

An Explicit 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. This extends the explicit three-player separations of Kelley, Lovett, and Meka (STOC 2024) and Kelley and Lyu (FOCS 2025) to every fixed number of players, and as a side product improves the three-player nondeterministic lower bound from $\Omega(n^{1/2})$ to the optimal $\Omega(n)$. Our construction is based on algebraic geometry codes.

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 connections to circuit lower bounds [HG91, BH12], pseudorandomness [BNS92, Vio07], and additive combinatorics [AS20, LS21, KLM24]. The extensive overlap among the players' views makes proving lower bounds for explicit functions challenging, even for a fixed number of 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. Non-disjointness asks whether k subsets of an n -element universe have a common element. For fixed k , such an element can be certified using $O(\log n)$ bits, whereas deciding whether one exists requires $n^{\Omega(1)}$ randomized communication [She16].

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. Obtaining a strong explicit separation

*Penn State University. Email: hjw5492@psu.edu.

[†]University of California, Los Angeles. Email: kobewang1359@ucla.edu.

[‡]Penn State University. Email: pei.wu@psu.edu.

proved difficult even for three players. Kelley, Lovett, and Meka [KLM24] gave an explicit three-player function testing equality among three pairwise inner products, with constant randomized cost and a nondeterministic lower bound of $\Omega(n^{1/3})$. Their lower-bound proof extends techniques from Kelley and Meka’s breakthrough on three-term arithmetic progressions [KM23] to a non-arithmetic setting. Kelley and Lyu [KL25] improved the lower bound for the same construction to $\Omega(n^{1/2})$. These results raise the question of whether the constant-versus-linear separation can be made explicit for every fixed number of players.

We give an explicit construction achieving constant randomized cost and linear nondeterministic complexity for every fixed number of players. For three players, this improves the lower bound from $\Omega(n^{1/2})$ to the optimal $\Omega(n)$; for every fixed $k \geq 3$, it matches the order of the nonconstructive separation.

1.1 Main results

Throughout, $k \geq 2$ is fixed, and implicit constants may depend on k . Write $R_k^{cc}(F)$ and $N_k^{cc}(F)$ for the public-coin randomized and nondeterministic communication costs, measured in bits. Public randomness and local computation are free. Randomized error is at most $1/3$ on every input; nondeterministic acceptance is possible exactly on the 1-inputs.

Theorem 1.1 (Explicit constant-cost separation). *For every fixed integer $k \geq 2$, there is an explicit family of total functions $F_n : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$ such that, for all sufficiently large n ,*

$$R_k^{cc}(F_n) = O_k(1), \quad N_k^{cc}(F_n) = \Omega_k(n).$$

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 communication classes coRP_k^{cc} , BPP_k^{cc} , and NP_k^{cc} allow cost polynomial in $\log n$; the first additionally requires perfect completeness.

Corollary 1.2. *For every fixed $k \geq 2$,*

$$\text{coRP}_k^{cc} \not\subseteq \text{NP}_k^{cc}, \quad \text{BPP}_k^{cc} \not\subseteq \text{NP}_k^{cc}.$$

1.2 Construction and proof outline

Our starting point is to test a vector identity by sampling a coordinate:

$$\sum_{\ell=1}^b \mathbf{a}_{1,\ell} \odot \cdots \odot \mathbf{a}_{k,\ell} = \sum_{i=1}^k \mathbf{c}_i. \tag{1.1}$$

Here $\odot$ denotes coordinatewise multiplication of vectors in $\mathbb{F}_q^N$, and b is a constant depending only on k . For this to work, every false identity should fail at a constant fraction of coordinates. We would therefore like to encode the inputs in a linear code $\mathcal{C} \subseteq \mathbb{F}_q^N$ whose coordinatewise products remain in another linear code $\mathcal{W} \subseteq \mathbb{F}_q^N$ of large relative distance:

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

Taking the factors from $\mathcal{C}$ and the shares from $\mathcal{W}$ then puts both sides of the identity in $\mathcal{W}$. Their difference is therefore either zero or nonzero on a constant fraction of coordinates.

Reed–Solomon codes already have this multiplicative property [Cou22]: multiplying low-degree polynomials produces a polynomial of correspondingly larger degree. However, their field size must grow with the block length, and the resulting evaluation protocol uses $O_k(\log n)$ communication bits. Algebraic geometry codes provide such pairs over a fixed field [Cou22, Ran15], allowing constant randomized communication. We can take q to be a power of two depending only on k , with $\dim \mathcal{C} = \Theta_k(N)$, $\dim \mathcal{W} = O_k(\dim \mathcal{C})$, and constant relative distance for $\mathcal{W}$.

Forehead i encodes b codewords $\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,b} \in \mathcal{C}$ and an additive share $\mathbf{c}_i \in \mathcal{W}$. The function accepts exactly when (1.1) holds at all N positions. At each coordinate, the left-hand side is a generalized inner product (GIP), and the shares on the right-hand side specify its claimed value.

Encoding codewords by their coefficients in fixed bases gives a total Boolean function with $n = \Theta_k(N)$ bits per forehead; padding extends the construction to every sufficiently large input length.

Random-coordinate verification. The difference between the two sides of (1.1) belongs to $\mathcal{W}$ by the product containment above. Thus a false identity fails at a constant fraction of coordinates. At a public random coordinate, player 2 sends the field symbols hidden from player 1, who checks the identity using $O_k(1)$ communication bits. Every true identity passes; a false one is rejected with constant probability. A constant number of repetitions reduces the error below $1/3$.

Nondeterministic lower bound. The defining identity already contains a generalized inner product at each coordinate. To exploit this structure, we need independent field symbols. An *information set* S for $\mathcal{C}$ provides $\dim \mathcal{C}$ coordinates whose values can be chosen freely and determine the codeword. Under the uniform distribution μ on all factors and shares, these coordinates are independent and uniform across all factors. Fix the shares. Summing (1.1) over S gives a necessary condition for acceptance:

$$\sum_{\ell=1}^b \sum_{r \in S} \prod_{i=1}^k \mathbf{a}_{i,\ell}(r) = \sum_{i=1}^k \sum_{r \in S} \mathbf{c}_i(r).$$

The left-hand side is GIP on $b \dim \mathcal{C}$ independent field symbols per player; the right-hand side is fixed.

For any accepting cylinder intersection $\mathcal{A}$, fixing the shares and passing to information coordinates preserve its cylinder-intersection structure: both changes act separately on each forehead. The resulting set lies in the displayed GIP level set. Applying the classical correlation bound [BNS92, BHK01, BDPW10, YZ25] and averaging over the shares gives

$$\mu(\mathcal{A}) \leq q^{-\Omega_k(b \dim \mathcal{C})}.$$

A nondeterministic protocol of cost d covers the accepting set by at most 2^d accepting cylinder intersections. The entire accepting set has measure $q^{-\dim \mathcal{W}}$: the factors and the first $k-1$ shares uniquely determine the last share. Increasing b strengthens the cylinder bound without changing this measure. Since $\dim \mathcal{W} = O_k(\dim \mathcal{C})$, a sufficiently large constant b forces every cover to contain $q^{\Omega_k(\dim \mathcal{C})}$ cylinder intersections. This gives $\mathsf{N}_k^{cc}(F_n) = \Omega_k(n)$. Our code parameters allow $b = 2^{k+1}k$.

2 Preliminaries

Throughout, $k \geq 2$ is fixed, and all communication models are k -player number-on-forehead models. All logarithms are to base 2, and implicit constants may depend on k . For a positive integer m ,

write $[m] = \{1, \dots, m\}$. Vectors are bold, and $\mathbf{u}(r)$ denotes the r th coordinate of $\mathbf{u}$. We write $\mathbf{1}\{E\}$ for the indicator of an event E and $\mathbf{1}_{\mathcal{A}}$ for the indicator of a set $\mathcal{A}$.

Write $[m] = \{1, \dots, m\}$, and let $\mathbb{F}_Q$ denote a field of prime-power order Q . Unless a distribution is specified, an expectation over a finite set is uniform. Several variables listed under the same expectation are sampled independently. For a tuple $x = (x_1, \dots, x_k)$, write x_{-i} for the tuple with coordinate i omitted. We use $\mathbf{1}_T$ for the indicator of a set T and $\mu(T)$ for its probability under a distribution μ . All logarithms are in base two. We use N for the code block length and n for the bit length of each input block. The notation $O_k(\cdot)$ and $\Omega_k(\cdot)$ allows the implied constants to depend on k .

When $Q = 2^L$, a basis of $\mathbb{F}_Q$ over $\mathbb{F}_2$ identifies field elements with L -bit strings. Every such string represents a field element, so the encoding imposes no promise on the inputs.

2.1 Linear codes and information sets

A *linear code* $\mathcal{C} \subseteq \mathbb{F}_Q^N$ is a subspace of $\mathbb{F}_Q^N$ with *block length* N , *dimension* $m = \dim \mathcal{C}$, and *rate* m/N . A generator matrix $G \in \mathbb{F}_Q^{m \times N}$ has full row rank and row span $\mathcal{C}$. The map $u \mapsto uG$ is a bijection from message vectors $u \in \mathbb{F}_Q^m$ to codewords, so a uniform message vector gives a uniform codeword.

The *hamming weight* of a vector $v \in \mathbb{F}_Q^N$, denoted $\text{wt}(v)$, is the number of its nonzero coordinates. The *minimum distance* $d_{\min}(\mathcal{C})$ of a code $\mathcal{C}$ is the minimum Hamming weight of the difference of two codewords in $\mathcal{C}$. Its *relative distance* is defined as $\delta(\mathcal{C}) := d_{\min}(\mathcal{C})/N$. Linear codes have the following nice property.

Fact 2.1 (Distance of linear codes). *Let $\mathcal{C} \subseteq \mathbb{F}_Q^N$ be a linear code. Then $d_{\min}(\mathcal{C}) = \min_{c \in \mathcal{C}} \text{wt}(c)$.*

Let $\mathcal{C} \subseteq \mathbb{F}_Q^N$ be a linear code of dimension m . An *information set* of $\mathcal{C}$ is a set $S \subseteq [N]$ of size m for which the coordinate projection

$$\pi_S : \mathcal{C} \longrightarrow \mathbb{F}_Q^S$$

is a linear bijection. These coordinates can be chosen freely; the other coordinates are determined by them.

Lemma 2.2 (Information coordinates). *Every m -dimensional linear code has an information set S of size m . On an information set, a uniform codeword has m independent uniform field coordinates. Every other coordinate is a linear function of these coordinates. Given a generator matrix, an information set and these linear functions can be computed by Gaussian elimination.*

Proof. Choose m linearly independent columns of a generator matrix G to form an invertible matrix G_S . A message u has information coordinates $v = uG_S$. This map is a bijection, so uniform u gives uniform $v \in \mathbb{F}_Q^m$ with independent entries. The full codeword is $vG_S^{-1}G$, which expresses every coordinate as a linear function of v . When $m = 0$, take the empty information set. $\square$

Definition 2.3 (Coordinate-wise product). For vectors $u, v \in \mathbb{F}_Q^N$, write $u \odot v = (u_1v_1, \dots, u_Nv_N)$ for their coordinate-wise product. For a linear code $\mathcal{C} \subseteq \mathbb{F}_Q^N$, let $\mathcal{C}^{\odot k}$ denote the linear span of all p -fold products, i.e.,

$$\mathcal{C}^{\odot k} := \text{span}(\{u_1 \odot \dots \odot u_k \mid u_1, \dots, u_k \in \mathcal{C}\})$$

2.2 Cylinder intersections and nondeterminism

We recall the connection between nondeterministic protocols and covers. A *cylinder in coordinate i* is a set whose membership is independent of x_i . A *cylinder intersection* is an intersection of one

cylinder in each coordinate. Equivalently, its indicator can be written as

$$\mathbf{1}_T(x_1, \dots, x_k) = \prod_{i=1}^k t_i(x_{-i}),$$

where x_{-i} omits x_i and each t_i takes values in $\{0, 1\}$. Each factor describes a condition that player i can check from its view. For two players, cylinder intersections are precisely combinatorial rectangles.

Call a cylinder intersection *accepting* if it is contained in $F^{-1}(1)$. Let $\text{cov}_1(F)$ be the minimum number of accepting cylinder intersections whose union is $F^{-1}(1)$. A nondeterministic protocol of cost d induces such a cover with at most 2^d members: each accepting transcript imposes consistency conditions on the players' views. Consequently,

$$N_k^{cc}(F) \geq \log \text{cov}_1(F).$$

We need only this direction of the cover characterization. It suffices to choose an input distribution under which each accepting cylinder intersection captures a small fraction of the 1-set.

Lemma 2.4 (Cylinder intersections and covers). *Let F be a Boolean function and μ a probability distribution on its input space. Suppose $\mu(F^{-1}(1)) = \alpha > 0$ and every cylinder intersection contained in $F^{-1}(1)$ has measure at most $\varepsilon > 0$. Then*

$$N_k^{cc}(F) \geq \log_2 \frac{\alpha}{\varepsilon}.$$

Proof. A union of M accepting cylinder intersections has measure at most $M\varepsilon$. Any cover therefore satisfies $\alpha \leq M\varepsilon$, and the cover characterization gives the result. $\square$

Changing the representation of each forehead or fixing some of its entries preserves cylinder intersections.

Lemma 2.5 (Restrictions and changes of coordinates). *Suppose T is a cylinder intersection on $\mathcal{X}_1 \times \dots \times \mathcal{X}_k$. For arbitrary maps $\phi_i : \mathcal{Y}_i \rightarrow \mathcal{X}_i$, the preimage*

$$\{(y_1, \dots, y_k) : (\phi_1(y_1), \dots, \phi_k(y_k)) \in T\}$$

is a cylinder intersection on $\mathcal{Y}_1 \times \dots \times \mathcal{Y}_k$.

Proof. Write $\mathbf{1}_T(x) = \prod_i t_i(x_{-i})$ and substitute $x_i = \phi_i(y_i)$. The i -th factor still does not depend on y_i . $\square$

For example, ϕ_i may encode the information coordinates on forehead i and attach a fixed share. Since each map acts on a single forehead, player i 's condition still does not depend on its hidden input.

3 The coded-GIP function

In this section, we formally define the function of interest. As mentioned in the introduction, the function is defined using a pair of code families. We first state the conditions on the codes needed for our construction. Then we define the function of interest. Lastly, we show there exist explicit constructions of the code families of the desired properties.

Definition 3.1 (Good pair of codes). Let $A \geq 1$ and $\delta \in (0, \delta]$. A pair of linear codes $\mathcal{C}, \mathcal{W} \subseteq \mathbb{F}_Q^N$ with dimension $m = \dim \mathcal{C} \geq 1$ and $M = \dim \mathcal{W}$ is a *good* pair of codes if they satisfy the following:

1. *Product containment*: $\mathcal{C}^{(p)} \subseteq \mathcal{W}$
2. *Dimension bound*: $M \leq A \cdot m$
3. *Distance requirement*: $\text{wt}(w) \geq \delta N$ for every $w \in \mathcal{W} \setminus \{0\}$.

Analogously, a pair of code families $\{(\mathcal{C}_n, \mathcal{W}_n)\}_{n=1}^\infty$ is a pair of *good* code families if $(\mathcal{C}_n, \mathcal{W}_n)$ is a good family of codes for all n .

In a good pair of code families, Q, A, δ may depend on p but are independent of N . As we discussed earlier, product containment ensures that the residual defined below belongs to $\mathcal{W}$. The dimension-ratio condition controls the number of field symbols used for an additive share relative to the number used for a factor. The distance condition makes every nonzero residual detectable at a constant fraction of coordinates, which is essential for proving our upper bound.

Given this, we can now define the *coded-GIP* function.

Definition 3.2 (The coded-GIP function). Let k be the number of players in the number-on-forehead model. Fix some $b \geq 1$ that depends only on k and a good pair of codes $(\mathcal{C}, \mathcal{W})$. Each input block is a tuple

$$\mathbf{x}_i = (\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,b}, \mathbf{c}_i) \in \mathcal{C}^b \times \mathcal{W}, \quad i \in [k].$$

In the number-on-forehead model, player i sees all blocks except x_i . The codewords $\mathbf{a}_{i,l}$ are the factors contributed by block i and $\mathbf{c}_i$ is its additive share. Define the *coded-GIP* function $\text{cGIP}_{k,\mathcal{C},\mathcal{W},b} : (\mathcal{C}^b \times \mathcal{W})^k \rightarrow \{0, 1\}$ by

$$\text{cGIP}_{p,\mathcal{C},\mathcal{W},b}(\mathbf{x}_1, \dots, \mathbf{x}_k) = 1 \iff \sum_{i=1}^k \mathbf{c}_i = \sum_{l=1}^b \mathbf{a}_{1,l} \odot \dots \odot \mathbf{a}_{k,l}. \quad (3.1)$$

Here the equality is an identity in $\mathbb{F}_Q^N$.

Write

$$\mathbf{z}(x) = \sum_{i=1}^k \mathbf{c}_i - \sum_{l=1}^b \mathbf{a}_{1,l} \odot \dots \odot \mathbf{a}_{k,l} \quad (3.2)$$

for the residual. Since $\mathcal{W}$ is linear and contains $\mathcal{C}^{\odot k}$, we have $R(x) \in \mathcal{W}$ for every input x . As a result, we know

$$\text{cGIP}_{k,\mathcal{C},\mathcal{W},b}(x) = 0 \implies |\{r \in [N] : R(x)(r) \neq 0\}| \geq \delta N. \quad (3.3)$$

Our $O(1)$ -randomized protocol relies heavily on this fact.

Total Boolean representation. Fix ordered bases of $\mathcal{C}$ and $\mathcal{W}$ and a basis of $\mathbb{F}_Q$ over $\mathbb{F}_2$. These choices identify

$$\mathcal{C}^b \times \mathcal{W} \cong \mathbb{F}_Q^{bm+M} \cong \{0, 1\}^{n_0}, \quad n_0 = (bm + M)L.$$

We can therefore regard $\text{cGIP}_{k,\mathcal{C},\mathcal{W},b}$ as a total function over $(\{0, 1\}^{n_0})^k$. Every bit string specifies a unique tuple of codewords, so there is no code-membership promise. These specifications act separately on each input block and preserve which block is hidden from each player as a result of the order on the bases. In particular, n_0 counts basis coefficients instead of all N coordinates of every codeword.

3.1 Explicit construction of good pairs of code families

We now show how to explicitly construct good pairs of code families using algebraic-geometry codes. The following lemma is a standard consequence of the second Garcia–Stichtenoth tower [GS95, GS96], the multiplicative property of AG evaluation codes [Cou22], and the generator-matrix construction of Shum et al [SAK⁺01]. We give the parameter substitution to specify exactly which consequence of these results is used.

Lemma 3.3 (Explicit multiplicative code pairs). *Fix $k \geq 2$, let s be a power of two with $s - 1 \geq 32k$, and put $Q = s^2$. There is a sequence of pairs of linear codes*

$$\mathcal{C}_j, \mathcal{W}_j \leq \mathbb{F}_Q^{N_j}, \quad N_j = (s^2 - s)s^j, \quad j \geq 0,$$

such that, for all sufficiently large j , writing $m_j = \dim \mathcal{C}_j$ and $M_j = \dim \mathcal{W}_j$, we have

$$\mathcal{C}_j^{\odot k} \subseteq \mathcal{W}_j, \tag{3.4}$$

$$\frac{3N_j}{16p} \leq m_j \leq \frac{N_j}{4p} + 1, \tag{3.5}$$

$$M_j \leq \frac{N_j}{4} + 1, \tag{3.6}$$

$$d_{\min}(\mathcal{W}_j) \geq \frac{3N_j}{4}. \tag{3.7}$$

For fixed p, s , a deterministic algorithm, given j , constructs bases of both codes in time polynomial in N_j . Information sets and all coordinates of codewords specified by their basis coefficients are computable within the same time bound.

Proof. The second Garcia–Stichtenoth tower supplies curves over $\mathbb{F}_{s^2}$ with $N_j = (s^2 - s)s^j$ explicitly specified rational evaluation points, a rational point $P_{\infty,j}$ outside this set, and genus

$$g_j < \frac{N_j}{s-1} \leq \frac{N_j}{32k}.$$

For these tower parameters, see [MP18, Section IV-D]; the evaluation points and the distinguished pole are described in [NW20, Sections 1.5 and 2]. Let D_j be the sum of the evaluation points, and set

$$d_j = \left\lfloor \frac{N_j}{4p} \right\rfloor, \quad G_j = d_j P_{\infty,j}.$$

Take the one-point AG evaluation codes

$$\mathcal{C}_j = C_L(D_j, G_j), \quad \mathcal{W}_j = C_L(D_j, kG_j),$$

using the same ordered list of evaluation points. Here $C_L(D, G)$ consists of the evaluations of functions in the Riemann–Roch space $L(G)$; in the one-point case, these are functions whose only allowed pole is at the distinguished point, with order bounded by the coefficient of that point in G .

Multiplying p such functions adds their pole bounds, giving $\mathcal{C}_j^{\odot k} \subseteq \mathcal{W}_j$ [Cou22, Theorem 5]. Since $kd_j < N_j$, the evaluation maps are injective. The standard dimension and designed-distance bounds for AG codes give

$$d_j + 1 - g_j \leq m_j \leq d_j + 1, \quad M_j \leq pd_j + 1, \quad d_{\min}(\mathcal{W}_j) \geq N_j - pd_j;$$

see [Cou22, Proposition 3]. Using the weaker bound $g_j \leq N_j/(16p)$, we obtain

$$m_j \geq \frac{N_j}{4p} - \frac{N_j}{16p} = \frac{3N_j}{16p}.$$

The remaining estimates follow from $d_j \leq N_j/(4p)$.

The algorithms of Shum et al. [SAK⁺01] construct the required generator matrices in polynomial time; see also [MP18, Section IV-D] and [NW20, Section 1.5]. Apply the construction to both pole bounds while keeping the evaluation-point order fixed. Gaussian elimination supplies information sets, and matrix–vector multiplication recovers codeword coordinates from basis coefficients. The field is fixed, so polynomially many field operations take polynomial time in bits. $\square$

Consequences for the construction. For all sufficiently large j , the displayed estimates give

$$M_j \leq 2p m_j, \quad d_{\min}(\mathcal{W}_j) \geq \frac{3}{4}N_j.$$

Thus Definition 3.2 holds with $A = 2k$ and $\delta = 3/4$. Moreover, for every fixed $b \geq 1$, the input block length satisfies

$$n_j = (bm_j + M_j) \log_2 Q = \Theta_{k,s,b}(N_j).$$

The functions $\text{cGIP}_{k,\mathcal{C}_j,\mathcal{W}_j,b}$ are therefore uniformly computable in time polynomial in their input length. Finally, $N_{j+1} = sN_j$ supplies the bounded growth of block lengths used when padding the construction to every sufficiently large input length.

4 Constant-Cost Randomized Protocol

In this section, we show that there exists a constant-cost randomized protocol for the coded-GIP function. The protocol is simple: we check the defining identity at one public random coordinate. The distance bound in Lemma 3.3 makes this check sufficient.

Theorem 4.1 (Constant-cost randomized protocol for the coded-GIP function). *The function $\text{cGIP}_{k,\mathcal{C},\mathcal{W},b}$ admits a public-coin randomized protocol that accepts every 1-input with probability 1, accepts every 0-input with probability at most $1/3$, and has $O(1)$ communication cost.*

Proof. Using $\lceil \log N \rceil + 4$ public random bits, sample U uniformly from $\{0, \dots, 2^{\lceil \log N \rceil + 4} - 1\}$ and set $r = 1 + (U \bmod N)$. Player 2 sees $\mathbf{x}_1$ and writes

$$\mathbf{a}_{1,1}(r), \dots, \mathbf{a}_{1,b}(r), \mathbf{c}_1(r)$$

on the blackboard. Player 1 sees $\mathbf{x}_2, \dots, \mathbf{x}_p$ and uses these values to compute

$$\mathbf{z}(\mathbf{x})(r) = \mathbf{c}_1(r) + \sum_{i=2}^p \mathbf{c}_i(r) - \sum_{\ell=1}^b \mathbf{a}_{1,\ell}(r) \prod_{i=2}^p \mathbf{a}_{i,\ell}(r).$$

Player 1 outputs $\mathbf{1}\{\mathbf{z}(\mathbf{x})(r) = 0\}$.

Correctness. If $\text{cGIP}_{k,\mathcal{C},\mathcal{W},b}(\mathbf{x}) = 1$, then $\mathbf{z}(\mathbf{x}) = \mathbf{0}$. So the protocol always accepts. Otherwise, $\mathbf{z}(\mathbf{x})$ is a nonzero word of $\mathcal{W}$ by the product containment property of Definition 3.1; and (3.7) implies that it has at most $N/4$ zero coordinates. Reduction modulo N samples each coordinate with probability at most $1/N + 2^{-\lceil \log N \rceil - 4} \leq 17/(16N)$. Thus

$$\Pr[\text{the protocol accepts } \mathbf{x}] \leq \frac{N}{4} \cdot \frac{17}{16N} < \frac{1}{3}.$$

Cost. Player 2 sends $b + 1$ field elements, each encoded by $\log q$ bits, and player 1 sends one bit. The total cost is $(b + 1) \log q + 1 = O_p(1)$. $\square$

5 Nondeterministic Communication Lower Bound

In this section, we prove lower bounds on the non-deterministic communication complexity of the function $\text{cGIP}_{k,\mathcal{C},\mathcal{W},b}$. Specifically, we prove the following theorem.

Theorem 5.1 (Non-deterministic lower bound). *Let $p \geq 2, Q = 2^L$ and $b \geq 1$. Suppose $\mathcal{C}, \mathcal{W} \subseteq \mathbb{F}_Q^N$ are linear codes of dimensions m, M respectively such that $\mathcal{C}^{\odot k} \subseteq W$. Then for the function $F = \text{cGIP}_{k,\mathcal{C},\mathcal{W},b}$, we have*

$$N_k^{\text{cc}}(F) \geq \frac{bm}{2^{p-1}} \log \frac{1}{\eta_{p,Q}} - M \cdot L, \quad \eta_{p,Q} := 1 - (1 - 1/Q)^{p-1}.$$

Recall *cylinder intersection* is a set T whose indicator factors as $\mathbf{1}_T(x) = \prod_{i=1}^p t_i(x_{-i})$. An *accepting cylinder intersection* is a cylinder intersection contained in $F^{-1}(1)$. Then a non-deterministic protocol of cost d yields a cover of $F^{-1}(1)$ by at most 2^d accepting cylinder intersections.

Hence, to prove Theorem 5.1, it suffices to show that *every accepting cylinder intersection has a small measure relative to the total accepting set*. It is easy to see that under independent uniform choices of all factors and shares, the accepting set has measure Q^{-K} . It remains to bound the measure of any accepting cylinder intersection.

Towards this end, notice that a m -dimensional code has m information coordinates that are jointly uniform (Lemma 2.2). After the shares are fixed, an accepting cylinder forces the sum of the bm corresponding products to equal one fixed field element. A classical generalized-inner-product bound makes every such cylinder small. Comparing its measure with that of the entire accepting set proves Theorem 5.1. We now make the argument precise.

Throughout the section, for simplicity of notation, we write $F := \text{cGIP}_{k,\mathcal{C},\mathcal{W},b}$. We first characterize the measure of the accepting set of F .

Lemma 5.2 (Measure of the accepting set). *For the function F and the uniform measure μ over the input space $(\mathcal{C}^b \times W)^k$, we have*

$$\mu(F^{-1}(1)) = Q^{-M}.$$

Proof. Fix all factors $a_{i,l}$ and the first $p-1$ shared. Acceptance is equivalent to

$$\mathbf{c}_p = \sum_{l=1}^b \mathbf{a}_{1,l} * \dots * \mathbf{a}_{k,l} - \sum_{i=1}^{k-1} \mathbf{c}_i.$$

Since $\mathcal{C}^{\odot k} \subseteq W$, we know the right-hand side belongs to W . Therefore, exactly one of the Q^K equally likely values of c_k therefore accepts. $\square$

Next, we state the main technical lemma we need for the proof. For $k \geq 2$ and $s \geq 1$, define the *generalized inner product (GIP)* function as:

$$\text{GIP}_{k,s,Q}(y_1, \dots, y_k) := \sum_{r=1}^s \prod_{i=1}^k y_{i,r}, \quad y_i \in \mathbb{F}_Q^s.$$

Lemma 5.3 (Bounding the pre-image of GIP [BHK01, BDPW10, YZ25]¹). *Let $p \geq 2, Q = 2^L$ and $s, L \geq 1$. Let ν be the uniform measure on $(\mathbb{F}_Q^s)^k$. For every cylinder intersection $T \subseteq (\mathbb{F}_Q^s)^k$ and every $t \in \mathbb{F}_Q$,*

$$\left| \nu(T \cap \text{GIP}_{k,s,Q}^{-1}(t)) - \frac{\nu(T)}{Q} \right| \leq \left(1 - \frac{1}{Q}\right) \eta_{k,Q}^{s/2^{k-1}}, \quad \eta_{k,Q} = 1 - (1 - 1/Q)^{k-1}. \quad (5.1)$$

¹This is the strong discrepancy estimate attributed to Babai, Hayes, and Kimmel [BHK01] and stated explicitly in [BDPW10]. Yang and Zhang [YZ25] use a similar character-correlation estimate for generalized inner product.

As an immediate corollary, we obtain the following bound on the measure of monochromatic cylinders for generalized inner product.

Corollary 5.4 (Monochromatic cylinders for generalized inner product). *With the parameters and measure in Lemma 5.3, suppose that T is a monochromatic cylinder intersection such that $\text{GIP}_{k,s,Q}(y) = t$ for every $y \in T$, for some $t \in \mathbb{F}_Q$. Then*

$$\nu(T) \leq \eta_{k,Q}^{s/2^{k-1}}.$$

Proof. Since $T \subseteq \text{GIP}_{k,s,Q}^{-1}(t)$, the left-hand side of (5.1) equals $(1 - 1/Q)\nu(T)$. Dividing both sides by $(1 - 1/Q)$ yields the desired result. $\square$

The next step is to translate this upper bound under measure ν to an upper bound under measure μ . The key bridge between these two measures is the existence of information set as in Lemma 2.2.

Lemma 5.5 (Every accepting cylinder is small). *Every cylinder intersection $T \subseteq (\mathcal{C}^b \times W)^k$ contained in $F^{-1}(1)$ satisfies*

$$\mu(T) \leq \eta_{k,Q}^{bm/2^{k-1}}.$$

Proof. Let T be an arbitrary accepting cylinder intersection. Choose an information set S for $\mathcal{C}$ using Lemma 2.2. Write $\mathbf{1}_T = \prod_{i=1}^k t_i(x_{-i})$ and fix the entire tuple of shares $c = (c_1, \dots, c_k) \in \mathcal{W}^k$. Let

$$T_c := \{(a_{i,l})_{i,l} \in (\mathcal{C}^b)^k : ((a_{i,1}, \dots, a_{i,b}, c_i))_{i=1}^k \in T\} \subseteq (\mathcal{C}^b)^k$$

be the resulting section in T . It is still a cylinder intersection on the k blocks of factors: fixing the shares does not introduce dependence on the missing block in any t_i . Moreover, the factors remain independent and uniform under this conditioning.

For each i , replace the block $(a_{i,1}, \dots, a_{i,b})$ by its information coordinates

$$y_i := (a_{i,l}(r))_{(l,r) \in [b] \times S} \in \mathbb{F}_Q^{bm}.$$

By Lemma 2.2, this is a bijection acting separately on each block. In particular, coordinates outside S depend only on information coordinates from the same block. Hence, the image $U_c \subseteq (\mathbb{F}_Q^{bm})^k$ of T_c under this map remains a cylinder intersection. Moreover, its measure under the uniform product distribution ν is precisely $\nu(U_c) = \mu(T|c)$.

Since every point of T_c is accepting, the residual must vanish at every coordinate in S . Summing these equations gives

$$\sum_{l=1}^b \sum_{r \in S} \prod_{i=1}^k a_{i,l}(r) = \sum_{i=1}^k \sum_{r \in S} c_i(r),$$

a necessary condition for acceptance. The left-hand side is $\text{GIP}_{k,bm,Q}(y_1, \dots, y_p)$ and the right-hand side is fixed given c . Hence U_c is monochromatic for this generalized inner product. By Corollary 5.4, we have

$$\mu(T|c) = \nu(U_c) \leq \eta_{k,Q}^{bm/2^{k-1}}.$$

This holds for every tuple of shares c . Averaging over c gives the desired statement. $\square$

We are now ready to prove Theorem 5.1.

Proof. (of Theorem 5.1) Let Π be a nondeterministic protocol for F with cost d . Then its accepting computations give a cover of the input domain by at most 2^d accepting cylinder intersections. By Lemma 5.2 and Lemma 5.5, we know

$$Q^{-M} = \mu(F^{-1}(1)) \leq 2^d \eta_{k,Q}^{bm/2^{k-1}}$$

where the inequality hold via union bound. The desired statement then follows by taking logarithms. $\square$

6 Putting everything together

We now prove Theorem 1.1 by combining the explicit codes from Lemma 3.3 with the communication bounds from Sections 4 and 5. Throughout this section, k denotes the number of players, while $m_j = \dim \mathcal{C}_j$ and $M_j = \dim \mathcal{W}_j$ denote code dimensions, as in Sections 3 and 5.

Fix $k \geq 2$, and set

$$s = 2^{\lceil \log_2(32k+1) \rceil}, \quad Q = s^2, \quad L = \log_2 Q, \quad b = 2^{k+1}k. \quad (6.1)$$

These parameters depend only on k , and $s - 1 \geq 32k$. Let $(\mathcal{C}_j, \mathcal{W}_j)$ be the code pairs supplied by Lemma 3.3, of block lengths $N_j = (s^2 - s)s^j$. Write

$$F^{(j)} = \text{cGIP}_{k, \mathcal{C}_j, \mathcal{W}_j, b}, \quad n_j = (bm_j + M_j)L.$$

Using the basis representation in Section 3, we regard $F^{(j)}$ as a total Boolean function on $(\{0, 1\}^{n_j})^k$.

Lemma 6.1 (Lower bound at the code lengths). *For all sufficiently large j ,*

$$\mathbf{N}_k^{cc}(F^{(j)}) \geq L \left(\frac{N_j}{8} - 1 \right) \geq \frac{LN_j}{16}.$$

Proof. Recall that $\eta_{k,Q} = 1 - (1 - 1/Q)^{k-1}$. Bernoulli's inequality and $k - 1 \leq s$ give

$$0 < \eta_{k,Q} \leq \frac{k-1}{Q} \leq \frac{1}{s} = Q^{-1/2}, \quad \log_2 \frac{1}{\eta_{k,Q}} \geq \frac{L}{2}.$$

Applying Theorem 5.1 and substituting the choice of b , we obtain

$$\begin{aligned} \mathbf{N}_k^{cc}(F^{(j)}) &\geq \frac{bm_j}{2^{k-1}} \log_2 \frac{1}{\eta_{k,Q}} - M_j L \\ &\geq \left(\frac{bm_j}{2^k} - M_j \right) L = (2km_j - M_j)L. \end{aligned}$$

The quantitative dimension bounds in Lemma 3.3 leave a positive linear margin:

$$2km_j - M_j \geq 2k \cdot \frac{3N_j}{16k} - \left(\frac{N_j}{4} + 1 \right) = \frac{N_j}{8} - 1.$$

For sufficiently large j we also have $N_j \geq 16$, which proves the second inequality. $\square$

Proof of Theorem 1.1. Choose j_0 so that Lemma 3.3 and Lemma 6.1 hold for every $j \geq j_0$. The dimension bounds in Lemma 3.3 imply

$$n_j = (bm_j + M_j)L = \Theta_k(N_j).$$

Since W_j has relative distance at least $3/4$, Theorem 4.1 gives

$$R_k^{cc}(F^{(j)}) \leq (b+1)L + 1 = O_k(1). \quad (6.2)$$

The protocol accepts every 1-input with probability one and accepts every 0-input with probability at most $17/64 < 1/3$. Together with Lemma 6.1, this already proves the desired separation at the input lengths n_j .

Padding to every sufficiently large input length. Set $B_k = (b+1)L$. Since $m_j, M_j \leq N_j$, we have $n_j \leq B_k N_j$. For each integer $n \geq B_k N_{j_0}$, choose the largest $j \geq j_0$ such that $B_k N_j \leq n$. The identity $N_{j+1} = sN_j$ gives

$$\frac{n}{sB_k} < N_j \leq \frac{n}{B_k}, \quad n_j \leq n. \quad (6.3)$$

Define $F_n : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$ by

$$F_n(x_1, \dots, x_k) = F^{(j)}(x_1|_{[n_j]}, \dots, x_k|_{[n_j]}), \quad (6.4)$$

where $x_i|_{[n_j]}$ is the prefix consisting of the first n_j bits of x_i . Thus F_n simply ignores the remaining bits on each forehead. Every prefix represents valid codewords, so F_n is total.

The randomized protocol ignores the padding and runs the protocol for $F^{(j)}$. It therefore has the cost in (6.2) and the same one-sided error guarantee. Conversely, fixing every padding bit to zero restricts F_n to $F^{(j)}$. Every nondeterministic protocol for F_n therefore induces one for $F^{(j)}$ of no greater cost. By Lemma 6.1 and (6.3),

$$N_k^{cc}(F_n) \geq N_k^{cc}(F^{(j)}) \geq \frac{LN_j}{16} > \frac{n}{16s(b+1)} = \Omega_k(n).$$

Explicitness. For fixed k , the parameters in (6.1) and the threshold j_0 are fixed. Given n , the selected level j can be found from the formula $N_j = (s^2 - s)s^j$ using $O_k(\log n)$ arithmetic steps. Lemma 3.3 constructs the code bases deterministically in time polynomial in N_j . Using these bases, we recover the codewords specified by the input prefixes and evaluate the residual $\mathbf{z}(x)$ in (3.2) at all N_j coordinates. Because $N_j = O_k(n)$ and k, b, Q are fixed, this evaluates F_n in time polynomial in n . Fixing the field basis and using the bases returned by this deterministic construction makes the family uniform.

For the finitely many smaller input lengths, define F_n to be identically zero. This completes the construction of an explicit total family satisfying all the assertions of Theorem 1.1. $\square$

Acknowledgement

The authors used GPT-5.6 for exploratory analysis, idea testing, editorial polishing. The authors reviewed and finalized all mathematical claims and proofs, and take full responsibility for the correctness, exposition, and attribution in the final 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.
- [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.
- [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.
- [Cou22] A. Couteur. How arithmetic and geometry make error-correcting codes better. arXiv:2110.11282, 2022.
- [GS95] A. Garcia and H. Stichtenoth. A tower of Artin–Schreier extensions of function fields attaining the Drinfeld–Vladut bound. *Inventiones Mathematicae*, 121:211–222, 1995.
- [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.
- [HG91] J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Computational Complexity*, 1:113–129, 1991.
- [KL25] Z. Kelley and X. Lyu. More efficient sifting for grid norms, and applications to multiparty communication complexity. arXiv:2505.01587, 2025.
- [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.
- [MP18] U. Martínez-Peñas. Communication efficient and strongly secure secret sharing schemes based on algebraic geometry codes. arXiv:1610.06082v2, 2018.

- [NW20] A. Narayanan and M. Weidner. Subquadratic time encodable codes beating the Gilbert–Varshamov bound. *IEEE Transactions on Information Theory*, 66(2):825–849, 2020. Also available as arXiv:1712.10052.
- [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.
- [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.
- [She16] A. A. Sherstov. The multiparty communication complexity of set disjointness. *SIAM Journal on Computing*, 45(4):1450–1489, 2016.
- [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.