

ARTICLE

A note on the Brown–Erdős–Sós conjecture in groups

Jason Long

Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, UK

Email: Jason.Long@maths.ox.ac.uk

(Received 9 April 2019; revised 4 September 2019; first published online 3 February 2020)

Abstract

We show that a dense subset of a sufficiently large group multiplication table contains either a large part of the addition table of the integers modulo some k , or the entire multiplication table of a certain large abelian group, as a subgrid. As a consequence, we show that triples systems coming from a finite group contain configurations with t triples spanning $O(\sqrt{t})$ vertices, which is the best possible up to the implied constant. We confirm that for all t we can find a collection of t triples spanning at most $t + 3$ vertices, resolving the Brown–Erdős–Sós conjecture in this context. The proof applies well-known arithmetic results including the multidimensional versions of Szemerédi’s theorem and the density Hales–Jewett theorem.

This result was discovered simultaneously and independently by Nenadov, Sudakov and Tyomkyn [5], and a weaker result avoiding the arithmetic machinery was obtained independently by Wong [11].

2010 MSC Codes: Primary 05C25; Secondary 05C35, 05C65

1. Introduction

A central open problem in extremal combinatorics is the Brown–Erdős–Sós conjecture [1]. We say that a subgraph H' of a hypergraph H is an (r, s) -configuration if $|E(H')| = s$ and $|V(H')| \leq r$. The Brown–Erdős–Sós conjecture states that, for any fixed positive integer $t \geq 3$, any 3-uniform hypergraph H on n vertices which does not contain a $(t + 3, t)$ -configuration has at most $o(n^2)$ edges. The number $t + 3$ cannot be decreased, since random constructions can achieve $\Omega(n^2)$ edges while avoiding any $(t + 2, t)$ -configurations [1]. The conjecture can be generalized to higher uniformity, but we shall focus on the 3-uniform case in this note.

Since its formulation in 1973 there has been a great deal of work on this problem. Ruzsa and Szemerédi [7] resolved the first non-trivial case ($t = 3$), but the conjecture remains open for all $t > 3$. The strongest result to date is due to Sárközy and Selkow [8], who showed that any 3-uniform hypergraph which does not contain a $(t + 2 + \lfloor \log_2 t \rfloor, t)$ -configuration has at most $o(n^2)$ edges.

When tackling the Brown–Erdős–Sós conjecture, we may additionally assume that the hypergraph H is linear (as noted in [9], for example). It is also clear that we may assume that H is tripartite, since, given a 3-graph H , we may obtain a tripartite 3-graph H' by taking three copies of the vertex set of H and placing edges between these partitions corresponding to the edges of H .

Given a linear, tripartite, 3-uniform hypergraph H on $n + n + n$ vertices we can associate a partially labelled $n \times n$ grid by labelling position (a, b) with label c if $(a, b, c) \in E(H)$. Thus the

Brown–Erdős–Sós conjecture can be formulated in terms of a quasigroup; this is noted in [9] and [10], for example.

Conjecture 1.1 (Brown–Erdős–Sós). Fix $t \in \mathbb{Z}^+$ and $\varepsilon > 0$. Then there exists $N = N(t, \varepsilon)$ such that, for any quasigroup G of order $n > N$ and any subset A of the multiplication table of G of density at least ε , we can find a $(t + 3, t)$ -configuration in A ; that is to say, a set of t triples in A spanning at most $t + 3$ vertices (i.e. rows, columns or labels).

In light of this formulation, it is natural to ask the same question when G is in fact a group, as the additional structure might provide greater local density than can be found in random constructions.

Conjecture 1.2 (Brown–Erdős–Sós for groups). Fix $t \in \mathbb{Z}^+$ and $\varepsilon > 0$. Then for any sufficiently large group G and any subset A of the multiplication table of G of density at least ε , we can find a $(t + 3, t)$ -configuration in A .

Since the Brown–Erdős–Sós conjecture is resolved for $t \leq 3$, the first interesting case of Conjecture 1.2 is $t = 4$. In 2015, Solymosi [9] resolved this case, showing that Conjecture 1.2 holds for $t = 4$.

Recently, Solymosi and Wong [10] showed that much more is true, proving that the Brown–Erdős–Sós threshold of $t + 3$ vertices can in fact be surpassed in the groups setting. In particular, they prove that dense subsets of sufficiently large group multiplication tables contain sets of t triples in A spanning asymptotically only $3t/4$ vertices. Since their result concentrates on the case of large t , they do not match Conjecture 1.2 for small t but prove that it holds for infinitely many t .

Given that the Brown–Erdős–Sós threshold can be surpassed in the groups setting, one may ask what the correct behaviour should be in this case. Since A corresponds to a linear hypergraph, we cannot find sets of t triples in A spanning fewer than \sqrt{t} vertices, but can we approach this lower bound?

Question 1.3 Let t be a fixed positive integer. What is the smallest number $F(t)$ such that we are guaranteed to find an $(F(t), t)$ -configuration in a dense subset of a sufficiently large group multiplication table?

In this note we answer this question up to a constant factor, and resolve Conjecture 1.2. By applying machinery from arithmetic combinatorics, including the multidimensional Szemerédi theorem and a multidimensional variant of the density Hales–Jewett theorem, we prove that any dense subset of a sufficiently large group multiplication table contains a large subgrid belonging to one of two families: either the subgrid matches part of the multiplication table of a cyclic group, or the subgrid matches the entire multiplication table of \mathbb{F}_p^m for some small prime p and large m . A precise statement appears in Theorem 2.2 following some notation.

This reduces Question 1.3 to a discrete optimization problem, in which we must find configurations with t edges spanning few vertices in each of the two cases resulting from our main theorem. We tackle this optimization problem in Section 4, showing that $F(t) = O(\sqrt{t})$ and resolving Conjecture 1.2 for all t .

2. Notation and statements

We write \mathbb{Z}_n for the group of integers modulo n under addition and we write $[k]$ for the set $\{0, 1, \dots, k - 1\}$. We begin with some definitions.

Definition. By the *multiplication table* of a group $G = (G, \circ)$ we mean the collection of triples $(a, b, a \circ b)$ for $a, b \in G$. The *vertex set* will be given by three disjoint copies of G called the *row vertices*, *column vertices* and *label vertices*. We shall refer to the triples as the edges or *faces* of the corresponding tripartite 3-uniform hypergraph. Typically, we will represent this as a labelled grid, with entry (a, b) given label $a \circ b$. In the case that $G = (G, +)$ is an abelian group, we will usually call the multiplication table an *addition table*.

Definition. By a *subgrid* of a labelled grid, we mean the labelled grid contained in the intersection of some subset of the rows and columns.

Definition. We say that a labelled grid A is *isomorphic* to another labelled grid B if we can biject the row sets, column sets and label sets of A and B in such a way that the resulting map is a graph isomorphism between the corresponding 3-graphs.

Using this notation we reformulate Question 1.3 in a precise way.

Question 2.1 Let t be a fixed positive integer and $\varepsilon > 0$. Let $F(t)$ be minimal such that, given any subset A of density at least ε of a sufficiently large (in terms of t and ε) group multiplication table, we may find an $(F(t), t)$ -configuration in A . How does $F(t)$ grow with t ? Is $F(t) \leq t + 3$ for all t ?

In order to answer this question, we prove the following structural result.

Theorem 2.2 Fix $k, m \in \mathbb{Z}^+$ and $\varepsilon > 0$. Then there exists $N = N(k, m, \varepsilon)$ such that, for any group G of order $n > N$ and any subset A of the multiplication table of G of density at least ε , A contains either a subgrid isomorphic to the addition table of $[k]$ as a subset of \mathbb{Z}_K for some $K \geq k$, or a subgrid isomorphic to the addition table of \mathbb{Z}_p^m for some $p < k$ prime.

Remark. This result is ‘best possible’ in terms of finding configurations with many edges spanned by few vertices, since if A is simply taken to be the addition table of $[n/2]$ as a subset of \mathbb{Z}_n , say, then any subgrid of A is isomorphic to part of a larger addition table and we cannot improve on the first case of the theorem. Similarly, if A is simply the addition table of \mathbb{Z}_p^t for small p and large t , then we cannot improve on the second case.

3. Proof of Theorem 2.2

We start by introducing the arithmetic machinery that we use later. We begin with a multidimensional version of Szemerédi’s theorem [3].

Theorem 3.1 (multidimensional Szemerédi theorem). Let $k, t \in \mathbb{Z}^+$ and let $\varepsilon > 0$. Then there exists $N = N(\varepsilon, k, t)$ such that, for any $n > N$ and any $A \subset \mathbb{Z}_n^t$ of density at least ε , we can find $a_1, a_2, \dots, a_t, d \in \mathbb{Z}_n$ such that

$$(a_1 + i_1 d, a_2 + i_2 d, \dots, a_t + i_t d) \in A$$

for each $i_j \in \{0, \dots, k-1\}$. In other words, A contains the Cartesian product of t arithmetic progressions of length k with the same common difference.

We shall also need a multidimensional version of the density Hales–Jewett theorem [4]. We recall the definition of a combinatorial line.

Definition. A combinatorial line in \mathbb{Z}_m^n is a set U of the form

$$U = \{(x_1, \dots, x_n) \mid x_i \text{ constant on } I, x_j = z_j \text{ for } j \notin I\}$$

for some indexing set $I \subset \{1, \dots, n\}$ and some $z \in \mathbb{Z}_m^n$. A combinatorial subspace of dimension k is a set U of the form

$$U = \{(x_1, \dots, x_n) \mid x_i \text{ constant on each } I_s, x_j = z_j \text{ for } j \notin \cup_s I_s\}$$

for some collection of k disjoint indexing sets $I_s \subset \{1, \dots, n\}$, and some $z \in \mathbb{Z}_m^n$.

Theorem 3.2 (density Hales–Jewett). Fix $m \in \mathbb{Z}^+$ and let $\varepsilon > 0$. Then there exists $N = N(\varepsilon, m)$ such that, for any $n > N$ and any $A \subset \mathbb{Z}_m^n$ of density at least ε , we can find a combinatorial line inside A .

The density Hales–Jewett theorem easily implies its own multidimensional variant; for a proof, see [2], for example.

Corollary 3.3 (multidimensional density Hales–Jewett). Let m, k be fixed positive integers and let $\varepsilon > 0$. There exists $N = N(\varepsilon, m, k)$ such that, for any $n > N$ and any $A \subset \mathbb{Z}_m^n$ of density at least ε , we can find an entire combinatorial subspace of dimension k inside A .

We will need a further variant of density Hales–Jewett, which follows easily from Corollary 3.3 by applying the same idea used to extend from Theorem 3.2 to Corollary 3.3.

Corollary 3.4 Let k, t be fixed positive integers, p a fixed prime, and let $\varepsilon > 0$. There exists $N = N(\varepsilon, p, k, t)$ such that, for any $n > N$ and any $A \subset (\mathbb{Z}_p^n)^t$ of density at least ε , we can find a subspace Γ of dimension k and $a_1, \dots, a_t \in \mathbb{Z}_p^n$ such that

$$(a_1 + \Gamma) \times (a_2 + \Gamma) \times \dots \times (a_t + \Gamma) \subset A.$$

Proof. We simply identify $(\mathbb{Z}_p^n)^t$ with \mathbb{Z}_p^{nt} in the obvious way. We can then apply Corollary 3.3 to find a combinatorial subspace of dimension k inside A , which gives us an affine subspace of dimension k . The result follows by translating back to $(\mathbb{Z}_p^n)^t$. \square

Lastly, we will need Pyber’s theorem [6], which provides us with a large abelian subgroup of G .

Theorem 3.5 (Pyber’s theorem). There is a universal constant $c > 0$ such that any group G of order n contains an abelian subgroup of order at least $e^c \sqrt{\log(n)}$.

We are now ready to prove Theorem 2.2.

Proof of Theorem 2.2. We begin by applying Theorem 3.5, which states that G contains an abelian subgroup G' of order at least $\exp(c\sqrt{N})$ for some absolute constant $c > 0$. In particular, $N' = |G'|$ tends to infinity with N .

Note that the multiplication table of G can be partitioned into the Cartesian products of left cosets of G' with right cosets of G' . Since A has density at least ε in the full multiplication table $G \times G$, we know that there exists $r, s \in G$ such that A has density at least ε in the Cartesian product $rG' \times G's$. The part of the multiplication table corresponding to this Cartesian product is isomorphic to the addition table of G' . Let $A' = A \cap (rG' \times G's)$ be the subset of $rG' \times G's$ of density at least ε obtained from A . Note that G' is a finite abelian group, and can therefore be written as a direct product of cyclic groups of prime power order.

Suppose that G' has a cyclic factor \mathbb{Z}_T . Then, as above, we can find a subset A'' which has density at least ε in a Cartesian product of two cosets of \mathbb{Z}_T in G , and this Cartesian product is isomorphic to the addition table of \mathbb{Z}_T . Thus A'' corresponds to a subset of the $T \times T$ addition table of density at least ε . By Theorem 3.1, if $T > T(k, \varepsilon)$ is sufficiently large then we can find a Cartesian product of two arithmetic progressions $(a, a + d, \dots, a + (k - 1)d)$ and $(b, b + d, \dots, b + (k - 1)d)$ in A'' . The labels in this subgrid belong to the set $\{a + b, a + b + d, \dots, a + b + 2d\}$. Indeed, this subgrid is isomorphic to the addition table $\{0, \dots, k - 1\} \times \{0, \dots, k - 1\} \subset \mathbb{Z}^2$ and so we are in the first case of the statement of the theorem.

So we are done if G' contains a cyclic factor \mathbb{Z}_T with $T > T(k, \varepsilon)$. Therefore we may assume that all factors of G' are cyclic groups with bounded (prime power) order. Since $|G'|$ tends to infinity with N , we see that for any positive integer M , if N is sufficiently large we may find (by the pigeonhole principle) a cyclic factor \mathbb{Z}_{p^a} which appears to the power M . In particular, G' contains \mathbb{Z}_p^M as a subgroup.

As above, we note that this means that we may find $A'' \subset A$ which has density at least ε in the Cartesian product of two cosets of \mathbb{Z}_p^M inside G , and this product is isomorphic to the multiplication table of \mathbb{Z}_p^M . If M is sufficiently large in terms of m , then by Corollary 3.4 we can find the complete Cartesian product of $a + \mathbb{Z}_p^m$ and $b + \mathbb{Z}_p^m$ inside A'' . This complete Cartesian product is isomorphic to the addition table of \mathbb{Z}_p^m . If $p \geq k$ then we can find the addition table of \mathbb{Z}_p and we are in the first case of the theorem, and otherwise we have $p < k$ and are in the second case. \square

We now see how Theorem 2.2 simplifies Question 2.1. We let $f(t)$ be minimal such that we can find an $(f(t), t)$ -configuration in the addition table of $[k] \subset \mathbb{Z}_K$ for any $K \geq k$ sufficiently large compared to t . Similarly, for each prime p we let $g_p(t)$ be minimal such that we can find an $(g_p(t), t)$ -configuration in the addition table of \mathbb{Z}_p^m for any sufficiently large m (in terms of t and p).

Corollary 3.6 *We have that $F(t) = \max_p (f(t), g_p(t))$.*

Proof. Clearly $F(t) \leq \max_p (f(t), g_p(t))$. For the other direction, we apply Theorem 2.2 for choices of k and m sufficiently large in terms of t . Given a subset A of density at least ε of the multiplication table of some sufficiently large group G , we may therefore find a subgrid isomorphic to the entire addition table of $[k]$ as a subset of \mathbb{Z}_K for some $K \geq k$, or a subgrid isomorphic to the entire addition table of \mathbb{Z}_p^m for some $p < k$ prime. If k and m are chosen large enough (in terms of t only), we deduce that A contains either an $(f(t), t)$ -configuration or a $(g_p(t), t)$ -configuration and so $F(t) \geq \max_p (f(t), g_p(t))$. \square

We have thus reduced Question 2.1 to the problem of finding $f(t)$ and $g_p(t)$. We will devote the next section to tackling this discrete optimization problem; providing an exact, closed form answer for all t is tricky because of certain divisibility considerations.

4. Finding locally dense configurations

In order to keep the note brief, we will not attempt to give the best possible bounds. We will instead show that $F(t) = O(\sqrt{t})$, and, because of the connection with Conjecture 1.1, we will separately confirm that $F(t) \leq t + 3$ for all t .

For the analysis of the discrete optimization problem arising from Corollary 3.6, it simplifies the calculations to try and maximize the number of faces induced by a fixed number ν of vertices rather than minimize the number of vertices spanned by a fixed number t of faces. Thus we let $f'(\nu)$ be the maximal number of faces that can be spanned by a set of ν vertices in the addition table of $[k] \subset \mathbb{Z}_K$ for any $K \geq k$ sufficiently large compared to ν , and observe that if $f'(\nu) \geq t$ then

$f(t) \leq v$. Similarly, we let $g'_p(v)$ be the maximal number of faces that can be spanned by a set of v vertices in the addition table of \mathbb{Z}_p^m for any m sufficiently large in terms of v and p , and observe that if $g'_p(v) \geq t$ then $g_p(t) \leq v$.

Proposition 4.1 *We have that $f'(v) \geq (1 + o(1))v^2/12$, and therefore $f(t) \leq (\sqrt{12} + o(1))\sqrt{t}$.*

Proof. We work in the addition table of $[k] \subset \mathbb{Z}_K$ for $K \geq k \geq v$. Given r rows and r columns, we can optimize the density of our configuration by including the s most numerous labels. The labels in the addition table are constant along falling diagonals. In the worst case, each falling diagonal corresponds to a different label, in which case the most numerous label occurs r times, the next two most numerous labels occur $r - 1$ times each, etc. Therefore, by including the s most numerous labels, we include a total of at least

$$r + (r - 1) + (r - 1) + (r - 2) + (r - 2) + \dots + (r - \lceil (s - 1)/2 \rceil) = sr - s(s - 1)/4 - \frac{1}{2} \lfloor s/2 \rfloor$$

different faces. The total number of vertices is $2r + s$, so we seek to maximize this expression with respect to the constraint that $2r + s \leq v$. Taking $r = \lfloor v/3 \rfloor$ and $s = \lceil v/3 \rceil$, and noting that $f'(v)$ is an increasing function of v , the proposition follows. □

Proposition 4.2 *We have that $g'_p(v) \geq (1 + o(1))v^2/49$ for all p , and therefore $g_p(t) \leq (7 + o(1))\sqrt{t}$.*

Proof. We work in the addition table T of \mathbb{Z}_p^m for m large. If $p \geq v/3$ then the construction in the proof of Proposition 4.1 finds a configuration in the addition table of \mathbb{Z}_p with $(1 + o(1))v^2/12$ faces and so we are done.

Otherwise, let l be minimal such that $3p^{l+1} > v$. For m sufficiently large, T contains a subgrid isomorphic to the multiplication table of \mathbb{Z}_p^{l+1} . We can partition this multiplication table into the Cartesian products of the cosets of \mathbb{Z}_p^l . These Cartesian products can be arranged into a $p \times p$ grid of blocks ($p^l \times p^l$ subgrids) corresponding to entries of the addition table $\mathbb{Z}_p \times \mathbb{Z}_p$.

We form our configuration by taking a union of these blocks. Let $v = \lambda p^l$, and so $\lambda \in [3, 3p)$. The number B of blocks that we can use is precisely the maximum number of faces induced by $\lfloor \lambda \rfloor$ vertices in the addition table of \mathbb{Z}_p . The number of vertices in the resulting configuration will be at most v , and the number of edges will $Bp^{2l} = Bv^2/\lambda^2$.

Since $p > \lambda/3$ we could use the construction idea from Proposition 4.1. Unfortunately, we cannot assume that λ is large (in which case we could take approximately $\lambda^2/12$ blocks and therefore approximately $v^2/12$ faces) and the worst cases for this construction will in fact be decided by the best options for small λ .

In order to minimize the calculation, we will instead simply take an $a \times a$ grid of these blocks, and we shall choose a maximal subject to our constraint on the number of vertices.

If we take the bottom left $a \times a$ grid of these Cartesian products we obtain a configuration with ap^l rows, ap^l columns and at most $(2a - 1)p^l$ labels. The configuration has a^2p^{2l} faces. Taking a maximal so that $4a - 1 \leq v/p^l = \lambda$, we obtain a configuration C with at most v vertices.

By the maximality of a we see that $a = \lfloor \lambda/4 + 1/4 \rfloor$, so in particular $a \geq \max(1, \lambda/4 - 3/4)$. The number of faces of the configuration C is a^2p^{2l} , which is therefore at least

$$\max\left(\frac{v^2}{\lambda^2}, \frac{(\lambda - 3)^2}{16\lambda^2}v^2\right),$$

which takes its minimal value of $v^2/49$ when $\lambda = 7$. □

Remark. It is not hard to show that Proposition 4.1 is in fact best possible, and $1/12$ is the correct constant in the limit. On the other hand, Proposition 4.2 does not give the correct constant. As mentioned in the proof, combining the construction in Proposition 4.1 with a careful analysis of small λ would allow improvements to be made quite easily. We can also make use of leftover vertices (when λ is not an integer, a union of blocks uses only $\lfloor \lambda \rfloor p^l < v$ vertices, leaving some unused) to interpolate between the constructions for integer values of λ . Using these techniques we can improve the constant from $1/49$ to $5/64$. However, the calculations are quite involved and the result would still not be the best possible, so we have tried to find a compromise between giving the best bounds that we can and providing a streamlined result.

Combining Propositions 4.1 and 4.2 with Corollary 3.6 gives the following result.

Corollary 4.3 $F(t) = O(\sqrt{t})$ (in fact, $F(t) \leq (7 + o(1))\sqrt{t}$).

Therefore, the Brown–Erdős–Sós threshold of $(t + 3, t)$ is far below what can be found given the extra group structure. Nevertheless, we will now confirm that we do indeed prove the Brown–Erdős–Sós conjecture in the context of group multiplication tables, which essentially involves checking that sufficiently dense configurations exist for the small values of t , as well as for large t as verified by Corollary 4.3.

Proposition 4.4 We have that $F(t) \leq t + 3$ for all $t \geq 3$.

Proof. Although much better bounds than $t + 3$ are possible for large t , it will be most convenient simply to find $(t + 3, t)$ -configurations in the addition table of $[k] \subset \mathbb{Z}_K$ for $K \geq k$ large, and also in the addition table of \mathbb{Z}_p^m for m large. The result will then follow by Corollary 3.6.

For the first case, working in the addition table of $[k] \subset \mathbb{Z}_K$, we note that taking the points in positions $(0, 0)$, $(0, 1)$ and $(1, 0)$ gives the configuration

$$\begin{matrix} 1 \\ 0 & 1 \end{matrix}$$

which has six vertices spanning three faces. Next, we can include the point in position $(1, 1)$, which introduces one new vertex (a new label, 2) and one new face. Then the point in position $(2, 0)$ introduces one new vertex (a new column) and one new face, and then the point in position $(2, 1)$ introduces one new vertex (a new label) and one new face. Continuing, we introduce the points in positions $(i, 0)$ and $(i, 1)$ for each i until we have t faces. At this point we have a configuration with t faces spanning $t + 3$ vertices.

In the second case, we are working in the addition table of \mathbb{Z}_p^m for m large. We can use the above argument to find an $(r + 3, r)$ -configuration for r up to $2p - 1$ by taking the bottom two rows, minus the final face, of the multiplication table of some copy of \mathbb{Z}_p . When we add in the final point in position $(p - 1, 1)$ we re-use the label in position $(0, 0)$ so we get an $(r + 2, r)$ -configuration. We can then start again in a new copy of \mathbb{Z}_p , including the corresponding points one by one in the same order as before. Our first point introduces two new vertices (a new row and new column) for just one more face, but since we are adding it to an $(r + 2, r)$ -configuration we get back to an $(r + 3, r)$ -configuration. Thereafter we add at most one new vertex with every new face. Once we finish the bottom two rows of the next copy of \mathbb{Z}_p we can start again in another copy, and we can continue until we have t faces. At that point we will span at most $t + 3$ vertices. □

Conjecture 1.2 follows immediately from Proposition 4.4.

5. Concluding remarks

We have shown that the Brown–Erdős–Sós conjecture is true for hypergraphs with an underlying group structure, and in fact a much stronger result is possible. We give a bound of $O(\sqrt{t})$ on the minimum size of a collection of vertices spanned by t edges, which is tight up to the implied constant. Theorem 2.2 provides an explanation for this local density by showing that bounded-size subgrids manifesting an abelian group structure can be found in any dense subset of a group multiplication table.

It is natural to wonder if the ability to find many configurations with density beating the Brown–Erdős–Sós threshold is in some way connected to group-like structure. Are there interesting structural constraints weaker than the group axioms that still provide local density beyond the Brown–Erdős–Sós threshold? Or does the existence of many (r, s) -configurations with r sufficiently small in terms of s require an underlying group structure?

Acknowledgement

The author would like to thank Tim Gowers for several helpful comments.

References

- [1] Brown, W. G., Erdős, P. and Sós, V. T. (1973) Some extremal problems on r -graphs. In *New Directions in the Theory of Graphs* (F. Harary, ed.), Academic Press.
- [2] Dodos, P., Kanellopoulos, V. and Tyros, K. (2014) A simple proof of the density Hales–Jewett theorem. *Int. Math. Res. Not.* **12** 3340–3352.
- [3] Furstenberg, H. and Katznelson, Y. (1978) An ergodic Szemerédi theorem for commuting transformations. *J. Analyse Math.* **38** 275–291.
- [4] Furstenberg, H. and Katznelson, Y. (1989) A density version of the Hales–Jewett Theorem for $k = 3$. *Discrete Math.* **75** 227–241.
- [5] Nenadov, R., Sudakov, B. and Tyomkyn, M. (2019) Proof of the Brown–Erdős–Sós conjecture in groups. *Math. Proc. Cam. Phil. Soc.* 1–11.
- [6] Pyber, L. (1997) How abelian is a finite group? In *The Mathematics of Paul Erdős I* (R. L. Graham et al., eds), Vol. 13 of *Algorithms and Combinatorics*, Springer, pp. 372–384.
- [7] Ruzsa, I. and Szemerédi, E. (1978) Triple systems with no six points carrying three triangles. *Coll. Math. Soc. J. Bolyai* **18** 939–945.
- [8] Sárközy, G. N. and Selkow, S. (2004) An extension of the Ruzsa–Szemerédi theorem. *Combinatorica* **25** 77–84.
- [9] Solymosi, J. (2015) The $(7, 4)$ -conjecture in finite groups. *Combin. Probab. Comput.* **24** 680–686.
- [10] Solymosi, J. and Wong, C. (2019) The Brown–Erdős–Sós conjecture in finite abelian groups. *Discrete Appl. Math.*
- [11] Wong, C. (2019) On the existence of dense substructures in finite groups. Preprint.