

Spanning subhypergraphs with degree constraints

Noga Alon^{*} Penny Haxell[†] Aleksa Milojević[‡] Jacques Verstraëte[§]

Abstract

An old result of Tutte states that any d -regular graph contains a spanning subgraph in which every vertex has degree k or $k+1$, for every $1 \leq k \leq d$. We generalize this statement to hypergraphs, showing, for example, that every 3-uniform d -regular hypergraph contains a subgraph in which all degrees are $k, k+1$ or $k+2$, for every $1 \leq k \leq d$. This statement is best possible in the sense that the corresponding statement with only two allowed consecutive values is not true. We provide generalizations of this statement to higher uniformities and discuss several open problems.

1 Introduction

Many classical problems in graph theory can be formulated as questions about finding a subgraph with prescribed degree conditions. For example, two basic theorems of graph theory, Hall's marriage theorem [11] and Tutte's matching theorem [22], give necessary and sufficient conditions for graphs to have a perfect matching, which is a spanning subgraph in which all degrees are one. If one asks the spanning subgraph to be 2-regular instead of 1-regular, then Petersen's theorem [16] guarantees that such subgraphs can be found in all even-regular graphs.

The general framework for studying questions about subgraphs with prescribed degrees was introduced by Lovász [14, 15]. Suppose that every vertex $v \in V(G)$ is given a list of allowed degrees $H(v)$. Under which conditions on the lists and on G does there exist a subgraph $G' \subseteq G$ for which $\deg_{G'} v \in H(v)$ for all $v \in V(G)$? This general problem was studied in a variety of contexts, including [1, 7, 8, 19].

Regular graphs have particularly nice properties in this context. For example, Hall's theorem shows that regular bipartite graphs always have perfect matchings, and Petersen's theorem shows that even-regular graphs always have 2-factors. Moreover, as Tutte [23] showed, any regular graph has a spanning subgraph in which all degrees are either 1 or 2. There is no hope of proving such statements without some regularity assumption, since the graph may be a star.

In this paper, we focus on a version of this problem for regular hypergraphs. The natural way to generalize Tutte's theorem is to ask for which values of t and r it is true that any regular r -uniform hypergraph $\mathcal{H}$ contains a spanning subhypergraph all of whose degrees are in the set $\{1, 2, \dots, t\}$.

The first interesting case of this question concerns regular 3-uniform hypergraphs, where the answer turns out to be $t = 3$. This is tight, as the example of the Fano plane shows.

^{*}Department of Mathematics, Princeton University and Schools of Mathematics, Computer Science and AI, Tel Aviv University, Israel. Research supported in part by NSF grant DMS-2553988 and by AFOSR grant 26RT0135. Email: nalon@math.princeton.edu

[†]Department of Combinatorics and Optimization, University of Waterloo, Waterloo ON, Canada. Research supported in part by NSERC. Email: pehaxell@uwaterloo.ca

[‡]Department of Mathematics, ETH, Zürich, Switzerland. Research supported in part by SNSF grant 200021-228014. Email: aleksa.milojevic@math.ethz.ch

[§]Department of Mathematics, University of California, San Diego, CA, USA. Research supported in part by NSF grant DMS-2347832. Email: jacques@ucsd.edu

Theorem 1.1. *Let $\mathcal{H}$ be a nonempty regular 3-uniform hypergraph. Then, there is a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ such that $\deg_{\mathcal{C}} v \in \{1, 2, 3\}$ for all $v \in V(\mathcal{H})$.*

A natural way to extend the above theorem is to ask about higher uniformities. In this direction, for $r \geq 2$, we define $f(r)$ to be the smallest value of t for which any nonempty regular r -uniform hypergraph contains a spanning subhypergraph with all degrees between 1 and t . Tutte's theorem shows that $f(2) = 2$ and Theorem 1.1 shows that $f(3) = 3$. The following result determines the asymptotics of $f(r)$ for growing r .

Theorem 1.2. *Let $r \geq 2$ be an integer, and let $\mathcal{H}$ be a nonempty regular r -uniform hypergraph. Then, there exists a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ such that $1 \leq \deg_{\mathcal{C}} v \leq 100 \log r$. In other words $f(r) \leq 100 \log r$.*

Planken and Ueckerdt [17] have constructed examples which give $f(r) \geq \lfloor \log_2(r+1) \rfloor$, showing that Theorem 1.2 is tight up to a constant factor.

Further, note that we always have $f(r+1) \geq f(r)$. To see this, suppose that $\mathcal{H}$ is an r -uniform hypergraph without a spanning subhypergraph all of whose degrees are in $\{1, \dots, t\}$, for some t . Then, one can construct an $(r+1)$ -uniform hypergraph $\mathcal{H}'$ with the same property. This can be done as follows: denote $S = V(\mathcal{H})$ and add a new vertex outside S to each edge of $\mathcal{H}$. Then, add further vertices and edges disjoint from S until one obtains a regular $(r+1)$ -uniform hypergraph $\mathcal{H}'$. If $\mathcal{H}'$ has a spanning subhypergraph all of whose degrees are between 1 and t , then intersecting the edges of this hypergraph with S would give a subhypergraph of $\mathcal{H}$ with the same property, which we assumed did not exist.

One can also attempt to generalize Theorem 1.1 in a different direction. To explain how, let us go back to graphs for a moment, where we have already mentioned that any regular graph contains a $\{1, 2\}$ -factor. An even stronger statement is true: Tutte showed that for any $d \geq k \geq 1$, any d -regular graph contains a $\{k, k+1\}$ -factor (see also [20] for a very nice proof). Our next result shows that a similar statement is true for 3-uniform hypergraphs.

Theorem 1.3. *Let $d \geq k \geq 1$ be integers, and let $\mathcal{H}$ be a d -regular 3-uniform hypergraph. Then, there is a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ such that $\deg_{\mathcal{C}} v \in \{k, k+1, k+2\}$ for all $v \in V(\mathcal{H})$.*

If one tries to extend this result to higher uniformities, one encounters a significantly different situation from Theorem 1.2. To be precise, for $r \geq 2$, we define $g(r)$ to be the smallest value of t such that for any $d \geq k \geq 1$ and any d -regular r -uniform hypergraph $\mathcal{H}$, there exists a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ such that for every vertex $v \in V(\mathcal{H})$ we have $\deg_{\mathcal{C}} v \in \{k, k+1, \dots, k+t-1\}$. Again, Tutte's theorem [24] and Theorem 1.3 show that $g(2) = 2$ and $g(3) = 3$. Observe that, using the same argument as in the case of $f(r)$, one can show that g is a nondecreasing function, i.e. that $g(r+1) \geq g(r)$ for all r . The following result does not describe the precise asymptotics of $g(r)$, but shows that for large r it is much larger than $f(r)$.

Theorem 1.4. *For any $r \geq 3$, we have $\lfloor \sqrt{r/2} \rfloor \leq g(r) \leq 2r - 3$.*

Note that the upper bound $g(r) \leq 2r - 3$ immediately implies Theorems 1.1 and 1.3 by setting $r = 3$. Thus, we will only focus on proving Theorems 1.2 and 1.4. Our proofs of the upper bounds in the above theorems are algorithmic in nature, meaning that they provide an efficient procedure to find a subhypergraph with required degree conditions.

The paper [10] describing an AI-generated proof of the Komlós Conjecture was posted in the arXiv essentially simultaneously with our manuscript. Corollary 1.2 in that paper implies that the hereditary discrepancy of any set system in which every point belongs to at most r sets is smaller than $3\sqrt{2\pi r}$. Combining it with the known relation between the linear discrepancy and

the hereditary discrepancy of set systems (see, e.g., [4], Theorem 13.3.2) this implies that the function $g(r)$ satisfies

$$g(r) \leq 6\sqrt{2\pi r}.$$

This improves the bound in Theorem 1.4 for all $r \geq 60$, and shows that for large r , $g(r) = \Theta(\sqrt{r})$. The proof in [10] does not provide an efficient algorithm.

One can also extend Theorem 1.4 to arbitrary sets of allowed degrees. Namely, if $\mathcal{H}$ is a d -regular r -uniform hypergraph and every vertex v is assigned a list $H(v)$ of more than $\lceil \frac{r-1}{r}d \rceil$ allowed degrees between 0 and d , then a simple modification of the Shirazi-Verstraëte argument from [19] shows that there is a spanning subhypergraph $\mathcal{H}' \subseteq \mathcal{H}$ for which $\deg_{\mathcal{H}'} v \in H(v)$ for all vertices $v \in V(\mathcal{H})$. This proof is based on the Combinatorial Nullstellensatz [3] and so it does not provide an efficient algorithm for finding this subhypergraph.

During the work on this paper we also found a proof of the weaker result that for $r = 3$ one can find a subhypergraph with all degrees in $\{1, 2, 3, 4\}$, using a completely different argument that is ultimately based on topological methods. For completeness, we present this different perspective as well.

1.1 Related work

A spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ in which all degrees are between 1 and t is closely related to the notion of t -shallow hitting sets introduced by Keszegh and Pálvölgyi [12]. They were initially interested in the polycolorability of geometric hypergraphs, where the question is to determine for a given r -uniform hypergraph $\mathcal{H}$ (usually coming from a geometric setting) the maximum k for which vertices of $\mathcal{H}$ can be colored using k colors such that every edge of $\mathcal{H}$ receives all colors. A natural way to construct such a coloring is to find a set $S \subseteq V(\mathcal{H})$ which intersects every hyperedge in at least 1 and at most t elements. Setting these vertices aside as one color class, we are then left with a hypergraph in which all edges have size at least $r - t$, and if one can iterate this procedure, it is possible to obtain a coloring using $\lfloor r/t \rfloor$ colors in which every hyperedge contains all colors.

A set S intersecting every hyperedge in at least 1 and at most t vertices is then called a *t -shallow hitting set* (see [17]). A natural dual notion is a set of edges $F \subseteq E(\mathcal{H})$ which covers every vertex at least once and at most t times. We call such an F a *t -shallow hitting edge set*.

Planken and Ueckerdt [17] have studied when t -shallow hitting edge sets exist in a variety of settings, including the setting of regular r -uniform hypergraphs. In particular, they have shown that $f(r) \geq \lfloor \log_2(r+1) \rfloor$ and $f(r) \leq O(r)$. Theorem 1.2 improves the upper bound and shows that $f(r) \leq O(\log r)$, thus obtaining a bound which is tight up to a constant factor.

Shallow hitting edge sets seem to be useful for finding other spanning structures in dense hypergraphs, see e.g. [13].

1.2 Further questions

Let us now discuss some questions in similar spirit that we could not answer precisely.

Firstly, it seems possible that the conclusion of Theorem 1.1 could be strengthened if one restricts to tripartite hypergraphs (observe that the Fano plane, which shows the tightness of Theorem 1.1, is not tripartite).

Conjecture 1.5. *Let $\mathcal{H}$ be a nonempty regular tripartite 3-uniform hypergraph. Then, $\mathcal{H}$ contains a spanning subhypergraph whose vertices each have degree either 1 or 2.*

Although we are unable to show this statement in general, as partial evidence towards it we provide a simple argument based on Seymour's work [18] that proves the conjecture for 3-regular tripartite hypergraphs.

Secondly, recall that Theorem 1.1 shows that any regular 3-uniform hypergraph has a spanning subgraph all of whose degrees are in $\{1, 2, 3\}$. What can be said about regular 4-uniform hypergraphs? What is the value of $f(4)$? We know that $f(4) \leq 2 \cdot 4 - 3 = 5$ from Theorem 1.4 and $f(4) \geq f(3) = 3$ by the monotonicity properties of f .

While we can say little about this question in general, we note that an old result of Thomassen shows that any 4-regular 4-uniform hypergraph has property B (see Theorem 5.1 in [21]). The dual of this statement says that the edge set of any 4-regular 4-uniform hypergraph $\mathcal{H}$ can be partitioned into $E(\mathcal{H}) = E_1 \cup E_2$, where each vertex is covered between 1 and 3 times. In other words, there exists a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ for which $\deg_{\mathcal{C}} v \in \{1, 2, 3\}$. This suggests that we perhaps have $f(4) = 3$, although this may be too specific a case to be an indicator of general conclusions.

Paper organization. We prove the upper bound from Theorem 1.4 in Section 2, and we note that this immediately implies Theorem 1.1 and 1.3. We also prove Theorem 1.2 in Section 2. Then we present an argument proving Conjecture 1.5 for 3-regular hypergraphs in Section 3. Our alternative argument for the weaker version of Theorem 1.1 is given in the appendix.

2 Proofs of Theorems 1.2 and 1.4

In this section, we show that every d -regular r -uniform hypergraph contains a spanning subhypergraph with all degrees in the set $\{k, \dots, k + 2r - 4\}$, for any $k \leq d$. Let us now briefly outline the proof ideas needed for this result. Our proof of the upper bound $g(r) \leq 2r - 3$ is inspired by the proof of the Beck-Fiala theorem [5], which states that if a set system $\mathcal{S}$ on the ground set X has the property that no element is contained in more than n sets, then one can color the elements of X red and blue so that the difference between the number of red and blue elements in any set $A \in \mathcal{S}$ is at most $2n - 1$. The proof is based on the floating-variables method, in which every element $v \in X$ is assigned a real number $x_v \in [-1, 1]$, with all numbers starting at 0. Then, the variables are gradually changed with the goal of each of them reaching either 1 (at which point we color the element red and do not change x_v further) or -1 (at which point we color the element blue and do not change x_v further). Throughout this process, one maintains $\sum_{v \in A} x_v = 0$ for all $A \in \mathcal{S}$, in order to ensure that the number of red and blue elements is balanced in each set. For the full proof of the Beck-Fiala theorem, see e.g. Section 13 of [4].

Our proof of Theorem 1.4 proceeds along very similar lines, and we explain it for 3-uniform hypergraphs for the sake of concreteness. The goal is to show that any d -regular hypergraph contains a spanning subhypergraph with degrees in the set $\{k, k + 1, k + 2\}$, if $d \geq k$. We begin by assigning a variable $x_e \in [0, 1]$ to every edge $e \in E(\mathcal{H})$, starting at $x_e = \frac{k+1}{d}$ (observe that if $d = k$ or $d = k + 1$, then the statement is trivial, so we may assume $d \geq k + 2$, so $0 < x_e < 1$). We proceed to move these values until one of them reaches either 0 or 1, at which point the edge is either excluded or included in our subhypergraph $\mathcal{C}$ and the variable x_e is not moved further. As before, throughout this process of moving the variables, we ensure that $\sum_{e \ni v} x_e = k + 1$ for all vertices $v \in V(\mathcal{H})$.

As long as the number of free variables coming from edges with $x_e \in (0, 1)$ is larger than the number of constraints coming from vertices, we can find a direction to move the variables until one of them reaches either 0 or 1. Crucially, in order to reduce the number of constraints as the process progresses, we observe that any vertex with two nonfixed edges remaining will, no matter what values are assigned to these nonfixed edges, have between k and $k + 2$ incident edges at the

end. Thus, as soon as all but two edges incident to a vertex are fixed, we can forget about the constraint corresponding to this vertex. As a consequence of a simple double counting, the only place where this process can get stuck is if we are left with a 3-regular hypergraph, from which we can then choose the remaining edges directly by a simple procedure. Finally, Theorem 1.2 can be derived from the above statement using a simple application of the Lovász Local Lemma.

This method is quite robust, and it allows for the assumptions to be significantly relaxed. For example, it is sufficient to assume that the hypergraphs in question have rank r instead of uniformity r (in other words, we can allow for edges of cardinality smaller than r). Additionally, the only place where the regularity assumption is used is to construct a fractional spanning subgraph with all fractional degrees $k + r - 2$, which is an assignment of weights $w_e \in [0, 1]$ to all edges $e \in E(\mathcal{H})$ satisfying $\sum_{e \ni v} w_e = k + r - 2$ for all $v \in V(\mathcal{H})$. Hence, Theorems 1.1–1.4 apply to all hypergraphs with such a fractional spanning subgraph, or equivalently, all r -uniform hypergraphs with a perfect fractional matching in which all weights are at most $1/(k + r - 2)$.

Proof of the upper bound in Theorem 1.4. Given $\mathcal{H}$, finding a subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ that witnesses $g(r) \leq 2r - 3$ for $\mathcal{H}$ corresponds to finding a vector of weights $x \in \{0, 1\}^{E(\mathcal{H})}$ such that for every vertex $v \in V(\mathcal{H})$ we have $\sum_{e \ni v} x_e \in \{k, \dots, k + 2r - 4\}$ (where $x_e = 1$ means that the edge e is included in the cover $\mathcal{C}$ and $x_e = 0$ means that it is not). Note that we may assume that $d \geq k + 2r - 3$, as otherwise taking $\mathcal{C} = \mathcal{H}$ already satisfies the needed conditions.

We find these values using the following algorithm. We start by setting $x_e = \frac{k+r-2}{d}$ for each $e \in E(\mathcal{H})$ and we will vary these variables until we arrive at a $\{0, 1\}$ -valued vector satisfying $\sum_{e \ni v} x_e \in \{k, \dots, k + 2r - 4\}$ for each $v \in V(\mathcal{H})$.

The algorithm maintains a partition of the edge set into *active* and *fixed* edges, $E(\mathcal{H}) = E_a \cup E_f$, and a partition of the vertex set into *safe* and *unsafe* vertices, $V(\mathcal{H}) = V_s \cup V_u$.

The active edges will be the ones having $0 < x_e < 1$, and these values will be changed throughout the rest of the algorithm. On the other hand, when a variable x_e reaches value 0 or 1, it becomes fixed and remains fixed for the rest of the process. If a vertex v has at most $r - 1$ active incident edges, then we declare v to be *safe* until the end of the algorithm.

Throughout the algorithm, we maintain the following set of constraints.

$$\sum_{e \ni v} x_e = k + r - 2 \text{ for every } v \in V_u \text{ and } 0 \leq x_e \leq 1 \text{ for every } e \in E(\mathcal{H}).$$

In a general step of the algorithm, as long as the number of active edges exceeds the number of vertex constraints, i.e. as long as $|E_a| > |V_u|$, we do the following. We find a solution to the system of equations

$$\sum_{\substack{e \ni v \\ e \in E_a}} y_e = 0 \text{ for every } v \in V_u.$$

If $|E_a| > |V_u|$, this system is underdetermined and so there exists a nontrivial solution $(y_e)_{e \in E_a}$ to this system. We then take this solution, define $y_e = 0$ for all fixed edges and update the variables $x_e \mapsto x_e + \lambda y_e$ for some real number λ . Observe that the new values of x_e still satisfy $\sum_{e \ni v} x_e = k + r - 2$ for every $v \in V_u$. Moreover, to choose the value of λ , we start at $\lambda = 0$ and increase it until one (or more) of the variables x_e becomes 0 or 1. At this point, we fix x_e , we check if there is any vertex v which becomes safe, and proceed to the next step of the algorithm. The algorithm terminates when $|E_a| = |V_u|$.

Let us make a couple of observations about this algorithm. First of all, we observe that if a vertex v is declared to be safe, then no matter whether the remaining active edges incident to v are included in the cover or not, v will have degree between k and $k + 2r - 4$ in the cover. To see this, we argue based on how many edges incident to v have been fixed with $x_e = 1$, at the moment when v is declared to be safe. If all edges incident to v are fixed when v is declared

safe, then v will surely have degree $k + r - 2$ in the final subhypergraph. So, suppose that this is not the case.

Since at least one and at most $r - 1$ active edges are incident to v , we have $0 < \sum_{\substack{e \ni v \\ e \in E_a}} x_e < r - 1$. Moreover, this quantity is an integer, and thus $\sum_{\substack{e \ni v \\ e \in E_f}} x_e = k + r - 2 - \sum_{\substack{e \ni v \\ e \in E_a}} x_e \in \{k, \dots, k + r - 3\}$. In other words, between k and $k + r - 3$ fixed edges containing v have $x_e = 1$. Thus, no matter what happens with the remaining $r - 1$ edges, v will have degree between k and $k + 2r - 4$ in the resulting subhypergraph.

Let $\mathcal{H}'$ be the hypergraph of rank at most r on the vertex set V_u , with edges given by $e \cap V_u$ for $e \in E_a$. Let us show that when the algorithm terminates with $|E_a| = |V_u|$, the hypergraph $\mathcal{H}'$ is r -regular and r -uniform. Note that every vertex of V_u is incident to at least r edges of E_a , so the average degree is at least r . On the other hand, the average degree of $\mathcal{H}'$ can be calculated as follows:

$$\bar{d}(\mathcal{H}') = \frac{1}{|V_u|} \sum_{e \in E_a} |e \cap V_u| \leq \frac{r|E_a|}{|V_u|} \leq r.$$

Hence, we have the equality all the way through, meaning that each edge of E_a is completely contained in V_u and that $\mathcal{H}'$ is an r -regular, r -uniform hypergraph.

So, when the algorithm terminates, we have obtained a partial subhypergraph $\mathcal{C}_0$ consisting of those fixed edges having $x_e = 1$. Let $f_v = \deg_{\mathcal{C}_0} v$ denote the degree already assigned to the vertex v . It can be computed from the equation $f_v + \sum_{\substack{e \ni v \\ e \in E_a}} x_e = k + r - 2$.

To complete the proof, we need to decide which of the remaining active edges we add to the set $\mathcal{C}_0$ to form the final hypergraph $\mathcal{C}$ having the property that $\deg_{\mathcal{C}} v \in \{k, \dots, k + 2r - 4\}$. We do that as follows: if an edge has $x_e \geq 1/2$ when the above algorithm terminates, we include it in the final hypergraph $\mathcal{C}$, and otherwise we leave it out.

The number of additional edges incident to v differs from $\sum_{\substack{e \ni v \\ e \in E_a}} x_e$ by at most $r/2$, since the rounding of every edge contributes at most $1/2$ to the difference. Hence, the degree of v in the final hypergraph $\mathcal{C}$ satisfies

$$|\deg_{\mathcal{C}} v - (k + r - 2)| \leq \left| \deg_{\mathcal{C}} v - f_v - \sum_{\substack{e \ni v \\ e \in E_a}} x_e \right| \leq r/2 < r - 1.$$

Since $\deg_{\mathcal{C}} v$ is an integer, we conclude $\deg_{\mathcal{C}} v \in \{k, \dots, k + 2r - 4\}$, thus completing the proof. $\square$

We now discuss the proof of the lower bound in Theorem 1.4. It is a simple consequence of the following well-known lower bound construction in discrepancy theory (see e.g. Corollary 2.10 in [6]).

Theorem 2.1. *Let $\mathcal{H}$ be a projective plane of order q , i.e. a $(q + 1)$ -uniform $(q + 1)$ -regular hypergraph on the ground set $[q^2 + q + 1]$ for which each pair of elements is contained together in exactly one set. Then, no matter how the vertices of $[q^2 + q + 1]$ are colored red and blue, there is an edge of $\mathcal{H}$ which contains at least $\sqrt{q}$ more vertices of one color than the other.*

Recall that the *dual* of a hypergraph $\mathcal{H}$ is the hypergraph with vertex set $E(\mathcal{H})$ and edge set $\{\{e \ni v\} : v \in V(\mathcal{H})\}$.

Proof of the lower bound in Theorem 1.4. Let q be a prime between $r/2$ and $r - 1$, and let $\mathcal{H}$ be a projective plane of order q , which is a $(q + 1)$ -uniform $(q + 1)$ -regular hypergraph. By Theorem 2.1, no matter how one colors the vertices of $\mathcal{H}$ red and blue, there always exists an edge in which there are at least $\sqrt{q}$ more red points than blue points (or vice versa).

If $\mathcal{H}'$ is the dual of the hypergraph $\mathcal{H}$, note that its edges cannot be partitioned into two sets $E(\mathcal{H}') = E_1 \cup E_2$ such that each $\deg_{E_1} v \in \{(q+1)/2 - \lfloor \sqrt{q}/2 \rfloor, \dots, (q+1)/2 + \lfloor \sqrt{q}/2 \rfloor\}$ (otherwise, this would provide a balanced coloring of the vertices of the hypergraph $\mathcal{H}$). This shows that $g(r) \geq g(q+1) \geq \lfloor \sqrt{q} \rfloor \geq \lfloor \sqrt{r}/2 \rfloor$. $\square$

We close this section by giving the proof of Theorem 1.2, which states that regular r -uniform hypergraphs contain spanning subhypergraphs with all degrees between 1 and $100 \log r$. By passing to the dual, we see that Theorem 1.2 is equivalent to the following statement.

Proposition 2.2. *If $\mathcal{F}$ is an r -regular d -uniform hypergraph, then there exists a set of vertices $S \subseteq V(\mathcal{F})$ such that for any $e \in E(\mathcal{F})$ we have $1 \leq |e \cap S| \leq 100 \log r$.*

We split the proof of this proposition into two parts: first, we discuss what happens when $d \leq O(r)$, in which an application of Lovász's local lemma [9] suffices to prove the statement. When d is much larger than r , a simple application of Theorem 1.4 allows us to reduce to the previous case. In the case $d \leq O(r)$, we prove the following statement.

Lemma 2.3. *Let $\mathcal{F}$ be an r -regular hypergraph, with all edges having size between k and $3k$, for some $1 \leq k \leq r$. Then there exists a set $S \subseteq V(\mathcal{F})$ such that $1 \leq |e \cap S| \leq 100 \log r$ for every $e \in E(\mathcal{F})$.*

Proof. Let S be a random set of vertices, obtained by including every $v \in V(\mathcal{F})$ independently with probability $\frac{10 \log r}{k}$ (we may assume that $k \geq 10 \log r$, as otherwise one may take $S = V(\mathcal{F})$). We aim to show that with positive probability one has $1 \leq |e \cap S| \leq 100 \log r$ for every $e \in E(\mathcal{F})$.

Our main goal will be to show that for every fixed hyperedge $e \in E(\mathcal{F})$ we have

$$\mathbb{P}[1 \leq |e \cap S| \leq 100 \log r] \geq 1 - \frac{1}{10kr}.$$

Suppose we have shown this. Then, a simple application of the local lemma suffices to complete the proof as follows. For every $e \in E(\mathcal{F})$, define the bad event B_e as the complement of the event $1 \leq |e \cap S| \leq 100 \log r$. Since the event B_e only depends on whether the vertices of e are included in S or not, it is independent of all $B_{e'}$ with $e \cap e' = \emptyset$. So, the maximum degree of the dependency graph is at most $3k \cdot r$, since $|e| \leq 3k$ and every vertex of e is contained in at most r other edges. Since $e(3kr+1) \cdot \frac{1}{10kr} < 1$, the local lemma implies that with positive probability no bad event B_e occurs, showing that we have $1 \leq |e \cap S| \leq 100 \log r$ for every $e \in E(\mathcal{F})$.

So, it remains to bound $\mathbb{P}[1 \leq |e \cap S| \leq 100 \log r]$. To do this, note first that

$$\mathbb{P}[e \cap S = \emptyset] \leq \left(1 - \frac{10 \log r}{k}\right)^k \leq e^{-10 \log r} = r^{-10}.$$

On the other hand, by the union bound, for any t we have

$$\mathbb{P}[|e \cap S| \geq t] \leq \binom{3k}{t} \left(\frac{10 \log r}{k}\right)^t \leq \left(\frac{3ek}{t}\right)^t \left(\frac{10 \log r}{k}\right)^t \leq \left(\frac{3e \cdot 10 \log r}{t}\right)^t.$$

Putting $t = 100 \log r$ gives

$$\mathbb{P}[|e \cap S| \geq 100 \log r] \leq \left(\frac{3e \cdot 10 \log r}{100 \log r}\right)^{100 \log r} \leq 0.9^{100 \log r} \leq r^{-10},$$

where we have used that $30e/100 < 0.9$ and $0.9^{100} \leq e^{-10}$. By a simple union bound, we obtain that $\mathbb{P}[1 \leq |e \cap S| \leq 100 \log r] \geq 1 - 2r^{-10} \geq 1 - \frac{1}{10kr}$, which completes the proof as we already explained. $\square$

Let us now give the proof of Proposition 2.2 (and thus of Theorem 1.2, which is its dual).

Proof of Theorem 1.2 and Proposition 2.2. Let $\mathcal{H}$ be the given d -regular r -uniform hypergraph and let $\mathcal{F}$ be its dual. If we have $d \geq 3r$, then by Theorem 1.4 applied with $k = r$ we find a spanning subhypergraph $\mathcal{C} \subseteq \mathcal{H}$ in which all degrees are in $\{r, \dots, 3r - 4\}$. Taking its dual hypergraph $\mathcal{F}'$ corresponds to restricting the set of vertices of $\mathcal{F}$ such that all hyperedges have size between r and $3r - 4$. Hence, by Lemma 2.3 applied with $k = r$ we find a subset $S \subseteq V(\mathcal{F}')$ with the property that $1 \leq |S \cap e| \leq 100 \log r$ for all $e \in E(\mathcal{F}')$. The set of hyperedges in $\mathcal{H}$ dual to S then gives a subhypergraph in which every vertex $v \in V(\mathcal{H})$ has degree between 1 and $100 \log r$.

In the case $d \leq 3r$ we argue similarly, by directly applying Lemma 2.3 with $k = \lceil d/3 \rceil$ to obtain a set S of vertices for which $1 \leq |S \cap e| \leq 100 \log r$. As we noted above, considering its dual gives the subhypergraph we need in $\mathcal{H}$. $\square$

3 Tripartite 3-regular hypergraphs

In this section, we present an argument which proves Conjecture 1.5 for 3-regular hypergraphs.

Proposition 3.1. *Let $\mathcal{H}$ be a 3-regular tripartite 3-uniform hypergraph. Then, $\mathcal{H}$ has a spanning subhypergraph whose degrees are all either 1 or 2.*

As in Section 2, it is natural to consider the dual hypergraph $\mathcal{F}$, whose vertex set is $E(\mathcal{H})$ and whose edges are given by the collections $\{e : e \ni v\}$ for each vertex $v \in V(\mathcal{H})$. Finding a spanning subhypergraph of $\mathcal{H}$ whose degrees are all 1 or 2 corresponds to finding a subset of vertices $S \subseteq V(\mathcal{F})$ such that every hyperedge $e \in E(\mathcal{F})$ contains one or two of these vertices. In other words, it corresponds to finding a bicoloring of $V(\mathcal{F})$ where no edge of $\mathcal{F}$ is monochromatic, which is saying precisely that $\mathcal{F}$ has Property B.

To prove Proposition 3.1, we use a simple adaptation of Seymour's argument from [18], by exploiting the fact that $\mathcal{F}$ has at least two (in fact, three) disjoint perfect matchings whenever $\mathcal{H}$ is a tripartite hypergraph. The heart of the argument is the following claim.

Claim 3.2. *Let $\mathcal{F}$ be a hypergraph, and suppose that every proper subhypergraph of $\mathcal{F}$ has property B. Let us assign a vector $x_v \in \mathbb{R}^{|E(\mathcal{F})|}$ to each vertex of the hypergraph $\mathcal{F}$, where $x_v(e) = 1$ if $v \in e$ and $x_v(e) = 0$ otherwise. If the vectors x_v are linearly dependent, then $\mathcal{F}$ has property B. In particular, if $|E(\mathcal{F})| < |V(\mathcal{F})|$, then $\mathcal{F}$ has Property B.*

Proof. If the vectors x_v are linearly dependent, there exists a linear dependency $\sum_v \alpha_v x_v = 0$ with not all α_v equal to zero. In other words, for every edge $e \in E(\mathcal{F})$ we have $\sum_{v \in e} \alpha_v = 0$.

Let $R = \{v : \alpha_v < 0\}$ and $B = \{v : \alpha_v > 0\}$, and color the vertices of R red and vertices of B blue. Note that some vertices may have $\alpha_v = 0$ and thus remain uncolored. Since $\sum_{v \in e} \alpha_v = 0$, any edge containing a red vertex also must contain a blue vertex. Let U be the set of uncolored vertices, and let $\mathcal{F}[U]$ be the induced subhypergraph on this set. By assumption, $\mathcal{F}[U]$ can be bicolored without monochromatic edges. Extending the coloring using the sets R, B gives a bicoloring of $\mathcal{F}$ without monochromatic edges, as we wanted. $\square$

Lemma 3.3. *Let $\mathcal{F}$ be a 3-uniform hypergraph of maximum degree 3 with two edge-disjoint perfect matchings. Then, $\mathcal{F}$ has Property B.*

Proof. Suppose that $\mathcal{F}$ does not have Property B, and let $\mathcal{C} \subseteq \mathcal{F}$ be the minimal subgraph of $\mathcal{F}$ which also does not have Property B.

To start, let us show that $|E(\mathcal{C})| \leq |V(\mathcal{C})|$. Indeed, note that the average degree of $\mathcal{C}$ is at most 3, since this is the maximum degree of $\mathcal{C}$. On the other hand, the average degree of $\mathcal{C}$ is

at least $3|E(\mathcal{C})|/|V(\mathcal{C})|$, since $\mathcal{C}$ is a 3-uniform hypergraph. Combining these two bounds gives $|E(\mathcal{C})| \leq |V(\mathcal{C})|$.

On the other hand, since $\mathcal{C}$ is a minimal subhypergraph not having Property B, by Claim 3.2, we must have $|E(\mathcal{C})| \geq |V(\mathcal{C})|$. Hence, all of the inequalities of the previous paragraph must be equalities, and in particular all vertices of $\mathcal{C}$ must have degree 3 in $\mathcal{C}$. This means that no edge contains vertices of both $\mathcal{C}$ and $\mathcal{F} \setminus \mathcal{C}$, and so every perfect matching of $\mathcal{F}$ restricts to a perfect matching of $\mathcal{C}$. Hence, $\mathcal{C}$ has two distinct perfect matchings.

This implies that the vectors $\{x_v : v \in V(\mathcal{C})\}$ cannot be linearly independent, as follows. Consider the vertex-edge incidence matrix A of the hypergraph $\mathcal{C}$, which has the vectors x_v as rows. If the rows of this square matrix are linearly independent, so are the columns. In other words, if $y_e \in \mathbb{R}^{|V(\mathcal{C})|}$ is the indicator vector of the edge e , the vectors $\{y_e : e \in E(\mathcal{C})\}$ are linearly independent.

However, this is not the case: if M, M' are the perfect matchings of $\mathcal{C}$, then $\sum_{e \in M} y_e = \sum_{e \in M'} y_e$ is a nontrivial linear dependence between the columns of A (both sums give an all-ones vector since M, M' are perfect). Hence, the rows of A must also be linearly dependent, which is what we aimed to show. By Claim 3.2 this means that $\mathcal{C}$ has property B, contradiction. $\square$

Proof of Proposition 3.1. To prove Proposition 3.1, it suffices to observe that the dual $\mathcal{F}$ of a tripartite 3-uniform hypergraph $\mathcal{H}$ has three disjoint perfect matchings. Indeed, the vertices inside each of the three partite sets of $\mathcal{H}$ are incident to distinct edges, and thus they correspond to a perfect matching in the dual. Hence, Proposition 3.1 follows directly by dualizing Lemma 3.3. $\square$

Acknowledgment. Part of the research on this project was done during a visit of the first and last authors at the Institute of Mathematical Sciences, National University of Singapore in 2026. We thank our hosts at the Institute for their hospitality.

AI disclosure. We have tried to use AI tools to resolve Conjecture 1.5. Although this attempt was ultimately unsuccessful, it pointed us to Seymour’s argument from [18], which resolves the 3-regular case and which we decided to include in the paper.

References

- [1] L. Addario-Berry, K. Dalal, B. A. Reed, Degree constrained subgraphs, *Discrete Appl. Math.*, 156 (2008) 1168–1174.
- [2] R. Aharoni, P. E. Haxell, Hall’s theorem for hypergraphs, *J. Graph Theory*, 35 (2000) 83–88.
- [3] N. Alon, Combinatorial Nullstellensatz, *Combin. Probab. Comput.*, 8(1999) 7–29.
- [4] N. Alon, J. H. Spencer, *The Probabilistic Method*, third ed., Wiley-Interscience, Hoboken, NJ, 2008.
- [5] J. Beck, T. Fiala, “Integer-making” theorems, *Discrete Appl. Math.*, 3 (1981) 1–8.
- [6] J. Beck, V. T. Sós, Discrepancy theory, in R. L. Graham, M. Grötschel, L. Lovász (eds.), *Handbook of Combinatorics*, Vol. 2, Elsevier, Amsterdam (1995) 1405–1446.
- [7] C. Berge, M. Las Vergnas, On the existence of subgraphs with degree constraints, *Nederl. Akad. Wetensch. Indag. Math.*, 40 (1978) 165–176.
- [8] G. Cornuéjols, General factors of graphs, *J. Combin. Theory Ser. B*, 45 (1988) 185–198.

- [9] P. Erdős, L. Lovász, Problems and results on 3-chromatic hypergraphs and some related questions, in: Infinite and Finite Sets, Colloq. Math. Soc. János Bolyai, vol. 10, North-Holland, Amsterdam, 1975, pp. 609–627.
- [10] S. Guo, E. X. Fang and J. Lu, Vector Balancing via Directional Total Variation, arXiv:2609.11189v1, 2026.
- [11] P. Hall, On representatives of subsets, J. Lond. Math. Soc., 10 (1935) 26–30.
- [12] B. Keszegh, D. Pálvölgyi, An abstract approach to polychromatic coloring: shallow hitting sets in ABA-free hypergraphs and pseudohalfplanes, J. Comput. Geom., 10 (2019) 1–26.
- [13] R. Lang, N. Sanhueza-Matamala, On sufficient conditions for spanning structures in dense graphs, Proc. Lond. Math. Soc., 127 (2023) 709–791.
- [14] L. Lovász, Subgraphs with prescribed valencies, J. Combin. Theory, 8 (1970) 391–416.
- [15] L. Lovász, The factorization of graphs. II, Acta Math. Acad. Sci. Hungar., 23 (1972) 223–246.
- [16] J. Petersen, Die Theorie der regulären Graphs, Acta Math., 15 (1891) 193–220.
- [17] T. Planken, T. Ueckerdt, Shallow hitting edge sets in uniform hypergraphs, preprint, arXiv:2307.05757 (2023).
- [18] P. D. Seymour, On the two-colouring of hypergraphs, Q. J. Math., 25 (1974) 303–311.
- [19] H. Shirazi, J. Verstraëte, A note on polynomials and f -factors of graphs, Electron. J. Combin., 15 (2008) Paper No. N22.
- [20] C. Thomassen, A remark on the factor theorems of Lovász and Tutte, J. Graph Theory, 5 (1981) 441–442.
- [21] C. Thomassen, The even cycle problem for directed graphs, J. Amer. Math. Soc., 5 (1992) 217–229.
- [22] W. T. Tutte, The factorization of linear graphs, J. Lond. Math. Soc., 22 (1947) 107–111.
- [23] W. T. Tutte, The 1-factors of oriented graphs, Proc. Amer. Math. Soc., 4 (1953) 922–931.
- [24] W. T. Tutte, The subgraph problem, Ann. Discrete Math., 3 (1978) 289–295.

A An alternative proof of $f(3) \leq 4$

Here we provide an alternative proof that regular hypergraphs contain shallow edge hitting sets. We first present an argument which gives a subhypergraph with maximum degree 5, and then modify it to show that 5 can be reduced to 4.

This proof is reminiscent of the proof of Petersen’s theorem on 2-factors of even-regular graphs. However, since we are working with hypergraphs instead of graphs, we substitute Hall’s theorem with a consequence of the Aharoni–Haxell hypergraph matching theorem [2]. Before we state it, let us recall that a 3-uniform hypergraph $\mathcal{H}$ is called *bipartite* if there is a partition $V(\mathcal{H}) = A \cup B$ such that each edge $e \in E(\mathcal{H})$ has $|e \cap A| = 1$ and $|e \cap B| = 2$. Also, let us recall that the link graph of a set S is the graph L_S on the vertex set $V(\mathcal{H})$ in which xy is an edge if and only if there exists $s \in S$ such that $\{x, y, s\} \in E(\mathcal{H})$. We denote by $\nu(G)$ the maximum size of a matching in a graph G , and by $\tau(G)$ the minimum size of a vertex set that intersects every edge of G .

Theorem A.1 (Corollary 1.2 in [2]). *Suppose that $\mathcal{H}$ is a bipartite 3-uniform hypergraph. If every set $S \subseteq A$ has $\nu(L_S) \geq 2|S|$, then there exists a matching M in $\mathcal{H}$ saturating A .*

Proposition A.2. *If $\mathcal{H}$ is a nonempty regular 3-uniform hypergraph, then $\mathcal{H}$ has a spanning subhypergraph with degrees in the set $\{1, 2, 3, 4, 5\}$.*

Proof. Consider the following auxiliary hypergraph $\mathcal{A}$. The vertex set of $\mathcal{A}$ consists of five copies of $V(\mathcal{H})$, i.e. for each vertex $v \in V(\mathcal{H})$ we add the vertices $v^{(1)}, v^{(2)}, v^{(3)}, v^{(4)}, v^{(5)}$ to $V(\mathcal{A})$. The edge set of $\mathcal{A}$ is defined as follows: for each edge $e = \{x, y, z\} \in E(\mathcal{H})$, we add to $E(\mathcal{A})$ the following six edges

$$\begin{aligned} &\{x^{(1)}, y^{(2)}, z^{(3)}\}, \quad \{y^{(1)}, z^{(2)}, x^{(3)}\}, \quad \{z^{(1)}, x^{(2)}, y^{(3)}\}, \\ &\{x^{(1)}, y^{(4)}, z^{(5)}\}, \quad \{y^{(1)}, z^{(4)}, x^{(5)}\}, \quad \{z^{(1)}, x^{(4)}, y^{(5)}\}. \end{aligned}$$

Note that the resulting hypergraph $\mathcal{A}$ is tripartite, since every edge contains one vertex from each of the three sets $A = \{v^{(1)} : v \in V(\mathcal{H})\}$, $B = \{v^{(2)}, v^{(4)} : v \in V(\mathcal{H})\}$ and $C = \{v^{(3)}, v^{(5)} : v \in V(\mathcal{H})\}$. See the illustration in Figure 1. Further, since $\mathcal{H}$ is d -regular, every vertex of A has degree $2d$, while every vertex of B and C has degree d .




Figure 1: The auxiliary hypergraph $\mathcal{A}$.

The goal is to show that there exists a matching covering A in this auxiliary hypergraph. To show this, note that for any $S \subseteq A$, the link graph L_S is a bipartite multigraph between B and C with maximum degree d and $2|S|d$ edges, since every vertex of S contributes at least $2d$ edges to the link graph. Hence, the cover number of this bipartite graph must be at least $2|S|$, and so by König's theorem we have $\nu(L_S) = \tau(L_S) \geq 2|S|$. By Theorem A.1, this shows that $\mathcal{A}$ contains a matching M saturating A .

In other words, for every $v \in V(\mathcal{H})$, there exists a unique edge of M containing $v^{(1)}$, say $\{v^{(1)}, u^{(2)}, w^{(3)}\}$ (or $\{v^{(1)}, u^{(4)}, w^{(5)}\}$). We construct the subhypergraph $\mathcal{C}$ by adding the edge $\{v, u, w\}$ for each v . We claim that every vertex appears in at most 5 edges of $\mathcal{C}$. Indeed, if a vertex v appears in an edge of $\mathcal{C}$, it must be that one of its copies in $\mathcal{A}$ was present in one of the edges of the matching M . But there are 5 copies of v in $\mathcal{A}$, and so v can appear in at most 5 edges of $\mathcal{C}$. This shows that $\mathcal{C}$ has maximum degree at most 5, as we wanted. $\square$

Proposition A.3. *If $\mathcal{H}$ is a nonempty regular 3-uniform hypergraph, then $\mathcal{H}$ has a spanning subhypergraph with degrees in the set $\{1, 2, 3, 4\}$.*

Proof. We start from the subhypergraph $\mathcal{C}$ obtained in Proposition A.2 and remove some edges to reduce its maximum degree by one. Note that we treat $\mathcal{C}$ as a multihypergraph.

The main observation is that if a vertex v has $\deg_{\mathcal{C}} v = 5$, then we may delete the edge containing $v^{(1)}$ in the matching M (which we included in order to cover v) and reduce the degree of v in this way. However, one must be careful in which order the vertices v are chosen, in order to ensure that every vertex is still covered after several edges get deleted in this process.

We call the edge containing $v^{(1)}$ in the matching M the *special edge* of v , and we use the same term to refer to the corresponding edge in $\mathcal{C}$.

Let D be an auxiliary directed graph on the set of vertices of degree 5. We put a directed edge in D from u to v if v is contained in the special edge of u and $v \neq u$. The goal is to delete some special edges associated to the vertices of D , and we will use D to guide the order in which we do this.

The maximum outdegree in D is 2, and hence there exists a vertex v of indegree at most 2. We now delete the special edge e_v of the vertex v from $\mathcal{C}$.

We claim that even though e_v was deleted, v will still be covered at the end of this procedure. This is simple to see: since v had at most 2 inedges, at most two additional edges covering v will be deleted. Since v had appeared in 5 edges of $\mathcal{C}$, even after all the deletions, v will still be covered.

We further claim that no other vertex w of e_v is left uncovered. If w never had degree 5, then the special edge e_w will not be deleted from $\mathcal{C}$ and w will remain covered. If w had degree 5 and the special edge e_w was deleted in the past, we have just argued in the previous paragraph that w will remain covered.

We iterate this process on the remaining vertices that are covered at least 5 times, until all vertices are covered at most 4 times, which completes the proof. $\square$