Separating Deterministic from Nondeterministic NOF Multiparty Communication Complexity
(Extended Abstract)

Paul Beame[1,⋆], Matei David[2,⋆⋆], Toniann Pitassi[2,⋆⋆⋆], and Philipp Woelfel[2,†]

[1] University of Washington
[2] University of Toronto

Abstract. We solve some fundamental problems in the number-on-forehead (NOF) k-party communication model. We show that there exists a function which has at most logarithmic communication complexity for randomized protocols with a one-sided error probability of 1/3 but which has linear communication complexity for deterministic protocols. The result is true for $k = n^{O(1)}$ players, where n is the number of bits on each players' forehead. This separates the analogues of RP and P in the NOF communication model. We also show that 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 and we do not know of any explicit function which allows such separations. However, for the 3-player case we exhibit an explicit function which has $\Omega(\log\log n)$ randomized complexity for private coins but only constant complexity for public coins.
It follows from our existential result that any function that is complete for the class of functions with polylogarithmic nondeterministic k-party 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 $f : X_1 \times \cdots \times X_k \to O$ when its input is distributed between k computationally unbounded players was first introduced in [17] and it has since been shown to have many diverse applications in complexity theory. The case of $k = 2$ players has been studied extensively [11]. For two or more players, we are interested in the "number-on-forehead" model (NOF), first introduced by

⋆ Supported by NSF grant CCR-0514870
⋆⋆ Supported by OGS
⋆⋆⋆ Supported by NSERC
† Supported by DFG grant Wo 1232/1-1

Chandra, Furst and Lipton in [7]. 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 number-on-forehead 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. Cylinder intersections appear difficult to understand combinatorially.

Lower bounds for multiparty complexity in the number-on-forehead model are connected to a major open problem in complexity theory: it has been 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 [6, 10]. The best lower bound obtained so far establishes a lower bound of $\Omega(n/2^k)$, which breaks down when the number of players is greater than logarithmic [3, 8, 16, 9]. 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 tradeoffs [3], circuit complexity bounds [10, 15, 14], and proof complexity bounds [4].

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-party 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? Surprisingly, the relationships between these complexity measures have not been resolved previously, even for $k = 3$.

Our main result is that for any k that is $n^{O(1)}$ there is a function with n bits on each players' forehead that is computable with a polylogarithmic complexity by a randomized k-party communication protocol with 1-sided error but which requires linear complexity for deterministic protocols. We obtain this result nonconstructively by showing that deterministic protocols for a certain class of *simple* functions have a nice normal form and then establishing a lower bound for such function via a counting argument over protocols in normal form. We thus separate the randomized 1-sided error and deterministic k-party NOF communication complexity classes RP_k^{cc} and P_k^{cc}. As a corollary of our lower bounds, we also establish an optimal separation between the public and private coin randomized NOF models.

These bounds are nonconstructive but, for k at most logarithmic in the input size, we can also give *explicit* families of simple functions with $\Omega(\log n)$ deterministic k-party complexity in the NOF model. (We believe that they have superpolylogarithmic deterministic complexity.) The best previous lower bound for any explicitly defined simple function is the $\Omega(\log \log n)$ lower bound from

[5] for the Exact-T function (originally investigated in [7]) in the special case of $k = 3$ players. As a corollary of our bound we obtain that our function families have $\Omega(\log \log n)$ complexity for randomized private coin protocols (with constant error probability) but only $O(1)$ complexity for public coin protocols.

The problem of separating deterministic from nondeterministic NOF complexity is particularly interesting because of its connection to proof complexity. In recent work [4], it has been shown that for $k = 3$, $(\log n)^{\omega(1)}$ lower bounds on the randomized NOF complexity of set intersection, which has nondeterministic NOF complexity $O(\log n)$, implies lower bounds for polynomial threshold proof systems, such as the Lovász-Schrijver proof systems, as well as the Chvátal cutting planes proof system. Moreover, it seems possible that these results can be modified to show that randomized lower bounds for *any* function with small NOF nondeterministic communication complexity would give lower bounds for related cutting planes proof systems.

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 and in [4] it had been assumed that the same holds for the number-on-forehead (NOF) model. (The number-in-hand 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.) However, we prove that under communication-free reductions, set intersection is not complete in the NOF model.

2 Definitions and Preliminaries

In the NOF multiparty communication complexity game [7] there are k parties (or players), numbered 1 to k, that are trying to collaborate to compute a function $f : X_1 \times \ldots \times X_k \to \{0, 1\}$ where each $X_i = \{0, 1\}^n$. The kn input bits are partitioned into k sets, each of size n. For $(x_1, \ldots, x_k) \in \{0, 1\}^{kn}$, and for each i, player i knows the values of all of the inputs except for x_i (which conceptually is thought of as being placed on player i's forehead).

The players exchange bits according to an agreed-upon protocol, by writing them on a public blackboard. A *protocol* specifies, for every possible blackboard contents, whether or not the communication is over, the output if over and the next player to speak if not. A protocol also specifies what each player writes 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. A *private-coin randomized* protocol is a protocol in which each player has access to a private random string. A *nondeter-*

ministic protocol is a randomized private coin protocol with 1-sided error (only false negatives) and an error probability less than 1.

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

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 in [2], who proved new lower bounds, as well as surprising upper bounds in this model.

Since any function f_n on kn bits can be computed using only n bits of communication, following [1], for sequences of functions $f = (f_n)_{n\in\mathbb{N}}$, algorithms are considered "efficient" or "polynomial" if only polylogarithmically many bits are exchanged. Accordingly, let P^{cc}_k denote the class of function families f for which $D_k(f_n)$ is $(\log n)^{O(1)}$, let NP^{cc}_k denote the class of function families f with nondeterministic complexity $(\log n)^{O(1)}$, and let RP^{cc}_k denote the class of function families f for which $R^1_k(f_n)$ is $(\log n)^{O(1)}$.

Multiparty communication complexity lower bounds are proven by analyzing properties of functions on *cylinder intersections*.

Definition 1. *An* i-cylinder C_i *in* $X_1 \times \ldots \times X_k$ *is a set such that for all* $x_1 \in X_1, \ldots, x_k \in X_k, x'_i \in X_i$ *we have* $(x_1, \ldots, x_i, \ldots, x_k) \in C_i$ *if and only if* $(x_1, \ldots, x'_i, \ldots, x_k) \in C_i$. *A* cylinder intersection *is a set of the form* $\bigcap_{i=1}^k C_i$ *where each* C_i *is an* i*-cylinder in* $X_1 \times \cdots \times X_k$.

3 Separating P^{cc}_k from RP^{cc}_k

3.1 Oblivious Players, Simple Functions, and a Normal Form

We will be interested in a special type of Boolean functions for which we can show, that without loss of generality, all but one of the players is oblivious. For

sets $X_1, \ldots, X_k$ a function $f : X_1 \times \cdots \times X_k \to \{0,1\}$ is *simple for player i* if for all $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots x_k) \in X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_k$ there exists at most one $x_i^* \in X_i$ such that $f(x_1, \ldots, x_{i-1}, x_i^*, x_{i+1}, \ldots, x_k) = 1$.

If f is simple for player i then it is reducible with no communication to 2-player n-bit equality EQ. Player i can compute the unique value for the input on its forehead for which the output could be 1 (if it exists), and any other player sees that input. All the players have to do is to decide whether these strings are equal. We know that $R^0_{2,1/n}(\mathrm{EQ})$ is $O(\log n)$ and $R_2^{0,\mathrm{pub}}(\mathrm{EQ})$ is $O(1)$. Therefore we get the following.

Lemma 2. *For all k and all simple functions f on kn bits, $R^0_{k,1/n}(f)$ is $O(\log n)$ and $R_k^{0,pub}(f)$ is $O(1)$.*

The following theorem shows that if a function is simple for one player then this player can act obliviously with only a small increase in the deterministic communication complexity.

Theorem 3. *Let $f : X_1 \times \cdots \times X_k \to \{0,1\}$ be a function that is simple for player i and has $D_k(f) = d$. Then there is a protocol P' for f in which player i first sends d bits and then all players $j \in \{1, \ldots, k\} - \{i\}$ simultaneously send exactly one bit b_j such that $f(x_1, \ldots, x_k) = 1$ if and only if all bits $b_j = 1$.*

Proof (Sketch). Let f be simple for player 1. Let P be a protocol for f with complexity d. We describe protocol P' on input $(x_1, \ldots, x_k)$. Assume that player 1 sees the partial input $(x_2, \ldots, x_k)$ on the other players' foreheads. Let x_1^* be the input in X_1 such that $f(x_1^*, x_2 \ldots, x_k) = 1$, if one exists, an arbitrary input in X_1, otherwise. Player 1 "simulates" protocol P for the input $(x_1^*, x_2, \ldots, x_k)$; i.e., she writes on the blackboard exactly the string I^* that would have been written by players $1, \ldots, k$ if protocol P were executed for that input. Then each player r, $2 \le r \le k$, verifies that I^* is consistent with what player r would have sent in protocol P if it had seen $(x_1, \ldots, x_{r-1}, x_{r+1}, \ldots, x_k)$ on the other players' foreheads. If player r does not find an error and the output of P is 1 for blackboard contents I^*, then he accepts by sending bit $b_r = 1$. Otherwise he sends $b_r = 0$. □

3.2 Representing Simple Functions by Colorings and Cylinder Intersections

Most lower bound proofs for $D_k(f)$ use the fact shown in [3] that any k-party protocol with complexity d for a function f yields a partitioning of the input into $O(2^d)$ disjoint cylinder intersections on which f is constant. For $k \ge 3$ players, the known techniques for proving lower bounds on the number of cylinder intersections needed for such a partitioning are discrepancy-based and inherently yield lower bounds even for nondeterministic and randomized protocols. Therefore, these techniques are not suitable for proving good lower bounds for functions with low nondeterministic communication complexity.

For simple functions we obtain different, although related, structures. These structures seem to be better suited for lower bound proofs for functions in RP_k^{cc}, as they will allow us to separate this class from P_k^{cc} and to prove $\Omega(\log n)$ lower bounds for explicit functions.

Throughout this section, $f : X_1 \times \cdots \times X_k \to \{0,1\}$ is simple for player 1. For any natural number D and a set S, a *D-coloring* of S is a mapping $c : S \to [D]$. Since f is simple for player 1 (Alice), there exists a function $g : X_2 \times \cdots \times X_k \to X_1 \cup \{\perp\}$, where $g(x_2, \ldots, x_k) = \perp$ if $f(x_1, \ldots, x_k) = 0$ for all $x_1 \in X_1$, and otherwise $g(x_2, \ldots, x_k) = x_1^*$, where x_1^* is the the unique element in X_1 with $f(x_1^*, x_2, \ldots, x_k) = 1$. In fact, any such mapping g uniquely defines the simple function f.

Assume that f can be computed by a d-bit protocol P. The special protocol P' for f, derived in Theorem 3, can be characterized by a coloring of $X_2 \times \cdots \times X_k$ and cylinder intersections in $X_2 \times \cdots \times X_k$: 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 if she sees $(x_2, \ldots, x_k)$. Consider a fixed message m 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 m is a cylinder intersection $I_{m,a}$. Each such cylinder intersection $I_{m,a}$ may also contain points that are not colored m. However, it is not possible that a point $p = (x_2, \ldots, x_k) \in I_{m,a}$ has color m but $g(p) \neq a$ because then Alice would send message m 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. We obtain the following.

Lemma 4. *Every function f that is simple for player 1 and has k-player communication complexity d can be uniquely represented by cylinder intersections $I_{m,a} \in X_2 \times \cdots \times X_k$, $m \in [2^d], a \in X_1$, and a 2^d-coloring c of $X_2 \times \cdots \times X_k$, such that $\forall a \in X_1, y \in X_2 \times \cdots \times X_k$: $f(a,y) = 1 \Leftrightarrow y \in I_{c(y),a}$.*
In particular, $I_{m,a}$ contains all points $y \in X_2 \times \cdots \times X_k$ with color $c(y) = m$ and $f(a,y) = 1$, but no point y' with color $c(y') = m$ and $f(a,y') = 0$.

Proof. We have already seen how to obtain c and the cylinder intersections $I_{m,a}$ from the function f. This representation is unique because for any input (a,p) with $a \in X_1$ and $p \in X_2 \times \cdots \times X_k$ we can retrieve the function value $f(a,p)$ by checking whether $p \in I_{c(p),a}$. □

3.3 The Lower Bound

In the following we consider a family of functions which have logarithmic communication complexity for randomized protocols with one-sided error and error probability bounded by 1/3. Using Lemma 4 we give an upper bound on the number of different deterministic protocols for the functions in that class in order to show that at least one such function requires at least linear deterministic communication complexity.

For positive integers n, m and t, let $G_{t,n,m}$ be the set of all mappings $g : \{0,1\}^{n \cdot t} \to \{0,1\}^m$. For any function $g \in G_{k-1,n,m}$, define $f_g : \{0,1\}^m \times$

$\{0,1\}^{n\cdot(k-1)}$ by $f_g(x_1,\ldots,x_k)=1$ if and only if $g(x_2,\ldots,x_k)=x_1$. By the proof of Lemma 2, randomized protocols for functions f_g, $g\in G_{k,n,m}$, have complexity at most $O(\log m)$. Hence, it follows that $f_g\in \mathsf{co\text{-}RP}_k^{cc}$ for all $g\in G_{k,n,n/2}$.

Theorem 5. *There is a $g\in G_{k-1,n,n/2}$ such that $D_k(f_g)$ is $\Omega(n-\log k)$.*

Corollary 6. *$\mathsf{P}_k^{cc}\neq\mathsf{RP}_k^{cc}$ for any k that is $n^{O(1)}$.*

Proof (of Theorem 5). Any function $g\in G_{k-1,n,m}$ has a domain of size $2^{(k-1)n}$ and a range of size 2^m. Therefore, it is not possible to encode every such function g with less than $m\cdot 2^{(k-1)n}$ bits. Note that if two functions g,g' are different, then f_g and $f_{g'}$ are different, too.

Clearly, any function f_g, $g\in G_{k-1,n,m}$, is simple for Alice. Assume that any such function f_g has $D_k(f_g)\le d$. Then by Lemma 4, every such function f_g can be uniquely represented by a 2^d-coloring of $\left(\{0,1\}^n\right)^{k-1}$ and $2^m\cdot 2^d$ cylinder intersections in $\left(\{0,1\}^n\right)^{k-1}$. The 2^d-coloring of $\left(\{0,1\}^n\right)^{k-1}$ can be encoded with $d\cdot 2^{(k-1)n}$ bits. The number of i-cylinders in $X_1\times\cdots\times X_t$ is $2^{\Pi_{j\neq i}|X_j|}$. Hence, $(k-1)\cdot 2^{(k-2)n}$ bits suffice for a unique encoding of any cylinder intersection in $\left(\{0,1\}^n\right)^{k-1}$. Thus, the total number of bits in which any function f_g, $g\in G_{k-1,n,m}$, can be encoded is bounded above by

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

As we have seen above, the number of bits needed to describe a function f_g for $g\in G_{k-1,n,m}$ is at least $m\cdot 2^{(k-1)n}$. Therefore, if for all f_g a protocol with complexity c exists then

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

This is equivalent to $2^d\ge 2^{n-m}\cdot(m-d)/(k-1)$. Hence, $d\ge \min\{m-1,\ n-m-\log(k-1)\}$, which for $m=\lfloor(n-\log k)/2\rfloor$ is at least $(n-\log k)/2-O(1)$. □

3.4 Separating Public from Private Coins

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, $R_k^{\mathrm{pub}}(f)\le R_k(f)$. In the other direction, Newman [13, 11] 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 works for any number of players.)

Proposition 7 ([13]). *There is a $c>0$ such that for every $k\ge 2$ and function $f:\{0,1\}^{kn}\to\{0,1\}$, $R_k(f)\le R_k^{pub}(f)+c\lceil\log_2 n\rceil$.*

We see that the maximum gap between the public-coin and private-coin randomized complexities of f is $\Theta(\log n)$, and it is achieved when $R_k^{\mathrm{pub}}(f)$ is $O(1)$ and $R_k(f)$ is $\Theta(\log n)$. The natural question arises, is there a function that achieves this gap? Our results allow us to answer this question affirmatively.

In order to obtain lower bounds, we need the following extension of Lemma 3.8 in [11] to k players. We omit the proof due to space constraints.

Lemma 8. *If $k^{1/\delta} < D_k(f)$ for some $\delta < 1$, then $R_k(f)$ is $\Omega(\log D_k(f))$.*

Corollary 9. *Let $\delta < 1$. For all k such that $k < n^\delta$, there exists a kn-bit function f such that $R_k^{pub}(f)$ is $O(1)$ and $R_k(f)$ is $\Theta(\log n)$.*

Proof. By Theorem 5 there is a function f that is simple for player 1 such that $D_k(f)$ is $\Omega(n)$. By Lemma 8, $R_k(f)$ is $\Omega(\log n)$. By Lemma 2, $R_k(f)$ is $O(\log n)$ and $R_k^{\text{pub}}(f)$ is $O(1)$. □

4 Lower Bounds for Explicit Simple Functions

The separations in Section 3.3 are nonconstructive. We conjecture that there also exists an *explicit* simple function that gives a linear or near linear separation between $D_k(f_n)$ and $R_k^0(f_n)$. In the following we prove $\Omega(\log n)$ bounds for the deterministic complexity of some explicit simple functions. This yields a separation between the deterministic and public coin randomized complexity for explicit simple functions, though this is much weaker than our conjecture.

We give two constructions of explicit functions, one for $k = 3$ and, one that holds for all $k \geq 3$. Write $\mathbb{F}_q$ for the field of q elements. Let $X = \mathbb{F}_{2^n}$, $Y = \mathbb{F}_{2^m}$ for some positive integers m and n. For $(a, b) \in X \times Y$ let the hash function $h_{a,b} : \mathbb{F}_{2^n} \to \mathbb{F}_{2^m}$ be defined as $h_{a,b}(x) = \phi(a \cdot x) + b$, where ϕ is a homomorphism from $\mathbb{F}_{2^n}$ to $\mathbb{F}_{2^m}$. Let H be the family of hash functions $h_{a,b}$ for $a \in X$, $b \in Y$. The explicit function for $k = 3$ is given by $f : Y \times X \times H \to \{0, 1\}$, where $f(y, x, h) = 1$ iff $h(x) = y$ which is clearly simple for player 1.

For the other construction let $k \geq 3$, let $X_1 = \mathbb{F}_{2^m}$ and $X_2 = \ldots = X_k = \mathbb{F}_{2^m}^n$ for positive integers n and m with $m \geq \log_2 n$. Let $\beta_1, \ldots, \beta_n$ be distinct elements of $\mathbb{F}_{2^m}$ and define $v_i = (\beta_1^{i-1}, \ldots, \beta_n^{i-1})$ for $1 \leq i \leq n$. The explicit function is f_g where $g(x_2, \ldots, x_k) = \sum_{i=1}^n \prod_{j=2}^k \langle v_i, x_j \rangle$ and operations are over $\mathbb{F}_{2^m}$.

Theorem 10. *There is a $\delta < 1$ such that*
(a) for $m = n^\delta$ and $f : Y \times X \times H \to \{0, 1\}$ defined as above, $D_3(f)$ is $\Omega(\log n)$,
(b) for any $k \geq 3$, $n \geq 4^{k+1}$, $m = n^\delta$, and f_g defined as above, $D_k(f_g)$ is $\Omega(\log n)$.

Proof. We give the proof for part (a): Let $d = D_3(f)$. Fix some ϵ with $0 < \epsilon < \delta$ and let $m = n^\delta$. Assume for contradiction that $d \leq \epsilon \cdot \log_2 n$.

Consider the $2^{n+m} \times 2^n$ matrix M where rows correspond to hash functions $h \in H$, columns correspond to inputs $x \in X$ and the entry $M_{h,x}$ is the hash function value $h(x)$. Cylinder intersections in $H \times X$ are rectangles. By Lemma 4, there is a 2^d-coloring c of M and there are 2^{d+m} rectangles $R_{\ell,y}$, $\ell \in [2^d]$, $y \in Y$, such that $\forall (y, x, h) \in Y \times X \times H :\ (h, x) \in R_{c(h,x),y} \Leftrightarrow h(x) = y$. Call an entry $M_{h,x}$ an (ℓ, y) *entry* iff $c(h, x) = \ell$ and $h(x) = y$. Correctness of the protocol implies that for $y' \neq y$, $R_{\ell,y}$ does not contain any (ℓ, y') entries.

Consider the coloring c of the matrix entries in M. The proof proceeds inductively decreasing the number of colors that are available and shrinking the matrix. During each such 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. This will contradict the existence of the initial coloring, hence of the d-bit protocol.

We prove by induction on $i \geq 0$ that, as long as $i \leq 2^d = n^\epsilon$, the following hold for large enough n:
– there exists a rectangle R_i such that $|R_i| \geq 2^{2n+m-i(d+2)}$,
– R_i contains at most $i \cdot 2^{2n}$ holes,
– non-hole entries in R_i can be colored with $2^d - i$ colors.

Assuming that we have established this inductive statement, letting $i = 2^d$ we see that there are no colors left for coloring the rectangle R_{2^d}. Moreover, for large enough n, this rectangle has size at least $2^{2n+m-n^\epsilon(2+\epsilon\cdot\log_2 n)}$. Since $m = n^\delta$ and $\delta > \epsilon$, $|R_{2^d}| > 2^{2n+d}$. The number of holes in this rectangle is bounded above by $2^d \cdot 2^{2n}$, so the rectangle is not empty, which is a contradiction.

We now prove the inductive statement. For $i = 0$, let $R_0 = M$. The existence of a d-bit protocol yields a coloring of R_0 (with no holes) using 2^d colors.

Now assume the inductive statement is true for some $0 \leq i < 2^d$. The number of non-hole entries in R_i is at least $2^{2n+m} \cdot \left(2^{-i(d+2)} - i \cdot 2^{-m}\right)$. Since $i < 2^d$ and $m - d > n^\epsilon(d+2) > i(d+2)$ (for large enough n), the number of non-hole entries in R_i is larger than $2^{2n+m-i(d+2)-1}$. Let (ℓ, y) be the most popular color-value pair from the non-hole entries in R_i and let $R_{i+1} = R_i \cap R_{\ell,y}$. The number of color-value pairs is at most 2^{d+m}, so the number of occurrences of the most popular pair (ℓ, y) in R_i is at least $2^{2n+m-i(d+2)-1-(m+d)} = 2^{2n-(i+1)(d+2)+1}$. By construction, $|R_{i+1}|$ is at least the number of such (ℓ, y) entries. Since R_{i+1} is a rectangle, by the Hash Mixing Lemma [12], for any $y \in Y$,

$$\Pr[h(x) = y] \leq \tfrac{1}{|Y|} + \sqrt{\tfrac{|H|}{|R_{i+1}|\cdot|Y|}} \leq 2^{-m} + 2^{((i+1)(d+2)-n-1)/2} \leq 2^{-m+1}$$

since $|H|/|R_{i+1}||Y| \leq 2^{n-2n+(i+1)(d+2)-1}$, $(i+1)(d+2) \leq n^\epsilon \log n < n^\delta = m$ and $n \geq 3m$ for sufficiently large n. Hence, the number of y-valued entries in R_{i+1} is at most $|R_{i+1}| \cdot 2^{-m+1}$. By the lower bound from above for the number of (ℓ, y)-pairs in R_{i+1}, we have $2^{2n-(i+1)(d+2)+1} \leq |R_{i+1}| \cdot 2^{-m+1}$ and thus we obtain the stronger bound $|R_{i+1}| \geq 2^{2n+m-(i+1)(d+2)}$ as required.

Since $R_{i+1} \subseteq R_{\ell,y}$, by Lemma 4 all ℓ-colored entries in R_{i+1} are (ℓ, y) entries. Define the holes in R_{i+1} to be its (ℓ, y) entries along with all holes in R_i. Thus, the number of colors available for non-hole entries has been reduced by at least 1. The number of extra holes we introduce is at most the number of entries in M with value y. Hence, at most 2^{2n} new holes can be introduced in a round. This completes the inductive step, and therefore the proof of (a).

The proof for part (b) is similar but requires the following property of g which is a natural analogue of the Hash Mixing Lemma [12] over cylinder intersections. Its proof is in the full paper.

Lemma 11. *Let $Z = \mathbb{F}_{2^m}^n$. For $k \geq 3$ and any cylinder intersection $I \subseteq Z^{k-1}$ and any $y \in \mathbb{F}_{2^m}$, choosing $(x_2, \ldots, x_k)$ from the uniform distribution on Z^{k-1}, $\left|\Pr[g(x_2, \ldots, x_k) = y \text{ and } (x_2, \ldots, x_k) \in I] - 2^{-m}|I|/|Z|^{k-1}\right| \leq 2^{-(m-2)n/4^{k-1}}$.* □

By Lemma 2, both f and f_g defined above have $O(1)$ public-coin randomized complexity but by Theorem 10 and Lemma 8, we obtain that $R_3(f)$ and $R_k(f_g)$ are both $\Omega(\log\log n)$. In fact, we conjecture that the $D_3(f)$ and $D_k(f_g)$ are both $\omega(\log n)$ or even $n^{\Omega(1)}$. Proving the latter would yield explicit examples of function in RP_3^{cc} but not in P_3^{cc}.

5 On Complete Problems for NP_k^{cc}

An alternative approach to separating P_k^{cc} from RP_k^{cc} 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 RP_k^{cc}, then by our separation result we can conclude that it is not in P_k^{cc}. Proving a lower bound for a function complete for NP_k^{cc} has the added benefit of potentially separating NP_k^{cc} from RP_k^{cc} as well. (Recall that simple functions are in RP_k^{cc} so they cannot separate NP_k^{cc} from RP_k^{cc}.) The set intersection function is complete for the class analogous to NP_k^{cc} in the number-in-hand (NIH) model, and thus also for NP_2^{cc}. In this section, we prove that this function is not complete for NP_k^{cc} for $k \geq 3$.

For sets $X_1, \ldots, X_k$, write $\overline{X}$ for $X_1 \times \cdots \times X_k$. We write $\overline{x} \in \overline{X}$ to denote a k-tuple $(x_1, \ldots, x_k)$ where $x_i \in X_i$ for all $i \in [k]$. Use $\overline{\varphi}$ to denote a k-tuple of functions $\varphi_1, \ldots, \varphi_k$. Furthermore, for $i \in [k]$, write $\overline{\alpha}_{-i}$ for the $(k-1)$-tuple obtained from $\overline{\alpha}$ by removing the i-th coordinate.

In two-party 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.

Definition 12. *For $k = 2$, let $f : \overline{X} \to \{0,1\}$ and $g : \overline{X'} \to \{0,1\}$. A pair of functions $\overline{\varphi}$ with $\varphi_i : X_i \to X_i'$ is a* rectangular reduction *of f to g, written $f \sqsubseteq g$, if and only if $f(x_1, x_2) = g(\varphi_1(x_1), \varphi_2(x_2))$.*

Furthermore, they defined an appropriate 'polynomially-bounded' version of rectangular reduction for function families.

Definition 13. *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 and only 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 14 ([1]). *Let f and g be function families. If $f \sqsubseteq_p g$ and $g \in \mathsf{P}_2^{cc}$ then $f \in \mathsf{P}_2^{cc}$. If $f \sqsubseteq_p g$ and $g \in \mathsf{NP}_2^{cc}$ then $f \in \mathsf{NP}_2^{cc}$.*

Definition 15. *A function family g is* complete for NP_2^{cc} under rectangular reductions *if and only if $g \in \mathsf{NP}_2^{cc}$ and for all $f \in \mathsf{NP}_2^{cc}$, $f \sqsubseteq_p g$.*

The *set intersection function* is $\textsc{Disj}_{k,n} : (\{0,1\}^n)^k \to \{0,1\}$ defined by $\textsc{Disj}_{k,n}(\overline{x}) = 1$ if and only if there is some $i \in [n]$ such that $x_{1,i} = \ldots = x_{k,i} = 1$. Clearly, $\textsc{Disj}_k \in \mathsf{NP}_k^{cc}$. Babai, Frankl and Simon observed the following:

Proposition 16 ([1]). $\textsc{Disj}_2$ *is complete for NP_2^{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 16 continues to hold in the NIH model under cubic reductions. One might conjecture that $\textsc{Disj}_k$ is also complete for NP_k^{cc} under a natural extension of rectangular reductions in the NOF model. Such a notion of reduction should not require any communication between the parties. This yields the following definition:

Definition 17. *Given $f : \overline{X} \to \{0,1\}$ and $g : \overline{X'} \to \{0,1\}$ we say that functions $\overline{\varphi}$ are a* cylindrical reduction *of f to g if and only if for every $\overline{x} \in \overline{X}$ there is an $\overline{x'} \in \overline{X'}$ such that for all $i \in [k]$, $\varphi_i(\overline{x}_{-i}) = \overline{x'}_{-i}$ and $f(\overline{x}) = g(\overline{x'})$. Thus each φ_i maps the NOF view of the i-th player on input $\overline{x}$ for f to the NOF view of the i-th player on input $\overline{x'}$ for g.*

We show that cylindrical reductions must be of a special form, given by the natural no-communication reductions associated with the number-in-hand model. $\overline{A} = A_1 \times \cdots \times A_k$ is a *cube*, if $A_i \subseteq X_i$ for all $i \in [k]$.

Lemma 18. *If there is a cylindrical reduction of $f : \overline{X} \to \{0,1\}$ to $\textsc{Disj}_{k,m}$ then $f^{-1}(1)$ is a union of m cubes.*

□

Theorem 19. *There is a function $f : \{0,1\}^{3n} \to \{0,1\}$ with deterministic 3-party NOF communication complexity at most 3 such that any cylindrical reduction of f to $\textsc{Disj}_{3,m}$ requires $m > 2^{n-3}$.*

Proof. For $x, y, z \in \{0,1\}^n$, define $f(x,y,z)$ to be 1 if and only if x, y, and z are pairwise orthogonal in $\mathbb{F}_2^n$. There is a trivial 3-party NOF protocol for f in which 3 bits are exchanged, namely, each party 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 $h(u,v) = 1$ iff $\langle x, y\rangle = 0$ 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 $1/2 + 1/2^{n+1} > 1/2$ (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 $\dim(\langle A_1\rangle) + \dim(\langle A_2\rangle) \leq n$. Therefore $|A_1 \times A_2| \leq |\langle A_1\rangle \times \langle A_2\rangle| \leq 2^{\dim(\langle A_1\rangle)+\dim(\langle A_2\rangle)} \leq 2^n$ so $|C| \leq 2^n |A_1 \times A_2| \leq 2^{2n}$ 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 NP_3^{cc} under efficient reductions that do not require communication.

References

[1] L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory (preliminary version). In *Proc. of 27th FOCS*, pp. 337–347. 1986.

[2] L. Babai, A. Gál, P. G. Kimmel, and S. V. Lokam. Communication complexity of simultaneous messages. *SIAM J. on Comp.*, 33:137–166, 2004.

[3] L. Babai, N. Nisan, and M. Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. of Comp. and Syst. Sci.*, 45:204–232, 1992.

[4] P. Beame, T. Pitassi, and N. Segerlind. Lower bounds for lovász-schrijver systems and beyond follow from multiparty communication complexity. In *Proc. of 32nd ICALP*, pp. 1176–1188. 2005.

[5] R. Beigel, W. Gasarch, and J. Glenn. The multiparty communication complexity of Exact-T: Improved bounds and new problems. In *Proc. of 31st MFCS*, pp. 146–156. 2006.

[6] R. Beigel and J. Tarui. On ACC. In *Proc. of 32nd FOCS*, pp. 783–792. 1991.

[7] A. K. Chandra, M. L. Furst, and R. J. Lipton. Multi-party protocols. In *Proc. of 15th ACM STOC*, pp. 94–99. 1983.

[8] F. R. K. Chung and P. Tetali. Communication complexity and quasi randomness. *SIAM J. Discrete Math.*, 6:110–125, 1993.

[9] J. Ford and A. Gál. Hadamard tensors and lower bounds on multiparty communication complexity. In *Proc. of 32nd ICALP*, pp. 1163–1175. 2005.

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

[11] E. Kushilevitz and N. Nisan. *Communication complexity.* Cambridge University Press, New York, NY, USA, 1997.

[12] Y. Mansour, N. Nisan, and P. Tiwari. The computational complexity of universal hashing. *Theor. Comp. Sci.*, 107:121–133, 1993.

[13] I. Newman. Private vs. common random bits in communication complexity. *IPL*, 39:67–71, 1991.

[14] N. Nisan. The communication complexity of threshold gates. In *Proceedings of "Combinatorics, Paul Erdos is Eighty"*, pp. 301–315. 1993.

[15] N. Nisan and A. Wigderson. Rounds in communication complexity revisited. *SIAM J. on Comp.*, 22:211–219, 1993.

[16] R. Raz. The BNS-Chung criterion for multi-party communication complexity. *Comp. Compl.*, 9:113–122, 2000.

[17] A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proc. of 11th ACM STOC*, pp. 209–213. 1979.