

# Approximate Degree and the Complexity of Depth Three Circuits

Mark Bun<sup>1</sup>

Princeton University, Princeton, NJ, USA  
mbun@cs.princeton.edu

Justin Thaler<sup>2</sup>

Georgetown University, Washington, DC, USA  
justin.thaler@georgetown.edu

---

## Abstract

Threshold weight, margin complexity, and Majority-of-Threshold circuit size are basic complexity measures of Boolean functions that arise in learning theory, communication complexity, and circuit complexity. Each of these measures might exhibit a *chasm* at depth three: namely, all polynomial size Boolean circuits of depth two have polynomial complexity under the measure, but there may exist Boolean circuits of depth three that have essentially maximal complexity  $\exp(\Theta(n))$ . However, existing techniques are far from showing this: for all three measures, the best lower bound for depth three circuits is  $\exp(\tilde{\Omega}(n^{2/5}))$ . Moreover, prior methods exclusively study *block-composed* functions. Such methods appear intrinsically unable to prove lower bounds better than  $\exp(\Omega(\sqrt{n}))$  even for depth four circuits, and have yet to prove lower bounds better than  $\exp(\tilde{\Omega}(\sqrt{n}))$  for circuits of any constant depth.

We take a step toward showing that all of these complexity measures indeed exhibit a chasm at depth three. Specifically, for any arbitrarily small constant  $\delta > 0$ , we exhibit a depth three circuit of polynomial size (in fact, an  $O(\log n)$ -decision list) of complexity  $\exp(\Omega(n^{1/2-\delta}))$  under each of these measures.

Our methods go beyond the block-composed functions studied in prior work, and hence may not be subject to the same barriers. Accordingly, we suggest natural candidate functions that may exhibit stronger bounds.

**2012 ACM Subject Classification** Theory of computation

**Keywords and phrases** approximate degree, communication complexity, learning theory, polynomial approximation, threshold circuits

**Digital Object Identifier** 10.4230/LIPIcs.APPROX-RANDOM.2018.35

**Related Version** The full version of this work is available at <https://eccc.weizmann.ac.il/report/2016/121/>.

## 1 Introduction

Let  $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$  be a Boolean function, and let  $\mathcal{C}: 2^{\{-1, 1\}^n} \rightarrow \mathbb{N}$  denote a measure of the complexity of  $f$ . We say that  $\mathcal{C}$  exhibits a *chasm* at depth three if all Boolean circuits<sup>3</sup> of depth two have polynomial complexity under the measure, but there exist circuits of depth

---

<sup>1</sup> This work was done while the author was at Harvard University and visiting Yale University, supported by an NDSEG Fellowship and NSF grant CNS-1237235.

<sup>2</sup> Parts of this work were performed while the author was at Yahoo Research.

<sup>3</sup> Throughout this paper, unless otherwise noted, all circuits under consideration are assumed to have polynomial size, and to be over the basis AND, OR and NOT.



© Mark Bun and Justin Thaler;

licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018).

Editors: Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer; Article No. 35; pp. 35:1–35:18



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

three that have essentially maximal complexity  $\exp(\Theta(n))$ . Examples of measures that may satisfy a chasm at depth three include:

- **Threshold Weight.** A polynomial  $p: \{-1, 1\}^n \rightarrow \mathbb{R}$  with integer coefficients is said to sign-represent  $f$  if  $p(x) \cdot f(x) > 0$  for all  $x \in \{-1, 1\}^n$ . The weight of  $p$ , denoted  $W(p)$ , is the sum of the absolute value of its coefficients. The *threshold weight* of  $f$  is the least weight of a sign-representing polynomial for  $f$ .

It is easy to see that all DNF and CNF formulae of size  $s$  have threshold weight  $O(ns)$ . The best known upper bound on the threshold weight of depth three circuits is the trivial  $2^{O(n)}$  bound. Hence, threshold weight may exhibit a chasm at depth three.

- **Discrepancy and Margin Complexity.** Discrepancy, defined formally in Appendix A.4, is a central quantity in communication complexity and circuit complexity.<sup>4</sup> For example, discrepancy is known to characterize the communication complexity class **PP** [13], and small discrepancy implies large communication complexity in nearly every communication model. The multiplicative inverse of discrepancy is also known to be equivalent to *margin complexity*, a central quantity in learning theory [20].

All DNF and CNF formulae have at least inverse-polynomial discrepancy. However, the best known lower bound on the discrepancy of depth three circuits is the trivial  $2^{-O(n)}$  bound. Hence, margin complexity and (the inverse of) discrepancy may exhibit a chasm at depth three.

- **Majority-of-Threshold Circuit Size.** Since OR and AND can each be computed by a single Majority gate, all DNF and CNF formulae are computed by Majority-of-Threshold (in fact, Majority-of-Majority) circuits of polynomial size. Meanwhile, the best known upper bound on the size of Majority-of-Threshold circuits computing depth three Boolean circuits is the trivial  $2^{O(n)}$  bound. Hence, Majority-of-Threshold circuit size may exhibit a chasm at depth three.

We discuss each of these measures, together with applications in communication complexity and learning theory, in more detail in Appendix A.

Unfortunately, we are currently quite far from proving that any of the above complexity measures actually exhibit such a chasm. For each measure, the best known lower bound for depth three circuits is  $\exp(\tilde{\Omega}(n^{2/5}))$  [8]. Moreover, as we explain in Section 1.2.4, existing techniques appear intrinsically unable to prove a lower bound better than  $\exp(\Omega(n^{1/2}))$  even for circuits of depth four. This barrier stems from the fact that previous work has focused exclusively on analyzing *block-composed* functions. Here, a function  $f: \{-1, 1\}^{N \cdot M} \rightarrow \{-1, 1\}$  is said to be block-composed if there are two functions  $h: \{-1, 1\}^N \rightarrow \{-1, 1\}$  and  $g: \{-1, 1\}^M \rightarrow \{-1, 1\}$  such that  $f = h \circ g := h(g, \dots, g)$ . That is, a function is block-composed if it interprets its input as a sequence of  $N$  blocks  $x_1, \dots, x_N \in \{-1, 1\}^M$ , applies a Boolean function  $g$  independently to each block  $x_i$ , and then feeds the  $N$  outputs into a different function  $h$ .

In this paper, we take a step toward showing that all three of these complexity measures indeed exhibit a chasm at depth three. Specifically, for any constant  $\delta > 0$ , we exhibit a depth three circuit of polynomial size (in fact, an  $O(\log n)$ -decision list) of complexity  $\exp(\Omega(n^{1/2-\delta}))$  under each of these measures. Our improvement over prior work stems from the fact that we move beyond block-composed functions, and hence our methods may not be subject to the same barriers. In particular, we suggest natural candidate functions that

<sup>4</sup> Discrepancy is often thought of as a matrix-analytic quantity, rather than as a Boolean function complexity measure. For a function  $f: \{-1, 1\}^n \times \{-1, 1\}^n \rightarrow \{-1, 1\}$ , when we refer to the discrepancy of  $f$ , we mean the discrepancy of the matrix  $[f(x, y)]_{x, y \in \{-1, 1\}^n}$ .

may exhibit stronger bounds, of the form  $\exp(\tilde{\Omega}(n))$ , where the  $\tilde{\Omega}$  notation hides factors polylogarithmic in  $n$  (cf. Section 3.3).

The functions that underly our analysis are rather complicated to define, but here we briefly highlight their novel features. Inspired by prior work of Podolskii [24] (see Section 1.2.2 for further discussion), we define our functions to be “almost” block-composed, but to have mild dependencies between blocks. Just like a block-composed function, each of our functions interprets its input as a sequence of  $N$  blocks  $x_1, \dots, x_N \in \{-1, 1\}^M$ , and applies a Boolean function  $g$  to each block, before feeding the  $N$  outputs into a different function  $h$ . However, for  $i \geq 2$ , before applying  $g$  to  $x_i$  we first pass  $x_i$  through a “pre-processing function” that depends on the preceding blocks  $x_1, x_2, \dots, x_{i-1}$ . This dependency is simple enough that the final function is computed by a circuit of depth three, but complicated enough that the function has much higher complexity than the block-composed functions considered in prior work.

## 1.1 Our Contributions: Details

The three complexity measures  $\mathcal{C}$  described above are intimately related to uniform approximability by low-degree polynomials, as we now explain. Roughly speaking, for each of the three measures, in order to construct a function of complexity at least  $2^d$  under  $\mathcal{C}$ , it suffices to identify a function  $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$  such that  $f$  cannot be uniformly approximated to error  $1 - 2^{-d}$  by polynomials of degree at most  $d$ . One can then apply known transformations [17, 30, 16] to transform  $f$  into a related function  $F: \{-1, 1\}^n \rightarrow \{-1, 1\}$  such that  $\mathcal{C}(F) \geq 2^{\Omega(d)}$ . Moreover, these transformations are simple in the following sense: if  $f$  is computed by a (polynomial size) depth  $d$  circuit with logarithmic bottom fan-in, so is  $F$ .

Accordingly, the main technical contribution of this paper is to prove a new lower bound on the approximability of suitable constant-depth circuits by low-degree polynomials.

► **Theorem 1.** *For any arbitrarily small constant  $\delta > 0$  and any arbitrarily large constant  $\Gamma > 1$ , there is an (explicitly given) function  $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$  that is computed by Boolean circuit of depth three, with logarithmic bottom fan-in, that satisfies the following property. For any polynomial  $p: \{-1, 1\}^n \rightarrow \mathbb{R}$  of total degree at most  $n^{1/2-\delta}$ , there exists some  $x \in \{-1, 1\}^n$  such that  $|p(x) - f(x)| > 1 - 2^{-n^\Gamma}$ .*

In fact, the function  $f$  in Theorem 1 is much simpler than an arbitrary depth three circuit with logarithmic bottom fan-in; it is an  $O(\log n)$ -decision list of polynomial length. An  $O(\log n)$ -decision list is a function whose output is determined by a very simple sequential decision process (essentially, a chain of “if-then-else” statements, where each “if” statement is a conjunction on  $O(\log n)$  variables—we give a precise definition in Section 2). Decision lists have been studied intensely in learning theory and complexity theory (see, e.g., [26, 15, 3, 35, 27, 16, 5, 11]).

By combining Theorem 1 with known transformations [17, 30], we obtain a depth three circuit  $F$  with very large complexity under the three measures described above. In fact,  $F$  itself is computed by an  $O(\log n)$ -decision list.

► **Corollary 2.** *For any constant  $\delta > 0$ , there is an (explicitly given) function  $F: \{-1, 1\}^n \rightarrow \{-1, 1\}$  computed by an  $O(\log n)$ -decision list of polynomial length (and hence also a depth-3 circuit with logarithmic bottom fan-in) such that:*

- *The threshold weight of  $F$  is  $\exp(\Omega(n^{1/2-\delta}))$ .*
- *The discrepancy of  $F$  is  $\exp(-\Omega(n^{1/2-\delta}))$ .*
- *Any Majority-of-Threshold circuit computing  $F$  has size  $\exp(\Omega(n^{1/2-\delta}))$ .*

Table 1 succinctly compares our results to prior work.

■ **Table 1** Comparison of our new bounds for  $AC^0$  to prior work. The circuit depth column lists the depth of the Boolean circuit used to exhibit the bound, and  $\delta$  denotes an arbitrarily small positive constant. All Boolean circuits are polynomial size.

Reference	Threshold Weight Bound	Discrepancy Bound	Majority-of-Threshold Circuit Size Bound	Circuit Depth
[17]	$\exp(\Omega(n^{1/3}))$	N/A	N/A	3
[29]	N/A	$\exp(-\Omega(n^{1/5}))$	$\exp(\Omega(n^{1/5}))$	3
[5, 30]	N/A	$\exp(-\Omega(n^{1/3}))$	$\exp(\Omega(n^{1/3}))$	3
[8]	$\exp(\Omega(n^{2/5}))$	$\exp(-\Omega(n^{2/5}))$	$\exp(\Omega(n^{2/5}))$	3
[33]	$\exp\left(\Omega\left(n^{\frac{k-1}{2k-1}}\right)\right)$	$\exp\left(-\Omega\left(n^{\frac{k-1}{2k-1}}\right)\right)$	$\exp\left(\Omega\left(n^{\frac{k-1}{2k-1}}\right)\right)$	$k + 1$ (for $k \geq 2$ )
[28]	$\exp(\Omega(n^{1/2}))$	$\exp(-\Omega(n^{1/2}))$	$\exp(\Omega(n^{1/2}))$	4
This work	$\exp(\Omega(n^{1/2-\delta}))$	$\exp(-\Omega(n^{1/2-\delta}))$	$\exp(\Omega(n^{1/2-\delta}))$	3

## 1.2 Prior Work

In order to discuss prior work, it is helpful to introduce the notions of approximate degree and threshold degree, which both capture the difficulty of pointwise approximation by low-degree polynomials. The  $\varepsilon$ -approximate degree of a function  $f$ , denoted  $\widetilde{\deg}_\varepsilon(f)$ , is the least degree of a real polynomial that pointwise approximates  $f$  to error  $\varepsilon$ . By convention,  $\widetilde{\deg}_{1/3}(f)$  is denoted simply as  $\widetilde{\deg}(f)$  and referred to without qualification as the approximate degree of  $f$  (the constant  $1/3$  is chosen for aesthetic reasons, and could be replaced with any other constant in  $(0, 1)$  without affecting the theory in any way). The threshold degree of  $f$ , denoted  $\deg_\pm(f)$ , is the least degree of a real polynomial that sign-represents  $f$  at all points (such a polynomial is called a polynomial threshold function (PTF) for  $f$ ). When appropriate, we also use subscripts to indicate the number of variables over which the function is defined. For example,  $OR_M$  denotes the OR function on  $M$  inputs.

### 1.2.1 Early Work on Approximating $AC^0$ Functions by Polynomials

Minsky and Papert [22] famously proved an  $\Omega(n^{1/3})$  lower bound on the threshold degree of the DNF formula  $OR_{n^{1/3}} \circ AND_{n^{2/3}}$ , now known as the Minsky-Papert DNF. Klivans and Servedio [14] proved an essentially matching upper bound of  $\tilde{O}(n^{1/3})$  on the threshold degree of *any* polynomial size DNF.

Beigel identified a DNF (in fact, a 1-decision list) known as OMB (short for ODD-MAX-BIT) that has threshold degree 1, but requires large degree to approximate to error bounded away from 1 [3]. OMB will play a central role in this paper, and we define it formally in Section 2. Quantitatively, Beigel showed<sup>5</sup> that for any  $d > 0$ , there is an  $\varepsilon = 1 - 2^{-\Omega(n/d^2)}$  such that  $\widetilde{\deg}_\varepsilon(OMB_n) \geq d$ . For any  $\varepsilon > 0$ , Klivans and Servedio [15] gave an optimal  $\varepsilon$ -approximating polynomial for any 1-decision list, showing that Beigel's lower bound is asymptotically tight for all  $d > 0$ .<sup>6</sup>

<sup>5</sup> Beigel describes his result as a lower bound on the *degree- $d$  threshold weight* of  $OMB_n$ , which refers to the least weight of a sign-representing polynomial  $p$  for  $f$  satisfying  $\deg(p) \leq d$ . However, his argument is easily seen to establish the claimed approximate degree lower bound.

<sup>6</sup> Like Beigel, Klivans and Servedio state their results in terms of degree- $d$  threshold weight. However, their construction is easily seen to imply the claimed upper bound on the approximate degree of  $OMB_n$ .

### 1.2.2 Prior Work of Podolskii

Podolskii pioneered a line of work devoted to proving approximate degree lower bounds that hold even when the error parameter  $\varepsilon$  is allowed to be *super-exponentially* close to 1 [24, 25].<sup>7</sup> In [25], he showed that for any constant  $d \geq 2$ , there exists a function of threshold degree  $d$  that cannot be uniformly approximated to error  $\varepsilon$  by polynomials of degree at most  $d$ , unless  $\varepsilon = 1 - n^{-\Omega(n^d)}$ . This result is tight, matching an upper bound proved by Buhrman et al. [5].

Our construction and analysis are inspired by another related result of Podolskii [24]. For any constant  $d > 0$ , Podolskii identified a function  $f$  of threshold degree  $d$  such that, even for  $D \gg d$ , the following holds:  $f$  cannot be uniformly approximated by degree  $D$  polynomials to error  $\varepsilon$ , unless  $\varepsilon$  is superexponentially close to 1. Quantitatively, he showed that for any constant  $d > 0$ , there exists a DNF (in fact, a  $d$ -decision list)  $f$  with threshold degree  $d$ , yet for any  $D < O(n^{1/5}/\log n)$ , there exists an  $\varepsilon \in 1 - \exp(-\Omega((n/D^4)^d))$  for which  $\deg_\varepsilon(f) \geq D$ . Unfortunately, Podolskii's construction does not yield any new bounds on the complexity measures we are interested in. By introducing new ideas, we are able to prove such improved bounds for depth three circuits.

### 1.2.3 Translating Approximate Degree Bounds to Complexity Bounds

Several works have focused on transforming approximate degree and threshold degree lower bounds into bounds on the complexity measures that we focus on in this paper (threshold weight, discrepancy/margin complexity, and Majority-of-Threshold circuit size). Krause and Pudlák [17] showed how to take a function  $f$  of threshold degree at least  $d$ , and turn  $f$  into a related function  $F$  of threshold weight<sup>8</sup> at least  $2^d$ . By applying this transformation to the Minsky-Papert DNF, Krause and Pudlák obtained a depth three circuit (with constant bottom fan-in) with threshold weight  $\exp(\Omega(n^{1/3}))$ .

Subsequent work by Krause [16] showed that for  $F$  to have threshold weight  $2^{\Omega(d)}$ , it is enough for  $f$  to satisfy  $\deg_{1-2^{-d}}(f) \geq d$ .<sup>9</sup> Krause applied his result to the function  $f = \text{OMB}$ , to obtain an  $\exp(\Omega(n^{1/3}))$  lower bound on the threshold weight of a specific 2-decision list.

Sherstov's pattern-matrix method [30] showed how to take a function  $f$  satisfying the same condition required by Krause (i.e.,  $\deg_{1-2^{-d}}(f) \geq d$ ), and turn it into a function  $F$  with discrepancy  $2^{-\Omega(d)}$ . By applying this transformation to the Minsky-Papert DNF or to OMB, Sherstov obtained a depth three circuit with discrepancy  $\exp(-\Omega(n^{1/3}))$ . Buhrman, Vereshchagin, and de Wolf independently proved an identical discrepancy bound via very different techniques [5]. These discrepancy bounds also implied corresponding lower bounds on Majority-of-Threshold Circuit Size, through standard transformations [23].

### 1.2.4 Recent Work on Approximating $\text{AC}^0$ Functions by Polynomials

A handful of recent works have established various forms of “hardness amplification” for approximate degree [28, 7, 8, 33, 28, 32, 19, 31]. Roughly speaking, these results show how to take a function  $g$  which is hard to approximate by degree  $d$  polynomials to error  $1/3$ , and

<sup>7</sup> Again, Podolskii describes his work in terms of degree- $d$  threshold weight, but his results hold for approximate degree as well.

<sup>8</sup> In fact, Krause and Pudlák showed that  $F$  has threshold length  $2^d$ , where threshold length is the least number of non-zero Fourier coefficients of any sign-representation for  $f$ . The threshold weight of  $f$  is always at least as large as its threshold length.

<sup>9</sup> Again, Krause phrased his lower bound in terms of the degree- $d$  threshold weight of OMB, but his result is easily seen to imply the statement here.

turn  $g$  into a related function  $f$  that is hard to approximate by degree  $d$  polynomials to error exponentially close to 1. Specifically, in these works,  $f$  is obtained from  $g$  by block-composing  $g$  with another function  $h$ .

Let  $\text{ED}_M$  denote the well-known Element Distinctness function and  $\overline{\text{ED}}_M$  its negation.  $\text{ED}_M$  played a central role in recent works on hardness amplification for approximate degree [8, 33, 28] because, at the time these works were written, it exhibited the largest known approximate degree lower bound for any function in  $\text{AC}^0$ :  $\widetilde{\text{deg}}(\text{ED}_M) = \widetilde{\Omega}(M^{2/3})$  [1].<sup>10</sup>

Our prior work [8] showed that the function  $f = \text{OR}_N \circ \text{ED}_M$  satisfies  $\widetilde{\text{deg}}_\varepsilon(f) \geq \widetilde{\Omega}(M^{2/3})$  for  $\varepsilon = 1 - 2^{-N}$ , and used this result to obtain a depth three circuit such that  $\mathcal{C}(F) = \exp(\widetilde{\Omega}(n^{2/5}))$  for the three complexity measures  $\mathcal{C}$  that we focus on in this work. Thaler [34] then showed that the function  $f = \text{OMB}_N \circ \overline{\text{ED}}_M$  satisfies an identical lower bound, yielding another depth three circuit  $F$  (in fact, an  $O(\log n)$ -decision list) with  $\mathcal{C}(F) = \exp(\widetilde{\Omega}(n^{2/5}))$ .

Sherstov [33] significantly strengthened the approach of [8] to obtain new *threshold degree* lower bounds for functions in  $\text{AC}^0$ . Specifically, in [33], for any  $k \geq 2$ , Sherstov exhibited a read-once formula of depth  $k$  (with polynomial bottom fan-in) that has threshold degree  $\Omega(n^{\frac{k-1}{2k-1}})$ . Applying the transformations of [17, 30, 16] to these circuits increases their depth by 1. In [28], Sherstov exhibited a depth four circuit of logarithmic bottom fan-in and threshold degree  $\Omega(n^{1/2})$ —applying the transformation of [17, 30, 16] to this circuit does not increase its depth, yielding a depth four circuit  $F$  satisfying  $\mathcal{C}(F) = \exp(\Omega(n^{1/2}))$ .

**The  $\exp(\Theta(\sqrt{n}))$  Barrier For Circuits of Depth Four.** Recall that for each of the three complexity measures  $\mathcal{C}$  in which we are interested, to construct a function of complexity at least  $2^d$  under  $\mathcal{C}$ , it suffices to identify a function  $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$  such that

$$f \text{ cannot be approximated to error } 1 - 2^{-d} \text{ by polynomials of degree at most } d. \quad (1)$$

We now argue that for circuits of depth 4, prior techniques cannot accomplish this for  $d \gg \sqrt{n}$ , *even if they assume the existence of a DNF  $g: \{-1, 1\}^M \rightarrow \{-1, 1\}$  with (one-sided) approximate degree  $\Omega(M)$* .<sup>11</sup>

The methods of [8, 34] start with a function  $g: \{-1, 1\}^M \rightarrow \{-1, 1\}$ , and assume nothing about  $g$  other than that  $g$  has (one-sided) approximate degree at least  $d$ . They show how to turn  $g$  into a “harder” function  $f = h \circ g$  by block-composing  $g$  with another function  $h \in \{\text{OR}_N, \text{OMB}_N\}$ . Quantitatively, the resulting bound is of the form  $\widetilde{\text{deg}}_{1-2^{-N}}(f) \geq d$ . Clearly, one must set  $N \geq d$  to obtain a bound of the form Eq. (1). Hence, even if  $g$  has the largest possible (one-sided) approximate degree,  $d = M$ , the best bound that can be obtained from the methods of [8, 34] is of the form  $\widetilde{\text{deg}}_{1-2^{-N}}(f) \geq N$ , obtained by setting  $M = N$ . In this case,  $f$  is a function over  $n = N^2$  variables, so these methods can only yield complexity bounds of the form  $\exp(\Omega(N)) = \exp(\Omega(\sqrt{n}))$ .

Both [8, 34] showed that their respective analyses are tight for many functions  $g$ . Hence, the  $\exp(\sqrt{n})$  barrier is not merely an artifact of the analysis in these works.

<sup>10</sup>In very recent work [9], which was performed after the work described in the present manuscript, we have proved a nearly optimal lower bound on the approximate degree of  $\text{AC}^0$ : for any constant depth  $d$ , it exhibited a depth- $d$   $\text{AC}^0$  circuit of approximate degree  $\Omega(n^{1-2^{-\Omega(d)}})$ . See Section 1.3 for details.

<sup>11</sup>One-sided approximate degree is a measure that is intermediate between approximate degree and threshold degree. One-sided approximate degree lower bounds is crucial to the analyses in [8, 33, 28]. However, we will not explicitly utilize one-sided approximate degree in our own results, so we do not formally define it here. While our work subsequent to this manuscript [9] has proved a nearly-linear lower bound on the approximate degree of  $\text{AC}^0$  functions, the best known one-sided approximate degree lower bound for an  $\text{AC}^0$  function currently remains  $\widetilde{\Omega}(M^{2/3})$ , exhibited by  $\text{ED}_M$  [8].

Indeed, at least in the case of  $h = \text{OR}_N$ , the barrier is inherent to any method that attempts to construct an  $f$  satisfying Eq. (1) by assuming nothing about a function  $g: \{-1, 1\}^M \rightarrow \{-1, 1\}$  other than that  $g$  has (one-sided) approximate degree at least  $d$ , and then block-composing  $g$  with  $h$ . To see this, first observe that if  $M \leq N$ , then for *any* function  $g: \{-1, 1\}^M \rightarrow \{-1, 1\}$ ,  $(1/N) \sum_{i=1}^N g(x_i) + (N-1)/N$  is a polynomial of degree at most  $M \leq n^{1/2}$  that approximates  $f = \text{OR}_N \circ g$  to error  $1 - 1/N$ .

Even if  $M \geq N$ , it is often the case that  $\text{OR}_N \circ g$  can be approximated to error  $1 - 2^{-\tilde{O}(n^{1/2})}$  by a polynomial of degree  $O(n^{1/2})$ . Indeed, many functions  $g$  with large approximate degree (such as  $\text{ED}_M$  for example) can be approximated to error  $1/3N^2$  by a *ratio*  $q_1(x)/q_2(x)$  of two polynomials of logarithmic degree and weight quasi-polynomial in  $M$  and  $N$ . One can use  $q_1, q_2$  to obtain a polynomial approximator  $p$  for  $f = h \circ g$  such that  $\deg(p) = O(N \log(M \cdot N))$ , and  $p$  uniformly approximates  $f$  to error  $1 - 2^{-O(N \cdot \text{polylog}(M))}$ . We omit the details for brevity, but the construction can be found in [4] (see also [28, Theorem 6.10]).

Sherstov [33, 28] introduced sophisticated and demanding methods that can prove stronger lower bounds for constant-depth circuits than [8, 34]. However, his methods apply block-composition multiple times, and crucially exploit alternation in the circuits computing the functions being composed; hence, in the context of Eq. (1), his analysis improves over [8, 34] only for circuits of greater depth or bottom fan-in than considered in those works. Furthermore, in [28, Section 9.4] Sherstov provides a detailed discussion on barriers facing his methods. In particular, he shows that the  $\exp(\Omega(n^{1/2}))$  barrier is inherent to the class of functions he considers in [28], and is not an artifact of the analysis. He does indicate that his methods might be extendable to break the  $\exp(\Omega(n^{1/2}))$  barrier by using circuits of depth 5 or greater.

### 1.3 Subsequent Work

Our recent works [9, 6], which were performed subsequent to the work described in this manuscript, introduced a very different method for performing hardness amplification in a manner that moves beyond block-composed functions. These works amplify the *degree* of the polynomials to which the lower bounds apply, thereby obtaining optimal or nearly optimal lower bounds for many functions in  $\text{AC}^0$ . In contrast, the techniques we introduce in the present manuscript amplify the *error* parameter for which the lower bounds apply. Combining the two methods to strengthen our lower bounds for depth three circuits, or circuits of larger (constant) depth, is an important direction for future work (see Section 3.3).

## 2 Preliminaries

**Notation.** We work with Boolean functions  $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ , where  $-1$  corresponds to logical TRUE and  $+1$  corresponds to logical FALSE. For a given Boolean function  $f$ , the function  $\bar{f} := -f$  denotes its negation. The notation  $[n]$  refers to the set  $\{0, 1, \dots, n\}$ . For any  $n \in \mathbb{N}$ , fix a canonical injection  $[n] \rightarrow \{-1, 1\}^{\lceil \log(n+1) \rceil}$ ; we refer to the image of any  $i \in [n]$  under this injection as the *binary representation* of  $i$ . All logarithms in this work are assumed to be taken in base 2.

**Decision Lists and OMB.** A  $k$ -decision list  $D$  of length  $L$  over the Boolean variables  $x_1, \dots, x_n$  is represented by a list of  $L$  pairs  $(C_0, b_0), (C_1, b_1), \dots, (C_{L-1}, b_{L-1})$  and a bit  $b_L$  where each  $C_i$  is a conjunction of at most  $k$  (possibly negated) variables, and each  $b_i$  is either  $-1$  or  $1$ . Given any  $x \in \{-1, 1\}^n$ , the value of  $D(x)$  is  $b_i$  if  $i$  is the smallest index such that  $C_i$  is made true by  $x$ ; if no  $C_i$  is true then  $D(x) = b_L$ .

Any  $k$ -decision list of length  $L$  is computed by a depth three circuit of size  $O(L)$  and bottom fan-in  $O(k)$ . Indeed, letting  $S = \{i \leq L : b_i = -1\}$ , the circuit is

$$(b_L \wedge \bar{C}_0(x) \wedge \cdots \wedge \bar{C}_{L-1}(x)) \vee \bigvee_{i \in S} (C_i(x) \wedge \bar{C}_0(x) \wedge \cdots \wedge \bar{C}_{i-1}(x)).$$

To see that this is indeed a circuit of depth three with bottom fan-in  $O(k)$ , observe that for any conjunction  $C_i$  of width  $k$ ,  $\bar{C}_i$  is computed by a disjunction of width  $k$ .

Let  $\text{OMB} : \{-1, 1\}^N \rightarrow \{-1, 1\}$  denote a specific 1-decision list known as ODD-MAX-BIT, defined as follows. For  $i = 1, 2, \dots, N$ , the conjunction  $C_i(x) = x_{N-i}$  and  $b_i = (-1)^{N-i}$ . Finally, define  $b_N = 1$ . OMB can be equivalently defined in the following manner. On input  $x = (x_1, \dots, x_N)$ , let  $\beta(x)$  denote the largest index  $i$  such that  $x_i = -1$ , and let  $\beta(x) = 0$  if no such index exists. Then

$$\text{OMB}(x_1, \dots, x_N) = \begin{cases} -1 & \text{if } \beta(x) \text{ is odd} \\ 1 & \text{otherwise.} \end{cases}$$

Beigel [3] showed that OMB has high approximate degree, even when the error parameter is exponentially close to 1. Specifically:

► **Theorem 3** (Beigel [3]). *There exists a constant  $c > 0$  for which the following holds. Let  $p(x)$  be a polynomial of degree at most  $d$  such that  $p(x) \cdot \text{OMB}_N(x) > 0$  for all  $x \in \{-1, 1\}^N$ . Then there exists an  $x \in \{-1, 1\}^N$  such that  $|p(x)| \geq 2^{cN/d^2} \cdot |p(1^N)|$ . In particular,  $\deg_\varepsilon(\text{OMB}_N) \geq d$  for some  $\varepsilon = 1 - 2^{-\Omega(N/d^2)}$ .*

To prove Theorem 3, Beigel iteratively constructs a sequence of inputs  $x^0, x^1, \dots, x^{cN/d^2}$  for which  $|p(x^{t+1})| \geq 2 \cdot |p(x^t)|$ . He obtains these inputs by repeatedly applying the following lemma, which we will also make use of directly.

► **Lemma 4** (Beigel [3]). *Let  $d, N \in \mathbb{N}$  and let  $\ell \geq 10d^2$  such that  $N/\ell$  is an integer. Consider the increasing family of sets  $S_1 \subset S_2 \subset \cdots \subset S_{N/\ell} \subseteq \{-1, 1\}^n$  defined by*

$$S_0 = \{1^N\}, S_1 = \{x : x_i = 1 \ \forall i > \ell\}, \dots, S_t = \{x : x_i = 1 \ \forall i > t\ell\}, \dots, S_{N/\ell} = \{-1, 1\}^N.$$

*Let  $p(x)$  be a polynomial of degree at most  $d$  such that  $p(x) \cdot \text{OMB}_N(x) > 0$  for all  $x \in S_{t+1} \setminus S_t$ . Let  $z \in S_t$ . Then there exists a  $z' \in S_{t+1} \setminus S_t$  such that  $|p(z')| \geq 2 \cdot |p(z)|$ .*

### 3 Intuition and Discussion of Theorem 1

#### 3.1 Overview of Our Function

As mentioned in Section 1, the function  $f$  that we exhibit in Theorem 1 is complicated to define. Hence, before formally defining  $f$ , we provide here some motivation for our definition. While the function we describe in this section differs from the  $f$  exhibited in Theorem 1, our description here highlights the main ideas underlying the construction of  $f$  itself.

Specifically, the function that we describe in this informal overview is a modification of the block-composed function  $\text{OMB}_N \circ \text{OR}_M$ . This is easily seen to be a sub-function of  $\text{OMB}_{2NM}$ , and in the formal statement and proof of Theorem 1, our construction relies on the function  $\text{OMB}_{2NM}$  directly. This is only for the sake of simplicity, and we remark that the proof of Theorem 1 carries over when one uses the function  $\text{OMB}_N \circ \text{OR}_M$  instead. We choose to present this overview using the function  $\text{OMB}_N \circ \text{OR}_M$  for two reasons. First, reasoning about this function gives the right intuition for both our lower bound and for the



approximating polynomial which it matches. Second, as discussed in Section 3.3, replacing the inner function  $\text{OR}_M$  with other functions of larger approximate degree yields natural candidates for improved lower bounds.

Recall that [34] showed that  $\widetilde{\text{deg}}_\varepsilon(\text{OMB}_N \circ \overline{\text{ED}}_M) = \tilde{\Omega}(M^{2/3})$  for  $\varepsilon = 1 - 2^{-N}$ . Moreover, this lower bound is essentially tight for  $\text{OMB}_N \circ \overline{\text{ED}}_M$ : there is in fact a polynomial  $p$  of degree  $O(\log M)$  that approximates  $\text{OMB}_N \circ \overline{\text{ED}}_M$  to error  $\varepsilon = 1 - 2^{-O(N \log M)}$ . The methods of [34] also show that  $\widetilde{\text{deg}}_\varepsilon(\text{OMB}_N \circ \text{OR}_M) = \Omega(M^{1/2})$  for  $\varepsilon = 1 - 2^{-N}$ , and there is a polynomial of degree  $O(1)$  that approximates  $\text{OMB}_N \circ \text{OR}_M$  to error  $\varepsilon = 1 - 2^{-O(N \log M)}$ .

Our goal is to modify  $\text{OMB}_N \circ \text{OR}_M$  to obtain an  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  that is much harder to approximate by low-degree polynomials, while still ensuring that  $f$  is computed by an  $O(\log n)$  decision list. Specifically, we will require, for some large constant  $k$  and small constant  $\delta > 0$ ,  $\widetilde{\text{deg}}_\varepsilon(f) = \Omega(M^{1/2-\delta})$  for  $\varepsilon = 1 - 2^{-N^k}$ .

A natural first attempt to construct such an  $f$  is to block-compose  $\text{OMB}_N \circ \text{OR}_M$  with the parity function on  $k$  variables. Specifically, let  $k$  be some constant, and consider the following function on  $k \cdot N \cdot M$  variables:  $\oplus_k \circ \text{OMB}_N \circ \text{OR}_M$ , where  $\oplus$  denotes the parity function. However, this function is still too easy to approximate: there is a polynomial of degree  $O(k \log M)$  that approximates  $\oplus_k \circ \text{OMB}_N \circ \text{OR}_M$  to error  $1 - 2^{-O(kN \log M)}$ . Indeed, letting  $p$  be the polynomial approximation to  $\text{OMB}_N \circ \text{OR}_M$  described above, the polynomial  $q(x_1, \dots, x_k) := 2^{-k} \prod_{i=1}^k p(x_i)$  does the trick.

We instead define  $f$  to be “just different enough” from  $\oplus_k \circ \text{OMB}_N \circ \text{OR}_M$  to foil this construction of an approximating polynomial. Specifically, our  $f$  will first “pre-process” its input  $(x_1, \dots, x_k)$ , before feeding it into  $\oplus_k \circ \text{OMB}_N \circ \text{OR}_M$ . The pre-processing step will introduce dependencies between blocks, so that an approximating polynomial for  $f$  will be unable to treat them independently in the manner of  $q$ .

In more detail,  $f$  will interpret its input  $x$  as  $k$  blocks,  $x_1, \dots, x_k$  (we will refer to  $x_1, \dots, x_k$  as “super-blocks”, since each  $x_i$  will itself be interpreted as consisting of  $N$  blocks, which will themselves each be interpreted as consisting of  $M$  “sub-blocks”). For expository purposes, we focus in the remainder of this section on the case  $k = 2$ , so that there are only two super-blocks  $x_1, x_2$  (The full construction is defined inductively, and described in the full version of this work). Assume for simplicity that  $N + 1$  is a power of 2. The two super-blocks will not contain the same number of bits:  $x_1$  will contain  $N \cdot M$  bits, while  $x_2$  will contain  $N \cdot M \cdot \log(N + 1)$  bits. We will ultimately treat  $x_1$  as an input to  $\text{OMB}_N \circ \text{OR}_M$ ; accordingly, let us interpret  $x_1$  as consisting of  $N$  blocks, each containing  $M$  bits, so that we can write  $x_1 = (x_{1,1}, \dots, x_{1,N}) \in (\{-1, 1\}^M)^N$ . Let  $\gamma(x_1) \in \{-1, 1\}^{\log N}$  be the binary representation of the largest integer  $j$  satisfying  $\text{OR}(x_{1,j}) = -1$ , and let  $\gamma(x_1) = 0$  if no such  $j$  exists. That is,  $\gamma(x_1)$  is the index of the “leading TRUE bit” that gets fed into  $\text{OMB}_N$  when evaluating  $(\text{OMB}_N \circ \text{OR}_M)(x_1)$ .

Similarly, we interpret  $x_2$  as consisting of  $N$  blocks. However, each block now contains  $M \log(N + 1)$  bits, and is comprised of  $M$  sub-blocks, each consisting of  $\log(N + 1)$  bits. Let  $\text{EQ}_{\gamma(x_1)} : \{-1, 1\}^{\log(N+1)} \rightarrow \{-1, 1\}$  denote the function that outputs  $-1$  if and only if its input equals  $\gamma(x_1)$ . Finally, let  $u = (u_1, \dots, u_N) \in (\{-1, 1\}^M)^N$  denote the vector obtained by applying  $\text{EQ}_{\gamma(x_1)}$  to each sub-block of  $x_2$ . That is,  $u$  is the vector obtained by first “pre-processing” each sub-block of  $x_2$  with an “equality test”  $\text{EQ}_{\gamma(x_1)}$  that is determined by  $x_1$ . Finally, we define

$$f = ((\text{OMB}_N \circ \text{OR}_M)(x_1)) \oplus ((\text{OMB}_N \circ \text{OR}_M)(u)). \tag{2}$$

Notice that the dependence of this pre-processing function on  $x_1$  is actually quite mild:  $u$  only depends on the “leading TRUE bit” that gets fed into  $\text{OMB}_N$  when evaluating  $\text{OMB}_N \circ \text{OR}_M(x_1)$ . This mild dependence is what allows  $f$  to be computed by an  $O(\log n)$  decision list.

It is not hard to see that an equivalent way to write  $f$  (that moreover helps reveal its structure as an  $O(\log n)$ -decision list) is:

$$\begin{aligned}
 f(x_1, x_2) = & \text{OMB}_{N^2+2N}(\text{OR}(u_1), \text{OR}(u_2), \dots, \text{OR}(u_N), \\
 & \text{OR}(x_{1,1}), \text{OR}(x_{1,1}) \wedge \text{OR}(u_1), \text{OR}(x_{1,1}) \wedge \text{OR}(u_2), \dots, \text{OR}(x_{1,1}) \wedge \text{OR}(u_N), \\
 & \vdots \\
 & \text{OR}(x_{1,N}), \text{OR}(x_{1,N}) \wedge \text{OR}(u_1), \dots, \text{OR}(x_{1,N}) \wedge \text{OR}(u_N)), \tag{3}
 \end{aligned}$$

where  $u$  is defined as above. It turns out that Representation (2) of  $f$  is useful for establishing lower bounds on the approximate degree of  $f$ , while Representation (3) is more useful for constructing approximating polynomials for  $f$ , and gaining intuition about  $f$ . In particular, Representation (3) suggests a natural method for approximating  $f$ : treat it as a 1-decision list over  $N^2 + 2N$  “derived” variables (and then use an optimal method of approximating 1-decision lists, which are well-understood). The proof of our lower bound (Theorem 1) implicitly shows that this simple approach is essentially optimal. The next subsection briefly explains the details of this approximation method.

### 3.2 A Nearly Matching Upper Bound

We begin by giving the well-known sign-representing polynomial for  $\text{OMB}_N$  itself. Define  $p: \{-1, 1\}^N \rightarrow \mathbb{R}$  via  $p(x_1, \dots, x_N) := 1 + \sum_{i=1}^N (-2)^i \cdot (1 - x_i)/2$ .

It is easy to see that  $\text{OMB}_N(x) \cdot p(x) > 0$  for all  $x \in \{-1, 1\}^N$ , and in fact  $2^{-N-1} \cdot p(x)$  approximates  $\text{OMB}_N$  to error  $\varepsilon = 1 - 2^{-N-1}$ . We now turn to constructing an approximant for the function  $\text{OMB}_N \circ \text{OR}_M$ . Our starting point is a polynomial  $q$  of degree  $O(M^{1/2})$  satisfying the following two properties (cf. [34]).

$$q(x) = 0 \text{ for all } x \in \text{OR}_M^{-1}(+1). \tag{4}$$

$$1 \leq q(x) \leq 2 \text{ for all } x \in \text{OR}_M^{-1}(-1). \tag{5}$$

Denoting an  $(N \cdot M)$ -bit input as  $(x_1, \dots, x_N) \in (\{-1, 1\}^M)^N$ , it is easy to check that

$$\text{OMB}_N \circ \text{OR}_M(x_1, \dots, x_N) = \text{sgn}(g(x_1, \dots, x_N)), \text{ where } g(x_1, \dots, x_N) = 1 + \sum_{i=1}^N (-3)^i \cdot q(x_i).$$

In fact,  $3^{-N-1} \cdot g(x)$  approximates  $\text{OMB}_N \circ \text{OR}_M$  to error  $1 - 3^{-N-1}$ , and has degree equal to that of  $q$ . Recall (cf. Eq. (3)) that in the case  $k = 2$ , our function can be written as

$$\text{OMB}_{N^2+2N}(\text{OR}(u_1), \text{OR}(u_2), \dots, \text{OR}(x_{1,N}) \wedge \text{OR}(u_N)).$$

Using techniques similar to the above, one can obtain a degree  $\tilde{O}(M^{1/2} \log N)$  polynomial that approximates this function to error  $1 - 2^{-O(N^2)}$ . Our lower bound shows this approximation is essentially optimal.

### 3.3 Prospects for Further Improved Lower Bounds

Fix a small constant  $\delta > 0$  and large constant  $\Gamma > 0$ . A remarkable feature of the function  $f$  exhibited in Theorem 1 is that it simultaneously satisfies the following three properties:

- It has threshold degree  $O(\log n)$ .

- It has approximate degree  $\tilde{\Theta}(n^{1/2})$ .
- It has  $\varepsilon$ -approximate degree  $\Omega(n^{1/2-\delta})$  for  $\varepsilon = 1 - 2^{-n^\Gamma}$ .

Hence, while  $f$  has very low threshold degree, it is essentially as hard to approximate  $f$  to error *super-exponentially* close to 1 (i.e., error as large as  $1 - 2^{-n^\Gamma}$  for any constant  $\Gamma > 0$ ), as it is to approximate to error  $1/3$ . To the best of our knowledge, ours is the first known function to exhibit these properties.

Clearly, improving the degree bound in Theorem 1 to be polynomially larger than  $\Omega(n^{1/2})$  will require considering functions of approximate degree larger than  $\Omega(n^{1/2})$ . A natural approach is to simply replace the function  $\text{OR}_M$  appearing in the construction of Section 3.1 with a function of approximate degree polynomially larger than  $\Omega(M^{1/2})$ .

For example, for any constant depth  $d > 0$ , our recent work [9] identifies a depth- $d$   $\text{AC}^0$  circuit  $F_d$  with approximate degree  $\Omega(n^{1-2^{-\Omega(d)}})$ . We conjecture that for any constant  $\delta > 0$ , replacing  $\text{OR}_M$  with  $F_d$  (for sufficiently large  $d(\delta)$ ) in the construction of Section 3.1 yields a function of complexity  $\exp(\Omega(n^{1-\delta}))$ . This would imply a nearly optimal lower bound on the complexity of  $\text{AC}^0$  under each of the complexity measures considered in this work.

## 4 Proof of Theorem 1

The proof of Theorem 1 is lengthy and technical, and deferred to the full version of this work. Here we consider an easier statement, the proof of which is much cleaner, while still capturing the main ideas of the general case.

### 4.1 Simplified Statement and its Proof

The function  $f$  that we exhibit in Theorem 1 is defined over  $k$  “superblocks”, where  $k$  is an arbitrarily large constant (see Section 3.1 for motivation for the superblock terminology). Here, we consider the simpler case of exactly  $k = 2$  superblocks.

**Notation.** Recall (cf. Section 2) that for  $x = (x_1, \dots, x_N) \in \{-1, 1\}^N$ ,  $\beta(x)$  denotes the largest index  $i = 1, \dots, N$  such that  $x_i = -1$  (or  $\beta(x) = 0$  if none exists). Assume for simplicity that  $N+1$  is a power of two. Given an input  $(x, y) \in \{-1, 1\}^N \times (\{-1, 1\}^{\log(N+1)})^N$ , we interpret  $y$  as consisting of  $N$  blocks  $y_1, \dots, y_N$ , each consisting of  $\log(N+1)$  bits. Let  $\text{EQ}_{\beta(x)} : \{-1, 1\}^{\log(N+1)} \rightarrow \{-1, 1\}$  denote the function that outputs  $-1$  if and only if its input equals the binary representation of  $\beta(x)$ . Finally, let  $u = (u_1, \dots, u_N) \in \{-1, 1\}^N$  denote the vector obtained by applying  $\text{EQ}_{\beta(x)}$  to each block of  $y$ . That is,  $u$  is the vector obtained by first “pre-processing” each block of  $y$  with an “equality test”  $\text{EQ}_{\beta(x)}$  that is determined by  $x$ .

**Function Definition.** Define  $F$  via:

$$F(x, y) = \text{OMB}_N(x_1, \dots, x_N) \oplus \text{OMB}_N(u_1(x, y_1), \dots, u_N(x, y_N)) \tag{6}$$

$$= \text{OMB}_N(x_1, \dots, x_N) \oplus \text{OMB}_N(\text{EQ}_{\beta(x)}(y_1), \dots, \text{EQ}_{\beta(x)}(y_N)). \tag{7}$$

► **Proposition 5.** *There exists a constant  $c$  for which the following holds. Let  $d, n \in \mathbb{N}$  where  $n = N + N \log(N+1)$ . Let  $p$  be a polynomial of degree at most  $d$  such that  $|p(x, y)| \geq 1$  and  $p(x, y) \cdot F(x, y) > 0$  for all  $(x, y) \in \{-1, 1\}^n$ . Then there exists an  $(x, y) \in \{-1, 1\}^n$  such that  $|p(x, y)| \geq 2^{(cN/d^2)^2}$ .*

To ease notation below, we will identify each block  $y_i \in \{-1, 1\}^{\log(N+1)}$  with the number in  $[N]$  for which  $y_i$  is the binary representation. That is, while we will write each  $y_i$  as though it were a number  $0, 1, \dots, N$ , it should always be thought of as the binary string representing that number.

**Proof Idea.** Let  $p(x, y)$  be a polynomial of degree at most  $d$  that agrees with  $F$  in sign. Building on Beigel’s proof of Theorem 3, we iteratively apply Lemma 4 to construct a sequence of inputs to the polynomial  $p$ , such that evaluating  $p$  on each point yields a value of (at least) twice the magnitude of the previous evaluation. By choosing these inputs carefully (and crucially exploiting the “pre-processing” step in the definition of  $F_2$  that transforms  $y$  into the vector  $u(x, y)$ , before feeding it into  $\text{OMB}_N$ ), we can apply Lemma 4 a total of  $(cN/d^2)^2$  times. This is a quadratic improvement over the number of times Beigel is able to apply Lemma 4 to  $\text{OMB}_N$  itself. Podolskii [24] used related ideas to obtain a lower bound for a different function, but we are able to avoid the significant quantitative losses that are inherent to his approach.

In more detail, recall that Beigel’s lower bound argument (cf. Theorem 3) for  $\text{OMB}_N$  started with the input  $x^0 = 1^N$ , and iteratively applied Lemma 4 to obtain inputs  $x^1, \dots, x^{cN/d^2}$  such that  $|p(x^t)| \geq 2|p(x^{t-1})|$  for all  $t \geq 1$ . Roughly speaking, the first input to  $F$  that we construct is a point  $(x^1, y^0)$  such that  $u(x^1, y^0) = 1^N$  and the last  $N - 10d^2$  bits of  $x^1$  are all set to 1. Since  $u(x^1, y^0)$  is fed into  $\text{OMB}_N$  in the definition of  $F$ , we are able to apply Beigel’s argument (Theorem 3) to obtain an input  $(x^1, y^1)$  such that  $|p(x^1, y^1)| \geq 2^{cn/d^2} \cdot |p(x^1, y^0)|$ . We then “use” the second block of  $10d^2$  bits of the first superblock to “clean up”  $u$ , in the following sense: we find an  $x^2$  whose last  $N - 20d^2$  bits are all equal to 1, such that  $u(x^2, y^1) = 1^N$  and  $|p(x^2, y^1)| \geq |p(x^1, y^1)|$ . This enables us to apply Beigel’s argument (Theorem 3) a second time, finding an input  $(x^2, y^2)$  such that  $|p(x^2, y^2)| \geq 2^{cn/d^2} \cdot |p(x^1, y^1)|$ . We then use the third  $10d^2$  bits of the first superblock to “clean up”  $u$  yet again, and repeat. We can continue the argument until we have “used up” all the bits of the first superblock, at which point we have obtained the desired lower bound.

**Proof of Proposition 5.** Let  $\ell = 10d^2$  and consider the increasing family of sets  $S_0 \subset S_1 \subset \dots \subset S_{N/\ell} \subseteq \{-1, 1\}^N$  defined as in Lemma 4. Let  $p$  be a polynomial of degree at most  $d$  such that  $p(x, y) \cdot F(x, y) > 0$  for all  $(x, y) \in \{-1, 1\}^n$ . We iteratively construct a sequence of inputs  $(x^0, y^0), (x^1, y^1), \dots, (x^{N/\ell}, y^{N/\ell})$  to  $p$  such that:

- Each  $x^t \in S_t$  and each  $y^t \in [t\ell]^N$ ,
- $|p(x^0, y^0)| \geq 1$ , and
- $|p(x^{t+1}, y^{t+1})| \geq 2^{cN/d^2} \cdot |p(x^t, y^t)|$  for each  $t = 0, \dots, N/\ell - 1$ .

At the conclusion of this process, we obtain an input  $(x^{N/\ell}, y^{N/\ell})$  such that  $|p(x^{N/\ell}, y^{N/\ell})| \geq 2^{(cN/d^2)^2}$ .

We may take as the first input  $(x^0, y^0)$  the point  $(1^N, 0^N)$ . We construct the remaining inputs  $(x^t, y^t)$  iteratively. The following claim formalizes this iterative process.

► **Claim 6.** *Let  $p$  be a polynomial of degree at most  $d$  and suppose  $p(x, y) \cdot F(x, y) > 0$  for all  $(x, y) \in \{-1, 1\}^n$ . Let  $(x^t, y^t)$  be an input with  $x^t \in S_t$  and  $y^t \in [t\ell]^N$ . Then there exists an input  $(x^{t+1}, y^{t+1})$  such that  $|p(x^{t+1}, y^{t+1})| \geq 2^{cN/d^2} \cdot |p(x^t, y^t)|$ , where  $x^{t+1} \in S_{t+1}$  and  $y^{t+1} \in [(t+1)\ell]^N$ .*

**Proof.** We prove the claim in two steps. First, we show that there exists an  $x^{t+1} \in S_{t+1}$  for which  $|p(x^{t+1}, y^t)| \geq |p(x^t, y^t)|$ . Second, we show that there exists a  $y^{t+1} \in [(t+1)\ell]^N$  such that  $|p(x^{t+1}, y^{t+1})| \geq 2^{cN/d^2} \cdot |p(x^{t+1}, y^t)|$ . Putting these steps together yields  $|p(x^{t+1}, y^{t+1})| \geq 2^{cN/d^2} \cdot |p(x^t, y^t)|$ .

**Step 1.** We examine the function  $F(x, y^t)$  (viewed as a function only of  $x$ ). By construction, each block  $y_i^t \leq t\ell$ . Thus,  $\text{EQ}_{\beta(x)}(y_i^t) = 1$  for all  $x \in S_{t+1} \setminus S_t$ , and hence

$$\text{OMB}_N(\text{EQ}_{\beta(x)}(y_1^t), \dots, \text{EQ}_{\beta(x)}(y_N^t)) = \text{OMB}_N(1^N) = 1$$

for all such inputs. As a result,  $F(x, y^t) = \text{OMB}_N(x)$  whenever  $x \in S_{t+1} \setminus S_t$ .

Now consider the polynomial  $q : \{-1, 1\}^N \rightarrow \mathbb{R}$  defined by  $q(x) = p(x, y^t)$ . Then  $q(x) \cdot \text{OMB}_N(x) > 0$  for all  $x \in S_{t+1} \setminus S_t$ . By assumption,  $x^t \in S_t$ . Thus, by Lemma 4, there exists an  $x^{t+1} \in S_{t+1}$  such that  $|q(x^{t+1})| \geq 2 \cdot |q(x^t)|$ . Unpacking the definition of  $q$ , we see that in particular,  $|p(x^{t+1}, y^t)| \geq |p(x^t, y^t)|$ .

**Step 2.** We now show that there exists a  $y^{t+1} \in [(t+1)\ell]^N$ , such that  $|p(x^{t+1}, y^{t+1})| \geq 2^{cN/d^2} \cdot |p(x^{t+1}, y^t)|$ . For  $w \in \{-1, 1\}^N$ , define the string  $y_w$  by  $(y_w)_i = \beta(x^{t+1})$  if  $w_i = -1$  and  $(y_w)_i = y_i^t$  if  $w_i = 1$ . Note that since  $\beta(x^{t+1}) \leq (t+1)\ell$  and each  $y_i^t \in [t\ell]$ , we have that  $y_w \in [(t+1)\ell]^N$  for every  $w \in \{-1, 1\}^N$ . Consider the function  $E : \{-1, 1\}^N \rightarrow \{-1, 1\}$  defined by  $E(w) := F(x^{t+1}, y_w)$ , and observe that

$$E(w) = \text{OMB}_N(x^{t+1}) \oplus \text{OMB}_N(w).$$

Now consider the polynomial  $r(w) := p(x^{t+1}, y_w)$ . Observe that  $y_w$  is an affine function of  $w$ , i.e., we can write

$$r(w) = p\left(x^{t+1}, \left(\frac{1-w_1}{2}\right) \cdot \beta(x^{t+1}) + \left(\frac{1+w_1}{2}\right) \cdot y_1^t, \dots, \left(\frac{1-w_N}{2}\right) \cdot \beta(x^{t+1}) + \left(\frac{1+w_N}{2}\right) \cdot y_N^t\right).$$

Thus  $r$  is a polynomial with  $\deg r \leq \deg p \leq d$ . Moreover,  $r(w) \cdot E(w) > 0$  for all  $w \in \{-1, 1\}^N$ . Since  $E$  is either the function  $\text{OMB}_N$  or its negation, we conclude by Theorem 3 that there exists a  $w^*$  such that  $|r(w^*)| \geq 2^{cN/d^2} \cdot |r(1^N)|$ . Setting  $y^{t+1} := y_{w^*}$  thus yields

$$|p(x^{t+1}, y^{t+1})| = |r(w^*)| \geq 2^{cN/d^2} \cdot |r(1^N)| = 2^{cN/d^2} \cdot |p(x^{t+1}, y^t)|,$$

as we wanted to show. ◀

With Claim 6 established, we conclude the proof of Proposition 5. ◀

---

## References

- 1 Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004. doi:10.1145/1008731.1008735.
- 2 László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory (preliminary version). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 337–347. IEEE Computer Society, 1986. doi:10.1109/SFCS.1986.15.
- 3 Richard Beigel. Perceptrons, PP, and the Polynomial Hierarchy. *Computational Complexity*, 4:339–349, 1994. doi:10.1007/BF01263422.
- 4 Richard Beigel, Nick Reingold, and Daniel A. Spielman. PP is closed under intersection. *J. Comput. Syst. Sci.*, 50(2):191–202, 1995. doi:10.1006/jcss.1995.1017.
- 5 Harry Buhrman, Nikolai K. Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *22nd Annual IEEE Conference on Computational Complexity (CCC 2007), 13-16 June 2007, San Diego, California, USA*, pages 24–32. IEEE Computer Society, 2007. doi:10.1109/CCC.2007.18.

- 6 Mark Bun, Robin Kothari, and Justin Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. In *Symposium on Theory of Computing (STOC)*, pages 297–310. ACM, 2018. doi:10.1145/3188745.3188784.
- 7 Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and markov-berstein inequalities. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *ICALP (1)*, volume 7965 of *Lecture Notes in Computer Science*, pages 303–314. Springer, 2013. doi:10.1007/978-3-642-39206-1\_26.
- 8 Mark Bun and Justin Thaler. Hardness amplification and the approximate degree of constant-depth circuits. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 268–280. Springer, 2015. Full version available at <http://eccc.hpi-web.de/report/2013/151>. doi:10.1007/978-3-662-47672-7\_22.
- 9 Mark Bun and Justin Thaler. A nearly optimal lower bound on the approximate degree of  $AC^0$ . In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 1–12, 2017. doi:10.1109/FOCS.2017.10.
- 10 Mikael Goldmann, Johan Håstad, and Alexander A. Razborov. Majority gates VS. general weighted threshold gates. *Computational Complexity*, 2:277–300, 1992. doi:10.1007/BF01200426.
- 11 Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 86:1–86:15, 2016. doi:10.4230/LIPIcs.ICALP.2016.86.
- 12 András Hajnal, Wolfgang Maass, Pavel Pudlák, Mario Szegedy, and György Turán. Threshold circuits of bounded depth. *J. Comput. Syst. Sci.*, 46(2):129–154, 1993. doi:10.1016/0022-0000(93)90001-D.
- 13 Hartmut Klauck. Lower bounds for quantum communication complexity. *SIAM J. Comput.*, 37(1):20–46, 2007. doi:10.1137/S0097539702405620.
- 14 Adam R. Klivans and Rocco A. Servedio. Learning DNF in time  $2^{\tilde{O}(n^{1/3})}$ . *J. Comput. Syst. Sci.*, 68(2):303–318, 2004. doi:10.1016/j.jcss.2003.07.007.
- 15 Adam R. Klivans and Rocco A. Servedio. Toward attribute efficient learning of decision lists and parities. *Journal of Machine Learning Research*, 7:587–602, 2006. URL: <http://www.jmlr.org/papers/v7/klivans06a.html>.
- 16 Matthias Krause. On the computational power of boolean decision lists. *Computational Complexity*, 14(4):362–375, 2006. doi:10.1007/s00037-005-0203-0.
- 17 Matthias Krause and Pavel Pudlák. On the computational power of depth-2 circuits with threshold and modulo gates. *Theor. Comput. Sci.*, 174(1-2):137–156, 1997. doi:10.1016/S0304-3975(96)00019-9.
- 18 Matthias Krause and Pavel Pudlák. Computing boolean functions by polynomials and threshold circuits. *Computational Complexity*, 7(4):346–370, 1998. doi:10.1007/s000370050015.
- 19 Troy Lee. A note on the sign degree of formulas. *arXiv*, abs/0909.4607, 2009. arXiv:0909.4607.
- 20 Nati Linial and Adi Shraibman. Learning complexity vs. communication complexity. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC 2008, 23-26 June 2008, College Park, Maryland, USA*, pages 53–63. IEEE Computer Society, 2008. doi:10.1109/CCC.2008.28.
- 21 Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1987. doi:10.1007/BF00116827.

- 22 Marvin Minsky and Seymour Papert. *Perceptrons - an introduction to computational geometry*. MIT Press, 1969.
- 23 Noam Nisan. The communication complexity of threshold gates. In *Combinatorics, Paul Erdos is Eighty*, pages 301–315, 1994.
- 24 Vladimir V. Podolskii. A uniform lower bound on weights of perceptrons. In Edward A. Hirsch, Alexander A. Razborov, Alexei L. Semenov, and Anatol Slissenko, editors, *Computer Science - Theory and Applications, Third International Computer Science Symposium in Russia, CSR 2008, Moscow, Russia, June 7-12, 2008, Proceedings*, volume 5010 of *Lecture Notes in Computer Science*, pages 261–272. Springer, 2008. doi:10.1007/978-3-540-79709-8\_27.
- 25 Vladimir Vladimirovich Podolskii. Perceptrons of large weight. *Problems of Information Transmission*, 45(1):46–53, 2009.
- 26 Ronald L. Rivest. Learning decision lists. *Machine Learning*, 2(3):229–246, 1987. doi:10.1007/BF00058680.
- 27 Rocco A. Servedio, Li-Yang Tan, and Justin Thaler. Attribute-efficient learning and weight-degree tradeoffs for polynomial threshold functions. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT*, volume 23 of *JMLR Proceedings*, pages 14.1–14.19. JMLR.org, 2012. URL: <http://www.jmlr.org/proceedings/papers/v23/servedio12/servedio12.pdf>.
- 28 A. A. Sherstov. The power of asymmetry in constant-depth circuits. In *FOCS*, 2015. Full version available at <http://eccc.hpi-web.de/report/2015/147/>.
- 29 Alexander A. Sherstov. Separating  $AC^0$  from depth-2 majority circuits. *SIAM J. Comput.*, 38(6):2113–2129, 2009. doi:10.1137/08071421X.
- 30 Alexander A. Sherstov. The pattern matrix method. *SIAM J. Comput.*, 40(6):1969–2000, 2011. doi:10.1137/080733644.
- 31 Alexander A. Sherstov. Approximating the and-or tree. *Theory of Computing*, 9(20):653–663, 2013.
- 32 Alexander A. Sherstov. The intersection of two halfspaces has high threshold degree. *SIAM J. Comput.*, 42(6):2329–2374, 2013. doi:10.1137/100785260.
- 33 Alexander A. Sherstov. Breaking the Minsky-Papert barrier for constant-depth circuits. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 223–232. ACM, 2014. doi:10.1145/2591796.2591871.
- 34 Justin Thaler. Lower bounds for the approximate degree of block-composed functions. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 17:1–17:15, 2016. doi:10.4230/LIPIcs.ICALP.2016.17.
- 35 Leslie G. Valiant. A theory of the learnable. In Richard A. DeMillo, editor, *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1984, Washington, DC, USA*, pages 436–445. ACM, 1984. doi:10.1145/800057.808710.

## **A** Applications

### **A.1** An Application in Communication Complexity

We briefly describe an additional application of our results. By combining Theorem 1 with standard machinery, we obtain an improved separation between the analogues of the complexity classes  $\mathbf{PP}$  and  $\mathbf{P}^{\mathbf{NP}}$  in communication and query complexity. Specifically, for a function  $F: \{-1, 1\}^n \times \{-1, 1\}^n \rightarrow \{-1, 1\}$ , let  $\mathbf{PP}(F)$  and  $\mathbf{P}^{\mathbf{NP}}(F)$  respectively denote the least cost of a  $\mathbf{PP}$  and  $\mathbf{P}^{\mathbf{NP}}$  communication protocol for  $F$ . In both the

communication and query complexity settings, for any constant  $\delta > 0$ , we exhibit an  $F$  satisfying  $\mathbf{PP}(F) = \Omega(n^{1/2-\delta})$  and  $\mathbf{P}^{\mathbf{NP}}(F) = O(\log^2 n)$ . This improves over prior work that gave an  $F$  satisfying  $\mathbf{PP}(F) = \Omega(n^{2/5})$  and  $\mathbf{P}^{\mathbf{NP}}(F) = O(\log n)$  [34] and earlier work of Buhrman et al. that gave an  $F$  satisfying  $\mathbf{PP}(F) = \Omega(n^{1/3})$  and  $\mathbf{P}^{\mathbf{NP}}(F) = O(\log n)$  [5]. We direct the interested reader to [34] for further details on this application.

## A.2 An Application to Learning Theory

A notorious open question in learning theory, dating back to Valiant’s seminal paper introducing the PAC model [35], is whether polynomial size DNF can be PAC learned in polynomial time. The fastest known algorithm, due to Klivans and Servedio [14] runs in time  $\exp(\tilde{O}(n^{1/3}))$ , and follows from their (tight)  $\tilde{O}(n^{1/3})$  upper bound on the threshold degree of any polynomial size DNF. Klivans and Servedio’s threshold degree upper bound makes use of polynomials with coefficients that are *doubly*-exponentially large (as large as  $2^{2^{\Theta(n^{1/3})}}$ ). They pose the following open question:

**Open Question (Klivans and Servedio [14]).** Is every polynomial size DNF computed by a polynomial threshold function with degree  $\tilde{O}(n^{1/3})$  and coefficients of magnitude at most  $2^{\tilde{O}(n^{1/3})}$ ?

A positive resolution of this open question would imply much *simpler* (though not necessarily faster) learning algorithms for DNFs. Our results can be viewed as a significant step toward resolving Klivans and Servedio’s question in the *negative* (in fact, achieving a full resolution of the open question was our original motivation for this work). This is because Theorem 1 implies that for any constants  $\Gamma, \delta > 0$  there is an  $O(\log n)$ -decision list  $F$  of polynomial length, such that any polynomial of degree  $O(n^{1/2-\delta})$  that sign-represents  $F$  requires coefficients of magnitude  $2^{n^\Gamma}$ . While  $O(\log n)$ -decision lists are not necessarily computed by DNFs, they seem to be only slightly more expressive than DNFs. In particular, they are computed by depth 3 circuits of logarithmic bottom fan-in.

Furthermore, our result points out an inherent limitation of Klivans and Servedio’s construction. Their threshold degree upper bound for DNFs works as follows. They first transform the DNF into a (slight generalization of) a  $k$ -decision list, for some  $k = O(n^{1/3})$ . Second, they prove a  $\tilde{O}(k)$  upper bound on the threshold degree of any  $k$ -decision list. Our result implies that the super-exponentially large coefficients arising in the second step of Klivans and Servedio’s construction are necessary, even for  $k$  bounded by  $O(\log n)$ . Hence, any affirmative resolution of Klivans and Servedio’s question cannot go through  $k$ -decision lists, even for  $k = O(\log n)$ .

## A.3 Threshold Weight of Depth Three Circuits

A *polynomial threshold function* (PTF) for a Boolean function  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  is a polynomial  $p : \{-1, 1\}^n \rightarrow \mathbb{R}$  with *integer coefficients* that agrees in sign with  $f$  on all Boolean inputs. The *weight* of an  $n$ -variate polynomial  $p$  is the sum of the absolute values of its coefficients. The *degree- $d$  threshold weight* of a Boolean function  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ , denoted  $W(f, d)$ , is defined to be the least weight of a degree- $d$  PTF for  $f$ . We let  $W(f)$  denote the quantity  $W(f, n)$ , i.e., the least weight of any threshold function for  $f$  regardless of its degree. Threshold weight upper bounds underly some of the most powerful techniques in computational learning theory based on the classic Perceptron [22] and Winnow [21] algorithms (see [8, Section 8.3] for a discussion). Thus, our threshold weight lower bounds impose limitations on how efficiently such algorithms can learn depth three circuits.



Degree- $d$  threshold weight is closely related to  $\varepsilon$ -approximate degree when  $\varepsilon$  is very close to 1:

► **Lemma 7.** *Let  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  be a Boolean function, and let  $w > 0$ . If  $\widetilde{\deg}_{1-\frac{1}{w}}(f) > d$ , then  $W(f, d) > w$ .*

**Proof.** We prove the contrapositive, i.e., that any PTF  $p$  for  $f$  having weight  $w$  and degree  $d$  can be transformed into a uniform approximation to  $f$  with error  $1 - \frac{1}{w}$ . Let  $p$  be such a PTF. Since  $p$  has integer coefficients and is nonzero on Boolean inputs,  $|p(x)| \geq 1$  on  $\{-1, 1\}^n$ . Moreover,  $|p(x)| \leq w$  by the weight bound, so the polynomial  $\frac{1}{w} \cdot p(x)$  satisfies  $|\frac{1}{w} \cdot p(x) - f(x)| \leq 1 - \frac{1}{w}$  for every  $x \in \{-1, 1\}^n$ . ◀

Thus, our main results yield new lower bounds on the degree- $d$  threshold weight of circuits of depth three.

► **Corollary 8.** *For any arbitrarily small constant  $\delta > 0$  and any arbitrarily large constant  $\Gamma > 1$ , there exists a depth three Boolean circuit  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  (with logarithmic bottom fan-in) such that  $W(f, n^{1/2-\delta}) > 2^{n^\Gamma}$ .*

Moreover, a result of Krause [16] allows us to translate each of these lower bounds into a degree independent threshold weight lower bound for a related function.

► **Lemma 9** ([16], Lemma 3.4). *Let  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  be a Boolean function, and define  $F : \{-1, 1\}^{3n} \rightarrow \{-1, 1\}$  by*

$$F(x_1, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_n) := f(\dots, (\bar{z}_i \wedge x_i) \vee (z_i \wedge y_i), \dots).$$

*Then  $W(F) \geq W(f, d)$  for all  $d$  for which  $2^d \geq W(f, d)$ .*

When this transformation is applied to a function  $f$  computed by a Boolean circuit of depth  $d$  with logarithmic bottom fan-in, the resulting function  $F$  is also computed by a depth  $d$  circuit with logarithmic bottom fan-in. To see this, note that if  $g$  is any function that depends on  $O(\log n)$  variables, then  $G(x, y, z) := g((\bar{z} \wedge x) \vee (z \wedge y))$  also depends on  $O(\log n)$  variables. Hence,  $G$  is computed by either a DNF or CNF of size  $\text{poly}(n)$  and bottom fan-in  $O(\log n)$ . So while  $F$  is naturally computed by a circuit of depth  $d + 2$ , the bottom three levels of gates can be replaced by such DNF or CNF formulae so as to merge a layer of gates and obtain a depth  $d$  circuit with logarithmic bottom fan-in.

► **Corollary 10.** *For any arbitrarily small constant  $\delta > 0$ , there exists a depth three Boolean circuit  $F : \{-1, 1\}^{3n} \rightarrow \{-1, 1\}$  (with logarithmic bottom fan-in) such that  $W(F) > \exp(\Omega(n^{1/2-\delta}))$ .*

While the weight bounds of Corollaries 8 and 10 are stated for polynomial threshold functions over  $\{-1, 1\}^n$  (i.e., for polynomials that are integer linear combinations of parities), a now standard transformation [18] shows that the same threshold weight lower bound also holds for polynomials over  $\{0, 1\}^n$  (i.e., for integer linear combinations of conjunctions) up to polynomial factors.

## A.4 Discrepancy of Depth Three Circuits

Discrepancy is a central quantity in communication complexity and circuit complexity. For instance, an upper bound on the discrepancy of a Boolean function  $f : X \times Y \rightarrow \{-1, 1\}$  yields lower bounds for computing  $f$  in essentially every model of communication complexity.

In particular, the discrepancy of  $f$  essentially characterizes its small-bias communication complexity in the **PP** model of Babai et al. [2]. Theorem 1 yields a new exponentially small upper bound on the discrepancy of a depth three circuit.

For a Boolean function  $f : X \times Y \rightarrow \{-1, 1\}$ , let  $M^{(f)}$  be its communication matrix  $M^{(f)} = [f(x, y)]_{x \in X, y \in Y}$ . A combinatorial rectangle of  $X \times Y$  is a set of the form  $A \times B$  with  $A \subseteq X$  and  $B \subseteq Y$ . For a distribution  $\mu$  over  $X \times Y$ , the discrepancy of  $f$  with respect to  $\mu$  is defined to be the maximum over all rectangles  $R$  of the *bias* of  $f$  on  $R$ . That is:

$$\text{disc}_\mu(f) = \max_R \left| \sum_{(x,y) \in R} \mu(x,y) f(x,y) \right|.$$

The discrepancy of  $f$ , denoted  $\text{disc}(f)$ , is defined to be  $\min_\mu \text{disc}_\mu(f)$ .

Sherstov's pattern matrix method [30] shows how to generically transform an  $\text{AC}^0$  function with high threshold degree or high threshold weight into another  $\text{AC}^0$  function with low discrepancy.

► **Theorem 11** (cf. [30], adapted from Corollary 1.2 and Theorem 7.3). *Let  $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$  be given, and define the communication problem  $F : \{-1, 1\}^{4n} \times \{-1, 1\}^{4n} \rightarrow \{-1, 1\}$  by*

$$F(x, y) = f(\dots, \bigvee_{j=1}^4 (x_{i,j} \wedge y_{i,j}), \dots).$$

*Then for every integer  $d \geq 0$ , we have*

$$\text{disc}(F)^2 \leq \max \left\{ \frac{2n}{W(f, d-1)}, 2^{-d} \right\}.$$

Recall that  $W(f, d-1)$  is the least weight of any degree  $d-1$  PTF for  $f$ . We apply the pattern matrix method to the function  $f$  of Corollary 8. By the same argument as in Section A.3, the pattern matrix method does not increase the depth of the circuits computing these functions. We thus obtain a new discrepancy upper bound for circuits of depth three:

► **Corollary 12.** *For any arbitrarily small constant  $\delta > 0$ , there exists a depth three Boolean circuit  $F : \{-1, 1\}^n \rightarrow \{-1, 1\}$  (with logarithmic bottom fan-in) such that  $\text{disc}(F) < \exp(-\Omega(n^{1/2-\delta}))$ .*

**Application to Circuit Complexity.** It is well-known that a discrepancy upper bound for a function  $F$  yields a lower bound on the size of Majority-of-Threshold circuits computing  $F$  [10, 12, 23, 29]. Indeed, the exponential Majority-of-Threshold circuit size lower bounds of [29, 30, 5, 8, 33, 28] for  $\text{AC}^0$  are all proved using discrepancy. Our discrepancy upper bound of Corollary 12 sharpens these previous lower bounds by yielding a depth three Boolean circuit  $F$  of polynomial size such that any Majority-of-Threshold circuit computing  $F$  requires size  $\exp(\Omega(n^{1/2-\delta}))$ .

► **Corollary 13.** *For any arbitrarily small constant  $\delta > 0$ , there exists a depth three Boolean circuit  $F : \{-1, 1\}^n \rightarrow \{-1, 1\}$  (with logarithmic bottom fan-in) such that any Majority-of-Threshold circuit computing  $F$  has size at least  $\exp(\Omega(n^{1/2-\delta}))$ .*

Combining Corollaries 10, 12, and 13 yields Corollary 2 from the introduction, noting that the function exhibited in Corollary 10 appears as a subfunction of the one in Corollaries 12 and 13.