Tight Cycles in Hypergraphs

P. Allen [a]   J. Böttcher [a]   O. Cooley [b]   R. 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, UK

Abstract

We apply a recent version of the Strong Hypergraph Regularity Lemma (see [1], [2]) to prove two new results on tight cycles in k-uniform hypergraphs.

The first result is an extension of the Erdős-Gallai Theorem for graphs: For every $\delta > 0$, every sufficiently large k-uniform hypergraph on n vertices with at least $(\alpha + \delta)\binom{n}{k}$ edges contains a tight cycle of length αn for any $\alpha \in [0, 1]$.

Our second result concerns k-partite k-uniform hypergraphs with partition classes of size n and for each $\alpha \in (0, 1)$ provides an asymptotically optimal minimum codegree requirement for the hypergraph to contain a cycle of length $\alpha k n$.

Keywords: Hypergraphs, Tight Cycles, Regularity Method

1   Introduction and main results

Dirac's Theorem states that a graph on n vertices with minimum degree at least $n/2$ contains a Hamilton cycle. This was subsequently strengthened by Bondy [3], who showed that a graph on n vertices with a Hamilton cycle and at least $n^2/4$ edges is either complete bipartite or *pancyclic*, i.e. contains cycles of every length $\ell \leq n$. The famous Erdős-Gallai Theorem [5] states that for any integer $d \geq 3$, any graph on n vertices which contains at least $(d-1)(n-1)/2+1$ edges contains a cycle of length at least d.

Comparatively little is known about tight cycles in k-uniform hypergraphs. A *k-uniform hypergraph* consists of a vertex set V and an edge set $E \subset \binom{V}{k}$. A *tight cycle* in a k-uniform hypergraph is a cyclically ordered set of vertices such that any k consecutive vertices form an edge. An approximate version of the analogue of Dirac's theorem was proved recently by Rödl, Ruciński and Szemerédi [10] (and the same authors proved an exact bound for 3-uniform hypergraphs in [11]). Towards a hypergraph analogue of the Erdős-Gallai Theorem, Győri, Katona and Lemons [6] proved that a k-uniform hypergraph on n vertices with more than $(\alpha n - k)\binom{n}{k-1}$ edges contains a tight path on αn vertices. Our first main result improves upon this by approximately a factor of k.

Theorem 1.1 *For every positive δ and every integer $k \geq 3$, there is an integer n_* such that the following holds for all $\alpha \in [0,1]$. If G is a k-uniform hypergraph on $n \geq n_*$ vertices with $e(G) \geq (\alpha + \delta)\binom{n}{k}$, then G contains a tight cycle of length ℓ for every $\ell \leq \alpha n$ that is divisible by k.*

We note that the asymptotic bound on the number of edges required is the best possible bound that is linear in α, and the divisibility condition is also necessary for general α.

There are also graph results for balanced k-partite graphs (in which all partition classes have the same size). For example, Moon and Moser [8] showed that in a bipartite graph G with classes of size n, a Hamilton cycle is guaranteed if $\delta(G) \geq n/2 + 1$ (i.e. almost halving the degree condition in Dirac's theorem). Best possible minimum degree bounds for Hamiltonicity in balanced k-partite graphs were proved in [4].

Regarding *k-partite*, k-uniform hypergraphs (which have k vertex classes and in which each edge contains exactly one vertex of each class), Rödl and Ruciński [9] proved an asymptotically approximate analogue of the result of

Moon and Moser: For every $\delta > 0$, every sufficiently large k-partite, k-uniform hypergraph with each class of size n, and such that each $(k-1)$-set of vertices containing at most one vertex from each partition class is contained in at least $(\frac{1}{2} + \delta)n$ edges, contains a tight Hamilton cycle.

We extend this result to an asymptotically best possible minimum codegree bound guaranteeing the existence of tight cycles of various lengths.

Theorem 1.2 *For every positive δ and every integer $k \geq 3$ there is an integer n_* such that the following holds for each $\alpha \in [0, 1]$. If G is a k-uniform k-partite hypergraph with parts of size $n \geq n_*$, such that any collection of $k - 1$ vertices, one in each of $k - 1$ parts of G, lies in at least $(\alpha + \delta)n$ edges of G, then*

  (i) *G contains a tight cycle of length ℓ for every $\ell \leq \alpha k n$ that is divisible by k, and*

  (ii) *if $\alpha \geq \frac{1}{2}$ then G contains a tight cycle of length ℓ for every $\ell \leq (1 - \delta)kn$ that is divisible by k.*

The minimum degree bounds in this theorem are easily seen to be asymptotically best possible, and the divisibility condition is certainly necessary since the hypergraph is k-partite.

2 The Regularity Method

The proofs of Theorems 1.1 and 1.2 rely on the regularity method, which we describe briefly in this section.

For graphs, the method is based on the seminal Szemerédi Regularity Lemma which, roughly speaking, says that any sufficiently large graph G may be partitioned into clusters of vertices in such a way that, after deleting a few edges, all of the bipartite graphs between two clusters are in some sense *regular*, meaning they satisfy certain *pseudo-random* properties. This lemma has turned out to be an extremely powerful result, underpinning a huge variety of different proofs in the forty years since it first appeared.

To aid in the application of the Regularity Lemma, we define the *reduced graph R*, an auxiliary graph whose vertices are the clusters of G as provided by the regularity lemma and with edges between clusters whenever there is a sufficiently dense regular bipartite graph between them. It is an easy consequence of the precise definitions that the reduced graph inherits many of the properties of the original graph, usually with only small error terms. In particular, minimum and average degree conditions (proportional to the size of the

respective graph) are inherited by R (with only slightly weakened parameters).

A typical proof strategy using the regularity method is as follows: Given a graph G with certain properties (e.g. large minimum degree), we seek some structure in G (e.g. a long cycle).

- We apply the Regularity Lemma to obtain a reduced graph R inheriting many of the properties of G (e.g. large minimum degree);
- We seek a related but simpler structure in R (e.g. a matching M in a connected component of R);
- Finally, we use the pseudo-randomness properties of R to transfer this simpler structure back to the more complicated structure we seek in G (e.g. by turning every edge of M into a long path in G and connecting up the endpoints using the fact that they all lie in one component of R).

To help complete these last steps, the Regularity Lemma is complemented by various embedding results including counting lemmas for the number of small subgraphs, embedding lemmas for large subgraphs with some suitable restrictions, and extension lemmas guaranteeing that embedded subgraphs can be extended to embed larger subgraphs.

For hypergraphs, there are various generalisations of the regularity lemma. However, the versions of these which are strong enough to help with the problem of embedding tight cycles present significant technical difficulty. In [2] we describe a new form of hypergraph regularity lemma, the Regular Slice Lemma, which avoids much of the complexity associated with previous variants, but which is still sufficiently strong for most applications. In particular, we show that the Regular Slice Lemma allows us to define a reduced hypergraph R inheriting many of the original properties of G (albeit in slightly weakened form); we may then apply corresponding forms of counting lemmas, embedding lemmas and extension lemmas for hypergraph regularity.

We do not go into details about the Regular Slice Lemma here, giving only an outline of how it can be used to prove our main theorems. For more details, see [1] and [2].

3 Proof of Theorem 1.1

We now sketch the proof of Theorem 1.1. Given a hypergraph G on n vertices with at least $(\alpha + \delta)\binom{n}{k}$ edges, we first apply the Regular Slice Lemma which yields a reduced k-uniform hypergraph R on r vertices with at least $(\alpha + \delta/2)\binom{r}{k}$ edges.

In the previous section, we described an outline for a graph proof in which to embed a cycle, we sought a large matching in a *component* of the reduced graph. Now to embed a tight cycle in a k-uniform hypergraph, the correct generalisation is to seek a large matching (i.e. a set of vertex-disjoint edges) in a *tight component* of R. Here a tight component is a maximal subset C of the edges such that for any pair of edges $e_1, e_2 \in C$ there is a *tight walk* $f_1 f_2 \ldots f_m$, where $f_i \in C$, where $|f_i \cap f_{i+1}| = k - 1$ and where $f_1 = e_1$ and $f_m = e_2$. In other words, we can walk between any pair of edges of C using edges which consecutively intersect in $k - 1$ vertices.

If we find a large matching in a tight component of R, then the various embedding and extension lemmas which accompany the Regular Slice Lemma will allow us to extend each edge of the matching to a long tight path which uses almost all of the vertices of the corresponding clusters, and to join all these paths together using short connecting paths within the walks of R that are guaranteed by the fact that the edges of the matching lie in the same tight component. The length of this cycle will be approximately the number of vertices of G lying in clusters of the matching in R. Thus the problem reduces to finding a matching of size at least $(\alpha + \delta/4)r/k$ in a tight component of R.

To do this, we first turn R into a k-*complex* $\mathcal{R}$ by taking the down-closure, i.e. for any edge e of R and any $f \subset e$, f also becomes an edge. We now denote by $E_i(\mathcal{R})$ the set of edges of size i in $\mathcal{R}$, and $e_i(\mathcal{R}) := |E_i(\mathcal{R})|$. Thus in particular $E_k(\mathcal{R}) = E(R)$. We observe that $e_{k-1}(\mathcal{R}) \leq \binom{r}{k-1}$ while

$$e_k(\mathcal{R}) \geq (\alpha + \delta/2) \binom{r}{k} \geq (\alpha + \delta/3)re_{k-1}(\mathcal{R})/k.$$

By averaging, there must exist a tight component $\mathcal{C}$ such that

$$e_k(\mathcal{C}) \geq (\alpha + \delta/3)re_{k-1}(\mathcal{C})/k.$$

We use this property to prove the existence of a large matching in $\mathcal{C}$.

Lemma 3.1 *Let k, s be any natural numbers and let $\mathcal{G}$ be a k-complex in which*

$$e_k(\mathcal{G}) \geq (s - 1)e_{k-1}(\mathcal{G}) + 1 \,.$$

Then $E_k(\mathcal{G})$ contains a matching with at least s edges.

Proof. We will assume that the vertices of $\mathcal{G}$ are the integers $1, \ldots, r$. An important element in the proof is the notion of a *compression* S_{ij}. For integers $i < j$, we define $S_{ij}(\mathcal{G})$ to be the complex obtained from $\mathcal{G}$ by replacing every edge e containing j but not i with the edge $S_{ij}(e) = E \cup \{i\} \setminus \{j\}$, unless

$S_{ij}(e)$ already exists in $\mathcal{G}$, in which case e remains unchanged. It follows directly from the definitions that $S_{ij}(\mathcal{G})$ is a k-complex with the same number of i-edges as $\mathcal{G}$ for each i and whose matching number in each level is at most as large as that of $\mathcal{G}$. We may therefore repeatedly apply compressions to $\mathcal{G}$ until any further compressions have no effect on the complex, and it suffices to prove the existence of a large matching in level k of the resulting complex. Equivalently, we may assume that $\mathcal{G}$ is already *fully compressed*.

A second critical element is the following proposition, which says that if some j-edge lies in few $(j+1)$-edges, we may delete this j-edge and all edges containing it, while maintaining the critical properties of the complex.

Proposition 3.2 *Let k, s be any natural numbers and let $\mathcal{G}$ be a k-complex in which $e_k(\mathcal{G}) \geq (s-1)e_{k-1}(\mathcal{G}) + 1$. Fix $0 \leq j \leq k-1$, and suppose that $e \in E_j(\mathcal{G})$ lies in fewer than $(k-j)s$ edges of $E_{j+1}(\mathcal{G})$. Let $\mathcal{G}'$ be the k-complex obtained by deleting e and any edge of $\mathcal{G}$ which contains e from $\mathcal{G}$. Then $e_k(\mathcal{G}') \geq (s-1)e_{k-1}(\mathcal{G}') + 1$.*

The proof of this proposition is a simple application of the local LYM inequality to the complex we delete – see [2] for details.

Combining this with our previous remarks, we observe that we may assume in Lemma 3.1 that the complex $\mathcal{G}$ is fully compressed and that every j-edge is contained in at least $(k-j)s$ edges of $E_{j+1}(\mathcal{G})$.

Now the proof of Lemma 3.1 is a simple induction: We consider the vertices $(k-1)s+i$, for $i = 1, \ldots, s$. Each of these vertices lies in at least $(k-1)s$ edges of $E_2(\mathcal{G})$ and in particular, since $\mathcal{G}$ is fully compressed, each forms a 2-edge with any vertex from $\{1, \ldots, (k-1)s\}$. Thus the edges $\{(k-2)s+i, (k-1)s+i\}$ form a matching in $E_2(\mathcal{G})$. Continuing in this way, we show that for $i = 1, \ldots, s$ the edges

$$e^{(i)} := \left\{ i, s+i, \ldots, (k-1)s+i \right\}$$

form a matching in $E_k(\mathcal{G})$. This proves Lemma 3.1, and therefore also Theorem 1.1.   □

4   Proof of Theorem 1.2

We first outline the proof of Theorem 1.2 i. Using the Regular Slice Lemma, once again it would be sufficient to prove the existence of a large matching in a tight component of the reduced graph R. In fact, we will prove the existence of a long tight path in R. The reduced graph does not quite inherit

the minimum degree conditions of G, but nevertheless *most* partite sets of $k-1$ vertices are contained in at least $(\alpha + \delta/2)r$ edges of R, where r denotes the size of a partition class in R. The most useful k-edges of R for finding a tight path are those in which every subset of size $k-1$ satisfies this degree condition. We call such edges *excellent*, and it follows from the Regular Slice Lemma that almost all edges are excellent.

We now simply begin from some excellent edge and successively extend the tight path P until this is no longer possible. More precisely, suppose we have a tight path P consisting of excellent edges. Then the last $(k-1)$-tuple of vertices satisfies the condition that it is contained in at least $(\alpha + \delta/2)r$ edges and therefore at least αr excellent edges. Thus providing we have used up less than αr vertices of the corresponding vertex class in R, we can extend P using one more excellent edge. This process certainly continues until the path has length $\alpha k r$. The matching which we require in a tight component then simply consists of every k-th edge of this tight path.

To prove Theorem 1.2 ii we have to be more careful, since the minimum degree condition and the same procedure as above would only guarantee a tight path of half the required length. Instead, we observe that it is enough to construct a perfect *fractional* matching in a tight component of R, i.e. a set of weighted edges such that the sum of the weights of the edges incident to each vertex is exactly one. However, the existence of a perfect fractional matching can also be deduced from the minimum degree conditions using the integer Farkas' Lemma, and in fact with this minimum degree there can be only one tight component (see [2] and [7] for details).

Thus in both cases we have a (fractional) matching of the required size in a tight component of R, and thus by hypergraph regularity arguments, we can also find a tight cycle of the required length.

References

[1] Allen, P., J. Böttcher, O. Cooley, and R. Mycroft, *Regular slices for hypergraphs*, Extended abstract submitted to Eurocomb 2015.

[2] Allen, P., J. Böttcher, O. Cooley, and R. Mycroft, *Tight cycles and regular slices in dense hypergraphs*, Submitted. Arxiv number 1411.4957, http://arxiv.org/pdf/1411.4957.

[3] Bondy, J. A., *Pancyclic graphs*, Journal of Combinatorial Theory, Series B **11** (1971), no. 1, 80–84.

[4] Chen, G., R. J. Faudree, R. J. Gould, M. S. Jacobson, and L. Lesniak, *Hamiltonicity in balanced k-partite graphs*, Graphs and Combinatorics **11** (1995), no. 3, 221–231.

[5] Erdős, P., and T. Gallai, *On maximal paths and circuits of graphs*, Acta Math. Acad. Sci. Hungar **10** (1959), 337–356 (unbound insert).

[6] Győri, E., G. Y. Katona, and N. Lemons, *Hypergraph Extensions of the Erdős-Gallai Theorem*, in preparation.

[7] Keevash, P., and R. Mycroft, *A geometric theory for hypergraph matching*, Memoirs of the American Mathematical Society **233** (2014).

[8] Moon, J., and L. Moser, *On hamiltonian bipartite graphs*, Isr. J. Math. **1** (1963), 163–165.

[9] Rödl, V., and A. Ruciński, *Dirac-type questions for hypergraphs—a survey (or more problems for Endre to solve)*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 561–590.

[10] Rödl, V., A. Ruciński, and E. Szemerédi, *An approximate Dirac-type theorem for k-uniform hypergraphs*, Combinatorica **28** (2008), no. 2, 229–260.

[11] Rödl, V., A. Ruciński, and E. Szemerédi, *Dirac-type conditions for Hamiltonian paths and cycles in 3-uniform hypergraphs*, Adv. Math. **227** (2011), no. 3, 1225–1299.