

There are no zero-hard problems in multiparty communication complexity

László Babai and Denis Pankratov
University of Chicago

DRAFT: 2015-04-29 7:00pm

Abstract

We study the feasibility of generalizing the notions of completeness and hardness in 2-party communication complexity (Babai, Frankl, Simon, FOCS 1986) to the “number-on-forehead” (NOF) model of k -party communication complexity with $k \geq 3$ (Chandra, Furst, Lipton, STOC 1983). Beame *et al.* (*Theory of Computing*, 2010) suggest that the k -party version of NP^{cc} may not have any complete families under the natural (cylindrical, polynomially bounded) reductions. We confirm this in a strong sense. Let us call a family $F = (f_n)$ of k -party communication problems (one problem for every input length n) *zero-hard* if all families of k -party communication problems solvable with zero communication (cylinders) admit polynomially bounded reductions to F . We show that there are no zero-hard families of k -party communication problems. Furthermore, we study reduction concepts that “allow communication.” We arrive at two definitions of reductions with t bits of *communication*, one of Karp-style, the other of Cook-style. Let us consider the class of problems solvable with t bits of communication. We show that no problem is hard for this class with respect to t -bit Karp-style reduction and $(t - 1)$ -bit Cook-style reduction. All the proofs are by simple counting arguments, much in the spirit of Beame *et al.*. The main purpose of the paper is clarification of the concepts.

1 Introduction

We write NP_k^{cc} to denote the analogue of NP in k -party communication complexity. Beame *et al.* point out that finding an explicit NP_k^{cc} -complete family of communication problems would turn lower bounds, obtained via counting, into lower bounds for explicit functions. They suggest [3, Section 5], however, that for $k \geq 3$ players, NP_k^{cc} -complete families may not exist. We confirm this in a strong form.

Throughout this paper we talk about the “number-on-forehead” (NOF) model of communication complexity, introduced by Chandra, Furst, and Lipton [4], unless the “number-in-hand” (NIH) model is specifically indicated. We

shall assume that the number of players is $k \geq 3$, unless the 2-player model is specifically discussed.

We use the notation and terminology of [3] with one slight difference in the concept of termination of the communication protocol (see the beginning of Section 3). Moreover, we shall be consistent in making a distinction between a “communication problem” (a finite object which refers to a communication function over a fixed finite domain) and a “family of communication problems,” consisting of one communication problem for each input length n .

Babai, Frankl, and Simon [1] introduced complexity classes of families of 2-party communication problems, analogous to Turing machine complexity classes such as P, NP, BPP, Σ_r , Π_r , #P; and defined “rectangular reductions” between families of 2-party communication problems. One of the results of [1] asserts that the “set intersection” problem (a family) is NP_2^{cc} -complete under “polynomially-bounded” (see definition in Section 3) rectangular reductions.

Beame *et al.* [3] generalize these complexity classes and the reduction concept to multiparty communication and observe that the NP_2^{cc} -completeness result of [1] extends to the NP_k^{cc} -completeness of k -party set-intersection in the NIH model, but not in the far more interesting NOF model. In fact they prove, for $k = 3$ players, that even certain functions of communication complexity at most 2 cannot be reduced to set-intersection without an exponential blowup in the input length. Based on this evidence and on further thoughts on functions with similarly low complexity, they suggest that “it seems unlikely that any function is complete for NP_3^{cc} under efficient reductions that do not require communication.”

Building on the insight of Beame *et al.*, we show that indeed, for $k \geq 3$, no family of k -party communication problems is *hard* even for a trivial subclass of NP_k^{cc} . This subclass consists of the families of *cylinders*, i. e., functions that do not depend on some player’s input, so that that player knows the answer with zero communication. We would call a family of functions hard for this class a “zero-hard family.” Alas, as we shall see (Corollary 3.6) no zero-hard families exist. It follows that no complexity class that contains the families of cylinders has a complete family of problems.

The absence of zero-hard families witnesses the weakness of the reduction concept under the consideration. One way to increase the power of the reduction, already suggested by the wording of the question of Beame *et al.*, is to allow communication. We introduce two natural definitions of such reductions - one permitting the precommunication (Karp-style), and one permitting the oracle access (Cook-style). We consider the class of families of problems (f_n) such that f_n can be computed with $t(n)$ bits of communication. We show that this class does not admit hard problems under the above reductions even if we allow $t(n)$ bits of communication for the Karp-style reduction (Corollary 4.15) and $t(n) - 1$ bits of communication for the Cook-style reduction (Corollary 5.5).

All the proofs are by simple counting arguments, much in the spirit of Beame *et al.* [3]. The main purpose of the paper is clarification of the concepts.

Organization. The rest of the paper is organized as follows. Section 2 introduces notation and definitions of communication complexity. In Section 3 we

study the reduction concepts defined by Beame *et al.* and prove the main result of this paper that there are no zero-hard families of problems. In Section 4 we introduce a Karp-style reduction and extend our results from Section 3 to the new reduction concept. In Section 5 we introduce a yet more powerful Cook-style reduction. We show that once again results similar to those in Section 3 hold even for the Cook-style reduction. In Section 6 we list a few open problems.

2 Preliminaries

Notation. For $k \in \mathbb{N}$ we write $[k] := \{1, \dots, k\}$. For a k -tuple $x = (x_1, \dots, x_k)$ and $i \in [k]$ we use $(x_1, \dots, \hat{x}_i, \dots, x_k)$ to denote the $(k-1)$ -tuple obtained from x by deleting the i -th coordinate.

A (k, n) -communication problem is a function $f : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$. Here k is the number of players and each player's input has n bits. A family of k -party communication problems is a sequence $(f_n \mid n \in \mathbb{N})$ where f_n is a (k, n) -communication problem.

In the k -party “number-on-the-forehead” (NOF) model [4], k players numbered 1 through k try to compute a given (k, n) -communication problem f on $x = (x_1, \dots, x_k)$. For each $i \in [k]$, player i sees $x_1, \dots, \hat{x}_i, \dots, x_k$, but not x_i . This can be visualized by imagining x_i written on the forehead of player i . We define the *NOF view* of the tuple x to be the k -tuple of $(k-1)$ -tuples:

$$((x_2, \dots, x_k), \dots, (x_1, \dots, \hat{x}_i, \dots, x_k), \dots, (x_1, x_2, \dots, x_{k-1}))$$

In other words, the i -th element of the NOF view of x is the $(k-1)$ -tuple of all the x_j that player i sees. We shall sometimes refer to the original k -tuple x as the *NIH view* of x .

The players can communicate with each other in order to compute $f(x_1, \dots, x_k)$. The players communicate by writing messages (binary strings) on a public blackboard in accordance with a *communication protocol*. The communication protocol specifies, as a function of the contents of the blackboard so far, which player writes the next message. The protocol also specifies, as a function of the contents of the blackboard and the writer's input, the next message to be written on the blackboard.

Following Yao's original convention [6] (see also Lovász [5]) and deviating from [3], we say that a communication protocol ends when *one of the players*, specified by the protocol, knows the answer. This is indeed the natural termination rule; this way, for instance, if each player has an n -bit input, the communication cost of the trivial protocol is n bits, rather than $n+1$ when a player has to broadcast the answer. In any case, the difference between the two conventions is at most 1 bit. But this difference will have a philosophical significance; under this definition, it is natural to consider the problems that require zero communication. These are the problems that do not depend on the input of one of the players. Following [2], we call such a problem a *cylinder*, or an *i-cylinder* if the i -th player's input is irrelevant.

The *cost* of the protocol is the maximum number of bits written on the public blackboard, where the maximum is taken over all possible inputs.

For more information on multiparty communication complexity, we refer the reader to [3] and references therein.

3 Reductions without Communication

Beame *et al.* [3, Section 5] introduce two reduction concepts, “cylindrical reductions,” which map the NOF view of the input to another such view, and “cubic reductions,” which map the NIH view of the input to another such view.

Definition 3.1 (Beame *et al.* [3]). Given the communication problems $f : X_1 \times \cdots \times X_k \rightarrow \{0, 1\}$ and $g : Y_1 \times \cdots \times Y_k \rightarrow \{0, 1\}$ we say that a k -tuple (ϕ_1, \dots, ϕ_k) of functions, $\phi_i : X_1 \times \cdots \times \widehat{X}_i \cdots \times X_k \rightarrow Y_1 \times \cdots \times \widehat{Y}_i \cdots \times Y_k$, is a *cylindric reduction* of f to g if for all $(x_1, \dots, x_k) \in X_1 \times \cdots \times X_k$ there exists $(y_1, \dots, y_k) \in Y_1 \times \cdots \times Y_k$ such that for all i we have $\phi_i(x_1, \dots, \widehat{x}_i, \dots, x_k) = (y_1, \dots, \widehat{y}_i, \dots, y_k)$ and $f(x_1, \dots, x_k) = g(y_1, \dots, y_k)$.

Definition 3.2 (Beame *et al.* [3]). Given the communication problems $f : X_1 \times \cdots \times X_k \rightarrow \{0, 1\}$ and $g : Y_1 \times \cdots \times Y_k \rightarrow \{0, 1\}$, we say that a k -tuple (ψ_1, \dots, ψ_k) of functions, $\psi_i : X_i \rightarrow Y_i$, is a *cubic reduction* of f to g if for all $(x_1, \dots, x_k) \in X_1 \times \cdots \times X_k$ we have $f(x_1, \dots, x_k) = g(\psi_1(x_1), \dots, \psi_k(x_k))$.

In our context, we have $X_i = \{0, 1\}^n$ and $Y_i = \{0, 1\}^q$.

Somewhat surprisingly, Beame *et al.* [3] prove that the two reduction concepts are equivalent; every cylindrical reduction can be expressed as a cubic reduction ([3, Lemma 5.7]).

Given a reduction concept between k -party communication problems, we say that a family (f_n) of k -party communication problems has a “polynomially-bounded” reduction to a family (g_n) of k -party communication problems if there is a *quasipolynomially bounded* ($\exp((\log n)^{O(1)})$) “stretch function” $q : \mathbb{N} \rightarrow \mathbb{N}$ such that for all sufficiently large n the communication problem f_n reduces to $g_{q(n)}$ in the given sense. (As argued in [1], $\log n$ is the natural notion of “input length,” and the analogue of “polynomial time” should allow increasing this quantity polynomially, hence the quasipolynomial increase in n .)

For a class \mathcal{C} of families of k -party communication problems, we say that a family (f_n) of k -party communication problems is “ \mathcal{C} -hard” if every family in \mathcal{C} has a polynomially bounded reduction to the family (f_n) . If, in addition, (f_n) belongs to the class \mathcal{C} , we say that it is “ \mathcal{C} -complete.”

We shall use these concepts relative to either of the two equivalent types of reductions between communication problems introduced in [3] (“cubic” and “cylindrical” reductions).

We show that if every 1-cylinder reduces to a given function g then the domain of g must be exponentially large. In fact, this remains true even if each 1-cylinder reduces to some member of a rather large set of functions g_i . We need this generalization for the subsequent diagonal argument.

Lemma 3.3. For all $k, n, q_i, m \in \mathbb{N}$, if there exists a set $\{g_i \mid i = 1, \dots, m\}$ of m communication problems $g_i : (\{0, 1\}^{q_i})^k \rightarrow \{0, 1\}$ such that every 1-cylinder $C : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$ (with n -bit inputs) has a cubic reduction to one of the g_i then $\max_i \{q_i\} \geq \frac{2^{n(k-2)}}{k} - \frac{\log m}{k2^n}$. In particular, if $m \leq 2^{2^{n(k-1)-1}}$ then $\max_i \{q_i\} \geq \frac{2^{n(k-2)}}{2k} > 2^{(n-1)(k-2)-2}$.

Proof. Let $q := \max_i \{q_i\}$. The number of 1-cylinders $C : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$ is $2^{2^{n(k-1)}}$. On the other hand, the number of cubic reductions (u_1, \dots, u_k) , where $u_i : \{0, 1\}^n \rightarrow \{0, 1\}^{q_i}$, is at most 2^{kq2^n} . They give rise to at most $m2^{kq2^n}$ functions $g_i(u_1(x_1), \dots, u_k(x_k))$. Since every 1-cylinder has a cubic reduction to some g_i , every 1-cylinder must occur among the functions $g_i(u_1(x_1), \dots, u_k(x_k))$. Consequently, $2^{2^{n(k-1)}} \leq m2^{kq2^n}$ and therefore $q \geq \frac{2^{n(k-2)}}{k} - \frac{\log m}{k2^n}$. \square

Definition 3.4. We let \mathcal{C}_1^k denote the class of families of k -party 1-cylinders.

Theorem 3.5. Let (f_n) be a family of communication problems such that every family from \mathcal{C}_1^k reduces to (f_n) . Then there exists a family $(C_n) \in \mathcal{C}_1^k$ that requires the corresponding stretch function $q : \mathbb{N} \rightarrow \mathbb{N}$ to satisfy $q(n) \geq 2^{(n-1)(k-2)-2}$.

Proof. We shall define the family (C_n) non-constructively for each n .

Consider the set of functions $S_n := \{f_\ell \mid \ell < 2^{(n-1)(k-2)-2}\}$. Clearly, we have $|S_n| < 2^{(n-1)(k-2)-2} < 2^{2^{n(k-1)-1}}$. Moreover, every 1-cylinder with n -bit inputs that has a cubic reduction to one of the functions from S_n has a corresponding stretch $q < 2^{(n-1)(k-2)-2}$. Thus by Lemma 3.3, there exists a cylinder C_n that cannot be cubically reduced to any of the $f_\ell \in S_n$. We combine thus defined cylinders into a single family (C_n) .

It follows that if a cubic reduction of the family $(C_n \mid n \in \mathbb{N})$ to the family (f_n) exists then the corresponding function q must satisfy $q(n) \geq 2^{(n-1)(k-2)-2}$. \square

Corollary 3.6. If $k \geq 3$ then there is no \mathcal{C}_1^k -hard family of communication problems.

The lower bound $\Omega(2^{n(k-2)})$ on the stretch function q provided by Theorem 3.5 should be compared with an upper bound of $O(2^{n(k-1)})$ we give next.

Proposition 3.7. There is a family of communication problems (f_ℓ) such that every family from \mathcal{C}_1^k reduces to (f_ℓ) with the stretch function $q : \mathbb{N} \rightarrow \mathbb{N}$ satisfying $q(n) \leq \lceil \frac{2^{n(k-1)}}{k-1} \rceil + n$. Moreover, the family (f_ℓ) is itself in \mathcal{C}_1^k .

Proof. The number of 1-cylinders $C_n : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ is $2^{2^{n(k-1)}}$. Thus each such cylinder has a name that can be described in binary by a string of length $2^{n(k-1)}$. We denote the name of C_n by $\langle C_n \rangle$. Now, we are ready to describe a cubic reduction $(u_{1,C_n}, \dots, u_{k,C_n})$ from C_n to f_ℓ (to be defined later):

- $u_{1,C_n}(x_1) = 0^{\lceil \frac{2^{n(k-1)}}{k-1} \rceil} \cdot x_1$, where \cdot denotes the concatenation,

- $u_{j,C_n}(x_j) = \langle C_n \rangle_j \cdot x_j$, where $j \in \{2, \dots, k\}$ and $\langle C_n \rangle_j$ denotes the substring of $\langle C_n \rangle$ spanning from index $(j-2)\lceil \frac{2^{n(k-1)}}{k-1} \rceil + 1$ to $(j-1)\lceil \frac{2^{n(k-1)}}{k-1} \rceil$ (except possibly for $j = k$, when the substring may be shorter, in which case we can inflate the substring by appending 0s to the end).

Now it is clear how the family (f_ℓ) should be defined. We define f_ℓ only for ℓ for which there exists n such that $\ell = \lceil \frac{2^{n(k-1)}}{k-1} \rceil + n$. For input (y_1, \dots, y_k) let \hat{y}_i be the substring of y_i consisting of the last n bits and \tilde{y}_i be the remaining part of y_i . Then define E to be the substring of $\tilde{y}_2 \cdot \tilde{y}_3 \cdots \tilde{y}_k$ consisting of the first $2^{n(k-1)}$ bits. Let C_n be the cylinder such that $\langle C_n \rangle = E$. Finally

$$f_\ell(y_1, \dots, y_k) := C_n(\hat{y}_1, \dots, \hat{y}_k).$$

Observe that $(f_\ell) \in \mathcal{C}_1^k$. In addition, for every 1-cylinder C_n we have

$$C_n(x_1, \dots, x_k) = f_\ell(u_{1,C_n}(x_1), \dots, u_{k,C_n}(x_k)),$$

where $\ell = \lceil \frac{2^{n(k-1)}}{k-1} \rceil + n$. □

4 Karp-Style Reduction with Communication

Let π be a protocol computing some k -party function on n -bit inputs. For each $(x_1, \dots, x_k) \in \{0, 1\}^{nk}$, we write $\pi(x_1, \dots, x_k)$ to denote the *communication transcript*, i. e., the concatenation of all messages exchanged during the execution of π on (x_1, \dots, x_k) . First we introduce the notion of “communication-constructible” functions, which plays the central role in our definition of the Karp-style reduction with communication.

Definition 4.1. Let $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}^t$ be a k -party communication problem with n -bit inputs. We say that f is *communication-constructible* if there exists a (NOF) protocol π such that for each $(x_1, \dots, x_k) \in \{0, 1\}^{nk}$ we have $f(x_1, \dots, x_k) = \pi(x_1, \dots, x_k)$.

If a function $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}^t$ is communication-constructible, then $f^i : \{0, 1\}^{nk} \rightarrow \{0, 1\}^i$, obtained by restricting output of f to a prefix of length i , is also communication-constructible. This follows by simply truncating protocol π for f after the first i bits of communication. This implies that the fraction of non-communication-constructible functions is monotonically non-decreasing in t . Thus most functions are not communication-constructible, as even for $t = 1$ most of the functions require n bits of communication by a classical counting argument. Using this observation one can derive a trivial upper bound on the number of communication constructible function. Next we give a weak, but sufficient for our purposes, upper bound on the number of communication-constructible functions.

Proposition 4.2. *The number of communication-constructible functions $\{0, 1\}^{nk} \rightarrow \{0, 1\}^t$ is at most*

$$2^{(2^{n(k-1)} + \log k)(2^t - 1)}.$$

Proof. Let $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}^t$ be a communication-constructible function. A t -bit protocol π for f can be viewed as a complete binary tree of depth t (see Figure 1). Each internal node u is labeled by a pair (P_u, f_u) , where $P_u \in [k]$ and $f_u : \{0, 1\}^{n(k-1)} \rightarrow \{0, 1\}$. The leaves do not have labels. The players use this




Figure 1: Representation of a k -party t -bit communication protocol π for a communication-constructible function as a labeled complete binary tree.

tree and the labeling for computing $f(x_1, \dots, x_k)$ in an obvious way. Starting at the root r , player P_r evaluates f_r on the given inputs and writes the result on the public blackboard. The players proceed to the corresponding subtree.

Since the set of communication-constructible functions can be mapped injectively into a set of the labeled trees as above, it suffices to count the number of such trees to establish the claim. Now, for each internal node u the number of possible labels of u is at most $k2^{n(k-1)} = 2^{2^{n(k-1)} + \log k}$. The number of internal nodes is $2^t - 1$. \square

Note that in spite of over-counting the above bound is meaningful for $t \leq n + \log \frac{n}{2}$. As we shall later see, for our purposes we may assume that $t \leq n$.

Remark 4.3. Consider t -bit k -party protocols with n -bit inputs that compute some Boolean function. The tree representation of such protocols would be similar to the tree representation from the proof of Proposition 4.2, except it would have the leaves labeled the same way as the internal nodes. Thus, we conclude that the number of t -bit k -party protocols with n -bit inputs is at most $2^{(2^{n(k-1)} + \log k)(2^{t+1} - 1)}$.

The notion of communication-constructible functions might be of interest in its own right; however, our motivation for introducing it is to give the following definition of the Karp-style reductions with communication.

Definition 4.4. Let $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ and $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ be two k -party Boolean functions. We say that f *cylindrically reduces to g with t*

bits of precommunication if there exist functions $h : \{0, 1\}^{nk} \rightarrow \{0, 1\}^t$ and $\psi_i : \{0, 1\}^{n(k-1)+t} \rightarrow \{0, 1\}^{q(k-1)}$ for $i \in [k]$ such that

1. (*precommunication*) h is communication-constructible;
2. (*consistency*) for each $(x_1, \dots, x_k) \in \{0, 1\}^{nk}$ and $z \in \{0, 1\}^t$ there exists $(y_1, \dots, y_k) \in \{0, 1\}^{qk}$ such that for each i we have $\psi_i(x_1, \dots, \widehat{x}_i, \dots, x_k, z) = (y_1, \dots, \widehat{y}_i, \dots, y_k)$, i. e., the i -th player computes $(y_1, \dots, \widehat{y}_i, \dots, y_k)$ given precommunication string z ;
3. (*validity*) for each $(x_1, \dots, x_k) \in \{0, 1\}^{nk}$, if $z = h(x_1, \dots, x_k)$ and (y_1, \dots, y_k) is defined as above then $f(x_1, \dots, x_k) = g(y_1, \dots, y_k)$.

Note: the vector (y_1, \dots, y_k) is unique.

We also refer to this reduction concept as “the Karp-style reduction with communication.”

This notion is a formal description of the following simple idea that generalizes the notion of reduction without communication: before mapping the NOF view of inputs of f to the NOF view of inputs of g , the players are allowed to perform some precommunication during which the players communicate t bits. These t bits are then available to the players in the computation of the new NOF view.

Definition 4.5. We define the *cubic reduction with t bits of precommunication* analogously to the cylindrical reduction by replacing in Definition 4.4 each function $\psi_i : \{0, 1\}^{n(k-1)+t} \rightarrow \{0, 1\}^{q(k-1)}$ with $\phi_i : \{0, 1\}^{n+t} \rightarrow \{0, 1\}^q$, where $\phi_i(x_i, z) = y_i$.

Remark 4.6. One could argue that it would be more natural to introduce and use the notion of NIH communication-constructible functions in the definition of cubic reductions with t bits of precommunication. With such a definition the cubic and cylindrical reductions would not be equivalent, and our main interest is in the NOF model. We are only interested in the NIH reduction concepts to the extent that they are equivalent to NOF reductions concepts and aid our counting arguments.

Remark 4.7. In Definition 4.4 we could have asked to enforce the consistency constraints only for $z = h(x_1, \dots, x_k)$. However, this does not seem to be enough to guarantee the equivalence of the two reduction concepts.

Now we show that the two reduction concepts, cubic and cylindrical, are equivalent, just as was the case for the reduction concepts without communication [3, Section 5].

Proposition 4.8. *Let $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ and $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ be two k -party functions. Then f cylindrically reduces to g with t bits of precommunication if and only if f cubically reduces to g with t bits of precommunication.*

Proof. Consider the following alternative characterization of a cylindrical reduction from f to g with t bits of precommunication: it is a set Ψ of at most 2^t cylindrical reductions from f to g without communication (each with stretch q) together with a communication-constructible function h that allows the players to agree (using t bits of precommunication) on a specific $(\psi_1, \dots, \psi_k) \in \Psi$ to use in the reduction step. Analogous characterization holds for the cubic reduction with t bits of precommunication - simply replace the set Ψ of cylindrical reductions without communication with a set of Φ of cubic reductions without communication. By the result of Beame *et al.* [3] we can transform each cylindrical reduction without communication $(\psi_1, \dots, \psi_k) \in \Psi$ into a cubic reduction without communication $(\phi_1, \dots, \phi_k) \in \Phi$, and vice versa. Function h is left unchanged. \square

From now on, we shall write “ f reduces to g with t bits of precommunication”, understanding that it could mean either cylindrical or cubic reduction, whichever we prefer in that particular case.

Definition 4.9. For $t \in \mathbb{N}$, let \mathcal{A}_t^k denote the class of k -party communication problems that are solvable by a t -bit protocol. When $t : \mathbb{N} \rightarrow \mathbb{N}$ is a function, we use \mathfrak{A}_t^k to denote the class of *families* (f_n) of k -party communication problems such that for each n we have $f_n \in \mathcal{A}_{t(n)}^k$.

In the rest of this section we show that reductions with precommunication exhibits the following sharp threshold phenomenon. While there is a trivial \mathfrak{A}_t^k -hard family under reductions with $t(n) + 1$ bits of precommunication, there is no \mathfrak{A}_t^k -hard family under reductions with $t(n)$ bits of precommunication.

Proposition 4.10. *For every $t : \mathbb{N} \rightarrow \mathbb{N}$ such that $t(n) \leq n$ there exists \mathfrak{A}_t^k -hard family of communication problems (τ_n) with respect to the reductions with $t(n) + 1$ bits of precommunication. Moreover, $(\tau_n) \in \mathcal{C}_1^k$ and the stretch function is $q(n) = 1$ for every $(f_n) \in \mathfrak{A}_t^k$.*

Proof. Define (τ_n) to be the family of projection functions on the second coordinate, i. e., $\tau_n : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ is defined by $(x_1, \dots, x_k) \mapsto x_2$. Let $(f_n) \in \mathfrak{A}_t^k$. To reduce (f_n) to (τ_n) the players compute the value of the function f_n and broadcast the answer using $t(n) + 1$ bits in the precommunication phase of the reduction. The functions ϕ_i simply map the players' inputs to the answer bit, which then gets output by τ_1 . In particular, we have $q(n) = 1$. Observe that from the entire family (τ_n) only τ_1 is used. \square

Lemma 4.11. *Let $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ be a k -party Boolean function. The number of distinct functions on n -bit inputs reducible with t bits of precommunication to g is at most*

$$2^{(2^{n(k-1)} + \log k)(2^t - 1) + qk2^{n+t}}.$$

Proof. A function *cubically* reducible to g with t bits of precommunication is completely determined by the two components h and the ϕ_i . The number of

distinct h is given by Proposition 4.2 and is at most $2^{(2^{n(k-1)} + \log k)(2^t - 1)}$. The total number of the $\phi_i : \{0, 1\}^{n+t} \rightarrow \{0, 1\}^q$ ($i \in [k]$) is $2^{qk2^{n+t}}$. \square

We present an analogue of Lemma 3.3 for the reductions with precommunication.

Lemma 4.12. *For all $k, n, q, m, t \in \mathbb{N}$, such that $t < n$, if there exists a set $\{g_i \mid i = 1, \dots, m\}$ of m communication problems $g_i : (\{0, 1\}^{q_i})^k \rightarrow \{0, 1\}$ such that every function $f : (\{0, 1\}^n)^k \rightarrow \{0, 1\}$ from \mathcal{A}_t^k (with n -bit inputs) is reducible to one of the g_i with t bits of precommunication then $\max_i \{q_i\} \geq \frac{2^{n(k-2)-t}}{k} - \frac{\log m}{k2^{n+t}} - 1$. In particular, if $m \leq 2^{2^{n(k-1)-1}}$ then $\max_i \{q_i\} \geq \frac{2^{n(k-2)-t}}{2k} - 1 \geq 2^{(n-1)(k-2)-t-2}$.*

Proof. Let $q := \max_i \{q_i\}$. By Lemma 4.11 the number of distinct functions reducible to one of the functions from the set $\{g_i\}$ is at most $m2^{(2^{n(k-1)} + \log k)(2^t - 1) + qk2^{n+t}}$. Observe that the functions $f : \{0, 1\}^t \times \{0, 1\}^{n(k-1)} \rightarrow \{0, 1\}$ are solvable with t bits of communication, thus they appear in \mathcal{A}_t^k . Thus, $|\mathcal{A}_t^k| \geq 2^{2^{n(k-1)+t}}$. Consequently, $2^{2^{n(k-1)+t}} \leq m2^{(2^{n(k-1)} + \log k)(2^t - 1) + qk2^{n+t}}$ and $q \geq \frac{2^{n(k-2)-t}}{k} - \frac{\log m}{k2^{n+t}} - \frac{(2^t - 1)\log k}{k2^{n+t}} \geq \frac{2^{n(k-2)-t}}{k} - \frac{\log m}{k2^{n+t}} - 1$. \square

The proof of the following theorem is analogous to the proof of Theorem 3.5.

Theorem 4.13. *Let $t : \mathbb{N} \rightarrow \mathbb{N}$ be such that $t(n) < n$, and let (f_n) be a family of communication problems such that every $(g_n) \in \mathfrak{A}_t^k$ reduces to (f_n) with $t(n)$ bits of precommunication. Then there exists a family $(g_n) \in \mathfrak{A}_t^k$ that requires the corresponding stretch function $q : \mathbb{N} \rightarrow \mathbb{N}$ to satisfy $q(n) \geq 2^{(n-1)(k-2)-t(n)-2}$.*

We immediately obtain the following corollaries.

Corollary 4.14. *For each $\epsilon \in (0, 1)$ and for each $t : \mathbb{N} \rightarrow \mathbb{N}$ such that $t(n) \leq (1 - \epsilon)n$, there is no \mathfrak{A}_t^3 -hard family of communication problems with respect to reductions with $t(n)$ bits of precommunication.*

Corollary 4.15. *For each $k \geq 4$ and for each $t : \mathbb{N} \rightarrow \mathbb{N}$ such that $t(n) < n$ there is no \mathfrak{A}_t^k -hard family of communication problems with respect to reductions with $t(n)$ bits of precommunication.*

Proposition 4.16. *For every $t : \mathbb{N} \rightarrow \mathbb{N}$, there is a family of communication problems (f_ℓ) such that every family from \mathfrak{A}_t^k reduces to (f_ℓ) without communication with the stretch function $q : \mathbb{N} \rightarrow \mathbb{N}$ satisfying $q(n) \leq 2^{n(k-1)+t(n)+1}$ for $n \geq 3$. Moreover, the family (f_ℓ) is itself in $\mathfrak{A}_{\frac{\log \ell}{k}}^k$.*

Proof. Let $(g_n) \in \mathfrak{A}_t^k$. Then there is a $t(n)$ -bit protocol π for computing g_n . In the cubic reduction, player i maps x_i to $(x_i, \langle \pi \rangle)$, where $\langle \pi \rangle$ is the name of the protocol π . By Remark 4.3, this stretches the input to size at most $(2^{n(k-1)} + \log k)(2^{t(n)+1} - 1) + n \leq 2^{n(k-1)+t(n)+1}$ for $n \geq 3$. The family to which we reduce is the ‘‘universal family’’ (U_ℓ) defined by $U_\ell((x_1, \langle \pi \rangle), \dots, (x_k, \langle \pi \rangle)) = \text{output of } \pi \text{ on } (x_1, \dots, x_k)$, where $\ell = (2^{n(k-1)} + \log k)(2^{t(n)+1} - 1) + n$.

Observe that the communication complexity of U_ℓ is bounded above by $t(n)$. Since $\ell \geq 2^{n(k-1)+t(n)} \geq 2^{t(n)k}$, we have $(U_\ell) \in \mathfrak{A}_{\frac{k}{\log \ell}}^k$. \square

Remark 4.17. Using a similar idea, if we allow $t(n)$ bits of precommunication in the above proposition then we can put the family (f_ℓ) in a class of problems solvable with 0 bits of communication. The players choose a $t(n)$ -bit protocol and run it on their inputs, as the precommunication step. In the reduction step, they append the name of the protocol together with the transcript to their inputs. The family to which we reduce is a “universal family with transcript” (UT_ℓ) defined by $UT_\ell((x_1, T, \langle \pi \rangle), \dots, (x_k, T, \langle \pi \rangle)) = \text{output of } \pi \text{ on } (x_1, \dots, x_k)$, where $T = \pi(x_1, \dots, x_k)$. Note that the stretch function satisfies $q(n) \leq (2^{n(k-1)} + \log k)(2^{t(n)+1} - 1) + t(n) + n \leq 2^{n(k-1)+t(n)+2}$ for $n \geq 3$.

5 Cook-style Reduction with Communication

Recall, the reduction concept described in Section 3 is a communication complexity analogue of the *Karp reduction* from the Turing machine world. In the current section we describe the reduction concept which is analogous to the *Cook reduction*. When f reduces to g via a Cook reduction it means that we can solve




Figure 2: Reduction of f to g with ℓ rounds of communication, i. e., ℓ queries to the g -oracle.

f with the help of an oracle for g . This has a rather intuitive interpretation in the communication complexity world (see Figure 2). More precisely:

Definition 5.1. We say that a function $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ *Cook-style reduces* to $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ with ℓ rounds and t bits of communication ($1 \leq \ell \leq t$) if there exists $t_0, \dots, t_\ell \in \mathbb{N}$, $h_i : \{0, 1\}^{nk+i+\sum_{j=0}^{i-1} t_j} \rightarrow \{0, 1\}^{t_i}$ ($0 \leq i \leq \ell - 1$), $h_\ell : \{0, 1\}^{nk+\ell+\sum_{j=0}^{\ell-1} t_j} \rightarrow \{0, 1\}$ and $\psi_j^{(i)} : \{0, 1\}^{n(k-1)+(i-1)+\sum_{s=0}^{i-1} t_s} \rightarrow \{0, 1\}^{q(k-1)}$ ($i \in [\ell], j \in [k]$) such that

1. $\ell + \sum_{i=0}^{\ell} t_i = t$,
2. for each $0 \leq i \leq \ell - 1$, h_i is communication-constructible, where the last $i + \sum_{j=0}^{i-1} t_j$ bits are interpreted as a shared string among the players,
3. communication complexity of h_ℓ is t_ℓ ,

4. for each $i \in [\ell]$ the functions $(\psi_1^{(i)}, \dots, \psi_k^{(i)})$ produce a consistent new NOF view on all inputs,
5. for each $(x_1, \dots, x_k) \in \{0, 1\}^{nk}$, after the computation is performed in the obvious way as per Figure 2 the protocol h_ℓ outputs $f(x_1, \dots, x_k)$.

Definition 5.2. We say that $f : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ *Cook-style reduces* to $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ with t bits of *communication* if there exists $0 \leq \ell \leq t$ such that f Cook-style reduces to g with ℓ rounds and t bits of total communication.

Observe that in Definition 5.1 each query to the g -oracle costs one bit of communication. In contrast, the reduction without rounds from Section 4 allows querying g for free. This difference is motivated by the following argument. If we expand g -oracle into a communication protocol for g , then after each query (= run of the protocol) we are only guaranteed that a single player knows the answer to the query. A query to g in the definition from Section 4 happens only once and at the end. Moreover, in that definition the output of the query matches the output of the function, which is sufficient to terminate the protocol for f without any extra bits of communication. However, if we allow rounds then the players should perform postcommunication with the assumption that *all of them* know the answer. This would require a player who knows the answer to a query from the run of a protocol for g to broadcast it. Thus, in case of rounds of communication we can visualize the g -oracle as writing the answers to the queries on the board, and everything that's written on the board is counted towards the communication cost.

In the rest of this section we show that the Cook-style reduction with $t-1$ bits of communication does not admit hard problems for the class of communication problems solvable with t bits of communication.

Lemma 5.3. *Let $g : \{0, 1\}^{qk} \rightarrow \{0, 1\}$ be a k -party communication problem. The number of functions on n inputs Cook-style reducible to g with t bits of communication is at most*

$$2^{2^{n(k-1)+t+1} + kq2^{n+t+1} + n2^n \log k + n - 2^{n(k-1)}}.$$

Proof. Fix $\ell \geq 1$ and t_0, \dots, t_ℓ such that

$$t_i \geq 0 \text{ and } \ell + \sum_{i=0}^{\ell} t_i = t. \tag{1}$$

We shall give a bound $b(t_0, \dots, t_\ell)$ on the number of functions that are reducible to g using t_i bits of communication in round $i+1$ of the ℓ -round reduction. This bound will be of the form $b(t_0, \dots, t_\ell) = \prod_{i=0}^{\ell} c_i$, where c_i is the number of choices available in the i -th round of reduction. Similarly to Lemma 4.11, we

have

$$\begin{aligned}
\log c_i &\leq (2^{n(k-1)+i+\sum_{j=0}^{i-1} t_j} + \log k)(2^{t_i} - 1) + kq2^{n+i+\sum_{j=0}^i t_j} \\
&= 2^{n(k-1)+i+\sum_{j=0}^i t_j} + 2^{t_i} \log k - 2^{n(k-1)+i+\sum_{j=0}^{i-1} t_j} - \log k \\
&\quad + kq2^{n+i+\sum_{j=0}^i t_j}
\end{aligned}$$

for $0 \leq i \leq \ell - 1$.

The last step in the reduction is postcommunication. Similarly to Remark 4.3, we have

$$\begin{aligned}
\log c_\ell &\leq (2^{n(k-1)+\ell+\sum_{i=0}^{\ell-1} t_i} + \log k)(2^{t_\ell+1} - 1) \\
&= 2^{n(k-1)+\ell+1} + 2^{t_\ell+1} \log k - 2^{n(k-1)+\ell+\sum_{i=0}^{\ell-1} t_i} - \log k.
\end{aligned}$$

Observe that the powers of 2 in the expressions for $\log c_i$ form the telescoping sums, thus

$$\begin{aligned}
\log b(t_0, \dots, t_\ell) &= \sum_{i=0}^{\ell} c_i \\
&\leq 2^{n(k-1)+\ell+1} - 2^{n(k-1)} + \sum_{i=0}^{\ell} 2^{t_i+1} \log k - (\ell + 1) \log k \\
&\quad + \sum_{i=0}^{\ell-1} kq2^{n+i+\sum_{j=0}^i t_j} \\
&\leq 2^{n(k-1)+\ell+1} - 2^{n(k-1)} + n2^n \log k + kq2^{n+\ell+1},
\end{aligned}$$

where $kq2^{n+\ell+1} \geq \sum_{i=0}^{\ell-1} kq2^{n+i+\sum_{j=0}^i t_j}$ since $i + \sum_{j=0}^i t_j$ are distinct for $i \in \{0, 1, \dots, \ell - 1\}$ and are all less than or equal to t . Also, $\ell \leq t \leq n$.

The number of functions that are reducible to g is $\leq \sum_{\ell, t_0, \dots, t_\ell} b(t_0, \dots, t_\ell)$, where ℓ, t_0, \dots, t_ℓ satisfy (1). For a fixed ℓ the number of solutions to (1) is $\binom{t}{\ell}$. Thus, overall the number of solutions to (1) is $\sum_{\ell=1}^t \binom{t}{\ell} < 2^t$. Finally, we obtain

$$\sum_{\ell, t_0, \dots, t_\ell} b(t_0, \dots, t_\ell) \leq 2^t 2^{n(k-1)+\ell+1} - 2^{n(k-1)} + n2^n \log k + kq2^{n+\ell+1},$$

the proof is now complete since $t \leq n$. \square

With this lemma, the following theorem follows easily using steps similar to the ones in Section 4.

Theorem 5.4. *Let $t : \mathbb{N} \rightarrow \mathbb{N}$ be such that $0 \leq t(n) \leq n$, and let (f_n) be a family of communication problems such that every $(g_n) \in \mathfrak{A}_t^k$ reduces to (f_n) with $t(n) - 1$ bits of total communication. Then there exists a family $(g_n) \in \mathfrak{A}_t^k$ that requires the corresponding stretch function $q : \mathbb{N} \rightarrow \mathbb{N}$ to satisfy $q(n) \geq 2^{(n-1)(k-2)-t(n)-2}$.*

Proof. Similar to the proof of Theorem 3.5. □

Corollary 5.5. *For each $t : \mathbb{N} \rightarrow \mathbb{N}$ such that $t(n) \leq n$ there is no \mathfrak{A}_t^k -hard family of communication problems with respect to reductions with $t(n) - 1$ bits of total communication.*

6 Open Questions

In Section 3 we showed $\Omega(2^{n(k-2)})$ lower bound on the stretch function (Theorem 3.5) and $O(2^{n(k-1)})$ upper bound (Proposition 3.7). We leave open finding the tight bound on the stretch function. This will likely close gaps in Sections 4 and 5, as well.

Problem 6.1. *Close the gap between the lower and upper bounds on the stretch function, $\Omega(2^{n(k-2)})$ vs. $O(2^{n(k-1)})$ (Theorem 3.5, Proposition 3.7). This is open even for $k = 3$.*

In Sections 4 and 5 we explored natural notions of reductions that extend the notion of Karp-style reductions in NOF communication complexity. These extensions did not add sufficient power to the reduction concept in order to permit zero-hard problems.

Problem 6.2. *Is there a reasonable notion of reduction that permits zero-hard problems?*

We expect the answer to be negative under any reasonable definition of “reasonable.”

References

- [1] László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Proc. 27th FOCS*, pages 337–347. IEEE Comp. Soc. Press, 1986.
- [2] László Babai, Noam Nisan, and Mária Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. Comp. and Syst. Sci.*, 45(2):204–232, 1992. Preliminary version appeared in 21st STOC, 1989.
- [3] Paul Beame, Matei David, Toniann Pitassi, and Philipp Woelfel. Separating deterministic from randomized multiparty communication complexity. *Theory of Computing*, 6(1):201–225, 2010.

- [4] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In *Proc. 15th STOC*, pages 94–99. ACM Press, 1983.
- [5] László Lovász. Communication complexity: a survey. In Bernhard Korte, László Lovász, Hans Jürgen Prömel, and Alexander Schrijver, editors, *Paths, Flows, and VLSI-Layout*, pages 235–265. Springer-Verlag New York, Inc., 1990.
- [6] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proc. 11th STOC*, pages 209–213. ACM Press, 1979.