Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 19 (2008)

ENTROPY OF OPERATORS OR WHY MATRIX MULTIPLICATION IS HARD FOR DEPTH-TWO CIRCUITS

STASYS JUKNA

Abstract. We consider unbounded fanin depth-2 circuits with *arbitrary* boolean functions as gates. We define the entropy of an operator $f : \{0,1\}^n \to \{0,1\}^m$ as the logarithm of the maximum number of vectors distinguishable by at least one special subfunction of f.

Our main result is that every depth-2 circuit for f requires at least entropy(f) wires. This gives an information-theoretic explanation of *why* some operators require many wires. We use this to prove a tight estimate $\Theta(n^3)$ of the smallest number of wires in any depth-2 circuit computing the product of two n by n matrices over any finite field. Previously known lower bound for this operator was $\Omega(n^2 \log n)$.

1. Introduction

One of the challenges in circuit complexity is to prove a nonlinear lower bound for log-depth circuits computing an explicitly given boolean operator $f : \{0,1\}^n \to \{0,1\}^n$. This corresponds to simultaneous computation of the sequence of boolean functions $f_j : \{0,1\}^n \to \{0,1\}$, where $f_j(\mathbf{x})$ is the j-th coordinate of the vector $f(\mathbf{x})$. An important result of Valiant [20] reduces this problem to proving a lower bound $\Omega(n^{1+\epsilon})$ on the number of wires in a depth-2 circuit computing a linear operator $y = A\mathbf{x}$ over GF_2, where we allow arbitrary boolean functions as gates. Note that in this case the phenomenon which causes complexity of circuits is *information transfer* instead of *information processing* in the case of single functions. It is therefore important to understand what properties of operators do force high information transfer in their depth-2 circuits.

A *depth-2 circuit* for $f : \{0,1\}^n \to \{0,1\}^m$ is a directed acyclic graph with n input nodes $x_1, \ldots, x_n$, and m output nodes $z_1, \ldots, z_m$. Every noninput node computes an *arbitrary* boolean function of its inputs, and there is no bound on the fanin or on the fanout. The *size* of a circuit is the total number of wires in it. Without loss of generality, we may assume that there are no direct wires from inputs to outputs: this can be easily achieved by adding at most n new wires.

Let $s_2(f)$ denote the minimum size of a depth-2 circuit computing f. Note that $s_2(f) \leq n^2$ for every operator $f : \{0,1\}^n \to \{0,1\}^n$.

Superlinear lower bounds of the form $s_2(f) = \Omega(n \log n)$ were obtained using graph-theoretic arguments by analyzing some superconcentration properties of the circuit as a graph [6, 9, 12, 11, 15]. Unfortunately, the approach based on superconcentrators cannot lead to

1991 *Mathematics Subject Classification.* 68Q17, 94C10.

Key words and phrases. Boolean circuits, bilinear forms, matrix multiplication, entropy.

Affiliation: Institute of Mathematics and Computer Science, Vilnius, Lithuania.

Current address: University of Frankfurt, Institute of Informatics, Robert-Mayer-Str. 11-15, D-60054 Frankfurt, Germany.

Email: jukna@thi.informatik.uni-frankfurt.de.

Research supported by the DFG grant SCHN 503/4-1.

ISSN 1433-8092

lower bounds for depth-2 circuits larger than $\Omega(n\log^2 n)$, since there are depth-2 superconcentrators with $O(n\log^2 n)$ [10], and even $O(n\log^2 n/\log\log n)$ [14] edges.

The (numerical) limitation of the graph-theoretic lower bounds comes from their power: they show much more than that the number of wires must be large—they also provide an information about the structure of the underlying graphs. It is therefore natural to expect to prove larger lower bounds, if we only care about the number of wires in a circuit, not about its structure. Such a direct approach has already led Cherukhin [5] to the highest known lower bound $s_2(f)=\Omega(n^{3/2})$ for an explicit boolean operator f—cyclic convolution computing n special *bilinear* forms $\mathbf{x}^\top A\mathbf{y}$ over GF_2. (Recall that such a bound for a *linear* operator $A\mathbf{x}$ would already imply nonlinear lower bound for log-depth circuits.)

In this paper we prove a general lower bound $s_2(f)\geq \text{entropy}(f)$, where the entropy of $f:\{0,1\}^n\to\{0,1\}^m$ is just the logarithm of the maximum number of vectors distinguishable by at least one special subfunction of f. This gives a simple explanation of what operators and, more importantly, *why* require many wires. The bound itself is reminiscent of a classical lower bound of Nechiporuk [8] on the formula size of a boolean function as the logarithm of the number of its subfunctions.

Since $\text{entropy}(f)$ is relatively easy to compute, this gives us a handy tool to prove large lower bounds for a whole string of explicit operators. We demonstrate this by a tight estimate $\Theta(n^3)$ of the smallest number of wires in any depth-2 circuit computing the product of two n by n matrices over any finite field. This improves the highest previously known lower bound $s_2(f)=\Omega(n^2\log n)$ for this operator derived in [15] using a technical lemma from [11] and graph-theoretic arguments.

2. Results

In this section we first introduce the notion of entropy of operators, and state some its basic properties. Then we prove our main result—a general lower bound on the number of wires in depth-2 circuits in terms of the entropy (Lemma 3 and Theorem 4).

2.1. **Entropy of function sets.** Let $F=\{f_1,\dots,f_m\}$ be a set of functions $f_j:\{0,1\}^n\to\{0,1\}$ on the same set of variables $x_1,\dots,x_n$. Say that a set of vectors $A\subseteq\{0,1\}^n$ is *separated* by F, if for every pair of vectors $a\neq b\in A$ there is a function $f\in F$ with $f(a)\neq f(b)$, that is, if the corresponding to F operator in injective on A. Define

$$\text{entropy}(F)=\max\{\log_2|A|\colon A\subseteq\{0,1\}^n \text{ and } F \text{ separates } A\}.$$

Say that a function f can be computed from a set of functions G if there exists a boolean function φ such that $f=\varphi(g_1,\dots,g_k)$ for some functions $g_1,\dots,g_k$ in G. We write $F\leq G$ if every function in F can be computed from the functions in G. Note that, in any circuit with arbitrary boolean functions as gates, every function is computed from the set of functions computed at its inputs. In particular, every set of functions F on variables $x_1,\dots,x_n$ is computatble from $G=\{x_1,\dots,x_n\}$.

Proposition 1. *Let F and G be some finite sets of boolean functions in n variables.*

(i) Upper bound: $\text{entropy}(F)\leq\min\{n,|F|\}$.
(ii) Lower bound: *if F contains r single variables, then* $\text{entropy}(F)\geq r$.
(iii) Main connection: *if $F\leq G$ then* $\text{entropy}(F)\leq\text{entropy}(G)\leq|G|$.

Proof. (i) The set $F=\{f_1,\dots,f_m\}$ defines a natural encoding of vectors $a\in\{0,1\}^n$ by vectors $F(a)=(f_1(a),\dots,f_m(a))$ in $\{0,1\}^m$. If a set $A\subseteq\{0,1\}^n$ is separated by F, then

each vector in A must receive its own code, implying that $|A| \leq 2^m = 2^{|F|}$, and hence, $\log_2 |A| \leq |F|$.

(ii) Suppose that F contains r single variables $x_1, \ldots, x_r$. Let $A \subseteq \{0,1\}^n$ be an arbitrary set of $|A| = 2^r$ vectors having the same values on all remaining $n-r$ variables. Since any pair of vectors $a \neq b \in A$ must differ in at least one of the first r coordinates, each such pair is separated by at least one of the variables $x_1, \ldots, x_r$.

(iii) Just observe that then $G(a) = G(b)$ implies $F(a) = F(b)$. Hence, any set separated by F must be also separated by G, implying that entropy$(F) \leq$ entropy$(G) \leq |G|$, where the last inequality follows from (i). ∎

2.2. **Entropy of subfunctions and the number of wires.** Let F and G be two sets of boolean functions. We can think of F as a set of functions computed by some circuit at its output nodes, and G as a set of functions computed at some intermediate nodes. Fix some set $\mathbf{x} = (x_1, \ldots, x_n)$ of variables, and call them *main variables*. Let $\mathbf{y} = (y_1, \ldots, y_r)$ be the set of the remaining *auxiliary* variables.

We say that a main variable x_i is *critical* for a function $g(\mathbf{x}, \mathbf{y})$ if $g(\mathbf{e}_i, \mathbf{y}) \neq g(\mathbf{0}, \mathbf{y})$, where $\mathbf{e}_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ is the vector of length n with precisely one 1 in the i-th coordinate.

Given a subset $X \subseteq \{x_1, \ldots, x_n\}$ of main variables, let $X(g)$ denote the set of all variables $x_i \in X$ which are critical for g. The number of variables in $X(g)$ is the *weight* of g with respect to the set of variables X. The *weight* of a set G of functions, denoted by weight$_X(G)$, is the sum of weights of all its functions. We will see soon (Lemma 3) that, in depth-2 circuits, weight$_X(G)$ lower bounds to the number of wires leaving the inputs in X.

Remark. Recall that a function $g(\mathbf{x}, \mathbf{y})$ *depends* on a variable x_i if $g(\mathbf{a} \oplus \mathbf{e}_i, \mathbf{y}) \neq g(\mathbf{a}, \mathbf{y})$ for at least one vector $\mathbf{a} \in \{0,1\}^n$. Hence, in general, the number $|X(g)|$ of variables in X that are critical for g may be much smaller than the total number of variables in X on which the function g depends. If, for example, $g(\mathbf{x}, \mathbf{y}) = x_1 x_2 \oplus y$ and $X = \{x_1, x_2\}$, then $g(1,0,1) = g(0,1,1) = 1 \neq 0 = g(1,1,1)$, implying that g depends on both variables x_1 and x_2. But $g(1,0,y) = g(0,1,y) = g(0,0,y) = y$ implies that $X(g) = \emptyset$.

If every function in F can be computed from the functions in G, then Proposition 1(iii) implies $|G| \geq$ entropy(F). To get a similar (entropic) lower bound on weight$_X(G)$ we consider the following set F_X of subfunctions of the functions in F.

We define the set F_X of subfunctions of F with respect to X to be the set of all boolean functions in variables Y that can be obtained from some function $f \in F$ by setting some variable $x_i \in X$ to 1 and all the remaining main variables to 0. That is,

$$F_X = \{f(\mathbf{e}_i, \mathbf{y}) \colon f \in F, x_i \in X\},$$

Note that F_X may contain up to $|X| \cdot |F|$ different functions.

Lemma 2 (Entropy and weight)**.** *If every function in F can be computed from the functions in G, then*

$$\text{weight}_X(G) \geq \text{entropy}(F_X) - |G|.$$

Proof. Since the functions in F can be computed from the functions in G, the subfunctions in F_X can be computed from the subfunctions in G_X, as well. By Proposition 1(iii), we have entropy$(F_X) \leq$ entropy(G_X). By Proposition 1(i), it remains to show that $|G_X| \leq$ weight$_X(G) + |G|$.

To show this, recall that G_X consists of all boolean functions $g(\mathbf{e}_i, \mathbf{y})$ obtained from some function $g \in G$ by setting some variable $x_i \in X$ to 1 and the remaining main variables to 0. If

$x_i \notin X(g)$, then $g(\mathbf{e}_i, \mathbf{y}) = g(\mathbf{0}, \mathbf{y})$. Hence, for each $g \in G$, the set $\{g(\mathbf{e}_i, \mathbf{y})\colon x_i \in X\}$ consist of at most $|X(g)|$ functions $g(\mathbf{e}_i, \mathbf{y})$ with $x_i \in X(g)$ and just one additional function $g(\mathbf{0}, \mathbf{y})$. Summing over all $g \in G$, we obtain that $|G_X| \le |G| + \sum_{g\in G} |X(g)| = |G| + \mathrm{weight}_X(G)$. ∎

Let now $f = (f_1, \ldots, f_m)$ be an operator and $F \subseteq \{f_1, \ldots, f_m\}$. Let also $X \subseteq \{x_1, \ldots, x_n\}$ be a subset of main variables. Lemma 2 yields the following basic relation between the entropy and the number of wires.

Lemma 3 (Entropy and depth-2 complexity)**.** *In any depth-2 circuit computing f, the number of wires leaving the inputs in X or entering the outputs in F must be at least* $\mathrm{entropy}(F_X)$.

Proof. Let M be the set of all nodes on the middle layer joined by a wire with at least one output in F. Then F must be computable from the set $G = \{g_v\colon v \in M\}$ of boolean functions computed at the nodes $v \in M$. Since we have $|M| \ge |G|$ wires entering the outputs in J, it remains, by Lemma 3, to show that at least $\mathrm{weight}_X(G)$ wires must leave the inputs in X.

Each node $v \in M$ must be connected by a wire with each input $x_i \in X$ of which the function g_v depends. Hence, at least $|X(g_v)|$ wires must go from X to the node v. Since no wire can go to more than one node, the total number of wires from X to M must be at least $\sum_{v\in M} |X(g_v)| = \mathrm{weight}_X(G)$. ∎

Define the *entropy* of an (n, m)-operator $f = (f_1, \ldots, f_m)$ as

$$\mathrm{entropy}(f) = \max \sum_{t=1}^{p} \mathrm{entropy}(\{f_j(\mathbf{e}_i, \mathbf{y})\colon i \in I_t, j \in J_t\}), \tag{2.1}$$

where the maximum is over all partitions $I_1, \ldots, I_p$ of inputs $[n]$ and all partitions $J_1, \ldots, J_p$ of outputs $[m]$. Since the total number of wires in a depth-2 circuit is just the number of wires incident to its input or output nodes, Lemma 2 directly yields the following

Theorem 4. *For every operator f, we have $s_2(f) \ge \mathrm{entropy}(f)$.*

Remark. Theorem 4 can be readily extended to sequences of functions $f : D^n \to D$ for any *finite* set D. For this, it is enough to take the logarithm to the basis $|D|$ in Definition **??** of the entropy. The rest is the same.

Remark. Taking *partitions* of inputs and outputs in the definition of $\mathrm{entropy}(f)$ is not crucial. For each natural number k, we can define $\mathrm{entropy}_k(f)$ as the maximum (2.1) over all subsets $I_1, \ldots, I_p$ of inputs and all subsets $J_1, \ldots, J_p$ of outputs such that no element belongs to more than k of these sets. Hence, taking partitions corresponds to $k = 1$. Now, if $d(i)$ is the number of wires leaving the input i, then the sum

$$\sum_{t=1}^{p} \sum_{i\in I_t} d(i) = \sum_{i=1}^{n} \sum_{t:i\in I_t} d(i) \le k \sum_{i=1}^{n} d(i)$$

is at most k times larger than the total number $\sum_{i=1}^{n} d(i)$ of wires leaving the inputs. Since the same also holds for the number of wires entering the output nodes, Lemma 2 implies

$$s_2(f) \ge \max_{k\ge 1} \frac{1}{k} \cdot \mathrm{entropy}_k(f).$$

3. Application: matrix multiplication

Theorem 4 allows one to show that $s_2(f)$ must be super-linear for many operators $f = (f_1, \ldots, f_m)$ on two sets of variables X and Y. For this, it is enough that we can split the set $F = \{f_1, \ldots, f_m\}$ of functions computed by this operator into some number p of disjoint sets $F_1, \ldots, F_p$ such that, for some partition $X_1, \ldots, X_p$ of the variables in X, and for each $t = 1, \ldots, p$, we can obtain each single variable $y \in Y$ by taking some function $f \in F_t$ and fixing one its variable $x \in X_t$ to 1 and the rest of X to 0. (We say in this case that f *isolates* the variable y.) Then, by Proposition 1(ii), the set of subfunctions in each F_t with respect to the corresponding set of variables X_t must have entropy at least $|Y|$. By Theorem 4, we then have $s_2(f) \geq p|Y|$.

One of the most natural functions isolating *all* its single variables is a scalar product function $f(\mathbf{x}, \mathbf{y}) = x_1y_1 + x_2y_2 + \cdots + x_ry_r$; then $f(\mathbf{e}_i, \mathbf{y}) = y_i$ for all $i = 1, \ldots, r$. Hence, natural examples of operators of large entropy are sequences of particular scalar products. Many operators computing sequences of bilinear functions, including that of cyclic n-convolution considered in [5], fall in this general (scalar product) frame. We illustrate this with one important example—matrix product.

Given two $r \times r$ boolean matrices $X = (x_{i,j})$ and $Y = (y_{i,j})$ over a finite field $\mathbb{F}$, our goal is to compute their product $Z = X \cdot Y$ over $\mathbb{F}$. The corresponding operator $f = mult_n(X, Y)$ has $n = 2r^2$ input variables, arranged in two matrices, and consists of $n = r^2$ scalar products $f_{i,j} = \sum_{k=1}^{r} x_{i,k}y_{k,j}$ corresponding to the entries of the product matrix $Z = (z_{i,j})$. (This time indexes of variables as well as of computed functions are *pairs* of numbers.)

Since $mult_n$ is just a sequence of r^2 scalar products on $2r$ variables, $(2r)r^2 = 2n^{3/2}$ is a trivial upper bound, even in depth-1. If we put no restrictions on the depth, then Strassen's algorithm [19], improved in [2], gives a circuit of size $O(n^{6/5})$. The only know lower bound in the unrestricted case, however, is the lower bound $2.5 \cdot n$ proved in [4]. A lower bound $s_2(mult_n) = \Omega(n \log n)$ for depth-2, as well as nonlinear lower bounds for any constant depth, were proved in [15] using superconcentrators. For depth-2, entropy arguments yield a tight estimate $s_2(mult_n) = \Theta(n^{3/2})$.

Lemma 5. $\text{entropy}(mult_n) \geq n^{3/2}$.

Proof. Let $f = mult_n$, and let $\mathbf{e}_{i,k}$ be the boolean $r \times r$ matrix with precisely one 1 in the position (i, k). Since $f_{i,j} = \sum_{k=1}^{r} x_{i,k}y_{k,j}$, we have that $f_{i,j}(\mathbf{e}_{i,k}, Y) = y_{k,j}$ for all $j = 1, \ldots, r$. That is, for each $i, k \in [r]$, the i-th row $f_{i,1}(\mathbf{e}_{i,k}, Y), \ldots, f_{i,r}(\mathbf{e}_{i,k}, Y)$ of the product matrix $\mathbf{e}_{i,k} \cdot Y$ is just the k-th row $y_{k,1}, \ldots, y_{k,r}$ of Y.

Hence, if we take $X_i = \{x_{i,1}, \ldots, x_{i,r}\}$ (the i-th row of X) and $F_i = \{f_{i,1}, \ldots, f_{i,r}\}$ (the i-th row of the product matrix), then the corresponding set of subfunctions of F_i with respect to the variables in X_i,

$$\begin{Bmatrix} f_{i,1}(\mathbf{e}_{i,1}, Y) & f_{i,2}(\mathbf{e}_{i,1}, Y) & \cdots & f_{i,r}(\mathbf{e}_{i,1}, Y) \\ f_{i,1}(\mathbf{e}_{i,2}, Y) & f_{i,2}(\mathbf{e}_{i,2}, Y) & \cdots & f_{i,r}(\mathbf{e}_{i,2}, Y) \\ \vdots & \vdots & & \vdots \\ f_{i,1}(\mathbf{e}_{i,r}, Y) & f_{i,2}(\mathbf{e}_{i,r}, Y) & \cdots & f_{i,r}(\mathbf{e}_{i,r}, Y) \end{Bmatrix} = \begin{Bmatrix} y_{1,1} & y_{1,2} & \cdots & y_{1,r} \\ y_{2,1} & y_{2,2} & \cdots & y_{2,r} \\ \vdots & \vdots & & \vdots \\ y_{r,1} & y_{r,2} & \cdots & y_{r,r} \end{Bmatrix}$$

contains all $r^2 = n$ variables of Y. Together with Proposition 1(ii), this implies that, for each $i = 1, \ldots, r$, the entropy of F_i with respect to X_i is at least n. By Theorem 4, $\text{entropy}(f) \geq rn = n^{3/2}$. ∎

Remark (Limitations). How large can entropy of operators be? Recall that in the definition of entropy(f) of an (n,m)-operator f, we first split the inputs into p blocks $I_1,\ldots,I_p$ of some sizes $a_1 \le a_2 \le \ldots \le a_p$, and the outputs into p blocks $J_1,\ldots,J_p$ of some sizes $b_1,\ldots b_p$. Then we just take the sum of the entropies of the corresponding (to these blocks) sets of subfunctions. Say that a partition is *balanced* if $b_1 \ge b_2 \ge \ldots \ge b_p$. Note that the partition (into the rows) which we used for the matrix product is balanced—there all b_i's were even equal.

Since each of the sets $\{f_j(\mathbf{e}_i,\mathbf{y})\colon i \in I_t, j \in J_t\}$ can have at most $|I_i \times J_i| = a_i b_i$ functions, Proposition 1(i) implies that the entropy of this set cannot exceed $a_i b_i$. If the partition is balanced, then Chebyshev's inequality yields

$$\text{entropy}(f) \le \sum_{i=1}^{p} a_i b_i \le \frac{1}{p}\Big(\sum_{i=1}^{p} a_i\Big)\Big(\sum_{i=1}^{p} b_i\Big) \le \frac{nm}{p}.$$

On the other hand, we have a trivial upper bound entropy(f) $\le pn$. Substituting $p \ge$ entropy(f)$/n$ in the previous inequality, we obtain that entropy(f) $\le n\sqrt{m}$. Thus, at least with respect to balanced partitions, the entropy of any (n,m)-operator does not exceed $n\sqrt{m}$. In particular, for such partitions, matrix multiplication has the largest possible entropy $\Theta(n^{3/2})$ among all (n,n)-operators.

4. Open problems

As mentioned in the introduction, a lower bound $\Omega(n^{1+\epsilon})$ on the number of wires in a depth-2 circuit, computing an explicit linear operator $A\mathbf{x}$ over GF_2, would yield a nonlinear lower bound for log-depth circuits. To approach this problem, it is natural to first prove such a bound for *linear* depth-2 circuits, where we only allow linear functions (sums mod 2) as gates. It is well known that matrices A requiring $\Omega(n^2/\log n)$ wires in this restricted model exist. The situation with *explicit* bounds is, however, much worse. For circuits over the real field a lower bound $\Omega(n^{3/2})$ was proved in [17]. However in their result it is essential that they use large integers in the matrix. It remains an open problem to prove such a bound for 0-1 matrices. For GF_2 the largest bound is $\Omega(n\log^{3/2} n)$ [1, 11, 13]. In would be therefore interesting to extend the entropic approach to depth-2 circuits computing *linear* operators.

A less famous problem about depth-2 circuits, related to another old problem in circuit complexity (proving lower bounds for ACC circuits), is the following one.

A *symmetric* depth two circuit is a depth two circuit, where the gates on the middle layer compute ORs of their inputs, and each output gate computes the same symmetric function of its inputs. That is, each output gate gives the value 1 iff the number of 1's in its input belongs to some specified (for the whole circuit) subset S of natural numbers. We also assume that there are no direct wires from an input to an output node.

Say that a circuit computing a set $F = \{f_1,\ldots,f_n\}$ of boolean functions *represents* a given boolean $n \times n$ matrix $A = (a_{ij})$ if, for every i and j, $f_i(\mathbf{e}_j) = a_{ij}$. That is, the circuit is only required to be correct on inputs with precisely one 1. Let $\text{sym}_2(A)$ be the minimum number of nodes on the middle layer in a symmetric depth-2 circuit representing A. That is, now we count nodes, not wires.

Remark. Note that the fact, that the circuit is allowed to output arbitrary values on inputs $\mathbf{x}$ with more than one 1, is crucial. So, for example, every circuit computing the linear operator $A\mathbf{x}$ over GF_2 or a set $F = \{f_1,\ldots,f_n\}$ of boolean functions $f_i(\mathbf{x}) = \bigvee_{j=1}^{n} a_{ij}x_j$ (a $(\wedge,\vee)$-boolean matrix-vector product $A\mathbf{x}$) represents the matrix A. By Proposition 1(iii), each such

circuit must have at least entropy(F) nodes on the middle layer. Hence, already the identity matrix A requires then n nodes even if arbitrary boolean functions can be used as gates.

Simple counting shows that matrices with $\mathrm{sym}_2(A) = \Omega(n)$ exist. The problem, due to Yao [21], is to exhibit an *explicit* boolean matrix A with large $\mathrm{sym}_2(A)$. In terms of set intersection representations of matrices, this problem was re-stated by Pudlák and Rödl in [13] (see Problem 10). To see the equivalence between $\mathrm{sym}_2(A)$ and their measure, just associate with each output node i and each input node j the sets U_i and V_j of all their neighbors on the middle layer. Then $a_{ij} = f_i(\mathbf{e}_j) = 1$ iff $|U_i \cap V_j| \in S$.

What we need is an explicit boolean $n \times n$ matrix A with $\mathrm{sym}_2(A) = 2^{(\log\log n)^{\omega(1)}}$. Together with the results of Yao [21], and Beigel and Tarui [3], this would yield a super-polynomial lower bound for *ACC circuits*. These are constant depth unbounded fanin circuits over a basis consisting of AND, OR and a finite number of modulo-counting functions: each such function gives the value 1 iff the number of 1's in the input is not divisible by p. When p is a prime, exponential lower bounds were proved by Razborov [16] and Smolensky [18]. However, the case of composite moduli p (even when one moduli $p = 6$ is allowed) remains widely open.

References

[1] N. Alon, M. Karchmer, A. Wigderson, Linear circuits over GF(2), SIAM. J. Comput. 19(6) (1990) 1064-1067.

[2] D. Coppersmith, S. Winograd, Matrix multiplications via arithmetic progressions, J. Symb. Comp. 9 (1990) 251-280.

[3] R, Beigel, J. Tarui On ACC, Computational Complexity 4 (1994) 350–366.

[4] N. H. Bshouty, A lower bound for matrix multiplication, SIAM J. Comput. 18 (1982) 759-765.

[5] D. Yu. Cherukhin, The lower estimate of complexity in the class of schemes of depth 2 without restrictions on a basis, Vestnik Moscow University, Ser. 1, Mathematika 60(4) (2005) 54-56 (in Russian).

[6] D. Dolev, C. Dwork, N. Pippenger, A. Wigderson, Superconcentrators, generalizer and generalized connectors with limited depth, in: Proc. 15th STOC (1983) 42-51.

[7] P. Erdős, Some remarks on the theory of graphs, Bull. Amer. Math. Soc. 51 (1947) 292-294.

[8] E. I. Nechiporuk, On a Boolean function, Soviet Math. Doklady 7(4) (1966) 999-1000.

[9] N. Pippenger, Superconcentrators, SIAM J. Comput. 6 (1977) 298-304.

[10] N. Pippenger, Superconcentrators of depth 2, J. Comput. Syst. Sci. 24 (1982) 82-90.

[11] P. Pudlák, Communication in bounded depth circuits, Combinatorica 14 (2) (1994) 203-216.

[12] P. Pudlák, P. Savický, On shifting networks, Theoretical Comput. Sci. 116 (1993) 415-419.

[13] P. Pudlák, V. Rödl, Some combinatorial-algebraic problems from complexity theory, Discrete Math. 136 (1994) 253-279.

[14] J. Radhakrishnan, A. Ta-Shma, Bounds for dispersers, extractors, and depth-two superconcentrators, SIAM J. Discrete Math. 13(1) (2000) 2-24.

[15] R. Raz, A. Shpilka, Lower bounds for matrix product in bounded depth circuits with arbitrary gates, SIAM J. Comput. 32(2) (2003) 488-513.

[16] A. A. Razborov, Bounded-depth formulae over the basis $\{\&, \oplus\}$ and some combinatorial problem, in: S.I. Adian (ed.), Problems of Cybernetics, Complexity Theory and Applied Mathematical Logic (1988) 149–166 (in Russian).

[17] V. Shoup, R. Smolensky, Lower bounds for polynomial evaluation and interpolation problems, Comput. Complexity 6(4) (1997) 301-311.

[18] R. Smolensky, Algebraic methods in the theory of lower bounds for Boolean circuit complexity, in: Proc. 19th STOC (1987) 77–82.

[19] V. Strassen, Die Berechnungskomplexität von elementarsymmetrischen Funktionen und von Interpoliationskoefizienten, Numer. Math. 20 (1973) 238-251.

[20] L. Valiant, Graph-theoretic methods in low-level complexity, in: Proc. 6th MFCS, Springer Lect. Notes in Comput. Sci. 53 (1977) 162-176.

[21] A. C. Yao, On ACC and threshold circuits, in: Proc. 31th FOCS (1990) 619–627.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/