

Dirac-type theorems in random hypergraphs

Asaf Ferber*

Matthew Kwan†

Abstract

For positive integers $d < k$ and n divisible by k , let $m_d(k, n)$ be the minimum d -degree ensuring the existence of a perfect matching in a k -uniform hypergraph. In the graph case (where $k = 2$), a classical theorem of Dirac says that $m_1(2, n) = \lceil n/2 \rceil$. However, in general, our understanding of the values of $m_d(k, n)$ is still very limited, and it is an active topic of research to determine or approximate these values. In this paper we prove a “transference” theorem for Dirac-type results relative to random hypergraphs. Specifically, for any $d < k$, any $\varepsilon > 0$ and any “not too small” p , we prove that a random k -uniform hypergraph G with n vertices and edge probability p typically has the property that every spanning subgraph of G with minimum d -degree at least $(1 + \varepsilon)m_d(k, n)p$ has a perfect matching. One interesting aspect of our proof is a “non-constructive” application of the absorbing method, which allows us to prove a bound in terms of $m_d(k, n)$ without actually knowing its value.

1 Introduction

Over the last few decades, there has been a great deal of interest in analogues of combinatorial theorems *relative to a random set*. To give a simple example, let us consider *Mantel’s theorem* [41], a classical theorem asserting that any subgraph of the complete n -vertex graph K_n with more than about half of the $\binom{n}{2}$ possible edges must contain a triangle. The random analogue of Mantel’s theorem says that if one considers a *random* subgraph $G \subseteq K_n$, obtained by including each edge independently at random with some suitable probability $0 < p < 1$, then typically G has the property that each subgraph with more than about half of the edges of G must contain a triangle. That is to say, Mantel’s theorem is “robust” in the sense that an analogous statement typically holds even in the “noisy environment” of a random graph. The study of combinatorial theorems relative to random sets has been closely related to several of the most exciting recent developments in probabilistic and extremal combinatorics, including the sparse regularity method, hypergraph containers and the absorbing method. See [11] for a general survey of this topic.

In the early history of this area, the available methods were somewhat ad-hoc, but recent years have seen the development of some very general tools and techniques that allow one to “transfer” a wide variety of combinatorial theorems to the random setting, without actually needing to know the details of their proofs. As an illustration of this, consider the hypergraph¹ *Turán problem*, which is a vast generalisation of Mantel’s problem. For a k -graph H , let $\text{ex}(n, H)$ be the maximum possible number of edges in an n -vertex k -graph which contains no copy of H , and define the *Turán density* of H as

$$\pi(H) = \lim_{n \rightarrow \infty} \frac{\text{ex}(n, H)}{\binom{n}{k}}$$

(a simple monotonicity argument shows that this limit exists). In the graph case, where $k = 2$, the values of each $\pi(H)$ are given by the celebrated Erdős–Stone–Simonovits theorem [15, 16], but for higher uniformities very little is known about the values of $\pi(H)$. Despite this, Conlon and Gowers [9] and independently Schacht [54] were able to prove an optimal theorem in the random setting. Let $G \sim \mathcal{H}^k(n, p)$ be an instance of the random k -graph with edge probability p ; they proved that if p is not too small then typically G has the property that every subgraph with at least $(\pi_k(H) + \varepsilon)\binom{n}{k}p$ edges has a copy of H .

*Department of Mathematics, University of California, Irvine. Email: asaff@uci.edu. Research supported in part by an NSF grant DMS-1954395.

†Department of Mathematics, Stanford University, Stanford, CA 94305. Email: mattkwan@stanford.edu. Research supported in part by SNSF project 178493 and NSF award DMS-1953990.

¹A *k*-uniform hypergraph, or *k*-graph for short, is a pair $H = (V, E)$, where V is a finite set of *vertices*, and E is a family of k -element subsets of V , referred to as the *edges* of H . Note that a 2-graph is just a graph.

Moreover, they were able to find the optimal range of p (that is, the essentially best possible definition of “not too small”) for which this holds.

In other words, whenever we are able to prove a Turán-type theorem for graphs or hypergraphs, we “automatically” get a corresponding theorem in the random setting. This was quite a striking development: before this work, similar theorems were known only for a few graphs H (and no higher-uniformity hypergraphs), despite rather a lot of effort. For a more detailed history of this problem we refer the reader to [9, 54] and the references therein.

The methods and tools developed by Conlon and Gowers, and by Schacht, were later supplemented by some further work by Conlon, Gowers, Samotij and Schacht [10]. The ideas developed by these authors are very powerful (and actually apply in much more general settings than just Turán-type problems), but a common shortcoming is that none of them are sensitive to “local” information about the individual vertices of a graph or hypergraph, and therefore they are not sufficient for proving relative versions of theorems in which one wishes to understand the presence of *spanning* substructures.

For example, *Dirac’s theorem* [13] famously asserts that every n -vertex graph with minimum degree at least $n/2$ has a *Hamiltonian cycle*: a cycle passing through all the vertices of the graph. A random analogue of this theorem was conjectured by Sudakov and Vu [55] and proved by Lee and Sudakov [39] (see also the refinements in [43, 46]): For any $\varepsilon > 0$, if p is somewhat greater than $\log n/n$, then a random graph $H^2(n, p)$ typically has the property that every spanning subgraph with minimum degree at least $(1 + \varepsilon)np/2$ has a Hamiltonian cycle. Since this work, there has been a lot of interest in random versions of Dirac-type theorems for other types of spanning or almost-spanning subgraphs (see for example [2, 3, 5, 6, 7, 19, 22, 28, 45, 48, 57]), introducing a large number of ideas and techniques that are quite independent of the aforementioned general tools. In this paper we are interested in Dirac-type problems for random *hypergraphs*. Before discussing this further, we take a moment to make some definitions and introduce the topic of (non-random) Dirac-type problems for hypergraphs.

Recall that Dirac’s theorem asserts that every n -vertex graph with minimum degree at least $n/2$ has a Hamiltonian cycle. If n is even then we can take every second edge on this cycle to obtain a *perfect matching*: a set of vertex-disjoint edges that covers all the vertices of our graph. So, Dirac’s theorem can also be viewed as a theorem about the minimum degree required to guarantee a perfect matching. While there are certain generalisations of (Hamiltonian) cycles to hypergraphs², the notion of a perfect matching generalises unambiguously, and we prefer to focus on perfect matchings when considering hypergraphs of higher uniformities.

One subtlety is that in the hypergraph setting there are actually multiple possible generalisations of the notion of minimum degree. For a k -graph $H = (V, E)$ and a subset $S \subseteq V$ of the vertices of H , satisfying $0 \leq |S| \leq k - 1$, we define the *degree* $\deg_H(S)$ of S to be the number of edges of H which include S . The *minimum d -degree* $\delta_d(H)$ of H is then defined to be the minimum, over all d -sets of vertices S , of $\deg_H(S)$. For integers n, k, d such that $1 \leq d \leq k - 1$ and n is divisible by k , let $m_d(k, n)$ be the smallest integer m such that every n -vertex k -graph H with $\delta_d(H) \geq m$ has a perfect matching. Dirac’s theorem says that $m_1(2, n) \leq \lceil n/2 \rceil$, and it is quite easy to see that this is tight.

The problem of determining or approximating the values of $m_d(k, n)$ is fundamental in extremal graph theory, and has attracted a lot of attention in the last few decades (see for example the surveys [50, 58] and the references therein). The main conjecture in this area is as follows.

Conjecture 1.1. *For fixed positive integers $d < k$, we have*

$$m_d(k, n) = \left(\max \left\{ \frac{1}{2}, 1 - \left(1 - \frac{1}{k} \right)^{k-d} \right\} + o(1) \right) \binom{n-d}{k-d},$$

where $o(1)$ represents some error term that tends to zero as n tends to infinity along some sequence of integers divisible by k .

There are constructions showing that the expression in Conjecture 1.1 is a lower bound for $m_d(k, n)$; the hard part is to prove upper bounds.

In much the same way that the Turán densities $\pi(H)$ encode the asymptotic behaviour of the extremal numbers $\text{ex}(H, n)$, it makes sense to define *Dirac thresholds* that encode the asymptotic behaviour of the values of $m_d(k, n)$. However, compared to the Turán case, convergence to a limit is nontrivial; in Section 3 we prove the following result (which will be helpful to state and prove our main result, but is also of independent interest).

²One of these generalisations is called a *Berge cycle*. Actually Clemens, Ehrenmüller and Person [8] recently proved a generalisation of Dirac’s theorem, and a random version of this theorem, for Hamiltonian Berge cycles.

Theorem 1.2. Fix positive integers $d < k$. Then the quantity $m_d(k, n) / \binom{n-d}{k-d}$ converges to a limit $\mu_d(k) \in [0, 1]$, as n tends to infinity along the positive integers divisible by k .

Note that Conjecture 1.1 can then be viewed as a conjecture for the values of the Dirac thresholds $\mu_d(k)$. This conjecture seems to be very difficult, but it has been proved in some special cases: namely, when $5d \geq 2k - 2$, and when $(d, k) \in \{(1, 4), (1, 5)\}$ (see [4, 26, 32, 40, 49, 52, 20]). A number of different upper and lower bounds have also been proved in various cases.

So, the situation is quite similar to the hypergraph Turán problem: the optimal theorems in the non-random setting are not known, but there is still some hope of proving a “transference” theorem, giving bounds in the random setting in terms of the (unknown) Dirac thresholds $\mu_d(k)$.

Of course, in order to prove a random analogue of any extremal theorem, in addition to having a handle on the extremal theorem one also needs to have a good understanding of random graphs and hypergraphs. This presents a rather significant obstacle when investigating perfect matchings, because the study of perfect matchings in random hypergraphs is notoriously difficult. Famously, *Shamir’s problem* asks for which p a random k -graph $H^k(n, p)$ has a perfect matching, and this was resolved only a few years ago in a tour-de-force by Johansson, Kahn and Vu [30] (see also the new simpler proof in [21], and the refinement in [31]). Roughly speaking, they proved that if p is large enough that $H^k(n, p)$ typically has no isolated vertices (the threshold value of p is about $n^{1-k} \log n$), then $H^k(n, p)$ typically has a perfect matching. All known proofs of this theorem are quite “non-constructive”, involving some ingenious way to show that a perfect matching is likely to exist without being able to say much about its properties or how to find it.

In any case, it is natural to make the following conjecture, “transferring” Dirac-type theorems to random hypergraphs.

Conjecture 1.3. Fix $\gamma > 0$ and positive integers $d < k$, and consider any $0 < p < 1$ (which may be a function of n). Suppose that n is divisible by k . Then a.a.s.³ $G \sim H^k(n, p)$ has the property that every spanning subgraph $G' \subseteq G$ with $\delta_d(G') \geq (\mu_d(k) + \gamma) \binom{n-d}{k-d} p$ has a perfect matching.

Note that in the above conjecture we do not make any assumption on p , though in some sense we are implicitly assuming $p = \Omega(n^{d-k} \log n)$, because otherwise one can show that a random k -graph $G \sim H^k(n, p)$ will a.a.s. have $\delta_d(G) = 0$ (meaning that there is no subgraph G' satisfying the condition in the conjecture). Due to the aforementioned difficulty of studying perfect matchings in random hypergraphs, we believe that Conjecture 1.3 will be extremely difficult to prove for small p (especially for $p \approx n^{1-k} \log n$), and therefore we believe that the hardest (and most interesting) case is where $d = 1$. On the other extreme, if $d = k - 1$ then it suffices to consider the regime where $p = \Omega(n^{-1} \log n)$, which is substantially easier due to certain techniques which allow one to reduce the problem of finding hypergraph perfect matchings to the problem of finding perfect matchings in certain bipartite graphs⁴. Using such a reduction, the $d = k - 1$ case of Conjecture 1.3 was proved by Ferber and Hirschfeld [17].

Our main result in this paper is the following substantial progress towards Conjecture 1.3, proving it for all $d < k$ under certain restrictions on p (even though the values of $\mu_d(k)$ are in general unknown).

Theorem 1.4. Fix $\gamma > 0$ and positive integers $d < k$. Then there is some $C > 0$ such that the following holds. Suppose that $p \geq \max\{n^{-k/2+\gamma}, Cn^{-k+2}\}$, and that n is divisible by k . Then a.a.s. $G \sim H^k(n, p)$ has the property that every spanning subgraph $G' \subseteq G$ with $\delta_d(G') \geq (\mu_d(k) + \gamma) \binom{n-d}{k-d} p$ has a perfect matching.

Recalling the implicit assumption $p = \Omega(n^{d-k} \log n)$, Theorem 1.4 actually resolves the $d > k/2$ case of Conjecture 1.3, and comes very close to resolving the case $d = k/2$ (if d is even). Also, note that except in the case where $k = 3$ and $d = 1$, the assumption $p \geq Cn^{-k+2} \log n$ is superfluous (being satisfied automatically when $p = \Omega(n^{d-k} \log n)$ and $p \geq n^{-k/2+\gamma}$). Actually, this particular assumption can be weakened quite substantially, but in the interest of presenting a clear proof, we discuss how to do this only informally, in Section 9.

There are a number of different ideas and ingredients that go into the proof of Theorem 1.4. Perhaps the most crucial one is a *non-constructive* way to apply the so-called *absorbing method*. To say just a few words about the absorbing method: in various different contexts, it is much easier to find *almost*-spanning substructures than genuine spanning substructures. For example, a perfect matching is a collection of

³By “asymptotically almost surely”, or “a.a.s.”, we mean that the probability of an event is $1 - o(1)$. Here and for the rest of the paper, asymptotics are as $n \rightarrow \infty$, unless stated otherwise.

⁴While hypergraph matchings are in general not well understood, there are a number of extremely powerful tools available for studying matchings in bipartite graphs (such as Hall’s theorem).

disjoint edges that cover all the vertices of a hypergraph, but it is generally much easier to find a collection of disjoint edges that cover *almost* all of the vertices of a hypergraph. The insight of the absorbing method is that one can sometimes find small “flexible” substructures called *absorbers*, arranged in a way that allows one to make local modifications to transform an almost-spanning structure into a spanning one. This method was pioneered by Erdős, Gyárfás and Pyber [14], and was later systematised by Rödl, Ruciński and Szemerédi [51, 52], in connection with their study of Dirac-type theorems in hypergraphs.

In previous work, the typical approach was to build absorbers in a “bare-hands” fashion, considering some set of vertices which we would like to be able to “absorb”, and reasoning about the possible incidences between edges close to these vertices in order to prove that an appropriate absorber is present. For this to be possible, one must define the notion of an absorber in a very careful way. In contrast, further developing some ideas that we introduced in [18], we are able to find absorbers using a “contraction” argument, together with one of the general tools developed by Conlon, Gowers, Samotij and Schacht [10]. This gives us an enormous amount of freedom, and in particular we can define absorbers in terms of the Dirac threshold $\mu_d(k)$ (without knowing its value!). This freedom is also crucial in allowing us to choose absorbers which exist in $H^k(n, p)$ for small p (that is, for p close to $n^{-k/2}$, which seems to be the limit of our approach).

The structure of the rest of the paper is as follows. First, in Section 2 we give an introduction to the absorbing method, and outline the proof of Theorem 1.4. Afterwards, we present a short proof of Theorem 1.2 in Section 3, as a warm-up to the absorbing method before we present the more sophisticated ideas in the proof of Theorem 1.4.

In Section 4 we discuss the so-called sparse regularity method, and in Section 5 we record some basic facts about concentration of the edge distribution in random hypergraphs. Everything in these sections will be quite familiar to experts. In Section 6 we explain how to find almost-perfect matchings in the setting of Theorem 1.4, in Section 7 we state a sparse absorbing lemma and explain how to use it to prove Theorem 1.4, and in Section 8 we present the proof of this sparse absorbing lemma.

Finally, in Section 9 we have some concluding remarks, including a discussion of how to weaken the assumption $p \geq Cn^{-k+2}$ in the case $(d, k) = (1, 3)$.

Remark 1.5 (added in proof). The general approach of defining absorbers in terms of a Dirac threshold has also appeared in earlier work by Glock, Kühn, Lo, Montgomery and Osthus [23]. We thank Stefan Glock for bringing this to our attention.

2 Outline of the proof of the main theorem

Suppose that $G \sim H^k(n, p)$ is a typical outcome of $H^k(n, p)$, and $G' \subseteq G$ is a spanning subgraph of G with minimum d -degree at least $(\mu_d(k) + \gamma)p \binom{n-d}{k-d}$. Our goal is to show that G' contains a perfect matching. Since the proof is quite involved, we break down the steps of the proof into subsections.

2.1 Almost-perfect matchings

The first observation is that our task is much simpler if we relax our goal to finding an *almost*-perfect matching (that is, a matching that covers all but $o(n)$ vertices). This is due to the existence of a powerful tool called the *sparse regularity lemma*. Roughly speaking, the sparse regularity lemma allows us to model the large-scale structure of the sparse k -graph G' using a small, dense k -graph \mathcal{R} called a *cluster k -graph*. Each edge of \mathcal{R} corresponds to a k -partite subgraph of G' where the edges are distributed in a “homogeneous” or “quasirandom” way⁵.

It is not hard to show that the degree condition on G' translates to a similar degree condition on \mathcal{R} , though small errors are introduced in the process: we can show that almost all of the d -sets of vertices in \mathcal{R} have degree at least say $(\mu_d(k) + \gamma/2) \binom{t-d}{k-d}$, where t is the number of vertices of \mathcal{R} . We then use the definition of $\mu_d(k)$ (without knowing its value!) to show that \mathcal{R} has an almost-perfect matching. This is not immediate, because \mathcal{R} may have a few d -sets of vertices with small degree, but it is possible to use a random sampling argument to overcome this difficulty. In any case, an almost-perfect matching in \mathcal{R} tells us how to partition most of the vertices of G' into subsets such that the subgraphs induced by these subsets each satisfy a certain quasirandomness condition. We can then take advantage of this

⁵Hypergraph regularity lemmas of the type we use here are sometimes known as *weak* regularity lemmas, to distinguish them from a much stronger and more complicated hypergraph regularity lemma which does not permit a description in terms of cluster k -graphs.

quasirandomness to find an almost-perfect matching in each of the subgraphs. Combining these matchings gives an almost-perfect matching in G' .

The details of this argument are in Section 6.

2.2 The absorbing method

It may not be obvious that being able to find almost-perfect matchings is actually useful, if our goal is to find a perfect matching. It is certainly not true that we can start from any almost-perfect matching and add a few edges to obtain a perfect matching. However, it turns out that something quite similar is often possible in problems of this type. Namely, in some hypergraph matching problems it is possible to find a small subset of vertices X which is very “flexible” in the sense that it can contribute to matchings in many different ways. We can then find an almost-perfect matching covering almost all the vertices outside X , and take advantage of the special properties of X to complete this into a perfect matching. This idea is now called the *absorbing method*. It was introduced as a general method by Rödl, Ruciński and Szemerédi [51, 52] (though similar ideas had appeared earlier, for example by Erdős, Gyárfás and Pyber [14] and by Krivelevich [36]). The absorbing method has been an indispensable tool for almost all work on hypergraph matching problems in the last decade.

To give a specific example, the *strong absorbing lemma* of Hán, Person and Schacht [26] (appearing here as Lemma 3.1) shows that in a very dense k -graph G we can find a small “absorbing” set of vertices X , with the special property that for any set W of $o(n)$ vertices outside X , the induced subgraph $G[X \cup W]$ has a perfect matching. So, if we can find an almost-perfect matching M_1 in $G - X$, we can take W as the set of unmatched vertices and use the special property of X to find a perfect matching M_2 in $G[X \cup W]$, giving us a perfect matching $M_1 \cup M_2$ in G .

It is much more difficult to prove absorbing lemmas in the sparse setting of Theorem 1.4. To explain why, we need to say a bit more about how absorbing lemmas are proved in the dense setting. Almost always, the idea is to build an absorbing set X using small subgraphs called *absorbers*⁶. In the context of matching problems in k -graphs, an absorber in a k -graph G rooted at a k -tuple of vertices x_1, \dots, x_k is a subgraph H whose edges can be partitioned into two matchings, one of which covers every vertex in $V(H)$ and the other of which covers every vertex except x_1, \dots, x_k . A single edge $\{x_1, \dots, x_k\}$ is a trivial absorber, and in the case $k = 2$ (that is, the case of graphs), an odd-length path between x_1 and x_2 is an absorber. See Figure 1 for a nontrivial example of a 3-uniform absorber.




Figure 1. An illustration of a 3-uniform absorber rooted on vertices x_1, x_2, x_3 . The dark edge covers all non-root vertices and the two light edges form a matching covering all the vertices of the absorber.

The details in the proofs of different absorbing lemmas vary somewhat, but a common first step is to show that there are many absorbers rooted at every k -tuple of vertices, using fairly “bare-hands” arguments that take advantage of degree assumptions. For example, suppose an n -vertex 3-graph G has $\delta_2(G) \geq (1/2 + \gamma)n$, consider any vertices x_1, x_2, x_3 , and suppose we are trying to find a copy of the absorber pictured in Figure 1. There are at least $(1/2 + \gamma)n$ choices for y_1 such that $\{x_1, x_2, y_1\} \in E(G)$. For any such y_1 , and any of the $n - 4$ remaining choices of y_2 , there are at least $(1/2 + \gamma)n$ choices for y_3 such that $\{y_1, y_2, y_3\} \in E(G)$, and at least $(1/2 + \gamma)n$ choices such that $\{x_3, y_2, y_3\} \in E(G)$, so by the inclusion-exclusion principle there are at least $2\gamma n$ choices for y_3 such that both $\{y_1, y_2, y_3\}$ and $\{x_3, y_2, y_3\}$ are in $E(G)$. All in all, this gives about $\gamma n^3 \approx 3\gamma \binom{n}{3}$ absorbers rooted at x_1, x_2, x_3 .

Having shown that every k -tuple of vertices supports many absorbers, one can then often use a straightforward probabilistic argument to construct an arrangement of absorbers that gives rise to an absorbing set X as in the strong absorbing lemma. Continuing with the previous example, if we choose a random set T of say $(\gamma/2)n$ disjoint triples of vertices, then for every choice of x_1, x_2, x_3 , there are typically about $(3\gamma)(\gamma/2)n$ triples in T which give an absorber rooted on x_1, x_2, x_3 . We can then take X

⁶The language in this field has still not been fully standardised. For example, in [26] the authors use the term “absorbing m -set” instead of “absorber”.

as the set of vertices in the triples in T . It is not hard to check that this set satisfies the assumptions of the strong absorbing lemma: given a set W of $o(n)$ vertices outside X , we can partition W into triples $\{x_1, x_2, x_3\}$ and iteratively “absorb” them into T to obtain a perfect matching in $G[X \cup W]$.

2.3 Finding absorbers in sparse graphs

Unfortunately, the ideas sketched above fail in many different ways in the sparse setting. First, there is the problem of how to actually find absorbers. It is in general very difficult to understand when one can find a copy of a specific k -graph in a subgraph G' of a random k -graph G . Indeed, this is the random Turán problem described in the introduction, and general results have become available only very recently. One of the most flexible tools in this area is the *sparse embedding lemma* proved by Conlon, Gowers, Samotij and Schacht [10] (previously, and sometimes still, known as the *KLR conjecture* of Kohayakawa, Łuczak and Rödl).

Roughly speaking, the sparse embedding lemma says that for any⁷ k -graph H , if p is large enough that a random k -graph $G \sim H^k(n, p)$ typically contains many copies of H (this depends on a “local sparseness” measure of H called *k-density*), then G satisfies the following property: If we apply the sparse regularity lemma to a spanning subgraph $G' \subseteq G$, and find a copy of H in the resulting cluster graph \mathcal{R} , then there is a corresponding copy of H in G' itself. Roughly speaking, the sparse embedding lemma allows us to work in the dense cluster graph, where it is much easier to reason directly about existence of subgraphs, and then “pull back” our findings to the original graph.

One may hope that we can just repeat the arguments in the proof of the strong absorbing lemma to find absorbers in the dense cluster graph, and somehow use the sparse embedding lemma to convert these into absorbers in the original graph. Unfortunately, life is not this simple, for (at least...) two reasons. The first issue is that we need our absorbers to satisfy some local sparseness condition, because otherwise we can only work with a very limited range of p . It is not obvious how to use existing “bare-hands” methods to find such absorbers.

The second issue is that the sparse embedding lemma is not suited for embedding *rooted* subgraphs. The cluster graph \mathcal{R} is just too rough a description of G' for it to be possible to deduce information about specific vertices in G' from information in \mathcal{R} .

To attack the first of these issues, we use a novel *non-constructive* method to find our absorbers. Namely, since an absorber is built out of matchings, we can use the definition of the Dirac threshold itself to find absorbers (even if we do not actually know its value). To be more specific, consider a k -graph G with $\delta_d(G) \geq (\mu_d(k) + \gamma) \binom{n-d}{k-d}$, and let M be a large constant. Using a concentration inequality, we can show that that almost all M -vertex induced subgraphs of G have minimum d -degree at least $(\mu_d(k) + \gamma/2) \binom{M-d}{k-d}$, so if M is large enough, then almost all M -vertex induced subgraphs have a perfect matching. We then have a lot of freedom to construct a locally sparse absorber using these matchings (specifically, we construct an absorber using an explicit locally sparse “pattern” graph).

To overcome the second of the aforementioned issues, we further develop a “contraction” technique we introduced in [18]. The problem is that the cluster graph does not “see” individual vertices; it can only see large sets of vertices. In the case $k = 2$ (that is, the graph case), an obvious fix would be to consider the set of *neighbours* (or perhaps neighbours-of-neighbours) of our desired roots, instead of the roots themselves. However, in the case $k \geq 3$, every edge containing a root vertex x_i contains $k - 1 > 1$ other vertices. That is to say, “neighbours” come grouped in sets of size $k - 1$ (the collection of all such sets is called the *link* $(k - 1)$ -*graph* of x_i). So, it seems we would need an embedding lemma that works with sets of $(k - 1)$ -sets of vertices, not just sets of vertices.

The way around this problem is to choose a large matching in the link $(k - 1)$ -graph of each x_i , and “contract” each of the edges in each of these matchings to a single vertex, to obtain a contracted graph G'_{cont} . If we do this carefully, the resulting graph can still be viewed as a subgraph of an appropriate random k -graph, so the sparse embedding lemma still applies. It then suffices to find a suitable “contracted absorber” in G'_{cont} , which would correspond to an absorber in the original k -graph G' . We can do this with the sparse embedding lemma.

The details of the arguments sketched in this section appear in Section 8.

⁷We are not being completely truthful here: strictly speaking, H must be a so-called *linear* k -graph, but this restriction turns out not to be particularly important for us.

2.4 Combining the absorbers

The above discussion gives a rough idea for how to find an absorber rooted at every k -set of vertices, in a suitable spanning subgraph of a random graph. However, it is still not at all obvious how to combine these to prove a sparse absorbing lemma. A simple probabilistic argument as sketched in Section 2.2 cannot suffice: unfortunately, there are just not enough absorbers.

We get around the issue as follows. Instead of using our absorbers to find a matching M which can “absorb” every k -set of vertices, we fix a specific “template” arrangement of only linearly many k -sets we would like to be able to absorb. It is easy to handle such a small number of k -sets: we can in fact greedily choose *disjoint* absorbers for each of these special k -sets, to obtain an “absorbing structure” H . Building on ideas due to Montgomery [42, 44], we show that it is possible to choose our template arrangement of k -sets in such a way that H has a very special kind of robust matching property: H has a “flexible set” of vertices Z such that H still has a perfect matching even after any constant fraction of the vertices in Z are deleted⁸.

We can then let $X = V(H)$, and prove that X gives a sparse absorbing lemma, as follows. For any small set W of vertices outside X , we can first find a matching M_1 covering W and a constant fraction of Z , using a hypergraph matching criterion due to Aharoni and Haxell [1]. Then, our robust matching property implies that $H - V(M_1)$ has a perfect matching M_2 , so $M_1 \cup M_2$ is a perfect matching of $G[W \cup X]$.

The details of this argument, along with the statement of our sparse absorbing lemma and the deduction of Theorem 1.4, are in Section 7.

3 Convergence of the Dirac threshold

In this section we prove Theorem 1.2, which will be a good warm-up for some of the ideas that we will develop further to prove Theorem 1.4.

The main ingredient in the proof of Theorem 1.2 is the *strong absorbing lemma* due to Han, Person and Schacht [26, Lemma 2.4] (building on ideas of Rodl, Rucinski and Szemeredi [51, 52]).

Lemma 3.1. *For any positive integers $d < k$, and any $\gamma > 0$, there is $n_0 \in \mathbb{N}$ such that for every $n > n_0$ the following holds. Suppose that G is a k -graph on n vertices with $\delta_d(G) \geq (1/2 + \gamma)\binom{n-d}{k-d}$. Then there is a set $X \subseteq V(G)$ such that*

- (i) $|X| \leq (\gamma/2)^k n$, and
- (ii) for every set $W \subseteq V(G) \setminus X$ of size at most $(\gamma/2)^{2k} n$ and divisible by k , there is a matching in G covering exactly the vertices of $X \cup W$.

Now, define

$$\tilde{\mu}_d(k) = \liminf_{n \rightarrow \infty} \frac{m_d(k, n)}{\binom{n-d}{k-d}},$$

where $n \rightarrow \infty$ along the integers n divisible by k . Our main goal is to prove the following lemma.

Lemma 3.2. *Fix positive integers $d < k$ and consider any $\gamma > 0$. Then for sufficiently large n divisible by k , every n -vertex k -graph G with $\delta_d(G) \geq (\tilde{\mu}_d(k) + \gamma)\binom{n-d}{k-d}$ has a perfect matching.*

Before we explain how to prove Lemma 3.2 we show how it implies Theorem 1.2.

Proof of Theorem 1.2, given Lemma 3.2. Lemma 3.2 implies that $m_d(k, n)/\binom{n-d}{k-d} \leq \tilde{\mu}_d(k) + \gamma$ for sufficiently large n (divisible by k), and since $\gamma > 0$ was arbitrary it follows that

$$\limsup_{n \rightarrow \infty} \frac{m_d(k, n)}{\binom{n-d}{k-d}} \leq \tilde{\mu}_d(k) = \liminf_{n \rightarrow \infty} \frac{m_d(k, n)}{\binom{n-d}{k-d}},$$

from which it follows that $m_d(k, n)/\binom{n-d}{k-d}$ converges to a limit $\mu_d(k) = \tilde{\mu}_d(k)$. □

⁸Various authors have coined different names for different ways to apply the absorbing method (though these names and their usage do not always seem to be completely consistent). The use of a flexible set Z which can optionally contribute to a desired structure is often called the *reservoir method*, where Z is called a *reservoir*. In particular, Montgomery’s approach, in which an absorbing structure is built using a template with a robust matching property, is often called *distributive absorption*, or sometimes the *absorber-template method*.

Now, our proof of Lemma 3.2 will consist of two steps. First, we prove that the conditions of Lemma 3.2 ensure an *almost-perfect* matching, then we will use the strong absorbing lemma (Lemma 3.1) to transform this into a perfect matching. The following lemma encapsulates the first of these steps.

Lemma 3.3. *Fix positive integers $d < k$ and consider any $\eta > 0$. Then for sufficiently large n , every n -vertex k -graph G with $\delta_d(G) \geq (\tilde{\mu}_d(k) + \eta) \binom{n-d}{k-d}$ has a matching covering all but $o(n)$ vertices.*

To prove Lemma 3.3 we need the following lemma showing that random subgraphs of hypergraphs typically inherit minimum-degree conditions. We state this in a slightly more general form than we need here, for later use.

Lemma 3.4. *There is $c = c(k) > 0$ such that the following holds. Consider an n -vertex k -graph G where all but $\delta \binom{n}{d}$ of the d -sets have degree at least $(\mu + \eta) \binom{n-d}{k-d}$. Let S be a uniformly random subset of $Q \geq 2d$ vertices of G . Then with probability at least $1 - \binom{Q}{d} (\delta + e^{-c\eta^2 Q})$, the random induced subgraph $G[S]$ has minimum d -degree at least $(\mu + \eta/2) \binom{Q-d}{k-d}$.*

Proof. Let $W^{(d)}$ be the collection of d -sets with degree less than $(\mu + \eta) \binom{n-d}{k-d}$ in G , and randomly order the vertices of G as v_1, \dots, v_n , so we may take $S = \{v_1, \dots, v_Q\}$. We will prove that

$$\Pr\left(\deg_S(\{v_1, \dots, v_d\}) < (\mu + \eta/2) \binom{Q-d}{k-d}\right) \leq \delta + e^{-c\eta^2 Q}$$

(where we abuse notation slightly and write $\deg_S(\{v_1, \dots, v_d\})$ for the number of edges which contain v_1, \dots, v_d and $k-d$ vertices of S). The desired result will then follow from symmetry and the union bound.

First note that the probability of the event $\{v_1, \dots, v_d\} \in W^{(d)}$ is at most δ . Now, condition on any outcome of $\{v_1, \dots, v_d\}$ which is not in $W^{(d)}$. Then $\{v_{d+1}, \dots, v_Q\}$ is a uniformly random subset of the vertices of G other than v_1, \dots, v_d , and $\mathbb{E} \deg_S(\{v_1, \dots, v_d\}) \geq (\mu + \eta) \binom{Q-d}{k-d}$. Also, making any ‘‘swap’’ to our subset $\{v_{d+1}, \dots, v_Q\}$ (that is, exchanging any element with an element outside this subset) affects $\deg_S(\{v_1, \dots, v_d\})$ by at most $\binom{Q-d-1}{k-d-1}$. So, by a concentration inequality such as [25, Corollary 2.2], conditioned on our outcome of $\{v_1, \dots, v_d\} \notin W^{(d)}$, the probability that $\{v_1, \dots, v_d\}$ has degree less than $(\mu + \eta/2) \binom{Q-d}{k-d}$ in $G[S]$ is at most

$$2 \exp\left(-\frac{2\left(\frac{\eta}{2} \binom{Q-d}{k-d}\right)^2}{(Q-d) \binom{Q-d-1}{k-d-1}^2}\right) \leq e^{-c\eta^2 Q}$$

(for say $c = 1/(4k^2)$, recalling that $Q \geq 2d$), as claimed. \square

Now we prove Lemma 3.3.

Proof of Lemma 3.3. Choose large Q , divisible by k , such that $m_d(k, Q) / \binom{Q-d}{k-d} \leq \tilde{\mu}_d(k) + \eta/2$ (this is possible by the definition of $\tilde{\mu}_d(k)$). Let $\lambda = \binom{Q}{d} e^{-c\eta^2 Q}$ be as in Lemma 3.4 (taking $\delta = 0$), and note that we can make λ arbitrarily small by making Q large.

Now, we randomly partition the vertex set into n/Q subsets of size Q . By Lemma 3.4, with positive probability all but a λ -fraction of the subsets have minimum degree at least $(\tilde{\mu}_d(k) + \eta/2) \binom{Q-d}{k-d}$. By our choice of Q , each of these Q -vertex subsets $S \subseteq V(G)$ has the property that $G[S]$ has a perfect matching, and we can combine these to find a matching covering all but λn vertices. Since λ could have been arbitrarily small, this implies that we can find a matching covering all but $o(n)$ vertices. \square

Now, it is straightforward to deduce Lemma 3.2 from Lemma 3.1 and Lemma 3.3, concluding our proof of Theorem 1.2.

Proof of Lemma 3.2. From the discussion in the introduction, note that $\tilde{\mu}_d(k) \geq 1/2$, so the assumptions in Lemma 3.1 are satisfied and we can find an ‘‘absorbing set’’ $X \subseteq V(G)$ of at most $(\gamma/2)^k n$ vertices. Let $n' = n - |X|$ and observe that since X is so small, we have $\delta_d(G - X) \geq (\tilde{\mu}_d(k) + \gamma/2) \binom{n'-d}{k-d}$. By Lemma 3.3, it follows that $G - X$ has a matching covering all vertices except a set W of size $o(n)$. By the defining property of the absorbing set X , it follows that G has a perfect matching. \square

4 The sparse regularity method

The proof of Theorem 1.4 makes heavy use of the sparse regularity method. So, we will need a sparse version of a hypergraph regularity lemma. There is a general hypergraph regularity lemma which is quite complicated to state and prove (see [24, 53]), but we will only need (a sparse version of) the so-called “weak” hypergraph regularity lemma (see [35]). Weak hypergraph regularity lemmas are suitable for embedding *linear* hypergraphs, which are hypergraphs in which no pair of edges share more than one vertex.

We remark that this section closely mirrors [18, Section 4], though some of the lemma statements are slightly more general.

To state our sparse hypergraph regularity lemma we first need to make some basic definitions.

Definition 4.1. Let $\varepsilon, \eta > 0$, $D > 1$ and $0 \leq p \leq 1$ be arbitrary parameters.

- **Density:** Consider disjoint vertex sets X_1, \dots, X_k in a k -graph G . Let $e(X_1, \dots, X_k)$ be the number of edges with a vertex in each X_i . Let

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

be the *density* between X_1, \dots, X_k .

- **Regular tuples:** A k -partite k -graph with parts V_1, \dots, V_k is (ε, p) -regular if, for every $X_1 \subseteq V_1, \dots, X_k \subseteq V_k$ with $|X_i| \geq \varepsilon|V_i|$, the density $d(X_1, \dots, X_k)$ of edges between X_1, \dots, X_k satisfies

$$|d(X_1, \dots, X_k) - d(V_1, \dots, V_k)| \leq \varepsilon p.$$

- **Regular partitions:** A partition of the vertex set of a k -graph into t parts V_1, \dots, V_t is said to be (ε, p) -regular if it is an equipartition (meaning that the sizes of the parts differ by at most one), and for all but at most $\varepsilon \binom{t}{k}$ of the k -sets $\{V_{i_1}, \dots, V_{i_k}\}$, the induced k -partite k -graph between V_{i_1}, \dots, V_{i_k} is (ε, p) -regular.
- **Upper-uniformity:** A k -graph G is (η, p, D) -upper-uniform if for any choice of disjoint subsets X_1, \dots, X_k with $|X_1|, \dots, |X_k| \geq \eta|V(G)|$, we have $d(X_1, \dots, X_k) \leq Dp$.

Now, our sparse weak hypergraph regularity lemma is as follows. We omit its proof since it is straightforward to adapt a proof of the sparse graph regularity lemma (see [34] for a sparse regularity lemma for graphs, and see [35, Theorem 9] for a weak regularity lemma for dense hypergraphs).

Lemma 4.2. *For every $\varepsilon, D > 0$ and every positive integer t_0 , there exist $\eta > 0$ and $T \in \mathbb{N}$ such that for every $p \in [0, 1]$, every (η, p, D) -upper-uniform k -graph G with at least t_0 vertices admits an (ε, p) -regular partition V_1, \dots, V_t of its vertex set into $t_0 \leq t \leq T$ parts.*

For us, the most crucial aspect of the sparse regularity lemma is that it can be used to give a rough description of a sparse k -graph in terms of a dense *cluster k -graph* which we now define.

Definition 4.3. Given an (ε, p) -regular partition V_1, \dots, V_t of the vertex set of a k -graph G , the *cluster hypergraph* is the k -graph whose vertices are the clusters V_1, \dots, V_t , with an edge $\{V_{i_1}, \dots, V_{i_k}\}$ if $d(V_{i_1}, \dots, V_{i_k}) > 2\varepsilon p$ and the induced k -partite k -graph between V_{i_1}, \dots, V_{i_k} is (ε, p) -regular.

If the sparse regularity lemma is applied with small ε and large t_0 , the cluster hypergraph approximately inherits minimum degree properties from the original graph G , as follows.

Lemma 4.4. *Fix positive integers $d < k$, $0 < \delta < 1$, some sufficiently small $\varepsilon > 0$ and some sufficiently large $t_0 \in \mathbb{N}$, and let G be an n -vertex $(o(1), p, 1 + o(1))$ -upper-uniform k -graph (in particular, we assume that n is sufficiently large). Let $G' \subseteq G$ be a spanning subgraph in which all but $o(n^d)$ of the d -sets of vertices have degree at least $\delta \binom{n-d}{k-d} p$. Let \mathcal{R} be the t -vertex cluster k -graph obtained by applying the sparse regularity lemma to G' with parameters t_0, p and ε . Then all but at most $\sqrt{\varepsilon} \binom{t}{d}$ of the d -sets of vertices of \mathcal{R} have degree at least $\delta \binom{t-d}{k-d} - (4\sqrt{\varepsilon} + k/t_0)t^{k-d}$.*

Lemma 4.4 can be proved with a standard counting argument. It is a special case of Lemma 4.7, which we will state and prove in the next subsection.

4.1 Refining an existing partition

We will need to apply the sparse regularity lemma to a k -graph whose vertices are already partitioned into a few different parts with different roles. It will be important that the regular partition in Lemma 4.2 can be chosen to be consistent with this existing partition.

Lemma 4.5. *Suppose that a k -graph G has its vertices partitioned into sets P_1, \dots, P_h . In the (ε, p) -regular partition guaranteed by Lemma 4.2, we can assume that all but at most εht of the clusters V_i are contained in some P_j .*

For the reader who is familiar with the proof of the regularity lemma, the proof of Lemma 4.5 is straightforward. Indeed, in order to prove the regularity lemma, one starts with an arbitrary partition and iteratively refines it. Therefore, one can start with the partition (P_1, \dots, P_h) and proceed in the usual way. For more details, see the reduction in [18, Lemma 4.6].

Next, we state a more technical version of Lemma 4.4, deducing degree conditions in the cluster graph from degree conditions between the P_i . To state this we first need to generalise the definition of a cluster graph.

Definition 4.6. Consider a k -graph G , and let P_1, \dots, P_h be disjoint sets of vertices. Also, consider a (ε, p) -regular partition V_1, \dots, V_t of the vertices of G . Then the *partitioned cluster graph* \mathcal{R} with *threshold* τ is the k -graph defined as follows. The vertices of \mathcal{R} are the clusters V_i which are completely contained in some P_j , with an edge $\{V_{i_1}, \dots, V_{i_k}\}$ if $d(V_{i_1}, \dots, V_{i_k}) > \tau p$ and the induced k -partite k -graph between V_{i_1}, \dots, V_{i_k} is (ε, p) -regular.

Lemma 4.7. *Let $d < k$ be positive integers, let $\varepsilon > 0$ be sufficiently small, and let $t_0 \in \mathbb{N}$ be sufficiently large. Let G be an n -vertex $(o(1), p, 1 + o(1))$ -upper-uniform k -graph with a partition P_1, \dots, P_h of its vertices into parts of sizes n_1, \dots, n_h , respectively. Let $G' \subseteq G$ be a spanning subgraph and let \mathcal{R} be the partitioned cluster k -graph with threshold τ obtained by applying Lemmas 4.2 and 4.5 to G' with parameters t_0, p and ε .*

For every $1 \leq i \leq h$, let \mathcal{P}_i be the set of clusters contained in P_i , and let $t_i = |\mathcal{P}_i|$. Also, for every $J \subseteq \{1, \dots, h\}$, write $P_J := \bigcup_{j \in J} P_j$, $n_J = |P_J|$, $\mathcal{P}_J = \bigcup_{j \in J} \mathcal{P}_j$ and $t_J = |\mathcal{P}_J|$. Then the following properties hold.

(1) *Each $t_i \geq (n_i/n)t - \varepsilon ht$.*

(2) *Consider some $i \leq h$ and some $J \subseteq \{1, \dots, h\}$, and suppose that all but $o(n^d)$ of the d -sets of vertices $X \subset P_i$ satisfy*

$$\deg_{P_J}(X) \geq \delta p \binom{n_J - d}{k - d}.$$

Then, all but at most $\sqrt{\varepsilon} \binom{t}{d}$ of the d -sets of clusters $\mathcal{X} \subset \mathcal{P}_i$ have

$$\deg_{\mathcal{P}_J}(\mathcal{X}) \geq \delta \binom{t_J - d}{k - d} - (\tau + \varepsilon h + \sqrt{\varepsilon} + k/t_0) t^{k-d}$$

in the cluster graph \mathcal{R} .

Note that Lemma 4.4 is actually a special case of Lemma 4.7 (taking $h = 1$ and threshold $\tau = 2\varepsilon$).

Proof. The clusters in \mathcal{P}_i comprise at most $t_i(n/t)$ vertices, so recalling the statement of Lemmas 4.2 and 4.5, we have $|P_i| = n_i \leq t_i(n/t) + \varepsilon ht(n/t)$. It follows that $t_i \geq (n_i/n)t - \varepsilon ht$, proving (1).

Now we prove (2). Let $\mathcal{W}^{(d)}$ be the collection of all d -sets of clusters $\mathcal{W} = \{W_1, \dots, W_d\}$ that are contained in more than $\sqrt{\varepsilon} \binom{t-d}{k-d}$ irregular (that is, non- (ε, p) -regular) k -sets $\{W_1, \dots, W_d, V_{i_1}, \dots, V_{i_{k-d}}\}$. Then

$$|\mathcal{W}^{(d)}| \leq \frac{\varepsilon \binom{t}{k}}{\sqrt{\varepsilon} \binom{t-d}{k-d}} \leq \sqrt{\varepsilon} \binom{t}{d}.$$

Now, consider any $i \leq h$ and $J \subseteq \{1, \dots, h\}$, and suppose that the condition in (2) holds. Consider any d -set \mathcal{X} of clusters in $\binom{\mathcal{P}_i}{d} \setminus \mathcal{W}^{(d)}$. We wish to estimate $\deg_{\mathcal{P}_J}(\mathcal{X})$.

Let E be the set of edges of G' which have a vertex in each of the d clusters in \mathcal{X} , and $k - d$ vertices in P_J . In order to estimate $\deg_{\mathcal{P}_J}(\mathcal{X})$, we count $|E|$ in two different ways. First, we break E into subsets depending on how they relate to the cluster graph \mathcal{R} .

Let $E_{\mathcal{R}}$ be the set of $e \in E$ which “arise from \mathcal{R} ” in the sense that the vertices of e come from distinct clusters that form an edge in \mathcal{R} containing \mathcal{X} . By upper-uniformity, we have

$$|E_{\mathcal{R}}| \leq (1 + o(1)) \deg_{\mathcal{P}_J}(\mathcal{X}) p(n/t)^k.$$

Let E_{irr} be the subset of edges in E arising from irregular k -sets. By the choice of $\mathcal{X} \in \binom{\mathcal{P}_i}{d} \setminus \mathcal{W}^{(d)}$, we have

$$|E_{\text{irr}}| \leq \sqrt{\varepsilon} \binom{t-d}{k-d} (1 + o(1)) p(n/t)^k.$$

Let E_{τ} be the subset of edges in E arising from k -sets of clusters (containing \mathcal{X}) that are regular but whose density is less than τ (and therefore do not appear in the cluster graph). We have

$$|E_{\tau}| \leq \tau \binom{t-d}{k-d} (1 + o(1)) p(n/t)^k.$$

Let E_{mul} be the set of edges in E which have multiple vertices in the same cluster of \mathcal{P}_J , and let E_Z be the set of edges in E which involve a vertex not in a cluster in \mathcal{P}_J (because its cluster was not completely contained in any P_j). Simple double-counting arguments give

$$|E_{\text{mul}}| \leq \frac{\binom{k-d}{2} t^{k-d-1}}{(k-d)!} (1 + o(1)) p(n/t)^k, \quad |E_Z| \leq \frac{(k-d)(\varepsilon h t) t^{k-d-1}}{(k-d)!} (1 + o(1)) p(n/t)^k.$$

All in all, we obtain

$$\begin{aligned} |E| &\leq |E_{\mathcal{R}}| + |E_{\text{irr}}| + |E_{\tau}| + |E_{\text{mul}}| + |E_Z| \\ &\leq (1 + o(1)) p \left(\frac{n}{t} \right)^k \left(\deg_{\mathcal{P}_J}(\mathcal{X}) + t^{k-d} (\sqrt{\varepsilon} + \tau + (k-d)/t_0 + \varepsilon h) \right). \end{aligned}$$

On the other hand, by the degree assumption in G' and the fact that $n_J \geq t_J(n/t)$, we have

$$|E| \geq \left(\left(\frac{n}{t} \right)^d - o(n^d) \right) \delta p \binom{n_J - d}{k-d} \geq (1 - o(1)) p \delta \binom{t_J - d}{k-d} \left(\frac{n}{t} \right)^k.$$

It follows that

$$\deg_{\mathcal{P}_J}(\mathcal{X}) \geq \delta \binom{t_J - d}{k-d} - (\sqrt{\varepsilon} + \tau + k/t_0 + \varepsilon h) t^{k-d}. \quad \square$$

4.2 A sparse embedding lemma

One of the most powerful aspects of the sparse regularity method is that, for a subgraph G' of a typical outcome of a random graph, if we find a substructure in the cluster graph (which is usually dense, therefore comparatively easy to analyse), then a corresponding structure must also exist in the original graph G' . For graphs, this was famously conjectured to be true by Kohayakawa, Łuczak and Rödl [33], and was proved by Conlon, Gowers, Samotij and Schacht [10]. We will need a generalisation to hypergraphs, which was already observed to hold in [10] and appears explicitly as [18, Theorem 4.12]. To state it we will need some definitions.

Definition 4.8. Consider a k -graph H with vertex set $\{1, \dots, r\}$ and let $\mathcal{G}(H, n, m, p, \varepsilon)$ be the collection of all k -graphs G obtained in the following way. The vertex set of G is a disjoint union $V_1 \cup \dots \cup V_r$ of sets of size n . For each edge $\{i_1, \dots, i_k\} \in E(H)$, we add to G an (ε, p) -regular k -graph with m edges between V_{i_1}, \dots, V_{i_k} . These are the only edges of G .

Definition 4.9. For $G \in \mathcal{G}(H, n, m, p, \varepsilon)$, let $\#_H(G)$ be the number of “canonical copies” of H in G , meaning that the copy of the vertex i must come from V_i .

Definition 4.10. The k -density $m_k(H)$ of a k -graph H is defined as

$$m_k(H) = \max \left\{ \frac{e(H') - 1}{v(H') - k} : H' \subseteq H \text{ with } v(H') > k \right\}.$$

Now, our sparse embedding lemma is as follows.

Theorem 4.11. *For every linear k -graph H and every $\tau > 0$, there exist $\varepsilon, \zeta > 0$ with the following property. For every $\kappa > 0$, there is $C > 0$ such that if $p \geq CN^{-1/m_k(H)}$, then with probability $1 - e^{-\Omega(N^k p)}$ the following holds in $G \in \mathbb{H}^k(N, p)$. For every $n \geq \kappa N$, $m \geq \tau pn^k$ and every subgraph G' of G in $\mathcal{G}(H, n, m, p, \varepsilon)$, we have $\#_H(G') > \zeta p^{e(H)} n^{v(H)}$.*

Note that the condition $m \geq \tau pn^k$ is precisely the condition that the corresponding $G' \in (H, n, m, p, \varepsilon)$ has density at least τ .

Theorem 4.11 may be proved using the methods in [10]. The necessary adaptations for the hypergraph setting are described in detail in [18, Appendix A].

5 Concentration lemmas

In this section we collect a number of basic facts about concentration of the edge distribution in random hypergraphs and random subsets of hypergraphs. First, we show that the upper-uniformity condition in the sparse regularity lemma is almost always satisfied in random hypergraphs.

Lemma 5.1. *Fix $k \in \mathbb{N}$, $D > 1$ and $0 < \eta < 1$, and consider $G \in \mathbb{H}^k(n, p)$. Then G is (η, p, D) -upper-uniform with probability at least $1 - 2^{kn} e^{-\Omega(n^k p)}$.*

Proof. Consider disjoint vertex sets X_1, \dots, X_k each having size at least ηn . Then $\mathbb{E}e(X_1, \dots, X_k) = p|X_1| \dots |X_k| = \Omega(n^k p)$, so by the Chernoff bound,

$$\Pr(e(X_1, \dots, X_k) \geq Dp|X_1| \dots |X_k|) = \exp(-\Omega((D-1)^2 n^k p)) = e^{-\Omega(n^k p)}.$$

We can then take the union bound over all choices of X_1, \dots, X_k . □

The following corollary is immediate, and will be more convenient in practice.

Corollary 5.2. *Fix $k \in \mathbb{N}$, let $p = \omega(n^{1-k} \log n)$, and consider $G \in \mathbb{H}^k(n, p)$. Then G is $(o(1), p, 1 + o(1))$ -upper-uniform with probability at least $1 - e^{-\omega(n \log n)}$.*

Next, the following lemma shows that in random hypergraphs all vertices have about the expected degree into any large enough set.

Lemma 5.3. *Fix $\lambda > 0$ and $k \geq 3$. Then there is $C > 0$ such that if $p \geq Cn^{2-k}$, a.a.s. $G \sim \mathbb{H}^k(n, p)$ has the following property. For every vertex w and every set $X \subseteq V(G)$ of at most λn vertices, there are at most $2(\lambda n)p \binom{n-2}{k-2}$ edges in G containing w and a vertex of X .*

Proof. For any vertex w and any set X of at most λn vertices, the expected number of edges containing both w and a vertex of X is at most $(\lambda n)p \binom{n-2}{k-2}$. Then the desired result follows from the Chernoff bound and the union bound over at most 2^n choices of X . □

We also need a slightly more sophisticated version of Lemma 5.3 that works for general d -degrees and has quite a weak assumption on p , but only gives a conclusion for almost all d -sets of vertices. For vertex sets S, X in a k -graph G , let $Z_G(S, X)$ be the number of edges $e \in E(G)$ that contain S and have non-empty intersection with X .

Lemma 5.4. *Fix $\lambda > 0$ and positive integers $d < k$. For $p \geq \omega(n^{d-k})$, a.a.s. $G \sim \mathbb{H}^k(n, p)$ has the following property. For every subset $X \subseteq V(G)$ of size $|X| \leq \lambda n$, there are at most $o(n^d)$ different d -sets $S \subseteq V(G)$ such that $Z_G(S, X) > 2 \binom{k}{d} (\lambda n) p \binom{n-d-1}{k-d-1}$.*

Proof. Let $\binom{[k]}{d}$ be the collection of all d -subsets of $\{1, \dots, k\}$. Consider the collection of all k -sets of vertices (i.e., all possible edges in G). For each such k -set e , let $\binom{e}{d}$ be the collection of d -subsets of e , and fix a (bijective) labelling $\phi_e : \binom{e}{d} \rightarrow \binom{[k]}{d}$. For each d -subset $I \subseteq \{1, \dots, k\}$ and set S of d vertices, let $Z_G^I(S, X)$ be the number of edges $e \in E(G)$ which include S , have non-empty intersection with X , and have $\phi_e(S) = I$.

Fix a vertex set X of size at most λn , and let \mathcal{E}_S^I be the event that $Z_G^I(S, X) > 2(\lambda n)p \binom{n-d-1}{k-d-1}$. By the Chernoff bound applied to $Z_G^I(S, X)$, we have $\Pr(\mathcal{E}_S^I) = o(1)$ for each S, I . For fixed I , each of the events \mathcal{E}_S^I (ranging over different S) are independent from each other, so, again using a Chernoff bound, with probability $1 - e^{-\omega(n^d)}$, all but $o(n^d)$ of the events \mathcal{E}_S^I are satisfied. Taking a union bound over $O(1)$ choices of I shows that with probability $1 - e^{-\omega(n^d)}$, for all but $o(n^d)$ of the sets of d vertices S , all the events of the form \mathcal{E}_S^I hold, meaning that $Z_G(S, X) \leq 2 \binom{k}{d} (\lambda n) p \binom{n-d-1}{k-d-1}$. The desired result then follows from a union bound over choices of X . □

We will also need the following lemma, showing that if we consider a high-degree spanning subgraph of a typical outcome of a random hypergraph, then random subsets are likely to inherit minimum degree properties.

Lemma 5.5. *Fix positive integers $d < k$, and fix any $0 \leq \mu \leq 1$ and $0 < \gamma, \sigma \leq 1$. Then for any $p \geq n^{d-k} \log^3 n$, a.a.s $G \sim \mathbb{H}^k(n, p)$ has the following property. Consider a spanning subgraph $G' \subseteq G$ with minimum d -degree at least $(\mu + \gamma)p \binom{n-d}{k-d}$, and let Y be a uniformly random subset of σn vertices of G' . Then a.a.s. every d -set of vertices has degree at least $(\mu + \gamma/2)p \binom{\sigma n-d}{k-d}$ into Y .*

Proof. For a d -set of vertices A and $1 \leq t \leq k-d$, define $p_t(A)$ to be the number of pairs of edges $(e, f) \in G$ such that e and f both include A , and $|e \cup f| = 2(k-d) - t$, and define the random variable

$$\bar{\Delta}_A = \sum_{t=1}^{k-d} p_t(A) \sigma^{2(k-d)-t}.$$

For each $(k-d)$ -set of vertices $e \subseteq V(G)$, let ξ_e be the indicator random variable for the event $e \cup A \in E(G)$. Note that $\bar{\Delta}_A$ is a quadratic polynomial in the ξ_e , and has expectation $E_0 := \mathbb{E} \bar{\Delta}_A = O(n^{2(k-d)-1} p^2)$. For any $(k-d)$ -sets of vertices $e, f \subseteq V(G)$, we compute the expected partial derivatives

$$\mathbb{E} \frac{\partial \bar{\Delta}_A}{\partial \xi_e} = O(n^{k-d-1} p), \quad \mathbb{E} \frac{\partial^2 \bar{\Delta}_A}{\partial \xi_e \partial \xi_f} = O(1).$$

Let E_1 be the common value of the $\mathbb{E}[\partial \bar{\Delta}_A / \partial \xi_e]$ and let E_2 be the maximum value of the $\mathbb{E}[\partial^2 \bar{\Delta}_A / \partial \xi_e \partial \xi_f]$. Let $E_{\geq 0} = \max\{E_0, E_1, E_2\}$ and let $E_{\geq 1} = \max\{E_1, E_2\}$. Given our assumption $p \geq n^{d-k}$ we have

$$E_{\geq 0} = O\left(1 + n^{2(k-d)-1} p^2\right), \quad E_{\geq 1} = O(1 + n^{k-d-1} p).$$

Applying a Kim–Vu-type polynomial concentration inequality (for example [56, Theorem 1.36], with $\lambda = \Omega((t/\sqrt{E_{\geq 0} E_{\geq 1}})^{1/(2-1/2)})$), we see that for all $t \geq 0$,

$$\Pr(\bar{\Delta}_A \geq E_0 + t) \leq \exp\left(-C_2 \left(\frac{t}{\sqrt{E_{\geq 0} E_{\geq 1}}}\right)^{1/(2-1/2)} + (2-1) \log \binom{n}{k-d}\right)$$

for some constant $C_2 > 0$. It follows that

$$\bar{\Delta}_A \leq \left(1 + n^{2(k-d)-1} p^2\right) \log^3 n \tag{1}$$

with probability $1 - n^{-\omega(1)}$. By a union bound, a.a.s. this holds for all d -sets A . This is the only property of G we require.

So, fix an outcome of G satisfying (1) for all A , and fix a spanning subgraph $G' \subseteq G$ with minimum d -degree at least $(\mu + \gamma)p \binom{n-d}{k-d}$. Instead of considering a uniformly random set of σn vertices, we consider a closely related “binomial” random vertex subset Y' obtained by including each vertex with probability $\sigma' := \sigma - n^{-2/3}$ independently. This suffices because Y' can a.a.s. be coupled as a subset of a uniformly random set of size σn (note that the standard deviation of the size of Y' is $O(\sqrt{n})$).

Let A be a d -set of vertices, and let $\Gamma(A)$ be the link $(k-d)$ -graph of A with respect to G' . Then $\deg_{Y'}(A)$ is the number of $e \in \Gamma(A)$ that are subsets of Y' , so it has expected value $|\Gamma(A)|(\sigma')^{k-d} \geq (\mu + \gamma)p \binom{\sigma' n-d}{k-d}$. By Janson’s inequality (see [29, Theorem 2.14]), our assumption on p , and (1), we have

$$\Pr\left(\deg_{Y'}(A) \leq (\mu + \gamma/2)p \binom{\sigma' n-d}{k-d}\right) \leq \exp\left(-\frac{\left((\gamma/2)p \binom{\sigma' n-d}{k-d}\right)^2}{2\bar{\Delta}_A}\right) = o(n^{-d}).$$

Then, take the union bound over all A . □

Finally, we also need the following “almost-all” version of Lemma 5.5.

Lemma 5.6. *Fix positive integers $d < k$, and fix any $0 \leq \mu \leq 1$ and $0 < \gamma, \sigma \leq 1$. Then for any $p = \omega(n^{d-k})$, a.a.s $G \sim \mathbb{H}^k(n, p)$ has the following property. Consider a spanning subgraph $G' \subseteq G$ such that all but at most $o(n^d)$ of the d -sets of vertices have degree at least $(\mu + \gamma)p \binom{n-d}{k-d}$, and let Y be a uniformly random subset of σn vertices of G' . Then a.a.s. all but $o(n^d)$ of the d -sets of vertices in G' have degree at least $(\mu + \gamma/2)p \binom{\sigma n-d}{k-d}$ into Y .*

Proof. Let $f = (n^{d-k}/p)^{1/4} = o(1)$. First, recall the definition of $\bar{\Delta}_A$ from Lemma 5.5. By Markov's inequality, with probability at least $1 - f = 1 - o(1)$ all but at most $fn^d = o(n^d)$ of the d -sets of vertices A in G' have $\bar{\Delta}_A \leq \mathbb{E}\bar{\Delta}_A/f = O(n^{2(k-d)-1}p^2/f)$ and have degree at least $(\mu + \gamma)p\binom{n-d}{k-d}$ (in which case say A is *good*). So, fix an outcome of G satisfying this property.

Now, let A be a good d -set of vertices, and let Y' be as in Lemma 5.5. By the same calculation as in Lemma 5.5 (using Janson's inequality), we have

$$\Pr\left(\deg_{Y'}(A) \leq (\mu + \gamma/2)p\binom{\sigma n - d}{k - d}\right) = o(1),$$

so by Markov's inequality, a.a.s. only $o(n^d)$ good d -sets fail to satisfy the degree condition (and there are only $o(n^d)$ non-good d -sets). \square

6 Almost-perfect matchings

The easier part of the proof of Theorem 1.4 is to show that a degree condition implies the existence of *almost-perfect* matchings, as follows.

Lemma 6.1. *Fix positive integers $d < k$ and consider any $\gamma > 0$. Suppose $p = \omega(n^{1-k} \log n)$. Then, with probability $1 - e^{-\Omega(pn^k)}$, $G \sim \mathbb{H}^k(n, p)$ has the following property. In every spanning subgraph $G' \subseteq G$ for which all but $o(n^d)$ of the d -sets of vertices have degree at least $(\mu_d(k) + \gamma)p\binom{n-d}{k-d}$, there is a matching covering all but $o(n)$ vertices.*

We will prove Lemma 6.1 with the sparse regularity lemma. It will be important that an almost-perfect matching in the cluster hypergraph \mathcal{R} can be translated to an almost-perfect matching in the original hypergraph G . This will be deduced from the following lemma.

Lemma 6.2. *Consider an (ε, p) -regular k -partite k -graph G with parts V_1, \dots, V_k of the same size m , and density at least $2\varepsilon p$. Then G has a matching of size at least $(1 - \varepsilon)m$.*

Proof. By the definition of (ε, p) -regularity, for any choice of $V'_1 \subseteq V_1, \dots, V'_k \subseteq V_k$, satisfying $|V'_i| \geq \varepsilon m$ for all i , we have

$$|d(V'_1, \dots, V'_k) - d(V_1, \dots, V_k)| \leq \varepsilon p,$$

so $d(V'_1, \dots, V'_k) \geq \varepsilon p$, meaning that the sets V'_1, \dots, V'_k span at least one edge.

Now, suppose for the purpose of contradiction that there is no matching of size $(1 - \varepsilon)m$ in G , and consider a maximum matching M . Since $|M| \leq (1 - \varepsilon)m$, the set of uncovered vertices V'_i in each V_i has size at least εm , so V'_1, \dots, V'_k span at least one edge, by the above discussion. But then this edge can be used to extend M to a larger matching, contradicting maximality. \square

Now we prove Lemma 6.1.

Proof of Lemma 6.1. Choose large Q , divisible by k , such that $m_d(k, Q)/\binom{Q-d}{k-d} \leq \mu_d(k) + \gamma/4$. Also, choose some small $\varepsilon > 0$ and let $\lambda = \binom{Q}{d}(\sqrt{\varepsilon} + e^{-cn^2Q})$ be as in Lemma 3.4 (taking $\eta = \gamma/2$ and $\delta = \sqrt{\varepsilon}$). Choosing large Q , and ε small relative to Q , we can make λ arbitrarily small.

Since $p = \omega(n^{1-k} \log n)$, by Corollary 5.2 G is a.a.s. $(o(1), p, 1 + o(1))$ -upper-uniform. So, by Lemma 4.4, if we apply the sparse regularity lemma (Lemma 4.2) to G' with small ε and large t_0 , we obtain a cluster k -graph \mathcal{R} such that all but $\sqrt{\varepsilon}\binom{t}{d}$ of the d -sets of clusters have degree at least $(\mu_d(k) + \gamma/2)\binom{n-d}{k-d}$.

Now, we randomly partition the t clusters into t/Q subsets of size Q . By Lemma 3.4, with positive probability all but a λ -fraction of the subsets have minimum degree at least $(\mu_d(k) + \gamma/4)\binom{Q-d}{k-d}$. By our choice of Q , each of these Q -subsets of clusters \mathcal{S} has the property that $\mathcal{R}[\mathcal{S}]$ has a perfect matching, and we can combine these to find a matching covering all but λt of the t vertices of the cluster graph.

Each edge of this matching corresponds to a k -tuple of clusters $(V_{i_1}, \dots, V_{i_k})$ in which we can find a matching with $(1 - \varepsilon)(n/t)$ vertices, by Lemma 6.2. We can combine these matchings to get a matching in G' covering at least $(1 - \varepsilon - \lambda)n$ vertices. Since ε and λ could have been arbitrarily small, this implies that we can find a matching covering all but $o(n)$ vertices. \square

7 Sparse absorption

The most challenging part of the proof of Theorem 1.4 is to prove a suitable sparse analogue of the strong absorbing lemma. Specifically, the lemma we will prove is as follows.

Lemma 7.1. *Fix positive integers $d < k$, and some $\gamma > 0$. There are $\lambda, C > 0$ such that the following holds. For p satisfying $p \geq \max\{n^{-k/2+\gamma}, Cn^{2-k}\}$, a.a.s. $G \sim \mathbb{H}^k(n, p)$ has the following property. For any spanning subgraph G' of G with minimum d -degree at least $(\mu_d(k) + \gamma)p \binom{n-d}{k-d}$, there is a set $X \subseteq V(G')$ such that*

- (i) $|X| \leq (\gamma/2)^k n$, and
- (ii) for every set $W \subseteq V(G) \setminus X$ of at most λn vertices, there is a matching in G' covering exactly the vertices of $X \cup W$ (provided $|X \cup W|$ is divisible by k).

Most of the rest of the paper will be devoted to proving Lemma 7.1, but first we give the simple deduction of Theorem 1.4 from Lemmas 5.4, 6.1 and 7.1.

Proof of Theorem 1.4. We may assume that γ is small relative to d and k (the lemma statement only gets stronger as we decrease γ). Let $G \sim \mathbb{H}^k(n, p)$ and let $G' \subseteq G$ be a spanning subgraph with $\delta_d(G') \geq (\mu_d(k) + \gamma)p \binom{n-d}{k-d}$. By Lemma 7.1, a.a.s. G' has an absorbing subset X of size at most $(\gamma/2)^k n$. Let $n' = n - |X|$ and observe that since X is so small, we have by Lemma 5.4 that all but $o(n^d)$ d -subsets in $G' - X$ have degree at least $(\mu_d(k) + \gamma/2) \binom{n'-d}{k-d} p$ in $G' - X$.

By Lemma 6.1, it follows that $G' - X$ has a matching covering all vertices of $V(G') \setminus X$ except a set W of size $o(n)$. By the defining property of the absorbing set X , it follows that G' has a perfect matching. \square

The crucial idea for the proof of Lemma 7.1 is to find many small subgraphs called *absorbers*, which can each contribute to a matching in two different ways.

Definition 7.2 (absorbers). An *absorber* rooted on a k -tuple of vertices (x_1, \dots, x_k) is a k -graph on some set of vertices including x_1, \dots, x_k , whose edges can be partitioned into:

- a perfect matching, in which each of x_1, \dots, x_k are in a unique edge (the *covering matching*), and;
- a matching covering all vertices except x_1, \dots, x_k (the *non-covering matching*).

We call x_1, \dots, x_k the *rooted vertices*, and the k edges containing each of x_1, \dots, x_k are called *rooted edges*. The *order* of the absorber is its number of vertices other than x_1, \dots, x_k .

Absorbers are the basic building blocks for a larger “absorbing structure”, whose vertex set we will take as the set X in Lemma 7.1. The relative positions of the absorbers in this structure will be determined by a “template” with a “resilient matching” property as will be described in the next few lemmas.

Lemma 7.3. *For fixed $k \in \mathbb{N}$ there is $L > 0$ such that the following holds. For any sufficiently large r , there exists a k -graph T with at most Lr vertices, at most Lr edges, and an identified set Z of r vertices, such that if we remove fewer than $r/2$ vertices from Z , the resulting hypergraph has a perfect matching (provided its number of vertices is divisible by k). We call T a resilient template and we call Z its flexible set. We say r is the order of the resilient template.*

We defer the proof of Lemma 7.3 to Section 7.1. It is a simple reduction from a random graph construction due to Montgomery [42, 44]. We will want to arrange absorbers in the positions prescribed by a resilient template, as follows.

Definition 7.4. An (r, Q) -*absorbing structure* is a k -graph H of the following form. Consider an order- r resilient template T and put externally vertex-disjoint absorbers of order at most Q on each edge of T (that is to say, the absorbers intersect only at their root vertices). We stress that the edges of T are not actually present in H , they just describe the relative positions of the absorbers. See Figure 2.

An absorbing structure H has the same crucial property as the resilient template T that defines it: if we remove fewer than half of the vertices of the flexible set Z then what remains of H has a perfect matching. Indeed, after this removal we can find a perfect matching M of T , then our perfect matching of H can be comprised of the covering matching of the absorber on each edge of M and the non-covering matching for the absorber on each other edge of T .




Figure 2. A cartoon of an absorbing structure. The dashed bubbles indicate template edges, which are not actually a part of the absorbing structure.

We will want to find an absorbing structure H whose flexible set Z has a certain “richness” property: essentially, we will want Z to have the property that for all small sets W disjoint from H , there is a matching covering W and a small portion of Z (and no other vertices). Sets Z with this property always exist in the setting of Lemma 7.1, as follows.

Lemma 7.5. *Fix positive integers $1 \leq d < k$, and fix $\rho, \delta > 0$. Then there is $\lambda > 0$ such that the following holds. For $p \geq n^{1-k} \log^3 n$, a.a.s. $G \sim \mathbb{H}^k(n, p)$ has the following property. Consider a spanning subgraph $G' \subseteq G$ with $\delta_d(G') \geq \delta p \binom{n-d}{k-d}$ for some $d \geq 1$. Then there is a set Z of ρn vertices such that for any $W \subseteq V(G) \setminus Z$ with $|W| = \lambda n$, there is a matching M in G' covering all vertices in W , each edge of which contains one vertex of W and $k - 1$ vertices of Z .*

The proof of Lemma 7.5 is not too difficult, but to avoid interrupting the flow of this section we defer its proof to Section 7.2. Briefly, the idea is to show that a random set Z typically does the job, using some concentration inequalities and a hypergraph matching criterion due to Aharoni and Haxell.

Having found a rich set Z as guaranteed by Lemma 7.5, we need to show that G' has an absorbing structure with Z as its flexible set. We will greedily construct such an absorbing structure using the following lemma, which says that absorbers can be found rooted on any triple of vertices, even if a few vertices are “forbidden”.

Lemma 7.6. *Fix positive integers $1 \leq d < k$ and some $\gamma > 0$. There is $Q \in \mathbb{N}$ and $C, \sigma > 0$ such that the following holds. For p satisfying $p \geq \max\{n^{-k/2+\gamma}, Cn^{2-k}\}$, a.a.s. $G \sim \mathbb{H}^k(n, p)$ has the following property. Every spanning subgraph G' of G with minimum d -degree at least $(\mu_d(k) + \gamma)p \binom{n-d}{k-d}$ has an absorber of order at most Q rooted on any k -tuple of vertices, even after deleting σn other vertices of G' .*

The proof of Lemma 7.6 is quite involved, and contains the most interesting new ideas in this paper. We defer it to Section 8. Finally, we deduce Lemma 7.1.

Proof of Lemma 7.1. Let $G' \subseteq G$ be a spanning subgraph with $\delta_d(G') \geq (\mu_d(k) + \gamma)p \binom{n-d}{k-d}$. Let L be as in Lemma 7.3 and let σ, Q be as in Lemma 7.6. We may assume $\sigma \leq (\gamma/2)^k$. Choose small ρ with $\rho(Q+1)L \leq \sigma$.

By Lemma 7.5, a.a.s. we can find a ρn -vertex “rich” set Z having the property that for some small $\lambda > 0$, every set of λn vertices can be “matched into” Z . We may assume $\lambda < \rho/(2k)$. By Lemma 7.3 there is an order- ρn resilient template T . Now, an absorbing structure on T would have at most $QL\rho n + L\rho n$ vertices, and since $\rho(Q+1)L \leq \sigma$, we can use Lemma 7.6 to greedily build a $(\rho n, Q)$ -absorbing structure $H \subseteq G'$ on the template T , having at most $(\gamma/2)^k n$ vertices, with Z as the flexible set. Let $X = V(H)$.

We now claim that X satisfies the assumptions of the lemma. First, by the assumption $\sigma \leq (\gamma/2)^k$, it has size at most $(\gamma/2)^k n$. For the second property, consider any $W \subseteq V(G') \setminus X$ with size at most λn , such that $W \cup X$ is divisible by k . By the defining property of our rich set Z , we can find a matching M_1 in G' covering W and $(k-1)|W| \leq (k-1)\lambda n < (\rho/2)n$ vertices of X . By the special property of our absorbing structure H , we can then find a matching M_2 covering all the remaining vertices in $G[X \cup W]$, and then $M_1 \cup M_2$ is the desired matching. \square

7.1 Constructing an absorbing template

In this subsection we prove Lemma 7.3. We will build our desired k -graph via a simple transformation from a bipartite graph with certain properties. The following lemma was proved by Montgomery, and appears as [44, Lemma 10.7]. We write \sqcup to indicate that a union of sets is disjoint.

Lemma 7.7. *For any sufficiently large s , there exists a bipartite graph R with vertex parts X and $Y \sqcup Z$, with $|X| = 3s$, $|Y| = |Z| = 2s$, and maximum degree 100, such that if we remove any s vertices from Z , the resulting bipartite graph has a perfect matching.*

From Lemma 7.7 we can deduce the following lemma (this is a k -uniform version of [37, Lemma 5.2]).

Lemma 7.8. *For any sufficiently large s , there exists a k -partite k -graph S with vertex parts X_1, \dots, X_{k-1} and $Y \sqcup Z$, with $|X_1| = \dots = |X_{k-1}| = 3s$, $|Y|, |Z| = 2s$, and maximum degree 100, such that if we remove any s vertices from Z , the resulting k -graph has a perfect matching (provided its number of vertices is even).*

Proof. Consider the bipartite graph R from Lemma 7.7 on the vertex set $X \sqcup (Y \sqcup Z)$, and let $X_{k-1} = X$. Obtain a k -partite graph R' by adding sets X_1, \dots, X_{k-2} each having $|X|$ new vertices, and for each $1 \leq i < k-1$ putting an arbitrary perfect matching between X_i and X_{i+1} . Now, our k -partite k -graph S has the same vertex set as R' , and an edge for every k -vertex path running through $X_1, \dots, X_{k-1}, Y \sqcup Z$ (call such paths *special paths*). Note that an edge in R can be uniquely extended to a special path in R' . Moreover, a matching in R can always be uniquely extended to a vertex-disjoint union of special paths in R' . \square

We also need the following simple lemma showing that there are sparse hypergraphs with no large independent sets.

Lemma 7.9. *For any k there is some $K \in \mathbb{N}$ such that the following holds. For sufficiently large r there is a k -graph G with r vertices and at most Kr edges, with no independent set of size $r/2$.*

Proof. Consider a random k -graph $G \sim H^k(r, p)$ for $p = (K/2)r / \binom{r}{k}$. For a set of $r/2$ vertices, the probability that there are no edges in that set is $(1-p)^{\binom{r/2}{k}} = e^{-\Omega(Kr)}$, so for large K , the union bound shows that a.a.s. every set of size $r/2$ induces at least one edge. Also, the Chernoff bound shows that a.a.s. G has at most Kn edges. \square

We are now ready to prove Lemma 7.3.

Proof of Lemma 7.3. Start with the k -graph S from Lemma 7.8, with $s = \lceil r/2 \rceil$, and delete at most one vertex from Z to make it have size r . Then, consider an r -vertex k -graph G as in Lemma 7.9 (which exists as long as r is large enough), and place G on the vertex set Z . Let T be the resulting k -graph. It has at most $(k-1)(3s) + 2s + 2s$ vertices and at most $100(4s) + K(2s)$ edges.

Now, consider any set W of fewer than $r/2$ vertices of Z , such that the number of vertices in $T - W$ is divisible by k . By the defining property of G , we can greedily build a matching M_1 in $T[Z \setminus W]$ covering all but s vertices, and then by the defining property of S , there is a perfect matching M_2 in $T - (W \cup V(M_1))$. Then $M_1 \cup M_2$ is the desired perfect matching of $T - W$. \square

7.2 Finding a rich set of vertices

In this subsection we prove Lemma 7.5. We will make use of the following Hall-type theorem for finding large matchings in hypergraphs, due to Aharoni and Haxell [1].

Theorem 7.10. *Let $\{L_1, \dots, L_t\}$ be a family of k' -uniform hypergraphs on the same vertex set. If, for every $\mathcal{I} \subseteq \{1, \dots, t\}$, the hypergraph $\bigcup_{i \in \mathcal{I}} L_i$ contains a matching of size greater than $k'(|\mathcal{I}| - 1)$, then there exists a function $g : \{1, \dots, t\} \rightarrow \bigcup_{i=1}^t E(L_i)$ such that $g(i) \in E(L_i)$ and $g(i) \cap g(j) = \emptyset$ for $i \neq j$.*

To apply Theorem 7.10, the following lemma will be useful, concerning the distribution of edges in random hypergraphs.

Lemma 7.11. *Fix $k, q \in \mathbb{N}$ and $\lambda > 0$, and suppose $p = \omega(n^{1-k} \log n)$. Then a.a.s. $G \sim H^k(n, p)$ has the following property. For every pair of vertex sets \mathcal{I}, U with $|\mathcal{I}| \leq \lambda n$ and $|U| \leq q|\mathcal{I}|$, there are at most $2\lambda qp|\mathcal{I}|n^{k-1}$ edges which contain a vertex from \mathcal{I} and a vertex from U .*

Proof. Fix \mathcal{I}, U as in the lemma statement. By the Chernoff bound, the number of edges intersecting \mathcal{I} and U is at most $2\lambda qp|\mathcal{I}|n^{k-1}$, with probability $1 - \exp(-\Omega(p|\mathcal{I}|n^{k-1}))$. So, by the union bound, the probability that the property in the lemma statement fails is at most

$$\sum_{i=1}^{\lambda n} n^i n^{qi} \exp(-\Omega(pin^{k-1})) = n^{-\omega(1)}. \quad \square$$

We will also need the (very simple) fact that minimum d -degree assumptions are strongest when d is large.

Lemma 7.12. *Let G be an n -vertex k -graph. If $d \geq d'$ and $\delta_d(G) \geq \alpha\binom{n-d}{k-d}$ then $\delta_{d'}(G) \geq \alpha\binom{n-d'}{k-d'}$.*

Proof. Suppose that $\delta_d(G) \geq \alpha\binom{n-d}{k-d}$, and fix any subset S of d' vertices of G . By assumption, for every subset $X \subseteq V(G) \setminus S$ of size $d - d'$, the number of edges containing $X \cup S$ is at least $\alpha\binom{n-d}{k-d}$. Since each edge containing S is being counted exactly $\binom{k-d'}{d-d'}$ times, we conclude that

$$\delta_{d'}(G) \geq \frac{\binom{n-d'}{d-d'}\alpha\binom{n-d}{k-d}}{\binom{k-d'}{d-d'}} = \alpha\binom{n-d'}{k-d'}. \quad \square$$

Now, we prove Lemma 7.5.

Proof of Lemma 7.5. Fix an outcome of G that satisfies the property in Lemma 7.11, for $q = (k-1)^2$ and small $\lambda > 0$ to be determined, and also satisfies the property in Lemma 5.5, for $\mu = \gamma = \delta/2$ and $d = 1$. Consider a spanning subgraph $G' \subseteq G$ with minimum 1-degree at least $\delta p\binom{n-1}{k-1}$. Let Z be a set of pn vertices such that every vertex outside Z has degree at least $(\delta/2)p\binom{n-1}{k-1}$ into Z (by Lemmas 5.5 and 7.12, almost every choice of Z will do).

Now, consider any $W \subseteq V(G') \setminus Z$ with size λn . For each $w \in W$ let L_w be the link $(k-1)$ -graph of w into Z (having an edge $e \subseteq Z$ whenever $e \cup \{w\}$ is an edge of G'). We claim that if λ is sufficiently small then for each $\mathcal{I} \subseteq W$, the $(k-1)$ -graph $H_{\mathcal{I}} := \bigcup_{w \in \mathcal{I}} L_w$ has a matching of size greater than $(k-1)(|\mathcal{I}| - 1)$. The desired result will then follow from Theorem 7.10. Actually, it will be convenient to view $H_{\mathcal{I}}$ as a multigraph (if an edge appears in multiple different L_w , among $w \in \mathcal{I}$, then we include that edge multiple times). This does not affect the existence of matchings, but the correspondence between edges of $H_{\mathcal{I}}$ and edges of G' will be more natural.

To prove the claim, suppose for the purpose of contradiction that there is some $\mathcal{I} \subseteq W$ for which a maximum matching M in $H_{\mathcal{I}}$ has size at most $(k-1)(|\mathcal{I}| - 1)$ (meaning that it has at most $(k-1)^2(|\mathcal{I}| - 1)$ vertices). By the degree condition defining Z , $H_{\mathcal{I}}$ has at least $|\mathcal{I}|(\delta/2)p\binom{n-1}{k-1}$ edges. On the other hand, all edges of $H_{\mathcal{I}}$ intersect $V(M)$ by maximality, so by the property in Lemma 7.11, $H_{\mathcal{I}}$ only has at most $2\lambda qp|\mathcal{I}|n^{k-1}$ edges. This is a contradiction if λ is sufficiently small. \square

8 Finding absorbers

Now we are finally ready to prove Lemma 7.6, showing that dense subgraphs of random graphs have absorbers and completing the proof of Lemma 7.1.

The main difficulty with finding absorbers is that they are *rooted* objects. It is not enough to find an absorber floating somewhere in our graph (which we could easily do with the sparse embedding lemma); what we want is to find an absorber on a specific k -tuple of vertices. In order to achieve this, we define a contraction operation that reduces the task of finding a rooted absorber to the task of finding “contracted absorbers” in a much more flexible setting.

Definition 8.1 (contractible absorbers). A *contractible absorber* rooted on a k -tuple of vertices (x_1, \dots, x_k) is an absorber obtained in the following way. Put k disjoint edges $e_1 = \{x_1, y_1^1, \dots, y_1^{k-1}\}, \dots, e_k = \{x_k, y_k^1, \dots, y_k^{k-1}\}$, then for each i put an externally vertex disjoint absorber H_i rooted on $\{y_1^i, \dots, y_k^i\}$ (we call each of these a *sub-absorber*). Note that the edges in the non-covering matching of a contractible absorber come from the covering matchings of its constituent sub-absorbers, and the non-rooted edges in the covering matching of the contractible absorber come from the non-covering matchings of its sub-absorbers.

A *contracted absorber* rooted at (x_1, \dots, x_k) is a hypergraph obtained as the union of $k-1$ absorbers (which we again call sub-absorbers) each rooted at (x_1, \dots, x_k) , disjoint except for their rooted vertices. One can show that a contracted absorber is always itself an absorber, but we will not need this fact. The *contraction* of a contractible absorber is the contracted absorber obtained by contracting each of its rooted edges to a single vertex. See Figure 3.

We will need our absorbers to satisfy a local sparsity condition to apply the sparse embedding lemma. We recall the definition of the girth of a hypergraph, and introduce a very closely related notion of local sparsity.




Figure 3. An illustration of an order-18 contractible absorber rooted on (x_1, x_2, x_3) , and its contraction. In this 3-uniform case the absorber construction involves two sub-absorbers. The dark hyperedges are the non-covering matching of the contractible absorber and the light hyperedges are the covering matching.

Definition 8.2. A (Berge) *cycle* in a hypergraph is a sequence of edges e_1, \dots, e_ℓ such that there exist distinct vertices v_1, \dots, v_ℓ with $v_i \in e_i \cap e_{i+1}$ for all i (where $e_{\ell+1} = e_1$). The *length* of such a cycle is its number of edges ℓ . The *girth* of a hypergraph is the length of the shortest cycle it contains (if the hypergraph contains no cycle we say it has infinite girth, or is *acyclic*). We say that an absorber rooted on x_1, \dots, x_k is K -*sparse* if it has girth at least K , even after adding the extra edge $\{x_1, \dots, x_k\}$ (that is to say, it has high girth and moreover the roots are “far from each other”).

(The fact that the roots are “far from each other” will allow us to “glue together” absorbers in various ways without worrying about the girth increasing). Recall the definition of the k -density $m_k(H)$ from Definition 4.10.

Lemma 8.3. Fix $\eta > 0$ and $k \in \mathbb{N}$. There is $K > 0$ such that the following holds. For any k -uniform contracted absorber H with girth at least K , we have $m_k(H) \leq 2/k + \eta$.

The proof of Lemma 8.3 is basically just a calculation so we defer it to Section 8.1.

We will also need the following lemma, showing that we can use the definition of $\mu_d(k)$ itself to find locally sparse absorbers. We will apply this lemma to a cluster k -graph obtained via the sparse regularity lemma.

Lemma 8.4. For any positive integers $d < k$ and any $\eta, K > 0$, there are $\delta, M > 0$ such that the following holds for sufficiently large n . Consider a k -graph G on n vertices such that all but $\delta \binom{n}{d}$ d -sets of vertices have degree at least $(\mu_d(k) + \eta) \binom{n-d}{k-d}$. Then for any vertices x_1, \dots, x_k which each have degree at least $\eta \binom{n-1}{k-1}$, there is a K -sparse absorber with order at most M , rooted on those vertices.

Without the local sparsity condition it would be fairly easy to prove Lemma 8.4, more or less by using the definition of the Dirac threshold twice. To deal with the sparseness condition we fix a locally sparse “pattern” (which is a high-girth bipartite graph), and use the definition of the Dirac threshold plus a random sampling trick to find an absorber “in line with the pattern”. We defer the details to Section 8.2.

As previously mentioned, for the proof of Lemma 7.6 we will use a contraction trick: we “contract” G' to obtain a k -graph $G'(\mathcal{F}, \mathcal{P})$, in such a way that if we can find contracted absorbers in $G'(\mathcal{F}, \mathcal{P})$ satisfying certain properties, then these correspond to rooted contractible absorbers in G' . Before proving Lemma 7.6 we define this “contraction” operation.

Definition 8.5. Consider any k -graph G , consider a family $\mathcal{F} \subseteq V(G)^{k-1}$ of disjoint $(k-1)$ -tuples, and consider a family \mathcal{P} of disjoint sets $U_1, \dots, U_{k-1} \subseteq V(G)$, such that the tuples in \mathcal{F} and the sets in \mathcal{P} do not share any vertices. Let $v(\mathcal{P})$ be the total number of vertices in the sets in \mathcal{P} .

Let $G(\mathcal{F}, \mathcal{P})$ be the k -graph obtained as follows. Start with the k -graph $G[U_1 \cup \dots \cup U_{k-1}]$, and for each $\mathbf{v} \in \mathcal{F}$ add a new vertex $w_{\mathbf{v}}$. For each tuple $\mathbf{v} = (v_1, \dots, v_{k-1}) \in \mathcal{F}$, each $1 \leq j \leq k-1$, and each $f \subseteq U_j$ such that $f \cup \{v_j\}$ is an edge of G , put an edge $f \cup \{w_{\mathbf{v}}\}$ in $G(\mathcal{F}, \mathcal{P})$.

One should visualise each $\mathbf{v} = (v_1, \dots, v_{k-1}) \in \mathcal{F}$ as being “contracted” to a single vertex $w_{\mathbf{v}}$, and all edges involving \mathbf{v} being deleted except those edges that contain some v_j and have all their other

vertices in the corresponding U_j . The reason for the edge deletion is to ensure that each edge in $G(\mathcal{F}, \mathcal{P})$ corresponds to exactly one edge in G . So, if $G \sim H^k(n, p)$ is a random k -graph, then $G(\mathcal{F}, \mathcal{P})$ may be interpreted as a subgraph of a random graph with the same edge probability p , and we may apply the sparse embedding lemma to it. In our proof of Lemma 7.6 we will choose \mathcal{F} and \mathcal{P} depending on our desired roots and the structure of G' .

Now we finally prove Lemma 7.6.

Proof of Lemma 7.6. First note that we can assume $p = \Omega(n^{d-k} \log n) = \omega(n^{d-k})$, because otherwise G itself a.a.s. has minimum d -degree zero and the lemma statement is vacuous. Also, we can assume that γ is sufficiently small with respect to k (the lemma statement only gets stronger as we decrease γ). Now, there are a number of constants in our proof that are defined in terms of each other (constants in the sense that they do not depend on n). First, let $\beta = 1/(3k)$. Second, let K be large enough to satisfy Lemma 8.3, applied with $\eta = \gamma/2$. Third, let M be large enough and $\delta > 0$ be small enough to satisfy Lemma 8.4, applied with $\eta = \gamma/4$ and the value of K just defined. Then, $\sigma, \tau > 0$ will be small relative to all constants defined so far (small enough to satisfy certain inequalities later in the proof), and we let $\alpha = \sigma/k^3$. Next, $\varepsilon > 0$ will be very small, and t_0 very large, even compared to σ . Finally, $\kappa > 0$ will be tiny compared to all other constants.

We record some properties that G and each of the $G(\mathcal{F}, \mathcal{P})$ a.a.s. satisfy.

Claim. G a.a.s. satisfies each of the following properties.

- (1) For each \mathcal{F}, \mathcal{P} as in Definition 8.5, such that $|\mathcal{F}| = k\alpha n$, the k -graph $G(\mathcal{F}, \mathcal{P})$ satisfies the conclusion of Theorem 4.11 (the sparse embedding lemma), for embedding all graphs H on at most $(k-1)(M+k)$ vertices which have $m_k(H) \leq 2/k + \gamma/2$. (The other parameters τ, κ with which we apply the sparse embedding lemma are as described at the beginning of the proof, and the lemma gives us an upper bound on ε in terms of τ).
- (2) Each $G(\mathcal{F}, \mathcal{P})$ as above is $(o(1), p, 1 - o(1))$ -upper-uniform.
- (3) For any set W' of $2\sigma n$ vertices and every vertex x , there are at most $4(\sigma n)p \binom{n-2}{k-2}$ edges in G containing x and a vertex of W' .
- (4) For any spanning subgraph $G' \subseteq G$ with minimum d -degree at least $(\mu_d(k) + \gamma)p \binom{n-d}{k-d}$, and any set W' of at most $2\sigma n$ vertices, let Y be a uniformly random subset of βn vertices of $V(G') \setminus W'$. Then a.a.s. all but $o(n^d)$ of the d -sets of vertices in G' have d -degree at least $(\mu_d + \gamma/2)p \binom{\beta n-d}{k-d}$ into Y .

Proof. Observe that each $G(\mathcal{F}, \mathcal{P})$ has at least $|\mathcal{F}| = \Omega(n)$ vertices and it can be coupled as a subset of the binomial random k -graph on its vertex set (with edge probability p). Indeed, for each possible edge e of $G(\mathcal{F}, \mathcal{P})$, there is a single k -set $\phi(e)$ whose presence as an edge in G determines whether e is in $G(\mathcal{F}, \mathcal{P})$.

For (1) and (2), there are at most $2^{nk} n^{k \cdot k\alpha n} = \exp(O(n \log n))$ ways to choose \mathcal{F} and \mathcal{P} . Since $p = \omega(n^{1-k} \log n)$, we may apply Theorem 4.11 and Corollary 5.2 and take the union bound over all possibilities for \mathcal{F} and \mathcal{P} . Note that the requirement on p in Theorem 4.11 is that p exceeds $n^{-1/m_k(H)}$ by a large constant, where $m_k(H) \leq 2/k + \gamma/2$. This is certainly satisfied since we are assuming $p \geq n^{-k/2+\gamma}$, and that γ is small with respect to k .

For (3) we simply apply Lemma 5.3 with $\lambda = 2\sigma$, and for (4) we apply Lemma 5.6 to $G' - W$ (recalling our assumption that $p = \omega(n^{d-k})$), after applying Lemma 5.4. \square

Consider an outcome of G satisfying all the above properties (for the rest of the proof, we can forget that G is an instance of a random graph and just work with these properties). Let $G' \subseteq G$ be a spanning subgraph with minimum d -degree at least $(\mu_d(k) + \gamma)p \binom{n-d}{k-d}$, consider any set W of σn vertices, and consider vertices x_1, \dots, x_k outside W . We will show that $G' - W$ has an absorber rooted on x_1, \dots, x_k of order at most $(k-1)(M+k)$.

We will accomplish this by studying $G'(\mathcal{F}, \mathcal{P})$ for certain \mathcal{F} and \mathcal{P} . First, \mathcal{F} will be defined in terms of the edges incident to x_1, \dots, x_k , using the following claim.

Claim. For each i let $\Gamma(x_i)$ be the link $(k-1)$ -graph of x_i with respect to $G' - W$. Then we can find matchings $M^i \subseteq E(\Gamma(x_i))$ of size αn , such that no two of these matchings share a vertex.

Proof. First, ignoring the disjointness condition, note that (3) implies that each $\Gamma(x_i)$ has a matching M_0^i of size $\sigma n/k$. Indeed, suppose a maximum matching in $\Gamma(x_i)$ were to have fewer than $\sigma n/k$ edges. Then the set of vertices of this matching would comprise a set W' of fewer than σn vertices such that

all the edges of G' which contain x intersect $W' \cup W$. But by the minimum degree condition on G' (and Lemma 7.12) there are at least $(\mu_d(k) + \gamma)p \binom{n-1}{k-1}$ such edges, contradicting (3) for small σ .

We can then delete some edges from the M_0^i to obtain the desired matchings M^i , recalling that $\alpha = \sigma/k^3$. \square

Now we can define \mathcal{F} : for each i , arbitrarily order the vertices in each edge of M^i to obtain a collection \mathcal{F}^i of αn disjoint $(k-1)$ -tuples. Let $\mathcal{F} = \bigcup_{i=1}^k \mathcal{F}^i$, and let $V(\mathcal{F})$ be the set of vertices in the tuples in \mathcal{F} .

Next, we will choose the sets U_1, \dots, U_{k-1} in \mathcal{P} randomly, so that they satisfy a certain degree condition. This is encapsulated in the following claim.

Claim. *There are disjoint βn -vertex sets $U_1, \dots, U_{k-1} \subseteq V(G) \setminus (W \cup V(\mathcal{F}))$ such that all but $o(n^d)$ d -sets of vertices in G' have degree at least $(\mu_d(k) + \gamma/3) \binom{\beta n - d}{k-d}$ into each U_j .*

Proof. By (4), almost any choice of $U_1, \dots, U_{k-1} \subseteq V(G)$ will do. \square

Now, in $G'(\mathcal{F}, \mathcal{P})$, let $X^i \subseteq V(G'(\mathcal{F}, \mathcal{P}))$ be the set of “newly contracted” vertices arising from tuples in \mathcal{F}^i , and let $X = \bigcup_{i=1}^k X^i$ be the set of all newly contracted vertices. Note that $G'(\mathcal{F}, \mathcal{P})$ is a subgraph of $G(\mathcal{F}, \mathcal{P})$. Our goal from now on is to prove the following claim.

Claim. *There is a set of vertices $y_1 \in X^1, \dots, y_k \in X^k$ such that for each $1 \leq j \leq k-1$ there is an absorber of order at most M in $G'(\mathcal{F}, \mathcal{P})$ rooted at y_1, \dots, y_k , whose other vertices lie entirely in $U_j \setminus (W \cup X)$.*

Recalling the definition of \mathcal{F} and $G(\mathcal{F}, \mathcal{P})$, the absorbers in the above claim then form sub-absorbers for some contractible absorber in G' of order at most $(k-1)M$, rooted at x_1, \dots, x_k . So, it suffices to prove the above claim to complete the proof of Lemma 7.6. Crucially, in our new goal, we have quite a lot of freedom to choose the roots y_1, \dots, y_k . This will allow us to use the sparse embedding lemma (that is, property (1)).

Proof of claim. Recalling property (2), we apply our sparse regularity lemma (Lemma 4.2) to $G'(\mathcal{F}, \mathcal{P})$, with small ε and large t_0 . Apply Lemma 4.5 and Lemma 4.7 with small threshold τ , to obtain a t -vertex cluster graph \mathcal{R} , with $t_0 \leq t$. Let \mathcal{U}_j be the set of clusters (vertices of \mathcal{R}) contained in each $U_j \setminus (W \cup X)$, let \mathcal{X}^i be the set of clusters contained in each X^i , and let \mathcal{W} be the set of clusters contained in W .

Since $\sigma, \tau, \varepsilon$ are very small, and t_0 is very large, Lemma 4.5 and Lemma 4.7 ensure that for each i , \mathcal{X}^i has about αt clusters, and almost all of those clusters have degree at least $(\mu_d(k) + \gamma/4) \binom{\beta t - 1}{k-1}$ into each \mathcal{U}_j . Fix such a cluster V^i , for each i .

Now, $\mathcal{R}[\mathcal{U}_j \cup \{V^1, \dots, V^k\}]$ satisfies the assumptions for Lemma 8.4, so it contains a K -sparse absorber in \mathcal{R} of order at most M rooted at V^1, \dots, V^k whose other vertices (clusters) lie in \mathcal{U}_j . These absorbers form the sub-absorbers for a contracted absorber in \mathcal{R} , which has girth at least K and therefore has k -density at most $2/k + \gamma/2$ by Lemma 8.3. Now, by property (1), a canonical copy of the same contracted absorber exists in $G'(\mathcal{F}, \mathcal{P})$ (in fact, there are many such copies). The sub-absorbers of this contracted absorber then satisfy the requirements of the claim. \square

This completes the proof of Lemma 7.6. \square

8.1 Locally sparse absorbers have low k -density

In this section we prove Lemma 8.3.

Proof of Lemma 8.3. Note that a cycle of length two corresponds to a pair of edges that intersect in more than one vertex (so a hypergraph has girth greater than two if and only if it is linear). The *line graph* $L(G)$ of a linear k -graph G has the edges of G as vertices, with an edge $e_1 e_2$ when e_1 and e_2 are incident in G . Note that $v(L(G)) = e(G)$ and $v(G) \geq ke(G) - e(L(G))$. Observe that a linear k -graph G is acyclic if and only if its line graph is a forest, in which case $e(L(G)) \leq v(L(G)) - 1 = e(G) - 1$, so $v(G) \geq (k-1)e(G) + 1$.

Now consider any subgraph $H' \subseteq H$ with $v(H') > k$. First, if H' is acyclic, then

$$\frac{e(H') - 1}{v(H') - k} \leq \frac{e(H') - 1}{(k-1)e(H') + 1 - k} = \frac{1}{k-1} \leq 2/k,$$

by the above discussion. Otherwise, H' has a cycle, which must have length at least K , so $v(H') \geq (k-1)K$. Note that every vertex of H has degree at most 2, except k vertices x_1, \dots, x_k which have degree $k-1$. So, $ke(H')$ can be bounded by $k(k-1) + 2(v(H') - k)$, and

$$\frac{e(H') - 1}{v(H') - k} \leq \frac{(k(k-1) + 2(v(H') - k))/k}{v(H') - k} \leq \frac{k(k+1)/(K(k-1) - k) + 2}{k} \leq \frac{2}{k} + \eta$$

for large K . □

8.2 Locally sparse absorbers in dense hypergraphs

In this subsection we prove Lemma 8.4. If we were to ignore the local sparseness condition it would be quite simple to find an absorber rooted on our desired vertices: recalling that an absorber essentially consists of two perfect matchings on some vertex set, we could simply apply the definition of the Dirac threshold twice, in appropriate subgraphs of G .

In order to deal with the sparseness condition, we fix a high-girth hypergraph L (with edges of size about $\log n$), which will form a “pattern” for our absorber. We consider a random injection from L into the vertex set of our graph G , so that (with positive probability) each of the edges of L corresponds to a subgraph of G with minimum d -degree exceeding its Dirac threshold, and therefore has a perfect matching. We will define L in such a way that the union of these perfect matchings gives us an absorber with the desired properties.

It is convenient to deduce Lemma 8.4 from a slightly simpler lemma where no vertices of exceptionally low degree are allowed. To state this lemma we define a slight generalisation of the notion of an absorber.

Definition 8.6. An r -absorber rooted at a rk -set of vertices y_1, \dots, y_{rk} is a hypergraph which can be partitioned into two matchings, one of which covers the entire vertex set and the other of which covers every vertex except y_1, \dots, y_{rk} . We say an r -absorber is K -sparse if it has girth at least K , even after adding an extra edge $\{y_1, \dots, y_{rk}\}$. The point of this definition is that if we have roots x_1, \dots, x_k and we pick arbitrary disjoint edges e_1, \dots, e_k containing the roots (whose other vertices are $y_1, \dots, y_{(k-1)k}$, say), then a K -sparse $(k-1)$ -absorber rooted at $y_1, \dots, y_{(k-1)k}$ (not containing the vertices x_1, \dots, x_k) gives us a K -sparse absorber rooted at x_1, \dots, x_k .

Now, the key lemma is as follows.

Lemma 8.7. *For any $\gamma > 0$ and $r, k, K \in \mathbb{N}$ the following holds for sufficiently large n . For any k -graph G on n vertices with $\delta_d(G) \geq (\mu_d(k) + \gamma) \binom{n-d}{k-d}$, there is a K -sparse r -absorber rooted on every rk -tuple of vertices.*

As outlined at the beginning of this subsection, to enforce our local sparsity condition we will apply the Dirac threshold to subgraphs arising from a high-girth “pattern”. This pattern will be constructed from the following bipartite graph.

Lemma 8.8. *Fix k and K . For sufficiently large n , there is a q -regular bipartite graph F with girth at least K and at most n edges, for some $q \geq \log^2 n$ which is divisible by k .*

Proof. There are many ways to prove this. For example, fix a prime p such that $n/2^{K+1} \leq p^{K+1} \leq n$ (which exists by Bertrand’s postulate), and consider the bipartite graph defined by Lazebnik and Ustimenko in [38], which is p -regular, has $2p^K$ vertices and has girth at least $K+5$. Then repeatedly delete perfect matchings at most $k-1$ times until we arrive at a q -regular bipartite graph with q divisible by k . (Note that nonempty regular bipartite graphs always have perfect matchings). For large n , we have $q \geq p - (k-1) \geq \log^2 n$. □

Another ingredient we will need is the following fact, which actually already appeared in the proof of Lemma 3.4.

Lemma 8.9. *There is $c = c(k) > 0$ (depending on k) such that the following holds. Consider an n -vertex k -graph G and consider a set A of d vertices with degree at least $(\mu + \gamma) \binom{n-d}{k-d}$. Let S be a random subset of $Q \geq 2d$ vertices of G . Then with probability at least $1 - e^{-c\gamma^2 Q}$, A has degree at least $(\mu + \gamma/2) \binom{Q-d}{k-d}$ into S .*

Now we are ready to prove Lemma 8.7.

Proof of Lemma 8.7. Consider a graph F as in Lemma 8.8, and for a vertex v , let $F(v)$ be the set of edges incident to v . Then define a q -uniform hypergraph L whose vertices are the edges of F and whose edges are the sets $F(v)$. Note that the girth condition on F transfers to L : the girth of L is at least K . Also, note that the two vertex parts of F correspond to two perfect matchings M_1 and M_2 partitioning the edges of L . Let L' be obtained by deleting rk vertices from one of the edges of M_2 , so that L' has rk vertices z_1, \dots, z_{rk} which have degree 1 in L' . Let L'' be the non-uniform hypergraph obtained by deleting each z_i from the edge it is contained in (so that L'' has rk edges with size only $q - 1$, in addition to the edge of L' with size only $q - rk$).

Now, consider an rk -tuple of vertices (y_1, \dots, y_{rk}) in G , and consider a uniformly random injection $\phi : V(L'') \rightarrow V(G) \setminus \{y_1, \dots, y_{rk}\}$. Extend ϕ to a map $V(L') \rightarrow V(G)$ by taking $\phi(z_i) = y_i$ for each i .

Then, for each edge $e \in E(L')$, note that $\phi(e)$ is ‘‘almost’’ a uniformly random subset of q vertices of G . To be precise, one can couple $\phi(e)$ with a uniformly random subset S of $q = \Omega(\log^2 n)$ vertices of G , in such a way that the size of the symmetric difference $|S \Delta \phi(e)|$ is at most $1 + 2rk$. By Lemma 8.9 and the union bound, with probability $1 - o(n^{-k})$ every d -set of vertices U satisfies

$$\deg_S(U) \geq (\mu_d(k) + \gamma/2) \binom{q-d}{k-d},$$

implying that $\deg_{\phi(e)}(U) \geq (\mu_d(k) + \gamma/3) \binom{q-d}{k-d}$. By the union bound, a.a.s. this holds for each $e \in E(L')$, so fix such an outcome of ϕ . Then for each $e \in E(L')$, $G[\phi(e)]$ has minimum d -degree at least $(\mu_d(k) + \gamma/3) \binom{q-d}{k-d}$, so has a perfect matching. The union of these perfect matchings gives a K -sparse r -absorber rooted at y_1, \dots, y_{rk} . \square

Now, we deduce Lemma 8.4 from Lemma 8.7.

Proof of Lemma 8.4. Consider x_1, \dots, x_k as in the theorem statement, and consider a random subset U of R vertices of $G - \{x_1, \dots, x_k\}$, for some large R to be determined. Then by Lemma 8.9 (with $d = 1$) and Lemma 3.4, with probability at least $1 - \binom{R}{d}(\delta + e^{-\Omega(R)}) - ke^{-\Omega(R)}$ each x_i has at least $(\eta/2)R$ neighbours in U , and $G[U]$ has minimum d -degree at least $(\mu_d(k) + \eta/2) \binom{R-d}{k-d}$. This probability is greater than zero for large R and small $\delta > 0$, so we may fix such a choice of U .

For each i , choose an edge containing x_i and k vertices in U , in such a way that these chosen edges form a matching M (we can do this greedily). Let $y_1, \dots, y_{(k-1)k}$ be the vertices in $V(M) \cap U$, and apply Lemma 8.7 to find a K -sparse $(k-1)$ -absorber H rooted at $y_1, \dots, y_{(k-1)k}$. Then $M \cup H$ is a K -sparse absorber of order at most R rooted at x_1, \dots, x_k . \square

9 Concluding remarks

We have proved that if $p \geq \max\{n^{-k/2+\gamma}, Cn^{-k+2}\}$, for any $\gamma > 0$ and sufficiently large C , then the random k -graph $G \sim \mathbb{H}^k(n, p)$ typically obeys a relative version of any Dirac-type theorem for perfect matchings in hypergraphs. There are a number of compelling further directions of research.

It is natural to try to improve our assumption on p , with the eventual goal of removing it entirely (as in Conjecture 1.3). First, as mentioned in the introduction, we observe that the assumption $p \geq Cn^{-k+2}$ can actually be weakened substantially (though this only affects the case $(d, k) = (1, 3)$). The reason for this assumption was to ensure that all vertices have linear degree, so that an absorbing structure of linear size could be built greedily. The reason we needed an absorbing structure of linear size was that Lemma 6.1 does not have effective bounds: it guarantees an almost-perfect matching covering all but $o(n)$ vertices, but since the regularity lemma is notorious for its extremely weak quantitative aspects, this $o(n)$ term is actually only very slightly sub-linear. However, it is possible to use a bootstrapping trick due to Nenadov and Škorić [47] to get a much stronger bound in the setting of Lemma 6.1, which allows us to make do with a much smaller absorbing structure. Using these ideas, it seems to be possible to take p to be as small as about $n^{-4/3}$, in the case $(d, k) = (1, 3)$. Actually, there is some hope of being able to remove the extra assumption $p \geq Cn^{-k+2}$ altogether, by using the Aharoni–Haxell matching criterion (Theorem 7.10) to build an absorbing structure, instead of building it greedily. We have not considered this in detail.

On the other hand, the assumption that p is somewhat larger than $n^{-k/2}$ seems to be much more crucial. An absorber has at least $k/2$ times more edges than unrooted vertices, so absorbers of constant size simply will not exist for smaller p . We imagine that completely new ideas will be required to bypass this barrier.

Another interesting direction would be to consider spanning subgraphs other than perfect matchings. For example, a *loose* cycle is a cyclically ordered collection of edges, such that only consecutive edges intersect, and then only in a single vertex. A *tight* cycle is a cyclically ordered collection of vertices, such that every k consecutive vertices form an edge. There is also a spectrum of different notions of cycles between these two extremes, and Dirac-type problems have been studied for Hamiltonian cycles of all these different types. We believe that it should be possible to adapt the methods in this paper to prove an analogue of Theorem 1.4 for loose Hamiltonian cycles, which are linear (no two edges intersect in more than one vertex) and behave in a very similar way to perfect matchings. It may also be possible to adapt our methods to study other types of Hamiltonian cycles, but this would probably require using different machinery from Theorem 4.11 (which only works for linear hypergraphs).

Finally, it may also be interesting to consider Dirac-type theorems relative to *pseudorandom* hypergraphs, which are not random but satisfy some characteristic properties of random hypergraphs. Certain extremal problems relative to pseudorandom hypergraphs have been studied by Conlon, Fox and Zhao [12] in connection with the Green–Tao theorem on arithmetic progressions in the prime numbers, and the existence of perfect matchings in pseudorandom hypergraphs has been studied by Hàn, Han and Morris [27]. It seems plausible that the methods in this paper can be adapted to work for hypergraphs satisfying some notion of pseudorandomness, but we have not explored this further.

Acknowledgements. We are grateful to the referee for their extremely careful reading of the paper, and a large number of useful comments and suggestions.

References

- [1] R. Aharoni and P. Haxell, *Hall’s theorem for hypergraphs*, J. Graph Theory **35** (2000), no. 2, 83–88.
- [2] P. Allen, J. Böttcher, J. Ehrenmüller, and A. Taraz, *The bandwidth theorem in sparse graphs*, Advances in Combinatorics (2020).
- [3] P. Allen, J. Böttcher, H. Hàn, Y. Kohayakawa, and Y. Person, *Blow-up lemmas for sparse graphs*, arXiv preprint arXiv:1612.00622 (2016).
- [4] N. Alon, P. Frankl, H. Huang, V. Rödl, A. Ruciński, and B. Sudakov, *Large matchings in uniform hypergraphs and the conjecture of Erdős and Samuels*, J. Combin. Theory Ser. A **119** (2012), no. 6, 1200–1215.
- [5] J. Balogh, B. Csaba, and W. Samotij, *Local resilience of almost spanning trees in random graphs*, Random Structures Algorithms **38** (2011), no. 1-2, 121–139.
- [6] S. Ben-Shimon, M. Krivelevich, and B. Sudakov, *On the resilience of Hamiltonicity and optimal packing of Hamilton cycles in random graphs*, SIAM J. Discrete Math. **25** (2011), no. 3, 1176–1193.
- [7] J. Böttcher, Y. Kohayakawa, and A. Taraz, *Almost spanning subgraphs of random graphs after adversarial edge removal*, Combin. Probab. Comput. **22** (2013), no. 5, 639–683.
- [8] D. Clemens, J. Ehrenmüller, and Y. Person, *A Dirac-type theorem for Berge cycles in random hypergraphs*, Electron. J. Combin. **27** (2020), no. 3, Paper No. 3.39, 23.
- [9] D. Conlon and W. T. Gowers, *Combinatorial theorems in sparse random sets*, Ann. of Math. (2) **184** (2016), no. 2, 367–454.
- [10] D. Conlon, W. T. Gowers, W. Samotij, and M. Schacht, *On the KLR conjecture in random graphs*, Israel J. Math. **203** (2014), no. 1, 535–580.
- [11] D. Conlon, *Combinatorial theorems relative to a random set*, Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV, Kyung Moon Sa, Seoul, 2014, pp. 303–327.
- [12] D. Conlon, J. Fox, and Y. Zhao, *A relative Szemerédi theorem*, Geom. Funct. Anal. **25** (2015), no. 3, 733–762.
- [13] G. A. Dirac, *Some theorems on abstract graphs*, Proc. London Math. Soc. (3) **2** (1952), 69–81.
- [14] P. Erdős, A. Gyárfás, and L. Pyber, *Vertex coverings by monochromatic cycles and trees*, J. Combin. Theory Ser. B **51** (1991), no. 1, 90–95.
- [15] P. Erdős and M. Simonovits, *A limit theorem in graph theory*, Studia Sci. Math. Hungar. **1** (1966), 51–57.
- [16] P. Erdős and A. H. Stone, *On the structure of linear graphs*, Bull. Amer. Math. Soc. **52** (1946), 1087–1091.

- [17] A. Ferber and L. Hirschfeld, *Co-degrees resilience for perfect matchings in random hypergraphs*, Electron. J. Combin. **27** (2020), no. 1, Paper No. 1.40.
- [18] A. Ferber and M. Kwan, *Almost all Steiner triple systems are almost resolvable*, Forum Math. Sigma **8** (2020), Paper No. e39, 24.
- [19] A. Ferber, R. Nenadov, A. Noever, U. Peter, and N. Škorić, *Robust Hamiltonicity of random directed graphs*, J. Combin. Theory Ser. B **126** (2017), 1–23.
- [20] P. Frankl and A. Kupavskii, *The Erdős matching conjecture and concentration inequalities*, arXiv preprint arXiv:1806.08855 (2018).
- [21] K. Frankston, J. Kahn, B. Narayanan, and J. Park, *Thresholds versus fractional expectation-thresholds*, Ann. of Math. (2) **194** (2021), no. 2, 475–495.
- [22] A. Frieze and M. Krivelevich, *On two Hamilton cycle problems in random graphs*, Israel J. Math. **166** (2008), 221–234.
- [23] S. Glock, D. Kühn, A. Lo, R. Montgomery, and D. Osthus, *On the decomposition threshold of a given graph*, J. Combin. Theory Ser. B **139** (2019), 47–127.
- [24] W. T. Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, Ann. of Math. (2) **166** (2007), no. 3, 897–946.
- [25] C. Greenhill, M. Isaev, M. Kwan, and B. D. McKay, *The average number of spanning trees in sparse graphs with given degrees*, European J. Combin. **63** (2017), 6–25.
- [26] H. Hàn, Y. Person, and M. Schacht, *On perfect matchings in uniform hypergraphs with large minimum vertex degree*, SIAM J. Discrete Math. **23** (2009), no. 2, 732–748.
- [27] H. Hàn, J. Han, and P. Morris, *Factors and loose Hamilton cycles in sparse pseudo-random hypergraphs*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2020, pp. 702–717.
- [28] H. Huang, C. Lee, and B. Sudakov, *Bandwidth theorem for random graphs*, J. Combin. Theory Ser. B **102** (2012), no. 1, 14–37.
- [29] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Cambridge University Press, 2000.
- [30] A. Johansson, J. Kahn, and V. Vu, *Factors in random graphs*, Random Structures Algorithms **33** (2008), no. 1, 1–28.
- [31] J. Kahn, *Asymptotics for Shamir’s problem*, arXiv preprint arXiv:1909.06834 (2019).
- [32] I. Khan, *Perfect matchings in 4-uniform hypergraphs*, J. Combin. Theory Ser. B **116** (2016), 333–366.
- [33] Y. Kohayakawa, T. Łuczak, and V. Rödl, *On K^4 -free subgraphs of random graphs*, Combinatorica **17** (1997), no. 2, 173–213.
- [34] Y. Kohayakawa and V. Rödl, *Szemerédi’s regularity lemma and quasi-randomness*, Recent advances in algorithms and combinatorics, CMS Books Math./Ouvrages Math. SMC, vol. 11, Springer, New York, 2003, pp. 289–351.
- [35] 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.
- [36] M. Krivelevich, *Triangle factors in random graphs*, Combin. Probab. Comput. **6** (1997), no. 3, 337–347.
- [37] M. Kwan, *Almost all Steiner triple systems have perfect matchings*, Proc. Lond. Math. Soc. (3) **121** (2020), no. 6, 1468–1495.
- [38] F. Lazebnik and V. A. Ustimenko, *Explicit construction of graphs with an arbitrary large girth and of large size*, vol. 60, 1995, ARIDAM VI and VII (New Brunswick, NJ, 1991/1992), pp. 275–284.
- [39] C. Lee and B. Sudakov, *Dirac’s theorem for random graphs*, Random Structures Algorithms **41** (2012), no. 3, 293–305.
- [40] A. Lo and K. Markström, *F-factors in hypergraphs via absorption*, Graphs Combin. **31** (2015), no. 3, 679–712.
- [41] W. Mantel, *Problem 28 (solution by H. Gouwentak, W. Mantel, J. Teixeira de Mattes, F. Schuh, and W. A. Wythoff)*, Wiskundige Opgaven (1907), no. 10, 60–61.
- [42] R. Montgomery, *Embedding bounded degree spanning trees in random graphs*, arXiv preprint arXiv:1405.6559 (2014).

- [43] R. Montgomery, *Hamiltonicity in random graphs is born resilient*, J. Combin. Theory Ser. B **139** (2019), 316–341.
- [44] R. Montgomery, *Spanning trees in random graphs*, Adv. Math. **356** (2019), 106793, 92.
- [45] R. Montgomery, *Hamiltonicity in random directed graphs is born resilient*, Combin. Probab. Comput. **29** (2020), no. 6, 900–942.
- [46] R. Nenadov, A. Steger, and M. Trujić, *Resilience of perfect matchings and Hamiltonicity in random graph processes*, Random Structures Algorithms **54** (2019), no. 4, 797–819.
- [47] R. Nenadov and N. Škorić, *On Komlós’ tiling theorem in random graphs*, Combin. Probab. Comput. **29** (2020), no. 1, 113–127.
- [48] A. Noever and A. Steger, *Local resilience for squares of almost spanning cycles in sparse random graphs*, Electron. J. Combin. **24** (2017), no. 4, Paper No. 4.8.
- [49] O. Pikhurko, *Perfect matchings and K_4^3 -tilings in hypergraphs of large codegree*, Graphs Combin. **24** (2008), no. 4, 391–404.
- [50] 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.
- [51] V. Rödl, A. Ruciński, and E. Szemerédi, *A Dirac-type theorem for 3-uniform hypergraphs*, Combin. Probab. Comput. **15** (2006), no. 1-2, 229–251.
- [52] V. Rödl, A. Ruciński, and E. Szemerédi, *Perfect matchings in large uniform hypergraphs with large minimum collective degree*, J. Combin. Theory Ser. A **116** (2009), no. 3, 613–636.
- [53] V. Rödl and J. Skokan, *Regularity lemma for k -uniform hypergraphs*, Random Structures Algorithms **25** (2004), no. 1, 1–42.
- [54] M. Schacht, *Extremal results for random discrete structures*, Ann. of Math. (2) **184** (2016), no. 2, 333–365.
- [55] B. Sudakov and V. H. Vu, *Local resilience of graphs*, Random Structures Algorithms **33** (2008), no. 4, 409–433.
- [56] T. Tao and V. Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge University Press, Cambridge, 2006.
- [57] N. Škorić, A. Steger, and M. Trujić, *Local resilience of an almost spanning k -cycle in random graphs*, Random Structures Algorithms **53** (2018), no. 4, 728–751.
- [58] Y. Zhao, *Recent advances on Dirac-type problems for hypergraphs*, Recent trends in combinatorics, IMA Vol. Math. Appl., vol. 159, Springer, [Cham], 2016, pp. 145–165.