Monotone Circuits for Connectivity Require Super-logarithmic Depth

Mauricio Karchmer
Avi Wigderson *

Inst. of Mathematics and Computer Science
Hebrew University
Jerusalem, Israel 91904

Abstract

We prove that every monotone circuit which tests *st*-connectivity of an undirected graph on n nodes has depth $\Omega(\log^2 n)$. This implies a superpolynomial ($n^{\Omega(\log n)}$) lower bound on the size of any monotone formula for *st*-connectivity.

The proof draws intuition from a new characterization of circuit depth in terms of communication complexity. It uses counting arguments and Extremal Set Theory.

Within the same framework, we also give a very simple and intuitive proof of a depth analogue of a theorem of Krapchenko concerning formula size lower bounds.

1 Introduction

The circuit complexity of Boolean functions has been studied for 40 years, but its main problem remains unsolved: we have no example of a simple function (say in NP) that requires super linear circuit size or super logarithmic (bounded fan-in) circuit depth. The reason is, perhaps, that although the circuit model is elegantly simple, our understanding of the way it computes is, at the most, vague.

In the last years, however, advance has been surprisingly fast. On the one hand, results of

*part of this work was done while the authors were visiting UC Berkeley supported by an NSF grant DCR-8612563. The second author also wishes to acknowledge support from the Alon Fellowship.

Andreev [An] and Razborov [Ra], improved by Alon and Boppana [AB], give exponential size lower bounds for monotone circuits. On the other hand, results of, by now a long list of authors (e.g. [Aj],[FSS],[Y1],[H]), give exponential size lower bounds for constant depth circuits. More than the results themselves, perhaps the main contribution of the mentioned papers has been the development of some general techniques for proving lower bounds, such as random restrictions and circuit approximation. These techniques, however, turn out to be hard to apply to other problems so that new ideas have been sought.

In this paper we show the equivalence between circuit depth and the communication complexity of a certain related problem [1]. We believe that the later model is much more appealing for both showing and understanding upper bounds, as well as for proving lower bounds. This characterization is reminiscent of, but somehow more explicit and intuitive than, the well known relationship between circuits and alternating machines [Ru]. This characterization allows us to view computation top-down (from output to input) and apply such techniques as random restrictions in that direction (rather than the common bottom-up approach). We argue the relevance of this model by presenting a very simple proof of a depth analogue of a theorem of Krapchenko, and by proving the first superlogarithmic (in the size of the circuit) depth lower bound for monotone circuits.

Though the mentioned results of Andreev and Razborov give exponential (in $\log n$) depth lower bounds for monotone circuits computing certain functions, the depth lower bound is always logarithmic in the size bound. That is, the techniques apply to size rather than to depth. We present a technique which captures, in a strong way, the essence of circuit depth. We give here a tight $\Omega(\log^2 n)$ depth bound for st-connectivity [2], a function which has $O(n^3 \log n)$ size, $O(\log^2 n)$ depth monotone circuits. As a consequence, we get non-polynomial ($n^{\Omega(\log n)}$) size lower bounds for monotone formulas computing st-connectivity and hence separate the monotone analogues of NC^1 and AC^1.

While our proof bears no obvious similarity to the methods of Razborov and Andreev, we point out the important role that (different) nontrivial results from extremal set theory play in both cases.

It is interesting to note here the different character of the connectivity and majority functions in the Boolean and arithmetic monotone circuits models. Shamir and Snir [ShS] showed an $\Omega(\log^2 n)$ depth bound for both functions in the arithmetic model. The difficulty in applying these techniques to Boolean circuits are the axioms $x \vee xy = x$ and its dual, which do not hold in fields. Indeed, Valiant [V] (by probabilistic methods) and [AKS] (by explicit constructions) showed that these axioms make a difference for the majority function which admits $O(logn)$ depth monotone Boolean circuits. Our result says that, unlike for majority, for connectivity the situation in the Boolean case is very similar to the arithmetic one.

It is worthwhile to mention that our results apply to undirected graph st-connectivity, a function that, in some models, is easier

[1] We were told that Yannakakis independently discovered this equivalence which is implicit in [KPPY].

[2] We present here and improved and simplified version of an early result of ours giving a $\Omega(\log^2 n/\log\log n)$ bound. This was possible after J. Hastad formulated and proved lemma 4.1. A similar improvement was independently discovered by R. Boppana.

than its directed version. For example, see [AKLLR] for some relevant evidence. More recently, Ajtai and Fagin [AF] show that, while undirected st-connectivity is definable in monadic second order logic, the directed case is not.

The paper is organized as follows: In §2 we define the communication game and show its equivalence to circuit depth; in §3 we give a simple proof of a theorem of Krapchenko. In §4 we give the lower bound for connectivity.

2 Communication Complexity and Circuit Depth

In this section we show the equivalence between circuit depth and a problem in communication complexity. We will be considering circuits over the basis $\{\vee, \wedge, \neg\}$ where $\{\vee, \wedge\}$-gates have fanin 2 and $\neg$-gates are only applied to input variables. For a function f, $d(f)$ is the minimum depth of a circuit computing f.

Let $B_0, B_1 \subseteq \{0,1\}^n$ such that $B_0 \cap B_1 = \emptyset$. Consider the following game between players I and II: Player I gets $x \in B_1$ while player II gets $y \in B_0$; their goal is to find a coordinate i such that $x_i \neq y_i$. Let $C(B_1, B_0)$ be the minimum number of bits they have to communicate in order for both to agree on such coordinate. Note that unlike standard problems in communication complexity [Y1], the task of the players here is to solve a search, rather than a decision, problem.

Theorem 2.1 *For every function $f : \{0,1\}^n \to \{0,1\}$ we have*

$$d(f) = C(f^{-1}(1), f^{-1}(0))$$

Proof: Follows from the following two lemmas. ∎

Lemma 2.1 *For all functions f and all $B_0, B_1 \subseteq \{0,1\}^n$ such that $B_0 \subseteq f^{-1}(0)$ and $B_1 \subseteq f^{-1}(1)$ we have*

$$C(B_1, B_0) \leq d(f)$$

Proof: By induction on $d(f)$.

If $d(f) = 0$ then f is either x_i or $\bar{x}_i$. In either case, we have that for all $x \in B_1$ and $y \in B_0$, $x_i \neq y_i$ so that i is always an answer and $C(B_1, B_0) = 0$.

For the induction step we suppose that $f = f_1 \wedge f_2$ (the case $f = f_1 \vee f_2$ is treated similarly) so that $d(f) = \max(d(f_1), d(f_2)) + 1$. Let $B_0^j = B_0 \cap f_j^{-1}(0)$ for $j = 1, 2$. By induction we have that $C(B_1, B_0^j) \leq d(f_j)$ for $j = 1, 2$. Consider the following protocol for B_1 and B_0: II sends a 0 if $y \in B_0^1$, otherwise he sends a 1; the players then follow the best protocol for each of the subcases. We have

$$\begin{aligned} C(B_1, B_0) &\leq 1 + \max_{j=1,2}(C(B_1, B_0^j)) \\ &\leq 1 + \max_{j=1,2}(d(f_j)) \\ &= d(f) \end{aligned}$$

∎

The converse is as follows:

Lemma 2.2 *Let $B_0, B_1 \subseteq \{0,1\}^n$ such that $B_0 \cap B_1 = \emptyset$. Then, there exists a function f with $B_0 \subseteq f^{-1}(0)$ and $B_1 \subseteq f^{-1}(1)$ such that*

$$d(f) \leq C(B_1, B_0)$$

Proof: By induction on $C(B_1, B_0)$.

If $C(B_1, B_0) = 0$ then there exists an i such that $\forall x \in B_1$ and $\forall y \in B_0$, $x_i \neq y_i$. It is clear that $\forall x', x'' \in B_1$ we have $x'_i = x''_i$ and the same holds for all $y', y'' \in B_0$. Without loss of generality $x_i = 1$ so that letting $f = x_i$ we have $B_0 \subseteq f^{-1}(0)$ and $B_1 \subseteq f^{-1}(1)$.

To prove the induction step, we assume that II sends the first bit (the other case is treated similarly). For some partition $B_0 = B_0^1 \cup B_0^2$, II sends a 0 if $y \in B_0^1$, a 1 otherwise; the players then continue with the best protocol for each of the subcases and

$$C(B_1, B_0) = 1 + \max_{j=1,2}(C(B_1, B_0^j))$$

By induction, there exist f_1, f_2 so that $B_0^j \subseteq f_j^{-1}(0)$, $B_1 \subseteq f_j^{-1}(1)$ and $d(f_j) \leq C(B_1, B_0^j)$ for $j = 1, 2$. Taking now $f = f_1 \wedge f_2$ we have

$$\begin{aligned} B_1 &\subseteq f_1^{-1}(1) \cap f_2^{-1}(1) = f^{-1}(1) \\ B_0 &= B_0^1 \cup B_0^2 \\ &\subseteq f_1^{-1}(0) \cup f_2^{-1}(0) = f^{-1}(0) \end{aligned}$$

and

$$\begin{aligned} d(f) &\leq 1 + \max_{j=1,2}(d(f_j)) \\ &\leq 1 + \max_{j=1,2}(C(B_1, B_0^j)) \\ &= C(B_1, B_0) \end{aligned}$$

∎

For monotone circuits we can give a modified version of theorem 2.1 which captures, in a nice way, the restrictions of monotone computation. A *minterm*, (*maxterm*) of a monotone function f is a minimal set of variables which if we set to 1 (0), f will be equal to 1 (0) regardless of the other variables. Define $min(f)$, $Max(f)$ as the set of minterms, respectively maxterms of f. It is easy to see that every minterm intersects every maxterm. We will look at minterms and maxterms as subsets of $[n]$ [3]. For a monotone function f, let $d_m(f)$ be the minimum depth of a monotone circuit computing f.

Consider the following communication game (the *monotone* game) between players I and II. Let $P, Q \subseteq 2^{[n]}$ such that $\forall p \in P$ and $\forall q \in Q$ we have $p \cap q \neq \emptyset$. Player I gets a $p \in P$ while player II gets a $q \in Q$; their goal is to find an element in $p \cap q$. Let $C_m(P, Q)$ be the minimum number of bits they have to communicate in order to find such element.

Theorem 2.2 *For every monotone function f we have*

$$d_m(f) = C_m(min(f), Max(f))$$

Proof: Note that in the base case of lemma 2.1, if the circuit is monotone, we always find a coordinate i such that $x_i = 1$ while $y_i = 0$. In the other hand, if the protocol always gives a coordinate i with the above property, lemma 2.2 gives a monotone circuit.

Let $x \in f^{-1}(1)$ be the characteristic vector of a subset $p \subseteq [n]$. Similarly, let $y \in f^{-1}(0)$ be the characteristic vector of the complement of a subset $q \subseteq [n]$ (i.e. $i \in q \iff y_i = 0$). By the above argument, it is clear that the answer of the protocol will be an element of $p \cap q$. The theorem follows by noticing that it is enough to give a protocol for $(min(f), Max(f))$ because the players, in case they get inputs p' and q', can always behave as if they got $p \subseteq p'$ and $q \subseteq q'$ where $p \in min(f)$ and $q \in Max(f)$. ∎

For proving lower bounds for the communication game, it may be convenient to have

[3] $[n] = \{1, ..., n\}$

more structure in the way the players behave. We would like to synchronize the protocol so that the players communicate in rounds where players I and II send messages of length at most a_I, a_{II} respectively:

Theorem 2.3 *Let D be a protocol in rounds where at each round player I sends a bits and player II responds with 2^a bits ($a \leq \log n$). Then for any function f, the number k of rounds satisfies*

$$k \leq \frac{d(f)}{a}$$

for the general game and

$$k \leq \frac{d_m(f)}{a}$$

for the monotone one.

Proof: Let C be the best circuit for f. The idea is to simulate a layers of C with a round of D. Divide C into stages of depth a each and look at the subcircuits of each stage. Each one computes a function which depends on at most 2^a wires and, thus, can be represented in CNF form with less than 2^{2^a} clauses, each of length 2^a. Following the proof of lemma 2.1, it is easy to see that such CNF representation can be simulated by a round where player II sends 2^a bits and player I sends a bits. The same holds for the monotone case. ∎

Of course, in theorem 2.3, the roles of I and II can be switched so that I and II send 2^a, a bits per round respectively.

3 Krapchenko's bound

As a nice application of theorem 2.1, we give a simple proof of a depth analogue of a theorem of Krapchenko. Let C_n be the graph of the n-cube with vertex set $\{0,1\}^n$ and two nodes adjacent iff they differ in one coordinate. Any subset A of edges induce a graph G_A of C_n in the natural way. For a graph G_A and a node x, we denote $d_A(x)$, $N_A(x)$ as the degree of x in A and the set of neighbors of x in A respectively. We drop the subindex A if no confusion arises. Let E denote expectation with uniform distribution.

Theorem 3.1 (Krapchenko) *Let $B_0, B_1 \subseteq \{0,1\}^n$ such that $B_0 \cap B_1 = \emptyset$. Let $A = C_n \cap (B_0 \times B_1)$. Then, for every function f with $B_0 \subseteq f^{-1}(0)$ and $B_1 \subseteq f^{-1}(1)$ we have*

$$d(f) \geq \log \frac{|A|^2}{|B_0||B_1|}$$

Proof: Fix a protocol D for the communication game and let $C(x,y)$ be the number of bits D uses on inputs x,y. We will prove that for (x,y) taken uniformly from A we have

$$\mathrm{E}(C(x,y)) \geq \log \frac{|A|^2}{|B_0||B_1|} \qquad (*)$$

By lemma 2.1, we get $d(f) \geq \mathrm{E}(C(x,y))$. We view $(*)$ as follows: Write

$$\log \frac{|A|^2}{|B_0||B_1|} = \log \frac{|A|}{|B_0|} + \log \frac{|A|}{|B_1|}$$

and notice that $|A|/|B_0|$ and $|A|/|B_1|$ are the average degree of nodes in B_0 and B_1 respectively. In what follows, we will claim that the number of bits player I sends is at least the logarithm of the average degree of nodes in B_1 (similarly with player II). Intuitively, this is so because even if player I knows y, he needs $\log d(y)$ bits to tell II which x he has.

We now proceed formally. For $(x,y) \in A$, let $b_{\mathrm{I}}(x,y)$, $b_{\mathrm{II}}(x,y)$ be the number of bits players I, II send when the input to the protocol

is (x, y). We have

$$\mathrm{E}(C(x,y)) = \frac{1}{|A|}\left[\sum_{(x,y)\in A}(b_{\mathrm{I}}(x,y)+b_{\mathrm{II}}(x,y))\right]$$

$$= \frac{1}{|A|}\left[\sum_{x\in B_1}\sum_{y\in N(x)} b_{\mathrm{II}}(x,y) + \sum_{y\in B_0}\sum_{x\in N(y)} b_{\mathrm{I}}(x,y)\right]$$

We claim:

- For any $x \in B_1$, $\sum_{y\in N(x)} b_{\mathrm{II}}(x,y) \geq d(x)\log d(x)$. This is so because, even if II knows x, he has to tell I which y he has.
- Similarly, $\forall y \in B_0$ we have $\sum_{x\in N(y)} b_{\mathrm{I}}(x,y) \geq d(y)\log d(y)$.

We now conclude,

$$\begin{aligned}
&\mathrm{E}(C(x,y)) \\
&\geq \frac{1}{|A|}\left(\sum_{x\in B_1} d(x)\log d(x) + \sum_{y\in B_0} d(y)\log d(y)\right) \\
&\geq \frac{1}{|A|}\left(\sum_{x\in B_1}\frac{|A|}{|B_1|}\log\frac{|A|}{|B_1|} + \sum_{y\in B_0}\frac{|A|}{|B_0|}\log\frac{|A|}{|B_0|}\right) \\
&= \log\frac{|A|^2}{|B_1||B_0|}
\end{aligned}$$

where the last inequality follows from the convexity of $x\log x$. ∎

4 A Lower Bound for Connectivity

In this section we give a $\Omega(\log^2 n)$ depth lower bound for monotone circuits computing undirected graph st-connectivity. This section is organized as follows: In §4.1 we give some intuition and we state the main theorem; In §4.2 we give some definitions and useful lemmas; finally, in §4.3 we give the proof of the theorem.

4.1 Intuition

The function st-connectivity receives as input an undirected graph and two distinguished vertices s and t, and tests whether there is a path from s to t or not. The function is obviously monotone with minterms corresponding to minimal st-paths, and maxterms corresponding to minimal st-cuts. We view st-paths as ordered sets of vertices excluding s and t. We view st-cuts as a partition of the set of vertices into two subsets, one containing s and the other containing t. The minimal cut contains all edges between the two subsets. This is the same as having a coloring $q : V \to \{0,1\}$ where $q(s) = 0$ and $q(t) = 1$. The game thus, is as follows: Player I gets an st-path while player II gets a coloring of the nodes. Their goal is to find a bichromatic edge in the path.

Let us look at the protocol based on the idea of raising the adjacency matrix of the graph to the n^{th} power: Player I sends the name of the middle vertex on his path, player II responds with the color of that vertex. The players then continue recursively on the half path where a bichromatic edge is ensured to exist. Note that the protocol requires $O(\log n)$ rounds in each of which I sends $\log n$ bits and II sends just one.

The crucial observation is that, even if player II would be allowed to send $O(n^\epsilon)$ bits each round (instead of one bit as in the protocol), the players will still need many rounds. Basically, this is because II doesn't know much about the nodes in I's path. If he sends $O(n^\epsilon)$ bits and the path is of length $O(n^\epsilon)$ then the

probability that I gets valuable information from II is negligible. If we could prove a $\Omega(\log n)$ lower bound for the number of rounds needed, we will be able to use theorem 2.3 to get the promised $\Omega(\log^2 n)$ depth lower bound for circuits.

Note the asymmetry between players I and II. Indeed, if the roles of both players were switched so that player I would be the one which sends $O(n^\epsilon)$ bits per round, they would be able to solve the problem in a constant number of rounds. This is consistent with the intuition gotten by Shamir and Snir in [ShS].

Define *stconn*(l) as the restriction of *st*-connectivity to the case where player I gets a path of length l. We state the main theorem of this section:

Theorem 4.1 *Let $l \leq n^{1/10}$. There exists an $0 < \epsilon < 1/2$ such that if D is a k-round protocol for stconn(l) where at each round, player I sends $\epsilon \log n$ bits and player II sends n^ϵ bits, then $k \geq \log l$.*

Corollary 4.1 *Any monotone circuit computing st-connectivity requires $\Omega(\log^2 n)$ depth.*

Proof: Follows from theorems 2.3 and 4.1 by taking $l = n^{1/10}$. ∎

Open 4.1 *Prove a similar lower bound for Connectivity.*

Corollary 4.2 *Any monotone formula for st-connectivity has size $n^{\Omega(\log n)}$.*

Proof: Follows by noting that the relation $d(f) = O(\log L(f))$ holds also in the monotone case. ∎

To prove theorem 4.1, we will assume, for contradiction, the existence of a k-round protocol ($k < \log l$) good for a large family of all possible paths and a large family of all possible colorings. We will pick a large subset of the paths and colorings for which players I and II sent the same message in the first round. We will give some extra information (by applying a random restriction to the coloring of the nodes) to both players so as to get smaller, yet nicer, subsets which are in 1-1 correspondence with a family of paths shorter in length (but of higher quality) and a family of colorings of fewer nodes. The fact that the original protocol had $(k-1)$ rounds to go, will allow us to find a $(k-1)$-round protocol for the smaller families. Repeating this $\Omega(\log n)$ times will give us a protocol without communication that solves a problem which cannot be solved without any messages.

Note the top-down structure of the proof; essentially, the argument shows that, if the output of a circuit is in some sense complex and, as long as we do not go too far down the circuit, there is a wire which computes a complex subfunction.

4.2 Notation and Definitions

Denote by $\Pi_l^n \subseteq [n]^l$ the set of all paths on $[n]$ of length l. $|\Pi_l^n| = (n)_l$ where

$$(n)_l = n(n-1)\cdots(n-l+1)$$

An *interval* $I \subseteq [l]$ is a subset of consecutive integers. For a path $p \in \Pi_l^n$ and an interval $I \subseteq [l]$, p_I is the subpath of p in the interval I. For $P \subseteq \Pi_l^n$, $P_I = \{p_I : p \in P\}$ is the projection of P into I. Note that $P_I \subseteq \Pi_{|I|}^n$. Conversely, for $p \in \Pi_{|I|}^n$, $P \subseteq \Pi_l^n$ and an interval I, let $Ext_{P,I}(p) = \{\tilde{p} \in P : \tilde{p}_I = p\}$ be the set

of extensions of p in P within I. We will drop the subindices P and I if no confusion arises. For $p \in \Pi_l^n$, the *support* of p, $supp(p)$, is defined as the set of nodes contained in p. When no confusion arises, we will denote $supp(p)$ by p. For $P \subseteq \Pi_l^n$, $supp(P) = \{supp(p) : p \in P\}$. Given a partition of $[l]$ into two intervals L and R, we will denote a path $p \in \Pi_l^n$ by (p_L, p_R) where each entry is the projection of p into the respective interval.

Similarly, for a coloring $q \in \{0,1\}^n$ and a subset $T \subseteq [n]$, q_T is the projection of q into T and for $Q \subseteq \{0,1\}^n$, $Q_T = \{q_T : q \in Q\}$ is the projection of Q into T. For $q \in \{0,1\}^{|T|}$, $Q \subseteq \{0,1\}^n$ and a subset $T \subseteq [n]$, let $Ext_{Q,T}(q) = \{\tilde{q} \in Q : \tilde{q}_T = q\}$ be the set of extensions of q in Q within T (again, we drop the subindices Q and T whenever possible). For a restriction $\rho : [n] \mapsto \{0,1,*\}$, we will denote by Q_ρ the set of colorings in Q consistent with ρ, (i.e. $\{q \in Q : \rho(i) \neq * \Rightarrow \rho(i) = q_i\}$).

For a subset A of a universe Ω, the density of A, $\mu(A)$, is defined as $|A|/|\Omega|$. In what follows, we will work with densities rather than with cardinalities.

We will need the following combinatorial lemma due to J. Hastad: Let $H \subseteq A_1 \times ... \times A_k$ and for $v \in A_i$, let $Ext_{A_i}(v) = \{\bar{u} \in H : u_i = v\}$. Note that, though $Ext_{A_i}(v) \subseteq H$, $Ext_{A_i}(v)$ may be considered as a subset of $H/A_i = A_1 \times ... \times A_{i-1} \times A_{i+1} \times ... \times A_k$. In what follows, we will consider $Ext_{A_i}(v)$ as a subset of H/A_i.

Lemma 4.1 *Let $H \subseteq A_1 \times ... \times A_k$. Let $B_i = \{u \in A_i : \mu(Ext_{A_i}(u)) \geq \mu(H)/2k\}$. Then*

$$\prod_{i=1}^{k} \mu(B_i) \geq \frac{\mu(H)}{2}$$

Proof: Say that a member $(u_1, ..., u_k)$ of H is bad if for some i, $u_i \notin B_i$. Let $\bar{H}$ be the set of bad elements in H. We have

$$\begin{aligned} \mu(\bar{H}) &\leq \sum_{i=1}^{k} \mu(\cup_{u \notin B_i} Ext_{A_i}(u)) \\ &< \sum_{i=1}^{k} \frac{\mu(H)}{2k} = \frac{\mu(H)}{2} \end{aligned}$$

the lemma follows immediately, by noting that

$$\prod_{i=1}^{k} \mu(B_i) \geq \mu(\bar{H}).$$

∎

Corollary 4.3 *If $k = 2$, then $\exists i$ such that $\mu(B_i) \geq \left(\frac{\mu(H)}{2}\right)^{1/2}$.* ∎

Corollary 4.4 $\Pr\left(\mu(B_i) < (\mu(H)/2)^{2/k}\right) < 1/2$ *for i chosen randomly from $\{1, ..., k\}$.* ∎

We will also need the following result in extremal graph theory which is based on the Kruskal-Katona theorem. Let $\mathcal{F} \subseteq 2^{[n]}$. $\mathcal{F}$ is called an *ideal* if $A \in \mathcal{F}$ and $B \subseteq A$ implies that $B \in \mathcal{F}$. For $0 \leq s \leq n$, $P_s(\mathcal{F})$ is the probability that a random s-subset S is in $\mathcal{F}$. We present the following lemma of Bollobas and Thomason:

Lemma 4.2 (BTh) *Let $\mathcal{F} \subseteq 2^{[n]}$ be an ideal. Then for $0 \leq s < r \leq n$ we have*

$$P_s(\mathcal{F})^r \geq P_r(\mathcal{F})^s$$

Intuitively, the lemma says that the probability of the different slices of an ideal decreases exponentially fast with their distance. Our main use of the lemma is the following corollary:

Corollary 4.5 *Let $S \subseteq [n]^{(s)}$ be a family of s-subsets of $[n]$ with $\mu(S) \geq \alpha$. Take a random t-subset T of $[n]$. The probability that T does not contain any element of S is at most $(1-\alpha)^{t/s}$.*

Proof: Say T is good if it contains an element of S, T is bad otherwise. It is easy to see that the ideal $\mathcal{I}$ generated by the bad T's misses all of S so that $P_s(\mathcal{I}) \leq (1-\alpha)$. By lemma 4.3 we get

$$\Pr(T \text{ is bad}) = P_t(\mathcal{I}) \leq (1-\alpha)^{t/s}$$

∎

4.3 The proof

Proof: [of theorem 4.1]

In what follows, all our protocols will be synchronized so that at each round I sends $\epsilon \log n$ bits and II responds with n^{ϵ} bits. The existence of ϵ will be clear from the proof, though one can check that $\epsilon = 1/10$ suffices. We will define a sequence of problems of different sizes as follows: We first define the parameters of the problems, let $t_{max} = \log l - 1$.

Let $V_0 = [n]$ and $V_{t+1} \subseteq V_t$ where $|V_t| = n_t$ and $n_{t+1} = n_t - 4n_t^{1/2}$. Note that

$$n/2 \leq n_t \leq n \qquad \text{for} \qquad t \leq t_{max}$$

Let $l_0 = l$ and $l_{t+1} = l_t/2$ and note that

$$2 \leq l_t \leq l \qquad \text{for} \qquad t \leq t_{max}$$

Consider the following property:

H(t, k): *There exist a collection of paths $P^t \subseteq \Pi_{l_t}^{n_t}$ of length l_t over V_t, and a collection of colorings $Q^t \subseteq \{0,1\}^{n_t}$ of V_t, with $\mu(P^t) \geq n^{-\epsilon}$ and $\mu(Q^t) \geq 2^{-2tn^{\epsilon}}$ such that there exists a k-round protocol D^t good for (P^t, Q^t).*

We will prove the following two claims:

Claim 4.1 *For $t \leq t_{max}$ $\neg H(t, 0)$.*

Claim 4.2 *For $t \leq t_{max}$ $H(t,k) \rightarrow H(t+1, k-1)$*

It is clear that the two claims imply $\neg H(0, t_{max})$ which in turn implies our theorem.

The first claim follows easily by noticing that there is not a single node which appears in every path of P^t so that player II cannot know any edge in it let alone the answer. To see this, note that the fraction of paths of length l_t out of n_t nodes which contain a given node is $l_t/n_t \ll n^{-\epsilon}$. This is enough for proving the claim as both players must know the answer. However, it can also be shown that, for most input pairs, player I will not know the color of a single node in its path.

The second claim will be proved by assuming $H(t,k)$ and constructing P^{t+1}, Q^{t+1} and D^{t+1} so as to satisfy $H(t+1, k-1)$. Take P^t, Q^t and D^t which satisfy $H(t,k)$. Let us look at the protocol after the first round. By the pigeonhole principle, there exist $P \subseteq P^t$ with $\mu(P) \geq n^{-2\epsilon}$ such that for every path in P, I sent the same message. Similarly, there exists $Q \subseteq Q^t$ with $\mu(Q) \geq 2^{-(2t+1)n^{\epsilon}}$ so that for every coloring of Q, II sent the same message.

Let $L = \{1, ..., l_t/2\}$ and $R = \{l_t/2 + 1, ..., l_t\}$ be a partition of the path's coordinates into left and right intervals of the same length. We say that P is L-good if "many" left projections of P have, each, "many" extentions to the right; that is, if

$$\mu\left(\{p_L : \mu(Ext_{P,L}(p_L)) \geq n^{-3\epsilon}\}\right) \geq 2n^{-\epsilon}$$

R-goodness is defined similarly. The following lemma says that if we shrink the length of the paths to half and we restricted our family P to one of the intervals, then we can improve the quality of our collection.

Lemma 4.3 *P is either L-good or R-good.*

Proof: We have $P \subseteq \prod_{l_{t+1}}^{n_t} \times \prod_{l_{t+1}}^{n_t}$. Note that $\mu(P) \geq n^{-2\epsilon}(1-n^{-4/5})$ by the choice of l. The lemma follows using corollary 4.3 and noting that $n^{-2\epsilon}(1-n^{-4/5})/4 \geq n^{-3\epsilon}$. ▓

Without loss of generality, assume that P is L-good and let A be the set of paths in P_L with many extensions. The next lemma is the heart of our argument:

Lemma 4.4 *There exists a restriction $\rho : V_t \mapsto \{0,1,*\}$ with $|\rho^{-1}(*)| = n_{t+1}$ such that the following properties hold:*

G1: *If $\bar{Q} = Q_\rho$ then $\mu(\bar{Q}) \geq 2^{-2(t+1)n^\epsilon}$.*

G2: *$\exists \bar{P} \subseteq P$ such that*

- *$\forall p \in \bar{P}, p_L \subseteq \rho^{-1}(*)$ and $p_R \subseteq \rho^{-1}(1)$*
- *$\forall p, p' \in \bar{P}$ $p_L \neq p'_L$*
- *$\mu(\bar{P}) \geq n^{-\epsilon}$.*

Assuming the lemma is true, we will finish the proof of the second claim:

Let $V^{t+1} = \rho^{-1}(*)$, $Q^{t+1} = \bar{Q}_{\rho^{-1}(*)}$ and $P^{t+1} = \bar{P}_L$. Note that there is a natural 1-1 correspondence between Q^{t+1} and $\bar{Q}$ and between P^{t+1} and $\bar{P}$. Also note that $\forall q \in \bar{Q}$ and $\forall p \in \bar{P}$ any bichromatic edge lies in the interval L. The protocol D^{t+1} on (P^{t+1}, Q^{t+1}) simulates the protocol D^t on $(\bar{P}, \bar{Q})$ by following the behaviour of the associated path and coloring. ▓

Proof: [of lemma 4.4]

In what follows, we denote $V = V_t$, $v = n_t$, $l = l_t$, $l' = l_{t+1}$ for simplicity sake. The existence of a good restriction will be shown by probabilistic methods. We will pick ρ uniformly from the set of all restrictions with $|\rho^{-1}(*)| = v - 4\sqrt{v}$ and $\Pr(\rho(x) = 0 | \rho(x) \neq *) = 1/2$, and show that, with positive probability, the conditions of the lemma are fulfilled. Specifically, we will show that $\Pr(\neg\mathbf{G1}) + \Pr(\neg\mathbf{G2}) \leq 1/2 + o(1)$.

Let us start with **G1**: Intuitively, the following lemma says that, with high probability, ρ does not give player I too much information about the colors of nodes in $\rho^{-1}(*)$.

Lemma 4.5 $\Pr\left(\mu(Q_\rho) < 2^{-2(t+1)n^\epsilon}\right) \leq \frac{1}{2} + o(1)$.

Proof:

Let $\alpha = 2^{-(2t+1)n^\epsilon}$. Picking ρ uniformly from all restrictions with $|\rho^{-1}(*)| = v - 4\sqrt{v}$, is equivalent to picking randomly $T = \rho^{-1}(1) \cup \rho^{-1}(0)$ among all $4\sqrt{v}$-subsets of V, and then picking the restriction of ρ to T randomly among all vectors x in $\{0,1\}^{4\sqrt{v}}$. Let $k = \sqrt{v}/4$. Say T is *bad* if

$$\mu\left(\left\{x : \mu(Ext_{Q,T}(x)) \geq \frac{\alpha}{2k}\right\}\right) < \left(\frac{\alpha}{2}\right)^{2/k}$$

T is *good* otherwise. We have

$$\begin{aligned}\Pr\left(\mu(Q_\rho) < \frac{\alpha}{2k}\right) &\leq \Pr(T \text{ is bad}) \\ &+ \Pr\left(\mu(Q_\rho) < \frac{\alpha}{2k} \mid T \text{ is good}\right)\end{aligned}$$

Note that $Q_\rho = Ext_{Q,T}(x)$. By the definition of goodness, and the choice of k, the second term is bounded by $1 - (\alpha/2)^{2/k} = o(1)$. Also note that $\alpha/2k \geq 2^{-2(t+1)n^\epsilon}$. It remains to bound the first term: We pick a random T by first picking a random partition of V into $4\sqrt{v}$-subsets and then picking a random subset from the partition. We must show that for

any partition, $\Pr(T \text{ is bad}) < 1/2$ for a random T in the partition. But this is precisely the content of corollary 4.4. ∎

Now we take care of **G2**:

Let $A^* = \{p \in A : p \subseteq \rho^{-1}(*)\}$. We say that ρ kills a path $p_L \in A$ if there is no $p_R \in Ext(p_L)$ with $p_R \subseteq \rho^{-1}(1)$. We will show that for every choice of ρ, $\mu(A^*)$ is large and hence $\Pr(\neg\mathbf{G2}) \leq \Pr(\exists p_L \in A \text{ killed by } \rho)$.

Claim 4.3 *For every ρ, $\mu(A^*) \geq n^{-\epsilon}$.*

Proof: Recall that $|\rho^{-1}(*)| = v - 4\sqrt{v}$ so that $|V \backslash \rho^{-1}(*)| = 4\sqrt{v}$. It is easy to see that at most a fraction $n^{-2/5}$ of the l'-subsets of V intersect $V \backslash \rho^{-1}(*)$ so that at least a fraction $2n^{-\epsilon} - n^{-2/5} \geq n^{-\epsilon}$ of them are in A^*. ∎

We now use some combinatorics to bound the probability that there exists a path in A killed by ρ:

Claim 4.4 $\Pr(\exists p_L \in A \text{ killed by } \rho) = o(1)$.

Proof: We have

$$\begin{aligned}&\Pr(\exists p_L \in A \text{ killed by } \rho)\\ &\quad\leq |A| \cdot \max_{p_L \in A}\{\Pr(p_L \text{ is killed by } \rho)\}\end{aligned}$$

so let us look at the worst possible $p_L \in A$. Note that $supp(Ext(p_L))$ contain at least a fraction $n^{-3\epsilon}$ of the l'-subsets of V. This is because the worst situation is when all possible orderings of a given subset are contained in the collection.

Let $F \equiv (\nexists p_R \in supp(Ext(p_L))$ with $p_R \subseteq \rho^{-1}(1))$. We will bound now $\Pr(F)$. We will break $\Pr(F)$ into two cases according to how small $|\rho^{-1}(1)|$ is. We overestimate $\Pr(F)$ as follows:

$$\begin{aligned}\Pr(F) &\leq \Pr(|\rho^{-1}(1)| < \sqrt{v})\\ &+ \Pr(F \mid |\rho^{-1}(1)| \geq \sqrt{v})\end{aligned}$$

The first term can be bounded by $\exp(-\sqrt{v}/2)$ using the Chernoff bound. For the second term, for every random ρ, we have $|\rho^{-1}(1)| \geq \sqrt{v}$. By choosing a random $\sqrt{v}$-subset $\bar{T}$ of $\rho^{-1}(1)$ we induce a marginal uniform distribution for $\bar{T}$ over all the $\sqrt{v}$-subsets of V. We can now apply corollary 4.4 to get

$$\Pr(F \mid |\rho^{-1}(1)| \geq \sqrt{v}) \leq (1 - n^{-3\epsilon})^{v^{1/2 - 1/10}}$$

so that $\Pr(F) \leq exp(-n^{1/4})$. Recalling that $|A|$ is less than $n^{n^{1/10}}$, we conclude easily our calculations and get

$$\begin{aligned}\Pr(\exists p \in A \text{ killed by } \rho) &\leq n^{n^{1/10}} \cdot e^{-n^{1/4}}\\ &= o(1)\end{aligned}$$

∎

We have $\Pr(\neg\mathbf{G1}) + \Pr(\neg\mathbf{G2}) \leq 1/2 + o(1)$ implying the existence of a good restriction. Take any consistent extension of each $p_L \in A^*$ not killed by ρ to form $\bar{P}$. We have $\mu(\bar{P}) \geq n^{-\epsilon}$ and lemma 4.4 is proved. ∎

ACKNOWLEDGMENTS

We are very grateful to Miklos Ajtai for a conversation which led to the proof of theorem 4.1. We are indebted to Johan Hastad for allowing us to present his improvements of our early results.

REFERENCES

[AB] N. Alon, R. Boppana, "The Monotone Circuit Complexity of Boolean Functions", *Combinatorica* 7, pp. 1-22 (1987).