An application of Hindman's theorem to a problem on communication complexity

Pavel Pudlák *

June 4, 2003

Abstract

We consider the k-party communication complexity of the problem to determine if a word w is of the form $w_0a_1w_1a_2\dots w_{k-1}a_kw_k$, for fixed letters $a_1,\dots,a_k$. Using the well-known theorem of Hindman (a Ramsey type result about finite subsets of natural numbers) we prove that for $k=4$ and 5, the communication complexity of the problem increases with the length of the word w.

1 Introduction

Let $f(x_1,x_2,\dots,x_k)$ be a function of k variables. The variables range over some finite domains, usually strings of bits, but the structure of the domains is not important. The *k-party communication complexity* of the function is defined as follows. Suppose there are k computers, numbered by $1,\dots,k$, that are to compute f. Suppose that, for $i=1,\dots,k$, computer i gets all inputs except for x_i. The k-party communication complexity of f is the amount of information, measured by bits, that any such computers need to exchange in order to compute the value of f for given inputs. More generally, one may consider the k-party communication complexity of any function $f(y_1,\dots,y_n)$ with $n\geq k$, y_j ranging over finite domains (eg. $\{0,1\}$, if f is a boolean function). In such a case one takes the maximum of the communication complexities over all partitions of the string $y_1y_2\dots y_n$ into k blocks.

We are mostly concerned with asymptotic behaviour of the complexity, thus we study sequences of functions and estimate the dependence of the complexity on the size of the inputs. We shall study the following problem.

The k-letter Problem. *Let $A=\{\eta,a_1,\dots,a_k\}$. For a given word $w\in A^*$ determine if it is of the form $w_0a_1w_1a_2\dots w_{k-1}a_kw_k$.*

Raymond, Tesson and Thérien studied multiparty communication complexity of regular languages [5]. In this regard Denis Thérien asked what is the k-party communication

*Supported by grants A1019901 of the Academy of Sciences of the Czech Republic, No. 201/01/1195 of the Grant Agency of the Czech Republic, and project No. LN00A056 of the Ministry of Education of the Czech Republic.

complexity of this problem. Very little is known about it; in particular, we do not know if, for some k, the communication complexity can be bounded by a constant independent of the size of the input. We shall show that the communication complexity is not constant for $k = 2, 3, 4$ and 5. For $k > 5$ it is an open problem; furthermore, the k-letter problem is only a special case of a more general problem posed by Thérien.

We shall use Hindman's theorem [1] to prove our lower bounds for $k = 4, 5$. The original version of this theorem speaks about sums of natural numbers. We need the set-theoretical version that talks about unions of finite sets.

Theorem 1 (Hindman's Theorem) *Let ϕ be a coloring of all finite subsets of natural numbers by a finite number of colors. Then there exists an infinite set $\mathcal{D}$ of finite pairwise disjoint sets such that all sets that are finite unions of sets of $\mathcal{D}$ have the same color.*

As we are going to prove a theorem about finite structures, we surely could do with the finite version of Hindman's theorem, but it seems that the proof would be more complicated, since we would have to control dependence among many parameters.

2 Multiparty communication complexity

Multiparty communication complexity is an important concept studied in complexity theory, see, eg., [4]. Since we are going only to prove that the communication complexity is not bounded by a constant we can consider a simplified model. It is well-known that in such a case it suffices to use the model in which all computers send independently of each other one message to a referee that determines the value of the function knowing only these messages. Such a model has a very simple combinatorial characterization.

Let D_j be the finite domain of possible values of x_j, for $j = 1, \ldots, k$. Thus the function that the k computers should cooperatively compute is a mapping of the form $f : D_1 \times \ldots \times D_k \to \{0, 1\}$. The function f determines a partition P of $D_1 \times \ldots \times D_k$ into two blocks which correspond to the two values of f (the kernel of f). The action of computer i (called the *protocol of computer* i) can be described as a mapping $g_i : D_1 \times \ldots \times D_{i-1} \times D_{i+1} \times \ldots \times D_k \to E_i$, where $g_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k)$ is the message that it sends when the computers get $(x_1, \ldots, x_k)$ as input. The function g_i determines a partition P_i of $D_1 \times \ldots \times D_{i-1} \times D_{i+1} \times \ldots \times D_k$, whose blocks correspond to the messages sent by computer i. Extend P_i to a partition Q_i of $D_1 \times \ldots \times D_k$ by ignoring the i-th coordinates. Then the referee can compute the value of the function f from the messages sent by the computers if and only if the least common refinement of $Q_1, \ldots, Q_k$ is a refinement of P.

Thus the question if a sequence of functions $f_n : D_{n,1} \times \ldots \times D_{n,k} \to \{0, 1\}$ has constant k-party communication complexity reduces to the following one: Is it possible to find partitions $P_{n,i}$ of $D_{n,1} \times \ldots \times D_{n,i-1} \times D_{n,i+1} \times \ldots \times D_{n,k}$ such that they have constant number of blocks and the least common refinement of their extensions to $D_{n,1} \times \ldots \times D_{n,k}$ is a refinement of the kernel of f_n?

For proving lower bounds on the k-letter problem the following version seems to be more convenient. For a given word $w \in A^*$ and $i = 1, \ldots, k$, computer i gets the entire word w, but all occurrences of letter a_i are replaced by η. A lower bound for this model easily translates to a lower bound for the standard model with at most constant factor reduction. For those who are used to the standard model of multiparty communication complexity, it is important to *keep in mind that throughout the paper we are using a different model*, because some arguments use essentially properties of this modification!

3 k=2 and 3

For $k = 2$ the problem is trivial. The property is equivalent to the condition that the first occurrence of a_1 precedes the last occurrence of a_2. Thus the communication problem is equivalent to the problem in which each of the two computers has a number and they are to determine if one is less than the other. It is a well-known and easy fact that they need $\log n$ bits, where n is the length of the word.

For $k = 3$ the property is equivalent to the condition that there is an occurrence of a_2 between the first occurrence of a_1 and the last occurrence of a_3. Also in this case we are able to determine the number of bits they need (up to a multiplicative constant). Though the proof is easy, it may be instructive to read it before studying more difficult lower bounds for $k = 4$ and 5.

First we prove a lower bound. Let the protocols of the three computers be given. As noted above, for a_1, only the first occurrence is important, and for a_3, it is only the last one. Thus we can restrict ourselves to words that contain only one occurrence of a_1 and only one occurrence of a_3; furthermore we shall assume that a_1 precedes a_3. Then the protocol for computer 2, the one that does not see letters a_2, is simply a mapping from pairs of numbers $i < j \leq n$ (n being the length of the word) into a finite set of messages. Using Ramsey's theorem for pairs (see, for example, [2]), we can find a set of numbers $X \subseteq [n]$ such that the message sent by computer 2 on words with a_1 on position i and a_3 on position j is the same for all $i, j \in X$, $i < j$ and such that the size of X is at least $\frac{1}{r} \log_r n$, where r denotes the number of messages. Let $X = \{x_0, \ldots, x_{q-1}\}$, $x_0 < \ldots < x_{q-1}$. W.l.o.g. assume q is even and the difference between every two elements is at least two. Consider words that for some $i, j \in X$, $i < j$, have a single occurrence of a_1 on position i, a single occurrence of a_3 on position j, and letter a_2 occurs on all positions between x_{2t-1} and x_{2t}, for $t = 1, \ldots, q/2 - 1$, but nowhere else. Now, consider only those words with $i = x_{2t}, j = x_{2t+1}$ for $t = 0, 1, \ldots, q/2 - 1$. All these should be rejected. If the number of pairs of messages that computer 1 and 3 can send is less than $q/2 - 1$, there are two such words on which they send the same messages. Let these two words be determined by x_{2t}, x_{2t+1} and $x_{2t'}, x_{2t'+1}$ respectively, $t < t'$. Computer 2 also sends the same message on these two words, as we have chosen positions from X. But if we now take the word determined by $x_{2t}, x_{2t'+1}$ they will send exactly the same messages again, because:

- computer 1 cannot distinguish it from the word determined by $x_{2t'}, x_{2t'+1}$,
- computer 3 cannot distinguish it from the word determined by x_{2t}, x_{2t+1},
- computer 2 sends the same message for every pair of positions $i, j \in X$, $i < j$.

Hence computers 1 and 3 must send at least $q/2-1$ different pairs of messages. Let ℓ be the number of possible messages that the three computers can send. Then we get $\ell = \Omega(\frac{1}{\ell}\log_\ell n)$, hence $\ell = \Omega(\sqrt{\log_2 n/\log_2\log_2 n})$. Thus the communication complexity is estimated to be $\Omega(\log\log n)$.

Now we show an upper bound; the idea of this proof is due to Jiří Sgall. W.l.o.g. assume that the length n of the words is a number of the form 2^k. The distance of two positions $i < j$ on the word will be the number $j - i$ (eg., the distance between the first and the last positions is $2^k - 1$). We shall use binary representation of distances. The protocol for computers will be as follows.

1. Computer 1 looks for the last occurrence of a_3 and the last occurrence of a_2 before it. If it fails to find such a pair, it sends a message saying that. Otherwise it computes the distance d_1 of the two positions and sends the position of the most significant bit of d_1.

2. Computer 2 looks for the first occurrence of a_1 and the last occurrence of a_3. If the former is after the latter, or one of the two letters does not occur, it sends a message saying that. Otherwise it computes the distance d_2 of the two positions and sends the position of the most significant bit of d_2.

3. Computer 3 looks for the first occurrence of a_1 and the first occurrence of a_2 after it. If it fails to find such a pair, it sends a message saying that. Otherwise it computes the distance d_3 of the two positions and sends the position of the most significant bit of d_3.

Given these messages one can determine if the word contains occurrences of a_1, a_2, a_3 in this order as follows. If any of the three computers sends an "error message", then the answer is *no*. Otherwise, we know the most significant digit of d_1 and d_3 and two most significant digits of d_2. Let $\tilde{d}_1, \tilde{d}_2, \tilde{d}_3$ be d_1, d_2, d_3 rounded down to the first, resp. first two, resp. first most significant digits. The information that we have is equivalent to knowing $\tilde{d}_1, \tilde{d}_2, \tilde{d}_3$. Then we say *yes* if $\tilde{d}_1 \le \tilde{d}_2/2$ or $\tilde{d}_3 \le \tilde{d}_2/2$, otherwise we say *no*. We shall prove that this is correct.

1. If the word has the property, then, clearly $d_1 + d_3 \le d_2$. Hence either $d_1 \le d_2/2$ or $d_3 \le d_2/2$. Then also $\tilde{d}_1 \le \tilde{d}_2/2$ or $\tilde{d}_3 \le \tilde{d}_2/2$.

2. If the word does not have the property, then $d_1 > d_2$ and $d_3 > d_2$. Hence $\tilde{d}_1 \ge \tilde{d}_2 > \tilde{d}_2/2$ and $\tilde{d}_3 \ge \tilde{d}_2 > \tilde{d}_2/2$.

Thus we have proved:

Theorem 2 *The communication complexity of the 3-letter problem is* $\Theta(\log\log n)$.

The same upper bound holds for $k > 3$. Our lower bounds for $k = 4, 5$ are much smaller than $\log\log n$ and we believe that the true value is also smaller, but we do not know of a better protocol than the one above.

4 Definitions and notation

We denote by $\mathbb{N} = \{0, 1, 2, \ldots\}$. For $n \in \mathbb{N}$, $n > 1$, $[n]$ will denote the set $\{1, \ldots, , n\}$; for a set X, we denote by $[X]^n$ the set of n-element subsets of X; $\mathcal{P}_{fin}(X)$ will denote the set of all finite subsets of X; we shall use $\mathcal{P}(X)$ if X is finite.

For a set of sets $\mathcal{A}$ we denote by $FU(\mathcal{A})$ the set of nonempty sets that are unions of finitely many subsets of $\mathcal{A}$.

For a number n and subsets X, Y of natural numbers we shall write $X < Y$ if $\max X < \min Y$; $n < X$ if $n < \min X$ etc.

A set $\mathcal{X}$ of sets will be called a *disjoint family*, if the sets in $\mathcal{X}$ are pairwise disjoint. We shall mostly consider sets of finite subsets of natural numbers. For such sets one can often replace the disjointness by the following stronger condition. We shall say that $\mathcal{X} \subseteq \mathcal{P}_{fin}(\mathbb{N})$ is *separated,* if for every $X, Y \in \mathcal{X}$, $X \neq Y$ either $X < Y$ or $Y < X$. It is an easy exercise to show that every infinite disjoint family of finite subsets of $\mathbb{N}$ contains an infinite separated family. Note that in a separated family the sets are linearly ordered by the relation $<$ defined above. Another easy fact is that for every two infinite disjoint families $\mathcal{A}_1, \mathcal{A}_2$ that contain only finite sets one can find infinite families $\mathcal{B}_1 \subseteq \mathcal{A}_1$, $\mathcal{B}_2 \subseteq \mathcal{A}_2$ such that $\mathcal{B}_1$ and $\mathcal{B}_2$ are disjoint.

Hindman's theorem, stated above, immediately implies the following formally stronger statement: for every $\phi : FU(\mathcal{C}) \to [\ell]$, with $\mathcal{C} \subseteq \mathcal{P}_{fin}(\mathbb{N})$ an infinite disjoint family, there exists an infinite disjoint family $\mathcal{D} \subseteq FU(\mathcal{C})$ such that ϕ is constant on $FU(\mathcal{D})$.

5 A reduction to a Ramsey type statement

Conjecture 1 *For every $r, \ell, m \in \mathbb{N}$, $r \geq 2$, there exists an $n \in \mathbb{N}$ such that for every*

$$\phi_1, \ldots, \phi_r : [n]^2 \times (\mathcal{P}([n]))^{r-1} \to [\ell]$$

there exists $X \subseteq [n]$, $|X| = m$, $B_1, \ldots, B_r \subseteq [n]$, $c_1, \ldots, c_r \in [\ell]$ such that

1. *$X, B_1, \ldots, B_r$ are pairwise disjoint;*
2. *for every two consecutive elements x, x' of X there exist $j_0 = x < j_1 < \ldots < j_{r-1} < j_r = x'$ such that*
   (a) *$B_i \cap (j_{i-1}, j_i] = \emptyset$ for $i = 1, \ldots, r$;*
   (b) *$B_{i'} \cap (j_{i-1}, j_i] \neq \emptyset$ for $i, i' = 1, \ldots, r, i \neq i'$;*
3. *for every $x, y \in X$, $x < y$, $i = 1, \ldots, r$,*

$$\phi_i(x, y, B_1, \ldots, B_{i-1}, B_{i+1}, \ldots, B_r) = c_i.$$

Proposition 3 *If the conjecture holds for $r \geq 2$, then for $k = r + 2$ the k-party communication complexity of the k-letter Problem is not constant.*

Proof. Let $r \geq 2$ be given, let $k = r+2$. Suppose that the k-letter problem can be solved by k computers each using at most ℓ different messages. Let $m = \ell^2 + 1$. Take n given to r, ℓ, m by the conjecture. We shall consider words of length n that contain exactly one a_1 and one a_k and a_1 is before a_k. Then the strategies of computers 2 to $k-1$ can be represented as mappings $\phi_1, \ldots, \phi_r : [n]^2 \times (\mathcal{P}_{fin}([n]))^{r-1} \to [\ell]$, where the two numerical inputs correspond to the positions of letters a_1 and a_k, and the set inputs correspond to the letters $a_2, \ldots, a_{k-1}$. More precisely, for ϕ_i, the input sets are sets of indices of occurrences of letters $a_2, \ldots, a_{i-1}, a_{i+1}, \ldots, a_{k-1}$. Let $X, B_1, \ldots, B_r$ be given by the conjecture. Then we consider words such that a_1, resp. a_k, is on the position x, resp. y, for $x, y \in X$, $x < y$, and for $i = 2, \ldots, k-1$, letter a_i occurs on the positions determined by the set B_{i-1}. Thus the positions of letters $a_2, \ldots, a_{k-1}$ are fixed, while there are several possibilities for letters a_1, a_k determined by a choice of $x, y \in X$, $x < y$. The condition on ϕ_i–s says that the messages sent by computers 2 to $k-1$ will be the same on all these words. The messages of computer 1 will depend only on the position of a_k and the messages of computer k will depend only on the position of a_1. Hence there are two pairs x, x' and y, y' of consecutive elements of X such that the messages sent by computer 1 and k will be the same for x, x' and y, y'. Thus they will send the same messages also for x, y'. Consequently, the messages of all computers will be the same for x, x' and x, y'.

To get a contradiction we shall show that the word determined by two consecutive elements of X should be rejected, while if $x, y \in X$ are not consecutive, it should be accepted. Let x, x' be two consecutive elements. Take the partition $j_0 = x < j_1 < \ldots < j_{r-1} < j_r = x'$ of the interval $[x, x']$. Suppose the word w determined by x, x' is of the form $w_0 a_1 w_1 a_2 \ldots w_{k-1} a_k w_k$. Then a_1 is on position j_0 and a_k on j_r. Since a_2 does not occur in the first interval of the partition, the occurrence of a_2 displayed in $w_0 a_1 w_1 a_2 \ldots w_{k-1} a_k w_k$ is in the second or later intervals. Since a_3 does not occur in the second interval of the partition, the occurrence of a_3 is in the third or later intervals etc. Thus a_{k-2} must occur in the $k-2$-nd, which is the last, the r-th interval. But then for a_{k-1} there is no more room. Now it is also clear that for non-consecutive elements x, y we can find $a_1, \ldots, a_k$ in this order in the word.

□

6 A common generalization of Ramsey's and Hindman's theorems

We shall prove a lemma that is in some sense a common generalization of Ramsey's and Hindman's theorem. A slightly weaker version of it was proved in [3].

We shall consider the types of arrangements of $n+1$ finite subsets of natural numbers. A *type* will be a word in the alphabet $\{x, a_1, \ldots, a_n\}$ that does not contain subwords $a_i a_i$ for $i = 1, \ldots, n$ (however, it may contain xx and every letter may occur more than once). For disjoint subsets of natural numbers $X, A_1, \ldots, A_n$, we define the type of this set arrangement, denoted by $type(X, A_1, \ldots, A_n)$, as follows.

1. $type(\emptyset, \emptyset, \ldots, \emptyset) = \Lambda$ (the empty word).

2. If $m < X, A_1, \ldots, A_n$, and $type(X, A_1, \ldots, A_n) = \tau$, then $type(\{m\} \cup X, A_1, \ldots, A_n) = x\tau$.

3. If $\emptyset \neq B < X, A_1, \ldots, A_n$, and $type(X, A_1, \ldots, A_n) = \tau$, then

   (a) $type(X, A_1, \ldots, A_{i-1}, B \cup A_i, A_{i+1}, \ldots, A_n) = \tau$, if τ starts with a_i,

   (b) $type(X, A_1, \ldots, A_{i-1}, B \cup A_i, A_{i+1}, \ldots, A_n) = a_i\tau$ otherwise.

Less formally, write the word with letters in order as indicated by the sets $X, A_1, \ldots, A_n$ and eliminate successive occurrences of letters other than x.

Example. Let $n = 3$, $X = \{1, 4, 6\}$, $A_1 = \{2, 3, 11\}$, $A_2 = \emptyset$ and $A_3 = \{8, 14, 15, 16\}$. Then $type(X, A_1, A_2, A_3) = xa_1xxa_3a_1a_3$.

Lemma 4 *Let $s, \ell, n \in \mathbb{N}$, let $X \subseteq \mathbb{N}$ be an infinite set, let $\mathcal{A}_1, \ldots, \mathcal{A}_n \subseteq \mathcal{P}_{fin}(\mathbb{N} \setminus X)$, be infinite separated families. Let*

$$\phi : [\mathbb{N}]^s \times \mathcal{P}_{fin}(\mathbb{N})^n \to [\ell],$$

and let τ be a type with s occurrences of letter x and at least one occurrence of every a_i, $i = 1, \ldots, n$. Then there exist an infinite subset $Y \subseteq X$, infinite separated pairwise disjoint families $\mathcal{B}_i \subseteq FU(\mathcal{A}_i)$, $i = 1, \ldots, n$, and $c \in [\ell]$ such that for every $z \subseteq Y$, $|z| = s$, $A_i \in FU(\mathcal{B}_i)$, $i = 1, \ldots, n$, if $type(z, A_1, \ldots, A_n) = \tau$, then $\phi(z, A_1, \ldots, A_n) = c$.

Note that Lemma 4 can be used repeatedly, thus one can ensure the homogeneity condition for any finite number of types.

Proof. We shall prove the lemma by induction on the length of the type. For the empty type the lemma is trivial.

1. Suppose $\tau = x\sigma$ and the lemma holds for the type σ. Let y_1 be the first element of X. Take $\phi_{y_1} : [\mathbb{N}]^{s-1} \times \mathcal{P}_{fin}(\mathbb{N})^n \to [\ell]$ obtained from ϕ by fixing the first argument to y_1. By the induction assumption, we can find an infinite subset $Y_1 \subseteq X$, infinite disjoint families $\mathcal{B}_i^1 \subseteq FU(\mathcal{A}_i)$, $i = 1, \ldots, n$, and $c_1 \in [\ell]$ such that for every $z \subseteq Y_1$, $|z| = s-1$, $A_i \in FU(\mathcal{B}_i^1)$, $i = 1, \ldots, n$, if $type(z, A_1, \ldots, A_n) = \sigma$, then $\phi(z, A_1, \ldots, A_n) = c_1$. For $i = 1, \ldots, n$, let B_i^1 be the first element of $\mathcal{B}_i$, that is after y_1, ie., $y_1 < B_i^1$.

Now suppose we have already chosen $y_1 < \ldots < y_m$, $X \supseteq Y_1 \supseteq \ldots \supseteq Y_m$, $\mathcal{B}_i^1, \ldots, \mathcal{B}_i^m$, such that $\mathcal{B}_i^{j+1} \subseteq FU(\mathcal{B}_i^j)$ for $i = 1, \ldots, n, j = 1, \ldots, m-1$, $c_m \in [\ell]$ and $B_i^1 < \ldots < B_i^m$, for $i = 1, \ldots, n$. Then we choose y_{m+1}, Y_{m+1}, $\mathcal{B}_i^{m+1} \subseteq FU(\mathcal{B}_i^m)$, $i = 1, \ldots, n$, using the lemma for the type σ, in the same way as we did for the base case, except for the following. We require, moreover, that y_{m+1} is after all B_i^m's, ie., $B_i^m < y_{m+1}$ for $i = 1, \ldots, n$. The sets B_i^{m+1}, $i = 1, \ldots, n$, are again the first elements of $\mathcal{B}_i^{m+1}$'s that are after y_{m+1}.

Let $Y' = \{y_1, y_2, \ldots\}$, $\mathcal{B}_i = \{B_i^1, B_i^2, \ldots\}$, $i = 1, \ldots, n$. Note that

$$y_1 < B_1^1, \ldots, B_n^1 < y_2 < B_1^2, \ldots, B_n^2 < y_3 < \ldots,$$

hence every $\mathcal{B}_i$ is infinite. The construction ensures that for every $z \subseteq Y'$, $|z| = s$, $A_i \in FU(\mathcal{B}_i)$, $i = 1, \ldots, n$, if $type(z, A_1, \ldots, A_n) = \tau$, then $\phi(z, A_1, \ldots, A_n)$ depends only on the

$\min z$. Namely, if $\min z = y_j$, then $\phi(z, A_1, \ldots, A_n) = \phi_{y_j}(z \setminus \{y_j\}, A_1, \ldots, A_n) = c_j$. Thus taking Y as those elements of Y' that correspond to a $c \in [\ell]$ that occurs infinitely many times, we get the statement of the lemma.

2. Suppose $\tau = a_r\sigma$ (thus the first letter of σ is not a_r) and the lemma holds for the type σ. We shall consider two cases according to whether or not a_r occurs in σ. First we assume that a_r does occur in σ. The proof is similar to the one above with the role of y_m's now played by B_r^m's.

Define

$$\phi_B(z, A_1, \ldots, A_n) = \phi(z, A_1, \ldots, A_{r-1}, B \cup A_r, A_{r+1}, \ldots, A_n). \tag{1}$$

We shall construct sets and numbers

$$B_r^1 < y_1, B_1^1, \ldots, B_{r-1}^1, B_{r+1}^1, \ldots, B_n^1 < B_r^2 < y_2, B_1^2, \ldots, B_{r-1}^2, B_{r+1}^2, \ldots, B_n^2 < B_r^3 < \ldots$$

and disjoint families $\mathcal{B}_1^m, \ldots, \mathcal{B}_n^m$, $m = 1, 2, \ldots$ as follows. B_r^1 is simply the first element of $\mathcal{A}_r$. Then we take $\phi_{B_r^1} : [\mathbb{N}]^s \times \mathcal{P}_{fin}(\mathbb{N})^n \to [\ell]$. By the induction assumption we get sets $Y_1, \mathcal{B}_1^1, \ldots, \mathcal{B}_n^1$ with the homogeneity property for $\phi_{B_r^1}$ with respect to the type σ. We define y_1 to be the first element of Y_1 after B_r^1, and B_i^1 to be the first element of $\mathcal{B}_i^1$ after B_r^1, for $i \in [n], i \neq r$. The next set B_r^2 is the first element of $\mathcal{B}_r^1$ that is after $y_1, B_1^1, \ldots, B_{r-1}^1, B_{r+1}^1, \ldots, B_n^1$.

In the m-th step we proceed similarly, but we do not use only $\phi_{B_r^m}$. We have to take all ϕ_B for $B \in FU(\{B_r^1, \ldots, B_r^m\})$ such that $B_r^m \subseteq B$. Since $FU(\{B_r^1, \ldots, B_r^m\})$ is finite, we can ensure the homogeneity property for all these colorings by repeated application of the induction assumption for the type σ.

Now put $Y = \{y_1, y_2, \ldots\}$, $\mathcal{B}_i = \{B_i^1, B_i^2, \ldots\}$, for $i \in [n], i \neq r$ and $\mathcal{B}_r' = \{B_r^1, B_r^2, \ldots\}$. Then the construction ensures that for every $z \subseteq Y$, $|z| = s$, $A_i \in FU(\mathcal{B}_i)$, $i \in [n], i \neq r$ and $A_r \in FU(\mathcal{B}_r')$, if $type(z, A_1, \ldots, A_n) = \tau$, $A_r = B \cup A_r'$, $A_r', B \in FU(\mathcal{B}_r)$, and

$$B < z, A_1, \ldots, A_{r-1}, A_r', A_{r+1}, \ldots, A_n,$$

then $\phi(z, A_1, \ldots, A_n)$ depends only on B. Indeed, if $\min z = y_m$, then $B \in FU(\{B_r^1, \ldots, B_r^m\})$. Also

$$\mathcal{B}_1 \subseteq \mathcal{B}_1^m, \ldots, \mathcal{B}_{r-1} \subseteq \mathcal{B}_{r-1}^m, \mathcal{B}_r' \subseteq \mathcal{B}_r^m, \mathcal{B}_{r+1} \subseteq \mathcal{B}_{r+1}^m, \ldots, \mathcal{B}_n \subseteq \mathcal{B}_n^m$$

and $A_r' \in \mathcal{B}_r^m$. Hence $\phi(z, A_1, \ldots, A_n)$ depends only on B by (1) and the construction of the families $\mathcal{B}_1^m, \ldots, \mathcal{B}_n^m$.

Finally, we color elements $A_r \in FU(\mathcal{B}_r')$ by $\phi(z, A_1, \ldots, A_n)$ for some $z \subseteq Y$, $|z| = s$, $A_i \in FU(\mathcal{B}_i)$, $i \in [n], i \neq r$ such that $type(z, A_1, \ldots, A_n) = \tau$ and take, using Hindman's theorem, an infinite separated family $\mathcal{B}_r \subseteq FU(\mathcal{B}_r')$ such that all sets in $FU(\mathcal{B}_r)$ have the same color.

3. Now suppose that $\tau = a_r\sigma$ and a_r does not occur in σ. Instead of (1), we define

$$\phi_B'(z, A_1, \ldots, A_{r-1}, A_{r+1}, \ldots, A_n) = \phi(z, A_1, \ldots, A_{r-1}, B, A_{r+1}, \ldots, A_n). \tag{2}$$

Then we argue in the same way as in part 2.

□

7 k=4 and 5

Theorem 5 *Conjecture 1 is true for $r = 2$ and 3. Hence for $k = 4$ and 5 the k-party communication complexity of the k-letter problem is not constant.*

Proof.
r=2. Suppose the conjecture is false. Let ℓ, m be such that for every n there exist mappings $\phi_1, \phi_2 : [n]^2 \times \mathcal{P}([n]) \to [\ell]$ that violate the conclusion of the conjecture. By König Lemma, there exist mappings $\phi_1, \phi_2 : \mathbb{N}^2 \times \mathcal{P}_{fin}(\mathbb{N}) \to [\ell]$ such that for every n their restrictions to $[n]^2 \times \mathcal{P}_{fin}([n])$ violate the conclusion of the conjecture. We shall use Lemma 4 with $s = 2$, ℓ as above, $n = 1$ and the type $axaxa$. (We shall use letters $a, b, \ldots$ instead of $a_1, a_2, \ldots$ for denoting types.) First we apply the theorem to ϕ_1 and $X = \mathbb{N}$, $\mathcal{A}_1 = \mathcal{P}_{fin}(\mathbb{N})$. Thus we obtain some $Y_1 \subseteq \mathbb{N}$ and $\mathcal{B}_1$. Then we apply it again to ϕ_2, and $X = Y_1$, $\mathcal{A}_1 = \mathcal{P}_{fin}(\mathbb{N})$. Thus we get some $Y_2 \subseteq Y_1$ and $\mathcal{B}_2$. Hence for both $i = 1, 2$, and for every $A_1 < x_1 < A_2 < x_2 < A_3$, $A_1, A_2, A_3 \in FU(\mathcal{B}_i)$, $x_1, x_2 \in Y_2$, the value of $\phi_i(x_1, x_2, A_1 \cup A_2 \cup A_3)$ is the same.

Now we choose $A_{i,0}, \ldots, A_{i,m+1} \in \mathcal{B}_i$, for $i = 1, 2$, and $x_1, \ldots, x_m \in Y_2$ such that

$$A_{2,0} < A_{1,0} < x_1 < A_{2,1} < A_{1,1} < x_2 < \ldots < x_m < A_{2,m+1} < A_{1,m+1}.$$

In terms of types it means that we choose a configuration of type $(bax)^m ba$, where a stands for sets in $\mathcal{B}_1$, b stands for sets in $\mathcal{B}_2$ and x for elements of Y_2. Then for $i = 1, 2$ and every $1 \le p < q \le m$, the value of $\phi_i(x_p, x_q, A_{i,0} \cup \ldots \cup A_{i,m+1})$ is the same, because the sets $A_{i,0} \cup \ldots \cup A_{i,p-1}, A_{i,p} \cup \ldots \cup A_{i,q-1}, A_{i,q} \cup \ldots \cup A_{i,m+1}$ are in $FU(\mathcal{B}_i)$. Hence ϕ_1, ϕ_2 restricted to the interval $[0, \max A_{1,m+1}]$, $X = \{x_1, \ldots, x_m\}$, $B_i = A_{i,0} \cup \ldots \cup A_{i,m+1}$, $i = 1, 2$ satisfy the statement of Conjecture 1, which is a contradiction.

r=3. Suppose the conjecture is false for $r = 3$ and some ℓ, m. Using the same argument as above we get mappings $\phi_1, \phi_2, \phi_3 : \mathbb{N}^2 \times (\mathcal{P}_{fin}(\mathbb{N}))^2 \to [\ell]$ such that for every n their restrictions to $[n]^2 \times (\mathcal{P}_{fin}([n]))^2$ violate the conclusion of the conjecture. Let t be the Ramsey number such that $t - 2 \to (m)^2_\ell$. We shall use Lemma 4 with $s = 2$, ℓ as above, $n = 2$ and all types of the form $(ab)^i x (ab)^{j-i} x (ab)^{t-j}$ with $1 < i < j < t$. First we apply it to ϕ_1 and $X = \mathbb{N}$, $\mathcal{A}_1 = \mathcal{A}_2 = \mathcal{P}_{fin}(\mathbb{N})$. Thus we get some Y, $\mathcal{C}$ and $\mathcal{D}$. Then we apply it to ϕ_2 and $X = Y$, $\mathcal{A}_1 = \mathcal{P}_{fin}(\mathbb{N})$, $\mathcal{A}_2 = \mathcal{D}$. Thus we get some Y', $\mathcal{B}$ and $\mathcal{D}'$. Finally we apply it to ϕ_3 and $X = Y'$, $\mathcal{A}_1 = \mathcal{B}$, $\mathcal{A}_2 = \mathcal{C}$. Thus we get some Y'', $\mathcal{B}'$ and $\mathcal{C}'$. Now we have to be more precise about the types that we are interested in, as eg., in $(ab)^i x (ab)^{j-i} x (ab)^{t-j}$ a may stand for elements of $FU(\mathcal{C})$ and b for elements of $FU(\mathcal{D})$, or b may stand for elements of $FU(\mathcal{C})$ and a for elements of $FU(\mathcal{D})$. Therefore, we shall use letters b, c, d in the types for $\mathcal{B}', \mathcal{C}', \mathcal{D}'$. The letter x stands, of course, for elements of Y''. Now the requirements can be stated briefly as follows.

1. For every type $(dc)^i x (dc)^{j-i} x (dc)^{t-j}$ with $1 < i < j < t$, ϕ_1 is constant on all arrangements of this type;

2. for every type $(db)^{i'} x (db)^{j'-i'} x (db)^{t-j'}$ with $1 < i' < j' < t$, ϕ_2 is constant on all arrangements of this type;

3. for every type $(cb)^{i''}x(cb)^{j''-i''}x(cb)^{t-j''}$ with $1 < i'' < j'' < t$, ϕ_3 is constant on all arrangements of this type.

Applying Ramsey's theorem to the pairs of indices in $[2, t-1]$ we get sets $I, I', I'' \subseteq [2, t-1]$, $|I| = |I'| = |I''| = m$ such that

1. ϕ_1 is constant on all arrangements of types $(dc)^i x(dc)^{j-i}x(dc)^{t-j}$ with $i < j$, $i, j \in I$;

2. ϕ_2 is constant on all arrangements of types $(db)^{i'}x(db)^{j'-i'}x(db)^{t-j'}$ with $i' < j'$, $i', j' \in I'$;

3. ϕ_3 is constant on all arrangements of types $(cb)^{i''}x(cb)^{j''-i''}x(cb)^{t-j''}$ with $i'' < j''$, $i'', j'' \in I''$.

Let us denote the elements of I resp. I' resp. I'' in increasing order as follows $i_1, \dots, i_m$ resp. $i'_1, \dots, i'_m$ resp. $i''_1, \dots, i''_m$.

Now we are ready to choose an arrangement from Y'', $\mathcal{B}', \mathcal{C}', \mathcal{D}'$ that will violate our initial assumption. We take an arrangement of type

$$(dc)^{i_1-1}(bd)^{i'_1-1}(cb)^{i''_1-1}x(dc)^{i_2-i_1-1}(bd)^{i'_2-i'_1-1}(cb)^{i''_2-i''_1-1}x\dots$$
$$\dots(dc)^{i_m-i_{m-1}-1}(bd)^{i'_m-i'_{m-1}-1}(cb)^{i''_m-i''_{m-1}-1}x(dc)^{t-i_m-1}(bd)^{t-i'_m-1}(cb)^{t-i''_m-1}$$

Let $X \subseteq Y''$, $|X| = m$, $B \in \mathcal{B}'$, $C \in \mathcal{C}'$ and $D \in \mathcal{D}'$ be the sets that form the arrangement of this type. If we take only sets X, C, D, then their type is

$$(dc)^{i_1}x(dc)^{i_2-i_1}x\dots(dc)^{i_m-i_{m-1}}x(dc)^{t-i_m}$$

Hence $\phi_1(x, y, C, D)$ has the same value for all $x, y \in X$, $x < y$. The same argument shows that also $\phi_2(x, y, B, D)$ and $\phi_3(x, y, B, C)$ have the same value for all $x, y \in X$, $x < y$. Thus, if we restrict ϕ_1, ϕ_2, ϕ_3 to $[n]^2 \times (\mathcal{P}([n]))^2$, for $n = \max B$, we get a contradiction with our initial assumption.

$\square$

Acknowledgment

I would like to thank to Denis Thérien for telling me the k-letter problem and for several discussion about things related to it, to Jiří Sgall for his upper bound for case k=3, to Vojtěch Rödl for the reference [3], to William Gasarch for his comments and to an anonymous referee for pointing out an error in the manuscript.

References

[1] N. Hindman, *Finite sums from sequences within cells of a partition of N.* J. Comb. Theor. (A) 17, (1974), pp. 1–11.

[2] L. Lovász, *Combinatorial Problems and Exercises.* Akadémiai Kiadó, Budapest, 1979.

[3] K.R. Milliken, *Ramsey's theorem with sums or unions.* J. Comb. Theor. (A) 18, (1975), pp. 276–290.

[4] E. Kushilevitz and N. Nisan, *Communication Complexity.* Cambridge Univ. Press, 1977.

[5] J.-F. Raymond, P. Tesson and D. Thérien, *An algebraic approach to communication complexity.* ICALP'98, Springer-Verlag LNCS 1443, (1998), pp. 29-40.