

Algorithmic Number On the Forehead Protocols Yielding Dense Ruzsa-Szemerédi Graphs and Hypergraphs

Noga Alon *

Princeton University
Princeton, NJ 08544, USA
and Tel Aviv University
Tel Aviv 69978, Israel
nalon@math.princeton.edu

Adi Shraibman

The Academic College of Tel-Aviv-Yaffo
Tel-Aviv, Israel
adish@mta.ac.il

Abstract

We describe algorithmic Number On the Forehead protocols that provide dense Ruzsa-Szemerédi graphs. One protocol leads to a simple and natural extension of the original construction of Ruzsa and Szemerédi. The graphs induced by this protocol have n vertices, $\Omega(n^2/\log n)$ edges, and are decomposable into $n^{1+O(1/\log \log n)}$ induced matchings. Another protocol is an explicit (and slightly simpler) version of the construction of [1], producing graphs with similar properties. We also generalize the above protocols to more than three players, in order to construct dense uniform hypergraphs in which every edge lies in a positive small number of simplices.

1 Introduction

For an integer n and a positive real c , let $h(n, c)$ denote the maximum number so that any n vertex graph with at least cn^2 edges in which every edge is contained in a triangle, must contain an edge lying in at least $h(n, c)$ triangles. Erdős and Rothschild asked to determine or estimate $h(n, c)$, see [5], [8], [9], [10]. Szemerédi observed that the triangle removal lemma (see [21]) implies that for every fixed $c > 0$, $h(n, c)$ tends to infinity with n , and Trotter and the first author noticed that for any $c < 1/4$ there is a c' so that $h(n, c) < c'\sqrt{n}$. A clever construction of Fox and Loh [13] shows that in fact for any fixed $c < 1/4$, $h(n, c) \leq n^{O(1/\log \log n)}$. While this is still very far from the lower bound based on the triangle removal lemma and its improved quantitative version in [12], which provides a lower bound exponential in $\log^* n$ for any fixed $c > 0$, it does show that $h(n, c) = n^{o(1)}$.

*Research supported in part by NSF grant DMS-1855464, ISF grant 281/17, BSF grant 2018267 and the Simons Foundation.

Note that the constant $1/4$ is tight, as it is known that any n -vertex graph with $\lfloor n^2/4 \rfloor + 1$ edges must contain an edge lying in at least $n/6$ triangles (see [15]).

The construction of Fox and Loh triggered another surprising result in the study of a closely related problem. The first author, Moitra and Sudakov [1] constructed (r, t) -Ruzsa-Szemerédi graphs on n vertices with $r = n^{1-o(1)}$ and $rt = (1 - o(1))\binom{n}{2}$. A graph is an (r, t) -Ruzsa-Szemerédi graph if its set of edges can be partitioned into t pairwise disjoint induced matchings, each of size r . These graphs were introduced in a paper by Ruzsa and Szemerédi [21]. They used these graphs, together with the regularity lemma of Szemerédi [23] to tackle the so called $(6, 3)$ -problem dealing with the maximum possible number of edges of a 3-uniform hypergraph on n vertices that contains no 3 edges spanning at most 6 vertices. Ruzsa-Szemerédi graphs have been studied extensively since, finding applications in Combinatorics, Complexity Theory and Information Theory. A natural line of research is to find dense graphs with relatively large r . One such construction is given by Birk, Linial and Meshulam [4], with $r = (\log n)^{\Omega(\log \log n / (\log \log \log n)^2)}$ and $t = \Omega(n^2/r)$. Meshulam conjectured that there are no (r, t) -Ruzsa-Szemerédi graphs with both $rt = \Theta(\binom{n}{2})$ and $r \geq n^{\Omega(1)}$. The construction from [1] disproved Meshulam's conjecture in a strong form, vastly improving the one in [4].

The first aim of the present short paper is to describe these results in communication complexity terms by providing algorithmic Number-On-the-Forehead (NOF, for short) protocols that entail them. Ruzsa-Szemerédi graphs are closely related to the NOF model in communication complexity, as observed in [18]. They are related to the communication complexity of 2-dimensional permutations and sub-permutations (see details in the sequel). We observe here that communication protocols in the NOF model for 2-dimensional permutations also imply upper bounds on $h(n, c)$.

We give algorithmic NOF protocols that derive the constructions of dense Ruzsa-Szemerédi graphs from [1] and also the results of Fox and Loh [13]. This makes the constructions strongly explicit and also somewhat simpler. Another advantage of this approach is that it provides a clear link between these results and the original results of Ruzsa and Szemerédi [21].

The second aim of this paper is to extend the above mentioned results to uniform hypergraphs. To do so we extend the protocols to any number $k > 3$ of players. Let $K_k = K_k^{(k-1)}$ denote the complete $(k-1)$ -uniform hypergraph ($(k-1)$ -graph, for short) on k vertices. For an integer n and a positive real c , let $h_{k-1}(n, c)$ denote the maximum number so that any n vertex $(k-1)$ -graph with at least cn^{k-1} edges, in which every edge is contained in a copy of K_k , must contain an edge lying in at least $h_{k-1}(n, c)$ such copies. By the hypergraph removal lemma proved in [14] and independently in [20], [19], for any fixed positive c , $h_{k-1}(n, c)$ tends to infinity with n . Indeed, for example, if G is an n -vertex 3-graph with at least cn^3 edges, and each edge is contained in at least 1 and at most $h = h_3(n, c)$ copies of $K = K_4$, then G must contain at least $\frac{cn^3}{4h}$ pairwise edge-disjoint copies of K . Hence at least that many edges have to be omitted from G in order to destroy all copies of K , and thus by the hypergraph removal lemma if h is a constant then G must contain at least $\Omega(n^4)$ copies of K , implying that some edges are contained in $\Omega(n)$ such copies, contradiction.

Unlike the graph case, the maximum possible number $ex_{k-1}(n, K_k)$ of edges of an n -vertex

$(k - 1)$ -graph with no copies of K_k is not known. The determination of this number is an old problem posed by Turán [22], and Erdős offered a significant award for its solution, see [7]. By a general result proved in [16], the limit of the ratio

$$\frac{ex_{k-1}(n, K_k)}{n^{k-1}}$$

as n tends to infinity exists. This is a positive number called the Turán density of K_k . Let $d_k = d(K_k)$ denote this number, which is conjectured to be $5/9$ for $k=4$. See [6] and its references for some of the work on this problem. Although d_k is not known, we can prove the following.

Theorem 1.1 *For any fixed $c < d_k$ there is some $b > 0$ so that $h_{k-1}(n, c) \leq n^{b/\log \log n}$.*

Note that by the results of [11] on supersaturated hypergraphs if $c > d_k$ then any $(k - 1)$ -graph on n vertices with at least cn^{k-1} edges contains $\Omega(n^k)$ copies of K_k . Therefore, for any such c there is a constant $b = b(c) > 0$ so that $h_{k-1}(n, c) \geq bn$, implying that the d_k bound in Theorem 1.1 is tight.

Our protocols also imply an extension of the main result of [1]. That is, it entails a construction of nearly complete $(k - 1)$ -graphs whose edges can be partitioned into a nearly linear number of induced subgraphs, each being a partial Steiner system. Recall that a $(k - 1)$ -graph is a partial Steiner system if no two of its edges share $k - 2$ common vertices. It is clear that any such graph on n vertices cannot contain more than $\frac{1}{k-1} \binom{n}{k-2} < n^{k-2}$ edges, and hence any $(k - 1)$ -graph with at least bn^{k-1} edges cannot be partitioned into less than $\Theta(n)$ partial Steiner systems. The hypergraph removal lemma shows here, too, that in fact the number of such systems cannot be $\Theta(n)$, that is, for any fixed positive b , this number divided by n must tend to infinity with n . The following result shows, however, that this number can be smaller than $n^{1+\epsilon}$ for any positive ϵ .

Theorem 1.2 *For every integer $k \geq 3$, there is an absolute constant $c > 0$ so that for sufficiently large n there is a $(k - 1)$ -graph on n vertices with at least*

$$(1 - o(1)) \binom{n}{k-1}$$

edges, whose edges can be decomposed into at most $n^{1+c/\log \log n}$ induced subgraphs, each being a partial Steiner system.

The rest of the paper contains the proofs of the above two theorems. The organization is as follows. Section 2 contains background on communication complexity and high-dimensional permutations, a recipe for proving Theorem 1.1 and Theorem 1.2 using communication protocols, and a simple application of this recipe to construct a graph on n vertices and $\Omega(n^2/\log n)$ edges, decomposable into $n^{1+O(1/\log \log n)}$ induced matchings. Section 3 contains the application of this recipe to prove Theorem 1.1 and Theorem 1.2. The details of the graphs and hypergraphs produced by this recipe, and the proof that it works correctly are given in Section 4. The final Section 5 contains a brief summary.

2 From communication to graphs and hypergraphs

2.1 Background and notation

General notation We let $[n] = \{1, 2, \dots, n\}$. A k -tuple is denoted either (x_1, \dots, x_k) or in abbreviated form \vec{x} .

Communication complexity We start with a few basic communication complexity notions. The definitions we give are a simplified version and adjusted to our needs. The interested reader can see [17] for a more comprehensive survey. In the NOF model k players wish to compute a function $f : X_1 \times X_2 \times \dots \times X_k \rightarrow \{0, 1\}$. The players agree on a communication *protocol* P . Then, an input (x_1, x_2, \dots, x_k) is presented to the players so player i sees all input except x_i , we sometimes refer to this player as the x_i -player. The players take turns to write messages on a blackboard according to the agreed protocol P . Each message of each player may depend on the part of the input seen by this player, and except for the last player it can also depend on the messages written so far on the blackboard. The message written by the last player depends only on the part of the input he sees, and is independent of the content of the blackboard. One way to visualize this is as if the last player wrote a message first and then did not participate in the rest of the transaction. The value of the function can be computed by all players from the content of the board at the end of the protocol. The *cost* of a protocol, denoted $C(P)$, is the maximal number of bits written on the board, over all inputs, by the first $k - 1$ players¹.

The string of bits written on the blackboard for a given input $\vec{x} = (x_1, \dots, x_k)$ is called a *transcript*, denoted $\mathcal{T}(\vec{x})$. We let $\mathcal{T}_i(\vec{x})$ for $i = 1, \dots, k$ be the part of this transcript that is written by player i . Let T be a transcript, the subset $S = S(T)$ of entries satisfying $\mathcal{T}(\vec{x}) = T$ and $f(\vec{x}) = 1$, is called a *cylinder intersection*². Note that a cylinder intersection is defined with respect to a function and a protocol for this function, we specify the function and protocol when it is necessary for a clear presentation and otherwise omit them.

We say that a subset of entries S is *symmetric* if membership in S does not depend on the order of the first $k - 1$ entries. That is, S is symmetric if $(x_1, \dots, x_{k-1}, x_k) \in S$ if and only if $(x_{\pi(1)}, \dots, x_{\pi(k-1)}, x_k) \in S$ for every permutation π on $\{1, 2, \dots, k - 1\}$.

High-dimensional permutations A *line* in $[n]^k$ is a subset $L \subset [n]^k$ such that $k - 1$ of the coordinates in L are fixed, and the remaining coordinate takes all possible values. Following is a simple example with $n = 5$ and $k = 3$:

$$L = \{(1, 1, 4), (1, 2, 4), (1, 3, 4), (1, 4, 4), (1, 5, 4)\}.$$

¹In the basic communication complexity definition all players can see each others messages, and the cost of the protocol depends also on the message of the last player. The version of communication complexity we gave here is from the *one-sided* model. Since we only need this version, we simplify our notations.

²The usual definition of cylinder intersection is more general, what we defined here is referred to as a 1-monochromatic cylinder intersection. Since we are only interested in 1-monochromatic cylinder intersections we abbreviate the notation.

In this example the first and third coordinates are fixed, and the second coordinate takes all possible values in $[5] = \{1, 2, 3, 4, 5\}$. There is a distinct line for every choice of unconstrained coordinate $i \in [k]$, and a choice of values to fix the remaining coordinates. A line in $[n_1] \times \cdots \times [n_k]$ is defined similarly. We say that the line is in the i th dimension if the unconstrained coordinate is i .

A $(k-1)$ -dimensional permutation is a function $f : [n]^k \rightarrow \{0, 1\}$ such that for every line L in $[n]^k$ there is exactly one $\vec{x} \in L$ such that $f(\vec{x}) = 1$. A sub-permutation is a function $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ such that every line in the k th dimension contains a single 1, and every other line contains at most one 1.

For example, let G be a group, define $f : G^k \rightarrow \{0, 1\}$ by $f(x_1, \dots, x_k) = 1$ if and only if $x_1 + x_2 + \cdots + x_k = 0$. Then f is a permutation. Let H be a subset of G , then the function $h : H^{k-1} \times G \rightarrow \{0, 1\}$ defined similarly to f , is a sub-permutation.

A weak permutation is a function $f : [n]^k \rightarrow \{0, 1\}$ such that every line contains at most one 1-entry, and a weak sub-permutation is defined similarly: it is an $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ with $N \geq n$ such that every line contains at most one 1-entry.

Ruzsa-Szemerédi graphs and hypergraphs As mentioned in the introduction, a graph is an (r, t) -Ruzsa-Szemerédi graph if its set of edges can be partitioned into t pairwise disjoint induced matchings, each of size r . Such a graph obviously has rt edges. A challenge in constructing Ruzsa-Szemerédi graphs is to make the density of edges as large as possible while keeping the number of matchings relatively low. We are therefore less concerned with the size of each matching, and only worry about the number of matchings and the density of the edges.

There is a natural way to extend the notion of Ruzsa-Szemerédi graphs to hypergraphs, by considering Steiner systems $S(k-2, k-1)$. A Steiner system $S(t-1, t)$ in a set V , is a family of t -element subsets of V (called *blocks*) such that each $(t-1)$ -element subset of V is contained in exactly one block. A partial Steiner system is defined similarly with the exception that each $(t-1)$ -element subset of V is contained in **at most** one block.

For a natural number $k > 2$, and a $(k-1)$ -graph $G = (V, E)$ we are interested in partitioning E into induced partial Steiner systems $S(k-2, k-1)$. Note that if V is the set of vertices of a graph, then a partial Steiner system $S(1, 2)$ in V is a matching. Thus, this definition extends the notion of a Ruzsa-Szemerédi graph.

2.2 A recipe

Given a function $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$, a protocol P for f , and a transcript T of the last player, denote

$$S_k(T) = \{(x_1, \dots, x_k) \in [n]^{k-1} \times [N] : \mathcal{T}_k(x_1, \dots, x_k) = T \text{ and } f(x_1, \dots, x_k) = 1\}.$$

Next we describe a recipe for generating Ruzsa-Szemerédi graphs and hypergraphs, as well as upper bounds on $h_{k-1}(n, c)$, from NOF protocols.

Recipe 1 - from protocols to graphs and hypergraphs

1. Choose a weak sub-permutation $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$, for natural numbers n , N and $k > 2$.
2. Construct a communication protocol P for f .
3. Pick a transcript T of the last player so that $S_k(T)$ is symmetric, and let $S = S_k(T)$.

The following theorem describes the outcome when following Recipe 1.

Theorem 2.1 *Let P be a protocol found in the second step of Recipe 1, and let S be the subset of inputs picked in the last step. Let $p = |S|/n^{k-1}$, $\gamma = C(P)$ and $N' = N \cdot 2^\gamma$, then*

1. *There is an (explicitly defined) $(k-1)$ -graph on n vertices whose edge density is p , that is the union of N' induced partial Steiner systems $S(k-2, k-1)$.*
2. *If $p = 1 - o(1)$, then $h_{k-1}(n, c) \leq (N'/n)^2$ for $c < d_k$. Here, the construction of the $(k-1)$ -graph that gives the bound is also explicit, given explicit constructions of $(k-1)$ -graphs of density $d_k - o(1)$ which contain no K_k .*

We defer the proof of Theorem 2.1 and the explicit definition of the graphs produced by Recipe 1 to Section 4. In the next section we give a simple example of how Theorem 2.1 can be applied, then in Section 3 we apply it to prove Theorems 1.1 and 1.2.

2.3 Applying Theorem 2.1 - an example

We apply Theorem 2.1 to prove:

Lemma 2.2 *There is a graph on n vertices with edge density $\Omega(1/\log n)$ that is the union of $n^{1+1/\Omega(\log \log n)}$ induced matchings.*

Proof We follow the steps of Recipe 1:

Choosing the function Let $q, d > 1$ be natural numbers, denote $n = q^d$, and define $Z_{q,d} = \{\frac{1}{2}(x+y) : x, y \in [q]^d\}$. Denote by $g_{q,d} : ([q]^d)^2 \times Z_{q,d} \rightarrow \{0, 1\}$ the function satisfying $g_{q,d}(x, y, z) = 1$ if and only if $x + y = 2z$ (here addition is in \mathbb{R}^d). It is not hard to verify that $g_{q,d}$ is a sub-permutation. Denote $N = N_{q,d} = |Z_{q,d}|$, then

$$N \leq (2q)^d = q^d \cdot 2^d = n^{1+1/\log q}.$$

Since $\log \log n = \log d + \log \log q$, we have that $N \leq n^{1+1/\Omega(\log \log n)}$ as long as $d \leq q^c$ for some constant c . We will later choose $d = q^4$.

The protocol Next we present a protocol for $g_{q,d}$.

Protocol 1 A protocol for $g_{q,d}$

1. The z -player computes $\|x - y\|_2^2$, and writes the result on the board.
2. The y -player writes 1 iff $\|x - y\|_2^2 = 4\|x - z\|_2^2$.
3. The x -player writes 1 iff $\|x - y\|_2^2 = 4\|y - z\|_2^2$.

At the end, all players know the value of the function. Indeed, the value of the function is 1 if the last two bits written on the board are both equal to 1, and 0 otherwise.

The cost of the protocol The cost of the protocol is $C(P) = 2$, as the first two players send only 2 verification bits.

The choice of S By the Chernoff-Hoeffding's inequality (c.f., e.g., [2]), the quantity $\|x - y\|_2^2$ computed by the third player satisfies

$$P(|\|x - y\|_2^2 - \mathbb{E}(\|x - y\|_2^2)| \geq t) \leq 2e^{-\frac{2t^2}{dq^4}}.$$

Thus, with constant probability, $\|x - y\|_2^2$ takes one of $\sqrt{dq^2}$ values. There is, therefore, a transcript T for the third player such that $|S_3(T)| \geq \Omega(n^2/\sqrt{dq^2})$. If we take $d = q^4$ we get $|S_3(T)| \geq \Omega(n^2/d) \geq \Omega(n^2/\log n)$. The fact that $S_3(T)$ is symmetric is easy to verify. Lemma 2.2 now follows from Theorem 2.1, part 1. ■

Note that we could improve the density of the graph in Lemma 2.2 to $\Omega(\log \log n / \log^\epsilon n)$ for any constant $\epsilon > 1/2$ by taking $d = q^c$ for an appropriately chosen large constant c . This seems to be the best one can get when using Protocol 1 though. In the next section we use a variant of this protocol in which the first two players participate more, in order to save communication bits of the last player. This will allow us to increase the density to near optimal.

3 Applying Theorem 2.1 to prove Theorems 1.1 and 1.2

3.1 The case $k = 3$

Choosing the function The function we choose is $g_{q,d}$, defined in Section 2.3. We later fix $d = q^5$.

The protocol For a natural number r let $G_r = (V, E_r)$ be the graph with $V = [q]^d$, where d is even, and $E_r = \{x, y : \|x - y\|_2^2 \leq r\}$ (later we take $r = \sqrt{d}$). The players agree on a proper coloring χ of G_{2r} by $d_{2r} + 1$ colors, where d_{2r} is its maximum degree. Let $\mu = \mathbb{E}(\|x - y\|_2^2) = \frac{1}{6}d(q^2 - 1)$, the players also agree on some partition P of $[0, dq^2]$ into intervals of length $r^2 + O(1)$. The players choose P that satisfy: the number of intervals in the partition is $\lceil dq^2/r^2 \rceil$, and the number μ is in the middle of the interval containing it. As an example, the players can choose a partition which is a translation of the partition induced by $DIV(L) = \lfloor \frac{L}{r^2} \rfloor$. Let $I_r : [0, dq^2] \rightarrow \{0, 1, \dots, dq^2/r^2\}$ map a number in $[0, dq^2]$ to the index of the interval containing it, according to P . Given an input (x, y, z) , the players then use the following protocol:

Protocol 2 A protocol for $g_{q,d}$

1. The z -player writes $I_r(\|x - y\|_2^2)$ on the board.
2. The y -player verifies that $I_r(\|x - y\|_2^2) = I_r(4\|x - z\|_2^2)$, and writes 1 on the board iff this is the case.
3. The x -player verifies that $I_r(\|x - y\|_2^2) = I_r(4\|y - z\|_2^2)$, and writes 1 on the board iff this is the case.
4. If one of the last two bits are equal to 0, reject and finish.
5. The x -player writes $\chi(2z - y)$ on the board.
6. The y -player writes the value of $g_{q,d}(x, y, z)$.

Theorem 3.1 *Protocol 2 is correct.*

For the proof of correctness, we use the following two observations (used also in [1]):

Lemma 3.2 (Parallelogram law) *Let $x, y, z \in \mathbb{R}^d$ then:*

$$\|x - y\|_2^2 + \|x + y - 2z\|_2^2 = 2\|x - z\|_2^2 + 2\|y - z\|_2^2$$

Lemma 3.3 ([1]) *For an even integer $d > 0$, the number of integral points contained in the ball of radius r in \mathbb{R}^d is at most:*

$$\frac{\pi^{d/2}(r + 0.5)^d}{(d/2)!} < \frac{(2\pi e)^{d/2}(r + 0.5\sqrt{d})^d}{(d)^{d/2}}$$

Proof [of Theorem 3.1] By Lemma 3.3, the maximum degree of G_r is at most

$$d_r = \frac{(2\pi e)^{d/2}(r + 0.5\sqrt{d})^d}{(d)^{d/2}}.$$

The chromatic number of G_{2r} is therefore at most $d_{2r} + 1$.

If $x + y = 2z$ then obviously the protocol reaches step 5. On the other hand, if the protocol reached step 5 then $\|x - y\|_2^2$, $4\|x - z\|_2^2$, and $4\|y - z\|_2^2$, all lie in the same interval of length r^2 . Thus, by the Parallelogram law

$$\begin{aligned} \|x + y - 2z\|_2^2 &= 2\|x - z\|_2^2 + 2\|y - z\|_2^2 - \|x - y\|_2^2 \\ &= \frac{1}{2} (4\|x - z\|_2^2 + 4\|y - z\|_2^2) - \|x - y\|_2^2 \\ &\leq r^2. \end{aligned}$$

Thus, $(2z - y)$ is in a ball $B(x, r)$ of radius r around x . Every other vector $v \in B(x, r)$ is in distance at most $2r$ from $(2z - y)$. The color of $(2z - y)$ in this ball is therefore unique. It follows that at step 6 the y -player knows the value of y and hence knows everything. \blacksquare

The cost of the protocol The number of bits used by the first two players is:

$$\log d_{2r} + \Theta(1) = \Theta \left(d + d \log \frac{2r + 0.5\sqrt{d}}{\sqrt{d}} \right).$$

If we take $r = \sqrt{d}$, the cost of the protocol is therefore bounded by

$$C(P) \leq O(d) = O \left(\frac{\log n}{\log q} \right).$$

The choice of S A transcript T of the z -player corresponds to a message $I_r(\|x - y\|_2^2)$. The size of $S_3(T)$ is therefore equal to the number of pairs $x, y \in [q]^d$ satisfying $I_r(\|x - y\|_2^2) = T$. Hoeffding's inequality implies that

$$P(|\|x - y\|_2^2 - \mu| \geq t) \leq 2e^{-\frac{2t^2}{dq^4}}.$$

In particular, the probability that $I_r(\|x - y\|_2^2) = I_r(\mu)$ is at least $(1 - 2e^{-\frac{r^4}{2dq^4}})$ since we chose the partition of the intervals so that μ lies in the middle of the interval containing it.

Take $r = \sqrt{d}$, and pick $S = S_3(T)$ for $T = I_{\sqrt{d}}(\mu)$, we have

$$|S| \geq (1 - 2e^{-\frac{d}{2q^4}})n^2.$$

Conclusion When applying Theorem 2.1 the parameters that we get are:

- $p = (1 - 2e^{-\frac{d}{2q^4}})$,
- $N' = n^{1+1/\Omega(\log \log n)} 2^{O(d)}$.

Taking $d = q^5$, and observing that S is symmetric, this proves the $k = 3$ case of Theorems 1.1 and 1.2.

3.2 The case $k > 3$

Choosing the function Let $Z_{m,q,d} = \{\frac{1}{m}(\sum_{i=1}^m x_i) : x_i \in [q]^d\}$ and define $g_{k,q,d} : ([q]^d)^{k-1} \times Z_{k-1,q,d} \rightarrow \{0, 1\}$ by $g_{k,q,d}(x_1, \dots, x_k) = 1$ if and only if $x_1 + \dots + x_{k-1} = (k-1)x_k$. It is easy to verify that $g_{k,q,d}$ is a sub-permutation, and

$$|Z_{k-1,q,d}| \leq (kq)^d = n^{1+1/\log_k q}.$$

The protocol The protocol is a simple reduction to the case $k = 3$.

Protocol 3 A protocol for $g_{k,q,d}$

1. The first player writes 1 on the board if and only if $\frac{1}{2}((k-1)x_k - x_3 - \dots - x_{k-1}) \in Z_{2,q,d}$.
2. If the last bit was equal to 0, the protocol ends with rejection.
3. Players 1, 2 and k run Protocol 2 for $g_{3,q,d}$ with $r = \sqrt{d}$ on input $x' = x_1, y' = x_2$, and $z' = \frac{1}{2}((k-1)x_k - x_3 - \dots - x_{k-1})$.

The correctness of the above protocol follows from the correctness of Protocol 2 and the fact that the equation $x_1 + x_2 + x_3 + \dots + x_{k-1} = (k-1)x_k$ holds if and only if $x_1 + x_2 = 2(\frac{1}{2}((k-1)x_k - x_3 - \dots - x_{k-1}))$. Note that the last equation cannot hold if $\frac{1}{2}((k-1)x_k - x_3 - \dots - x_{k-1})$ does not belong to $Z_{2,q,d}$.

The cost of the protocol Outside the reduction to Protocol 2, the players send only one more bit. The cost of the protocol thus satisfy $C(P) \leq O(d) \leq O(\frac{\log n}{\log q})$, as before.

The choice of S We can choose, as in Section 3.1, the set $S = S_k(T)$ for $T = I_{\sqrt{d}}(\mu)$. By Hoeffding's inequality, the size of S is $(1 - o(1))n^{k-1}$ as long as $d \gg q^4$. The only problem is that S is not symmetric. To remedy that, just add to the protocol a test whether $I_r(\|x_i - x_j\|_2^2) = I_r(\mu)$ for every $1 \leq i < j < k$. These tests can all be carried out by the last player, so this adds only one more communication bit, which for simplicity we assume is the last bit. Now pick the transcript $T' = (T, 1)$ which imply that $I_r(\|x_i - x_j\|_2^2) = I_r(\mu)$ for all $1 \leq i < j < k$. The corresponding set $S_k(T')$ is now symmetric, and as long as k is a constant, Hoeffding's inequality still implies that the size of $S_k(T')$ is at least $(1 - o(1))n^{k-1}$.

4 Proof of Theorem 2.1

We first rephrase Theorem 2.1 slightly.

Theorem 4.1 Let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, and let S be a symmetric cylinder intersection (w.r.t. f). Let $p = |S|/n^{k-1}$, then

1. There is an (explicitly defined) $(k-1)$ -graph on n vertices whose edge density is p , that is the union of N induced partial Steiner systems $S(k-2, k-1)$.
2. If $p = 1 - o(1)$, then $h_{k-1}(n, c) \leq (N/n)^2$ for $c < d_k$. Here, the construction of the $(k-1)$ -graph that gives the bound is explicit, given explicit constructions of $(k-1)$ -graphs of density $d_k - o(1)$ which contain no K_k .

Lemma 4.2 *Theorem 4.1 implies Theorem 2.1.*

Proof The difference between Theorem 4.1 and Theorem 2.1 lies in the different properties of the subset S . In Theorem 2.1 S is defined by

$$S = S_k(T_k) = \{(x_1, \dots, x_k) \in [n]^{k-1} \times [N] : \mathcal{T}_k(x_1, \dots, x_k) = T_k \text{ and } f(x_1, \dots, x_k) = 1\},$$

for some transcript T_k of **the last player**. In Theorem 4.1 on the other hand, S is a cylinder intersection, that is

$$S = S(T) = \{(x_1, \dots, x_k) \in [n]^{k-1} \times [N] : \mathcal{T}(x_1, \dots, x_k) = T \text{ and } f(x_1, \dots, x_k) = 1\},$$

for some transcript T of **all players**.

This difference is easily bridged though. Let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, P a protocol for f , T_k a transcript of the last player, and $S = S_k(T_k)$ a subset, found using Recipe 1. Let $\gamma = C(P)$, and denote $N' = N \cdot 2^\gamma$. For simplicity identify $[N']$ with $[N] \times \{0, 1\}^\gamma$.

Define $g : [n]^{k-1} \times [N'] \rightarrow \{0, 1\}$ by $g(x_1, \dots, x_{k-1}, (x_k, T_{1\dots k-1})) = 1$ if and only if $f(x_1, \dots, x_{k-1}, x_k) = 1$ and $T_{1\dots k-1} = \mathcal{T}_1(x_1, \dots, x_k) \circ \dots \circ \mathcal{T}_{k-1}(x_1, \dots, x_k)$. That is, $T_{1\dots k-1}$ is the message written on the board by the first $k-1$ players, according to protocol P , on input (x_1, \dots, x_k) .

It is not hard to verify that g is a weak sub-permutation. We use the following protocol P' for g , on input $(x_1, \dots, x_{k-1}, (x_k, T_{1\dots k-1}))$: the last player sends his message as in P , then each of the other players verifies (using one bit of communication each) that his part in $T_{1\dots k-1}$ agrees with P . Obviously P' is correct if and only if P is correct. The subset

$$S' = \{(x_1, \dots, (x_k, T_{1\dots k-1})) \in [n]^{k-1} \times [N'] : \mathcal{T}_k(x_1, \dots, x_k) = T_k \text{ and } f(x_1, \dots, x_k) = 1\}$$

is a cylinder intersection with respect to P' and g , and $|S'|/n^{k-1} = |S|/n^{k-1}$. Theorem 4.1 can now be applied to prove Theorem 2.1. ■

In the rest of this section we prove Theorem 4.1. For simplicity we first prove it for the case of graphs ($k = 3$) and then explain the necessary adjustments for the general case ($k \geq 3$).

4.1 The case $k = 3$

We prove the first conclusion of Theorem 4.1, concerning Ruzsa-Szemerédi graphs, in Section 4.1.1. The upper bound on $h(n, c)$ is proved in Section 4.1.2. We use the following simple fact proved in [18].

Lemma 4.3 ([18]) *Let $f : [n] \times [n] \times [N] \rightarrow \{0, 1\}$ be a function satisfying that every line in the third dimension contains at most a single 1, and let S be a cylinder intersection (w.r.t f). Then, S does not contain stars: triplets of the form $(x', y, z), (x, y', z), (x, y, z')$ where $x \neq x', y \neq y'$ and $z \neq z'$.*

4.1.1 Ruzsa-Szemerédi graphs

The relation between Ruzsa-Szemerédi graphs and the communication complexity of 2-dimensional permutations was observed in [18]. The graphs constructed in [18] are bipartite though, and we need slightly different settings. Let $S \subseteq [n] \times [n] \times [N]$ be symmetric, define

$$E_S = \{(x, y), (x, z), (y, z) : (x, y, z) \in S\}.$$

Let $G_S = (V, E_S)$ be the graph with vertex set $V = V_A \cup V_B$, where $V_A = [n]$ and $V_B = [N]$, and edge set E_S . We allow self loops in E_S , and consider a collection of self loops as a matching. Note that when S is a cylinder intersection with respect to a weak sub-permutation there is always at most one edge between a pair of vertices. The following lemma implies the first conclusion in Theorem 4.1.

Lemma 4.4 *Let $f : [n] \times [n] \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, and let S be a symmetric cylinder intersection. Let $H = ([n], F)$ be the subgraph of G_S induced on V_A . That is:*

$$F = \{(x, y) : \exists z \in V_B \text{ s.t. } (x, y, z) \in S\}.$$

Then, the edges of $|F|$ can be partitioned into N induced matchings.

Proof Partition the edge set F as follows, for every $z \in B$ let

$$F_z = \{(x, y) : (x, y, z) \in S\}.$$

This is a partition of F since f a sub-permutation, and therefore there is at most a single z such that $(x, y, z) \in S$ for every $(x, y) \in [n]^2$.

The fact that F_z is an induced matching follows from Lemma 4.3. Assume in contradiction that F_z is not an induced matching, then there is an edge $(x, y) \in F_{z'}$ for $z' \neq z$ such that $(x, y'), (x', y)$ are in F_z . We then get a star $(x', y, z), (x, y', z), (x, y, z') \in S$, contradicting Lemma 4.3. Note that the fact that f is a sub-permutation also implies that $x' \neq x$ and $y' \neq y$. ■

4.1.2 An upper bound on $h(n, c)$

Consider the same graph G_S as in the previous section. A basic observation is:

Lemma 4.5 *Let $f : [n] \times [n] \times [N] \rightarrow \{0, 1\}$ be a function satisfying that every line in the third dimension contains at most a single 1, and let S be a symmetric cylinder intersection (w.r.t f). Then, a triangle (x, y, z) where $x, y \in V_A$ and $z \in V_B$ exists in G_S if and only if $(x, y, z) \in S$.*

Proof The fact that a triangle $(x, y), (x, z), (y, z)$ where $x, y \in V_A$ and $z \in V_B$ exists in G_S for every $(x, y, z) \in S$ follows immediately from the definition of E_S . Assume in contradiction that there is also such a triangle in G_S for $(x, y, z) \notin S$. Then necessarily there are $x', y' \in V_A$ and $z' \in V_B$ such that $(x', y, z), (x, y', z), (x, y, z') \in S$. But then S contains a star, in contradiction to Lemma 4.3. \blacksquare

Lemma 4.6 *Let $f : [n] \times [n] \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, and let S be a symmetric cylinder intersection satisfying $|S| = (1 - o(1))n^2$. Then $h(n, c) \leq N^2/n^2$ for $c < 1/4$.*

Proof Consider the graph G_S again. By lemma 4.5, and the fact that f is a weak sub-permutation, an edge in G_S appears in exactly one triangle (x, y, z) with $x, y \in V_A$ and $z \in V_B$. Therefore, if we take a bipartite subgraph inside V_A , we will have every edge lie in exactly one triangle, which is optimal. But, the density of edges in G_S is relatively small, since there are $n + N$ vertices and order of $(1 - o(1))n^2$ edges. To remedy this, we define a product function, aiming to increase the density of edges. The price we pay is that the number of triangles an edge can lie in increases.

Let $t \geq 2$ be a natural number, define $f^t : ([2^t] \times [n])^2 \times [N] \rightarrow \{0, 1\}$ by $f^t((\alpha, x), (\beta, y), z) = 1$ if and only if $f(x, y, z) = 1$. Let

$$S^t = \{((\alpha, x), (\beta, y), z) : (x, y, z) \in S\}.$$

It is not hard to verify that S^t is a symmetric cylinder intersection with respect to f^t . By Lemma 4.5 a triangle $((\alpha, x), (\beta, y), z)$ where $(\alpha, x), (\beta, y) \in ([2^t] \times [n])$ and $z \in [N]$ exists in G_{S^t} if and only if $(x, y, z) \in S$. Thus, every edge of G_{S^t} lies in at most 2^t triangles of this sort. To remove other kind of triangles let $H = ([2^t] \times [n], E_H)$ be a bipartite graph with density $1/4$. Now define

$$E'_{S^t} = \{((\alpha, x), (\beta, y)), ((\alpha, x), z), ((\beta, y), z) : (x, y, z) \in S, ((\alpha, x), (\beta, y)) \in E_H\}.$$

Then every edge in E'_{S^t} lies in at least one triangle and at most 2^t triangles. The number of edges satisfy $|E'_{S^t}| \geq (1 - o(1))(2^t n)^2/4$. The density of edges is thus

$$(1 - o(1)) \frac{1}{4} \frac{(2^t n)^2}{(2^t n + N)^2}.$$

If we take $t = 2 \log(N/n)$ this becomes

$$(1 - o(1)) \frac{1}{4} \frac{(N^2/n)^2}{(N^2/n + N)^2}.$$

Recall that S is a cylinder intersection of size $(1 - o(1))n^2$. It therefore follows from the graph removal lemma (and the hypergraph removal lemma for larger k) - see Theorem 34 in [18] for details - that necessarily $n = o(N)$. The density is thus $(1 - o(1))\frac{1}{4}$. Since every edge is in at most $2^t = N^2/n^2$ triangles, this completes the proof. ■

4.2 The general case

We outline the proof of Theorem 4.1 for $k \geq 3$. Since the general case is very similar to the proof of the $k = 3$ case, we do not repeat all the details here.

For $\vec{x} = (x_1, \dots, x_k) \in [n]^{k-1} \times [N]$ denote by $[\vec{x}]_{k-1}$ the family of all subsets of size $k - 1$ of entries of \vec{x} . That is:

$$[\vec{x}]_{k-1} = \binom{\{x_1, \dots, x_k\}}{k-1}.$$

Let $S \subseteq [n]^{k-1} \times [N]$ be a symmetric subset of entries, define

$$E_S = \bigcup_{\vec{x} \in S} [\vec{x}]_{k-1}.$$

Let $G_S = (V, E_S)$ be the $(k - 1)$ -graph with vertex set $V = V_A \cup V_B$, where $V_A = [n]$ and $V_B = [N]$, and edge set E_S .

The generalized version of Lemma 4.3 is:

Lemma 4.7 ([18]) *Let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a function satisfying that every line in the k th dimension contains at most a single 1, and let S be a cylinder intersection (w.r.t f). Then, S does not contain stars: k entries of the form $(x'_1, x_2, \dots, x_k), (x_1, x'_2, \dots, x_k), (x_1, x_2, \dots, x'_k)$ where $x'_i \neq x_i$ for $i = 1 \dots k$.*

This immediately gives:

Lemma 4.8 *Let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a function satisfying that every line in the k th dimension contains at most a single 1, and let S be a symmetric cylinder intersection (w.r.t f). Then, we have that $[\vec{x}]_{k-1}$ is a copy of K_k in G_S with $x_1, \dots, x_{k-1} \in V_A$ and $x_k \in V_B$, if and only if $\vec{x} = (x_1, \dots, x_k) \in S$.*

Proof Similar to the proof of Lemma 4.5, but using Lemma 4.7 instead of Lemma 4.3. ■

The following two lemmas generalize Lemma 4.4 and Lemma 4.6:

Lemma 4.9 For an integer $k \geq 3$, let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, and let S be a symmetric cylinder intersection. Let $G' = ([n], E')$ be the subgraph of G_S induced on V_A . Then, the edges of $|E'|$ can be partitioned into N partial Steiner systems $S(k-2, k-1)$.

Proof The proof is similar to the proof of Lemma 4.4, we rewrite the main points. The edges of G' are:

$$E' = \{(x_1, \dots, x_{k-1}) : \exists x_k \in V_B \text{ s.t. } (x_1, \dots, x_{k-1}, x_k) \in S\}.$$

Partition the edge set E' as follows, for every $z \in V_B$ let

$$E'_z = \{(x_1, \dots, x_{k-1}) : (x_1, \dots, x_{k-1}, z) \in S\}.$$

This is a partition of E' since f a sub-permutation, and the fact that E'_z is a partial Steiner system follows from Lemma 4.7. ■

Lemma 4.10 For an integer $k \geq 3$, let $f : [n]^{k-1} \times [N] \rightarrow \{0, 1\}$ be a weak sub-permutation, and let S be a symmetric cylinder intersection satisfying $|S| = (1 - o(1))n^{k-1}$. Then $h_{k-1}(n, c) \leq (N/n)^2$ for $c < d_k$.

Proof The proof is very similar to the proof of Lemma 4.6, just instead of taking the subgraph $H = ([2^t] \times [n], E_H)$ to be a bipartite graph with density $1/4$, take a subhypergraph with no copies of K_k and density d_k . Note that we do not need to know d_k or H , we just need to know that d_k is finite and that H exists. ■

5 Summary

As mentioned in the introduction, there is a link between the main construction of [1] and the original construction of Ruzsa and Szemerédi [21]. We describe this link here, starting with a new construction, equivalent to the one of Ruzsa and Szemerédi, derived using the recipe in Section 2.2. Our approach avoids the use of Behrends construction of a large set of integers without a three-term arithmetic progression [3], which was the heart of the construction of Ruzsa and Szemerédi.

Lemma 5.1 ([21]) *There exists a graph on n vertices, with $n^2/2^{O(\sqrt{\log n})}$ edges, that is the union of $\Theta(n)$ induced matchings.*

Proof We follow the steps of Recipe 1. The details are very similar to those in Section 2.3, with slight modifications.

Choosing the function Let $q, d > 1$ be natural numbers and denote $n = q^d$. Let $f_{q,d} : ([q]^d)^3 \rightarrow \{0, 1\}$ be the function satisfying $f_{q,d}(x, y, z) = 1$ if and only if $x + y = 2z$. It is not hard to verify that $f_{q,d}$ is a weak sub-permutation, in fact it is a weak permutation. We later set q to be even and $d = \log(q) = \Theta(\sqrt{\log n})$.

The protocol The protocol is identical to the protocol for $g_{q,d}$ in Section 2.3.

The cost of the protocol The cost of the protocol is $C(P) = 2$.

The choice of S By Hoeffding's inequality, with constant probability, $\|x - y\|_2^2$ takes one of $\sqrt{d}q^2$ values. There is, therefore, a transcript T for the third player such that $|S_k(T)| \geq \Omega(|f_{q,d}^{-1}(1)|/\sqrt{d}q^2)$. Where $|f_{q,d}^{-1}(1)|$ is the number of 1's of the function $f_{q,d}$. That is, it is the number of $x, y \in [q]^d$ such that $(x + y)/2$ is also in $[q]^d$. Assume for simplicity that q is even, then $|f_{q,d}^{-1}(1)| \geq q^d \cdot (q/2)^d$. Therefore

$$|S_k(T)| \geq \Omega(q^d \cdot (q/2)^d / \sqrt{d}q^2) \geq \Omega(n^2 / 2^d \sqrt{d}q^2).$$

Taking $d = \log q = \Theta(\sqrt{\log n})$ we get $|S_k(T)| \geq n^2 / 2^{O(\sqrt{\log n})}$. $S_k(T)$ is symmetric, thus Lemma 2.2 follows from Theorem 2.1. ■

We can now describe the relation between the construction of Ruzsa and Szemerédi [21] and that of [1]. Call the construction above A , the simple construction of Section 2.3 B , and the construction of Section 3.1 (providing the graphs similar to [1]) C . The table below compares these constructions.

	A	B	C
Function	Domain: $([q]^d)^3$ Def. rule: $x+y=2z$	Dom.: $([q]^d)^2 \times Z_{q,d}$ Def. rule: $x+y=2z$	Dom.: $([q]^d)^2 \times Z_{q,d}$ Def. rule: $x+y=2z$
Protocol idea	Third player sends $\ x - y\ _2^2$.	Third player sends $\ x - y\ _2^2$.	Third player sends some bits of $\ x - y\ _2^2$, then the first two players compute the rest.
Number of vertices	$n = q^d$	$n = q^d$	$n = q^d$
Edge density	$2^{-O(\sqrt{\log n})}$	$\Omega(\log \log n / \log^\epsilon n)$ for any constant $\epsilon > 1/2$	$1-o(1)$
Number of matchings	$\Theta(n)$	$n^{1+O(1/\log \log n)}$	$n^{1+O(1/\log \log n)}$

References

- [1] N. Alon, A. Moitra and B. Sudakov, Nearly complete graphs decomposable into large induced matchings and their applications, Proc. of the 44th ACM STOC (2012), 1079-1089. Also: J. Eur. Math. Soc. (JEMS) 15 (2013), no. 5, 1575–1596.

- [2] N. Alon and J. Spencer, *The Probabilistic Method* (4th edition), Wiley Interscience, 2016.
- [3] F. A. Behrend. On sets of integers which contain no three terms in arithmetic progression, *Proc. National Academy of Sciences USA* 32 (1946), 331–332.
- [4] Y. Birk, N. Linial, and R. Meshulam. On the uniform-traffic capacity of single-hop interconnections employing shared directional multichannels. *IEEE Transactions on Information Theory*, 39(1):186–191, 1993.
- [5] F. R. K. Chung and R. L. Graham, *Erdős on Graphs; his legacy of unsolved problems*, A K Peters, Ltd., 1998.
- [6] F. R. K. Chung and L. Lu, An upper bound for the Turán number $t_3(n, 4)$, *J. Combin. Theory Ser. A* 87 (1999), 381–389.
- [7] P. Erdős, On the combinatorial problems which I would most like to see solved, *Combinatorica* 1 (1981), 25–42.
- [8] P. Erdős, Some problems on finite and infinite graphs, *Logic and combinatorics* (Arcata, Calif., 1985), 223–228. *Contemp. Math.*, 65, Amer. Math. Soc., Providence, RI, 1987.
- [9] P. Erdős, Problems and results in combinatorial analysis and graph theory, *Proceedings of the First Japan Conference on Graph Theory and Applications* (Hakone, 1986), *Discrete Math.* 72 (1988), 81–92.
- [10] P. Erdős, Some of my favourite problems in various branches of combinatorics, *Combinatorics* 92 (Catania, 1992). *Matematiche (Catania)* 47 (1992), no. 2, 231–240.
- [11] P. Erdős and M. Simonovits, Supersaturated graphs and hypergraphs, *Combinatorica* 3 (1983), 181–192.
- [12] J. Fox, A new proof of the graph removal lemma, *Ann. of Math. (2)* 174 (2011), no. 1, 561–579.
- [13] J. Fox and P. Loh, On a Problem of Erdős and Rothschild on Edges in Triangles, *Combinatorica* 32 (2012), no. 6, 619–628.
- [14] W. T. Gowers, Hypergraph regularity and the multidimensional Szemerédi theorem, *Ann. of Math. (2)* 166 (2007), no. 3, 897–946.
- [15] N. G. Hadziivanov and S. V. Nikiforov, Solution of a problem of P. Erdős about the maximum number of triangles with a common edge in a graph, (in Russian), *C. R. Acad. Bulgare Sci.* 32 (1979), no. 10, 1315–1318.
- [16] G. Katona, T. Nemetz, and M. Simonovits, On a problem of Turán in the theory of graphs, *Mat. Lapok* 15 (1964), 228–238.
- [17] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [18] N. Linial, T. Pitassi, and A. Shraibman. On the communication complexity of high-dimensional permutations. *arXiv preprint arXiv:1706.02207*, 2017.
- [19] B. Nagle, V. Rödl and M. Schacht, The counting lemma for regular k -uniform hypergraphs, *Random Structures and Algorithms*, 28 (2006), 113–179.

- [20] V. Rödl and J. Skokan, Regularity lemma for k -uniform hypergraphs, *Random Structures and Algorithms*, 25 (2004), 1–42.
- [21] I. Rusza and E. Szemerédi, Triple Systems with no Six Points Carrying Three Triangles, *Colloquia Mathematica Societatis János Bolyai* (1978), 939–945.
- [22] P. Turán, Research problems. *MTA Mat. Kutató Int. Közl.* 6 (1961), 417–423.
- [23] E. Szemerédi, Regular partitions of graphs. In: *Proc. Colloque Inter. CNRS*, (J. C. Bermond, J. C. Fournier, M. Las Vergnas and D. Sotteau, eds.), (1978), 399–401.