Electronic Colloquium on Computational Complexity, Report No. 129 (2017)

An Efficient Randomized Protocol for every Karchmer-Wigderson Relation with Two Rounds

Or Meir*

August 27, 2017

Abstract

One of the important challenges in circuit complexity is proving strong lower bounds for constant-depth circuits. One possible approach to this problem is to use the framework of Karchmer-Wigderson relations: Karchmer and Wigderson [KW90] observed that for every Boolean function f there is a corresponding communication problem KW_f, called the Karchmer-Wigderson relation of f, whose *deterministic* communication complexity is tightly related to the depth complexity of f. In particular, if we could prove that every *deterministic constant-round* protocol for KW_f must transmit at least c bits, then this would imply a lower bound of $2^{\Omega(c)}$ on the size of constant-depth circuits computing f.

In this work, we observe that there is a *randomized* two-round protocol that solves every Karchmer-Wigderson relation KW_f by transmitting only $O(\log^2 n)$ bits. This means that if we wish to use Karchmer-Wigderson relations in order to prove exponential lower bounds for constant-depth circuits, then we cannot use techniques that work against randomized protocols.

1 Introduction

Proving circuit lower bounds is a central challenge of complexity theory. Unfortunately, proving even super-linear lower bounds for general circuits seems to be beyond our reach at this stage. In order to make progress and develop new proof techniques, much of the current research focuses on proving lower bounds for restricted models of circuits. One of the simplest restricted models that are not yet fully understood is circuits of constant depth (with unbounded fan-in).

By a standard counting argument, we know that there exists a non-explicit function that requires such circuits of size $\Omega(2^n)$. On the other hand, the strongest lower bound we have for an explicit function [Ajt83, FSS84, Hås86] says that circuits of depth d computing the parity of n bits must be of size $2^{\Omega(n^{1/(d-1)})}$. Hence, while strong lower bounds are known in this model, there is still a significant gap in our understanding. In particular, it is an outstanding open problem to prove a lower bound of $\Omega(2^n)$ even for depth-3 circuits computing an explicit function (or, indeed, any lower bound that is better than $2^{\Omega(\sqrt{n})}$).

One possible approach for attacking this problem is a framework due to Karchmer and Wigderson [KW90]. This framework was originally developed for proving lower bounds on the depth of circuits with *bounded fan-in*. Given a function f, we define the *depth complexity* of f to be the smallest depth of a circuit with fan-in 2 that computes f. Karchmer and Wigderson observed that for every Boolean function f there is a corresponding communication problem KW_f, called the

[*] Department of Computer Science, Haifa University, Haifa 31905, Israel. ormeir@cs.haifa.ac.il. Partially supported by the Israel Science Foundation (grant No. 1445/16).

ISSN 1433-8092

Karchmer-Wigderson relation of f, such that the deterministic communication complexity of KW_f is exactly equal to the depth complexity of f. Hence, one can prove lower bounds on the depth complexity of a function f by proving lower bounds on the communication complexity of KW_f. This approach has proved very fruitful in the setting of *monotone* circuits [KW90, GS91, RW92, KRW95].

The framework of Karchmer and Wigderson could also be used to prove lower bounds on constant-depth circuits with unbounded fan: it is implicit[1] in the work of [KW90] that lower bounds on the deterministic communication complexity of *constant-round* protocols for KW_f imply lower bounds for constant-depth circuits. More specifically, if every deterministic r-round protocol for KW_f must transmit at least c bits, then every depth-r circuit (with unbounded fan-in) that computes f must be of size at least $2^{c/r}$. Hence, if we could find an explicit function f such that every constant-round protocol for KW_f must transmit $\Omega(n)$ bits, we would obtain a lower bound of $\Omega(2^n)$ on the size of circuits computing f.

Soon after the introduction of Karchmer-Wigderson relations, Karchmer observed a severe limitation of this framework (see [RW89]): there is a randomized protocol that solves every Karchmer-Wigderson relation by transmitting $O(\log n)$ bits. This means that if one wishes to use Karchmer-Wigderson relations in order to prove super-logarithmic lower bounds on depth complexity, then one has to use proof techniques that cannot prove lower bounds against randomized protocols. Since the most powerful techniques in the field of communication complexity are effective against randomized protocols, this limitation makes the use of Karchmer-Wigderson relations quite difficult.

Karchmer's protocol uses a logarithmic number of rounds, so it is not clear a priori that this limitation applies to proving constant-depth lower bounds. In this work, we observe that a similar limitation applies in the setting of constant-depth lower bounds as well. Specifically, we show that there is a randomized *two-round* protocol that solves every Karchmer-Wigderson relation KW_f by transmitting only $O(\log^2 n)$ bits. This means that proof techniques that are effective against randomized protocols can only prove lower bounds of at most $n^{O(\log n)}$ for constant-depth circuits, and in particular, cannot prove exponential lower bounds.

2 Preliminaries and Our Result

For $n \in \mathbb{N}$, we denote $[n] \overset{\text{def}}{=} \{1, \dots n\}$. For a string $x \in \{0,1\}^n$ and a set of coordinates $S \subseteq [n]$, we denote by $x|_S$ the projection of x to the coordinates in S. We use the standard definitions of communication complexity — see the book of Kushilevitz and Nisan [KN97] for more details.

Our proof uses the Hamming code, which we present next. Given two strings $x, y \in \{0,1\}^n$, the *(Hamming) distance* between x and y is the number of coordinates on which they differ. Given a string $x \in \{0,1\}^n$ and $r \in \mathbb{N}$, the *Hamming ball of radius r around x* is the set of all strings whose Hamming distance from x is at most r. The *Hamming code* is a partition of $\{0,1\}^n$ to balls of radius 1, and it exists for every $n \in \mathbb{N}$ for which $n+1$ is a power of 2 (see, e.g., Lecture 2 in [Sud01] for the construction of the Hamming code).

2.1 Karchmer-Wigderson relations and the universal relation

Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. The Karchmer-Wigderson relation KW_f is defined as follows: Alice gets an input $x \in f^{-1}(0)$, and Bob gets as input $y \in f^{-1}(1)$. Clearly, it holds that $x \neq y$. The goal of Alice and Bob is to find a coordinate i such that $x_i \neq y_i$. Note that there may be more than one possible choice for i, which means that KW_f is a relation rather than a function.

[1] It follows more explicitly from the discussions of Karchmer-Wigderson relations in [Raz90, KKN95], and a similar observation was also made in [KPPY84].

As noted above, Karchmer and Wigderson observed that the communication complexity of KW_f is exactly equal to the depth complexity of f.

In order to study Karchmer-Wigderson relations, Karchmer, Raz, and Wigderson defined the *universal relation*, which is the following computational problem: Alice and Bob as inputs two distinct strings $x, y \in \{0,1\}^n$ respectively, and their goal is to find a coordinate $i \in [n]$ such that $x_i \neq y_i$. It is easy to see that every Karchmer-Wigderson relation reduces to the universal relation. Thus, in order to prove our result, it suffices to devise an efficient randomized protocol for the universal relation[2]. Our result can now be stated as follows.

Theorem 1. *There is a randomized two-round protocol that solves the universal relation over n bits with probability at least $\frac{2}{3}$ by transmitting at most $O(\log^2 n)$.*

We prove this theorem in two steps: In the first step, described in Section 3, we devise an efficient *deterministic* two-round protocol that solves the universal relation in the special case where the inputs differ only on one coordinate (i.e., the Hamming distance between the inputs is 1). In the second step, described in Section 4, we reduce the general case to the foregoing special case along the lines of the Valiant-Vazirani reduction [VV86].

3 The Case of Hamming Distance 1

In this section, we present a *deterministic* two-round protocol that solves the universal relation in the special case where Alice and Bob get inputs that disagree on exactly one coordinate. The protocol will transmit $O(\log n)$ bits. Let us denote by $x, y \in \{0,1\}^n$ the inputs of Alice and Bob respectively. Without loss of generality, we may assume that $n+1$ is a power of 2, so the Hamming code exists over $\{0,1\}^n$: otherwise, the players pad their inputs with 0s in order to satisfy this restriction, and this increases n by a factor of at most 2. The protocol is as follows:

1. In the beginning of the protocol, Alice finds the ball in the Hamming code to which x belongs, and denotes its center by $c_x \in \{0,1\}^n$. Bob does similarly for y, thus obtaining a center $c_y \in \{0,1\}^n$.

2. Alice sends the first message in the protocol, which is an integer from 0 to n that she determines as follows:

   (a) If $x = c_x$, then Alice sends 0.

   (b) Otherwise, Alice sends the unique coordinate j in $[n]$ on which x and c_x disagree.

3. If Alice sent 0:

   (a) Observe that in this case it holds that $c_y = x = c_x$ (since x and y are within distance 1 and x is the center of a ball).

   (b) Thus, Bob sends to Alice the unique coordinate i on which y and $c_y = x$ disagree, and this is the output of the protocol.

4. If Alice sent $j \in [n]$ and $y = c_y$:

   (a) Observe that in this case, it holds that $c_x = y = c_y$ (since x and y are within distance 1 and y is the center of a ball).

[2] We note that the aforementioned protocol of Karchmer solves the universal relation as well.

   (b) Thus, Bob can deduce that j is the coordinate on which x and y differ.
   (c) Hence, Bob sends j back to Alice, and this is the output of the protocol.

5. If Alice sent $j \in [n]$ and $y \neq c_y$:
   (a) Let us denote by $i \in [n]$ the unique coordinate on which x and y disagree.
   (b) Observe that $i \neq j$, since otherwise it would follow that $y = c_y$.
   (c) This implies that y and c_x disagree exactly on the coordinates i and j
   (d) Bob computes the string y' obtained by flipping the j-th coordinate of y, so y' disagrees with c_x only the coordinate i.
   (e) Then, y' is within Hamming distance 1 of c_x, and therefore must be in the ball around c_x in the Hamming code.
   (f) Bob now determines c_x by finding the ball of y' in the Hamming code, and deduces i by finding the unique coordinate on which y' and c_x disagree.
   (g) Bob sends i to Alice, and this is the output of the protocol.

The correctness of the protocol is explained within the foregoing description, and it is not hard to see that it indeed sends $O(\log n)$ bits, as required.

Remark 2. By personal communication, we know that this protocol was discovered independently by Avi Wigderson and Mauricio Karchmer, and by Benjamin Rossman. However, to the best of our knowledge, this is the first time this protocol is published.

4 Proof of Theorem 1

In this section we describe a randomized two-round protocol that solves the universal relation in the general case, thus proving Theorem 1. We describe a public-coin protocol, and it can be converted into a private-coin protocol using Newman's lemma [New91].

The idea is to reduce the general case to the special case of Section 3 along the lines of the Valiant-Vazirani reduction: For start, suppose that the parties knew that their inputs disagree on exactly ℓ coordinates. In this case, the parties could choose a random set of coordinates of size $\frac{n}{2\ell}$, and with constant probability this set would contain exactly one coordinate on which they disagree. Thus, the parties could project their inputs to this set and use the protocol of Section 3. The next step in the argument is to observe that this idea works even if the parties only have an estimate of ℓ up to a factor of 2. Finally, since the parties do not have such an estimate of ℓ, they try different values of $\ell = 1, 2, 4, 8, \ldots, n$ and apply the foregoing protocol in parallel for each of those values. Details follow.

Formally, the protocol is defined as follows. Suppose that Alice and Bob get as inputs the strings $x, y \in \{0,1\}^n$ respectively, so $x \neq y$. They perform the following steps for every $t \in \{1, \ldots, \lceil \log n \rceil + 1\}$:

1. Using the public coins, choose a random set of coordinates $S \subseteq [n]$ by putting each coordinate in S with probability 2^{-t} independently.

2. Execute the protocol of Section 3 on $x|_S$ and $y|_S$, thus obtaining a coordinate $i \in [n]$.

3. If $x_i \neq y_i$, output the coordinate i and end the protocol.

It is easy to see that the protocol indeed transmits $O(\log^2 n)$ bits. Moreover, the protocol can be implemented in two rounds, since the above steps can be performed in parallel for all values of t. It remains to show that it outputs a coordinate $i \in [n]$ on which $x_i \neq y_i$ with good probability.

Fix specific inputs to the players $x, y \in \{0,1\}^n$, and let $I \subseteq [n]$ be the set of coordinates on which x and y differ. Observe that the protocol succeeds whenever, in the foregoing steps, the random set S contains exactly one coordinate on which x and y differ (i.e., $|S \cap I| = 1$). When $t = \lceil \log |I| \rceil + 1$, the probability that the latter event happens is

$$\begin{aligned}
\Pr\left[|S \cap I| = 1\right] &= \sum_{i \in I} \Pr\left[S \cap I = \{i\}\right] \\
&= \sum_{i \in I} \Pr\left[i \in S\right] \cdot \prod_{j \in I \setminus \{i\}} \Pr\left[j \notin S\right] \\
&= |I| \cdot \frac{1}{2^t} \cdot \left(1 - \frac{1}{2^t}\right)^{|I|-1} \\
\text{(Since } |I| \geq 2^{t-2}\text{)} &\geq \frac{1}{4} \cdot \left(1 - \frac{1}{2^t}\right)^{|I|} \\
\text{(Since } |I| \leq 2^{t-1}\text{)} &\geq \frac{1}{4} \cdot \left(1 - \frac{1}{2 \cdot |I|}\right)^{|I|} \\
&\geq \frac{1}{4} \cdot \left(1 - \frac{|I|}{2 \cdot |I|}\right) \\
&\geq \frac{1}{8}.
\end{aligned}$$

It follows that the protocol succeeds with probability at least $\frac{1}{8}$. Note that the protocol is a zero-error protocol, i.e., the parties always know whether they succeeded or not. Hence, th success probability can be amplified to $\frac{2}{3}$ by repeating the protocol a constant number of times, while maintaining a communication complexity of $O(\log^2 n)$. This concludes the proof.

Acknowledgement. We would like to thank Avi Wigderson, Benjamin Rossman and Daniel Kane for valuable discussions and ideas.

References

[Ajt83] Miklós Ajtai. Σ^1_1-formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.

[FSS84] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.

[GS91] Michelangelo Grigni and Michael Sipser. Monotone separation of Logspace from NC. In *Structure in Complexity Theory Conference*, pages 294–298, 1991.

[Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *STOC*, pages 6–20, 1986.

[KKN95] Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. *SIAM J. Discrete Math.*, 8(1):76–92, 1995.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity.* Cambridge University Press, 1997.

[KPPY84] Maria M. Klawe, Wolfgang J. Paul, Nicholas Pippenger, and Mihalis Yannakakis. On monotone formulae with restricted depth (preliminary version). In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1984, Washington, DC, USA*, pages 480–487, 1984.

[KRW95] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.

[KW90] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990.

[New91] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67–71, 1991.

[Raz90] Alexander A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.

[RW89] Ran Raz and Avi Wigderson. Probabilistic communication complexity of boolean relations (extended abstract). In *FOCS*, pages 562–567, 1989.

[RW92] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992.

[Sud01] Madhu Sudan. Algorithmic introduction to coding theory (lecture notes), 2001. Available from `http://theory.csail.mit.edu/~madhu/FT01/`.

[VV86] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.