Electronic Colloquium on Computational Complexity, Revision 2 of Report No. 187 (2026)

An Explicit and Optimal Separation of BPP from NP in Number-on-Forehead Communication Complexity

Haoyu Wang* Yimeng Wang† Pei Wu‡

Abstract

For every fixed $k \geq 3$, we construct an explicit total Boolean function in the k-player number-on-forehead model with public-coin randomized communication complexity $O_k(1)$ and nondeterministic communication complexity $\Omega_k(n)$, where n is the number of bits on each forehead. For three players, our construction improves the strongest previously known explicit separation, due to Kelley and Lyu (FOCS 2025) building on Kelley, Lovett, and Meka (STOC 2024), from constant randomized versus $\Omega(n^{1/2})$ nondeterministic communication to the optimal constant versus linear separation. Our construction uses algebraic geometry codes to preserve the nondeterministic hardness of generalized inner product while enabling a randomized protocol with constant communication.

*Penn State University. Email: hjw5492@psu.edu.
†University of California, Los Angeles. Email: kobewang1359@ucla.edu.
‡Penn State University. Email: pei.wu@psu.edu.

ISSN 1433-8092

1 Introduction

The number-on-forehead (NOF) model, introduced by Chandra, Furst, and Lipton [CFL83], is a central model of multiparty communication complexity. An input consists of k strings of n bits, and player i sees every string except the one on forehead i. The players communicate on a public blackboard to compute a Boolean function of the input, with the goal of minimizing the number of bits communicated. The model has many applications, for example to circuit lower bounds [HG91, BH12], pseudorandomness [BNS92, Vio07], proof complexity [BPS07], data structures [Păt10], and additive combinatorics [AS20, LS21, KLM24]. The extensive overlap among the players' views makes proving lower bounds for explicit functions challenging, even for three players.

A basic question in this model is how randomness and nondeterminism affect the amount of communication required. Randomized protocols use shared random bits and may err with small probability, while nondeterministic protocols verify certificates of acceptance without error. The intersection problem shows that nondeterminism can be exponentially more powerful than randomness. It asks whether k subsets of an n-element universe have a common element, and is the complement of set disjointness. For fixed k, a common element can be certified using $O(\log n)$ bits, whereas deciding whether one exists requires $\Omega_k(\sqrt{n})$ randomized communication [She14].

We study the reverse direction: how much more efficient can randomized protocols be than nondeterministic protocols? For two players, an elementary example gives the optimal separation: equality has a constant-cost public-coin randomized protocol but requires $\Omega(n)$ nondeterministic communication. For three players, however, equality of all three strings can be decided deterministically with two bits: two players each report whether the two strings they see are equal.

Whether a comparable separation holds for three or more players remained open until Beame, David, Pitassi, and Woelfel [BDPW10] proved its existence nonconstructively. For every fixed number of players, they obtained total NOF functions with constant public-coin randomized communication, one-sided error, and linear nondeterministic communication. They also gave explicit functions with constant randomized cost and an $\Omega(\log n)$ nondeterministic lower bound for every fixed $k \geq 3$.

Making the optimal separation explicit proved much harder. For three players, Kelley, Lovett, and Meka [KLM24] constructed a function with constant randomized complexity and nondeterministic complexity $\Omega(n^{1/3})$. Their proof builds on techniques from Kelley and Meka's breakthrough on three-term arithmetic progressions [KM23], extending them to a non-arithmetic setting. Kelley and Lyu [KL25] subsequently improved the lower bound for the same function to $\Omega(n^{1/2})$. Thus, even for three players, strong explicit separations required substantial advances, while an explicit construction achieving the optimal separation for every fixed number of players remained open.[1]

We resolve this problem. For every fixed $k \geq 3$, we construct an explicit total function with constant randomized complexity and linear nondeterministic complexity, matching the nonconstructive separation of Beame, David, Pitassi, and Woelfel [BDPW10]. For three players, this strengthens the previous $\Omega(n^{1/2})$ lower bound to the optimal $\Omega(n)$. Our function is different from the function studied in [KLM24, KL25], and our proof is algebraic rather than combinatorial.

[1] Related progress concerns Exactly-N [CFL83], which asks whether k input integers in $[N]$ sum to N. With $n = \Theta(\log N)$ bits per forehead, this problem has constant randomized complexity. Jaber, Liu, Lovett, Ostuni, and Sawhney [JLL+25] proved nondeterministic lower bounds of $n^{\Omega(1)}$ for three players and $\Omega(\log\log\log n)$ for four players. Their three-player bound inherits exponent $1/600$ from their corners theorem; the authors do not attempt to optimize it.

1.1 Main results

Write $\mathsf{R}_k(\mathsf{F})$, $\mathsf{N}_k(\mathsf{F})$, and $\mathsf{D}_k(\mathsf{F})$ for the public-coin randomized, nondeterministic, and deterministic NOF communication complexity of F, respectively. Randomized protocols must err with probability at most $1/3$ on every input. See Section 2 for the formal definitions. The number of players k is fixed, and implicit constants may depend on k.

Theorem 1.1 (Explicit constant-cost separation)**.** *For every fixed integer $k \geq 3$, there is an explicit family of total functions $\mathsf{F}_n : (\{0,1\}^n)^k \to \{0,1\}$ such that, for all sufficiently large n,*

$$\mathsf{R}_k(\mathsf{F}_n) = O_k(1), \qquad \mathsf{N}_k(\mathsf{F}_n) = \Omega_k(n).$$

The randomized protocol is one-sided: it accepts every 1-input with probability one and accepts every 0-input with probability at most $1/3$.

The linear lower bound is optimal up to a constant factor. For any total function, one player can send the n-bit input hidden from another player, who can then compute and announce the answer. The functions F_n are highly explicit, computable in polynomial time.

Following [BDPW10], let BPP^{cc}_k, NP^{cc}_k, and P^{cc}_k be the classes of function families with public-coin randomized, nondeterministic, and deterministic NOF protocols of cost $(\log n)^{O(1)}$, respectively. Let $\mathsf{coRP}^{cc}_k \subseteq \mathsf{BPP}^{cc}_k$ be the subclass whose protocols accept every 1-input with probability one. Theorem 1.1 gives the following separation.

Corollary 1.2. *For every fixed $k \geq 3$,*

$$\mathsf{coRP}^{cc}_k \not\subseteq \mathsf{NP}^{cc}_k, \qquad \textit{and hence} \qquad \mathsf{BPP}^{cc}_k \not\subseteq \mathsf{NP}^{cc}_k.$$

Since deterministic communication is at least nondeterministic communication, Theorem 1.1 and Corollary 1.2 also give the following optimal separations, which were unknown previously for general NOF protocols.[2]

Corollary 1.3 (Randomized versus deterministic communication)**.** *For every fixed $k \geq 3$, the explicit family (F_n) from Theorem 1.1 satisfies, for all sufficiently large n,*

$$\mathsf{R}_k(\mathsf{F}_n) = O_k(1), \qquad \mathsf{D}_k(\mathsf{F}_n) = \Omega_k(n).$$

Furthermore,

$$\mathsf{coRP}^{cc}_k \not\subseteq \mathsf{P}^{cc}_k, \qquad \textit{and hence} \qquad \mathsf{BPP}^{cc}_k \not\subseteq \mathsf{P}^{cc}_k.$$

The main difficulty in establishing Theorem 1.1 is that proving explicit nondeterministic lower bounds in the NOF model is highly nontrivial for many candidate functions one can come up with. We take advantage of the generalized inner product (GIP) function, for which a strong lower bound is known [BNS92]. Our construction embeds GIP inside an identity between codewords. Code distance makes the identity cheap to test at random coordinates. On a suitable set of coordinates, the codewords are unrestricted, allowing us to use the GIP bound to show that every accepting cylinder intersection is small.

[2] For one-way protocols, an optimal explicit separation was obtained in [YZ25].

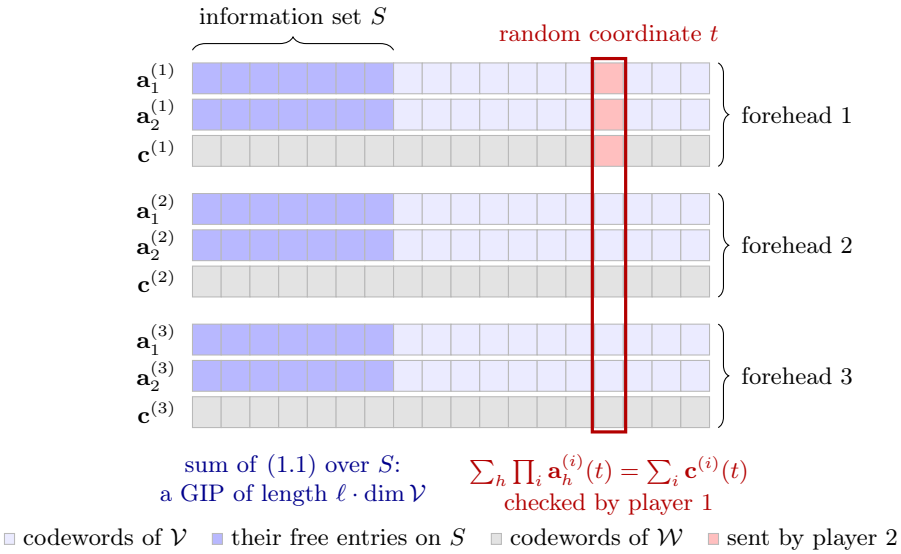


Figure 1: The construction for $k = 3$ players and $\ell = 2$. Each row is a vector in $\mathbb{F}_q^r$, and $\mathbf{a}(t)$ denotes the entry of row $\mathbf{a}$ in column t. Forehead i carries the rows $\mathbf{a}_1^{(i)}, \mathbf{a}_2^{(i)}, \mathbf{c}^{(i)}$. The randomized protocol checks (1.1) at a single random column t, where player 2 sends player 1 the three entries of forehead 1. The lower bound uses the columns of an information set S, where the factor entries are free and the sum of (1.1) over them is a GIP.

1.2 Construction and proof outline

In our NOF function construction, forehead i carries ℓ factor vectors $\mathbf{a}_1^{(i)}, \dots, \mathbf{a}_\ell^{(i)}$ and one vector $\mathbf{c}^{(i)}$, all in $\mathbb{F}_q^r$. The function accepts exactly when

$$\sum_{h=1}^{\ell} \mathbf{a}_h^{(1)} \odot \cdots \odot \mathbf{a}_h^{(k)} = \sum_{i=1}^{k} \mathbf{c}^{(i)}. \tag{1.1}$$

Here $\odot$ denotes coordinatewise multiplication of vectors in $\mathbb{F}_q^r$, and ℓ is a constant depending only on k; see Figure 1. Our starting point is to test (1.1) by checking a randomly sampled coordinate. For this to work, every false identity should fail at a constant fraction of coordinates. We would therefore like to take the factors from a linear code $\mathcal{V} \subseteq \mathbb{F}_q^r$ whose coordinatewise products remain in another linear code $\mathcal{W} \subseteq \mathbb{F}_q^r$ of large relative distance:

$$\mathbf{a}_1 \odot \cdots \odot \mathbf{a}_k \in \mathcal{W} \qquad \text{for all } \mathbf{a}_1, \dots, \mathbf{a}_k \in \mathcal{V}.$$

If each $\mathbf{a}_h^{(i)}$ belongs to $\mathcal{V}$ and each $\mathbf{c}^{(i)}$ belongs to $\mathcal{W}$, then both sides of (1.1) lie in $\mathcal{W}$. Their difference is therefore either zero or nonzero on a constant fraction of coordinates.

Reed–Solomon codes already have this property [Cou22]: multiplying k polynomials of degree below $r/(2k)$ gives a polynomial of degree below $r/2$. So we can take $\mathcal{V}$ and $\mathcal{W}$ to be the evaluations of such polynomials on r points of $\mathbb{F}_q$, with $q \geq r$. Then $\mathcal{W}$ has relative distance at least $1/2$ and $\dim \mathcal{W} \leq k \cdot \dim \mathcal{V}$. With these codes, the construction above already gives an explicit function with $\mathsf{R}_k = O_k(\log n)$ and $\mathsf{N}_k = \Omega_k(n)$, which suffices for Corollary 1.2. The only cost is that the field must grow with the block length, so each field symbol takes $O(\log n)$ bits.

To establish the optimal separation, we use algebraic geometry codes. Algebraic geometry codes have the same property over a fixed field [Cou22, Ran15], which brings the randomized cost down to

$O_k(1)$. We can take q to be a power of two depending only on k, which we treat as a constant from now on, with $\dim \mathcal{V} = \Theta_k(r)$, $\dim \mathcal{W} = O_k(\dim \mathcal{V})$, and constant relative distance for $\mathcal{W}$.

We take every $\mathbf{a}_h^{(i)} \in \mathcal{V}$ and every $\mathbf{c}^{(i)} \in \mathcal{W}$. At each public random coordinate t, player 2 sends the $\ell + 1$ field symbols $\mathbf{a}_1^{(1)}(t), \ldots, \mathbf{a}_\ell^{(1)}(t), \mathbf{c}^{(1)}(t)$ hidden from player 1, who checks the identity at t and announces the result. This costs $O_k(1)$ bits, and a constant number of repetitions rejects every 0-input with probability at least $2/3$. The function remains total if each forehead holds the messages of its codewords as a bit string, so every bit string is a valid input.

Nondeterministic lower bound. The hardness is inherited from GIP. Choose an *information set* S for $\mathcal{V}$: a set of $\dim \mathcal{V}$ coordinates on which every assignment of field symbols extends uniquely to a codeword. The ℓ factors on each forehead are then determined by an unrestricted vector in $\mathbb{F}_q^{\ell \cdot \dim \mathcal{V}}$, namely their values on S. Fix $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)}$ and sum (1.1) over S:

$$\underbrace{\sum_{t \in S} \sum_{h=1}^{\ell} \prod_{i=1}^{k} \mathbf{a}_h^{(i)}(t)}_{\text{generalized inner product}} = \sum_{t \in S} \sum_{i=1}^{k} \mathbf{c}^{(i)}(t).$$

Thus on all 1-inputs with these $\mathbf{c}^{(i)}$, the k-player GIP of length $\ell \cdot \dim \mathcal{V}$ over $\mathbb{F}_q$ takes one fixed value.

A cost-d nondeterministic protocol covers the 1-inputs by at most 2^d accepting cylinder intersections. To lower bound d, we compare the measure of one such intersection with the total measure of the 1-inputs under the uniform distribution on inputs. After fixing the $\mathbf{c}^{(i)}$, passing from each factor codeword to its coordinates on S is a bijection on each forehead, preserving cylinder intersections and the conditional uniform distribution. By the GIP bound [BNS92, BHK01, BDPW10], a cylinder intersection on which GIP is constant has measure at most $2^{-\gamma \ell \cdot \dim \mathcal{V}}$ for some constant $\gamma > 0$. So every accepting cylinder intersection of F_n has measure at most $2^{-\gamma \ell \cdot \dim \mathcal{V}}$, conditioned on any fixing of the $\mathbf{c}^{(i)}$ and therefore also unconditionally. The 1-inputs have measure $q^{-\dim \mathcal{W}} = 2^{-O_k(\dim \mathcal{V})}$, since the factors and $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k-1)}$ determine $\mathbf{c}^{(k)}$. Comparing these measures gives

$$2^d \geq 2^{\gamma \ell \cdot \dim \mathcal{V} - O_k(\dim \mathcal{V})}.$$

The density of the 1-inputs does not depend on ℓ, whereas the GIP bound strengthens as ℓ increases. Taking ℓ to be a sufficiently large constant therefore gives $d = \Omega_k(\dim \mathcal{V})$, which is $\Omega_k(n)$ because each forehead holds $n = (\ell \cdot \dim \mathcal{V} + \dim \mathcal{W}) \log q = \Theta_k(\dim \mathcal{V})$ bits.

Organization. Section 2 gives the preliminaries. Section 3 defines the coded GIP function and the codes we use. Section 4 gives the randomized protocol, and Section 5 proves the nondeterministic lower bound. The appendix contains the construction of the Garcia–Stichtenoth codes.

2 Preliminaries

All logarithms are to base 2. Let $[r] = \{1, \ldots, r\}$. For a prime power q and integer $r \geq 1$, write $\mathbb{F}_q^r$ for the vector space of length-r vectors over $\mathbb{F}_q$. For $\mathbf{a} \in \mathbb{F}_q^r$, write $\mathbf{a}(t)$ for its t-th coordinate and $\mathbf{a}|_S$ for its restriction to a subset $S \subseteq [r]$. The notation $O_k(\cdot)$, $\Omega_k(\cdot)$, and $\Theta_k(\cdot)$ allows implicit constants to depend on the number of players k.

Linear codes. We call $\mathcal{V} \subseteq \mathbb{F}_q^r$ a *linear code* if it is a linear subspace of $\mathbb{F}_q^r$. For $\mathbf{u} \in \mathbb{F}_q^r$, let $\|\mathbf{u}\|_0$ be the number of its nonzero entries. The *minimum distance* of a nonzero linear code $\mathcal{V}$ is

$$\mathsf{dist}(\mathcal{V}) := \min_{\mathbf{u} \in \mathcal{V} \setminus \{\mathbf{0}\}} \|\mathbf{u}\|_0.$$

Given a linear code $\mathcal{V}$, define its k-product by

$$\mathcal{V}^{\odot k} := \mathrm{span}_{\mathbb{F}_q}\{\mathbf{u}_1 \odot \cdots \odot \mathbf{u}_k : \mathbf{u}_i \in \mathcal{V},\ i \in [k]\},$$

where $\odot$ denotes coordinatewise multiplication in $\mathbb{F}_q^r$.

Information sets. Let $\mathcal{V} \subseteq \mathbb{F}_q^r$ have dimension m. A set $S \subseteq [r]$ is an *information set* if $|S| = m$ and the map sending each codeword $\mathbf{a} \in \mathcal{V}$ to $\mathbf{a}|_S \in \mathbb{F}_q^m$ is a bijection. A *generator matrix* of $\mathcal{V}$ is a matrix whose rows form a basis of $\mathcal{V}$. It has rank m, so choosing m linearly independent columns gives an information set. Let $\rho_S : \mathbb{F}_q^m \to \mathcal{V}$ be the inverse of the restriction map $\mathbf{a} \mapsto \mathbf{a}|_S$. For $\mathbf{a} \in \mathcal{V}$ and $\mathbf{y} \in \mathbb{F}_q^m$, $\rho_S(\mathbf{y})|_S = \mathbf{y}$ and $\rho_S(\mathbf{a}|_S) = \mathbf{a}$. If $\mathbf{a}$ is uniform in $\mathcal{V}$, the bijection gives $\Pr[\mathbf{a}|_S = \mathbf{y}] = q^{-m}$ for every $\mathbf{y} \in \mathbb{F}_q^m$.

NOF communication. For a Boolean function $\mathsf{F} : \mathcal{X}_1 \times \cdots \times \mathcal{X}_k \to \{0,1\}$, player i sees every input block except the ith. The players communicate on a public blackboard, and local computation is free. In a deterministic protocol, each message is determined by the blackboard so far and the inputs visible to the speaker. A public-coin randomized protocol also has access to shared random bits, which are free. Its cost is the maximum communication over inputs and random choices, and $\mathsf{R}_k(\mathsf{F})$ is the minimum cost with error at most $1/3$ on every input. A protocol has one-sided error if it accepts every 1-input with probability one.

A nondeterministic protocol provides the players with a common certificate of a fixed length, which they verify by a deterministic NOF protocol. A 1-input has some certificate causing the protocol to accept, whereas no certificate is accepted on a 0-input. The cost of a nondeterministic protocol is the certificate length plus the worst-case communication cost of the verification protocol. We write $\mathsf{N}_k(\mathsf{F})$ for the minimum such cost. We write $\mathsf{D}_k(\mathsf{F})$ for the minimum worst-case communication cost of a deterministic protocol for F.

Cylinder intersections. Given a product space $\mathcal{X}_1 \times \cdots \times \mathcal{X}_k$, a set $\mathcal{A}_i$ is an *i-cylinder* if membership in $\mathcal{A}_i$ does not depend on the ith input: whenever $(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(k)}) \in \mathcal{A}_i$, changing $\mathbf{x}^{(i)}$ to any other element of $\mathcal{X}_i$ keeps the input in $\mathcal{A}_i$. A set $\mathcal{A}$ is a *cylinder intersection* if $\mathcal{A} = \bigcap_{i=1}^k \mathcal{A}_i$ for cylinders $\mathcal{A}_1, \ldots, \mathcal{A}_k$. It is *accepting* for a Boolean function F if $\mathcal{A} \subseteq \mathsf{F}^{-1}(1)$.

Generalized inner product. The generalized inner product (GIP) is

$$\mathsf{GIP}_{q,w}^k(\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)}) = \sum_{t=1}^{w} \prod_{i=1}^{k} \mathbf{y}^{(i)}(t), \qquad \mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)} \in \mathbb{F}_q^w.$$

3 The coded GIP problem

Definition 3.1 (Compatible codes)**.** Fix an integer $k \geq 1$ and $\varepsilon, \delta \in (0, 1]$. Two linear codes $\mathcal{V}, \mathcal{W} \subseteq \mathbb{F}_q^r$ are (ε, δ)-*compatible* for k-products if the following hold:

1. *Product containment:* $\mathcal{V}^{\odot k} \subseteq \mathcal{W}$.
2. *Dimension of the small code:* $\dim(\mathcal{V}) \geq \varepsilon r$.
3. *Distance of the large code:* $\mathsf{dist}(\mathcal{W}) \geq \delta r$.

Our function tests whether a sum of products of $\mathcal{V}$-codewords is equal to a sum of $\mathcal{W}$-codewords. Product containment puts both sides in $\mathcal{W}$, and distance ensures that a uniformly random coordinate detects a false identity with probability at least δ. The dimension bound provides a large information set on which we embed GIP for the nondeterministic lower bound.

Garcia–Stichtenoth codes provide such pairs over a field whose size is independent of the code length. This allows each sampled coordinate to be checked with $O_k(1)$ communication.

Theorem 3.2 (Garcia–Stichtenoth compatible codes, [GS96, SAK+01])**.** *Let $k \geq 2$ be a fixed integer, and let q be a power of 4 satisfying $(32k)^2 \leq q \leq (64k)^2$. For every integer $j \geq 0$, let $r_j = (q - \sqrt{q})q^{j/2}$. Then there exist linear codes $\mathcal{V}_j, \mathcal{W}_j \subseteq \mathbb{F}_q^{r_j}$ that are $(6/\sqrt{q}, 3/4)$-compatible for k-products.*

Furthermore, for fixed q, generator matrices for both codes can be constructed and codewords can be encoded deterministically in $\mathrm{poly}(r_j)$ *time.*

Fix $k \geq 2$, take q and the code family from Theorem 3.2. For a fixed level $j \geq 0$, write

$$\mathcal{V} := \mathcal{V}_j, \quad \mathcal{W} := \mathcal{W}_j, \quad r := r_j, \quad m := m_j := \dim(\mathcal{V}_j), \quad \widetilde{m} := \widetilde{m}_j := \dim(\mathcal{W}_j).$$

By Definition 3.1, Theorem 3.2 and the Singleton bound, the dimensions satisfy

$$6r/\sqrt{q} \leq m \leq \widetilde{m} \leq r/4 + 1. \tag{3.1}$$

Since q depends only on k, both code dimensions are $\Theta_k(r)$.

The problem. Now we describe our NOF problem: each forehead will encode ℓ codewords from $\mathcal{V}$ and one from $\mathcal{W}$. Set

$$\ell := 2^{k+2}k, \qquad n := n_j := (\ell m + \widetilde{m}) \log q.$$

Then $n = \Theta_k(r)$ and $m = \Theta_k(n)$. Arbitrary input lengths are handled by standard padding. Fix bases of $\mathcal{V}$ and $\mathcal{W}$ and a binary representation of $\mathbb{F}_q$. Split each $\mathbf{x} \in \{0,1\}^n$ into blocks

$$\mathbf{x} = (\mathbf{u}_1, \ldots, \mathbf{u}_\ell, \mathbf{v}), \qquad \forall h \in [\ell],\ \mathbf{u}_h \in \{0,1\}^{m \log q}, \quad \mathbf{v} \in \{0,1\}^{\widetilde{m} \log q}.$$

Interpreting consecutive groups of $\log q$ bits as field elements, use these blocks as the basis coefficients of codewords $\mathbf{a}_1, \ldots, \mathbf{a}_\ell \in \mathcal{V}$ and $\mathbf{c} \in \mathcal{W}$. This gives a bijection between $\{0,1\}^n$ and $\mathcal{V}^\ell \times \mathcal{W}$. We identify $\mathbf{x}$ with $(\mathbf{a}_1, \ldots, \mathbf{a}_\ell, \mathbf{c})$ under this bijection.

Definition 3.3 ($\mathsf{CodedGIP}$ function)**.** For inputs $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(k)} \in \{0,1\}^n$, represented as $\mathbf{x}^{(i)} = (\mathbf{a}_1^{(i)}, \ldots, \mathbf{a}_\ell^{(i)}, \mathbf{c}^{(i)})$, define the coded generalized inner product function $\mathsf{CodedGIP} : (\{0,1\}^n)^k \to \{0,1\}$ by

$$\mathsf{CodedGIP}(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(k)}) = \mathbf{1}\left\{ \sum_{h=1}^{\ell} \mathbf{a}_h^{(1)} \odot \cdots \odot \mathbf{a}_h^{(k)} = \sum_{i=1}^{k} \mathbf{c}^{(i)} \right\}.$$

In the NOF model, player i sees all input blocks except $\mathbf{x}^{(i)}$. We prove the randomized upper bound in Section 4 and the nondeterministic lower bound in Section 5.

4 Constant-cost randomized communication protocol

In this section, we give a randomized protocol with $O_k(1)$ communication for $\mathsf{CodedGIP}$. Let

$$\mathbf{b} := \sum_{h=1}^{\ell} \mathbf{a}_h^{(1)} \odot \cdots \odot \mathbf{a}_h^{(k)} - \sum_{i=1}^{k} \mathbf{c}^{(i)}.$$

Thus $\mathsf{CodedGIP} = \mathbf{1}\{\mathbf{b} = \mathbf{0}\}$. By Theorem 3.2, $\mathbf{b} \in \mathcal{W}$ and if $\mathbf{b} \neq \mathbf{0}$, at least $3r/4$ of its coordinates are nonzero. Therefore checking one random coordinate suffices.

Theorem 4.1. *There exists a randomized protocol for* $\mathsf{CodedGIP}$ *in the* k*-player NOF model with cost* $O_k(1)$*. When* $\mathsf{CodedGIP} = 1$*, it always outputs* 1*; when* $\mathsf{CodedGIP} = 0$*, it outputs* 0 *with probability at least* $3/4$*.*

Proof. We first describe the protocol. The players view their visible input blocks as codewords.

1. Choose a uniformly random coordinate $t \in [r]$ by public randomness.
2. Player 2 sends $\mathbf{a}_1^{(1)}(t), \ldots, \mathbf{a}_\ell^{(1)}(t), \mathbf{c}^{(1)}(t)$.
3. Player 1 uses these values and the visible blocks $\mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(k)}$ to compute

$$\mathbf{b}(t) = \sum_{h=1}^{\ell} \prod_{i=1}^{k} \mathbf{a}_h^{(i)}(t) - \sum_{i=1}^{k} \mathbf{c}^{(i)}(t),$$

   and sends the output bit $\mathbf{1}\{\mathbf{b}(t) = 0\}$.

If $\mathsf{CodedGIP} = 1$, $\mathbf{b} = \mathbf{0}$. So the protocol always outputs 1. If $\mathsf{CodedGIP} = 0$, then $\mathbf{b} \in \mathcal{W} \setminus \{\mathbf{0}\}$, and at least $3r/4$ coordinates of $\mathbf{b}$ are nonzero. Therefore the protocol outputs 0 with probability at least $3/4$. Player 2 sees $\mathbf{x}^{(1)}$ and sends $\ell + 1$ field elements, each represented by $\log q$ bits. Since $\ell = 2^{k+2}k$ and $q = \Theta(k^2)$, the total cost is $(\ell + 1)\log q + 1 = O_k(1)$. □

5 Nondeterministic lower bound

In this section, we prove the nondeterministic communication lower bound. A cost-d protocol covers all 1-inputs by at most 2^d accepting cylinder intersections. Using a known correlation bound for GIP, we show that each accepting cylinder intersection contains only a $2^{-\Omega_k(n)}$ fraction of the 1-inputs, which forces $d = \Omega_k(n)$.

Theorem 5.1. *The nondeterministic communication complexity of* $\mathsf{CodedGIP}$ *satisfies*

$$\mathsf{N}_k(\mathsf{CodedGIP}) = \Omega_k(n).$$

5.1 Cylinder intersections and GIP

We first recall the cylinder-cover bound and a correlation bound for GIP.

Fact 5.2 ([BDPW10, proof of Theorem 3.4])**.** *If* F *has a nondeterministic NOF protocol with cost* d*, then* $\mathsf{F}^{-1}(1)$ *can be covered by at most* 2^d *accepting cylinder intersections.*

For any positive integer w, let ν be the uniform distribution on $(\mathbb{F}_q^w)^k$. The following estimate first appears in [BHK01]; we use the form stated in [BDPW10].

Lemma 5.3 ([BDPW10, Fact 4.11])**.** *Fix positive integers $k \geq 2$ and $w \geq 1$. Let q be a power of 2. For every cylinder intersection $\mathcal{B} \subseteq (\mathbb{F}_q^w)^k$ and $\tau \in \mathbb{F}_q$,*

$$\left| \nu\left(\mathcal{B} \cap (\mathsf{GIP}_{q,w}^k)^{-1}(\tau)\right) - \frac{\nu(\mathcal{B})}{q} \right| \leq \left(1 - \frac{1}{q}\right)\left(1 - \left(1 - \frac{1}{q}\right)^{k-1}\right)^{w/2^{k-1}}.$$

After fixing the codewords $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)}$, we use an information set of $\mathcal{V}$ to express each accepting cylinder intersection in coordinates on which GIP is constant. Recall that $m = \dim(\mathcal{V})$, and let μ and ν be uniform on $(\{0,1\}^n)^k$ and $(\mathbb{F}_q^{\ell m})^k$, respectively.

Lemma 5.4. *For every cylinder intersection $\mathcal{A} \subseteq \mathsf{CodedGIP}^{-1}(1)$ and fixed $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)} \in \mathcal{W}$, there exist a cylinder intersection $\mathcal{B} \subseteq (\mathbb{F}_q^{\ell m})^k$ and $\tau \in \mathbb{F}_q$ such that*

$$\nu(\mathcal{B}) = \mu\big(\mathcal{A} \mid \mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)}\big), \qquad \mathcal{B} \subseteq (\mathsf{GIP}_{q,\ell m}^k)^{-1}(\tau).$$

Proof. Fix $\mathcal{A} = \bigcap_{i=1}^k \mathcal{A}_i \subseteq \mathsf{CodedGIP}^{-1}(1)$ and $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)} \in \mathcal{W}$, where each $\mathcal{A}_i$ is an i-cylinder. Choose an information set $S \subseteq [r]$ of $\mathcal{V}$. The entries on S give unrestricted coordinates for each factor codeword; summing the acceptance identity over S fixes the GIP of these coordinates. For $i \in [k]$ and $h \in [\ell]$, write

$$\mathbf{y}^{(i)} = (\mathbf{y}_1^{(i)}, \ldots, \mathbf{y}_\ell^{(i)}) \in \mathbb{F}_q^{\ell m}, \qquad \mathbf{y}_h^{(i)} = \mathbf{a}_h^{(i)}|_S \in \mathbb{F}_q^m.$$

To write the original cylinder conditions using these entries, recover each input with the inverse ρ_S defined in Section 2:

$$\Phi_i : \mathbb{F}_q^{\ell m} \longrightarrow \mathcal{V}^\ell \times \{\mathbf{c}^{(i)}\}, \qquad \Phi_i(\mathbf{y}^{(i)}) = \big(\rho_S(\mathbf{y}_1^{(i)}), \ldots, \rho_S(\mathbf{y}_\ell^{(i)}), \mathbf{c}^{(i)}\big).$$

Set $\Phi(\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)}) = (\Phi_1(\mathbf{y}^{(1)}), \ldots, \Phi_k(\mathbf{y}^{(k)}))$ and

$$\mathcal{B}_i = \Phi^{-1}(\mathcal{A}_i), \qquad \mathcal{B} = \Phi^{-1}(\mathcal{A}) = \bigcap_{i=1}^k \mathcal{B}_i.$$

Changing only $\mathbf{y}^{(i)}$ changes only $\Phi_i(\mathbf{y}^{(i)})$, which does not affect membership in $\mathcal{A}_i$. Hence each $\mathcal{B}_i$ is an i-cylinder, so $\mathcal{B}$ is a cylinder intersection.

Next, we verify the conditional measure identity. The map Φ is a bijection onto the inputs with the $\mathbf{c}^{(i)}$ fixed, which are uniform under μ conditioned on $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)}$. Therefore,

$$\nu(\mathcal{B}) = \frac{|\mathcal{B}|}{q^{k\ell m}} = \mu\big(\mathcal{A} \mid \mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(k)}\big).$$

It remains to show that GIP has a fixed value on $\mathcal{B}$. For every $(\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)}) \in \mathcal{B}$, summing the acceptance condition in Definition 3.3 over $t \in S$ gives

$$\mathsf{GIP}_{q,\ell m}^k(\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)}) = \sum_{h=1}^{\ell} \sum_{t \in S} \prod_{i=1}^k \mathbf{a}_h^{(i)}(t) = \sum_{t \in S} \sum_{i=1}^k \mathbf{c}^{(i)}(t) =: \tau.$$

Since the $\mathbf{c}^{(i)}$ are fixed, $\tau \in \mathbb{F}_q$ is independent of $(\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(k)})$, and $\mathcal{B} \subseteq (\mathsf{GIP}_{q,\ell m}^k)^{-1}(\tau)$. □

5.2 The lower bound proof

We now compare the measure of all 1-inputs with the measure of a single accepting cylinder intersection to finish the lower bound proof.

Proof of Theorem 5.1. Under the uniform distribution μ, the codewords $\mathbf{a}_h^{(i)}$ and $\mathbf{c}^{(i)}$ are mutually independent and uniform in $\mathcal{V}$ and $\mathcal{W}$, respectively.

Step 1. Compute the measure of all 1-inputs. Fix all $\mathbf{a}_h^{(i)}$ and $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(k-1)}$. The input is a 1-input exactly when

$$\mathbf{c}^{(k)} = \sum_{h=1}^{\ell} \mathbf{a}_h^{(1)} \odot \cdots \odot \mathbf{a}_h^{(k)} - \sum_{i=1}^{k-1} \mathbf{c}^{(i)}.$$

Since $\mathcal{V}^{\odot k} \subseteq \mathcal{W}$ and $\mathcal{W}$ is linear, the right-hand side lies in $\mathcal{W}$, so exactly one choice of $\mathbf{c}^{(k)}$ gives a 1-input. Recalling that $\widetilde{m} = \dim(\mathcal{W})$, we obtain

$$\mu\big(\mathsf{CodedGIP}^{-1}(1)\big) = q^{-\widetilde{m}}. \tag{5.1}$$

Step 2. Bound the measure of one accepting cylinder intersection. Let $\mathcal{A}$ be an accepting cylinder intersection and fix $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(k)}$. Lemma 5.4 gives a cylinder intersection $\mathcal{B} \subseteq (\mathbb{F}_q^{\ell m})^k$ with $\nu(\mathcal{B}) = \mu(\mathcal{A} \mid \mathbf{c}^{(1)}, \dots, \mathbf{c}^{(k)})$ and a value $\tau \in \mathbb{F}_q$ such that $\mathcal{B} \subseteq (\mathsf{GIP}_{q,\ell m}^k)^{-1}(\tau)$. Since q is a power of 2, applying Lemma 5.3 with $w = \ell m$ we have

$$\nu(\mathcal{B}) = \frac{\big|\nu\big(\mathcal{B} \cap (\mathsf{GIP}_{q,\ell m}^k)^{-1}(\tau)\big) - \nu(\mathcal{B})/q\big|}{1 - 1/q} \le \left(1 - \left(1 - \frac{1}{q}\right)^{k-1}\right)^{\ell m / 2^{k-1}} \le q^{-4km}.$$

The last inequality uses $\ell = 2^{k+2}k$ and $k - 1 \le \sqrt{q}$. By the measure identity of Lemma 5.4 and averaging over $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(k)}$, we have

$$\mu(\mathcal{A}) \le q^{-4km}. \tag{5.2}$$

Step 3. Derive the communication lower bound. We compare (5.1) and (5.2) to bound the number of accepting cylinder intersections in any protocol. Suppose a nondeterministic protocol has cost d. By Fact 5.2, its 1-inputs are covered by accepting cylinder intersections $\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(s)}$ with $s \le 2^d$. Hence

$$q^{-\widetilde{m}} \le \sum_{a=1}^{s} \mu(\mathcal{A}^{(a)}) \le 2^d q^{-4km}.$$

Recall the parameter choice $(32k)^2 \le q \le (64k)^2$ in Theorem 3.2 and the dimensions $6r/\sqrt{q} \le m \le \widetilde{m} \le r/4 + 1$ in (3.1). Together with $n = (\ell m + \widetilde{m}) \log q$, we obtain

$$d \ge (4km - \widetilde{m}) \log q > \left(\frac{r}{8} - 1\right) \log q = \Omega_k(n). \qquad \square$$

Appendix A Garcia–Stichtenoth codes

In this appendix, we prove Theorem 3.2 using the second Garcia–Stichtenoth tower [GS96] and the basis algorithm of [SAK+01]. Throughout, fix k and q as in Theorem 3.2.

A.1 Functions, orders, and evaluation

Let t be a formal variable, and let $\mathbb{F}_q(t)$ consist of the ratios of polynomials over $\mathbb{F}_q$ with nonzero denominators. We work with functions in a field $\mathcal{K}$ containing $\mathbb{F}_q(t)$ and finite-dimensional over it. A point P is described by an order map $v_P : \mathcal{K} \to \mathbb{Z} \cup \{\infty\}$ satisfying

$$v_P(0) = \infty, \qquad v_P(\mathcal{K} \setminus \{0\}) = \mathbb{Z}, \qquad v_P(c) = 0 \quad (c \in \mathbb{F}_q \setminus \{0\}),$$

and, for nonzero $f, g \in \mathcal{K}$,

$$v_P(fg) = v_P(f) + v_P(g), \qquad v_P(f+g) \geq \min\{v_P(f), v_P(g)\}. \tag{A.1}$$

A positive order means a zero, and a negative order means a pole, whose order is $-v_P(f)$. We call P a *rational point* if every f with $v_P(f) \geq 0$ has a unique $c \in \mathbb{F}_q$ such that $v_P(f-c) > 0$; we then define $f(P) = c$. The rules in (A.1) imply that evaluation preserves addition and multiplication.

For example, in $\mathbb{F}_q(t)$ the point $t = c$ records the exponent of $t - c$ in a rational expression after cancellation. To describe infinity, use $u = 1/t$ and examine $u = 0$. For nonzero polynomials $A, B \in \mathbb{F}_q[t]$,

$$\frac{A(t)}{B(t)} = u^{\deg B - \deg A} \frac{u^{\deg A} A(1/u)}{u^{\deg B} B(1/u)}, \qquad v_\infty(A/B) = \deg B - \deg A.$$

The numerator and denominator of the last fraction have nonzero values at $u = 0$, namely the leading coefficients of A and B. In particular, $v_\infty(t) = -1$ and $v_\infty(u) = 1$.

Fix a rational point P_∞. For an integer $a \geq 0$, define the space of functions with no poles except possibly at P_∞, and with pole order at most a there:

$$\mathcal{L}(aP_\infty) = \{f \in \mathcal{K} : v_{P_\infty}(f) \geq -a,\ v_P(f) \geq 0 \text{ for every } P \neq P_\infty\}.$$

This is a vector space over $\mathbb{F}_q$. The standard Riemann–Roch dimension formula associates to $\mathcal{K}$ a nonnegative integer g, called its *genus*, such that

$$a \geq 2g - 1 \quad \Longrightarrow \quad \dim_{\mathbb{F}_q} \mathcal{L}(aP_\infty) = a + 1 - g. \tag{A.2}$$

We also use the standard bound that a nonzero function in this space has at most a distinct rational zeros away from P_∞:

$$0 \neq f \in \mathcal{L}(aP_\infty) \quad \Longrightarrow \quad \big|\{P \neq P_\infty : P \text{ is rational and } f(P) = 0\}\big| \leq a. \tag{A.3}$$

A.2 The GS tower and its computation

The second Garcia–Stichtenoth tower starts with a formal variable t_0 and is defined recursively by

$$\mathcal{K}_0 = \mathbb{F}_q(t_0), \qquad \mathcal{K}_{h+1} = \mathcal{K}_h(t_{h+1}), \qquad t_{h+1}^{\sqrt{q}} + t_{h+1} = \frac{t_h^{\sqrt{q}}}{t_h^{\sqrt{q}-1} + 1} \quad (h \geq 0). \tag{A.4}$$

At level j, let $\mathcal{P}_j$ consist of the tuples $(a_0, \dots, a_j) \in \mathbb{F}_q^{j+1}$ satisfying

$$a_0^{\sqrt{q}} + a_0 \neq 0, \qquad a_{h+1}^{\sqrt{q}} + a_{h+1} = \frac{a_h^{\sqrt{q}}}{a_h^{\sqrt{q}-1} + 1} \quad (0 \le h < j). \tag{A.5}$$

Here $\mathbb{F}_{\sqrt{q}}$ denotes the subfield of $\mathbb{F}_q$ with $\sqrt{q}$ elements. All denominators in (A.5) are nonzero. Indeed, if $a^{\sqrt{q}} + a \neq 0$, then

$$\frac{a^{\sqrt{q}}}{a^{\sqrt{q}-1} + 1} = \frac{a^{\sqrt{q}+1}}{a^{\sqrt{q}} + a} \in \mathbb{F}_{\sqrt{q}} \setminus \{0\}.$$

The additive map $y \mapsto y^{\sqrt{q}} + y$ has exactly $\sqrt{q}$ roots and takes its values in $\mathbb{F}_{\sqrt{q}}$, so each element of $\mathbb{F}_{\sqrt{q}}$ is attained exactly $\sqrt{q}$ times. Thus there are $q - \sqrt{q}$ choices for a_0 and $\sqrt{q}$ extensions of every prefix, each still satisfying $a_{h+1}^{\sqrt{q}} + a_{h+1} \neq 0$. Hence

$$|\mathcal{P}_j| = (q - \sqrt{q}) q^{j/2} = r_j.$$

Garcia and Stichtenoth [GS96] prove that each tuple $\mathbf{a} \in \mathcal{P}_j$ gives a distinct rational point $P_{\mathbf{a}}$ of $\mathcal{K}_j$ with $t_h(P_{\mathbf{a}}) = a_h$. They also prove that t_0 has a unique pole in $\mathcal{K}_j$, at a rational point denoted by P_∞, and that

$$v_{P_\infty}(t_h) = -q^{(j-h)/2} \quad (0 \le h \le j), \qquad g \le \frac{r_j}{\sqrt{q} - 1}. \tag{A.6}$$

To compute at P_∞, put $u_h = 1/t_h$ in (A.4), obtaining

$$u_h + u_h^{\sqrt{q}} = \frac{u_{h+1}^{\sqrt{q}}}{1 + u_{h+1}^{\sqrt{q}-1}}, \qquad v_{P_\infty}(u_h) = q^{(j-h)/2}.$$

Every u_h has value 0 at P_∞, the denominator has value 1, and $v_{P_\infty}(u_j) = 1$.

Fix an order of $\mathcal{P}_j$ and write its points as $P_1, \dots, P_{r_j}$. For functions with no poles at these points, define

$$\mathrm{ev}(f) = (f(P_1), \dots, f(P_{r_j})) \in \mathbb{F}_q^{r_j}.$$

To compute a generator matrix, we need a basis of the chosen function space and its values at these points. For fixed q, given $j \ge 1$ and an integer $0 \le a < r_j$, the algorithm in Section IV-A and Theorem 7 of [SAK+01] computes a basis of $\mathcal{L}(aP_\infty)$ and its evaluations deterministically in $\mathrm{poly}(r_j)$ time. Its even-characteristic assumption holds because q is a power of 4. For $j = 0$, use the basis $1, t_0, \dots, t_0^a$ and evaluate at the elements $c \in \mathbb{F}_q$ with $c^{\sqrt{q}} + c \neq 0$.

A.3 The compatible codes

We choose pole bounds $8r_j/\sqrt{q} - 1$ and $r_j/4 - 1$ to allow k-products while keeping at least $3r_j/4$ nonzero evaluations.

Proof of Theorem 3.2. **Construction.** Fix $j \ge 0$. In $\mathcal{K}_j$, set

$$\mathcal{V}_j = \mathrm{ev}\big(\mathcal{L}((8r_j/\sqrt{q} - 1)P_\infty)\big), \qquad \mathcal{W}_j = \mathrm{ev}\big(\mathcal{L}((r_j/4 - 1)P_\infty)\big). \tag{A.7}$$

These are linear codes. Both pole bounds are integers, and (A.6) together with the choice of q in Theorem 3.2 gives

$$g \le \frac{2r_j}{\sqrt{q}}, \qquad 2g - 1 \le \frac{8r_j}{\sqrt{q}} - 1 \le \frac{r_j}{4} - 1 < r_j. \tag{A.8}$$

Products. Since $k \le \sqrt{q}/32$ from Theorem 3.2, we have

$$k\left(\frac{8r_j}{\sqrt{q}} - 1\right) \le \frac{r_j}{4} - k \le \frac{r_j}{4} - 1.$$

Thus the product of any $f_1, \ldots, f_k \in \mathcal{L}((8r_j/\sqrt{q} - 1)P_\infty)$ has no poles away from P_∞ and pole order at most $r_j/4 - 1$ there, by (A.1). Since evaluation preserves multiplication,

$$\operatorname{ev}(f_1) \odot \cdots \odot \operatorname{ev}(f_k) = \operatorname{ev}\left(\prod_{i=1}^{k} f_i\right) \in \mathcal{W}_j.$$

Taking the linear span proves $\mathcal{V}_j^{\odot k} \subseteq \mathcal{W}_j$.

Distance and dimension. For $0 \le a < r_j$ and nonzero $f \in \mathcal{L}(aP_\infty)$, (A.3) gives

$$\|\operatorname{ev}(f)\|_0 \ge r_j - a > 0.$$

Evaluation is therefore injective on both spaces, and $\mathsf{dist}(\mathcal{W}_j) \ge 3r_j/4$. By (A.2) and (A.8),

$$\frac{6r_j}{\sqrt{q}} \le m_j = \frac{8r_j}{\sqrt{q}} - g \le \frac{8r_j}{\sqrt{q}}, \qquad \widetilde{m}_j = \frac{r_j}{4} - g \le \frac{r_j}{4}.$$

Thus $\mathcal{V}_j, \mathcal{W}_j$ are $(6/\sqrt{q}, 3/4)$-compatible for k-products.

Computation. Fix an order of $\mathbb{F}_q$. Enumerate a_0 with $a_0^{\sqrt{q}} + a_0 \ne 0$, then extend each prefix by trying all q choices for the next coordinate in (A.5). Apply the cited basis algorithm to the two function spaces in (A.7), using the polynomial bases above when $j = 0$. The evaluations of each basis form the rows of a generator matrix. Encoding multiplies the basis coefficients by this matrix. For fixed q, these constructions and encoding take $\operatorname{poly}(r_j)$ time. See [SAK+01] for more details. □

A.4 Garcia–Stichtenoth codes in TCS works

Rubinstein [Rub18] uses GS-based codes in an MA protocol for Set Disjointness to prove SETH-based hardness of approximate nearest-neighbor search. [KLM18] extend this approach to Merlin-assisted multiparty protocols using the k-product property, obtaining hardness of approximating Dominating Set.

For proof systems, [BSKK+16] use GS-derived codes to obtain non-uniform PCPs with linear proof length and sublinear query complexity. [BSCG+17] combine GS codes with interaction to obtain three-round IOPs with linear proof length and constant bit-query complexity. [BHI+24] use AG multiplication codes to construct linear-length dot-product proofs over fixed fields, verified by one linear query to the input and proof together.

Other applications include secret sharing and secure computation over small fields [CC06, Che19, HVW24]. Variants of GS codes support improved list decoding and multipoint local decoding [GP08, GX22, CXY20]. They also yield pseudorandom generators and codes resilient to computationally bounded channels [Shp09, CTS13, SS21]. Their recursive structure enables subquadratic-time encoding after preprocessing [NW19].

We bring this tool to NOF communication. The randomized protocol uses k-product containment and constant relative distance to detect a false identity by checking one random coordinate. GS codes combine a fixed field with linearly many independent coefficients, giving $O_k(1)$ communication; Reed–Solomon codes require a growing field and give $O_k(\log n)$ in the same construction. A nondeterministic protocol must verify the global vector identity, although every player is missing one factor in each product. After fixing the $\mathcal{W}$-codewords, summing over an information set gives a generalized inner product on linearly many independent field symbols; Lemma 5.3 then bounds the measure of each accepting cylinder intersection and yields the $\Omega_k(n)$ lower bound.

Acknowledgement

The authors would like to thank Shachar Lovett, Raghu Meka and Guangxu Yang for valuable feedback on an earlier version of this paper. The authors also thank the Simons Institute for the Theory of Computing for hosting the ICM 2026 Satellite Conference, where the authors initiated this collaboration.

AI methodology. We used GPT-5.6 Pro in developing the construction. The project grew out of an earlier work on a two-party separation between BPP and P^{NP}, inspired by Gavinsky's work [Gav26]. Motivated by that development, the authors proposed an algebraic approach to separating BPP from NP in the NOF model using the PCP machinery. During the discussions with AI, the AI suggested a construction using the k-product structure of Reed–Solomon codes which has been exploited in the PCP literature, unknown to the authors at the time. It later suggested using algebraic geometry codes to obtain the optimal separation. This process produced the construction described in Section 1 and developed in Section 3 and Appendix A.

The authors prepared initial drafts and developed the technical exposition. Codex GPT-6 and Claude Opus-5.5 assisted with polishing, restructuring, and iterative discussion of the presentation, particularly the introduction for this version of the manuscript. The authors reviewed and finalized the mathematical arguments and exposition and take responsibility for the correctness, originality, and attribution of the entire manuscript.

References

[AS20] 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.

[BDPW10] P. Beame, M. David, T. Pitassi, and P. Woelfel. Separating deterministic from randomized multiparty communication complexity. *Theory of Computing*, 6:201–225, 2010.

[BH12] P. Beame and T. Huynh. Multiparty communication complexity and threshold circuit size of AC^0. *SIAM Journal on Computing*, 41(3):484–518, 2012.

[BHI+24] Nir Bitansky, Prahladh Harsha, Yuval Ishai, Ron D. Rothblum, and David J. Wu. Dot-product proofs and their applications. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 806–825. IEEE, 2024.

[BHK01] László Babai, Thomas P. Hayes, and Peter G. Kimmel. The cost of the missing bit: Communication complexity with help. *Combinatorica*, 21(4):455–488, 2001.

[BNS92] 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(2):204–232, 1992.

[BPS07] Paul Beame, Toniann Pitassi, and Nathan Segerlind. Lower bounds for Lovász–Schrijver systems and beyond follow from multiparty communication complexity. *SIAM Journal on Computing*, 37(3):845–869, 2007.

[BSCG+17] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. Interactive oracle proofs with constant rate and query complexity. In *44th*

International Colloquium on Automata, Languages, and Programming (ICALP 2017), volume 80 of *Leibniz International Proceedings in Informatics*, pages 40:1–40:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017.

[BSKK+16] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. Constant rate PCPs for Circuit-SAT with sublinear query complexity. *Journal of the ACM*, 63(4):32:1–32:57, 2016.

[CC06] Hao Chen and Ronald Cramer. Algebraic geometric secret sharing schemes and secure multi-party computations over small fields. In *Advances in Cryptology—CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 521–536. Springer, 2006.

[CFL83] A. K. Chandra, M. L. Furst, and R. J. Lipton. Multi-party protocols. In *Proceedings of the 15th ACM Symposium on Theory of Computing*, pages 94–99, 1983.

[Che19] Mahdi Cheraghchi. Nearly optimal robust secret sharing. *Designs, Codes and Cryptography*, 87(8):1777–1796, 2019.

[Cou22] A. Couvreur. How arithmetic and geometry make error-correcting codes better. arXiv:2110.11282, 2022.

[CTS13] Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. Technical Report TR13-155, Electronic Colloquium on Computational Complexity, 2013. Revision 2, March 2014.

[CXY20] Ronald Cramer, Chaoping Xing, and Chen Yuan. Efficient multi-point local decoding of Reed–Muller codes via interleaved codex. *IEEE Transactions on Information Theory*, 66(1):263–272, 2020.

[Gav26] Dmytro Gavinsky. On the quantum communication complexity of total functions. *arXiv preprint arXiv:2608.18784*, 2026.

[GP08] Venkatesan Guruswami and Anindya C. Patthak. Correlated algebraic-geometric codes: Improved list decoding over bounded alphabets. *Mathematics of Computation*, 77(261):447–473, 2008.

[GS96] A. Garcia and H. Stichtenoth. On the asymptotic behaviour of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248–273, 1996.

[GX22] Venkatesan Guruswami and Chaoping Xing. Optimal rate list decoding over bounded alphabets using algebraic-geometric codes. *Journal of the ACM*, 69(2):10:1–10:48, 2022.

[HG91] J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Computational Complexity*, 1:113–129, 1991.

[HVW24] Carmit Hazay, Muthuramakrishnan Venkitasubramaniam, and Mor Weiss. The price of active security in cryptographic protocols. *Journal of Cryptology*, 37(3):30, 2024.

[JLL+25] Michael Jaber, Yang P. Liu, Shachar Lovett, Anthony Ostuni, and Mehtaab Sawhney. Quasipolynomial bounds for the corners theorem. In *Proceedings of the 66th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2025. arXiv:2504.07006.

[KL25] Zander Kelley and Xin Lyu. More efficient sifting for grid norms, and applications to multiparty communication complexity. In *Proceedings of the 66th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2025. arXiv:2505.01587.

[KLM18] Karthik C. S., Bundit Laekhanukit, and Pasin Manurangsi. On the parameterized complexity of approximating dominating set. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1283–1296. ACM, 2018.

[KLM24] Z. Kelley, S. Lovett, and R. Meka. Explicit separations between randomized and deterministic Number-on-Forehead communication. In *Proceedings of the 56th ACM Symposium on Theory of Computing*, pages 1299–1310, 2024. Also available as arXiv:2308.12451.

[KM23] Z. Kelley and R. Meka. Strong bounds for 3-progressions. In *Proceedings of the 64th IEEE Symposium on Foundations of Computer Science*, pages 933–973, 2023. Also available as arXiv:2302.05537.

[LS21] N. Linial and A. Shraibman. An improved protocol for the Exactly-N problem. In *36th Computational Complexity Conference (CCC 2021)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:8. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

[NW19] Anand Kumar Narayanan and Matthew Weidner. Subquadratic time encodable codes beating the Gilbert–Varshamov bound. *IEEE Transactions on Information Theory*, 65(10):6010–6021, 2019.

[Păt10] Mihai Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 603–610, 2010.

[Ran15] H. Randriambololona. On products and powers of linear codes under componentwise multiplication. In *Arithmetic, Geometry, Cryptography and Coding Theory*, volume 637 of *Contemporary Mathematics*, pages 3–78. American Mathematical Society, 2015.

[Rub18] Aviad Rubinstein. Hardness of approximate nearest neighbor search. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1260–1268. ACM, 2018.

[SAK+01] K. W. Shum, I. Aleshnikov, P. V. Kumar, H. Stichtenoth, and V. Deolalikar. A low-complexity algorithm for the construction of algebraic-geometric codes better than the Gilbert–Varshamov bound. *IEEE Transactions on Information Theory*, 47(6):2225–2241, 2001.

[She14] A. A. Sherstov. Communication lower bounds using directional derivatives. *Journal of the ACM*, 61(6):34:1–34:71, 2014.

[Shp09] Amir Shpilka. Constructions of low-degree and error-correcting ε-biased generators. *Computational Complexity*, 18(4):495–525, 2009. With an appendix by Venkatesan Guruswami.

[SS21] Ronen Shaltiel and Jad Silbak. Explicit list-decodable codes with optimal rate for computationally bounded channels. *Computational Complexity*, 30(1):3, 2021.

[Vio07] E. Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM Journal on Computing*, 36(5):1387–1403, 2007.

[YZ25] Guangxu Yang and Jiapeng Zhang. Deterministic lifting theorems for one-way Number-on-Forehead communication. arXiv:2506.12420, 2025.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il