Electronic Colloquium on Computational Complexity, Report No. 49 (2026)

ECCC

No Constant-Cost Protocol for Point–Line Incidence

Mika Göös
EPFL

Nathaniel Harms
UBC

Florian K. Richter
EPFL

Anastasia Sofronova
EPFL

April 7, 2026

Abstract

Alice and Bob are given n-bit integer pairs (x, y) and (a, b), respectively, and they must decide if $y = ax + b$. We prove that the randomised communication complexity of this *Point–Line Incidence* problem is $\Theta(\log n)$. This confirms a conjecture of Cheung, Hatami, Hosseini, and Shirley (CCC 2023) that the complexity is super-constant, and gives the first example of a communication problem with constant support-rank but super-constant randomised complexity.

Contents

ISSN 1433-8092

1 Introduction

Given n-bit integers x, y, a, b, how hard is it to check whether

$$y = ax + b?$$

In communication complexity, we suppose Alice has (x, y) and Bob has (a, b). How many bits of communication are required for them to check this equation? This is the simplest arithmetic (in)equality whose randomised communication complexity is not yet known. For comparison, the randomised complexities of $y = ax$, $y \geqslant b$, and $y \geqslant ax + b$ are fully understood; see Table 1.

The problem of deciding $y = ax + b$ is known as Point–Line Incidence (PL). To our annoyance, it is *a priori* unclear if this can be solved by even a *constant-cost* randomised protocol, that is, with cost independent of n. Indeed, there is no known structural criterion that rules out a constant-cost protocol for PL, despite it being such a natural problem (see Section 1.3 for the ongoing quest to find such structural criteria). Cheung, Hatami, Hosseini, and Shirley [CHHS23] conjectured that it does not have a constant-cost protocol. We confirm their conjecture:

Theorem 1. *The public-coin randomised communication complexity of* PL *is* $\Omega(\log n)$.

This is tight: there is an $O(\log n)$-bit randomised protocol as observed by [CLV19]. The players choose a random prime $p \leqslant O(n)$, Alice sends x and y modulo p, and Bob checks whether $y = ax + b \pmod p$. This protocol has one-sided error, because if $y = ax + b$ then this remains true modulo p.

The same idea gives a protocol for the more general problem of checking $\sum_{i=1}^{k} x_i y_i = 0$ for n-bit integers x_i, y_i. This is Integer Inner Product [CLV19, CHHS23, CHH+25]. As a corollary, we get tight bounds on this problem as well, stated in Section 1.2. These results witness the first separation between randomised communication and *support-rank*, which we discuss in Section 1.1.

Techniques. Any lower bound argument for PL needs to rule out using further number-theoretic tricks to speed up the protocol. The lower bound boils down to the following (informal) statement. Consider the $N \times N$ grid, $N = 2^n$. Let $(\boldsymbol{x}, \boldsymbol{y}) \sim [N]^2$ be a uniformly random point and $\boldsymbol{\ell}$ a line with random slope $\boldsymbol{p}$ through $(\boldsymbol{x}, \boldsymbol{y})$ ($\boldsymbol{p}$ will be a random small prime). Let $\boldsymbol{\ell}^{\downarrow d}$ be the same line but shifted down by a small offset d (which will be the product of all small numbers). We will show that any sufficiently dense set $A \subseteq [N]^2$ cannot distinguish $\boldsymbol{\ell}$ from $\boldsymbol{\ell}^{\downarrow d}$ in the following sense:

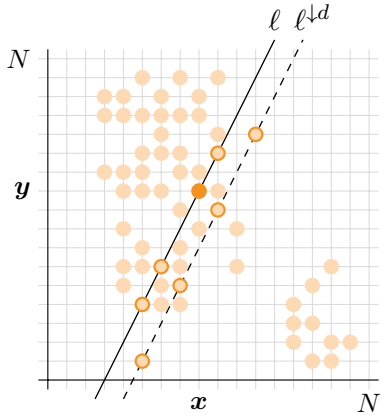

Line Lemma (Lemma 8, informal):
With high probability, $|A \cap \boldsymbol{\ell}| \approx |A \cap \boldsymbol{\ell}^{\downarrow d}|$.

Problem	Task	R	rank_0	$\mathrm{rank}_\pm$	Reference
Equality	$y = b$	$\Theta(1)$	$2^\dagger$	3	
Greater-Than	$y \geqslant b$	$\Theta(\log n)$	$\exp(\Theta(n))$	2	[BW15, Vio15, SY23]
Halfplane Membership	$y \geqslant ax + b$	$\Theta(n)$	$\exp(\Theta(n))$	3	[ACHS24]
Point–Line Incidence	$y = ax + b$	$\Theta(\log n)$	$3^\dagger$	$\Theta(1)$	This paper (Thm. 1)

Table 1: Arithmetic problems together with their randomised communication complexity (R), support-rank (rank_0), and sign-rank ($\mathrm{rank}_\pm$). Here Alice and Bob get n-bit integers (x, y) and (a, b), respectively. When indicated with †, the support-rank is given for the *negated* problem.

Our proof of this lemma relies on a new type of decomposition lemma (Lemma 9) that expresses any bounded function $f\colon \mathbb{Z}_N^2 \to [0,1]$ as a sum of a structured component (which is locally periodic) and a pseudorandom component (which is unbiased over lines). We discuss our proof techniques in more detail in Section 2. For now, we spend the rest of this introduction motivating the study of Point–Line Incidence and explaining the implications of Theorem 1.

1.1 Consequences for Rank Measures

Since $y = ax + b$ is a simple polynomial equation, it has low algebraic complexity as formalised by *support-* and *sign-rank*. These are defined for a two-party function $f\colon \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ as follows.

- *Support-rank* $\mathrm{rank}_0(f)$ is the smallest dimension $r \in \mathbb{N}$ in which f can be represented by linear equalities; that is, such that there exist embeddings $\phi\colon \mathcal{X} \to \mathbb{R}^r$ and $\psi\colon \mathcal{Y} \to \mathbb{R}^r$ with

$$\forall x, y\colon \qquad f(x,y) = 0 \iff \langle \phi(x), \psi(y) \rangle = 0.$$

  In other words, the support-rank of f, viewed as a $|\mathcal{X}|$-by-$|\mathcal{Y}|$ boolean matrix, is the least rank of a matrix $M \in \mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$ that has the same support as f. Previously, support-rank has been called *nondeterministic rank* in quantum communication complexity [dW03], *equality rank* in circuit complexity [HP10], and *minimum rank* in graph theory [FH07].

- *Sign-rank* $\mathrm{rank}_\pm(f)$ is the smallest dimension $r \in \mathbb{N}$ in which f can be represented by linear inequalities; that is, such that there exist embeddings $\phi\colon \mathcal{X} \to \mathbb{R}^r$ and $\psi\colon \mathcal{Y} \to \mathbb{R}^r$ with

$$\forall x, y\colon \qquad \langle \phi(x), \psi(y) \rangle \cdot (-1)^{f(x,y)} > 0.$$

  Sign-rank is known to be equivalent to *unbounded-error randomised communication*, where the two parties have private randomness and must succeed with probability $> 1/2$ [PS86]. Sign-rank is much more powerful than support-rank: Every problem of support-rank r has sign-rank $O(r^2)$ (see, e.g., [GHIS25]), while there are problems over $n := \log\max(|\mathcal{X}|, |\mathcal{Y}|)$ input bits with sign-rank 2 and support-rank 2^n (e.g., Greater-Than).

For Point–Line Incidence, the embeddings $\phi(x,y) = (x, y, 1)$ and $\psi(a,b) = (a, -1, b)$ show that the support-rank of the negation $\neg\mathrm{PL}$ is at most 3. Consequently, Theorem 1 implies the following separation of support-rank and public-coin randomised communication complexity $\mathrm{R}(f)$; previously, it was not known whether $\mathrm{rank}_0(f) = O(1)$ implies $\mathrm{R}(f) = O(1)$.

Corollary 2. *There is an f with $\mathrm{rank}_0(f) \leqslant 3$ and $\mathrm{R}(f) \geqslant \omega(1)$.*

The analogous separation for sign-rank has been long known: Greater-Than has sign-rank 2 but randomised complexity $\Omega(\log n)$ [BW15, Vio15, SY23]. For support-rank, the separation in Corollary 2 is qualitatively the best possible, in the sense that all problems of support-rank 2 reduce to Equality and therefore have randomised complexity $O(1)$. On the other hand, proving a more dramatic quantitative separation for an n-bit problem remains open:

Open Problem 3. Is there an f with $\mathrm{rank}_0(f) \leqslant O(1)$ and $\mathrm{R}(f) \geqslant n^{\Omega(1)}$?

The analogous separation for sign-rank was recently obtained by [HHL20, ACHS24]. They exhibit problems with sign-rank 3 that have randomised complexity $\Theta(n)$. In fact, this holds for the Halfplane Membership problem of deciding whether $y \geqslant ax + b$.

1.2 Integer Inner Product, and Motivation from Oracle Protocols

In the Integer Inner Product (IIP_n^k) problem Alice and Bob receive n-bit integers $x_1, \ldots, x_k$ and $y_1, \ldots, y_k$, respectively, and they want to check whether $\sum_{i=1}^k x_i y_i = 0$. As already mentioned, this problem has randomised complexity $O(k \log n)$ [CLV19]. We get a matching lower bound:

Corollary 4. *The public-coin randomised communication complexity of* IIP_n^k *is* $\Omega(k \log n)$ *for every* $3 \leqslant k \leqslant n^\varepsilon$ *where* $\varepsilon > 0$ *is some constant.*

The condition $k \leqslant n^\varepsilon$ is most likely an artefact of our technique. Nevertheless, it is only a mild restriction since for large k there is also an $\Omega(k)$ lower bound by a reduction from Disjointness.

Integer Inner Product was introduced by [CLV19] to show that efficient randomised protocols cannot be simulated by the Equality oracle. Specifically, they showed that IIP_n^6 has no efficient protocol if we are only allowed to use randomness to solve instances of Equality. In their words, "Equality alone does not simulate randomness." In fact, Equality doesn't help in this problem at all: a deterministic protocol with access to an Equality oracle still requires $\Omega(n)$ bits of communication. We write this result as $\mathrm{D}^{\mathrm{EQ}}(\mathrm{IIP}_n^6) \geqslant \Omega(n)$, where $\mathrm{D}^{\mathrm{EQ}}(f)$ is the number of Equality queries required to compute f. This was improved to hold for PL by [CHHS23] and simplified by [CHH+25]. It remains open whether it is possible to push this type of separation further: the papers [CHHS23, GHR25] ask whether a similar lower bound against Equality-oracle protocols holds for some problem with constant randomised complexity:

Open Problem 5. Is there a problem f with $\mathrm{R}(f) \leqslant O(1)$ but $\mathrm{D}^{\mathrm{EQ}}(f) \geqslant \Omega(n)$?

This was the initial context for the conjecture $\mathrm{R}(\mathrm{PL}) \geqslant \omega(1)$ of [CHHS23], because if Point–Line Incidence had a constant-cost protocol, it would have resolved this question. In light of our Theorem 1, the question remains open. Currently, the best lower bound against Equality-oracle protocols for a problem with $\mathrm{R}(f) \leqslant O(1)$ is $\mathrm{D}^{\mathrm{EQ}}(f) \geqslant \Omega(\sqrt{n})$ [GHR25].

1.3 Motivation from Constant-Cost Communication

Our results fit more broadly within a recent line of work on "constant-cost" communication. Traditionally, communication complexity tries to classify n-bit communication problems into *easy problems* that can be solved with $(\log n)^{O(1)}$ bits of communication and *hard problems* that require $n^{\Omega(1)}$ bits. The papers [HHH23, HWZ22] proposed to consider the *constant-cost* complexity dichotomy $O(1)$-vs-$\omega(1)$ instead of the traditional $(\log n)^{O(1)}$-vs-$n^{\Omega(1)}$. This provides a new "sandbox" where we can study old questions from a novel perspective, helping us develop new lower-bound methods, and also to discover new questions interesting in their own right.

Class BPP_0. A central mystery in this new theory asks to characterise all total problems f that have constant randomised complexity, $\mathrm{R}(f) \leqslant O(1)$. This class of problems, denoted BPP_0, captures the most extreme ways in which randomness can benefit protocols. It has been heavily investigated in recent work [HHH22, EHK22, HHP+22, HZ24, FHHH24, FGHH25, GHR25, BHT25]. The class has many equivalent definitions (constant discrepancy, point–halfspace arrangements with constant margin, PAC learnable under pure differential privacy [FX14]), but all of them so far are "semantic," hiding a promise. Since constant-cost protocols seem very restrictive, one may hope for a simple "syntactic" characterisation. Various natural candidates for a characterisation have been disproved: the class admits no complete problem [FHHH24]; nor certain complete hierarchies [FGHH25, GHR25]. Theorem 1 poses a challenge for any such characterisation, which must somehow generalise our highly-tailored lower bound argument (at least qualitatively).

Class UPP_0. The second-most prominent constant-cost class is UPP_0 that contains all total problems with a constant-cost unbounded-error protocol, that is, problems of constant sign-rank [PS86]. As discussed above, Greater-Than shows that $\mathsf{UPP}_0 \not\subseteq \mathsf{BPP}_0$ and our Corollary 2 upgrades this to a separation $\mathsf{SUPP}_0 \not\subseteq \mathsf{BPP}_0$ where $\mathsf{SUPP}_0 \subseteq \mathsf{UPP}_0$ is the class of problems with constant support-rank, defined in [GHIS25]. The converse separation has grown into a nagging problem:

Open Problem 6. Show that $\mathsf{BPP}_0 \not\subseteq \mathsf{UPP}_0$.

This separation has been shown for *partial* functions by Hatami, Hosseini, and Meng [HHM23]. Towards a separation for total functions, one might first ask to show the weaker separation $\mathsf{BPP}_0 \not\subseteq \mathsf{SUPP}_0$. But this already follows by considering the Equality problem, which has support-rank 2^n on n-bit inputs. A better first step is the following, asked in [GHIS25].

Open Problem 7. Show that $\mathsf{BPP}_0 \not\subseteq \mathsf{P}_0^{\mathsf{SUPP}}$.

Here $\mathsf{P}_0^{\mathsf{SUPP}}$ is the class of problems expressible as boolean combinations of constantly many problems of constant support-rank. This dispenses with Equality as a separating example since its complement has constant support-rank. It is known that $\mathsf{P}_0^{\mathsf{SUPP}} \subseteq \mathsf{UPP}_0$ [GHIS25] so this is truly a necessary first step towards Open Problem 6. For more discussion on constant-cost communication complexity, we recommend the excellent survey of Hatami and Hatami [HH24].

2 Proof Overview

We formally view PL as a function where Alice is given a point $(x, y) \in [N]^2$, $N \coloneqq 2^n$, and Bob is given line parameters $(a, b) \in \{-N^2, \ldots, N^2\}^2$ and they need to compute $\mathrm{PL}((x, y), (a, b)) = 1$ iff $y = ax + b$. Our lower bound uses the textbook discrepancy method [RY20, §6] (and thus it will hold for quantum protocols as well). We define a pair of input distributions $(\mathcal{D}_0, \mathcal{D}_1)$, where $\mathcal{D}_i$ is supported over $\mathrm{PL}^{-1}(i)$, such that every rectangle $A \times B$ (where A, B are subsets of Alice's and Bob's inputs, respectively) has small discrepancy,

$$\left|\mathcal{D}_0(A \times B) - \mathcal{D}_1(A \times B)\right| \leqslant n^{-\Omega(1)}. \tag{1}$$

Theorem 1 then follows immediately from this discrepancy bound (see Section 4). The distributions are defined as follows:

– Choose a uniformly random point $(\boldsymbol{x}, \boldsymbol{y}) \in [N]^2$, and a uniformly random prime $\boldsymbol{p} \leqslant W$; we explain how to choose W later.
– Let $\boldsymbol{\ell}_1$ be the line with slope $\boldsymbol{p}$ through $(\boldsymbol{x}, \boldsymbol{y})$.
– Let $\boldsymbol{\ell}_0$ be the same as $\boldsymbol{\ell}_1$ but shifted down by an offset parameter d, which will be the product of all "small" numbers.
– Let $\mathcal{D}_i$ be the distribution of $((\boldsymbol{x}, \boldsymbol{y}), \boldsymbol{\ell}_i) \in \mathrm{PL}^{-1}(i)$.

To see why we choose the offset d as stated, observe that if there is a small number q not dividing d, Alice can send her coordinates modulo q and Bob is able to perfectly distinguish $\mathcal{D}_0$ from $\mathcal{D}_1$ by testing his equation for the line modulo q.

Showing the discrepancy bound in Equation (1) boils down to analysing rectangles $A \times B$ where Alice's set A has large marginal probability, that is,

$$\mathbb{P}\left[(\boldsymbol{x}, \boldsymbol{y}) \in A\right] = |A|/N^2 \geqslant 1/n^{0.01}.$$

Conditioned on the event $(\boldsymbol{x}, \boldsymbol{y}) \in A$, Bob's input line in $\mathcal{D}_i$ becomes $\boldsymbol{\ell}'_i \coloneqq (\boldsymbol{\ell}_i \mid (\boldsymbol{x}, \boldsymbol{y}) \in A)$. For every choice of A, the maximum discrepancy (1) over all choices of B is then characterised by the *total variation distance*[1], $\mathrm{dist}_{\mathrm{TV}}$, between $\boldsymbol{\ell}'_0$ and $\boldsymbol{\ell}'_1$. Hence our goal becomes to show

$$\mathrm{dist}_{\mathrm{TV}}(\boldsymbol{\ell}'_0, \boldsymbol{\ell}'_1) \leqslant n^{-\Omega(1)}. \tag{2}$$

To this end, we note that for all lines ℓ,

$$\frac{\mathbb{P}[\boldsymbol{\ell}'_1 = \ell]}{\mathbb{P}[\boldsymbol{\ell}'_0 = \ell]} = \frac{|A \cap \ell|}{|A^{\uparrow d} \cap \ell|} = \frac{|A \cap \ell|}{|A \cap \ell^{\downarrow d}|}, \tag{3}$$

where for a set $A \subseteq \mathbb{Z}^2$ we denote by $A^{\uparrow d} \coloneqq A + (0, d)$ and $A^{\downarrow d} \coloneqq A - (0, d)$ its translations up/down by d. To show (2), we need to prove that the ratio (3) is very close to 1 for typical lines ℓ. To carry out this plan, we proceed to develop tools to understand intersection sizes as in (3).

2.1 Input Grid Embedding

It will be technically convenient for us to work in the abelian group $\mathbb{Z}_N^2 = \mathbb{Z}_N \times \mathbb{Z}_N$. Thus, we'll actually think of the input domain of PL as corresponding to a smaller grid $[M]^2 \subseteq \mathbb{Z}_N^2$ where $M \leqslant N/2$. This input grid is then naturally embedded in the lower-left corner of $\mathbb{Z}_N^2$ (this is drawn in blue in the illustration below). We define a line with slope $a \in \mathbb{N}$ through the origin that is capped to the square $(-M, M)^2$ (and then reduced modulo N) as

$$\ell_{a,M} \coloneqq \{(x, y) \bmod N : \ y = ax;\ x, y \in \{-M+1, \ldots, M-1\}\} \subseteq \mathbb{Z}_N^2. \tag{4}$$

We consider these lines anchored at an input grid point $(x, y) \in [M]^2$:

$$\ell \coloneqq \ell_{a,M} + (x, y).$$

Lines of this type are *safe*: when restricted to the input grid, $\ell \cap [M]^2$, they precisely correspond to Bob's inputs in PL. Since $M \leqslant N/2$, these lines do not "wrap around" inside $[M]^2$ despite the modulo-N arithmetic. Here is an illustration of a line with slope $a = 2$ anchored at (x, y).

[1] For distributions $\mu_0, \mu_1 \in [0,1]^\Omega$ over a set Ω, we define $\mathrm{dist}_{\mathrm{TV}}(\mu_0, \mu_1) \coloneqq \max_{B \subseteq \Omega} |\mu_0(B) - \mu_1(B)| = \frac{1}{2}\|\mu_0 - \mu_1\|_1$.

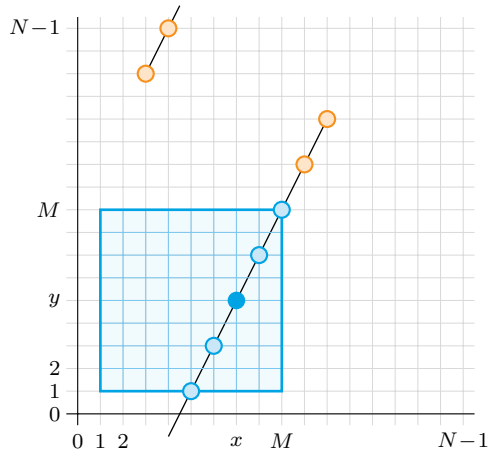

2.2 Notation

We equip functions $f\colon \mathbb{Z}_N^2 \to \mathbb{R}$ with the L^p-norm,

$$\|f\|_p := \left(\mathbb{E}_{z\in\mathbb{Z}_N^2}\left[|f(z)|^p\right]\right)^{1/p}.$$

For $f, g\colon \mathbb{Z}_N^2 \to \mathbb{R}$ we define their *cross-correlation* $f \star g\colon \mathbb{Z}_N^2 \to \mathbb{R}$ by

$$(f\star g)(z) := \mathbb{E}_{z'\in\mathbb{Z}_N^2}\left[f(z')g(z+z')\right].$$

This is related to convolution "$*$" by $f \star g := f' * g$ where $f'(z) := f(-z)$. For a set $A \subseteq \mathbb{Z}_N^2$, we write $\mathbb{1}_A\colon \mathbb{Z}_N^2 \to \{0,1\}$ for its indicator function. By a slight abuse of notation, we identify A with its *density function* $A\colon \mathbb{Z}_N^2 \to [0, N^2]$ defined by

$$A(z) := \tfrac{N^2}{|A|}\mathbb{1}_A(z).$$

Under this convention, A has L^1-norm

$$\|A\|_1 = \tfrac{N^2}{|A|}\mathbb{E}_{z\in\mathbb{Z}_N^2}[\mathbb{1}_A(z)] = 1,$$

and the expression $(A\star f)(z)$ computes the average of f over the translate $A + z$:

$$(A\star f)(z) = \mathbb{E}_{z'\in A}\left[f(z+z')\right].$$

Young's inequality says that such a smoothing operation can only decrease the L^2-norm:

$$\|A\star f\|_2 \leqslant \|A\|_1 \cdot \|f\|_2 \leqslant \|f\|_2. \tag{Young}$$

2.3 Technical Lemmas

To understand intersection sizes $|A\cap\ell|$ as they appear in Equation (3), we may now reformulate these quantities more abstractly using the notation introduced above. The density of a set A relative to a line $\ell := \ell_{a,M} + z$ is expressed as

$$(\ell_{a,M}\star\mathbb{1}_A)(z) = \mathbb{E}_{z'\in\ell_{a,M}}[\mathbb{1}_A(z+z')] = \tfrac{1}{|\ell|}\left|A\cap\left(\ell_{a,M}+z\right)\right|.$$

The following lemma states that these densities typically change very little when we shift a line down by d. Here we use $\mathcal{P}\subseteq\mathbb{N}$ to denote the set of prime numbers and we set $\mathcal{P}_W := \mathcal{P}\cap[W]$.

Lemma 8 (Line Lemma)**.** *For every $n \in \mathbb{N}$ there exist $2 \leqslant d, W \leqslant \exp(n^{0.4})$ such that for every $N, M \geqslant 2^n$ and function $f\colon \mathbb{Z}_N^2 \to [0,1]$, we have*

$$\mathbb{E}_{p\in\mathcal{P}_W}\left[\left\|\ell_{p,M} \star f - \ell_{p,M}^{\downarrow d} \star f\right\|_2\right] \leqslant O(1/n^{0.1}). \tag{5}$$

Our claimed discrepancy bound (1) follows from this lemma; see Section 4 for the proof.

Our proof of Line Lemma relies on a new decomposition lemma, which is our main technical contribution. It asserts that any function $f\colon \mathbb{Z}_N^2 \to [0,1]$ can be split into a *structured* component that has a local period and a *pseudorandom* component that is unbiased in typical lines. Decompositions of a similar nature arise frequently in number theory, Fourier analysis, and ergodic theory; we refer the reader to [Tao07] for a survey of this perspective.

Lemma 9 (Decomposition Lemma)**.** *For every $\varepsilon > 0$ there exist $2 \leqslant d, W \leqslant \exp(1/\varepsilon^4)$ such that for every $N, M \geqslant \exp(1/\varepsilon^5)$, every function $f\colon \mathbb{Z}_N^2 \to [0,1]$ can be written as $f = f_{\mathrm{str}} + f_{\mathrm{psd}}$ where*

(D1) *$f_{\mathrm{str}}\colon \mathbb{Z}_N^2 \to \mathbb{R}$ satisfies $\|f_{\mathrm{str}} - f_{\mathrm{str}}^{\uparrow d}\|_2 \leqslant O(\varepsilon)$.*
(D2) *$f_{\mathrm{psd}}\colon \mathbb{Z}_N^2 \to \mathbb{R}$ satisfies $\mathbb{E}_{p\in\mathcal{P}_W}\left[\|\ell_{p,M} \star f_{\mathrm{psd}}\|_2\right] \leqslant O(\varepsilon)$.*

This lemma readily implies Line Lemma; see Section 2.4 below for the short proof. The proof of Decomposition Lemma, on the other hand, uses Fourier analysis and some number theory; see Section 3 for an overview and the proof.

2.4 Proof of Line Lemma

Apply Decomposition Lemma with $\varepsilon := 1/n^{0.1}$ to obtain the decomposition $f = f_{\mathrm{str}} + f_{\mathrm{psd}}$ with associated parameters $d, W \leqslant \exp(n^{0.4})$. Write $\ell_p := \ell_{p,M}$, $\ell^{\downarrow} := \ell^{\downarrow d}$ for short. We need to verify Equation (5). We expand it via the triangle inequality:

$$\mathrm{LHS}(5) \leqslant \mathbb{E}_p\left[\|(\ell_p - \ell_p^{\downarrow}) \star f_{\mathrm{str}}\|_2\right] + \mathbb{E}_p\left[\|(\ell_p - \ell_p^{\downarrow}) \star f_{\mathrm{psd}}\|_2\right]. \tag{6}$$

We'll show that each of these two terms is $O(\varepsilon)$, which will complete the proof. To bound the first term, we use Young's inequality, $\|\ell_p\|_1 = 1$, and the property (D1):

$$\|(\ell_p - \ell_p^{\downarrow}) \star f_{\mathrm{str}}\|_2 = \|\ell_p \star (f_{\mathrm{str}} - f_{\mathrm{str}}^{\uparrow})\|_2 \leqslant \|\ell_p\|_1 \cdot \|f_{\mathrm{str}} - f_{\mathrm{str}}^{\uparrow}\|_2 \leqslant O(\varepsilon).$$

To bound the second term, consider it for a fixed $p \in \mathcal{P}_W$:

$$\|(\ell_p - \ell_p^{\downarrow}) \star f_{\mathrm{psd}}\|_2 \leqslant \|\ell_p \star f_{\mathrm{psd}}\|_2 + \|\ell_p^{\downarrow} \star f_{\mathrm{psd}}\|_2 = 2\|\ell_p \star f_{\mathrm{psd}}\|_2,$$

where we used that $\|\ell_p \star f_{\mathrm{psd}}\|_2 = \|\ell_p^{\downarrow} \star f_{\mathrm{psd}}\|_2$. Taking expectation over $p \in \mathcal{P}_W$ and applying (D2) shows that the second term is $O(\varepsilon)$, as desired.

3 Proof of Decomposition Lemma

3.1 Fourier-Analytic Preliminaries

Let $e(t)$ be shorthand for $e^{2\pi i t}$ where $t \in \mathbb{R}$. The *Fourier transform* of a function $f\colon \mathbb{Z}_N^2 \to \mathbb{C}$ is the function $\hat{f}\colon \mathbb{Z}_N^2 \to \mathbb{C}$ given by

$$\hat{f}(\xi,\eta) := \mathbb{E}_{(x,y)\in\mathbb{Z}_N^2}[f(x,y)\overline{e_{\xi,\eta}(x,y)}] \qquad \text{where} \qquad e_{\xi,\eta}(x,y) := e\left(\tfrac{\xi x + \eta y}{N}\right).$$

The *Fourier inversion formula* tells us that any f can be recovered from its Fourier coefficients:

$$f \;=\; \sum_{(\xi,\eta)\in\mathbb{Z}_N^2} \hat{f}(\xi,\eta) e_{\xi,\eta}. \tag{7}$$

We also have the following *Parseval's identity*:

$$\|f\|_2^2 \;=\; \sum_{(\xi,\eta)\in\mathbb{Z}_N^2} |\hat{f}(\xi,\eta)|^2. \tag{8}$$

For real-valued f, this gives (using $\widehat{f\star g}(\xi,\eta) = \hat{f}(-\xi,-\eta)\hat{g}(\xi,\eta) = \overline{\hat{f}(\xi,\eta)}\hat{g}(\xi,\eta)$),

$$\|f\star g\|_2^2 \;=\; \sum_{(\xi,\eta)\in\mathbb{Z}_N^2} |\hat{f}(\xi,\eta)|^2 \cdot |\hat{g}(\xi,\eta)|^2. \tag{9}$$

3.2 Proof Outline

Choosing the decomposition. Our decomposition $f = f_{\mathrm{str}} + f_{\mathrm{psd}}$ will be obtained by partitioning the Fourier coefficients of f into two parts. The partitioning strategy is inspired by the classic *circle method* in number theory (due to Hardy and Littlewood) and proceeds as follows in our setting. When should a Fourier term $\hat{f}(\xi,\eta)e_{\xi,\eta}$ be included in the structured part? It is when $e_{\xi,\eta}$ is nearly unchanged under a shift by d, that is, when the following difference is small:

$$\Big|e_{\xi,\eta}(x,y-d) - e_{\xi,\eta}(x,y)\Big| = \left|\frac{e_{\xi,\eta}(x,y-d)}{e_{\xi,\eta}(x,y)} - 1\right| = \left|e\left(\frac{\eta d}{N}\right) - 1\right|. \tag{10}$$

Note that when $\eta d/N$ is an integer, this difference is 0. The idea is to include $e_{\xi,\eta}$ in the structured part when $\eta d/N$ is *nearly* an integer so that (10) is *nearly* 0—here we are using the estimate

$$\forall t\in\mathbb{R}\colon \quad |e(t)-1| = \Theta(\|t\|_{\mathbb{T}}) \qquad \text{where} \qquad \|t\|_{\mathbb{T}} \coloneqq \min_{n\in\mathbb{Z}} |t-n|. \tag{11}$$

We will set d to be the product of all "small" numbers, $d \coloneqq Q!$, where Q is a parameter. Then $\eta d/N$ is nearly an integer when the frequency $\eta/N \in [0,1)$ is near a rational number a/q with a small denominator $q \leqslant Q$. Formally, we classify $\eta\in\mathbb{Z}_N$ as a *major arc* ($\mathfrak{M}_T$) if the frequency η/N is near (distance $\leqslant 1/T$) a small-denominator rational; otherwise it is a *minor arc* ($\mathfrak{m}_T$):

$$\mathfrak{M}_T \coloneqq \bigcup_{\substack{1\leqslant q\leqslant Q\\ \gcd(a,q)=1}} \left\{\eta\in\mathbb{Z}_N : \left|\frac{\eta}{N} - \frac{a}{q}\right| \leqslant \frac{1}{T}\right\} \qquad \text{and} \qquad \mathfrak{m}_T \coloneqq \mathbb{Z}_N\backslash\mathfrak{M}_T.$$

Our candidate decomposition is then $f_{\mathrm{str}} \coloneqq f_{\mathfrak{M}_T}$ and $f_{\mathrm{psd}} \coloneqq f_{\mathfrak{m}_T}$, where

$$f_{\mathfrak{M}_T} \coloneqq \sum_{\xi\in\mathbb{Z}_N,\eta\in\mathfrak{M}_T} \hat{f}(\xi,\eta)e_{\xi,\eta} \qquad \text{and} \qquad f_{\mathfrak{m}_T} \coloneqq \sum_{\xi\in\mathbb{Z}_N,\eta\in\mathfrak{m}_T} \hat{f}(\xi,\eta)e_{\xi,\eta}. \tag{12}$$

Verifying (D1) and (D2). We now verify that (12) satisfies Decomposition Lemma when the parameters are set as

$$Q \coloneqq 10\varepsilon^{-3}, \qquad d \coloneqq Q!, \qquad W \coloneqq 10e^{Q}, \qquad T \coloneqq 10\varepsilon^{-1}d.$$

Consider a major arc $\eta \in \mathfrak{M}_T$. By definition of $\mathfrak{M}_T$, we have that $a/q - 1/T \leqslant \eta/N \leqslant a/q + 1/T$ for some $q \leqslant Q$. Multiplying this by $d = Q!$ shows that

$$\left\| \frac{\eta d}{N} \right\|_{\mathbb{T}} \leqslant \frac{d}{T} \leqslant \varepsilon. \tag{13}$$

Write $\delta_{(x,y)} \coloneqq N^2 \mathbb{1}_{(x,y)}$ for the point mass density at $(x,y) \in \mathbb{Z}_N^2$. Note that $\hat{\delta}_{(x,y)}(\xi,\eta) = \overline{e_{\xi,\eta}(x,y)}$. We can now verify (D1):

$$\begin{aligned}
\|f_{\mathrm{str}} - f_{\mathrm{str}}^{\uparrow d}\|_2^2 &= \|(\delta_{(0,0)} - \delta_{(0,d)}) \star f_{\mathrm{str}}\|_2^2 \\
&= \textstyle\sum_{\xi \in \mathbb{Z}_N, \eta \in \mathfrak{M}_T} \left|1 - e\left(\frac{\eta d}{N}\right)\right|^2 \cdot |\hat{f}(\xi,\eta)|^2 && \text{(Using (9))} \\
&\leqslant \textstyle\sum_{\xi \in \mathbb{Z}_N, \eta \in \mathfrak{M}_T} O\Big(\left\| \frac{\eta d}{N} \right\|_{\mathbb{T}} \Big)^2 \cdot |\hat{f}(\xi,\eta)|^2 && \text{(Using (11))} \\
&\leqslant O(\varepsilon^2). && \text{(Using (13), (8), } \|f\|_2 \leqslant 1)
\end{aligned}$$

To verify (D2), we apply the next lemma, where plugging in our parameter values yields the desired $O(\varepsilon)$ bound. The remainder of this section is dedicated to its proof.

Lemma 10 (Minor Arc Bound)**.** *For Q, W, T such that $e^Q < W < \sqrt{T}/2$:*

$$\mathbb{E}_{p \in \mathcal{P}_W}\left[\|\ell_{p,M} \star f_{\mathfrak{m}_T}\|_2^2\right] \leqslant O\left(\frac{\log\log Q}{Q} + \frac{T^2 W^2}{M^2}\right).$$

3.3 Proof of Minor Arc Bound

Writing $\ell_p \coloneqq \ell_{p,M}$ for short, we have

$$\begin{aligned}
\mathbb{E}_{p \in \mathcal{P}_W}\left[\|\ell_p \star f_{\mathfrak{m}_T}\|_2^2\right] &= \sum_{\xi \in \mathbb{Z}_N, \eta \in \mathfrak{m}_T} \mathbb{E}_{p \in \mathcal{P}_W}\left[|\hat{\ell}_p(\xi,\eta)|^2\right] \cdot |\hat{f}(\xi,\eta)|^2 && \text{(Using (9))} \\
&\leqslant \max_{\xi \in \mathbb{Z}_N, \eta \in \mathfrak{m}_T} \mathbb{E}_{p \in \mathcal{P}_W}\left[|\hat{\ell}_p(\xi,\eta)|^2\right]. && \text{(Using (8), } \|f\|_2 \leqslant 1)
\end{aligned}$$

We'll bound this for each fixed $\xi \in \mathbb{Z}_N$ and $\eta \in \mathfrak{m}_T$. Let us start by computing $|\hat{\ell}_p(\xi,\eta)|^2$ for a fixed choice of slope $p \in \mathcal{P}_W$. Define $X_p \coloneqq \{x \in (-M, M) : px \in (-M, M)\}$ as the set of all x-coordinates of points in the line ℓ_p and note that $|X_p| = |\ell_p| \geqslant \Omega(M/W)$. Then

$$\begin{aligned}
|\hat{\ell}_p(\xi,\eta)|^2 &= \left| \frac{1}{|X_p|} \sum_{x \in X_p} e\big(\tfrac{\xi x + \eta p x}{N}\big) \right|^2 \\
&= \left| \frac{1}{|X_p|} e\big(\tfrac{\xi t + \eta p t}{N}\big) \sum_{x=0}^{|X_p|-1} e\big(\tfrac{\xi x + \eta p x}{N}\big) \right|^2 && \text{(where } t \coloneqq \min X_p) \\
&= \frac{1}{|X_p|^2} \left| \frac{1 - e(\alpha_p |X_p|)}{1 - e(\alpha_p)} \right|^2 && \text{(where } \alpha_p \coloneqq \tfrac{\xi + \eta p}{N}) \\
&\leqslant O\big(\tfrac{W^2}{M^2}\big) \frac{1}{|1 - e(\alpha_p)|^2}. && (14)
\end{aligned}$$

Here we have two cases depending on the magnitude $|1 - e(\alpha_p)| = \Theta(\|\alpha_p\|_{\mathbb{T}})$. We define the set of primes for which this magnitude is small ($\ll 1/T$) and large ($\gg 1/T$) as

$$\mathcal{P}_{\xi,\eta,T} \coloneqq \left\{ p \in \mathcal{P} : \left\| \frac{\xi + p\eta}{N} \right\|_{\mathbb{T}} \leqslant \frac{1}{2T} \right\} \qquad \text{and} \qquad \overline{\mathcal{P}}_{\xi,\eta,T} = \mathcal{P} \setminus \mathcal{P}_{\xi,\eta,T}.$$

We now justify the following calculation that will finish the proof:

$$\begin{aligned}
\mathbb{E}_{p\in\mathcal{P}_W}\left[|\hat{\ell}_p(\xi,\eta)|^2\right] &= \mathbb{E}_{p\in\mathcal{P}_W}\left[\mathbb{1}_{\overline{\mathcal{P}}_{\xi,\eta,T}}(p)\cdot|\hat{\ell}_p(\xi,\eta)|^2\right] + \mathbb{E}_{p\in\mathcal{P}_W}\left[\mathbb{1}_{\mathcal{P}_{\xi,\eta,T}}(p)\cdot|\hat{\ell}_p(\xi,\eta)|^2\right] \\
&\leqslant \mathbb{E}_{p\in\overline{\mathcal{P}}_{\xi,\eta,T}}\left[|\hat{\ell}_p(\xi,\eta)|^2\right] + |\mathcal{P}_{\xi,\eta,T}\cap[W]|/|\mathcal{P}_W| \qquad (|\hat{\ell}_p(\xi,\eta)|\leqslant 1) \\
&\leqslant O\left(\frac{T^2W^2}{M^2}\right) + O\left(\frac{\log\log Q}{Q}\right).
\end{aligned}$$

In the last inequality, every $p \in \overline{\mathcal{P}}_{\xi,\eta,T}$ satisfies $|1 - e(\alpha_p)| \geqslant \Omega(\|\alpha_p\|_{\mathbb{T}}) \geqslant \Omega(1/T)$ and plugging this in (14) gives the bound on the first term. On the other hand, the following lemma (proved in the remainder of this section) gives the bound on the second term.

Lemma 11 (Prime Bound)**.** *Let $\xi \in \mathbb{Z}_N$, $\eta \in \mathfrak{m}_T$, and suppose $e^Q < W < \sqrt{T}/2$. Then*

$$\frac{|\mathcal{P}_{\xi,\eta,T}\cap[W]|}{|\mathcal{P}_W|} \leqslant O\left(\frac{\log\log Q}{Q}\right). \tag{15}$$

3.4 Proof of Prime Bound

For this proof, we need the Siegel–Walfisz theorem [MV06, Corollary 11.19] that shows an upper bound on the density of primes in arithmetic progressions.

Theorem 12 (Siegel–Walfisz)**.** *There exists a constant $C > 0$ such that the following holds: for any $c > 0$ and any a, q, W with $\gcd(a, q) = 1$ and $q \leqslant (\log W)^c$ it holds that*

$$|\mathcal{P}_W \cap (q\mathbb{Z} + a)| = O\left(\frac{W}{\phi(q)\log W}\right) + O_c\left(W\exp(-C\log W)\right),$$

where $\phi(\cdot)$ is Euler's totient function.

Below, we will establish the next claim:

Claim 13. *For every $\xi, \eta \in \mathbb{Z}_N$ there exist $q \in \mathbb{N}$, $a, r \in \{0, 1, \ldots, q-1\}$ with $\gcd(a, q) = 1$ s.t.*

$$\mathcal{P}_{\xi,\eta,T}\cap[W] \subseteq q\mathbb{Z} + r \qquad \text{and} \qquad \left|\frac{\eta}{N} - \frac{a}{q}\right| \leqslant \frac{1}{T}.$$

With this claim, we proceed to show the upper bound (15). If $\mathcal{P}_{\xi,\eta,T} = \emptyset$, the bound is trivial. Suppose it is non-empty. We use Claim 13 to find a, q, r associated with ξ, η. Note that we must have $q \geqslant Q$, because otherwise $\eta \in \mathfrak{M}_T$, contradicting our choice of $\eta \in \mathfrak{m}_T$. Suppose first that $q \geqslant (\log W)^2$. Then $|\mathcal{P}_{\xi,\eta,T}\cap[W]| \leqslant W/q$ since $\mathcal{P}_{\xi,\eta,T} \subseteq q\mathbb{Z} + r$. The bound (15) follows from

$$|\mathcal{P}_{\xi,\eta,T}\cap[W]| \leqslant \frac{W}{q} \leqslant \frac{W}{(\log W)^2} \leqslant O\left(\frac{|\mathcal{P}_W|}{\log W}\right) \leqslant O\left(\frac{|\mathcal{P}_W|}{Q}\right).$$

Otherwise, suppose then that $Q \leqslant q \leqslant (\log W)^2$. Here we can apply Theorem 12:

$$|\mathcal{P}_{\xi,\eta,T} \cap [W]| \leqslant O\left(\frac{W}{\phi(q)\log W}\right). \tag{16}$$

It is known (e.g. [MV06, Thm 2.9]) that $\phi(q) = \Omega(q/\log\log q)$. This gives $\phi(q) \geqslant \Omega(Q/\log\log Q)$ since $q \geqslant Q$. Plugging this in (16) yields the desired bound (15), completing the proof.

Proof of Claim 13. If $\mathcal{P}_{\xi,\eta,T} \cap [W]$ is the empty set or contains exactly one element then we can simply take $q = \frac{N}{\gcd(N,\eta)}$ and $a = \frac{\eta}{\gcd(N,\eta)}$ and the first part of the claim follows readily. So for the remainder of this proof, let us assume that $\mathcal{P}_{\xi,\eta,T} \cap [W]$ contains at least two elements.

It follows from the definition of $\mathcal{P}_{\xi,\eta,T}$ and the triangle inequality that for all $p_1, p_2 \in \mathcal{P}_{\xi,\eta,T} \cap [W]$ with $p_1 < p_2$ we have

$$\left\| \frac{p_2\eta}{N} - \frac{p_1\eta}{N} \right\|_{\mathbb{T}} \leqslant \frac{1}{T}.$$

This means there exists some integer $a(p_1, p_2)$ such that

$$\left| \frac{(p_2 - p_1)\eta}{N} - a(p_1, p_2) \right| \leqslant \frac{1}{T},$$

which implies

$$\left| \frac{\eta}{N} - \frac{a(p_1, p_2)}{(p_2 - p_1)} \right| \leqslant \frac{1}{T}.$$

Using the triangle inequality again, we conclude that for all $p_1, p_2, p_3, p_4 \in \mathcal{P}_{\xi,\eta,T} \cap [W]$ with $p_1 < p_2$ and $p_3 < p_4$ we have

$$\left| \frac{a(p_1, p_2)}{(p_2 - p_1)} - \frac{a(p_3, p_4)}{(p_4 - p_3)} \right| \leqslant \frac{2}{T}.$$

Since $T/2 > W^2$, whereas $(p_2 - p_1)$ and $(p_4 - p_3)$ are no larger than W, we must have

$$\frac{a(p_1, p_2)}{(p_2 - p_1)} = \frac{a(p_3, p_4)}{(p_4 - p_3)}.$$

In other words, there exists a fraction a/q with $\gcd(a, q) = 1$ such that for all $p_1, p_2 \in \mathcal{P}_{\xi,\eta,T} \cap [W]$ with $p_1 < p_2$,

$$\frac{a(p_1, p_2)}{(p_2 - p_1)} = \frac{a}{q}.$$

This implies that for all $p_1, p_2 \in \mathcal{P}_{\xi,\eta,T} \cap [W]$ the difference $p_2 - p_1$ is divisible by q and hence for some $r \in \{0, 1, \ldots, q-1\}$ we have $\mathcal{P}_{\xi,\eta,T} \cap [W] \subseteq q\mathbb{Z} + r$, which yields the claim. □

4 Communication Lower Bounds

Our communication lower bounds use the discrepancy method. The *discrepancy* of $f\colon \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ relative to a distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$ is defined by

$$\operatorname{disc}(f, \mathcal{D}) := \max_R |\mathcal{D}_0(R) - \mathcal{D}_1(R)|$$

where the maximum is over all rectangles $R = A \times B$ with $A \subseteq \mathcal{X}$, $B \subseteq \mathcal{Y}$, and

$$\mathcal{D}_b(R) = \mathbb{P}_{\boldsymbol{z}\sim\mathcal{D}}[\boldsymbol{z} \in R \wedge f(\boldsymbol{z}) = b].$$

The *discrepancy bound* for a function f is

$$\mathrm{Disc}(f) \coloneqq \max_{\mathcal{D}} \log(1/\operatorname{disc}(f, \mathcal{D})),$$

where the maximum is over all distributions over $\mathcal{X} \times \mathcal{Y}$. We have the following basic fact.

Fact 14 ([RY20, §6])**.** $\mathrm{R}(f) \geqslant \Omega(\mathrm{Disc}(f))$ *for all* f*.*

4.1 Point–Line Incidence

We use the Line Lemma to prove a discrepancy bound. We restate the lemma here for convenience.

Lemma 8 (Line Lemma)**.** *For every* $n \in \mathbb{N}$ *there exist* $2 \leqslant d, W \leqslant \exp(n^{0.4})$ *such that for every* $N, M \geqslant 2^n$ *and function* $f\colon \mathbb{Z}_N^2 \to [0,1]$*, we have*

$$\mathbb{E}_{p\in\mathcal{P}_W}\left[\left\|\ell_{p,M} \star f - \ell_{p,M}^{\downarrow d} \star f\right\|_2\right] \leqslant O(1/n^{0.1}). \tag{5}$$

Theorem 1 follows from the next discrepancy bound, combined with Fact 14.

Lemma 15. $\mathrm{Disc}(\mathrm{PL}) = \Omega(\log n)$.

Proof. Let $M \coloneqq 2^n$ so that Alice's inputs are $(x, y) \in [M]^2$. Let $N = 4M$. Let $d, W \leqslant \exp(n^{0.4})$ be obtained from the Line Lemma. We define distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ following the instructions of Section 2, and define $\mathcal{D} \coloneqq (\mathcal{D}_0 + \mathcal{D}_1)/2$.

Let $A \times B$ be any rectangle (so $A \subseteq [M]^2$ is a set of points, and B a set of lines). Let $\boldsymbol{\ell}_1$ be a random line obtained by choosing $(\boldsymbol{x}, \boldsymbol{y}) \sim A$ uniformly at random, choosing $\boldsymbol{p} \sim \mathcal{P}_W$ uniformly at random, and taking the line with slope $\boldsymbol{p}$ through $(\boldsymbol{x}, \boldsymbol{y})$. Let $\boldsymbol{\ell}_0$ be independently distributed in a similar way, but taking the line through $(\boldsymbol{x}, \boldsymbol{y} - d)$ instead. We may upper bound the discrepancy on $A \times B$ using the TV distance:

$$\begin{aligned} |\mathcal{D}_0(A \times B) - \mathcal{D}_1(A \times B)| &= \frac{|A|}{M^2} \left|\mathbb{P}\left[\boldsymbol{\ell}_1 \in B\right] - \mathbb{P}\left[\boldsymbol{\ell}_0 \in B\right]\right| \\ &\leqslant \frac{|A|}{M^2} \max_C \left|\mathbb{P}\left[\boldsymbol{\ell}_1 \in C\right] - \mathbb{P}\left[\boldsymbol{\ell}_0 \in C\right]\right| \\ &= \frac{|A|}{M^2} \cdot \mathrm{dist}_{\mathrm{TV}}(\boldsymbol{\ell}_1, \boldsymbol{\ell}_0), \end{aligned} \tag{17}$$

where the maximum is over all sets C of lines. To establish the TV distance bound, define the function $f\colon \mathbb{Z}_N^2 \to [0,1]$ as $f(x, y) = \mathbb{1}_A[x, y]$. Recall from Equation (4) that

$$\ell_{p,M} \coloneqq \{(h, ph) \in [-M+1, M-1]^2 \colon h \in \mathbb{Z}\}.$$

For a fixed line $\ell \subseteq [M]^2$ with slope $p \leqslant W$, we extend it into $\ell^* \subseteq \mathbb{Z}_N^2$ as

$$\ell^* \coloneqq (u, v) + \ell_{p,M},$$

where $(u, v) \in \ell$ is chosen arbitrarily. Observe that $\ell = \ell^* \cap [M]^2$ since $\ell \subseteq (u, v) + \ell_{p,M}$ for all $(u, v) \in \ell$, and $\ell^* \subseteq [-M, 2M]^2$, so there are no "wraparounds" in $\mathbb{Z}_N^2$. For every $(x, y) \in \ell^*$,

$$\ell_{p,M} \star f(x, y) = \mathbb{E}_{(h,ph)\in\ell_{p,M}}[f(x + h, y + ph)] = \frac{1}{|\ell_{p,M}|}|A \cap \ell|,$$

since $\ell \subseteq (x, y) + \ell_{p,M}$ by definition. Similarly,

$$\ell^{\downarrow}_{p,M} \star f(x, y) = \ell_{p,M} \star f(x, y + d) = \frac{1}{|\ell_{p,M}|}|A \cap \ell^{\downarrow}|.$$

Observe that $|\ell^*| = |\ell_{p,M}|$. Therefore, we may write

$$\mathbb{P}[\boldsymbol{\ell}_1 = \ell] = \mathbb{P}[\boldsymbol{p} = p] \cdot \frac{|A \cap \ell|}{|A|} = \mathbb{P}[\boldsymbol{p} = p] \frac{1}{|A|} \cdot \sum_{(x,y)\in\ell^*} \ell_{p,M} \star f(x, y),$$

and, similarly,

$$\mathbb{P}[\boldsymbol{\ell}_0 = \ell] = \mathbb{P}[\boldsymbol{p} = p] \cdot \frac{|A \cap \ell^{\downarrow}|}{|A|} = \mathbb{P}[\boldsymbol{p} = p] \frac{1}{|A|} \cdot \sum_{(x,y)\in\ell^*} \ell^{\downarrow}_{p,M} \star f(x, y).$$

Then we can use the Line Lemma to bound the TV distance by

$$\begin{aligned}
2 \cdot \mathrm{dist}_{\mathrm{TV}}(\boldsymbol{\ell}_1, \boldsymbol{\ell}_0) &= \sum_{p\in\mathcal{P}_W} \sum_{\text{lines } \ell \text{ of slope } p} |\mathbb{P}[\boldsymbol{\ell}_1 = \ell] - \mathbb{P}[\boldsymbol{\ell}_0 = \ell]| \\
&= \frac{1}{|A|} \cdot \mathbb{E}_{\boldsymbol{p}} \left[\sum_{\ell \text{ of slope } \boldsymbol{p}} \left| \sum_{(x,y)\in\ell^*} \ell_{\boldsymbol{p},M} \star f(x, y) - \ell^{\downarrow}_{\boldsymbol{p},M} \star f(x, y) \right| \right]
\end{aligned}$$

For each p, the line segments ℓ^* are disjoint, so:

$$\begin{aligned}
&\leqslant \frac{1}{|A|} \cdot \mathbb{E}_{\boldsymbol{p}} \left[\sum_{(x,y)\in[-M,2M]^2} \left| \ell_{\boldsymbol{p},M} \star f(x, y) - \ell^{\downarrow}_{\boldsymbol{p},M} \star f(x, y) \right| \right] \\
&\leqslant \frac{3M}{|A|} \cdot \mathbb{E}_{\boldsymbol{p}} \left[\left(\sum_{(x,y)\in[-M,2M]^2} \left(\ell_{\boldsymbol{p},M} \star f(x, y) - \ell^{\downarrow}_{\boldsymbol{p},M} \star f(x, y) \right)^2 \right)^{1/2} \right] && \text{(Cauchy–Schwarz)} \\
&\leqslant \frac{3M}{|A|} \cdot N \cdot \mathbb{E}_{\boldsymbol{p}} \left[\|\ell_{\boldsymbol{p},M} \star f - \ell^{\downarrow}_{\boldsymbol{p},M} \star f\|_2 \right] \\
&\leqslant O\left(\frac{MN}{|A| n^{0.1}} \right). && \text{(Line Lemma)}
\end{aligned}$$

Using Equation (17), this gives a bound of

$$|\mathcal{D}_0(A \times B) - \mathcal{D}_1(A \times B)| \leqslant O\left(\frac{N}{M n^{0.1}} \right) = O(1/n^{0.1}). \qquad \square$$

4.2 Integer Inner Product

We now prove Corollary 4. Recall that PL is a special case of $\mathrm{IIP}^3_n\colon \mathcal{Z}^3\times\mathcal{Z}^3 \to \{0,1\}$, where $\mathcal{Z}$ is the set of n-bit integers, in the sense that PL is a submatrix of IIP^3_n. Consider the function $\textsc{And}_k\circ\mathrm{IIP}^3_n$ that first evaluates k copies of IIP^3_n and then outputs their logical-AND, that is,

$$(\textsc{And}_k \circ \mathrm{IIP}^3_n)(x,y) \coloneqq \textsc{And}_k(\mathrm{IIP}^3_n(x^1,y^1),\dots,\mathrm{IIP}^3_n(x^k,y^k)),$$

where $x \coloneqq (x^1,\dots,x^k)$ and $x^i \coloneqq (x^i_1,x^i_2,x^i_3) \in \mathcal{Z}^3$ and similarly for y. We claim that $\textsc{And}_k \circ \mathrm{IIP}^3_n$ reduces to IIP^{3k}_n via a randomised reduction, which will show that

$$\mathrm{R}(\textsc{And}_k \circ \mathrm{IIP}^3_n) \leqslant O(\mathrm{R}(\mathrm{IIP}^{3k}_n)). \tag{18}$$

Indeed, suppose (x,y) are the inputs to $\textsc{And}_k \circ \mathrm{IIP}^3_n$. We let Alice replace her input x by

$$\boldsymbol{z} \odot x \coloneqq (\boldsymbol{z}_1 x^1,\dots,\boldsymbol{z}_k x^k),$$

where Alice chooses $\boldsymbol{z} \in \{-1,1\}^k$ uniformly at random. Then:

- If $\langle x^i,y^i\rangle = 0$ for all $i\in[k]$, then $\langle \boldsymbol{z}\odot x, y\rangle = 0$ with probability 1.
- If $\langle x^i,y^i\rangle \neq 0$ for some $i\in[k]$, then $\langle \boldsymbol{z}\odot x, y\rangle \neq 0$ with probability $\geqslant 1/2$.

These two properties show that any randomised protocol for IIP^{3k}_n can be used to derive a randomised protocol for $\textsc{And}_k \circ \mathrm{IIP}^3_n$, proving Equation (18).

It remains to show that for every $k \leqslant n^\varepsilon$ where $\varepsilon > 0$ is a sufficiently small constant,

$$\mathrm{R}(\textsc{And}_k \circ \mathrm{PL}) \geqslant \Omega(k \log n). \tag{19}$$

To this end, we employ the following And-composition lemma from [GJPW18, Lemma 10]. The lemma there is originally stated with a measure 2WAPP (aka "smooth rectangle bound") in place of Disc, but the latter is a lower bound on the former [JK10].

Lemma 16 ([GJPW18])**.** $\mathrm{Disc}(f) \leqslant O(\mathrm{R}(\textsc{And}_k \circ f)/k + \log \mathrm{R}(\textsc{And}_k \circ f))$ *for all* f*.*

Instantiating this with $f \coloneqq \mathrm{PL}$ gets us

$$\begin{aligned}
\log n &\leqslant O(\mathrm{Disc}(\mathrm{PL})) && \text{(Lemma 15)}\\
&\leqslant O(\mathrm{R}(\textsc{And}_k \circ \mathrm{PL})/k + \log \mathrm{R}(\textsc{And}_k \circ \mathrm{PL})) && \text{(Lemma 16)}\\
&\leqslant O(\mathrm{R}(\textsc{And}_k \circ \mathrm{PL})/k + \log(k \log n))\\
&\leqslant O(\mathrm{R}(\textsc{And}_k \circ \mathrm{PL})/k + \varepsilon \log n). && (k \leqslant n^\varepsilon)
\end{aligned}$$

Choosing $\varepsilon > 0$ small enough and rearranging gives Equation (19), as desired.

Acknowledgements

We thank Hamed Hatami, Kaave Hosseini, Shachar Lovett, and Raghu Meka for discussions. Special thanks to Oliver Göös for serving as an uncritical sounding board during the writing of the paper. M.G. and A.S. are supported by the Swiss State Secretariat for Education, Research, and Innovation (SERI) under contract number MB22.00026. F.K.R. was supported by the Swiss National Science Foundation grant TMSGI2-211214.

References

[ACHS24] Manasseh Ahmed, Tsun-Ming Cheung, Hamed Hatami, and Kusha Sareen. Communication complexity and discrepancy of halfplanes. In *Proceedings of the Symposium on Computational Geometry (SoCG)*, pages 5:1–5:17. Schloss Dagstuhl, 2024. doi:10.4230/LIPIcs.SoCG.2024.5.

[BHT25] Igor Balla, Lianna Hambardzumyan, and István Tomon. Factorization norms and an inverse theorem for MaxCut. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 947–963. IEEE, 2025. doi:10.1109/focs63196.2025.00049.

[BW15] Mark Braverman and Omri Weinstein. A discrepancy lower bound for information complexity. *Algorithmica*, 76(3):846–864, 2015. doi:10.1007/s00453-015-0093-8.

[CHH+25] Tsun-Ming Cheung, Hamed Hatami, Kaave Hosseini, Aleksandar Nikolov, Toniann Pitassi, and Morgan Shirley. A lower bound on the trace norm of boolean matrices and its applications. In *Proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*, volume 325 of *LIPIcs*, pages 37:1–37:15. Schloss Dagstuhl, 2025. doi:10.4230/LIPIcs.ITCS.2025.37.

[CHHS23] Tsun-Ming Cheung, Hamed Hatami, Kaave Hosseini, and Morgan Shirley. Separation of the factorization norm and randomized communication complexity. In *Proceedings of the Conference on Computational Complexity (CCC)*, volume 264 of *LIPIcs*, pages 1:1–1:16. Schloss Dagstuhl, 2023. doi:10.4230/LIPIcs.CCC.2023.1.

[CLV19] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In *Proceedings of the Conference on Computational Complexity (CCC)*, pages 14:1–14:11. Schloss Dagstuhl, 2019. doi:10.4230/LIPIcs.CCC.2019.14.

[dW03] Ronald de Wolf. Nondeterministic quantum query and communication complexities. *SIAM Journal on Computing*, 32(3):681–699, 2003. doi:10.1137/s0097539702407345.

[EHK22] Louis Esperet, Nathaniel Harms, and Andrey Kupavskii. Sketching distances in monotone graph classes. In *International Conference on Randomization and Computation (RANDOM)*, pages 18–1. Schloss Dagstuhl, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.18.

[FGHH25] Yuting Fang, Mika Göös, Nathaniel Harms, and Pooya Hatami. Constant-cost communication does not reduce to k-Hamming distance. In *Proceedings of the Symposium on Theory of Computing (STOC)*, 2025. doi:10.48550/arXiv.2407.20204.

[FH07] Shaun Fallat and Leslie Hogben. The minimum rank of symmetric matrices described by a graph: A survey. *Linear Algebra and its Applications*, 426(2–3):558–582, 2007. doi:10.1016/j.laa.2007.05.036.

[FHHH24] Yuting Fang, Lianna Hambardzumyan, Nathaniel Harms, and Pooya Hatami. No complete problem for constant-cost randomized communication. In *Proceedings of the Symposium on Theory of Computing (STOC)*, 2024. doi:10.48550/arXiv.2404.00812.

[FX14] Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. In *Proceedings of the Conference on Learning Theory (COLT)*, pages 1000–1019. PMLR, 2014. doi:10.48550/arXiv.1402.6278.

[GHIS25] Mika Göös, Nathaniel Harms, Valentin Imbach, and Dmitry Sokolov. Sign-rank of k-Hamming distance is constant. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 2353–2368. IEEE, December 2025. doi:10.1109/focs63196.2025.00123.

[GHR25] Mika Göös, Nathaniel Harms, and Artur Riazanov. Equality is far weaker than constant-cost communication. In *International Conference on Randomization and Computation (RANDOM)*, volume 353 of *LIPIcs*, pages 58:1–58:14. Schloss Dagstuhl, 2025. doi:10.4230/LIPIcs.APPROX/RANDOM.2025.58.

[GJPW18] Mika Göös, T.S. Jayram, Toniann Pitassi, and Thomas Watson. Randomized communication vs. partition number. *ACM Transactions on Computation Theory*, 10(1):4:1–4:20, 2018. doi:10.1145/3170711.

[HH24] Hamed Hatami and Pooya Hatami. Guest column: Structure in communication complexity and constant-cost complexity classes. *ACM SIGACT News*, 55(1):67–93, 2024. doi:10.1145/3654780.3654788.

[HHH22] Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. A counter-example to the probabilistic universal graph conjecture via randomized communication complexity. *Discrete Applied Mathematics*, 322:117–122, 2022. doi:10.1016/j.dam.2022.07.023.

[HHH23] Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. Dimension-free bounds and structural results in communication complexity. *Israel Journal of Mathematics*, 253(2):555–616, 2023. doi:10.1007/s11856-022-2365-8.

[HHL20] Hamed Hatami, Kaave Hosseini, and Shachar Lovett. Sign rank vs discrepancy. In *Proceedings of the Conference on Computational Complexity (CCC)*, pages 18–1. Schloss Dagstuhl, 2020. doi:10.4230/LIPIcs.CCC.2020.18.

[HHM23] Hamed Hatami, Kaave Hosseini, and Xiang Meng. A Borsuk-Ulam lower bound for sign-rank and its applications. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 463–471, 2023. doi:10.1145/3564246.3585210.

[HHP+22] Hamed Hatami, Pooya Hatami, William Pires, Ran Tao, and Rosie Zhao. Lower bound methods for sign-rank and their limitations. In *International Conference on Randomization and Computation (RANDOM)*, pages 22–1. Schloss Dagstuhl, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.22.

[HP10] Kristoffer Arnsfelt Hansen and Vladimir Podolskii. Exact threshold circuits. In *Proceedings of the Conference on Computational Complexity (CCC)*, pages 270–279. IEEE, 2010. doi:10.1109/ccc.2010.33.

[HWZ22] Nathaniel Harms, Sebastian Wild, and Viktor Zamaraev. Randomized communication and implicit graph representations. In *Proceedings of the Symposium on Theory of Computing (STOC)*, 2022. doi:10.1145/3519935.3519978.

[HZ24] Nathaniel Harms and Viktor Zamaraev. Randomized communication and implicit representations for matrices and graphs of small sign-rank. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*, pages 1810–1833. SIAM, 2024. doi:10.1137/1.9781611977912.72.

[JK10] Rahul Jain and Hartmut Klauck. The partition bound for classical communication complexity and query complexity. In *Proceedings of the Conference on Computational Complexity (CCC)*, pages 247–258. IEEE, 2010. doi:10.1109/CCC.2010.31.

[MV06] Hugh Montgomery and Robert Vaughan. *Multiplicative Number Theory I: Classical Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2006. doi:10.1017/CBO9780511618314.

[PS86] Ramamohan Paturi and Janos Simon. Probabilistic communication complexity. *Journal of Computer and System Sciences*, 33(1):106–123, 1986. doi:10.1016/0022-0000(86)90046-2.

[RY20] Anup Rao and Amir Yehudayoff. *Communication Complexity: and Applications*. Cambridge University Press, January 2020. doi:10.1017/9781108671644.

[SY23] Srikanth Srinivasan and Amir Yehudayoff. The discrepancy of greater-than. Technical report, arXiv, 2023. doi:10.48550/ARXIV.2309.08703.

[Tao07] Terence Tao. Structure and randomness in combinatorics. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 3–15. IEEE, 2007. doi:10.1109/focs.2007.17.

[Vio15] Emanuele Viola. The communication complexity of addition. *Combinatorica*, 35(6):703–747, 2015. doi:10.1007/s00493-014-3078-3.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il