



XI Latin and American Algorithms, Graphs and Optimization Symposium

An Approximate Blow-up Lemma for Sparse Hypergraphs

Peter Allen^{a,1}, Julia Böttcher^{a,1}, Eng Keat Hng^{a,1}, Jozef Skokan^{a,1}, Ewan Davies^{b,2}

^a*Department of Mathematics, London School of Economics and Political Science
Houghton Street, London WC2A 2AE, United Kingdom*

^b*Department of Computer Science, University of Colorado Boulder, Boulder CO, USA*

Abstract

We obtain an approximate sparse hypergraph version of the blow-up lemma, showing that partite hypergraphs with sufficient regularity of small subgraph counts behave as if they were complete partite for the purpose of embedding bounded degree hypergraphs.

© 2021 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the XI Latin and American Algorithms, Graphs and Optimization Symposium

Keywords: hypergraphs; embeddings; blow-up lemma.

1. Introduction

The blow-up lemma is a powerful tool developed by Komlós, Sárközy and Szemerédi [12] for embedding spanning subgraphs of bounded degree. Roughly speaking, it says that sufficiently regular graphs with no atypical vertices behave almost like complete partite graphs for the purpose of embedding bounded degree graphs. The regularity method of combining Szemerédi's regularity lemma [17] with the blow-up lemma has led to many significant breakthroughs in extremal graph theory.

There are at least two natural directions in which to generalise the above results: to sparse host graphs and to hypergraphs. Sparse analogues of the regularity lemma were independently established by Kohayakawa [10] and Rödl, while Allen, Böttcher, Hàn, Kohayakawa and Person [1] proved blow-up lemmas for sparse pseudorandom graphs and random graphs. In the direction of hypergraphs, there are various generalisations of the regularity lemma to hypergraphs (e.g. [15]) and Keevash [8] proved a hypergraph analogue of the blow-up lemma.

There has been significant interest in combining both directions. In the study of sparse graphs, it is somewhat standard to consider suitably well-behaved graphs satisfying a pseudorandomness condition. A well-known and extensively studied pseudorandomness condition is one known as jumbledness. This is a somewhat strong condition requiring significant control over edges between very small sets of vertices. For a variety of number theoretic appli-

¹ Email:p.d.allen@lse.ac.uk, j.boettcher@lse.ac.uk, e.hng@lse.ac.uk, j.skokan@lse.ac.uk

² Email:ewan.davies@gmail.com

cations however, it is desirable to obtain results with a weaker notion of pseudorandomness given in terms of small subgraph counts. This is sometimes referred to as linear forms conditions. Conlon, Fox and Zhao [5] proved a sparse weak regularity lemma for hypergraphs with linear forms conditions. This was recently extended by Allen, Davies and Skokan [3] to a sparse regularity lemma for hypergraphs in which they employed a concept of regularity known as octahedron-minimality.

Here we continue this line of research by proving an almost-spanning embedding lemma for hypergraphs with sufficient regularity of small subgraph counts, thereby establishing an approximate sparse hypergraph version of the blow-up lemma. In order to formally state the result we require further definitions.

A *complex* is a hypergraph whose edge set E is down-closed: if $e \in E$ and $f \subseteq e$ then $f \in E$. A k -*complex* is a complex whose edges are of size at most k . Our main focus is to embed a given k -complex H into another k -complex \mathcal{G} . We allow singletons and \emptyset to be edges in hypergraphs. This represents a departure from the norm, but it is consistent with the convention used by Allen, Davies and Skokan [3] and turns out to be convenient. For our purposes it will be convenient to consider \mathcal{G} as a weighted hypergraph. A *weighted hypergraph* on a vertex set V is a function from the power set of V to the non-negative real numbers. A *weighted- k -graph* is a weighted hypergraph \mathcal{G} with a weight function g such that $g(e) = 1$ for all $e \subseteq V(\mathcal{G})$ with $|e| > k$. The *weighted analogue* of a k -complex \mathcal{G} is the weighted- k -graph on $V(\mathcal{G})$ with weight function

$$g(e) = \begin{cases} 1 & \text{if } |e| > k \text{ or } e \in E(\mathcal{G}), \\ 0 & \text{otherwise.} \end{cases}$$

We use the calligraphic letters \mathcal{D} , \mathcal{G} and \mathcal{H} for weighted hypergraphs, and the corresponding lower case letters d , g and h for their weight functions. For positive real numbers w , x , y and z , we write $w = (x \pm y)z$ to mean $(x - y)z \leq w \leq (x + y)z$. Given $\ell \in \mathbb{N}$ and a complex H we write $H^{(\ell)}$ to mean the ℓ -uniform hypergraph on $V(H)$ with edge set $\{e \in E(H) : |e| = \ell\}$. For a graph H we write $\Delta(H)$ to mean the maximum degree of H .

We will focus on a *partite setting* as follows. For an index set J , we will consider a k -complex H , a weighted- k -graph \mathcal{G} , a partition $\mathcal{X} = \{X_j\}_{j \in J}$ of $V(H)$ indexed by J , and a partition $\mathcal{V} = \{V_j\}_{j \in J}$ of $V(\mathcal{G})$ indexed by J . We call the sets X_j and V_j the *parts* of H and \mathcal{G} respectively. We say that a set of vertices in H (resp. \mathcal{G}) is J -*partite* if it contains at most one vertex from each part of H (resp. \mathcal{G}). We say that H is J -*partite* if all its edges are J -partite and say that \mathcal{G} is J -partite to mean that \mathcal{G} is associated with a partition of $V(\mathcal{G})$ indexed by J . For $x \in X_j$ we write V_x to mean V_j . For a J -partite subset $S \subseteq V(H)$ we write $V_S = \prod_{x \in S} V_x$ for the collection of J -partite $|S|$ -subsets of $V(\mathcal{G})$ with vertices in $\bigcup_{x \in S} V_x$. For J -partite $S \subseteq V(H)$ the *index* of S is $i(S) := \{j \in J : S \cap X_j \neq \emptyset\}$. For $\kappa \geq 1$ we say that \mathcal{V} is κ -*balanced* if there exists $m \in \mathbb{N}$ such that $m \leq |V_j| \leq km$ for all $j \in J$. We say that \mathcal{V} and \mathcal{X} are *size-compatible* if $|V_j| = |X_j|$ for all $j \in J$. We need the following notion of the ‘rough’ structure of H . Let R be a complex on a finite set J and let H be a J -partite complex. For a partition $\mathcal{X} = \{X_j\}_{j \in J}$ of $V(H)$, we say that (H, \mathcal{X}) is an R -*partition* if the set $I \subseteq J$ is an edge of R whenever there are edges in H with index I .

A *homomorphism* from a complex H to a weighted hypergraph \mathcal{G} is a map $\phi : V(H) \rightarrow V(\mathcal{G})$ such that $|\phi(e)| = |e|$ for each $e \in E(H)$. The *weight* of ϕ is

$$\mathcal{G}(\phi) := \prod_{e \in E(H)} g(\phi(e)).$$

Given the partite setting, we say that a homomorphism from H to \mathcal{G} is a *partite homomorphism* if it maps each X_j into V_j . Define

$$\mathcal{G}(H) := \left(\prod_{j \in J} |V_j|^{-|X_j|} \right) \sum_{\phi} \prod_{e \in E(H)} g(\phi(e)),$$

where the sum is over all partite homomorphisms ϕ from H to \mathcal{G} . This is the expected weight of a uniformly random partite homomorphism from H to \mathcal{G} . Our main theorem reads as follows.

Theorem 1.1. *Given $k, \Delta, \Delta_R \in \mathbb{N}$, $\mu > 0$ and $\kappa \geq 1$, there exist $\eta > 0$ and $c \in \mathbb{N}$ such that for all finite sets J there exists $n_0 \in \mathbb{N}$ such that the following holds for all $n \geq n_0$. Let R be a k -complex on J . Let H and \mathcal{G} be J -partite k -complexes with κ -balanced vertex partitions $\mathcal{X} = \{X_j\}_{j \in J}$ and $\mathcal{V} = \{V_j\}_{j \in J}$ respectively such that \mathcal{G} has n vertices. Let \mathcal{D} be a weighted- k -graph on J with $d(\emptyset) = 1$, $d(\{j\}) = 1$ for all $j \in J$ and $d(e) > 0$ for all $e \subseteq J$. Suppose that the following hold.*

(BUL1) $\Delta(R^{(2)}) \leq \Delta_R$, $\Delta(H^{(2)}) \leq \Delta$ and (H, X) is an R -partition.

(BUL2) $|X_j| \leq (1 - \mu)|V_j|$ for each $j \in J$.

(BUL3) For all J -partite k -complexes F on at most c vertices we have

$$\mathcal{G}(F) = (1 \pm \eta)\mathcal{D}(F).$$

Then there is an embedding $\phi : V(H) \rightarrow V(\mathcal{G})$ of H into \mathcal{G} such that $\phi(x) \in V_x$ for each $x \in V(H)$.

To provide some context, we shall state the blow-up lemma of Komlós, Sárközy and Szemerédi [12] using our terminology. We shall provide some additional definitions for this statement. Note that 2-complexes are graphs. Let \mathcal{G} be a J -partite graph with vertex partition $\mathcal{V} = \{V_j\}_{j \in J}$. Let A and B be two disjoint nonempty subsets of $V(\mathcal{G})$. Let $a \in V(\mathcal{G}) \setminus B$. Write $e_{\mathcal{G}}(A, B)$ for the number of edges between A and B . Write $\deg_{\mathcal{G}}(a; B)$ for the number of neighbours of a in B within the graph \mathcal{G} . For a pair $i, j \in J$ such that $i \neq j$, we say that (V_i, V_j) is (η, δ) -super-regular if for all $A \subseteq V_i$ and $B \subseteq V_j$ satisfying $|A| > \eta|V_i|$ and $|B| > \eta|V_j|$, we have $e_{\mathcal{G}}(A, B) > \delta|A||B|$, and furthermore we have $\deg_{\mathcal{G}}(a; V_j) > \delta|B|$ for all $a \in A$ and $\deg_{\mathcal{G}}(b; V_i) > \delta|A|$ for all $b \in B$.

Theorem 1.2 (Blow-up Lemma [12]). *Given $\Delta \in \mathbb{N}$, $\delta > 0$ and a finite set J , there exists $\eta > 0$ such that the following holds for any graph R on J . Let H and \mathcal{G} be J -partite graphs with size-compatible vertex partitions $\mathcal{X} = \{X_j\}_{j \in J}$ and $\mathcal{V} = \{V_j\}_{j \in J}$ respectively. Suppose that*

(BL1) $\Delta(H) \leq \Delta$, (H, X) is an R -partition and $|V_j| > 0$ for all $j \in J$,

(BL2) (V_i, V_j) is (η, δ) -super-regular for each $ij \in E(R)$.

Then there is an embedding $\phi : V(H) \rightarrow V(\mathcal{G})$ of H into \mathcal{G} such that $\phi(x) \in V_x$ for each $x \in V(H)$.

The blow-up lemma of Komlós, Sárközy and Szemerédi [12] is a powerful tool for finding large subgraphs in dense graphs. In the following two subsections we shall present two applications of our result which involve finding large substructures in sparse hypergraphs.

1.1. Application: hypergraph Maker-Breaker games

In a $(1 : b)$ biased Maker-Breaker game, we are given a ground set X and a collection $\mathcal{F} \subset \mathcal{P}(X)$ of *winning sets*. Alternately, Maker claims up to 1, and then Breaker up to b , of the elements of X , until no unclaimed elements of X remain. Maker wins if she has claimed all the elements of any winning set (and perhaps some further elements), and Breaker wins otherwise. Since this is a finite game of perfect information, it is determined: one of the two players has a winning strategy with best play. The *threshold bias* b^* of the game is defined to be the smallest natural number b such that Breaker wins the $(1 : b)$ -game; assuming $\emptyset \notin \mathcal{F}$, this number is well-defined.

In particular, given k and n , if H is any k -graph, we can take X to be the edges of $K_n^{(k)}$ and \mathcal{F} to be the edge sets of all isomorphic copies of H in $K_n^{(k)}$. Thus Maker wins this H -game if the edges of $K_n^{(k)}$ she eventually claims contain an isomorphic copy of H .

The H -game is fairly well understood when H is a fixed 2-graph and n is large — in particular, Bednarska and Łuczak [4] determined the order of magnitude of the threshold bias (though even for $H = K_3$, where the threshold bias is $\Theta(n^{1/2})$, we do not know the constant multiplying $n^{1/2}$), and their methods extend to give a lower bound on the threshold bias also for fixed k -graphs. However when H depends on n , much less is known. The threshold bias for the Hamiltonicity game in graphs was determined by Krivelevich [13], and recently Liebenau and Nenadov [14] found asymptotically the threshold bias for the K_r -factor (that is, $\frac{n}{r}$ vertex-disjoint copies of K_r). There is also a general lower bound on the threshold bias for any bounded-degree graph on up to n vertices, which is a consequence of the Sparse Blow-up Lemma, due to Allen, Böttcher, Kohayakawa, Naves and Person [2].

As an application of Theorem 1.1, we prove the following general lower bound on the threshold bias for the H -game, where H is any almost-spanning bounded-degree k -graph.

Theorem 1.3. *Given integers $\Delta, k \geq 2$ and $\gamma > 0$ there exists a constant $\nu > 0$ such that the following holds for all sufficiently large n . Let H be any $(1 - \gamma)n$ -vertex k -graph, and let $b = n^\nu$. Then Maker wins the $(1 : b)$ H -game on $K_n^{(k)}$.*

The proof of this theorem is rather similar to the deduction of the $k = 2$ version of this result in [1]. Namely, we show that Maker has a randomised strategy that wins against any given Breaker strategy with positive probability. If Breaker had a winning strategy, then this would be impossible (Maker would always lose against Breaker's winning strategy) and hence Breaker does not have a winning strategy. Since the game is determined, it follows that Maker has a deterministic winning strategy.

Proof. Given Δ and k , we let c, η be returned by Theorem 1.1 for input $k, (k - 1)\Delta, \mu = \frac{1}{2}\gamma$ and $\kappa = 2$. Let $J = [(k - 1)\Delta + 1]$. We set $\nu = \frac{1}{2} \binom{c+k}{k}^{-1}$. Let $b = n^\nu$, and let $\varepsilon \ll \eta$ be sufficiently small for the following calculations. Observe that a strategy for Breaker is simply a rule which, given the edges claimed by respectively Maker and Breaker in their previous turns, outputs the edges that Breaker should claim in the current turn; in particular, to define a Breaker strategy we do not need to specify the winning sets of the game. Fix any Breaker strategy.

Fix any partition $V(K_n^{(k)}) = V_1 \cup \dots \cup V_{(k-1)\Delta+1}$ with parts of sizes differing by at most one. Let Q denote the complete partite k -graph with parts $V_1, \dots, V_{(k-1)\Delta+1}$. Maker's strategy is now the following. She randomly orders the edges of Q , and in her i th turn tries to claim the i th edge in her list; if this edge was previously claimed by Breaker, she claims no edge in that turn.

Let $\ell = \varepsilon \binom{n}{k} b^{-1}$. Let Γ be the graph of the first ℓ edges in Maker's list, and G the subgraph of edges which Maker successfully claimed. Observe that Γ is distributed as the uniform random ℓ -edge subgraph of Q . We claim that $e(G) \geq (1 - 3\varepsilon)\ell$. To see this, observe that in the i th turn, Maker chooses uniformly at random from the edges of Q which she has not previously chosen. Of these, at most $ib \leq \ell b = \varepsilon \binom{n}{k}$ were chosen by Breaker in previous rounds, and hence Maker's probability of picking a claimed edge is at most 2ε . The total number of edges Maker fails to claim is therefore stochastically dominated by $\text{Bin}(\ell, 2\varepsilon)$, which with high probability by Chernoff's inequality is at most $3\varepsilon\ell$.

Let $p = \ell/e(Q)$; let $p^+ = (1 + \varepsilon)p$, and let $p^- = (1 - \varepsilon)p$. Note that $p = \Theta(n^{-\nu})$. With high probability, when we choose edges of Q independently with probability p^- to obtain Q_{p^-} , we obtain less than ℓ edges; when we choose with probability p^+ to obtain Q_{p^+} , we obtain more than ℓ edges. There is then a standard coupling $Q_{p^-} \subseteq \Gamma \subseteq Q_{p^+}$ which succeeds with high probability. Namely, choose $e(Q_{p^-})$ and $e(Q_{p^+})$ from the binomial distributions $\text{Bin}(e(Q), p^-)$ and $\text{Bin}(e(Q), p^+)$ respectively, and fail if we do not obtain $e(Q_{p^-}) \leq \ell \leq e(Q_{p^+})$. If we do not fail, then choose Q_{p^-} by selecting $e(Q_{p^-})$ edges uniformly at random, Γ by adding $\ell - e(Q_{p^-})$ further edges uniformly at random, and Q_{p^+} by adding a further $e(Q_{p^+}) - \ell$ uniform random edges.

We view each of these four partite k -graphs as k -complexes by adding all partite edges of uniformity less than k , and let \mathcal{D} denote the weighted- k -graph on $[(k - 1)\Delta + 1]$ in which all edges of uniformity less than k have weight 1, and all edges of uniformity k have weight p . Let F be any k -complex as in (BUL3). By a minor modification of a theorem of Kim and Vu [9, Theorem 4.3.1], the number of embeddings of F into each of Q_{p^-} and Q_{p^+} is within a $(1 \pm \frac{1}{2}\eta)$ -factor of their expectations. Specifically, the theorem as stated there refers to a random subgraph of $K_n^{(k)}$; however all of the expectations computed in the proof there are upper bounds for the corresponding expectations in our setting, and in addition the expected number of F -copies in Q_{p^-} and Q_{p^+} is of the same order of magnitude as in a p^+ -random subgraph of $K_n^{(k)}$, so the proof applies in our setting also. Thus, by definition of \mathcal{D} , with high probability we have $Q_{p^-}(F), Q_{p^+}(F) = (1 \pm \frac{1}{2}\eta)\mathcal{D}(F)$ for all the k -complexes F of (BUL3), and so the same applies to Γ .

On the other hand, for any given such F , the number of embeddings of F using a given k -edge e in Q_{p^+} is stochastically dominated by the number of copies of F using e in the p^+ -random subgraph of $K_n^{(k)}$. As proved by Kim and Vu [9, Theorem 4.2.4], with high probability for all edges e this quantity is at most $(p^+)^s n^t$, where F has $s + 1$ k -edges and $t + k$ vertices.

Suppose that all the above mentioned likely events occur. Since G has at most $3\varepsilon\ell \leq 3\varepsilon p^+ n^k$ edges fewer than Γ , it has at most $3\varepsilon(p^+)^{s+1} n^{t+k}$ fewer embeddings of F than Γ . We claim that the embeddings of F into G make up almost all of the homomorphic copies of F counted by $G(F)$: to see this, observe that the number of homomorphic copies of F is $\Theta(n^{\nu(F)} p^{e_k(F)})$, where $e_k(F) \leq \binom{c}{k}$ counts the number of k -edges of F . By choice of ν , this is $\Omega(n^{\nu(F)-1/2})$, whereas trivially any homomorphic copy of F in G which is not an embedding uses at most $\nu(F) - 1$ vertices of G , and so there are at most $n^{\nu(F)-1}$ such. By choice of ε , and since n is sufficiently large, we see that $G(F) = (1 \pm \eta)\mathcal{D}(F)$, as required by (BUL3).

Now let H be any k -graph with vertex degree at most Δ and at most $(1 - \gamma)n$ vertices. We view this as a k -complex by taking the down-closure. Note that $\Delta(H^{(2)}) \leq (k - 1)\Delta$. By the Hajnal-Szemerédi Theorem [7], there is a partition of $V(H)$ into $(k - 1)\Delta + 1$ parts $X_1, \dots, X_{(k-1)\Delta+1}$ which differ in size by at most one and such that all the edges of $H^{(2)}$ (and so all the edges of H) are partite; this gives (BUL1). Since we chose $\mu = \frac{1}{2}\gamma$, we have $|X_i| \leq (1 - \mu)|V_i|$ for each i , verifying (BUL2).

By Theorem 1.1, we conclude that H is a subgraph of G , as desired. □

Note that this proof actually gives a slightly stronger conclusion than Theorem 1.3 claims: Maker actually ends up claiming a k -graph which contains not just any one H satisfying the conditions of the theorem, but all of them simultaneously (i.e. it is *universal*). To the best of our knowledge, previous to this result it was not even known that Maker has a winning strategy in the $(1 : b)$ H -game for any connected hypergraph H with $v(H) = \Theta(n)$ and any b growing with n (of course, b constant follows from Keevash [8]).

1.2. Application: size Ramsey numbers for bounded degree hypergraphs

The *size Ramsey number* $\hat{r}_\ell(H)$ of a k -graph H is defined to be the minimum of $e(\Gamma)$ over k -graphs Γ with the following property: however the edges of Γ are ℓ -coloured, one of the colour classes contains a subgraph isomorphic to H . We have the trivial bound $\hat{r}_\ell(H) \leq \binom{r_\ell(H)}{k}$, where $r_\ell(H)$ is the usual ℓ -colour Ramsey number, since a complete graph on $r_\ell(H)$ vertices by definition has the desired property. It was proved by Cooley, Fountoulakis, Kühn and Osthus [6] that when H is an n -vertex k -graph with maximum degree at most Δ , there is a constant C depending on k , ℓ and Δ such that $r_\ell(H) \leq Cn$, from which it follows $\hat{r}_\ell(H) = O(n^k)$.

For $k = 2$, i.e. graphs, Rödl and Szemerédi [16] proved that for some graphs H with $\Delta(H) = 3$ we have $\hat{r}_2(H) = \omega(n)$, and conjectured that there is $\varepsilon > 0$ such that for some H we have $\hat{r}_2(H) \geq n^{1+\varepsilon}$, and for all H with $\Delta(H) \leq \Delta$ we have $\hat{r}_\ell(H) = O(n^{2-\varepsilon})$, where ε depends on Δ only. The former conjecture remains open, but the latter was proved by Kohayakawa, Rödl, Schacht and Szemerédi [11], who showed it holds with any $\varepsilon < \frac{1}{\Delta}$. This bound, which is generally believed to be rather far from optimal, nevertheless remains the state of the art.

For $k \geq 3$, to the best of our knowledge there was no result improving on the trivial bound. We prove the following polynomial improvement.

Theorem 1.4. *For every k and Δ , there exists $\varepsilon > 0$ such that the following holds for each constant ℓ and all sufficiently large n . For any n -vertex k -graph H with $\Delta(H^{(2)}) \leq \Delta$, we have $\hat{r}_\ell(H) \leq n^{k-\varepsilon}$.*

Much of this proof is a fairly standard application of hypergraph regularity (though in a sparse setting) and we skip the details familiar from the dense setting. In this proof we refer to the notion of an (η^*, c^*) -THC graph for H , which we define in Definition 2.2; for the purposes of this proof, it is not essential to understand this definition as it is a black-box technical requirement for another theorem which will give us the counting condition we need.

Proof. Given k and Δ , let $\eta > 0$ and c be returned by Theorem 1.1 for input $k, \Delta, \mu = \frac{1}{2}$ and $\kappa = 2$. Let $c^* = \max(2c - 1, 4k^2 + k)$. Let $\rho > 0$ be sufficiently small such that [3, Lemma 4], with $p = Cn^{-\rho}$, can handle k -graphs with maximum degree and degeneracy at most c^* on at most c^* vertices. Let $p = Cn^{-\rho}$.

Given ℓ , we let $d_k = \frac{1}{2\ell}$. Let ε_k be small enough for [3, Theorem 6] with input $\eta_k = \frac{\eta}{2c}$ and d_k , and let $\varepsilon(t)$ be sufficiently small, for each $t \in \mathbb{N}$, to play the role of each of $\varepsilon_1, \dots, \varepsilon_{k-1}$ for densities $d_1, \dots, d_k \geq t^{-1}$. Let integers t_1 and n_0 , and $\eta^* > 0$, be returned by [3, Lemma 23] for input $k, q = \Delta + 1, s = \ell, \varepsilon_k$ and ε . Suppose that additionally η^* is small enough to play the role of η_0 in [3, Theorem 6] with input as above and with any $t \leq t_1$. Finally we choose $C = 2t_1$.

Let $\Gamma = G^{(k)}(Cn, p)$. We apply [3, Lemma 4], with input η^*, c^*, c^*, c^*, k , and with $J = [\Delta + 1]$, for each J -partite k -complex H on at most c vertices. Since there are boundedly many such (independent of n), the good event of [3, Lemma 4] w.h.p. holds for each such H simultaneously. That is, for any partition $V(\Gamma) = \bigcup_{j \in J} V_j$ with parts of size at least $n/\log n$, Γ is (η^*, c^*) -THC for H , with density graph \mathcal{P} whose edges of uniformity smaller than k all have weight 1, and whose k -edges have weight p . Suppose that this good event occurs. Note that in particular by choice of c^* we have $\Gamma(H) = (1 \pm \eta^*)p^{e(H)}$ for any k -graph H on at most $4k$ vertices, which is the condition required for [3, Lemma 23].

We claim that however $E(\Gamma)$ is ℓ -coloured, there is some colour class which contains H as a subgraph. Since $E(\Gamma) = O(n^{k-\rho})$, this proves the theorem.

Given a colouring of $E(\Gamma)$, we apply the regularity lemma [3, Lemma 23], with input as above, to this colouring. We select uniformly at random a collection of $r_\ell(K_{\Delta+1}^{(k)})$ clusters from the resulting returned family of partitions; between each pair of clusters, we select one of the 2-cells in the family of partitions uniformly at random, and so on up to $(k - 1)$ -cells.

Since the fraction of k -polyads in the family of partitions which are not ε_k -oct-regular in any one of the ℓ colours is at most $\ell\varepsilon_k$, the expected number of irregular polyads we selected randomly is at most $\binom{r_\ell(K_{\Delta+1}^{(k)})}{k} \ell\varepsilon_k < 1$, and in particular with positive probability we selected no irregular polyads. Fix such a choice. We now draw an auxiliary complete k -graph F whose vertices are the selected clusters, and where we colour any given k -edge with a majority colour appearing on the corresponding selected polyad. By definition of $r_\ell(K_{\Delta+1}^{(k)})$, there is a colour χ and a set $V_1, \dots, V_{\Delta+1}$ of clusters such that every k -polyad among these clusters has density at least $\frac{1}{\ell}$ in colour χ .

We now let $J = [\Delta + 1]$ and \mathcal{G} be the J -partite k -complex on $V_1, \dots, V_{\Delta+1}$ obtained by taking all the selected i -cells for each $2 \leq i \leq k - 1$, and the supported k -edges of Γ of colour χ . The good event of [3, Lemma 4] holding means that the partite subgraph of Γ on these clusters is (η^*, c^*) -THC for each J -partite H with at most c vertices. This, together with the oct-regularity and density of the cells of uniformity less than k (guaranteed by [3, Lemma 23]) and density at least d and ε_k -oct-regularity for uniformity k (by construction) is what we need to apply [3, Theorem 6], obtaining that for each J -partite k -complex F with at most c vertices, we have $\mathcal{G}(F) = (1 \pm \eta)\mathcal{D}(F)$ where $\mathcal{D}(F)$ is the density graph of \mathcal{G} (which is the same as the relative density graph mentioned in [3, Theorem 6] except that the weights differ on k -edges by a factor p).

Given H , we consider it as a k -complex by down-closure, and greedily J -partition vertices of H into parts $(X_i)_{i \in [\Delta+1]}$ such that no edge of $H^{(2)}$ is in any one part. Since each cluster V_i has at least $Cn/t_1 = 2n$ vertices, in particular $|X_i| \leq \frac{1}{2}|V_i|$ as required for (BUL2), and by the F -counting results we just established, Theorem 1.1 gives an embedding of H into \mathcal{G} . In particular, we find an isomorphic copy of the k -graph H in the colour χ edges of Γ . \square

Again, note that this proof (as with the proof of Kohayakawa, Rödl, Schacht and Szemerédi [11]) actually gives a stronger conclusion: the graph Γ has a colour class which contains simultaneously all n -vertex k -graphs H with $\Delta(H^{(2)}) \leq \Delta$. This property is called *partition universality*.

2. Proving Theorem 1.1

2.1. Pseudorandomness for weighted hypergraphs

While our main theorem is stated in terms of small subgraph counts, our proof method involves a related notion of pseudorandomness known as typically hereditary counting, which we define in Definition 2.2. This was introduced by Allen, Davies and Skokan [3]. We need some additional definitions to give a formal definition of this notion. We require the definition of the link graph of a vertex v in a weighted hypergraph \mathcal{G} . Let J be an index set and let \mathcal{G} be a weighted hypergraph with vertex sets $\{V_j\}_{j \in J}$. Let $i \in J$. For a vertex $v \in V_i$ we let \mathcal{G}_v be the weighted hypergraph on the vertex sets $\{V_j\}_{j \in J \setminus \{i\}}$ with weight function g_v , defined as follows. For $f \subseteq J \setminus \{i\}$ and $e \in V_f$, we set

$$g_v(e) = g(e) \cdot g(e \cup \{v\}).$$

We call \mathcal{G}_v the *link graph* of the vertex v in the weighted hypergraph \mathcal{G} . Since we are working in the weighted setting, we need to work with the sum of the weights of the vertices in a set instead of the size of that set. In particular, it will be convenient to work with a normalised version of this notion. Let \mathcal{G} be a weighted hypergraph with vertex sets $\{V_j\}_{j \in J}$. For a subset $U \subseteq V_j$ we write $\|U\|_{\mathcal{G}} := |V_j|^{-1} \sum_{u \in U} g(u)$.

It will be convenient to work with partite homomorphisms which send exactly one vertex of H to each part of \mathcal{G} . The following ‘copying process’ enables a reduction to this setting from the general partite setting.

Definition 2.1 (Standard construction). *Let J be an index set, H be a k -complex with its vertex set partitioned into $\mathcal{X} = \{X_j\}_{j \in J}$ and \mathcal{G} be a weighted- k -graph with its vertex set partitioned into $\mathcal{V} = \{V_j\}_{j \in J}$. The standard construction for $(\mathcal{G}, \mathcal{V})$ with respect to (H, \mathcal{X}) is a $V(H)$ -partite weighted- k -graph \mathcal{G}' with vertex sets $\{V'_x\}_{x \in V(H)}$ where for each*

$x \in V(H)$ the set V'_x is a copy of the set V_j such that $x \in X_j$, and where for each $f \subseteq V(H)$ and each $e' \in V'_f$ we define

$$g'(e') = \begin{cases} g(e) & \text{if } f \in E(H), \\ 1 & \text{if } f \notin E(H), \end{cases}$$

where e is the natural projection of e' to $V(\mathcal{G})$. We will omit mention of the complex H and the partitions \mathcal{X} and \mathcal{V} when they are clear from context.

Here we state the definition of typically hereditary counting.

Definition 2.2 (Typically hereditary counting (THC) [3]). *Given $k \in \mathbb{N}$, a vertex set J endowed with a linear order τ and a weighted- k -graph \mathcal{D} on J , we say that the J -partite weighted- k -graph \mathcal{G} is an (η, c) -THC graph if the following two properties hold.*

(THC1) *For each J -partite k -complex R with at most c vertices, we have*

$$\mathcal{G}(R) = (1 \pm \nu(R)\eta) \frac{g(\emptyset)}{d(\emptyset)} \mathcal{D}(R).$$

(THC2) *If $|J| \geq 2$ and x is the first vertex of J , there is a set $V'_x \subseteq V_x$ with $\|V'_x\|_{\mathcal{G}} \geq (1 - \eta)\|V_x\|_{\mathcal{G}}$ such that for each $v \in V'_x$ the graph \mathcal{G}_v is an (η, c) -THC graph on $J \setminus \{x\}$ with density graph \mathcal{D}_x .*

We say that \mathcal{D} is a density weighted- k -graph of \mathcal{G} .

Let H be a J -partite k -complex with vertex partition \mathcal{X} , \mathcal{G} be a J -partite weighted- k -graph with vertex partition \mathcal{V} and \mathcal{D} be a density weighted- k -graph for \mathcal{G} . We say that \mathcal{G} is an (η, c) -THC graph for H if the standard construction for $(\mathcal{G}, \mathcal{V})$ with respect to (H, \mathcal{X}) is an (η, c) -THC graph with the standard construction for \mathcal{D} with respect to (H, \mathcal{X}) as its density weighted- k -graph.

In broad terms, being THC means having the following two properties. Firstly, we can accurately count copies of small complexes and find that these counts are close to their expected values in a ‘random graph’ setting. Secondly, we are also able to accurately count in the link graph \mathcal{G}_v for most vertices v , in ‘typical’ link graphs of \mathcal{G}_v , and so on. In particular, THC goes beyond ‘we can accurately count all small subgraphs’ in that we may take ‘typical’ links a very large number of times while still preserving THC. Quantitatively, note that the parameters of THC remain unchanged in (THC2). This hereditary property turns out to be extremely useful for our proof of Theorem 1.1 as it allows us to maintain this pseudorandomness condition in the course of a vertex-by-vertex embedding procedure which embeds vertices in a ‘typical’ and ‘valid’ manner.

It may seem that being THC is a strictly stronger condition than having small subgraph counts alone, because we also require a hereditary property in the form of (THC2). It turns out that these two notions of pseudorandomness are highly related. The following theorem of Allen, Davies and Skokan [3] tells us that (weighted) hypergraphs with certain small subgraph counts are in fact THC graphs. It implies a qualitative equivalence between the two notions of pseudorandomness and quantifies the additional counts required to obtain the hereditary property (THC2). Note that the conclusion does not specify the linear order on the indexing set which forms part of the definition of THC. In fact, the proof of Theorem 2.3 by Allen, Davies and Skokan [3] is valid for any choice of linear order on $V(H)$.

Theorem 2.3 ([3]). *For all $k, \Delta \geq 2, c \geq \Delta + 2$ and $0 < \eta < 1/2$, there exists $\eta_0 > 0$ such that whenever $0 < \eta' < \eta_0$ the following holds. Let H be a k -complex with $\Delta(H^{(2)}) \leq \Delta$. Suppose that \mathcal{G} is a $V(H)$ -partite weighted- k -graph with vertex sets $\{V_x\}_{x \in V(H)}$ which is identically 1 on any V_f such that $f \notin E(H)$, and \mathcal{D} is a density weighted- k -graph on $V(H)$ such that for all $V(H)$ -partite k -complexes F on at most $(\Delta + 2)c$ vertices we have*

$$\mathcal{G}(F) = (1 \pm \eta') \frac{g(\emptyset)}{d(\emptyset)} \mathcal{D}(F).$$

Then \mathcal{G} is an (η, c) -THC graph.

The qualitative equivalence between being THC and having small subgraph counts plays a central role in our proof of Theorem 1.1 for technical reasons. While the recursive nature of the hereditary property of THC is extremely useful for analysing the evolution of a vertex-by-vertex embedding procedure, it also makes direct verification of the property of being THC rather difficult. Conversely, small subgraph counts are usually easier to obtain, but they are harder to analyse in relation to a vertex-by-vertex embedding procedure. Furthermore, the interplay between these two notions allows us to extract additional information about subgraph counts.

2.2. Proof outline for Theorem 1.1

Our proof draws on ideas used in the proofs of the graph blow-up lemma in [12] and of sparse graph blow-up lemmas in [1]. We also incorporate ideas from [3] in relation to the notions of pseudorandomness previously discussed.

We want to embed a k -complex H with a partition X into a sparse pseudorandom k -complex \mathcal{G} with a compatible partition \mathcal{V} , where the notion of sparse pseudorandomness is that the small subgraph counts in \mathcal{G} are roughly the same as those obtained if \mathcal{G} were truly random, up to a small constant relative error. Theorem 2.3 tells us that \mathcal{G} is a THC graph for H . Our proof strategy involves embedding H into \mathcal{G} by using a random greedy algorithm similar to those of Komlós, Sárközy and Szemerédi [12] and Allen, Böttcher, Hàn, Kohayakawa and Person [1]. The following definitions will be useful. Let ϕ be a partial partite homomorphism of H into \mathcal{G} . That is, a homomorphism with domain $\text{Dom}(\phi) \subseteq V(H)$ such that $\phi(x) \in V_x$ for all $x \in \text{Dom}(\phi)$. Let the *candidate graph* \mathcal{G}_ϕ be the weighted- k -graph obtained from \mathcal{G} by repeatedly taking links of vertices in $\text{Im}(\phi)$. For each $x \in V(H) \setminus \text{Dom}(\phi)$ we define the *candidate set* $C_\phi(x) := \{v \in V_x : g_\phi(v) = 1\}$. Since \mathcal{G}_ϕ is a $\{0, 1\}$ -weighted hypergraph, we have $|C_\phi(x)| = \mathcal{G}_\phi(x)|V_x|$. Our algorithm will generate a sequence $\phi_0, \dots, \phi_{|V(H)|}$ of partial partite homomorphisms of H into \mathcal{G} and we will use a subscript t to mean that the relevant object is with reference to ϕ_t . For example, we write $C_t(x)$ to mean the candidate set of x with reference to ϕ_t .

Our random greedy algorithm proceeds as follows. We have k -complexes H and \mathcal{G} which satisfy the conditions of Theorem 1.1 and we select an arbitrary linear order τ on $V(H)$. We also partition the vertices of \mathcal{G} into two parts at random, a small set V^q and the remainder V^{main} . We choose vertices into V^q independently with probability which is significantly larger than η but significantly smaller than μ . We embed the elements of $V(H)$ vertex-by-vertex according to the order given by τ . At time $\tau(x)$, when we need to embed x , we look at the vertices of $C_{\tau(x)-1}(x) \cap V^{\text{main}}$. Provided that not too many of them have been previously used to embed a vertex, we pick a vertex uniformly at random from the unused vertices, which preserves the THC property, and embed x there. If on the other hand exceptionally many vertices of $C_{\tau(x)-1}(x) \cap V^{\text{main}}$ have been used previously, we pick instead a vertex of $C_{\tau(x)-1}(x) \cap V^q$ which preserves the THC property uniformly at random, and embed x there. In this latter case we say x goes to the queue.

This is much the same as the approach of [1] so far, although making it work in hypergraphs is rather more complicated, and the THC property is critical here. It is the THC property which guarantees for us that any vertex which does not go to the queue can be embedded (and we choose its image from a reasonably large set of vertices). As in [1], it is not too hard to prove a Queue Lemma, which states that with high probability the number of vertices in any X_i which go to the queue is much smaller than $|V_i \cap V^q|$.

What is much more difficult than in [1] is to show that with high probability we succeed in embedding all the vertices which go to the queue. Let us first explain how it is that we could fail. We have some $x \in X_i$ which goes to the queue, and which we therefore want to embed to $C_x := C_{\tau(x)-1}(x) \cap V^q$. This set is guaranteed to be quite well behaved, however it is (because we are working in sparse hypergraphs) very tiny compared to the set of all vertices in X_i which go to the queue. If before $\tau(x)$ we have already used too many (say more than half) of the vertices of C_x , then our strategy will fail.

What we therefore need to do is to argue that for each $x' \in X_i$ which goes to the queue, it is rather unlikely that we will embed x' to C_x . What we would ideally like is to say that the chance of embedding x' to C_x should be about $\frac{|C_x|}{|V_i \cap V^q|}$, i.e. the probability we would get if x' were embedded to $V_i \cap V^q$ uniformly at random. Provided that we can establish such a bound, then the Queue Lemma together with a standard martingale concentration argument will tell us that our strategy is very unlikely to fail.

This too is rather similar to the approach of [1]. However, establishing the desired probability bound is much harder in our setting. In [1], a rather strong quasirandomness condition is used at this point, which says that we cannot have exceptionally many 2-edges between pairs of sets in \mathcal{G} , even when the pairs of sets are much smaller than linear-sized. Even in the $k = 2$ case of Theorem 1.1, such a statement need not be true; we can only control density between linear-sized pairs of sets.

We now outline how we can provide an upper bound for the probability of embedding x' to C_x . Suppose (for purposes of illustration) that we are at some time τ such that we already embedded all neighbours of x (so C_x is fixed), and we also embedded all vertices whose H -distance from x' is exactly 3, but none of the vertices closer to x' . Let F denote the subgraph of H induced by the set of (at most $(\Delta + 1)^3$) vertices at distance 3 or smaller to x' in H ; we say the *roots* of F are the vertices at distance exactly 3 from x' , which are already embedded. When we go on to embed

the remaining vertices of F , it turns out that we will choose a rooted J -partite copy of F in \mathcal{G} more or less uniformly at random. Again for illustration, suppose we in fact choose one uniformly at random. What we would like to know, then, is the fraction of rooted F -copies in \mathcal{G} such that the vertex x' is in C_x ; this is the probability of embedding x' to C_x .

The THC property tells us (up to a small relative error) what the total number of rooted F -copies is: this is precisely what (THC1) with F' consisting of F with the roots removed gives. However it cannot tell us how many of these copies put x' into C_x . Similarly, it follows from the THC property with some work (because C_x is a candidate set) that we can also count copies of any graph F^2 on up to c vertices, some of which are required to be embedded to C_x and the rest of which are embedded to specified parts of \mathcal{G} . However we cannot use this to count rooted copies of F with x' in C_x , because we cannot specify that the roots of F have to be embedded to any subsets of parts of \mathcal{G} , let alone to specific vertices.

The solution here is the following. We let F_1 denote the subgraph of F induced by vertices at distance 1, 2 or 3 from x' , with again the vertices at distance 3 from x' being the roots. We let F_2 denote the subgraph of F induced by x' together with vertices at distance 1 from x' . Note that there can be edges of F which are contained in both F_1 and F_2 , namely those all of whose vertices are at distance 1 from x' . But every edge of F is contained in at least one of F_1 and F_2 . We say the vertices of F at distance exactly 1 from x' are the *middle* of F . See Figure 2.2 for an example (with uniformity 2).




Figure 1. The graphs F , F_1^* and F_2^* , with roots grey diamonds and middle solid diamonds.

Fix an embedding ψ of the middle of F . Then the number of ways to extend this to a rooted embedding of F with x' embedded to C_x is given by $c(\psi)d(\psi)$, where $c(\psi)$ is the number of rooted embeddings of F_1 whose middle is embedded according to ψ ; and $d(\psi)$ is the number of embeddings of F_2 such that x' is mapped to C_x and the middle is embedded according to ψ . It follows that the quantity we are trying to upper bound — the number of rooted embeddings of F such that x' is embedded to C_x — is given by

$$\sum_{\psi} c(\psi)d(\psi) \leq \sqrt{(\sum_{\psi} c(\psi))^2 (\sum_{\psi} d(\psi))^2}, \tag{1}$$

where the inequality is the Cauchy-Schwarz inequality. Now, $(\sum_{\psi} c(\psi))^2$ counts pairs of rooted copies of F_1 with middle ψ , summed over ψ . This is the same thing as counting rooted homomorphic copies of F_1^* , where F_1^* is obtained from F_1 by duplicating all the vertices at distance exactly 2 from x' and all the edges containing them. Similarly, letting F_2^* be obtained from F_2 by duplicating the vertex x' and all the edges it contains, $(\sum_{\psi} d(\psi))^2$ counts homomorphic copies of F_2^* in which x' and its duplicate are embedded to C_x .

Good estimates on these two homomorphism counts are things we can obtain from the THC property. For the former, rooted copies of F_1^* , this is given directly from (THC1) with the graph F_1^* with the roots removed; again, for F_2^* we need to do a bit more work, but as above it follows from the fact that C_x is a candidate set. These two observations plus (1) gives us an upper bound on the desired number of rooted copies of F with x' embedded to C_x . This upper bound turns out to be roughly what one would expect if the edges were distributed at random, and in particular what we obtain is that the probability of embedding x' to C_x is roughly the same as if x' were uniformly embedded to $V_i \cap V^q$, as we wanted.

We should stress that the above sketch is an oversimplification. We will not generally know what C_x is before embedding vertices close to x' . We cannot generally assume that all vertices at distance 3 from x' are embedded before any at distance 2. And (which we skipped entirely in the above) we might be worried that the event of a vertex x' going to the queue is correlated with x' being exceptionally likely to be embedded to C_x (so that we cannot simply multiply the probability we calculated above by the number of vertices which go to the queue). But all of these caveats

can be dealt with: what we finally conclude is that for each x , it is very unlikely that more than half of $C_{\tau(x)-1}(x) \cap V^q$ is used before time $\tau(x)$, and hence the greedy embedding succeeds with high probability.

References

- [1] Allen, P., Böttcher, J., Hàn, H., Kohayakawa, Y., Person, Y., 2016. Blow-up lemmas for sparse graphs. [arXiv:1612.00622](#). arXiv:1612.00622.
- [2] Allen, P., Böttcher, J., Kohayakawa, Y., Naves, H., Person, Y., 2017. Making spanning graphs. [arXiv:1711.05311](#). arXiv:1711.05311.
- [3] Allen, P., Davies, E., Skokan, J., 2019. Regularity inheritance in hypergraphs. [arXiv:1901.05955](#). arXiv:1901.05955.
- [4] Bednarska, M., Łuczak, T., 2000. Biased positional games for which random strategies are nearly optimal. *Combinatorica* 20, 477–488. doi:[10.1007/s004930070002](#).
- [5] Conlon, D., Fox, J., Zhao, Y., 2015. A relative Szemerédi theorem. *Geom. Funct. Anal.* 25, 733–762. doi:[10.1007/s00039-015-0324-9](#).
- [6] Cooley, O., Fountoulakis, N., Kühn, D., Osthus, D., 2009. Embeddings and Ramsey numbers of sparse k -uniform hypergraphs. *Combinatorica* 29, 263–297. doi:[10.1007/s00493-009-2356-y](#).
- [7] Hajnal, A., Szemerédi, E., 1970. Proof of a conjecture of P. Erdős, in: *Combinatorial theory and its applications, II (Proc. Colloq., Balatonfüred, 1969)*, North-Holland, Amsterdam. pp. 601–623.
- [8] Keevash, P., 2011. A hypergraph blow-up lemma. *Random Structures Algorithms* 39, 275–376. doi:[10.1002/rsa.20362](#).
- [9] Kim, J.H., Vu, V.H., 2000. Concentration of multivariate polynomials and its applications. *Combinatorica* 20, 417–434.
- [10] Kohayakawa, Y., 1997. Szemerédi’s regularity lemma for sparse graphs, in: *Foundations of computational mathematics (Rio de Janeiro, 1997)*. Springer, Berlin, pp. 216–230.
- [11] Kohayakawa, Y., Rödl, V., Schacht, M., Szemerédi, E., 2011. Sparse partition universal graphs for graphs of bounded degree. *Advances in Mathematics* 226, 5041–5065.
- [12] Komlós, J., Sárközy, G.N., Szemerédi, E., 1997. Blow-up lemma. *Combinatorica* 17, 109–123. doi:[10.1007/BF01196135](#).
- [13] Krivelevich, M., 2011. The critical bias for the Hamiltonicity game is $(1 + o(1))n / \ln n$. *J. Amer. Math. Soc.* 24, 125–131.
- [14] Liebenau, A., Nenadov, R., 2020. The threshold bias of the clique-factor game. [arXiv:2002.02578](#). arXiv:2002.02578.
- [15] Rödl, V., Skokan, J., 2004. Regularity lemma for k -uniform hypergraphs. *Random Structures Algorithms* 25, 1–42. doi:[10.1002/rsa.20017](#).
- [16] Rödl, V., Szemerédi, E., 2000. On size Ramsey numbers of graphs with bounded degree. *Combinatorica* 20, 257–262.
- [17] Szemerédi, E., 1978. Regular partitions of graphs, in: *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*. CNRS, Paris. volume 260 of *Colloq. Internat. CNRS*, pp. 399–401.