Larger Corner-Free Sets from Better NOF Exactly-N Protocols

Nati Linial [*] Adi Shraibman [†]

Abstract

A subset of the integer planar grid $[N] \times [N]$ is called *corner-free* if it contains no triple of the form $(x, y), (x + \delta, y), (x, y + \delta)$. It is known that such a set has a vanishingly small density, but how large this density can be remains unknown. The best previous construction was based on Behrend's large subset of $[N]$ with no 3-term arithmetic progression. Here we provide the first substantial improvement to this lower bound in decades. Our approach to the problem is based on the theory of communication complexity.

In the 3-players exactly-N problem the players need to decide whether $x + y + z = N$ for inputs x, y, z and fixed N. This is the first problem considered in the multiplayer Number On the Forehead (NOF) model. Despite the basic nature of this problem, no progress has been made on it throughout the years. Only recently have explicit protocols been found for the first time, yet no improvement in complexity has been achieved to date. The present paper offers the first improved protocol for the exactly-N problem. This is also the first significant example where algorithmic ideas in communication complexity bear fruit in additive combinatorics.

1 Introduction

Van der Waerden's well known theorem [19] states that for every r, k and every large enough N, if the elements of $[N] := \{1, \ldots, N\}$ are colored by r colors, then there must exist a length-k monochromatic arithmetic progression. Erdős and Turán introduced the density version of this theorem. Let $\rho_k(N)$ be the largest density of a subset of $[N]$ without an arithmetic progression of length k. Szemerédi's famous theorem [18] shows[1] that $\rho_k(N) = o(1)$ for every $k \geq 3$.

Extending van der Waerden's theorem, Gallai proved that in every finite coloring of $\mathbb{Z}^2$ some color contains arbitrarily large square subarrays. In search of a density version of Gallai's theorem,

[*] Hebrew University of Jerusalem, Jerusalem, Israel. nati@cs.huji.ac.il.

Supported in part by Grant 659/18 "High dimensional combinatorics" of the Israel Science Foundation.

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

[1] Unless otherwise specified, all asymptotic statements are taken with $N \to \infty$.

Erdős and Graham asked about the largest density of a subset of the integer grid $[N] \times [N]$ without a *corner*, i.e., a triple $(x, y), (x + \delta, y), (x, y + \delta)$ for some $\delta \neq 0$. Denote this quantity by $\rho_3^\angle(N)$.

Ajtai and Szemerédi [2] proved the first *corners theorem*, showing that $\rho_3^\angle(N) = o(1)$. This theorem easily yields that $\rho_3(N) = o(1)$, namely, the $k = 3$ case of Szemerédi's theorem, due to Roth [15]. Later on Solymosi [17] showed how to derive Ajtai and Szemerédi's corners Theorem from the Triangle Removal Lemma [16].

The quantitative aspects of all these results: Szemerédi's theorem, the corner theorem, the $(6, 3)$ Theorem (e.g., [16]) and the triangle removal lemma remain unfortunately poorly understood. The upper bounds have gradually improved over the years and the current "world record" of Bloom and Sisask [8] is

$$\rho_3(N) \leq (\log N)^{-1-c} \quad \text{for some absolute constant} \quad c > 0.$$

In contrast, not much has happened with lower bounds in these problems. Behrend [6] has constructed a large subset of $[N]$ without a 3-term arithmetic progression, whence[2]

$$\rho_3(N) \geq 2^{-2\sqrt{2}\sqrt{\log N} + o(\sqrt{\log N})}.$$

Elkin's modification [10] of Behrend's construction, has improved only the little-o term. Behrend's construction also yields the previously best known lower bound on $\rho_3^\angle(N)$, viz.

$$\rho_3^\angle(N) \geq 2^{-2\sqrt{2}\sqrt{\log N} + o(\sqrt{\log N})} = 2^{-2.828\ldots\sqrt{\log N} + o(\sqrt{\log N})}. \tag{1}$$

Here we improve these bounds, using a new approach to these problems that is based on communication complexity. We prove

Theorem 1.1

$$\rho_3^\angle(N) \geq 2^{-2\sqrt{\log e}\sqrt{\log N} + o(\sqrt{\log N})} = 2^{-2.4022\ldots\sqrt{\log N} + o(\sqrt{\log N})}.$$

There is an explicit corner-free subset of $[N] \times [N]$ *of size*

$$N^2 / 2^{2\sqrt{\log e}\sqrt{\log N} + o(\sqrt{\log N})}.$$

1.1 The Computational Perspective

The multiplayer Number On the Forehead (NOF) model of communication complexity was introduced by Chandra, Furst and Lipton [9]. Given a function $f : [N]^k \to \{0, 1\}$, the k players in this scenario should jointly find out $f(x_1, \ldots, x_k)$. We think of x_i as being placed on player i's forehead, so that each player sees the whole input bar one argument. Players communicate by writing bits

[2] All logarithms in this paper are in base 2.

on a shared blackboard according to an agreed-upon protocol. This model is intimately connected to several key problems in complexity theory. E.g., lower bounds on the size of ACC^0 circuits for a natural function in P [20, 11], branching programs, time-space tradeoffs for Turing machines [12], and proof complexity [5]. In addition, progress in the NOF model, even for specific problems and for $k = 3$, would have profound implications in graph theory and combinatorics [13, 3].

Much of Chundra, Furst and Lipton's seminal paper [9] is dedicated to the exactly-N function $f : [N]^k \to \{0, 1\}$, where $f(x_1, \ldots, x_k) = 1$ iff $\sum x_i = N$. They discovered a connection between the communication complexity of this function and well-known problems in additive combinatorics and Ramsey theory. They used Ramsey's theory to prove a (rather weak) lower bound on the NOF communication complexity of this function. Using the connection to additive number theory, they showed that a $O(\sqrt{\log N})$ protocol exists, although they have not made this protocol explicit.

Our main concern here is with the 3-player NOF Exactly-N problem (or, the essentially equivalent addition problem where the players need to decide whether $x + y = z$). The three players jointly design a communication protocol $\mathcal{CP}$. Then they get separated, and are given access to inputs x, y, z, as described above. Namely, player P_x gets to see inputs y and z, P_y sees x and z, and P_z sees x and y. According to their chosen $\mathcal{CP}$, they take turns writing messages on a blackboard that is visible to all three players. The game ends when every player can deduce whether $x + y + z = N$ (resp. $x + y = z$) based on the inputs she sees, and the contents of the blackboard (called *transcript*). The *complexity* of $\mathcal{CP}$ is the maximal length of a transcript over all instances x, y, z. As a function of N, the *communication complexity* of the problem is the minimal complexity of a protocol $\mathcal{CP}$ that solves the problem correctly for all inputs. A *one-round* protocol starts with P_z who writes a message on the board. Subsequently P_x and P_y each write a single verification bit. Let us spell out the connection between the communication complexity of the 3-players NOF addition problem and the corners theorems:

Claim 1.2 ([9], implicit) *1. There is an optimal one-round protocol for exactly-N.*

2. Let $T = \mathbb{T}(x, y)$ be the message that P_z posts on inputs (x, y) in a one-round protocol for exactly-N. Then the set

$$S(T) = \{(x, y) : \mathbb{T}(x, y) = T\}$$

is corner-free.

See [9, 7, 1, 13, 3] for more details about the above claim and the relation between communication complexity and additive combinatorics.

There are several reasons why it is highly significant to determine the communication complexity of the exactly-N function, aside of the very fundamental nature of the problem:

- Our poor understanding of this question is manifested by the huge gap between the upper and lower bounds that we currently have on the communication complexity of this problem. This gap is double exponential for three players, and even worse for $k > 3$ players.

- Despite the significance of the NOF model, we still know very little about it. The rich web of mathematical and computational concepts surrounding the exactly-N function suggests that it may open the gate to progress in understanding numerous other NOF functions.
- The k-player exactly-N function is a *graph function* [4]. For most functions in this class the deterministic and randomized communication complexity differ substantially, but no explicit function with deterministic complexity larger than polylogarithmic is presently known.
- This problem is *equivalent* to corner theorems in additive combinatorics (e.g., [2]), and is closely related to other important problems such as constructing Ruzsa-Szemerédi graphs and the triangle removal lemma [13, 3].

As mentioned, the existence of a protocol for the exactly-N problem has already been known since [9]. However, this was just an existential statement and no actual protocol was provided. This lacuna was recently remedied with two protocols [13, 3] of the exact same complexity as the one whose existence was proven in [9]

$$2\sqrt{2}\sqrt{\log N} + o(\sqrt{\log N}) = 2.828...\sqrt{\log N} + o(\sqrt{\log N}). \tag{2}$$

Here we give the first improved protocol for the exactly-N problem, and prove

Theorem 1.3 *There is an explicit protocol for* 3*-players exactly-*N *of complexity*

$$2\sqrt{\log e}\sqrt{\log N} + o(\sqrt{\log N}) = 2.4022...\sqrt{\log N} + o(\sqrt{\log N}). \tag{3}$$

2 Proof of Theorem 1.3

We recall that the three players called P_x, P_y and P_z get to see the inputs $(y, z), (x, z)$ and (x, y) respectively. Given integers $q, d > 1$, define $g = g_{q,d}(\alpha, \beta, \gamma)$ to be 1 if $\alpha + \beta = \gamma$ and 0 otherwise. Here $\alpha, \beta \in [q]^d$, $\gamma \in [2q]^d$ and addition is vector addition in $\mathbb{R}^d$. The following one-round protocol [3] for g is correct because the inequality $\|2\alpha - \gamma\|^2 + \|2\beta - \gamma\|^2 \geq 2\|\alpha - \beta\|^2$ holds always and is an equality iff $\gamma = \alpha + \beta$.

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

1. *P_z computes $\|\alpha - \beta\|_2^2$, and writes the result on the board.*
2. *P_y writes 1 iff $\|\alpha - \beta\|_2^2 = \|2\alpha - \gamma\|_2^2$.*
3. *P_x writes 1 iff $\|\alpha - \beta\|_2^2 = \|2\beta - \gamma\|_2^2$.*

The cost of this protocol is $2 + \log dq^2$.

The above is an efficient method to decide high-dimensional vector addition, but our objective is to decide the integer addition relation $X+Y+Z=N$. We let $x=X, y=Y$ and $z=N-Z$, so the relation we need to consider is $x+y=z$.

Our protocol to decide whether $x+y=z$ builds on the protocol for $g_{q,d}$. It is the issue of carry in integer addition that makes this decision problem harder. The integers $q,d>1$ are chosen so that

$$2qN > q^d \geq 2N. \tag{4}$$

the specific choice is made below so as to minimize the cost of the protocol.
We denote by w_q the vector that corresponds to the base q representation of the integer w.

As usual, e_i is the d-dimensional vector with 1 in the i-th coordinate and zeros elsewhere. Let $C(x,y) \in \{0,1\}^d$ be the carry vector when x and y are added in base q. The relation $x+y=z$ among integers is equivalent to the vector relation

$$x_q + y_q = \zeta,$$

where the i-th coordinate of ζ is

$$\zeta_i = z_i + q \cdot C(x,y)_i - C(x,y)_{i-1}$$

(Here $C(x,y)_0 = 0$). The protocol from [3] now suggests itself: P_z posts $C(x,y)$, and Protocol 1 is used to decide the relation $x_q + y_q = \zeta$. This yields again the estimate (2).

The alternative approach that we adopt here considers instead the equivalent vector relation

$$x_q + \eta = z_q$$

where

$$\eta = (x+y)_q - x_q.$$

Concretely, the i-th coordinate of η is:

$$\eta_i = y_i - q \cdot C(x,y)_i + C(x,y)_{i-1}.$$

In order to run Protocol 1, P_z needs to know η and x_q, and he does. With P_y it's even simpler, since he needs to know x_q and z_q which are his inputs. The only difficulty is with P_x who needs to know z_q (which he does) and η. The latter is not part of his input and P_z fills in the missing information for him.

The obvious solution is for P_z to reveal $C(x,y)$ to P_x using d bits of information. However, we can save communication by exploiting the fact that P_x and P_z share some information, i.e., they both know y for every $y \neq 0$.

By a standard argument in this area which we detail below (Proposition 2.1), a protocol that works for *typical* pairs x, y can be easily modified to work in *all* cases. So, let us pick x and y uniformly at random from among the d-digit numbers in base q and think of $C = C(x, y)$, the vector of carry bits as a random variable on this probability space. The number of bits that P_z needs to post so that P_x gets to know C, and therefore know η, is $H(C|y)$, the entropy of C given y. The gain is clear, since $H(C) > H(C|y)$.

It remains to estimate $H(C|y)$. Let X be the random variable that is a uniformly sampled subset of $[s]$ of cardinality $\geq t$, for some integers $s \geq t \geq 0$. We recall that $H(X) = (1 + o_s(1)) \cdot s \cdot h(t/s)$, where $h(\cdot)$ is the univariate entropy function, and the same holds also if we consider subsets of $[s]$ of cardinality $\leq t$. Let r be an integer in the range $d \gg r \gg 1$, e.g., $r \approx \sqrt{d}$. For $j = 1, \ldots, r$, let

$$S_j = \{i \mid \frac{qj}{r} > y_i \geq \frac{q(j-1)}{r}\},$$

where $q > y_i \geq 0$ is the i-th digit of y. A carry occurs in digit $i \in S_j$ only if $x_i > \frac{q(r-j)}{r}$, where x_i is the i-th digit of x. Then

$$H(C|y) \leq (1 + o_r(1)) \sum_{j=1}^{r} \frac{|S_j|}{d} h(\frac{j}{r}).$$

Since y is chosen at random, $|S_j| \leq (1 + o_r(1))\frac{d}{r}$, and so

$$H(C|y) \leq (1 + o_r(1)) \sum_{j=1}^{r} \frac{1}{r} h(\frac{j}{r}).$$

The limit of this expression as $r \to \infty$ is

$$\lambda = \int_0^1 h(u)du = \frac{\log e}{2} = 0.721...$$

It is left to optimize on q and d. The complexity of our protocol is

$$\lambda d + \log dq^2 + 2,$$

where recall that $2qN > q^d \geq 2N$. It is not hard to verify that choosing

$$d = \sqrt{\frac{2}{\lambda} \log 2N} \qquad q = 2^{\sqrt{\frac{\lambda}{2} \log 2N}}, \tag{5}$$

we get a protocol with complexity

$$2\sqrt{2\lambda \log N} + o(\sqrt{\log N}),$$

and this is asymptotically optimal in our setting.

To sum up, here is the protocol which proves Theorem 1.3:

Protocol 2 A protocol for exactly-N, for typical pairs x, y

For d, q as in Equation (5)

1. *P_z publishes the vector $\eta = (x+y)_q - x_q$ in a way that P_x can read.*

2. *The players run protocol 1 for $g_{q,d}$ on inputs x_q, η, z_q. That is:*

   (a) *P_z writes $\|\eta - x_q\|_2^2$ on the board*

   (b) *P_y writes 1 iff $\|\eta - x_q\|_2^2 = \|2x_q - z_q\|_2^2$.*

   (c) *P_x writes 1 iff $\|\eta - x_q\|_2^2 = \|2\eta - z_q\|_2^2$.*

Proposition 2.1 *Let $\mathcal{P}$ be an NOF protocol for the exactly-N that works correctly for an $\Omega(1)$-fraction of the input pairs x, y (and every z) with communication complexity $\Phi(N)$. Then there is an NOF protocol that works for all inputs with communication complexity $\Phi(N) + O(\log\log N)$.*

Proof Let $S \subseteq [N] \times [N]$ be the set of input pairs x, y on which $\mathcal{P}$ succeeds. We claim that there is a collection F of $O(\log N)$ vectors $\Delta \in [N] \times [N]$ such that

$$\cup_{\Delta \in F}(S + \Delta) \supseteq [N] \times [N].$$

In the modified protocol, P_z sees x, y and announces the index of some $\Delta = (\Delta_1, \Delta_2) \in F$ for which $(x - \Delta_1, y - \Delta_2) \in S$. Then the players run Protocol 2 with inputs $(x - \Delta_1, y - \Delta_2, z - \Delta_1 - \Delta_2)$.

The construction of F uses a standard fact about the set-cover problem. For a family of finite sets $\mathcal{X} \subseteq 2^\Omega$ we denote by $c(\mathcal{X})$ the least number of members in $\mathcal{X}$ whose union is Ω. Also $c^*(\mathcal{X})$ is the minimum cost of a fractional cover. Namely,

$$c^*(\mathcal{X}) = \min \sum_{\mathcal{X}} \omega_{\mathrm{X}}, \text{ where } \omega_{\mathrm{X}} \geq 0 \text{ for every } \mathrm{X} \in \mathcal{X} \text{ and } \sum_{\mathrm{x} \in \mathrm{X}} \omega_{\mathrm{X}} \geq 1 \text{ for every } \mathrm{x} \in \Omega.$$

Then

$$c(\mathcal{X}) \leq \log(|\Omega|) \cdot c^*(\mathcal{X})$$

(e.g., Lovász [14]) and actually the greedy algorithm yields a set cover that meets this bound.

In our case $\Omega = [N] \times [N]$, and

$$\mathcal{X} = \{(S + \Delta) \cap ([N] \times [N]) \mid \Delta \in [-N, N] \times [-N, N]\}.$$

It is easily verified that the weights $\omega_x = \frac{10}{N^2}$ constitute a fractional cover, so that $c^*(\mathcal{X}) \leq 40$ and hence $c(\mathcal{X}) \leq 80 \log N$, as claimed. ∎

3 Discussion

The strong relation between the exactly-N problem in the NOF model and questions in additive combinatorics has been discovered decades ago, in the seminal paper of Chundra, Furst and Lipton [9]. However, this subject remains under-developed. We believe that there is a lot to be done here, and many interesting avenues of research that this study can take. One obvious candidate is to seek further improvement is the addition problem. We conjecture:

Conjecture 3.1 *The NOF communication complexity of exactly-N is $o(\sqrt{\log N})$. Possibly it is much smaller, even as small as $(\log \log N)^{O(1)}$.*

In the realm of additive combinatorics these conjectures translate to:

Conjecture 3.2

$$\rho_3^{\angle}(N) \geq 2^{-o(\sqrt{\log N})}.$$

and possibly even

$$\rho_3^{\angle}(N) \geq 2^{-(\log \log N)^{O(1)}}.$$

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 and A. Shraibman. Number on the forehead protocols yielding dense ruzsa–szemerédi graphs and hypergraphs. *Acta Mathematica Hungarica*, 161(2):488–506, 2020.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] T. F. Bloom and O. Sisask. Breaking the logarithmic barrier in Roth's theorem on arithmetic progressions. *arXiv preprint arXiv:2007.03528*, 2020.

[9] 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.

[10] 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.

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

[12] E. Kushilevitz and N. Nisan. *Communication Complexity.* Cambridge University Press, 1997.

[13] N. Linial, T. Pitassi, and A. Shraibman. On the communication complexity of high-dimensional permutations. In *10th Innovations in Theoretical Computer Science Conference, ITCS San Diego, California, USA*, volume 124, pages 54:1–54:20, 2019.

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

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

[16] 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.

[17] J. Solymosi. Note on a generalization of Roth's theorem. *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, pages 825–827, 2003.

[18] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith*, 27(199-245):2, 1975.

[19] B. L. van der Waerden. Beweis einer Baudetschen Vermutung. *Nieuw Arch. Wiskunde*, 15:212–216, 1927.

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