Lower Bounds on Threshold and Related Circuits via Communication Complexity

V. P. Roychowdhury * A. Orlitsky † K. Y. Siu ‡

Abstract

Communication-complexity definitions and arguments are used to derive linear ($\Omega(n)$) and almost-linear ($\Omega(n/\log n)$) lower bounds on the size of circuits implementing certain functions. The techniques utilize only basic features of the gates used and of the functions implemented, hence apply to a large class of gates (including unbounded fan-in AND/OR, threshold, symmetric, and generalized symmetric) and to a large class of functions (including equality, comparison, and inner product mod 2). Each of the bounds derived is shown to be tight for some functions. Some applications to threshold-circuit complexity are indicated. The results generalize and in some cases strengthen results in [1, 2].

Index Terms: Linear/Almost-Linear Circuit-size Lower Bounds; Communication Complexity; Threshold gates/circuits; Symmetric gates/circuits; Equality, Comparison and Inner Product mod 2 Boolean functions.

*School of Electrical Eng., Purdue University, West Lafayette, IN 47907.

†AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 07974.

‡Department of Electrical and Computer Eng., University of California at Irvine, Irvine, CA 92717.

1 Introduction

We describe the model, review known results, and introduce techniques and results presented in this paper.

Gates, Circuits, and Complexity

An *n-variable Boolean function* is a mapping

$$f : \{0,1\}^{\{1,\ldots,n\}} \to \{0,1\} \ .$$

An element of $\{1,\ldots,n\}$ is a *variable*. An element of $\{0,1\}^{\{1,\ldots,n\}}$, viewed as a value assignment to the variables, is an *input*. If x is an input, then $f(x)$ is the corresponding *output* of f.

An *n-variable gate* is a physical device computing a single n-variable function. The input variables of a gate can be permuted, omitted, or repeated, hence we identify the gate with the set of functions derived by such operations. For example, the set of functions implementable by a gate computing the 4-variable function $(x \bigwedge y) \bigvee (z \bigwedge w)$, where $\bigwedge$ is logical "AND" and $\bigvee$ logical "OR," includes functions such as $(x \bigwedge z) \bigvee (y \bigwedge w)$, $(x \bigwedge y) \bigvee (x \bigwedge y) \equiv x \bigwedge y$ and $(y \bigwedge y) \bigvee (y \bigwedge y) \equiv y$.

We usually consider a set, or a *family*, of gates. We identify the family with the union of the function sets corresponding to each of its gates.

Let $\mathcal{G}$ be a family of gates. A circuit whose gates are all from $\mathcal{G}$ is a *$\mathcal{G}$-circuit*. The *size* of a circuit is the number of gates it contains and its *depth* is the maximum number of gates along a path from an input to an output. The *$\mathcal{G}$-circuit complexity* $C_{\mathcal{G}}(f)$ of f is the size of the smallest $\mathcal{G}$-circuit that computes f. In principle, some function may not be computed by a $\mathcal{G}$-circuit. However, every gate family considered here forms a complete basis, and hence $C_{\mathcal{G}}(f)$ is always defined.

The circuit complexity of functions has many theoretic and practical applications. Therefore, several gate families have been extensively investigated. They include:

AND/OR/NOT gates ($\mathcal{AON}$) These gates perform logical "AND" or "OR" of their, possibly negated, inputs. AND/OR/NOT gates come in two varieties: constant fan-in gates and unbounded fan-in gates. The bounds we prove apply to both.

Symmetric gates ($\mathcal{SYM}$) Gates of the form $g(\sum_{i=1}^{n} x_i)$ for arbitrary binary functions g. These gates compute some binary function of their input sum.

One type of a symmetric gate is a mod_m gate. It computes a binary function of the form $g((\sum_{i=1}^{n} x_i) \bmod m)$ for some constant integer m.

Threshold gates ($\mathcal{TH}$) Gates of the form $\mathrm{sgn}(\sum_{i=1}^{n} w_i x_i - T)$ where T is an arbitrary *threshold*, the w_i's are integer *weights*, and $\mathrm{sgn}(x)$ is 1 if $x \geq 0$ and 0 otherwise.

In the analysis we distinguish between general (arbitrary weight) threshold gates and *polynomial-weight* threshold gates where the w_is are restricted to be polynomial in n.

Generalized symmetric gates ($\mathcal{GS}$) Gates of the form $g(\sum_{i=1}^{n} w_i x_i)$ for arbitrary function g and *weights* w_i that are polynomial in n.

The weights are restricted to be polynomial because every function can be computed by a single generalized symmetric gate with arbitrary weights.

Note that every AND/OR/NOT gate is also a polynomial-weight threshold gate and that any polynomial-weight threshold gate as well as any symmetric gate is also a generalized-symmetric gate.

Related Results and Motivation

Much research has gone into estimating $C_{\mathcal{G}}(f)$ for various functions and gate families [3]. The strongest results apply to bounded-depth circuits. For constant depth AND/OR/NOT circuits and mod_p circuits (where p is prime), [4, 5, 6] established exponential-size lower bounds for specific functions such as the parity. For more powerful circuits, less is known. For example, [7] proved an exponential-size lower bound on the size of *depth-2* threshold circuits implementing the n-variable *inner product mod 2* function:

$$\mathrm{IP}(x_1,\ldots,x_{\frac{n}{2}},y_1,\ldots,y_{\frac{n}{2}}) = \begin{cases} 1 & \text{if } \sum_{i=1}^{\frac{n}{2}} x_i \wedge y_i \text{ is odd,} \\ 0 & \text{otherwise.} \end{cases}$$

However, this bound applies only when the weights in the second layer are restricted to be polynomial. No superlinear lower bounds are known for depth-2 threshold circuits with exponential weights in the second layer, or for depth-3 threshold circuits with polynomial weights.

For unrestricted-depth unbounded-fan-in circuits even weak lower bounds, such as linear or logarithmic in the number of input variables, are considered difficult to prove [3, 1]. For example, an $\Omega(\log n)$ lower bound on the size of threshold circuits computing the parity of n bits is shown in [3]. Only recently have linear/almost-linear lower bounds been established for circuits with gates of unbounded fan-in. A linear-size lower bound on circuits where each gate computes a commutative and associative function, was given in [8]. However, the family of gates is too restrictive to apply to symmetric or threshold circuits.

Recently, [1] established an $\Omega(n/\log n)$ lower bound on the size of symmetric-gate circuits computing the n-variable *equality* function:

$$\mathrm{EQ}(x_1,\ldots,x_{\frac{n}{2}},y_1,\ldots,y_{\frac{n}{2}}) = \begin{cases} 1 & \text{if } x_i = y_i \text{ for all } 1 \le i \le \frac{n}{2}, \\ 0 & \text{otherwise.} \end{cases}$$

Novel techniques such as analytic-function interpolation of Boolean functions and the differential dimension were used. More recently [2] proved a linear lower bound $(n/4)$ on the size of arbitrary-weight threshold circuits computing the n-variable IP .

Techniques and Results in this Paper

Using communication-complexity concepts and techniques, we derive linear and almost-linear lower bounds on the size of circuits implementing certain functions. This approach utilizes only basic features of the gates used, hence the bounds hold for general families of gates of which the symmetric and threshold gates considered in [1, 2] are special cases. Thus communication complexity arguments serve to generalize known lower bounds and unify their proofs.

In the next section we define the *decomposition number* and the *largest monochromatic rectangle* of a function. These are simple attributes that have proven useful in analyzing the communication complexity of various functions.

In Section 3 we consider *polynomially-rectangular* gates. These gates, which include symmetric, generalized symmetric, and polynomial-weight threshold gates, compute functions with small decomposition numbers. We show that functions computed by small-size circuits of polynomially-rectangular gates have small decomposition numbers. It follows that functions with high decomposition numbers require large circuits. We then use some effective techniques that have been developed to lower bound decomposition numbers to prove almost-linear lower bounds on the circuit complexity of several functions.

In Section 4 we strengthen the results for *triangular gates*. These gates, which include all threshold gates, compute functions with large monochromatic rectangles. We show that any function computed by a small circuit of triangular gates contains a large monochromatic rectangle. Therefore, functions with only small monochromatic rectangles require large, in some cases linear-size, circuits.

We illustrate the results using the equality and the inner product mod 2 functions defined earlier in this section. The bounds we derive imply:

1. Any implementation of n-variable EQ or IP by generalized symmetric gates requires about $n/\log n$ gates. Namely, if the weights are bounded by n^k, then
$$\frac{1}{4(k+1)}\frac{n}{\log n} \le C_{\mathcal{GS}}(\text{EQ}), C_{\mathcal{GS}}(\text{IP}) \le \frac{\log 3}{2k}\frac{n}{\log n} \ .$$

2. Any implementation of n-variable EQ or IP by symmetric gates requires at least $\frac{n}{\log n}$ gates:
$$C_{\mathcal{SYM}}(\text{EQ}), C_{\mathcal{SYM}}(\text{IP}) \ge \frac{n}{4\log n} \ .$$

3. Any implementation of n-variable EQ or IP by AND/OR/NOT gates requires about n gates:
$$\frac{n}{2\log 3} \le C_{\mathcal{AON}}(\text{EQ}), C_{\mathcal{AON}}(\text{IP}) \le 2n \ .$$

4. Any implementation of n-variable IP by threshold gates requires about n gates.
$$\frac{1}{4}n \le C_{\mathcal{TH}}(\text{IP}) \le \frac{3}{4}n + 1 \ .$$

   Both upper and lower bounds apply to arbitrary- and polynomial-weight threshold circuits.

Note that the bounds in (1), (2), and (3) are tight up to a small multiplicative factor.

Related to EQ is the n-variable *comparison* function:
$$\text{COMP}(x_1, \ldots, x_{\frac{n}{2}}, y_1, \ldots, y_{\frac{n}{2}}) = \begin{cases} 1 & \text{if } \sum_{i=1}^{\frac{n}{2}} 2^i x_i \ge \sum_{i=1}^{\frac{n}{2}} 2^i y_i, \\ 0 & \text{otherwise.} \end{cases}$$

Although we do not discuss COMP explicitly, it shares the same size bounds as EQ .

2 Communication Complexity Arguments

As before, let $f : \{0,1\}^{\{1,\ldots,n\}} \to \{0,1\}$ be an n-variable Boolean function. Recall that an element of $\{1,\ldots,n\}$ is a *variable* and an element of $\{0,1\}^{\{1,\ldots,n\}}$ is an *input*. If X is a set of variables then an element of $\{0,1\}^X$ is a value assignment to the variables in X and is called an X-*input*.

Let $\{X,Y\}$ partition the set of variables ($X \bigcup Y = \{1,\ldots,n\}$ and $X \bigcup Y = \emptyset$). An X-input x together with a Y-input y correspond in an obvious way to an input which we call the *joint input* and denote by (x,y). In the same way, the set of all inputs corresponds to the Cartesian product $\{0,1\}^X \times \{0,1\}^Y$. We can therefore associate with the function f and the partition $\{X,Y\}$ a matrix $M_{f,X,Y}$. It has $2^{|X|}$ rows, each indexed by an X-input, $2^{|Y|}$ columns, each indexed by a Y-input, and

$$M_{f,X,Y}(x,y) = f(x,y) \ .$$

An $\{X,Y\}$-*rectangle* is a Cartesian product $A \times B$ where A is a set of X-inputs and B is a set of Y-inputs. The *size* of the rectangle is $|A| \cdot |B|$, the number of inputs it contains. An $\{X,Y\}$-*decomposition* is a partition of $\{0,1\}^X \times \{0,1\}^Y$ into $\{X,Y\}$-rectangles. The *size* of the decomposition is the number of rectangles in the partition. A set of inputs is *f-constant* if f assigns the same value to all its elements. An *f-constant* $\{X,Y\}$-*decomposition* is an $\{X,Y\}$-decomposition whose rectangles are all f-constant.

Rectangles play a major role in the following *communication complexity* problem. As before, let f be an n-variable Boolean function and $\{X,Y\}$ a partition of the variables. A person $P_{\mathcal{X}}$ knows an X-input, a person $P_{\mathcal{Y}}$ knows a Y-input, and they communicate according to a predetermined protocol in order to find the value of f on their joint input. We are interested in $\hat{C}(f,X,Y)$, the number of bits $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ must transmit for the worst input.

As shown by [9],

1. Every protocol induces an $\{X,Y\}$-decomposition.

2. If the protocol always produces the correct answer, this decomposition is f-constant.

3. The number of bits required by the protocol for the worst input is at least the logarithm[1]

[1] All logarithms are to the base 2.

of the size of the decomposition.

Let $\rho_{f,X,Y}$ be the smallest size of an f-constant $\{X, Y\}$-decomposition. From the above,

$$\hat{C}(f, X, Y) \geq \log \rho_{f,X,Y} \ . \tag{1}$$

Aho, Ullman, and Yanakakis [10] showed that this bound is not far from being tight:

$$\hat{C}(f, X, Y) \leq \log^2 \rho_{f,X,Y} \ .$$

For that reason, several simple methods were introduced to lower bound $\rho_{f,X,Y}$ for arbitrary f, X, and Y.

Largest f-constant rectangle

Let $L_{f,X,Y}$ be the size of the largest f-constant $\{X, Y\}$-rectangle. Clearly,

$$\rho_{f,X,Y} \geq \frac{2^n}{L_{f,X,Y}} \ .$$

Fooling set

An f-constant subset S of $\{0,1\}^X \times \{0,1\}^Y$ is an $\{X,Y\}$-*fooling set* if $(x_1, y_1), (x_2, y_2) \in S$ implies that either $f(x_1, y_2)$ or $f(x_2, y_1)$ differs from the common value of f over S. Let $F_{f,X,Y}$ be the size of the largest $\{X, Y\}$-fooling set. An f-constant $\{X, Y\}$-rectangle contains at most one element of a given $\{X, Y\}$-fooling set, hence:

$$\rho_{f,X,Y} \geq F_{f,X,Y} \ .$$

Rank

The matrix representing the indicator function of a rectangle has rank 1, and ranks are subadditive under matrix addition. Melhorn and Schmidt [11] concluded that under any field

$$\rho_{f,X,Y} \geq \mathrm{rank}(M_{f,X,Y}) \ .$$

In our applications, we can choose the most advantageous partition of the input variables. We therefore define the *decomposition number* of f,

$$\rho_f \stackrel{\mathrm{def}}{=} \max\{\rho_{f,X,Y} : \{X, Y\} \text{ partitions } \{1, \ldots, n\}\} \ ,$$

to be the number of rectangles needed in the variable partition that yields the strongest bound in (1). We use the methods above to lower bound the decomposition number of our two functions.

Example 1 We show that the decomposition numbers of both EQ and IP are larger than $2^{\frac{n}{2}}$. In the following, $X = \{1, \ldots, \frac{n}{2}\}$ and $Y = \{\frac{n}{2}+1, \ldots, n\}$. Every $\frac{n}{2}$-bit sequence corresponds in an obvious way to an X-input and to a Y-input. We can therefore talk about the joint input (x, x) where $x \in \{0,1\}^{\frac{n}{2}}$.

Equality The set $\{(x,x) : x \in \{0,1\}^{\frac{n}{2}}\}$ is an $\{X,Y\}$-fooling set of size $2^{\frac{n}{2}}$, implying that $\rho_{\mathrm{EQ},X,Y} \geq 2^{\frac{n}{2}}$. In fact, $\rho_{eq} = \rho_{\mathrm{EQ},X,Y} = 2^{\frac{n}{2}+1}$.

Inner product mod 2 $M_{\mathrm{IP},X,Y}$ has full rank over the reals, hence $\rho_{\mathrm{IP}} \geq 2^{\frac{n}{2}}$. □

3 Rectangular gates

The last section was motivated by the notion that a function with a high decomposition number is "complicated." To show that computing such a function requires many gates, we now show that the gates used are "simple," that is, they can be decomposed into a small number of rectangles.

A function f is *r-rectangular* for some integer r if for every variable partition $\{X,Y\}$ there is an f-constant $\{X,Y\}$-decomposition consisting of at most r rectangles. Namely, if

$$\rho_f \leq r \ .$$

Let $p : \mathcal{Z}^+ \to \mathcal{Z}$. A family $\mathcal{G}$ of functions is *p-rectangular* if for every $m \leq n$, all m-variable functions in $\mathcal{G}$ are $p(n)$-rectangular. The family is *polynomially-rectangular* if it is p-rectangular for some polynomial p. These definition apply to gates and families of gates via their underlying functions. The next lemma, its simple proof omitted, provides a basic tool for proving that a function is r-rectangular.

Lemma 1 Let f be a Boolean function and let $\{X,Y\}$ partition the set of variables. If $f(x,y)$ can be expressed as $h(g_1(x), g_2(y))$ then

$$\rho_{f,X,Y} \leq |g_1| \cdot |g_2|$$

where $|g_i|$ is the size of the range of g_i. □

To prove that a function is r-rectangular we apply the lemma to *all* possible partitions of the variables.

Example 2 We show that the gate families mentioned in the introduction are polynomially rectangular. In the following, $\{X,Y\}$ is an arbitrary partition of $\{1,\ldots,n\}$.

AND/OR/NOT gates

$$\bigwedge_{i\in\{1,\ldots,n\}} x_i = (\bigwedge_{i\in X} x_i)\bigwedge(\bigwedge_{i\in Y} x_i) ,$$

hence the lemma implies that every AND gate is 4-rectangular (three rectangles suffice). The same holds for NOT gates.

Symmetric gates

$$f(x,y) = h\left(\sum_{i\in X} x_i, \sum_{i\in Y} y_i\right),$$

hence

$$\rho_{f,X,Y} \leq (|X|+1)\cdot(|Y|+1) \leq \left(\frac{n}{2}+2\right)^2 .$$

Generalized symmetric gates

$$f(x,y) = h\left(\sum_{i\in X} w_i x_i + \sum_{i\in Y} w_i x_i\right) .$$

where the w_i's are bounded by some polynomial $p(n)$. The first sum attains at most $(|X|+1)\cdot p(n)$ values and likewise for the second, hence f is $(\frac{n}{2}+1)^2\cdot p^2(n)$-rectangular. It follows that the family of generalized symmetric functions (and in particular, of polynomial-weight threshold circuits) is polynomially rectangular. $\Box$

Lemma 2 Let $\mathcal{G}$ be a p-rectangular family of gates. If an $\mathcal{G}$-circuit consisting of k gates computes an n-variable function f, then

$$\rho_f \leq (p(n))^k .$$

Proof: Order the gates in the circuit so that if $i < j$ then gate i does not follow gate j. Let g_j denote the function computed by gate j. We prove by induction on j that the vector-valued function $G_j \stackrel{\text{def}}{=} (g_1, g_2, \ldots, g_j)$ has $\rho_{G_j,X,Y} \leq (p(n))^j$ for all variable partitions $\{X,Y\}$. The lemma will follow.

The induction basis holds by definition; suppose it holds for j, and consider the $(j+1)$st gate. Let $\{X,Y\}$ be a variable partition. There is a G_j-constant $\{X,Y\}$-decomposition consisting of at most $(p(n))^j$ rectangles. Let R be a rectangle in this decomposition. Over R, all of $g_1,\ldots,g_j$ are constant, hence the $(j+1)$st gate coincides with a $p(n)$ rectangular function of the original variables. Therefore R can be partitioned into $p(n)$ G_{j+1}-constant $\{X,Y\}$ rectangles, and the induction step follows. $\Box$

Corollary 1 Let $\mathcal{G}$ be a p-rectangular family of gates. For every n-variable function f,

$$C_{\mathcal{G}}(f) \geq \frac{\log \rho_f}{\log p(n)} . \qquad \Box$$

We apply the corollary to lower bound the number of gates needed to implement our two functions.

Corollary 2

1. For circuits consisting of AND, OR, and NOT gates:

$$\frac{n}{2\log 3} \leq C_{\mathcal{AON}}(\mathrm{EQ}), C_{\mathcal{AON}}(\mathrm{IP}) \leq 2n .$$

2. For circuits consisting of generalized symmetric gates:

$$C_{\mathcal{GS}}(\mathrm{EQ}), C_{\mathcal{GS}}(\mathrm{IP}) \in \Theta\left(\frac{n}{\log n}\right) .$$

   More specifically, if the weights are bounded by n^k, then

$$\frac{1}{4(k+1)}\frac{n}{\log n} \leq C_{\mathcal{GS}}(\mathrm{EQ}), C_{\mathcal{GS}}(\mathrm{IP}) \leq \frac{\log 3}{2k}\frac{n}{\log n} .$$

3. For circuits consisting of symmetric gates:

$$C_{\mathcal{SYM}}(\mathrm{EQ}), C_{\mathcal{SYM}}(\mathrm{IP}) \geq \frac{n}{4\log n} .$$

Proof: All six lower bounds follow from Corollary 1 as both EQ and IP have decomposition numbers of at least $2^{\frac{n}{2}}$. The upper bounds in (1) follow from a simple construction. To prove the upper bounds in (2) we implement EQ as a depth-2 threshold circuit, yielding a simple circuit with slightly more gates than the upper bound. We implement IP as a depth-3 generalized symmetric circuit (the next section shows it cannot be implemented using less than n threshold gates).

Let $m = 2\lfloor k \log n \rfloor$. Clearly, m-variable COMP can be written as

$$\mathrm{COMP}(x_1, \ldots, x_{\frac{m}{2}}, y_1, \ldots, y_{\frac{m}{2}}) = \mathrm{sgn}\left(\sum_{i=1}^{m/2} 2^i (x_i - y_i)\right), \tag{2}$$

thus can be implemented by a single threshold gate with weights of at most n^k. For $i = 1, \ldots, \lceil n/m \rceil$, let $x^i = x_{(i-1)\cdot m/2+1}, \ldots, x_{i \cdot m/2}$ and $y^i = y_{(i-1)\cdot m/2+1}, \ldots, y_{i \cdot m/2}$. Then,

$$\mathrm{EQ}(x^i, y^i) = \mathrm{COMP}(x^i, y^i) + \mathrm{COMP}(y^i, x^i) - 1 \ .$$

Hence, m-variable EQ can be implemented by a depth-2 threshold circuit with weights of at most n^k and where the top gate is just a weighted sum of the first-level outputs (without a threshold). Finally, observe that

$$\mathrm{EQ}(x_1, \ldots, x_{\frac{n}{2}}, y_1, \ldots, y_{\frac{n}{2}}) = \bigwedge\nolimits_{i=1}^{\lceil n/m \rceil} \mathrm{EQ}(x^i, y^i) \ .$$

Since any AND is just the sum of its variables with an appropriate threshold, this gate can be combined with the second layer above to derive a depth-2 circuit for EQ of size $2\lceil n/2k \log n \rceil + 1$. When generalized symmetric gates are used instead of threshold gates, the number of gates can be reduced to $\lceil n/2k \log n \rceil + 1$.

When trying to meet the lower bound for IP, we cannot use threshold gates as we did for EQ. The next section shows that any threshold circuit for IP (even with exponential weights) has at least linear size. Yet, we can use the circuit structure applied to EQ. Every $(k \log n)$-variable function, in particular $\mathrm{IP}(x_1, \ldots, x_{k \log n/2}, y_1, \ldots, y_{k \log n/2})$, can be computed by a single generalized symmetric gate with weights of at most n^k. Use $\lceil n/k \log n \rceil$ generalized symmetric gates to compute the partial IP's, then use a single (symmetric) gate to compute their parity. $\Box$

A note on COMP **:** Equation (2) shows that n-variable COMP can be computed by a single threshold gate with exponential weights. However, if the weights are polynomially bounded, then as noted in the introduction, the lower bound on EQ can be modified to show that $C_{\mathcal{GS}}(\mathrm{COMP}) \geq \Omega(n/\log n)$. Thus a single threshold gate with polynomial weights cannot compute COMP . We next show that the lower bound can be met by a depth-3 polynomial-weight threshold circuit. It is not known whether the lower bound can be met by a polynomial-weight threshold circuit of depth two.

Let $m = 2\lceil \log n \rceil$. For $i = 1, \ldots, \lceil n/m \rceil$, let

$$C_i = \mathrm{sgn}\left(\sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j (x_j - y_j) \right),$$

and

$$\tilde{C}_i = \mathrm{sgn}\left(\sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j (x_j - y_j) - 1 \right).$$

Note that both C_i and $\tilde{C}_i$ can be computed with threshold gates of polynomially bounded weights. Further,

$$C_i = 1 \text{ iff } \sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j x_j \geq \sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j y_j$$

and

$$\tilde{C}_i = 1 \text{ iff } \sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j x_j > \sum_{j=(i-1)\cdot m/2+1}^{i\cdot m/2} 2^j y_j$$

Define Boolean expressions

$$B_{\lceil n/m \rceil} = \tilde{C}_{\lceil n/m \rceil}$$

$$B_k = \tilde{C}_k \bigwedge_{j=k+1}^{\lceil n/m \rceil} C_j \qquad \text{for } k = 2, ..., \lceil n/m \rceil - 1$$

$$\text{and} \qquad B_1 = \bigwedge_{j=1}^{\lceil n/m \rceil} C_j$$

It is straightforward to see that

$$\mathrm{COMP}(x_1, \ldots, x_n, y_1, \ldots, y_n) = \bigvee_{j=1}^{n} B_j.$$

The first layer of our circuit for the COMP function has $O(n/\log n)$ gates computing the C_i and $\tilde{C}_i$. With these computed values as inputs, the second layer has $O(n/\log n)$ gates each computing the B_j. Finally the output gate computes the OR ($\bigvee$) of all the B_j's. The total number of gates is $O(n/\log n)$.

4 Triangular Gates

A matrix is *strictly triangular* if all its rows and columns are nondecreasing. In a strictly triangular Boolean matrix, the sets of 1's and 0's resemble a (possibly truncated) triangle, hence the name. A matrix is *triangular* if its rows and columns can be permuted so that the resulting matrix is strictly triangular.

Lemma 3 (Alternative Definition) A binary matrix is triangular if and only if it contains no 2 by 2 rectangle of the form

$$\begin{matrix} 0 & 1 \\ 1 & 0 \end{matrix} \quad \text{or} \quad \begin{matrix} 1 & 0 \\ 0 & 1 \end{matrix}$$

(recall that a rectangle need not be contiguous).

Proof: Row and column permutations preserve this non-containment property, so "only if" is clear. For the other direction, permute the rows so that the number of 1's in each row is non decreasing, then permute the columns so that the number of 1's in each column is non decreasing. The resulting matrix is strictly triangular for if in some column a 1 appears above a 0, then, as the numbers of 1's does not decrease with the rows, there must be another column where in the same locations a 0 appears above a 1, contradicting the non-containment assumption. $\square$

Some properties of triangular matrices are apparent:

1. Every submatrix of a triangular matrix is triangular.

   Proof: Obvious by either definition.

2. Every triangular matrix contains a constant rectangle of 1/4 the size.

   Proof: Permute the rows/columns till you get a strictly-triangular matrix. Consider the mid point (x, y). If the (x, y)th element of the matrix is 0 then the rectangle above and to the left of (x, y) is all 0, otherwise, the rectangle to the right and below (x, y) is all 1.

3. Every submatrix of a triangular matrix contains a constant rectangle of 1/4 its size.

   Proof: Combine properties 1 and 2.

An n-variable function f is *triangular* if $M_{f,X,Y}$ is triangular for *all* $\{X, Y\}$-partitions of the variables. A family of functions is triangular if all the functions in the family are. The definition applies to gates and families of gates via the underlying functions.

Example 3 ([2]) Threshold gates (and in particular AND and OR gates) are triangular. We use the Lemma 3. Let $f(x,y) = \mathrm{sgn}(\sum u_i x_i + \sum v_i y_i)$. Suppose that $f(x,y) = f(x',y') = 1$ and that $f(x,y') = f(x',y) = 0$. Then $\sum u_i x_i + \sum v_i y_i > \sum u_i x_i + \sum v_i y'_i$ while $\sum u_i x'_i + \sum v_i y_i < \sum u_i x'_i + \sum v_i y'_i$. Impossible. $\square$

Recall that $L_{f,X,Y}$ was defined to be the size of the largest f-constant $\{X,Y\}$-rectangle. Define L_f to be $L_{f,X,Y}$ for the most advantageous partition of the variables:

$$L_f \stackrel{\text{def}}{=} \min\{L_{f,X,Y} : \{X,Y\} \text{ partitions } \{1,\ldots,n\}\ \}$$

Lemma 4 If a circuit consisting of k triangular gates computes a function f then

$$L_f \geq \frac{2^n}{4^k} \ .$$

Proof: As in Lemma 2, order the gates in the circuit so that if $i < j$ then gate i does not follow gate j. Let g_j denote the function computed by gate j. We prove by induction on j that the vector-valued function $G_j \stackrel{\text{def}}{=} (g_1, g_2, \ldots, g_j)$ has $L_{G_j,X,Y} \geq \frac{2^n}{4^k}$ for all variable partitions $\{X,Y\}$. The lemma will follow.

The induction basis holds by property (2) above. Suppose it holds for j, and consider the $(j+1)$st gate. Let $\{X,Y\}$ be a variable partition. By induction hypothesis, there is a G_j-constant $\{X,Y\}$-rectangle R of size $2^n/4^k$. Over R, the outputs of the first j gates are fixed, hence the input to the $(j+1)$st gate varies only with the original inputs. It follows that over R the $(j+1)$st gate coincides with a triangular function whose inputs are the original inputs. By property (3), there must be a subrectangle of R of size $\geq |R|/4$ over which the $(j+1)$st gate has a constant output. $\square$

Corollary 3 For every function f and every family $\mathcal{G}$ of triangular gates,

$$C_{\mathcal{G}}(f) \geq \frac{n - \log L_f}{2} \ . \qquad \square$$

Example 4 Let $X = \{1, \ldots, \frac{n}{2}\}$ and $Y = \{\frac{n}{2}+1, \ldots, n\}$. Lindsey [12] showed that the largest IP-constant $\{X,Y\}$-rectangles are of size at most $2^{\frac{n}{2}}$. Hence

$$C_{\mathcal{TH}}(\mathrm{IP}) \geq \frac{n}{4} \ . \qquad \square$$

The bound on $C_{\mathcal{TH}}(\mathrm{IP})$ is asymptotically tight too. A simple depth-3 circuit computes IP using $\frac{3}{4}n+1$ polynomial-weight threshold gates. In a sense, this circuit is depth optimal too. [7] showed that every depth-2 threshold circuit for IP has exponential size if the weights at the second layer are polynomial. It is not known whether there is a polynomial size depth-2 threshold circuit for IP when exponential weights are allowed at the second layer.

5 Application To Threshold Circuits

We briefly discuss some applications of the results and techniques discussed in the previous sections to threshold-circuit complexity.

Depth-Weight Tradeoffs in Threshold Circuits

Recent results [13] have shown that any depth-d threshold circuit (with arbitrary weights) can be simulated by a depth-$(d+1)$ polynomial-weight threshold circuit with only a polynomial factor increase in size (for fixed d). However, no upper- or lower-bounds have been shown for the degree of this polynomial.

One can implement the n-variable EQ using only 3 threshold gates in depth-2. Yet Corollary 2 gave a lower bound of $\Omega(n/\log n)$ on the size of any polynomial-weight threshold circuits for EQ . We therefore have:

Corollary 4 There are n-variable functions whose polynomial-weight threshold-circuit complexity (regardless of depth) is at least $n/\log n$ times larger than their unrestricted-weight depth-3 threshold-circuit complexity. $\square$

Weighted-Sum gates

In our discussions, we often observed that the output gate of a given threshold circuit does not always require the sgn function usually associated with a threshold gate. A gate that computes a linear combination $\sum w_i x_i$ of its inputs (without taking a threshold) is a *weighted-sum gate*. No explicit function is known that requires super-polynomial size when implemented by a depth-2 arbitrary-weight threshold circuit with a weighted-sum gate at the output. This is a special case of the more difficult open problem of proving that some given function requires super-polynomial size when implemented by a depth-2 arbitrary-weight threshold circuit (with a threshold allowed in the output gate). We prove a partial result regarding weighted-sum gates in the context of the equality and other related functions.

As mentioned earlier, the n-variable EQ can be implemented by a depth-2 circuit consisting of 2 threshold gates with exponential weights in the first layer and a weighted-sum gate in the second layer. We show that any circuit for EQ that consists of polynomial-weight threshold gates at the first layer and of a weighted-sum gate at the second layer (possibly with exponential weights) has exponential size.

Lemma 5 Suppose that a depth-2 circuit consisting of $p(n)$-rectangular gates in the first layer and a weighted-sum gate (possibly with exponential weights) at the output computes the n-variable EQ. Then the size of the circuit is at least $2^{\frac{n}{2}}/p(n)$.

Proof: Let $g_1, \ldots, g_k$ be the output functions of the k gates in the first layer of the circuit. Consider the 'natural' partition $X = \{1, \ldots, \frac{n}{2}\}$ and $Y = \{\frac{n}{2}+1, \ldots, n\}$ of the input variables. Since the output function is a weighted sum of g_i's we have

$$M_{\mathrm{EQ},X,Y} = \sum_{i=1}^{k} w_i M_{g_i,X,Y} \ .$$

By subadditivity of ranks,

$$\mathrm{rank}(M_{\mathrm{EQ},X,Y}) \leq \sum_{i=1}^{k} \mathrm{rank}(M_{g_i,X,Y}) \ .$$

But

$$\mathrm{rank}(M_{\mathrm{EQ},X,Y}) = 2^{\frac{n}{2}}$$

and for all $i \in \{1, \ldots, k\}$,

$$p(n) \geq \rho_{g_i,X,Y} \geq \mathrm{rank}(M_{g_i,X,Y}) \ .$$

The lemma follows. $\Box$

Corollary 5 Suppose that a depth-2 circuit consisting of polynomial-weight threshold gates in the first layer and a weighted-sum gate (possibly with exponential weights) at the output, computes the n-variable EQ. Then the size of the circuit is $\Omega(2^{\frac{n}{2}-\epsilon})$ for every $\epsilon > 0$.

Proof: Example 2 implies that any threshold gate with weights bounded by $p(n)$ is $(\frac{n}{2}+1)^2 p^2(n)$ rectangular. $\Box$

The above result holds for all functions f, (e.g., COMP) for which $\mathrm{rank}(M_{f,X,Y})$ is exponentially large for some partition $\{X, Y\}$ of the input variables.

6 Concluding Remarks

Several problems remain unresolved.

1. The best symmetric-gates lower bounds for EQ and IP are $\Omega(n/\log n)$ while the best upper bounds are linear.

2. Is there a two-layer polynomial-weight threshold circuit for COMP that meets the lower bound of $\Omega(n/\log n)$?

3. The set of polynomially-rectangular gates, introduced in Section 3, includes the set of generalized symmetric gates. Are the two sets the same? Similarly the set of triangular gates, introduced in Section 4, includes the set of threshold gates. Are these sets the same?

APPENDIX

A lower bound on the Differential Dimension of Boolean Functions

Smolensky [1] used the differential dimension of Boolean functions to lower-bound symmetric-circuit complexity. In this paper we used communication-complexity arguments to simplify the proofs. We now show that similar communication-complexity arguments can be used to lower bound the differential dimension of Boolean functions

Let S be a finite set of points in the n-dimensional complex vector space $\mathcal{C}^n$. Let V denote the space of functions from S to $\mathcal{C}$.

Differential Dimension

The differential dimension of an analytic function $g : \mathcal{C}^n \to \mathcal{C}$ over S is the dimension of the subspace of V spanned by the restrictions to S of g and all of its partial derivatives.

Since we are concerned with functions that interpolate Boolean functions, we assume without loss of generality that $S = \{0,1\}^{\{1,\ldots,n\}}$.

Differential Dimension of Boolean Functions

The differential dimension of a Boolean function $f : \{0,1\}^{\{1,\ldots,n\}} \to \{0,1\}$ is the minimal differential dimension over $S = \{0,1\}^{\{1,\ldots,n\}}$ of any analytic function $g : \mathcal{C}^n \to \mathcal{C}$ that interpolates f.

Let $g : \mathcal{C}^n \to \mathcal{C}$ be an analytic function and let $v \in \mathcal{C}^n$. The *shifted* function g_v is defined by: $g_v(x) = g(x - v)\ \forall\ x \in \mathcal{C}^n$.

Proposition 1 ([1]) The subspace of V spanned by all the partial derivates of all orders of g restricted to S coincides with the subspace of V spanned by all the shifts of g restricted to S. $\square$

Thus if g interpolates a given Boolean function f, then the dimension of the space spanned by the shifts g_v for all $v \in S = \{0,1\}^{\{1,\ldots,n\}}$ lower bounds the differential dimension of g.

Any function g restricted to the set $S = \{0,1\}^{\{1,\ldots,n\}}$ can be viewed as a 2^n-dimensional vector in $\mathcal{C}^n$; each coordinate of the vector is the value of g at a distinct point in S. For any $v_i \in S$, we shall represent the shift g_{v_i} restricted to S as a 2^n-dimensional vector, and denote it as $g_{v_i,S}$. Then the dimension spanned by the shifts $g_{v_1}, g_{v_2}, \cdots, g_{v_k}$, $v_i \in S$, is the *rank* of the following $2^n \times k$ matrix:

$$[g_{v_1,S}\ g_{v_2,S}\ \cdots\ g_{v_k S}]\ .$$

Lemma 6 The differential dimension of a Boolean function f is $\Omega(r)$, where

$$r = \max\{\text{rank}\ (M_{f,X,Y})\ :\ \{X,Y\} \text{ partitions } \{1,\ldots,n\}\}$$

Proof: Let $\{X,Y\}$ partition $\{1,\ldots,n\}$ and let $M_{f,X,Y}$ be the corresponding function matrix. Choose $k = \text{rank}\ (M_{f,X,Y})$ linearly independent columns of M_{fXY}, and let $\{y_1, y_2, \cdots, y_k\}$ be the set of Y-inputs corresponding to the chosen columns. Let $g(x,y)$ interpolate f and consider the following k shifts: $g_{(0,-y_1)}, g_{(0,-y_2)}, \cdots, g_{(0,-y_k)}$. One can show the following for the shifts $g_{(0,-y_i)}$, restricted to $S = \{0,1\}^{\{1,\ldots,n\}}$: 1) $g_{(0,-y_i)}(x,0) = g(x,y_i) = f(x,y_i)$, is known for every $x \in \{0,1\}^X$ and the values of $g_{(0,-y_i)}(x,y)$ might be undetermined if $y \neq 0$; 2) If the entries of the vector $g_{(0,-y_i),S}$ are arranged so that the first $2^{|X|}$ entries correspond to $(x,0) \in \{0,1\}^{\{1,\ldots,n\}}$, then in the following $2^n \times k$ matrix

$$Y_k = [g_{(0,-y_1),S}\ g_{(0,-y_2),S}\ \cdots\ g_{(0,-y_k)S}]$$

the sub-matrix defined by the first $2^{|X|}$ rows are the k linearly independent columns (corresponding to Y-inputs $(y_1, \cdots, y_k)$ chosen from $M_{f,X,Y}$. Thus rank $(Y_k) = k$. Hence by Proposition 1, the differential dimension of any function g interpolating the Boolean function f is $\Omega(\text{rank}\ (M_{f,X,Y}))$. □

The above result implies, for example, that the differential dimensions of the n-variable EQ and COMP are $\Omega(2^{n/2})$.

References

[1] R. Smolnesky. On interpolation by analytical functions with special properties and some weak lower bounds on the size of circuits with symmetric gates. In *Proc. of the 31st Annual Symposium on Foundations of Computer Science*, pages 628–631, 1990.

[2] H. D. Gröger and G. Turán. On linear decision trees computing boolean functions. In *Proc. of the ICALP*, pages 277–288, 1991.

[3] I. Wegener. *The Complexity of Boolean Functions.* New York: Wiley, 1987.

[4] M. Furst, J. B. Saxe, and M. Sipser. Parity circuits and the polynomial time hierarchy. In *Proc. of the 22nd Annual Symposium on Foundations of Computer Science*, pages 260–270, 1981.

[5] A. A. Razborov. Lower bounds for the size of circuits of bolunded depth with basis $\{\wedge, \oplus\}$. In *Math. Notes of the Academy of Science of the USSR, 41:4*, pages 333–338, 1987.

[6] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proc. of the 19th Annual ACM Symposium on Theory of Computing*, pages 77–82, 1987.

[7] A. Hajnal, W. Mass, P. Pudlák, M. Szegedy, and G. Turán. Threshold circuits of bounded depth. In *Proc. of the 28th Annual Symposium on Foundations of Computer Science*, pages 99–110, 1987.

[8] J. Hromkovic. Linear Lower Bounds on unbounded fan-in Boolean Circuits. *Information Processing Letters*, 21:71–74, 1985.

[9] A.C. Yao. Some complexity questions related to distributive computing. In *Proc. of the 11th Annual ACM Symposium on Theory of Computing*, pages 209–213, 1979.

[10] A.V. Aho, J.D. Ullman, and M. Yanakakis. On notions of information transfer in VLSI circuits. In *Proc. of the 15th Annual ACM Symposium on Theory of Computing*, 1983.

[11] K. Melhorn and E.M. Schmidt. Las vegas is better than determinism in vlsi and distributed computing. In *Proc. of the 14th Annual ACM Symposium on Theory of Computing*, 1982.

[12] J. Spencer. *Ten Lectures on the Probabilistic Method.* SIAM, 1987.

[13] M. Goldmann, J. Håstad, and A. Razborov. Majority gates vs. general weighted threshold gates. In *Proc. of the 7th Annual Conference on Structure in Complexity Theory*, 1992.