

$AC^0[p]$ Lower Bounds Against MCSP via the Coin Problem

Alexander Golovnev

Harvard University, Cambridge, USA
alexgolovnev@gmail.com

Rahul Ilango

Rutgers University, New Brunswick, USA
rahul.ilango@rutgers.edu

Russell Impagliazzo

University of California San Diego, USA
russell@cs.ucsd.edu

Valentine Kabanets

Simon Fraser University, Burnaby, Canada
kabanets@cs.sfu.ca

Antonina Kolokolova

Memorial University of Newfoundland, St. John's, Canada
kol@mun.ca

Avishay Tal

Stanford University, USA
avishay.tal@gmail.com

Abstract

Minimum Circuit Size Problem (MCSP) asks to decide if a given truth table of an n -variate boolean function has circuit complexity less than a given parameter s . We prove that MCSP is hard for constant-depth circuits with mod p gates, for any prime $p \geq 2$ (the circuit class $AC^0[p]$). Namely, we show that MCSP requires d -depth $AC^0[p]$ circuits of size at least $\exp(N^{0.49/d})$, where $N = 2^n$ is the size of an input truth table of an n -variate boolean function. Our circuit lower bound proof shows that MCSP can solve the coin problem: distinguish uniformly random N -bit strings from those generated using independent samples from a biased random coin which is 1 with probability $1/2 + N^{-0.49}$, and 0 otherwise. Solving the coin problem with such parameters is known to require exponentially large $AC^0[p]$ circuits. Moreover, this also implies that MAJORITY is computable by a non-uniform AC^0 circuit of polynomial size that also has MCSP-oracle gates. The latter has a few other consequences for the complexity of MCSP, e.g., we get that any boolean function in NC^1 (i.e., computable by a polynomial-size formula) can also be computed by a non-uniform polynomial-size AC^0 circuit with MCSP-oracle gates.

2012 ACM Subject Classification Theory of computation \rightarrow Circuit complexity; Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases Minimum Circuit Size Problem (MCSP), circuit lower bounds, $AC^0[p]$, coin problem, hybrid argument, MKTP, biased random boolean functions

Digital Object Identifier 10.4230/LIPIcs.ICALP.2019.66

Category Track A: Algorithms, Complexity and Games

Related Version A full version of the paper is available at <https://eccc.weizmann.ac.il/report/2019/018/>.

Funding *Alexander Golovnev*: Supported by a Rabin Postdoctoral Fellowship.

Russell Impagliazzo: Work supported by a Simons Investigator Award from the Simons Foundation.

Valentine Kabanets: Supported in part by an NSERC Discovery grant.



© Alexander Golovnev, Rahul Ilango, Russell Impagliazzo, Valentine Kabanets, Antonina Kolokolova, and Avishay Tal; licensed under Creative Commons License CC-BY

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).
Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;
Article No. 66; pp. 66:1–66:15



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Antonina Kolokolova: Supported in part by an NSERC Discovery grant.

Avishay Tal: Supported by a Motwani Postdoctoral Fellowship and by NSF grant CCF-1763299. Part of this work was done while the last four authors were visiting Simons Institute for the Theory of Computing.

Acknowledgements This work was partly carried out while many of the authors were visiting the Simons Institute for the Theory of Computing in association with the DIMACS/Simons Collaboration on Lower Bounds in Computational Complexity, which is conducted with support from the National Science Foundation. We also thank Chris Umans and Ronen Shaltiel for helpful discussions in the early stages of this project (during the Dagstuhl 2018 workshop on “Algebraic Methods in Complexity”). We thank Eric Allender and Shuichi Hirahara for their comments, and special thanks to Eric for pointing us to the paper of Dančik [9] and the discussion of various circuit and formula complexity measures for constant-depth circuit models. We are grateful to our anonymous reviewers for helpful comments on this paper.

1 Introduction

Minimum Circuit Size Problem (MCSP) asks to decide if a given boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ (presented by its truth table of length $N = 2^n$) can be computed by a boolean circuit of size at most s , for a given parameter $0 \leq s \leq 2^n$. There is no nontrivial algorithm currently known for MCSP other than the “brute force” enumeration of all circuits of size up to s and checking if any one of them computes f . On the other hand, while MCSP is obviously in NP, it is a major open question to decide if MCSP is NP-complete (and there is a growing list of research papers providing arguments for and against various NP-completeness reductions to MCSP [17, 2, 4, 3, 6, 13, 12, 24]).

Another natural question is to prove circuit lower bounds (for restricted circuit models) for MCSP. Here some results are known. Allender et al. [2] showed that MCSP requires super-polynomial-size AC^0 circuits (constant-depth circuits with AND, OR, and NOT gates). Hirahara and Santhanam [11] proved that MCSP requires almost quadratic-size De Morgan formulas.

It was an open question [5, 25] to prove that MCSP requires super-polynomial $AC^0[p]$ circuits (constant-depth circuits with AND, OR, NOT and mod p counting gates), for a prime $p > 0$. We resolve this question in the present paper. Our main result is that MCSP requires d -depth $AC^0[p]$ circuits of size at least $\exp(N^{0.49/d})$, where $N = 2^n$ is the size of an input truth table of an n -variate boolean function.

Previous proof methods of circuit lower bounds for MCSP

The lack of NP-completeness reductions to MCSP and the scarcity of circuit lower bounds for MCSP underscore the general phenomenon that there are very few known reductions to MCSP. The main (if not the only one) use of MCSP inside known reductions is to “break” pseudorandom function generators, an idea going back to the celebrated paper of Razborov and Rudich [28] on “natural proofs”. The point is that known candidate constructions of pseudorandom function generators produce pseudorandom functions that do have “small” circuit complexity, whereas truly random functions are known to require “high” circuit complexity. Thus, an assumed efficient MCSP algorithm can distinguish between the truth tables of such pseudorandom functions and those of truly random functions, thereby “breaking” the pseudorandom function generator.

Both previously known circuit lower bounds for MCSP by Allender et al. [2] and by Hirahara and Santhanam [11] used MCSP's ability to break pseudorandom function generators together with the existence of known pseudorandom (function) generators that are provably secure against AC^0 and quadratic-size De Morgan formulas, respectively. The same approach cannot be applied to the case of $AC^0[p]$ circuits as we currently do not have any strong enough pseudorandom generators secure against $AC^0[p]$!

Our approach

Our approach instead is to reduce the Majority function to MCSP, and use the known $AC^0[p]$ lower bounds against Majority [27, 31]. In fact, we give a reduction to MCSP from the coin problem where one is asked to distinguish between an N -bit uniformly random string, and an N -bit string sampled so that each bit is independently set to 1 with probability $1/2 - \epsilon$ (and to 0 otherwise), for some parameter $\epsilon > 0$. We then use the result of Shaltiel and Viola [29] showing that any algorithm solving such a coin problem yields an efficient algorithm for computing the Majority function on inputs of length $1/\epsilon$. To conclude super-polynomial-size $AC^0[p]$ circuit lower bounds for MCSP from the known lower bounds for the Majority function, we need to be able to solve the coin problem for N -bit strings with the parameter $\epsilon < 1/\text{poly}(\log N)$.

Here is some intuition why MCSP could be useful for solving the coin problem. For $N = 2^n$, an N -bit random string has binary entropy N . On the other hand, N -bit strings sampled using a biased coin with probability $p = 1/2 - \epsilon$ of being 1 would likely have close to $pN < N/2$ ones only, and so come from a smaller set of about $\binom{N}{pN} \approx 2^{H(p) \cdot N}$ strings of size N , where H is the binary entropy function. Information-theoretically, we can describe each string with at most pN ones using about $H(p) \cdot N$ bits. For $p \ll 1/2$, we have that $H(p) \cdot N \ll N$, and so most biased functions have a description of bit complexity much less than N . If somehow we could extend this information-theoretic argument to show that most biased functions will have *circuit complexity* noticeably smaller than that of random functions, we'd be done because MCSP would be able to distinguish between random functions (of higher circuit complexity) and random biased functions (of lower circuit complexity).

Lupanov in 1965 [21] proved that, indeed, biased random functions have circuit complexity smaller than that of random boolean functions. However, Lupanov's result applies only to the case of bias probability $p = 1/2 - \epsilon$ for large (close to constant) ϵ only, and doesn't give anything useful for our case of $\epsilon < 1/\text{poly}(\log N)$. To circumvent the lack of tighter circuit upper bounds for slightly biased random functions, we employ two new ideas.

First, we show that the circuit complexity of q -random functions is very *tightly concentrated* around its expectation, for every probability q . This can be proved using a simple martingale argument (McDiarmid's Inequality [22]). The point is that the circuit complexity of a given n -variate boolean function f changes by at most $O(n)$ when we change the value of f on exactly one n -bit input (which can be simply hard-wired into the new circuit for the modified function).

Secondly, we use a *hybrid argument*. By Lupanov's result [21], one can show that for $p = 0.01$, almost all p -biased random functions will have circuit complexity noticeably smaller than $2^n/n$, and in particular, the *expected* circuit complexity of a p -biased random function is at most $0.1 \cdot 2^n/n$. On the other hand, for $p' = 1/2$, well-known counting arguments show that almost all such random functions have circuit complexity very close to $2^n/n$, and in particular, the expected circuit complexity of a random function is at least $0.9 \cdot 2^n/n$. Imagine we have t equally spaced probabilities between $p = 0.01$ and $p' = 1/2$, for some number t to be chosen. Then by the hybrid argument, there will exist two successive probabilities q and $q' \approx q + 1/t$, where the *expected circuit sizes* for q -random and q' -random functions differ by at least $\Omega(2^n/n)/t$.

By the “concentration around the expected circuit size” result mentioned above, we conclude that MCSP is able to distinguish between most q -random boolean functions and most $q' \approx q + 1/t$ -random ones. By re-scaling, we conclude that MCSP is able to distinguish between $1/2 - 1/t$ -random function and $1/2$ -random ones, i.e., that MCSP solves the coin problem for the bias $\epsilon = 1/t$. Finally, our concentration result is strong enough to allow us to choose $t = 2^{\Omega(n)}$, which yields an exponential $AC^0[p]$ circuit lower bound for MCSP. (To get the quantitatively strongest circuit lower bound for MCSP, we use the recent $AC^0[p]$ lower bounds for the coin problem by [18], rather than apply the reduction from the Majority function to the coin problem from [29].)

Other results

We are able to generalize our $AC^0[p]$ circuit lower bounds to several natural variants of MCSP. For a circuit class \mathcal{C} , let \mathcal{C} -MCSP denote the MCSP problem asking about the \mathcal{C} -type circuit complexity of a given truth table. For example, \mathcal{C} can be the class AC^0 , where we ask about the gate complexity of a smallest AC^0 circuit computing a given boolean function, or \mathcal{C} can be the class of boolean formulas, where we ask about the formula (leaf) complexity of a smallest formula. We show that for both such cases of \mathcal{C} , the problem \mathcal{C} -MCSP requires exponential-size $AC^0[p]$ circuits. This generalization requires us to re-visit Lupanov’s general circuit upper bounds for biased random functions. We provide new “hashing-based” arguments for such circuit constructions that apply to the case of formulas as well as constant-depth circuits and formulas.

As a corollary of our reduction of the coin problem to MCSP and some previous results, we obtain that every function in NC^1 can be computed by a non-uniform AC^0 circuit with MCSP oracle gates. We also show that at least one of the following lower bounds must be true: either $NEXP \not\subseteq P/poly$, or $MCSP \notin ACC^0$.

Finally, we give a new coin-problem based proof of $AC^0[p]$ circuit lower bounds for MKTP, a Kolmogorov-complexity variant of MCSP, re-proving the result of [5].

The rest of the paper

We give the necessary definitions and facts in Section 2. We prove the aforementioned circuit complexity concentration result for random biased functions in Section 3, and then prove our $AC^0[p]$ circuit lower bound for MCSP in Section 4. In Section 5, we give corollaries of our main result. The lower bound for MKTP is given in Section 7. We generalize our lower bounds to the case of \mathcal{C} -MCSP in Section 6. Section 8 lists some open questions.

2 Preliminaries

2.1 Complexity basics

For a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, let $size(f)$ denote the size of a smallest circuit (using AND, OR, and NOT gates) that computes f . Let $size(n)$ be the maximum of $size(f)$, over all n -variate boolean functions f . Finally, define $s_n = \mathbf{E}[size(f)]$ to be the average of circuit complexities over all n -variate boolean functions.

Below, all logarithms are base 2, unless stated otherwise.

Building on the work by Shannon [30], Lupanov [19] proved the following lower and upper bounds for $size(n)$.

► **Theorem 2.1** (Lupanov [19]). *For all sufficiently large $n \in \mathbb{N}$,*

$$\text{size}(n) \leq \frac{2^n}{n} + O\left(\frac{2^n \log n}{n^2}\right).$$

Moreover, all but $o(1)$ fraction of uniformly random n -variate boolean functions f require

$$\text{size}(f) \geq \frac{2^n}{n} + \Omega\left(\frac{2^n \log n}{n^2}\right).$$

For the case of n -variate functions with a fixed fraction of 1 inputs, Lupanov [21] proved the following generalization of his earlier bounds.

► **Theorem 2.2** (Lupanov [21]). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be any boolean function that has the value 1 on $k \leq 2^{n-1}$ inputs, where $k \in \Omega(2^n)$. Then, for all sufficiently large $n \in \mathbb{N}$,*

$$\text{size}(f) \leq \frac{\log \binom{2^n}{k}}{\log \log \binom{2^n}{k}} + O\left(\frac{2^n \log n}{n^2}\right).$$

Moreover, all but $o(1)$ fraction of random such functions f require

$$\text{size}(f) \geq \frac{\log \binom{2^n}{k}}{\log \log \binom{2^n}{k}} + \Omega\left(\frac{2^n \log n}{n^2}\right).$$

2.2 Probability basics

► **Theorem 2.3** (McDiarmid's Inequality [22]). *Let $X_1, \dots, X_N \in \{0, 1\}$ be independent random variables. Let $f: \{0, 1\}^n \rightarrow \mathbb{R}$ be any function such that, for some function $c = c(N)$, for all $1 \leq i \leq N$ and for all $b_1, \dots, b_N, \tilde{b}_i \in \{0, 1\}$, it holds that*

$$\left| f(b_1, \dots, b_N) - f(b_1, \dots, b_{i-1}, \tilde{b}_i, b_{i+1}, \dots, b_N) \right| \leq c.$$

Then, for any $\lambda > 0$,

$$\Pr [|f(X_1, \dots, X_N) - \mathbf{E}[f(X_1, \dots, X_N)]| \geq \lambda] \leq 2 \cdot \exp\left(-\frac{2\lambda^2}{Nc^2}\right).$$

Roughly speaking, the inequality states that with high probability the value of f will be of distance at most $O(\sqrt{N} \cdot c)$ around its mean.

2.3 Coin problem

A coin problem is the problem to distinguish between two coins (boolean-valued random variables), where one coin has probability p of being 1, and the other coin probability $q > p$. Usually, $p = 1/2 - \epsilon$ and $q = 1/2 + \epsilon$, for some $\epsilon > 0$; or, $p = 1/2 - \epsilon$ and $q = 1/2$. By a simple “translation argument”, it is possible to show that all of these problems are essentially equivalent. For completeness, we state this argument next.

▷ **Claim 2.4 (Translation Argument)**. Let $0 < p \leq q < 1$, $\epsilon > 0$. Suppose C is a circuit of size S that solves the $p(1 - \epsilon)$ versus p coin problem on inputs of length N with advantage α . Then, there exists a circuit \tilde{C} of size at most S that solves the $q(1 - \epsilon)$ versus q coin problem on inputs of length N with advantage at least α .

66:6 AC⁰[p] Lower Bounds Against MCSP via the Coin Problem

Proof. For $p \in (0, 1)$, we denote by μ_p the product distribution on $\{0, 1\}^N$ where each bit is independently sampled to be 1 with probability p , and 0 otherwise. Let C be a circuit of size S that solves the $p(1 - \varepsilon)$ versus p coin problem with advantage

$$\alpha := \Pr_{x \sim \mu_{p(1-\varepsilon)}} [C(x)] - \Pr_{x \sim \mu_p} [C(x)].$$

We construct a distribution over circuits C' that achieves the same advantage on the $q(1 - \varepsilon)$ versus q coin problem. The randomized circuit C' is defined as follows: “For each input bit x_i , let $x'_i = x_i$ with probability p/q and $x'_i = 0$ otherwise. Apply C on (x'_1, \dots, x'_N) .” Note that, by design, if x is distributed according to a μ_q then x' is distributed according to μ_p , and if x is distributed according to $\mu_{q(1-\varepsilon)}$, then x' is distributed according to $\mu_{p(1-\varepsilon)}$. We get a distribution over circuits C' that solves the $q(1 - \varepsilon)$ versus q coin problem with advantage at least α . By averaging, there must exist a deterministic circuit \tilde{C} that solves the $q(1 - \varepsilon)$ versus q coin problem with advantage at least α . Note that \tilde{C} is obtained from C' by fixing the internal randomness that C' used to decide for each $1 \leq i \leq N$ whether to set $x'_i = x_i$ or $x'_i = 0$. With those choices fixed, $\tilde{C}(x_1, \dots, x_N)$ is just a restriction of $C(x_1, \dots, x_N)$ where some x_i 's are set to 0. Hence, the size of \tilde{C} is at most that of C , as claimed. \triangleleft

► **Theorem 2.5** ([29]). *Let A be an algorithm that distinguishes, with constant distinguishing probability, between n -bit uniformly random strings, and n -bit strings sampled so that each bit is independently set to 1 with probability $1/2 - \varepsilon$ (and to 0 otherwise). Then there is a non-uniform AC⁰ circuit of size $\text{poly}(n/\varepsilon)$ that computes the majority function on binary inputs of length $1/\varepsilon$, using A -oracle gates.*

Using the theorem above as well as the well-known lower bound for the majority function against AC⁰[p] circuits, for any constant prime p , we can deduce that any algorithm solving the coin problem with bias ε on n -bit inputs requires AC⁰[p] depth d circuits of size at least $\exp((1/\varepsilon)^{1/O(d)})$. This lower bound has been recently sharpened.

► **Theorem 2.6** ([18]). *Let A be a boolean function that distinguishes, with constant distinguishing probability, between n -bit uniformly random strings, and n -bit strings sampled so that each bit is independently set to 1 with probability $1/2 - \varepsilon$ (and to 0 otherwise). Then any depth d AC⁰[p] circuit computing A must have size at least $\exp(\Omega((1/\varepsilon)^{1/(d-1)}))$.*

3 Concentration of Circuit Complexity

For every $n \geq 1$, let μ be any product distribution over $\{0, 1\}^N$, where $N = 2^n$. Recall that $\text{size}(f)$ is the size of a smallest circuit computing a boolean function f . For each integer $n \geq 1$, define

$$s_n^\mu = \mathbf{E}_{f \sim \mu} [\text{size}(f)],$$

the expectation of $\text{size}(f)$ over n -variate boolean functions f whose N -bit truth tables are sampled according to μ . We show that a random n -variate boolean function sampled from μ is likely to have its circuit complexity very close to the expected circuit complexity s_n^μ .

► **Theorem 3.1.** *For μ and s_n^μ as defined above, if a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is chosen at random according to distribution μ , then, with probability at least $1 - 2^{-n}$,*

$$|\text{size}(f) - s_n^\mu| \leq \sqrt{N} \cdot n^2.$$

Proof. Observe that for any two n -variate boolean functions f and g that differ on exactly one input $a \in \{0, 1\}^n$, we have that $|size(f) - size(g)| \leq O(n)$. Indeed, suppose, without loss of generality, that $size(f) \leq size(g)$. Then we can compute g as follows: “Given $z \in \{0, 1\}^n$, check if $z = a$. If so, then output the bit $b = g(a)$. Otherwise, use the circuit for f to output $f(z)$.” More precisely, if C is a circuit computing f , then the circuit D for g can be described as

$$D(z) := (b \wedge (\wedge_{i=1}^n (z_i = a_i))) \vee (C(z) \wedge (\vee_{i=1}^n (z_i \neq a_i))), \tag{1}$$

where $(z_i = a_i)$ is defined to be z_i for $a_i = 1$, and the negation \bar{z}_i for $a_i = 0$; and $(z_i \neq a_i)$ is the negation of $(z_i = a_i)$. Clearly, the size of D is that of C plus $O(n)$.

Let $X_1, \dots, X_N \in \{0, 1\}$ be independent random variables sampled from the product distribution μ . We apply McDiarmid’s Inequality of Theorem 2.3 to the function $size: \{0, 1\}^N \rightarrow \mathbb{N}$, which, as we just argued, differs by at most $c = O(n)$ on any two truth tables that agree in all but one coordinate. We get the desired concentration result by choosing $\lambda = \sqrt{N} \cdot n^2$. That is, all but $\exp(-n)$ fraction of μ -random n -variate boolean functions f have their circuit size $size(f)$ within $\sqrt{N} \cdot n^2$ of the expected circuit size s_n^μ . ◀

4 Main theorem

► **Theorem 4.1.** *Let $p \geq 2$ be any prime. For any depth $d > 0$ and large enough input size $N = 2^n$, MCSP on N -bit truth tables requires depth d $AC^0[p]$ circuits of size $\exp(\Omega(N^{0.49/(d-1)}))$.*

Proof. Let $t = \lceil 2^{0.49n} \rceil$. Consider an arithmetic sequence of probabilities $(p_0, p_1, p_2, \dots, p_t)$ with $p_0 = 0.01$, $p_t = 0.5$ and $p_i = p_0 + i \cdot 0.49/t$. For each $i = 0, \dots, t$, let μ^i be the product distribution on $\{0, 1\}^N$, where each bit is independently sampled to be 1 with probability p_i , and 0 with probability $1 - p_i$. Let

$$s^i = s_n^{\mu^i} = \mathbf{E}_{f \sim \mu^i} [size(f)].$$

By Lupanov’s estimates of Theorem 2.1, we have

$$s^t \geq (1 - o(1)) \cdot \frac{2^n}{n} \geq 0.9 \cdot \frac{2^n}{n}.$$

By the Chernoff bound, almost all n -variate boolean functions sampled according to μ^t will assume the value 1 on at most $k = 0.011 \cdot N$ inputs. By Theorem 2.2, we have

$$s^0 \leq H(0.011) \cdot \frac{2^n}{n} + o\left(\frac{2^n}{n}\right) \leq 0.1 \cdot \frac{2^n}{n},$$

where $H()$ is the binary entropy function. It follows that

$$s^t - s^0 \geq 0.8 \cdot \frac{2^n}{n}.$$

This means that there must be an i such that $s^{i+1} \geq s^i + \Omega(2^{0.51n}/n)$. Let $s^* = (s^i + s^{i+1})/2$. By circuit complexity concentration given in Theorem 3.1, we get that, with high probability, n -variate random boolean functions sampled from μ^i have circuit size smaller than s^* , and those sampled from μ^{i+1} have circuit size larger than s^* . Hence, $MCSP(x, s^*)$ can distinguish between μ^i and μ^{i+1} , with a constant distinguishing probability.

Finally, assume by contradiction that $\text{MCSP} \in \text{AC}^0[p]$ of size S and depth d . Let n be large enough, and let $N = 2^n$. Then, by fixing the second input of MCSP to s^* , we get an $\text{AC}^0[p]$ circuit that on N -bit inputs solves the coin problem distinguishing between p_i and p_{i+1} . By Claim 2.4, there exists a circuit of size at most S and depth at most d that solves the $(1 - \varepsilon)/2$ versus $1/2$ coin problem for $\varepsilon = 1 - p_i/p_{i+1} = \Theta(1/t) = \Theta(N^{-0.49})$. However, the latter implies by Theorem 2.6 that $S \geq \exp(\Omega(\varepsilon^{-1/(d-1)})) = \exp(\Omega(N^{0.49/(d-1)}))$. ◀

5 Consequences

Below, whenever we talk about circuit classes such as AC^0 , ACC^0 , and TC^0 , we mean *non-uniform* circuit classes.

Using Theorem 2.5, we get the following corollary to Theorem 4.1 regarding the MAJORITY function, denoted MAJ.

► **Corollary 5.1.** $\text{MAJ} \in (\text{AC}^0)^{\text{MCSP}}$.

Combined with the inclusion $\text{NC}^1 \subseteq (\text{TC}^0)^{\text{MCSP}}$ of [25], Corollary 5.1 yields the following.

► **Corollary 5.2.** $\text{NC}^1 \subseteq (\text{AC}^0)^{\text{MCSP}}$.

In fact, using the same techniques, we can prove something more general.

► **Theorem 5.3.** *Let $\mathcal{C} \subseteq \text{P/poly}$ be any complexity class that has a complete problem under TC^0 -computable reductions that is also random-self-reducible via a TC^0 -computable reduction. Then we have*

$$\mathcal{C} \subseteq (\text{AC}^0)^{\text{MCSP}}.$$

Proof sketch. It follows from [7] that any function $f \in \text{P/poly}$ has a non-uniform $(\text{TC}^0)^{\text{MCSP}}$ circuit C of polynomial size that agrees with f on all but an inverse polynomial fraction of inputs. If f is random-self-reducible via a TC^0 reduction, we can recover from C a new polynomial-size $(\text{TC}^0)^{\text{MCSP}}$ circuit computing f exactly (on all inputs). Applying Corollary 5.1 concludes the proof. ◀

As the class GapL has Determinant as a complete problem under AC^0 reductions (see [1] for a survey on logspace counting complexity classes), we get the following.

► **Corollary 5.4.** $\text{GapL} \subseteq (\text{AC}^0)^{\text{MCSP}}$.¹

The following is a non-uniform version of a similar “Karp-Lipton”-style “collapse” theorem from [14], which we state just for the class EXP .

► **Theorem 5.5.** *If $\text{EXP} \subseteq \text{P/poly}$, then $\text{EXP} \subseteq (\text{AC}^0)^{\text{MCSP}}$.*

Proof. Using TC^0 -computable locally list-decodable binary codes of [10], we get that EXP contains a complete language that is random-self-reducible via a TC^0 -computable reduction. We then appeal to Theorem 5.3. ◀

¹ The potentially bigger class DET is the class of languages that are NC^1 -Turing reducible to computing the determinant of an integer matrix [8]. It is not immediately clear if $\text{DET} \subseteq (\text{AC}^0)^{\text{MCSP}}$. Perhaps, one can use the techniques of [5] who showed such a result for MKTP .

We do not know if MCSP is NP-complete. There is a line of research showing that MCSP (or its variants) can't be NP-complete under very restricted kinds of reductions (e.g., “local” reductions of [24]). One corollary of Theorem 5.5 is that it will be difficult to rule out non-uniform Turing AC^0 reductions from SAT to MCSP.

► **Corollary 5.6.** *If $SAT \notin (AC^0)^{MCSP}$, then $EXP \not\subseteq P/poly$.*

Finally, while we don't know how to disprove that $MCSP \in ACC^0$, we get that at least some lower bound must be true.

► **Corollary 5.7.** *Either $NEXP \not\subseteq P/poly$, or $MCSP \notin ACC^0$.*

Proof. Towards a contradiction, suppose that both (1) $NEXP \subseteq P/poly$, and (2) $MCSP \in ACC^0$. By the Easy Witness Lemma of [15], (1) implies that $NEXP = EXP$. Then by Theorem 5.5, we get that $NEXP \subseteq (AC^0)^{MCSP}$. Combining this with (2) yields that $NEXP \subseteq ACC^0$. But the latter contradicts the known lower bound of [32]. ◀

6 Generalizations

Here we show that, for a number of typical circuit classes \mathcal{C} , our lower bound proof (and a reduction from MAJORITY) works also for \mathcal{C} -MCSP. In particular, we will show that both AC^0 -MCSP and Formula-MCSP require exponential $AC^0[p]$ circuit lower bounds.

Our lower bound for MCSP used two main ingredients: (1) circuit size concentration for random (biased) boolean functions, and (2) a noticeable difference between most likely circuit sizes for uniformly random and biased boolean functions (where each bit of the truth table is 1 with a small constant probability, say 0.01).

For property (1), we note that the concentration argument only needs the Lipschitz property of a given circuit size measure, which comes from the fact that changing a boolean function on a single n -bit input may change the circuit size of the function by at most $O(n)$ additive term. This holds for virtually every reasonable circuit model, as the proof of Theorem 3.1 shows; there is a potential increase in depth for constant-depth circuits, but this can be avoided for the case where the circuit size is defined to be the total number of gates in the constant-depth circuit (see the proof of Corollary 6.2).

For property (2), we need a Shannon-style counting argument to show that most random n -variate functions have at least certain size S in a given circuit model \mathcal{C} , as well as a Lupanov-style argument that (most) boolean functions with very few (a small constant fraction α of) 1s have \mathcal{C} -circuit complexity at most some constant fraction δ of S , for some $0 < \delta < 1$ (dependent on α).

Property (2) is known for the case of boolean circuits, as implied by Lupanov's Theorem 2.2, and is known for formulas, by the work of Pippenger [26]. Moreover, it is possible to use the celebrated constructions of Lupanov, giving tight upper bounds for circuit complexity [19] and formula complexity [20] for all boolean functions, to show that biased random functions have relatively small circuits as well as small formulas (see the full version of this paper on ECCC for more details). However, we give a different argument (based on some hashing ideas) that will allow us to reduce the problem of showing small circuit complexity of a random biased boolean function to the known worst-case upper bounds for boolean function on fewer variables. Using such known worst-case upper bounds for the classes of circuits, formulas, and constant-depth circuits (counting the number of gates), we then obtain the required upper bounds for the circuit complexity of constant-biased random functions in these circuit models.

► **Theorem 6.1.** *Let $0 < \alpha < 1/2$ be any constant. For all but $o(1)$ fraction of α -biased n -variate random functions f , we have*

1. $circuit\text{-}size(f) \leq O(\alpha \cdot \log 1/\alpha) \cdot \frac{2^n}{n}$,
2. $(AC^0)\text{-}formula\text{-}size(f) \leq O(\alpha \cdot \log 1/\alpha) \cdot \frac{2^n}{\log n}$, and
3. $AC^0\text{-}circuit\text{-}size(f) \leq O(\alpha \cdot \log 1/\alpha) \cdot 2^{n/2}$ (where consider the gate complexity of a given constant-depth circuit); moreover, the upper bound is for some fixed depth $d_0 > 0$ (independent of f).

We omit the proof of Theorem 6.1 due to space constraints (see the full version of this paper on ECCC for more details).

Using Theorem 6.1, we conclude the following.

► **Corollary 6.2.** *Let \mathcal{C} be the class of general circuits, or formulas, or constant-depth AC^0 circuits. For any prime $p \geq 2$ and any depth $d > 0$ and large enough input size $N = 2^n$, \mathcal{C} -MCSP on N -bit truth tables requires depth d $AC^0[p]$ circuits of size at least $\exp(\Omega(N^{0.49/(d-1)}))$.*

Proof. As observed earlier, our lower bound proof for MCSP requires three ingredients: the Lipschitz property of the circuit complexity measure, a Shannon-style lower bound on the complexity measure for random n -variate boolean functions, and a $O(\alpha \log(1/\alpha))$ factor smaller upper bound on the complexity measure for random α -biased boolean functions (which can be made an arbitrary constant factor ϵ smaller than the corresponding Shannon-style upper bound by choosing the constant bias $\alpha > 0$ to be sufficiently small).

The Lipschitz property is easily seen to hold for both general circuits and formulas. For constant-depth circuits, where we count the number of gates, it also holds, provided the depth of our circuits is at least 3. We sketch the argument next.

We may assume that the circuit has alternating levels of ANDs and ORs, with negations on the bottom variables level. Without loss of generality, the top gate is an OR. (The other case is symmetric.) Case 1. We want to flip the value on $a \in \{0, 1\}^n$ from 0 to 1. Add an AND of $x_i^{a_i}$'s, and feed this AND into the top OR gate. (Use just one extra gate.) Case 2: We want to flip from 1 to 0 on $a \in \{0, 1\}^n$. Add an OR of $x_i^{1-a_i}$'s, and feed this OR into every AND-gate just one level below the top OR-gate. (Use just one extra gate.) Note that the depth doesn't change, if the original circuit is of depth $d \geq 3$.

The Shannon-style lower bounds for random n -variate functions are known for general circuits $\Omega(2^n/n)$, formulas $\Omega(2^n/(\log n))$, and constant-depth circuits $\Omega(2^{n/2})$ (see, e.g., [16]). Finally, Theorem 6.1 gives matching upper bounds for biased functions. The proof follows. ◀

7 Circuit lower bounds for MKTP via the coin problem

Here we show how to re-prove a known $AC^0[p]$ circuit lower bound for MKTP [5], using the coin problem. This was the starting point in our attempt to prove an $AC^0[p]$ lower bound for MCSP. For MKTP, we managed to show that biased random strings have a noticeably smaller KT complexity than that of uniformly random strings, where the bias $1/2 - \epsilon$ can be chosen for a sufficiently small ϵ so that we immediately get an $AC^0[p]$ circuit lower bound for MKTP using Theorem 2.6. We provide the details next.

We first define the KT complexity [2]. Fix a universal random-access Turing machine U . The KT complexity of a string $x \in \{0, 1\}^N$ is defined as the

$$\min\{|d| + t \mid \forall 0 \leq i \leq N + 1 \ U^d(i) = x_i \text{ in at most } t \text{ steps}\},$$

and $x_{N+1} = \perp$. In other words, x can be computed at every position i in time at most t by a TM U that has random access to some binary string d , and we want to minimize the total sum of the length of such an auxiliary string d and the time bound t .

MKTP is then naturally defined as follows: Given $x \in \{0, 1\}^N$ and a parameter s , decide if the KT complexity of x is at most s .

To prove a super-polynomial $AC^0[p]$ circuit lower bound for MKTP via the coin problem approach (using Theorem 2.6), it would suffice to show that the MKTP oracle can distinguish between uniformly random N -bit strings and those where each bit is sampled, independently, with probability $1/2 - \epsilon$, for some $\epsilon \ll 1/\text{poly log } N$. We'll show how to do this for $\epsilon = N^{-\gamma}$, for some $\gamma > 0$ (in fact, for $\gamma \approx 1/6$).

What we need is to show that a random biased N -bit string can be compressed to have its KT complexity noticeably smaller than that of a uniformly random N -bit. Let $q = 1/2 - \epsilon$ be the probability for sampling each bit to be 1 in a random biased string. By standard concentration bounds, we know that a random biased N -bit string will have, with very high probability, the number of 1s very close to $K = qN$. For the simplicity of exposition, we will assume that our random biased strings have at most $K = qN$ ones in them. We then show that every N -bit string with (at most) K ones has its KT complexity much less than N , which is the lower bound on the KT complexity of a uniformly random N -bit string.

It is natural to think of an N -bit string with K ones as a subset of size K in the universe of size N . As there are exactly $\binom{N}{K}$ such subsets, the minimal bit complexity to represent any one of such subsets is $OPT = \log_2 \binom{N}{K}$. Such an information-theoretically optimal encoding of K -size subsets of the N -size universe is known, and is achieved by using the combinatorial number system where we represent each such subset by the unique number of the form

$$\binom{c_K}{K} + \dots + \binom{c_2}{2} + \binom{c_1}{1}.$$

In more detail, suppose we have $x \in \{0, 1\}^N$ where X has exactly K ones. For the base case, when $K = 0$, we output 0. For $K > 0$, we associate with x an integer number using the following recursive procedure: if $x_N = 1$, then output $\binom{N-1}{K-1} + R_1$, where R_1 is the recursively computed integer associated with $N - 1$ -bit prefix of x and the parameter $K - 1$; if $x_N = 0$, then output R_0 , which is the integer associated with the $N - 1$ -bit prefix of x and the parameter K .

Note that the final integer associated with a given K -size subset $x \in \{0, 1\}^N$ has value at most $\binom{N}{K}$ (using Pascal's identity that $\binom{N}{K} = \binom{N-1}{K-1} + \binom{N-1}{K}$), and so has the optimal bit complexity. The encoding is efficient (as outlined above). The decoding is also efficient: given an integer B encoding some unknown K -size subset $x \in \{0, 1\}^N$, if $B \geq \binom{N-1}{K-1}$, then set $x_N = 1$, and recursively decode $B - \binom{N-1}{K-1}$ for a $K - 1$ -size subset of the $N - 1$ -size universe; otherwise, set $x_N = 0$, and recursively decode B for a K -size subset of the $N - 1$ -size universe. The running time of such a decoding algorithm is clearly $\text{poly}(N)$, and can be shown to be about $O(N^2)$.

For the KT complexity of a string $x \in \{0, 1\}^N$ with K ones, we could define d to be the integer encoding this K -size subset, and then define $U^d(i)$ to be an algorithm that does the decoding of d to get all the bits of x . The problem with this is that the runtime to do a complete decoding of d into x is more than N , which is too much. However, the decoding we do is *global*, recovering all bits x_i simultaneously, whereas we just need to give a *local* decoding algorithm: given i , recover x_i .

Encoding

To get a locally decodable representation for our string $x \in \{0, 1\}^N$, we partition x into blocks of size b (for b to be determined). For each block j , $1 \leq j \leq N/b$, let K_j be the number of ones in that block. Note that $\sum_{j=1}^{N/b} K_j = K$. Given K_j , encode each block j information-theoretically optimally as described above, using at most $\log_2 \binom{b}{K_j}$ bits. Write the resulting encodings of all blocks one after the other; add N/b pointers (of at most $\log N$ bits each) that point to the beginnings of the encodings of the blocks; add N/b numbers K_j 's to the encoding. We get that the total bit size of the overall encoding of x is at most

$$O(N/b)(\log N + \log b) + \sum_{j=1}^{N/b} \lceil \log_2 \binom{b}{K_j} \rceil.$$

The latter sum is at most

$$\sum_{j=1}^{N/b} \log_2 \binom{b}{K_j} + N/b = \log_2 \prod_{j=1}^{N/b} \binom{b}{K_j} + N/b \leq \log_2 \binom{N}{K} + N/b,$$

since the number of sets with K_j ones in block j is at most the number of all subsets of $[N]$ with K ones. Overall, the encoding size is $OPT + O(N \log N/b)$.

Decoding

Given $i \in [N]$, we first figure out which block j it is in, and then decode that entire block (after looking up its number of ones in K_j and its compressed image). As discussed earlier, the decoding runs in time about $O(b^2)$.

Upper-bounding the KT complexity

To keep the KT complexity of x low, we choose b to be $N^{1/3}$. Then the KT complexity of x is at most $OPT + O(N^{2/3})$.

Finally, we show that $MKTP \notin AC^0[p]$ as follows. Recall that we consider random biased strings of length N where the bias probability is $q = 1/2 - \epsilon$. Let $K = qN$ be the expected number of ones in a typical biased string, and let's assume that most biased strings have at most K ones in them (for simplicity). We get that for such a biased string, its KT complexity is at most

$$\log_2 \binom{N}{K} + O(N^{2/3}) \approx H(q) \cdot N + O(N^{2/3}),$$

where H is the binary entropy function. For $q = 1/2 - \epsilon$, we can estimate $H(q) \approx 1 - O(\epsilon^2)$. Thus, to have the KT complexity of a typical q -biased N -bit string to be strictly less than N , we need $N \cdot O(\epsilon^2) - O(N^{2/3}) > 0$, which implies that we need $\epsilon > \Omega(N^{-1/6})$. This implies by Theorem 2.6 that $MKTP$ on inputs of size N requires $AC^0[p]$ circuits of depth d of size at least $\exp(\Omega(N^{1/6(d-1)}))$.

8 Open questions

We managed to find a work-around the lack of a tighter Lupanov-style upper bound on the circuit complexity of just slightly biased random functions where the bias probability is arbitrarily close to $1/2$ (as opposed to the bias probability being a small constant bounded

away from $1/2$). Our proof would be much more direct and constructive if we had such refinements of Lupanov's circuit upper bounds for biased boolean functions (since then we could have proceeded similarly to our proof for the MKTP case in the previous section). Can one prove such tighter circuit upper bounds?

Our current circuit lower bound applies to MCSP, but doesn't seem to apply for its average-case version, the Razborov-Rudich natural property [28]. Can one show such an extension?

Finally, we showed (Corollary 5.7) that either $\text{NEXP} \not\subseteq \text{P/poly}$ or $\text{MCSP} \notin \text{ACC}^0$. For the proof, we used the original Easy Witness Lemma of [15], the existence of random-self-reducible problems in EXP, plus the known lower bound that $\text{NEXP} \not\subseteq \text{ACC}^0$ [32]. Given the new Easy Witness Lemma and the improved circuit lower bound that $\text{NQP} = \text{NTIME}[n^{\text{poly} \log n}] \not\subseteq \text{ACC}^0$ [23], it is natural to ask the following: Can we show that either $\text{NQP} \not\subseteq \text{P/poly}$ or $\text{MCSP} \notin \text{ACC}^0$? Our current proof techniques rely on the existence of a random-self-reducible problem complete for EXP, and no such problem is known for the class NQP.

References

- 1 Eric Allender. Arithmetic Circuits and Counting Complexity Classes. In Jan Krajíček, editor, *Complexity of Computations and Proofs, Quaderni di Matematica*, volume 13, pages 33–72. Seconda Università di Napoli, 2004.
- 2 Eric Allender, Harry Buhrman, Michal Koucký, Dieter van Melkebeek, and Detlef Ronneburger. Power from Random Strings. *SIAM J. Comput.*, 35(6):1467–1493, 2006. doi:10.1137/050628994.
- 3 Eric Allender and Bireswar Das. Zero Knowledge and Circuit Minimization. In *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II*, pages 25–32, 2014. doi:10.1007/978-3-662-44465-8_3.
- 4 Eric Allender, Lisa Hellerstein, Paul McCabe, Toniann Pitassi, and Michael E. Saks. Minimizing Disjunctive Normal Form Formulas and AC^0 Circuits Given a Truth Table. *SIAM J. Comput.*, 38(1):63–84, 2008. doi:10.1137/060664537.
- 5 Eric Allender and Shuichi Hirahara. New Insights on the (Non-)Hardness of Circuit Minimization and Related Problems. In Kim G. Larsen, Hans L. Bodlaender, and Jean-François Raskin, editors, *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21-25, 2017 - Aalborg, Denmark*, volume 83 of *LIPICs*, pages 54:1–54:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.MFCS.2017.54.
- 6 Eric Allender, Dhiraj Holden, and Valentine Kabanets. The Minimum Oracle Circuit Size Problem. In *32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, March 4-7, 2015, Garching, Germany*, pages 21–33, 2015. doi:10.4230/LIPICs.STACS.2015.21.
- 7 Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Learning Algorithms from Natural Proofs. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, pages 10:1–10:24, 2016.
- 8 Stephen A. Cook. A Taxonomy of Problems with Fast Parallel Algorithms. *Information and Control*, 64(1-3):2–21, 1985. doi:10.1016/S0019-9958(85)80041-3.
- 9 Vladimir Dančik. Complexity of Boolean functions over bases of unbounded fan-in gates. *Information Processing Letters*, 57:31–34, 1996.
- 10 Shafi Goldwasser, Dan Gutfreund, Alexander Healy, Tali Kaufman, and Guy N. Rothblum. Verifying and decoding in constant depth. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 440–449. ACM, 2007. doi:10.1145/1250790.1250855.

- 11 Shuichi Hirahara and Rahul Santhanam. On the Average-Case Complexity of MCSP and Its Variants. In *CCC*, pages 7:1–7:20, 2017. doi:10.4230/LIPIcs.CCC.2017.7.
- 12 Shuichi Hirahara and Osamu Watanabe. Limits of Minimum Circuit Size Problem as Oracle. In *31st Conference on Computational Complexity, CCC*, pages 18:1–18:20, 2016. doi:10.4230/LIPIcs.CCC.2016.18.
- 13 John M. Hitchcock and A. Pavan. On the NP-Completeness of the Minimum Circuit Size Problem. In *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS*, pages 236–245, 2015. doi:10.4230/LIPIcs.FSTTCS.2015.236.
- 14 Russell Impagliazzo, Valentine Kabanets, and Ilya Volkovich. The Power of Natural Properties as Oracles. In Rocco A. Servedio, editor, *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, volume 102 of *LIPIcs*, pages 7:1–7:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. doi:10.4230/LIPIcs.CCC.2018.7.
- 15 Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *J. Comput. Syst. Sci.*, 65(4):672–694, 2002. doi:10.1016/S0022-0000(02)00024-7.
- 16 Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-24508-4.
- 17 Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *STOC*, pages 73–79, 2000. ECCC:TR99-0045.
- 18 Nutan Limaye, Karteek Sreenivasaiyah, Srikanth Srinivasan, Utkarsh Tripathi, and S. Venkitesh. The Coin Problem in Constant Depth: Sample Complexity and Parity Gates. *arXiv preprint*, 2018. arXiv:1809.04092.
- 19 Oleg B. Lupanov. On the synthesis of switching circuits. *Doklady Akademii Nauk SSSR*, 119(1):23–26, 1958. English translation in *Soviet Mathematics Doklady*.
- 20 Oleg B. Lupanov. Complexity of formula realization of functions of logical algebra. *Problemy Kibernetiki*, 3:782–811, 1962.
- 21 Oleg B. Lupanov. On a certain approach to the synthesis of control systems — the principle of local coding. *Problemy Kibernetiki*, 14:31–110, 1965. (in Russian).
- 22 Colin McDiarmid. *On the method of bounded differences*, pages 148–188. London Mathematical Society Lecture Note Series. Cambridge University Press, 1989. doi:10.1017/CB09781107359949.008.
- 23 Cody Murray and R. Ryan Williams. Circuit lower bounds for nondeterministic quasi-polytime: an easy witness lemma for NP and NQP. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 890–901. ACM, 2018. doi:10.1145/3188745.3188910.
- 24 Cody D. Murray and R. Ryan Williams. On the (Non) NP-Hardness of Computing Circuit Complexity. *Theory of Computing*, 13(1):1–22, 2017. doi:10.4086/toc.2017.v013a004.
- 25 Igor C. Oliveira and Rahul Santhanam. Conspiracies between learning algorithms, circuit lower bounds, and pseudorandomness. In *Proceedings of the 32nd Computational Complexity Conference*, page 18. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- 26 Nicholas Pippenger. Information theory and the complexity of boolean functions. *Mathematical systems theory*, 10:129–167, 1976.
- 27 Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Math. Notes*, 41(4):333–338, 1987.
- 28 Alexander A. Razborov and Steven Rudich. Natural Proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997. doi:10.1006/jcss.1997.1494.
- 29 Ronen Shaltiel and Emanuele Viola. Hardness Amplification Proofs Require Majority. *SIAM J. Comput.*, 39(7):3122–3154, 2010. doi:10.1137/080735096.
- 30 Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell System Tech. J.*, 28:59–98, 1949.

- 31 Roman Smolensky. Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*, pages 77–82, 1987.
- 32 Ryan Williams. Nonuniform ACC Circuit Lower Bounds. *J. ACM*, 61(1):2:1–2:32, 2014. doi:10.1145/2559903.