THEORY OF COMPUTING, Volume 6 (2010), pp. 201–225
www.theoryofcomputing.org

Separating Deterministic from Randomized Multiparty Communication Complexity

Paul Beame[*] Matei David[†] Toniann Pitassi[‡] Philipp Woelfel[§]

Received: September 5, 2008; published: November 15, 2010.

Abstract: We solve some fundamental problems in the number-on-forehead (NOF) k-player communication model. We show that there exists a function which has at most logarithmic communication complexity for randomized protocols with one-sided false-positives error probability of 1/3, but which has linear communication complexity for deterministic protocols and, in fact, even for the more powerful nondeterministic protocols. The result holds for every $\varepsilon > 0$ and every $k \leq 2^{(1-\varepsilon)n}$ players, where n is the number of bits on each player's forehead. As a consequence, we obtain the NOF communication class separation $\mathsf{coRP} \not\subset \mathsf{NP}$. This in particular implies that $\mathsf{P} \neq \mathsf{RP}$ and $\mathsf{NP} \neq \mathsf{coNP}$. We also show that for every $\varepsilon > 0$ and every $k \leq n^{1-\varepsilon}$, there exists a function which has constant randomized complexity for public coin protocols but at least logarithmic complexity for private coin protocols. No larger gap between private and public coin protocols is possible.

Our lower bounds are existential; no explicit function is known to satisfy nontrivial lower bounds for $k \geq \log n$ players. However, for every $\varepsilon > 0$ and every $k \leq (1-\varepsilon) \cdot \log n$ players, the $\mathsf{NP} \neq \mathsf{coNP}$ separation (and even the $\mathsf{coNP} \not\subset \mathsf{MA}$ separation) was obtained independently by Gavinsky and Sherstov (2010) using an explicit construction. In this work, for $k \leq (1/9) \cdot \log n$ players, we exhibit an explicit function which has communication

[*] Supported by NSF grant CCR-0514870
[†] Supported by NSF grant CCF-0832797
[‡] Supported by NSERC
[§] Supported by NSERC and DFG grant WO 1232/1-1

ACM Classification: F.1.3

AMS Classification: 68Q15

Key words and phrases: communication complexity, number on forehead

DOI: 10.4086/toc.2010.v006a009

complexity $O(1)$ for public coin protocols and $\Omega(\log n)$ for deterministic protocols. This improves the best previously known deterministic lower bound for a function with efficient randomized protocols, which was $\Omega(\log\log n)$, given by Beigel, Gasarch, and Glenn (2006).

It follows from our existential result that any function that is complete for the class of functions with polylogarithmic nondeterministic k-player communication complexity does not have polylogarithmic deterministic complexity. We show that the set intersection function, which is complete in the number-in-hand model, is not complete in the NOF model under cylindrical reductions.

1 Introduction

The question of how much communication is necessary in order to compute a function when its input is distributed between several computationally unbounded players was introduced by Yao [26], and it has since been shown to have many applications in diverse areas of computer science. The case of $k = 2$ players has been studied extensively. In this paper we are interested in the case of $k \geq 3$ players, specifically in the ”number-on-forehead” (NOF) model, which was first considered by Chandra, Furst, and Lipton [12]. In this model, the input is partitioned into k parts, so that player i can see all parts except for the i^{th} part (since it is “written on his forehead”). The standard reference on communication complexity is the text by Kushilevitz and Nisan [18].

The NOF communication model is a fascinating and complex model that is not well understood when $k \geq 3$. The complexity of the situation arises from the fact that every part of the input is seen by multiple players. As the number of players increases, the sharing becomes increasingly generous. During the execution of a protocol, the set of inputs consistent with a particular message sequence is described by a so-called cylinder intersection. The combinatorial structure of cylinder intersections appears difficult to understand.

Lower bounds for multiparty complexity in the NOF model are connected to a major open problem in complexity theory; in a series of papers by Yao [27], Beigel and Tarui [10], and Håstad and Goldman [16], it was established that superlogarithmic communication complexity lower bounds in the NOF model for any explicit function with polylogarithmically many players would imply explicit lower bounds for ACC^0. The best lower bound known for an explicit function was given by Babai, Nisan, and Szegedy [4], and it is of the form $\Omega\left(n/2^k\right)$, so it breaks down when the number of players is greater than logarithmic. Lower bounds in this model have many other important applications as well, including: constructions of pseudorandom generators for space bounded computation, constructions of universal traversal sequences, time-space trade-offs [4], further circuit complexity bounds [16, 23, 22], and proof complexity bounds [8, 6].

The motivation for our work is to pursue a broader understanding of the NOF complexity model. In particular, we would like to answer some of the basic questions that are still open for this model, but have well-known solutions in the 2-player model. For $k \geq 3$, we consider the three usual versions of communication complexity: deterministic, randomized, and nondeterministic complexity. Are there functions separating these three different complexity measures? Nothing was known in this sense prior to the preliminary version of this work [5], even for $k = 3$ players.

Our main result is that for every $\varepsilon > 0$ and every $k \leq 2^{(1-\varepsilon)n}$, there is a function with n bits on each players' forehead that is computable with logarithmic communication by a randomized k-player NOF communication protocol with bounded one-sided false-positives error, but which requires linear amount of communication for deterministic protocols, and in fact, even for nondeterministic (unbounded false-negatives error) protocols. We obtain this result nonconstructively by showing that nondeterministic protocols for a certain class of "graph functions" have a nice normal form and then establishing a lower bound for such functions via a counting argument over protocols in normal form. Communication complexity classes such as $\mathsf{P}_k^{\mathrm{cc}}$ and $\mathsf{NP}_k^{\mathrm{cc}}$ were introduced for $k = 2$ players by Babai, Frankl, and Simon [1]; we generalize these concepts to $k \geq 3$ players, and we infer the following separation: $\mathsf{coRP}_k^{\mathrm{cc}} \not\subseteq \mathsf{NP}_k^{\mathrm{cc}}$. In particular, this implies $\mathsf{P}_k^{\mathrm{cc}} \neq \mathsf{RP}_k^{\mathrm{cc}}$ and $\mathsf{NP}_k^{\mathrm{cc}} \neq \mathsf{coNP}_k^{\mathrm{cc}}$. As a corollary of our lower bound, we also establish an optimal separation between the powers of public- and private-coin randomized NOF protocols, albeit only for $k \leq n^{1-\varepsilon}$ (for every $\varepsilon > 0$).

The lower bound above is nonconstructive, but for k at most logarithmic in the input size, we also give *explicit* families of graph functions with NOF communication complexity $O(1)$ for public-coin randomized protocols and $\Omega(\log n)$ for deterministic protocols. (We believe that the latter is, in fact, super-polylogarithmic.) The best previously known deterministic lower bound for a function with efficient randomized NOF protocols is the $\Omega(\log\log n)$ lower bound given by Beigel, Gasarch, and Glenn [9] for the Exact-T function, which was originally investigated in [12] in the special case of $k = 3$ players. As a corollary, we also obtain the fact that the function families we define have $\Omega(\log\log n)$ private-coin randomized NOF communication complexity but only $O(1)$ public-coin randomized NOF communication complexity.

The problem of separating deterministic from nondeterministic NOF communication complexity is particularly interesting because of its connection to proof complexity. It was shown by Beame, Pitassi, and Segerlind [8] that for $k = 3$, $(\log n)^{\Omega(1)}$ lower bounds on the randomized NOF complexity of set intersection, which has nondeterministic NOF complexity $O(\log n)$, would imply lower bounds for polynomial threshold proof systems, such as the Lovász-Schrijver proof systems, as well as the Chvátal cutting planes proof system. This brings us to our second question: is there a "complete" problem for the class of problems with efficient NOF nondeterministic algorithms under a suitable notion of reduction? Given our separation result, such a function would automatically be hard for deterministic protocols. Following [1], it is not hard to see that set intersection is complete under communication-free reductions for the number-in-hand (NIH) model. (The NIH model is an alternative generalization of the 2-player model in which each player gets his part of the input in his hand, and thus each player sees only his own part.) We prove that under communication-free reductions, set intersection is not complete in the NOF model.

Subsequent to the preliminary version of this paper [5], in a series of works by Lee and Shraibman [19], Chattopadhyay and Ada [13], and Beame and Huynh [7], lower bounds were obtained for the randomized complexity of the set disjointness function for $k \leq O\left((\log n)^{1/3}\right)$ players, implying the NOF communication complexity class separation $\mathsf{NP}_k^{\mathrm{cc}} \not\subseteq \mathsf{BPP}_k^{\mathrm{cc}}$, and therefore $\mathsf{RP}_k^{\mathrm{cc}} \neq \mathsf{NP}_k^{\mathrm{cc}}$. David, Pitassi, and Viola [14] improve this separation using a different function, for every $\varepsilon > 0$ and every $k \leq (1-\varepsilon)\log n$ players. See the excellent survey article by Sherstov [25] for more details on this line of work.

Independently of our work and using different techniques, Gavinsky and Sherstov [15] obtained the

$\mathsf{coNP}_k^{\mathrm{cc}} \not\subseteq \mathsf{MA}_k^{\mathrm{cc}}$ separation, which also implies the $\mathsf{NP}_k^{\mathrm{cc}} \neq \mathsf{coNP}_k^{\mathrm{cc}}$ separation. Their results are based on an explicit construction, and they hold for every $\varepsilon > 0$ and every $k < (1-\varepsilon)\log n$ players.

2 Definitions and preliminaries

Let $k : \mathbb{N} \to \mathbb{N}$ be a non-decreasing function controlling the number of players as a function of the input size n. In general, we write k instead of $k(n)$ when the input size n is clear from the context.

In the NOF multiparty communication complexity model of computation [12] there are k players, numbered 1 through k, that compute a function $f_{k,n} : X_1 \times \cdots \times X_k \to \{0,1\}$, where each X_i is a set of size at most 2^n. For each $1 \leq i \leq k$, an input $x_i \in X_i$ is (metaphorically) placed on the forehead of player i. Thus, player i knows the values of all of the inputs *except for* x_i.

The players exchange bits according to an agreed-upon protocol, by writing them on a public blackboard. A *protocol* specifies solely as a function of the blackboard contents: whether or not the communication is over; if it is over, the output of the protocol; and if it is not over, the next player to speak. A protocol also specifies what each player communicates as a function of the blackboard contents and of the inputs seen by that player. The *cost* of a protocol is the maximum number of bits written on the blackboard.

In a *deterministic protocol*, the blackboard is initially empty. A *public-coin randomized protocol* of cost c is simply a probability distribution over deterministic protocols of cost c, which can be viewed as a protocol in which the players have access to a shared random string that is *not* counted towards the cost. A *private-coin randomized protocol* is a protocol in which each player has access to a private random string. A *nondeterministic protocol* is a randomized private coin protocol with one-sided error (only false negatives) and error probability less than 1.

The *deterministic communication complexity of* $f_{k,n}$, written $D_k(f_{k,n})$, is the minimum cost of a deterministic protocol for $f_{k,n}$ that always outputs the correct answer. For $0 \leq \varepsilon < 1/2$, let $R_{k,\varepsilon}^{\mathrm{pub}}(f_{k,n})$ denote the minimum cost of a public-coin randomized protocol for $f_{k,n}$ which, for every input, makes an error with probability at most ε (over the choice of the deterministic protocols). The *public-coin randomized communication complexity of* $f_{k,n}$ is $R_k^{\mathrm{pub}}(f_{k,n}) = R_{k,1/3}^{\mathrm{pub}}(f_{k,n})$. Let $R_{k,\varepsilon}(f_{k,n})$ denote the minimum cost of a private-coin randomized protocol for $f_{k,n}$ which, for every input, makes an error with probability at most ε (over the choice of the private random strings). The *private-coin randomized communication complexity of* $f_{k,n}$ is $R_k(f_{k,n}) = R_{k,1/3}(f_{k,n})$. For both public-coin and private-coin complexities we add a superscript 1 if we require that the protocol errs only on 1-inputs (i. e., false-negatives), and superscript 0 if we require that the protocol errs only on 0-inputs (i. e., false-positives). For example, $R_{k,\varepsilon}^{0,\mathrm{pub}}(f_{k,n})$ is the minimum cost of a k-player public-coin protocol for $f_{k,n}$ which is always correct on 1-inputs and errs with probability at most ε on every 0-input. The *nondeterministic communication complexity of* $f_{k,n}$, written $N_k(f_{k,n})$, is the minimum cost of a nondeterministic protocol for $f_{k,n}$.

In general, we consider a function family $f = (f_{k(n),n})$ defined for all but finitely many n, and we are interested in how the communication complexity of the functions in the family grows with the input size n. When n is clear from the context, we drop the subscripts and we write f instead of $f_{k,n}$. Thus, for any of the complexity measures C defined above, we write $C_k(f)$ for $C_k(f_{k(n),n})$, which is itself a function of n.

Since the general model laid out above is very powerful, we are also interested in communication restrictions. A player is *oblivious* in a certain protocol if the message he writes on the board is a function of the inputs he sees, but not a function of the messages sent by other players. Since we are interested in the best protocol, we may safely assume that all oblivious players write first, and then non-oblivious players continue to communicate using the information written by the former. A protocol in which all players are oblivious is called *simultaneous*. The simultaneous multiparty model was studied by Babai et al. [2], who proved new lower bounds, as well as surprising upper bounds in this model.

Since any function $f_{k,n}$ can be computed using only $n+1$ bits of communication, following [1], for a family of functions $f = (f_{k,n})$, communication protocols are considered "efficient" or "polynomial" if only polylogarithmically many bits are exchanged. Accordingly, let $\mathsf{P}_k^{\mathrm{cc}}$ denote the class of function families f for which $D_k(f) \le (\log n)^{O(1)}$, let $\mathsf{NP}_k^{\mathrm{cc}}$ denote the class of function families f for which $N_k(f) \le (\log n)^{O(1)}$, and let $\mathsf{RP}_k^{\mathrm{cc}}$ denote the class of function families f for which $R_k^1(f) \le (\log n)^{O(1)}$. The classes $\mathsf{BPP}_k^{\mathrm{cc}}$, $\mathsf{coRP}_k^{\mathrm{cc}}$ and $\mathsf{coNP}_k^{\mathrm{cc}}$ can be defined similarly to their computational complexity counterparts.

The following are some important function families.

Definition 2.1. The *equality* function family is $\mathrm{EQ} = \left(\mathrm{EQ}_{2,n}\right)$, where $\mathrm{EQ}_{2,n} : (\{0,1\}^n)^2 \to \{0,1\}$ is defined by setting $\mathrm{EQ}_{2,n}(x_1,x_2) = 1$ if $x_1 = x_2$. The *set intersection* function family is $\mathrm{SetInt} = (\mathrm{SetInt}_{k,n})$, where $\mathrm{SetInt}_{k,n} : (\{0,1\}^n)^k \to \{0,1\}$ is defined by setting $\mathrm{SetInt}_{k,n}(x_1,\dots,x_k) = 1$ if there exists some $j \in [n]$ such that $x_{1,j} = \cdots = x_{k,j} = 1$.

Following [4], *cylinder intersections* have played an important role in multiparty communication complexity lower bounds.

Definition 2.2. An *i-cylinder* $C_i \subseteq X_1 \times \cdots \times X_k$ is a set such that for all $x_1 \in X_1,\dots,x_k \in X_k, x_i' \in X_i$ we have $(x_1,\dots,x_i,\dots,x_k) \in C_i$ if and only if $(x_1,\dots,x_i',\dots,x_k) \in C_i$. For a set $S \subseteq X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_k$, we say that *$S$ is the foot of the i-cylinder C_i* if for every $x_{-i} = (x_1,\dots,x_{i-1},x_{i+1},\dots,x_k)$, $x_{-i} \in S$ if and only if there exists $x_i \in X_i$ such that $(x_1,\dots,x_k) \in C_i$.

A *cylinder intersection* is a set of the form $C = \bigcap_{i=1}^k C_i$ where each C_i is an i-cylinder in $X_1 \times \cdots \times X_k$.

3 Non-constructive separations

3.1 Graph functions, representations, and a normal form

We are interested in a special type of Boolean functions for which we can show that deterministic protocols can be put in a restrictive "normal form," where player 1 is oblivious and every other player sends only 1 bit.

Definition 3.1. Let $X_1,\dots,X_k$ be sets. Let $g : X_2 \times \cdots \times X_k \to X_1$ be a function. Let $\mathrm{graph}^g : X_1 \times X_2 \times \cdots \times X_k \to \{0,1\}$ be the function defined by setting $\mathrm{graph}^g(x_1,x_2,\dots,x_k) = 1$ if $g(x_2,\dots,x_k) = x_1$. We say that graph^g *is the graph function of* g. In general, we say that f *is a graph function* if $f = \mathrm{graph}^g$ for some g.

We observe the following easy facts.

Fact 3.2. *A function $f : X_1 \times \cdots \times X_k \to \{0,1\}$ is a graph function if and only if for all $(x_2,\ldots,x_k) \in X_2 \times \ldots \times X_k$, there exists* exactly one $x_1 \in X_1$ *such that* $f(x_1,x_2,\ldots,x_k) = 1$. *Furthermore, a function g completely determines the graph function* graph^g*, and conversely, a graph function f completely determines a function g such that* $f = \mathrm{graph}^g$.

If f is a graph function, then it is reducible with no communication to 2-player n-bit equality $\mathrm{EQ}_{2,n}$: player 1 computes the unique value x_1^* for the input on its forehead for which the output is 1, and player 2 sees the real input x_1; players 1 and 2 now run an equality testing protocol on (x_1^*, x_1). We know that $R^0_{2,1/n}(\mathrm{EQ}_{2,n}) \le O(\log n)$ and $R_2^{0,\mathrm{pub}}(\mathrm{EQ}_{2,n}) = \Theta(1)$ [18]. Therefore, we obtain the following.

Lemma 3.3. *For every graph function family $f = (f_{k,n})$, we have $R^0_{k,1/n}(f) \le O(\log n)$ and $R_k^{0,\mathrm{pub}}(f) = \Theta(1)$. In particular, $f \in \mathsf{coRP}_k^{\mathrm{cc}}$.*

The following theorem shows that a nondeterministic protocol for a graph function can be put in a special normal form: the nondeterminism is removed, player 1 acts obliviously, and all other players simply send a check bit.

Theorem 3.4. *Let f be a graph function. Let P be a nondeterministic protocol for f, with cost d. Then, there exists a deterministic protocol P' for f such that:*

- *player 1 first sends d bits obliviously;*
- *next, all other players simultaneously send one "check" bit each;*
- *the output of the protocol is* 1 *if and only if all check bits are* 1.

Proof. Let $X_1 \times \cdots \times X_k$ be the domain of f. By Fact 3.2, there is a unique function g such that $f = \mathrm{graph}^g$. This can be computed by all players without communication, as it does not depend on the input. Furthermore, it is well known that if f has a d-bit nondeterministic protocol, then there exists a *cover* of the 1-inputs of f by a family of 2^d monochromatic cylinder intersections. Let $\left(I_\ell \mid \ell \in \{0,1\}^d\right)$ be such a cover. For every ℓ, let $I_\ell = I_\ell^1 \cap \cdots \cap I_\ell^k$, where I_ℓ^i is an i-cylinder.

We describe the protocol P' on input $x = (x_1,x_2,\ldots,x_k)$. First, player 1 computes $x_1^* = g(x_2,\ldots,x_k)$. Next, player 1 computes *some* index $\ell \in \{0,1\}^d$ such that $(x_1^*,x_2,\ldots,x_k) \in I_\ell$. Such an index exists because $(x_1^*,x_2,\ldots,x_k)$ is a 1-input of f. Player 1 communicates ℓ on d bits. Then, every other player i communicates a check bit which is equal to 1 if and only if the input $(x_1,x_2,\ldots,x_{i-1},x_{i+1},\ldots,x_k)$ seen by player i belongs to the foot of the i-cylinder I_ℓ^i. The output of P' is 1 if and only if all check bits are 1.

We argue that P' is correct. Assume $f(x) = 1$. By definition of a graph function, $x_1^* = x_1$. Then, clearly $(x_1,x_2,\ldots,x_k) = (x_1^*,x_2,\ldots,x_k) \in I_\ell$, so all check bits are 1, and P' outputs 1. Conversely, assume all check bits are 1. We claim $x \in I_\ell = I_\ell^1 \cap I_\ell^2 \cap \cdots \cap I_\ell^k$. For every $i \ge 2$, the fact that $x \in I_\ell^i$ follows from the fact that the check bit sent by player i is 1. Lastly, observe that $(x_1,x_2,\ldots,x_k) \in I_\ell^1$ if and only if $(x_1^*,x_2,\ldots,x_k) \in I_\ell^1$, because membership in I_ℓ^1 cannot depend on the first coordinate. But the latter holds by construction of I_ℓ. Hence, $x \in I_\ell$ and thus $f(x) = 1$. □

3.2 Separating randomized and deterministic protocols

In the following, we consider a class of functions which have logarithmic communication complexity for private-coin randomized protocols with bounded one-sided error. Using Theorem 3.4, we give an upper bound on the number of functions that have an efficient deterministic protocol. Since this number is smaller than the total number of functions, we conclude that some function in that class has linear deterministic communication complexity.

For integers $k \geq 2$ and $m,n \geq 1$, let $G_{k-1,n,m}$ be the set of all functions $g : (\{0,1\}^n)^{k-1} \to \{0,1\}^m$. The following Lemma is the heart of our counting argument.

Lemma 3.5. *Let $\varepsilon > 0$ and let $k = k(n)$ satisfy $k \leq 2^{(1-\varepsilon)n}$ for all large n. Let $m = m(n) := \lfloor (n - \log k)/2 \rfloor$. There exists a function $g_{k-1,n} \in G_{k-1,n,m}$ such that the function family $f = (f_{k,n})$ defined by $f_{k,n} = \mathsf{graph}^{g_{k-1,n}}$ satisfies $m-1 \leq N_k(f) \leq R^1_k(f) \leq D_k(f) \leq n+1$ for all large n.*

Observe that for the function family above, $(\varepsilon/2)n - 1 \leq m \leq n$, so $N_k(f) \geq \Omega(n)$, and $f \notin \mathsf{NP}^{\mathrm{cc}}_k$. Also, $f_{k,n}$ is a graph function on $\{0,1\}^m \times (\{0,1\}^n)^{k-1}$, so by Lemma 3.3, $f \in \mathsf{coRP}^{\mathrm{cc}}_k$. Our main result follows.

Corollary 3.6. *For every $\varepsilon > 0$ and for every $k = k(n)$ satisfying $k \leq 2^{(1-\varepsilon)n}$ for all large n, we have $\mathsf{coRP}^{\mathrm{cc}}_k \not\subseteq \mathsf{NP}^{\mathrm{cc}}_k$. In particular, $\mathsf{P}^{\mathrm{cc}}_k \neq \mathsf{RP}^{\mathrm{cc}}_k$ and $\mathsf{NP}^{\mathrm{cc}}_k \neq \mathsf{coNP}^{\mathrm{cc}}_k$.*

Proof of Lemma 3.5. A function $g \in G_{k-1,n,m}$ has a domain $(\{0,1\}^n)^{k-1}$ and range $\{0,1\}^m$. Therefore, it is impossible to encode every such function on less than $m \cdot 2^{(k-1)n}$ bits. Moreover, by Fact 3.2, for $g \neq g'$, we have $\mathsf{graph}^g \neq \mathsf{graph}^{g'}$.

Let $d = d(n)$ be the minimum value such that for all $g \in G_{k-1,n,m}$, graph^g has nondeterministic communication complexity $N(\mathsf{graph}^g) \leq d$. By Theorem 3.4, for every such g, there exists a *deterministic* protocol P' over $\{0,1\}^m \times (\{0,1\}^n)^{k-1}$ computing graph^g in the normal form given in Theorem 3.4. Every such protocol is completely specified by a function governing player 1 with domain $(\{0,1\}^n)^{k-1}$ and range $\{0,1\}^d$, and $k-1$ other functions, each one governing a distinct player $i > 1$, with domain $\{0,1\}^m \times (\{0,1\}^n)^{k-2} \times \{0,1\}^d$ and range $\{0,1\}$. These functions uniquely determine the protocol P', which in turn uniquely determines graph^g (because P' is correct on every input), which in turn uniquely determines g. Therefore, we can encode every $g \in G_{k-1,n,m}$ on at most $d \cdot 2^{(k-1)n} + (k-1) \cdot 2^{m+d+(k-2)n}$ bits.

Putting the upper and lower bounds together, we obtain

$$d \cdot 2^{(k-1)n} + (k-1) \cdot 2^{m+d+(k-2)n} \geq m \cdot 2^{(k-1)n}.$$

This is equivalent to $2^d \geq (m-d) \cdot 2^{n-m-\log(k-1)}$. Hence, $d \geq \min\{m,\ n-m-\log(k-1)\}$. Using $m = \lfloor (n-\log k)/2 \rfloor$, we get $d \geq m-1$. Trivially, $d \leq R^1_k(f) \leq D_k(f) \leq n+1$. □

3.3 Separating public-coin and private-coin protocols

We now consider the difference between public-coin and private-coin randomized protocols. Trivially, any private-coin protocol can be simulated by tossing the coins in public, so for all f and k, $R^{\mathrm{pub}}_k(f) \leq R_k(f)$. In the other direction, Newman [21] provides a simulation of a public-coin protocol by a private-coin protocol. (Although it is stated for the special case of 2 players, the proof still works if the number of players satisfies $k = k(n) \leq n^{O(1)}$.)

Proposition 3.7 ([21]). *Suppose that $k = k(n) \leq n^{O(1)}$. For every function $f : \{0,1\}^{kn} \to \{0,1\}$, we have $R_k(f) \leq R_k^{\text{pub}}(f) + O(\log n)$.*

We see that the maximum possible gap between the public-coin and private-coin randomized complexities of f is an additive $\Theta(\log n)$. A natural question that arises is whether there is a function that achieves this gap. (In the special case of $k = 2$ players, this is achieved by the equality function.) Our results allow us to answer this question affirmatively. To obtain the lower bounds, we need the following extension of Lemma 3.8 in [18] to k players.

Lemma 3.8. *For every $\varepsilon > 0$, for every $k = k(n)$, and for every function family $f = (f_{k,n})$ such that $D_k(f_{k,n}) \geq \omega(1)$ and $k = k(n) \leq O\left(D_k(f)^{1-\varepsilon}\right)$, we have $R_k(f) \geq \Omega(\log D_k(f))$.*

Before proving Lemma 3.8, we complete the separation of public-coin and private-coin protocols.

Corollary 3.9. *For every $\varepsilon > 0$ and for every $k = k(n)$ satisfying $k \leq n^{1-\varepsilon}$ for all large n, there exists a function family $f = (f_{k,n})$ such that $R_k^{\text{pub}}(f) = \Theta(1)$ and $R_k(f) = \Theta(\log n)$.*

Proof of Corollary 3.9. By Lemma 3.5 there is a family of graph functions f such that

$$\lfloor (n - \log k)/2 \rfloor - 1 \leq D_k(f) \leq n + 1\,.$$

Since $k \leq n$ for all large n, we get that $n/2 \leq D_k(f)$ for all large n. Now using the stronger assumption $k \leq n^{1-\varepsilon}$, we get $k \leq 2^{1-\varepsilon} \cdot D_k(f)^{1-\varepsilon}$ for all large n. Observe that the $2^{1-\varepsilon}$ factor does not depend on n, hence $k \leq O\left(D_k(f)^{1-\varepsilon}\right)$. We can now apply Lemma 3.8, obtaining $R_k(f) \geq \Omega(\log n)$. By Lemma 3.3, $R_k(f) \leq O(\log n)$ and $R_k^{\text{pub}}(f) = \Theta(1)$. □

Proof of Lemma 3.8. We claim the following holds. For every n,

$$D_k(f) \leq (k-1) \cdot 2^{R_{k,\varepsilon'}(f)} \cdot \left(1 + \log(k-1) - \log\left(\frac{1}{2} - \varepsilon'\right) + R_{k,\varepsilon'}(f)\right)\,. \tag{3.1}$$

We first complete the proof of the Lemma using the statement above. Assume it is not the case that $R_k(f) \geq \Omega(\log D_k(f))$, so $R_k(f) \leq \varepsilon/2 \cdot \log D_k(f)$ for infinitely many n. Using the fact that $R_k(f) = R_{k,1/3}(f)$, the assumption that $k \leq c \cdot D_k(f)^{1-\varepsilon}$ for some constant c and all large n, and Equation 3.1 above, we get that for infinitely many n,

$$\begin{aligned} D_k(f) &\leq c \cdot D_k(f)^{1-\varepsilon} \cdot D_k(f)^{\varepsilon/2} \cdot \left(1 + (1-\varepsilon) \cdot \log D_k(f) + \frac{\varepsilon}{2} \cdot \log D_k(f)\right) \\ &\leq 3 \cdot c \cdot D_k(f)^{1-(\varepsilon/2)} \cdot \log D_k(f)\,. \end{aligned}$$

Since $\varepsilon > 0$ and $D_k(f) \geq \omega(1)$, this is a contradiction. Hence, $R_k(f) \geq \Omega(\log D_k(f))$.

The proof of Equation 3.1 follows the same idea as the proof of Lemma 3.8 in [18]. Let $c = R_{k,\varepsilon'}(f_{k,n})$ and consider the ε'-error randomized protocol P that achieves communication cost c. Without loss of generality, assume that P always halts after exactly c bits are communicated. Let

$$t = 1 + \log(k-1) - \log(1/2 - \varepsilon') + c\,.$$

We construct a deterministic protocol P' for $f_{k,n}$ of cost $(k-1)\cdot 2^c \cdot t$ as follows.

For every player i and every string $\ell \in \{0,1\}^c$, let $p_i(\ell)$ denote the probability player i responds consistently with the transcript (blackboard content) ℓ (over their own private random strings). Player i (for $1 \le i \le k-1$) computes, for every ℓ, the real number $p_i(\ell)$ and publishes $p_i^*(\ell)$, a t-bit approximation of $p_i(\ell)$. This introduces an error of at most $\phi = 2^{-t}$ for every such value.

Player k now computes, for every ℓ, the value $\left(\prod_{i=1}^{k-1} p_i^*(\ell)\right) \cdot p_k(\ell)$. Since

$$p_i^*(\ell) \in \{p_i(\ell)-\phi, p_i(\ell)+\phi\}$$

and since $p_i(\ell) \le 1$, we get

$$\left(\prod_{i=1}^{k-1} p_i^*(\ell)\right) \cdot p_k(\ell) \in \left\{\prod_{i=1}^{k} p_i(\ell) - ((1+\phi)^{k-1}-1), \prod_{i=1}^{k} p_i(\ell) + ((1+\phi)^{k-1}-1)\right\}.$$

Each transcript of the randomized protocol is associated with an output (0 or 1). Player k now estimates the probability of an output in the randomized protocol by summing over the estimates of the probabilities for each transcript corresponding to that output. Finally, player k decides to output in P' the value that has a probability higher than 1/2.

In the original protocol P, an error is made with probability at most ε'. For the deterministic simulation to work, we need to make sure that the extra error introduced by the rounding process is less than $1/2-\varepsilon'$. Since each estimate has error at most $(1+\phi)^{k-1}-1$ and there are at most 2^c leaves, we need to make sure that $2^c \cdot ((1+\phi)^{k-1}-1) < (1/2-\varepsilon')$. Equivalently, we need $(1+\phi)^{k-1} < 1+(1/2-\varepsilon')/2^c$.

We choose

$$t = 1+\log(k-1)-\log(1/2-\varepsilon')+c\,, \quad \text{so} \quad \phi = 2^{-t} = \frac{1/2-\varepsilon'}{2\cdot(k-1)\cdot 2^c}\,.$$

We know that, for $0 \le x < 1/2$, we have

$$\left(1+\frac{x}{k-1}\right)^{k-1} \le e^x \le 1+2\cdot x\,.$$

Moreover, $(1/2-\varepsilon')/(2\cdot 2^c) < 1/2$. Therefore,

$$(1+\phi)^{k-1} < 1+2\cdot\frac{1/2-\varepsilon'}{2\cdot 2^c} = 1+(1/2-\varepsilon')/2^c.$$

This shows P' is correct, completing the proof of the Lemma. □

4 Lower bounds for explicit graph functions

The separation $\mathsf{coRP}_k^{\mathrm{cc}} \not\subseteq \mathsf{NP}_k^{\mathrm{cc}}$ in Section 3.2 is nonconstructive. It would be interesting to achieve this separation, or even the weaker $\mathsf{coRP}_k^{\mathrm{cc}} \not\subseteq \mathsf{P}_k^{\mathrm{cc}}$, with an explicit family of functions, even for $k=3$ players. In this section, we give two constructions of explicit families of graph functions (thus, in $\mathsf{coRP}_k^{\mathrm{cc}}$ by Lemma 3.3) that, we believe, might be outside $\mathsf{P}_k^{\mathrm{cc}}$ for $k \ge 3$. While we are unable to prove the super-polylogarithmic ($(\log n)^{\omega(1)}$) deterministic communication complexity lower bounds required to place them outside $\mathsf{P}_k^{\mathrm{cc}}$, we are able to prove much weaker logarithmic ($\Omega(\log n)$) lower bounds. As a corollary, we obtain a separation between deterministic and public-coin randomized complexities for explicit families of graph functions, though this is much weaker than our conjecture.

We begin in Section 4.1 by developing a method for representing graph functions with efficient deterministic protocols. In Section 4.2, we give a family of graph functions for the case where we have only $k = 3$ players. This construction does not immediately generalize for $k > 3$ players, but we present it first because it is both simple enough and illustrative of a certain "mixing" property (made precise later) that plays a crucial role in our arguments. We then prove the lower bound $D_k(f) \geq \Omega(\log n)$ for, in fact, *any* function family f that satisfies this mixing property. In Section 4.3, we present a different family of graph functions, that can be defined for every $k = k(n)$ with $3 \leq k \leq (1/9) \cdot \log n$, and we show that this family satisfies the same mixing property used in Section 4.2, thus yielding a similar lower bound.

4.1 Representing graph functions by colorings and cylinder intersections

Most lower bound proofs for $D_k(f)$ use the fact shown in [4] that any k-player protocol with complexity d for a function f yields a partition of the input into $O\left(2^d\right)$ cylinder intersections on which f is constant. For $k \geq 3$ players, the known techniques for proving lower bounds on the number of cylinder intersections needed for such a partition are discrepancy-based and inherently yield lower bounds even for randomized protocols. Therefore, these techniques are not suitable for proving good lower bounds for functions with low randomized communication complexity. For graph functions we obtain different, although related structures. These structures seem to be better suited for lower bound proofs for functions in $\mathsf{RP}_k^{\mathrm{cc}}$, as they allow us to prove an $\Omega(n)$ non-explicit lower bound in Lemma 3.5, as well as an $\Omega(\log n)$ explicit lower bound in Section 4.

Throughout this section, $f : X_1 \times \cdots \times X_k \to \{0,1\}$ is a graph function. Using Fact 3.2, let $g : X_2 \times \cdots \times X_k \to X_1$ be the unique function such that $f = \mathrm{graph}^g$. For any natural number D and a set S, a *D-coloring of S* is a mapping $c : S \to [D]$. We identify $\left[2^d\right] = \{1, \ldots, 2^d\}$ with $\{0,1\}^d$.

Assume that f can be computed by a d-bit protocol P. The special protocol P' for f, derived in Theorem 3.4, can be characterized by a coloring of $X_2 \times \cdots \times X_k$ and cylinder intersections in $X_2 \times \cdots \times X_k$, as follows. Let c be the 2^d-coloring of $X_2 \times \cdots \times X_k$, where $c(x_2, \ldots, x_k)$ is the message Alice sends in P' if she sees $(x_2, \ldots, x_k)$. Consider a fixed message ℓ from Alice and a fixed value $a \in X_1$ on Alice's forehead. The subset of points in $X_2 \times \cdots \times X_k$ for which all other players accept if they see a on Alice's forehead and receive message ℓ is a cylinder intersection $I_{\ell,a}$. Note, each such cylinder intersection $I_{\ell,a}$ may also contain points that are not colored ℓ. However, it is not possible that a point $p = (x_2, \ldots, x_k) \in I_{\ell,a}$ has color ℓ but $g(p) \neq a$ because then Alice would send message ℓ if she saw p and the other players would all accept if they saw a on Alice's forehead. Hence, $(a, x_2, \ldots, x_k)$ would be accepted by P', a contradiction. This proves the following.

Lemma 4.1. *Let $f : X_1 \times \cdots \times X_k \to \{0,1\}$ be a graph function. Assume f has a deterministic protocol of cost d. Then, there exist:*

- *a 2^d-coloring c of $X_2 \times \cdots \times X_k$ and*
- *a collection of cylinder intersections $\left(I_{\ell,a} \subseteq X_2 \times \cdots \times X_k \mid \ell \in \{0,1\}^d, a \in X_1\right)$,*

such that

$$\forall x = (x_1, x_2, \ldots, x_k) \in X_1 \times X_2 \times \cdots \times X_k \qquad f(x) = 1 \iff (x_2, \ldots, x_k) \in I_{c(x_2,\ldots,x_k),x_1} .$$

In particular, $I_{\ell,a}$ contains all points $y \in X_2 \times \cdots \times X_k$ with color $c(y) = \ell$ and $f(a,y) = 1$, but no point y' with color $c(y') = \ell$ and $f(a,y') = 0$.

4.2 A function family for 3 players

In the case of $k = 3$ players, our construction is based on universal families of hash functions, introduced by Carter and Wegman [11]. We begin with an informal description of our arguments.

Definition 4.2 ([11])**.** Let A, B be sets, and let $H = \{h : A \to B\}$ be a family of functions from A to B. We say that H is a *universal family of hash functions* if for every $x_1 \neq x_2 \in A$ and for every $y_1, y_2 \in B$, we have

$$\Pr_{h \in H}[h(x_1) = y_1 \text{ and } h(x_2) = y_2] = \frac{1}{|B|^2} .$$

Given a universal family H of hash functions from A to B, our three-player function f is defined on the set $B \times H \times A$ as follows. The second player holds a hash function $h \in H$, the third player holds an input $x \in A$, and $h(x) \in B$ is the unique value for the input of the first player that makes f evaluate to 1. The key "mixing" property that this construction satisfies is closely related to the Hash Mixing Lemma obtained by Mansour, Nisan, and Tiwari [20].

Lemma 4.3 ([20, Lemma 13])**.** *Let $H = \{h : A \to B\}$ be a universal family of hash functions. Let $A' \subseteq A$, $B' \subseteq B$, and $H' \subseteq H$. Then,*

$$\left| \Pr_{x \in A', h \in H'} \left[h(x) \in B' \right] - \frac{|B'|}{|B|} \right| \leq \sqrt{\frac{|H| \cdot |B'|}{|A'| \cdot |H'| \cdot |B|}} .$$

Informally, this says that, for every large rectangle $R \subseteq H \times A$ and every $b \in B$, if we visualize R as a matrix where the entry at location (h, x) has value $h(x)$, then the number of b-valued entries in R is very close to uniform, that is, $|R|/|B|$. After making these definitions precise, we use the mixing property combined with the characterization of a deterministic protocol for f from Lemma 4.1, to give some evidence as to why we believe $D_k(f)$ might be large. We subsequently prove in Theorem 4.8 that $D_k(f) \geq \Omega(\log n)$ for any function f that satisfies the mixing property.

Definition of the family We write $\mathbb{F}_q$ for the finite field of q elements when q is a prime power. Let $n \geq 4$ be a positive integer, and let $m = \lfloor n^{1/2} \rfloor$. For $x \in \{0,1\}^n$ and $a \in \{0,1\}^{n+m-1}$, let $a \circ x$ (the *convolution* of a and x) be the m-bit string z whose i-th bit is defined by $z_i = \sum_{j=1}^{n} x_j a_{(i-1)+j} \bmod 2$. For two bit strings z and b, let $z \oplus b$ be the bitwise exclusive-or of the two strings. For $(a,b) \in \mathbb{F}_{2^{n+m-1}} \times \mathbb{F}_{2^m}$, let $h_{a,b} : \mathbb{F}_{2^n} \to \mathbb{F}_{2^m}$ be defined by $h_{a,b}(x) = (a \circ x) \oplus b$. Then, $\mathcal{H}^{n,m} = \{h_{a,b} \mid a \in \mathbb{F}_{2^{n+m-1}}, b \in \mathbb{F}_{2^m}\}$ is a universal family of hash functions from $\mathbb{F}_{2^n}$ to $\mathbb{F}_{2^m}$ [20]. We are now ready to define our family of graph functions for $k = 3$ players.

Definition 4.4. Let $f = (f_{3,n})$ be the family of functions defined as follows. For n large enough, let $n' = \lfloor n/2 \rfloor$ and let $m = \lfloor n^{1/2} \rfloor$. Let $f_{3,n} : \mathbb{F}_{2^m} \times \mathcal{H}^{n',m} \times \mathbb{F}_{2^{n'}} \to \{0,1\}$ be defined by setting $f_{3,n}(y,h,x) = 1$ if $y = h(x)$.

Observe that $|\mathbb{F}_{2^m}| = 2^m \le 2^n$, $|\mathcal{H}^{n',m}| = 2^{n'+2m-1} \le 2^n$ for large n, and $|\mathbb{F}_{2^{n'}}| = 2^{n'} \le 2^n$, so they can all be embedded as subsets of $\{0,1\}^n$. Clearly, $f_{3,n}$ is a graph function, so by Lemma 3.3,

$$R^0_{3,1/n}(f) \le O(\log n) \quad \text{and} \quad R^{0,\text{pub}}_3(f) = \Theta(1)\,.$$

Mixing property The following is a “mixing” property of f, that we use to prove lower bounds on $D_3(f)$.

Definition 4.5 (Mixing Property)**.** Let $f = (f_{k,n})_{n\in\mathbb{N}}$ be a family of graph functions. Let $m = \lceil \log |X_1| \rceil$. Note that, in general, $k = k(n)$ and $m = m(n)$. Let $Z = X_2 \times \cdots \times X_k$. Writing f for $f_{k,n}$, let g be the unique function such that $f = f^g$, as given by Fact 3.2.

We say that the family f has a *mixing property* if the following holds. For large enough n, for every cylinder intersection $I \subseteq Z$ with $|I| \ge |Z| \cdot 2^{-2m}$, and for every $x_1 \in X_1$, when $(x_2,\ldots,x_k)$ is drawn uniformly from I,

$$\Pr[g(x_2,\ldots,x_k) = x_1] \le 2 \cdot \frac{1}{2^m} = 2^{-m+1}.$$

Intuitively, if g were a true random function, we’d have $\Pr[g(x_2,\ldots,x_k) = x_1] = 2^{-m}$ for every x_1 and every S. The mixing property says that f^g looks “almost random” on any “large” cylinder intersection.

Note. In the condition $|I| \ge |Z| \cdot 2^{-2m}$, the choice of $-2m$ as an exponent might seem arbitrary. As it is, Definition 4.5 allows us to prove Theorem 4.8. But, in fact, for both the construction in this section, and for the construction in Section 4.3, we could prove an even stronger mixing property, one that would apply to even smaller rectangles: for every α, for large enough n (now a function of α), for every cylinder intersection I with $|I| \ge |Z| \cdot 2^{-\alpha m}$, the same property as above would be required of I. However, we do not know how to use this stronger mixing property to prove a larger lower bound than the one in Theorem 4.8.

In the special case when f is the family from Definition 4.4, we have $k = 3$, so we can visualize $Z = \mathcal{H}^{n',m} \times \mathbb{F}_{2^{n'}}$ as a 2-dimensional matrix in which rows correspond to hash functions $h \in \mathcal{H}^{n',m}$, columns correspond to inputs $x \in \mathbb{F}_{2^{n'}}$, and the entry at row h and column x is $Z_{h,x} = h(x) \in \mathbb{F}_{2^m}$. Furthermore, cylinder intersections in 2 dimensions are rectangles. With this interpretation, the mixing property says that for every large rectangle $R \subseteq Z$, more precisely, when $|R| \ge |Z| \cdot 2^{-m}$, the number of y-entries in R is at most twice the expected number if R was filled at random with values from $\mathbb{F}_{2^m}$. The fact that f has this property follows directly from Lemma 4.3.

Lemma 4.6. *The function family $f = (f_{3,n})_{n\in\mathbb{N}}$ from Definition 4.4 has the mixing property from Definition 4.5.*

Proof. Since $H^{n',m}$ is a universal family of hash functions, by Lemma 4.3, for every rectangle $R \subseteq Z$ and for every $y \in \mathbb{F}_{2^m}$, when (h,x) are drawn uniformly from R, we have

$$\Pr[h(x) = y] \le \frac{1}{|\mathbb{F}_{2^m}|} + \left(\frac{|H^{n',m}|}{|R| \cdot |\mathbb{F}_{2^m}|} \right)^{1/2} = 2^{-m} + \left(\frac{|Z|}{|R|} \cdot 2^{-n'-m} \right)^{1/2}.$$

When $|R| \geq |Z| \cdot 2^{-2m}$, as in Definition 4.5, we have $\Pr[h(x) = y] \leq 2^{-m} + 2^{-n'/2+m/2}$. Since $n'/2 \geq 3m/2$ for large enough n, we get $\Pr[h(x) = y] \leq 2^{-m+1}$. □

Evidence towards a conjecture Let $f = (f_{3,n})$ be the function family from Definition 4.4. The following is not a precise argument, but here is why we believe that the deterministic communication complexity of f might be large.

We write f for $f_{3,n}$. Let $d = D_3(f)$. By Lemma 4.1, there exists a 2^d-coloring c of Z and there are 2^{d+m} rectangles $R_{\ell,y} \subseteq M$, for $\ell \in [2^d]$ and $y \in \mathbb{F}_{2^m}$, such that

$$\forall (y,h,x) \in \mathbb{F}_m \times H^{n,m} \times \mathbb{F}_n : \quad (h,x) \in R_{c(h,x),y} \Leftrightarrow h(x) = y\,.$$

In keeping with the matrix interpretation of Z, we say that (h,x) *is an* (ℓ,y)*-entry in* Z if $c(h,x) = \ell$ and $h(x) = y$.

Definition 4.7. For a subset $S \subseteq Z$, let $\#_{(\ell,y)}(S)$ denote the number of (ℓ,y)-entries in S and let $\rho_{(\ell,y)}(S) = \#_{(\ell,y)}(S)/|S|$ denote the density of (ℓ,y)-entries in S. We use the notation $(\ell,\cdot)$ and $(\cdot,y)$ to refer to all entries with color ℓ and all entries with value y, respectively.

For every $(\ell,y) \in [2^d] \times \mathbb{F}_{2^m}$, let

$$B_{\ell,y} = R_{\ell,y} \setminus \bigcup_{y' \neq y} R_{\ell,y'}\,.$$

We call this set the *boundary* of the rectangle $R_{\ell,y}$. Note that, by definition of the coloring, all (ℓ,y)-entries in Z are in $B_{\ell,y}$.

The entries in Z are colored with 2^d colors. Let red be a "popular" color, in the sense that $\rho_{(\text{red},\cdot)}(Z) \geq 1/2^d$. When y is chosen uniformly at random from $\mathbb{F}_{2^m}$,

$$\mathrm{E}\left[\#_{(\cdot,y)}(B_{\text{red},y})\right] \geq \mathrm{E}\left[\#_{(\text{red},y)}(B_{\text{red},y})\right] = \mathrm{E}\left[\#_{(\text{red},y)}(Z)\right] = \frac{\#_{(\text{red},\cdot)}(Z)}{2^m} \geq \frac{|Z|}{2^{m+d}}\,.$$

Furthermore, for various y, the sets $B_{\text{red},y}$ are disjoint, so

$$\mathrm{E}\left[|B_{\text{red},y}|\right] \leq \frac{|Z|}{2^m}\,.$$

If we fixed a y for which both quantities above are within constant factors of their expectations (this is imprecise, but we believe it could be made precise), we would obtain a pair (red,y) such that $\rho_{(\cdot,y)}(B_{\text{red},y}) \geq \Omega\left(1/2^d\right)$, whereas the mixing property (Definition 4.5) says that, if $R_{\text{red},y}$ is large enough, $\rho_{(\cdot,y)}(R_{\text{red},y}) \leq O(1/2^m)$. The smaller d is, the larger the gap between these numbers. The reason this situation does not immediately translate into a lower bound for d in terms of m is that $B_{\text{red},y}$ is far from a rectangle, so its density of $(\cdot,y)$-entries could be much larger than that of the rectangle $R_{\text{red},y}$. However, we also consider this as evidence that d might have to be large.

A weaker lower bound In the following Theorem, we show that every family of graph functions $f = (f_{k,n})$ that satisfies the mixing property from Definition 4.5 has deterministic communication complexity at least logarithmic in $m = m(n) = \lceil \log |X_1| \rceil$, the size of the input of player 1.

We describe the proof idea for the case of $k = 3$ players, when we can view Z as a matrix, but the proof itself works in general for $k \geq 3$. Let $d = D_3(f)$ and consider the 2^d-coloring c of the matrix Z given by Lemma 4.1. The proof proceeds by inductively decreasing the number of colors available and shrinking the matrix. During each step, we introduce a number of "holes" in the matrix (entries that are colored in the original matrix with one of the removed colors). We show that eventually there are no colors left to use, but the matrix still does not consist only of holes.

To see the limitation of this technique, note that even if we were able to shrink the matrix by, say, a factor of 2 in order to remove every color, there are still 2^d colors to remove, so at the end we would have shrunk the matrix by a factor of 2^{2^d}. Since the matrix Z has size $2^{m^{O(1)}}$, this technique can only produce lower bounds of the form $d \geq \Omega(\log m)$.

Theorem 4.8. *Let $f = (f_{k,n})$ be a family of graph functions, with $f_{k,n} : X_1 \times \cdots \times X_k \to \{0,1\}$, that satisfies the mixing property in Definition 4.5. Let $m = m(n) = \lceil \log |X_1| \rceil$ be the size of the input of player 1. Then, $D_k(f_{k,n}) \geq \Omega(\log m)$.*

Proof. We write f for $f_{k,n}$. Let $d = D_k(f)$. Let $Z = X_2 \times \cdots \times X_k$. Using Fact 3.2, let $g : Z \to X_1$ be the unique function such that $f = f^g$. By Lemma 4.1, there is a 2^d-coloring c of Z and there are 2^{d+m} cylinder intersections $I_{\ell,y}$, for $(\ell, y) \in [2^d] \times X_1$, such that

$$\forall (y, x_2, \ldots, x_k) \in X_1 \times Z, \qquad (x_2, \ldots, x_k) \in I_{c(x_2,\ldots,x_k),y} \iff g(x_2, \ldots, x_k) = y\,.$$

We say that a point $(x_2, \ldots, x_k)$ has *color* $c(x_2, \ldots, x_k)$ and *value* $g(x_2, \ldots, x_k)$. For a set $S \subseteq Z$, let $\#_{(\ell,y)}(S)$ denote the number of points $(x_2, \ldots, x_k) \in S$ with color ℓ and value y.

Assume there exists some $\varepsilon < 1$ such that for all m, $d \leq \varepsilon \cdot \log m$. We will derive a contradiction.

For large enough m, we prove by induction that for all $0 \leq i \leq 2^d$, there exists a cylinder intersection $I_i \subseteq Z$ and a set of "holes" $H_i \subseteq I_i$ such that:

- $|I_i| \geq |Z| \cdot 2^{-i(d+2)}$;
- $|H_i| \leq i \cdot |Z| \cdot 2^{-m+1}$; and
- the initial coloring induces a coloring of the points in $I_i \setminus H_i$ with $2^d - i$ colors.

Assuming we have established this inductive statement, letting $i = 2^d$, we see that we must have $I_{2^d} \setminus H_{2^d} = \emptyset$, for any points in this set would have been uncolored in the original coloring. For large enough m,

$$|I_{2^d}| \geq |Z| \cdot 2^{-2^d(d+2)} \geq |Z| \cdot 2^{-m^{\varepsilon}(\varepsilon \cdot \log m + 2)} > |Z| \cdot 2^{-m + \varepsilon \cdot \log m + 1} \geq |Z| \cdot 2^{-m+d+1}\,.$$

Since $|H_{2^d}| \leq 2^d \cdot |Z| \cdot 2^{-m+1}$, we get $I_{2^d} \setminus H_{2^d} \neq \emptyset$, which is a contradiction.

We now prove the inductive statement. For $i = 0$, let $I_0 = Z$ and let $H_0 = \emptyset$. Then, c is a coloring of $I_0 \setminus H_0$ with 2^d colors.

Next, assume the inductive statement is true for some $0 \leq i < 2^d$. We have

$$|I_i \setminus H_i| \geq |Z| \cdot \left(2^{-i(d+2)} - i \cdot 2^{-m+1}\right) > |Z| \cdot \left(2^{-i(d+2)} - 2^{-m+d+1}\right) > |Z| \cdot 2^{-i(d+2)-1}\,,$$

where in the last inequality we used $m-d-1 > m^{\varepsilon}(\varepsilon \cdot \log m + 2) > i(d+2)$, for large enough m.

Let (ℓ, y) be the most popular color-value pair in $I_i \setminus H_i$. There are at most 2^{m+d} such pairs, so

$$\#_{(\ell,y)}(I_i \setminus H_i) \geq |I_i \setminus H_i| \cdot 2^{-m-d} \geq |Z| \cdot 2^{-i(d+2)-1-m-d} = |Z| \cdot 2^{-(i+1)(d+2)-m+1} .$$

Let $I_{i+1} = I_i \cap I_{\ell,y}$, which is a cylinder intersection in Z because both I_i and $I_{\ell,y}$ are. By the property of the coloring c in Lemma 4.1, all points in Z with color-value (ℓ, y) are in $I_{\ell,y}$. Hence,

$$|I_{i+1}| \geq \#_{(\ell,y)}(I_i \setminus H_i) \geq |Z| \cdot 2^{-(i+1)(d+2)-m+1} .$$

Note that $(i+1)(d+2)-1 \leq m^{\varepsilon}(\varepsilon \log m + 2) - 1 < m$, so $|I_{i+1}| \geq |Z| \cdot 2^{-2m}$. Then, by the mixing property in Definition 4.5, $\#_{(\cdot,y)}(I_{i+1})/|I_{i+1}| \leq 2^{-m+1}$, so $|I_{i+1}| \geq \#_{(\cdot,y)}(I_{i+1}) \cdot 2^{m-1}$. Also, $\#_{(\cdot,y)}(I_{i+1}) \geq \#_{(\ell,y)}(I_{i+1}) \geq \#_{(\ell,y)}(I_i \setminus H_i)$. Putting these together, we get

$$|I_{i+1}| \geq \#_{(\ell,y)}(I_i \setminus H_i) \cdot 2^{m-1} \geq |Z| \cdot 2^{-(i+1)(d+2)-m+1} \cdot 2^{m-1} = |Z| \cdot 2^{-(i+1)(d+2)} ,$$

establishing the first part of the inductive statement.

Let $H_{i+1} = H_i \cup \{(x_2, \ldots, x_k) \in I_{i+1} \mid c(x_2, \ldots, x_k) = \ell\}$. By the induction hypothesis, points in $I_i \setminus H_i$ are colored with at most $2^d - i$ colors. Since ℓ is no longer available, we see that all points left in $I_{i+1} \setminus H_{i+1}$ must now be colored with at most $2^d - i - 1$ colors, establishing the third part of the inductive statement.

Finally, by the property of c in Lemma 4.1, all points in $I_{\ell,y}$ that have color ℓ must have value y. In the entire set Z, by the mixing property, there are at most $|Z| \cdot 2^{-m+1}$ points with value y. Then,

$$|H_{i+1}| \leq |H_i| + |Z| \cdot 2^{-m+1} \leq (i+1) \cdot |Z| \cdot 2^{-m+1} ,$$

establishing the second part of the inductive statement. □

Corollary 4.9. *The function family f from Definition 4.4 satisfies $D_3(f) \geq \Omega(\log n)$. Furthermore, f provides a separation between public-coin and private-coin complexities, as $R_3^{\text{pub}}(f) \leq O(1)$ and $R_3(f) \geq \Omega(\log \log n)$.*

Proof. By Lemma 4.6, f has the mixing property. By Theorem 4.8, $D_3(f) \geq \Omega(\log m)$. By definition of f, $m = (n/3)^{1/2}$, so $D_3(f) \geq \Omega(\log n)$. Since the family f consists of graph functions, by Lemma 3.3, $R_3^{\text{pub}}(f) \leq O(1)$. Finally, by Lemma 3.8, $R_3(f) \geq \Omega(\log \log n)$. □

4.3 A function family for 3 or more players

In the case of $k \geq 3$ players, the family of graph functions we construct is based on the generalized inner product function $\text{GIP}^{\mathbb{F}}$ over a finite field $\mathbb{F}$. It turns out that the crucial "mixing property" in Definition 4.5 follows quite easily from a bound on the "strong discrepancy" of $\text{GIP}^{\mathbb{F}}$ obtained by Babai, Hayes, and Kimmel [3].

Following [3], for a *non-Boolean* function $f : Z \to B$ and a set $S \subseteq Z$, the *strong discrepancy* of f on S is defined as follows:[1]

$$\begin{aligned}\mathrm{disc}^*(f,S) &:= \max_{y\in B} \frac{1}{|Z|} \cdot \left| |f^{-1}(y) \cap S| - \frac{|S|}{|B|} \right| \\ &= \max_{y\in B} \frac{|S|}{|Z|} \cdot \left| \Pr_{z\sim S}[f(z) = y] - \frac{1}{|B|} \right| .\end{aligned}$$

This measure considers *all* possible function values $y \in B$, and it reflects the largest difference between the actual number of y-inputs in S and the average number of such inputs if f were perfectly balanced on S. It is not hard to see that strong discrepancy is connected to the mixing property from Definition 4.5. Before making this connection formal, we define the function we will be using.

Definition 4.10. For integers $n \geq 1$, $k \geq 2$, and for a finite field $\mathbb{F}$, the generalized inner product function over $\mathbb{F}$, $\mathrm{GIP}^{\mathbb{F}}_{k,n} : (\mathbb{F}^n)^k \to \mathbb{F}$, is defined as follows. For $1 \leq i \leq k$, let $x_i = (x_{i,1}, \ldots, x_{i,n}) \in \mathbb{F}^n$, where for $1 \leq j \leq n$, $x_{i,j} \in \mathbb{F}$. Then,

$$\mathrm{GIP}^{\mathbb{F}}_{k,n}(x_1, \ldots, x_k) := \sum_{j=1}^{n} \left(\prod_{i=1}^{k} x_{i,j} \right) .$$

In [3], the following upper bound is obtained on the strong discrepancy of this function.

Fact 4.11 ([3][2]). *Let $n, m \geq 1$ and $p \geq 2$. For every cylinder intersection $I \subseteq ((\mathbb{F}_{2^m})^n)^p$,*

$$\mathrm{disc}^*\left(\mathrm{GIP}^{\mathbb{F}_{2^m}}_{p,n}, I\right) \leq \left(1 - \frac{1}{2^m}\right) \cdot \left(1 - \left(1 - \frac{1}{2^m}\right)^{p-1}\right)^{n \cdot 2^{1-p}} .$$

The following is the family of graph functions we use for $k \geq 3$ players.

Definition 4.12. Let $k = k(n) \geq 3$ be some function, let $n' = n'(n) := \lfloor n^{1/2} \rfloor$, and let $m = m(n) := \lfloor n^{1/4} \rfloor$. Let $V = (\mathbb{F}_{2^m})^{n'}$. Let $f = (f_{k,n})$ be the function family where $f_{k,n} : \mathbb{F}_{2^m} \times V^{k-1} \to \{0,1\}$ is defined by setting $f_{k,n}(x_1, x_2, \ldots, x_k) = 1$ if $\mathrm{GIP}^{\mathbb{F}_{2^m}}_{k-1,n'}(x_2, \ldots, x_k) = x_1$.

We claim that the small strong discrepancy of $\mathrm{GIP}^{\mathbb{F}_{2^m}}_{k-1,n'}$ implies that f satisfies the mixing property from Definition 4.5.

Lemma 4.13. *If $k = k(n)$ satisfies $3 \leq k \leq (1/9) \cdot \log n$, then the function family f from Definition 4.12 has the mixing property from Definition 4.5.*

Proof of Lemma 4.13. Let $I \subseteq V^{k-1}$ be a cylinder intersection with $|I| \geq |V^{k-1}| \cdot 2^{-2m}$ and let $y \in \mathbb{F}_{2^m}$. We want to show that

$$\Pr_{(z_2,\ldots,z_k)\sim I} \left[\mathrm{GIP}^{\mathbb{F}_{2^m}}_{k-1,n'}(z_2, \ldots, z_k) = y \right] \leq 2^{-m+1}.$$

[1] We note that strong discrepancy when $|B| = 2$ and regular (Boolean) discrepancy are off by a factor of 2. To get them to agree, we'd have to multiply strong discrepancy by $|B|$. In here, we simply use the definition given in [3].

[2] In [3], this fact is not stated as a stand-alone lemma, but it is proved and used as part of the proof of Corollary 4.12.

Writing z for $(z_2,\dots,z_k)$ and g for $\mathrm{GIP}^{\mathbb{F}_{2^m}}_{k-1,n'}$, from the definition of strong discrepancy we get

$$\begin{aligned}\left|\Pr_{z\sim I}[g(z)=y]-\frac{1}{2^m}\right| &\leq \frac{|V^{k-1}|}{|I|}\cdot \mathrm{disc}^*(g,I)\\ &\leq 2^{2m}\cdot \mathrm{disc}^*(g,I)\,.\end{aligned}$$

By Fact 4.11 (with $p=k-1$),

$$\begin{aligned}\mathrm{disc}^*(g,I) &\leq \left(1-\frac{1}{2^m}\right)\cdot\left(1-\left(1-\frac{1}{2^m}\right)^{k-2}\right)^{n'\cdot 2^{2-k}}\\ &\leq \left(1-\left(1-\frac{1}{2^m}\right)^{k-2}\right)^{n'\cdot 2^{2-k}} && \text{(since } m\geq 1\text{)}\\ &\leq \exp\left(-\left(1-\frac{1}{2^m}\right)^{k-2}\cdot n'\cdot 2^{2-k}\right) && \text{(using } 1-x\leq e^{-x}\text{)}\\ &\leq \exp\left(-\frac{1}{2^{k-2}}\cdot n'\cdot 2^{2-k}\right) && \text{(since } 1-1/2^m\geq 1/2\text{)}\\ &= 2^{-16\cdot(\log e)\cdot n'\cdot 2^{-2k}}\,.\end{aligned}$$

Putting everything together,

$$\begin{aligned}\Pr_{z\sim\mathcal{U}(I)}[g(z)=y] &\leq \frac{1}{2^m}+2^{2m}\cdot \mathrm{disc}^*(g,I)\\ &\leq 2^{-m}+2^{2m-16\cdot(\log e)\cdot n'\cdot 2^{-2k}}\\ &\leq 2^{-m}+2^{-m}=2^{-m+1} && \text{(for large enough } n\text{).}\end{aligned}$$

Above, we need $n'\geq(3\cdot m\cdot 2^{2k})/(16\cdot\log e)$. Since $\log n'\geq(1/2)\cdot\log n-1$, $\log m\leq(1/4)\cdot\log n$ and $2k\leq(2/9)\cdot\log n<(1/4)\cdot\log n$, the inequality holds for large enough n. □

Corollary 4.14. *For every function $k=k(n)$ such that $3\leq k\leq(1/9)\cdot\log n$, the function family f from Definition 4.12 satisfies $D_k(f_{k,n})\geq\Omega(\log n)$. Furthermore, f separates public-coin and private-coin randomized protocols, as $R^{\mathrm{pub}}_k(f_{k,n})\leq O(1)$ and $R_k(f_{k,n})\geq\Omega(\log\log n)$.*

Proof. By Lemma 4.13 and Theorem 4.8, $D_k(f_{k,n})\geq\Omega(\log m)$. Since $m=\lfloor n^{1/4}\rfloor$, $D_k(f_{k,n})\geq\Omega(\log n)$. Since $F_{k,n}$ is a graph function, by Lemma 3.3, $R^{\mathrm{pub}}_k(f_{k,n})\leq O(1)$. Finally, by Lemma 3.8, $R_k(f_{k,n})\geq\Omega(\log\log n)$. □

5 On complete problems

An alternative approach to separating $\mathsf{P}^{\mathrm{cc}}_k$ from $\mathsf{RP}^{\mathrm{cc}}_k$ with an explicit function is to find a function that is complete in some sense. If we can prove for some explicit function that it is "at least as hard" as any function in $\mathsf{RP}_k^{\mathrm{cc}}$, then by our separation result we can conclude that it is not in $\mathsf{P}_k^{\mathrm{cc}}$. The set intersection function is complete for the class analogous to $\mathsf{NP}_k^{\mathrm{cc}}$ in the number-in-hand (NIH) model, and thus also for $\mathsf{NP}_2^{\mathrm{cc}}$. In this section, we prove that this function is not complete for $\mathsf{NP}_k^{\mathrm{cc}}$ for $k \geq 3$.

In two-player communication complexity, Babai, Frankl, and Simon [1] defined a natural notion of a reduction between problems called a "rectangular" reduction that does not require any communication to compute, as well as an appropriate "polynomially-bounded" version of rectangular reduction for function families.

Definition 5.1 ([1])**.** Let $f : X_1 \times X_2 \to \{0,1\}$ and $g : X_1' \times X_2' \to \{0,1\}$. A pair of functions (φ_1, φ_2) with $\varphi_i : X_i \to X_i'$ is a *rectangular reduction* of f to g, written $f \sqsubseteq g$, if $f(x_1,x_2) = g(\varphi_1(x_1), \varphi_2(x_2))$. For function families $f = \{f_n\}$ and $g = \{g_n\}$ where $f_n, g_n : (\{0,1\}^n)^2 \to \{0,1\}$, we write $f \sqsubseteq_p g$ if there is a function $m : \mathbb{N} \to \mathbb{N}$ such that for every n, $f_n \sqsubseteq g_{m(n)}$ and $m(n)$ is $2^{(\log n)^{O(1)}}$.

Proposition 5.2 ([1])**.** *Let f and g be function families. If $f \sqsubseteq_p g$ and $g \in \mathsf{P}_2^{\mathrm{cc}}$ then $f \in \mathsf{P}_2^{\mathrm{cc}}$. If $f \sqsubseteq_p g$ and $g \in \mathsf{NP}_2^{\mathrm{cc}}$ then $f \in \mathsf{NP}_2^{\mathrm{cc}}$.*

Definition 5.3. A function family g is *complete for $\mathsf{NP}_2^{\mathrm{cc}}$ under rectangular reductions* if $g \in \mathsf{NP}_2^{\mathrm{cc}}$ and for all $f \in \mathsf{NP}_2^{\mathrm{cc}}$, $f \sqsubseteq_p g$.

Recall the set intersection function SetInt from Definition 2.1 Clearly, SetInt $\in \mathsf{NP}_k^{\mathrm{cc}}$. In [1] it was observed that:

Proposition 5.4 ([1])**.** SetInt *is complete for $\mathsf{NP}_2^{\mathrm{cc}}$ under rectangular reductions.*

For $k \geq 3$, rectangular reductions extend to *cubic reductions* in the NIH model of communication complexity. Moreover, it is easy to see that the completeness result of Proposition 5.4 continues to hold in the NIH model under cubic reductions. One might expect that SetInt is also complete for $\mathsf{NP}_k^{\mathrm{cc}}$ under a natural extension of rectangular reductions in the NOF model. Such a notion of reduction should not require any communication between the players. This yields the following definition:

Definition 5.5. Given $f : X_1 \times \cdots \times X_k \to \{0,1\}$ and $g : X_1' \times \cdots \times X_k' \to \{0,1\}$, we say that a k-tuple of functions $(\varphi_1, \ldots, \varphi_k)$ is a *cylindrical reduction* of f to g if for every $(x_1,\ldots,x_k) \in X_1 \times \cdots \times X_k$ there is an $(x_1',\ldots,x_k') \in X_1' \times \cdots \times X_k'$ such that $\varphi_i(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_k) = (x_1',\ldots,x_{i-1}',x_{i+1}',\ldots,x_k')$ for all $i \in [k]$, and $f(x_1,\ldots,x_k) = g(x_1',\ldots,x_k')$. Thus, each φ_i maps the NOF view of the i-th player on input $(x_1,\ldots,x_k)$ for f to the NOF view of the i-th player on input $(x_1',\ldots,x_k')$ for g.

We show that cylindrical reductions must be of a very special form, given by the natural no-communication reductions associated with the number-in-hand model.

Definition 5.6. Given $f : X_1 \times \cdots \times X_k \to \{0,1\}$ and $g : X_1' \times \cdots \times X_k' \to \{0,1\}$, we say that a k-tuple of functions $(\psi_1, \ldots, \psi_k)$ is a *cubic reduction* of f to g if $\psi_i : X_i \to X_i'$ for every i, and $f(x_1,\ldots,x_k) = g(\psi_1(x_1), \ldots, \psi_k(x_k))$.

Lemma 5.7. *If $(\varphi_1, \ldots, \varphi_k)$ is a cylindrical reduction of f to g then there is a cubic reduction $(\psi_1, \ldots, \psi_k)$ of f to g such that, for all i,*

$$\varphi_i(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_k) = (\psi_1(x_1), \ldots, \psi_{i-1}(x_{i-1}), \psi_{i+1}(x_{i+1}), \ldots, \psi_k(x_k)).$$

Proof. We prove by induction on k that any consistent cylindrical reduction (whether or not it correctly reduces f to g) must be cubic. The claim is trivial for $k = 2$. Assume that $k > 2$. Consider $(x_1, \ldots, x_k)$ and let $(x'_1, \ldots, x'_k)$ be the output of the cylindrical reduction on x. Let $y_k \in X_k$. The requirement that $\varphi_k(x_1, \ldots, x_{k-1}) = (x'_1, \ldots, x'_{k-1})$ and the fact that the views output by the φ_i for $i < k$ must be consistent with this output implies that for $i < k$,

$$\varphi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}, x_k) = (x'_1, \ldots, x'_{i-1}, x'_{i+1}, \ldots, x'_{k-1}, x'_k)$$

and

$$\varphi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}, y_k) = (x'_1, \ldots, x'_{i-1}, x'_{i+1}, \ldots, x'_{k-1}, y'_k)$$

for some $y'_k \in X_k$. Thus the first $k-2$ coordinates of the output of φ_i are independent of the last coordinate of its input. Since $x = (x_1, \ldots, x_k)$ and y_k were chosen arbitrarily, for any such input we can define functions $(\varphi'_1, \ldots, \varphi'_{k-1})$ where $\varphi'_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1})$ consists of the first $k-2$ coordinates of $\varphi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}, y_k)$ for any $y_k \in X_k$. These form a consistent map on $k-1$ coordinates and therefore by the inductive hypothesis there are $(\psi_1, \ldots, \psi_{k-1})$ such that

$$\varphi'_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}) = (\psi_1(x_1), \ldots, \psi_{i-1}(x_{i-1}), \psi_{i+1}(x_{i+1}), \ldots, \psi_{k-1}(x_{k-1}))$$

and therefore for $i < k$ and any $x_1, \ldots, x_k \in X_k$,

$$\varphi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{k-1}, x_k) = (\psi_1(x_1), \ldots, \psi_{i-1}(x_{i-1}), \psi_{i+1}(x_{i+1}), \ldots, \psi_{k-1}(x_{k-1}), x'_k)$$

for some $x'_k = \phi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k)$ for some function ϕ_i; i. e., φ_i acts componentwise on all but the k-th coordinate. Now, since $k > 2$ there is some $j \notin \{i, k\}$ and, by symmetry, the same inductive argument can be applied to characterize φ_i for all $i \neq j$ so that φ_i acts componentwise on all but the j-th coordinate. Moreover, the inductive argument implies that there are functions $\psi'_1, \ldots, \psi'_{j-1}, \psi'_{j+1}, \ldots, \psi'_k$ that give this componentwise behavior. Defining $\psi'_j = \psi_j$ and $\psi_k = \psi'_k$ we see that we must have $\psi'_i = \psi_i$ for all $i \in [k]$. Therefore,

$$\varphi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k) = (\psi_1(x_1), \ldots, \psi_{i-1}(x_{i-1}), \psi_{i+1}(x_{i+1}), \ldots, \psi_k(x_k))$$

for $i < k$. Since k was arbitrarily chosen, the same applies for $i = k$ and the result follows by induction. □

Definition 5.8. The set $A \subseteq X_1 \times \cdots \times X_k$ is a *cube* if $A = A_1 \times \cdots \times A_k$ for some sets $A_i \subseteq X_i$, for all $i \in [k]$.

Lemma 5.9. *If there is a cylindrical reduction of $f : X_1 \times \cdots \times X_k \to \{0,1\}$ to* $\mathrm{SetInt}_{k,m}$ *then $f^{-1}(1)$ is a union of m cubes.*

Proof. By Lemma 5.7, there are functions $(\psi_1, \ldots, \psi_k)$ such that

$$f(x_1, \ldots, x_k) = \mathrm{SetInt}_{k,m}(\psi_1(x_1), \ldots, \psi_k(x_k)).$$

Thus $(x_1, \ldots, x_k) \in f^{-1}(1)$ if and only if there is some $i \in [m]$ such that the i-th coordinate of each of $\psi_1(x_1), \ldots, \psi_k(x_m)$ is 1. For $j \in [k]$ let $A_{i,j} = \{x_j \mid \text{the } i\text{-th coordinate of } \psi_j(x_j) \text{ is } 1\} \subseteq X_j$. Therefore $A_{i,1} \times \cdots \times A_{i,k}$ is a cube for each $i \in [m]$ and $f^{-1}(1) = \bigcup_{i \in [m]} A_{i,1} \times \cdots \times A_{i,k}$ as required. □

Theorem 5.10. *There is a function $f : (\{0,1\}^n)^3 \to \{0,1\}$ with deterministic 3-player NOF communication complexity at most 3 such that any cylindrical reduction of f to* $\mathrm{SetInt}_{3,m}$ *requires $m > 2^{n-3}$.*

Proof. For $x,y,z \in \{0,1\}^n$, define $f(x,y,z)$ to be 1 if x, y, and z are pairwise orthogonal in $\mathbb{F}_2^n$ under the standard dot product. There is a trivial 3-player NOF protocol for f in which 3 bits are exchanged, namely, each player checks that the inputs it sees are orthogonal. We now show that any way to write $f^{-1}(1)$ as a union of cubes must contain exponentially many cubes since each cube can only cover an exponentially small portion of $f^{-1}(1)$.

For $u,v \in \{0,1\}^n$, let $h(u,v) = 1$ if $\langle x,y\rangle = 0$ in $\mathbb{F}_2^n$. Then $f(x,y,z) = h(x,y)h(y,z)h(x,z)$. Consider the uniform distribution μ over $\{0,1\}^{3n}$.

We first show that $f^{-1}(1)$ is a set of probability more than 1/8. Under μ, for each pair $u,v \in \{x,y,z\}$, the probability that $h(u,v) = 1$ is $1/2 + 1/2^n > 1/2$ (consider whether or not $u = 0^n$). We claim that the probability that $f(x,y,z) = 1$ is at least 1/8. Suppose that $x \neq 0^n$. Then the probability that y is orthogonal to x is precisely 1/2. Now, z is orthogonal to the span $\langle\{x,y\}\rangle$ with probability at least 1/4. So, conditioned on $x \neq 0^n$, the probability that $f(x,y,z) = 1$ is at least 1/8. If $x = 0^n$ then the probability that $f(x,y,z) = 1$ is precisely the probability that y and z are orthogonal which is at least $1/2$. Therefore the probability that $f(x,y,z) = 1$ is more than $1/8$ overall.

Now since $f(x,y,z) = h(x,y)h(y,z)h(x,z)$, any cube $C = A_1 \times A_2 \times A_3$ with $C \subseteq f^{-1}(1)$ must, in particular, have, $A_1 \times A_2 \subseteq h^{-1}(1)$. Thus every $x \in A_1$ must be orthogonal to every $y \in A_2$ and so the dimensions of their spans must satisfy $\dim(\langle A_1\rangle) + \dim(\langle A_2\rangle) \le n$. Therefore

$$|A_1 \times A_2| \le |\langle A_1\rangle \times \langle A_2\rangle| \le 2^{\dim(\langle A_1\rangle)+\dim(\langle A_2\rangle)} \le 2^n$$

so $|C| \le 2^n \cdot |A_3| \le 2^{2n}$ and the probability that $(x,y,z) \in C$ is at most 2^{-n}. The claimed result follows immediately. □

This argument can be extended to other functions $h : \{0,1\}^{2n} \to \{0,1\}$ that have only small 1-monochromatic rectangles. It suffices that $h(x,y)h(y,z)h(x,z)$ be 1 on a large fraction of inputs. Also, although the above Lemma is stated only for $k = 3$ it is easy to see that the same bounds hold for larger k.

Given that any function $f(x,y,z)$ of the form $h_1(x,y)h_2(x,z)h_3(y,z)$ has communication complexity at most 3, it seems unlikely that any function is complete for $\mathsf{NP}_3^{\mathrm{cc}}$ under efficient reductions that do not require communication.

6 Discussion

In this paper we give a nonconstructive separation of the NOF communication complexity classes $\mathsf{P}_k^{\mathrm{cc}}$ and $\mathsf{RP}_k^{\mathrm{cc}}$ for up to $k \le n^{O(1)}$ players. We leave it as an open problem to exhibit an explicit function achieving this separation, even for as few as $k = 3$ players.

To put the limitation on the number of players in perspective, we note that the only method for proving any strong lower bound for an explicit function is the discrepancy method, and we only know how to bound discrepancy in the case of up to $k = \log n$ players. Moreover, both the original discrepancy method [4], and its recent generalizations [17, 24, 25] automatically yield lower bounds for randomized protocols. It is an open problem to adapt these methods to proving deterministic lower bounds for functions with efficient randomized protocols.

Prior to this work, the best deterministic lower bound for a function with efficient randomized protocols was $\Omega(\log\log n)$ for the Exact-T function [9]. In this paper, we improve this to an $\Omega(\log n)$ lower bound for two families of explicit functions, one based on universal families of hash functions, the other on the generalized inner product function.

Perhaps surprisingly, some of the technical arguments we use in this paper bear a similarity to the ones used by Babai, Hayes, and Kimmel [3]. In that paper, the goal was to obtain lower bounds for k-player NOF protocols "with help." These are protocols for non-Boolean functions $g : (\{0,1\}^n)^k \to \{0,1\}^m$, where at the start of the protocol, the players receive an $(m-1)$-bit message from a benevolent helper that sees the entire input. The result proved in that paper is discrepancy-based, so it works even for randomized protocols with help, and it says that there are functions for which the best protocol with help requires $\Omega(n/c^k)$ bits of communication, for some constant c.

Even though the results in [3] and the ones presented in this paper seem rather different, it turns out that they are somewhat similar at the technical level. We have seen that if a function f^g has a deterministic protocol of cost c, then there is a 2^c-coloring of the inputs to g, and cylinder intersections $I_{\ell,x}$ such that, for every (color) ℓ, $(I_{\ell,x} \mid x \in \{0,1\}^m)$ *partition the ℓ-colored inputs to g*. In contrast, if a function g has a protocol with c bits of help and communication cost s, it can be shown that there are cylinder intersections $T_{\ell,x}$ such that, for every (message/color) ℓ, $(T_{\ell,x} \mid x \in \{0,1\}^s)$ *partition the entire set of inputs to g*. Notably, the latter partition is stricter than the former, because the sets corresponding to color ℓ can no longer intersect on inputs which are not ℓ-colored. In [3], a lower bound was obtained for the cost of a protocol with help directly from the strong discrepancy of the function being computed. In this paper, we are unable to derive a similar direct connection; instead, we use strong discrepancy indirectly, through the Mixing Property in Definition 4.5, and we obtain weaker bounds.

It would also be interesting to study whether counting arguments can be used to separate other types of protocols, for example, randomized from nondeterministic for the same type of one-sided error (i. e., $\mathsf{RP}_k^{\mathrm{cc}} \neq \mathsf{NP}_k^{\mathrm{cc}}$). Such a separation exists [14], but in [14], the number of players is limited to $k < \log n$ because of its ultimate reliance on the iterated Cauchy-Schwarz scheme used to estimate discrepancy [4]. Potentially, a counting argument might avoid this limitation.

7 Acknowledgments

Originally, we stated Theorem 3.4 only for deterministic protocols. We are grateful to the anonymous referees for pointing out that the original proof easily extends even to nondeterministic protocols.

References

[1] LÁSZLÓ BABAI, PÉTER FRANKL, AND JÁNOS SIMON: Complexity classes in communication complexity theory (preliminary version). In *Proc. 27th FOCS*, pp. 337–347. IEEE Comput. Soc. Press, 1986. [doi:10.1109/SFCS.1986.15] 203, 205, 218

[2] LÁSZLÓ BABAI, ANNA GÁL, PETER KIMMEL, AND SATYANARAYANA LOKAM: Communication complexity of simultaneous messages. *SIAM J. Comput.*, 33(1):137–166, 2004. [doi:10.1137/S0097539700375944] 205

[3] LÁSZLÓ BABAI, THOMAS HAYES, AND PETER KIMMEL: The cost of the missing bit: Communication complexity with help. *Combinatorica*, 21(4):455–488, 2001. 215, 216, 221

[4] LÁSZLÓ BABAI, NOAM NISAN, AND MÁRIÓ SZEGEDY: Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. Comput. System Sci.*, 45(2):204–232, 1992. [doi:10.1016/0022-0000(92)90047-M] 202, 205, 210, 220, 221

[5] PAUL BEAME, MATEI DAVID, TONIANN PITASSI, AND PHILIPP WOELFEL: Separating deterministic from nondeterministic NOF multiparty communication complexity. In *Proc. 34th Intern. Colloq. Autom. Lang. Program. (ICALP)*, volume 4596 of *LNCS*, pp. 134–145. Springer, 2007. 202, 203

[6] PAUL BEAME, TRINH HUYNH, AND TONIANN PITASSI: Hardness amplification in proof complexity. In *Proc. 42nd STOC*, pp. 87–96. ACM Press, 2010. [doi:10.1145/1806689.1806703] 202

[7] PAUL BEAME AND DANG-TRINH HUYNH-NGOC: Multiparty communication complexity and threshold circuit size of AC^0. In *Proc. 50th FOCS*, pp. 53–62. IEEE Comput. Soc. Press, 2009. [doi:10.1109/FSCS.2009.12] 203

[8] PAUL BEAME, TONIANN PITASSI, AND NATHAN SEGERLIND: Lower bounds for Lovász–Schrijver systems and beyond follow from multiparty communication complexity. *SIAM J. Comput.*, 37(3):845–869, 2007. [doi:10.1137/060654645] 202, 203

[9] RICHARD BEIGEL, WILLIAM GASARCH, AND JAMES GLENN: The multiparty communication complexity of Exact-T: Improved bounds and new problems. In *Proc. 31st. Intern. Symp. Math. Found. Comput. Sci. (MFCS)*, volume 4162 of *LNCS*, pp. 146–156. Springer, 2006. [doi:10.1007/11821069_13] 203, 221

[10] RICHARD BEIGEL AND JUN TARUI: On ACC. *Comput. Complexity*, 4(4):350–366, 1994. [doi:10.1007/BF01263423] 202

[11] J. LAWRENCE CARTER AND MARK N. WEGMAN: Universal classes of hash functions. *J. Comput. System Sci.*, 18(2):143–154, 1979. [doi:10.1016/0022-0000(79)90044-8] 211

[12] ASHOK K. CHANDRA, MERRICK L. FURST, AND RICHARD J. LIPTON: Multi-party protocols. In *Proc. 15th STOC*, pp. 94–99. ACM Press, 1983. [doi:10.1145/800061.808737] 202, 203, 204

[13] ARKADEV CHATTOPADHYAY AND ANIL ADA: Multiparty communication complexity of disjointness. Technical Report TR08-002, Electron. Colloq. on Comput. Complexity (ECCC), 2008. 203

[14] MATEI DAVID, TONIANN PITASSI, AND EMANUELE VIOLA: Improved separations between nondeterministic and randomized multiparty communication. *ACM Trans. Comput. Log.*, 1(2):1–20, 2009. [doi:10.1145/1595391.1595392] 203, 221

[15] DMITRY GAVINSKY AND ALEXANDER A. SHERSTOV: A separation of NP and coNP in multiparty communication complexity. *Theory of Computing*, 6:227–245, 2010. [doi:10.4086/toc2010.v006a010] 203

[16] JOHAN HÅSTAD AND MIKAEL GOLDMANN: On the power of small-depth threshold circuits. *Comput. Complexity*, 1(2):113–129, 1991. [doi:10.1007/BF01272517] 202

[17] HARTMUT KLAUCK: Lower bounds for quantum communication complexity. *SIAM J. Comput.*, 37(1):20–46, 2007. [doi:10.1137/S0097539702405620] 220

[18] EYAL KUSHILEVITZ AND NOAM NISAN: *Communication Complexity*. Cambridge University Press, 1997. 202, 206, 208

[19] TROY LEE AND ADI SHRAIBMAN: Disjointness is hard in the multi-party number-on-the-forehead model. In *Proc. 23rd Comput. Complex. Conf. (CCC)*, pp. 81–91. IEEE Comput. Soc. Press, 2008. [doi:10.1109/CCC.2008.29] 203

[20] YISHAY MANSOUR, NOAM NISAN, AND PRASOON TIWARI: The computational complexity of universal hashing. *Theoret. Comput. Sci.*, 107(1):121–133, 1993. [doi:10.1016/0304-3975(93)90257-T] 211

[21] ILAN NEWMAN: Private vs. common random bits in communication complexity. *Inform. Process. Lett.*, 39(2):67–71, 1991. [doi:10.1016/0020-0190(91)90157-D] 207, 208

[22] NOAM NISAN: The communication complexity of threshold gates. In *Combinatorics, Paul Erdős is Eighty*, number 1 in Bolyai Society Mathematical Studies, pp. 301–315. J. Bolyai Math. Soc., Budapest, 1993. 202

[23] NOAM NISAN AND AVI WIGDERSON: Rounds in communication complexity revisited. *SIAM J. Comput.*, 22(1):211–219, 1993. [doi:10.1137/0222016] 202

[24] A. A. RAZBOROV: Quantum communication complexity of symmetric predicates. *Izv. Math.*, 67(1):145–159, 2003. [doi:10.1070/IM2003v067n01ABEH000422] 220

[25] ALEXANDER A. SHERSTOV: Communication lower bounds using dual polynomials. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS*, 95:59–93, 2008. 203, 220

[26] ANDREW CHI-CHIH YAO: Some complexity questions related to distributive computing (preliminary report). In *Proc. 11th STOC*, pp. 209–213. ACM Press, 1979. [doi:10.1145/800135.804414] 202

[27] ANDREW CHI-CHIH YAO: On ACC and threshold circuits. In *Proc. 31st FOCS*, volume 2, pp. 619–627, 1990. [doi:10.1109/FSCS.1990.89583] 202

AUTHORS

Paul Beame
professor
University of Washington, Seattle, WA
beame@cs.washington.edu
http://www.cs.washington.edu/homes/beame/

Matei David
postdoctoral research fellow
Princeton University, Princeton, NJ
mateid@cs.princeton.edu
http://www.cs.toronto.edu/~matei/

Toniann Pitassi
professor
University of Toronto, Toronto, ON
toni@cs.toronto.edu
http://www.cs.toronto.edu/~toni/

Philipp Woelfel
assistant professor
University of Calgary, Calgary, AB
woelfel@cpsc.ucalgary.ca
http://pages.cpsc.ucalgary.ca/~woelfel/

ABOUT THE AUTHORS

PAUL BEAME is a Professor of Computer Science & Engineering at the University of Washington. He received his Ph. D. in Computer Science from the University of Toronto in 1987 under the supervision of Stephen A. Cook. He is currently Chair of the IEEE CS Technical Committee on the Mathematical Foundations of Computing. His research has primarily focused on the complexity of concrete computational problems and proof complexity, with a particular emphasis on complexity lower bounds.

MATEI DAVID recently graduated from the University of Toronto; his advisor was fellow coauthor Toniann Pitassi. In writing this, Matei realizes the relativity of the term "recently."

TONIANN PITASSI is a professor at the University of Toronto, who still feels lucky to get paid to do this. Hobbies include sculpting (where she is enthusiastic, but lacking in talent), running (same skill set), and spending quality family time, which currently means YouTube videos or Rummikub.

PHILIPP WOELFEL graduated in December 2003 from the University Dortmund under the supervision of Ingo Wegener. In 2005 the German Research Foundation admitted him to the Emmy-Noether Programme, which allowed him to meet two of his coauthors during his Postdoctoral Fellowship (2005-2007) at the University of Toronto. Currently, he is an Assistant Professor at the University of Calgary. His research interests include computational complexity, randomized algorithms, and distributed computing.