Reprinted (without any changes) from:
Mathematical Logics and Its Applications, vol. 5 (1987), pp.22-31

ISSN 0207-8961

Lower Bounds on Communication Complexity[*]

Stasys P. Jukna

Abstract

A notion of "communication complexity" is used to formally measure the degree to which a Boolean function is "global". An explicit combinatorial lower bound for this complexity measure is presented. In particular, this leads to an $\exp(\Omega(\sqrt{n}))$ lower bound on the complexity of depth-restricted contact schemes computing some natural Boolean functions in NP.

1 Introduction

Suppose that a Boolean function $f : \{0,1\}^n \to \{0,1\}$ must be computed by two distinct computers. Each computer receives half of the input bits, and the computation proceeds using some protocol for communication between the two computers. The minimum number of bits that has to be exchanged in order to successfully compute f, minimized over all partitions of the input into two equal parts, is called the *communication complexity* of f. This model of communication was introduced by Ch. Papadimitriou and M. Sipser [5]. The motivation for this complexity measure is that it provides a direct lower bound for the minimum bisection width of any chip that recognizes f.

The paper is divided as follows. Section 2 involves the definition of communication complexity. The basic result concerning the lower bound for this complexity measure is given in Section 3. In Section 4 the communication in bounded-depth contact-gating schemes is involved and a lower bound for such a schemes is provided. Section 5 contains an example of a Boolean function with hight communication complexity.

2 The Model

Fix some set of Boolean variables $X = \{x_1, \ldots, x_n\}$ with $n \equiv 0 \,(\mathrm{mod} 2)$. An *assignment* on X is a mapping δ from X into $X \cup \{0,1\}$ such that $(\forall x \in X)\ \delta(x) \notin \{0,1\} \to \delta(x) = x$; $\mathrm{dom}(\delta) = \delta^{-1}(0) \cup \delta^{-1}(1)$ is the *domain* of δ. For $Y \subseteq X$, let $[Y]$ denote the set of all

---

[*] Some of these results were presented at the Second All Union Seminar on Discrete Mathematics and Its Applications (Moscow, January 1987).

assignments γ on X with $\mathrm{dom}(\gamma) = Y$. Note that $[X] = \{0,1\}^X$. A *restriction* of $\delta \in [X]$ to $Y \subseteq X$ is an assignment $\delta\rceil_Y$ in $[Y]$ that coincides (on Y) with δ. A *partition* of X is a pair $\pi = (X_0, X_1)$ of its subsets with $X = X_0 \cup X_1$, $X_0 \cap X_1 = \emptyset$ and $|X_0| = |X_1|$ (throughout, $|A|$ is the cardinality of A).

A (nondeterministic) *protocol on input* X is a pair $P = (\pi, \Phi)$, where

(a) $\pi = (X_0, X_1)$ is a partition of X.

(b) Φ is a relation

$$\Phi \subseteq ([X_0] \cup [X_1]) \times \{0,1,\#\}^* \times (\{0,1\}^* \cup \{\text{accept}, \text{reject}\})$$

Intuitively, the first argument of Φ is the local part of the input, while the second argument is the sequence of all previous messages. The third argument is the next message. For a given string $w \in \{0,1,\#\}^*$, the relation Φ has the following *prefix-freeness* property: for any two $\delta, \gamma \in \{0,1\}^{n/2}$, if (δ, w, u), $(\gamma, w, v) \in \Phi$ then u is not a prefix of v (for a motivation of such a restriction see, e.g. [6]). The protocol $P = (\pi, \Phi)$ is called *deterministic* if Φ is a function from $([X_0] \cup [X_1]) \times \{0,1,\#\}^*$ to $(\{0,1\}^* \cup \{\text{accept}, \text{reject}\})$.

A *computation* of $P = (\pi, \Phi)$ *on input* $\delta \in [X]$ is a string $w = w_1\#w_2\#\ldots\#w_k$, where $k \geq 1$, $w_1, w_2, \ldots, w_k \in (\{0,1\}^* \cup \{\mathit{accept}, \mathit{reject}\})$, and such that, for each i, $0 \leq i \leq k$, we have

$$(\delta_i, w_1\#w_2\#\ldots\#w_i, w_{i+1}) \in \Phi$$

where δ_i is the restriction of δ to $X_{i\,(\mathrm{mod} 2)}$. A computation w *accepts* δ if $w_1, \ldots, w_{k-1} \notin \{\text{accept}, \text{reject}\}$ and $w_k = \text{accept}$. We say P *computes* a Boolean function $f : [X] \to \{0,1\}$ if, for all inputs δ in $[X]$, $f(\delta) = 1$ iff there is a computation of P on input δ that accepts δ.

The *depth* of a computation w is the number of messages in w, i.e. $\mathrm{depth}(w)$ is the number of #'s plus one. The *width* of w is the length of its maximal message (the messages 0 and 1 are supposed to have zero length). If P computes f then $\mathrm{depth}(P)$ ($\mathrm{width}(P)$) is the minimum of $\mathrm{depth}(w)$ (of $\mathrm{width}(w)$ resp.) over all computations w that accept δ, maximized over all inputs δ from $f^{-1}(1)$.

For $k \geq 1$, we define the *communication complexity* of f by

$$\mathrm{comm}_k(f) = \min\{\mathrm{width}(w) \,:\, P \text{ computes } f \text{ and } \mathrm{depth}(P) \leq k\}.$$

Notice that for any $f : \{0,1\}^n \to \{0,1\}$ and $k \geq 1$,

$$0 \leq \mathrm{comm}_{k+1}(f) \leq \mathrm{comm}_k(f) \leq n/2,$$

and for $k = n/2$,

$$0 \leq \mathrm{comm}_k(f) \leq 1.$$

3 The Lower Bound

For a $(0,1)$-matrix A, let $\mathrm{per}(A)$ denote the permanent of A and let $\langle A, q\rangle$ denote the set of all $q\times q$-submatrices of A. The *term-rank*, $\mathrm{tr}(A)$, and the *clique-number*, $\mathrm{cl}(A)$, of A are defined by

$$\mathrm{tr}(A) = \max\{q \,:\, \mathrm{per}(B) > 0 \text{ for some } B \text{ in } \langle A, q\rangle\}$$

and

$$\mathrm{cl}(A) = \max\{q \,:\, \mathrm{per}(B) = q! \text{ for some } B \text{ in } \langle A, q\rangle\}.$$

Given a Boolean function $f(X)$, $X = \{x_1, \ldots, x_n\}$, and an assignment δ on X, denote by f^δ the function we get by composing f and δ, i.e. $f^\delta = f(\delta(x_1), \ldots, \delta(x_n))$. Note that f^δ is a function of $n - |\mathrm{dom}(\delta)|$ variables. For a partition $\pi = (Y, Z)$ of X, we define the following $(0,1)$-matrix $M(f,\pi)$ of order $2^{n/2}\times 2^{n/2}$:

$$M(f,\pi) = \left\{ f^{\delta\gamma} \,:\, \delta \in [Y] \text{ and } \gamma \in [Z] \right\}.$$

For $f \neq$ *const*, define the *dispersion*, $\Theta(f)$, of f by

$$\Theta(f) = \min\left\{ \frac{\mathrm{tr}(M(f,\pi))}{\mathrm{cl}(M(f,\pi))} \,:\, \pi \text{ is a partition of } X \right\}.$$

Note that

$$1 \leq \Theta(f) \leq 2^{n/2}.$$

Theorem 3.1 *For any $k \geq 1$ and a Boolean function $f \neq$ const the following bound holds*

$$\mathrm{comm}_k(f) \geq k^{-1} \cdot \log \Theta(f).$$

Proof. Choose some protocol $P = (\pi, \Phi)$ computing f, and such that $\mathrm{depth}(P) \leq k$ and $\mathrm{comm}_k(f) = \mathrm{width}(P)$. Let $\pi = (Y, Z)$.

Choose some maximal subset of assignments $D \subseteq \{\delta \in [X] \,:\, f^\delta = 1\}$ such that, for all $\delta \neq \gamma$ in D, $\delta\rceil_Y \neq \gamma\rceil_Y$ and $\delta\rceil_Z \neq \gamma\rceil_Z$. Then $|D| = \mathrm{tr}(M(f,\pi))$.

Now, let $t = \mathrm{width}(P)$. Define the computation $w = w_1\#w_2\#\ldots\#w_k$ of P inductively as follows. To define the message w_{i+1} consider the set $D(i)$ of all assignments δ in D for which $w_1\#\ldots\#w_i$ is a prefix of an accepting δ computation of depth k. (Hence $D(0) = D$). Let w_{i+1} be a message in $\{0,1\}^* \cup \{0,1\}$ for which $w_1\#\ldots\#w_i\#w_{i+1}$ is the prefix of computations (of depth k) that accept at least $|D(i)|\cdot 2^{-t}$ assignments in $D(i)$.

Since the computation w accepts at least $|D(k)| \geq |D|\cdot 2^{-tk}$ assignments in D and $t = \mathrm{comm}_k(f)$, it remains to show that $\mathrm{cl}(M(f,\pi)) \geq |D(k)|$.

Indeed, if $\delta \neq \gamma$ are in $D(k)$ then both δ and γ are accepted by w. Then by cutting and pasting argument, w accepts both $(\delta\rceil_Y, \gamma\rceil_Z)$ and $(\delta\rceil_Z, \gamma\rceil_Y)$. So,

$$f^\delta = f^{\delta\rceil_Y,\gamma\rceil_Z} = f^{\delta\rceil_Z,\gamma\rceil_Y} = f^\gamma,$$

and the proof follows. $\square$

4 Communication in Contact Schemes

A contact-gating scheme over the set of Boolean variables $X = \{x_1, \ldots, x_n\}$ is a finite acyclic digraph (multiple edges allowed) with edges labeled by $x_1, \ldots, x_n$, $\bar{x}_1, \ldots, \bar{x}_n$ (cf. [4]). One of the nodes is a *source* (has fan-in zero), some other nodes are *leafs* (fan-out zero). A *branching program* is a contact-gating scheme such that

(i) every node has outdegree at most 2, and

(ii) for every node v with outdegree=2, one of the edges leaving v is labeled by a variable $x \in X$ and the other is labeled by its complement $\bar{x}$ (see, e.g. [2,3,7]).

A scheme computes a Boolean function in a natural way: $S(X)$ computes $f : \{0,1\}^X \to \{0,1\}$ if for any δ in $\{0,1\}^X$, it holds that $f(\delta) = 1$ iff $S(\delta)$ contains a path from the source to a leaf. The *size* of a scheme is the number of edges.

A set of nodes V of S is called a *cut* if each path from the source to a leaf contains exactly one node in V. For cuts U and V we shall write $U \leq V$ if there is a path from each node in U to some node in V. For $U \leq V$, let $S[U, V]$ denote the sub-scheme of S between U and V (including U and V).

For a scheme $S(X)$, let depth(S) denote the minimal number k for which there exists a partition $\pi = (X_0, X_1)$ of X and a sequence of cuts

$$V_0 \leq V_1 \leq \ldots \leq V_k$$

such that $V_0 = \{source\}$, $V_k = \{leafs\}$ and, for each $i = 0, \ldots, k-1$, the function computed by $S[V_i, V_{i+1}]$ does not depend on variables in $X_{i+1\,(\mathrm{mod} 2)}$. For a Boolean function f and $k \geq 1$, denote

$$C_k(f) = \min\{\mathrm{size}(S) : S \text{ computes } f \text{ and } \mathrm{depth}(S) \leq k\}.$$

In case of branching programs the corresponding measure is denoted by $BP_k(f)$. Obviously, $BP_k(f) \geq C_k(f)$.

Remark: The depth k contact scheme model is quite powerful even for $k = const$. There are Boolean functions f_n that require nearly-exponential (up to $\exp(n/\log n)$) size to be computed by any sufficiently "local" scheme (see [2], [3]), and $BP_4(f_n) = O(n^2)$. On the other hand, constant-depth schemes are also quite powerful for almost all functions (the term "almost all" refers to a $(1 - o(1))$ fraction of the $\exp\exp(n)$ possible choices of n-variable Boolean functions). Namely, the method by O.B. Lupanov [4] implies that for almost all $f_n : \{0,1\}^n \to \{0,1\}$, the following asymptotic holds

$$C(f_n) \sim C_4(f_n) \sim 2^n/n.$$

Every contact-gating scheme $S(X)$ of depth k defines the following protocol $P_S = (\pi, \Phi_S)$. Set $t = \max |V_i|$ and fix some injection ν from $\{0, 1, \ldots, t-1\}$ to $\{0,1\}^{\log t}$. The relation Φ_S is defined as follows. For any $i \in \{0, \ldots, k-1\}$, $m_1, \ldots, m_{i+1} \in \{0, \ldots, t-1\}$ and an assignment δ in $X_{i\,(\mathrm{mod} 2)}$, let

$$(\delta, \nu(m_1)\# \ldots \#\nu(m_i), \nu(m_{i+1})) \in \Phi_S$$

iff there exists a path in $S(\delta)$ from the m_i-th node in V_i to m_{i+1}-th node in V_{i+1}. Notice that only the last message $\nu(m_i)$ is essential for Φ_S.

Remark: If S is a branching program then the corresponding protocol P_S is deterministic.

The scheme S and the protocol P_S both compute the same Boolean function. Moreover,

$$\mathrm{size}(S) \geq 2^{\mathrm{width}(P_S)}. \tag{4.1}$$

It is known (see, e.g. [7]) that for $k = \infty$, the contact gating scheme complexity and the branching program complexity are polynomially related. Namely, there exists a constant $c \geq 1$ such that for any Boolean function f, it holds that

$$BP(f) \leq C(f) \leq (BP(f))^c.$$

However, for depth-restricted schemes the picture changes drastically.

Proposition 4.1 *There is a sequence of Boolean functions* $\{f_n\}_{n=1}^{\infty}$ *such that*

$$C_2(f_n) \leq n^{O(1)}$$

and for any $k = k(n) \geq 1$,

$$BP_k(f_n) \geq 2^{\Omega(n/k)}.$$

Proof. For $m \geq 2$, let $T_n(X)$ denote the function of $n = \binom{m}{2}$ Boolean variables X, whose value is 1 iff X represents the adjacency matrix of an undirected graph of m nodes containing a triangle. It is easy to check that $C_2(T_n) \leq n^{O(1)}$. On the other hand, it is known ([6]) that for any $k \geq 1$,

$$k \cdot \text{det-com}_k(T_n) \geq \Omega(n)$$

where "det-com" stands for *deterministic* communication complexity. It remains to use (4.1). ▯

Theorem 4.1 *For any* $k \geq 1$ *and any Boolean function* $f \neq$ *const, it holds that*

$$C_k(f) \geq \Theta(f)^{1/k}.$$

Proof. Follows directly from (4.1) and Theorem 3.1. ▯

5 Example

Let $GF(q)$ be the finite Galois field of order q, where q is a prime power and $q \equiv 0\,(\mathrm{mod} 2)$. Define $\mathrm{POL}_n(X)$ to be the function of $n = q^2$ Boolean variables $X = \{x_{a,b} : a, b \in GF(q)\}$, whose value is 1 iff there exists a polynomial Q of degree at most $d = q/2 - 1$ over $GF(q)$ such that for all a, b in $GF(q)$,

$$x_{a,b} = 1 \iff b = Q(a).$$

Remark: POL_n is the characteristic function for the set of all "lower ones" (i.e. of prime implicants) of an NP-complete monotone Boolean function investigated by A.E. Andreev [1].

Lemma 5.1 $\Theta(\mathrm{POL}_n) \geq 2^{\sqrt{n}/2}$.

To prove the lemma, we need some combinatorial properties of POL_n.

Lemma 5.2 *For any partition $\pi = (Y, Z)$ of X and an assignment δ in $[Y]$, put*

$$\Gamma^\delta = \{\gamma \in [Z] \,:\, \mathrm{POL}_n^{\delta,\gamma} = 1\}.$$

Then

$$|\Gamma^\delta| \leq \max\{0,\, H(\|\delta\|)\}$$

Proof. If $\|\delta\| \geq d+1$ then $\Gamma^\delta = \emptyset$ since any two distinct polynomials of degree at most d over $GF(q)$ differ in at least $q - d$ points.

Let $\|\delta\| = t \leq d$. Denote by $\mathbb{C}_\delta$ the set of all columns C of X such that $C - \delta^{-1}(1) \neq \emptyset$. The either $\Gamma^\delta = \emptyset$ or $|\mathbb{C}_\delta| = q - \|\delta\} = q - t$. Let $\mathbb{C}_\delta = \{C_1, \ldots, C_{q-t}\}$ and put $s_i = |C_i - \mathrm{dom}(\delta)|$. Then

$$s_1 + \ldots + s_{q-t} \leq |Z| = n/2 \tag{5.1}$$

where w.l.o.g.

$$s_1 \leq s_2 \leq \ldots \leq s_{q-t}. \tag{5.2}$$

Set $r = d + 1 - t$, and let $h(s_1, \ldots, s_r)$ denote the number of all r-tuples $(j_1, \ldots, j_r)$, where $1 \leq j_i \leq s_i$, $i = 1, \ldots, r$. Then obviously,

$$|\Gamma^\delta| \leq h(s_1, \ldots, s_r), \tag{5.3}$$

where by (5.1) and (5.2)

$$s_1 + \ldots + s_r \leq n/2 - (d+1)s_r \tag{5.4}$$

Since the sum of $s_1, \ldots, s_r$ is bounded, the maximum of $h(s_1, \ldots, s_r)$ is achieved for $s_1 = \ldots = s_r$. Hence by (5.4), $s_r \leq n/(2(q-t))$, and so, $|\Gamma^\delta| \leq H(t) = h(n/(2(q-t)), \ldots, n/(2(q-t)))$. ▯

Proof of Lemma 5.1. Let $f = \mathrm{POL}_n$ and $\pi = (Y, Z)$ be a partition of X such that $\Theta(\mathrm{POL}_n) = \mathrm{tr}(M)/\mathrm{cl}(M)$, where $M = M(f, \pi)$. The matrix M contains exactly $|f^{-1}(1)| = q^{d+1}$ ones. Since $H(t) \leq H(0)$, by Lemma 5.2 we have that the minimal number of lines (columns and rows) we need to cover all the 1's of M is no less than $q^{d+1} \cdot H(0) = 2^{d+1}$. By König-Egervary theorem (see, e.g. [5]) this minimal number of lines coincides with the term-rank of M. Therefore,

$$\Theta(\mathrm{POL}_n) \geq 2^{d+1} \cdot \mathrm{cl}(\mathrm{POL}_n)^{-1},$$

where by Lemma 5.2, $\mathrm{cl}(\mathrm{POL}_n) = 1$. ▯

Corollary 5.1 *For any* $k = k(n)$,

$$\mathrm{comm}_k(\mathrm{POL}_n) \geq \sqrt{n}/2.$$

Corollary 5.2 *If* $k = O(n^{1/2-\epsilon})$ *for some* $0 \leq \epsilon \leq 1/2$ *then*

$$C_k(\mathrm{POL}_n) \geq 2^{n^{\epsilon - o(1)}}.$$

In particular, for any constant k,

$$C_k(\mathrm{POL}_n) \geq 2^{\Omega(\sqrt{n})}.$$

Finally, notice that

$$\mathrm{comm}_1(\mathrm{POL}_n) \leq \sqrt{n}\log n/4.$$

Indeed, let some $(0,1)$-matrix A of order $q \times q$ be given. Divide A into two submatrices A_0 and A_1 of order $q \times (d+1)$ each. To compute POL_n, the first computer transmits either reject or the binary code $bin(Q)$ of a polynomial Q of degree at most d over $GF(q)$ such that (the graph of) Q corresponds to A_0. The second computer then has enough information to decide acceptance. The length of $bin(Q)$ is at most $(d+1)\log q$. Note that this protocol is even deterministic (i.e. with Φ, a function).

References

[1] А.Е. Андреев. Об одном методе получения нижних оценок сложности индивидуалных функции. ДАН СССР. 1985. Т. 282. N. 5. С. 1033-1037.

[2] S.P. Jukna. Lower bounds on the complexity of local circuits. *Proc. 12th Symp. Math. Foundations of Comput. Sci.* Bratislava/Czech. 1986. Lect. Notes in Comput. Sci. 1986. V. 233. P. 440-448.

[3] S.P. Jukna. Entropy of contact circuits and lower bounds on their complexity. *Theoretical Computer Science.* 1987 (to appear)

[4] О.Б. Лупанов. О вентилных и контактно-вентилных схемах. ДАН СССР. 1956. Т. III. N. 6. С. 1171-1174.

[5] H. Minc. *Permanents.* Addison-Wesley. 1978.

[6] Ch.H. Papadimitriou and M. Sipser. Communication complexity. *Journal of Comput. Syst. Sci.* 1984. V. 28. P. 26-269.

[7] P. Pudlák. The hierarchy of Boolean circuits. Inst. of Math. Prague. Preprint N. 20. 1986.

[8] С.П. Юкна. Об одном методе получения нижних оценок сложности былевых функции. ДАН СССР. 1987.