

HYPERGRAPH CUTS ABOVE THE AVERAGE

BY

DAVID CONLON*

*Mathematical Institute, University of Oxford
Andrew Wiles Building, Radcliffe Observatory Quarter
Woodstock Road, Oxford, OX2 6GG, UK
e-mail: david.conlon@maths.ox.ac.uk*

AND

JACOB FOX** AND MATTHEW KWAN†

*Department of Mathematics, Stanford University, Stanford, CA 94305, USA
e-mail: jacobfox@stanford.edu, mattkwan@stanford.edu*

AND

BENNY SUDAKOV††

*Department of Mathematics, ETH Zürich—Zentrum
CH-8092 Zürich, Switzerland
e-mail: benjamin.sudakov@math.ethz.ch*

* Research supported by a Royal Society University Research Fellowship and by ERC Starting Grant 676632.

** Research supported by a Packard Fellowship and by NSF Career Award DMS-1352121.

† This research was done while the author was working at ETH Zurich.

†† Research supported in part by SNSF grant 200021-175573.

Received March 27, 2018 and in revised form October 22, 2018

ABSTRACT

An r -cut of a k -uniform hypergraph H is a partition of the vertex set of H into r parts and the size of the cut is the number of edges which have a vertex in each part. A classical result of Edwards says that every m -edge graph has a 2-cut of size $m/2 + \Omega(\sqrt{m})$ and this is best possible. That is, there exist cuts which exceed the expected size of a random cut by some multiple of the standard deviation. We study analogues of this and related results in hypergraphs. First, we observe that similarly to graphs, every m -edge k -uniform hypergraph has an r -cut whose size is $\Omega(\sqrt{m})$ larger than the expected size of a random r -cut. Moreover, in the case where $k = 3$ and $r = 2$ this bound is best possible and is attained by Steiner triple systems. Surprisingly, for all other cases (that is, if $k \geq 4$ or $r \geq 3$), we show that every m -edge k -uniform hypergraph has an r -cut whose size is $\Omega(m^{5/9})$ larger than the expected size of a random r -cut. This is a significant difference in behaviour, since the amount by which the size of the largest cut exceeds the expected size of a random cut is now considerably larger than the standard deviation.

1. Introduction

The **max-cut** of a graph G is the maximum number of edges in a bipartite subgraph of G . Equivalently, it is the maximum size of a cut, where a **cut** of a graph is a partition of its vertex set into two parts and the size of such a cut is the number of edges with one vertex in each part. The max-cut parameter has been studied extensively over the last 50 years, both from the algorithmic perspective emphasised in computer science and optimisation and from the extremal perspective taken in combinatorics. For a thorough (though now somewhat outdated) overview of the subject, see the survey of Poljak and Tuza [30].

In computer science, the problem of computing the max-cut of a graph already featured in Karp's famous list of 21 NP-complete problems [23]. However, efficient algorithms were subsequently found for computing the max-cut of certain restricted classes of graphs, including planar graphs [20, 28]. More recently, the problem has attracted attention due to the approximation algorithm of Goemans and Williamson [18] and its connections to the unique games conjecture [24].

In combinatorics, it is an important problem to estimate the max-cut of a graph in terms of its number of edges. As a first approximation, it is easy to see that the max-cut of any m -edge graph is at least $m/2$. Indeed, this is the expected size of a uniformly random cut. Moreover, though there are certain classes of graphs for which the max-cut is significantly larger (see, for example, [22]), the max-cut of an m -edge graph is typically quite close to $m/2$. As such, much of the focus has been on maximising the excess of a cut, defined to be its size minus $m/2$.

There are several interesting results about the excess of a graph with m edges. Answering a conjecture of Erdős, Edwards [11] proved that every m -edge graph has a cut with excess at least $(\sqrt{8m+1}-1)/8$. This bound is exact for complete graphs and so best possible when m is of the form $\binom{n}{2}$. Nevertheless, answering another conjecture of Erdős, Alon [2] showed that for infinitely many m this bound can be improved by $\Omega(m^{1/4})$. We refer the reader to [4, 7] for further results in this direction and to [2, 3, 5, 12, 29, 32] for improved bounds when the graph is known to be H -free.

For graphs without isolated vertices, it is also interesting to ask for bounds on the excess in terms of the number of vertices. Edwards [11] proved that every connected n -vertex graph has a cut with excess at least $(n-1)/4$. Erdős, Gyárfás and Kohayakawa [13] gave a simplified proof of this fact (and of the Edwards bound) and they also proved that every graph with no isolated vertices has a cut with excess at least $n/6$. Both of these bounds are best possible, attained by a clique and a disjoint union of triangles, respectively.

The goal of this paper is to extend these classical results, where possible, from graphs to hypergraphs. We now describe our results in detail.

1.1. MAX-CUT FOR HYPERGRAPHS. A **hypergraph** $H = (V, E)$ consists of a vertex set V and a collection E of subsets of V known as the edges of H . A **k -uniform hypergraph** (or **k -graph** for short) is a hypergraph where all of the edges have size exactly k . In particular, a 2-graph is simply a graph.

There are several possible notions of max-cut for hypergraphs that generalise the notion of max-cut for graphs. One obvious generalisation is to define a cut of a hypergraph H to be a partition of its vertex set into two parts and the size of such a cut to be the number of edges which have nonempty intersection with each of the two parts. As in graphs, we then define the max-cut of H to be the maximum size of a cut. An equivalent definition of the same notion is that the

max-cut of H is the size of its largest 2-colourable subgraph (see, for example, [26]). In computer science, the problem of computing this type of max-cut is commonly called max set splitting and the restriction to k -uniform hypergraphs is called max Ek -set splitting (see, for example, [19, 21]).

A second very natural generalisation to k -graphs H is to define a cut to be a partition of the vertices into k parts and the size of such a cut to be the number of edges with exactly one vertex in each part. The corresponding notion of max-cut can also be defined as the size of the largest k -partite subgraph of H (note that k -partite k -graphs have special significance in combinatorics, for example, in Ryser’s conjecture [1] and in hypergraph Turán theory [16, Section 14]).

One can also interpolate between these two notions of max-cut. For a k -graph H and $2 \leq r \leq k$, we define an r -cut of H to be a partition of the vertex set of H into r parts and the size of the r -cut to be the number of edges which have a vertex from every part (we say such edges are **multicoloured**). Then the **max- r -cut** of H is the maximum size of an r -cut. This generalised notion of hypergraph max-cut was first considered by Erdős and Kleitman [14], who observed that the expected size of a uniformly random r -cut is

$$\frac{S(k, r)r!}{r^k}m,$$

where $S(k, r)$ is a Stirling number of the second kind (the number of unlabelled partitions of $\{1, \dots, k\}$ into r nonempty sets). For the convenience of the reader, we remark that

$$S(k, 2)2!/2^k = 1 - 2^{1-k} \quad \text{and} \quad S(k, k)k!/k^k = k!/k^k$$

are the constants corresponding to the aforementioned special cases of 2-cuts in k -graphs and k -cuts in k -graphs. Just as for graphs, we define the **excess** of an r -cut in a k -graph to be its size minus $(S(k, r)r!/r^k)m$. We will chiefly be interested in proving Edwards-type bounds for r -cuts of a k -graph.

For fixed $2 \leq r \leq k$, there are at least two natural conjectures regarding the maximum excess of an r -cut of a k -graph. The first conjecture, inspired by the graph case, is that if $m = \binom{n}{k}$, then complete k -graphs have the smallest max- r -cut among all m -edge k -graphs. This conjecture was proposed by Scott [31, Problem 4.1] in the case $r = 2$. As a corollary of our Theorem 3.1, this would imply a lower bound of $\Omega(m^{(k-1)/k})$ on the maximum excess of an m -edge k -graph.

The second plausible conjecture is that, for any fixed $2 \leq r \leq k$, the smallest maximum r -cut over all m -edge k -graphs has excess $\Theta(\sqrt{m})$. There are several reasons why this is a natural guess. First, one of the most basic and effective means for proving bounds on max-cut in graphs is via the chromatic number, and, as we will see in Section 3, these methods generalise to prove a lower bound of $\Omega(\sqrt{m})$ on the maximum excess of a hypergraph cut. Second, the standard deviation of the size of a uniformly random cut can be as small as $\Theta(\sqrt{m})$. In particular, this is the case when H is a **linear** k -graph, meaning no pair of vertices is involved in more than one edge. Third, we have already seen that $\Theta(\sqrt{m})$ is the correct order of magnitude if $r = k = 2$ and, moreover, we will next see that it is also the correct order of magnitude when $r = 2$ and $k = 3$ (this disproves Scott’s conjecture in the case $k = 3$). A (k -uniform) **perfect matching** is a k -graph in which every vertex appears in exactly one edge. A **Steiner triple system** is a 3-graph in which every pair of vertices appears in exactly one edge. It is a very classical theorem in design theory [25] that an n -vertex Steiner triple system exists whenever $n \equiv 1$ or $n \equiv 3 \pmod 6$.

THEOREM 1.1: *Every m -edge 3-graph has a 2-cut with excess at least*

$$(\sqrt{24m + 1} - 1)/16.$$

Every n -vertex connected 3-graph H has a 2-cut with excess at least

$$(n - 1)/8.$$

Both of these bounds are best possible, being attained by any Steiner triple system.

If H is a 3-graph with n vertices, none of which are isolated, then it has a 2-cut with excess at least $n/12$. This bound is also best possible, being attained by a perfect matching.

Proof. It is easy to check (see [21]) that a 2-cut with size z in an m -edge 3-graph H corresponds precisely to a 2-cut with size $2z$ in the $3m$ -edge multigraph $G(H)$ obtained by replacing each edge of H with a triangle on the same vertex set. Each of the standard bounds for graph max-cut are also valid for multigraphs, so the desired lower bounds immediately follow. Now, observe that if H is a perfect matching then $G(H)$ is a disjoint union of triangles and if H is a Steiner triple system then $G(H)$ is a complete graph. These are precisely the graphs which are extremal for the standard graph max-cut theorems. ■

In light of Theorem 1.1, it is tempting to study highly regular sparse hypergraphs (such as those coming from design theory) as potential examples of k -graphs whose maximum excess of an r -cut is $\Theta(\sqrt{m})$. However, surprisingly, our main theorem shows that one can always find r -cuts of significantly greater excess, except when $r = 2$ and $k = 2$ or 3 .

THEOREM 1.2: *For any fixed $2 \leq r \leq k$ with $r \geq 3$ or $k \geq 4$, every m -edge k -graph has an r -cut with excess $\Omega(m^{5/9})$.*

We prove Theorem 1.2 in Sections 5 to 7. The idea is to first randomly choose the parts for some of the vertices in our hypergraph H , in such a way as to reduce the problem to finding a large 2-cut for a certain random hypergraph derived from H . Although the edges of this random hypergraph will not be independent, we will be able to reduce to the case where the dependencies are in a certain sense sparse, meaning that a typical small subset of vertices is not likely to witness any dependencies. The idea then is to split the vertex set into many such subsets, to greedily find a good cut for each of these vertex subsets individually, and then to combine these cuts randomly to give a cut of H with large excess. We remark that the general idea of partitioning a random graph into subsets with few dependencies also appeared in [6].

Regarding upper bounds, we show that random k -graphs can have smaller max- r -cut than complete graphs with the same number of edges. This disproves Scott's conjecture for $k > 3$.

THEOREM 1.3: *For $k > 3$, in the binomial random k -graph $\mathbb{G}_k(n, n^{3-k})$, a.a.s.¹ every r -cut has excess at most $O(n^2) = O(m^{2/3})$.*

As in graphs, it is also interesting to consider bounds for the excess of a k -graph r -cut in terms of the number of vertices. Using methods developed by Crowston, Fellows, Gutin, Jones, Rosamond, Thomassé and Yeo [10], Giannopoulou, Kolay and Saurabh [17] proved that n -vertex k -graphs which are partition connected (a stronger requirement than being connected) have 2-cuts with excess at least $(n-1)/(k2^{k-1})$. We give improved bounds in a much more general setting.

¹ By **asymptotically almost surely** or **a.a.s.**, we mean that the probability of the event is $1 - o(1)$. Asymptotics will always be as $n \rightarrow \infty$ or as $m \rightarrow \infty$ —it should be clear from the context which is meant. The constants implied by all asymptotic notation may depend on fixed values of k and r .

THEOREM 1.4: *For any fixed $2 \leq r \leq k$, consider an n -vertex k -graph H with no isolated vertices. Then H has an r -cut with excess $\Omega(n)$. In particular, if $k > 2$, then H has a 2-cut with excess at least $n/(k2^{k-1})$. This is best possible, being tight for a perfect matching.*

We prove Theorem 1.4 in Section 4 by analysing a variant of the method of conditional probabilities introduced by Erdős and Selfridge [15]. We note that the proof easily gives rise to a deterministic polynomial-time algorithm for constructing cuts of the required size.

1.2. STRUCTURE OF THE PAPER. The rest of the paper is organised as follows. In Section 2, we prove Theorem 1.3 and in Section 3, we give some simple lower bounds on max- r -cut based on the (strong) chromatic number. The bulk of the paper is then spent proving Theorems 1.2 and 1.4. We remark that these two proofs are independent of each other.

The proof of Theorem 1.4 is presented first, in Section 4. It is convenient to first show the required bound for max-2-cut, then use a simple reduction based on conditional expectations to deduce the general result for max- r -cut. Next, we turn to the proof of Theorem 1.2. Again, we make some reductions and then prove the theorem in a few special cases. In Section 5, we give some simple arguments to treat the case where H is “far” from being a linear hypergraph and we give a few different arguments which allow us to reduce our consideration to the case $k = r = 3$ and the case $r = 2$. Then, in Section 6, we treat the case of 3-cuts in almost-linear 3-graphs, while in Section 7 we treat the case of 2-cuts in almost-linear k -graphs. Both of these proofs follow the same basic strategy, but the former case is much simpler and can be viewed as a “warm-up”.

1.3. NOTATION. We use standard graph-theoretic notation throughout. The vertex set and edge set of a hypergraph H are denoted $V(H)$ and $E(H)$, respectively, and we write

$$e(H) = |E(H)|$$

for the number of edges in H . For a subset U of the vertex set of H , we write $H[U]$ for the subgraph of H **induced** by U , that is, the subgraph of H whose vertex set is U and whose edge set consists of all edges in $E(H)$ which are entirely contained in U .

Less standardly, a **mixed** k -graph is a hypergraph where all edges have size at most k and a **multihypergraph** is a hypergraph where multiple copies of each edge are allowed (formally, a multihypergraph H is defined by a vertex set $V(H)$ and a multiset $E(H)$ of subsets of $V(H)$). We also define k -multigraphs and mixed k -multigraphs, in the obvious way. In addition, for a (multi)hypergraph H , we let $G(H)$ be the multigraph obtained by replacing each edge $e \in E(H)$ with a clique $K_{|e|}$.

For a real number x , the floor and ceiling functions are denoted

$$\lfloor x \rfloor = \max\{i \in \mathbb{Z} : i \leq x\} \quad \text{and} \quad \lceil x \rceil = \min\{i \in \mathbb{Z} : i \geq x\}.$$

We will however omit floor and ceiling symbols whenever they are not crucial, for the sake of clarity of presentation. For real numbers x, y , we write $x \vee y$ to denote $\max\{x, y\}$ and $x \wedge y$ to denote $\min\{x, y\}$. All logs are base e .

Finally, we use standard asymptotic notation throughout, as follows. For functions $f = f(n)$ and $g = g(n)$, we write $f = O(g)$ to mean that there is a constant C such that $|f| \leq C|g|$, $f = \Omega(g)$ to mean that there is a constant $c > 0$ such that $f \geq c|g|$, $f = \Theta(g)$ to mean that $f = O(g)$ and $f = \Omega(g)$, and $f = o(g)$ to mean that $f/g \rightarrow 0$ as $n \rightarrow \infty$.

2. Upper bounds

Let $K_n^{(k)}$ be the complete k -graph on n vertices. In this section, we prove upper bounds on the max- r -cut in $K_n^{(k)}$ and in random k -graphs, in the process proving Theorem 1.3.

LEMMA 2.1: *Every r -cut in $K_n^{(k)}$ has excess at most $O(n^{k-1})$.*

Proof. Consider an r -cut with parts of size n_1, \dots, n_r , where $\sum_i n_i = n$. The size of this r -cut is

$$\sum_{(s_1, \dots, s_r)} \prod_{i=1}^r \binom{n_i}{s_i},$$

where the sum runs over all $(s_1, \dots, s_r) \in (\mathbb{N}^+)^r$ with $\sum_i s_i = k$. Recall that the expected size of a random r -cut is

$$\frac{S(k, r)r!}{r^k} \binom{n}{k} = \sum_{(s_1, \dots, s_r)} \binom{k}{s_1, \dots, s_r} \frac{1}{r^k} \binom{n}{k}.$$

Now, as in [14], one may verify that the max- r -cut is obtained when each $\lfloor n/r \rfloor \leq n_i \leq \lceil n/r \rceil$ (we say such a cut is **equitable**). Therefore, the excess of the r -cut is

$$\begin{aligned} \sum_{(s_1, \dots, s_r)} \left(\prod_{i=1}^r \binom{n_i}{s_i} - \binom{k}{s_1, \dots, s_r} \frac{1}{r^k} \binom{n}{k} \right) \\ = \sum_{(s_1, \dots, s_r)} \left(\prod_{i=1}^r \frac{1}{s_i!} \right) \left(\prod_{i=1}^r n_i^{s_i} - \frac{n^k}{r^k} \right) + O(n^{k-1}) \\ = O(n^{k-1}), \end{aligned}$$

as desired. ■

Now we can prove Theorem 1.3.

Proof of Theorem 1.3. Let $p = n^{3-k}$ and $G \in \mathbb{G}_k(n, p)$, so that a.a.s. G has $p \binom{n}{k} + o(n^2)$ edges. Consider any r -cut of $K_n^{(k)}$. By Lemma 2.1, it has size at most $(S(k, r)r! / r^k) \binom{n}{k} + O(n^{k-1})$, so the size of the corresponding r -cut in G is

$$p \left(\frac{S(k, r)r!}{r^k} \right) \binom{n}{k} + O(pn^{k-1} + t)$$

with probability $e^{-\Omega(t^2/pn^k)}$. If t is a sufficiently large multiple of n^2 then this probability is $o(r^{-n})$, which allows us to take a union bound over all r -cuts. ■

3. Basic chromatic number lower bounds

If an r -cut of a hypergraph H is very “imbalanced” (for example, if all but one part is empty), then its size is small. Therefore, a natural way to obtain an r -cut above the average size is to choose a random cut which is biased towards being “balanced”. For example, we could choose a uniformly random r -cut of H into parts of equal sizes, but it turns out that we can improve on this by choosing a random r -cut which is balanced in a certain way with respect to a fixed strong colouring of H .

We recall that a **strong colouring** of a (multi)hypergraph is a colouring in which all the vertices in each edge receive a distinct colour. The **strong chromatic number** of a (multi)hypergraph is then the minimum number of colours required for a strong colouring.

THEOREM 3.1: *For any fixed $2 \leq r \leq k$ and any n -vertex, m -edge k -multigraph H with strong chromatic number χ , there is an r -cut with excess $\Omega(m/\chi)$.*

Proof. Consider the cut obtained by splitting the colour classes in a strong χ -colouring of H randomly into r groups, each with equally many classes. (We can assume that χ is divisible by r , by adding at most $r - 1$ additional isolated vertices each in their own colour class.) We claim that the probability an edge e is multicoloured is

$$\frac{S(k, r)r!}{r^k} + \Omega(1/\chi),$$

which clearly suffices to prove the theorem. To prove our claim, consider some ordering on the k vertices of e (we assume the vertices are in fact the numbers $\{1, \dots, k\}$) and, for each $0 \leq i \leq k$, let $\omega_i : e \rightarrow \{1, \dots, r\}$ be the random (“partial”) cut where the parts of the first i vertices of e are coloured according to our special “balanced” random cut and the last $k - i$ vertices are coloured uniformly at random. We will prove that for each $1 \leq i \leq k$ the probability e is multicoloured with respect to ω_i is no less than the corresponding probability for ω_{i-1} , and that for $2 \leq i \leq k$ these probabilities differ by $\Omega(1/\chi)$.

First note that ω_0 and ω_1 actually have the same distribution. Next, note that ω_i and ω_{i-1} can be coupled to differ only on vertex i . Expose the value of each $\omega_i(v) = \omega_{i-1}(v)$, for $v \neq i$. With probability $\Omega(1)$, in e we have exposed a vertex in every part except some part j (if this situation does not occur, then it has already been determined whether e is multicoloured or not). Note that the conditional probability that e is multicoloured with respect to ω_{i-1} is $\Pr(\omega_{i-1}(i) = j) = 1/r$, whereas the corresponding probability for ω_i is

$$\Pr(\omega_i(i) = j) \geq \frac{\chi/r}{\chi - i + 1} \geq \frac{1}{r} + \Omega\left(\frac{i - 1}{\chi}\right).$$

This proves the desired claim. ■

As the strong chromatic number of H is trivially at most n , we have the following immediate corollary.

COROLLARY 3.2: *For any fixed $2 \leq r \leq k$ and any n -vertex, m -edge k -multigraph H , there is an r -cut with excess $\Omega(m/n)$.*

Next, the following corollary was mentioned in the introduction.

COROLLARY 3.3: *For any fixed $2 \leq r \leq k$ and any m -edge k -multigraph H , there is an r -cut with excess $\Omega(\sqrt{m})$.*

Proof. Let χ be the strong chromatic number of H and consider a strong colouring with χ colours. Since every pair of distinct colours must appear in an edge,

$$\binom{\chi}{2} \leq m \binom{k}{2}$$

and, hence, $\chi = O(\sqrt{m})$. ■

We end this section with a slight generalisation of Corollary 3.2. Recall that $G(H)$ is the multigraph obtained by replacing each edge $e \in E(H)$ with a clique $K_{|e|}$.

LEMMA 3.4: *For any fixed $2 \leq r \leq k$, consider a k -multigraph H . Suppose there is an n' -vertex subset W inducing m' edges of $G(H)$. Then H has an r -cut of excess $\Omega(m'/n')$.*

Proof. Let ω be the random r -cut where the parts of the vertices outside W are chosen uniformly at random and the parts of the vertices inside W are chosen with a uniformly random equipartition of W . Virtually the same proof as for Theorem 3.1 shows that for each of the $\Omega(m')$ edges e of H with $|e \cap W| \geq 2$, the probability that e is multicoloured is $S(k, r)r!/r^k + \Omega(1/n')$. Indeed, for each $0 \leq i \leq |e \cap W|$, let ω_i be the random “partial” cut where the first i vertices of $e \cap W$ are coloured according to ω , and the last $|e \cap W| - i$ are coloured uniformly at random. As in the proof of Theorem 3.1, the probability e is multicoloured with respect to ω_i is no less than the corresponding probability for ω_{i-1} , and for $2 \leq i \leq |e \cap W|$ these probabilities differ by $\Omega(1/n')$.

Also, all edges e of H with $|e \cap W| < 2$ are multicoloured in ω with probability exactly $S(k, r)r!/r^k$. Therefore, the expected excess of ω is $\Omega(m'/n')$. ■

4. Proof of Theorem 1.4

In this section, we prove Theorem 1.4. By first choosing which vertices are in the last $r - 2$ parts, we will be able to reduce the problem of finding a large-excess r -cut in a k -graph to the problem of finding a large 2-cut in a certain mixed k -graph. Note that we can generalise the notion of excess to a mixed multihypergraph H in an obvious way: if Z is the size of a uniformly random r -cut of H and an r -cut has size z , then we say the excess of that r -cut is $z - \mathbb{E}Z$. Theorem 1.4 will be a consequence of the following lemma.

LEMMA 4.1: *For any $k \geq 3$, consider an n -vertex mixed k -multigraph H with its vertices ordered. Let W be the set of vertices which, according to this ordering, are among the first two vertices of some edge of H which has size at least three. Then H has a 2-cut with excess at least $|W|/2^k$.*

Proof. We assume that the vertex set V consists of the integers $1, \dots, n$ with the natural ordering.

We greedily build a cut $\omega : V \rightarrow \{1, 2\}$ using the method of conditional probabilities, with the slight adaptation that we temporarily mark some vertices as “undetermined” if their choice of part has no effect on the subsequent conditional expectation. We will determine the part of an undetermined vertex v later, when we reach a vertex that shares an edge with v . The details are as follows.

Following our chosen order, we iteratively do the following for each vertex v . Let U be the set of undetermined vertices that precede v . For each edge e , let E_e be the event that e is multicoloured (2-coloured), conditioning on the choices $\omega(v')$ we have made so far for the determined vertices v' , and choosing the parts $\omega(v'')$ of the other vertices v'' uniformly at random. Let $Z_v = \sum_e \mathbb{1}_{E_e}$ be the size of this random cut (so Z_1 is just the unconditional size of a uniformly random cut). Say that an edge e is **uncertain** if $\Pr(E_e) \notin \{0, 1\}$ and let $U_v \subseteq U$ be the set of undetermined vertices $u \in U$ such that u and v share an uncertain edge. For each of the $2^{|U|+1}$ possible assignments of parts for the vertices in $U \cup \{v\}$, consider the conditional expectation of Z_v given that assignment. The average of these conditional expectations is $\mathbb{E}Z_v$. If all these conditional expectations are exactly $\mathbb{E}Z_v$, mark v as undetermined. Otherwise, for an assignment that maximizes this conditional expectation, fix the parts of the vertices in $U_v \cup \{v\}$ according to that assignment. The vertices in U_v are no longer undetermined. (We emphasise that we are considering an assignment of parts for all the vertices in $U \cup \{v\}$, including the vertices in $U \setminus U_v$, but we are only fixing the parts of the vertices in $U_v \cup \{v\}$.) Finally, at the end of the procedure, after going through all the vertices, choose the parts for the remaining undetermined vertices arbitrarily.

Now we show that this procedure produces a cut with the desired excess. First, we make the important observation that each of the conditional expectations compared during the algorithm is an integer multiple of $1/2^{k-1}$. Next, we prove the following claim.

CLAIM: At each step v of the process,

(1) there is an assignment A of the vertices in $U \cup \{v\}$ such that

$$\mathbb{E}[Z_v \mid A] \geq \mathbb{E}Z_v + \frac{|U_v|}{2^{k-1}};$$

(2) for any assignment A^v of the vertices in $U_v \cup \{v\}$ and any assignment A^U of the vertices in $U \setminus U_v$,

$$\mathbb{E}[Z_v \mid A^v, A^U] = \mathbb{E}[Z_v \mid A^v].$$

Proof. Consider some step v and assume (1) and (2) held for all preceding steps. First note that for any assignment A of the vertices in U , we have

$$(3) \quad \mathbb{E}[Z_v \mid A] = \mathbb{E}Z_v.$$

If $v - 1$ is undetermined, then this immediately follows from the definition of an undetermined vertex, and otherwise it follows from the fact that (2) held in step $v - 1$ (note that no $u \in U$ can have been in U_{v-1} or otherwise it would have become determined). Note also that there is no uncertain edge which involves more than one vertex of U , because the later vertex would have become determined at its step by (1).

Now, to prove (2), note that any conditional expectation of Z_v is the sum of the conditional probabilities of each edge being multicoloured. The set of uncertain edges containing a vertex in $U \setminus U_v$ is disjoint from the set of uncertain edges containing a vertex in $U_v \cup \{v\}$, so we can express Z_v as the sum of two terms Z_v^U and Z_v^v , the former of which depends on $(\omega(v'))_{v' \in U \setminus U_v}$ but not $(\omega(v'))_{v' \in U_v \cup \{v\}}$, and the latter of which depends on $(\omega(v'))_{v' \in U_v \cup \{v\}}$ but not $(\omega(v'))_{v' \in U \setminus U_v}$ (strictly speaking there is also a third term which doesn't depend on any vertex in U , but we can view this as being part of say Z_v^v). Applying (3), for any assignment A^U of the vertices in $U \setminus U_v$, we have

$$\mathbb{E}Z_v^U + \mathbb{E}Z_v^v = \mathbb{E}Z_v = \mathbb{E}[Z_v \mid A^U] = \mathbb{E}[Z_v^U \mid A^U] + \mathbb{E}Z_v^v.$$

So, $\mathbb{E}[Z_v^U \mid A^U] = \mathbb{E}Z_v^U$. Therefore, for any A^U, A^v as in the statement of (2),

$$\mathbb{E}[Z_v \mid A^v, A^U] = \mathbb{E}[Z_v^U \mid A^U] + \mathbb{E}[Z_v^v \mid A^v] = \mathbb{E}Z_v^U + \mathbb{E}[Z_v^v \mid A^v] = \mathbb{E}[Z_v \mid A^v],$$

proving (2). Finally, we prove (1) for step v . Since no edge contains multiple vertices of U , there are at least $|U_v|$ uncertain edges containing both v and a vertex $u \in U$. Let A_1 be the assignment of all vertices in U to part 1 and A_2

the assignment of all such vertices to part 2. Recalling (3), we will now give a lower bound for

$$\begin{aligned} & \mathbb{E}[Z_v \mid A_2, \omega(v) = 1] - \mathbb{E}[Z_v \mid A_1, \omega(v) = 1] \\ &= (\mathbb{E}[Z_v \mid A_2, \omega(v) = 1] - \mathbb{E}[Z_v \mid A_2]) - (\mathbb{E}[Z_v \mid A_1, \omega(v) = 1] - \mathbb{E}[Z_v \mid A_1]). \end{aligned}$$

To this end, for each edge e , define

$$P_e = (\Pr[E_e \mid A_2, \omega(v) = 1] - \Pr[E_e \mid A_2]) - (\Pr[E_e \mid A_1, \omega(v) = 1] - \Pr[E_e \mid A_1]).$$

We consider all possible cases for e , as follows:

- if e is not uncertain, then $P_e = 0 - 0 = 0$;
- if e does not contain v , then $P_e = 0 - 0 = 0$;
- if e does not contain a vertex of U , then

$$P_e = (\Pr[E_e \mid \omega(v) = 1] - \Pr[E_e]) - (\Pr[E_e \mid \omega(v) = 1] - \Pr[E_e]) = 0;$$

- if e contains v and a vertex $u \in U$, then:
 - if e contains no determined vertices, then

$$P_e \geq 2^{-(k-1)} - (-2^{-(k-1)}) \geq 2^{-(k-2)};$$

- if e contains a determined vertex in part 1, but no determined vertex in part 2, then

$$P_e \geq 0 - (-2^{-(k-2)}) \geq 2^{-(k-2)};$$

- if e contains a determined vertex in part 2, but no determined vertex in part 1, then

$$P_e \geq 2^{-(k-2)} - 0 \geq 2^{-(k-2)}.$$

We have shown that $P_e \geq 0$ and $P_e \geq 2^{-(k-2)}$ if e is uncertain and contains both v and a vertex $u \in U$. Therefore,

$$\mathbb{E}[Z_v \mid A_2, \omega(v) = 1] - \mathbb{E}[Z_v \mid A_1, \omega(v) = 1] \geq |U_v|/2^{k-2}.$$

This means that there is a choice $A \in \{A_1, A_2\}$, such that $\mathbb{E}[Z_v \mid A, \omega(v) = 1]$ differs by at least $|U_v|/2^{k-1}$ from $\mathbb{E}Z_v = \mathbb{E}[Z_v \mid A]$. This choice of A , and some choice of a part for v , proves (1). ■

Now, at the end of the procedure, let D be the set of vertices which did not get marked as undetermined during their step. Then the total number of determined vertices at the end of the procedure is $|D| + \sum_v |U_v|$, because each vertex in U_v becomes determined at step v by part (1) of the above claim.

Note that each vertex in W , being among the first two vertices in some edge e , becomes determined at the point when we reach the third vertex in e , at the latest. Indeed, when we reach the third vertex v of an edge e , either both of the first two vertices have already been determined, or else e has at most one determined vertex and is therefore uncertain (e could still either receive two different colours or not). In the latter case, U_v contains all the edges of e that were previously marked as undetermined, and therefore they become determined at this step. It follows that

$$|D| + \sum_v |U_v| \geq |W|.$$

Let Z_{n+1} be the eventual size of the 2-cut produced by the procedure. Recalling part (2) of the above claim, note that $\mathbb{E}Z_{v+1} = \max_A \mathbb{E}[Z_v \mid A]$, where the maximum is over all assignments A of the vertices in $U \cup \{v\}$. The expected excess of our 2-cut is

$$\mathbb{E}Z_{n+1} - \mathbb{E}Z_1 = \sum_{v=1}^n (\mathbb{E}Z_{v+1} - \mathbb{E}Z_v) = \sum_{v=1}^n (\max_A \mathbb{E}[Z_v \mid A] - \mathbb{E}Z_v).$$

Now, for each $v \in D$, we have $\max_A \mathbb{E}[Z_v \mid A] - \mathbb{E}Z_v \neq 0$, so, recalling that all conditional expectations are an integer multiple of $1/2^{k-1}$, the expected excess is at least $|D|/2^{k-1}$. By part (1) of the above claim, this expected excess is also at least $\sum_v |U_v|/2^{k-1}$. Taking the average of these two lower bounds, the expected excess is therefore at least

$$\frac{|D| + \sum_v |U_v|}{2^k} \geq \frac{|W|}{2^k}.$$

The desired result follows. ■

The 2-cut case of Theorem 1.4 is a very basic corollary of Lemma 4.1, as follows.

Proof of the 2-cut case of Theorem 1.4. Suppose that $k > 2$ and H is a k -multigraph with no isolated vertices. Consider a random ordering of the vertices of H . For any vertex v , there is an edge containing v , and the probability that v is among one of the first two vertices of that edge is $2/k$. Therefore, in the notation of Lemma 4.1,

$$\mathbb{E}|W| = \sum_v \Pr(v \in W) \geq (2/k)n.$$

Applying Lemma 4.1 with an ordering satisfying $|W| \geq 2n/k$, we conclude that there is a 2-cut with excess at least $n/(k2^{k-1})$, as desired. ■

Note that with essentially the same proof (considering a random ordering), we have a corresponding result for 2-cuts of mixed k -graphs.

LEMMA 4.2: *For any fixed $k \geq 3$, consider a mixed k -multigraph H such that at least n vertices are contained in edges of size at least three. Then H has a 2-cut with excess at least $n/(k2^{k-1})$.*

Combining this with the fact (mentioned in the introduction) that a multigraph with n non-isolated vertices has a cut with excess at least $n/6$, we have the following result, which we need below.

COROLLARY 4.3: *For any fixed $k \geq 2$, consider a mixed k -multigraph H such that at least n vertices are contained in edges of size k . Then H has a 2-cut with excess at least $n/(k2^k)$.*

Now we use Lemma 4.2 to prove Theorem 1.4 in the case where $r > 2$.

Proof of the $r > 2$ case of Theorem 1.4. Let $V = V(H)$. As we have seen, an r -cut can be represented as a function $\omega : V \rightarrow \{1, \dots, r\}$. Consider the ‘‘information restriction’’ function P , defined as follows. For any $\omega : V \rightarrow \{1, \dots, r\}$, let $P(\omega) = P\omega$ be the function $V \rightarrow \{*, 3, \dots, r\}$ with

$$P\omega(v) = \begin{cases} * & \text{if } \omega(v) \in \{1, 2\}, \\ \omega(v) & \text{otherwise.} \end{cases}$$

For $\rho : V \rightarrow \{*, 3, \dots, r\}$, let H_ρ^{part} be the mixed $(k - r + 2)$ -multigraph with edge multiset

$$\{e \cap \rho^{-1}(*) : e \in E(H), \rho(e) \supseteq \{3, \dots, r\}, |e \cap \rho^{-1}(*)| \geq 2\}$$

and vertex set $\rho^{-1}(*)$. Note that a 2-cut of H_ρ^{part} with size z corresponds to an r -cut ω of H with the same size z satisfying $P\omega = \rho$.

Now, let ω be a uniformly random r -cut of H and let Z be its size. Let $\mathbb{E}[Z \mid P\omega]$ be the conditional expectation of Z given $P\omega$, which is precisely the average size of a uniformly random 2-cut of $H_{P\omega}^{\text{part}}$. Let Q be the number of vertices in $H_{P\omega}^{\text{part}}$ which are contained in an edge of size $k - r + 2$. By Corollary 4.3, for any outcome of $P\omega$, the multihypergraph $H_{P\omega}^{\text{part}}$ has a 2-cut with size at least

$$\mathbb{E}[Z \mid P\omega] + \frac{Q}{((k - r + 2)2^{k-r+2})}.$$

Each vertex v of H is contained in at least one edge e of size k and the probability that e corresponds to a size- $(k - r + 2)$ edge in $H_{P\omega}^{\text{part}}$, containing v , is $\Omega(1)$. For example, this occurs if v and $k - r + 1$ other vertices are put in part 1, and the remaining $r - 2$ vertices are put in parts $3, \dots, r$. It follows that $\mathbb{E}Q = \Omega(n)$ and, therefore,

$$\mathbb{E}\left[\mathbb{E}[Z \mid P\omega] + \frac{Q}{(k - r + 2)2^{k-r+2}}\right] = \mathbb{E}Z + \Omega(n).$$

We deduce that there is an outcome of $P\omega$ such that $H_{P\omega}^{\text{part}}$ has a 2-cut with size $\mathbb{E}Z + \Omega(n)$, which corresponds to an r -cut of H with the same size, and therefore with excess $\Omega(n)$. ■

5. Reductions for the proof of Theorem 1.2

We now turn to the proof of Theorem 1.2. In this section we make a few simple reductions in preparation for Sections 6 and 7, which will comprise the heart of the proof of Theorem 1.2.

First, note that if H is very dense, then it trivially has a large-excess r -cut by Corollary 3.2. In fact, it is not hard to show that H has a large-excess r -cut whenever it is “far” from being a linear hypergraph in various senses. In Section 5.1 we collect a number of lemmas of this type.

Second, recall that in the proof of Theorem 1.4, we used a partial exposure trick to deduce a general theorem about r -cuts in k -graphs from a corresponding result for 2-cuts in mixed multihypergraphs. In Section 5.2 we make some reductions of this type for Theorem 1.2, essentially showing that it suffices to consider the case of 3-cuts in 3-multigraphs and the case of 2-cuts in mixed multihypergraphs.

By the end of the section, we will have shown that in order to prove Theorem 1.2, it suffices to prove two lemmas (Lemmas 5.8 and 5.9) concerning 3-cuts in almost-linear 3-graphs and concerning 2-cuts in almost-linear mixed hypergraphs. These lemmas will be proved in Sections 6 and 7.

5.1. REDUCTION TO THE ALMOST-LINEAR CASE. In a linear hypergraph, all degrees are $O(n)$ and the joint degree (or 2-degree) of any pair of vertices is at most 1. We first show how to reduce to the case where, except for a small subset

of bad vertices, the degrees and 2-degrees are not too large, that is, H is almost linear. We use notation of the form $\deg_H(u)$ and $\deg_H(u, v)$ to emphasise that degrees and 2-degrees are relative to H and not to any induced subgraph $H[U]$.

LEMMA 5.1: *For any fixed $2 \leq r \leq k$, consider an n -vertex, m -edge k -multi-graph H . For any q, g, Δ , at least one of the following holds: there is an r -cut with excess $\Omega(gq)$ or there is a vertex subset U with $|U| \geq n - 2q - km/\Delta$ such that for distinct $u, v \in U$, $\deg_H(v) \leq \Delta$ and $\deg_H(u, v) \leq g$.*

For the proof of Lemma 5.1 we will need the simple observation that, for a uniformly random r -cut in a multihypergraph, conditioning on the event that some disjoint edges are “partially multicoloured” increases the probability that an edge containing them is multicoloured.

LEMMA 5.2: *Consider a uniformly random r -cut of a multihypergraph H and, for any $f \subseteq V(H)$ and $\ell \leq r$, let $E_{f,\ell}$ be the event that f has vertices in at least ℓ different parts (so that $E_f := E_{f,r}$ is the event that f is multicoloured). Then, for any $e \in E(H)$ of size k , any disjoint $f_1, \dots, f_s \subseteq e$ of size at least 2, and any $\ell_1, \dots, \ell_s \geq 2$,*

$$\Pr \left(E_e \mid \bigcap_{i=1}^s E_{f_i, \ell_i} \right) = \frac{S(k, r)r!}{r^k} + \Omega(1).$$

Proof. We prove this by induction on s , assuming that it is true for all smaller values of s (so if $s = 1$ we are making no assumption). Let

$$E' = \bigcap_{i=1}^{s-1} E_{f_i, \ell_i}$$

(if $s = 1$ this event is the entire probability space). Either trivially or by the inductive hypothesis, we have $\Pr(E_e \mid E') \geq S(k, r)r!/r^k$.

Given this observation, it suffices to show that

$$\Pr(E_e \mid E' \cap E_{f_s, \ell_s}) > \Pr(E_e \mid E' \cap \overline{E_{f_s, \ell_s}}),$$

because each individual colouring has probability at least $1/r^k = \Omega(1)$. This is equivalent to showing that

$$\Pr(E_e \mid E_{f_s, \ell_s}) > \Pr(E_e \mid \overline{E_{f_s, \ell_s}})$$

in the conditional probability space where we are conditioning on E' . So, for the rest of the proof, we work in this conditional probability space. Let $f = f_s$ and $\ell = \ell_s$.

For each particular choice $\rho : f \rightarrow \{1, \dots, r\}$ of the parts of vertices in f , let A_ρ be the event that our random cut agrees with ρ on f . Then $E_{f,\ell}$ is the disjoint union of all A_ρ with $|\rho(f)| \geq \ell$ and $\overline{E_{f,\ell}}$ is the disjoint union of all $A_{\rho'}$ with $|\rho'(f)| < \ell$. Now, consider the individual conditional probabilities $\Pr(E_e | A_\rho)$. If $\rho'(f) \subsetneq \rho(f)$, then a simple coupling argument shows that $\Pr(E_e | A_{\rho'}) < \Pr(E_e | A_\rho)$. But, by symmetry, the probabilities $\Pr(E_e | A_\rho)$ only depend on $|\rho(f)|$.

For every ρ with $|\rho(f)| \geq \ell$ and every ρ' with $|\rho'(f)| < \ell$, we have $|\rho'(f)| < |\rho(f)|$, so by the above discussion

$$\Pr(E_e | A_{\rho'}) < \Pr(E_e | A_\rho).$$

Observe that $\Pr(E_e | E_{f,\ell})$ is the average of all the $\Pr(E_e | A_\rho)$ with $|\rho(f)| \geq \ell$ and, similarly, $\Pr(E_e | \overline{E_{f,\ell}})$ is the average of all the $\Pr(E_e | A_{\rho'})$ with $|\rho'(f)| < \ell$, so the desired conclusion follows. ■

Now we prove Lemma 5.1.

Proof of Lemma 5.1. Let $G_g(H)$ be the graph of pairs of vertices u, v such that $\deg_H(u, v) > g$. If $G_g(H)$ has an independent set of size $n - 2q$, then we can remove the (at most km/Δ) vertices with degree greater than Δ and we are done. Otherwise, $G_g(H)$ has a matching M of size q . Now, consider a uniformly random r -cut and condition on the event that each edge of the matching has its vertices in two different parts. By Lemma 5.2, due to the conditioning, each of the $\Omega(gq)$ edges of H which contains an edge of M is multicoloured with probability

$$\frac{S(k, r)r!}{r^k} + \Omega(1)$$

and the probability that every other edge is multicoloured is unaffected by the conditioning. Hence, the expected excess of this random r -cut is $\Omega(gq)$ and so there is a particular r -cut with this excess. ■

In addition to having few bad vertices, we would also like to know that a good proportion of the edges do not involve the bad vertices. For 2-cuts, this is easy to obtain.

LEMMA 5.3: Fix $k \geq 2$ and consider an n -vertex, m -edge k -multigraph H . Suppose U is a vertex subset with $|U| = n - b$. Then, at least one of the following holds: there is a 2-cut with excess $\Omega(m/b)$ or $e(H[U]) = \Omega(m)$.

Proof. Let $V = V(H)$. If there are $(1 - c)m$ edges intersecting both U and $V \setminus U$, for sufficiently small $c > 0$, then the 2-cut with parts U and $V \setminus U$ itself has excess $\Omega(m)$. Otherwise, at least one of U or $V \setminus U$ induces $\Omega(m)$ edges. If $V \setminus U$ induces $\Omega(m)$ edges, then $e(G(H)[V \setminus U]) = \Omega(m)$ as well, so, by Lemma 3.4, there is a 2-cut with excess $\Omega(m/b)$. ■

Next, we observe that the case of 3-cuts in 3-graphs reduces to the case of 2-cuts in 3-graphs.

LEMMA 5.4: *If a 3-multigraph has a 2-cut with excess x , then it has a 3-cut with excess $\Omega(x)$.*

Proof. Note that $S(3, 2)2!/2^3 = 3/4$ and $S(3, 3)3!/3^3 = 6/27$. Given a 2-cut of size $(3/4)m + x$, create a third part by including each vertex independently with probability $1/3$. The probability an edge that spans the 2-cut is multicoloured in the 3-cut is $2(1/3)(2/3)^2 = 8/27$, so the expected size of our random 3-cut is $(6/27)m + \Omega(x)$. ■

As a consequence, the next lemma immediately follows from Lemma 5.3.

LEMMA 5.5: *Consider an n -vertex, m -edge 3-multigraph H . Suppose U is a vertex subset with $|U| = n - b$. Then, at least one of the following holds: there is a 3-cut with excess $\Omega(m/b)$ or $e(H[U]) = \Omega(m)$.*

For general r -cuts in k -graphs, it is not clear how to reduce to the case where $H[U]$ has many edges, but we can reduce instead to the case where there are many edges of H which have all but at most one of their vertices in U .

LEMMA 5.6: *Fix $2 \leq r \leq k$ and consider an n -vertex, m -edge k -multigraph H . Suppose U is a vertex subset with $|U| = n - b$. Then, at least one of the following holds: there is an r -cut with excess $\Omega(m/b)$ or there are $\Omega(m)$ edges which have at least $k - 1$ of their vertices in U .*

Proof. Let $V = V(H)$. If $e(G(H)[V \setminus U]) = \Omega(m)$, then, by Lemma 3.4, there is an r -cut with excess $\Omega(m/b)$. Otherwise, almost all of the m edges have at most one of their vertices in $V \setminus U$. ■

To summarise, it only remains to treat the case where H contains a large subset U of vertices with the property that many edges are almost completely contained in U and the degrees and 2-degrees in U are not too high.

5.2. REDUCING TO SOME SPECIAL VALUES OF k AND r . Recall that in the proof of Theorem 1.4, we were able to reduce a general r -cut problem to a 2-cut problem by making a choice for the vertices in parts $3, \dots, r$ and then passing to an auxiliary mixed $(k - r + 2)$ -multigraph H^{part} . We will use the same idea here, but since 2-graphs and 3-graphs might only have small 2-cuts, it will only be useful when $r \leq k - 2$. To be more specific, we will use this idea to deduce the following lemma from another result.

LEMMA 5.7: *Fix $2 \leq r \leq k - 2$ and consider an m -edge, n -vertex k -multigraph H with a distinguished vertex subset U . Suppose that H has $\Omega(m)$ edges which have at least four of their vertices in U , $\deg_H(u) \leq \Delta$ for all $u \in U$, and $\deg_H(u, v) \leq g$ for all distinct $u, v \in U$. Let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3})$ and suppose $g = o(pm)$ and $g \log n = o(p\Delta)$. Then H has an r -cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta})$.*

Lemma 5.7 basically says that almost-linear k -graphs have large-excess r -cuts. Combined with the results of Section 5.1, this will suffice to prove the case $r \leq k - 2$ of Theorem 1.2. The details will be given at the end of this subsection. As promised, we will now use the H^{part} trick to show that Lemma 5.7 is a consequence of the following 2-cut lemma.

LEMMA 5.8: *Fix $k \geq 4$ and consider an m -edge, n -vertex mixed k -multigraph H with a distinguished vertex subset U . Suppose that $H[U]$ has $\Omega(m)$ edges of size at least four, $\deg_H(u) \leq \Delta$ for all $u \in U$, and $\deg_H(u, v) \leq g$ for all distinct $u, v \in U$. Let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3})$ and suppose $g = o(pm)$ and $g \log n = o(p\Delta)$. Then H has a 2-cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta})$.*

Proof of Lemma 5.7 given Lemma 5.8. We will use almost exactly the same approach as in the proof of Theorem 1.4. Recall the definition of the operator P and the mixed $(k - r + 2)$ -multigraphs H_ρ^{part} . (Informally, we remind the reader that P restricts the information of whether a vertex is in part 1 or 2: if $\omega(v) \in \{1, 2\}$ then $P\omega(v) = *$. The multihypergraph H_ρ^{part} on the vertex set $\rho^{-1}(*)$ is defined such that a 2-cut of H_ρ^{part} corresponds to an r -cut of H with the same size.)

Let ω be a uniformly random r -cut of H and let Z be its size. Let Q be the number of edges in $H_{P\omega}^{\text{part}}[U]$ of size at least four. By Lemma 5.8, there is $c = \Omega(\sqrt{p/\Delta})$ such that for each outcome of $P\omega$, the multihypergraph $H_{P\omega}^{\text{part}}$ has a 2-cut of size at least $\mathbb{E}[Z | P\omega] + cQ$ (and therefore H has an r -cut of this size).

Now, the probability that an edge in H with at least four vertices in U corresponds to a size-4 edge in $H_{P\omega}^{\text{part}}[U]$ is $\Omega(1)$, so $\mathbb{E}Q = \Omega(m)$ and $H_{P\omega}^{\text{part}}$ has a 2-cut of size at least

$$\mathbb{E}[\mathbb{E}[Z \mid P\omega] + cQ] = \mathbb{E}Z + \Omega(cm).$$

This corresponds to an r -cut of H which has the same size and, therefore, has excess $\Omega(cm) = \Omega(\sqrt{pm}/\sqrt{\Delta})$. ■

We next use a slight variation of the H^{part} trick to directly show that the $r \geq k - 1$ case of Theorem 1.2 follows from the $r = k = 3$ case.

Proof of the $r \geq k - 1$ case of Theorem 1.2 given the $r = k = 3$ case. First, consider the case where $k = r \geq 3$. In this case we redefine P and H_{ρ}^{part} as follows. For any $\omega : V \rightarrow \{1, \dots, k\}$, let $P(\omega) = P\omega$ be the function $V \rightarrow \{*, 4, \dots, k\}$ with

$$P\omega(v) = \begin{cases} * & \text{if } \omega(v) \in \{1, 2, 3\}, \\ \omega(v) & \text{otherwise.} \end{cases}$$

For $\rho : V \rightarrow \{*, 4, \dots, k\}$, let H_{ρ}^{part} be the 3-multigraph with edge multiset

$$\{e \cap \rho^{-1}(*) : e \in E(H), \rho(e) \supseteq \{4, \dots, k\}, |e \cap \rho^{-1}(*)| = 3\}$$

and vertex set $\rho^{-1}(*)$. Note that a 3-cut of H_{ρ}^{part} with size z corresponds to an r -cut ω of H with the same size z satisfying $P\omega = \rho$.

We now proceed in essentially the same way as in the proofs of Theorem 1.4 and Lemma 5.7. Let ω be a uniformly random k -cut of H , let Z be its size, and let $Q = e(H_{P\omega}^{\text{part}})$. By assumption, there is $c = \Omega(1)$ such that $H_{P\omega}^{\text{part}}$ has a 3-cut with size $\mathbb{E}[Z \mid P\omega] + cQ^{5/9}$.

The probability that an edge in H corresponds to an edge in $H_{P\omega}^{\text{part}}$ is $(k!/k^k)/(3!/3^3) = \Omega(1)$, so $\mathbb{E}Q = \Omega(m)$ and, applying Markov's inequality to $m - Q$, we have

$$\Pr(Q < \mathbb{E}Q/2) = \Pr(m - Q > (1 + \Omega(1))\mathbb{E}[m - Q]) = 1 - \Omega(1).$$

This implies that $\mathbb{E}Q^{5/9} \geq \Pr(Q \geq \mathbb{E}Q/2)(\mathbb{E}Q/2)^{5/9} = \Omega(m^{5/9})$. So, $H_{P\omega}^{\text{part}}$ has a 2-cut of size at least

$$\mathbb{E}[\mathbb{E}[Z \mid P\omega] + cQ^{5/9}] = \mathbb{E}Z + \Omega(m^{5/9}).$$

This corresponds to an r -cut of H with the same size and, therefore, with excess $\Omega(m^{5/9})$, as desired.

Finally, it remains to consider the case where $r = k - 1$, for $k \geq 4$. In this case we observe that our k -uniform r -cut problem corresponds to an r -multigraph r -cut problem in much the same way that 3-graph 2-cut problems correspond to multigraph cut problems (recall that this was important for Theorem 1.1). Let H' be the km -edge r -multigraph obtained from H by taking each k -edge and replacing it with an r -graph k -clique on its vertex set. Now, there is essentially only one way that an edge of H can be multicoloured (with a single repeated part), so an r -cut of H' with size z corresponds precisely to an r -cut of H with size $z/2$. Since we have already proved that $\Omega(m)$ -edge r -multigraphs have r -cuts with excess $\Omega(m^{5/9})$, the desired result follows. ■

To prove the $r = k = 3$ case of Theorem 1.2, we will combine the lemmas in Section 5.1 with the following result.

LEMMA 5.9: *Consider an m -edge, n -vertex 3-multigraph H with a distinguished vertex subset U . Suppose that $H[U]$ has $\Omega(m)$ edges, $\deg_H(u) \leq \Delta$ for all $u \in U$, and $\deg_H(u, v) \leq g$ for all distinct $u, v \in U$. Let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3})$ and suppose $g = o(pm)$ and $g \log n = o(p\Delta)$. Then H has a 3-cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta})$.*

We now show how to deduce the $r = k = 3$ case and the $r \leq k - 2$ case of Theorem 1.2 from Lemmas 5.7 and 5.9 and the lemmas in Section 5.1.

Proof of Theorem 1.2 given Lemmas 5.7 and 5.9. We can assume H has no isolated vertices (otherwise we could delete them without changing the max-cut or the number of edges). This means $n = O(m)$. Define

$$\Delta = m^{5/9}, \quad g = m^{3/5} \Delta^{-4/5} = m^{7/45} \quad \text{and} \quad q = \sqrt{m/g} = m^{19/45}.$$

Note that $m^{5/9} = o(gq)$, so by Lemma 5.1 it suffices to consider the case where there is a set U with $|U| \geq n - O(q + m/\Delta)$ such that, for all distinct $u, v \in U$, $\deg(v) \leq \Delta$ and $\deg(u, v) \leq g$.

If $r = k = 3$, we may now apply Lemma 5.5. If H has a 3-cut with excess $\Omega(m/(q + m/\Delta)) = \Omega((m/q) \wedge \Delta) = \Omega(m^{5/9})$, then we are done. Otherwise, $e(H[U]) = \Omega(m)$. In this case, let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3}) = m^{-1/3}$. Note that $g = m^{7/45} = o(pm)$ and $p\Delta = m^{2/9}$, so $g \log n = o(p\Delta)$. We may therefore apply Lemma 5.9, from which it follows that H has a 3-cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta}) = \Omega(m^{5/9})$.

If $r \leq k - 2$, we may proceed in basically the same way. If $r = 2$ (therefore $k \geq 4$), apply Lemma 5.3: if H has a 3-cut with excess $\Omega(m/(q + m/\Delta))$, we are done, so we may assume $e(H[U]) = \Omega(m)$. Otherwise, if $r \geq 3$ (therefore $k \geq 5$), apply Lemma 5.6: if H has an r -cut with excess $\Omega(m/(q + m/\Delta))$, we are done, so we may assume H has $\Omega(m)$ edges with at least $k - 1 \geq 4$ of their vertices in U . In either case, we may then apply Lemma 5.7 with $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3}) = m^{-1/3}$ to prove that H has an r -cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta}) = \Omega(m^{5/9})$. ■

It remains to prove Lemmas 5.8 and 5.9.

6. 3-cuts in almost-linear 3-multigraphs

In this section, we prove Lemma 5.9. The idea is to consider a uniformly random 3-cut but to expose only the information about whether each vertex is in part 3 or not, meaning that for vertices which are not chosen to be in part 3, we are free to choose whether they should be in part 1 or part 2. (This is exactly as in the proofs of Theorem 1.4 and Lemma 5.7.) We thereby reduce the problem of finding a large-excess 3-cut in H to the problem of finding a large 2-cut in the random multigraph G^{part} obtained by considering each hyperedge of H which has exactly one vertex in part 3, deleting that vertex, and adding the resulting pair of vertices as an edge in G^{part} .

Now, we expect that in a random multigraph we can find a large-excess 2-cut in a greedy fashion: we simply go through the vertices one-by-one and place each vertex in the part that maximizes the number of multicoloured edges between that vertex and the preceding vertices. Doing this will result in at least half the edges being multicoloured, with equality only if at each step both choices of the part for the current vertex are equally good. Since we have some variation due to randomness, it's unlikely that both choices will be equally good at most steps.

We cannot immediately use this method to find a large-excess cut of the random multigraph G^{part} , because the edges of G^{part} are not completely independent. In particular, if two edges of H intersect, then the edges of G^{part} that can arise from H depend on each other. But since we are assuming H is almost linear, there are not too many of these dependencies, so if we look at the subgraph of G^{part} induced by a small random set of vertices, we expect most of

the edges in this induced subgraph to be independent. Therefore, we can find a large-excess 2-cut in many small induced subgraphs of G^{part} and then randomly combine these 2-cuts to obtain a large-excess 2-cut of G^{part} itself.

In preparation for proving Lemma 5.9, we need a few lemmas. First, we will need to be able to split our vertex set into many small subsets such that the subgraph of G^{part} induced by each of these vertex sets has many edges but few dependencies.

Definition 6.1: Consider a 3-multigraph H with a distinguished vertex subset U . Consider a partition $\mathcal{V} = (V_1, \dots, V_t)$ of U . We say that \mathcal{V} is (m', Δ') -good with respect to (H, U) if:

- (i) $\bigcup_i G(H[U])[V_i]$ has at least m' edges;
- (ii) $\bigcup_i G(H)[V_i]$ has maximum degree at most Δ' ;
- (iii) no edge of H is completely contained in any V_i ;
- (iv) for each i , any two edges $e_1, e_2 \in E(H)$ that both intersect V_i in two vertices must satisfy $e_1 \cap e_2 \cap U \setminus V_i = \emptyset$, unless they are coincident edges on the same set of 3 vertices.

We say that \mathcal{V} is y -almost (m', Δ') -good with respect to (H, U) if properties (i) and (ii) hold, and if at most $y/2$ edges of H fail to satisfy property (iii) and at most $y/2$ pairs of edges of H fail to satisfy property (iv).

To explain property (iv), the idea is that if $e_1, e_2 \in E(H)$ both intersect V_i in two vertices, then the third vertex of e_1 is different to the third vertex in e_2 , except for complications caused by multiple coincident edges and complications caused by the few high-degree vertices outside U .

Next, the following lemma allows us to find an almost-good partition in the setting of Lemma 5.9 where degrees and codegrees are controlled.

LEMMA 6.2: *Consider an m -edge, n -vertex 3-multigraph H with a distinguished vertex subset U . Suppose that $H[U]$ has $\Omega(m)$ edges, $\deg_H(u) \leq \Delta$ for all $u \in U$, and $\deg_H(u, v) \leq g$ for all distinct $u, v \in U$. Let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3})$ and suppose $g = o(pm)$ and $g \log n = o(p\Delta)$. Then, for any $c > 0$, there are $m' = \Omega(pm)$ and $\Delta' = O(p\Delta)$ such that there exists a $(cm'/\sqrt{\Delta'})$ -almost (m', Δ') -good partition $\mathcal{V} = (V_1, \dots, V_t)$ with respect to (H, U) .*

To prove Lemma 6.2, we may simply take a random partition and show that it satisfies the desired almost-goodness property with positive probability. We

defer the details of the proof to Section 7, where we will prove a generalisation, Lemma 7.3, of Lemma 6.2.

The next lemma allows us to combine large-excess cuts on many vertex subsets into a single large-excess cut.

LEMMA 6.3: *Consider a graph G with m edges and suppose there are disjoint vertex subsets V_1, \dots, V_t , with V_i inducing a graph G_i with a 2-cut of excess x_i . Then G has a 2-cut with excess at least $\sum_i x_i$.*

Proof. We can assume V_1, \dots, V_t form a partition of $V(G)$, adding singleton parts if necessary. Consider a cut $\omega_i : V_i \rightarrow \{1, 2\}$ of each G_i with excess x_i . For each i , let $\sigma_i : \{1, 2\} \rightarrow \{1, 2\}$ be a uniformly random permutation and define the random cut $\omega : V(G) \rightarrow \{1, 2\}$ by $\omega(v) = \sigma_i(\omega_i(v))$ for $v \in V_i$. Among the

$$m' = e\left(\bigcup_i G[V_i]\right)$$

edges within the parts, $\sum_i (e(G[V_i])/2 + x_i) = m'/2 + \sum_i x_i$ of them are multicoloured, while the probability every other edge is multicoloured is $1/2$. Therefore, the expected excess of ω is $\sum_i x_i$, so there must be an outcome of ω with at least this excess. ■

Next, we need the fact that random graphs have large-excess cuts. This is a corollary of a more general result, Lemma 7.6, which we will prove in Section 7.

LEMMA 6.4: *Fix constant $p \in (0, 1)$ and let G and F be multigraphs on the same set of n vertices such that G has m edges and maximum degree Δ . Suppose, moreover, that the edges of G are partitioned into groups, each only containing edges between the same two vertices. Let R be the random multigraph obtained by including each group independently with probability p and let X^{\max} be the maximum excess of a 2-cut of $F \cup R$. Then*

$$\mathbb{E}X^{\max} = \Omega(m/\sqrt{\Delta}).$$

Now we can prove Lemma 5.9.

Proof of Lemma 5.9. Let $V = V(H)$. Recall the operator P and the multihypergraphs H_ρ^{part} defined in the proof of Theorem 1.4. (Informally, we remind the reader that P restricts the information of whether a vertex is in part 1 or 2: if $\omega(v) \in \{1, 2\}$, then $P\omega(v) = *$.) In the case of 3-cuts in 3-graphs H_ρ^{part} is

always a multigraph, so we write $G_\rho^{\text{part}} = H_\rho^{\text{part}}$. As outlined at the beginning of this section, G_ρ^{part} may be obtained by considering each hyperedge of H which has exactly one vertex in part 3, deleting that vertex, and adding the resulting pair of vertices as an edge in G_ρ^{part} . Recall that a 2-cut of G_ρ^{part} with size z corresponds to a 3-cut ω of H with the same size z satisfying $P\omega = \rho$.

Consider some small $c > 0$ and apply Lemma 6.2 to find a cx -almost $(2m', \Delta')$ -good partition $\mathcal{V} = (V_1, \dots, V_i)$ with respect to (H, U) , where $m' = \Omega(pm)$, $\Delta' = O(p\Delta)$ and $x = m'/\sqrt{\Delta'} = \Omega(\sqrt{pm}/\sqrt{\Delta})$. Delete at most cx edges of H (causing corresponding changes in the G_ρ^{part}) so that \mathcal{V} becomes an (m', Δ') -good partition with respect to (H, U) . We will show that the resulting 3-graph H has a 3-cut with excess $\Omega(x)$, where the implied constant does not depend on c . This suffices to prove the lemma, because, for sufficiently small $c > 0$, restoring the deleted edges will not have too significant an impact on the excess of the 3-cut.

Let ω be a uniformly random 3-cut of H , so that $G_{P\omega}^{\text{part}}$ is a random multigraph. For each i let

$$V'_i = V_i \cap V(G_{P\omega}^{\text{part}}) = \{v \in V_i : \omega(v) \in \{1, 2\}\}.$$

The point of the good partition \mathcal{V} is that while $G_{P\omega}^{\text{part}}$ has dependencies between its edges, if we condition on some V'_i and individually look at the induced subgraph $G_{P\omega}^{\text{part}}[V'_i]$, it is a random subgraph of $G(H)[V'_i]$ with essentially independent edges. This is because for every edge $e \in E(G(H)[V'_i])$ (arising from the edge $e \cup \{v_e\} \in E(H)$, say), the presence of e in $G_{P\omega}^{\text{part}}[V'_i]$ solely depends on the random value $P\omega(v_e)$. By property (iii) of goodness the v_e are outside V'_i and by property (iv), except for repetitions due to multiple coincident edges, all the v_e inside U are distinct. That is to say, each of the edges of $G(H[U])[V'_i]$ have their own independent sources of randomness. Let

$$G_i^{\text{good}} = G(H[U])[V'_i]$$

and let G_i^{bad} contain the edges e in $G(H)[V'_i]$ which are not in G_i^{good} , meaning that $v_e \in V \setminus U$ (these v_e may not be distinct). See Figure 1 for an illustration of the above considerations.

For each i , let R_i^{good} (respectively, R_i^{bad}) be the random multigraph of edges $e \in E(G_i^{\text{good}})$ (respectively, $e \in E(G_i^{\text{bad}})$) such that $P\omega(v_e) = 3$. Taking into account the above discussion,

$$G_{P\omega}^{\text{part}}[V'_i] = R_i^{\text{good}} \cup R_i^{\text{bad}}$$

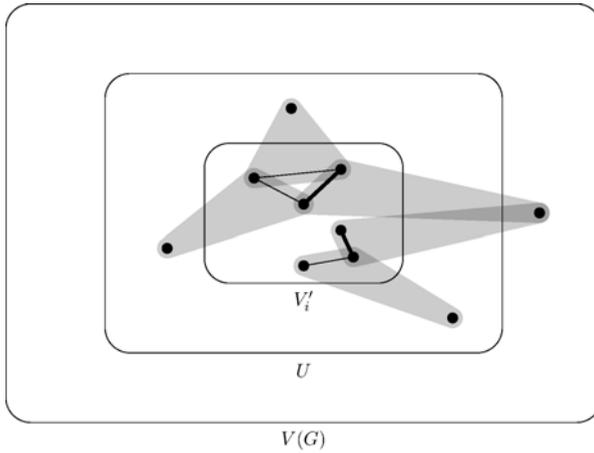


Figure 1. A close-up picture of some V'_i . Each edge of H intersecting V'_i in two vertices contributes an edge to $G(H)[V'_i]$. The thinner edges correspond to $G_i^{\text{good}} = G(H[U])[V'_i]$ and the thicker edges correspond to G_i^{bad} . Crucially, all edges $e \in E(G_i^{\text{good}})$ have different v_e . Note that some of the edges in $G(H[U])[V'_i]$ may be multiple coincident edges, but if so the corresponding edges of H also coincide.

and if we condition on V'_i , then R_i^{good} and R_i^{bad} are independent. Also, let X_i be the maximum excess of a 2-cut of $G_{P\omega}^{\text{part}}[V'_i]$ and write $\mathbb{E}[X_i \mid V'_i]$ for the conditional expectation of X_i given the random set V'_i . We will next use Lemma 6.4 to estimate $\mathbb{E}[X_i \mid V'_i]$. Towards this end, condition on a particular outcome of V'_i , so that in what follows we may consider it a fixed set of vertices. Since R_i^{good} and R_i^{bad} are independent, we may also condition on any particular outcome of R_i^{bad} , without affecting the distribution of R_i^{good} . By property (ii) of goodness, we are now in a position to apply Lemma 6.4, with $F = R_i^{\text{bad}}$, $G = G_i^{\text{good}}$, $R = R_i^{\text{good}}$, $p = 1/3$, $\Delta = \Delta'$ and the groups corresponding to edges of G_i^{good} coming from multiple coincident edges of H on the same 3 vertices. It follows that

$$\mathbb{E}[X_i \mid V'_i] = \Omega\left(\frac{e(G(H[U])[V'_i])}{\sqrt{\Delta'}}\right).$$

The above is true for any outcome of V'_i . Now, we return to viewing V'_i as a random set (depending on $P\omega$, where ω is a uniformly random 3-cut of H). Each edge of $G(H[U])[V'_i]$ is present in $G(H[U])[V_i]$ with probability $(2/3)^2 = \Omega(1)$, so, by property (i) of goodness, with $X = \sum_i X_i$ we have

$$\mathbb{E}X = \sum_i \mathbb{E}[\mathbb{E}[X_i \mid V'_i]] = \sum_i \Omega\left(\frac{e(G(H[U])[V_i])}{\sqrt{\Delta_i}}\right) = \Omega\left(\frac{m'}{\sqrt{\Delta_i}}\right) = \Omega(x).$$

By Lemma 6.3, $G_{P\omega}^{\text{part}}$ has a cut with size at least $e(G_{P\omega}^{\text{part}})/2 + X$. Each of the $3m$ edges of $G(H)$ corresponds to an edge of $G_{P\omega}^{\text{part}}$ with probability $(1/3)(2/3)^2 = 4/27$, so $\mathbb{E}e(G_{P\omega}^{\text{part}}) = (4/9)m$ and there is an outcome of $P\omega$ such that $G_{P\omega}^{\text{part}}$ (and therefore H) has a cut of size at least

$$\mathbb{E}[e(G_{P\omega}^{\text{part}})/2 + X] = (2/9)m + \Omega(x).$$

Observing that a random 3-cut of H has expected size $(2/9)m$, we are done. ■

7. 2-cuts in almost-linear hypergraphs

In this section, we prove Lemma 5.8. Although the proof is overall more complicated than the proof of Lemma 5.9, there is one simplifying reduction that we can make in this case: we do not need to worry about the “bad” vertices outside U .

LEMMA 7.1: *To prove Lemma 5.8, it suffices to consider the case where $U = V(H)$.*

To prove Lemma 7.1 we will again use a variant of the H^{part} trick, but this time there will be an extra twist. The basic idea is to consider a random 2-cut, but only expose the parts of the vertices outside U . One might hope that the problem of completing this partial cut to a large cut of H can be directly formulated as a max-cut problem on some auxiliary multihypergraph H^{part} with vertex set U . We would then hope to apply the $V(H) = U$ case of Lemma 5.8 and show that the resulting large-excess cut of H^{part} corresponds, on average, to a large-excess cut of H . Unfortunately, fixing the parts of vertices outside U introduces some asymmetry: there are now some edges e which so far have vertices only in one specific part, meaning that for e to be multicoloured we need to put one of the vertices of $e \cap U$ in the other part. This asymmetry cannot be expressed in a max-cut problem.

We overcome this issue with a simple averaging trick. Observe that the choices $\phi : U \rightarrow \{1, 2\}$ of parts for vertices in U can be paired up in a natural way: we pair ϕ with its “opposite” $\bar{\phi} : U \rightarrow \{1, 2\}$ defined by

$$\bar{\phi}(u) = 3 - \phi(u).$$

Now, given a choice for the parts of the vertices outside U , instead of trying to find $\phi : U \rightarrow \{1, 2\}$ completing our 2-cut such that the size of the resulting cut of H is maximised, we aim to find $\phi : U \rightarrow \{1, 2\}$ such that the average size of the two cuts completed by ϕ and by $\bar{\phi}$ is maximised. This problem can in fact be formulated as a 2-cut problem, as follows. For an edge e intersecting U , which so far has vertices only in one specific part, there are two relevant possibilities for $\phi(e \cap U)$. Either $\phi(e \cap U) = \{1, 2\}$, in which case e is multicoloured in both of the two cuts completed by ϕ and by $\bar{\phi}$, or else $|\phi(e \cap U)| = 1$, in which case e will be multicoloured in exactly one of the two completions given by ϕ and $\bar{\phi}$. That is to say, the contribution of e to the average size of the two cuts under consideration depends only on whether $e \cap U$ is multicoloured by ϕ .

Of course, there are also edges e of H which are completely contained in U and, for these edges, the averaging has no effect: either e is multicoloured by both ϕ and $\bar{\phi}$ or it is multicoloured by neither. Our auxiliary multihypergraph H^{part} will therefore consist of edges of two different types: first, the edges in $H[U]$, and second, derived edges $e \cap U$ of the type described in the previous paragraph. These two types of edges need to effectively be assigned different weights, because they contribute differently to the average size of a pair of opposite cuts. We will accomplish this by making multiple copies of the edges in $H[U]$. The details follow.

Proof of Lemma 7.1. We will prove Lemma 5.8 assuming its statement holds for any multihypergraph H and $U = V(H)$. Let $V = V(H)$ and let $\bar{U} = V \setminus U$. For a 2-cut $\omega : V \rightarrow \{1, 2\}$, let $P_\omega : \bar{U} \rightarrow \{1, 2\}$ be the restriction of ω to \bar{U} .

For $\rho : \bar{U} \rightarrow \{1, 2\}$, let $H'_\rho \subseteq H$ consist of those edges e in H but not in $H[U]$ with $|\rho(e \setminus U)| < 2$ and $|e \cap U| > 0$. This means the edges e in H'_ρ are precisely those edges that are not contained in U , such that given the restricted information P_ω it is not yet determined whether e is multicoloured by ω . Let H_ρ^{part} be the mixed k -multigraph with vertex set U , containing two copies of each of the edges in $H[U]$ and also containing an edge $e \cap U$ for each $e \in E(H'_\rho)$. Let N_ρ^{multi} be the number of edges of H which can already be seen to be multicoloured by ω given the information $P_\omega = \rho$.

Now, consider any $\rho : \bar{U} \rightarrow \{1, 2\}$ and $\phi : U \rightarrow \{1, 2\}$. Let z^U be the size of ϕ as a 2-cut of $H[U]$, let z' be the size of ϕ as a 2-cut of H'_ρ , and let

$$z^{\text{part}} = 2z^U + z'$$

be the size of ϕ as a 2-cut of H_ρ^{part} . Let $\omega : V \rightarrow \{1, 2\}$ be the 2-cut which is equal to ρ outside U and equal to ϕ inside U , and let $\bar{\omega} : V \rightarrow \{1, 2\}$ be the 2-cut which is equal to ρ outside U and equal to $3 - \phi$ inside U .

For each edge $e \in E(H'_\rho)$ whose corresponding edge in H_ρ^{part} is not already multicoloured by ϕ , e is multicoloured by exactly one of ω and $\bar{\omega}$. Therefore, the average size of ω and $\bar{\omega}$, as 2-cuts of H , is

$$(4) \quad z^U + z' + (e(H'_\rho) - z')/2 + N_\rho^{\text{multi}} = z^{\text{part}}/2 + e(H'_\rho)/2 + N_\rho^{\text{multi}}.$$

Now, let $\omega : V \rightarrow \{1, 2\}$ be a uniformly random 2-cut of H , let Z be the size of this 2-cut, and let Z^{part} be the size of ω restricted to U as a 2-cut of $H_{P\omega}^{\text{part}}$. Taking the expectation of (4) we have

$$(5) \quad \mathbb{E}[Z \mid P\omega = \rho] = \mathbb{E}[Z^{\text{part}} \mid P\omega = \rho]/2 + e(H'_\rho)/2 + N_\rho^{\text{multi}}.$$

Conditioning on the event $P\omega = \rho$, note that the restriction of ω to U is a uniformly random 2-cut of H_ρ^{part} , so $\mathbb{E}[Z^{\text{part}} \mid P\omega = \rho]$ is the average size of a random 2-cut of H_ρ^{part} . Combining (4) and (5), if H_ρ^{part} has a 2-cut with some excess x' , then H has a 2-cut with excess at least

$$(6) \quad (x' + \mathbb{E}[Z^{\text{part}} \mid P\omega = \rho])/2 + e(H'_\rho)/2 + N_\rho^{\text{multi}} - \mathbb{E}Z \\ = x'/2 + \mathbb{E}[Z \mid P\omega = \rho] - \mathbb{E}Z.$$

Finally, observe that $H_{P\omega}^{\text{part}}$ always contains (two copies of) all $\Omega(m)$ edges of $H[U]$ with size at least four and that there is an outcome ρ of $P\omega$ such that

$$\mathbb{E}[Z \mid P\omega = \rho] \geq \mathbb{E}Z.$$

Since $U = V(H_\rho^{\text{part}})$, by assumption, H_ρ^{part} has a 2-cut with excess $\Omega(\sqrt{pm}/\sqrt{\Delta})$, corresponding to a 2-cut of H with at least half this excess. ■

Now, our proof of Lemma 5.8 will follow the same general approach as for Lemma 5.9, but with a few added complications. Before we discuss the proof further, we give some generalisations of the lemmas in Section 6. First, we generalise the notion of goodness from Definition 6.1.

Definition 7.2: Consider a mixed k -multigraph H and a (not necessarily spanning) subgraph H' of H . Consider a partition $\mathcal{V} = (V_1, \dots, V_t)$ of $V(H')$. We say that \mathcal{V} is (m', Δ') -**good** with respect to (H, H') if:

- (i) $\bigcup_i G(H')[V_i]$ has at least m' edges;
- (ii) $\bigcup_i G(H)[V_i]$ has maximum degree at most Δ' ;
- (iii) each edge e of H has vertices in at least $|e \cap V(H')| - 1$ different parts of \mathcal{V} ;
- (iv) there is no pair of edges $e_1, e_2 \in E(H)$ and part V_i such that both e_1 and e_2 have two vertices in V_i , comprising a total of at least 3 vertices, and also e_1 and e_2 intersect in $V(H') \setminus V_i$.

We say that \mathcal{V} is y -**almost** (m', Δ') -good with respect to (H, H') if properties (i) and (ii) hold, and if at most $y/2$ edges of H fail to satisfy property (iii) and at most $y/2$ pairs of edges of H fail to satisfy property (iv).

Note that for a 3-multigraph H , if \mathcal{V} is $(y$ -almost) (m', Δ') -good with respect to $(H, H[U])$, then \mathcal{V} is $(y$ -almost) (m', Δ') -good with respect to (H, U) as defined in Definition 6.1. Therefore, the following lemma is a generalisation of Lemma 6.2.

LEMMA 7.3: *Fix k and consider an m -edge, n -vertex mixed k -multigraph H . Let H' be a (not necessarily spanning) subgraph of H with $\Omega(m)$ edges. Suppose that $\deg_H(u) \leq \Delta$ for all $u \in V(H')$ and $\deg_H(u, v) \leq g$ for all distinct $u, v \in V(H')$. Let $p = \Delta^{-3/5} \wedge (g^{-2/3} \Delta^{-1/3})$ and suppose $g = o(pm)$ and $g \log n = o(p\Delta)$. Then, for any $c > 0$, there are $m' = \Omega(pm)$ and $\Delta' = O(p\Delta)$ such that there exists a $(cm'/\sqrt{\Delta'})$ -almost (m', Δ') -good partition $\mathcal{V} = (V_1, \dots, V_t)$ with respect to (H, H') .*

Proof. Let $t = 1/(c'p)$ for some small $c' > 0$ (we will assume that c' is small enough to satisfy certain inequalities later in the proof). Let $\mathcal{V} = (V_1, \dots, V_t)$ be a random partition of $V(H')$ obtained by choosing the part of each vertex uniformly at random, so that the probability a vertex falls into a particular part is $p' := c'p$. We will show that \mathcal{V} satisfies the necessary almost-goodness properties with positive probability.

Let M be the number of edges in $\bigcup_i G(H')[V_i]$. Then we have $\mathbb{E}M = p'e(G(H'))$ and $\mathbb{E}M^2 \leq ((p'e(G(H'))))^2 + p'q$, where $q = O(mg)$ is the number of pairs of (possibly not distinct) edges in $G(H')$ on the same two vertices. (Note that for any two edges of $G(H')$ which do not coincide, their presence in $\bigcup_i G(H')[V_i]$ is

independent.) Since we are assuming $g = o(pm)$, we have $\text{Var } M = o((\mathbb{E}M)^2)$, so, by Chebyshev's inequality, property (i) holds a.a.s. with $m' = p'e(H')/2$.

For each vertex $v \in V(H')$ and $1 \leq i \leq t$, let

$$Q_v^i = \sum_{w \neq v} \deg_H(v, w) \mathbb{1}_{w \in V_i},$$

which is an upper bound on the number of edges of $G(H)$ between v and a vertex of V_i . Note that $\mathbb{E}Q_v^i \leq p'k\Delta$ and

$$\sum_{w \neq v} \text{Var}(\deg_H(v, w) \mathbb{1}_{w \in V_i}) \leq p' \sum_{w \neq v} \deg_H(v, w)^2 = O(p\Delta g),$$

and note that each $|\deg_H(v, w) \mathbb{1}_{w \in V_i}|$ is bounded by g . If we set

$$\Delta' = 2p'k\Delta = \Theta(p\Delta),$$

Bernstein's inequality (see, for example, [9, Equation (2.10)]) implies that

$$\Pr(Q_v^i > \Delta') \leq \exp\left(-\frac{(\Delta')^2}{O(p\Delta g + g\Delta')}\right) = e^{-\Omega(p\Delta/g)} = o(1/n^2).$$

(Here we have used the assumption that $g \log n = o(p\Delta)$.) Therefore, by the union bound, a.a.s. $Q_v^i \leq \Delta'$ for each v, i . That is, property (ii) a.a.s. holds.

Next, the number of pairs of edges of H which share multiple vertices in $V(H')$ is $O(mg)$. The probability that such a pair violates property (iv) is $O((p')^2)$. The number of pairs of edges of H which share only one vertex in $V(H')$ is $O(m\Delta)$ and the probability that such a pair violates property (iv) is $O((p')^3)$. Therefore, recalling the definition of p , the expected number of pairs of edges violating property (iv) is

$$O(mg(p')^2 + m\Delta(p')^3) \leq c(m'/\sqrt{\Delta'})/8$$

for c' sufficiently small in terms of c . So, with probability at least $3/4$, the number of such edges is at most $c(m'/\sqrt{\Delta'})/2$.

Finally, we consider (iii). The expected number of edges f of H which do not have vertices in at least $|f \cap V(H')| - 1$ different parts of \mathcal{V} is

$$O((p')^2 m) \leq c(m'/\sqrt{\Delta'})/8$$

(this inequality is not tight). Therefore, with probability at least $3/4$, the number of edges violating property (iii) is at most $c(m'/\sqrt{\Delta'})/2$. ■

We will also need a generalisation of Lemma 6.3. In Lemma 6.3 we only considered multigraphs, but here we will need to consider higher-uniformity multihypergraphs. Unfortunately, it is not in general true that large-excess 2-cuts of induced subgraphs of a multihypergraph H can be combined into a large-excess 2-cut of H . We therefore introduce the notion of average excess, which does allow us to combine multiple cuts under certain conditions, as follows. For a subset V' of the vertices of a multihypergraph H , a V' -**partial** 2-cut of H is an assignment of parts to the vertices in V' . The **average size** of a V' -partial 2-cut $\omega' : V' \rightarrow \{1, 2\}$ is the expected size of the random 2-cut of H obtained by starting with ω' and then choosing the part of each vertex $v \notin V'$ uniformly at random. The **average excess** of ω' is its average size minus the expected size of a uniformly random cut of H .

LEMMA 7.4: *Let H be a mixed multihypergraph and let V_1, \dots, V_t be disjoint subsets of its vertex set such that each edge f of H has vertices in at least $|f| - 1$ different parts. Suppose that for each $i \leq t$ there is a V_i -partial 2-cut of H with average excess x_i . Then there is a 2-cut of H with excess at least*

$$\sum_i x_i.$$

Proof. As in the proof of Lemma 6.3, we can assume V_1, \dots, V_t actually form a partition of $V(H)$, adding singleton parts if necessary. For each $1 \leq i \leq t$, let $H_i \subseteq H$ be the multihypergraph of edges which intersect V_i in two vertices and let H_0 contain the edges not in any other H_i . By assumption, no edge of H can appear in two different H_i .

For each $i \in \{1, \dots, t\}$, let $\omega_i : V_i \rightarrow \{1, 2\}$ be a V_i -partial 2-cut with average excess x_i . For each edge e in some H_i , let p_e be the probability that that edge is multicoloured in the random 2-cut obtained by starting with ω_i and choosing the part of each vertex $v \notin V_i$ uniformly at random. Let $p_e^* = 1 - 2^{1-|e|}$ be the probability that edge is multicoloured in a uniformly random 2-cut. By the definition of average excess, for each i , we have

$$\sum_{e \in E(H_i)} (p_e - p_e^*) = x_i.$$

Now let $\sigma_i : \{1, 2\} \rightarrow \{1, 2\}$ be a uniformly random permutation and define the random 2-cut $\omega : V(H) \rightarrow \{1, 2\}$ by $\omega(v) = \sigma_i(\omega_i(v))$ for $v \in V_i$. Crucially, for each $e \in E(H_i)$ with $i \neq 0$, the restriction of ω to $e \setminus V_i$ is a uniformly random

function $e \setminus V_i \rightarrow \{1, 2\}$, because each vertex of $e \setminus V_i$ is in a different V_j . Therefore, the expected size of ω is

$$\sum_{i \neq 0} \sum_{e \in E(H_i)} p_e + \sum_{e \in E(H_0)} p_e^*,$$

so its expected excess is

$$\sum_{i \neq 0} \sum_{e \in E(H_i)} (p_e - p_e^*) = \sum_i x_i.$$

Therefore, there is an outcome of ω with at least this excess. ■

In order to work with the notion of average excess, we will need an additional lemma allowing us to reduce a large-average-excess problem to a large-excess problem for a certain auxiliary edge-weighted graph. We define the size and excess of a cut of an edge-weighted multigraph in the obvious way: the size of a cut is the sum of the weights of multicoloured edges and the excess of a cut is its size minus the expected size of a uniformly random cut.

LEMMA 7.5: *Consider a multihypergraph H and a vertex subset $V' \subseteq V(H)$ such that no edge of H intersects V' in more than two vertices. Consider the weighted graph R on the vertex set V' , defined such that the weight between a pair of vertices u, v is*

$$\eta_{u,v} = \sum_e 2^{-|e \setminus V'|} = \sum_e 2^{2-|e|},$$

where the sum is over all edges $e \in E(H)$ intersecting V' in u and v . Then for any $\omega : V' \rightarrow \{1, 2\}$, the excess of ω as a 2-cut of R is exactly equal to the average excess of ω as a V' -partial 2-cut of H .

Proof. Let z be the size of ω as a 2-cut of R and let z' be the average size of ω as a V' -partial 2-cut of H . It suffices to show that $z - z'$ does not depend on ω .

First, consider the edges e intersecting V' in at most one vertex. The contribution of each such e to z is zero and the contribution to z' is $p_e^* = 1 - 2^{1-|e|}$, independently of ω .

Next, for each u, v let $N_{u,v}$ be the number of terms in the sum defining $\eta_{u,v}$. If $\omega(u) = \omega(v)$, then for any edge $e \in E(H)$ intersecting V'_i in u and v , the probability that e is not multicoloured in a uniformly random extension of ω is $2^{2-|e|}$. So, the contribution of all such edges to z' is $N_{u,v} - \eta_{u,v}$, while the contribution to z is zero. On the other hand, if $\omega(u) \neq \omega(v)$, then the contribution of $\{u, v\}$ to z' is $N_{u,v}$, while the contribution to z is $\eta_{u,v}$. ■

Next, we give the necessary generalisation of Lemma 6.4. The **kurtosis** of a random variable X is $\mathbb{E}(X - \mathbb{E}X)^4 / (\text{Var } X)^2$.

LEMMA 7.6: *Let F and G be multigraphs on a set V of n vertices. Suppose G has m edges and maximum degree Δ and let $g_{u,v}$ be the number of edges between two vertices u and v in G . Consider a random weighted graph R on the same vertex set V , where the weight between vertices u and v is a random variable $\eta_{u,v}$ with variance $\Omega(g_{u,v})$ and kurtosis $O(1)$. Suppose, moreover, that the weights $\eta_{u,v}$ are independent. Let X^{\max} be the maximum excess of a 2-cut of $F \cup R$. Then $\mathbb{E}X^{\max} = \Omega(m/\sqrt{\Delta})$.*

Before proving Lemma 7.6 we briefly explain why Lemma 6.4 follows from it. For this we will need the simple observation that sums of independent low-kurtosis random variables themselves have low kurtosis.

LEMMA 7.7: *Let X_1, \dots, X_n be independent random variables each with kurtosis $O(1)$. Then $X = \sum_i X_i$ has kurtosis $O(1)$.*

Proof. We may suppose, without loss of generality, that X_1, \dots, X_n have mean zero. Let $\sigma_i^2 = \text{Var } X_i = \mathbb{E}X_i^2$ and let $\sigma^2 = \sum_i \sigma_i^2 = \text{Var } X$. By the independence of the X_i and the kurtosis assumptions, we have

$$\mathbb{E}X^4 = \sum_i \mathbb{E}X_i^4 + 6 \sum_{i \neq j} \sigma_i^2 \sigma_j^2 = O(\sigma^4). \quad \blacksquare$$

Proof of Lemma 6.4 given Lemma 7.6. Recall the definition of the random multigraph R arising from G in the statement of Lemma 6.4. The number of edges $\eta_{u,v}$ in R between any two vertices u, v can be represented in the form $\eta_{u,v} = \eta_{u,v}^{(1)} + \dots + \eta_{u,v}^{(q)}$, where the $\eta_{u,v}^{(i)}$ are independent, and each $\eta_{u,v}^{(i)}$ satisfies $\Pr(\eta_{u,v}^{(i)} = g_{u,v}^{(i)}) = p$ and $\Pr(\eta_{u,v}^{(i)} = 0) = 1 - p$ for some $g_{u,v}^{(1)}, \dots, g_{u,v}^{(q)} \in \mathbb{N}$ with $g_{u,v}^{(1)} + \dots + g_{u,v}^{(q)} = g_{u,v}$. Note that

$$\text{Var}(\eta_{u,v}^{(i)}) = (g_{u,v}^{(i)})^2 p(1 - p) = \Omega(g_{u,v}^{(i)}),$$

so $\text{Var } \eta_{u,v} = \Omega(g_{u,v})$. Also, one may compute that the kurtosis of each $\eta_{u,v}^{(i)}$ is $(1 - 3p + 3p^2)/(p - p^2) = O(1)$, so by Lemma 7.7, $\eta_{u,v}$ has kurtosis $O(1)$ as well. We may therefore apply Lemma 7.6. \blacksquare

We will also need an additional lemma showing that for sums of independent low-kurtosis random variables we are fairly likely to see fluctuations of size comparable to the standard deviation.

LEMMA 7.8: *Let X_1, \dots, X_n be a sequence of independent random variables with kurtosis $O(1)$ and let $X = \sum_i X_i$. Then there is a positive constant c (not depending on n) such that*

$$\Pr(|X - \mathbb{E}X| \geq c\sqrt{\text{Var } X}) \geq c.$$

Proof. Recalling Lemma 7.7, it suffices to show that the desired conclusion holds whenever X itself has kurtosis $O(1)$. Also, without loss of generality, it suffices to consider the case where $\mathbb{E}X = 0$. The desired conclusion now follows from [27, Proposition 9.4], but we include the short proof for completeness. By the Paley–Zygmund inequality,

$$\begin{aligned} \Pr\left(|X| \geq \frac{1}{2}\sqrt{\text{Var } X}\right) &= \Pr\left(X^2 \geq \frac{1}{4}\mathbb{E}X^2\right) \\ &\geq (1 - 1/4)^2 \frac{(\mathbb{E}X^2)^2}{\mathbb{E}X^4} = \Omega(1). \quad \blacksquare \end{aligned}$$

Now we prove Lemma 7.6.

Proof of Lemma 7.6. Consider a random ordering of the vertices of G and let $d_{<}(v)$ be the number of edges between v and the vertices that precede it (in G). For each v , because a random ordering has the same distribution as its reverse, note that $d_{<}(v)$ has the same distribution as $d_G(v) - d_{<}(v)$, so

$$\mathbb{E}\sqrt{d_{<}(v)} = \mathbb{E}[(1/2)(\sqrt{d_{<}(v)} + \sqrt{d_G(v) - d_{<}(v)})].$$

By concavity, $\sqrt{d_{<}(v)} + \sqrt{d_G(v) - d_{<}(v)} \geq \sqrt{d_G(v)}$, so in fact

$$\mathbb{E}\sqrt{d_{<}(v)} \geq \sqrt{d_G(v)}/2.$$

Hence, we can fix an ordering with

$$\sum_{v \in V} \sqrt{d_{<}(v)} \geq \sum_{v \in V} \frac{\sqrt{d_G(v)}}{2} \geq \sum_{v \in V} \frac{d_G(v)}{2\sqrt{\Delta}} = \Omega\left(\frac{m}{\sqrt{\Delta}}\right).$$

We now build a cut with parts A and B greedily, as follows. For each vertex v in the chosen order in turn, expose the edges of R between v and the preceding vertices. Consider the total weight $e_A(v)$ (respectively, $e_B(v)$) of edges between v and the vertices previously placed in part A (respectively, part B) in $F \cup R$. If $e_A(v) \geq e_B(v)$, put v in part B . Otherwise, put v in part A . The outcome of this procedure is a cut with excess at least $\sum_v |e_A(v) - e_B(v)|/2 \leq X^{\max}$.

For each step v , let A_v be the set of vertices already placed in A and let B_v be the set of vertices already placed in B . Conditioning on these previous choices, note that

$$e_A(v) - e_B(v) = \sum_{u \in A_v} \eta_{u,v} - \sum_{u \in B_v} \eta_{u,v}$$

is a sum of independent random variables each with kurtosis $O(1)$ and the variance of this sum is

$$\sum_{u \in A_v} \Omega(g_{u,v}) + \sum_{u \in B_v} \Omega(g_{u,v}) = \Omega(d_{<}(v)).$$

By Lemma 7.8, there is $c = \Omega(1)$ such that

$$\Pr(|e_A(v) - e_B(v)| \geq c\sqrt{d_{<}(v)}) \geq c.$$

This means that

$$\mathbb{E}X^{\max} \geq \mathbb{E} \sum_{v \in V} \frac{|e_A(v) - e_B(v)|}{2} \geq \frac{c^2}{2} \sum_{v \in V} \sqrt{d_{<}(v)} = \Omega\left(\frac{m}{\sqrt{\Delta}}\right),$$

as desired. ■

We need one more ingredient to prove Lemma 5.8. The following lemma was first proved by Bonami [8] (for a short proof see [27], where it is given the name ‘‘Bonami Lemma’’). It can be viewed as a simple case of the Bonami–Beckner hypercontractive inequality for Boolean functions.

LEMMA 7.9: *Let ξ_1, \dots, ξ_n be independent Rademacher random variables, that is $\Pr(\xi_i = 1) = \Pr(\xi_i = -1) = 1/2$. Let X be a degree- b polynomial in the variables ξ_1, \dots, ξ_n . Then $\mathbb{E}X^4 \leq 9^b(\mathbb{E}X^2)^2$.*

Now, we can finally prove Lemma 5.8, recalling from Lemma 7.1 that it suffices to consider the case $U = V(H)$ (we will therefore not mention U in the proof).

Proof of Lemma 5.8. Let $V = V(H)$ and, for $W \subseteq V$, let $\overline{W} = V \setminus W$. As in the proof of Lemma 5.9, we need some way to introduce randomness into the problem. For Lemma 5.9 we exposed one of the three parts of a random cut, but here in the 2-cut case we have no parts to spare. Our approach is to choose some subset $W \subseteq V$ and expose the parts of the vertices outside W . Using the same averaging trick as in Lemma 7.1, this information will allow us to define a random hypergraph H^{part} on the vertex set W , such that a large cut of H^{part} yields a large cut of H .

Specifically, first define the “information restriction” operator P as follows. For $\omega : V \rightarrow \{1, 2\}$, let $P\omega : \overline{W} \rightarrow \{1, 2\}$ be the restriction of ω to \overline{W} . Then, for $\omega : V \rightarrow \{1, 2\}$, let $H'_{P\omega} \subseteq H$ consist of those edges e in H but not in $H[W]$ with $|P\omega(e \setminus W)| < 2$ and $|e \cap W| > 0$, and let $H_{P\omega}^{\text{part}}$ be the mixed k -multigraph with vertex set W , two copies of every edge in $H[W]$, and an edge $e \cap W$ for each $e \in E(H'_{P\omega})$. This definition is as in Lemma 7.1, but with W in place of U .

If ω is a uniformly random 2-cut, then $H_{P\omega}^{\text{part}}$ is a random hypergraph. In particular, for any edge e of H that intersects W and \overline{W} , that edge will contribute to $H_{P\omega}^{\text{part}}$ with probability $2^{1-|e \setminus W|}$. If $|e \setminus W| = 1$, then this event is not actually random and if $|e \cap W| = 1$, then the corresponding edge in $H_{P\omega}^{\text{part}}$ will have just one vertex. So, the number of edges of H which meaningfully contribute a random edge to $H_{P\omega}^{\text{part}}$ is of the same order of magnitude as the number of edges of H which intersect both W and \overline{W} in at least two vertices. For a typical choice of W , this is of the same order of magnitude as the number of edges of H which have size at least 4, which explains why we need this quantity to be large.

We will later go into more detail about our choice of W and the properties of $H_{P\omega}^{\text{part}}$, but the basic idea is that it suffices to show that the random hypergraph $H_{P\omega}^{\text{part}}$ has a large-excess 2-cut. The first step towards this goal will be to use Lemma 7.3 to obtain a partition V'_1, \dots, V'_t of W satisfying the condition in the statement of Lemma 7.4. Then, let X_i be the maximum average excess of a V'_i -partial 2-cut of $H_{P\omega}^{\text{part}}$, and let $X = \sum_i X_i$. By Lemma 7.4, $H_{P\omega}^{\text{part}}$ has a 2-cut of excess X . As in the proof of Lemma 5.9, we may therefore turn our attention towards estimating the $\mathbb{E}X_i$.

Towards this end, we will use Lemma 7.5 to interpret X_i as the maximum excess of a 2-cut in a certain weighted graph R_i on the vertex set V'_i . Since $H_{P\omega}^{\text{part}}$ is random, this weighted graph R_i will be random as well, and in fact, due to certain properties of the V'_i , the edges of R_i will be independent. We will finally apply Lemma 7.6 to R_i to estimate $\mathbb{E}X_i$. In particular, we will verify the necessary kurtosis condition with Lemma 7.9.

Now we turn to the details of the proof, which are presented in a slightly different order to the outline above. First, we apply Lemma 7.3 to obtain a suitable good partition. Let $H^{(\geq 4)} \subseteq H$ be the subgraph of edges with size at least four (with vertex set V). Apply Lemma 7.3 (with some small $c = \Omega(1)$) to obtain a cx -almost $(2m', \Delta')$ -good partition $\mathcal{V} = (V_1, \dots, V_t)$ with respect

to $(H, H^{(\geq 4)})$, where $m' = \Omega(pm)$, $\Delta' = O(p\Delta)$ and $x = m'/\sqrt{\Delta'}$. As in the proof of Lemma 5.9, delete cx edges from H (causing corresponding changes in $H^{(\geq 4)}$) so that \mathcal{V} is an (m', Δ') -good partition with respect to $(H, H^{(\geq 4)})$. We will show that the resulting multihypergraph H has a 2-cut with excess $\Omega(x)$, where the implied constant does not depend on c . This will suffice to prove Lemma 5.8.

Next we choose W . Start by letting W be a random set of vertices obtained by including each vertex independently with probability $1/2$. For each i , let H_i be the subgraph of $H^{(\geq 4)}$ consisting of all edges that intersect V_i in two vertices. By properties (i) and (iii) of Definition 7.2, the H_i are disjoint and $\sum_i e(H_i) \geq m'$. Now, let G_i be the random multigraph with vertex set $V_i \cap W$ and an edge $\{u, v\}$ for every edge of H_i containing u, v and at least two vertices outside W . Each edge e of H_i contributes to G_i with probability $\Omega(1)$, so $\mathbb{E}e(G_i) = \Omega(e(H_i))$ and we can fix an outcome of W with

$$(7) \quad \sum_i e(G_i) = \sum_i \Omega(e(H_i)) = \Omega(m').$$

This outcome of W will be fixed for the rest of the proof, so the G_i are now fixed graphs. Let $V'_i = V_i \cap W$.

Now, let $\omega : V \rightarrow \{1, 2\}$ be a uniformly random 2-cut of H , and let Z be the size of this 2-cut. Recall the definition of P and $H_{P\omega}^{\text{part}}$ from the outline at the start of the proof. As in the proof of Lemma 7.1, if, for some outcome ρ of $P\omega$, H_{ρ}^{part} has a 2-cut ϕ with some excess x' , then we may consider the average size of two different cuts related to ϕ to show that H has a 2-cut with excess at least

$$(8) \quad x'/2 + \mathbb{E}[Z \mid P\omega = \rho] - \mathbb{E}Z.$$

(The proof of (8) is exactly the same as the proof of (6), replacing “ U ” with “ W ”.)

Recall that we defined X_i to be the maximum average excess of a V'_i -partial 2-cut of $H_{P\omega}^{\text{part}}$ and set $X = \sum_i X_i$. As in the proof of Lemma 5.9, it suffices to show that $\mathbb{E}X = \Omega(x)$. Indeed, by Lemma 7.4 and property (iii) of Definition 7.2, the multihypergraph $H_{P\omega}^{\text{part}}$ has a cut with excess X , and by (8), it follows that H has a cut with excess at least $X/2 + \mathbb{E}[Z \mid P\omega] - \mathbb{E}Z$. If we could prove that $\mathbb{E}X = \Omega(x)$ we could conclude that H has a cut with excess at least

$$\mathbb{E}[X/2 + \mathbb{E}[Z \mid P\omega] - \mathbb{E}Z] = \Omega(x),$$

as desired.

We now wish to estimate each $\mathbb{E}X_i$. Let R_i be the random edge-weighted graph such that the weight between any pair of distinct vertices $u, v \in V'_i$ is the random variable

$$\eta_{u,v} = \sum_f 2^{2-|f|},$$

where the sum is over all edges $f \in E(H_{P\omega}^{\text{part}})$ intersecting V'_i in u and v . By Lemma 7.5, X_i is the maximum excess of a 2-cut of R_i .

Now, recalling the definition of $H_{P\omega}^{\text{part}}$, the direct interpretation of $\eta_{u,v}$ (in terms of the random outcome of $P\omega$) is as follows. An edge $e \in H$ with $e \not\subseteq W$ and $u, v \in e$ contributes a term $2^{2-|e \cap W|}$ to $\eta_{u,v}$ if $|P\omega(e \setminus W)| < 2$ (this event is only random if $|e \setminus W| \geq 2$, meaning that e is in $H_i \subseteq H^{(\leq 4)}$). Edges $e \in H[W]$ with $u, v \in e$ contribute two non-random terms both equal to $2^{2-|e|}$ (because we included two copies of each such edge in $H_{P\omega}^{\text{part}}$).

An immediate consequence of the above discussion is that the $\eta_{u,v}$ are independent. Indeed, by (iv) of Definition 7.2, if $\{u_1, v_1\} \neq \{u_2, v_2\}$ are different pairs of vertices in V'_i , e_1 is an edge of H containing u_1 and v_1 and e_2 is an edge of H containing u_2 and v_2 , then $e_1 \setminus V'_i \supseteq e_1 \setminus W$ and $e_2 \setminus V'_i \supseteq e_2 \setminus W$ do not intersect. However, the edges of H containing u and v may intersect each other in complicated ways, so the $\eta_{u,v}$ may have rather complicated distributions. We next need to estimate the variance and kurtosis of these distributions. In particular, to apply Lemma 7.6 we need to show that each $\eta_{u,v}$ has kurtosis $O(1)$ and we estimate the variances of the $\eta_{u,v}$ in terms of the edge multiplicities of G_i . Recall that G_i has an edge $\{u, v\}$ for every edge of H_i containing u, v and at least two vertices outside W . For each pair of distinct vertices $u, v \in V_i$, let $g_{u,v}$ be the number of edges between u and v in G_i .

CLAIM: For each pair of distinct vertices $u, v \in V_i$, $\eta_{u,v}$ has kurtosis $O(1)$ and $\text{Var } \eta_{u,v} = \Omega(g_{u,v})$.

Proof. For $u, v \in V'_i$, let $H_{u,v}$ be the set of all $g_{u,v}$ edges $e \in E(H^{(\geq 4)})$ containing u, v and at least two vertices outside W . For each $e \in E(H_{u,v})$, define

$$Q_e = 2^{2-|e \cap W|} \mathbb{1}_{|P\omega(e \setminus W)| < 2}.$$

Observe that $\eta_{u,v}$ is the sum of all such Q_e (strictly speaking $\eta_{u,v}$ is a translation of this sum: the edges which have fewer than two vertices outside W each contribute additional non-random terms).

If $e_1 \setminus W$ and $e_2 \setminus W$ are disjoint, then Q_{e_1} and Q_{e_2} are independent. Otherwise, note that for any $h \subseteq \overline{W}$ we have $\Pr(|P\omega(h)| < 2) = 2^{1-|h|}$, so

$$\text{Cov}(Q_{e_1}, Q_{e_2}) = 2^{2-|e_1 \cap W|} 2^{2-|e_2 \cap W|} (2^{1-|e_1 \cup e_2 \setminus W|} - 2^{1-|e_1 \setminus W|} 2^{1-|e_2 \setminus W|}).$$

That is, $\text{Cov}(Q_{e_1}, Q_{e_2}) = 0$ unless $e_1 \setminus W$ and $e_2 \setminus W$ share at least two vertices, in which case

$$\text{Cov}(Q_{e_1}, Q_{e_2}) = \Theta(1).$$

This implies that

$$\text{Var}(\eta_{u,v}) = \Theta(N_2),$$

where $N_2 \geq g_{u,v}$ is the number of ordered pairs of (not necessarily distinct) edges of $H_{u,v}$ intersecting in at least two vertices outside W .

Next, for each $w \in \overline{W}$, let $\xi_w = (-1)^{P\omega(w)}$, so that the ξ_w are independent Rademacher random variables as in Lemma 7.9. For each $h \subseteq \overline{W}$, let y be an arbitrary vertex in h and observe that

$$\mathbb{1}_{|P\omega(h)| < 2} = \frac{1}{2^{|h|-1}} \prod_{w \in h \setminus \{y\}} (1 + \xi_y \xi_w).$$

This means that each Q_e is a polynomial in the ξ_w with degree $O(1)$ (actually, using the fact that $\xi^2 = 1$ for $\xi \in \{-1, 1\}$, the degree can be bounded by $|e \setminus W| \leq k - 2$). Therefore, $\eta_{u,v} - \mathbb{E}\eta_{u,v}$ is a polynomial in the ξ_w with degree $O(1)$ as well and, by Lemma 7.9, it follows that $\eta_{u,v}$ has kurtosis $O(1)$, as desired. ■

By the above claim and property (ii) of Definition 7.2, we may apply Lemma 7.6 with $F = \emptyset$, $G = G_i$, $R = R_i$, $\Delta = \Delta'$ to see that

$$\mathbb{E}X_i = \Omega\left(\frac{e(G_i)}{\sqrt{\Delta'}}\right).$$

Recalling (7), it follows that

$$\mathbb{E} \sum_i X_i = \Omega\left(\frac{m'}{\sqrt{\Delta'}}\right) = \Omega(x),$$

completing the proof. ■

8. Concluding remarks

For fixed $2 \leq r \leq k$ with $k > 3$ or $r > 2$, we have shown that every k -graph has max- r -cut at least

$$\frac{S(k, r)r!}{r^k}m + \Omega(m^{5/9}),$$

while there exist k -graphs with max- r -cut only

$$\frac{S(k, r)r!}{r^k}m + O(m^{2/3}).$$

The most interesting problem left open by this paper is to close the gap between these bounds. We make the following conjecture.

CONJECTURE 8.1: *For fixed $2 \leq k \leq r$ with $k > 3$ or $r > 2$, every k -graph has an r -cut of size*

$$\frac{S(k, r)r!}{r^k}m + \Omega(m^{2/3}).$$

We remark that our specific bound $\Omega(m^{5/9})$ is simply what arises from the tradeoff between the lemmas in Section 5.1 and Lemmas 5.8 and 5.9, and it's possible that further improvements could come from both directions. In particular, for almost-linear hypergraphs, our use of Lemma 7.3 seems quite wasteful, as we are dividing our hypergraph into many small parts and only dealing with the small fraction of edges which intersect a part in two vertices. The reason we do this is to reduce our problem to a collection of random subproblems each with few dependencies, but one might hope that a more sophisticated argument could tolerate dependencies somehow.

We have also shown that every m -edge, n -vertex k -graph with no isolated vertices has an r -cut of size

$$\frac{S(k, r)r!}{r^k}m + c_{r,k}n$$

for some positive constant $c_{r,k}$. When $r = 2$, we were able to determine the best possible value for the constants $c_{2,k}$. It remains an open problem to do the same for $r > 2$.

Finally, we remark that in graphs there is a more general notion of a cut than we discussed in the introduction. An ℓ -cut is a partition of the vertex set into ℓ parts and the size of such an ℓ -cut is the number of edges which have both their vertices in different parts. Extending this more general notion of a cut to hypergraphs would result in a very general two-parameter family of cut problems. Indeed, for $r \leq \ell, k$, we can define the r -size of an ℓ -cut to be the

number of edges which have vertices in at least r of the ℓ parts and the **max**-(r, ℓ)-**cut** of a k -graph H to be the maximum r -size among all ℓ -cuts of H . We imagine that our methods could be fairly straightforwardly adapted to prove generalised counterparts of Theorems 1.2 and 1.4, but we have not explored this further.

References

- [1] R. Aharoni, *Ryser's conjecture for tripartite 3-graphs*, *Combinatorica* **21** (2001), 1–4.
- [2] N. Alon, *Bipartite subgraphs*, *Combinatorica* **16** (1996), 301–311.
- [3] N. Alon, B. Bollobás, M. Krivelevich and B. Sudakov, *Maximum cuts and judicious partitions in graphs without short cycles*, *Journal of Combinatorial Theory. Series B* **88** (2003), 329–346.
- [4] N. Alon and E. Halperin, *Bipartite subgraphs of integer weighted graphs*, *Discrete Mathematics* **181** (1998), 19–29.
- [5] N. Alon, M. Krivelevich and B. Sudakov, *Maxcut in H -free graphs*, *Combinatorics, Probability and Computing* **14** (2005), 629–647.
- [6] N. Alon, A. Pokrovskiy and B. Sudakov, *Random subgraphs of properly edge-coloured complete graphs and long rainbow cycles*, *Israel Journal of Mathematics* **222** (2017), 317–331.
- [7] B. Bollobás and A. D. Scott, *Better bounds for max cut*, in *Contemporary Combinatorics*, Bolyai Society Mathematical Studies, Vol. 10, János Bolyai Mathematical Society, Budapest, 2002, pp. 185–246.
- [8] A. Bonami, *Ensembles $\Lambda(p)$ dans le dual de D^∞* , Université de Grenoble. *Annales de l'Institut Fourier* **18** (1968), 193–204.
- [9] S. Boucheron, G. Lugosi and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*, Oxford University Press, Oxford, 2013.
- [10] R. Crowston, M. Fellows, G. Gutin, M. Jones, F. Rosamond, S. Thomassé and A. Yeo, *Simultaneously satisfying linear equations over \mathbb{F}_2 : MaxLin2 and Max- r -Lin2 parameterized above average*, in *31st International Conference on Foundations of Software Technology and Theoretical Computer Science*, Leibniz International Proceedings in Informatics, Vol. 13, Schloss Dagstuhl Leibniz-Zentrum für Informatik, Wadern, 2011, pp. 229–240.
- [11] C. S. Edwards, *Some extremal properties of bipartite subgraphs*, *Canadian Journal of Mathematics* **3** (1973), 475–485.
- [12] P. Erdős, *Problems and results in graph theory and combinatorial analysis*, in *Graph Theory and Related Topics (Proc. Conf., Univ. Waterloo, Waterloo, Ont., 1977)*, Academic Press, New York–London, 1979, pp. 153–163.
- [13] P. Erdős, A. Gyárfás and Y. Kohayakawa, *The size of the largest bipartite subgraphs*, *Discrete Mathematics* **177** (1997), 267–271.
- [14] P. Erdős and D. J. Kleitman, *On coloring graphs to maximize the proportion of multi-colored k -edges*, *Journal of Combinatorial Theory* **5** (1968), 164–169.
- [15] P. Erdős and J. L. Selfridge, *On a combinatorial game*, *Journal of Combinatorial Theory. Series A* **14** (1973), 298–301.

- [16] P. Frankl, *Extremal set systems*, in *Handbook of Combinatorics, Volume II*, Elsevier, Amsterdam, 1995, pp. 1293–1329.
- [17] A. C. Giannopoulou, S. Kolay and S. Saurabh, *New lower bound on max cut of hypergraphs with an application to r -set splitting*, in *LATIN 2012: Theoretical Informatics*, Lecture Notes in Computer Science, Vol. 7256, Springer, Heidelberg, 2012, pp. 408–419.
- [18] M. X. Goemans and D. P. Williamson, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, *Journal of the Association for Computing Machinery* **42** (1995), 1115–1145.
- [19] V. Guruswami, *Inapproximability results for set splitting and satisfiability problems with no mixed clauses*, *Algorithmica* **38** (2004), 451–469.
- [20] F. Hadlock, *Finding a maximum cut of a planar graph in polynomial time*, *SIAM Journal on Computing* **4** (1975), 221–225.
- [21] J. Håstad, *Some optimal inapproximability results*, *Journal of the Association for Computing Machinery* **48** (2001), 798–859.
- [22] F. Kardoš, D. Král’ and J. Volec, *Maximum edge-cuts in cubic graphs with large girth and in random cubic graphs*, *Random Structures & Algorithms* **41** (2012), 506–520.
- [23] R. M. Karp, *Reducibility among combinatorial problems*, in *Complexity of Computer Computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*, Plenum, New York, 1972, pp. 85–103.
- [24] S. Khot, G. Kindler, E. Mossel and R. O’Donnell, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, *SIAM Journal on Computing* **37** (2007), 319–357.
- [25] T. P. Kirkman, *On a problem in combinations*, *Cambridge and Dublin Mathematical Journal* **2** (1847), 191–204.
- [26] L. Lovász, *Coverings and colorings of hypergraphs*, in *Proceedings of the Fourth South-eastern Conference on Combinatorics, Graph Theory, and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1973)*, Utilitas Mathematica, Winnipeg, MB, 1973, pp. 3–12.
- [27] R. O’Donnell, *Analysis of Boolean Functions*, Cambridge University Press, New York, 2014.
- [28] G. Orlova and Y. Dorfman, *Finding the maximum cut in a graph*, *Engineering Cybernetics* **10** (1972), 502–506.
- [29] S. Poljak and Zs. Tuza, *Bipartite subgraphs of triangle-free graphs*, *SIAM Journal on Discrete Mathematics* **7** (1994), 307–313.
- [30] S. Poljak and Zs. Tuza, *Maximum cuts and large bipartite subgraphs*, in *Combinatorial Optimization (New Brunswick, NJ, 1992–1993)*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 20, American Mathematical Society, Providence, RI, 1995, pp. 181–244.
- [31] A. D. Scott, *Judicious partitions and related problems*, in *Surveys in Combinatorics 2005*, London Mathematical Society Lecture Note Series, Vol. 327, Cambridge University Press, Cambridge, 2005, pp. 95–117.
- [32] J. B. Shearer, *A note on bipartite subgraphs of triangle-free graphs*, *Random Structures & Algorithms* **3** (1992), 223–226.