



Regular slices for hypergraphs

Peter Allen,^a Julia Böttcher,^a Oliver Cooley,^b
Richard Mycroft^c

^a *Department of Mathematics, London School of Economics,
Houghton Street, London WC2A 2AE, U.K.*

^b *Graz University of Technology, Institute of Optimization and Discrete
Mathematics, Steyrergasse 30, 8010 Graz, Austria.*

^c *School of Mathematics, University of Birmingham,
Edgbaston, Birmingham B15 2TT, U.K.*

Abstract

We present a ‘Regular Slice Lemma’ which, given a k -graph \mathcal{G} , returns a regular $(k - 1)$ -complex \mathcal{J} with respect to which \mathcal{G} has useful regularity properties. We believe that many arguments in extremal hypergraph theory are made considerably simpler by using this lemma rather than existing forms of the Strong Hypergraph Regularity Lemma, and advocate its use for this reason.

Keywords: Hypergraphs, Regularity Lemma.

¹ Email: p.d.allen@lse.ac.uk, j.boettcher@lse.ac.uk, cooley@math.tugraz.at, r.mycroft@bham.ac.uk

² PA was partially supported by FAPESP (Proc. 2010/09555-7); JB by FAPESP (Proc. 2009/17831-7); PA and JB by CNPq (Proc. 484154/2010-9); OC by the DFG (TA 309/2-2); The cooperation of the authors was supported by a joint CAPES-DAAD project (415/ppp-probral/po/D08/11629, Proj. no. 333/09). The authors are grateful to NUMEC/USP, Núcleo de Modelagem Estocástica e Complexidade of the University of São Paulo, and Project MaCLinC/USP, for supporting this research.

1 Introduction

The Szemerédi Regularity Lemma [10] is a powerful tool in extremal graph theory; a great number of advances in this area over recent decades either rely on, or at least were inspired by, the Regularity Lemma. Finding the right extension of this result for uniform hypergraphs turned out to be a challenging endeavour, which culminated in the proof of the Strong Hypergraph Regularity Lemma together with a corresponding Counting Lemma (see [4,6,7,8,9]), which provide an analogous machinery for extremal problems in hypergraphs. The difficulty with these tools is their technical intricacy, which leads to significant additional complexity in applications of the regularity method in extremal hypergraph theory.

We argue that in many cases much of this complexity can be avoided by using a structure which we call a *regular slice* instead of the more complicated structure returned by the Strong Hypergraph Regularity Lemma. Our main result is a Regular Slice Lemma, derived from the Strong Hypergraph Regularity Lemma, which asserts the existence of regular slices which inherit enough structure from the original hypergraph to be useful for embedding problems.

2 Regular Complexes

In this section we give key definitions, including the notion of a regular k -complex (this idea plays a key role in describing regularity for hypergraphs). For a more expository introduction to regular k -complexes we recommend [4].

A *hypergraph* $\mathcal{H} = (V, E)$ consists of a vertex set V and an edge set E , where each edge $e \in E$ is a subset of V . We often identify a hypergraph with its edge set, writing $e \in \mathcal{H}$ to mean $e \in E$ and writing $|\mathcal{H}|$ for the number of edges of \mathcal{H} . Similarly, given two hypergraphs \mathcal{G} and \mathcal{H} with vertex set V , we write $\mathcal{G} \cup \mathcal{H}$ for the hypergraph on V with edge set $E(\mathcal{G}) \cup E(\mathcal{H})$.

We say that a hypergraph \mathcal{H} is *k -uniform* if every edge has size k , and abbreviate ‘ k -uniform hypergraph’ to *k -graph*. Also, we say that \mathcal{H} is a *k -complex* if every edge of \mathcal{H} has size at most k , and moreover for any $e \in \mathcal{H}$ and $e' \subseteq e$ we have $e' \in \mathcal{H}$. We informally think of a k -complex \mathcal{H} as having ‘layers’, where ‘layer’ i is the i -graph formed by edges of \mathcal{H} of size i . Given a $(k-1)$ -complex \mathcal{H} with vertex set V , we say that a k -set $S \subseteq V$ is *supported on* \mathcal{H} if $S' \in \mathcal{H}$ for any $S' \subsetneq S$, and similarly that a k -graph \mathcal{G} on V is *supported on* \mathcal{H} if every edge of \mathcal{G} is supported on \mathcal{H} . So, informally, in a k -complex \mathcal{H} the i th layer is supported on the $(i-1)$ -complex formed by the edges of all lower layers. Finally, given a vertex set V and a partition \mathcal{P} of V , we say that

a set $S \subseteq V$ is \mathcal{P} -partite if S contains at most one vertex from any part of \mathcal{P} , and we say that a hypergraph \mathcal{H} with vertex set V is \mathcal{P} -partite if every edge of \mathcal{H} is \mathcal{P} -partite.

For the rest of this section we fix a vertex set V and a partition \mathcal{P} of V into parts V_1, \dots, V_t , which we call *clusters*. Let \mathcal{H} be a k -complex on V . For any $\ell \geq 2$ and any $A \in \binom{[t]}{\ell}$, we define $V_A := \bigcup_{i \in A} V_i$, and write \mathcal{P}_A for the partition of V_A inherited from \mathcal{P} (so \mathcal{P}_A has $|A| = \ell$ parts). Similarly, we write \mathcal{H}_A for the \mathcal{P}_A -partite ℓ -graph with vertex set V_A and whose edges are precisely the edges of \mathcal{H} which have ℓ vertices, one in each part of \mathcal{P}_A . We also denote by \mathcal{H}_A^* the \mathcal{P}_A -partite ℓ -graph with vertex set V_A whose edges are precisely those \mathcal{P}_A -partite sets $S \in \binom{V_A}{\ell}$ such that every proper subset $S' \subsetneq S$ is an edge of \mathcal{H} . We then define the *relative density of \mathcal{H} at A* to be

$$d_A(\mathcal{H}) := \frac{|\mathcal{H}_A \cap \mathcal{H}_A^*|}{|\mathcal{H}_A^*|}$$

if $|\mathcal{H}_A^*| > 0$, so $d_A(\mathcal{H})$ is the proportion of \mathcal{P}_A -partite sets $S \in \binom{V_A}{\ell}$ which could possibly be edges of \mathcal{H} (in the sense that S is supported on the ‘lower levels’ of \mathcal{H}) which are in fact edges of \mathcal{H} . If instead $|\mathcal{H}_A^*| = 0$ then for convenience we define $d_A(\mathcal{H}) := 0$. In the same way, if $\mathbf{Q} := (\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_r)$ is a collection of r not-necessarily-disjoint subcomplexes of \mathcal{H} , we define

$$d_A(\mathcal{H}|\mathbf{Q}) := \frac{|\mathcal{H}_A \cap \bigcup_{i \in [r]} (\mathcal{J}_i)_A^*|}{|\bigcup_{i \in [r]} (\mathcal{J}_i)_A^*|}$$

if $|\bigcup_{i \in [r]} (\mathcal{J}_i)_A^*| > 0$, and take $d_A(\mathcal{H}|\mathbf{Q}) := 0$ otherwise. We say that \mathcal{H} is (d_i, ε, r) -regular at A if we have $d_A(\mathcal{H}|\mathbf{Q}) = d_i \pm \varepsilon$ for every r -set \mathbf{Q} of subcomplexes of \mathcal{H} such that $|\bigcup_{i \in [r]} (\mathcal{J}_i)_A^*| > \varepsilon |\mathcal{H}_A^*|$. We refer to $(d_i, \varepsilon, 1)$ -regularity simply as (d_i, ε) -regularity. Moreover, for constants d_2, \dots, d_k we say that \mathcal{H} is $(d_k, \dots, d_2, \varepsilon_k, \varepsilon, r)$ -regular if

- (a) \mathcal{H} is (d_i, ε) -regular at A for any $2 \leq i \leq k-1$ and any $A \in \binom{[s]}{i}$, and
- (b) \mathcal{H} is (d_k, ε_k, r) -regular at A for any $A \in \binom{[s]}{k}$.

This definition of a regular k -complex provides the best generalisation of the notion of regularity in graphs to the hypergraph setting. Indeed, for regular k -complexes we have a Counting Lemma (see [4,6,7,8,9]), which gives the approximate number of copies of any small fixed k -complex within a regular k -complex, as well as an Extension Lemma [3], an Embedding Lemma [3] and (under some additional conditions) a Blow-up Lemma [5], each of which functions similarly as in the graph case.

3 Regular Slices for Hypergraphs

To make use of the definitions of the previous section in solving embedding problems in k -graphs, we need a form of ‘hypergraph regularity lemma’. Informally this should, given a k -graph \mathcal{G} , return one or more $(k-1)$ -complexes \mathcal{J} on $V(\mathcal{G})$ such that adding edges of \mathcal{G} as the ‘ k th layer’ of \mathcal{J} results in a regular k -complex.

To formalise this idea, we make the following further definitions (maintaining the notation of the previous section). Suppose that \mathcal{J} is a \mathcal{P} -partite $(k-1)$ -complex on V , and that \mathcal{G} is a k -graph on V . Then the *restriction of \mathcal{G} to \mathcal{J}* , denoted $\mathcal{G}[\mathcal{J}]$, is the subgraph of \mathcal{G} consisting of all edges of \mathcal{G} which are supported on \mathcal{J} . It follows that $\mathcal{J} \cup \mathcal{G}[\mathcal{J}]$ is a k -complex on V . Moreover, for any k -set X of clusters of \mathcal{P} , we say that \mathcal{G} is (ε_k, r) -regular with respect to X if the restriction of $\mathcal{J} \cup \mathcal{G}[\mathcal{J}]$ to the clusters of X forms a $(d, d_{k-1}, \dots, d_2, \varepsilon_k, \varepsilon, r)$ -regular k -complex for some d ; we refer to this value of d as the *relative density of \mathcal{G} with respect to \mathcal{J} at X* , denoted by $d^*(X)$ if \mathcal{G} and \mathcal{J} are clear from the context.

Ideally a hypergraph regularity lemma would, given a k -graph \mathcal{G} , return a $(k-1)$ -complex \mathcal{J} on $V(\mathcal{G})$ such that most of \mathcal{G} is supported on \mathcal{J} and $\mathcal{J} \cup \mathcal{G}[\mathcal{J}]$ is a regular k -complex; sadly, this is not possible. Instead, existing forms of the hypergraph regularity lemma say that (very roughly speaking) given a k -graph \mathcal{G} with vertex set V , we can find the following. First, a partition \mathcal{P} of V into a bounded number of parts of equal size, called *clusters*. Second, for each $2 \leq \ell \leq k-1$ a partition of the \mathcal{P} -partite ℓ -sets of vertices into a bounded number of parts, called *cells*, with the property that for any \mathcal{P} -partite k -set S of vertices of V , the hypergraph \mathcal{J} whose edge set is the union of all cells containing subsets of S is a regular $(k-1)$ -complex. Moreover, for almost all choices of S , the k -complex $\mathcal{J} \cup \mathcal{G}[\mathcal{J}]$ should be regular (so in particular, all but a few edges of \mathcal{G} lie in the ‘ k th layer’ of a regular k -complex whose ‘lower layers’ are unions of cells). The partitions into cells are often collectively referred to as a *partition k -complex*; much of the technical complexity involved in applications of the Strong Hypergraph Regularity Lemma arises when working with this structure.

Our Regular Slice Lemma is quite different. Indeed, given a k -graph \mathcal{G} it returns a single $(k-1)$ -complex \mathcal{J} for which $\mathcal{J} \cup \mathcal{G}[\mathcal{J}]$ has desirable regularity properties, as in the following definition.

Definition 3.1 Given $\varepsilon, \varepsilon_k > 0$, $r, t_0, t_1 \in \mathbb{N}$ and a k -graph \mathcal{G} with vertex set V , a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice for \mathcal{G} is a $(k-1)$ -complex \mathcal{J} on V such that

- (a) \mathcal{J} is \mathcal{P} -partite for some partition \mathcal{P} of V into t parts of equal size, where $t_0 \leq t \leq t_1$. We call \mathcal{P} the *ground partition* of \mathcal{J} , and call the parts of \mathcal{P} the *clusters* of \mathcal{J} .
- (b) There exists a *density vector* $\mathbf{d} = (d_{k-1}, \dots, d_2)$ such that for each $2 \leq i \leq k-1$ we have $d_i \geq 1/t_1$ and $1/d_i \in \mathbb{N}$, and \mathcal{J} is $(d_{k-1}, \dots, d_2, \varepsilon, \varepsilon, 1)$ -regular.
- (c) \mathcal{G} is (ε_k, r) -regular with respect to all but at most $\varepsilon_k \binom{t}{k}$ of the k -sets of clusters of \mathcal{J} .

Having obtained a regular slice \mathcal{J} for a k -graph \mathcal{G} , we define a weighted reduced k -graph according to the relative density $d^*(X)$ of \mathcal{G} with respect to \mathcal{J} at each k -set X of clusters of \mathcal{J} .

Definition 3.2 [Weighted reduced k -graph] Given a k -graph \mathcal{G} and a $(t_0, t_1, \varepsilon, \varepsilon_k, r)$ -regular slice \mathcal{J} for \mathcal{G} , the *reduced k -graph* $\mathcal{R}_{\mathcal{J}}(\mathcal{G})$ of \mathcal{G} and \mathcal{J} is the complete weighted k -graph whose vertices are the clusters of \mathcal{J} , and where each edge X is given weight $d^*(X)$ (in particular, the weight is in $[0, 1]$). When \mathcal{J} is clear from the context we simply write $\mathcal{R}(\mathcal{G})$ instead of $\mathcal{R}_{\mathcal{J}}(\mathcal{G})$.

In general, it is not very helpful to know that \mathcal{J} is a regular slice for a k -graph \mathcal{G} . Indeed, $\mathcal{G}[\mathcal{J}]$ will usually contain only a tiny fraction of the edges of \mathcal{G} , which need not be representative, so the reduced k -graph of \mathcal{G} with respect to \mathcal{J} does not necessarily resemble \mathcal{G} in the way that the reduced 2-graph of a 2-graph H with respect to a Szemerédi partition resembles H . However, our Regular Slice Lemma states that there exists a regular slice \mathcal{J} for which $\mathcal{R}(\mathcal{G})$ *does* resemble \mathcal{G} , in the sense that densities of small subgraphs (part (a) of the Regular Slice Lemma) and degree conditions (part (b)) are preserved. In order to make this precise, we make the following further definitions.

Fix a weighted k -graph \mathcal{G} , whose weight function we denote by d^* . Given a set $S \subseteq V(\mathcal{G})$ of size j for some $1 \leq j \leq k-1$, the *relative degree* $\overline{\text{deg}}(S; \mathcal{G})$ of S is defined to be

$$\overline{\text{deg}}(S; \mathcal{G}) := \frac{\sum_{e \in \mathcal{G}: S \subseteq e} d^*(e)}{\binom{|V(\mathcal{G}) \setminus S|}{k-j}}.$$

So if every edge of \mathcal{G} has weight one, then $\overline{\text{deg}}(S; \mathcal{G})$ is simply the proportion of k -sets of vertices of \mathcal{G} containing S which are in fact edges of \mathcal{G} . Likewise, for any k -graph \mathcal{H} , we define the \mathcal{H} -density of \mathcal{G} as

$$d_{\mathcal{H}}(\mathcal{G}) := \frac{\sum_{\phi: V(\mathcal{H}) \rightarrow V(\mathcal{G})} \prod_{e \in E(\mathcal{H})} d^*(\phi(e))}{\binom{v(\mathcal{G})}{v(\mathcal{H})} \cdot v(\mathcal{H})!},$$

where ϕ ranges over all injective maps from $V(\mathcal{H})$ to $V(\mathcal{G})$ and d^* is the weight function on $E(\mathcal{G})$. So if every edge of \mathcal{G} has weight one, then the numerator is simply the number of labelled copies of \mathcal{H} in \mathcal{G} , justifying the use of the term \mathcal{H} -density.

We can now give the statement of the Regular Slice Lemma.

Lemma 3.3 (Regular Slice Lemma) *Let $k \geq 3$ be a fixed integer. For any $t_0 \in \mathbb{N}$ and $\varepsilon_k > 0$ and any functions $r : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow (0, 1]$, there are integers t_1 and n_0 such that the following holds for all $n \geq n_0$ which are divisible by $t_1!$. Let \mathcal{G} be a k -graph on n vertices; then there exists a $(k - 1)$ -complex \mathcal{J} on $V(\mathcal{G})$ which is a $(t_0, t_1, \varepsilon(t_1), \varepsilon_k, r(t_1))$ -regular slice for \mathcal{G} with the following additional properties.*

(a) *For any k -graph \mathcal{H} with $v(\mathcal{H}) \leq 1/\varepsilon_k$ we have*

$$|d_{\mathcal{H}}(\mathcal{R}(\mathcal{G})) - d_{\mathcal{H}}(\mathcal{G})| < \varepsilon_k.$$

(b) *For any $1 \leq j \leq k - 1$ and any set Y of j clusters of \mathcal{J} we have*

$$|\overline{\deg}(Y; \mathcal{R}(\mathcal{G})) - \overline{\deg}(\mathcal{J}_Y; \mathcal{G})| < \varepsilon_k.$$

In fact, the version of the Regular Slice Lemma stated in [1] is stronger in several useful ways. First, given several k -graphs $\mathcal{G}_1, \dots, \mathcal{G}_s$ on the same vertex set, it allows us to find a single \mathcal{J} which is simultaneously a regular slice for each \mathcal{G}_i . Second, given an initial partition \mathcal{Q} of $V(\mathcal{G})$ into parts of equal size, it allows us to insist that the ground partition \mathcal{P} of \mathcal{J} is a refinement of \mathcal{Q} . Third, we may insist that the reduced graph inherits degrees and \mathcal{H} -densities within linear-size subsets of $V(\mathcal{G})$, rather than simply within all of $V(\mathcal{G})$. Finally, we can also insist that for any $S \subseteq V(\mathcal{G})$ of size at most $k - 1$, the neighbourhood of S in $\mathcal{G}[\mathcal{J}]$ is an accurate representation of the neighbourhood of S in \mathcal{G} (this final point is particularly useful for embedding spanning subgraphs). However, the form stated above is sufficient for the principal application given in [1], namely to prove a hypergraph analogue of the Erdős-Gallai theorem (see also [2]).

4 Advantages of the Regular Slice Lemma

We believe that many applications of hypergraph regularity can be simplified considerably by using the Regular Slice Lemma, for a number of reasons. Firstly, use of the Regular Slice Lemma avoids the need to introduce and work with the notion of a ‘partition complex’ or related structure; this in itself yields

significant reductions in length and notational complexity. Furthermore, the reduced k -graph as defined in Definition 3.2 is the correct indexing structure for a regular slice, and this fact allows arguments which are much closer in style to arguments using regularity in graphs. By contrast, Keevash [5] observed that the correct indexing structure for a partition complex arising from the Strong Hypergraph Regularity Lemma is a so-called ‘multicomplex’, a more technical structure which is less straightforward to handle.

Due to this, various attempts have been made to define and work with a reduced k -graph following an application of the Strong Hypergraph Regularity Lemma; typically these have a vertex corresponding to each cluster, with k clusters forming an edge if \mathcal{G} is both regular and dense with respect to some $(k - 1)$ -complex formed by cells on these clusters. The drawback of this approach is that edges may intersect in the reduced k -graph without any corresponding intersection in the corresponding complexes. That is, taking $k = 3$ for simplicity, even if both $V_1V_2V_3$ and $V_2V_3V_4$ are edges of the reduced 3-graph, there need be no pair $\{v_2, v_3\}$ with $v_2 \in V_2$ and $v_3 \in V_3$ which is an edge of both the 2-complexes indicated by this fact.

By contrast, if \mathcal{G} is a k -graph, and \mathcal{J} is a regular slice for \mathcal{G} , then sub-complexes of \mathcal{J} corresponding to intersecting edges of the reduced k -graph $\mathcal{R}(\mathcal{G})$ always do share edges in the common clusters. More specifically, again taking $k = 3$, let V_1, V_2, V_3 and V_4 be clusters of \mathcal{J} for which $V_1V_2V_3$ and $V_2V_3V_4$ are both edges of the reduced 3-graph $\mathcal{R}(\mathcal{G})$ which correspond to regular complexes of high density. Then most edges $\{v_2, v_3\} \in \mathcal{J}$ with $v_2 \in V_2$ and $v_3 \in V_3$ are contained in many edges of $\mathcal{G}[V_1, V_2, V_3]$ and also in many edges of $\mathcal{G}[V_2, V_3, V_4]$. Using this fact we can, for example, easily embed a long tight path in \mathcal{G} including vertices from both V_1 and V_4 (where a *tight path* is a sequence of distinct vertices so that any three consecutive vertices form an edge of \mathcal{G}). Indeed, we embed a tight path in $\mathcal{G}[V_1, V_2, V_3]$ using the fact that \mathcal{G} is regular and dense with respect to these clusters, ending at some pair $\{v_2, v_3\} \in \mathcal{J}$ with v_2 in V_2 and $v_3 \in V_3$. Following this we similarly embed a tight path in $\mathcal{G}[V_2, V_3, V_4]$ beginning with the same pair $\{v_2, v_3\}$; these two paths consequently form a single tight path as required. Proceeding in this manner across linearly many edges of $\mathcal{R}(\mathcal{G})$ is the essence of the proof of the Cycle Embedding Lemma proved in [1], which was a key ingredient in the hypergraph analogue of the Erdős-Gallai theorem proved in the same paper (also see [2]). As described above, it is much less straightforward to proceed in this manner following an application of the Strong Hypergraph Regularity Lemma.

References

- [1] P. Allen, J. Böttcher, O. Cooley, and R. Mycroft, *Tight cycles and regular slices in dense hypergraphs*, arXiv:1411.4957.
- [2] ———, *Tight cycles in hypergraphs*, Extended abstract submitted to Eurocomb 2015.
- [3] O. Cooley, N. Fountoulakis, D. Kühn, and D. Osthus, *Embeddings and Ramsey numbers of sparse k -uniform hypergraphs*, *Combinatorica* **29** (2009), no. 3, 263–297.
- [4] W. T. Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, *Ann. of Math. (2)* **166** (2007), no. 3, 897–946.
- [5] P. Keevash, *A hypergraph blow-up lemma*, *Random Structures Algorithms* **39** (2011), no. 3, 275–376.
- [6] B. Nagle, V. Rödl, and M. Schacht, *The counting lemma for regular k -uniform hypergraphs*, *Random Structures Algorithms* **28** (2006), no. 2, 113–179.
- [7] V. Rödl and M. Schacht, *Regular partitions of hypergraphs: counting lemmas*, *Combin. Probab. Comput.* **16** (2007), no. 6, 887–901.
- [8] ———, *Regular partitions of hypergraphs: regularity lemmas*, *Combin. Probab. Comput.* **16** (2007), no. 6, 833–885.
- [9] V. Rödl and J. Skokan, *Regularity lemma for k -uniform hypergraphs*, *Random Structures Algorithms* **25** (2004), no. 1, 1–42.
- [10] E. Szemerédi, *Regular partitions of graphs*, *Problèmes combinatoires et théorie des graphes (Orsay, 1976)*, *Colloques Internationaux CNRS*, vol. 260, CNRS, 1978, pp. 399–401.