

Exact supported co-degree bounds for Hamilton cycles

Shoham Letzter*

Arjun Ranganathan[†]

December 9, 2025

Abstract

For any $k \geq 3$ and $\ell \in [k-1]$ such that $(k, \ell) \neq (3, 1)$, we show that any sufficiently large k -graph G must contain a Hamilton ℓ -cycle provided that it has no isolated vertices and every set of $k-1$ vertices contained in an edge is contained in at least $\left(1 - \frac{1}{\lfloor \frac{k}{k-\ell} \rfloor (k-\ell)}\right) n - (k-3)$ edges. We also show that this bound is tight for infinitely many values of k and ℓ and is off by at most 1 for all others, and is hence essentially optimal. This improves an asymptotic version of this result due to Mycroft and Zárate-Guerén [31], and the case $\ell = k-1$ completely resolves a conjecture of Illingworth, Lang, Müyesser, Parczyk and Sgueglia [15].

These results support the utility of *minimum supported co-degree* conditions in a k -graph, a recently introduced variant of the standard notion of minimum co-degree applicable to k -graphs with non-trivial strong independent sets. Our proof techniques involve a novel blow-up tiling framework introduced by Lang [23], avoiding traditional approaches using the regularity and blow-up lemmas.

*Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK. Email: s.letzter@ucl.ac.uk. Research supported by the Royal Society.

[†]Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK. Email: s.letzter@ucl.ac.uk.

Contents

1	Introduction	3
2	Overview	7
2.1	Extremal lower bound constructions	7
2.2	Proof of the main result using two key theorems	9
2.3	Proof overviews	10
3	Preliminaries	11
3.1	Notation	11
3.2	Preliminary results	13
3.3	Probabilistic tools	14
4	Non-extremal hypergraphs	14
4.1	Proof sketch for Theorem 2.5	14
4.2	Main lemmas and proof of Theorem 2.5	16
5	Basic tools	19
6	Structural inheritance	21
7	Perfect weighted fractional matchings	25
8	Finding an absorbing and connecting walk	29
9	Proof of Lemma 4.3	35
10	Extremal hypergraphs	39
10.1	Proof sketch for Theorem 2.6	39
11	Key lemmas	40
11.1	Graph theoretic lemmas	41
11.2	Supported ℓ -paths	42
11.3	Cleaning lemmas	44
11.4	Perfect matchings in almost complete k -partite k -graphs	47
12	Structure of extremal hypergraphs	51
13	Balancing the sizes of A' and B'	54
14	Building a Hamilton cycle	59
15	Open problems	63
A	Proof of Lemma 4.1	67
B	Improving Theorem 1.1	69

1 Introduction

A widespread research theme in extremal graph theory is to determine sufficient conditions that ensure the existence of specified spanning structures in graphs and hypergraphs. A classic result in this vein, due to Dirac [4], states that any graph on $n \geq 3$ vertices with minimum degree at least $n/2$ contains a Hamilton cycle, and it is not too difficult to see that the degree condition is the best possible. Dirac’s theorem has been generalised to various distinct settings over the years (see the surveys [6, 10, 21, 22, 34, 40]) and this has contributed to the development of several powerful techniques, such as regularity, absorption, and rotation-extension methods. In this paper, we will focus on hypergraph extensions of Dirac’s theorem.

A k -uniform hypergraph or simply a k -graph G consists of a set of vertices $V(G)$ and a set of edges $E(G)$, where each edge consists of exactly k vertices. The most common generalisation of a minimum degree condition from graphs to hypergraphs is accomplished via the notion of a *minimum co-degree* – which is the minimum d such that every set of $k - 1$ vertices in G is contained in at least d edges – and we denote it here by $\delta(G)$. The larger uniformity k allows for multiple distinct ways to define a cycle in a k -graph. The cyclic structures we search for are ℓ -cycles, which are perhaps the most common extensions of graph cycles. Intuitively, we think of an ℓ -cycle as a spanning path formed by a cyclic set of edges such that each edge has exactly ℓ vertices in common with the preceding edge (it could intersect other edges as well). More formally, given an n -vertex k -graph G and any $\ell \in [k - 1]$ such that $k - \ell$ divides n , we define a *Hamilton ℓ -cycle* in G to be an ordering of the vertices of G , say $v_1 \dots v_n$, such that the vertices of the subsequence $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$ form an edge for all $i \geq 0$, where we view the indices cyclically modulo n . Since each edge of an ℓ -cycle contains $k - \ell$ vertices that were not in the previous edge, we trivially require that $k - \ell$ divides n so that when we “cycle around” the vertices of G , the original sequence of edges gets repeated. We call $(k - 1)$ -cycles and 1-cycles *tight* and *loose* cycles respectively.

The question of determining sufficient minimum co-degree conditions to ensure a tight Hamilton cycle was first raised by Katona and Kierstead [16] who conjectured that if $\delta(G) \geq (n - k + 2)/2$ then G has a tight Hamilton cycle, and they showed that this bound is the best possible. Rödl, Ruciński and Szemerédi [35, 37] first proved this asymptotically for $k \geq 3$ and $\delta(G) \geq n/2 + o(n)$, and later obtained an exact result for $k = 3$ [39]. Additionally, Liu and Liu [25] made progress towards an exact result for $k = 4$. In fact, for any $\ell \in [k - 1]$ such that $k - \ell$ divides k , a tight Hamilton cycle will contain such a Hamilton ℓ -cycle (provided the necessary divisibility condition holds). If we further suppose that k divides n , then this Hamilton ℓ -cycle will imply the existence of a perfect matching which necessitates $\delta(G) \geq n/2 - k$ as shown in [19, 38], and hence the aforementioned tight cycle bound of $\delta(G) \geq n/2 + o(n)$ is tight. If k does not divide n , then Rödl, Ruciński and Szemerédi [38] have determined the minimum codegree threshold for a near-perfect matching of size $\lfloor n/k \rfloor$ to be roughly n/k . It is therefore plausible that the correct minimum codegree threshold for a spanning ℓ -cycle, when $k - \ell$ divides k but is not 1 and n is not divisible by k , is significantly

lower than $n/2$. Conversely, if $k - \ell$ does not divide k , a series of works by Kühn and Osthus [20], Keevash, Kühn, Mycroft and Osthus [17], Hàn and Schacht [13] and Kühn, Mycroft and Osthus [19] have established that

$$\delta(G) \geq \frac{n}{\left\lceil \frac{k}{k-\ell} \right\rceil (k-\ell)} + o(n)$$

suffices and is optimal up to the $o(n)$ term, yielding a much lower threshold. Exact bounds, however, have proven notoriously difficult to obtain and are known only in a few special cases. To the best of our knowledge, these are $(k, \ell) = (3, 2)$ by Rödl, Ruciński and Szemerédi [39], $(k, \ell) = (3, 1)$ by Czygrinow and Molla [3], $(k, \ell) = (4, 2)$ by Garbe and Mycroft [7], and $k \geq 3$ and $\ell < k/2$ by Han and Zhao [14]. We refer the reader to the surveys [22, 34, 43] for more detailed discussions.

A drawback of a minimum co-degree condition is that it tends to be a fairly strong requirement. For instance, if we start with a complete k -graph and remove all edges that contain a fixed pair of vertices, it is immediate that the resulting hypergraph has minimum co-degree equal to zero. However, if $k \geq 3$, it is not hard to see that even fairly small hypergraphs of this form are quite dense, and trivially contain Hamilton cycles. In fact, there are several natural hypergraph classes that are quite dense, but have minimum co-degree equal to zero (multipartite hypergraphs, for instance), and hence the previous theorems are not immediately applicable. This motivates the notion of the *minimum supported co-degree* of a k -graph G , which we denote $\delta^*(G)$, and define as the maximum integer d such that every $(k - 1)$ -set that is contained in at least one edge is contained in at least d edges. A minimum supported co-degree requirement immediately handles the previous issues of zero co-degree sets, although it does not rule out isolated vertices, and hence we will usually impose the (fairly weak) requirement of all vertices being non-isolated.

Multiple recent studies have investigated questions that arise from replacing a minimum co-degree with a minimum supported co-degree requirement in classical extremal hypergraph theoretic problems. The notion of supported co-degrees was introduced by Balogh, Lemons and Palmer [2] (they use the term “positive co-degree” instead) to formulate a generalisation of the Erdős-Ko-Rado theorem, providing bounds on the sizes of intersecting families satisfying a minimum supported co-degree condition. Subsequently, Frankl and Wang [5] improved these bounds in almost all cases, and Spiro [41] has extended these studies to the more general case of t -intersecting families. Other papers have worked on generalisations of the Andrásfai-Erdős-Sós theorem [27] and unique colourability of hypergraphs [26].

Halfpap, Lemons and Palmer [11] investigated variants of hypergraph Turán problems, and established the asymptotic supported co-degree threshold for a host 3-graph to contain a copy of a fixed graph F for several distinct 3-graphs F , and Wu [42] has consequently addressed this question for other examples. Halfpap, Lemons and Palmer [11] also study “jumps” in this threshold, which has further been tackled by Balogh, Halfpap, Lidický, and Palmer [1]. Pikhurko [33] recently proved that these problems are well-defined for a generalisation to the “minimum positive/supported ℓ -co-degree” for any $\ell \in [k - 1]$, providing analogous results to those Lo and Markström [28] regarding

the usual notion of ℓ -co-degree.

The question of finding minimum supported co-degree conditions that ensure the existence of spanning structures in hypergraphs was first raised by Halfpap and Magnan [12]. They prove an exact optimal lower bound on the minimum supported co-degree required to guarantee a perfect matching in 3-graphs and a slightly weaker bound for all higher uniformities, which was later improved to an exact tight bound by Mycroft and Zárate-Guerén [32]. Furthermore, Halfpap and Magnan also establish an exact best-possible minimum supported co-degree condition for Hamilton Berge cycles and an asymptotic one for 3-uniform loose cycles. In a different direction, Illingworth, Lang, Müyesser, Parczyk and Sgueglia [15] showed that an n -vertex k -graph with $\delta^*(G) \geq n/2 + o(n)$ contains a spanning k -sphere, asymptotically confirming a conjecture of Georgakopoulos, Haslegrave, Montgomery and Narayanan [8].

The problem of determining the optimal minimum supported co-degree for k -uniform ℓ -cycles was initially suggested by Halfpap and Magnan in [12, Section 5], and reiterated by Illingworth, Lang, Müyesser, Parczyk and Sgueglia [15] for the special case of tight cycles. This question was recently tackled by Mycroft and Zárate-Guerén [31], who prove an asymptotically optimal result for all $k \geq 3$ and $\ell \in [k - 1]$.

Our main result provides an improved sufficient minimum supported co-degree condition for k -uniform ℓ -cycles for all $k \geq 3$ and $\ell \in [k - 1]$ except $(k, \ell) = (3, 1)$, which is the best possible for infinitely many values of k and ℓ , and is off by at most 1 for all values of k and ℓ . We discuss the optimality of our result in Section 2.1. Importantly, we emphasise that our results are (essentially) exact and not asymptotic, a rather stark difference from the minimum co-degree version of the problem where very few cases have been resolved exactly.

Theorem 1.1. *Let $1 \leq \ell \leq k - 1$ be such that $k \geq 3$ and $(k, \ell) \neq (3, 1)$, let $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$, and let n be sufficiently large and divisible by $k - \ell$. Then every k -uniform n -vertex hypergraph without isolated vertices and having minimum supported co-degree at least $(1 - 1/t)n - (k - 3)$ will contain a Hamilton ℓ -cycle.*

This improves the aforementioned asymptotic results of Mycroft and Zárate-Guerén [31], who showed that $\delta^*(G) \geq (1 - 1/t)n + o(n)$ suffices. Our methods are largely different from theirs. Our analysis splits into two regimes: the extremal and non-extremal (see Section 2.2 for details). The extremal regime does not feature in [31] at all, and for this we use ad-hoc structural analysis with lemmas about random matchings in bipartite graphs and about finding almost spanning subhypergraphs with large minimum supported codegree in almost complete hypergraphs. For the non-extremal case, while Mycroft and Zárate-Guerén rely on a version of the regularity lemma for hypergraphs (called the weak regularity lemma) along with the absorption method, our approach stems from Lang's breakthrough work [23] on hypergraph tilings, which allows us to completely avoid the regularity lemma and greatly simplifies the use of the absorption method. For the latter

part, however, we use a novel idea from [31] about weighted fractional matchings that, coupled with Farkas' linear-algebraic lemma, allows one to find spanning ℓ -cycles in certain k -partite k -graphs, which is one of the steps in our proof.

We point out that Illingworth, Lang, Müyesser, Parczyk and Sgueglia [15] conjectured that for all $k \geq 3$, any sufficiently large k -graph with $\delta^*(G) \geq (1 - 1/k)n$ contains a tight Hamilton cycle, which is a special case of Theorem 1.1 (up to the $k - 3$ term). Observe that if k divides n , then a tight Hamilton cycle in G will contain a perfect matching. Thus, our result implies that $\delta^*(G) \geq (1 - 1/k)n - (k - 3)$ ensures a perfect matching, which recovers a result of Mycroft and Zárate-Guerén [32] up to an additive constant of one, who show that a bound of $(1 - 1/k)n - (k - 2)$ suffices.

As alluded to above, the bound in Theorem 1.1 is off by one in most cases. In fact, in many cases (namely whenever $t \geq \ell + 2$ or when $\ell = k - 1$, that is, we are seeking a tight Hamilton cycle), our techniques can be adapted to prove Theorem 1.1 with the improved bound of $\delta^*(G) \geq (1 - 1/t)n - (k - 2)$ instead (which, for instance, would immediately imply the previously discussed optimal perfect matching bound). However, since this will require some technical modifications, we choose to present a single unified proof for $\delta^*(G) \geq (1 - 1/t)n - (k - 3)$ that will work for all cases instead. We point out the part of our proof that requires the exact degree condition in Remark 13.3, and briefly discuss how to suitably alter our proof and remove the extra one in the bound in all relevant cases in Appendix B.

Finally, we remark that in a recent personal communication, Mycroft and Zárate-Guerén informed us that they are preparing a manuscript where they prove that any sufficiently large 3-graph with no isolated vertices and $\delta^*(G) \geq n/2$ contains a loose Hamilton cycle. This improves their previous general asymptotic bound of $n/2 + o(n)$ in this particular case [31] (which our methods can recover as well) and extends Theorem 1.1 to $(k, \ell) = (3, 1)$, the only case we do not handle.

Organisation of the paper. In Section 2, we first provide lower bound constructions to prove the optimality of Theorem 1.1. We then split the proof of Theorem 1.1 into two complementary cases and handle them separately in Theorems 2.5 and 2.6. We also provide brief proof overviews for both these theorems in Section 2.3. Then in Section 3 we introduce some notation and definitions and provide some basic tools.

Sections 4 to 9 are dedicated to proving Theorem 2.5, and Sections 10 to 14 prove Theorem 2.6. These two parts are treated largely independently, and can be read as such. We conclude with some open problems in Section 15.

2 Overview

In this section we first give three extremal constructions (see Section 2.1). We then state, in Section 2.2, two theorems that will split the task of finding a Hamilton ℓ -cycle in a k -graph with appropriate minimum supported codegree into two cases: non-extremal and extremal, and observe that these theorems together imply our main result. Finally, we give brief proof sketches for each theorem in Section 2.3.

2.1 Extremal lower bound constructions

We first discuss the optimality of the lower bound in Theorem 1.1. To begin with, we say that a subset U of at most k vertices in a k -graph G is *supported* if there is an edge containing U . Furthermore, if $|U| \leq k-1$, set $d_G^1(U)$ to be the number of vertices $v \in V(G)$ such that $U \cup \{v\}$ is supported. We say that a subset A of vertices is a *strong independent set* if every edge intersects A in at most one vertex.

We will provide three lower bound constructions. The first one will show that, at the very least, we need $\delta^*(G) \geq (1-1/t)n - (k-2)$ in Theorem 1.1 for any $k \geq 3$, $\ell \in [k-1]$ and n which is divisible by $k-\ell$. The second one, which is a slight modification of the first, shows that if k is odd, $\ell = (k-1)/2$, and n satisfies some divisibility conditions, then we require $\delta^*(G) \geq (1-1/t)n - (k-3)$, which is why our main result is essentially tight and is the strongest possible bound that applies to all relevant k , ℓ and n . The final one will be an entirely different construction that applies only to the $(k, \ell) = (3, 1)$ case and shows that we need $\delta^*(G) \geq n/2$, which is in line with the second example.

Example 2.1. Let $k \geq 3$, $\ell \in [k-1]$ and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Let n be chosen such that $k-\ell$ divides n . Now, consider an n -vertex k -graph G with a partition $A \sqcup B$ of its vertex set such that $|A| = \lfloor n/t \rfloor + 1$ and $E(G)$ consists of all k -sets $e \in \binom{V(G)}{k}$ that satisfy $|e \cap A| \leq 1$ (and so A is a strong independent set). Suppose S is any supported $(k-1)$ -set. If $|S \cap A| = 0$, then S supports all vertices outside itself, and so $d_G^1(S) = n - k + 1$. If $|S \cap A| = 1$, then it supports all vertices in $B \setminus S$, meaning that $d_G^1(S) = n - \lfloor n/t \rfloor - (k-2)$, and this is equal to $\delta^*(G)$ since there is no supported set with $|S \cap A| \geq 2$. The following observation shows that G cannot contain a Hamilton ℓ -cycle.

Observation 2.2. *Let H be an n -vertex k -uniform ℓ -cycle. Then the maximum size of a strong independent set in H is at most $\lfloor n/t \rfloor$.*

Proof. As noted previously, since H is an ℓ -cycle, $k-\ell$ divides n . Recall that we can label the vertices in H as $v_1 \dots v_n$ such that $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$ is an edge for all $i \geq 0$, with indices seen cyclically modulo n . Define the set of *segments* to be the collection of subsequences $\mathcal{S} = \{v_{i(k-\ell)+1} \dots v_{i(k-\ell)+t} : i \geq 0\}$ and the set of *intervals* $\mathcal{I} = \{v_{i(k-\ell)+1} \dots v_{(i+1)(k-\ell)} : i \geq 0\}$, where we view indices cyclically modulo n as usual. Let X be a strong independent set.

First, we observe that every segment contains exactly $t/(k-\ell)$ intervals, and that every interval is part of $t/(k-\ell)$ segments. Since the intervals partition $V(H)$, we see that every vertex is contained in $t/(k-\ell)$ segments. The crucial observation is that, due to the definition of ℓ -cycles and since trivially $t \leq k$, every segment is contained in an edge of H , and consequently every segment contains at most one vertex of X . Now, consider the set of vertex-segment pairs $\{(v, P) : v \in X, P \in \mathcal{S}, x \in P\}$. Since each vertex belongs to $t/(k-\ell)$ segments, we see that the size of this set is exactly $t|X|/(k-\ell)$. However, as observed previously, each segment $P \in \mathcal{S}$ can contain at most one vertex $v \in X$, and hence the number of such pairs is at most $|\mathcal{S}| = n/(k-\ell)$. Comparing the two shows $|X| \leq n/t$ and hence $|X| \leq \lfloor n/t \rfloor$. \square

We remark that Mycroft and Zárate-Guerén [31, Section 1.4] prove the same bound with this extremal construction, but we include our proof as it is simpler. We also wish to point out that Illingworth, Lang, Müyesser, Parczyk and Sgueglia [15] provide a version of this construction for tight cycles.

Example 2.3. Next, we provide a specialised modification of the previous example. Let $k \geq 3$ be odd, $\ell = (k-1)/2$, and consequently $t = \lfloor \frac{k}{k-\ell} \rfloor (k-\ell) = (k+1)/2 = \ell + 1 = k - \ell$. Let n be such that t divides n and $n/t + 1$ is even. Yet again, we consider an n -vertex k -graph G with a partition $A \sqcup B$ of its vertex set with $|A| = n/t + 1$. Write $A = \{a_1, \dots, a_{|A|}\}$ and let $E(G)$ consist of all $e \in \binom{V(G)}{k}$ such that $e \cap A \subseteq \{a_{2i-1}, a_{2i}\}$ for some $i \in [|A|/2]$ (so the graph $\partial^2[A]$ is a perfect matching). Clearly $|e \cap A| \leq 2$ for all $e \in E(G)$. If S is any supported $(k-1)$ -set, then, similar to the previous example, it is not too hard to see that

$$d_G^1(S) = \begin{cases} n - (k-1) & \text{if } |e \cap A| = 0, \\ |B| - (k-2) + 1 = |B| - (k-3) & \text{if } |e \cap A| = 1, \\ |B| - (k-3) & \text{if } |e \cap A| = 2, \end{cases}$$

and so $\delta^*(G) = |B| - (k-3) = n - n/t - (k-2)$. We now argue that G cannot contain a Hamilton ℓ -cycle, which shows that the bound in Theorem 1.1 cannot be improved.

To the contrary, suppose that there exists such a cycle $C = v_1 \dots v_n$, and let $E(C)$ denote the edges of G that are part of the cycle C . Every edge $e \in E(C)$ shares $\ell = (k-1)/2$ vertices with the preceding and succeeding edges of the cycle, and hence contains a unique vertex v_e at position $\ell + 1 = t$ (the “middle” vertex) that is not part of any other edge of $E(C)$. Let C^t denote the set of these vertices, so that $|C^t| = n/t$. Another important observation we require is that if $u \in A$ and $uv_1 \dots v_{k-1} \in E(G)$, then $vv_1 \dots v_{k-1} \in E(G)$ for any $v \in B \setminus \{v_1, \dots, v_{k-1}\}$.

Consider any edge $e \in E(C)$ such that $e \cap A \neq \emptyset$. If $v_e \in B$, then there exists some $u \in e \cap A$ that is distinct from v_e . From our observation above, and since e is the only edge of C that contains v_e , the vertex sequence obtained by swapping u and v_e in the sequence corresponding to C still yields a Hamilton ℓ -cycle. Hence, we may assume that $v_e \in A$ for all $e \in E(C)$ that intersect A .

Next, suppose there exists some $u \in A \setminus C^t$. The discussion in the previous paragraph implies that there are two edges $e, f \in E(C)$ that contain u . By choice of C , as A intersects e and f , we know that $v_e, v_f \in A$. Hence, we may conclude that uv_e and uv_f are both supported pairs due to the edges e and f respectively, meaning that u is a vertex of degree at least two in $\partial^2[A]$, a contradiction. Thus, we see that $A \subseteq C^t$, which implies $|A| \leq |C^t| = n/t$, providing the desired contradiction.

Example 2.4. Suppose $k = 3$ and $\ell = 1$, so that $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell) = 2$, and suppose $n \equiv 2 \pmod{4}$. Let G be an n -vertex 3-graph that consists of the union of two complete 3-graphs H_1 and H_2 , each on $n/2 + 1$ vertices, such that H_1 and H_2 have precisely two vertices in common, say a and b . Since any supported 2-set must be contained entirely in H_i for some i , we see that $\delta^*(G) = |H_i| - 2 = n/2 - 1$.

Suppose G has a 3-uniform loose cycle $C = v_1 \dots v_n$. Then, since any two consecutive vertices are supported, we see that any $v \in V(H_i) \setminus \{a, b\}$ must be preceded and succeeded only by vertices of $V(H_i)$ in the cycle C . From this, it is easy to see that C must contain a sequence of consecutive vertices corresponding to a loose Hamilton path in H_i for each i . However, any 3-uniform loose path must have odd order (since the first edge has three vertices and each consecutive edge adds two new vertices), and $|H_i| = n/2 + 1$ is even, which is a contradiction.

We point out that the last example was brought to our attention by Richard Mycroft and Camila Zárate-Guerén, and we are thankful for that.

2.2 Proof of the main result using two key theorems

We prove our main result, Theorem 1.1, by splitting it into two disjoint and complementary cases. Roughly speaking, we handle hypergraphs that are structurally “near-extremal” and “non-extremal” separately. We classify hypergraphs into one of these two types based on the observation that both extremal examples contain a large “sparse” set of size roughly n/t with very few supported pairs.

The first theorem treats k -graphs that are far from extremal, that is, where every set of size at least n/t induces several supported pairs. Note that here we allow for a minimum supported co-degree which is slightly lower than the exact bound given in Theorem 1.1.

Theorem 2.5. *Let $k \geq 3$ and $\ell \in [k - 1]$ be such that $(k, \ell) \neq (3, 1)$, and set $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$. Let $1/n \ll \varepsilon \ll \mu \ll 1/k$ such that n is divisible by $k - \ell$. If G is an n -vertex k -graph with no isolated vertices and $\delta^*(G) \geq (1 - 1/t - \varepsilon)n$ such that every set of $\lfloor n/t \rfloor$ vertices of G contains at least μn^2 supported pairs, then G has a Hamilton ℓ -cycle.*

The second theorem deals with k -graphs that are close to the two extremal examples above, namely where there is a set of size $\lfloor n/t \rfloor$ which contains few supported pairs. This result requires the exact bound on the minimum supported co-degree.

Theorem 2.6. *Let $k \geq 3$ and $\ell \in [k-1]$ be such that $(k, \ell) \neq (3, 1)$, and set $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Let $1/n \ll \varepsilon \ll 1/k \leq 1/3$ be such that n is divisible by $k-\ell$. Suppose G is an n -vertex k -graph with no isolated vertices and $\delta^*(G) \geq (1 - 1/t)n - (k-3)$. If there is a subset $A \subseteq V(G)$ with $|A| = \lfloor n/t \rfloor$ that contains at most εn^2 supported pairs, then G has a Hamilton ℓ -cycle.*

Notice that the proof of Theorem 1.1 follows immediately from Theorems 2.5 and 2.6. For the most part, the proofs of the two theorems use separate arguments and tools. In the rest of the section, we provide short overviews for the proofs of Theorems 2.5 and 2.6.

2.3 Proof overviews

2.3.1 Proof overview for Theorem 2.5

In this section, we provide a brief overview of the key ideas we use to prove Theorem 2.5, and include a more detailed proof sketch in Section 4.1.

The approach we use is built on a blow-up tiling technique introduced by Lang [23], and extended in [15, 24] (we make precise what we mean by a blow-up in Section 4.2). Our “blow-up tiling lemma”, namely Lemma 4.1, allows us to tile G with almost balanced blow-ups of smaller k -graphs that obey approximate versions of the degree condition and the non-extremal structure of G . These tiles will have a special cyclic linkage property in terms of common edges, which will then reduce our problem to finding a Hamilton ℓ -path within each blown up tile, because these can be “chained together” to form a Hamilton ℓ -cycle in G .

Hence, given any “non-extremal” F and a sufficiently large nearly balanced blow-up F^* , we want a Hamilton ℓ -path in F^* . We will first partition F^* into complete k -partite k -graphs with carefully selected part sizes to ensure each has an almost spanning ℓ -path. These k -partite k -graphs will be formed by splitting up the complete k -partite k -graphs corresponding to the blow-ups of edges of F . We determine how to split up these k -graphs based on edge weights provided by a perfect fractional matching (defined in Section 7) in F . Technically, we work with a “vertex weighted” version of fractional matchings due to Mycroft and Zárate-Guerén [31], which is crucial for ℓ -cycles that are not necessarily tight.

Finally, we use the minimum supported co-degree condition to find a path in F^* (obtained by blowing up a suitable walk in F) that can be used to link together the nearly spanning ℓ -paths found in the aforementioned k -partite subgraphs and absorb a few uncovered vertices.

2.3.2 Proof overview for Theorem 2.6

We provide a short overview of our proof strategy for Theorem 2.6 for the special case of tight cycles (which is what motivates the general strategy) and include a more detailed sketch including

ℓ -cycles in Section 10.1. Suppose G is an “extremal” n -vertex k -graph, meaning it contains a set A of size roughly n/k (since $t = k$ for tight cycles) within which there are few supported pairs. Let $B = V(G) \setminus A$. Example 2.1 suggests that a tight Hamilton cycle has a strong independent set of size n/k , whose vertices are exactly the k th vertices along the cycle. We build a Hamilton cycle in G based on this observation.

We first show that we can transfer a few vertices across A and B to ensure high co-degrees for supported sets across the partition while maintaining $|A| \approx n/k$. We then replace a few vertices A with short tight paths, and this will recover $|B| = (k-1)|A|$.

We then partition B into $k-1$ sets B_1, \dots, B_{k-1} such that $|A| = |B_i|$ for all i . We define an auxiliary k -partite k -graph G^+ with parts $B_1 \times \dots \times B_{k-1} \times A$ whose edges correspond to vertex sequences that form tight paths, and randomly construct a perfect matching M . The randomness will maintain the high co-degree conditions so that the edges of M can be connected into a tight Hamilton cycle in G .

3 Preliminaries

In this section we provide notation and a few simple preliminary results and standard probabilistic tools that will be used in the proofs of both the non-extremal and extremal theorems. More specialised notation and results, pertaining to only one of the theorems, will be mentioned at the beginning of the relevant part of the paper.

3.1 Notation

For any $m_1, m_2 \in \mathbb{N}$ with $m_1 \leq m_2$, we let $[m_1]$ denote the set $\{1, 2, \dots, m_1\}$ and often use $[m_1, m_2]$ to denote $\{m_1, m_1 + 1, \dots, m_2\}$. Given any set S and any $0 \leq m \leq |S|$, we use $\binom{S}{m}$ to denote the set of all m -subsets of S . Throughout the paper, we frequently use “ \ll ” notation. We write $a \ll b$ to mean that for any choice of $b > 0$, there exists some $a_0 > 0$ such that the statement in question holds for all $a \leq a_0$. We may sometimes write $b \gg a$ for the same.

For a k -graph $G = (V(G), E(G))$, we use $|G|$ to denote the number of vertices and $e(G)$ to denote the number of edges. We often simply write $v_1 \dots v_k$ for an edge $\{v_1, \dots, v_k\} \in E(G)$. We let $E^*(G)$ denote the set of ordered edges of G , that is, the set of all ordered tuples (v_1, \dots, v_k) such that $\{v_1, \dots, v_k\} \in E(G)$ (and hence each edge of $E(G)$ leads to $k!$ ordered edges in $E^*(G)$). We sometimes denote this ordered edge as $v_1 \dots v_k$ as well when it is clear from context (or explicitly specified) that we are dealing with ordered edges.

For any $U \subseteq V(G)$, we let $G[U]$ denote the subgraph of G induced by U . We say that a vertex subset $S \subseteq V(G)$ with $|S| \leq k$ is *supported in G* if there exists some edge of G containing S . Analogous

to ordered edges, we define an ordered supported set to be an ordered tuple $v_1 \dots v_s = (v_1, \dots, v_s)$ such that the set $\{v_1, \dots, v_s\}$ is supported in G . For any $i \in [k-1]$, we let define the i -shadow of G , denoted $\partial^i G$, to be the i -graph with $V(\partial^i G) = V(G)$ whose edges are precisely the supported i -sets of G . We define $\partial_G^i[U]$ to be the subgraph induced by a vertex subset U , and will often drop the subscript G and write $\partial^i[U]$, as long as it does not lead to any ambiguity.

Given any $n \gg k \geq 3$ and $\ell \in [k-1]$, we set $t = t(k, \ell) = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$, and we refer to this parameter t throughout the paper. Define $\mathcal{E}_\varepsilon^{\text{ext}}(n, k)$ to be the set of all n -vertex k -graphs G with n divisible by $k-\ell$ that have no isolated vertices, satisfy $\delta^*(G) \geq n - \lfloor n/t \rfloor - (k-3)$, and for which there is a set of $\lfloor n/t \rfloor$ vertices containing at most εn^2 supported pairs. We define the family $\mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$ as the collection of all n -vertex k -graphs with n divisible by $k-\ell$ that have no isolated vertices, satisfy $\delta^*(G) \geq (1 - 1/t - \varepsilon)n$ and for which every subset of at least n/t vertices contains at least μn^2 supported pairs.

For a k -graph G and $\ell \in [k-1]$, a k -uniform ℓ -walk W in G is a sequence of (not necessarily distinct) vertices $v_1 \dots v_r$ such that $r \equiv k \pmod{k-\ell}$ and $v_{s(k-\ell)+1} \dots v_{s(k-\ell)+k}$ is an edge for every integer $s \in \left[0, \frac{r-k}{k-\ell}\right]$. In particular, consecutive edges $v_{s(k-\ell)+1} \dots v_{s(k-\ell)+k}$ and $v_{(s+1)(k-\ell)+1} \dots v_{(s+1)(k-\ell)+k}$ share the ℓ vertices $v_{(s+1)(k-\ell)+1}, \dots, v_{(s+1)(k-\ell)+\ell}$ (they possibly share more vertices as vertices are allowed to repeat). We say that the ℓ -walk $W = v_1 \dots v_r$ has order r , and will denote its order as $|W|$. We call such a W an ℓ -path if all the v_i are distinct. We call W a *tight walk* if $\ell = k-1$ and call it a *loose walk* if $\ell = 1$. We define tight and loose paths analogously.

A k -uniform ℓ -cycle C is a cyclic sequence of distinct vertices $v_1 \dots v_r$, where r is divisible by $k-\ell$ and $v_{s(k-\ell)+1} \dots v_{s(k-\ell)+k}$ is an edge for every $s \in \left[0, \frac{r-k}{k-\ell}\right]$ (addition of indices is taken modulo r). A k -uniform $(k-1)$ -cycle is called a *tight cycle*, and a 1-cycle is called a *loose cycle*. Finally, we say that an ℓ -path or an ℓ -cycle is *Hamiltonian* if every vertex of G appears in the corresponding sequence exactly once, and will call these structures Hamilton ℓ -paths and Hamilton ℓ -cycles respectively.

Additionally, we also introduce some non-standard notation that is motivated by our proof strategy. For any supported i -set $S \subseteq V(G)$, we define the *vertex neighbourhood* of S , denoted $N_G^1(S)$, to be the set of all vertices $v \in V(G)$ such that $S \cup \{v\}$ is a supported $(i+1)$ -set, and set $d_G^1(S) = |N_G^1(S)|$ to be the *vertex co-degree* of S . For $S = \{v_1, \dots, v_i\}$, we will simply write $N_G^1(v_1 \dots v_i)$ and $d_G^1(v_1 \dots v_i)$ for $N_G^1(S)$ and $d_G^1(S)$ respectively. As always, we may sometimes drop the subscript G if there is no ambiguity in the host graph. Furthermore, for any $U \subseteq V(G)$, we will denote $N_G^1(S) \cap U$ as $N_U^1(S)$ and $d_U^1(S) = |N_U^1(S)|$. Finally, for an arbitrary i -set S (which may not be supported), we define the *degree of S* , denoted $d(S)$, to be the number of edges of G which contain S . The *minimum i -degree* of G , denoted $\delta_i(G)$, is the maximum integer d such that every set of i vertices is contained in at least d edges.

3.2 Preliminary results

First, we mention a few quick facts about $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Since these are fairly simple, we may often use these without quoting the observation below.

Observation 3.1. *Let $k \geq 3$, $\ell \in [k-1]$, and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Then $\ell+1 \leq t \leq k$, $t \geq \frac{k+1}{2}$, and $t \geq 3$ unless $(k, \ell) = (3, 1)$.*

Proof. We prove the claims sequentially. For the first one, we see that

$$t > \left(\frac{k}{k-\ell} - 1 \right) (k-\ell) = \ell,$$

and the fact that $t \leq k$ is immediate from the definition.

We prove the second claim in two cases. If $\ell \geq k/2$, then we are done by the first part. Otherwise, we will have $k-\ell > k/2$ which implies $t \geq k-\ell \geq (k+1)/2$.

The final claim is easy since $t \in \{1, 2\}$ requires at least one of $\left\lfloor \frac{k}{k-\ell} \right\rfloor$ and $(k-\ell)$ to be equal to one and the other to be equal to either one or two depending on the desired value of t . Then, as $k \geq 3$, it is simple to check that $t = 2$ for $(k, \ell) = (3, 1)$ and $t \geq 3$ otherwise. \square

Next, we show a quick result about certain vertices in ℓ -paths that will be required for both the extremal and non-extremal cases. Intuitively speaking, these vertices form the “sparse” set of size n/t in the extremal constructions. In a vertex sequence corresponding to an ℓ -path or ℓ -cycle, an edge always starts at a vertex immediately after any of these special vertices. This motivates the proof strategy described in Section 2.3.2 for near-extremal hypergraphs.

Proposition 3.2. *Let $k \geq 3$, $\ell \in [k-1]$ and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Suppose G is a k -graph and $W = v_1 \dots v_r$ is a vertex sequence, where r is divisible by $k-\ell$. Let $E_W = \{v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k} : i \geq 0\}$ be a collection of subsequences of W and $A = \{v_j : j \equiv 0 \pmod{t}\}$, where we view all indices cyclically modulo r . Then every $e \in E_W$ contains exactly one vertex from A .*

Proof. The edges of W are of the form $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$. Since $t \leq k$, each edge contains at least one vertex from A . We will show that no edge has two vertices from A .

Fix an edge $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$. Observe that if $v_j \in A$, then j is a multiple of $k-\ell$. Consequently, if j is the smallest index such that $v_j \in A \cap \{v_{i(k-\ell)+1}, \dots, v_{i(k-\ell)+k}\}$, then $j \geq (i+1)(k-\ell)$. Now, we see that

$$\begin{aligned} j+t &\geq (i+1)(k-\ell) + \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell) \\ &> (i+1)(k-\ell) + \left(\frac{k}{k-\ell} - 1 \right) (k-\ell) \\ &= i(k-\ell) + k, \end{aligned}$$

which means there is no other vertex of A in the edge $\{v_{i(k-\ell)+1}, \dots, v_{i(k-\ell)+k}\}$. \square

3.3 Probabilistic tools

Next, we state (without proof) standard Chernoff bounds for binomial and hypergeometric random variables.

Lemma 3.3. *Let X_1, \dots, X_n be random variables that take values in $\{0, 1\}$ such that $\mathbb{P}[X_i = 1] = p$ for all i . Suppose $X = X_1 + \dots + X_n$ has either a binomial or a hypergeometric distribution with sample size n and success probability p . Then for any $a > 0$,*

$$\mathbb{P}[|X - \mathbb{E}X| > a] = \mathbb{P}[|X - np| > a] \leq 2 \exp\left(-\frac{a^2}{3np}\right)$$

Another tool we will require is McDiarmid's inequality.

Lemma 3.4 (McDiarmid's inequality [30]). *Let X_1, \dots, X_s be independent random variables taking values in sets $\Omega_1, \dots, \Omega_s$, respectively. Let $f : \Omega_1 \times \dots \times \Omega_s \rightarrow \mathbb{R}$ be a function. Suppose there exist constants $c_1, \dots, c_s \in \mathbb{R}$ such that, for any $i \in [s]$,*

$$|f(x_1, \dots, x_s) - f(x'_1, \dots, x'_s)| \leq c_i$$

for any $(x_1, \dots, x_s), (x'_1, \dots, x'_s) \in \Omega_1 \times \dots \times \Omega_s$ that differ only in the i th coordinate. Then the random variable $Z = f(X_1, \dots, X_s)$ satisfies the following. For every $\lambda \geq 0$,

$$\mathbb{P}[|Z - \mathbb{E}Z| \geq \lambda] \leq 2 \exp\left(\frac{-2\lambda^2}{c_1^2 + \dots + c_s^2}\right).$$

4 Non-extremal hypergraphs

We will first prove Theorem 2.6 (the non-extremal case) across the next six sections. In this section, we first provide a detailed proof sketch, and then briefly describe the principal results and techniques that we utilise. We finally state three key lemmas – Lemmas 4.1 to 4.3 – and show how they imply Theorem 2.6. The remaining five sections and Appendix A then prove these three chief lemmas.

4.1 Proof sketch for Theorem 2.5

Our goal is to show every $G \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$ has a Hamilton ℓ -cycle. For the sake of clarity, we only sketch the proof for tight cycles ($\ell = k - 1$, in which case $t = k$), and briefly discuss the necessary modifications for ℓ -cycles at the end. Our proof has three main steps

Tiling. As mentioned in Section 2.3.1, the main lemma behind our strategy, namely Lemma 4.1, provides a tiling of G with almost balanced blow-ups (defined in Section 4.2) of smaller k -graphs that obey (almost) the same minimum supported co-degree condition as G and also have an approximate version of the non-extremal structure of G . Moreover, these blown-up tiles will be linked in a cyclic fashion, where each tile shares a common edge with the previous and next tile in the chain, and is disjoint from the rest. Consequently, within each blown-up tile, we find a tight Hamilton path that starts at the edge shared with the previous tile, and ends at the edge shared with the next tile. The tight Hamilton paths in consecutive tiles can then be “linked together”, and repeating this process along the chain of cyclically linked tiles forms a tight Hamilton cycle in G .

Blowing up a fractional matching. Thus (morally speaking), our problem is reduced to finding a tight Hamilton path in a k -graph F^* that is an almost balanced blow-up of some $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, k-1)$, with the start and end edges of the path specified a priori. Our strategy starts by (nearly) partitioning F^* into a collection of balanced complete k -partite k -graphs (which will be suitable subgraphs of the k -graphs formed by blow-ups of edges of F) that cover almost all of F^* . To this end, we first find a perfect fractional matching in F , which is a set of non-negative edge weights such that the total weight on each vertex is precisely one. We do this via Farkas’ lemma, a standard linear algebraic technique. We do not describe how we use it to find the perfect fractional matching here as this part is fairly disjoint from the rest of the proof, we just mention that we make use of both the degree condition and the non-extremal structure of F . Now, for each vertex $v \in V(F)$, let $B_v \subseteq V(F^*)$ denote the blown-up part corresponding to v . We partition each B_v into parts corresponding to each edge $e \in E(F)$ that contains v such that the proportion of B_v covered by this part B_v^e is given by the weight of the edge e in the fractional matching. Thus, we end up with a subgraph T_e (whose parts are the sets B_v^e for all vertices $v \in e$) of the original k -partite k -graph formed by the blow-up of e . All the parts of T_e will have the same size (governed by the weight of e), and it is easy to see T_e has a spanning tight cycle. Ideally, since the fractional matching is perfect, these disjoint subgraphs will span all the vertices of G . However, we may end up having to round down the sizes of these parts B_v^e to integer values, meaning that we lose some vertices and consequently the k -graphs $\{T_e\}_{e \in E(F)}$ only cover almost all the vertices of G .

Connection and absorption. Finally, we find a tight path in F^* that can be used to link up the spanning tight paths found in the aforementioned k -partite subgraphs, and can also “absorb” into itself the remaining few uncovered vertices (technically, we first find such a path and then apply the strategy outlined in the previous paragraph to the leftover). In order to build such a tight path P , we will heavily rely on the fact that $\delta^*(F) > n/2$ (which will hold even in the ℓ -cycle case as $(k, \ell) \neq (3, 1)$ implies $t \geq 3$). Since $\delta^*(F) > n/2$, it is easy to argue that we can find a tight walk W in F that contains (several copies of) every edge of F . We can also specify the start and end edges, which, as discussed, is important to our proof. This is blown up to a tight path P in F^*

by walking along W and picking a new copy of a vertex whenever we see a repeated vertex, which is possible if the blow-up is sufficiently large relative to $|F|$. We then show that we can replace any segment of P formed by a blow-up of an edge e by a spanning path in the k -graph T_e , which allows us to connect the necessary tight paths. All that is left is to imbue W , and consequently P , with some absorption features. Given a vertex $v \in V(F)$, the “simple absorber” we use will consist of two supported sets W_1 and W_2 such that both W_1W_2 and W_1vW_2 are tight paths. In a manner similar to what we describe above, we can exploit the connectivity property due to the degree condition to ensure that the tight walk W , which contains every edge of F as a subwalk, also has several absorber subwalks W_1W_2 for each vertex v . The blown-up path P in F^* then retains these absorption attributes, allowing it to incorporate into itself the few unused vertices in F^* .

Adjustments for ℓ -cycles. The broad proof structure for ℓ -cycles remains the same, but the details need more care. To begin with, the complete k -partite subgraphs T_e will now be modified to have carefully chosen part sizes (described in Proposition 8.2) that will ensure a spanning ℓ -cycle instead. As mentioned in Section 2.3.1, this is done using a vertex-weighted generalisation of fractional matchings introduced by Mycroft and Zárate-Guerén [31], which will lead to the “correct” imbalance in the parts of T_e . The connectivity aspect of the proof remains fairly similar. In fact, we usually form the required ℓ -walks by building tight walks whose orders satisfy specific divisibility conditions such that the same vertex sequences will also correspond to ℓ -walks and ℓ -paths (see Observation 8.1). Additionally, while the starting point for the absorbers we use for ℓ -cycles are the same as the simple absorbers in the tight case, we need to be careful with the length of the walks and we need to be able to absorb $(k - \ell)$ vertices at a time, as detailed in Section 8.

4.2 Main lemmas and proof of Theorem 2.5

Before we state our three chief lemmas, we introduce necessary notation. Given a k -graph F , we say that a k -graph F^* is a *blow-up* of F if it can be obtained by replacing each vertex $v \in V(F)$ with an independent set B_v and each edge $v_1 \dots v_k \in E(F)$ with every possible k -partite edge across B_{v_1}, \dots, B_{v_k} , that is, with a complete k -partite k -graph across the corresponding parts. We call F^* a (γ, m) -*regular blow-up* if the size of each part lies in the interval $[(1 - \gamma)m, (1 + \gamma)m]$. We say that F^* is a (γ, m) -*nearly-regular blow-up* if all but at most one part has size in the interval $[(1 - \gamma)m, (1 + \gamma)m]$, and the remaining part (if it exists) has size exactly one. In particular, every (γ, m) -regular blow-up is a (γ, m) -nearly-regular blow-up. If F^* is a blow-up of F , we define the *projection map* $\phi : V(F^*) \rightarrow V(F)$ defined by $\phi(u) = v$ if u is a blow-up of v , that is, $u \in B_v$. For any set $U \subseteq V(F^*)$, we denote $\phi(U) = \{\phi(u) : u \in U\}$. Finally, we will often refer to blow-ups of ordered edges, which we formalise now. For an arbitrary ordered edge $e = v_1 \dots v_k \in E^*(F)$, we say that an ordered edge of F^* , say $e^* = u_1 \dots u_k \in E^*(F^*)$, is a *blow-up of e* if $u_i \in B_{v_i}$ for every $i \in [k]$.

Here is the general “tiling by blow-ups” lemma. It is very similar to Lemma 2.1 from [15], with the only difference being that here the lemma is stated for general hypergraph properties \mathcal{G} and \mathcal{S} , whereas in [15] \mathcal{G} is the family of n -vertex k -graphs G with $\delta^*(G) \geq (1/2 + \varepsilon)n$ and \mathcal{S} is the family of s -vertex k -graphs S with $\delta^*(S) \geq (1/2 + \varepsilon/2)s$. The more general statement given here follows from the proof of Lemma 2.1 in [15], by replacing the specific properties used there by \mathcal{G} and \mathcal{S} and using the assumption about the number of ways to extend a vertex in any $G \in \mathcal{G}$ to a hypergraph in \mathcal{S} ; we elaborate a bit more on this in Appendix A. We remark that Lang and Sanhueza-Matamala [24] provide a more general result (see [24, Proposition 6.3]), but with slightly stronger assumptions and a different conclusion.

Lemma 4.1. *Let $1/n \ll 1/m_2 \ll 1/m_1 \ll 1/s, \gamma \ll \varepsilon, 1/k \leq 1/2$. Let \mathcal{G} and \mathcal{S} be families of n -vertex and s -vertex k -graphs such that for every $G \in \mathcal{G}$ and every $W \subseteq V(G)$ of size at most $2k$, there are at least $(1 - 1/s^2) \binom{n-|W|}{s-|W|}$ sets $U \subseteq V(G)$ of size s that contain W , such that $G[U]$ is a copy of a graph in \mathcal{S} . Suppose that $G \in \mathcal{G}$. Then there exists a sequence of s -vertex k -graphs F_1, \dots, F_r and a sequence of subgraphs $F_1^*, \dots, F_r^* \subseteq G$ such that the following hold for all $i, j \in [r]$ with addition modulo r :*

- F_i is a copy of a k -graph in \mathcal{S} ,
- F_i^* is a (γ, m_i^*) -nearly-regular blow-up of F_i for some $m_i^* \in [m_1, m_2]$,
- $V(F_1^*) \cup \dots \cup V(F_r^*) = V(G)$,
- $V(F_i^*) \cap V(F_j^*) = \emptyset$ unless $j \in \{i-1, i, i+1\}$,
- $V(F_i^*) \cap V(F_{i+1}^*)$ consists of exactly k vertices that induce an edge in F_i^* and F_{i+1}^* that is disjoint from the singleton parts of both blow-ups (if they exist).

In order to use Lemma 4.1 in our setting, we will prove the following “structural inheritance” lemma.

Lemma 4.2. *Let $1/n \ll 1/s \ll \varepsilon \ll \mu' \ll \mu \ll 1/k \leq 1/3$ and let $\ell \in [k-1]$. Let $G \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$, let $W \subseteq V(G)$ be a subset of size at most $2k$, and let S be a uniformly random s -set from $V(G)$ that contains W . Then, with probability at least $1 - 1/s^2$, the subgraph $G[S]$ is in $\mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(s, k, \ell)$.*

As explained before, roughly speaking, our plan is to find a spanning ℓ -path in each blow-up in the structure given by Lemma 4.1, and then join up these paths. The following lemma shows that, under a necessary divisibility condition, every such blow-up indeed contains a spanning ℓ -path.

Lemma 4.3. *Let $1/m \ll 1/s \ll \gamma \ll \varepsilon \ll \mu \ll 1/k \leq 1/3$, let $\ell \in [k-1]$, and suppose $(k, \ell) \neq (3, 1)$. Suppose $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(s, k, \ell)$ and F^* is a (γ, m) -nearly-regular blow-up of F satisfying $|F^*| \equiv k$*

(mod $k - \ell$). Let $e_1, e_2 \in E(F^*)$ be two ordered edges that are vertex-disjoint from each other and from the singleton part of F^* (if it exists). Then there exists a spanning ℓ -path in F^* from e_1 to e_2 .

Assuming Lemmas 4.1 to 4.3, it is fairly straightforward to prove Theorem 2.5.

Proof of Theorem 2.5. Let us set up the chain of parameters $1/n \ll 1/m_2 \ll 1/m_1 \ll 1/s \ll \gamma \ll \varepsilon \ll \mu' \ll \mu \ll 1/k \leq 1/3$. Let $\mathcal{G} = \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$ and $\mathcal{S} = \mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(s, k, \ell)$. Then, by Lemma 4.2, we can apply Lemma 4.1 to $G \in \mathcal{G}$, and let $F_1, \dots, F_r \in \mathcal{S}$ and F_1^*, \dots, F_r^* be the k -graphs so obtained, where each F_i^* is a (γ, m_i^*) -nearly-regular blow-up of F_i for some $m_i^* \in [m_1, m_2]$.

For every $i \in [r]$, let $e_i = E(F_i^*) \cap E(F_{i+1}^*)$, and fix any ordering $v_{i,1} \dots v_{i,k}$ of e_i . In particular, we have $e_r = e_0 = E(F_r^*) \cap E(F_1^*)$ (as we think of the indices $i \in [r]$ cyclically modulo r). We define integers $s_1, \dots, s_r \in [k - \ell]$ and $t_1, \dots, t_r \in [\ell + 1, k]$, as follows. Take $s_1 = 1$. Given s_i , let t_i be the unique integer in $[\ell + 1, k]$ such that

$$|F_i^*| - (s_i - 1) - (k - t_i + 1) \equiv k \pmod{k - \ell}. \quad (4.4)$$

Finally, set $s_{i+1} = t_i - \ell$. Define $F_i^{**} := F_i^* - \{v_{i-1,1}, \dots, v_{i-1,s_i-1}, v_{i,t_i}, \dots, v_{i,k}\}$. Since e_{i-1} and e_i are both disjoint of the singleton part of F_i^* , it is immediate that F_i^{**} is a $(2\gamma, m_i^*)$ -nearly-regular blow-up of F_i , using $s_i + k - t_i \leq 2k \ll \gamma m_i^*$. Additionally, we remark that $\bigcup_{i \in [r]} V(F_i^{**}) = V(G)$. Indeed, let us define $U_i = V(F_i^*) \setminus V(F_i^{**})$. By choice, $U_i \subseteq V(e_{i-1}) \sqcup V(e_i)$. But then

$$U_i \cap V(e_i) = \{v_{i,t_i}, \dots, v_{i,k}\} \subseteq \{v_{i,s_{i+1}}, \dots, v_{i,k}\} \subseteq V(F_{i+1}^{**}),$$

and similarly $U_i \cap V(e_{i-1}) \subseteq V(F_{i-1}^{**})$. Hence, $\bigcup_{i \in [r]} V(F_i^{**}) = \bigcup_{i \in [r]} V(F_i^*) = V(G)$.

Now, our goal is to find a spanning ℓ -path, say P_i , in F_i^{**} from $S_i := v_{i-1,s_i} \dots v_{i-1,k}$ to $T_i := v_{i,1} \dots v_{i,t_i-1}$. To this end, we first find ordered edges $e_{S_i}, e_{T_i} \in E(F_i^{**})$ such that e_{S_i} starts with S_i , e_{T_i} ends with T_i , and e_{S_i} and e_{T_i} are vertex-disjoint from each other and from the singleton part of F_i^{**} , if it exists. Applying the supported co-degree condition $\delta^*(F_i) \geq (1 - 1/k - 2\varepsilon)s$ on $\phi(S_i)$ and $\phi(T_i)$ provides edges e'_{S_i} , and e'_{T_i} in F which satisfy the above requirements (with $\phi(S_i)$ and $\phi(T_i)$ in place of S_i and T_i respectively), and we let $e_{S_i}, e_{T_i} \in E(F_i^{**})$ be arbitrary blow-ups of e'_{S_i} , and e'_{T_i} respectively that contains S_i and T_i , respectively. Since $|F^{**}| \equiv k \pmod{k - \ell}$ (see (4.4)), Lemma 4.3 then provides the desired path P_i .

The key point to note is that, for each $i \in [r]$, the sequence of the final ℓ vertices of P_i exactly coincides with the sequence of the first ℓ vertices of P_{i+1} , namely $v_{i,t_i-\ell} \dots v_{i,t_i-1}$. Set Q_i to be the vertex sequence obtained by removing the last ℓ vertices of P_i . Let Q be the cyclic vertex sequence obtained by sequentially concatenating the sequences $Q_1 \dots Q_r$. Then, by the observation above, we may conclude that the vertex sequence Q forms an ℓ -cycle in G . Since each P_i spans F_i^{**} and $\bigcup_{i \in [r]} V(F_i^{**}) = V(G)$, the ℓ -cycle Q spans $V(G)$, completing the proof. \square

Organisation of the next five sections. Since Lemma 4.1 is based on [15, Lemma 2.1] with largely the same proof, we provide only a sketch and postpone it to Appendix A. We prove some useful facts about $\mathcal{E}_{n,k,\ell}^{\text{non-ext}}$ in Section 5, and use them to prove Lemma 4.2 in Section 6. Then, as discussed in Section 2.3.1, we describe how we find a suitable fractional matching and an absorbing and connecting walk in Sections 7 and 8, respectively. We put it all together to prove Lemma 4.3 in Section 9.

5 Basic tools

In this section we prove a few preliminary results that we will use later. Recall that $\mathcal{E}_{\varepsilon,\mu}^{\text{non-ext}}(n, k, \ell)$ is the family of n -vertex graphs with minimum positive co-degree at least $(1 - 1/t - \varepsilon)n$ where every $\lfloor n/t \rfloor$ vertices span at least μn^2 supported pairs. To begin with, we prove an approximate alternate description for $\mathcal{E}_{\varepsilon,\mu}^{\text{non-ext}}(n, k, \ell)$ that we use when proving Lemma 4.2. The aim of Lemma 4.2 is to show that, roughly speaking, the properties of graphs in $\mathcal{E}_{\varepsilon,\mu}^{\text{non-ext}}(n, k, \ell)$ are preserved under taking a random subset of vertices S . A naive approach to guaranteeing that the non-extremal properties are preserved takes a union bound over all subsets of S of size $|S|/t$, which would not yield an effective bound. Instead, using also the minimum supported co-degree condition, we show that it suffices to guarantee that non-neighbourhoods of vertices with vertex co-degree close to $(1 - 1/t)n$ have many supported pairs. This is formalised in the following proposition.

Proposition 5.1. *Let $k \geq 3$ and $\ell \in [k - 1]$ be integers, $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$, and let $1/n \ll \varepsilon \ll \mu' \ll \mu \ll 1/k$.*

- (i) *Suppose $G \in \mathcal{E}_{\varepsilon,\mu}^{\text{non-ext}}(n, k, \ell)$. Then for all $v \in V(G)$, either $d^1(v) \geq (1 - 1/t + \mu')n$ or $V(G) \setminus N^1(v)$ contains at least $\mu'n^2$ supported pairs.*
- (ii) *Let G be an n -vertex k -graph with no isolated vertices and $\delta^*(G) \geq (1 - 1/t - \varepsilon)n$. Suppose that for all $v \in V(G)$, either $d^1(v) \geq (1 - 1/t + 3\mu't)n$ or $V(G) \setminus N^1(v)$ contains at least $\mu'n^2$ supported pairs. Then every set of $\lfloor n/t \rfloor$ vertices in G contains at least $\mu'n^2$ supported pairs.*

Proof. Item (i) is fairly simple to prove. Suppose there exists some $v \in V(G)$ such that $d^1(v) \leq (1 - 1/t + \mu')n$, and hence $|V(G) \setminus N_G^1(v)| \geq n/t - \mu'n$. Let A be any set of the form $(V(G) \setminus N_G^1(v)) \sqcup B$, where B is an arbitrary set of distinct vertices from $N_G^1(v)$ such that $|B|$ is the smallest possible value required to make $|A| \geq \lfloor n/t \rfloor$; then $|B| \leq \mu'n$ (note that B is potentially empty). Since $G \in \mathcal{E}_{\varepsilon,\mu}^{\text{non-ext}}(n, k, \ell)$, the set A contains at least μn^2 supported pairs. As $|B| \leq \mu'n$, at most $\mu'n^2$ of these supported pairs have a vertex in B , and thus at least $(\mu - \mu')n^2 \geq \mu'n^2$ lie entirely within $V(G) \setminus N_G^1(v)$.

Now for (ii). Suppose there exists some $U \subseteq V(G)$ with $|U| \geq \lfloor n/t \rfloor \geq n/t - 1$ that contains less than $\mu'n^2$ supported pairs. Then, by the handshaking lemma on the graph of supported pairs in

U , there exists some $v \in U$ such that $d_U^1(v) \leq 3\mu'tn$. Since $|V(G) \setminus U| \leq n - n/t$, we see that $d_G^1(v) \leq (1 - 1/t + 3\mu't)n$, and hence, by assumption on G , the set $T = V(G) \setminus N_G^1(v)$ contains at least μn^2 supported pairs. Since G has no isolated vertices and $\delta^*(G) \geq (1 - 1/t - \varepsilon)n$, we know that $d_G^1(v) \geq (1 - 1/t - \varepsilon)n$, and hence $|T| \leq n/t + \varepsilon n$. By our choice of v , we have $|T \cap U| \geq |U| - 3\mu'tn \geq n/t - 4\mu'tn$, and consequently $|T \setminus U| \leq n/t + \varepsilon n - (n/t - 4\mu'tn) \leq 8\mu'tn$. Hence, the number of supported pairs in T which contain at least one vertex outside U is at most $8\mu'tn^2$. Consequently, $T \cap U \subseteq U$ contains at least $(\mu - 8\mu't)n^2 \geq \mu'n^2$ supported pairs, a contradiction to our assumption on U . \square

In a similar vein, but somewhat in the “opposite direction”, we also show that taking a blow-up approximately preserves non-extremal structure. Here we use the alternative characterisation of non-extremal graphs given by the previous proposition.

Proposition 5.2. *Let $0 \leq \gamma \ll \varepsilon \ll \mu' \ll \mu \ll 1/k \leq 1/3$, let $m, s \geq 1$ be integers, and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$. Suppose $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(s, k, \ell)$ and F^* is a (γ, m) -regular blow-up of F . Then $F^* \in \mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(|F^*|, k, \ell)$.*

Proof. Set $|F^*| = N$. By the definition of F^* , it is straightforward to see that $(1 - \gamma)ms \leq N \leq (1 + \gamma)ms$. First, notice that since F has no isolated vertices, neither does F^* . Second, since $\delta^*(F) \geq (1 - 1/t - \varepsilon)s$ and F^* is a (γ, m) -regular blow-up of F , we see that

$$\delta^*(F^*) \geq \left(1 - \frac{1}{t} - \varepsilon\right) (1 - \gamma)ms \geq \left(1 - \frac{1}{t} - 2\varepsilon\right) (1 + \gamma)ms \geq \left(1 - \frac{1}{t} - 2\varepsilon\right) N,$$

using $\gamma \ll \varepsilon$.

Finally, we need to show that any subset of $V(F^*)$ of size at least N/t contains at least $\mu'N^2$ supported pairs. We do this via Proposition 5.1. Fix η such that $\mu' \ll \eta \ll \mu$. Let $w \in V(F^*)$ be arbitrary, and let $\phi(w) = v \in V(F)$ (recall that $\phi : V(F^*) \rightarrow V(F)$ maps all copies in F^* of a vertex v in F to v). Since $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(s, k, \ell)$, by Proposition 5.1 (i), we know that either $d^1(v) \geq (1 - 1/t + \eta)s$ or $V(F) \setminus N^1(v)$ contains at least ηs^2 supported pairs. If the former is true, then since every vertex neighbour of v is blown up to at least $(1 - \gamma)m$ vertices in F^* , we have

$$d_{F^*}^1(w) \geq \left(1 - \frac{1}{t} + \eta\right) (1 - \gamma)ms \geq \left(1 - \frac{1}{t} + 3\mu't\right) (1 + \gamma)ms \geq \left(1 - \frac{1}{t} + 3\mu't\right) N,$$

where we use $N \leq (1 + \gamma)ms$ and $\mu' \ll \eta \ll 1/t$. In the other case, we conclude that the number of supported pairs in $V(F^*) \setminus N^1(w)$ is at least

$$\eta s^2 \cdot (1 - \gamma)^2 m^2 \geq \frac{\eta^2}{4} (1 + \gamma)^2 m^2 s^2 \geq \frac{\eta^2}{4} N^2.$$

Since $\mu' \ll \eta^2$, Proposition 5.1 (ii) implies that every set of at least N/t vertices contains at least $\mu' N^2$ supported pairs, as desired. \square

Next, we also present a simple proposition that helps provide a minimum degree condition based on the minimum supported co-degree. This is Proposition 1.4 of [31]. We include the proof here for completeness.

Proposition 5.3. *Let G be a k -graph. Then every supported set S in G with $|S| \leq k-1$ satisfies*

$$d(S) \geq \frac{\delta^*(G)^{k-|S|}}{(k-|S|)!}.$$

Proof. By definition, for every supported set S of size at most $k-1$, there are at least $\delta^*(G)$ vertices v such that $v \in N^1(S)$. By repeatedly applying this observation, we see that there must be at least $\delta^*(G)^{k-|S|}$ ordered sequences of vertices $(v_1, \dots, v_{k-|S|})$ such that $S \cup \{v_1, \dots, v_{k-|S|}\}$ is an edge of G . Hence, we have

$$d(S) \geq \frac{\delta^*(G)^{k-|S|}}{(k-|S|)!}. \quad \square$$

6 Structural inheritance

We now prove our structural inheritance lemma. As noted above, the crucial observation here is the alternative characterisation of non-extremal graphs given by Proposition 5.1 (ii).

Proof of Lemma 4.2. Write $W = \{w_1, \dots, w_r\}$, where $r = |W| \leq 2k$. Let y_1, \dots, y_{s-r} be $s-r$ vertices of $V(G)$ picked independently and uniformly at random (with repetition). Set $T = \{x_1, \dots, x_s\}$, where $x_i = w_i$ for $i \leq r$ and $x_i = y_{i-r}$ for $i > r$. Let $T_0 = T \setminus W = \{x_{r+1}, \dots, x_s\}$ and recall that S is a randomly uniform set of s vertices in $V(G)$ that contains W . Let \mathcal{A} be the event that x_1, \dots, x_s are distinct. Let \mathcal{B}_1 and \mathcal{B}_2 be the events that $G[S] \in \mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(s, k, \ell)$ and $G[T] \in \mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(s, k, \ell)$ respectively. Then, we first see that

$$\mathbb{P}[\mathcal{A}^c] \leq \frac{s^2}{n} = o(1).$$

Consequently, we have

$$\mathbb{P}[\mathcal{B}_1^c] = \mathbb{P}[\mathcal{B}_2^c \mid \mathcal{A}] = \frac{\mathbb{P}[\mathcal{B}_2^c \cap \mathcal{A}]}{\mathbb{P}[\mathcal{A}]} \leq \mathbb{P}[\mathcal{B}_2^c](1 + o(1)).$$

Hence, it will suffice to show that $\mathbb{P}[\mathcal{B}_2^c] \leq \frac{1}{2s^2}$.

Let $H = G[T]$ and fix parameters μ_1, μ_2 such that $\mu' \ll \mu_1 \ll \mu_2 \ll \mu$. We define the following two events.

- $\mathcal{C}_1 = \left\{ \frac{d_H^1(U)}{s} \geq \frac{d_G^1(U)}{n} - \varepsilon \text{ for every } U \subseteq T \text{ with } |U| \leq k-1 \text{ which is supported in } G \right\},$
- $\mathcal{C}_2 = \left\{ \text{if } x \in T \text{ satisfies } d_G^1(x) \leq (1 - 1/t + \mu_2)n \text{ then } T \setminus N_H^1(x) \text{ has at least } \mu_1 s^2 \text{ supported pairs} \right\}.$

First, we will argue that $\mathcal{C}_1 \cap \mathcal{C}_2$ implies \mathcal{B}_2 . It will subsequently suffice to show that each \mathcal{C}_i occurs with high probability, and then use a union bound.

Indeed, from \mathcal{C}_1 and from $G \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$, we know that H has no isolated vertices and that $\delta^*(H) \geq (1 - 1/t - 2\varepsilon)s$. Now, consider any $x \in V(H)$ such that $d_H^1(x) \leq (1 - 1/t + 3\mu't)s$. From \mathcal{C}_1 , we see that $d_G^1(x) \leq (1 - 1/t + 3\mu't + \varepsilon)n \leq (1 - 1/t + \mu_2)n$, and from \mathcal{C}_2 we conclude that $V(H) \setminus N_H^1(x)$ contains at least $\mu_1 s^2$ supported pairs. Then by Proposition 5.1 (ii), we know that every $\lfloor s/t \rfloor$ -subset in $V(H)$ has at least $\mu' s^2$ supported pairs. Thus, overall, we have $H \in \mathcal{E}_{2\varepsilon, \mu'}^{\text{non-ext}}(s, k, \ell)$.

Let us start with \mathcal{C}_1 . Let $I \subseteq [s]$ be any subset such that $|I| \in [k-1]$. Define the random variables $X_I = \{x_i : i \in I\}$ and $d_I = |N_H^1(X_I) \cap T_0|$. Consider the event

$$\mathcal{E}_I = \{X_I \text{ is supported in } G \text{ and } d_H^1(X_I)/s \leq d_G^1(X_I)/n - \varepsilon\}.$$

Let $\mathcal{U}_I = \{U \subseteq V(G) : |U| \leq |I|, \mathbb{P}[X_I = U] > 0\}$. Then, observe that

$$\begin{aligned} \mathbb{P}[\mathcal{E}_I] &= \sum_{\substack{U \subseteq V(G) \\ |U| \leq |I|}} \mathbb{P}[\mathcal{E}_I \cap \{X_I = U\}] \\ &= \sum_{U \in \mathcal{U}_I} \mathbb{P}[\mathcal{E}_I \mid X_I = U] \cdot \mathbb{P}[X_I = U]. \end{aligned}$$

In particular, to show that $\mathbb{P}[\mathcal{E}_I] \leq e^{-\Omega(s)}$, we may consider any arbitrary $U \in \mathcal{U}_I$, condition on $X_I = U$, and argue that $\mathbb{P}[\mathcal{E}_I \mid X_I = U] \leq e^{-\Omega(s)}$ for any such choice of U .

If U is not supported in G , then $\mathbb{P}[\mathcal{E}_I \mid X_I = U] = 0$, and hence we may assume that it is supported. We may lower bound the expected value of d_I as

$$\begin{aligned} \mathbb{E}d_I &= d_G^1(U) \cdot \left(1 - \left(1 - \frac{1}{n} \right)^{s - |U \cup W|} \right) \\ &\geq \frac{(s - 3k) \cdot \delta^*(G)}{2n} \\ &\geq \frac{s - 3k}{2} \cdot \left(1 - \frac{1}{t} - \varepsilon \right) \geq \frac{s}{4}, \end{aligned}$$

using $t \geq 3$ and $k \ll s$. Notice that changing the outcome of any single x_i changes d_I by at most

1, and so by McDiarmid's inequality (Lemma 3.4),

$$\mathbb{P}\left[d_I - \mathbb{E}d_I \leq -\frac{\varepsilon \mathbb{E}d_I}{2}\right] \leq 2 \exp\left(\frac{-2\varepsilon^2(\mathbb{E}d_I)^2}{s}\right) = e^{-\Omega(s)}.$$

Thus, since $d_H(X_I) \geq d_I$, with probability at least $1 - e^{-\Omega(s)}$, we have

$$\frac{d_H(X_I)}{s} \geq \frac{d_I}{s} \geq \frac{1 - \varepsilon/2}{s} \cdot (s - r) \cdot \frac{d}{n} \geq \left(1 - \frac{\varepsilon}{2}\right) \left(1 - \frac{k}{s}\right) \cdot \frac{d}{n} \geq \frac{d}{n} - \varepsilon,$$

where we use that $k/s \ll \varepsilon$. Hence, as argued previously, we now have $\mathbb{P}[\mathcal{E}_I] \leq e^{-\Omega(s)}$ for every $I \subseteq [s]$ with $|I| \in [k-1]$. Taking a union bound over all possibilities of I , we see that $\mathbb{P}[\mathcal{C}_1] \geq 1 - (k-1) \binom{s}{k-1} e^{-\Omega(s)} \geq 1 - \frac{1}{4s^2}$.

Next, we will show that \mathcal{C}_2 occurs with high probability. For each $i \in [s]$, define the event

$$\mathcal{F}_i = \{T \setminus N_H^1(x_i) \text{ has less than } \mu_1 s^2 \text{ supported pairs and } d_G^1(x_i) < (1 - 1/t + \mu_2)n\}.$$

Fix any integer $i \in [r+1, s]$. Similar to our previous calculations, we note that

$$\mathbb{P}[\mathcal{F}_i] = \sum_{u \in V(G)} \mathbb{P}[\mathcal{F}_i \cap \{x_i = u\}] = \sum_{u \in V(G)} \mathbb{P}[\mathcal{F}_i \mid x_i = u] \cdot \mathbb{P}[x_i = u].$$

Consider any arbitrary $u \in V(G)$. We will argue that, for each fixed $i \in [r+1, s]$, we have $\mathbb{P}[\mathcal{F}_i \mid x_i = u] \leq e^{-\Omega(s)}$, and then the above equation then guarantees that $\mathbb{P}[\mathcal{F}_i] \leq e^{-\Omega(s)}$. We comment at the end of the proof on how to handle \mathcal{F}_i for $i \in [r]$.

If $d_G^1(u) \geq (1 - 1/t + \mu_2)n$, then $\mathbb{P}[\mathcal{F}_i \mid x_i = u] = 0$, and we are done. Hence, we may suppose $d_G^1(u) \leq (1 - 1/t + \mu_2)n$. Since $G \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$, by Proposition 5.1 (i), there must be at least $\mu_2 n^2$ supported pairs in $V(G) \setminus N_G^1(u)$. Let us call this set of supported pairs P . From Proposition 5.3, any supported pair $w_1 w_2 \in P$ satisfies

$$d_G(w_1 w_2) \geq \frac{\delta^*(G)^{k-2}}{(k-2)!} \geq \left(\frac{\delta^*(G)}{k-2}\right)^{k-2} \geq \left(\frac{n - n/t - \varepsilon n}{k-2}\right)^{k-2} \geq \left(\frac{n}{3k}\right)^{k-2}$$

since $t \geq 2$. Let E_P be the set of all edges of G containing at least one pair of P . We can bound the number of such edges from below as

$$|E_P| \geq \frac{1}{\binom{k}{2}} \sum_{w_1 w_2 \in P} d_G(w_1 w_2) \geq \frac{2\mu_2}{(3k)^{k-2} \cdot k^2} \cdot n^k \geq \frac{2\mu_2}{(3k)^k} \cdot n^k.$$

For each edge $e \in E_P$, the probability that $e \in E(H[T_0 \setminus \{x_i\}])$ is

$$\begin{aligned} & \mathbb{P} \left[\bigcup_{J \subseteq [r+1, s] \setminus \{i\}, |J|=k} \{\{x_j : j \in J\} = e\} \right] \\ & \geq \sum_{\substack{J \subseteq [r+1, s] \setminus \{i\}, \\ |J|=k}} \mathbb{P}[\{x_j : j \in J\} = e] - \sum_{\substack{J, J' \subseteq [r+1, s] \setminus \{i\}, \\ |J|=|J'|=k, J \neq J'}} \mathbb{P}[\{x_j : j \in J\} = \{x_{j'} : j' \in J'\} = e] \\ & \geq \frac{(s-r-1)(s-r-2) \dots (s-r-k)}{n^k} - \frac{((s-r-1)(s-r-2) \dots (s-r-k))^2}{n^{k+1}} \geq \frac{s^k}{2n^k}, \end{aligned}$$

where we used that if J, J' are two distinct sets of size k then $|J \cup J'| \geq k+1$. We thus obtain the following lower bound on $Z := |E_P \cap E(H[T_0 \setminus \{x_i\}])|$.

$$\mathbb{E}Z \geq \sum_{e \in E_P} \mathbb{P}[e \in E(H[T_0 \setminus \{x_i\}])] \geq \frac{\mu_2}{(3k)^k} \cdot s^k = \beta \mu_2 s^k.$$

where $\beta = 1/(3k)^k$ is a function of k . Next, observe that if a single vertex in $T_0 \setminus \{x_i\}$ is replaced by an arbitrary vertex of $V(G)$, then the number of edges of E_P in $H[T_0 \setminus \{x_i\}]$ changes by at most s^{k-1} since this bounds the maximum degree of a vertex in H . Using McDiarmid's inequality (Lemma 3.4), we have

$$\mathbb{P}[Z - \mathbb{E}Z \leq -\frac{1}{2} \cdot \mathbb{E}Z] \leq 2 \exp \left(\frac{-\beta^2 \mu_2^2 s^{2k}}{2 \cdot s \cdot s^{2k-2}} \right) \leq \exp(-\Omega(s)).$$

It readily follows that $Z \geq \beta \mu_2 s^k / 2$ with probability at least $1 - e^{-\Omega(s)}$. Each edge of E_P counted in Z contains at least one pair of P , and each pair of P is counted at most s^{k-2} times. Hence, with probability at least $1 - e^{-\Omega(s)}$, there are at least $\beta \mu_2 s^2 / 2 \geq \mu_1 s^2$ supported pairs of P in H . By choice, all of these lie outside $N_H^1(x_i) = N_H^1(u)$.

Thus, as noted initially, we conclude that $\mathbb{P}[\mathcal{F}_i] \leq e^{-\Omega(s)}$ for all $i \in [r+1, s]$. Finally, if $i \in [r]$, we have $\mathbb{P}[\mathcal{F}_i \cap \{x_i = u\}] = 0$ for all $u \neq w_i$, which means that $\mathbb{P}[\mathcal{F}_i] = \mathbb{P}[\mathcal{F}_i \mid x_i = w_i]$, and hence we can perform the same calculations as above with $u = w_i$. Additionally, here we will instead have $Z = E_P \cap E(H[T_0])$, but the bounds are identical anyway. Hence, by taking a union bound over all $i \in [s]$, we conclude that $\mathbb{P}[\mathcal{C}_2] \geq 1 - se^{-\Omega(s)} \geq 1 - \frac{1}{4s^2}$.

Since $\mathcal{C}_1 \cap \mathcal{C}_2$ implies \mathcal{B}_2 , we use a union bound and see that

$$\mathbb{P}[\mathcal{B}_2^c] \leq \mathbb{P}[\mathcal{C}_1^c \cup \mathcal{C}_2^c] \leq \frac{1}{2s^2},$$

and thus $\mathbb{P}[\mathcal{B}_2] \geq 1 - 1/2s^2$, as required. \square

7 Perfect weighted fractional matchings

Our overall approach will rely on finding a large, nearly perfect matching in a regular blow-up F^* of some $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$, and finding a way to connect the edges of this matching via an ℓ -path. Towards that end, we make use of a recently introduced idea of Mycroft and Zárate-Guerén [31] that they call *perfect weighted fractional matchings*, and we borrow their definitions and notation.

To begin with, we define a *fractional matching* in a k -graph H to be a weight function that maps the edges of H to the unit interval, say $q : E(H) \rightarrow [0, 1]$, such that for every vertex $v \in V(H)$ the sum of $q(e)$ for all edges e containing v is at most 1, that is,

$$\sum_{\substack{e \in E(H) \\ e \ni v}} q(e) \leq 1.$$

We say that q is a *perfect fractional matching* if this sum is precisely 1 for every vertex $v \in V(H)$. Note that a (perfect) matching is a special type of a (perfect) fractional matching where the edge weights are restricted to the set $\{0, 1\}$.

Now, suppose we have k weights $w_1, \dots, w_k \in [0, 1]$. Then a (w_1, \dots, w_k) -*fractional matching* is similar to a fractional matching q , except that any edge e places weight $w_i q(e)$ on its i th vertex. Of course, we now need to consider edges with a pre-defined vertex ordering. Hence, instead, we will simply consider all possible permutations of the vertices of each edge.

Recall from Section 3.1 that, for a given a k -graph H , we let $E^*(H)$ denote the set of ordered edges of H . For every vertex $v \in V(H)$ and $i \in [k]$, we write $E_i^v(H) \subseteq E^*(H)$ for the set of all ordered edges with v as the i th vertex. Fix $\mathbf{w} = (w_1, \dots, w_k) \in \mathbb{R}_{\geq 0}^k$. A \mathbf{w} -*weighted fractional matching* is a weight function $q : E^*(H) \rightarrow \mathbb{R}$ such that $q(e) \geq 0$ for all $e \in E^*(H)$, and

$$\sum_{i \in [k]} \sum_{e \in E_i^v(H)} w_i \cdot q(e) \leq 1$$

for every vertex $v \in V(H)$. We call it *perfect* if we have equality in the above inequality for all $v \in V(H)$. Observe that a (perfect) fractional matching is simply a (perfect) $(1, \dots, 1)$ -weighted fractional matching. Following [31], we will focus on perfect weighted fractional matchings with weights $\mathbf{w}^* = (w_1, \dots, w_k)$, defined by

$$w_i = \begin{cases} k-1 & i = 1 \\ t-1 & 2 \leq i \leq k, \end{cases}$$

where

$$t := t(k, \ell) = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell).$$

The reason for this particular choice of weights is that any complete k -uniform k -partite hypergraph

with parts A, B_1, \dots, B_{k-1} , such that $|A|/(k-1) = |B_i|/(t-1)$ for every $i \in [k-1]$ (and an additional divisibility condition holds), has a spanning ℓ -cycle (see Proposition 8.2 and Proposition 1.8 in [31]). This means that if an ordered edge $e = v_1 \dots v_k$ receives some weight $q(e)$, then there is an ℓ -cycle that we can use cover almost all the vertices in the blow-up of e where v_i is blown-up by $w_i \lfloor q(e) \rfloor$ vertices, for $i \in [k]$, a key step in our proof.

A convenient and often elegant method to find perfect fractional matchings is using Farkas' lemma, a linear algebraic result that is quite handy for this purpose – see [18, 29, 36] for examples. In order to state the lemma, we introduce some simple notation. Given a vector $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$ and any $a \in \mathbb{R}$, we write $\mathbf{v} \geq a$ to mean that $v_i \geq a$ for all $i \in [n]$. We use analogous notation for $\mathbf{v} \leq a$, $\mathbf{v} > a$ and $\mathbf{v} < a$. Finally, we use $\mathbf{1}$ to denote the all ones vector.

Lemma 7.1 (Farkas' lemma). *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Then either there exists $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq 0$, or there exists $\mathbf{y} \in \mathbb{R}^m$ such that $\mathbf{A}^\top \mathbf{y} \geq 0$ and $\mathbf{b}^\top \mathbf{y} < 0$.*

Using this, we can show that non-extremal k -graphs have a \mathbf{w}^* -weighted perfect fractional matching.

Lemma 7.2. *Let $k \geq 3$ and $\ell \in [k-1]$ be integers, $t = \lfloor \frac{k}{k-\ell} \rfloor (k-\ell)$, and let $1/n \ll \varepsilon \ll 1/k$, where n is divisible by t . Let $\mathbf{w}^* = (w_1, \dots, w_k)$ where $w_1 = k-1$ and $w_i = t-1$ for all $2 \leq i \leq k$. Then every $H \in \mathcal{E}_{\varepsilon, 4\varepsilon}^{\text{non-ext}}(n, k, \ell)$ has a \mathbf{w}^* -weighted perfect fractional matching.*

Proof. Let $V(H) = \{v_1, \dots, v_n\}$. For each edge $e \in E(H)$ and each $v_i \in e$, define $\chi_i(e) \in \mathbb{R}^n$ to be the vector with i th coordinate $k-1$, j th coordinate $t-1$ for all j such that $v_j \in e \setminus \{v_i\}$ and all other coordinates 0. Let $\mathcal{X} = \{\chi_i(e) : e \in E(H), v_i \in e\}$.

Suppose H does not have a \mathbf{w}^* -weighted perfect fractional matching. This means that there is no set of weights $\{q(\mathbf{x}) : q(\mathbf{x}) \geq 0, \mathbf{x} \in \mathcal{X}\}$ which satisfies $\sum_{\mathbf{x} \in \mathcal{X}} q(\mathbf{x}) \cdot \mathbf{x} = \mathbf{1}$, for otherwise, defining for each edge $e = \{v_{i_1}, \dots, v_{i_k}\}$ and $j \in [k]$,

$$q'(v_{i_j} v_{i_1} \dots v_{i_{j-1}} v_{i_{j+1}} \dots v_{i_k}) = q(\chi_{i_j}(e)),$$

(and $q'(e) = 0$ for all other ordered edges corresponding to e), would constitute a \mathbf{w}^* -weighted perfect fractional matching. Hence, by Lemma 7.1, there exists $\mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{y} \cdot \mathbf{1} < 0$ and $\mathbf{y} \cdot \mathbf{x} \geq 0$ for every $\mathbf{x} \in \mathcal{X}$.

Let $\mathbf{y} = (y_1, \dots, y_n)$. Then $\mathbf{y} \cdot \mathbf{1} = \sum y_i < 0$. By potentially reordering the vertices, we may assume without loss of generality that $y_i \leq y_j$ for all $i \leq j$. Let r be the smallest index such that $v_r \in N^1(v_1)$, that is, $v_1 v_r$ is a supported pair. Since $H \in \mathcal{E}_{\varepsilon, 4\varepsilon}^{\text{non-ext}}(n, k, \ell)$, we know that $\delta^*(H) \geq (1 - 1/t - \varepsilon)n$, and hence $r \leq n/t + \varepsilon n + 1 \leq n/t + 2\varepsilon n$.

Overall, we hope to obtain a contradiction by finding a “low-weight” edge of H , say $e = v_{i_1} \dots v_{i_k}$, which satisfies the following two key properties.

- $\frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \cdot ((k-1)y_{i_1} + (t-1)y_{i_2}) \leq y_1 + \dots + y_{2n/t}$, and

- $i_j \leq 2n/t + 1$ for all $j \geq 3$.

Indeed, given any such edge e , we let $\mathbf{x} \in \mathbb{R}^n$ be the vector with $x_{i_1} = k - 1$, $x_{i_j} = t - 1$ for all $2 \leq j \leq k - 1$, and zero otherwise. Then clearly $\mathbf{x} = \chi_{i_1}(e) \in \mathcal{X}$. Consequently, we obtain a contradiction as

$$\begin{aligned}
0 &\leq \frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \cdot \mathbf{y} \cdot \mathbf{x} \\
&= \frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \left((k-1)y_{i_1} + \sum_{j \geq 2} (t-1)y_{i_j} \right) \\
&= \frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \cdot ((k-1)y_{i_1} + (t-1)y_{i_2}) + \frac{n}{t} \cdot \left(\sum_{j \geq 3} \frac{t-2}{k-2} \cdot y_{i_j} \right) \\
&\leq \left(y_1 + \cdots + y_{\frac{2n}{t}} \right) + \frac{n}{t} \cdot (t-2) \cdot y_{\frac{2n}{t}+1} \\
&\leq \left(y_1 + \cdots + y_{\frac{2n}{t}} \right) + \left(y_{\frac{2n}{t}+1} + \cdots + y_n \right) \\
&< 0,
\end{aligned}$$

where the penultimate inequality uses $2 \leq t \leq k$ and the first and final ones follow from the assumptions on \mathbf{y} .

Let us try to find such an e . First, suppose $r \leq n/t$. Set $i_1 = 1$ and $i_2 = r$. By repeatedly using the minimum supported co-degree condition, we can find vertices v_{i_3}, \dots, v_{i_k} such that $v_{i_j} \in N^1(v_{i_1}, \dots, v_{i_{j-1}})$ and $i_j \leq n/t + \varepsilon n + 1$ for all $j \geq 3$. Observe that $e = v_{i_1} \dots v_{i_k} \in E(G)$. Furthermore, we have

$$\begin{aligned}
\frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \cdot ((k-1)y_{i_1} + (t-1)y_{i_2}) &= \frac{n}{t} \left(\frac{(t-2)(k-1)}{(t-1)(k-2)} \cdot y_1 + \frac{t-2}{k-2} \cdot y_{\frac{n}{t}+1} \right) \\
&\leq \frac{n}{t} \left(y_1 + y_{\frac{n}{t}+1} \right) \\
&\leq \left(y_1 + \cdots + y_{\frac{n}{t}} \right) + \left(y_{\frac{n}{t}+1} + \cdots + y_{\frac{2n}{t}} \right),
\end{aligned}$$

where we use $2 \leq t \leq k$. Thus, e meets the required criteria.

Hence, we may assume that $n/t + 1 \leq r \leq n/t + 2\varepsilon n$. Let $A = \{v_1, \dots, v_r\}$ and $B = \{v_{r+1}, \dots, v_{2n/t}\}$. Since $H \in \mathcal{E}_{\varepsilon, 4\varepsilon}^{\text{non-ext}}(n, k, \ell)$ and $|A| = r \geq n/t$, there are at least $4\varepsilon n^2$ supported pairs in A . This implies that we can greedily find a matching M_1 of supported pairs in A that contains $v_1 v_r$ and which spans at least $4\varepsilon n$ vertices.

Now, we note that $|B| = 2n/t - r$ and $|A \setminus V(M_1)| \leq r - 4\varepsilon n$. Since $r \leq n/t + 2\varepsilon n$, we see that $|B| \geq |A \setminus V(M_1)|$. Hence, we can pair up each vertex of $|A \setminus V(M_1)|$ with a distinct vertex of B arbitrarily (note that these vertex pairs need not be supported in H). Add this collection of pairs M_1 and call the resulting set of vertex pairs M_2 .

For our final step, we consider all vertices in $C = \{v_1, \dots, v_{2n/t}\} \setminus V(M_2)$. By construction every vertex in $\{v_1, \dots, v_r\}$ is part of some pair in M_2 . Thus, we may deduce that $C \subseteq \{v_{r+1}, \dots, v_{2n/t}\} = B$. Since $|V(M_2)|$ is even, the same must be true of $|C| = 2n/t - |V(M_2)|$. Hence, we can arbitrarily pair up all vertices in C . We add this collection of vertex pairs to M_2 , and let the resulting set be M_3 . Since every vertex in $v_1, \dots, v_{2n/t}$ is contained in exactly one pair of M_3 , we have $|M_3| = n/t$.

Let $v_{i_1}v_{i_2}$ to be a supported pair in M_1 such that $(k-1)y_{i_1} + (t-1)y_{i_2} \leq (k-1)y_i + (t-1)y_j$ for all $v_iv_j \in M_1$ with $i < j$. Since $v_1v_r \in M_1$ and $i_1 \geq 1$, we have $i_2 \leq r$. By construction, for any pair $v_iv_j \in M_3 \setminus M_1$ with $i < j$, we have $v_j \in B$, and hence $j > r$. This implies

$$(k-1)y_{i_1} + (t-1)y_{i_2} \leq (k-1)y_1 + (t-1)y_r \leq (k-1)y_i + (t-1)y_j$$

for all $v_iv_j \in M_3$, where the first inequality uses that $v_1v_r \in M_1$ and the definition of $v_{i_1}v_{i_2}$.

Since $v_{i_1}v_{i_2}$ is a supported pair, by repeatedly using the minimum supported co-degree condition, we can find an edge $v_{i_1}v_{i_2} \dots v_{i_k} \in E(H)$ such that $v_{i_j} \in N^1(v_{i_1}, \dots, v_{i_{j-1}})$ and $i_j \leq n/t + \varepsilon n + 1 \leq 2n/t + 1$ for all $j \geq 2$. Additionally, we have

$$\begin{aligned} \frac{n}{t} \cdot \frac{t-2}{(t-1)(k-2)} \cdot ((k-1)y_{i_1} + (t-1)y_{i_2}) &= \frac{t-2}{(t-1)(k-2)} \cdot |M_3| \cdot ((k-1)y_{i_1} + (t-1)y_{i_2}) \\ &\leq \frac{t-2}{(t-1)(k-2)} \sum_{v_iv_j \in M_3} ((k-1)y_i + (t-1)y_j) \\ &\leq \sum_{v_iv_j \in M_3} \left(\frac{(t-2)(k-1)}{(t-1)(k-2)} y_i + \frac{t-2}{k-2} y_j \right) \\ &\leq \sum_{v_iv_j \in M_3} (y_i + y_j) \\ &= y_1 + \dots + y_{\frac{2n}{t}}, \end{aligned}$$

where we have used $2 \leq t \leq k$ as before. Since e satisfies the desired conditions, the proof is complete. \square

Notice that in the previous lemma we assume (for convenience) that n is divisible by t . One can remove the divisibility assumption by taking an appropriate blow-up, as follows.

Corollary 7.3. *Let $k \geq 3$ and $\ell \in [k-1]$ be integers, $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$, and let $1/n \ll \varepsilon \ll \mu \ll 1/k$. Let $\mathbf{w}^* = (w_1, \dots, w_k)$ where $w_1 = k-1$ and $w_i = t-1$ for all $2 \leq i \leq k$. Then every $H \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$ has a \mathbf{w}^* -weighted perfect fractional matching.*

Proof. Let H^* be the t -blow-up of H , namely where each vertex is replaced by an independent set of size t . Then, clearly, $|H^*|$ is divisible by t . Moreover, by Proposition 5.2 and because H^* is a $(0, t)$ -blow-up of H , we have that $H^* \in \mathcal{E}_{\varepsilon, 4\varepsilon}^{\text{non-ext}}(nt, k, \ell)$. Thus, by Lemma 7.2 the graph H^*

has a \mathbf{w}^* -weighted perfect fractional matching q . This can be translated to a \mathbf{w}^* -weighted perfect fractional matching q' in H as follows. For any ordered edge $e \in E^*(H)$, let $F_e \subseteq E^*(H^*)$ be the collection of all ordered edges of H^* that are blow-ups of e . Define

$$q'(e) = \frac{1}{t} \sum_{e^* \in F_e} q(e^*).$$

It is fairly straightforward to verify that q' is a \mathbf{w}^* -weighted perfect fractional matching in H . \square

8 Finding an absorbing and connecting walk

We first recall the definitions of ℓ -walks and ordered supported sets from Section 3.1. We begin this section by introducing some convenient notation and terminology. Given a tight walk $W = v_1 \dots v_r$, we say that two subwalks $W_1 = v_i \dots v_{i+t}$ and $W_2 = v_j \dots v_{j+s}$ are *disjoint* if $j > i + t$ or $i > j + s$, i.e. if their index sets are disjoint (hence, they need not be disjoint as vertex sets). Given vertex sequences (not necessarily ℓ -walks) W_1, \dots, W_s , we let $W_1 \dots W_s$ denote the concatenation of W_1, \dots, W_s in order. Next, if $S = u_1 \dots u_s$ and $T = v_1 \dots v_t$ are ordered supported sets, we say that an ℓ -walk W *joins* S and T if its first s vertices are $u_1 \dots u_s$ and its last t vertices are $v_1 \dots v_t$.

We first note that any tight walk of suitable order contains an ℓ -walk.

Observation 8.1. *Let $k \geq 2$ and $\ell \in [k-1]$. Suppose $W = v_1 \dots v_r$ is a k -uniform tight walk in a k -graph G such that $r \equiv k \pmod{k-\ell}$. Then $v_1 \dots v_r$ is an ℓ -walk in G comprised of edges of the form $v_{i+1} \dots v_{i+k}$, whenever $i \equiv 0 \pmod{k-\ell}$ and $i \in [0, r-k]$.*

Our strategy for proving Lemma 4.3 will first involve replacing each edge $e \in E(F)$ with a large complete k -partite k -graph (the blow-up of e) in F^* , where the sizes of the parts will be determined based on a perfect weighted fractional matching that we obtain via Corollary 7.3. We will consider a suitable ℓ -path in F^* that spans all the vertices of this k -partite k -graph. Then, we will use a “connecting path” to move to the next k -partite k -graph, and repeat this procedure. The following proposition will be central to finding spanning ℓ -paths in suitable k -partite k -graphs.

Proposition 8.2. *Let H be a complete k -partite k -graph on r vertices with parts A, B_1, \dots, B_{k-1} , where r is at least 1 and is divisible by $t(k-1)$. Let $|A| = r/t$ and $|B_i| = r(t-1)/t(k-1)$ for all $i \in [k-1]$. Then there exists a spanning ℓ -cycle in H which contains an ordered edge $e \in B_1 \times \dots \times B_{t-1} \times A \times B_t \times \dots \times B_{k-1}$.*

Proof. Set $q = (t-1)(k-1)$. We define a sequence $X_1 \dots X_{q+k-1}$ of (non-distinct) parts belonging to the set $\{A, B_1, \dots, B_{k-1}\}$ as

$$X_1 \dots X_{q+k-1} = B_1 \dots B_{t-1} A^{(1)} B_t \dots B_{2t-2} A^{(2)} \dots A^{(k-2)} B_{q-t+2} \dots B_q A^{(k-1)},$$

where $A^{(i)} = A$, and, $B_i = B_j$ for the unique $j \in [k-1]$ such that $i \equiv j \pmod{k-1}$, for every i (that is, we think of the indices $B_1 \dots B_q$ cyclically modulo $k-1$). To begin with, we argue that we can sequentially list all the vertices of H as $v_1 \dots v_r$ with $v_i \in X_i$ for all i , where we view the indices of X_i cyclically modulo $q+k-1$.

Indeed, we first note that $r \equiv 0 \pmod{q+k-1}$, since $q+k-1 = t(k-1)$ divides r by assumption. This means that the sequence $X_1 \dots X_r$ is formed by repeating the sequence $X_1 \dots X_{q+k-1}$ exactly $r/(q+k-1) = r/t(k-1)$ times. The set A appears $k-1$ times in $X_1 \dots X_{q+k-1}$, and hence appears exactly $r/t = |A|$ times in $X_1 \dots X_r$. Similarly, each B_i appears precisely $t-1$ times in $X_1 \dots X_{q+k-1}$, and hence there are $r(t-1)/t(k-1) = |B_i|$ instances of B_i in $X_1 \dots X_r$. Thus, we can find the desired sequence $v_1 \dots v_r$.

Since r is divisible by t , we know that $r \equiv 0 \pmod{k-\ell}$. We claim that $v_1 \dots v_r$ is the required spanning ℓ -cycle in H . It suffices to show that $v_{i+1} \dots v_{i+k}$ forms an edge for all $i \leq r-k$ such that $i \equiv 0 \pmod{k-\ell}$.

Let us first fix any subsequence $v_{i+1} \dots v_{i+k}$ of k consecutive vertices such that $i \equiv 0 \pmod{k-\ell}$. By construction, for each $i \in [k-1]$, any k consecutive vertices in $v_1 \dots v_r$ can contain at most one vertex from B_i . By Proposition 3.2, we know that $v_{i+1} \dots v_{i+k}$ contains at most one vertex of the form v_j with $j \equiv 0 \pmod{t}$, and consequently at most one vertex from A . Hence, $v_i \dots v_{i+k-1}$ contains exactly one vertex from each of A, B_1, \dots, B_{k-1} , implying that it forms an edge of H . For the final part of the proposition, notice that $v_1 \dots v_k$ is an ordered edges in the desired set that is contained in the ℓ -cycle. \square

When proving Lemma 4.3, we will see that the divisibility condition in the previous proposition might lead to a few leftover vertices in the blow-up of each edge e of F . We thus require that the connecting path in F^* has some ‘‘absorption’’ features to incorporate into itself these leftover vertices. We start by building a tight walk in F that has all the required features, in a manner such that intermediate edges can be deleted to form an ℓ -walk with suitable start and end vertices. We argue how this can be blown up to a tight path in F^* by selecting distinct vertices along the walk.

As preparation, we prove a handy result which is based on the proof of [15, Lemma 3.3]. Given any two ordered supported sets, we wish to find a tight walk from one ordered edge to the other.

Proposition 8.3. *Suppose G is an n -vertex k -graph with $\delta^*(G) > n/2$ and no isolated vertices, where $n \geq 4k$. Let $S = u_1 \dots u_s$ and $T = v_1 \dots v_t$ be ordered supported sets, where with $s, t \in [k]$. Then there exists a tight walk in G joining S and T , and has at least k vertices between them.*

Proof. Consider any edge $f = x_1 \dots x_k$ that is vertex disjoint from $S \cup T$ (which exists since $n \geq 4k$, G has no isolated vertices, and $\delta^*(G) > n/2$). We will construct a tight walk joining $u_1 \dots u_s$ and $x_1 \dots x_k$, and another one that joins $x_1 \dots x_k$ and $v_1 \dots v_t$.

Since $\delta^*(G) > n/2 \geq 2k$, there exists an edge $e_S = u_1 \dots u_s u_{s+1} \dots u_k$ that contains S and is vertex disjoint from $\{x_1, \dots, x_k\} \cup \{v_1, \dots, v_t\}$. We sequentially find vertices z_1, \dots, z_k such that

$$z_i \in N^1(u_{i+1} \dots u_k z_1 \dots z_{i-1}) \cap N^1(x_{i+1} \dots x_k z_1 \dots z_{i-1}).$$

Define the following tight walks.

$$W_i = z_1 \dots z_{i-1} x_i \dots x_k$$

for all $i \in [k]$, and

$$Q = u_1 \dots u_k z_1 \dots z_k.$$

By choice of z_1, \dots, z_k , it is not hard to see that these sequences are indeed tight walks (in fact, W_1, \dots, W_k are edges). Additionally, we crucially note that $W_{i+1}W_i$ is a tight walk for all $i \in [k-1]$ —intuitively, we “cycle through” the edge $W_{i+1} = z_1 \dots z_i x_i \dots x_k$ one vertex at a time for $i-1$ steps, then replace z_i with x_i , and cycle through $x_{i+1} \dots x_k z_1 \dots z_{i-1} x_i$ for $k-i$ steps. Then $R = QW_k \dots W_1$ forms a tight walk that joins $S = u_1 \dots u_s$ and $x_1 \dots x_k$. The same arguments can be used to construct a tight walk R' joining $x_1 \dots x_k$ and $v_1 \dots v_t$. Then RR' is the required tight walk. \square

Next, we state an easy observation which will serve as the building block for the absorption.

Observation 8.4. *Let $k \geq 3$ and suppose G is an n -vertex k -graph with $\delta^*(G) > n/2$, and $v \in V(G)$ is an arbitrary vertex. Then there exist ordered supported sets W_1 and W_2 in G such that W_1W_2 and W_1vW_2 are tight walks and $|W_i| = k-1$ for $i \in [2]$.*

Proof. Consider any edge $e = u_1 \dots u_{k-1}v$ containing v . Since $\delta^*(G) > n/2$, we may sequentially pick vertices w_1, \dots, w_{k-1} such that

$$w_i \in N^1(u_{i+1} \dots u_{k-1}vw_1 \dots w_{i-1}) \cap N^1(u_i \dots u_{k-1}w_1 \dots w_{i-1}).$$

Then $W_1 = u_1 \dots u_{k-1}$ and $W_2 = w_1 \dots w_{k-1}$ are the required ordered supported sets. \square

As mentioned previously, proving Lemma 4.3 will require a connecting path in F^* that can absorb a few uncovered vertices, and the number of such vertices will turn out to be a multiple of $k-\ell$. Hence, our next step is to build larger tight walks that can integrate into themselves multisets of $k-\ell$ vertices. We would also like to be able to delete edges from them to convert them to ℓ walks when necessary (via Observation 8.1), and thus will also have a restriction on the order of such a walk.

Suppose we have some fixed $k \geq 3$ and $\ell \in [k-1]$. Given a sequence $U = v_1 \dots v_s$ of (not necessarily distinct) vertices, we define a *tight U -absorber* to be a sequence of tight walks (W_0, \dots, W_s) such that $|W_i| \geq k-1$ for every i and $W_0 \dots W_s$ and $W_0v_1W_1 \dots W_{s-1}v_sW_s$ are both tight walks in G .

We define a U -absorber to be a sequence of finite vertex sequences (W_0, \dots, W_s) with $|W_i| \geq k-1$ for every i , such that the vertex sequences formed by $W_0 \dots W_s$ and $W_0 v_1 W_1 \dots W_{s-1} v_s W_s$ are both ℓ -walks. Note that for tight walks, this is exactly the same as the notion of a tight U -absorber. We will often simply refer to the walk $W_0 \dots W_s$ as a (tight) U -absorber.

Given this, we argue that we can find absorbers of any desired length.

Proposition 8.5. *Suppose G is an n -vertex k -graph with $\delta^*(G) > n/2$ and no isolated vertices, where $n \geq 4k$, and let $\ell \in [k-1]$. Let $U = v_1 \dots v_s$ be any finite vertex sequence. Then, for any $q \in \{0, \dots, k-\ell-1\}$, G contains a tight U -absorber $W = W_0 \dots W_s$ such that $|W| \equiv q \pmod{k-\ell}$.*

Proof. Using Observation 8.4 for each $i \in [s]$, we can find tight walks A_i and B_i such that $A_i B_i$ and $A_i v_i B_i$ are tight walks, and $|A_i|, |B_i| = k-1$. Furthermore, since $\delta^*(G) > n/2$ and $n \geq 4k$, by Proposition 8.3, we can find a tight walk F_i joining B_i and A_{i+1} for all $i \in [s-1]$. Hence, $R = A_1 B_1 F_1 A_2 B_2 \dots F_{s-1} A_s B_s$ and $R' = A_1 v_1 B_1 F_1 A_2 v_2 B_2 \dots F_{s-1} A_s v_s B_s$ are both tight walks.

Next, suppose that $|R| \equiv r \pmod{k-\ell}$. Set $p = q - r \pmod{k-\ell}$ and suppose that $A_1 = x_1 \dots x_{k-1}$ (viewed as an ordered supported set). Then, using minimum supported co-degree condition, we successively pick vertices y_1, \dots, y_p where $y_i \in N^1(y_{i-1} \dots y_1 x_1 \dots x_{k-i})$, which ensures that $y_p \dots y_1 A_1$ is a tight path. Set $W_0 = y_p \dots y_1 A_1$, $W_i = B_i F_i A_{i+1}$ for all $i \in [s-1]$ and $W_s = B_s$. Since $|A_i|, |B_i| = k-1$ for every $i \in [s]$, we have $|W_i| \geq k-1$ for every $i \in [s]$. Then the $W_0 \dots W_s$ constitutes the desired tight U -absorber. \square

Next, we use the absorbers we just found to modify the above walk obtained in Proposition 8.3 to control its order, which will allow us to convert it into a suitable ℓ -walk by deleting edges.

Proposition 8.6. *Suppose G is an n -vertex k -graph with $\delta^*(G) > n/2$ and no isolated vertices, where $n \geq 4k$, and let $\ell \in [k-1]$. Suppose S and T are ordered supported sets of size s and t respectively, with $s, t \in [k]$. Then, for any $q \in \{0, \dots, k-\ell-1\}$, there exists a tight walk in G joining S and T whose order is at least $2k$ and is congruent to $q \pmod{k-\ell}$.*

Proof. Consider any sequence of $k-\ell$ vertices $U = v_1 \dots v_{k-\ell}$. Applying Proposition 8.5, we can find a tight U -absorber $W_0 \dots W_{k-\ell}$ whose order is congruent to $k \pmod{k-\ell}$. Crucially, by construction, the same is true of the tight walk $W_0 v_1 W_1 \dots W_{k-\ell-1} v_{k-\ell} W_{k-\ell}$. Let S' be the ordered supported set corresponding to the first $k-1$ vertices of W_0 and T' be the ordered supported set corresponding to the last $k-1$ vertices of $W_{k-\ell}$. By Proposition 8.3, we have tight walks R_S and R_T of order at least k that join S to S' and T' to T respectively. Consequently, the tight walk $R_S W_0 \dots W_{k-\ell} R_T$ joins S and T . Suppose that the order of this walk is congruent to $r \pmod{k-\ell}$, and set $p = q - r \pmod{k-\ell}$. Then the walk $R W_0 v_1 W_1 \dots W_{p-1} v_p W_p W_{p+1} W_{p+2} \dots W_{k-\ell} Q_{k-\ell} R'$ meets the conditions we need. \square

We now put it all together. As earlier, we let $E^*(G)$ denote the set of ordered edges of G . We wish to find a tight walk that contains every ordered edge, $e \in E^*(G)$, has several (tight) absorbers for each $(k - \ell)$ -sequence, and whose order is congruent to $k \pmod{k - \ell}$.

Lemma 8.7. *Let G be an n -vertex k -graph with $\delta^*(G) > n/2$ and no isolated vertices, where $n \geq 4k$. Let $\ell \in [k - 1]$ and let d be any positive integer. Suppose $S = u_1 \dots u_s$ and $T = v_1 \dots v_t$ are ordered supported sets of order s and t respectively, where $s, t \in [k]$. Then there exists a tight walk W that satisfies the following conditions.*

- (T1) W joins S and T ,
- (T2) $|W| \equiv k \pmod{k - \ell}$,
- (T3) for every ordered edge $e \in E^*(G)$, the walk W contains a subwalk W_e corresponding to e , such that these subwalks are pairwise disjoint, and the first vertex of each such subwalk is at a position congruent to $1 \pmod{k - \ell}$ in W ,
- (T4) for each $(k - \ell)$ -sequence U , the walk W contains at least d subwalks that are tight U -absorbers, all such subwalks are pairwise disjoint from each other, and the first vertex of each of these subwalk is at a position congruent to $1 \pmod{k - \ell}$ in W , and
- (T5) the subwalks in (T3) and (T4) are pairwise disjoint from each other.

Proof. Let $\mathcal{P} = \{P_1, \dots, P_x\}$ be a collection of tight walks satisfying the following properties.

- $|P_i| \equiv k \pmod{k - \ell}$ for all $i \in [x]$,
- for each $e \in E^*(G)$, there exist some distinct $i_e \in [x]$ such that $P_{i_e} = e$,
- for each $(k - \ell)$ -sequence U , there exists some $I_U \subseteq [x]$ such that $|I_U| \geq d$ and every P_i with $i \in I_U$ is a tight U -absorber, and
- for any two $(k - \ell)$ -sequences $U \neq U'$, the sets I_U and $I_{U'}$ are disjoint, and $i_e \notin I_U$ for any ordered edge e and $(k - \ell)$ -sequence U .

Observe that we can always find such a collection \mathcal{P} using Proposition 8.5. Then, by Proposition 8.6, we may find a set of tight walks $\mathcal{Q} = \{Q_0, \dots, Q_x\}$ such that Q_0 joins S and the first $k - 1$ vertices of P_1 , Q_i joins the last $k - 1$ vertices of P_i and the first $k - 1$ vertices of P_{i+1} for all $i \in [x - 1]$, Q_x joins the last $k - 1$ vertices of P_x and T and $|Q_i| \geq 2k$ for all integers $i \in [0, x]$. Additionally, we also impose that $|Q_0| \equiv k - 1 \pmod{k - \ell}$, $|Q_i| \equiv k - 2 \pmod{k - \ell}$ for every $i \in [x - 1]$, and $|Q_x| \equiv 2k - 1 \pmod{k - \ell}$. Let Q'_0 be obtained by removing the last $k - 1$ vertices from Q_0 , so $|Q'_0| \equiv 0 \pmod{k - \ell}$; let Q'_i be obtained by removing the first and last $k - 1$ vertices in Q_i for $i \in [x - 1]$, so $|Q'_i| \equiv -k \pmod{k - \ell}$; and let Q'_x be obtained by removing the first $k - 1$ vertices

from Q_x , so $|Q'_x| \equiv k \pmod{k-\ell}$. Let W be the tight walk $Q'_0 P_1 Q'_1 \dots P_x Q'_x$. Clearly W joins S and T by construction. It is also straightforward to see that

$$|W| \equiv \sum_{j \in [0, x]} |Q'_j| + \sum_{j \in [x]} |P_j| \equiv k \pmod{k-\ell}.$$

Furthermore, for every $i \in [x]$, we see that the first vertex of P_i is at position

$$|Q'_0| + \sum_{j=1}^{i-1} |P_j| + \sum_{j=1}^{i-1} |Q'_j| + 1 \equiv 1 \pmod{k-\ell}.$$

in W . Thus, W is the desired walk. \square

Next, by applying Observation 8.1, we immediately obtain the corresponding statement for ℓ walks.

Corollary 8.8. *Let G be an n -vertex k -graph with $\delta^*(G) > n/2$ and no isolated vertices, where $n \geq 4k$. Let $\ell \in [k-1]$ and let d be any positive integer. Suppose $S = u_1 \dots u_s$ and $T = v_1 \dots v_t$ are ordered supported sets of order s and t respectively, where $s, t \in [k]$. Then there exists an ℓ -walk W that satisfies the following conditions.*

- (W1) W joins S and T ,
- (W2) for every ordered edge $e \in E^*(G)$, the walk W contains a subwalk corresponding to e and all such subwalks are pairwise disjoint,
- (W3) for each $(k-\ell)$ -sequence U , the walk W contains at least d subwalks that are U -absorbers and all such subwalks are pairwise disjoint, and
- (W4) the subwalks from claims (W2) and (W3) are pairwise disjoint from each other

Finally, we note that in a regular blow-up of F with $\delta^*(F) > 5s/9$ (say), one can find an ℓ -path (rather than walk) with similar properties to the above.

Lemma 8.9. *Let $1/m \ll 1/s \ll \gamma \ll 1/k \leq 1/3$. Suppose F is an s -vertex k -graph with $\delta^*(F) > 5s/9$ and let F^* be a (γ, m) -regular blow-up of F . Let $e_1, e_2 \in E(F^*)$ be any two ordered edges. Then there exists an ℓ -path Q in F^* from e_1 to e_2 that satisfies the following conditions.*

- (Q1) Q joins e_1 and e_2 ,
- (Q2) $|Q| \leq \gamma m$,
- (Q3) for every ordered edge $e \in E^*(F)$, the path Q contains a subpath corresponding to blow-up of e , and these blow-ups are all pairwise disjoint,

(Q4) let $\mathcal{U} = \{U_1, \dots, U_r\}$ be any collection of pairwise disjoint $(k - \ell)$ -subsets of $V(F^*)$ such that $r \leq k^2 s^k$. Then the path Q contains at least one subpath that is a U_i -absorber for each $i \in [r]$, and all these subpaths are pairwise disjoint, and

(Q5) the subpaths from claims (Q3) and (Q4) are pairwise disjoint from each other.

Proof. Since $\delta^*(F) > 5s/9 > s/2$, we can use Corollary 8.8 to find an ℓ -walk W in F that satisfies (W1)–(W4) with $s, k^2 s^k, \phi(e_1)$ and $\phi(e_2)$ in place of n, d, S and T respectively.

Let $|W| = x$, where $x = x(s, k, \ell)$. We now blow this up to an ℓ -path Q in F^* . Based on our parameter choices, we may suppose that $x \leq (1 - \gamma)m$. Setting $W = w_1 \dots w_x$, for each $i \in [x]$, we choose distinct vertices $v_i \in V(F^*)$ such that $\phi(v_i) = w_i$, which is possible since $x \leq (1 - \gamma)m \leq |\phi^{-1}(w)|$ for every $w \in V(F)$. Then, by choice of W , it is clear that $Q = v_1 \dots v_x$ is an ℓ -path in F^* which satisfies (Q1)–(Q3). To see (Q4), we note that W has at least $k^2 s^k$ subwalks that are $\phi(U_i)$ -absorbers for each $i \in [r]$. Since $r \leq k^2 s^k$ the ℓ -path Q will contain a distinct absorber for each U_i . Then property (W4) of W immediately implies (Q5) for Q . \square

9 Proof of Lemma 4.3

Our goal in this section is to prove Lemma 4.3, which asserts that nearly-regular blow-ups of non-extremal graphs with a certain divisibility condition have a Hamilton ℓ -path (with specified start and end). To avoid having to worry about the singleton part throughout the proof, we now state a version of Lemma 4.3 for regular blow-ups (i.e. where there is no singleton part) and show that it implies Lemma 4.3. The rest of the section will then be dedicated to the proof of this new version. As always, given k and ℓ , we set $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$ and recall that $t \geq 3$ for all $k \geq 3$ and $\ell \in [k - 1]$ except $(k, \ell) \neq (3, 1)$, as this will be fairly important to our proof.

Lemma 9.1. *Let $1/m \ll 1/s \ll \gamma \ll \varepsilon \ll \mu \ll 1/k \leq 1/3$, let $\ell \in [k - 1]$, and suppose $(k, \ell) \neq (3, 1)$. Suppose $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(s, k, \ell)$ and F^* is a (γ, m) -regular blow-up of F satisfying $|F^*| \equiv k \pmod{k - \ell}$. Let $e_1, e_2 \in E(F^*)$ be two disjoint ordered edges. Then there exist a spanning ℓ -path in F^* from e_1 to e_2 .*

Proof of Lemma 4.3 using Lemma 9.1. If F^* does not contain a singleton part, then we are done, so we may assume that it does. Let $u^* \in V(F), V(F^*)$ be the vertex corresponding to the singleton part. Since $F \in \mathcal{E}_{\varepsilon, \mu}^{\text{non-ext}}(n, k, \ell)$, we know $\delta^*(F) \geq (1 - 1/t - \varepsilon)s$, where $t \geq 3$ as $(k, \ell) \neq (3, 1)$. Let $H = F - u^*$. It is easy to check that $H \in \mathcal{E}_{2\varepsilon, \mu/2}^{\text{non-ext}}(n, k, \ell)$. In particular, $\delta^*(H) \geq \delta^*(F) - 1 > s/2$.

We claim that there is a tight walk W in F such that W starts in $\phi(e_1)$, ends in some edge f in H , contains exactly one instance of u^* , and has order $k \pmod{k - \ell}$. Indeed, let $e^* = \{u^*, u_1, \dots, u_{k-1}\}$ be any edge in F that contains u^* , and let $f = \{u_1, \dots, u_k\}$ be any edge in F that contains

$\{u_1, \dots, u_{k-1}\}$ and avoids u^* (which exists by the assumption on $\delta^*(F)$). By Proposition 8.6, we can find a tight walk W' in H from $\phi(e_1)$ to $u_1 \dots u_{k-1}$ such that $|W'| \equiv -1 \pmod{k-\ell}$. Then $W = W'u^*u_1 \dots u_k$ satisfies the requirements. Write $W = w_1 \dots w_x$ where $|W| = x$ can be bounded as a function of s, k and ℓ , and we may thus assume $x \leq \gamma m$.

Our next step is to convert this into a tight path in F^* that avoids e_2 . For each $i \in [x]$, we choose a vertex $v_i \in V(F^*) \setminus V(e_2)$ such that $\phi(v_i) = w_i$, all the v_i are distinct, and the first k vertices correspond to e_1 . This is possible since $x \leq \gamma m - k \leq |\phi^{-1}(w) - V(e_2)|$ for every $w \in V(F) \setminus \{u^*\}$, and recalling that u^* appears precisely once in W . Then $P = v_1 \dots v_x$ is a tight path in $F^* - V(e_2)$ from e_1 to f^* (where $\phi(f^*) = f$), it contains u^* , and $|P| \equiv k \pmod{k-\ell}$. By Observation 8.1, the sequence $v_1 \dots v_x$ can be thought of as an ℓ -path from e_1 to f^* .

Set $H^* = F^* - \{v_1, \dots, v_{x-k}\}$. Notice that if H^* contains a spanning ℓ -path from f^* to e_2 , then F^* contains a spanning ℓ -path from e_1 to e_2 (obtained by adding $v_1 \dots v_{x-k}$ at the beginning). By Lemma 9.1, it will suffice to show that $|H^*| \equiv k \pmod{k-\ell}$ and that H^* is a $(2\gamma, m)$ -regular blow-up of H , since $f^*, e_2 \in E(H^*)$ are disjoint ordered edges by choice. The first condition is immediate as $|H^*| = |F^*| - (x - k) \equiv |F^*| \pmod{k-\ell} \equiv k \pmod{k-\ell}$. Next, notice that H^* is indeed a blow-up of H , as $u^* \notin V(H^*)$. Moreover, each $w \in V(H)$ is blown up to at least $(1 - \gamma)m$ vertices in F^* , meaning that there are at least $(1 - \gamma)m - x \geq (1 - 2\gamma)m$ copies of w left in H^* . \square

Before turning to the proof of Lemma 9.1, we recall some notation. As defined in Section 4, we let $\phi : V(F^*) \rightarrow V(F)$ denote the projection map, that is, $\phi(v') = v$ for all v' in the independent set corresponding to v . As in Section 7, we use $E^*(H)$ to denote the set of ordered edges of a given k -graph H , and, for every $v \in V(H)$ and $i \in [k]$, we write $E_i^v(H) \subseteq E^*(H)$ for the set of ordered edges with v as their i th vertex. Finally, for any ordered edge $e = v_1 \dots v_k \in E^*(F)$, we use $e^{+t} \in E^*(F)$ to be the ordered edge $v_2 \dots v_{t-1}v_1v_t \dots v_k$.

To begin with, we will integrate blow-ups of edges of F into the ℓ -path Q in F^* obtained by Lemma 8.9. The idea is as follows. Every ordered edge $e \in E^*(F)$ is blown up to a k -partite k -graph in F^* . We will consider a subgraph of this k -graph (obtained by paring down the sizes of the parts based on the \mathbf{w}^* -weighted perfect fractional matching) that will be almost perfectly balanced except for one part, and will find suitable a spanning ℓ -path in this subgraph using Proposition 8.2. The idea will then be to replace the edge of Q that corresponds to e^{+t} with this spanning ℓ -path. Overall, we will partition almost all of $V(F^*)$ into such k -partite k -graphs, and replacing the relevant edges of Q with the corresponding spanning ℓ -paths will provide an ℓ -path that covers all but at most a few vertices of $V(F^*)$. We will then use the absorption features of Q to absorb these into the path as well.

Proof of Lemma 9.1. Let e_1 and e_2 be as in the statement. As $F \in \mathcal{E}_{\varepsilon, mu}^{\text{non-ext}}(s, k, \ell)$ and $t \geq 3$, we can apply Lemma 8.9 to find an ℓ -path Q in F^* from e_1 to e_2 satisfying (Q1)–(Q5) with e_1 and e_2 . By (Q2), we know that $|V(Q)| \leq \gamma m$. Write $F^{**} := F^* \setminus V(Q)$; so F^{**} is a $(2\gamma, m)$ -regular blow-up

of F . By Proposition 5.2 and Corollary 7.3, there is a \mathbf{w}^* -weighted perfect fractional matching in F^{**} , where we recall that $\mathbf{w}^* = (w_1, \dots, w_k)$ with $w_1 = k - 1$ and $w_i = t - 1$ for all $2 \leq i \leq k$. Denote this \mathbf{w}^* -weighted perfect fractional matching q .

For each vertex $v \in V(F)$, we let $B_v \subseteq V(F^{**})$ denote the vertex subset corresponding to the blow-up of v . For every ordered edge $e \in E^*(F)$, define

$$\hat{q}(e) = \sum_{\substack{e' \in E^*(F^{**}) \\ \phi(e')=e}} q(e').$$

Observe that

$$\begin{aligned} \sum_{i \in [k]} \sum_{e \in E_i^v(F)} w_i \hat{q}(e) &= \sum_{i \in [k]} \sum_{e \in E_i^v(F)} \sum_{\substack{e' \in E^*(F^{**}) \\ \phi(e')=e}} w_i q(e') \\ &= \sum_{v' \in B_v} \sum_{i \in [k]} \sum_{e' \in E_i^{v'}(F^{**})} w_i q(e') \\ &= \sum_{v' \in B_v} 1 = |B_v|. \end{aligned}$$

For every vertex $v \in V(F)$, we arbitrarily partition B_v into sets B_v^e , for all ordered edges e in $E^*(F)$ that contain v , and an additional set B_v^* , so that $|B_v^e| = w_i \lfloor \hat{q}(e) \rfloor$ for every $e \in E_i^v(F)$ and $i \in [k]$, and B_v^* consists of all the remaining vertices. We see that

$$\begin{aligned} |B_v^*| &= |B_v| - \sum_{e \ni v} |B_v^e| \\ &= \sum_{i \in [k]} \sum_{e \in E_i^v(F)} w_i (\hat{q}(e) - \lfloor \hat{q}(e) \rfloor) \\ &\leq \sum_{i \in [k]} \sum_{e \in E_i^v(F)} (k - 1) \\ &\leq k^2 s^{k-1}, \end{aligned}$$

since there are at most s^{k-1} ordered edges in F containing v in the i th position. Finally, for any ordered edge $e = v_1 \dots v_k \in E^*(F)$, define $T_e \subseteq F^{**}$ to be the complete k -partite k -graph with parts $B_{v_1}^e, \dots, B_{v_k}^e$.

We now modify Q into the desired Hamilton ℓ -path in F^* . We will first replace appropriate edges of Q with spanning ℓ -paths in each T_e , which will form an ℓ -path that spans all the vertices except those in $\bigsqcup_{v \in V(F)} B_v^*$. We then absorb these remaining few vertices into Q using property (Q4).

We start with the latter part, detailing the absorption mechanism. Define $B^* = \bigcup_{v \in V(F)} B_v^*$. We will need the following simple claim.

Claim 9.2. $|B^*|$ is divisible by $k - \ell$ and is at most $k^2 s^k$.

Proof. Observe that

$$\bigsqcup_{v \in V(F)} \bigsqcup_{\substack{e \in E^*(F) \\ e \ni v}} B_v^e = \bigsqcup_{e \in E^*(F)} T_e,$$

implying that $|B^*| = |F^{**}| - \sum_{e \in E^*(F)} |T_e|$. Thus, for divisibility, it suffices to show that each $|T_e|$ and $|F^{**}|$ are divisible by $k - \ell$. For T_e , this follows from the following calculation.

$$|T_e| = \sum_{i \in [k]} |B_{v_i}^e| = \sum_{i \in [k]} w_i \lfloor \hat{q}(e) \rfloor = \lfloor \hat{q}(e) \rfloor (k - 1 + (k - 1)(t - 1)) = \lfloor \hat{q}(e) \rfloor t(k - 1). \quad (9.3)$$

Since Q is an ℓ -path, we know $|V(Q)| \equiv k \pmod{k - \ell}$, and by assumption the same is true of $|F^*|$, which means $|F^*| - |V(Q)| \equiv 0 \pmod{k - \ell}$. This establishes that $|B^*| \equiv 0 \pmod{k - \ell}$. For the upper bound, notice that $|B_v^*| \leq k^2 s^{k-1}$ for each $v \in V(F)$. Thus $|B^*| \leq |V(F)| \cdot k^2 s^{k-1} = k^2 s^k$. \square

By the last claim we can construct a partition $B^* = \bigsqcup_{i=1}^r W_i$, where each $|W_i| = k - \ell$, and so that $r \leq k^2 s^k$. Then, by (Q4), we can modify Q to an ℓ -path Q' from e_1 to e_2 such that $V(Q') = V(B^*) \cup V(Q)$ and, using (Q5), we deduce that Q' retains property (Q3). Hence, we have

$$V(F^*) \setminus V(Q') = \bigsqcup_{v \in V(F)} \bigsqcup_{\substack{e \in E^*(F) \\ e \ni v}} B_v^e = \bigsqcup_{e \in E^*(F)} T_e.$$

We move on to the final step, where we replace certain edges of Q' with spanning paths in each T_e , forming a spanning ℓ -path. The following claim is key to this.

Claim 9.4. For every ordered edge $e = v_1 \dots v_k \in E^*(F)$ for which T_e is non-empty, there exists a spanning ℓ -cycle $C_e \subseteq T_e$ which contains an ordered edge that is a blow-up of the ordered edge $e^{+t} \in E^*(F)$.

Proof. As noted in (9.3), we know $|T_e| = \lfloor \hat{q}(e) \rfloor t(k - 1)$, and hence $|T_e|$ is divisible by $t(k - 1)$. Additionally, $|B_{v_1}^e| = (k - 1) \lfloor \hat{q}(e) \rfloor = |T_e|/t$ and $|B_{v_i}^e| = (t - 1) \lfloor \hat{q}(e) \rfloor = |T_e|(t - 1)/t(k - 1)$. Hence, by setting $A = B_{v_1}^e$ and $(B_1, \dots, B_{k-1}) = (B_{v_2}^e, \dots, B_{v_k}^e)$ and applying Proposition 8.2, we can find a spanning ℓ -cycle $C_e \subseteq T_e$ that contains an ordered edge from $B_{v_2}^e \times \dots \times B_{v_t}^e \times B_{v_1}^e \times B_{v_{t+1}}^e \times \dots \times B_{v_k}^e$, that is, a copy of the ordered edge e^{+t} . \square

Now, consider an arbitrary ordered edge $e \in E^*(F)$ for which T_e is non-empty. From Claim 9.4, we know that there exists a spanning ℓ -cycle $C_e \subseteq T_e$ containing an ordered edge, say c , which is a copy of the ordered edge $e^{+t} \in E^*(F)$. Let $P_e = C_e - c$ be the ℓ -path in T_e obtained by removing the edge c (that is, we just think of P_e as consisting of all the edges of C_e except c , ordered in a fashion that makes it an ℓ -path). By (Q3), we know that Q' has a blow-up of every ordered edge of F , and

thus in particular contains an ordered edge d which is a blow-up of e^{+t} . Let f_1 and f_2 be the edges that precede and succeed d in Q' . Since c and d are blow-ups of the same ordered edge of F , and by definition of the ℓ -path P_e , we may replace the subpath f_1df_2 in Q' with the subpath $f_1dP_e cf_2$ (here we view an ℓ -path as a sequence of edges, without explicitly describing the vertex sequence as we did previously). This yields an ℓ -path on vertex set $V(Q') \cup V(T_e)$. Since the k -graphs T_e are pairwise disjoint for each $e \in E^*(F)$ and, by (Q3), Q' contains disjoint subpaths corresponding to the set of edges $\{e^{+t} : e \in E^*(F)\}$, we may repeat this procedure for every $e \in E^*(F)$ (that is, for every T_e), and let Q'' be the resulting ℓ -path. Then Q'' is an ℓ -path from e_1 to e_2 that spans $V(F^*)$, precisely as required. \square

10 Extremal hypergraphs

Now we turn our attention to hypergraphs that obey the exact minimum supported co-degree condition and whose structure is close to those of the extremal examples, complementing Theorem 2.5. More precisely, we will prove Theorem 2.6 (restated below) over the next four sections. In this section, we first provide a detailed proof sketch and then introduce various tools that we will need for the proof.

Theorem 2.6. *Let $k \geq 3$ and $\ell \in [k-1]$ be such that $(k, \ell) \neq (3, 1)$, and set $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Let $1/n \ll \varepsilon \ll 1/k \leq 1/3$ be such that n is divisible by $k-\ell$. Suppose G is an n -vertex k -graph with no isolated vertices and $\delta^*(G) \geq (1 - 1/t)n - (k-3)$. If there is a subset $A \subseteq V(G)$ with $|A| = \lfloor n/t \rfloor$ that contains at most εn^2 supported pairs, then G has a Hamilton ℓ -cycle.*

10.1 Proof sketch for Theorem 2.6

As in Section 2.3.2, in this sketch we focus on tight cycles, namely on the case $\ell = k-1$ (and $t = k$), and comment on the necessary modifications for general ℓ -cycles at the end. Suppose $G \in \mathcal{E}_\varepsilon^{\text{ext}}(n, k, k-1)$. Recall that this implies that G has a vertex subset A of size $\lfloor n/k \rfloor$ with at most εn^2 supported pairs, and set $B = V(G) \setminus A$.

First, since there are so few supported pairs in A , there are very few vertices which have large vertex degree in A . We transfer these to B and argue that, after this modification, every supported set that intersects A has high co-degree into B . The extremal structure also implies that almost every $(k-1)$ -subset of B is A -rich – meaning that it supports almost all of A (Definition 12.3). We then use a “cleaning lemma” (Lemma 11.7) to find a nearly spanning subgraph of A -rich sets in B with nearly complete minimum supported co-degree into B , and transfer the few uncovered vertices into A . Throughout, we maintain $|A| \approx n/k$. These modifications are summarised in Proposition 12.1.

Next, we “balance” the sizes of A and B by replacing a few vertices of A with short tight paths. The exact vertex degree condition (via Observation 11.1) provides us with a few paths of length two

in $\partial^2[A]$, and we extend these to tight paths, including one special path of some specified length to ensure that $|B| = (k-1)|A|$. Additionally, these paths will start and end with supported sets that have high co-degree across the $A \sqcup B$ partition. This balancing is achieved in Proposition 13.1.

The set A now consists mostly of single vertices and a few tight paths. As explained in Section 2.3.2, we arbitrarily partition B into $k-1$ sets B_1, \dots, B_{k-1} such that $|A| = |B_i|$ for all i , and define an auxiliary k -partite k -graph G^+ with parts B_1, \dots, B_{k-1}, A whose edges correspond to vertex sequences in $B_1 \times \dots \times B_{k-1} \times A$ that form tight paths. We use a random greedy procedure, based on a simple version for almost complete balanced bipartite graphs (see Lemmas 11.13 and 11.14), to build a perfect matching M in G^+ by starting with vertices in B_1 , extending them to ordered supported sets in $B_1 \times B_2$, and so on. We then form an auxiliary digraph whose vertices correspond to the vertex sequences of M , and for any $P \in M$, we let its in-neighbours (respectively, out-neighbours) be all the $P' \in M$ such that $P'P$ (respectively, PP') forms a tight path. We argue that the randomness preserves the high supported co-degree conditions across the partition $A \sqcup B$ so that this digraph has high minimum semi-degree. This will ensure a directed Hamilton cycle, which corresponds to a tight Hamilton cycle in G . This final step of the proof is done in Section 14.

The major difference for ℓ -cycles is that in our balancing step, instead of replacing vertices with tight paths, we use vertex sequences of specified order that “support ℓ -paths” (see Definition 11.5) – intuitively, these consist of ℓ -paths for “as long as possible” and have some small supported sets on either end. The upshot is that concatenating two such sequences forms another sequences that supports an ℓ -path. We do this to make $|B| = (t-1)|A|$. In the next step, we define G^+ to be an auxiliary t -partite t -graph on $B_1 \times \dots \times B_{t-1} \times A$ whose edges are vertex sequences that support ℓ -paths. The rest of the proof is largely similar.

Structure of the next four sections. In the next section, we prove fundamental results for our proof strategy, starting with basics about the minimum supported co-degree and some simple graph theoretic statements in Section 11.1. Then Section 11.2 describes key properties about vertex sequences that support ℓ -paths. Section 11.3 provides a general mechanism for finding a nearly spanning subgraph that approximately preserves supported co-degree conditions in a dense host graph. We then use this to analyse our random matching procedure in Section 11.4.

Once we have set up the necessary machinery, as described in our proof sketch, we modify G to ensure desired co-degree properties in Section 12, balance the part sizes in Section 13, and subsequently construct a Hamilton cycle in Section 14.

11 Key lemmas

We start with a quick observation based on the minimum supported co-degree condition.

Observation 11.1. Let $k \geq 3$, $\ell \in [k-1]$, and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$ and let G be a k -graph with $\delta^*(G) \geq n - \lfloor n/t \rfloor - (k-3)$. For all $i \in [k-1]$ and for any supported i -set $S \in V(G)$, we have

$$d_G^1(S) \geq n - \left\lfloor \frac{n}{t} \right\rfloor - (i-2) \geq \left(1 - \frac{1}{t}\right)n - k.$$

Proof. Let $S = \{v_1, \dots, v_i\}$. Since S is supported, there must exist a set $W = \{w_{i+1}, \dots, w_k\}$ of $(k-i)$ vertices such that $S \cup W \in E(G)$. In particular, $S' = \{v_1, \dots, v_i, w_{i+1}, \dots, w_{k-1}\}$ is a supported $(k-1)$ -set. It is clear that $W, N_G^1(S') \subseteq N_G^1(S)$. Thus, we have

$$\begin{aligned} d_G^1(S) &\geq d_G^1(S') + |W| - 1 \\ &\geq \delta^*(G) + k - i - 1 \\ &\geq n - \left\lfloor \frac{n}{t} \right\rfloor - (i-2) \geq \left(1 - \frac{1}{t}\right)n - k, \end{aligned}$$

where we subtract one in the first inequality to prevent counting $w_k \in N_G^1(S') \cap W$ twice. \square

11.1 Graph theoretic lemmas

We now prove a simple graph theoretic result about short disjoint paths. A *cherry* in a graph will refer to a path of length two.

Lemma 11.2. Suppose H is a graph on n vertices with minimum degree $\delta(H) \geq x \geq 1$, and suppose that A is a subset of $V(H)$ satisfying $|A| > n/2 + 9x/2$. Then H contains a linear forest F that consists of $x-1$ cherries with leaves in A .

Proof. Let F' be a disjoint union of cherries with leaves in A , such that $|F'|$ is maximised. Let $A' := A \setminus V(F')$ and let E' be the set of edges with exactly one end in A' . Then, $H[A']$ has maximum degree at most 1 (otherwise we could add a cherry with vertices in A' to F' , contradicting maximality), showing that $|E'| \geq |A'|(x-1)$. We now wish to upper bound $|E'|$, using the following points.

- For each cherry uvw in F' , at most one of u , v and w has at least four neighbours in A' . Indeed, otherwise we can find two disjoint cherries with centres in $\{u, v, w\}$ and leaves in A , and replace uvw by these cherries, contradicting the maximality of F' . It follows that the number of edges from A' to $\{u, v, w\}$ is at most $|A'| + 6$.
- There is at most one edge from A' to each vertex in $V(H) \setminus (A \cup V(F'))$, as otherwise we could add an additional cherry to F .

Altogether, denoting by y the number of cherries in F' (so that $|F'| = 3y$) we find that

$$|E'| \leq y(|A'| + 6) + (n - |A|). \quad (11.3)$$

It remains to show that $y \geq x - 1$. Otherwise, by combining (11.3) with $|E'| \geq (x - 1)|A'|$, we get $|A'| + |A| \leq 6x + n$. This, along with the observation that $|A| - |A'| \leq |F'| \leq 3x$, implies that $|A| \leq n/2 + 9x/2$, a contradiction. \square

Our next preliminary result is a directed version of Dirac's theorem from [9]. Given a digraph D with minimum out-degree $\delta^+(D)$ and minimum in-degree $\delta^-(D)$, we define the *minimum semi-degree* of D as $\delta^0(D) = \min\{\delta^+(D), \delta^-(D)\}$.

Lemma 11.4 (Ghouila-Houri, 1960 [9]). *If D is an n -vertex digraph with $\delta^0(D) \geq n/2$, then it has a directed Hamilton cycle.*

11.2 Supported ℓ -paths

In this subsection we define the notion of vertex sequences that “support an ℓ -path”. Intuitively, we have an ℓ -path for as long as possible, and require the final few vertices to form a supported set. These are set up so that it is easy to join them up to longer sequences that support an ℓ -path, and ultimately to a spanning ℓ -cycle. These will play an important role in our main strategy of splitting our graph into a perfect matching of sequences that support ℓ -paths that satisfy useful degree conditions, allowing them to combine into a spanning ℓ -cycle. We also define a sequence that supports an extended ℓ -path, which we will use to construct the required perfect matching of sequences supporting ℓ -paths.

Definition 11.5. Let $k \geq 3$, $\ell \in [k - 1]$, $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$. Let $P := w_1 \dots w_r$ be a sequence of distinct vertices in a k -graph G . We say that P *supports an ℓ -path* in G if $r \geq t$, r is divisible by $k - \ell$, $w_{i(k-\ell)+1} \dots w_{i(k-\ell)+k}$ is an edge for all integers $i \in [0, \frac{r-t}{k-\ell} - 1]$, and $\{w_{r-t+1}, \dots, w_r\}$ is a supported set.

Let $P = w_0 \dots w_r$ be a sequence of distinct vertices in G . We say that P *supports an extended ℓ -path* if $r \geq t$, r is divisible by $k - \ell$, the sequence $w_1 \dots w_r$ supports an ℓ -path, and $w_0 \dots w_{t-1}$ is a supported set.

We summarise the facts about constructing and connecting sequences supporting ℓ -paths and extended ℓ -paths in the following proposition.

Proposition 11.6. *Let G be a k -graph, where $k \geq 3$, $\ell \in [k - 1]$ and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k - \ell)$.*

- (i) *Let P and Q be two sequences of vertices in G of length at least t , both of which support an ℓ -path, and whose vertices are disjoint. If the sequence consisting of the last t vertices in P followed by the first $t - 1$ vertices in Q is a tight path, then PQ supports an ℓ -path.*

- (ii) Let P_1, \dots, P_m be sequences of vertices in G of length at least t , each of which supports an ℓ -path, and whose vertices are pairwise disjoint. If $P_i P_{i+1}$ supports an ℓ -path for every $i \in [m]$ (addition of indices taken modulo m) then $P_1 \dots P_m$ is an ℓ -cycle.
- (iii) Let $P = v_0 \dots v_r$ be a sequence of distinct vertices where r is divisible by t . If every subsequence of at most k consecutive vertices that contains at most one vertex v_i with i divisible by t is supported by an edge, then P supports an extended ℓ -path.
- (iv) Let P be a sequence of distinct vertices of length r such that $r \equiv -1 \pmod{k-\ell}$ and $r \geq t-1$, and let u, v be two vertices that do not appear in P , such that uP and Pv are tight paths (or supported sets if $r+1 \leq k$). Then uPv supports an extended ℓ -path.

Proof. For (i), say that the length of P is p and that of Q is q . Let $PQ = v_1 \dots v_{p+q}$, where $P = v_1 \dots v_p$ and $Q = v_{p+1} \dots v_{p+q}$. Since P and Q both support ℓ -paths, we know that p and q are divisible by $k-\ell$, implying the same is true of $p+q$. Furthermore, as $q \geq t$, we also know that the last t vertices of PQ correspond to the last t vertices of Q , and hence are supported. All that is left to show is that $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$ is an edge for all $i \leq \frac{p+q-t}{k-\ell} - 1$. Since P and Q support ℓ -paths (and in particular their length is divisible by $k-\ell$), we only need to restrict our attention to $i \in [\frac{p-t}{k-\ell}, \frac{p}{k-\ell} - 1]$. For this, it would suffice to argue that the subsequence $v_{i_1} \dots v_{i_2}$ forms a tight path, where $i_1 = p-t+1$ and $i_2 = p+\ell$. By assumption, we know that $v_{p-t+1} \dots v_{p+t-1}$ is a tight path, and we know $t-1 \geq \ell$ (see Observation 3.1), which completes the proof.

Next, for (ii), let $v_1 \dots v_p$ denote the sequence $P = P_1 \dots P_m$. Since each P_i supports an ℓ -path, we know that p is divisible by $k-\ell$. It suffices to show that $e_j := v_{j(k-\ell)+1} \dots v_{j(k-\ell)+k}$ is an edge for all $j \geq 0$, where we view the indices cyclically modulo p . Crucially, since $|P_i|$ is divisible by $k-\ell$ and at least t for every i , for each e_j , if e_j starts in P_i then the suffix of the sequence $P_i P_{i+1}$ starting at e_j has length at least $k-\ell+t \geq k$, showing that e_j is a subsequence of $P_i P_{i+1}$. Furthermore, as the length of $P_1 \dots P_{i-1}$ is divisible by $k-\ell$, and since $P_i P_{i+1}$ supports an ℓ -path by item (i), we conclude that e_j is indeed an edge.

We now move to the (iii). By Proposition 3.2, we know that every subsequence of the form $v_{i(k-\ell)+1} \dots v_{i(k-\ell)+k}$ contains at most one vertex v_i with i divisible by t and is thus forms an edge, by assumption. The sets $v_0 \dots v_{t-1}$ and $v_{r-t+1} \dots v_r$ each have length t and thus have only one v_i with i divisible by t , showing that they are supported in G . Thus, P supports an extended ℓ -path.

The final item (iv) is also straightforward. Since Pv is a tight path of length divisible by $k-\ell$ and at least t , it clearly supports an ℓ -path. Furthermore, as $r \geq t-1$, the tight path uP contains at least t vertices, and the first t of these are supported. \square

11.3 Cleaning lemmas

In this subsection, we will prove important results about finding substructures with large supported co-degrees in very dense hypergraphs. Broadly speaking, we will show that we can delete a few edges and find a nearly spanning subgraph that has no isolated vertices, and approximately preserves degree conditions. Similar ideas have appeared other papers regarding exact co-degree bounds for Hamilton cycles, such as [25, 39], but these were presented in an ad-hoc manner as intermediate steps in their proofs. We provide more general and formalised results about finding large, dense subgraphs with good degree conditions that we believe could prove to be independently useful in other settings as well.

Let us the begin with the first result. Given an almost complete k -graph, we will drop to an almost complete subgraph that spans almost all the vertices, and has high supported co-degree conditions. For $k = 2$ this can be done easily by removing all vertices of small degree, and it turns out that a similar but more subtle strategy works for higher uniformities. We point out that Halfpap, Lemons and Palmer [11, Lemma 18] have a similar result with essentially the same proof. However, our other two results in this section, namely Corollaries 11.10 and 11.11, are proved using the first one, and yield a stronger conclusion than that of [11, Lemma 18]. We expect these to be applicable in other settings as well.

Lemma 11.7. *Let $1/n \ll \varepsilon \ll \delta \ll 1/k \leq 1/2$. Suppose F is an n -vertex k -graph with $e(F) \geq (1-\varepsilon)\binom{n}{k}$. Then there exists a subgraph $F' \subseteq F$ with no isolated vertices such that $\delta^*(F') \geq (1-\delta)n$. In particular, $|F'| \geq (1-\delta)n$ and $e(F') \geq (1-k\delta)\binom{n}{k}$.*

Proof. Let μ satisfy $\varepsilon \ll \mu \ll \delta$. Let us describe how we obtain F' from F . We construct a sequence of hypergraphs F_k, \dots, F_1 , where each F_i is i -uniform. First, we define $F_k = F$. Then, for $i \geq 2$, given F_{i+1} , we let F_i be the subgraph of $\partial^i F_{i+1}$ with $V(F_i) = V(F)$ which contains every i -edge S satisfying $d_{F_{i+1}}^1(S) \geq (1-\mu)n$. Finally, in the graph F_2 , we remove all vertices v such that $d_{F_2}^1(v) \leq (1-\mu)n$, and let F_1 be the remaining vertex set. Now, we let $V(F') = F_1$, and set $E(F') \subseteq \binom{V(F')}{k}$ to consist of all k -sets S which have $\binom{S}{i} \subseteq E(F_i)$ for all $i \in [k]$, that is, every i -subset of S is an edge of F_i .

We show that F' satisfies the statement of Lemma 11.7 in two parts.

Claim 11.8. *We have $|F'| \geq (1 - 2^{k-1}\varepsilon/\mu^{k-1})n \geq (1-\delta)n$.*

Proof. Our proof proceeds by successively bounding the number of edges in each F_i from below. We know that $e(F_k) = e(F) \geq (1-\varepsilon)\binom{n}{k}$. We define $\mu_i = 2^{k-i}\varepsilon/\mu^{k-i}$ for all $i \in [k]$. We claim that, for all such i , we have

$$e(F_i) \geq (1 - \mu_i) \binom{n}{i}.$$

Indeed, suppose we know that this is true for F_k, \dots, F_{i+1} . Write $f := e(F_i)$, so that f is the number of supported i -sets S in F_{i+1} that have $d_{F_{i+1}}^1(S) \geq (1 - \mu)n$. Consequently, by considering a degree sum over all supported i -sets in F_{i+1} and splitting the sum into i -sets that are preserved in F_i and those that are not, we see that

$$\binom{i+1}{i}(1 - \mu_{i+1})\binom{n}{i+1} \leq fn + \left(\binom{n}{i} - f\right)(1 - \mu)n.$$

Using $\binom{i+1}{i}\binom{n}{i+1} = n\binom{n-1}{i}$ and rearranging, we see that

$$\begin{aligned} \mu f &\geq \mu\binom{n}{i} - \mu_{i+1}\binom{n-1}{i} - \left(\binom{n}{i} - \binom{n-1}{i}\right) \\ &= \mu\binom{n}{i} - \mu_{i+1}\binom{n-1}{i} - \binom{n-1}{i-1} \\ &\geq \mu\binom{n}{i} - 2\mu_{i+1}\binom{n}{i}, \end{aligned}$$

where we use $\binom{n}{i} = \binom{n-1}{i} + \binom{n-1}{i-1}$ and $\binom{n-1}{i-1} \leq \mu_{i+1}\binom{n}{i}$ for $n \gg k \geq i$. Dividing throughout by μ completes the proof of the claim. In particular, these calculations show that $|F'| = |F_1| \geq (1 - \mu_1)n \geq (1 - 2^{k-1}\varepsilon/\mu^{k-1})n$ as desired. \square

Claim 11.9. *All supported sets S in F' of size at most $k-1$ satisfy $d_{F'}^1(S) \geq (1 - 2^k\mu)n \geq (1 - \delta)n$.*

Proof. Consider a supported set S of size at most $k-1$. Notice that a vertex $s \in V(F') \setminus S$ is in $N_{F'}^1(S)$ if and only if $s \in N_{F_{|S'|+1}}^1(S')$ for every non-empty subset $S' \subseteq S$. Since every non-empty subset $S' \subseteq S$ satisfies $S' \in E(F_{|S'|})$, by definition of $F_{|S'|}$ it satisfies $d_{F_{|S'|+1}}^1(S') \geq (1 - \mu)n$. This means that every subset $S' \subseteq S$ disallows at most μn vertices, and moreover we discard at most $(2^{k-1}\varepsilon/\mu^{k-1})n$ vertices to obtain F_1 . Hence, the number of infeasible vertices s is at most $(2^{k-1}\varepsilon/\mu^{k-1})n + 2^{k-1}\mu n \leq 2^k\mu n \leq \delta n$, as claimed. \square

The last claim shows that $\delta^*(F') \geq (1 - \delta)n$, and by choice of F' it has no isolated vertices. This completes the proof of the main assertion of the lemma. For the lower bound on $e(F')$, we use Proposition 5.3 and find that every vertex in F' is in at least $(1 - \delta)^{k-1} \frac{n^{k-1}}{(k-1)!} \geq (1 - k\delta)\binom{n}{k-1}$ edges, showing that $e(F') \geq \frac{1}{k}(1 - k\delta)n\binom{n}{k-1} \geq (1 - k\delta)\binom{n}{k}$, as claimed. \square

Next, we prove an analogue of Lemma 11.7 where the host hypergraph is any arbitrary k -graph. Specifically, we show that if the host graph F has linear minimum supported co-degree, and F' is formed by removing few edges from F , then there is a subgraph F'' of F , which is obtained by removing few edges, and where for every supported set S in F'' of size at most $k-1$, the degree of S in F'' is almost as large as it is in F .

Corollary 11.10. *Let $1/n \ll \varepsilon \ll \delta \ll 1/k \leq 1/2$. Let F be an n -vertex k -graph with and let $F' \subseteq F$ be a subgraph satisfying $e(F') \geq e(F) - \varepsilon n^k$. Then there is a subgraph $F'' \subseteq F'$ such that the following properties hold.*

- F'' has no isolated vertices,
- $|F''| \geq (1 - \delta)n$,
- $e(F'') \geq e(F) - \delta n^k$,
- $d_{F''}(S) \geq d_F(S) - \delta n^{k-|S|}$ for every set S of size at most $k - 1$ which is supported in F'' .

Proof. Let ε' and μ satisfy $\varepsilon \ll \varepsilon' \ll \mu \ll \delta$. Let H be the complement of F , that is, the k -graph on $V(F)$ whose edges are the k -sets that are non-edges in F . Write $G := H \cup F$ and $G' := H \cup F'$. Apply Lemma 11.7 to G' with parameters ε' and μ , noting that $e(G') \geq \binom{n}{k} - (e(F) - e(F')) \geq \binom{n}{k} - \varepsilon n^k \geq (1 - \varepsilon')\binom{n}{k}$. Denote the resulting hypergraph by G'' . Let $V(F'') = V(G'')$ and $E(F'') = E(G'') \setminus E(H)$. We claim that F'' satisfies the requirements of the corollary. Indeed, let S be a supported set in F'' of size $s \leq k - 1$. Then, from Proposition 5.3, we get

$$d_{G''}(S) \geq \frac{((1 - \mu)n)^{k-s}}{(k-s)!} \geq \binom{n-s}{k-s} - \delta n^{k-s} = d_H(S) + d_F(S) - \delta n^{k-s}.$$

Since $d_{F''}(S) = d_{G''}(S) - d_H(S)$ we get $d_{F''}(S) \geq d_F(S) - \delta n^{k-s}$, as required for the fourth item. The other items follow directly from the choice of G'' and F'' . \square

Finally, we obtain a similar corollary which shows that if the host graph F has linear minimum supported co-degree, and F' is formed by removing few edges from F , then there is a subgraph F'' of F , which is obtained by removing few edges, and where for every supported set S in F'' of size at most $k - 1$, the vertex co-degree of S in F'' is almost as large as it is in F .

Corollary 11.11. *Let $1/n \ll \varepsilon \ll \delta \ll \alpha \ll 1/k \leq 1/2$. Let F be an n -vertex k -graph with minimum supported co-degree at least αn , and let $F' \subseteq F$ be a subgraph satisfying $e(F') \geq e(F) - \varepsilon n^k$. Then there is a subgraph $F'' \subseteq F'$ such that the following properties hold.*

- F'' has no isolated vertices,
- $|F''| \geq (1 - \delta)n$,
- $e(F'') \geq e(F) - \delta n^k$,
- $d_{F''}^1(S) \geq d_F^1(S) - \delta n$, for every set S of at most $k - 1$ vertices that is supported in F'' .

Proof. Apply Corollary 11.10 to F' with parameters ε and μ , where $\varepsilon \ll \mu \ll \delta$, and denote the resulting graph F'' . The graph F'' clearly satisfies the first three properties above, so it suffices to prove the fourth item. Suppose that for some set of vertices S of size $s \leq k-1$, which is supported in F'' , we have $d_{F''}^1(S) \leq d_F^1(S) - \delta n$. This means that there is a set X of size δn such that, for every $x \in X$, $S \cup \{x\}$ is supported in F but not in F'' . Thus, using Proposition 5.3 alongside the minimum supported co-degree assumption on F , we have

$$\begin{aligned} d_F(S) - d_{F''}(S) &\geq |\{e \in E(F) : S \cup \{x\} \subseteq e \text{ for some } x \in X\}| \geq \frac{|X| \cdot (\alpha n)^{k-s-1}}{(k-s)!} \\ &\geq \delta \cdot \alpha^{k-s-1} \binom{n}{k-s} > \mu n^{k-s}, \end{aligned}$$

a contradiction to the choice of F'' according to Corollary 11.10. \square

11.4 Perfect matchings in almost complete k -partite k -graphs

Our main aim in this subsection is to prove Corollary 11.15 below, which asserts that, if given a balanced k -partite k -graph H with large minimum supported co-degree and no isolated vertices, and a not-too-large family \mathcal{F} of subgraphs of H that are almost complete, there is a perfect matching M such that for every $F \in \mathcal{F}$, almost every edge of M is in F .

We will use the following consequence of Hall's theorem.

Proposition 11.12. *Let $G = X \sqcup Y$ be a bipartite graph with $|X| = |Y| = n$. If $\delta(G) \geq n/2$, then G has a perfect matching.*

Proof. Consider any subset $S \subseteq X$. If we can show that $N_G(S)$ has at least $|S|$ vertices, then we are done by Hall's theorem.

Clearly $|N_G(S)| \geq \delta(G) \geq n/2$ for all S . So if $|S| \leq n/2$, we are done. Hence, we may suppose $|S| > n/2$. Consider any $y \in Y$. Since $d_G(y) \geq n/2 > |X \setminus S|$, the vertex y must have a neighbour in S . Since y is simply an arbitrary vertex of Y , we conclude that $N_G(S) = Y$, and clearly $|Y| = |X| \geq |S|$. \square

The main content of the proof of the aforementioned result about matchings in k -partite k -graph is the following lemma, which resolves the graph case, i.e. when $k = 2$. The matching is found in randomly, thus allows us to deduce that almost all edges of M are in F , for a not-too-large collection of subgraphs \mathcal{F} of H that are almost complete.

Lemma 11.13. *Let $1/n \ll \varepsilon \ll \delta \ll 1$. Let H be a balanced bipartite graph on $2n$ vertices with minimum degree at least $(1 - \varepsilon)n$, and let \mathcal{F} be a collection of subgraphs of H , each with at least $(1 - \varepsilon)n^2$ edges, such that $|\mathcal{F}| = e^{o(\sqrt{n})}$. Then, there is a perfect matching M in H such that $|M \cap E(F)| \geq (1 - \delta)n$ for every $F \in \mathcal{F}$.*

Proof. Fix parameters α and β such that $\varepsilon \ll \alpha \ll \beta \ll \delta$. Denote the two parts of H by X and Y . Let x_1, \dots, x_m be a sequence of $m := (1 - \beta)n$ distinct vertices in X . We sequentially pick vertices $y_1, \dots, y_m \in Y$ such that each y_i is a uniformly random neighbour of x_i in Y , and is distinct from y_1, \dots, y_{i-1} . This is always possible since $d(x) \geq (1 - \varepsilon)n \geq m$ for all $x \in X$.

Let X' and Y' be the remaining unmatched vertices of X and Y , respectively. By choice, we have $|X'| = |Y'| = \beta n$. Notice that the minimum degree in $H[X', Y']$ is at least $(1 - \varepsilon)n - (1 - \beta)n = (\beta - \varepsilon)n \geq \beta n/2 = |X'|/2$. Thus, by Proposition 11.12, there is a perfect matching M' in $H[X', Y']$. Let M be the union of the matching $\{x_1 y_1, \dots, x_m y_m\}$ with M' . Then M is a perfect matching in H .

Fix $F \in \mathcal{F}$. We will show that $|M \cap E(F)| \geq (1 - \delta)n$, with probability at least $1 - e^{-\Omega(\sqrt{n})}$. A union bound will then conclude the proof.

Colour the edges of F blue and the remaining edges in H red. By assumption, there are at most εn^2 red edges. We say that a vertex $x \in X$ is *scared* if its red degree into Y is at least αn , and we will call it *happy* otherwise. Summing the red degrees of these vertices and using the bound on the number of red edges, we see that the number of scared vertices in X is at most $\varepsilon n/\alpha$.

For each $j \leq m$ such that x_j is happy, let p_j be the probability that $x_j y_j$ is a red edge. Since x_j is happy, we have

$$p_j \leq \frac{\alpha n}{(1 - \varepsilon)n - (j - 1)} \leq \frac{\alpha n}{(1 - \varepsilon)n - m} = \frac{\alpha n}{(1 - \varepsilon)n - (1 - \beta)n} \leq \frac{2\alpha}{\beta}.$$

If we let R denote the random variable corresponding to the number of red edges in $\{x_j y_j : j \in [m], x_j \text{ is happy}\}$, then clearly R is stochastically dominated by a binomial random variable \hat{R} with parameters $(n, 2\alpha/\beta)$. Hence, using Chernoff bounds (Lemma 3.3), we have

$$\begin{aligned} \mathbb{P}[R \geq 2\alpha n/\beta + n^{3/4}] &\leq \mathbb{P}[\hat{R} \geq 2\alpha n/\beta + n^{3/4}] \\ &= \mathbb{P}[\hat{R} \geq \mathbb{E}\hat{R} + n^{3/4}] \\ &\leq 2 \exp\left(-\frac{n^{3/2}}{6n\alpha/\beta}\right) = e^{-\Omega(\sqrt{n})}. \end{aligned}$$

Since the number of red edges in M is at most $R + \beta n + \varepsilon n/\alpha$, we have that, with probability at least $1 - e^{-\Omega(\sqrt{n})}$, the number of red edges in M is at most

$$\frac{2\alpha n}{\beta} + n^{3/4} + \beta n + \frac{\varepsilon n}{\alpha} \leq \delta n.$$

Thus, with probability at least $1 - e^{-\Omega(\sqrt{n})}$, $|M \cap E(F)| \geq (1 - \delta)n$. By a union bound, with high probability this holds for all $F \in \mathcal{F}$. \square

The next lemma deals with k -partite k -graphs. In particular, given a collection of subgraphs

satisfying some high supported co-degree conditions, we wish to find a matching that shares almost all its edges with each subgraph in the collection.

Lemma 11.14. *Let $1/n \ll \varepsilon \ll \delta \ll 1/k \leq 1/3$. Let H be a k -partite k -graph with parts X_1, \dots, X_k , each having size $|X_i| = n$. Suppose that H has no isolated vertices and, for all $i \in [k]$, every supported set S in H which is disjoint from X_i satisfies $d_{X_i}^1(S) \geq (1 - \varepsilon)n$. Let \mathcal{F} be a collection of subgraphs satisfying the following conditions:*

- *There exists a constant $c = c(k)$ such that $|\mathcal{F}| \leq n^c$,*
- *each $F \in \mathcal{F}$ has no isolated vertices, and*
- *for any $F \in \mathcal{F}$, every supported set S of F which is disjoint from X_i satisfies $d_{X_i}^1(S) \geq (1 - \varepsilon)n$.*

Then there exists a perfect matching M in G such that $|M \cap E(F)| \geq (1 - \delta)n$ for all $F \in \mathcal{F}$.

We construct each edge one vertex at a time. In the first step, we form a matching of supported pairs between X_1 and X_2 . We then extend this to a matching of supported triples across X_1 , X_2 and X_3 , and so on. In each step, we will extend our existing matching using Lemma 11.13 and maintain certain properties that will allow us to continue the process in the next iteration.

Proof. Let $\gamma_1, \dots, \gamma_k$ be parameters satisfying $\varepsilon = \gamma_1 \ll \dots \ll \gamma_k = \delta$. We will first prove that for every $i \in [k - 1]$ there is a perfect matching M_i in $\partial^i H[X_1 \cup \dots \cup X_i]$ such that for every $F \in \mathcal{F} \cup \{H\}$ and $v \in V(F) \setminus (X_1 \cup \dots \cup X_i)$, there are at least $(1 - \gamma_i)n$ edges S in M_i such that $S \cup \{v\}$ is supported in F .

We prove this by induction on i . Notice that we can take $M_1 = X_1$, using that F has no isolated vertices, which implies that any $v \in V(F) \setminus X_1$ is supported and thus $d_{X_1}^1(v) \geq (1 - \varepsilon)n = (1 - \gamma_1)n$. For $i \in [2, k - 1]$, suppose that M_{i-1} is a perfect matching in $\partial^i H[X_1 \cup \dots \cup X_{i-1}]$ satisfying the requirements.

Let T_i be the bipartite graph on $M_{i-1} \sqcup X_i$, where Sx , with $S \in M_{i-1}$ and $x \in X_i$, is an edge if $S \cup \{x\}$ is supported in H . We claim that $\delta(T_i) \geq (1 - \gamma_{i-1})n$. Indeed, first notice that, by choice of M_{i-1} , every $S \in M_{i-1}$ is supported in H and thus satisfies $d_{X_i}^1(S) \geq (1 - \varepsilon)n$, which exactly means that $d_{T_i}(S) \geq (1 - \varepsilon)n \geq (1 - \gamma_{i-1})n$. Now, given $x \in X_i$, recall that M_{i-1} was chosen so that for at least $(1 - \gamma_{i-1})n$ sets $S \in M_{i-1}$ we have that $S \cup \{x\}$ is supported in H , which precisely means that $d_{T_i}(x) \geq (1 - \gamma_{i-1})n$.

For every $F \in \mathcal{F} \cup \{H\}$ and every $v \in V(F) \setminus (X_1 \cup \dots \cup X_i)$ let $T_{i,F,v}$ be the subgraph of T_i where Sx , with $S \in M_{i-1}$ and $x \in X_i$, is an edge if $S \cup \{x, v\}$ is supported in F . We claim that $e(T_{i,F,v}) \geq (1 - 2\gamma_{i-1})n^2$. Indeed, by choice of M_{i-1} , we have that at least $(1 - \gamma_{i-1})n$ sets S in M_{i-1} are such that $S \cup \{v\}$ is supported in F . For every such S , since $S \cup \{v\}$ avoids X_i and by

assumption on F , this implies that $d_{X_i}^1(S) \geq (1 - \varepsilon)n$, meaning that $d_{T_{i,F,v}}(S) \geq (1 - \varepsilon)n$. Thus $e(T_{i,F,v}) \geq (1 - \gamma_{i-1})n \cdot (1 - \varepsilon)n \geq (1 - 2\gamma_{i-1})n^2$.

Apply Lemma 11.13 with the graph T_i and the family $\{T_{i,F,v} : F \in \mathcal{F} \cup \{H\}, v \in V(F) \setminus (X_1 \cup \dots \cup X_i)\}$ (which has size at most $(|\mathcal{F}| + 1) \cdot kn \leq n^{c+2}$). This yields a perfect matching M in T_i such that $|M \cap E(T_{i,F,v})| \geq (1 - \gamma_i)n$ for every relevant F and v . Form M_i by including $S \cup \{x\}$ for every edge Sx in M (with $S \in M_{i-1}$ and $x \in X_i$). It is easy to check that M_i is a perfect matching in $\partial^i H[X_1 \cup \dots \cup X_i]$ that has at least $(1 - \gamma_i)n$ edges S such that $S \cup \{v\}$ is supported in F , for every relevant F and v , as required.

Finally, suppose that we found a matching M_{k-1} with the desired properties. We proceed similarly to the above to find a matching M_k that satisfies the requirements of the lemma. Define T_k as above, so that $\delta(T_k) \geq (1 - \gamma_{k-1})n$. For every $F \in \mathcal{F}$, let $T_{k,F}$ be the subgraph of T_k where Sx (with $S \in M_{k-1}$ and $x \in X_k$) is an edge if $S \cup \{x\}$ is an edge in F . By choice of M_{k-1} , for every $x \in V(F) \cap X_k$, it contains at least $(1 - \gamma_{k-1})n$ edges S such that $S \cup \{x\}$ is supported in F , which shows (using that $|V(F) \cap X_k| \geq (1 - \varepsilon)n$ which follows implicitly by the third and fourth items in the statement) that $e(T_{k,F}) \geq (1 - \varepsilon)n \cdot (1 - \gamma_{k-1})n$. As above, apply Lemma 11.13 to obtain a perfect matching M in T_k such that $|M \cap T_{k,F}| \geq (1 - \gamma_k)n$ for every $F \in \mathcal{F}$. This corresponds to a perfect matching M_k in H satisfying $|M_k \cap E(F)| \geq (1 - \gamma_k)n = (1 - \delta)n$ for every $F \in \mathcal{F}$, as required. \square

Finally, we deduce the following corollary, where the requirements on each $F \in \mathcal{F}$ are replaced by F being an almost complete k -partite k -graph. This is done by applying the ‘cleaning’ lemma Corollary 11.11. The concrete upper bound given on the size of \mathcal{F} is based on what we need in a later section.

Corollary 11.15. *Let $1/n \ll \varepsilon \ll \delta \ll 1/k \leq 1/3$. Let H be a k -partite k -graph with parts X_1, \dots, X_k , each having size $|X_i| = n$. Suppose that H has no isolated vertices and, for all $i \in [k]$, every supported set S in H which is disjoint from X_i satisfies $d_{X_i}^1(S) \geq (1 - \varepsilon)n$. Let \mathcal{F} be a collection of at most $2n^k$ subgraphs of H satisfying $e(F) \geq (1 - \varepsilon)n^k$. Then there exists a perfect matching M in G such that $|M \cap E(F)| \geq (1 - \delta)n$ for all $F \in \mathcal{F}$.*

Proof. Let μ be a parameter satisfying $\varepsilon \ll \mu \ll \delta$. For each $F \in \mathcal{F}$, apply Corollary 11.11 with the complete k -partite k -graph on $X_1 \sqcup \dots \sqcup X_k$ playing the role of F , with F playing the role of F' , and parameters ε and μ . Denote the resulting graph F_{clean} . Then F_{clean} has no isolated vertices and $|N_{F_{\text{clean}}}^1(S) \cap X_i| \geq |X_i| - \mu n$ for every supported set S that misses X_i . Thus, by Lemma 11.14, applied with parameter μ in place of ε and family $\{F_{\text{clean}} : F \in \mathcal{F}\}$, a matching with the desired properties exists. \square

12 Structure of extremal hypergraphs

Our aim in this section is to modify the set A from the definition of extremal hypergraphs to obtain better control over co-degrees across and within the parts. Our main result in this section is the following.

Proposition 12.1. *Let $1/n \ll \varepsilon \ll \mu \ll 1/k \leq 1/3$, let $\ell \in [k-1]$ and set $t = \lfloor \frac{k}{k-\ell} \rfloor (k-\ell)$. Let $G \in \mathcal{E}_\varepsilon^{\text{ext}}(n, k, \ell)$. Then there exist subsets $A, A', B' \subseteq V(G)$ and a subgraph $F_B^+ \subseteq \partial^{k-1}[B']$ with vertex set B' that satisfy the following properties.*

- (i) $V(G) = A' \sqcup B'$ and $A \subseteq A'$,
- (ii) $n/t - \mu n \leq |A|, |A'| \leq n/t + \mu n$,
- (iii) any supported set S of size at most $k-1$ that intersects A satisfies $d_{B'}^1(S) \geq |B'| - \mu n$, and
- (iv) F_B^+ has no isolated vertices, $\delta^*(F_B^+) \geq (1-\mu)|B'|$, and $d_A^1(S) \geq |A| - \mu n$ for every $S \in E(F_B^+)$.

This section will be devoted to the proof of Proposition 12.1. Throughout this section, we fix parameters

$$1/n \ll \varepsilon \ll \varepsilon_A \ll \varepsilon_{k-1} \ll \mu \ll 1/k \leq 1/3,$$

and we subsequently fix ℓ and t and a k -graph G as in the statement of Proposition 12.1.

Recall that the extremality of G implies that there is a set A on $\lfloor n/t \rfloor$ vertices that has at most εn^2 supported pairs. In the next claim we modify A so that every vertex of A supports only a few other vertices of A .

Claim 12.2. *There exist subsets $A, B \subseteq V(G)$ such that the following properties hold.*

- (i) $V(G) = A \sqcup B$,
- (ii) $n/t - 4\varepsilon n/\varepsilon_A \leq |A| \leq n/t$,
- (iii) G has at most εn^2 supported pairs with both vertices in A , and
- (iv) every supported set S of size at most $k-1$ that intersects A satisfies $d_A^1(S) \leq \varepsilon_A n$ and $d_B^1(S) \geq |B| - \varepsilon_A n$.

Proof. Let A be a set of $\lfloor n/t \rfloor$ vertices with at most εn^2 supported pairs (which exists since $G \in \mathcal{E}_\varepsilon^{\text{ext}}(n, k, \ell)$ being extremal) and let $B = V(G) \setminus A$. Let x_A be the number of vertices $u \in A$ with $d_A^1(u) \geq \varepsilon_A n/2$. Then the number of supported pairs in A is at least $x_A \varepsilon_A n/4$, but also at most εn^2 , showing that $x_A \leq \frac{4\varepsilon n}{\varepsilon_A}$. Transfer these x_A vertices from A to B . We claim that the modified sets A and B satisfy the requirements of the claim.

The first three parts follow directly from the choice of A . For the final part, let S be a supported set of size at most $k - 1$ with $u \in S \cap A$. Clearly u must support every vertex that S supports, that is, we have $N_A^1(S) \subseteq N_A^1(u)$. This implies

$$d_A^1(S) \leq d_A^1(u) \leq \frac{\varepsilon_A n}{2} \leq \varepsilon_A n,$$

where the final inequality comes from the modification of A . Then, since $d_G^1(S) = d_A^1(S) + d_B^1(S)$ and $d_G^1(S) \geq (1 - 1/t)n - k$ (by Observation 11.1), we have

$$d_B^1(S) \geq d_G^1(S) - d_A^1(S) \geq (1 - 1/t)n - k - \frac{\varepsilon_A n}{2} \geq |B| - \frac{4\varepsilon n}{\varepsilon_A} - (k + 1) - \frac{\varepsilon_A n}{2} \geq |B| - \varepsilon_A n,$$

so the final item holds too. \square

Next, we tackle supported $(k - 1)$ -sets in B which have high vertex co-degree into A .

Definition 12.3. We say that a supported $(k - 1)$ -set $S \subseteq B$ in B is *A-scarce* if $d_A^1(S) \leq |A| - \varepsilon_{k-1}n$. Otherwise, we say that it is *A-rich*.

Let x_{k-1} be the number of *A-rich* $(k - 1)$ -sets in B . In the next claim we show that almost all $(k - 1)$ -sets in B are *A-rich*.

Claim 12.4. *We have $x_{k-1} \geq (1 - \varepsilon_{k-1})\binom{|B|}{k-1}$.*

Proof. Let m denote the number of edges of G with one vertex in A and $(k - 1)$ vertices in B . We will double count m , obtain a lower and an upper bound, and then compare the two.

First, let us lower bound m . We will construct an edge with $(k - 1)$ vertices in B by repeatedly using Claim 12.2 to build larger and larger supported sets until we form an edge. Start with any vertex $u_1 \in A$. Suppose that we have constructed a supported i -set $\{u_1, \dots, u_i\}$, with $u_2, \dots, u_i \in B$ and $i \leq k - 1$. By Claim 12.2, there are at least $|B| - \varepsilon_A n$ choices for a vertex in B that can be used to extend our set to a supported $(i + 1)$ -set. Thus, we see that

$$\begin{aligned} m &\geq \frac{|A|}{(k - 1)!} (|B| - \varepsilon_A n)^{k-1} \\ &\geq \frac{|A|}{(k - 1)!} \left(|B|^{k-1} - (k - 1)\varepsilon_A n |B|^{k-2} \right). \end{aligned}$$

Next, we bound m from above. We do this by looking at all possible supported $(k - 1)$ -sets in B , and then counting the maximum possible number of vertices in A that each such $(k - 1)$ -set could

support. Hence, we have

$$\begin{aligned} m &\leq x_{k-1}|A| + \left(\binom{|B|}{k-1} - x_{k-1} \right) (|A| - \varepsilon_{k-1}n) \\ &\leq x_{k-1}\varepsilon_{k-1}n + (|A| - \varepsilon_{k-1}n) \cdot \frac{|B|^{k-1}}{(k-1)!}. \end{aligned}$$

Finally, we compare the two bounds

$$\begin{aligned} x_{k-1}\varepsilon_{k-1}n &\geq \frac{\varepsilon_{k-1}n|B|^{k-1} - (k-1)\varepsilon_A n|A||B|^{k-2}}{(k-1)!} \\ &\geq n \cdot \frac{|B|^{k-1}}{(k-1)!} \cdot (\varepsilon_{k-1} - (k-1)\varepsilon_A) \\ &= n \cdot \binom{|B|}{k-1} \cdot (\varepsilon_{k-1} - (k-1)\varepsilon_A) + O(n^{k-1}) \\ &\geq n \cdot \binom{|B|}{k-1} \cdot (\varepsilon_{k-1} - k\varepsilon_A), \end{aligned}$$

where we use $|A| \leq |B|$ (since $|A| \leq n/t \leq n/2$). Simplifying further yields

$$x_{k-1} \geq \binom{|B|}{k-1} \left(1 - \frac{k\varepsilon_A}{\varepsilon_{k-1}} \right) \geq \binom{|B|}{k-1} (1 - \varepsilon_{k-1}),$$

proving the claim. \square

We are now ready to prove Proposition 12.1, using the last claim and one of the cleaning lemma from the previous section, and moving a few exceptional vertices from B to A .

Proof of Proposition 12.1. Apply Lemma 11.7 to the subgraph F' of $\partial^{k-1}[B]$ consisting of A -rich supported $(k-1)$ -sets with parameters ε_{k-1} and $\mu/2$ in place of ε and δ respectively, to obtain a subgraph $F_B^+ \subseteq \partial^{k-1}[B]$ that has no isolated vertices and satisfies $\delta^*(F_B^+) \geq (1 - \mu/2)|B|$; the lemma is applicable due to Claim 12.4 above. In particular, we have $|F_B^+| \geq (1 - \mu/2)|B|$. Define $B' = V(F_B^+)$ and $A' := V(G) \setminus B'$. We now prove that Properties (i) to (iv) hold.

Item (i) is immediate from the choice of A' . For the lower bound in (ii), note that $|A'| \geq |A| \geq n/t - 4\varepsilon n/\varepsilon_A \geq n/t - \mu n$, using Claim 12.2. For the upper bound we have $|B'| \geq (1 - \mu)|B| \geq |B| - \mu n$, and thus $|A'| = n - |B'| \leq n - |B| - \mu n = |A| - \mu n \leq n/t - \mu n$, using Claim 12.2 again. For (iii), let S be a supported set of size at most $k-1$ that intersects A . Then $d_A^1(S) \leq \varepsilon_A n$ and $d_B^1(S) \geq |B| - \varepsilon_A n$ by Claim 12.2. Because $|A' \setminus A| \leq \mu n/2$, it follows that $d_A^1(S) \leq \varepsilon_A n + \mu n/2 \leq \mu n$ and $d_{B'}^1(S) \geq |B'| - \varepsilon_A n - \mu n/2 \geq |B'| - \mu n$. Finally, (iv) follows from the choice of F_B^+ . \square

13 Balancing the sizes of A' and B'

In the following two sections (especially in the next one) we will use the notions of vertex sequences that support an ℓ -path or an extended ℓ -path; see Definition 11.5. We will also use the convention that a tight path on $r < k$ vertices is just an ordered supported set of size r .

Our aim in this section is to prove the following proposition that finds a collection \mathcal{P} of vertex-disjoint sequences that support extended ℓ -paths, cover all of A' , and ‘balance’ the sizes of A' and B' , where A' and B' are as given by Proposition 12.1.

Proposition 13.1. *Let $1/n \ll \varepsilon \ll \mu \ll 1/k \leq 1/3$, $\ell \in [k-1]$ such that $(k, \ell) \neq (3, 1)$, and $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$. Suppose that $G \in \mathcal{E}_\varepsilon^{\text{ext}}(n, k, \ell)$ and let A, A', B' be as given by Proposition 12.1. Then there is a collection \mathcal{P} of pairwise vertex-disjoint sequences that support extended ℓ -paths in G such that:*

- (i) Every $P \in \mathcal{P}$ starts and ends in A ,
- (ii) $A' \subseteq V(\mathcal{P})$,
- (iii) $|\mathcal{P}| \geq n/t - 4\mu n$,
- (iv) $|V(\mathcal{P})| \leq n/t + 4t\mu n$,
- (v) $(t-1)|\mathcal{P}| = |B' \setminus V(\mathcal{P})|$,
- (vi) every $P \in \mathcal{P}$ satisfies $|P| = 1$ (so P is a singleton from A) or $|P| \geq t+1$.

Throughout this section, fix parameters $n, \varepsilon, \mu, k, \ell, t$, a k -graph G , the sets A, A', B' and the $(k-1)$ -graph F_B^+ as in the statement of Proposition 12.1, and additionally require $(k, \ell) \neq (3, 1)$. As an intermediate step towards proving Proposition 13.1, we find a collection \mathcal{P}' with slightly weaker requirements.

Proposition 13.2. *There is a collection \mathcal{P}' of pairwise disjoint vertex sequences in $V(G)$ that supported extended ℓ -paths such that:*

- (i) Every $P \in \mathcal{P}'$ is either a singleton from A' or satisfies $|P| = t+1$ and starts and ends in A , and the number of non-singleton sequences P is at most μn ,
- (ii) $A' \subseteq V(\mathcal{P}')$,
- (iii) $|\mathcal{P}'| \geq n/t - \mu n$,
- (iv) $|V(\mathcal{P}')| \leq n/t + t\mu n$,
- (v) $0 \leq |B' \setminus V(\mathcal{P}')| - (t-1)|\mathcal{P}'| \leq t-1$.

In order to prove Proposition 13.2, we first need to establish some preliminary machinery. Let us write $|A'| = \lfloor n/t \rfloor + x$, which implies $|B'| = n - \lfloor n/t \rfloor - x$. If $x \leq 0$ then we can just let \mathcal{P} consist of the set of all vertices in A' , and this would prove Proposition 13.2 (where we use Proposition 12.1 (ii) with parameter $\mu/2$ to bound x and conclude (v)). So we may suppose that $x > 0$. From Proposition 12.1 (ii), we know $x \leq \mu n/2$. We claim that $\partial^2[A']$ has minimum degree at least $x+1$. Indeed, consider $u \in A'$, and notice that u 's degree in $\partial^2[A']$ is $d_{A'}^1(u)$, which satisfies the following, using Observation 11.1.

$$d_{A'}^1(u) = d_G^1(u) - d_{B'}^1(u) \geq \left(n - \left\lfloor \frac{n}{t} \right\rfloor + 1\right) - |B'| = |A'| - \left\lfloor \frac{n}{t} \right\rfloor + 1 = x + 1,$$

as claimed. Hence, by Lemma 11.2 and $|A| \geq |A'| - \mu n > |A'|/2 + 5x$, there is a subgraph $M_{A'}$ of $\partial^2[A']$ that consists of x pairwise vertex-disjoint cherries with leaves in A .

Remark 13.3. A critical point to note is that this is the only part of the entire proof that utilises the exact bound $\delta^*(G) \geq n - \lfloor \frac{n}{t} \rfloor - (k-3)$ via Observation 11.1 and Lemma 11.2. As mentioned in Section 1, we can prove Theorem 2.6 and consequently Theorem 1.1 with the weaker bound of $\delta^*(G) \geq n - \lfloor \frac{n}{t} \rfloor - (k-2)$ instead, for most values of k and ℓ (namely when $t \geq \ell + 2$ or when $\ell = k-1$) which is achieved by modifying Lemma 11.2 to potentially allow for an additional cherry or some edges. We discuss how to suitably adapt the rest of the proof in Appendix B.

Moving on, we extend the cherries in $M_{A'}$ into sequences of order $t+1$ that support extended ℓ -paths, as follows. Let $u_1 u_2 u_3$ denote a cherry of $M_{A'}$ with middle vertex u_2 . We wish to find distinct vertices $v_1, \dots, v_{t-2} \in B'$ such that $\{u_1, u_2, v_1, \dots, v_{t-2}\}$ and $\{u_2, u_3, v_1, \dots, v_{t-2}\}$ are supported sets in G . Given such v_1, \dots, v_{t-2} , we will call the sequence $u_1 u_2 v_1 \dots v_{t-2} u_3$ an *extended cherry* built from the cherry $u_1 u_2 u_3$ of $M_{A'}$ with *extension* $v_1 \dots v_{t-2}$. Notice that, by Proposition 11.6 (iv), extended cherries support extended ℓ -paths.

Claim 13.4. *For each cherry $u_1 u_2 u_3 \in M_{A'}$, there exists an extension $v_1 \dots v_{t-2}$ consisting of vertices in B' such that the extensions for distinct cherries are all pairwise vertex-disjoint.*

Proof. For any cherry $u_1 u_2 u_3 \in M_{A'}$ with middle vertex u_2 , we know that $u_1, u_3 \in A$, and therefore $\{u_1, u_2\}$ and $\{u_2, u_3\}$ are supported sets that each contain at least one vertex of A . Hence, Proposition 12.1 (iii) implies

$$d_{B'}^1(u_i u_{i+1}) \geq |B'| - \mu n, \quad i \in \{1, 2\}.$$

The inequality above tells us that the number of common vertex neighbours of $\{u_1, u_2\}$ and $\{u_2, u_3\}$ in B' satisfies

$$|N_{B'}^1(u_1 u_2) \cap N_{B'}^1(u_2 u_3)| \geq |B'| - 2\mu n.$$

Pick any common neighbour $v_1 \in N_{B'}^1(u_1 u_2) \cap N_{B'}^1(u_2 u_3)$. We will build the extension one vertex at a time for $(t-2)$ steps – after the i th step, we have vertices $v_1, \dots, v_i \in B'$ such that $\{u_1, u_2, v_1, \dots, v_i\}$

and $\{u_2, u_3, v_1, \dots, v_i\}$ are supported $(i+2)$ -sets, each containing at least one vertex from A . Then, using the same rationale as above, they have at least $|B'| - 2\mu n$ common vertex neighbours in B' , and we will pick v_{i+1} from this common vertex neighbourhood. By construction, $\{u_1, u_2, v_1, \dots, v_{t-2}\}$ and $\{u_2, u_3, v_1, \dots, v_{t-2}\}$ are supported sets in G as desired.

All that is left is to show that we can make distinct choices each time to ensure that the extensions are disjoint. We know that $M_{A'}$ consists of $x \leq \mu n$ pairwise disjoint cherries. As each extension requires $t-2$ vertices and since we have at least $|B'| - 2\mu n \geq t\mu n$ choices for a vertex in each step of extending a cherry, we can find disjoint extensions in B' for each cherry of $M_{A'}$. \square

We are now ready to prove Proposition 13.2.

Proof of Proposition 13.2. As before, we write $|A'| = \lfloor n/t \rfloor + x$, where we can assume $0 < x \leq \mu n/2$. Let $M_{A'}$ be the subgraph of $\partial^2[A']$ guaranteed by the preceding argument; i.e. it consists of x pairwise disjoint cherries. Then, invoking Claim 13.4, we extend the cherries of $M_{A'}$ so that the extensions are pairwise vertex-disjoint. Let \mathcal{P}' be the collection consisting of: a single vertex for each vertex in $A' \setminus V(M_{A'})$, and an extended cherry for each cherry in $M_{A'}$ (so that the extensions are pairwise vertex-disjoint). Then item (i) in statement of Proposition 13.2 clearly holds, using that singletons and extended cherries support extended ℓ -paths (see Proposition 11.6 (iv)) and that $M_{A'}$ consists of $x \leq \mu n$ cherries. Item (ii) is immediate by construction. For (iii) we have

$$|\mathcal{P}'| = |A'| - 2x = \left\lfloor \frac{n}{t} \right\rfloor - x \geq \frac{n}{t} - \mu n. \quad (13.5)$$

For (iv), notice that

$$|V(\mathcal{P}')| = |A'| + (t-2)x = \left\lfloor \frac{n}{t} \right\rfloor + (t-1)x \leq \frac{n}{t} + t\mu n.$$

as required. Also, using (13.5),

$$|B' \setminus V(\mathcal{P}')| - (t-1)|\mathcal{P}'| = n - \left\lfloor \frac{n}{t} \right\rfloor - x - (t-2)x - (t-1) \left(\left\lfloor \frac{n}{t} \right\rfloor - x \right) = n - t \left\lfloor \frac{n}{t} \right\rfloor.$$

Since $0 \leq n - t \left\lfloor \frac{n}{t} \right\rfloor < t$, this proves (v), the final claim of Proposition 13.2. \square

To prove Proposition 13.1, we modify the family given by Proposition 13.2 by first extending single path by an appropriate amount, and then extending all other paths as needed so that they start and end in A . The following claim will allow us to make these extensions.

Claim 13.6. *Let $a \in A'$, let W be a set of at most $10t\mu n$ forbidden vertices. Then there are distinct vertices $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1} \in B' \setminus W$ and $a_1, a_2 \in A \setminus W$ such that the sequence $a_1 v_1 \dots v_{t-1} a u_1 \dots u_{t-1} a_2$ supports an extended ℓ -path.*

Proof. By Proposition 11.6 (iii), it suffices to find distinct $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1} \in B' \setminus W$ and $a_1, a_2 \in A \setminus W$ so that: $v_1 \dots v_{t-1} a u_1 \dots u_{t-1}$ is a tight path and the sets $\{a_1, v_1, \dots, v_{t-1}\}$ and $\{a_2, u_1, \dots, u_{t-1}\}$ are supported. Indeed, in the sequence $a_1 v_1 \dots v_{t-1} a u_1 \dots u_{t-1} a_2$, the only subsequences of consecutive vertices that have at most one index divisible by t (when we number the indices by 0 to $2t$) are contained in $a_1 v_1 \dots v_{t-1}$, $v_1 \dots v_{t-1} a u_1 \dots u_{t-1}$ or $u_1 \dots u_{t-1} a_2$, and it is easy to check that for each of them every sequence of at most k consecutive edges forms a supported set.

First, observe that if P is a tight path of order $p \leq 4k$, and W' is a set of size at most $11t\mu n$, then the number of vertices $v \in B' \setminus W'$ such that Pv (similarly vP) is a tight path of order $p+1$ is the number of vertex neighbours of the last $\min\{p, k-1\}$ vertices in P (or the first $\min\{p, k-1\}$ vertices in P if we are looking for vP) that lie outside W' , which, by Observation 11.1 and Proposition 12.1, is at least

$$n - \frac{n}{t} - k - |A'| - |W'| - p \geq \frac{n}{2t},$$

using that $t \geq 3$ (which follows from $(k, \ell) \neq (3, 1)$; see Observation 3.1). Starting with a and applying this inequality repeatedly to obtain the vertices $u_1, \dots, u_{t-1}, v_{t-1}, \dots, v_1$ in that order, we find that the number of sequences $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1} \in B' \setminus W$ such that $v_1 \dots v_{t-1} a u_1 \dots u_{t-1}$ is a tight path is at least $(\frac{n}{2t})^{2(t-1)}$.

We argue that there is such a sequence such that $\{v_1, \dots, v_{t-1}\}$ and $\{u_1, \dots, u_{t-1}\}$ are supported sets in F_B^+ . Indeed, as $V(F_B^+) = B'$ and F_B^+ has no isolated vertices, by the minimum positive co-degree assumption, the number of sequences $x_1, \dots, x_{t-1} \in B'$ such that $\{x_1, \dots, x_{t-1}\}$ is supported in F_B^+ is at least $(1 - \mu)^{t-2} |B'|^{t-1}$. Thus, the number of sequences $x_1, \dots, x_{t-1} \in B'$ such that $\{x_1, \dots, x_{t-1}\}$ is not supported in F_B^+ is at most $(1 - (1 - \mu)^{t-2}) |B'|^{t-1} \leq t\mu n^{t-1}$, implying that the number of sequences $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1} \in B'$ such that one of $\{v_1, \dots, v_{t-1}\}$ and $\{u_1, \dots, u_{t-1}\}$ is not supported in F_B^+ is at most $2t\mu n^{2(t-1)} < (\frac{n}{2t})^{2(t-1)}$. Thus, there is a choice of $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1}$ as above such that the v_i 's and the u_i 's form supported sets in F_B^+ .

To finish, recall that supported sets of size at most $k-1$ in F_B^+ are A -rich. This means that $|N^1(\{v_1, \dots, v_{t-1}\}) \cap A| \geq |A| - \mu n \geq n/t - 2\mu n$, which implies that there is a vertex $a_1 \in (A \setminus (W \cup \{a\})) \cap N^1(\{v_1, \dots, v_{t-1}\})$. Similarly, there is a vertex $a_2 \in (A \setminus (W \cup \{a, a_1\})) \cap N^1(\{u_1, \dots, u_{t-1}\})$. The sequence $v_1, \dots, v_{t-1}, u_1, \dots, u_{t-1}, a_1, a_2$ satisfies the requirement of the claim. \square

Claim 13.7. *Let W be a set of at most $10t\mu n$ forbidden vertices and let r satisfy $t+1 \leq r \leq 2t$ and $r \equiv 1 \pmod{k-\ell}$. Then there is a sequence of distinct r vertices in $V(G) \setminus W$ that starts and ends in A , its other vertices are in B' , and it supports an extended ℓ -path.*

Proof. We follow a similar strategy to the previous proof. Pick some $a_1 \in A \setminus W$. By the second paragraph in the previous proof, there are at least $(\frac{n}{2t})^{r-2}$ sequences $v_1 \dots v_{r-1}$ of distinct vertices in $B' \setminus W$ such that $a v_1 \dots v_{r-1}$ is a tight path in G . By the third paragraph in the same proof, there is such a sequence where the final $\min\{r-1, k-1\}$ vertices form a supported set in F_B^+ . This

implies that there is $a_2 \in A \setminus (W \cup \{a_1\})$ such that $v_1 \dots v_{r-1} a_2$ is a tight path in G . Finally, by Proposition 11.6 (iv), we conclude that $a_1 v_1 \dots v_{r-2} a_2$ is the desired sequence. \square

Proof of Proposition 13.1. Let \mathcal{P}' be the output of Proposition 13.2. We modify \mathcal{P}' in two steps as described above. Let A_0 be the set of all vertices in $A' \setminus A$ that form singleton sequences in \mathcal{P}' . Then, if we invoke Proposition 12.1 with $\mu/2$ in place of μ (note that this makes all its conclusions strictly stronger, so all preceding results in this section still hold), we have $|A_0| \leq |A' \setminus A| \leq \mu n$. By repeatedly applying Claim 13.6, we find, for each $a \in A_0$, a sequence S_a consisting of a , two vertices from A and $2(t-1)$ vertices from B' (in some order), so that: S_a starts and ends in A ; it supports an extended ℓ -path; and the S_a 's are pairwise vertex-disjoint. Define \mathcal{P}'' to be the collection of vertex sequences P such that either $P \in \{S_a : a \in A_0\}$ or $P \in \mathcal{P}'$ is not a singleton vertex of A that appears in some sequence S_a with $a \in A_0$. Write $r := |B' \setminus V(\mathcal{P}'')| - (t-1)|\mathcal{P}''|$. Then $r = |B' \setminus V(\mathcal{P}')| - (t-1)|\mathcal{P}'|$, because replacing the three singleton paths from A' that appear in S_a by S_a decreases $|B' \setminus V(\mathcal{P}')|$ by $2(t-1)$ and $|\mathcal{P}'|$ by 2, showing that $0 \leq r \leq t$ via Proposition 13.2 (v). Moreover, by Proposition 13.2 (ii), we know $A' \subseteq V(\mathcal{P}')$, so we see that $|B' \setminus V(\mathcal{P}')| + |V(\mathcal{P}')| = |A'| + |B'| = n$. Hence, we have

$$n = r + (t-1)|\mathcal{P}'| + |V(\mathcal{P}')| \equiv r + (t-1)|\mathcal{P}'| + |\mathcal{P}'| \equiv r \pmod{k-\ell},$$

using that all sequences in \mathcal{P}' are pairwise disjoint and have length $1 \pmod{t}$, and that t is divisible by $k-\ell$. Since n is divisible by $k-\ell$, so is r . In particular, $r+t+1 \equiv 1 \pmod{k-\ell}$. By Claim 13.7, we can find a sequence S of $r+t+1$ vertices such that: it starts and ends in A ; its other vertices are in B' ; it supports an extended ℓ -path; and it is disjoint from all non-singleton sequences in \mathcal{P}'' , where we use the fact that the sequences $\{S_a : a \in A_0\}$ span at most $2t|A_0| \leq 2t\mu n$ vertices and Proposition 13.2 (i) implies that the other non-singleton sequences in \mathcal{P}'' span at most $(t+1)\mu n$ vertices. Let \mathcal{P} be obtained from \mathcal{P}'' by removing the two singleton sequences in \mathcal{P}'' that appear in S and adding S .

It remains to verify that the five desired properties in the statement of Proposition 13.1 indeed hold. Throughout, we will utilise Proposition 13.2 items (i) to (v). The first two are immediate from the construction of \mathcal{P} . For the property (iii), we have

$$|\mathcal{P}| = |\mathcal{P}''| - 1 \geq |\mathcal{P}'| - 2|A' \setminus A| \geq |\mathcal{P}'| - 2\mu n - 1 \geq \frac{n}{t} - 4\mu n.$$

For the item (iv),

$$\begin{aligned} |V(\mathcal{P})| &\leq |V(\mathcal{P}'')| + (r+t+1) \leq |V(\mathcal{P}')| + (2t+1)|A' \setminus A| + (r+t+1) \\ &\leq \frac{n}{t} + (t+2t+1)\mu n + 2t \leq \frac{n}{t} + 4t\mu n. \end{aligned}$$

Next, by the choice of \mathcal{P} , we have $|\mathcal{P}| = |\mathcal{P}''| - 1$ and $|B' \setminus V(\mathcal{P})| = |B' \setminus V(\mathcal{P}'')| - (r+t-1)$, and

therefore

$$|B' \setminus V(\mathcal{P})| - (t-1)|\mathcal{P}| = |B' \setminus V(\mathcal{P}'')| - (t-1)|\mathcal{P}''| - (r+t-1) + t-1 = 0,$$

proving (v). Finally, (vi) is immediate from the construction: all sequences have length 1 (singletons from A), $t+1$ (extended cherries), $2t+1$ (sequences S_a with a being a singleton from A' in \mathcal{P}'), and $r+t+1 \geq t+1$ for the final sequence S . \square

14 Building a Hamilton cycle

Rather than constructing a Hamilton ℓ -cycle in G directly, we first build a perfect matching in an auxiliary balanced t -partite t -graph (see G^+ , defined below). We will do so in a semi-random manner using Lemma 11.14, and argue that the co-degrees in the auxiliary graph are large enough to connect this matching into a Hamilton cycle, and the auxiliary t -graph is chosen such that this will correspond to a Hamilton ℓ -cycle in G .

The t -partite t -graph will have one part corresponding to the collection \mathcal{P} found in the previous section, and the other $t-1$ parts will correspond to an arbitrary equipartition of $B'' := B' \setminus V(\mathcal{P})$. That these parts have the same size follows from the fifth (and crucial) property (v) of \mathcal{P} , which asserts that $|B''| = (t-1)|\mathcal{P}|$. Because we intend to use Lemma 11.14, which applies to t -partite t -graphs with very high minimum supported co-degree, we need to prove that G^+ indeed has high minimum supported co-degree, and this will follow from the choice of B' as the vertex set of the hypergraph F_B^+ (see Proposition 12.1), whose edges are A -rich sets, and from property (i) of \mathcal{P} , which guarantees that each sequence in \mathcal{P} starts and ends in A .

Throughout this section, we will frequently recall the notion of A -rich sets (recall Definition 12.3) and, like in the previous section, that of vertex sequences that support (extended) ℓ -paths (recall Definition 11.5). Again, we say that a tight path on $r < k$ vertices is simply an (ordered) supported set of size r . As in the previous two sections, we fix

$$1/n \ll \varepsilon \ll \mu \ll 1/k \leq 1/3,$$

and we subsequently set $\ell \in [k-1]$ such that $(k, \ell) \neq (3, 1)$, $t = \left\lfloor \frac{k}{k-\ell} \right\rfloor (k-\ell)$, and let $G \in \mathcal{E}_\varepsilon^{\text{ext}}(n, k, \ell)$ be a k -graph. We fix sets $A, A', B \subseteq V(G)$ as in Propositions 12.1 and 13.1 and let \mathcal{P} be the collection of vertex sequences output by Proposition 13.1.

Write $m := |\mathcal{P}|$, define $B'' := B' \setminus V(\mathcal{P})$ (so $|B''| = (t-1)m$ by Proposition 13.1 (v)) and let $B_1 \sqcup \dots \sqcup B_{t-1}$ be an arbitrary equipartition of B'' . We define a t -partite t -graph G^+ as follows. Let $V(G^+) = B_1 \sqcup \dots \sqcup B_{t-1} \sqcup \mathcal{P}$ and $E(G^+)$ be the subset of $B_1 \times \dots \times B_{t-1} \times \mathcal{P}$ such that any $(v_1, \dots, v_{t-1}, P) \in B_1 \times \dots \times B_{t-1} \times \mathcal{P}$ forms an edge in G^+ if and only if the vertex sequence corresponding to $v_1 \dots v_{t-1}P$ (which, we note, has length congruent to 0 modulo $k-\ell$ by

Proposition 13.1 and Definition 11.5) supports an ℓ -path in G and $\{v_1, \dots, v_{t-1}\}$ is supported in F_B^+ . The partite structure of G^+ lends an inherent ordering to every t -edge. Specifically, we will always think of edges of G^+ with vertices as ordered tuples from $B_1 \times \dots \times B_{t-1} \times \mathcal{P}$, and will often write them as vertex sequences with this canonical ordering.

As mentioned above, we would like to apply Lemma 11.14 to G^+ . This requires showing that G^+ has high minimum supported co-degree, which is what we show in the next claim. For notational convenience, we define $B_t := \mathcal{P}$ and let \mathcal{P}_A be the set of singleton paths in \mathcal{P} (which are, by property (i) of \mathcal{P} , vertices in A). Notice that the number of non-singleton paths in \mathcal{P} is at most $\frac{1}{t+1}(|V(\mathcal{P})| - |\mathcal{P}|) \leq 4\mu n$ (using properties (iii) and (iv) of \mathcal{P} , the fact non-singleton paths have length at least $t+1$ from (vi), and $t \geq 2$). Hence,

$$|\mathcal{P}_A| \geq |\mathcal{P}| - 4\mu n. \quad (14.1)$$

Claim 14.2. *Let $i \in [t]$ and let S be a supported set in G^+ that avoids B_i , then $d_{B_i}^1(S) \geq |B_i| - 5t\mu n$ (where d^1 is taken with respect to G^+).*

Proof. Since S is supported, it is contained in some edge e , which has the form $b_1 \dots b_t$, with $b_i \in B_i$. Writing $S' = \{b_1, \dots, b_t\} \setminus \{b_i\}$, it suffices to show $d_{B_i}^1(S') \geq |B_i| - 5t\mu n$ in G^+ , as $d_{B_i}^1(S) \geq d_{B_i}^1(S')$. We consider two cases: $i = t$ and $i \in [t-1]$. Let us begin with the first case. By definition of G^+ , we know that $b_1 \dots b_{t-1}$ is supported by some edge of F_B^+ , which means that $b_1 \dots b_{t-1}$ contained in some A -rich set, implying that $b_1 \dots b_{t-1}a$ is supported in G for all but at most μn vertices $a \in A$, and hence for all but at most μn vertices $a \in \mathcal{P}_A$. Since every (ordered) supported t -set supports an ℓ -path, it follows that $b_1 \dots b_{t-1}a \in E(G^+)$ for all but at most μn vertices $a \in \mathcal{P}_A$. Thus, we find that

$$d_{B_i}^1(S') \geq d_{\mathcal{P}_A}^1(b_1 \dots b_{t-1}) \geq |\mathcal{P}_A| - \mu n \geq |\mathcal{P}| - 5\mu n \geq |B_t| - 5t\mu n, \quad (14.3)$$

using (14.1) and $B_t = \mathcal{P}$.

For the second case, recall, from Proposition 13.1, that $b_t = w_1 \dots w_r$ is a vertex sequence with $w_1, w_r \in A$. For convenience, define $w_j = w_r$ for $j \geq r+1$. For $j \in [i]$, write

$$e_j = \{b_j, \dots, b_{i-1}, b_{i+1}, \dots, b_{t-1}, w_1, \dots, w_{k-t+j}\}.$$

We claim that, for any $b_i \in B_i$, the sequence $b_1 \dots b_t$ is an edge in G^+ if and only if the following conditions hold.

- $b_i \in N_{F_B^+}^1(b_1 \dots b_{i-1} b_{i+1} \dots b_{t-1})$,
- $b_i \in N_G^1(e_j)$ for $j = 1$ and for every other $j \in [i]$ such that $j \equiv 1 \pmod{k-\ell}$ and $k-t+j \leq r$.

Indeed, the first condition is equivalent to $b_1 \dots b_{t-1}$ being supported in F_B^+ (using that S' is supported in G^+ , which implies that $b_1 \dots b_{i-1} b_{i+1} \dots b_{t-1}$ is supported in F_B^+), so it suffices to

check that the second condition is equivalent to $b_1 \dots b_t w_1 \dots w_r$ supporting an ℓ -path. If $r = 1$ one just needs to check when $b_1 \dots b_{t-1} w_1$ is supported in G , which is exactly when $b_i \in N_G^1(e_1)$ (again using that S' is supported in G^+ , showing that e_1 is supported in G), so we may assume that $r \geq 2$, which implies $r \geq t + 1$ by (vi) in Proposition 13.1. Recall that for $b_1 \dots b_{t-1} w_1 \dots w_r$ to support an ℓ -path we need that every k consecutive vertices in this sequence, that start at an index which is 1 (mod $k - \ell$), form an edge, and that the last t vertices in the sequence form a supported set in G . The latter holds due to $w_1 \dots w_r$ supporting an extended ℓ -path and $r \geq t + 1$. The former is equivalent to the second item above (noting that the sequence has length at least $2t \geq k$).

In particular, since S' is supported in G^+ we know that there is some $b_i \in B_i$ that satisfies the above requirements, showing that $b_1 \dots b_{i-1} b_{i+1} \dots b_{t-1}$ is supported in F_B^+ and that e_j is supported in G for $j = 1$ and every other $j \leq \min\{i, r - k + t\}$ with $j \equiv 1 \pmod{k - \ell}$; denote the set of such j 's by J .

Since, for every $j \in J$, we have that e_j is a supported set in G that contains $w_1 \in A$, it satisfies $d_{B'}^1(e_j) \geq |B'| - \mu n$ by Proposition 12.1 (iii), and thus $d_{B_i}^1(e_j) \geq |B_i| - \mu n$. Moreover, by Proposition 12.1 (iv), we have $d_{F_B^+}^1(b_1 \dots b_{i-1} b_{i+1} \dots b_{t-1}) \geq |B'| - \mu n$ in B' . Altogether, we have $d_{B_i}^1(S') \geq |B_i| - t\mu n \geq |B_i| - 5t\mu n$ in G^+ . \square

Before proceeding further, it would be useful to describe an outline of our strategy for the remaining stages of the proof. Our next step will be to find a perfect matching M in G^+ , which corresponds to a collection of pairwise disjoint vertex sequences in G that support ℓ -paths, and also span G . We then form an auxiliary digraph D^+ whose vertices are the edges of M , and we direct an edge from e to f whenever the vertex sequence ef supports an ℓ -path. In particular, a directed Hamilton cycle in D^+ will correspond to a spanning ℓ -cycle in G (using Proposition 11.6 (ii)). Thus, we need to show that D has large minimum semi-degree, and then Lemma 11.4 will provide the desired cycle. We will show that each edge of G^+ has many potential in-neighbour and out-neighbour edges in G^+ . We then find the matching M in G^+ using Lemma 11.14, and argue that the digraph D^+ formed as a result of this matching maintains high semi-degree conditions.

To formalise this, for an edge $e = b_1 \dots b_{t-1} P \in E(G^+)$ (where $b_i \in B_i$ and $P \in \mathcal{P}$), let F_e^+ be the subgraph of G^+ consisting of edges $f = b'_1 \dots b'_{t-1} P'$ (where $b'_i \in B_i$ and $P' \in \mathcal{P}$) such that the vertex sequence $ef = b_1 \dots b_{t-1} P b'_1 \dots b'_{t-1} P'$ supports an ℓ -path in G , and similarly define F_e^- as the subgraph of G^+ consisting of edges f such that fe supports an ℓ -path in G .

Claim 14.4. $e(F_e^+) \geq |B_1|^t - 5t\mu n^t$ for every $e \in E(G^+)$.

Proof. Write $e = b_1 \dots b_{t-1} P$. We wish to find edges $f = b'_1 \dots b'_{t-1} P' \in G^+$ such that ef supports an ℓ -path in G . By assumption, e supports an ℓ -path, which implies that the last t vertices of e , say $w_1 \dots w_t$, constitute a supported set in G . Since we know that f supports an ℓ -path in G as well, if we can show that $w_1 \dots w_t b'_1 \dots b'_{t-1}$ is a tight path, then by Proposition 11.6 (i) we have $f \in E(F_e^+)$.

From Proposition 13.1 (i), we know that $w_t \in A$, and then Proposition 12.1 (iii) implies that $d_{B'}^1(S) \geq |B'| - \mu n$ and consequently $d_{B_i}^1(S) \geq |B_i| - \mu n$ for any G -supported set S containing w_t and for any $i \in [t-1]$. We show by induction that, for $i \in [t-1]$, the number of vertex sequences $b'_1 \dots b'_i \in B_1 \times \dots \times B_i$ such that $w_1 \dots w_t b'_1 \dots b'_i$ is a tight path in G and $\{b'_1, \dots, b'_i\}$ is supported in F_B^+ is at least $(|B_1| - 2\mu n)^i$. This holds for $i = 1$ by the arguments above. Let $i \in [2, t-1]$ and suppose the assertion holds up to $i-1$. Let $b'_1 \dots b'_{i-1} \in B_1 \times \dots \times B_{i-1}$ be such that $w_1 \dots w_t b'_1 \dots b'_{i-1}$ is a tight path and $\{b'_1, \dots, b'_{i-1}\}$ is supported in F_B^+ . Then, for $b'_i \in B_i$ to be such that $w_1 \dots w_t b'_1 \dots b'_i$ is a tight path and $\{b'_1, \dots, b'_i\}$ is supported in F_B^+ , we exactly need

$$b'_i \in N_G^1(w_{s_i} \dots w_t b'_1 \dots b'_{i-1}) \cap N_{F_B^+}^1(b'_1 \dots b'_{i-1}) \cap B_i, \quad s_i := \max\{t - (i - 1) + 1, 1\}.$$

By the previous remark and the structure of F_B^+ given by Proposition 12.1 (iv), there are at least $|B_i| - 2\mu n = |B_1| - 2\mu n$ options for b'_i , showing that the number of sequences b'_1, \dots, b'_i as required is at least $(|B_1| - 2\mu n)^i$, as claimed. To finish off, we recall that $b'_1 \dots b'_{t-1} a$ is an edge in G^+ for every $a \in \mathcal{P}_A \cap N_G^1(b'_1 \dots b'_{t-1})$. Since $\{b'_1, \dots, b'_{t-1}\}$ is supported in F_B^+ , it is A -rich and thus the number of such a is at least $|\mathcal{P}_A| - \mu n \geq |\mathcal{P}| - 5\mu n = |B_1| - 5\mu n$, by (14.1). In total we conclude that the number of edges f in G^+ that are in F_e^+ is at least $(|B_1| - 2\mu n)^{t-1}(|B_1| - 5\mu n) \geq |B_1|^t - (2(t-1)\mu n + 5\mu n)|B_1|^{t-1} \geq |B_1|^t - 5t\mu n^t$. \square

Claim 14.5. $e(F_e^-) \geq |B_1|^t - 5t\mu n^t$ for every $e \in E(G^+)$.

Proof. Again, set $e = b_1 \dots b_t P$. Recall that we want edges $f = b'_1 \dots b'_{t-1} P' \in G^+$ such that fe supports an ℓ -path in G . From (14.1), we know that most elements of $B_t = \mathcal{P}$ consist of singleton vertices from A . For simplicity, we will only try to find edges in F_e^- that are of the form $f = b'_1 \dots b'_{t-1} a$, where $a \in \mathcal{P}_A$. Analogous to the situation above, since e and f support ℓ -paths in G , in order for fe to do so as well, it will suffice to argue that $b'_1 \dots b'_{t-1} a b_1 \dots b_{t-1}$ is a tight path in G .

We inductively argue that, for each $i \in [0, t-1]$, there are at least $(|B_1| - 2\mu n)^i(|B_1| - 5\mu n)$ sequences $b'_{t-i} \dots b'_{t-1} a \in B_{t-i} \times \dots \times B_{t-1} \times \mathcal{P}_A$ such that $b'_{t-i} \dots b'_{t-1} a b_1 \dots b_{t-1}$ is a tight path in G and $\{b'_{t-i}, \dots, b'_{t-1}\}$ is supported in F_B^+ . Start with $i = 0$. By definition of G^+ , we know that $\{b_1, \dots, b_{t-1}\}$ is A -rich, meaning that $a b_1 \dots b_{t-1}$ is supported in G for all but at most μn vertices $a \in \mathcal{P}_A$ and hence, by (14.1), for at least $|\mathcal{P}_A| - \mu n \geq |B_1| - 5\mu n$ vertices $a \in \mathcal{P}_A$. Let $i \in [1, t-1]$ and suppose the induction hypothesis holds for $i-1$. Let $b'_{t-i+1} \dots b'_{t-1} a \in B_{t-i+1} \times \dots \times B_{t-1} \times \mathcal{P}_A$ be such that $b'_{t-i+1} \dots b'_{t-1} a b_1 \dots b_{t-1}$ is a tight path in G and $\{b'_{t-i+1}, \dots, b'_{t-1}\}$ is supported in F_B^+ . Then, for b'_{t-i} to be a valid choice, we require exactly that

$$b'_{t-i} \in N_G^1(b'_{t-i+1} \dots b'_{t-1} a b_1 \dots b_{s_i}) \cap N_{F_B^+}^1(b'_{t-i+1} \dots b'_{t-1}) \cap B_i, \quad s_i := \min\{k - i - 1, t - 1\}.$$

Crucially, as $a \in \mathcal{P}_A \subseteq A$, we can use the same rationale as in the above proof of Claim 14.4 to conclude that there are at least $|B_{t-i}| - 2\mu n = |B_1| - 2\mu n$ choices for b'_{t-i} . This proves the

inductive statement, and the same calculations as in the proof of Claim 14.4 complete the proof of this claim. \square

We are now finally ready to prove Theorem 2.6.

Proof of Theorem 2.6. Fix a constant μ' such that $\mu \ll \mu' \ll 1/k \leq 1/3$, let G^+ be the auxiliary t -partite t -graph defined above, and let $\mathcal{F} := \{F_e^+, F_e^- : e \in E(G^+)\}$ be as defined above. Notice that G^+ and \mathcal{F} satisfy the requirements of Corollary 11.15 (with parameters $|B_1| = m$, μ' and t in place of n , ε and k respectively), using Claims 14.2, 14.4 and 14.5. Thus, by Corollary 11.15 there is a perfect matching M in G^+ such that $|M \cap E(F_e^+)|, |M \cap E(F_e^-)| \geq m/2$ for every $e \in E(G^+)$.

As discussed previously, we now define a corresponding auxiliary digraph D^+ . Set $V(D^+) = E(M)$, so the vertices of D^+ correspond to edges of G^+ . For any $e \in V(D^+)$, define its in and out neighbourhoods to be

$$N^+(e) = \{f \in V(D^+) : f \in F_e^+\} \text{ and } N^-(e) = \{f \in V(D^+) : f \in F_e^-\}.$$

From the choice of M and D^+ , we have $\delta^0(D^+) \geq m/2 = |D^+|/2$. In particular, if we let $M = \{P_1, \dots, P_m\}$ where we think of each $P_i = (b_1, \dots, b_{t-1}, P) \in B_1 \times \dots \times B_{t-1} \times \mathcal{P}$ as the vertex sequence $b_1 \dots b_{t-1}P$, then Lemma 11.4 can be applied to find a directed Hamilton cycle $P_1 \dots P_m$ in D^+ .

The crucial observation we make here is that M consists of pairwise disjoint vertex sequences in G^+ that span G , each of which supports an ℓ -path in G . By Proposition 11.6 (ii), we conclude that the vertex sequence $P_1 \dots P_m$ is a Hamilton tight cycle in G . \square

15 Open problems

Recall that our main result determined (up to an additive error of 1) the minimum supported co-degree condition guaranteeing a spanning ℓ -cycle in a k -graph on n vertices (when $k - \ell$ divides n , which is a necessary divisibility condition), for all $k \geq 3$ and $\ell \in [k - 1]$ except when $(k, \ell) = (3, 1)$. As mentioned in Section 1, an upcoming paper of Mycroft and Zárate-Guérén determines the threshold for 3-uniform loose Hamilton cycles.

However, the broader research theme of determining the optimal minimum supported co-degree to find a specified spanning substructure poses some tantalising questions. Our original motivation for working on this problem arose from Illingworth, Lang, Műyesser, Parczyk and Sgueglia [15, Conjecture 1.4], who conjectured that $\delta^*(G) \geq (1 - 1/k)n$ guarantees a tight Hamilton cycle. The main focus of their paper is the minimum supported co-degree required to ensure a spanning $(k - 1)$ -sphere – a subset of edges in a k -graph that induces a homogeneous simplicial complex

homeomorphic to the $(k - 1)$ -sphere \mathbb{S}^{k-1} , and whose 0-skeleton spans the entire vertex set. They tackle the following conjecture of Georgakopoulos, Haslegrave, Montgomery and Narayanan [8] (paraphrased).

Conjecture 15.1. *Every k -graph G of order $n > k \geq 2$ with $\delta^*(G) \geq n/2$ and with no isolated vertices contains a spanning copy of \mathbb{S}^{k-1} .¹*

The chief result in [8] provides an asymptotically optimal minimum co-degree condition for a 3-graph to contain a vertex-spanning copy of any *surface*, meaning an arbitrary connected, closed 2-manifold. In particular, this addresses the minimum co-degree necessary to guarantee a spanning 2-sphere in a 3-graph.

Theorem 15.2 ([8, Theorem 1.4]). *Let \mathcal{S} be an arbitrary surface and $\varepsilon > 0$. Then any sufficiently large n -vertex 3-graph with $\delta(G) \geq n/3 + \varepsilon n$ contains a spanning copy of \mathcal{S} . Moreover, for any $n \in \mathbb{N}$, there exists an n -vertex 3-graph H with $\delta(H) = \lfloor \frac{n}{3} \rfloor - 1$ such that there are at most $2 \lceil \frac{n}{3} \rceil$ vertices in the 0-skeleton of a copy of any surface in H .*

The authors of [15] utilise the same blow-up framework we use to prove Theorem 2.5 (as mentioned, Lemma 4.1 is based on Lemma 2.1 of their paper) to prove an asymptotic weakening of this conjecture, that is, that $\delta^*(G) \geq n/2 + o(n)$ suffices. We believe that the methods used here can potentially be used to improve this to a stability result for Conjecture 15.1, and possibly achieve the same improvement for Theorem 15.2 as well. Perhaps one could build on such results to prove an exact bounds, thereby fully settling these conjectures, though this is likely to require some careful extremal analysis.

References

- [1] J. Balogh, A. Halfpap, B. Lidický, and C. Palmer, *Positive co-degree densities and jumps*, arXiv:2412.08597 (2024). 4
- [2] J. Balogh, N. Lemons, and C. Palmer, *Maximum size intersecting families of bounded minimum positive co-degree*, SIAM J. Discrete Math. **35** (2021), no. 3, 1525–1535. 4
- [3] A. Czygrinow and T. Molla, *Tight codegree condition for the existence of loose Hamilton cycles in 3-graphs*, SIAM J. Discrete Math. **28** (2014), no. 1, 67–76. 4
- [4] G. A. Dirac, *Some theorems on abstract graphs*, Proc. Lond. Math. Soc. **3** (1952), no. 1, 69–81. 3
- [5] P. Frankl and J. Wang, *Intersecting families with large shadow degree*, Acta Math. Hungar. (2025), 1–11. 4
- [6] A. Frieze, *Hamilton cycles in random graphs: a bibliography*, arXiv:1901.07139 (2019). 3

¹In the original phrasing of the conjecture G is required to be tightly connected – meaning that any two edges can be joined by a tight walk – but this follows from the other assumptions.

- [7] F. Garbe and R. Mycroft, *Hamilton cycles in hypergraphs below the Dirac threshold*, J. Combin. Theory Ser. B **133** (2018), 153–210. [4](#)
- [8] A. Georgakopoulos, J. Haslegrave, R. Montgomery, and B. Narayanan, *Spanning surfaces in 3-graphs*, J. Eur. Math. Soc. **24** (2022), no. 1. [5](#), [64](#)
- [9] A. Ghoulahouri, *A sufficient condition for the existence of a Hamiltonian circuit*, C. R. Acad. Sci. **251** (1960), no. 4, 495–497. [42](#)
- [10] R. J. Gould, *Recent advances on the Hamiltonian problem: Survey III*, Graphs Combin. **30** (2014), no. 1, 1–46. [3](#)
- [11] A. Halfpap, N. Lemons, and C. Palmer, *Positive co-degree density of hypergraphs*, J. Graph Theory **110** (2025), no. 2, 209–222. [4](#), [44](#)
- [12] A. Halfpap and V. Magnan, *Positive co-degree thresholds for spanning structures*, arXiv:2409.09185 (2024). [5](#)
- [13] H. Hàn and M. Schacht, *Dirac-type results for loose Hamilton cycles in uniform hypergraphs*, J. Combin. Theory Ser. B **100** (2010), no. 3, 332–346. [4](#)
- [14] J. Han and Y. Zhao, *Minimum codegree threshold for Hamilton ℓ -cycles in k -uniform hypergraphs*, J. Combin. Theory Ser. A **132** (2015), 194–223. [4](#)
- [15] F. Illingworth, R. Lang, A. Műyesser, O. Parczyk, and A. Sgueglia, *Spanning spheres in Dirac hypergraphs*, Combinatorica **45** (2025), no. 4, 1–29. [1](#), [5](#), [6](#), [8](#), [10](#), [17](#), [19](#), [30](#), [63](#), [64](#), [67](#), [68](#), [69](#)
- [16] G. Y. Katona and H. A. Kierstead, *Hamiltonian chains in hypergraphs*, J. Graph Theory **30** (1999), no. 3, 205–212. [3](#)
- [17] P. Keevash, D. Kühn, R. Mycroft, and D. Osthus, *Loose Hamilton cycles in hypergraphs*, Discrete Math. **311** (2011), no. 7, 544–559. [4](#)
- [18] P. Keevash and R. Mycroft, *A geometric theory for hypergraph matching*, vol. 233, American Mathematical Society, 2015. [26](#)
- [19] D. Kühn, R. Mycroft, and D. Osthus, *Hamilton ℓ -cycles in uniform hypergraphs*, J. Combin. Theory Ser. A **117** (2010), no. 7, 910–927. [3](#), [4](#)
- [20] D. Kühn and D. Osthus, *Loose Hamilton cycles in 3-uniform hypergraphs of high minimum degree*, J. Combin. Theory Ser. B **96** (2006), no. 6, 767–821. [4](#)
- [21] ———, *A survey on Hamilton cycles in directed graphs*, Europ. J. Combin. **33** (2012), 750–766. [3](#)
- [22] ———, *Hamilton cycles in graphs and hypergraphs: an extremal perspective*, arXiv:1402.4268 (2014). [3](#), [4](#)
- [23] R. Lang, *Tiling dense hypergraphs*, arXiv:2308.12281 (2023). [1](#), [5](#), [10](#), [68](#)
- [24] R. Lang and N. Sanhueza-Matamala, *A hypergraph bandwidth theorem*, arXiv:2412.14891 (2024). [10](#), [17](#), [67](#), [68](#)
- [25] G. Liu and X. Liu, *Hamiltonian paths and cycles in some 4-uniform hypergraphs*, Graphs Combin. **38** (2022), no. 4, 122. [3](#), [44](#)
- [26] X. Liu, J. Ma, T. Wang, and T. Zhu, *Uniquely colorable hypergraphs*, arXiv:2409.01654 (2024).

- 4
- [27] X. Liu, S. Ren, and J. Wang, *Positive codegree Andrásfai–Erdős–Sós theorem for the generalized triangle*, arXiv:2411.07090 (2024). 4
 - [28] A. Lo and K. Markström, *ℓ -degree Turán density*, SIAM J. Discrete Math. **28** (2014), no. 3, 1214–1225. 4
 - [29] A. Lo and E. Williams, *Towards an edge-coloured Corrádi–Hajnal theorem*, arXiv:2408.10651 (2024). 26
 - [30] C. McDiarmid, *On the method of bounded differences*, Surveys in combinatorics **141** (1989), 148–188. 14
 - [31] R. Mycroft and C. Zárate-Guerén, *Positive codegree thresholds for Hamilton cycles in hypergraphs*, arXiv:2505.11400 (2025). 1, 5, 6, 8, 10, 16, 21, 25, 26
 - [32] ———, *Positive codegree thresholds for perfect matchings in hypergraphs*, arXiv:2505.17981 (2025). 5, 6
 - [33] O. Pikhurko, *On the Limit of the Positive ℓ -Degree Turán Problem*, Electron. J. Combin. (2023), P3–25. 4
 - [34] V. Rödl and A. Ruciński, *Dirac-type questions for hypergraphs—a survey (or more problems for Endre to solve)*, An Irregular Mind: Szemerédi is 70, Springer, 2010, pp. 561–590. 3, 4
 - [35] V. Rödl, A. Ruciński, and E. Szemerédi, *A Dirac-type theorem for 3-uniform hypergraphs*, Combin. Probab. Comput. **15** (2006), no. 1-2, 229–251. 3
 - [36] ———, *Perfect matchings in uniform hypergraphs with large minimum degree*, European J. Combin. **27** (2006), no. 8, 1333–1349. 26
 - [37] ———, *An approximate Dirac-type theorem for k -uniform hypergraphs*, Combinatorica **28** (2008), no. 2, 229–260. 3
 - [38] ———, *Perfect matchings in large uniform hypergraphs with large minimum collective degree*, J. Combin. Theory Ser. A **116** (2009), no. 3, 613–636. 3
 - [39] ———, *Dirac-type conditions for Hamiltonian paths and cycles in 3-uniform hypergraphs*, Adv. Math. **227** (2011), no. 3, 1225–1299. 3, 4, 44
 - [40] M. Simonovits and E. Szemerédi, *Embedding graphs into larger graphs: results, methods, and problems*, Building Bridges II: Mathematics of László Lovász, Springer, 2020, pp. 445–592. 3
 - [41] S. Spiro, *On t -intersecting hypergraphs with minimum positive codegrees*, SIAM J. Discrete Math. **37** (2023), no. 2, 675–683. 4
 - [42] Z. Wu, *Positive Co-Degree Turán Number for C_5 and C_5^-* , J. Graph Theory **109** (2025), no. 1, 25–30. 4
 - [43] Y. Zhao, *Recent advances on Dirac-type problems for hypergraphs*, Recent trends in combinatorics (2016), 145–165. 4

A Proof of Lemma 4.1

In this section we sketch the proof of Lemma 4.1, restated here.

Lemma 4.1. *Let $1/n \ll 1/m_2 \ll 1/m_1 \ll 1/s, \gamma \ll \varepsilon, 1/k \leq 1/2$. Let \mathcal{G} and \mathcal{S} be families of n -vertex and s -vertex k -graphs such that for every $G \in \mathcal{G}$ and every $W \subseteq V(G)$ of size at most $2k$, there are at least $(1 - 1/s^2) \binom{n-|W|}{s-|W|}$ sets $U \subseteq V(G)$ of size s that contain W , such that $G[U]$ is a copy of a graph in \mathcal{S} . Suppose that $G \in \mathcal{G}$. Then there exists a sequence of s -vertex k -graphs F_1, \dots, F_r and a sequence of subgraphs $F_1^*, \dots, F_r^* \subseteq G$ such that the following hold for all $i, j \in [r]$ with addition modulo r :*

- F_i is a copy of a k -graph in \mathcal{S} ,
- F_i^* is a (γ, m_i^*) -nearly-regular blow-up of F_i for some $m_i^* \in [m_1, m_2]$,
- $V(F_1^*) \cup \dots \cup V(F_r^*) = V(G)$,
- $V(F_i^*) \cap V(F_j^*) = \emptyset$ unless $j \in \{i-1, i, i+1\}$,
- $V(F_i^*) \cap V(F_{i+1}^*)$ consists of exactly k vertices that induce an edge in F_i^* and F_{i+1}^* that is disjoint from the singleton parts of both blow-ups (if they exist).

As mentioned earlier, this lemma can be proved by following the arguments of the proof of [15, Lemma 2.1] and making a few small modifications. We note that Lang and Sanhueza-Matamala [24, Lemma 6.3] provide a somewhat more general lemma that could be used for our purpose instead, but both assumption and conclusion are slightly different to the ones in Lemma 4.1. We now state [15, Lemma 2.1].

Lemma A.1 (Lemma 2.1 in [15]). *Let $1/n \ll 1/m_2 \ll 1/m_1 \ll 1/s, \gamma \ll \varepsilon, 1/k \leq 1/2$. Let G be an n -vertex k -graph with no isolated vertices and with $\delta^*(G) \geq (1/2 + \varepsilon)n$, and let \mathcal{S} be the family of s -vertex k -graphs with no isolated vertices and with $\delta^*(S) \geq (1/2 + \varepsilon/2)s$. Then there exists a sequence of k -graphs F_1, \dots, F_r and a sequence of subgraphs $F_1^*, \dots, F_r^* \subseteq G$ such that the following hold for all $i, j \in [r]$:*

- F_i is a copy of a k -graph in \mathcal{S} ,
- F_i^* is a (γ, m_i^*) -nearly-regular blow-up of F_i for some $m_i^* \in [m_1, m_2]$,
- $V(F_1^*) \cup \dots \cup V(F_r^*) = V(G)$,
- $V(F_i^*) \cap V(F_j^*) = \emptyset$ unless $j \in \{i-1, i, i+1\}$,
- $V(F_i^*) \cap V(F_{i+1}^*)$, with $i \in [r-1]$, consists of exactly k vertices that induce an edge in F_i^* and F_{i+1}^* that is disjoint from the singleton parts of both blow-ups (if they exist).

For convenience, we make the following definition (inspired by terminology from [24]).

Definition A.2. Let G be an n -vertex k -graph and let \mathcal{S} be a family of s -vertex k -graphs. We say that G *satisfies \mathcal{S} r -robustly* if for every set W of r vertices in G , there are at least $(1 - 1/s^2)^{\binom{n-r}{s-r}}$ sets U of s vertices in G such that $W \subseteq U$ and $G[U] \in \mathcal{S}$.

Notice that Lemma 4.1 and Lemma A.1 differ on two points. First, instead of considering general families \mathcal{G} and \mathcal{S} of n -vertex and s -vertex k -graphs (where, in the terminology we just defined, every $G \in \mathcal{G}$ satisfies \mathcal{S} r -robustly for every $r \leq 2k$), the latter lemma is stated for specific choice \mathcal{G}_0 and \mathcal{S}_0 , namely where \mathcal{G}_0 is the family of n -vertex k -graphs with no isolated vertices and with minimum supported codegree at least $(1/2 + \varepsilon)n$, and \mathcal{S}_0 is the family of s -vertex k -graphs with no isolated vertices and with minimum supported codegree at least $(1/2 + \varepsilon/2)s$. Notice that, by [15, Lemma 3.8], it is indeed the case that every $G \in \mathcal{G}_0$ satisfies \mathcal{S} r -robustly for every $r \leq 2k$.² The second difference is more minor: while in Lemma 4.1 we require the sequence F_1^*, \dots, F_r^* to form a cyclic chain structure, namely that every two consecutive graphs intersect in one edge, etc., in Lemma A.1 they are only required to form a chain, meaning that the ‘linking edge’ between F_r^* and F_1^* is missing.

In short, in order to modify the proof of Lemma A.1 from [15] to prove Lemma 4.1, we can simply replace each mention of the specific properties \mathcal{G}_0 and \mathcal{S}_0 by general properties \mathcal{G} and \mathcal{S} such that every $G \in \mathcal{G}$ satisfies \mathcal{S} r -robustly for every $r \leq 2k$. Additionally, we simply add one more ‘link’ to the chain F_1^*, \dots, F_r^* in order to make it cyclic.

In the rest of the section, we give a more detailed proof sketch of Lemma 4.1. Let η and $m_{1.5}$ be constants such that $1/n \ll \eta \ll 1/m_2 \ll 1/m_{1.5} \ll m_1$. Fix families \mathcal{G} and \mathcal{S} as above and let $G \in \mathcal{G}$. We proceed in three steps. Here a *packing* in a graph G is a collection of pairwise vertex-disjoint subgraphs of G .

Finding an almost perfect packing of G by regular blow-ups of graphs in \mathcal{S} . In this step we find a collection of vertex-disjoint $(0, m_2)$ -regular blow-ups of graphs in \mathcal{S} within G that cover all but at most ηn vertices in G . That this is possible follows immediately by Lemma 4.4 from the breakthrough paper of Lang [23], providing very general conditions for the existence of a perfect packing of a graph G by copies of a hypergraph F . The lemma is applicable for a graph G and a family \mathcal{S} exactly when for every vertex $w \in V(G)$ there are at least $(1 - 1/s + \mu)^{\binom{n-1}{s-1}}$ many s -sets of vertices in G that contain w and induce a graph in \mathcal{S} , for some constant μ , and this is implied by G satisfying \mathcal{S} 1-robustly.

²In fact, they prove a stronger bound, namely that for every $G \in \mathcal{G}_0$ and set W of r vertices in G , where $r \leq 2k$, there are at least $(1 - e^{-\sqrt{s}})^{\binom{n-r}{s-r}}$ sets U of s vertices such that $W \subseteq U$ and $G[U] \in \mathcal{S}_0$. The weaker bound of $(1 - 1/s^2)^{\binom{n-r}{s-r}}$ suffices for the purpose of the proof of Lemma A.1.

Finding a perfect packing of G by nearly-regular blow-ups of graphs in \mathcal{S} . Here the idea is to start with an almost perfect packing as in the previous step, and, while there is an uncovered vertex, greedily finding a $(0, m_{1.5})$ -nearly-regular blow-up of a graph in \mathcal{S} , in such a way that these new blow-ups are pairwise vertex-disjoint and together they cover only a small proportion of each blow-up from the previous step. The desired perfect packing is obtained by taking the new nearly-regular blow-ups together with the old regular blow-ups but with the vertices from the new blow-ups removed. To accomplish this, we use that the Turán number of any k -partite k -graph is zero, implying that any n -vertex k -graph with $\Omega(n^k)$ edges contains a $(0, m_3)$ -regular blow-up of an edge when $m_3 \ll n$, and a pigeonhole argument, to conclude that every vertex v is contained in many $(0, m_{1.5})$ -nearly-regular blow-ups of a graph in \mathcal{S} . To find the desired perfect packing, we use that G contains \mathcal{S} 1-robustly, and this is the only property of \mathcal{G}_0 and \mathcal{S}_0 that is used in [15] for this step.

Finding a cyclic chain. Denote by B_1^*, \dots, B_r^* the perfect packing of G by nearly regular blow-ups of graphs in \mathcal{S} , found in the previous step. For each $i \in [r]$ we let e_{left}^i and e_{right}^i be two disjoint edges in B_i that avoid the vertex in B_i corresponding to the singleton part (if it exists), and let b_{left}^i and b_{right}^i be the blow-ups of e_{left}^i and e_{right}^i in B_i^* . Now, using [15, Lemma 4.3] and that G satisfies \mathcal{S} $2k$ -robustly (corresponding to the collection of copies of graphs in \mathcal{S} containing one copy of e_{right}^i and one copy of e_{left}^{i+1}), we find that there are many $(0, m_1)$ -regular blow-ups of graphs in \mathcal{S} that contain a copy of e_{right}^i and e_{left}^{i+1} . One can then conclude that we can pick such blow-ups C_i^* so that they are pairwise vertex-disjoint, and together cover only a small proportion of each B_i^* . Now, for each $i \in [r]$, remove from B_i^* all vertices participating in some C_j^* except for the vertices of one copy of e_{right}^i contained in C_i^* and one copy of e_{left}^i contained in C_{i-1}^* . We thus obtain a cyclic chain $B_1^*, C_1^*, \dots, B_r^*, C_r^*$. Again, all we need to know about \mathcal{G} and \mathcal{S} here is that every $G \in \mathcal{G}$ satisfies \mathcal{S} $2k$ -robustly. Also, note that the only difference between getting a chain and getting a cyclic chain is that for the latter we also need to define C_r^* , and for the former it suffices to find C_1^*, \dots, C_{r-1}^* .

B Improving Theorem 1.1

Recall that our main theorem asserts that every n -vertex k -graph with minimum supported codegree at least $(1 - 1/t)n - (k - 3)$ has a spanning ℓ -cycle. This condition is tight when $\ell = \frac{k-1}{2}$ (and t divides n and $n/t + 1$ is even) and is off by at most 1 in general, as can be seen in Example 2.3 and Example 2.1, respectively. We now briefly describe how to improve the bound in Theorem 2.6 and consequently in Theorem 1.1 to $\delta^*(G) \geq (1 - 1/t)n - (k - 2)$ when $t \geq \ell + 2$ or when $\ell = k - 1$ (i.e. we are seeking a Hamilton tight cycle). Notice that we always have $t \geq \ell + 1$ (see Observation 3.1). This leaves open the question of whether the bound $(1 - 1/t)n - (k - 2)$ is tight or off by 1 when $t = \ell + 1$ and $\ell \neq \frac{k-1}{2}$, or when $\ell = \frac{k-1}{2}$ and n does not satisfy certain divisibility conditions mentioned above.

The case $t \geq \ell + 2$. The crucial difference here, as alluded to in Remark 13.3, is that Observation 11.1 now instead becomes $d_G^1(S) \geq n - \lfloor n/t \rfloor - (i - 1)$. This implies that $\partial^2[A']$ has minimum degree at least x , not $x + 1$ as before, and we can no longer guarantee the existence of x cherries with leaves in A . Instead, we then modify Lemma 11.2 to instead obtain a linear forest F consisting of edges and cherries, both with leaves in A , such that $|F| \in \{3x, 3x + 1\}$. We extend a cherry $u_1u_2u_3$ of F to $u_1u_2v_1 \dots v_{t-2}u_3$ as before. In a similar manner, we also extend any edge $u_1u_2 \in E(F)$ to a sequence $uv_1 \dots v_{t-2}u_1u_2$ for $v_i \in B'$ and $u \in A$ (using a strategy similar to Claims 13.6 and 13.7) such that $uv_1 \dots v_{t-2}$ and $v_1 \dots v_{t-2}u_1u_2$ are both supported. Observe that this sequence does not support an extended ℓ -path according to Definition 11.5, but if we change the definition so as to additionally allow a sequence $v_0 \dots v_t$ to be considered to support an extended ℓ -path if $v_1 \dots v_t$ and $v_0 \dots v_{t-2}$ are supported, then the rest of the analysis carries through. Indeed, the crucial point is that if $U = u_1 \dots u_{2t}$ is a sequence of distinct vertices such that $u_1 \dots u_{2t-2}$ is a tight path and $u_{t+1} \dots u_{2t}$ is supported, then U supports an ℓ -path. Indeed, clearly the last t vertices in U form a supported set, so it remains to check that for every $i \in [0, \frac{t}{k-\ell} - 1]$, the subsequence $u_{i(k-\ell)+1} \dots u_{i(k-\ell)+k}$ is an edge. Notice that if $i \leq \frac{t}{k-\ell} - 1$ then $i(k-\ell) + k \leq t - (k-\ell) + k = t + \ell \leq 2t - 2$, using $\ell \leq t - 2$, so $u_{i(k-\ell)+1} \dots u_{i(k-\ell)+k}$ is contained in $u_1 \dots u_{2t-2}$, which is a tight path, and thus it is an edge. This allows us to form edges $b_1 \dots b_{t-1}P \in E(G^+)$ with $b_i \in B_i$ and $P \in \mathcal{P}$ corresponding to an “extended edge of F ” of the form $uv_1 \dots v_{t-2}u_1u_2$, as described above.

The case $\ell = k - 1$ (tight cycles). In this case the strategy we have outlined above will work by extending only the cherries in F but not the edges of F (and simply treating these as supported sets of size two in \mathcal{P}). The reason this modification works for tight cycles but fails in general is that it is important that the sequences in the family \mathcal{P} in Proposition 13.1 have length 1 (mod $k - \ell$) and indeed this holds for sequences of length 2 when $\ell = k - 1$, but fails otherwise.