View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Dagstuhl Research Online Publication Server

On the Communication Complexity of High-Dimensional Permutations

Nati Linial[1]
Hebrew University of Jerusalem, Jerusalem, Israel
nati@cs.huji.ac.i

Toniann Pitassi
University of Toronto, Toronto, Canada and IAS, Princeton, U.S.A.
toni@cs.toronto.edu

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

Abstract

We study the multiparty communication complexity of high dimensional permutations in the Number On the Forehead (NOF) model. This model is due to Chandra, Furst and Lipton (CFL) who also gave a nontrivial protocol for the Exactly-n problem where three players receive integer inputs and need to decide if their inputs sum to a given integer n. There is a considerable body of literature dealing with the same problem, where $(\mathbb{N}, +)$ is replaced by some other abelian group. Our work can be viewed as a far-reaching extension of this line of research. We show that the known lower bounds for that group-theoretic problem apply to all high dimensional permutations. We introduce new proof techniques that reveal new and unexpected connections between NOF communication complexity of permutations and a variety of well-known problems in combinatorics. We also give a direct *algorithmic* protocol for Exactly-n. In contrast, all previous constructions relied on large sets of integers without a 3-term arithmetic progression.

2012 ACM Subject Classification Theory of computation → Communication complexity

Keywords and phrases High dimensional permutations, Number On the Forehead model, Additive combinatorics

Digital Object Identifier 10.4230/LIPIcs.ITCS.2019.54

Related Version A full version of the paper is available at `https://arxiv.org/abs/1706.02207`.

1 Introduction

The multiplayer Number On the Forehead (NOF) model of communication complexity was first introduced by Chandra, Furst and Lipton [13]. Here k players need to evaluate a given function $f : [n]^k \to \{0, 1\}$, where we think of f as having k arguments $x_1, \ldots, x_k$, each comprised of $\log n$ bits. The i-th input vector x_i is placed metaphorically on player i's forehead, so that every player sees the whole input but one argument. Players communicate by writing bits on a shared blackboard in order to compute f.

The NOF communication model has turned out to be a fascinating, though exceedingly difficult object of study. Indeed, good lower bounds in the NOF model would resolve several longstanding open problems in complexity theory, such as lower bounds on the size of ACC^0

[1] Supported in part by ERC grant 339096, High-dimensional combinatorics.

© Nati Linial, Toniann Pitassi, and Adi Shraibman;
licensed under Creative Commons License CC-BY

circuits for a natural function in P [40, 23]. They also imply lower bounds for branching programs, time-space tradeoffs for Turing machines [26], and proof complexity lower bounds [8]. The implications of good NOF lower bounds go in other, less expected directions as well. E.g., knowing the communication complexity of specific natural functions, even for $k = 3$, would have profound implications in graph theory and combinatorics. Finally, the search for nontrivial protocols in this area is a wonderful challenge for algorithms designers. There is a short list of such beautiful examples [13, 22] which begs to be extended.

Furthermore, our understanding of NOF communication complexity, even for $k = 3$ players, lags well behind our understanding of the standard model ($k = 2$ players). This gap is usually attributed to the dearth of proof techniques in the NOF setting. In the 2-party setting, many measures of complexity allow us to prove both upper and lower bounds. Such measures include matrix rank, various matrix norms, nonnegative rank, discrepancy, corruption bounds and information complexity. Most of these measures are computationally simple and admit dual characterizations which are very helpful in proving both upper and lower bounds. On the other hand, in the NOF setting for $k \geq 3$, the key combinatorial objects are cylinder intersections (rather than combinatorial rectangles) and tensor norms. These are much more complex, and thus far have resisted a workable characterization.

A case in point is the separation of randomized from deterministic communication complexity. The 2-party *equality function* has a randomized protocol of bounded cost, whereas a simple rank argument shows that every deterministic protocol must incur linear cost. This provides an optimal separation of deterministic and randomized communication complexity [26]. On the other hand, for $k \geq 3$, the best *explicit* separation between nondeterministic and randomized NOF complexity is logarithmic, even though counting arguments yield linear separations [7]. The Exactly-n function is defined as follows: Input $x_1, \ldots, x_k \in [n]$ is accepted iff $\sum_i x_i = n$. In their seminal paper, Chandra, Furst and Lipton [13] conjectured that Exactly-n achieves a strong separation, and also connected the communication complexity of this function to well-known problems in additive combinatorics. But thus far, despite considerable research effort, the lower bounds for Exactly-n are much weaker even than the best (logarithmic) explicit separations.

The main goal of our work is to further investigate the connections between NOF complexity of functions and questions in additive combinatorics, with the hope of stimulating further research to make progress in both directions. A large and rapidly growing body of work establishes interesting relationships between problems in additive combinatorics and complexity theory. For example, the study of expander graphs and extractors, pseudorandomness, and property testing is closely related, some time even synonymous with similar notions in additive combinatorics. Moreover, techniques from complexity theory have been useful in additive combinatorics and vice versa. Some recent examples include the proof of the cap-set conjecture [14, 17] and Dvir's resolution [15] of the finite field Kakeya problem, as well as the beautiful interplay between dense model theorems in additive combinatorics and the notions of boosting and hardcore sets from complexity theory [11, 38, 29].

Here we consider a broad class of functions called high dimensional permutations. We uncover strong connections between the NOF communication complexity of these functions and several fundamental problems in additive combinatorics. Originally defined in [27], a $(k-1)$-dimensional permutation is a function $f : [n]^k \to \{0,1\}$ such that for every index $k \geq i \geq 1$ and for every choice of $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k \in [n]$, there is exactly one value of $x_i \in [n]$ for which $f(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_k) = 1$. This class of functions generalizes many well-studied functions in communication complexity. It is also closely related to many other functions such as the Exactly-n function mentioned above.

We will show that many well-studied problems in NOF complexity are not just related, but are in fact identical or nearly identical to central problems in additive combinatorics. We feel that this mutual relation deserves much more attention, and that progress in this area is likely to greatly advance both domains. Specifically we believe that the study of the communication complexity of high dimensional permutations and related graph functions (defined in [7]) is a worthwhile undertaking that will help us develop new lower bounds proof techniques for the notoriously difficult NOF model. Using these connections, we make modest progress on several upper and lower bounds in NOF communication complexity.

1.1 Our Contributions

As mentioned above, our main goal and contribution is to unveil the strong relationships between the NOF complexity of high dimensional permutation problems and central problems in additive combinatorics and Ramsey theory. Already the founding paper of Chandra, Furst and Lipton [13] makes a connection between the NOF complexity of Exactly-n and the areas of Ramsey theory and additive combinatorics. A more general framework was introduced in [10]: Given an abelian group G and $T \in G$, the function $f^G_{k,T}$ evaluates to 1 on input $x_1, \ldots, x_k \in G$ iff $\sum_i x_i = T$ (this expression is well-defined since G is abelian). The functions $f^G_{k,T}$ are high dimensional permutations. (Note that this holds as well for non-abelian G, though we need to specify the order at which $\prod_i x_i$ is evaluated). Another strong connection is that the Hales-Jewett theorem, a cornerstone of Ramsey theory, can be interpreted in terms of communication complexity [35].

We establish a new and close connection between the NOF communication complexity of high dimensional permutations and dense Ruzsa-Szemerédi graphs. These graphs appear in various contexts in Combinatorics, Computer Science and Information Theory, thus highlighting new connections between communication complexity and these various problems. For example, an efficient deterministic communication protocol for any permutation yields an efficient wiring scheme for shared directional multi-channels. For more on this, see e.g., [12] and [4]. In the classical, $k = 2$ case, monochromatic submatrices play a key role in the theory. For higher k this is replaced by the much more poorly understood monochromatic *cylinder intersection*. Naturally, much of our work here revolves around these complicated objects. However, in certain simple cases we are able to get a grip on the largest size of a cylinder intersection that contains only 1-inputs of f. As we show, in this case knowledge of this quantity essentially determines the NOF communication complexity of f (see more on this in the next section). The case in question is $k = 3$ and the group $G = \mathbb{Z}_2^n$. As we show, the size of the largest cylinder intersection containing only 1-inputs of f is the largest cardinality of a subset $W \subseteq \mathbb{Z}_4^n$ such that for every three distinct members $\mathbf{x}, \mathbf{y}, \mathbf{z} \in W$ there is an index $1 \leq i \leq n$ for which $(x_i, y_i, z_i) \notin X$, where

$$X = \{(0,0,0), (1,1,1), (2,2,2), (3,3,3), (0,1,2), (1,0,3), (2,3,0), (3,2,1)\}.$$

This parameter may seem artificial, but in fact, this framework includes several important problems in combinatorics, for different choices of X. Thus, if we take $X := \{(0,0,0), (1,1,1), (2,2,2), (3,3,3), (0,1,2)\}$, then this becomes precisely the density Hales-Jewett problem, solved in [19]. Also, if X is comprised of all triplets $(a,b,c) \in \mathbb{Z}_4^3$ with $a + c = 2b$, we arrive at the cap-set problem for $\mathbb{Z}_4^n$ which was recently settled in breakthrough papers by Croot, Lev, and Pach, and by Ellenberg and Gijswijt [14, 17]. In the next subsection we list our new results that stem from these connections.

1.1.1 Upper Bounds

We give a new algorithm for Exactly-n as well as several other instances of $f^G_{k,T}$. All previous upper bounds for these functions crucially depend on Behrend's famous construction [9] of a large set of integers with no 3-term arithmetic progressions. This yields a large monochromatic cylinder intersection, and a simple probabilistic translation lemma then shows how to cover the whole space by monochromatic cylinder intersections, thus providing an efficient protocol. To show that this indeed yields a large monochromatic cylinder intersection, we appeal to the notion of *corner-free* sets [2, 36] which here, too, plays a key role. We cannot realistically hope to improve the bounds by finding a construction better than Behrend's, in view of the many such failed attempts throughout the past 70 years (but note [16]). However, Behrend's construction is actually more than we need. The solution of $f^G_{k,T}$ only requires corner-free sets. That is, 3-term AP freeness implies corner-freeness, but we do not expect that the two concepts are equivalent. We take a first step in this direction and give a new algorithm which is not dependent on 3-term AP freeness. We hope that this indicates a viable approach to improved protocols for the Exactly-n function. We also describe a nontrivial protocol for the $f^G_{3,T}$ problem for $G = \mathbb{Z}_2^n$.

1.1.2 Lower Bounds

We give a counting argument which shows that almost every k-dimensional permutation has communication complexity $\Omega(\frac{\log n}{k})$. Clearly, up to the $\frac{1}{k}$ factor, this is as high as this quantity can get. Our proof relies on a recent lower bound of Keevash [25] on the number of high-dimensional permutations. This method resembles the counting argument for graph functions of [7], which does not apply, though, to permutations.

Regarding bounds on explicit functions, we prove a weak upper bound on the size of a 1-monochromatic cylinder intersection for any permutation (in fact our result holds for a wider family of functions that we call linjections). This bound uses a graph theoretic characterization of the communication complexity of permutations, connecting it also to Ruzsa-Szemerédi graphs. Not unexpectedly, our proof mirrors a similar result for Ruzsa-Szemerédi graphs: Solymosi [36] showed that the multidimensional Szemeréedi theorem follows from the triangle removal lemma. We adapt Solymosi's proof to our context. The main tools in the proof are thus the graph and the hypergraph removal lemmas.

We note that previous results were limited to the $f^G_{k,T}$ function for abelian groups with many factors, whereas ours works for general permutations. To emphasize the significance of the last point, consider the NOF complexity of following three classes of functions: (i) Permutations that come from Abelian groups, (ii) Those that come from general groups, (iii) Latin squares. We consider each such function up to an arbitrary renaming of rows and columns. The sizes of these three classes differ very substantially. For a given order n the size of the relevant class is (i) $\exp(O(\sqrt{\log n}))$, (ii) At most $\exp((\frac{2}{27} + o(1)) \log^3 n)$, and (iii) $((1 + o(1))\frac{n}{e^2})^{n^2}$.

For $k = 3$ we can say more: The communication complexity of every 2-dimensional permutation $[n]^3 \to \{0, 1\}$ is $\Omega(\log \log \log n)$. This extends the lower bound of [10] from the realm of abelian groups to all permutations. The proof of the this lower bound uses only elementary counting arguments, and is closely related to the result of [20] on monochromatic corners on the integer grid.

The above lower bound also implies a result of Meshulam that was derived toward the study of shared directional multi-channels. Meshulam's result appears as Proposition 4.3 in [4], where further background can be found.

1.2 Related Work

The NOF model was introduced by Chandra, Furst and Lipton in [13]. One of the functions they consider is Exactly-n : $[n]^3 \to \{0,1\}$. For $x, y, z \in [n]$, we let Exactly-n $(x, y, z) = 1$ if and only if $x + y + z = n$. Surprisingly, they proved that the communication complexity of this function is only $O(\sqrt{\log n})$, but their proof yields no explicit protocol. Although this function is not a permutation, the proofs go through as well for the mod$n$ permutation [10], and thus far this is the most efficient protocol found for any permutation. The protocol of [13] is based on Behrend's famous construction [9] of a large subset of $[n]$ with no three-term arithmetic progression. In addition, they prove an inexplicit lower bound of $\omega_n(1)$ on the complexity of Exactly-n . This is based on Gallai's result [21, p. 38] that every finite coloring of a Euclidean space contains a monochromatic homoteth of every finite set in that space.

Beigel, Gasarch and Glenn [10] considered the more general $f^G_{k,T}$ problem, where G is an abelian group, T is an element of G and $k \geq 2$ an integer. In this scenario k players need to decide whether $x_1 + x_2 + \cdots + x_k = T$. They show that the communication complexity of $f^G_{3,T}$ is at least $\Omega(\log\log\log n)$ for every abelian group G and any $T \in G$. For the case $G = \mathbb{Z}_n$, this follows as well from [20] and a recent result of Shkerdov [34] also yields a similar lower bound for every abelian group G. For general $k \geq 3$ and for an abelian group G that is the product of t cyclic groups, it is shown in [10] that the deterministic NOF complexity of $f^G_{k,T}$ is $\omega_t(1)$. The proof is by reduction to a lower bound from [37], that is based on the Hales-Jewett Theorem (see [21]).

Note that $f^G_{k,T}$ can be defined as well in non-abelian groups G. Namely, $f^G_{k,T}(x_1, \ldots, x_k) = 1$ iff $x_1 \cdot x_2 \cdot \ldots \cdot x_k = T$, where now the order of multiplication matters. Note also that the function $f^G_{k,T}$ is a permutation for every group G, every $T \in G$ and $k \geq 2$.

As mentioned above, [7] studies graph functions and give a nonexplicit strong separation between randomized and deterministic NOF complexity. To be precise, this counting argument shows that most graph functions $f : [n]^{k-1} \times [N] \to \{0,1\}$ with $N \cong \sqrt{\frac{n}{k}}$ have deterministic communication complexity $\Omega(\log \frac{n}{k})$. Still, even for $k = 3$ it remains open to find *explicit* graph functions with high deterministic communication complexity. Currently, the best lower bound on the deterministic communication complexity of a graph function $f : [n]^{k-1} \times [N] \to \{0,1\}$ for $k \geq 3$ is $\Omega(\log\log n)$ proved in [7] (using also results from [5]). We note that the discrepancy method, used to establish NOF lower bounds (e.g., [6]), cannot be utilized here since it also applies to randomized communication complexity.

Lastly, we comment on the Hales-Jewett theorem, a pillar of Ramsey theory. It was previously applied in the study of the combinatorial problems mentioned above. It turns out that this theorem has an equivalent formulation in the language of communication complexity [35], and is tightly coupled with the NOF multiparty communication complexity of high dimensional permutations.

2 Basics

2.1 NOF Communication Complexity

In the Number On the Forehead (NOF) multiparty communication complexity game, k players collaborate to compute a function $f : X_1 \times \ldots \times X_k \to \{0,1\}$. Usually, $X_i = [n]$ for all $i \in [k]$, but we also consider occasionally a variation where the last player is exceptional and $X_k = [N]$ for some integer N that is not necessarily equal to n.

For $(x_1, \ldots, x_k) \in X_1 \times \ldots \times X_k$, and for each $i \in [k]$, player i receives $x^{-i} \in X_1 \times \ldots \times X_{i-1} \times X_{i+1} \times \ldots \times X_k$; that is, all but x_i. The players exchange bits according to

an agreed-upon protocol, by writing them on a publicly visible blackboard. The protocol specifies, for every possible blackboard contents, whether or not the communication is ongoing. It shows the final output when the communication is over, and shows the next player to speak if the communication is still ongoing. The protocol also specifies what each player writes as a function of the blackboard contents and of the inputs seen by that player. The *cost* of the protocol is the maximum number of bits written on the blackboard.

The *deterministic communication complexity* of f, $D_k(f)$, is the minimum cost of a deterministic protocol for f that always outputs the correct answer. A randomized protocol of cost c is just a distribution over deterministic protocols each of cost at most c. For $0 \leq \epsilon < 1/2$, the *randomized communication complexity* of f, $R_{k,\epsilon}(f)$, is the minimum cost over randomized protocols such that for every input, err with probability at most ϵ (over the distribution of deterministic protocols).

In the $k = 2$ players case, the key combinatorial objects of study are combinatorial rectangles: Every cost-c communication protocol for $f : X_1 \times X_2 \to \{0,1\}$ partitions $X_1 \times X_2$ into 2^c monochromatic combinatorial rectangles. For k-party NOF communication, a cost-c protocol induces a partition of $X_1 \times \ldots \times X_k$ into 2^c *monochromatic cylinder intersections*:

▶ **Definition 1.** A *cylinder in dimension i* is a subset $S \subseteq \prod X_i$ such that if $(x_1, \ldots, x_k) \in S$, then $(x_1, \ldots, x_{i-1}, x'_i, x_{i+1}, \ldots, x_k) \in S$ for all x'_i. A *cylinder intersection* is a set of the form $\cap_{i=1}^{k} T_i$, where T_i is a cylinder in dimension i.

2.2 Graph Functions, Permutations and Linjections

▶ **Definition 2.** The line $L \subseteq [n]^k$, defined by a pair (a, i), where $a \in [n]^{k-1}$, $i \in [k]$, is the set of vectors $v \in [n]^k$ such that $v^{-i} = a$ and v_i is an arbitrary element in $[n]$.

▶ **Definition 3.** A function $f : [n]^{k-1} \times [N] \to \{0,1\}$ is a *graph function* if for every $(x_1, \ldots, x_{k-1})$ there is a unique $b \in [N]$ such that $f(x_1, \ldots, x_{k-1}, b) = 1$. In other words, every line in the kth dimension, $L = (a, k)$, intersects $f^{-1}(1)$ in exactly one point.

Associated with every graph function $f : [n]^{k-1} \times [N] \to \{0,1\}$ is a map $A(f) : [n]^{k-1} \to [N]$, where $A(f)(x_1, \ldots, x_{k-1}) = y$ if and only if $f(x_1, \ldots, x_{k-1}, y) = 1$. We consider the two as one and the same object and freely switch back and forth between the two descriptions.

▶ **Definition 4.** Let $f : [n]^{k-1} \times [N] \to \{0,1\}$ be a graph function. We denote by $\alpha_k(f)$ the largest size of a cylinder intersection that is contained in $f^{-1}(1)$. In other words, the largest cardinality of 1-monochromatic cylinder intersection with respect to f. Also, let $\chi_k(f)$ be the least number of 1-monochromatic cylinder intersections whose union is $f^{-1}(1)$. We omit the subscript k when it is clear from context.

Given a graph function f, the measure $\chi(f)$ corresponds to the nondeterministic NOF communication complexity of f, since it is a covering of the 1's of f by cylinder intersections [26]. In general, the nondeterministic NOF communication complexity of a Boolean function can be much smaller than the deterministic complexity – in fact, for the set disjointness function, nondeterministic complexity is logarithmic in the deterministic complexity (for constant k). However, graph functions are special; the following lemma shows that for graph functions, the two notions basically coincide. The proof is an adaptation of a proof from [13]; see also [10, 7] for similar arguments.

▶ **Theorem 5.** *For every graph function $f : [n]^{k-1} \times [N] \to \{0,1\}$, $\log \chi_k(f) \leq D_k(f) \leq \lceil \log \chi_k(f) \rceil + k - 1$.*

A communication protocol in which players write only one message on the board, of arbitrary length is called a *one-way* protocol. This applies to the protocol in the proof of Theorem 5. The restriction to a single message per player may make one-way protocols much weaker than standard protocols [30, 5]. However for graph functions, one-way protocols and regular protocols are equally powerful:

▶ **Corollary 6.** *For every graph function* $f : [n]^{k-1} \times [N] \to \{0,1\}$, $D_k(f) \leq D_k^1(f) \leq D_k(f) + k$ *where* $D_k^1(f)$ *is the one-way communication complexity of* f.

Implicit in the proofs of the above statements is the fact that for graph functions, monochromatic cylinder intersections can be nicely characterized by forbidden (dual) objects called *stars*, which we define next. We will see in the next section that stars are very closely connected to corners (and higher dimensional generalizations) in Ramsey theory.

▶ **Definition 7.** A *star* $Star(\mathbf{x}, \mathbf{x}')$ is a subset of $[n]^{k-1} \times [N]$ of the form

$$\{(x'_1, x_2, \ldots, x_k), (x_1, x'_2, \ldots, x_k), \ldots, (x_1, x_2, \ldots, x'_k)\},$$

where $x_i \neq x'_i$ for each i. We refer to $\mathbf{x} = (x_1, x_2, \ldots, x_k)$ as the star's *center*, and note that the center does *not* belong to the star.

▶ **Lemma 8.** *Let* $f : [n]^{k-1} \times [N] \to \{0,1\}$ *be a graph function, and let* $S \subseteq f^{-1}(1)$. *Then* S *is a (*1*-monochromatic) cylinder intersection with respect to* f *if and only if it does not contain a star.*

Next we define high dimensional permutations and linjections.

▶ **Definition 9.** A $(k-1)$-dimensional permutation of order-n is a map $f : [n]^k \to \{0,1\}$ with the property that for every line $L = (a, i)$ in $[n]^k$, $\left|L \cap f^{-1}(1)\right| = 1$.

In other words, f is a permutation function if and only if every line contains a unique 1 entry. This property is easily seen to be equivalent to the property that for every choice of $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k \in [n]$, there is exactly one value, $A_i(x^{-i})$ for $x_i \in [n]$ such that $f(x_1, \ldots, x_{i-1}, A_i(x^{-i}), x_{i+1}, \ldots, x_k) = 1$.

▶ **Example 10.** For the sake of gaining better intuition we often consider the important special case $k = 3$. This is insightful, since 2-dimensional permutations $f : [n]^3 \to \{0,1\}$ are synonymous with Latin squares. In this case $A(f)$ is an $n \times n$ matrix with entries in $[n]$ where every row and column contains each element in $[n]$ exactly once. Here we see an elementary but important connection with additive combinatorics; stars coincide with the well-studied notion of *corners* [2, 36]. A star is a triplet of entries in $f^{-1}(1)$, $(x, y, z'), (x', y, z), (x, y', z)$, which corresponds to the "corner" or "A-star", $(x, y), (x', y), (x, y')$, where $A(x', y)$ and $A(x, y')$ have the same value (z), but $A(x, y)$ has a different value z'.

▶ **Example 11.** High dimensional permutations generalize the family of functions $f_{k,T}^G$ for abelian groups G. For this communication problem, each player receives (on his/her forehead) an element $x_i \in G$, and they want to decide whether or not $x_1 + \ldots + x_k = T$, that is, whether the sum of the elements is exactly T.

We can further generalize the notion of a permutation function as follows.

▶ **Definition 12.** A *linjection* is a graph function $f : [n]^{k-1} \times [N] \to \{0,1\}$ with $N \geq n$ such that $|f^{-1}(1)| = n^{k-1}$ and every line contains **at most** one point at which $f = 1$. A function f is a linjection if and only if the restriction of $A(f)$ to any line is an injection.

Injections are graph functions where $N \geq n$, (with permutation functions corresponding to $n = N$), but not vice versa. Determining the least possible communication complexity of a linjection in certain dimensions is an interesting and challenging problem. Henceforth, we use and study the following two notions.

▶ **Definition 13.** Define $\alpha_k(n, N) = \max_f \alpha_k(f)$, and $\chi_k(n, N) = \min_f \chi_k(f)$, both taken over all linjections $f : [n]^{k-1} \times [N] \to \{0, 1\}$.

Note that $\chi_k(n, N) \geq n^{k-1}/\alpha_k(n, N)$.

3 High Dimensional Permutations and Additive Combinatorics

3.1 A Graph-theoretic Characterization

In this section we give a new characterization of α_k which will turn out to be a variant of the maximum density of Ruzsa-Szemerédi graphs. We start with the case $k = 3$.

Recall that we can view a linjection $f : [n]^2 \times [N] \to \{0, 1\}$ as an $n \times n$ matrix, $A = A(f)$ with entries from $[N]$. Alternatively we view it as a tripartite graph $G(A)$ with parts $R = [n], C = [n]$ and $W \subseteq [N]$. Its edge set is defined as follows: for every triple $(x, y, b) \in f^{-1}(1)$, we add the triangle $(x, y), (y, b), (x, b)$, $x \in R$, $y \in C$, $b \in W$ to $G(A)$. In particular, $R \cup C$ span a complete bipartite subgraph of $G(A)$ and (i, b), $i \in R$, $b \in W$ is an edge iff there is a b entry in row i of A, likewise for columns.

Let us consider the triangles $< x, y, b >$, $x \in R, y \in C, b \in W$, in $G(A)$. A triangle $< x, y, b >$ in G is *trivial* if $A(x, y) = b$. However, there can also be nontrivial (induced) triangles in G, which correspond to centers of stars. We define a *G-star* to be a triple of triangles in G of the form $< x, y, b' >, < x', y, b >, < x, y', b >$. The point is that while these (trivial) triangles are edge-disjoint, their union contains the additional induced triangle $< x, y, b >$. Define $\overline{\alpha}(G)$ to be the largest cardinality of a family of edge-disjoint triangles that contains no G-star. In other words, a family of edge-disjoint triangles the union of which contains no additional triangle.

Let $\overline{\alpha}(n, N) = \max_G \overline{\alpha}(G)$ where the maximum is over subgraphs of $K_{n,n,N}$. Then:

▶ **Theorem 14.** *For every two integers $n, N > 0$, if $n \leq N$ then $\alpha_3(n, N) \leq \overline{\alpha}(n, N)$. If $N \geq 2n - 1$, then $\alpha_3(n, N) = \overline{\alpha}(n, N)$.*

Proof. We show first that $\alpha_3(n, N) \leq \overline{\alpha}(n, N)$. Let $f : [n] \times [n] \times [N] \to \{0, 1\}$ be a linjection and let $S \subseteq [n] \times [n] \times [N]$ be a star-free subset of $f^{-1}(1)$. We prove the claim by constructing a G-star-free family T of $|S|$ edge-disjoint triangles in $G = G(A(f))$. Let

$$T = \{< x, y, b > | (x, y, b) \in S\}.$$

The claim follows, since stars $\{(x', y, b), (x, y', b), (x, y, b)\}$ correspond to G-stars in T. Next we prove the reverse inequality $\alpha_3(n, N) \geq \overline{\alpha}(n, N)$ when $N \geq 2n - 1$.

Given a G-star-free family T of edge-disjoint triangles in a subgraph G of $K_{n,n,N}$, we find a linjection $A : [n] \times [n] \to [N]$ that contains an A-star-free subset $S \subset [n]^2$ of size $|T|$. In the proof we actually first construct S and only then proceed to define A in full.

We define S to be the projection of T to its first two coordinates. Namely,

$$S = \{(x, y) \mid \ < x, y, b > \in T \text{ for some } b\}.$$

To define A, we first let $A(x, y) = b$ for every $< x, y, b > \in T$.

Since T is G-star-free, it follows that S is A-star-free. What is missing is that A is only partially defined. We show that when $N \geq 2n-1$ this partial definition can be extended to a linjection. Since the triangles in T are edge-disjoint it follows that in the partially defined A, no value appears more than once in any row or column. It remains to define A on all the entries outside of S and maintain this property. Indeed this can be done entry by entry. At worst there are $2n-2$ values that are forbidden for the entry of A that we attempt to define next, and therefore there is always an acceptable choice. ◀

General k. The construction for general k is a natural extension of the case $k=3$. We associate with every linjection $A : [n]^{k-1} \to [N]$ a k-partite $(k-1)$-uniform hypergraph $H(A)$. The parts of the vertex set are denoted $Q_1, \ldots, Q_{k-1}$ and W. Each Q_i is a copy of $[n]$ and, as above, W is the range of A. There is a complete $(k-1)$-partite hypergraph on the $k-1$ parts $Q_1, \ldots, Q_{k-1}$. Given $x_1 \in Q_1, \ldots, x_{i-1} \in Q_{i-1}, x_{i+1} \in Q_{i+1}, \ldots, x_{k-1} \in Q_{k-1}$ and $w \in W$, we put the hyperedge $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}, w$ in $H(A)$ iff there is a (necessarily unique) $x_i^* \in [n]$ for which $A(x_1, \ldots, x_{i-1}, x_i^*, x_{i+1}, \ldots, x_{k-1}) = w$.

We proceed to investigate *cliques* in $H(A)$, i.e., sets of k vertices, every $k-1$ of which form an edge. For $k=3$, we distinguished between those triangles in $G(A)$ that correspond to an entry in $[n]^2$ and those that form a star, and a similar distinction applies for general k.

It is easy to see that if $A(x_1, \ldots, x_{k-1}) = w$, then $x_1, \ldots, x_k, w$ from a clique. Such a clique is considered *trivial*. In contrast, $x_1, \ldots, x_{k-1}, w$ is a nontrivial clique iff for every i there exists an $x_i' \neq x_i$ such that $A(x_1, \ldots, x_{i-1}, x_i', x_{i+1}, \ldots, x_{k-1}) = w$.

As above, we define for $H = H(A)$ the parameter $\overline{\alpha}_k(H)$. It is the largest size of a family K of cliques in H such that: (i) No two share a hyperedge, and (ii) The hypergraph comprised of all cliques in K contains no additional cliques. Let $\overline{\alpha}_k(n, N) = \max_H \overline{\alpha}_k(H)$ over all k-partite $(k-1)$-uniform hypergraphs H. Then

▶ **Theorem 15.** *For every two integers $n \leq N$, $\alpha_k(n, N) \leq \overline{\alpha}_k(n, N)$, and if $N > (k-1)(n-1)$ then $\alpha_k(n, N) = \overline{\alpha}_k(n, N)$.*

The proof is similar to the proof of Theorem 14 and appears in the full paper.

As the proofs show, $\overline{\alpha}_k(n, N)$ is the largest cardinality of a star-free subset of $[n]^{k-1} \times [N]$ that meets every line in $[n]^{k-1} \times [N]$ at most once. To qualify for $\alpha_k(n, N)$ this subset must, in addition, be extendable to a linjection, so clearly $\overline{\alpha}_k(n, N) \geq \alpha_k(n, N)$. We wonder whether this additional requirement creates a substantial difference between the two parameters. Specifically, how are $\overline{\alpha}_k(n, N)$ and $\alpha_k(n, N)$ related in the range $n \leq N \leq (k-1)(n-1)$? These two parameters need not be equal in this range, since $\alpha_3(4, 4) = 8$ and $\overline{\alpha}_3(4, 4) = 9$, as we show in Section 4.3.

Connection to Ruzsa-Szemerédi Graphs. A graph is called an (r, t)-Ruzsa-Szemerédi graph if its edge set can be partitioned into t edge-disjoint induced matchings, each of size r. These graphs were introduced in 1978 and have been extensively studied since then. Of particular interest are dense Ruzsa-Szemerédi graphs, with r and t large, in terms of n, the number of vertices. Such graphs have applications in Combinatorics, Complexity theory and Information theory. Also, there are several known interesting constructions, relying on different techniques.

Let G be a tripartite graph with parts R, C, W of cardinalities n, n, N respectively. Let T be a G-star-free family of edge disjoint triangles in G. Let F be the bipartite graph with parts R and C where there is an edge between $r \in R$ and $c \in C$ iff there is some $b \in W$ such that $(r, c, b) \in T$. Then F is the union of at most N edge-disjoint induced matchings, since all the edges that correspond to a given $b \in W$ form an induced matching.

This construction can easily be reversed: Let F be a subgraph of $K_{n,n}$ that is the union of N edge disjoint induced matchings, with a total of $\overline{\alpha}$ edges. We can construct a tripartite G (a subgraph of $K_{n,n,N}$) that contains a family of $\overline{\alpha}$ pairwise disjoint triangles, and has no G-stars. We conclude that

▶ **Observation 16.** *Let $n \leq N$ be positive integers, then $\overline{\alpha}_3(n, N)$ is the largest number of edges in a union of N edge-disjoint induced matchings in $K_{n,n}$.*

This observation exhibits a strong connection between (i) The problem of constructing dense (r, t)-Ruzsa-Szemerédi graphs, and (ii) The construction of a large star-free subset $S \subseteq [n] \times [n] \times [t]$ that meets every line at most once. The two problems differ only slightly. In one, the underlying graph is bipartite and in the other all induced matching must have the same cardinality. But these differences can be bridged quite easily, as observed in the following lemma.

▶ **Lemma 17.** **1.** *(r, t)-Ruzsa-Szemerédi graphs on n vertices imply $\overline{\alpha}_3(\frac{n}{2}, t) \geq \frac{rt}{2}$.*
2. *$\overline{\alpha}_3(n, t) \geq rt$ implies that there exists a $(\frac{r}{2}, t)$-Ruzsa-Szemerédi graph on n vertices.*

Proof. For the first claim, let $G = (V, E)$ be a (r, t)-Ruzsa-Szemerédi graph on n vertices, and let $E_1, E_2, \ldots, E_t$ be the partition of E into induced matchings. We can find (e.g., by a random choice) a subset $A \subset V$ of $\lfloor \frac{n}{2} \rfloor$ vertices, so that at least $|E|/2$ edges are in the cut $C = (A, \bar{A})$. Also, $C \cap E_1, C \cap E_2, \ldots, C \cap E_t$ is a partition of the edges of the bipartite graph $(A, \bar{A}, C)$ into t disjoint induced matchings. Therefore, $\overline{\alpha}_3(\frac{n}{2}, t) \geq \frac{rt}{2}$.

For the second part, suppose that $\overline{\alpha}_3(n, t) \geq rt$. Namely, there is a collection of disjoint induced matchings $M_1, \ldots M_t \subseteq E(K_{n,n})$ with $\sum_1^t |M_i| \geq rt$. We split each M_i into $\lfloor \frac{2|M_i|}{r} \rfloor$ sets of $\geq r/2$ edges each. Note that $\sum_1^t a_i \geq rt$ implies that $\sum_1^t \lfloor \frac{2a_i}{r} \rfloor \geq t$ and a subset of an induced matching is an induced matching, so we finally have a family of at least t disjoint induced matchings each of size $\frac{r}{2}$. ◀

3.1.1 Application to Shared Directional Multi-channels

Ruzsa-Szemerédi graphs have various applications in several fields [36, 4, 33, 3, 24, 12]. In [12] they are applied to Information Theory, and the study of *shared directional multi-channels*, a subject that is strongly related to communication complexity. Such a channel is comprised of a set of inputs and a set of outputs. to which are connected transmitters and receivers respectively. Associated with each input is a set of outputs, that receive any signal placed at that input. A message is received successfully at an output of the channel if and only if it is addressed to the receiver connected to that output and no other signals concurrently reach that output. Therefore, when communicating over a shared channel, we want the edges (corresponding to messages sent in one round) to form an induced matching. The challenge is to partition $K_{n,n}$ into families of pairwise disjoint induced matchings. The number of parts correspond to the number of receivers allowed at each output, and the number of matchings in each partition corresponds to the number of rounds.

The relation to communication complexity is as follows: A c-bit communication protocol for any linjection $A : [n] \times [n] \to [N]$ induces a partition of $K_{n,n}$ into c such families of disjoint induced matchings. Thus, such a communication protocol, gives an N round protocol for the shared directional multi-channel, with c receivers per station, and vice-versa.

In constructing a shared directional multi-channel, we seek to minimize the number of rounds required for a given number of transmitters. Alon, Moitra, and Sudakov [4] showed that for any $\epsilon > 0$ there is partition of $K_{n,n}$ into at most $2^{O(\frac{1}{\epsilon})}$ graphs each of which is a family of at most $O(n^{1+\epsilon})$ induced matchings. This gives an $O(n^{1+\epsilon})$ round protocol for shared directional multi-channel with $2^{O(\frac{1}{\epsilon})}$ receivers.

Translated to the language of NOF protocols and combining with Corollary 34 (see Section 5.3 in the sequel), we conclude:

▶ **Theorem 18.** *For all $\epsilon > 0$ and all large enough n, $2^{O(\frac{1}{\epsilon})} \geq \chi_3(n, n^{1+\epsilon}) \geq \Omega(\log \frac{1}{\epsilon})$.*

3.2 A Characterization of $\alpha_k(f_{k,T}^{\mathbb{Z}_2^n})$

In this section we focus on the problem $f_{k,T}^G$ for the abelian group $\mathbb{Z}_2^n$. In other words, we study the permutation $f_{k,T}^{\mathbb{Z}_2^n}$. We give an alternative characterization of $\alpha_3(f_{3,T}^{\mathbb{Z}_2^n})$ which brings forth the relation between this problem and several known combinatorial objects. The complexity of $f_{k,T}^{\mathbb{Z}_2^n}$ is independent of T, so we will omit the subscript T in this section. Also, throughout this section we let $A_k^G = A(f_k^G)$.

Let $X \subset \mathbb{Z}_4^3$. We call a subset of $W \subseteq \mathbb{Z}_4^n$ *X-free* if for every three distinct members $\mathbf{x}, \mathbf{y}, \mathbf{z} \in W$ there is an index $1 \leq i \leq n$ for which $(x_i, y_i, z_i) \notin X$.

▶ **Theorem 19.** *Let*

$$X = \{(0,0,0), (1,1,1), (2,2,2), (3,3,3), (0,1,2), (1,0,3), (2,3,0), (3,2,1)\} \subset \mathbb{Z}_4^3,$$

then $\alpha_3(A_3^{\mathbb{Z}_2^n})$ is the largest cardinality of an X-free subset of $\mathbb{Z}_4^n$.

Proof. Recall that $\alpha_3(A_3^{\mathbb{Z}_2^n})$ is the largest cardinality of an A_n-star free subset of $(\mathbb{Z}_2^n)^2$, where $A_n = A_3^{\mathbb{Z}_2^n}$. So it suffices to find a bijection ψ from $(\mathbb{Z}_2^n)^2$ to $\mathbb{Z}_4^n$ such that $S \subseteq (\mathbb{Z}_2^n)^2$ is mapped to an X-free set if and only if S is A_n-star free.

We define ψ for $n = 1$ and extend is entry-wise to a mapping from $(\mathbb{Z}_2^n)^2$ to $\mathbb{Z}_4^n$. The definition for $n = 1$ is as follows: $\psi(0,0) = 0$, $\psi(0,1) = 1$, $\psi(1,0) = 2$ and $\psi(1,1) = 3$.

We need to show that if $(x_1, y_1), (x_2, y_2), (x_3, y_3) \in (\mathbb{Z}_2^n)^2$ is a A_n-star, then every coordinate in $(\psi(x_1, y_1), \psi(x_2, y_2), \psi(x_3, y_3))$ belongs to X, and vice versa. Since the map ψ is defined coordinate-wise it suffices to check this for $n = 1$. A triple $(x_1, y_1), (x_1, y_1 + d), (x_1 + d', y_1)$ is a (trivial or non-trivial) star in A_1 iff $x_1 + (y_1 + d) = (x_1 + d') + y_1$, i.e., $d = d'$, and thus an A_1-star is a triple of the form $(x_1, y_1), (x_1, y_1 + d), (x_1 + d, y_1)$. If $d = 0$ then obviously $(\psi(x_1, y_1), \psi(x_1, y_1 + d), \psi(x_1 + d, y_1)) \in \{(0,0,0), (1,1,1), (2,2,2), (3,3,3)\} \subset X$. When $d = 1$ there are four cases to check:

1. $x_1 = 0$ and $y_1 = 0$ then $(\psi(x_1, y_1), \psi(x_1, y_1 + d), \psi(x_1 + d, y_1)) = (0,1,2) \in X$.
2. $x_1 = 0$ and $y_1 = 1$ then $(\psi(x_1, y_1), \psi(x_1, y_1 + d), \psi(x_1 + d, y_1)) = (1,0,3) \in X$.
3. $x_1 = 1$ and $y_1 = 0$ then $(\psi(x_1, y_1), \psi(x_1, y_1 + d), \psi(x_1 + d, y_1)) = (2,3,0) \in X$.
4. $x_1 = 1$ and $y_1 = 1$ then $(\psi(x_1, y_1), \psi(x_1, y_1 + d), \psi(x_1 + d, y_1)) = (3,2,1) \in X$.

On the other hand it is not hard to check that for each $(a, b, c) \in X$ the triplet $\psi^{-1}(a), \psi^{-1}(b), \psi^{-1}(c)$ is a star in A_1 or $a = b = c$. This proves the claim. ◀

Fix an integer $s \geq 2$ and let $HJ(n, s)$ denote the largest size of a Y_s-free subset of $[s]^n$, where Y_s is the following set of s-tuples: $\{(1, \ldots, s)\} \cup \{(i, i, \ldots, i) | i = 1, 2, \ldots, s\}$. The density Hales-Jewett theorem states that $HJ(n, s) = o(s^n)$ for every fixed s [19, 31]. Theorem 19, and the observation that the first three coordinates of the 4-tuples in Y_4 all belong to X, imply that $\alpha_3(A_3^{\mathbb{Z}_2^n}) \leq HJ(n, 4)$.

The cap-set problem for $\mathbb{Z}_4^n$ also belongs to the same circle of problems. It concerns the largest size of an arithmetic-triple-free set in $\mathbb{Z}_4^n$. We mention in passing the recent breakthrough [14, 17] in this area which showed that this size is at most $4^{(\gamma + o(1)) \cdot n}$ with $\gamma \approx 0.926$. Let $Z \subset \mathbb{Z}_4^3$ be the set of all ordered triplets $(a, b, c) \in \mathbb{Z}_4^3$ satisfying $a + c = 2b$. The cap set problems concerns exactly the largest possible cardinality of a Z-free subset of $\mathbb{Z}_4^n$. Since $X \subset Z$ it follows that this size is bounded by $\alpha_3(A_3^{\mathbb{Z}_2^n})$.

The proof of Theorem 19 extends verbatim to general $k \geq 3$. It yields a subset $X \subset \mathbb{Z}_{2^{k-1}}^k$ such that $\alpha_k(A_k^{\mathbb{Z}_2^n})$ is the largest cardinality of an X-free subset of $\mathbb{Z}_{2^{k-1}}^n$.

By taking X that includes all vectors $(a, a, \ldots, a) \in \mathbb{Z}_{2^{k-1}}^k$ for $a \in \mathbb{Z}_{2^{k-1}}$ and the vector $(0, 1, 2, 4, \ldots, 2^{k-2})$ we can maintain the relation between $\alpha_k(A_k^{\mathbb{Z}_2^n})$ and the density Hales-Jewett theorem for every k.

4 Upper Bounds

4.1 An Algorithmic Protocol for Exact-T over $\mathbb{Z}^d$

The aim of this section is to give the first algorithmic protocol for Exactly-n as well as other Exact-T functions. Our protocol is explicit, and does not rely on a construction of a large set without a 3-term AP. We only appeal to the elementary fact that no sphere can contain three equally spaced colinear points. The algorithm has two main steps. We first provide a very efficient protocol for Exact-T over $\mathbb{Z}^d$, whose cost grows only logarithmically with d.

Let $f : ([m]^d)^3 \to \{0, 1\}$ be defined via $f(x, y, z) = 1$ if and only if $x + y + z = T$, where $T \in \mathbb{Z}^d$ is some fixed vector. We provide an explicit NOF protocol for f whose cost is only $O(\log md)$. In words, players try to compute the vector $x + 2y + 3z$ "to the best of their knowledge" and then they compare notes.

1. Player 1 computes $v_x = T - y - z + 2y + 3z$.
2. Player 2 computes $v_y = x + 2(T - x - z) + 3z$.
3. Player 3 computes $v_z = x + 2y + 3(T - x - y)$.
4. Player 1 writes $\|v_x\|_2^2$ on the blackboard.
5. Player 2 writes 1 or 0 on the blackboard depending on whether $\|v_y\|_2^2 = \|v_x\|_2^2$.
6. Player 3 writes 1 or 0 on the blackboard depending on whether $\|v_z\|_2^2 = \|v_x\|_2^2$.
7. The protocol outputs 1 if the last two bits were both equal to 1, and 0 otherwise.

The cost of the above protocol is essentially determined by the largest possible value of $\|v_x\|_2^2$ in step 4 which is at most $O(m^2 d)$. Therefore, this cost does not exceed $O(\log md)$. We turn to prove correctness.

▶ **Lemma 20.** *The above protocol is correct.*

Proof. First note that the protocol outputs 1 if and only if $\|v_x\|_2^2 = \|v_y\|_2^2 = \|v_z\|_2^2$. Also, $v_x + v_z = 2v_y$, so that this condition holds only if all three vectors are equal, in which case $T - x - y - z = 0$. ◀

▶ Remark (More general protocols). Several variations on the above theme suggest them selves. Fix integers $a, b, c \in \mathbb{Z}$ and a $d \times d$ positive definite matrix D with integer entries. The players compute $a(T - y - z) + by + cz$, $ax + b(T - x - z) + cz$ and $ax + by + c(T - x - y)$, and rather than comparing the values of $\|v\|_2$, they consider the values of vDv^t. We wonder if these, or similar variations can together improve the complexity of the protocol.

4.2 Algorithmic Protocols for Exactly-N and $f_{k,T}^G$ over $\mathbb{Z}_m^n$

We seek algorithmically explicit protocols for Exactly-N [2], namely for the function $f : [N]^3 \to \{0, 1\}$ such that $f(x, y, z) = 1$ if and only if $x + y + z = N$, i.e., the exact-T problem over $\mathbb{Z}$

[2] Since n is used as an exponent in this section we use N for the size of input in Exactly-N .

or equivalently over $\mathbb{Z}_N$. We can give an efficient protocol to this problem by reduction to the protocol in the previous section, even though when applied directly to $\mathbb{Z}_N$ they give no improvement over the trivial protocol.

First fix a base m and let $n = 1 + \lceil \log_m N \rceil$. Consider the base-$m$ representation on the elements of $[N]$. Given a representation $x \in \{0, 1, \ldots, m-1\}^n$ of a number base m, for convenience we consider x_1 as the least significant digit. Note that all representations are of length n, if a number is small its representation is padded with zeros. The following protocol solves the Exactly-N problem in these settings. Let T be the base-m representation of N.

1. Player 1 computes the vector $C \in \{0, 1, 2\}^m$ defined as follows: the i-th entry of C is equal to $k \in \{0, 1, 2\}$ satisfying

$$T_i + (k-1)m < y_i + z_i + C_{i-1} \leq T_i + km,$$

   where addition is over $\mathbb{Z}$, and we define $C_0 = 0$.
2. Denote by C_x the carry vector computed by Player 1 in step 1. Player 2 and 3 compute corresponding vectors C_y and C_z, in a similar way.
3. Player 1 writes $C = C_x$ on the board.
4. Player 2 and 3, in turn, write 1 on the board if and only if their vector C_y (C_z) is equal to C_x.
5. If the last two bits written on the board are equal to 1, continue. Otherwise output 0 and terminate.
6. All players compute (in private) the vector $T'_i = T_i + mC_i - C_{i-1}$, for $i = 1, \ldots, n$.
7. The players run a protocol for the exact-T problem over $\mathbb{Z}^n$ with x, y, z and T'.

The cost of the above protocol is $O(n+2)$ for steps 1-5, plus the cost of the protocol used in step 7. The cost is thus $O(n + \log mn)$ if the players use the protocol from Section 4.1 in the last step. We prove next that this protocol is correct.

▶ **Lemma 21.** *The above protocol is correct.*

Proof. First assume $N = x + y + z$ over $\mathbb{Z}$. It is easy to verify the correctness of the protocol in this case, except maybe step 5. The correctness of step 5 follows from the following simple observation: assume $x_i + y_i + z_i + C_{i-1} = T_i + km$ (over $\mathbb{Z}$) for $k \in \{0, 1, 2\}$, then it must be that the sum of any pair of x_i, y_i, z_i and C_{i-1} is larger than $T_i + (k-1)m$ and at most $T_i + km$. Now consider the case $T \neq x + y + z$. If the protocol rejects on step 5 then obviously this is correct. If it does not reject then all players compute the same vector T', and $x + y + z = N$ over $\mathbb{Z}$ if and only if $x + y + z = T'$ over $\mathbb{Z}^n$. The correctness now follows from the correctness of the protocol over $\mathbb{Z}^n$. ◀

The above protocol for Exactly-N is correct for any choice of base m. To get an efficient protocol we optimize the choice of m. The running time of the protocol is $O(n + \log mn) = O(n + \log m)$. Since $m^n = N$, we get that $\log N = n \log m$, and thus the optimal choice is roughly $m = 2^{\sqrt{\log N}}$ which gives a running time of $O(\sqrt{\log N})$.

▶ Remark (The group $\mathbb{Z}_m^n$). The above protocol can also be adapted for $\mathbb{Z}_m^n$ (with addition modulo m). The idea is very similar, the only difference is that in the first steps Player 1 computes the vector $I_x \in \{0, 1, 2\}^n$ defined as follows: the i-th entry of I_x is equal to $k \in \{0, 1, 2\}$ satisfying

$$T_i + (k-1)m < y_i + z_i \leq T_i + km,$$

where addition is over $\mathbb{Z}$. The other two players compute analogous vectors.

4.3 A Protocol for $f^G_{k,T}$ over $\mathbb{Z}_2^n$

In this section we focus on the exact-T problem for the abelian group $\mathbb{Z}_2^n$. In other words, we study the permutation $f^{\mathbb{Z}_2^n}_{k,T}$. First we prove a lower bound on $\alpha_3(f^{\mathbb{Z}_2^n}_{3,T})$, and then show that this lower bound implies the existence of an efficient protocol for $f^{\mathbb{Z}_2^n}_{3,T}$. The complexity of $f^{\mathbb{Z}_2^n}_{k,T}$ is independent of T, so we can and will omit the subscript T in this section. Throughout this subsection we let $A^G_k = A(f^G_k)$.

First we prove that $A^{\mathbb{Z}_2^n}_t$-star freeness is preserved under tensor product.
Let $S \subset (\mathbb{Z}_2^n)^{k-1}$, denote by $S \otimes S$ the subset of $(\mathbb{Z}_2^{2n})^{k-1}$ comprised of all vectors $(x_1, y_1, \ldots, x_{k-1}, y_{k-1})$ such that $x_i, y_i \in S$ for $i = 1, \ldots, k-1$.

▶ **Lemma 22.** *If S is $A^{\mathbb{Z}_2^n}_k$-star free then $S \otimes S$ is $A^{\mathbb{Z}_2^{2n}}_k$-star free.*

Proof. Let $A = A^{\mathbb{Z}_2^{2n}}_k$ and let

$$(z_1, \ldots, z_{k-1}), (z_1 + d, \ldots, z_{k-1}), \ldots, (z_1, \ldots, z_{k-1} + d)$$

be an A-star in $S \times S$, where for each $1 \le i \le k-1$, $z_i = (x_i, y_i)$ with $x_i, y_i \in S$. Denote also $d = (d^1, d^2)$ where $d^1, d^2 \in \mathbb{Z}_2^n$. Then either

$$(x_1, \ldots, x_{k-1}), (x_1 + d^1, \ldots, x_{k-1}), \ldots, (x_1, \ldots, x_{k-1} + d^1)$$

is an $A^{\mathbb{Z}_2^n}_k$-star in S, or

$$(y_1, \ldots, y_{k-1}), (y_1 + d^2, \ldots, y_{k-1}), \ldots, (y_1, \ldots, y_{k-1} + d^2)$$

is an $A^{\mathbb{Z}_2^n}_k$-star in S, since either $d^1 \neq 0$ or $d^2 \neq 0$. ◀

It follows that if, for some fixed m, we can find a large $A^{\mathbb{Z}_2^m}_k$-star free subset S, then tensor powers of S are large $A^{\mathbb{Z}_2^n}_k$-star free sets. We show:

▶ **Lemma 23.** $\alpha_3(A^{\mathbb{Z}_2^2}_3) = \alpha_3(n, n) = 8$.

Together with Lemma 22 this yields:

▶ **Corollary 24.** *For every integer $n \ge 2$, there holds $\alpha_3(A^{\mathbb{Z}_2^n}_3) \ge 2^{3n/2}$.*

Proof. Let S be a star-free subset in $A^{\mathbb{Z}_2^2}_3$ of cardinality $8 = 4^{3/2}$ as in Lemma 23. The claim follows by taking the tensor powers of S as in Lemma 22. ◀

Proof of Lemma 23. We denote the elements of $\mathbb{Z}_2^2$ as follows $(0,0) = 0$, $(0,1) = 1$, $(1,0) = 2$ and $(1,1) = 3$. The matrix associated with $A^{\mathbb{Z}_2^2}_3$ is:

0	**1**	2	3
1	0	**3**	**2**
2	**3**	0	1
3	2	**1**	**0**

The 8 entries in bold form a star-free set, so that $\alpha_3(A^{\mathbb{Z}_2^2}_3) \ge 8$, and consequently $\alpha_3(4,4) \ge 8$. One can verify that in fact $\alpha_3(4,4) = \alpha_3(A^{\mathbb{Z}_2^2}_3) = 8$. To see this first notice that if there is a star-free subset of cardinality 9 then one of the values must appear three times which already determines 10 out of the 16 entries. One can now rule out the existence of a size 9 star-free subset by exhaustive search. ◀

It is interesting to determine $\alpha_k(n,n)$ for some small values of n. For example:

- Determine $\alpha_3(8,8)$, in particular compute $\alpha_3(A_3^{\mathbb{Z}_2^3})$.
- Determine $\alpha_k(4,4)$, in particular compute $\alpha_k(A_k^{\mathbb{Z}_2^2})$, for $k>3$.

It is interesting to note that, while as shown, $\alpha_3(4,4)=8$, there holds $\overline{\alpha}_3(4,4)=9$. The fact that $\overline{\alpha}_3(4,4)\leq 9$ is easy to verify, and the following example shows the equality:

$$\begin{array}{cccc} \mathbf{1} & * & * & \mathbf{3} \\ * & \mathbf{1} & * & \mathbf{4} \\ * & * & \mathbf{1} & \mathbf{2} \\ \mathbf{2} & \mathbf{3} & \mathbf{4} & * \end{array}$$

Thus, continuing the discussion at the end of Section 3.1, $\overline{\alpha}_3(n,N)$ and $\alpha_3(n,N)$ need not be equal when $N<2n-1$.

The following theorem shows that for groups, α_k (the size of the largeset 1-monochromatic cylinder intersection) completely characterizes χ_k (the minimum number of cylinder intersections that partition the 1's). The proof is a simple generalization of Theorem 4.3 in [13].

▶ **Theorem 25.** *If G is a group of order n, then*

$$\chi_k(f_k^G)\leq O\left(\frac{kn^{k-1}\log n}{\alpha_k(f_k^G)}\right).$$

Proof. The proof is in two steps:

Step I: A-star freeness is preserved under translation, where $A=A_k^G$. Indeed, let $S\subset G^{k-1}$ and let $\mathbf{a}=(a_1,\ldots,a_{k-1})\in G^{k-1}$. If

$$(x_1,\ldots,x_{k-1}),(x_1+d,\ldots,x_{k-1}),\ldots,(x_1,\ldots,x_{k-1}+d)$$

is an A-star in $S+\mathbf{a}$, then

$$(x_1,\ldots,x_{k-1})-\mathbf{a},(x_1+d,\ldots,x_{k-1})-\mathbf{a},\ldots,(x_1,\ldots,x_{k-1}+d)-\mathbf{a}$$

is an A-star in S.

Step II: Every $S\subset G^{k-1}$ has $O(\frac{kn^{k-1}\log n}{|S|})$ translates whose union covers all of G^{k-1}. This follows from the integrality gap for covering [28], but for completeness here is a proof. Pick at random t translates $\mathbf{a}_1,\ldots,\mathbf{a}_t\in[n]^{k-1}$ of S. The probability that a given element $\mathbf{x}\in[n]^{k-1}$ is covered by a random translate of S is exactly $\frac{|S|}{n^{k-1}}$. Therefore, and since the translates are picked independently uniformly at random, the expected number of uncovered elements of G^{k-1} is

$$n^{k-1}\cdot\left(1-\frac{|S|}{n^{k-1}}\right)^t.$$

Taking $t=O(\frac{kn^{k-1}\log n}{|S|})$ makes the expectation less than 1, which proves the lemma. ◀

▶ **Corollary 26.** $\chi_3(f_3^{\mathbb{Z}_2^m})\leq O\left(m\cdot 2^{m/2}\right)$.

The bound in Corollary 26 is similar to the bound of Ada, Chattopadhyay, Fawzi and Nguyen [1] for the case $k=3$, with slight improvement in the log factors. Ada et al. proved $\chi_3(f_3^{\mathbb{Z}_2^m})\leq O(m^{k+1}2^{m/2^{k-2}})$, by observing that this function is a composed function of the form $NOR\circ XOR$ and giving non trivial protocols for such cases.

Note that the proof of Theorem 25 yields a cover of $[n]^{k-1}$ by A-star free sets, but this is easily turned into a partition, since a subset of an A-star free set is also A-star free. Therefore, any lower bound on $\alpha_k(f_k^G)$ can be translated into an upper bound on $\chi_k(f_k^G)$ which in turn implies an efficient (non-explicit) protocol for f_k^G (By Theorem 5). Another interesting consequence of Theorem 25 is that any lower bound on $\chi_k(f_k^G)$ significantly larger than $\log n$ improves the known bounds for the size of a corner-free subset of G. This clearly boosts our interest in the multiparty communication complexity of f_k^G.

We wonder whether there are analogs of Theorem 25 for every permutation.

▶ **Question 27.** *How large can* $\chi_k(A) \cdot \alpha_k(A)/n^{k-1}$ *be for an arbitrary permutation* A*?*

5 Lower Bounds

5.1 Nonconstructive Lower Bounds

We first prove a nearly tight but nonconstructive lower bound on the communication complexity of random high-dimensional permutations.

▶ **Theorem 28.** *For every integer* $k \geq 3$*, and for most* $(k-1)$*-dimensional permutations* $f : [n]^k \to \{0,1\}$*,*

$$\log \chi_k(f) \geq \Omega(\frac{\log n}{k}).$$

Proof. The lower bound on the number of high-dimensional permutations was recently improved by Keevash [25] who showed that there are at least $2^{\Omega(n^d \log n)}$ d-dimensional permutations. If we view a permutation as a map $[n]^k \to \{0,1\}$, this means at least $2^{\Omega(n^{k-1} \log n)}$ permutations. In the spirit of the proof of Lemma 3.5 in [7], we now estimate the number of such permutation for which $\chi_k(f)$ is bounded. Note that we cannot simply use the estimate from [7] since it only works for functions $f : [n]^{k-1} \times [N] \to \{0,1\}$ with N that is much smaller than n, roughly $N \leq \sqrt{\frac{n}{k}}$.

Let $f : [n]^k \to \{0,1\}$ be a $(k-1)$-dimensional permutation, and let $\{C_1, \ldots, C_\chi\}$ be a partition of $f^{-1}(1)$ into $\chi = \chi_k(f)$ cylinder intersections. For $i \in [k]$ define a function $A_i : [n]^{k-1} \to [\chi]$ as follows: For $a = (a_1, \ldots, a_{k-1}) \in [n]^{k-1}$, let $L = (a, i)$ be a line in $[n]^{k-1}$. There is a unique 1 entry in L and this entry is in exactly one of the cylinder intersections $\{C_1, \ldots, C_\chi\}$, say C_j. In this case we define $A_i(a_1, \ldots, a_{k-1}) = j$.

As seen in the proof of Theorem 5, it is possible to recover f from knowledge of the functions $A_1, \ldots, A_k$. Namely, $f(x_1, \ldots, x_k) = 1$ if and only if all the values $A_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1})$ for $i = 1, \ldots, k$ are equal. But for every $i \in [k]$ there are $\chi^{n^{k-1}}$ possible functions $A_i : [n]^{k-1} \to [\chi]$. Thus, the number of $(k-1)$-dimensional permutations $f : [n]^k \to \{0,1\}$ with $\chi_k(f) \leq \chi$ is at most $(\chi^{n^{k-1}})^k = 2^{kn^{k-1} \cdot \log \chi}$. Combining this with Keevash's lower bound achieves our result. ◀

A simple corollary of Theorem 28, and Theorem 5 is:

▶ **Corollary 29.** *For every integer* $k \geq 2$*, almost all* $(k-1)$*-dimensional permutations* $f : [n]^k \to \{0,1\}$ *satisfy* $D_k(f) \geq \Omega(\frac{\log n}{k})$.

Theorem 28 proves the lower bound $\chi_k(f) \geq 2^{\Omega(\frac{\log n}{k})}$ for a random permutation $f : [n]^k \to \{0,1\}$. It is interesting to find out how this extends for a random linjection $f : [n]^{k-1} \times [N] \to \{0,1\}$ with $n < N$. It is also interesting to see whether the dependency on k can be removed.

Finally we turn to the case $k = 3$. The number of 2-dimensional permutations (aka Latin squares) is known to be $((1 + o(1))\frac{n}{e^2})^{n^2}$ (see [39]). It follows that for most 2-dimensional permutations f there holds $\log \chi_3(f) \geq \frac{1}{3} \log n - \Theta(1)$.

5.2 Lower Bounds for $\chi_k(n, N)$

We prove an upper bound on $\alpha_k(n, N)$, using its graph theoretic interpretation from Section 3.1, which implies the corresponding lower bound on χ_k. We start with $k = 3$:

▶ **Theorem 30.** *Let $A : [n] \times [n] \to [N]$ be a linjection, where $N \leq n \cdot 2^{c \log^*(n)}$. Then there exists $c > 0$ such that $\alpha_3(A) \leq O\left(\frac{n^2}{2^{c \log^*(n)}}\right)$.*

The proof of Theorem 30 is an adaptation of Solymosi's [36] simplification of Ajtai and Szemerédi's [2] Corners Theorem. We will use the improved version of the triangle removal lemma [33] due to Fox [18]:

▶ **Lemma 31** (Triangle removal lemma). *For every $\epsilon > 0$ there is a $\delta > 0$ such that every n-vertex graph with at most δn^3 triangles can be made triangle-free by removing ϵn^2 edges. Specifically δ^{-1} can be taken as a tower of twos of height $405 \log \epsilon^{-1}$.*

Proof of Theorem 30. Let $G = G(A)$, $V = V(G)$. Notice that $|V| = 2n + N$. Let $S \subset [n]^2$ be an A-star free subset of size $\alpha_3(A)$. As in the proof of Theorem 14 we let $T = \{< x, y, A(x, y) > | (x, y) \in S\}$ be the family of triangles in G that corresponds to S. Let F be that subgraph of G whose edge set is the union of all triangles in T. This graph contains the $|S|$ edge-disjoint triangles in T, and no additional triangles.

Thus, if we denote $\delta = |S|/|V|^3$ and $\epsilon = |S|/|V|^2$, then F contains exactly $\delta|V|^3$ triangles and it cannot be made triangle free by removing fewer than $\epsilon|V|^2$ edges. Lemma 31 yields $\log^*(\delta^{-1}) \leq 405 \log(\epsilon^{-1})$, and since $\delta < \frac{n^2}{(2n+N)^3} < \frac{1}{N}$ we conclude that

$$\epsilon \leq 2^{\frac{-1}{405} \log^*(N)}.$$

But $|S| = \epsilon|V|^2 \leq 9\epsilon N^2$, so that for $N \leq 2^{c \log^*(n)} n$, with $c = (3 \cdot 405)^{-1}$, $|S| \leq O\left(\frac{n^2}{2^{c \log^*(n)}}\right)$. ◀

We now state the case of general k, proved in the full version.

▶ **Theorem 32.** *For every natural numbers $k \geq 3$,n and N it holds that*

$$\alpha_k(n, N) \leq O\left(\frac{kn^{k-2}N}{\log^*(n)}\right).$$

5.3 A Lower Bound on $\chi_3(n, N)$

In this section we state our better lower bound for the case $k = 3$. The proofs appears in the full version of our paper.

▶ **Lemma 33.** *Let $L = \chi_3(n, N)$ for some integers $N \geq n$, then $\log n < (2^{L+1} - 1) \cdot \log(4NL/n)$. In particular for $k = 3$, we have $\chi_3(n, n) \geq \log \log n - O(\log \log \log n)$.*

Another simple corollary of Lemma 33 is due to Meshulam and is reproduced in [4].

▶ **Corollary 34.** *If $\chi_3(n, N) \leq L$ for some integers $N \geq n$, then $N \geq \frac{1}{4L} \cdot n^{1+1/(2^L-1)}$.*

A note on the case $k > 3$. As we have just seen $\chi_3(A) \geq \Omega(\log\log n)$ for every 2-dimensional permutation A. It is conceivable that a similar bound holds for higher dimensions as well. This was previously conjectured in [10] for the Exact-T problem. If we try to adapt the proof of Lemma 33 to higher k, exactly one difficulty arises which we formulate as a question.

▶ **Question 35.** *Let $S \subseteq [n]^k$ be a set of cardinality m that meets every line at most once. Determine, or estimate $\phi_k(n,m)$, the least possible cardinality $|\bar{S}|$ of its closure. We use the shorthand $\phi_k(m)$ when appropriate.*

For $k = 2$ the answer is easy: $\phi_2(m) = m^2$, since $|\bar{S}| = |S|^2$. But for $k > 2$ the problem becomes very hard and no lower bound is known. In fact, for $k \geq 3$, and for large enough m there holds $\phi_k(m) = m$. In other words, unlike the case $k = 2$ it may happen that $\bar{S} = S$ for large S. For example, as shown in [13], $\phi_3(m) = m$ when $m = n^2/2^{\Omega(\sqrt{\log n})}$, whereas it is shown in [34] that $\phi_3(m) > m$ when $m \geq n^2/(\log\log n)^{\frac{1}{22}}$. For $k > 3$ the situation is even worse, and all we have are the very weak lower bounds from Section 5.2. Namely, it follows from Theorem 32 that $\phi_k(m)$ must be larger than m when $m \geq \Omega\left(\frac{kn^{k-1}}{\log^*(n)}\right)$. Proving any non-trivial bounds on $\phi_k(m)$ is a very interesting challenge. We raise the following conjecture in an attempt to improve the lower bounds on $\chi_3(n,n)$:

▶ **Conjecture 36.** *There are constants $c_1, c_2 > 0$ such that if $S \subseteq [n]^3$ meets every line at most once, and if $|S| \geq n^2/(\log\log n)^{c_1}$, then $|\bar{S}| \geq n^3/(\log\log n)^{c_2}$.*

6 Conclusion and Open Problems

This paper raises numerous open problems. Below we collect some of the major ones and explain some implications that would follow from progress on these questions.

▶ **Question 37.** *Improve the lower bound $\chi_3(n,n) \geq \Omega(\log\log n)$.*

Any lower bound $\chi_3(n,n) \geq \omega(\log\log n)$ yields an improvement to the best known bound on the number of colors required to color the $n \times n$ grid with no monochromatic equilateral right triangles. This subject goes back to Ajtai and Szemerédi's corners theorem [2] and its implications in additive combinatorics due to Solymosi [36]. A lower bound $\chi_3(n,n) \geq \omega(\log n)$ would improve the best known gap between randomized and deterministic communication complexity in the 3-players NOF model. A lower bound $\chi_3(n,n) \geq \Omega(\log n \cdot \log\log n)$ will improve the best known upper bound on the size of corner-free subsets of G^2 for any abelian group G. And finally, a lower bound $\chi_3(n,n) \geq \Omega(\log^2 n)$ will improve the best bounds on the size of a subset of $\mathbb{Z}_n$ with no three-term arithmetic progression. This is a classic problem that goes back at least to the 1950's [32].

▶ **Question 38.** *Improve the upper bound $\chi_3(n,n) \leq 2^{O(\sqrt{\log n})}$.*

The construction of denser Ruzsa-Szemerédi graphs than currently known. Namely, n-vertex graphs which are the disjoint union of n induced matchings, all of the same size r. This, in turn, reflects on the many applications of these. This would also improve our understanding regarding the limits of the triangle removal lemma; note that the current gaps between the bound in this lemma are huge.

▶ **Question 39.** *Improve the bounds on $\chi_k(n,n)$ for $k > 3$.*

▶ **Question 40.** *Improve the bounds on $\alpha_k(n,n)$ for $k > 3$.*

That would improve our state of knownledge regarding the bounds for the hypergraph removal lemma. It is also interesting to determine $\alpha_k(n,n)$ for some small values of n. For example: Determine $\alpha_3(8,8)$, and in particular compute $\alpha_3(A_3^{\mathbb{Z}_2^3})$, or Determine $\alpha_k(4,4)$, and in particular compute $\alpha_k(A_k^{\mathbb{Z}_2^2})$ for $k > 3$.

▶ **Question 41.** *What is the relationship between $\overline{\alpha}_k(n,N)$ and $\alpha_k(n,N)$ in the whole range $n \leq N \leq (k-1)(n-1)$?*

References

1. A. Ada, A. Chattopadhyay, O. Fawzi, and P. Nguyen. The NOF multiparty communication complexity of composed functions. *computational complexity*, 24(3):645–694, 2015.
2. M. Ajtai and E. Szemerédi. Sets of lattice points that form no squares. *Stud. Sci. Math. Hungar*, 9(1975):9–11, 1974.
3. N. Alon. Testing subgraphs in large graphs. *Random Structures & Algorithms*, 21(3-4):359–370, 2002.
4. N. Alon, A. Moitra, and B. Sudakov. Nearly complete graphs decomposable into large induced matchings and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1079–1090. ACM, 2012.
5. L. Babai, T. Hayes, and P. Kimmel. The cost of the missing bit: communication complexity with help. *Combinatorica*, 21:455–488, 2001.
6. L. Babai, N. Nisan, and M. Szegedy. Multiparty Protocols, Pseudorandom Generators for Logspace, and Time-Space Trade-offs. *Journal of Computer and System Sciences*, 45:204–232, 1992.
7. P. Beame, M. David, T. Pitassi, and P. Woelfel. Separating deterministic from randomized NOF multiparty communication complexity. In *Proceedings of the 34th International Colloquium On Automata, Languages and Programming*, Lecture Notes in Computer Science. Springer-Verlag, 2007.
8. P. Beame, T. Pitassi, and N. Segerlind. Lower bounds for Lovász-Schrijver systems and beyond follow from multiparty communication complexity. *SIAM Journal on Computing*, 37(3):845–869, 2006.
9. F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences*, 32(12):331–332, 1946.
10. R. Beigel, W. Gasarch, and J. Glenn. The multiparty communication complexity of Exact-T: Improved bounds and new problems. In *International Symposium on Mathematical Foundations of Computer Science*, pages 146–156. Springer, 2006.
11. K. Bibak. Additive Combinatorics: With a View Towards Computer Science and Cryptography - An Exposition. In *Number Theory and Related Fields, In Memory of Alf van der Poorten*, pages 99–128. Springer, 2013.
12. 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.
13. A. Chandra, M. Furst, and R. Lipton. Multi-party protocols. In *Proceedings of the 15th ACM Symposium on the Theory of Computing*, pages 94–99. ACM, 1983.
14. E. Croot, V. Lev, and P. Pach. Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small. *arXiv preprint*, 2016. `arXiv:1605.01506`.
15. Z. Dvir. On the size of Kakeya sets in Finite Fields. *J. Amer. Math Soc.*, pages 1093–1097, 2009.
16. M. Elkin. An improved construction of progression-free sets. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 886–905. Society for Industrial and Applied Mathematics, 2010.

17 J. S. Ellenberg and D. G. On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *Annals of Mathematics*, 185(1):339–343, 2017.

18 J. Fox. A new proof of the graph removal lemma. *Annals of Mathematics*, pages 561–579, 2011.

19 H. Furstenberg and Y. Katznelson. A density version of the Hales-Jewett theorem. *Journal d'Analyse Mathematique*, 57(1):64–119, 1991.

20 R. Graham and J. Solymosi. Monochromatic equilateral right triangles on the integer grid. In *Topics in discrete mathematics*, pages 129–132. Springer, 2006.

21 R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*, volume 20. John Wiley & Sons, 1990.

22 V. Grolmusz. The BNS lower bound for multi-party protocols is nearly optimal. *Information and computation*, 112(1):51–54, 1994.

23 J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Computational Complexity*, 1:113–129, 1991.

24 J. Håstad and A. Wigderson. Simple analysis of graph tests for linearity and PCP. *Random Structures & Algorithms*, 22(2):139–160, 2003.

25 P. Keevash. The existence of designs II. *arXiv preprint*, 2018. `arXiv:1802.05900`.

26 E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

27 N. Linial and Z. Luria. An upper bound on the number of high-dimensional permutations. *Combinatorica*, 34(4):471–486, 2014.

28 L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.

29 S. Lovett. Additive Combinatorics and its Applications in Theoretical Computer Science. *Theory of Computing, Graduate Surveys*, 8:1–55, 2017.

30 N. Nisan and A. Widgerson. Rounds in communication complexity revisited. In *Proceedings of ACM STOC*, pages 419–429. ACM, 1991.

31 D.H.J. Polymath. A new proof of the density Hales-Jewett theorem. *arXiv preprint*, 2009. `arXiv:0910.3926`.

32 K. F. Roth. On certain sets of integers. *Journal of the London Mathematical Society*, 1(1):104–109, 1953.

33 I. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Keszthely, 1976), Coll. Math. Soc. J. Bolyai*, 18:939–945, 1978.

34 I. D. Shkredov. On a two-dimensional analogue of Szemerédi's theorem in Abelian groups. *Izvestiya: Mathematics*, 73(5):1033–1075, 2009.

35 A. Shraibman. A Note on Multiparty Communication Complexity and the Hales-Jewett Theorem. *arXiv preprint*, 2017. `arXiv:1706.02277`.

36 J. Solymosi. Note on a Generalization of Roth's Theorem. In *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, pages 825–827. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.

37 P. Tesson. An application of the Hales-Jewett Theorem to multiparty communication complexity, 2004.

38 L. Trevisan. Guest column: additive combinatorics and theoretical computer science. *SIGACT News*, 40(2):50–66, 2009.

39 J. H. van Lint and R. M. Wilson. *A course in combinatorics*. Cambridge university press, 2001.

40 A. Yao. On ACC and threshold circuits. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science*, pages 619–627. IEEE, 1990.