and changing the cell on Z affects only those containing $Y \cup Z$, so we have

$$c_Z \leq \frac{\binom{|X \setminus Z|}{k - |Y \cup Z|}}{\binom{|X \setminus Y|}{k - |Y|}} \leq 2k! |X|^{-|Z \setminus Y|},$$

where the last inequality is obtained by using $|X| \geq t/q$ and the choice of $t \geq t_0^* \geq 10q^k k^k \varepsilon_k^{-1}$ in (3). Note that $|Z \setminus Y| \geq 1$ by assumption, so summing over all choices of Z (where we write $|Z \setminus Y| = j$) we have

$$\sum_Z c_Z^2 \leq \sum_{j=1}^{k-1} 2^{|Y|} \binom{|X|}{j} 4k!^2 |X|^{-2j} \leq 4k!^2 k 2^{|Y|} |X|^{-1} \leq \frac{4k!^2 k 2^k q}{t},$$

where the final inequality uses $|X| \geq t/q$ and $|Y| \leq k$.

By Theorem 26 we find that the probability that $\overline{\deg}(Y; R(G), X)$ fails to be within $\varepsilon_k/2$ of $\mathbb{E}_{\mathcal{J}}[\overline{\deg}(Y; R(G), X)]$ (and therefore also the probability that $\overline{\deg}(\mathcal{J}_Y; G, \bigcup X)$ differs from $\overline{\deg}(Y; R(G), X)$ by more than ε_k) is at most

$$2 \exp\left(\frac{-\varepsilon_k^2 t}{32k!^2 k 2^k q}\right) \stackrel{(3)}{\leq} 2^{-2\sqrt{t}} \stackrel{(3)}{\leq} 2^{-\sqrt{t}} / (2^{q_s}).$$

For property (b) there are only at most 2^q possibilities for X (since we required that $\bigcup X$ is a union of parts of \mathcal{Q}), and at most $\binom{t}{1} + \dots + \binom{t}{k-1} \leq t^{k-1}$ possibilities for Y . So we can take a union bound over all $G = G_1, \dots, G_s$, all sets Y of between 1 and $k-1$ clusters of \mathcal{J} , and all sets X of clusters such that $\bigcup X$ is a union of parts of \mathcal{Q} , to deduce that property (b) holds with probability at least $1 - t^{k-1} 2^{-\sqrt{t}}$.

To complete the proof, we need to show that property (c) holds with high probability; the argument for this splits into two parts. The ‘probabilistic’ part is to show that \mathcal{J} takes about the same fraction of *every* slice (regular or otherwise) through $\tilde{\mathcal{P}}^*$. Let us briefly sketch how this helps us. We want to show that the density of rooted copies of H is about equal to the density of rooted copies of H supported on \mathcal{J} . The former density is easily seen to be close to the average over slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ of the density of rooted copies of H supported on $\tilde{\mathcal{J}}$. Now Lemma 22(f) implies that if $\tilde{\mathcal{J}}$ is any slice through $\tilde{\mathcal{P}}^*$, and X is any $\tilde{\mathcal{P}}$ -consistent collection of \tilde{t} clusters of \mathcal{P} , then the density of rooted copies of H supported on $\tilde{\mathcal{J}}$ is close to the density of rooted copies of H supported on $\tilde{\mathcal{J}}[X]$. Finally, since \mathcal{J} contains about the same number of subcomplexes of the form $\tilde{\mathcal{J}}[X]$ for

each $\tilde{\mathcal{J}}$, the density of rooted copies of H supported on \mathcal{J} is close to the average over $\tilde{\mathcal{J}}$ of the density of rooted copies of H supported on $\tilde{\mathcal{J}}$, which is what we want to show.

The proof of the following ‘probabilistic part’ follows the same pattern as we saw for parts (a) and (b), showing that a certain random variable has the desired expectation, and then using McDiarmid’s inequality to establish concentration. We remark that the latter depends on the fact that the number of clusters of \mathcal{J} is much greater than the number of slices through $\tilde{\mathcal{P}}^*$ by choice of $p(\cdot)$.

Claim 31. *With probability at least $1 - 2^{-t}$, for each slice $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$, the number of sets X of \tilde{t} clusters of \mathcal{P}^* , one in each cluster of $\tilde{\mathcal{P}}$, such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$, is*

$$(17) \quad \left(1 \pm \frac{\varepsilon_k}{16}\right) p(t_1^*)^{\tilde{t}} \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}}.$$

Proof. Let $\tilde{\mathcal{J}}$ be a slice through $\tilde{\mathcal{P}}^*$. If X is a fixed set of \tilde{t} clusters of \mathcal{P}^* , one in each cluster of $\tilde{\mathcal{P}}^*$, then, by Lemma 22(b), $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$ exactly if for each $2 \leq j \leq k-1$ and set Y of $\binom{X}{j}$, we happened to choose the one j -cell on Y which is a subset of the corresponding j -cell in $\tilde{\mathcal{J}}$. Conditioning on having done this for j' -cells for $j' < j$, the probability of doing so for any given Y is d_j by Claim 28, and these choices for different members of $\binom{X}{j}$ are independent. It follows that the probability that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$ is

$$\prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}}.$$

The total number of choices of X is $p(t_1^*)^{\tilde{t}}$, since each cluster of $\tilde{\mathcal{P}}^*$ is split into $p(t_1^*)$ clusters of \mathcal{P}^* by Lemma 22(c). By linearity of expectation, the expected number of sets X such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$ is

$$p(t_1^*)^{\tilde{t}} \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}}$$

as desired. Similarly as for parts (a) and (b), we can write the number of $\tilde{\mathcal{P}}$ -consistent sets X of \tilde{t} clusters of \mathcal{P}^* such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$ in the form $f''(\mathbf{x})$, where \mathbf{x} represents the choice of j -cell on Y for each $Y \in \binom{\mathcal{P}}{j}$ and each $2 \leq j \leq k-1$. As before, we let c_Y be the maximum change in $f''(\mathbf{x})$ which can be obtained by changing the choice of cell on Y . Observe that if Y contains two or more clusters from the same part of $\tilde{\mathcal{P}}$ then we have $c_Y = 0$, while otherwise changing the cell on Y can only affect those \mathcal{P}^* -consistent sets X such that $Y \subseteq X$, of which there are exactly $p(t_1^*)^{\tilde{t}-|Y|}$. So we have

$$\sum_Y c_Y^2 \leq \sum_{j=2}^{k-1} \binom{\tilde{t}}{j} p(t_1^*)^j (p(t_1^*)^{\tilde{t}-j})^2 \leq 2^{\tilde{t}} p(t_1^*)^{2\tilde{t}-2}$$

and hence by Theorem 26, the probability that the number of sets X with $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$ deviates by more than

$$\frac{\varepsilon_k}{16} p(t_1^*)^{\tilde{t}} \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}}$$

from its expectation is at most

$$2 \exp\left(\frac{-\varepsilon_k^2 p(t_1^*)^{2\tilde{t}} \prod_{j=2}^{k-1} d_j^{2\binom{\tilde{t}}{j}}}{2 \cdot 16^2 \cdot 2^{\tilde{t}} p(t_1^*)^{2\tilde{t}-2}}\right) \leq 2 \exp\left(-2^{-t_1^*-9} \varepsilon_k^2 p(t_1^*)^2 \prod_{j=2}^{k-1} \left(\frac{1}{t_1^*}\right)^{2\binom{t_1^*}{j}}\right),$$

where the inequality follows from Lemma 22(c) and (d), which imply that $d_j \geq 1/t_1^*$ for each j and $\tilde{t} \leq t_1^*$. Since we chose $p(\cdot)$ to satisfy (5) (which we apply with t_1^* in place of x), and we have $t \leq t_1 = p(t_1^*)t_1^*$ by Lemma 22(c), this probability is smaller than

$$2^{-p(t_1^*)t_1^*} \prod_{j=2}^{k-1} \left(\frac{1}{t_1^*}\right)^{\binom{t_1^*}{j}} \leq 2^{-t} \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}}.$$

We can therefore take a union bound over the (by Claim 28) $\prod_{j=2}^{k-1} d_j^{-\binom{\tilde{t}}{j}}$ choices of $\tilde{\mathcal{J}}$ to obtain the desired conclusion. \square

We want to show that if the likely event of Claim 31 holds, then using Lemma 22(f) we can deduce property (c). So fix $G = G_i$ for some $i \in [s]$. Now for any fixed $1 \leq \ell \leq 1/\varepsilon_k$, any fixed k -graph H equipped with a set of distinct root vertices x_1, \dots, x_ℓ such that $h = v(H) \leq 1/\varepsilon_k$ and any fixed set of distinct vertices v_1, \dots, v_ℓ in V , we would like to show that the density $d_H(G_i; v_1, \dots, v_\ell)$ of rooted H -copies in G is within ε_k of the density $d_H(G_i; v_1, \dots, v_\ell, \mathcal{J})$ of rooted H -copies in G supported on \mathcal{J} . Recall that the vertex set of \mathcal{H}^{skel} consists of all vertices of H except for x_1, \dots, x_ℓ . Choose uniformly at random a map $\psi : V(\mathcal{H}^{skel}) \rightarrow V(G)$, and define $\psi' : V(H) \rightarrow V(G)$ by taking $\psi'(x_j) = v_j$ for any $j \in [\ell]$ and $\psi'(x) = \psi(x)$ for any $x \in V(\mathcal{H}^{skel})$ (note carefully that, unlike for the previous cases, we do not insist that ψ is injective, and that even if ψ is injective the same may not be true of ψ'). Let INJ denote the event that ψ' is injective, and let EMB denote the event that ψ' is an embedding of H into G (so EMB is a subset of INJ). Next, let CROSS $_{\mathcal{P}}$ denote the event that ψ is \mathcal{P} -crossing, meaning as before that each vertex of \mathcal{H}^{skel} is mapped to a distinct part of \mathcal{P} , and similarly let CROSS $_{\tilde{\mathcal{P}}}$ denote the event that ψ is $\tilde{\mathcal{P}}$ -crossing. Finally, for any $(k-1)$ -complex \mathcal{L} on V let SL $_{\mathcal{L}}$ denote the event that ψ is an embedding of \mathcal{H}^{skel} into \mathcal{L} . Then by definition we have

$$(18) \quad d_H(G_i; v_1, \dots, v_\ell) = \mathbb{P}(\text{EMB} \mid \text{INJ}),$$

and

$$(19) \quad d_H(G_i; v_1, \dots, v_\ell, \mathcal{J}) = \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}} \cap \text{CROSS}_{\mathcal{P}}),$$

and our aim is to show that these two conditional probabilities differ by at most ε_k . We shall frequently refer to events of the form $\text{SL}_{\mathcal{L}} \cap \text{CROSS}_{\tilde{\mathcal{P}}}$ for some $(k-1)$ -complex \mathcal{L} on V , so for brevity we denote this event by $\text{SL}_{\mathcal{L}}^*$.

Our first goal is to show that conditioning on $\text{CROSS}_{\tilde{\mathcal{P}}}$ instead of on INJ in (18) and on $\text{CROSS}_{\mathcal{P}}$ in (19) has an insignificant effect on the probabilities expressed in these equations, since these events are all highly probable (even after conditioning on $\text{SL}_{\mathcal{J}}$). Indeed, similarly as before we have

$$(20) \quad \mathbb{P}(\text{CROSS}_{\tilde{\mathcal{P}}}^c) \leq \frac{\binom{n}{h-\ell-1} h(n/\tilde{t})}{\binom{n}{h-\ell}} = \frac{h(h-\ell)n/\tilde{t}}{n-h+\ell+1} \leq \frac{2h^2}{t_0} \stackrel{(3)}{\leq} \frac{\varepsilon_k}{100}.$$

Now observe that ψ can be obtained by selecting a uniformly-random image $\psi(x)$ for each $x \in V(\mathcal{H}^{skel})$ in turn. The event INJ will occur unless for some x we select $\psi(x)$

to be one of the at most $h - \ell$ previously-chosen images or one of the vertices v_1, \dots, v_ℓ . Taking a union bound, it follows that for any $X \in \mathcal{X}$ we have

$$(21) \quad \mathbb{P}(\text{INJ}^c) \leq (h - \ell) \frac{h}{n} \leq \frac{h^2}{n} \leq \frac{\varepsilon_k}{100}.$$

By (20) we may apply Proposition 27 with EMB and $\text{CROSS}_{\tilde{\mathcal{P}}}$ in place of E and C respectively, and by (21) we may apply the same proposition with EMB and INJ in place of E and C respectively (in each case we take S to be the event which always occurs and $\frac{\varepsilon_k}{100}$ in place of ε). This gives our approximation for (18), namely

$$(22) \quad \mathbb{P}(\text{EMB} \mid \text{INJ}) = \mathbb{P}(\text{EMB}) \pm \frac{\varepsilon_k}{50} = \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\tilde{\mathcal{P}}}) \pm \frac{2\varepsilon_k}{50}.$$

We turn now to (19), beginning with the next claim. We define \mathcal{X} to be the set of all $\tilde{\mathcal{P}}$ -consistent sets $X \in \binom{\mathcal{P}}{\tilde{\mathcal{P}}}$. (Recall that this means that X is a set of \tilde{t} clusters in \mathcal{P} , one contained in each part of $\tilde{\mathcal{P}}$).

Claim 32. *Let \mathcal{L} be a slice through \mathcal{P}^* . Then*

- (a) $\mathbb{P}(\text{SL}_{\mathcal{L}} \mid \text{CROSS}_{\mathcal{P}}) = (1 \pm \frac{\varepsilon_k}{100}) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}$,
- (b) $\mathbb{P}(\text{SL}_{\mathcal{L}} \mid \text{CROSS}_{\tilde{\mathcal{P}}}) = (1 \pm \frac{\varepsilon_k}{100}) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}$, and
- (c) for any $X \in \mathcal{X}$ we have $\mathbb{P}(\text{SL}_{\mathcal{L}[X]}^* \mid \text{SL}_{\mathcal{L}}^*) = (1 \pm \frac{\varepsilon_k}{20}) p(t_1^*)^{\ell-h}$.

Furthermore, for any slice $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ we have $\mathbb{P}(\text{SL}_{\tilde{\mathcal{J}}} \mid \text{CROSS}_{\tilde{\mathcal{P}}}) = (1 \pm \frac{\varepsilon_k}{100}) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}$.

Proof. Identify the vertices of \mathcal{H}^{skel} with the integers of $[h - \ell]$. Then for any $(h - \ell)$ -tuple $\mathcal{C} = (C_1, \dots, C_{h-\ell})$ of distinct clusters of \mathcal{L} (i.e. parts of \mathcal{P}), we say that a copy of \mathcal{H}^{skel} in \mathcal{L} is \mathcal{C} -distributed if for each $j \in [h - \ell]$ vertex j of \mathcal{H}^{skel} lies in cluster C_j of \mathcal{L} . Likewise, we say that ψ is \mathcal{C} -distributed if $\psi(j) \in C_j$ for each $j \in [h - \ell]$, an event which we denote by $\text{DIST}_{\mathcal{C}}$. Note that if we condition on the event $\text{DIST}_{\mathcal{C}}$, then for each $j \in [h - \ell]$ the image $\psi(j)$ is a uniformly-random vertex in C_j . It follows that $\mathbb{P}(\text{SL}_{\mathcal{L}} \mid \text{DIST}_{\mathcal{C}})$ is equal to the number of \mathcal{C} -distributed labelled copies of \mathcal{H}^{skel} in \mathcal{L} divided by $m^{h-\ell}$ (the number of possibilities for ψ given that $\text{DIST}_{\mathcal{C}}$ occurs). So by Lemma 23 we obtain the estimate

$$(23) \quad \mathbb{P}(\text{SL}_{\mathcal{L}} \mid \text{DIST}_{\mathcal{C}}) = (1 \pm \frac{\varepsilon_k}{100}) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}.$$

Since the events $\text{DIST}_{\mathcal{C}}$ partition the event $\text{CROSS}_{\mathcal{P}}$, this implies (a). Similarly we obtain (b) since the event $\text{CROSS}_{\tilde{\mathcal{P}}}$ is partitioned by the events $\text{DIST}_{\mathcal{C}}$ for those $(h - \ell)$ -tuples \mathcal{C} for which each C_j is a subset of a distinct part of $\tilde{\mathcal{P}}$. Now, fix $X \in \mathcal{X}$, and let DIST_X be the event that ψ maps each vertex of \mathcal{H}^{skel} to a distinct member of X (so $\text{DIST}_X \subseteq \text{CROSS}_{\tilde{\mathcal{P}}}$). Then the event DIST_X is partitioned by the events $\text{DIST}_{\mathcal{C}}$ for $(h - \ell)$ -tuples \mathcal{C} of distinct sets in X , each of which is equally likely to occur. Summing (23) over all such k -tuples \mathcal{C} , we obtain

$$\mathbb{P}(\text{SL}_{\mathcal{L}[X]}^* \mid \text{DIST}_X) = \mathbb{P}(\text{SL}_{\mathcal{L}}^* \mid \text{DIST}_X) = (1 \pm \frac{\varepsilon_k}{100}) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}.$$

Together with the fact that $\mathbb{P}(\text{DIST}_X \mid \text{CROSS}_{\tilde{\mathcal{P}}}) = p(t_1^*)^{-(h-\ell)}$ (since ψ was chosen uniformly at random), this gives

$$\mathbb{P}(\text{SL}_{\mathcal{L}[X]}^* \mid \text{CROSS}_{\tilde{\mathcal{P}}}) = (1 \pm \frac{\varepsilon_k}{100}) p(t_1^*)^{-(h-\ell)} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}.$$

Since $\text{SL}_{\mathcal{L}[X]}^* \subseteq \text{SL}_{\mathcal{L}}^*$, this equation and part (a) together prove (c). Finally, the proof of the final statement is almost identical to the proof of (a); the only differences are that we instead consider k -tuples \mathcal{C} of parts of $\tilde{\mathcal{P}}$, and that the term $m^{h-\ell}$ is replaced by $(n/\tilde{t})^{h-\ell}$ (but this term is, as above, cancelled by the output from Lemma 23). \square

Note that $\mathbb{P}(\text{CROSS}_{\mathcal{P}} \cap \text{SL}_{\mathcal{J}}) = \mathbb{P}(\text{SL}_{\mathcal{J}} \mid \text{CROSS}_{\mathcal{P}})\mathbb{P}(\text{CROSS}_{\mathcal{P}}) \leq \mathbb{P}(\text{SL}_{\mathcal{J}} \mid \text{CROSS}_{\mathcal{P}})$, and Claim 32 part (a) gives an approximation for this probability. Similarly we have $\mathbb{P}(\text{CROSS}_{\tilde{\mathcal{P}}} \cap \text{SL}_{\mathcal{J}}) = \mathbb{P}(\text{SL}_{\mathcal{J}} \mid \text{CROSS}_{\tilde{\mathcal{P}}})\mathbb{P}(\text{CROSS}_{\tilde{\mathcal{P}}})$, which we can approximate by Claim 32 part (b) and (20). So, since $\text{CROSS}_{\tilde{\mathcal{P}}} \subseteq \text{CROSS}_{\mathcal{P}}$, we obtain

$$\begin{aligned} \mathbb{P}(\text{CROSS}_{\tilde{\mathcal{P}}} \mid \text{CROSS}_{\mathcal{P}} \cap \text{SL}_{\mathcal{J}}) &= \frac{\mathbb{P}(\text{CROSS}_{\tilde{\mathcal{P}}} \cap \text{SL}_{\mathcal{J}})}{\mathbb{P}(\text{CROSS}_{\mathcal{P}} \cap \text{SL}_{\mathcal{J}})} \\ &\geq \frac{\left(1 - \frac{\varepsilon_k}{100}\right) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})} \left(1 - \frac{\varepsilon_k}{100}\right)}{\left(1 + \frac{\varepsilon_k}{100}\right) \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}} \\ &\geq 1 - \frac{\varepsilon_k}{20}. \end{aligned}$$

We may therefore apply Proposition 27 with EMB, $\text{CROSS}_{\mathcal{P}} \cap \text{SL}_{\mathcal{J}}$ and $\text{CROSS}_{\tilde{\mathcal{P}}}$ in place of E , S and C respectively to obtain

$$\mathbb{P}(\text{EMB} \mid \text{CROSS}_{\mathcal{P}} \cap \text{SL}_{\mathcal{J}}) = \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\tilde{\mathcal{P}}} \cap \text{SL}_{\mathcal{J}}) \pm \frac{\varepsilon_k}{10}.$$

Combining this equation with (18), (19), (22) and the definition of $\text{SL}_{\mathcal{J}}^*$ as $\text{SL}_{\mathcal{J}} \cap \text{CROSS}_{\tilde{\mathcal{P}}}$, we conclude that it is sufficient to prove that

$$(24) \quad \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\tilde{\mathcal{P}}}) = \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}}^*) \pm \frac{\varepsilon_k}{2}.$$

We begin with the left hand side of (24). Recall that there are precisely $\prod_{j=2}^{k-1} d_j^{-\binom{i}{j}}$ slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$, and observe that if the event $\text{CROSS}_{\tilde{\mathcal{P}}}$ occurs then by Claim 28 there are precisely $\prod_{j=2}^{k-1} d_j^{e_j(\mathcal{H}^{skel}) - \binom{i}{j}}$ slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ for which the event $\text{SL}_{\tilde{\mathcal{J}}}^*$ occurs. By definition we have $\text{SL}_{\tilde{\mathcal{J}}}^* \subseteq \text{CROSS}_{\tilde{\mathcal{P}}}$ for any such slice, so summing over all slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ we obtain

$$\begin{aligned} (25) \quad \mathbb{P}(\text{EMB} \mid \text{CROSS}_{\tilde{\mathcal{P}}}) &= \prod_{j=2}^{k-1} d_j^{\binom{i}{j} - e_j(\mathcal{H}^{skel})} \sum_{\tilde{\mathcal{J}}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \mathbb{P}(\text{SL}_{\tilde{\mathcal{J}}}^* \mid \text{CROSS}_{\tilde{\mathcal{P}}}) \\ &= \left(1 \pm \frac{\varepsilon_k}{100}\right) \prod_{j=2}^{k-1} d_j^{\binom{i}{j}} \sum_{\tilde{\mathcal{J}}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*), \end{aligned}$$

where the final equality holds by the final part of Claim 32.

We next show that for any slice $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ and any $X \in \mathcal{X}$ the probabilities $\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*)$ and $\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}[X]}^*)$ are roughly equal. To do this, fix any $X \in \mathcal{X}$, and for any set \tilde{Y} of $h - \ell$ parts of $\tilde{\mathcal{P}}$, let Y be the (unique) \tilde{Y} -consistent subset of X . Then for any choice of \tilde{Y} , Lemma 22(f) states that we have

$$\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}[\tilde{Y}]}^*) = \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}[Y]}^*) \pm \varepsilon_k^*.$$

Since the events $\text{SL}_{\tilde{\mathcal{J}}[\tilde{Y}]}^*$ for $\tilde{Y} \in \binom{\tilde{\mathcal{P}}}{h-\ell}$ partition the event $\text{SL}_{\tilde{\mathcal{J}}}^*$, the events $\text{SL}_{\tilde{\mathcal{J}}[Y]}^*$ partition the event $\text{SL}_{\tilde{\mathcal{J}}[X]}^*$, and $\varepsilon_k^* \leq \frac{\varepsilon_k}{100}$ by (4), we may sum over all $\tilde{Y} \in \binom{\tilde{\mathcal{P}}}{h-\ell}$ to obtain

$$(26) \quad \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) = \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}[X]}^*) \pm \frac{\varepsilon_k}{100}.$$

Turning to the right hand side of (24), we next show that $\mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}}^*)$ is close to the average of $\mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}[X]}^*)$ over all of the $p(t_1^*)^{\tilde{t}}$ sets $X \in \mathcal{X}$. Observe that the events $\text{SL}_{\mathcal{J}[X]}^*$ for $X \in \mathcal{X}$ cover $\text{SL}_{\mathcal{J}}^*$, with each $\psi \in \text{SL}_{\mathcal{J}}^*$ covered exactly $p(t_1^*)^{\tilde{t}-h+\ell}$ times. We thus obtain

$$(27) \quad \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}}^*) = p(t_1^*)^{-\tilde{t}+h-\ell} \sum_{X \in \mathcal{X}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}[X]}^*) \mathbb{P}(\text{SL}_{\mathcal{J}[X]}^* \mid \text{SL}_{\mathcal{J}}^*) \\ = \left(1 \pm \frac{\varepsilon_k}{20}\right) \left(p(t_1^*)^{-\tilde{t}} \sum_{X \in \mathcal{X}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}[X]}^*)\right).$$

where the second equality holds by Claim 32 part (c).

Finally, we want to connect the averages in (25) and (27). Observe first that for any $X \in \mathcal{X}$ there is precisely one slice $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$, and for this $\tilde{\mathcal{J}}$ we have $\mathcal{J}[X] = \tilde{\mathcal{J}}[X]$, which implies

$$(28) \quad \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}[X]}^*) = \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}[X]}^*) \stackrel{(26)}{=} \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \pm \frac{\varepsilon_k}{100}$$

Since the last term does not depend on X , we only need to know, for each $\tilde{\mathcal{J}}$, how many sets $X \in \mathcal{X}$ satisfy $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}$. If \mathcal{J} satisfies the good event of Claim 31, then the answer is given by (17), and so by summing (28) over all $X \in \mathcal{X}$ we obtain

$$\sum_{X \in \mathcal{X}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}[X]}^*) = \sum_{\tilde{\mathcal{J}}} \sum_{X \in \mathcal{X}: \mathcal{J}[X] \subseteq \tilde{\mathcal{J}}} \left(\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \pm \frac{\varepsilon_k}{100} \right) \\ \stackrel{(17)}{=} \sum_{\tilde{\mathcal{J}}} \left(1 \pm \frac{\varepsilon_k}{16}\right) p(t_1^*)^{\tilde{t}} \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}} \left(\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \pm \frac{\varepsilon_k}{100} \right),$$

where the sum is over all slices $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$. Together with (27), this implies that

$$\mathbb{P}(\text{EMB} \mid \text{SL}_{\mathcal{J}}^*) \pm \frac{\varepsilon_k}{20} = \sum_{\tilde{\mathcal{J}}} \left(1 \pm \frac{\varepsilon_k}{16}\right) \prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}} \left(\mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \pm \frac{\varepsilon_k}{100} \right) \\ = \left(\prod_{j=2}^{k-1} d_j^{\binom{\tilde{t}}{j}} \sum_{\tilde{\mathcal{J}}} \mathbb{P}(\text{EMB} \mid \text{SL}_{\tilde{\mathcal{J}}}^*) \right) \pm \frac{\varepsilon_k}{8},$$

which together with (25) implies (24). This completes the proof that if \mathcal{J} has the good event of Claim 31 then it satisfies property (c) of Lemma 6.

Finally, we conclude that the probability that \mathcal{J} satisfies all the conclusions of Lemma 6 is at least $1/2 - s2^{-t} - t^{k-1}2^{-\sqrt{\tilde{t}}} - 2^{-t}$. Since $t \geq t_0^*$, by choice of t_0^* in (3) this probability is strictly greater than zero, and so some \mathcal{J} exists as required. \square

9. EMBEDDING TIGHT CYCLES

Our aim in this section is to prove Lemma 9. We start with some definitions.

Recall the definition of a tight walk W from Section 4.3. As we did for tight paths, we define the length $\ell(W)$ of W to be the number of edges in W (i.e. the number of consecutive k -sets of vertices). We refer to the first s vertices of W , ordered as they appear in W , as the *initial s -tuple* of W , and similarly to the final s vertices as the *terminal s -tuple* of W . Initial and terminal $(k-1)$ -tuples have particular importance, as given tight walks W and W' for which the terminal $(k-1)$ -tuple of W is identical to the initial $(k-1)$ -tuple of W' , we may *concatenate* W and W' to form a new tight walk,

which we denote $W + W'$; when doing so, we always include the common $(k - 1)$ -tuple only once, so the edges of $W + W'$ are precisely the edges of W followed by the edges of W' . In particular, we have $\ell(W + W') = \ell(W) + \ell(W')$. Note that this definition includes the case where W and W' are tight paths, in which case $W + W'$ is also a tight path provided that W and W' have no common vertices outside this common $(k - 1)$ -tuple.

Before giving the full proof of Lemma 9, we briefly sketch our approach. We show that for any k -tuple $X = (X_1, \dots, X_k)$ of clusters of \mathcal{J} which forms an edge of $R := R_{d_k}(G)$, we can find a long tight path ‘winding around’ the k clusters in $G[\bigcup_{i \in [k]} X_i]$ by repeated use of the Extension Lemma. Indeed, we proceed in steps, at each time $j \geq 1$ keeping track of a tight path $P^{(j)}$ and a set $\mathcal{P}^{(j)}$ of ‘possible extensions’ of $P^{(j)}$: the latter are short tight paths whose initial $(k - 1)$ -tuples are equal to the terminal $(k - 1)$ -tuple of $P^{(j)}$, and whose terminal $(k - 1)$ -tuples are all distinct. The Extension Lemma tells us that at each time step j there must exist some $P \in \mathcal{P}^{(j)}$ so that, taking $P^{(j+1)}$ to be $P^{(j)} + P$, there is a family $\mathcal{P}^{(j+1)}$ which are possible extensions of $P^{(j+1)}$ (this argument is formalised in Claim 33). This procedure, which we refer to as ‘filling the edge X ’, can be continued until only few vertices remain in each cluster of X .

Of course, for large ℓ we need to cover more vertices than are contained in any k clusters; this is the point at which the connectedness of the fractional matching comes into play. By assumption we can find a tight walk W in R which visits every edge $e \in R$ whose weight w_e is non-zero. Starting at some such edge e , we proceed to ‘fill’ this edge as described above, stopping when our tight path $P^{(j)}$ covers around a w_e -proportion of the vertices of each cluster. We then extend $P^{(j)}$ by ‘traversing the walk W ’ to the next edge e' of non-zero weight. Since this extension of $P^{(j)}$ is short, few vertices are used in total in walk-traversing steps. So the final proportion of vertices covered by our tight path in any cluster X_i is approximately the sum of the weights of edges of R containing X_i . Overall, this gives a path covering sufficiently many vertices for the bound given in the lemma. To obtain a shorter tight path, we simply stop ‘filling’ each edge at an earlier stage.

Finally, it remains to ‘join the ends’ of our tight path to form a tight cycle. For this, at the very start of the argument we set aside some large subsets Z_1, \dots, Z_k of some clusters X_1, \dots, X_k which form an edge of R . We then choose our first path $P^{(0)}$ so that, as well as there being many extensions $\mathcal{P}^{(0)}$ suitable for use to form $P^{(1)}$, there are also many $(k - 1)$ -tuples f in Z_1, \dots, Z_{k-1} which can be extended to the initial $(k - 1)$ -tuple of $P^{(0)}$.

We now present the full details of the proof.

Proof of Lemma 9. We set

$$(29) \quad \alpha = \frac{\psi}{5} \quad \text{and} \quad \beta = \frac{1}{200}.$$

We will want to apply Lemma 24 to k -partite k -complexes (i.e. with $s = k$) with α as given in (29) and with ε_k playing the same role here as there. Also, we will want to apply Lemma 25 to both k -partite and $(2k - 1)$ -partite k -complexes (i.e. with $s = k$ and with $s = 2k - 1$) whose top levels have relative density at least d_k , with each choice of $1 \leq b' < b \leq 3k$, with β as given in (29), and with $\sqrt{\varepsilon_k}$ in place of ε_k . We require $\varepsilon_k > 0$ to be small enough for each of these applications.

Given d_2, \dots, d_{k-1} , the various applications of Lemma 25 mentioned above require various sufficiently small positive values for c . We take $c > 0$ to be the minimum of these values (we will have a bounded number of choices of parameters, hence this minimum is well-defined).

We now require $\varepsilon < c^2$ to be small enough for each of the applications of Lemmas 24 and 25 mentioned above to k -complexes whose underlying $(k-1)$ -complex is $(\cdot, \cdot, \sqrt{\varepsilon})$ -equitable with density vector $\mathbf{d} = (d_{k-1}, \dots, d_2)$. In addition, we require $\varepsilon \leq \frac{1}{2} \prod_{i=2}^{k-1} d_i^{\binom{k-1}{i}}$ to be small enough that we can apply Fact 3 with $\frac{1}{2}$ in place of β to $(\cdot, \cdot, \sqrt{\varepsilon})$ -equitable k -partite k -complexes with density vector \mathbf{d} . We require r to be large enough for the above mentioned applications of Lemma 25. We also choose $m_0 \geq 16(k-1)/\varepsilon$ to be large enough so that any $m \geq \alpha m_0$ is acceptable for all of these applications.

Given t , we set

$$(30) \quad n_0 = t \cdot \max \left(m_0, \frac{200k}{\varepsilon}, \frac{8k}{\alpha\sqrt{\varepsilon}}, \frac{2(k+1)t^{2k}}{\alpha} \right).$$

Now let G be an n -vertex k -graph, where $n \geq n_0$, and let \mathcal{J} be a $(\cdot, \cdot, \varepsilon, \varepsilon_k, r)$ -regular slice for G with t clusters and density vector \mathbf{d} . Let $R := R_{d_k}(G)$, and let $m := n/t$, so each cluster of \mathcal{J} has size m . We write \mathcal{G} for the k -complex obtained from \mathcal{J} by adding all edges of G supported on $\mathcal{J}^{(k-1)}$ as the ‘ k -th level’ of \mathcal{G} . So for any edge $X \in R$, $\mathcal{G}[\bigcup_{X' \in X} X']$ is a $(d^*(X), d_{k-1}, \dots, d_2, \varepsilon, \varepsilon_k, r)$ -regular k -partite k -complex with $d^*(X) \geq d_k$. Furthermore, for convenience of notation, for any s -tuple $X = (X_1, \dots, X_s)$ of clusters of \mathcal{J} and any subsets $Y_j \subseteq X_j$ for $j \in [s]$ we write $\mathcal{G}(Y_1, \dots, Y_s)$ for the s -partite s -graph $\mathcal{G}_X[\bigcup_{j \in [s]} Y_j]$, that is, whose edges are the edges of $\mathcal{G}^{(s)}$ with one vertex in each Y_j . In addition, we say that an s -tuple (v_1, \dots, v_s) of vertices of G is an *ordered edge of $\mathcal{G}(Y_1, \dots, Y_s)$* if $\{v_1, \dots, v_s\}$ is an edge of \mathcal{G} and $v_j \in Y_j$ for each $j \in [s]$.

Since \mathcal{J} is a regular slice for G , for any set X of k clusters X_1, \dots, X_k of \mathcal{J} the k -partite $(k-1)$ -complex $\mathcal{J}[\bigcup_{j \in [k]} X_j]$ is $(\mathbf{d}, \varepsilon, \varepsilon, \mathbf{1})$ -regular. By adding all sets of k vertices supported on the polyad $\hat{\mathcal{J}}_X$ as a ‘ k -th level’, we may obtain a $(1, d_{k-1}, \dots, d_2, \varepsilon, \varepsilon_k, r)$ -regular k -partite k -complex, whose restriction to any subsets $Y_j \subseteq X_j$ of size $|Y_j| = \alpha m$ for each $j \in [k]$ is then $(1, d_{k-1}, \dots, d_2, \sqrt{\varepsilon}, \sqrt{\varepsilon_k}, r)$ by Lemma 24. We conclude by Fact 3 that for any subsets Y_1, \dots, Y_{k-1} of distinct clusters of \mathcal{J} , each of size αm , we have

$$(31) \quad e(\mathcal{G}(Y_1, \dots, Y_{k-1})) \geq \varepsilon m^{k-1}.$$

The heart of our embedding lemma is the following claim. We will use it in steps when we fill an edge of R with $i = 0$, and in walk-traversing steps with $i = 1$.

Claim 33. *Let $\{X_1, \dots, X_k\}$ be an edge of R , and choose any $Y_j \subseteq X_j$ for each $j \in [k]$ so that $|Y_1| = \dots = |Y_k| = \alpha m$. Let \mathcal{P} be a collection of at least $\frac{1}{2} e(\mathcal{G}(Y_1, \dots, Y_{k-1}))$ tight paths in G (not necessarily contained in $\bigcup_{j \in [k]} Y_j$) each of length at most $2k+1$ and whose terminal $(k-1)$ -tuples are distinct members of $\mathcal{G}(Y_1, \dots, Y_{k-1})$. Then for each $i \in \{0, 1\}$ there is a path $P \in \mathcal{P}$ and a collection \mathcal{P}' of $\frac{9}{10} e(\mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1}))$ tight paths in G , each of length $k+i$, all of whose initial $(k-1)$ -tuples are the terminal $(k-1)$ -tuple of P , whose terminal $(k-1)$ -tuples are distinct members of $\mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1})$, and where the j th vertex of each path in \mathcal{P}' lies in Y_j and, if $j \geq k$, is not contained in P .*

One might expect to prove this claim by trying to extend the terminal $(k-1)$ -tuple of a path in \mathcal{P} in many different ways using Lemma 25. But it turns out to be hard to show that these many different ways really go to many different terminal $(k-1)$ -tuples, and so what we actually do is show the stronger statement that most pairs of disjoint $(k-1)$ -tuples are joined by many paths. That is, we apply Lemma 25 with \mathcal{H} the k -complex formed by a tight path and \mathcal{H}' the subcomplex induced by its initial and terminal $(k-1)$ -tuples.

Proof. We take \mathcal{H} to be the k -complex generated by the down-closure of a tight path of length $k + i$ (so \mathcal{H} has $2k - 1 + i$ vertices), and consider its k -partition in which the i th vertex of the path lies in the vertex class V_j with $j = i$ modulo k . We take \mathcal{H}' to be the subcomplex of \mathcal{H} (with empty k th level) induced by the first and last $k - 1$ vertices of \mathcal{H} . By our assumptions on our various constants and by Lemma 24, $\mathcal{G}[\bigcup_{j \in [k]} Y_j]$ satisfies the conditions to apply Lemma 25. Consider the pairs (e, f) , where e is an ordered $(k - 1)$ -edge of $\mathcal{G}(Y_1, \dots, Y_{k-1})$ and f is an ordered $(k - 1)$ -edge of $\mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1})$. For any such ordered edge e there are at most km^{k-2} such ordered edges f which intersect e , so by (30) and (31), at most a $1/200$ -proportion of the pairs (e, f) are not disjoint. On the other hand, if e and f are disjoint, then (the down-closure of) the pair (e, f) forms a labelled copy of \mathcal{H}' in $\mathcal{G}[\bigcup_{j \in [k]} Y_j]$, so by Lemma 25 with $b = 2k - 1 + i$ and $b' = 2k - 2$, for all but at most a $1/200$ -proportion of the disjoint pairs (e, f) there are at least $c(\alpha m)^{i+1} \geq \sqrt{\varepsilon}(\alpha m)^{i+1}$ extensions to copies of \mathcal{H} in $\mathcal{G}[\bigcup_{j \in [k]} Y_j]$. Each such copy of \mathcal{H} corresponds to a tight path in G of length $k + i$ with all vertices in the desired clusters. We conclude that at least a $99/100$ -proportion of all pairs (e, f) of ordered edges are disjoint and are linked by at least this many tight paths in G of the desired type; we call these pairs *extensible*.

Let us call an ordered edge $e \in \mathcal{G}(Y_1, \dots, Y_{k-1})$ *good* if at most one-twentieth of the ordered edges $f \in \mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1})$ do not make an extensible pair with e . Then at most one-fifth of the ordered edges in $\mathcal{G}(Y_1, \dots, Y_{k-1})$ are not good. In particular, there must exist a path $P \in \mathcal{P}$ whose terminal $(k - 1)$ -tuple is a good ordered edge e . Fix such a P and e . Given any ordered edge f in $\mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1})$ which is disjoint from P , suppose the pair (e, f) is an extensible pair. By definition there are at least $\sqrt{\varepsilon}(\alpha m)^{i+1}$ tight paths in G from e to f where the j th vertex of each path lies in Y_j . We claim that at least one of these paths has the further property that if $j \geq k$, then the j th vertex is not contained in P (and we can therefore put this path in \mathcal{P}'). Indeed, as f is disjoint from P , if $i = 0$ it suffices to show that one of these paths has the property that its k th vertex is in $Y_k \setminus V(P)$. This is true because there are only $V(P) \leq 2k + 1 < \sqrt{\varepsilon}(\alpha m)$ (where the last inequality is by (30)) paths which do not have this property. If on the other hand $i = 1$, then we need a path whose k th and $(k + 1)$ st vertices are not in $V(P)$, which is possible since $2(2k + 1)(\alpha m) < \sqrt{\varepsilon}(\alpha m)^2$ by (30).

Finally, consider the ordered edges $f \in \mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1})$. Since by (30) and (31) we have $20|P|(\alpha m)^{k-2} \leq \varepsilon m^{k-1} \leq e(\mathcal{G}(Y_{i+1}, \dots, Y_{i+k-1}))$, at most one-twentieth of these edges f intersect P , and by choice of e at most one-twentieth of these edges f are such that (e, f) is not extensible. This leaves at least nine-tenths of the edges f remaining; choosing a tight path for each such f as described above gives the desired set \mathcal{P}' . \square

Now, let e_1, \dots, e_s be the edges of non-zero weight in our fractional matching in R , and let w_1, \dots, w_s be the corresponding weights. For each $i \in [s]$ let n_i be any integer with

$$(32) \quad 0 \leq n_i \leq (1 - 3\alpha)w_i m.$$

We next construct tight walks W_i in R for ‘traversing’ from e_i to e_{i+1} for each $1 \leq i \leq s-1$. Since our fractional matching is tightly connected, we may choose a minimum length tight walk W_1 from e_1 to e_2 . Then for each $2 \leq i \leq s-1$, we may take a minimum length tight walk W_i from e_i to e_{i+1} whose initial $(k - 1)$ -tuple is the terminal $(k - 1)$ -tuple of W_{i-1} .

Finally, we construct a tight walk W_s from e_s to e_1 whose initial $(k - 1)$ -tuple is the terminal $(k - 1)$ -tuple of W_{s-1} and whose terminal $(k - 1)$ -tuple is the initial $(k - 1)$ -tuple of W_1 . We do this construction differently in order to ensure that our final cycles have length divisible by k . Indeed, let W' be the concatenation $W_1 + \dots + W_{s-1}$, so $W' = (A_1, \dots, A_{\ell'})$ is a tight walk in R from e_1 to e_s . Now we let W_s be the sequence of clusters given by

writing down the terminal $(k-1)$ -tuple of W' , that is, $A_{\ell'-k+2}, A_{\ell'-k+3}, \dots, A_{\ell'}$, followed by the penultimate $(k-1)$ -tuple $A_{\ell'-k+1}, A_{\ell'-k+2}, \dots, A_{\ell'-1}$, and so on until we eventually write the initial $(k-1)$ -tuple. Now any k consecutive vertices of W_s come from two $(k-1)$ -tuples in W' which share $k-2$ vertices. Since W' is a tight walk in R , the vertices of the two $(k-1)$ -tuples make an edge of R . In other words, W_s is a tight walk in R . Note that by construction the terminal $(k-1)$ -tuple of W_s is the initial $(k-1)$ -tuple of W_1 , and also that we have

$$\ell(W_s) = (k-1)\ell(W') = (k-1) \sum_{i=1}^{s-1} \ell(W_i).$$

Note that any given $(k-1)$ -tuple can appear at most once (consecutively) in a minimum length tight walk between two edges of R , or else we could contract the walk. It follows that each of the walks W_i has length $\ell(W_i) \leq t^{k-1}$. Since R has at most $\binom{t}{k}$ edges we conclude that $\sum_{i=1}^s \ell(W_i) = k \sum_{i=1}^{s-1} \ell(W_i)$ is a multiple of k and is at most t^{2k} .

Let (X_1, \dots, X_k) be the initial k -tuple of W_1 (so X_1, \dots, X_k are the clusters of e_1 in the order in which they appear in W_1). Given any subsets $Y_j \subseteq X_j$ of size $|Y_j| = \alpha m$ for $j \in [k]$, we say that an edge $e \in \mathcal{G}(X_1, \dots, X_{k-1})$ is *well-connected to (Y_1, \dots, Y_{k-1}) via Y_k* if for at least nine-tenths of the $(k-1)$ -tuples f in $\mathcal{G}(Y_1, \dots, Y_{k-1})$ there exist distinct vertices $u, v \in Y_k$ such that the concatenations $e + (u) + f$ and $e + (v) + f$ are tight paths in G of length k . Now fix a subset $Z_j \subseteq X_j$ of size αm for each $j \in [k]$, and write $Z = \bigcup_{j \in [k]} Z_j$. We reserve the vertices of Z for joining together the ends of the tight path we will construct to obtain a cycle; the following claim establishes the properties we will need to do this.

Claim 34. *For any subsets $Y_j \subseteq X_j$ of size αm for $j \in [k]$ such that each Y_j is disjoint from Z_j the following statements hold.*

- (a) *At least nine-tenths of the $(k-1)$ -tuples e in $\mathcal{G}(Z_1, \dots, Z_{k-1})$ are well-connected to (Z_1, \dots, Z_{k-1}) via Z_k .*
- (b) *At least nine-tenths of the $(k-1)$ -tuples e in $\mathcal{G}(Z_1, \dots, Z_{k-1})$ are well-connected to (Y_1, \dots, Y_{k-1}) via Y_k .*
- (c) *At least nine-tenths of the $(k-1)$ -tuples e in $\mathcal{G}(Y_1, \dots, Y_{k-1})$ are well-connected to (Z_1, \dots, Z_{k-1}) via Y_k .*

Proof. Observe that (a) was in fact proved in the proof of Claim 33 (in the case $i = 0$), with Y_1, \dots, Y_k there corresponding to Z_1, \dots, Z_k here. We here modify this argument to prove (b); a near-identical argument proves (c). We apply Lemma 25 as in Claim 33, with \mathcal{H} being the k -complex generated by the down-closure of a tight path of length k (that is, with $2k-1$ vertices), and \mathcal{H}' the subcomplex induced by its initial and terminal $(k-1)$ -tuples. However, we now regard \mathcal{H} as a $(2k-1)$ -partite k -complex, with one vertex in each vertex class. The role of \mathcal{G} in Lemma 25 is played by the $(2k-1)$ -partite subcomplex of \mathcal{G} with vertex classes $Z_1, \dots, Z_{k-1}, Y_k, Y_1, \dots, Y_{k-1}$; the first vertex of \mathcal{H} is to be embedded in Z_1 , the second in Z_2 , and so forth. So by Lemmas 24 and 25 the proportion of pairs (e, f) for which there is no path as in (b) is at most $1/200$, and the remainder of the argument then follows exactly as in Claim 33. \square

We are now ready to construct our cycle. Arbitrarily choose a subset $X_j^{(0)} \subseteq X_j$ of size αm which is disjoint from Z_j for each $j \in [k]$. By Claim 34(a) and (b) we may fix a $(k-1)$ -tuple $e \in \mathcal{G}(Z_1, \dots, Z_{k-1})$ such that e is both well-connected to (Z_1, \dots, Z_{k-1}) via Z_k and well-connected to $(X_1^{(0)}, \dots, X_{k-1}^{(0)})$ via $X_k^{(0)}$. Set $P^{(0)}$ to be the tight path with no k -edges consisting simply of the vertices of e in their given order. By choice of e there is

a set $\mathcal{P}^{(0)}$ of tight paths of the form $e + (v) + f$ for $v \in X_k^{(0)}$ and $f \in \mathcal{G}(X_1^{(0)}, \dots, X_{k-1}^{(0)})$ for which the terminal $(k-1)$ -tuples of paths in $\mathcal{P}^{(0)}$ are all distinct and constitute at least half of the ordered edges of $\mathcal{G}(X_1^{(0)}, \dots, X_{k-1}^{(0)})$. We now describe the algorithm we use to construct the desired cycle. Set the initial state to be ‘filling the edge e_1 ’. We proceed for each time $j \geq 1$ as follows, maintaining the following property (\dagger) .

(\dagger) The terminal $(k-1)$ -tuples of the paths $\mathcal{P}^{(j)}$ constitute at least half of the ordered edges $\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})$.

Suppose first that our current state is ‘filling the edge e_i ’ for some i . If we have previously completed n_i steps in this state, then we do nothing, and immediately change state to ‘position 1 in traversing the walk W_i ’. Otherwise, since (\dagger) holds for $j-1$, we may apply Claim 33 with $i=0$ to obtain a path $P \in \mathcal{P}^{(j-1)}$ and a collection $\mathcal{P}^{(j)}$ of $\frac{9}{10} e(\mathcal{G}(X_1^{(j-1)}, \dots, X_{k-1}^{(j-1)}))$ tight paths of length k , all of whose initial $(k-1)$ -tuples are the terminal $(k-1)$ -tuple of P , whose terminal $(k-1)$ -tuples are distinct members of $\mathcal{G}(X_1^{(j-1)}, \dots, X_{k-1}^{(j-1)})$ and are disjoint from $V(P)$, and whose remaining vertex lies in $X_k^{(j-1)} \setminus V(P)$. We define $P^{(j)}$ to be the concatenation $P^{(j-1)} + P$. For each $1 \leq p \leq k$ we generate $X_p^{(j)}$ from $X_p^{(j-1)}$ by removing the (at most two) vertices of $P^{(j)}$ in $X_p^{(j-1)}$ and replacing them by vertices from the same cluster which do not lie in Z or in $P^{(j)}$. We will prove in Claim 35 that this is possible and that (\dagger) is maintained.

Now suppose that our current state is ‘position q in traversing the walk W_i ’ for some i . Since (\dagger) holds for $j-1$, we may apply Claim 33 with $i=1$, which returns a path $P \in \mathcal{P}^{(j-1)}$ and a collection $\mathcal{P}^{(j)}$ of $\frac{9}{10} e(\mathcal{G}(X_2^{(j-1)}, \dots, X_k^{(j-1)}))$ tight paths of length $k+1$, all of whose initial $(k-1)$ -tuples are the terminal $(k-1)$ -tuple of P , whose terminal $(k-1)$ -tuples are distinct members of $\mathcal{G}(X_2^{(j-1)}, \dots, X_k^{(j-1)})$ which are disjoint from $V(P)$, and whose two remaining vertices lie in $X_k^{(j-1)} \setminus V(P)$ and $X_1^{(j-1)} \setminus V(P)$ respectively. Exactly as before we define $P^{(j)}$ to be the concatenation $P^{(j-1)} + P$. We also form $X_p^{(j)}$ from $X_{p+1}^{(j-1)}$ for each $1 \leq p \leq k-1$ by removing the vertices of $P^{(j-1)}$ in $X_{p+1}^{(j-1)}$ and replacing them by vertices from the same cluster which do not lie in Z or $P^{(j)}$. If we have now reached the end of W_i , meaning that the $(k-1)$ -tuple of clusters containing $X_1^{(j)}, \dots, X_{k-1}^{(j)}$ (in that order) is the terminal $(k-1)$ -tuple of W_i , then we choose the set $X_k^{(j)}$ as follows. If $i < s$ then we let $X_k^{(j)}$ be a subset of the remaining cluster of e_{i+1} (that is, the cluster not included in the terminal $(k-1)$ -tuple) which has size αm and is disjoint from $P^{(j)} \cup Z$. If $i = s$ we let $X_k^{(j)}$ be a subset of X_k (which is the remaining cluster of e_1) of size αm disjoint from $P^{(j)} \cup Z$. We then change our state to ‘filling the edge e_{i+1} ’ if $i < s$, or ‘completing the cycle’ if $i = s$. Alternatively, if we have not yet reached the end of W_i , we instead choose $X_k^{(j)}$ to be a subset of size αm of the cluster at position $q+k$ in the sequence W_i , again chosen so that $X_k^{(j)}$ does not intersect $P^{(j)}$ or Z . This ensures that for each $p \in [k]$, $X_p^{(j)}$ is a subset of the cluster at position $p+q$ in the sequence W_i ; in particular, these clusters form an edge of R since W_i is a tight walk. In this case we now change our state to ‘position $q+1$ in traversing W_i ’. Again, we prove in Claim 35 that these choices are all possible and that (\dagger) is maintained.

Finally, if our state is ‘completing the cycle’ then $X_1^{(j-1)}, \dots, X_k^{(j-1)}$ must be subsets of X_1, \dots, X_k respectively. So by (\dagger) and Claim 34(c) we may choose a path $P \in \mathcal{P}^{(j-1)}$ such that the terminal $(k-1)$ -tuple $f \in \mathcal{G}(X_1^{(j-1)}, \dots, X_{k-1}^{(j-1)})$ of P is well-connected to (Z_1, \dots, Z_{k-1}) via $X_k^{(j-1)}$. Let $P^{(j)}$ be the concatenation $P^{(j-1)} + P$, and recall that we

chose e , the initial $(k-1)$ -tuple of $P^{(j-1)}$, to be well-connected to (Z_1, \dots, Z_{k-1}) via Z_k . Together with the well-connectedness of f , and the fact that $V(P) \cap X_k = 1$, this implies that we may choose a $(k-1)$ -tuple e' in $\mathcal{G}(Z_1, \dots, Z_{k-1})$ and vertices $v \in X_k^{(j-1)} \setminus V(P)$ and $v' \in Z_k$ such that e' is disjoint from e and both $Q := f + (v) + e'$ and $Q' := e' + (v') + e$ are tight paths in G . Return $P^{(j)} + Q + Q'$ as the output tight cycle in G .

Claim 35. *The algorithm described above is well-defined (that is, it is always possible to construct the sets $X_p^{(j)}$), maintains (\dagger) , and returns a tight cycle of length*

$$\left(3 + \sum_{i \in [s]} n_i\right) \cdot k + \left(\sum_{i \in [s]} \ell(W_i)\right) \cdot (k+1).$$

In particular, this length is divisible by k .

Proof. To see that the output is indeed a tight cycle, recall that we always chose $X_1^{(j)}, \dots, X_k^{(j)}$ to be disjoint from Z and $P^{(j)}$, and the vertices added to $P^{(j)}$ to form $P^{(j+1)}$ are always taken from $\bigcup_{p \in [k]} X_p^{(j-1)}$. So by construction, the final $P^{(j)}$ is a tight path which only meets Z in its initial $(k-1)$ -tuple e and does not contain the vertex v used in the ‘completing the cycle’ step. Since when completing the cycle e' is chosen to be disjoint from e , no vertices are repeated, and so the output is indeed a tight cycle.

To see that (\dagger) is maintained, observe that applying (31) we have $e(\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})) \geq \varepsilon m^{k-1}$ for each j . Fix some j . By construction, for either $A_p := X_p^{(j-1)}$ or $A_p := X_{p+1}^{(j-1)}$ (according to our current state) we obtain sets A_1, \dots, A_{k-1} each of size αm such that the terminal $(k-1)$ -tuples of $\mathcal{P}^{(j)}$ constitute at least nine-tenths of the ordered edges of $\mathcal{G}(A_1, \dots, A_{k-1})$, and for each $1 \leq i \leq k-1$ the set $X_i^{(j)}$ is formed from A_i by removing at most two vertices and replacing them with the same number of vertices. Since each vertex is in at most m^{k-2} ordered $(k-1)$ -edges of either $\mathcal{G}(A_1, \dots, A_{k-1})$ or $\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})$, we conclude that the fraction of ordered $(k-1)$ -edges of $\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})$ which are terminal $(k-1)$ -tuples of paths in $\mathcal{P}^{(j)}$ is at least

$$\begin{aligned} & \frac{\frac{9}{10} e(\mathcal{G}(A_1, \dots, A_{k-1})) - 2(k-1)m^{k-2}}{e(\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})} \\ & \geq \frac{\frac{9}{10}(e(\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})) - 2(k-1)m^{k-2}) - 2(k-1)m^{k-2}}{e(\mathcal{G}(X_1^{(j)}, \dots, X_{k-1}^{(j)})} \\ & \geq \frac{9}{10} - \frac{4(k-1)m^{k-2}}{\varepsilon m^{k-1}} \geq \frac{1}{2}, \end{aligned}$$

where the final inequality is because we have $m \geq m_0 \geq 16(k-1)/\varepsilon$. Thus (\dagger) holds for j as desired.

To see that we can always construct the sets $X_p^{(j)}$, observe that it is enough to check that at termination every cluster still has at least $2\alpha m$ vertices not in $P^{(j)}$, as then there are at least αm such vertices outside Z . Observe that at each step we choose a set of paths $\mathcal{P}^{(j)}$, precisely one member of which is then used to extend $P^{(j)}$ to $P^{(j+1)}$ in the subsequent step. In each walk-traversing step each path in $\mathcal{P}^{(j)}$ contains precisely $k+1$ new vertices (i.e. vertices outside $P^{(j)}$), and the total number of walk-traversing steps is precisely $\sum_{i \in [s]} \ell(W_i)$. Recalling that this number is at most t^{2k} , and using $t \leq t_1$, by (30) we have $(k+1)t^{2k} < \alpha m/2$, so in particular fewer than $\alpha m/2$ vertices are added to $P^{(j)}$ as a result of walk-traversing steps. On the other hand, we remain in the state

‘filling the edge e_i ’ for precisely n_i steps, and in each of these steps each path in $\mathcal{P}^{(j)}$ contains precisely k new vertices, one from each cluster of e_i . So for any cluster C , the number of vertices of C which are added to $P^{(j)}$ as a result of edge-filling steps is $\sum_{i: C \in e_i} n_i \leq \sum_{i: C \in e_i} (1 - 3\alpha)w_i m \leq (1 - 3\alpha)m$, where the first inequality holds by (32) and the second by definition of a fractional matching. Together with the $k - 1$ vertices of $P^{(0)}$, and the k vertices of the chosen path in $\mathcal{P}^{(0)}$, we conclude that in total at most $(1 - 2\alpha)m$ vertices of any cluster are contained in the path $P^{(j)}$ at termination, as desired.

Finally, the total length of the cycle is equal to the number of vertices it contains, which we can calculate similarly. Recall that our initial $P^{(0)}$ contained $k - 1$ vertices, to which k vertices were added from some member of $\mathcal{P}^{(0)}$ in the first step to form P^1 . Each of the $\sum_{i \in [s]} n_i$ edge-filling steps resulted in k new vertices being added to $P^{(j)}$ in the subsequent step, and each of the $\sum_{i \in [s]} \ell(W_i)$ walk-traversing steps resulted in $k + 1$ new vertices being added to $P^{(j)}$ in the subsequent step. Finally, when completing the cycle we used $k + 1$ vertices which were not contained in our final path $P^{(j)}$ (namely v, v' and the vertices of e'). In summation, the cycle formed has length

$$(k - 1) + k + \left(\sum_{i \in [s]} n_i \right) \cdot k + \left(\sum_{i \in [s]} \ell(W_i) \right) \cdot (k + 1) + (k + 1),$$

giving the claimed expression. Since $\sum_{i \in [s]} \ell(W_i)$ is divisible by k , the same is true of this length. \square

Recall that $\sum_{i \in [s]} \ell(W_i) \leq t^{2k}$. So if we take $n_i = 0$ for every $i \in [s]$ then we obtain a tight cycle of length at most $2kt^{2k}$. On the other hand, if we take the n_i to be as large as permitted, so $n_i = (1 - 3\alpha)w_i m$ for each $i \in [s]$, then, since $\sum_{i=1}^s w_i = \mu$, we obtain a tight cycle of length at least $(1 - 3\alpha)\mu km \geq (1 - \psi)k\mu n/t$. Clearly by choosing the n_i appropriately we may obtain tight cycles of any length between these two extremes which is divisible by k . So it remains only to prove the lemma for cycle lengths $\ell < 2kt^{2k}$. Note that $2kt^{2k} < m$, so the number of vertices in such a cycle is fewer than the number of vertices in any cluster. Hence, we need only use one edge of the reduced graph to find the desired tight cycle in this case.

More precisely, fix any edge $e_1 = \{X_1, \dots, X_k\}$ of R . If $\ell \geq 3k$, then we may proceed as before with $n_1 := \ell/k - 3$. That is, we choose a $(k - 1)$ -tuple e , disjoint sets $Z_j, X_j^{(0)} \subseteq X_j$ of size αm , and a set of extensions $\mathcal{P}^{(0)}$ of $P^{(0)} = e$ as previously, then enter state ‘filling the edge e_1 ’, where we remain for n_1 steps. Following this, we move directly to the state ‘completing the cycle’ (since there is no walk to traverse), which proceeds as before. By similar arguments as before, this process gives a tight cycle of length $kn_1 + 3k = \ell$, as required.

Finally, if $\ell = 2k$, we simply apply Lemma 25 with \mathcal{H} the complex generated by the down-closure of a tight cycle of length $2k$, and \mathcal{H}' the subcomplex induced by any $k - 1$ consecutive vertices of this cycle, to $\mathcal{G}(X_1, \dots, X_k)$, where X_1, \dots, X_k are the clusters of any edge of R . This gives a cycle of length $2k$ in G . \square

To conclude this section, we note that by a similar approach it is possible to find paths of specified lengths whose initial and terminal $(k - 1)$ -tuples lie in specified sets of $(k - 1)$ -tuples. We will not need this result here, but we state it for future convenience.

Lemma 36. *Let k, r, n_0, t, B be positive integers, and $\psi, d_2, \dots, d_k, \varepsilon, \varepsilon_k, \nu$ be positive constants such that $1/d_i \in \mathbb{N}$ for each $2 \leq i \leq k-1$, and such that $1/n_0 \ll 1/t$,*

$$\frac{1}{n_0}, \frac{1}{B} \ll \frac{1}{r}, \varepsilon \ll \varepsilon_k, d_2, \dots, d_{k-1} \quad \text{and} \quad \varepsilon_k \ll \psi, d_k, \nu, \frac{1}{k}.$$

Then the following holds for all integers $n \geq n_0$.

Let G be a k -graph on n vertices, and \mathcal{J} be a $(\cdot, \cdot, \varepsilon, \varepsilon_k, r)$ -regular slice for G with t clusters and density vector (d_{k-1}, \dots, d_2) . Let M be a tightly connected fractional matching of weight μ in $R := R_{d_k}(G)$. Also let X and Y be $(k-1)$ -tuples of clusters which lie in the same tight component of R as M , and let S_X and S_Y be subsets of \mathcal{J}_X and \mathcal{J}_Y of sizes at least $\nu|\mathcal{J}_X|$ and $\nu|\mathcal{J}_Y|$ respectively. Finally, let W be a tight walk in R from X to Y of length at most t^{2k} , and let $p = \ell(W)$.

Then for any ℓ divisible by k with $3k \leq \ell \leq (1-\psi)k\mu n/t$ there is a tight path P in G of length $\ell + p(k+1)$ whose initial $(k-1)$ -tuple forms an edge of S_X and whose terminal $(k-1)$ -tuple forms an edge of S_Y , where the orders of these tuples are given by the orders of X and Y respectively. Furthermore P uses at most $\mu(C)n/t + B$ vertices from any cluster C , where $\mu(C)$ denotes the total weight of edges of M containing C .

Proof. We need to be able to apply Lemma 25 with $\beta = \min(\nu/40, 1/100)$ rather than $1/100$, and we need $\alpha \ll \nu$, but the remaining constants only have to satisfy the given order of magnitude hierarchy. We use the same notation as in the proof of Lemma 9. So, writing $X = (X_1, \dots, X_{k-1})$ and $Y = (Y_1, \dots, Y_{k-1})$ we can rewrite our assumption on S_X and S_Y in a more familiar form: that S_X constitutes at least a ν -proportion of $\mathcal{G}(X_1, \dots, X_{k-1})$ and S_Y constitutes at least a ν -proportion of $\mathcal{G}(Y_1, \dots, Y_{k-1})$. Let X_k be the cluster following X in W , and let Y_k be the cluster preceding Y in W , so $\{X_1, \dots, X_k\}$ and $\{Y_1, \dots, Y_k\}$ are edges of R . Then we choose subsets $Z_j \subseteq Y_j$ and $X_j^{(0)} \subseteq X_j$ of size αn for each $j \in [k]$. A similar argument as in the proof of Lemma 9 implies that we may choose $e \in S_Y$ so that for at least half the members f of $\mathcal{G}(Z_1, \dots, Z_k)$ there is a tight path of length k in G from f to e whose remaining vertex lies in Z_k . Similarly, we may choose a $(k-1)$ -tuple $P^{(0)}$ in S_X and a set $\mathcal{P}^{(0)}$ of tight paths of the form $P^{(0)} + v + e'$ with $v \in Y_k$ and $e' \in \mathcal{G}(Y_1, \dots, Y_{k-1})$ so that the members of $\mathcal{P}^{(0)}$ have distinct terminal $(k-1)$ -tuples which together occupy at least nine-tenths of $\mathcal{G}(Y_1, \dots, Y_{k-1})$. We then proceed by exactly the same algorithm as in the proof of Lemma 9 to repeatedly extend $P^{(j)}$ whilst avoiding Z_1, \dots, Z_k and e . The only difference is that now in the ‘completing the cycle’ state, when we identify tight paths Q and Q' which together connect our final $P^{(j)}$ to e , this does not yield a tight cycle but rather a tight path whose initial $(k-1)$ -tuple lies in S_X and whose final $(k-1)$ -tuple lies in S_Y , as claimed.

It remains only to check the lengths of tight paths which can be obtained in this way. As in Lemma 9 the shortest tight path is achieved by never entering the state of ‘filling an edge’, in which case we obtain a tight path of length $3k + p(k+1)$. On the other hand, exactly as for Lemma 9, by extending W to include all edges of R of non-zero weight before we implement the algorithm, we can obtain tight paths of length up to $(1-\psi)k\mu n/t$, but without using more than $\mu(C)n/t + B$ vertices from any cluster, where $B = B(t, k)$ does not depend on n . \square

10. CONCLUDING REMARKS

Towards an extremal theorem for tight paths and cycles. We mentioned in the introduction that our result is an approximate analogue of the Erdős-Gallai Theorem. The exact analogue is the following, which we conjecture.

Conjecture 37. *For any ℓ , all n -vertex k -graphs with more than $\frac{\ell-k}{k} \binom{n}{k-1}$ edges contain a tight path on ℓ vertices, and all n -vertex k -graphs with more than $\frac{\ell-1}{k} \binom{n-1}{k-1}$ edges have a tight cycle of length at least ℓ .*

As we saw in Section 3, for any fixed ℓ this conjecture is sharp for n satisfying certain divisibility conditions. We also observed there that if $p \ll n^{-(\ell-k)/(\ell-1)}$ then a random hypergraph of density p has fewer ℓ -vertex tight cycles than edges, so we can easily delete all short cycles without significantly altering the density. Hence we cannot ask for the existence of cycles of lengths up to ℓ .

Recall that Győri, Katona and Lemons [12] proved that more than $(\ell-k) \binom{n}{k-1}$ edges suffices to guarantee the existence of an ℓ -vertex tight path, which is weaker than the conjecture both in that it only deals with tight paths and in that a factor k more edges than conjectured are required. However their result holds for all $\ell \leq n$. Our result is off by a factor only $(1+o(1))$ from the conjectured number of edges, but we require ℓ to be linear in n . Possibly one could even prove the conjecture exactly in our range of ℓ using the Stability Method, but we do not believe an attempt to do so is worthwhile, for the following reason.

For large cycles we do not believe Conjecture 37 is best possible. It is easy to check that if $\alpha \gg n^{-1/2}$ then no designs (or even set systems with nearly as many edges as a design) with sets of size αn and without pairwise intersections of size two or greater exist: The lines of \mathbb{F}_p^2 for prime $p \approx n^{1/2}$ form a design with all pairs covered exactly once, so this is optimal.

The best lower bound we know of in this range of α is the simple construction presented in Section 3 in which we take all edges meeting a fixed $(\alpha n/k - 1)$ -vertex set. There is a substantial gap between this lower bound and Conjecture 37. It would be very interesting—but also, we suspect, very difficult—to close this gap. In fact, we think that it is already difficult to solve the following problem for some fixed $\alpha \in (0, 1)$.

Problem 38. *Determine the limit of the maximum edge density of n -vertex 3-graphs which do not contain αn -vertex tight paths.*

We have no suggestions of good candidates for extremal structures, which would also be interesting to obtain. For α close to 1, it even seems possible that the construction presented in Section 3 in which we take an $(\alpha n - 1)$ -vertex clique and all further edges which meet it in less than $k - 1$ vertices could be extremal.

Spanning structures and stability. Let us support our claim that Lemma 6 property (c) can be useful for proving results involving spanning subgraphs. For this purpose we first briefly sketch a proof of Rödl, Ruciński and Szemerédi [29] giving the following Dirac-type condition for tight Hamilton cycles in k -graphs, and then explain how their approach can be simplified with the help of Lemma 6.

Theorem 39 ([29], Theorem 1.1). *For each $k \geq 3$ and $\gamma > 0$ there exists n_0 such that for each $n \geq n_0$ the following holds. If G is a k -graph on n vertices and each $(k-1)$ -set of vertices of G is contained in at least $(1/2 + \gamma)n$ edges of G , then G has a tight Hamilton cycle.*

The strategy of Rödl, Ruciński and Szemerédi is as follows. First, they show that the codegree condition implies that for each vertex u of G , there are $\Theta(n^{2k-2})$ ‘absorbing structures for u ’, that is, vertex tuples (v_1, \dots, v_{2k-2}) such that both of the tuples (v_1, \dots, v_{2k-2}) and $(v_1, \dots, v_{k-1}, u, v_k, \dots, v_{2k-2})$ form tight paths in G . Second, they use a probabilistic argument to show that there is a collection of pairwise-disjoint $(2k-2)$ -vertex tight paths which cover a small fraction of $V(G)$ and which have the property that

for each $u \in V(G)$, $\Theta(n)$ of these tight paths are absorbing structures for u . Third, they establish a Connecting Lemma which allows them to connect the collection of $(2k - 2)$ -vertex tight paths into one ‘absorbing path’, which is a tight path P such that for any not too large set $S \subseteq V(G) \setminus V(P)$, there is a tight path on the vertices $V(P) \cup S$ with the same endpoints as P . Fourth, they show that P can be extended to an almost-spanning tight cycle. Then setting $S = V(G) \setminus V(C)$ and using the absorbing property of P completes their proof.

In their approach, the Regularity Lemma is used only to complete the fourth step, and most of the work is in proving the Connecting Lemma. We can use Lemma 6 to simplify this approach. Property (c) implies that since each $u \in V(G)$ has many absorbing structures, so each u also has many absorbing structures supported on the regular slice \mathcal{J} returned by Lemma 6, and the same probabilistic argument as in their second step then gives a collection of $(2k - 2)$ -vertex tight paths supported on \mathcal{J} with otherwise the same properties. But now their third step can be replaced by showing that the reduced graph $R(G)$ is (after deleting a few vertices) tightly connected. This, however, is easy and can be obtained as in our proof of Theorem 2. Their fourth step can be carried out in more or less the same way on \mathcal{J} , and the proof is then completed exactly as before.

Besides simplifying the proof, we also believe it would be much easier to prove a ‘stability version’ of Theorem 39 using our approach, because the tight connectivity of $R(G)$ we use is a comparatively easy concept to work with. It appears more challenging to prove a stability version of the Connecting Lemma in their approach.

Furthermore, it would be interesting to prove similar results for more general ‘tight-path-like’ hypergraphs. To do this one will need a suitable ‘embedding lemma’ which can complete the fourth step (in either approach) and this does not currently exist. Given such an ‘embedding lemma’, it would be easy to modify our approach to complete the proof of such a result.

Entropy preserving regular slices. We note that one can ask for the slice \mathcal{J} provided by Lemma 6 to satisfy further properties. In particular, one property which is useful in enumeration is that entropy is preserved.

The *binary entropy* of a number $x \in (0, 1)$ is defined to be $\mathbf{H}(x) = -x \log_2 x - (1 - x) \log_2(1 - x)$, and we define $\mathbf{H}(0) = \mathbf{H}(1) = 0$. Given a weighted reduced k -graph R , we define the binary entropy $\mathbf{H}(R)$ of R to be the average over all $e \in \binom{V(R)}{k}$ of $\mathbf{H}(d^*(e))$. Then we can ask for the following additional property of Lemma 6. The number of n -vertex k -graphs G whose reduced graph $R(G)$ given by Lemma 6 has $\mathbf{H}(R(G)) \leq x$ is at most $2^{x \binom{n}{k} + \varepsilon_k n^k}$.

We define the binary entropy $\mathbf{H}(G, \mathcal{P}^*)$ of G with respect to a family of partitions \mathcal{P}^* to be the average of $\mathbf{H}(d(G | \hat{P}(Q)))$ over the polyads $\hat{P}(Q)$ of \mathcal{P}^* (where $d(G | \hat{P}(Q))$ is the relative density of G with respect to $\hat{P}(Q)$ as defined in Section 4.1). Then the method used in the proof of Lemma 6(a) can be used to show that $\mathbf{H}(R(G))$ is with high probability close to $\mathbf{H}(G, \mathcal{P}^*)$, where \mathcal{P}^* is as given in the proof of Lemma 6. Hence we may ask that Lemma 6 returns a regular slice \mathcal{J} such that $\mathbf{H}(R(G)) \approx \mathbf{H}(G, \mathcal{P}^*)$. Then a standard and easy enumeration argument yields the desired conclusion.

Regular slices and the Hypergraph Blow-up Lemma. Recently, Keevash [17] provided a new major tool for extremal hypergraph theory, the Hypergraph Blow-up Lemma. However, while our Regular Slice Lemma, Lemma 6, is based on the Strong Hypergraph Regularity Lemma, Lemma 20, the Hypergraph Blow-up Lemma requires the stronger regularity properties given by the Regular Approximation Lemma, Lemma 21. Hence Lemma 6 is not suitable for an application together with the Hypergraph Blow-up Lemma.

We remark though that Lemma 6 can be modified appropriately to allow for such an application. The proof of such a modified version is almost identical to the proof of Lemma 6, but uses a version of the Strengthened Regularity Lemma, Lemma 22, which gives regularity properties of comparable strength to those in Lemma 21. It is easy to modify the proof of Lemma 22 to obtain this.

APPENDIX A. DERIVATION OF LEMMAS 23 AND 25

We first describe the differences between Lemma 23 and the Lemma provided by Cooley, Fountoulakis, Kühn, and Osthus [8, Lemma 4]. Firstly, their version also allows to count copies of H in which multiple vertices of H may be embedded within the same cluster. We do not need this strengthening, so we omit it for notational convenience.

Secondly, their lemma includes an additional constant d_k with $\varepsilon_k \ll d_k$ and $1/d_k \in \mathbb{N}$, and requires for any edge $e \in H$ that G is (d_k, ε_k, r) -regular with respect to $\hat{\mathcal{J}}_X$, where $X = \{V_j : j \in e\}$, whereas our lemma only requires that G is (d, ε_k, r) -regular with respect to $\hat{\mathcal{J}}_X$ for some d which may depend on X . We now describe how this apparently stronger result may be derived.

Introduce a constant d_k with $d_k \ll \beta$, and also a constant γ with $\varepsilon \ll \gamma \ll \varepsilon_k, d_2, \dots, d_{k-1}$. Let X be any k -set of clusters of \mathcal{J} . The version of the Dense Extension Lemma [8, Lemma 7], applied with \mathcal{J} in place of \mathcal{G} and $\mathcal{H} \setminus \mathcal{H}^{(k-1)}$ in place of \mathcal{H} , implies that at most $\gamma|K_i(\hat{\mathcal{J}}_X)|$ edges of G_X are contained in more than $\prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}) - \binom{k}{i}}$ copies of \mathcal{H} . Also, the Slicing Lemma [8, Lemma 8] implies that we can partition G_X into subgraphs $G_X^0, G_X^1, \dots, G_X^{p_X}$ for some integer p_X , so that $|G_X^0| \leq d_k|K_i(\hat{\mathcal{J}}_X)|$ and G_X^i is $(d_k, 2\varepsilon_k, r)$ -regular with respect to $\hat{\mathcal{J}}_X$ for each $i \geq 1$. However we choose $q_X \in [p_X]$ for each X , we may then apply [8, Lemma 4] to establish that there are $(1 \pm \beta/3)m^s \prod_{i=2}^k d_i^{e_i(\mathcal{H})}$ copies of \mathcal{H} in $\bigcup_X G_X^{q_X}$. Different choices of q_X give distinct copies of \mathcal{H} , so summing over all possible choices we obtain $(\prod_{e \in \mathcal{H}} d^*(e) \pm 2\beta/3)m^s \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H})}$ copies of \mathcal{H} in G , having counted all copies except for those which contain an edge from some G_X^0 . Fact 3 tells us that for any X there are at most $d_k|K_i(\hat{\mathcal{J}}_X)| \leq d_k \prod_{i=2}^{k-1} d_i^{\binom{k}{i}} m^k$ such edges, each of which lies in at most $\prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}) - \binom{k}{i}} m^{s-k}$ copies of \mathcal{H} , except for at most $\gamma|K_i(\hat{\mathcal{J}}_X)| \leq d_k \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H})} m^k$ edges, which each lie in up to m^{s-k} copies of \mathcal{H} . Overall this adds at most $(\beta/3)m^s \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H})}$ further copies of \mathcal{H} in G , giving Lemma 23.

Similarly, [21, Lemma 4.6] differs from our Lemma 25 in that it requires that for each $A \in \binom{[s]}{k}$ such that there is an edge of \mathcal{H} with index A , the graph $\mathcal{G}^{(k)}[V_A]$ is regular with density exactly d_k with respect to $\mathcal{G}^{(k-1)}[A]$, whereas we require density at least d_k . Moreover, [21, Lemma 4.6] also gives a formula for the typical number of extensions, whereas we give only a lower bound. This latter difference makes it elementary to reduce our lemma to theirs by applying the Slicing Lemma [8, Lemma 8].

APPENDIX B. PROOF OF LEMMA 22.

Let us first justify why we need Lemma 22. As we saw in the proof of Lemma 6, Lemma 22 is used to ensure that the link of *every* vertex of G is represented in the regular slice. It may not be clear why Lemma 20 is not good enough for this purpose. But observe that if G is a k -graph and \mathcal{P}^* a family of partitions obtained by applying Lemma 20, then we can construct a k -graph G' from G by associating to each slice \mathcal{J} through \mathcal{P}^* a vertex $v_{\mathcal{J}}$ of $V(G)$, and then removing all edges of G of the form $\{v_{\mathcal{J}}\} \cup e$ where $e \in \mathcal{J}^{(k-1)}$. Then G' is also regular with respect to \mathcal{P}^* , but whatever slice \mathcal{J}

through \mathcal{P}^* is chosen, the link of $v_{\mathcal{J}}$ in G' is not represented on \mathcal{J} , so \mathcal{J} does not satisfy Lemma 6(c).

Now we turn to the proof of Lemma 22. The basic idea is simple. We will apply the Regular Approximation Lemma, Lemma 21, to our k -graph G to generate a family of partitions $\tilde{\mathcal{P}}^*$ with $\tilde{t} \leq t_1^*$ clusters. We will then take a random equipartition of each cluster into $p(t_1^*)$ parts to obtain a partition \mathcal{P} , and let \mathcal{P}^* be a refinement of $\tilde{\mathcal{P}}^*$ generated by \mathcal{P} . We will show that deterministically G has the desired regularity with respect to \mathcal{P}^* , and that with high probability, for any set \tilde{X} of clusters of $\tilde{\mathcal{P}}^*$ and any \tilde{X} -consistent set X of clusters of \mathcal{P}^* , the rooted H -densities on X of any slice through \mathcal{P}^* are close to the densities on \tilde{X} of the corresponding slice through $\tilde{\mathcal{P}}^*$.

Let us explain why it is necessary in this proof to apply Lemma 21 rather than Lemma 20 to obtain $\tilde{\mathcal{P}}^*$. The reason is that typically $p(\cdot)$ grows fast and we must therefore split each cluster into parts which are too small to control using any regularity obtainable from Lemma 20.

To analyse the random equipartition we need the following standard concentration inequality for the hypergeometric distribution.

Theorem 40 (see [15, Theorem 2.10]). *Given a set V and a subset W of V , let V' be chosen uniformly at random from all ℓ -subsets of V . We have $\mathbb{P}\left(|V' \cap W| = \frac{\ell|W|}{|V|} \pm t\right) \geq 1 - 2 \exp\left(-\frac{t^2}{2|V|}\right)$. \square*

Proof of Lemma 22. Assume we are given integers q, t_0 and s , a constant ε_k , functions $r : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$ and a monotone increasing function $p : \mathbb{N} \rightarrow \mathbb{N}$. Without loss of generality, we may assume that r is also monotone increasing while ε is monotone decreasing. We set

$$\nu = \varepsilon_k^3 / (16k!) \quad \text{and} \quad t_0^* = \max(t_0, 4k^2 / \varepsilon_k).$$

We let $\varepsilon^* : \mathbb{N} \rightarrow (0, 1]$ be a monotone decreasing function such that for each x we have

$$(33) \quad \varepsilon^*(x) \leq \min\left(\varepsilon(p(x)x)^2, \frac{\varepsilon_k^4}{16r(p(x)x)^2}\right).$$

We further require $\varepsilon^*(\cdot)$ to be small enough that we can apply Lemma 24 with $\alpha = 1/p(x)$, and Lemma 23 with $\beta = \frac{\varepsilon_k}{10}$ and any $s \leq 1/\varepsilon_k$, to $(\mathbf{d}, \varepsilon^*(x), \varepsilon^*(x), 1)$ -regular k -complexes, provided that each d_i is at least $1/x$.

Let t_1^* and n_0^* be returned by Lemma 21 with inputs q, t_0^*, s, ν and $\varepsilon^*(\cdot)$. Let $t_1 := p(t_1^*)t_1^*$, $\varepsilon^* := \varepsilon^*(t_1^*)$, $\varepsilon := \varepsilon(t_1)$ and $r := r(t_1)$, so in particular we have $\varepsilon^* \leq \varepsilon^2$ by definition of ε^* . Let $m_0 \geq n_0^*$ be large enough for the above applications of Lemmas 23 and 24 with $x = t_1^*$ and clusters of size $m \geq m_0$. We then choose $n_0 \geq m_0 t_1$ sufficiently large for the union bound at the end of the proof.

Let V be a set of $n \geq n_0$ vertices, where n is divisible by $t_1!$, and let \mathcal{Q} partition V into at most q parts of equal size. Let G_1, \dots, G_s be edge-disjoint k -graphs on the vertex set V . We start by applying Lemma 21 (with the inputs stated above), which yields k -graphs G'_1, \dots, G'_s with $|G'_i \Delta G_i| \leq \nu n^k$ for each $1 \leq i \leq s$, and a $(k-1)$ -family of partitions $\tilde{\mathcal{P}}^*$ on V which is $(t_0^*, t_1^*, \varepsilon^*)$ -equitable, whose ground partition $\tilde{\mathcal{P}}$ refines \mathcal{Q} , and which is such that each G'_i is perfectly $(\varepsilon^*, 1)$ -regular with respect to $\tilde{\mathcal{P}}^*$. Let the density vector of $\tilde{\mathcal{P}}^*$ be $\mathbf{d} = (d_2, \dots, d_{k-1})$. By definition of 'equitable', the entries of \mathbf{d} are all at least $1/t_1^*$.

We let the partition \mathcal{P} of V be obtained by partitioning each cluster X of $\tilde{\mathcal{P}}^*$ uniformly at random into $p(t_1^*)$ parts of equal size. We let \tilde{t} be the number of clusters of $\tilde{\mathcal{P}}$, and t

be the number of clusters of \mathcal{P} . Thus we have $t = p(t_1^*)\tilde{t}$, which is property (c). Since $\tilde{\mathcal{P}}$ refines \mathcal{Q} , so also \mathcal{P} refines \mathcal{Q} , giving property (a).

We now obtain a family of partitions \mathcal{P}^* with ground partition \mathcal{P} as follows. For each $2 \leq i \leq k-1$ and each i -cell C in $\tilde{\mathcal{P}}^*$ on the i clusters \mathcal{C} of $\tilde{\mathcal{P}}$, we put into \mathcal{P}^* each of the i -uniform induced subgraphs of C obtained by choosing i clusters of \mathcal{P} , one in each cluster of \mathcal{C} . We then need to add further 2-cells between pairs of clusters in \mathcal{P} which were both contained in one cluster of $\tilde{\mathcal{P}}$. We do this by choosing for each such pair X_1, X_2 of clusters in \mathcal{P} an arbitrary equipartition of the complete bipartite graph with vertex classes X_1 and X_2 into $1/d_2$ cells which are each $(d_2, \sqrt{\varepsilon^*}, 1)$ -regular with respect to the 1-graph formed by the vertices of X_1 and X_2 (for example, it is not hard to show that a random partition has this property with high probability). We then need to choose all those 3-cells whose supporting clusters do not lie in distinct parts of $\tilde{\mathcal{P}}$, and so on. We do this in the same way, while also ensuring that we are consistent with the 2-cells we just chose, and so on, so that we obtain a family of partitions. By construction \mathcal{P}^* is generated from $\tilde{\mathcal{P}}^*$ by \mathcal{P} , giving property (b). Now consider any $\tilde{\mathcal{P}}$ -partite set $Q \in \binom{V}{k}$, let $\tilde{\mathcal{J}}(Q)$ be the k -partite $(k-1)$ -complex whose edge set is the union of the cells \tilde{C}_Q of $\tilde{\mathcal{P}}^*$ which contain proper subsets of Q , and similarly let $\mathcal{J}(Q)$ be the k -partite $(k-1)$ -complex whose edge set is the union of the cells C_Q of \mathcal{P}^* which contain proper subsets of Q . The fact that $\tilde{\mathcal{P}}^*$ is $(t_0^*, t_1^*, \varepsilon^*)$ -equitable implies that $\tilde{\mathcal{J}}(Q)$ is $(\mathbf{d}, \varepsilon^*, \varepsilon^*, \mathbf{1})$ -regular, whereupon Lemma 24 implies that $\mathcal{J}(Q)$ is $(\mathbf{d}, \sqrt{\varepsilon^*}, \sqrt{\varepsilon^*}, \mathbf{1})$ -regular. It follows that for any $2 \leq i \leq k-1$ and any i -cell C of \mathcal{P}^* , if C was obtained as an induced subgraph of an i -cell of $\tilde{\mathcal{P}}^*$, then C is $(d_i, \sqrt{\varepsilon^*}, 1)$ -regular with respect to the $(i-1)$ -cells of \mathcal{P}^* on which it is supported. On the other hand, if C was not obtained in this manner, then by construction C is $(d_i, \sqrt{\varepsilon^*}, 1)$ -regular with respect to these $(i-1)$ -cells. Since $t_0^* \geq t_0$, $t_1 = p(t_1^*)t_1^* \geq t_1^*$ and $\sqrt{\varepsilon^*} \leq \varepsilon$, we conclude that \mathcal{P}^* is (t_0, t_1, ε) -equitable with density vector \mathbf{d} . This establishes property (d).

We next want to establish property (e). Fix a graph $G = G_i$, and let $G' = G'_i$. The next claim shows that if G is not (ε_k, r) -regular with respect to some $\tilde{\mathcal{P}}$ -partite polyad of \mathcal{P}^* , then $G\Delta G'$ is dense on that polyad.

Claim 41. *For any $\tilde{\mathcal{P}}$ -partite k -set $Q \in \binom{V}{k}$, if G is not (ε_k, r) -regular with respect to the polyad $\hat{P} = \hat{P}(Q; \mathcal{P}^*)$ of \mathcal{P}^* , then $d(G'\Delta G | \hat{P}) > (\varepsilon_k/2)^2$.*

Proof. Let $\mathbf{W} = (W_1, \dots, W_r)$ be any collection of r subgraphs of \hat{P} such that $|K_k(\mathbf{W})| \geq \varepsilon_k |K_k(\hat{P})|$. We split these r subgraphs into *large* subgraphs W_j with $|K_k(W_j)| \geq \sqrt{\varepsilon^*} |K_k(\hat{P})|$, and the remaining *small* subgraphs.

Since Q is $\tilde{\mathcal{P}}$ -partite, Lemma 24 implies that G' is $(\sqrt{\varepsilon^*}, 1)$ -regular with respect to \hat{P} , so on any large W_j we have $d(G'|W_j) = d(G'|\hat{P}) \pm \sqrt{\varepsilon^*}$. On the other hand the total number of k -cliques supported on small subgraphs W_j is by definition of ‘small’ less than $r\sqrt{\varepsilon^*} |K_k(\hat{P})|$. It follows that

$$d(G'|\mathbf{W}) = d(G'|\hat{P}) \pm \frac{2r\sqrt{\varepsilon^*}}{\varepsilon_k}.$$

If $d(G'\Delta G|\hat{P}) \leq (\varepsilon_k/2)^2$, then by definition we have $d(G'|\hat{P}) = d(G|\hat{P}) \pm (\varepsilon_k/2)^2$. Furthermore, since we assumed $|K_k(\mathbf{W})| \geq \varepsilon_k |K_k(\hat{P})|$, we have

$$d(G'|\mathbf{W}) = d(G|\mathbf{W}) \pm d(G'\Delta G|\hat{P}) \frac{|K_k(\hat{P})|}{|K_k(\mathbf{W})|} = d(G|\mathbf{W}) \pm \varepsilon_k/4.$$

Putting these together we have

$$d(G|\mathbf{W}) = d(G|\hat{P}) \pm \left(\frac{\varepsilon_k^2}{4} + \frac{2r\sqrt{\varepsilon_k^*}}{\varepsilon_k} + \frac{\varepsilon_k}{4} \right) \stackrel{(33)}{=} d(G|\hat{P}) \pm \varepsilon_k,$$

hence G is (ε_k, r) -regular with respect to \hat{P} . \square

Let \mathcal{B} be the set of all polyads $\hat{P} = \hat{P}(Q; \mathcal{P}^*)$ of \mathcal{P}^* for which $Q \in \binom{V}{k}$ is $\tilde{\mathcal{P}}$ -partite and G is not (ε_k, r) -regular with respect to \hat{P} . Now using the fact that $e(G\Delta G') \leq \nu n^k$ and Claim 41 we have

$$\nu n^k \geq \sum_{\hat{P} \in \mathcal{B}} d(G\Delta G' | \hat{P}) |K_k(\hat{P})| > (\varepsilon_k/2)^2 \sum_{\hat{P} \in \mathcal{B}} |K_k(\hat{P})|,$$

and hence the number of $\tilde{\mathcal{P}}$ -partite sets $Q \in \binom{V}{k}$ such that G is not (ε_k, r) -regular with respect to $\hat{P}(Q; \mathcal{P}^*)$ is at most $\nu n^k (\varepsilon_k/2)^{-2} \leq \frac{1}{2} \varepsilon_k \binom{n}{k}$, where the inequality is by choice of ν and n_0 . Moreover, the number of k -sets $Q \in \binom{V}{k}$ which are not $\tilde{\mathcal{P}}$ -partite is at most $\binom{n}{k-1} \frac{(k-1)n}{t_0^*}$, which by choice of t_0^* is at most $\frac{1}{2} \varepsilon_k \binom{n}{k}$. Putting these together we see that there are at most $\varepsilon_k \binom{n}{k}$ \mathcal{P} -partite sets $Q \in \binom{V}{k}$ such that G is not (ε_k, r) -regular with respect to $\hat{P}(Q; \mathcal{P}^*)$. In other words, G is (ε_k, r) -regular with respect to \mathcal{P}^* , proving property (e).

It remains to show that property (f) holds with positive probability. To that end, fix a k -graph $G = G_i$, $1 \leq i \leq 1/\varepsilon_k$, a k -graph H equipped with roots x_1, \dots, x_ℓ such that $v(H) \leq 1/\varepsilon_k$, vertices v_1, \dots, v_ℓ in V , a slice $\tilde{\mathcal{J}}$ through $\tilde{\mathcal{P}}^*$ and a $v(\mathcal{H}^{skel})$ -tuple of clusters \tilde{X} of $\tilde{\mathcal{P}}^*$, and let $h := v(\mathcal{H}^{skel})$. Now for any permutation ϕ of $1, \dots, h$, the rooted copies of H in G supported on $\tilde{\mathcal{J}}[\tilde{X}]$, with the i th vertex of \mathcal{H}^{skel} in the $\phi(i)$ th cluster of \tilde{X} , correspond to a h -uniform $\tilde{\mathcal{P}}$ -partite hypergraph F_ϕ . Let X be an \tilde{X} -consistent h -tuple of clusters of \mathcal{P}^* . Then we want to show that for any fixed slice \mathcal{J} through \mathcal{P}^* such that $\mathcal{J}[X] \subseteq \tilde{\mathcal{J}}[\tilde{X}]$, with high probability we have

$$d_H(G; v_1, \dots, v_\ell, \tilde{\mathcal{J}}[\tilde{X}]) = d_H(G; v_1, \dots, v_\ell, \mathcal{J}[X]) \pm \varepsilon_k.$$

By Lemma 23 the number of labelled \mathcal{H}^{skel} -copies in $\tilde{\mathcal{J}}[\tilde{X}]$ for which no two vertices of \mathcal{H}^{skel} lie in the same cluster of $\tilde{\mathcal{J}}$ is

$$n'_{\mathcal{H}^{skel}}(\tilde{\mathcal{J}}[\tilde{X}]) = h! \left(1 \pm \frac{\varepsilon_k}{10}\right) \left(\frac{n}{t}\right)^{v(\mathcal{H}^{skel})} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})},$$

where the $h!$ term is the number of possible allocations of vertices of \mathcal{H}^{skel} to clusters of $\tilde{\mathcal{J}}$. Similarly

$$n'_{\mathcal{H}^{skel}}(\mathcal{J}[X]) = h! \left(1 \pm \frac{\varepsilon_k}{10}\right) \left(\frac{n}{t}\right)^{v(\mathcal{H}^{skel})} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}.$$

It therefore suffices to show that with high probability, for each of the $h!$ permutations ϕ we have

$$(34) \quad d(F_\phi) = d(F_\phi[X]) \pm \frac{\varepsilon_k}{2(h!)} \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})}.$$

Fix any permutation ϕ , and for each $0 \leq j \leq h$, let F_j denote the subgraph of F_ϕ given by taking the first j clusters from X and the last $h - j$ clusters from \tilde{X} . Thus we have $F_0 = F_\phi$ and $F_h = F_\phi[X]$. Observe that for each j , the distribution of the j th cluster of X is the uniform distribution over all n/t -subsets of the j th cluster of \tilde{X} . We can

thus use the following claim to show that for each $1 \leq j \leq h$, with high probability the densities $d(F_{j-1})$ and $d(F_j)$ are close.

Claim 42. *For each $s, p \in \mathbb{N}$ and each $\delta > 0$, let B be an s -partite s -uniform hypergraph with parts V_1, \dots, V_s . Choose uniformly at random a subset $V' \subseteq V_1$ of size $|V_1|/p$, and let $B' = B[V', V_2, \dots, V_s]$. Then with probability at least*

$$1 - \frac{4}{\delta} \exp\left(-\frac{\delta^4 |V_1|}{32p^2}\right),$$

we have $d(B') = d(B) \pm \delta$.

Proof. We partition the vertices V_1 into sets $W_0, W_1, \dots, W_{2/\delta}$, with the property that all vertices in W_j have degree in B in the interval

$$\left(\frac{(j-1)\delta|V_2|\cdots|V_s|}{2}, \frac{j\delta|V_2|\cdots|V_s|}{2}\right].$$

Let V' be a random subset of V_1 of size $\ell = |V_1|/p$, and $B' = B[V', V_2, \dots, V_s]$. By Theorem 40, for each $1 \leq j \leq 2/\delta$ we have

$$\mathbb{P}(|V' \cap W_j| = \frac{|W_j|}{p} \pm t) \geq 1 - 2 \exp\left(-\frac{t^2}{2|V_1|}\right).$$

Taking $t = \delta^2|V_1|/(4p)$ and using a union bound, we conclude that we have

$$|V' \cap W_j| = \frac{|W_j|}{p} \pm \frac{\delta^2|V_1|}{4p}$$

for each $1 \leq j \leq 2/\delta$ with probability at least

$$1 - \frac{4}{\delta} \exp\left(-\frac{\delta^4 |V_1|}{32p^2}\right).$$

Conditioning on this likely event, we have

$$\begin{aligned} e(B') &= \sum_{j=1}^{2/\delta} |V' \cap W_j| \cdot \left(j - \frac{1}{2} \pm \frac{1}{2}\right) \frac{\delta|V_2|\cdots|V_s|}{2} \\ &= \sum_{j=1}^{2/\delta} \left(\frac{|W_j|}{p} \pm \frac{\delta^2|V_1|}{4p}\right) \cdot \left(j - \frac{1}{2} \pm \frac{1}{2}\right) \frac{\delta|V_2|\cdots|V_s|}{2} \\ &= \sum_{j=1}^{2/\delta} \frac{|W_j|}{p} \cdot \left(j - \frac{1}{2} \pm \frac{1}{2}\right) \frac{\delta|V_2|\cdots|V_s|}{2} \pm \frac{\delta}{2p} |V_1| \cdots |V_s| \\ &= \frac{1}{p} e(B) \pm \left(\sum_{j=1}^{2/\delta} \frac{|W_j|}{p} \cdot \frac{\delta|V_2|\cdots|V_s|}{2}\right) \pm \frac{\delta}{2p} |V_1| \cdots |V_s| \\ &= \frac{1}{p} e(B) \pm \frac{\delta}{p} |V_1| \cdots |V_s|, \end{aligned}$$

which gives the desired conclusion. □

Applying Claim 42 to $B = F_{j-1}$ and $B' = F_j$ for each $1 \leq j \leq h$, with

$$p = p(t_1^*) \quad \text{and} \quad \delta = \frac{1}{2h(h!)} \varepsilon_k \prod_{i=2}^{k-1} d_i^{e_i(\mathcal{H}^{skel})},$$

we conclude that with probability at least

$$1 - \frac{4h}{\delta} \exp\left(-\frac{\delta^4 n}{32p^2 t_1^*}\right),$$

we have $d(F_0) = d(F_h) \pm \delta h$. By choice of δ this implies that the equation (34) holds for the permutation ϕ .

Now we take a union bound over the s choices of $G = G_i$, the $1/\varepsilon_k$ choices of ℓ , the at most $2^{\binom{1/\varepsilon_k}{k}}$ choices of rooted H , the at most n^{1/ε_k} choices of v_1, \dots, v_ℓ , the $\prod_{j=2}^{k-1} d_j^{-\binom{\tilde{t}}{j}}$ choices of $\tilde{\mathcal{J}}$, the at most $\binom{\tilde{t}}{1/\varepsilon_k}$ choices of \tilde{X} , the $\prod_{j=2}^{k-1} d_j^{-\binom{\tilde{t}}{j}}$ choices of \mathcal{J} , the at most $\binom{\tilde{t}}{1/\varepsilon_k}$ choices of X , and the $h! \leq (1/\varepsilon_k)!$ permutations ϕ of $1, \dots, h$. Since we have the lower bound $d_j \geq 1/t_1^*$ for each j , the total number of events over which we take a union bound is polynomial in n . The failure probability of each good event is exponentially small in n , so if n_0 was chosen sufficiently large, with positive probability all good events hold. In other words, with positive probability (f) holds, which completes the proof. \square

REFERENCES

- [1] P. Allen, *Minimum degree conditions for cycles*, in preparation.
- [2] N. Alon, *The largest cycle of a graph with a large minimal degree*, J. Graph Theory **10** (1986), 123–127.
- [3] B. Bollobás, *Combinatorics*, Cambridge University Press, 1986.
- [4] J. A. Bondy, *Pancyclic graphs. I*, J. Combinatorial Theory Ser. B **11** (1971), 80–84.
- [5] G. Chen, R. J. Faudree, R. J. Gould, M. S. Jacobson, and L. Lesniak, *Hamiltonicity in balanced k -partite graphs*, Graphs Combin. **11** (1995), no. 3, 221–231.
- [6] D. Christofides, J. Hladký, and A. Máthé, *Hamilton cycles in dense vertex-transitive graphs*, J. Combin. Theory Ser. B **109** (2014), 34–72.
- [7] F. R. K. Chung, *Regularity lemmas for hypergraphs and quasi-randomness*, Random Structures Algorithms **2** (1991), no. 2, 241–252.
- [8] O. Cooley, N. Fountoulakis, D. Kühn, and D. Osthus, *Embeddings and Ramsey numbers of sparse k -uniform hypergraphs*, Combinatorica **29** (2009), no. 3, 263–297.
- [9] G. A. Dirac, *Some theorems on abstract graphs*, Proc. London Math. Soc. **2** (1952), 69–81.
- [10] P. Erdős and T. Gallai, *On maximal paths and circuits of graphs*, Acta Math. Acad. Sci. Hungar **10** (1959), 337–356.
- [11] W. T. Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, Ann. of Math. (2) **166** (2007), no. 3, 897–946.
- [12] E. Györi, G. Y. Katona, and N. Lemons, *Hypergraph Extensions of the Erdős-Gallai Theorem*, in preparation.
- [13] P. E. Haxell, T. Łuczak, Y. Peng, V. Rödl, A. Ruciński, and J. Skokan, *The Ramsey number for 3-uniform tight hypergraph cycles*, Combin. Probab. Comput. **18** (2009), no. 1-2, 165–203.
- [14] H. Huang, P.-S. Loh, and B. Sudakov, *The size of a hypergraph and its matching number*, Combin. Probab. Comput. **21** (2012), no. 3, 442–450.
- [15] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Wiley-Interscience, New York, 2000.
- [16] P. Keevash, *The existence of designs*, submitted (arXiv:1401.3665).
- [17] P. Keevash, *A hypergraph blow-up lemma*, Random Structures Algorithms **39** (2011), no. 3, 275–376.
- [18] P. Keevash and R. Mycroft, *A geometric theory for hypergraph matching*, Mem. Amer. Math. Soc. **233** (2014), no. 1098.
- [19] Y. Kohayakawa, B. Nagle, V. Rödl, and M. Schacht, *Weak hypergraph regularity and linear hypergraphs*, J. Combin. Theory Ser. B **100** (2010), no. 2, 151–160.
- [20] Y. Kohayakawa, V. Rödl, and J. Skokan, *Hypergraphs, quasi-randomness, and conditions for regularity*, J. Combin. Theory Ser. A **97** (2002), no. 2, 307–352.
- [21] D. Kühn, R. Mycroft, and D. Osthus, *Hamilton ℓ -cycles in uniform hypergraphs*, J. Combin. Theory Ser. A **117** (2010), no. 7, 910–927.
- [22] T. Łuczak, *$R(C_n, C_n, C_n) \leq (4 + o(1))n$* , J. Combin. Theory Ser. B **75** (1999), no. 2, 174–187.
- [23] C. McDiarmid, *On the method of bounded differences*, Surveys in Combinatorics **141** (1989), 148–188.
- [24] J. Moon and L. Moser, *On Hamiltonian bipartite graphs*, Israel J. Math. **1** (1963), 163–165.
- [25] B. Nagle, V. Rödl, and M. Schacht, *The counting lemma for regular k -uniform hypergraphs*, Random Structures Algorithms **28** (2006), no. 2, 113–179.
- [26] V. Nikiforov and R. H. Schelp, *Cycle lengths in graphs with large minimum degree*, J. Graph Theory **52** (2006), no. 2, 157–170.
- [27] V. Rödl, *On a packing and covering problem*, European J. Combin. **6** (1985), 69–78.

- [28] V. Rödl and A. Ruciński, *Dirac-type questions for hypergraphs—a survey (or more problems for Endre to solve)*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 561–590.
- [29] V. Rödl, A. Ruciński, and E. Szemerédi, *An approximate Dirac-type theorem for k -uniform hypergraphs*, Combinatorica **28** (2008), no. 2, 229–260.
- [30] V. Rödl, A. Ruciński, and E. Szemerédi, *Dirac-type conditions for Hamiltonian paths and cycles in 3-uniform hypergraphs*, Adv. Math. **227** (2011), no. 3, 1225–1299.
- [31] V. Rödl and M. Schacht, *Regular partitions of hypergraphs: counting lemmas*, Combin. Probab. Comput. **16** (2007), no. 6, 887–901.
- [32] ———, *Regular partitions of hypergraphs: regularity lemmas*, Combin. Probab. Comput. **16** (2007), no. 6, 833–885.
- [33] V. Rödl and J. Skokan, *Regularity lemma for k -uniform hypergraphs*, Random Structures Algorithms **25** (2004), no. 1, 1–42.
- [34] E. Schmeichel and J. Mitchem, *Bipartite graphs with cycles of all even lengths*, J. Graph Theory **6** (1982), 429–439.
- [35] E. Szemerédi, *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Orsay, 1976), Colloques Internationaux CNRS, vol. 260, CNRS, 1978, pp. 399–401.

DEPARTMENT OF MATHEMATICS, LONDON SCHOOL OF ECONOMICS, HOUGHTON STREET, LONDON WC2A 2AE, U.K.

E-mail address: p.d.allen@lse.ac.uk

DEPARTMENT OF MATHEMATICS, LONDON SCHOOL OF ECONOMICS, HOUGHTON STREET, LONDON WC2A 2AE, U.K.

E-mail address: j.boettcher@lse.ac.uk

GRAZ UNIVERSITY OF TECHNOLOGY, INSTITUTE OF OPTIMIZATION AND DISCRETE MATHEMATICS, STEYRERGASSE 30, 8010 GRAZ, AUSTRIA

E-mail address: cooley@math.tugraz.at

SCHOOL OF MATHEMATICS, UNIVERSITY OF BIRMINGHAM, EDGBASTON, BIRMINGHAM B15 2TT, UK

E-mail address: r.mycroft@bham.ac.uk