Hardness Amplification Proofs Require Majority

Ronen Shaltiel*
Department of Computer Science
University of Haifa, Israel
ronen@cs.haifa.ac.il

Emanuele Viola†
Department of Computer Science
Columbia University, USA
viola@cs.columbia.edu

ABSTRACT

Hardness amplification is the fundamental task of converting a δ-hard function $f : \{0,1\}^n \to \{0,1\}$ into a $(1/2 - \epsilon)$-hard function $Amp(f)$, where f is γ-hard if small circuits fail to compute f on at least a γ fraction of the inputs. Typically, ϵ, δ are small (and $\delta = 2^{-k}$ captures the case where f is worst-case hard). Achieving $\epsilon = 1/n^{\omega(1)}$ is a prerequisite for cryptography and most pseudorandom-generator constructions.

In this paper we study the complexity of black-box proofs of hardness amplification. A class of circuits $\mathcal{D}$ *proves* a hardness amplification result if for any function h that agrees with $Amp(f)$ on a $1/2 + \epsilon$ fraction of the inputs there exists an oracle circuit $D \in \mathcal{D}$ such that D^h agrees with f on a $1 - \delta$ fraction of the inputs. We focus on the case where every $D \in \mathcal{D}$ makes *non-adaptive* queries to h. This setting captures most hardness amplification techniques. We prove two main results:

1. The circuits in $\mathcal{D}$ "can be used" to compute the majority function on $1/\epsilon$ bits. In particular, these circuits have large depth when $\epsilon \leq 1/\text{poly} \log n$.
2. The circuits in $\mathcal{D}$ must make $\Omega\left(\log(1/\delta)/\epsilon^2\right)$ oracle queries.

Both our bounds on the depth and on the number of queries are tight up to constant factors.

Our results explain why hardness amplification techniques have failed to transform known lower bounds against constant-depth circuit classes into strong average-case lower bounds. When coupled with the celebrated "Natural Proofs" result by Razborov and Rudich (J. CSS '97) and the pseudorandom functions by Naor and Reingold (J. ACM '04), our results show that *standard techniques for hardness amplification can only be applied to those circuit classes for which standard techniques cannot prove circuit lower bounds.*

Our results reveal a contrast between *Yao's XOR Lemma* ($Amp(f) := f(x_1) \oplus \ldots \oplus f(x_t) \in \{0,1\}$) and the *Direct-Product Lemma* ($Amp(f) := f(x_1) \circ \ldots \circ f(x_t) \in \{0,1\}^t$; here $Amp(f)$ is non-Boolean). Our results (1) and (2) apply to Yao's XOR lemma, whereas known proofs of the direct-product lemma violate both (1) and (2).

One of our contributions is a new technique to handle "non-uniform" reductions, i.e. the case when $\mathcal{D}$ contains many circuits.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: GENERAL; F.1.3 [**Theory of Computation**]: Complexity Measures and Classes—*Relations among complexity measures*

General Terms

Theory

Keywords

average-case complexity, hardness amplification, majority, constant-depth circuits, black-box, natural proofs.

1. INTRODUCTION

Proving circuit lower bounds is a major goal of Complexity Theory. However, the celebrated "Natural Proofs" result by Razborov and Rudich [50], coupled with the pseudorandom functions by Naor and Reingold [45], marks the class of polynomial-size constant-depth circuits *with majority gates* (TC^0) as a fundamental limit for most currently available lower bounding techniques. This limitation already applies to *worst-case* lower bounds, where one seeks a function that small circuits fail to compute on *at least one* input. In particular, it applies to *average-case* lower bounds, where one seeks a function that small circuits fail to compute on *many* inputs. Average-case hard functions are especially important as they are a prerequisite for most modern cryptography and can be used to construct pseudorandom generators [46] which in turn have a striking variety of applications (see, e.g., the books by Goldreich [17, 18]). We stress that both these applications require *strongly* average-case hard functions. That is functions that small circuits cannot compute

*This research was supported by BSF grant 2004329 and ISF grant 686/07.

†Supported by grants NSF award CCF-0347282 and NSF award CCF-0523664. Research partially done while at Harvard University, supported by grants NSF CCR-0133096, US-Israel BSF 2002246, and ONR N-00014-04-1-0478, and at the Institute for Advanced Study, supported by NSF grant CCR-0324906.

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.

with even a small advantage over random guessing, for a randomly chosen input. (For concreteness, the reader may think of a function $f : \{0,1\}^n \to \{0,1\}$ that any small circuit fails to compute with probability $1/2 - 1/n^{\omega(1)}$ over the choice of the input).

As we do not know how to prove unconditional lower bounds for general circuit classes, a long line of research has focused on *hardness amplification*. This is the task of transforming worst-case hard functions (or sometimes *mildly* average-case hard functions) into average-case hard functions [65, 40, 8, 6, 7, 32, 20, 15, 35, 36, 12, 55, 59, 51, 56, 47, 60, 58, 29, 52, 24, 33, 34, 22]. This research was largely successful in its goal. In particular, it provided worst-case to average-case connections within many complexity classes. Many of these connections give *strongly* average-case hard functions. This research also spurred fruitful interaction with coding theory (see, e.g., the survey by Trevisan [57]).

Complexity theory has produced many exciting and useful lower bounds for restricted computational models, most notably against classes of circuits with unbounded fan-in and constant depth with various gates [16, 66, 27, 49, 54, 28, 25, 5, 26]. In some of these classes we in fact can prove worst-case lower bounds, but cannot prove *strongly* average-case lower bounds (e.g. [49, 54, 5]). Several such examples are surveyed in the full version of this paper and in [62, Chapter 6]; for concreteness, an example is the lower bound against constant-depth circuits with And, Or and Parity gates [49, 54]. One would expect that hardness amplification techniques could be used to produce strongly average-case lower bounds from the known lower bounds (which would in turn give pseudorandom generators for these classes [46]). But in fact "standard hardness amplification techniques" fail.

In this paper we show that:

"standard hardness amplification techniques" only apply when starting with hardness against circuits that can compute the majority function.

This explains the following *"lose-lose" phenomenon:* For classes that are weaker than TC^0 (e.g. constant-depth circuits, or constant-depth circuits with parity gates) we can prove lower bounds, however we do not have hardness amplification theorems, while for classes at least as powerful as TC^0 we have hardness amplification theorems but cannot prove circuit lower bounds; see Figure 1.

A couple of remarks is in order. First, our results likely do not apply to "every conceivable" class of circuits, but rather they apply to the most well-studied ones. Second, we note that, just like Razborov and Rudich's result [50] is not claiming that it is impossible to prove lower bounds for classes like TC^0, but rather that certain techniques will not do, this work is not claiming that it is impossible to prove strong average-case hardness results for circuit classes weaker than TC^0, but that we cannot obtain such results by "standard hardness amplification techniques." We elaborate on these techniques next.

1.1 Hardness amplification

In this section we review the notion of hardness amplification. Let us start by formalizing our notion of hardness.

Definition 1.1 (Average-case hardness). *A function $f : \{0,1\}^k \to \{0,1\}$ is δ-hard for a class of circuits $\mathcal{C}$ (e.g., all circuits of size s) if for every circuit $C \in \mathcal{C}$ we have $\Pr_{x \in \{0,1\}^k}[C(x) \neq f(x)] \geq \delta$.*

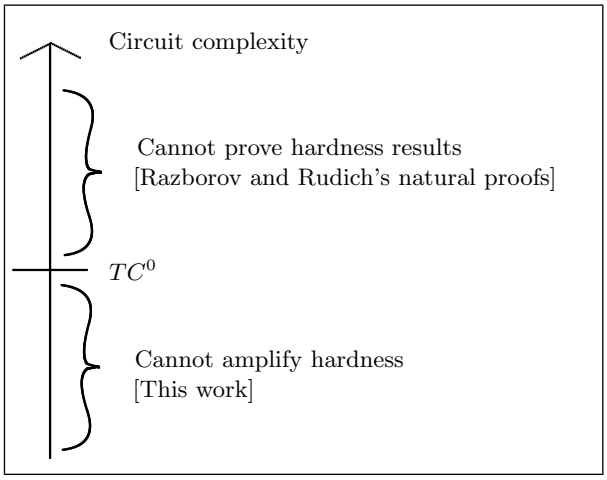

Figure 1: Reach of "standard techniques." Recall TC^0 is the class of polynomial-size constant-depth circuits with majority gates.

Hardness amplification is the generic task of transforming a given function $f : \{0,1\}^k \to \{0,1\}$ that is δ-hard for a class of circuits $\mathcal{C}$ into another function $Amp(f) : \{0,1\}^n \to \{0,1\}$ that is $(1/2 - \epsilon)$-hard for a related class of circuits $\mathcal{C}'$, where one wants ϵ as small as possible and n not much larger than k. The first and most important example of hardness amplification is *Yao's XOR lemma* (cf. [20]), which works as follows. We let $n := t \cdot k$ for a parameter t and on input $(x_1, \ldots, x_t) \in \left(\{0,1\}^k\right)^t = \{0,1\}^n$ we define

$$Amp(f)(x_1, \ldots, x_t) := f(x_1) \oplus \cdots \oplus f(x_t),$$

where $\oplus$ denotes exclusive OR. The lemma states that if f is δ-hard for (the class of) circuits of size s, then choosing $t := O(\log(1/\epsilon)/\delta)$ one has that $Amp(f)$ is $(1/2 - \epsilon)$-hard for circuits of size $s \cdot \mathrm{poly}(\epsilon \cdot \delta/k)$. In particular, if f is 1/3-hard for circuits of superpolynomial size $s = n^{\omega(1)}$, then by choosing a suitable $t := \omega(\log n)$ we obtain a $(1/2 - 1/n^{\omega(1)})$-hard function, (recall that such a function is a prerequisite of most cryptography and can be used to construct pseudorandom generators [46]).

Yao's XOR lemma is not useful when starting from worst-case hard functions, i.e., when $\delta = 2^{-k}$. Hardness amplification from worst-case hardness is still possible (e.g., [40, 8, 6, 7, 15, 12, 55, 59]) but is more difficult. This distinction is not relevant to our work which, jumping ahead, proves limitations on hardness amplification that already apply when amplifying from constant hardness $\delta = \Omega(1)$ (and in particular apply when amplifying from worst-case hardness $\delta = 2^{-k}$).

1.2 Black-box hardness amplification

We now explain what we mean by "standard techniques" for proving hardness amplification theorems. To explain this, we use the classical notion of an *oracle circuit* $D^h(x)$, where $h : \{0,1\}^n \to \{0,1\}$. This is simply a circuit with special oracle gates that on input $y \in \{0,1\}^n$ return the value $h(y) \in \{0,1\}$. We note that this notion also makes sense when restricting the depth of the circuit D. It has been observed several times (see, e.g., [56]) that most proofs of hardness amplification in the literature are *black-box* in the following sense.

DEFINITION 1.2. (BLACK-BOX HARDNESS AMPLIFICATION). *A $\delta \to (1/2-\epsilon)$ black-box hardness amplification with input lengths k and n is a pair $(Amp, \mathcal{D})$ such that Amp is a map from functions $f : \{0,1\}^k \to \{0,1\}$ to functions $Amp(f) : \{0,1\}^n \to \{0,1\}$, $\mathcal{D}$ is a class of oracle circuits on k input bits (e.g., all oracle circuits of size s), and the following holds:*

For every function $f : \{0,1\}^k \to \{0,1\}$ and every function $h : \{0,1\}^n \to \{0,1\}$ such that

$$\Pr_{y \in \{0,1\}^n} [h(y) \neq Amp(f)(y)] < 1/2 - \epsilon$$

there is an oracle circuit $D \in \mathcal{D}$ such that

$$\Pr_{x \in \{0,1\}^k} \left[D^h(x) \neq f(x) \right] < \delta.$$

The black-box hardness amplification is non-adaptive q-query *if every circuit $D \in \mathcal{D}$ makes q non-adaptive queries to h. Finally, we say that a class of circuits $\mathcal{D}$* proves *a black-box hardness amplification (with certain parameters) if there is a map Amp such that $(Amp, \mathcal{D})$ is a black-box hardness amplification (with the same parameters).*

Why black-box hardness amplification lets us amplify hardness.

It is instructive to verify that black-box hardness amplification indeed lets us amplify hardness. To see this, suppose that $(Amp, \mathcal{D})$ is a q-query $\delta \to (1/2-\epsilon)$ black-box hardness amplification where $\mathcal{D}$ is the class of circuits of size s. Now let $f : \{0,1\}^k \to \{0,1\}$ be δ-hard for (the class of) circuits of size $t \geq 2 \cdot s$. Observe that indeed the function $Amp(f) : \{0,1\}^n \to \{0,1\}$ is $(1/2-\epsilon)$-hard for circuits of size $t/(2 \cdot q)$ (where recall q is the number of oracle queries made by circuits in $\mathcal{D}$). This is proven by a standard counterpositive argument. Suppose for the sake of contradiction that there exists a circuit h of size $t/(2 \cdot q)$ that computes $Amp(f)$ on more than a $1/2+\epsilon$ fraction of the inputs. Then by definition of black-box hardness amplification there is a circuit $D \in \mathcal{D}$ such that D^h computes f on more than a $1-\delta$ fraction of the inputs. Since D has size s and makes q oracle queries, by replacing each query with a copy of the circuit for h we see that D^h can be computed by a circuit of size $q \cdot t/(2 \cdot q) + s \leq t/2 + s \leq t$, contradicting our assumption that f was δ-hard for circuits of size t.

It is also instructive to remark that, in the language of Definition 1.2, Yao's XOR lemma is a $\delta \to (1/2-\epsilon)$ black-box hardness amplification $(Amp, \mathcal{D})$ with input lengths k and n, where $n = O(k \cdot \log(1/\epsilon)/\delta)$ and $\mathcal{D}$ is the class of circuits of size $\text{poly}(k/(\epsilon \cdot \delta))$.

The complexity of $\mathcal{D}$.

We want to stress that the complexity of the class $\mathcal{D}$ plays a crucial role when deriving average-case hardness results using a black-box hardness amplification. Specifically, to obtain hardness amplification the initial function $f : \{0,1\}^k \to \{0,1\}$ must be hard for a class of circuits that contains $\mathcal{D}$. This is a key point for our results which will essentially show that $\mathcal{D}$ has to be at least as powerful as TC^0, the class of constant-depth circuits with majority gates. Thus for hardness amplification we need to start from a lower bound against TC^0.

Non-uniformity.

Another aspect we wish to stress is the *non-uniformity* of the notion of black-box hardness amplification. In Definition 1.2 the circuit $D \in \mathcal{D}$ is allowed to depend arbitrarily on both the δ-hard function f and the function h that approximates $Amp(f)$. It can be shown that some non-uniformity is *necessary* for black-box hardness amplification: $|\mathcal{D}| \geq (1/\epsilon)^{\Omega(1)}$ [59]. Establishing hardness amplification results with small non-uniformity (e.g. $|\mathcal{D}| = \text{poly}(1/\epsilon)$) is important for achieving "uniform hardness amplification within NP" and is the focus of a lot of recent attention (see Section 1.4 on related work). In this work we give impossibility results for black-box hardness amplification and therefore are interested in handling *any* black-box hardness amplification, including ones which use large non-uniformity (e.g. $|\mathcal{D}| = \exp(1/\epsilon)$).

1.3 Our results

The main result of this paper applies to *non-adaptive* black-box hardness amplification and can be stated informally as follows:

If a set of circuits $\mathcal{D}$ proves non-adaptive $\delta \to (1/2-\epsilon)$ black-box hardness amplification then $\mathcal{D}$ "can be used" to compute majority on $1/\epsilon$ bits.	$(\star)$

The formal statement of the above result requires a bit of notation, and is deferred to Theorem 1.6 at the end of this section where, intuitively, we show how oracle access to the circuits $\mathcal{D}$ is sufficient to compute majority. For now we state a qualitatively weaker result which requires less notation. Specifically, the next theorem shows that if $\mathcal{D}$ proves non-adaptive $\delta \to (1/2-\epsilon)$ black-box hardness amplification, then the depth of the circuits in $\mathcal{D}$ must be large whenever ϵ is small (cf. Definition 1.2 for the definition of "proves"). This weak form of the theorem intuitively follows from $(\star)$ by using the well-known fact that computing the majority function on $m := 1/\epsilon$ bits by circuits of depth d requires size $s \geq \exp\left(m^{\Omega(1/d)}\right) = \exp\left((1/\epsilon)^{\Omega(1/d)}\right)$, i.e. exponential in $1/\epsilon$ [27, 49, 54].

THEOREM 1.3. (DECODING REQUIRES MAJORITY, STATED IN TERMS OF CIRCUIT DEPTH). *Suppose that a class of non-adaptive oracle circuits $\mathcal{D}$ proves a $(\delta = 1/3) \to (1/2-\epsilon)$ black-box hardness amplification $(Amp, \mathcal{D})$ with input lengths k and n.*

Suppose that every circuit $D \in \mathcal{D}$ has size s and depth d. Then

$$s \geq \min\left\{ \exp\left((1/\epsilon)^{\Omega(1/d)} \right), 2^{\Omega(k)} \right\}.$$

In particular, Theorem 1.3 implies that $\text{poly}(n)$-size constant-depth circuits (i.e., d is fixed and $s = \text{poly}(n)$ grows) can only prove hardness amplification up to $1/2-\epsilon \leq 1/2 - 1/\text{poly}\log n$. This should be contrasted with standard hardness amplifications (e.g., [20]) that show that if we do not put any restriction on the depth of the circuits in $\mathcal{D}$ then circuits of size $\text{poly}(n)$ can prove hardness amplification up to $1/2 - 1/n$.

We remark that, for constant-depth circuits, the size bound in Theorem 1.3 is tight. This follows easily from Impagliazzo's beautiful hard-core set theorem [32] when amplifying from constant hardness $\delta = \Omega(1)$. Moreover, Impagliazzo's result [32] conceptually matches our result $(\star)$

by showing that computing majority on $\text{poly}(1/\epsilon)$ bits is "all that is needed" for proving hardness amplification. Precisely this feature was exploited a few times in complexity theory, for example in Klivans' elegant work [38]. When amplifying from worst-case hardness $\delta = 2^{-k}$, the construction by Goldwasser et al. [21][1] again matches the size bound in our Theorem 1.3.

Our second main result is a lower bound on the number of queries made by circuits $\mathcal{D}$ in any black-box hardness amplification $(Amp, \mathcal{D})$. One reason for studying the number of queries necessary for proving hardness amplification is the loss in circuit size, i.e. the difference between the circuit sizes that come up in the assumption and conclusion of the hardness amplification theorem. The question of how much loss is necessary has been raised a number of times (see, e.g., [20, 37]) but was never answered in generality until this paper. Additional motivation is given in the full version of this paper.

THEOREM 1.4 (DECODING REQUIRES MANY QUERIES). *There is a universal constant $C > 1$ such that the following holds. Let $(Amp, \mathcal{D})$ be a non-adaptive q-query $\delta \to (1/2 - \epsilon)$ black-box hardness amplification. Suppose that $\log |\mathcal{D}| \leq 2^{k/C}$, and $n, k \geq C^2$, and that both δ and ϵ are between $2^{-k/C}$ and $1/3$.*

Then

$$q \geq \frac{1}{C} \cdot \frac{\log(1/\delta)}{\epsilon^2}.$$

We also note that the lower bound of Theorem 1.4 is tight (up to constants) even when only considering XOR-lemmas. This is because Impagliazzo's proof of the XOR-lemma [32] can be made to work with $q = O\left(\log(1/\delta)/\epsilon^2\right)$ queries matching our lower bound.[2]

It has been observed (see e.g. [56]) that black-box hardness amplification is closely related to list-decodable codes. Using this connection our results can be seen as lower bounds on the "complexity of decoding" locally-decodable codes. We explain this view in the full version of this paper.

XOR lemma vs. direct product: A qualitative difference.

So far we have discussed hardness amplification where the amplified function $Amp(f)$ is Boolean, i.e. its range is $\{0,1\}$, and our leading example was Yao's XOR lemma which recall is defined as $Amp(f) := f^{\oplus t}(x_1, \ldots, x_t) = f(x_1) \oplus \ldots \oplus f(x_t) \in \{0,1\}$, where $\oplus$ denotes exclusive-or.

Hardness amplification where the amplified function $Amp(f) : \{0,1\}^n \to \{0,1\}^t$ is *not* Boolean, i.e. $t \geq 1$, is also widely studied. The first and most important example of this is the *direct product* which is defined as follows $Amp(f) := f^{\circ t}(x_1, \ldots, x_t) = f(x_1) \circ \ldots \circ f(x_t) \in \{0,1\}^t$, where $\circ$ denotes concatenation. Recall that in XOR-lemmas we are interested in amplifying hardness from δ to $1/2 - \epsilon$, whereas in direct-product lemmas we are interested in amplifying from δ to $1 - \epsilon$.

The direct product and the XOR lemma, and more generally Boolean and non-Boolean hardness amplification, have often been regarded as essentially interchangeable. In fact, many proofs of Boolean hardness amplification proceed by proving the direct product first and then transforming the amplified function $f^{\circ t}$ into a Boolean function (see, e.g., [20, 35, 55, 47, 56, 29]), often using the remarkable Goldreich-Levin Theorem [19]. The converse, proving a direct product lemma from an XOR lemma, is much easier [64].

By contrast, *our results show that Yao's XOR lemma and the direct product lemma are qualitatively different.*

The main difference is that the proof of Yao's XOR lemma requires majority, whereas the proof of the direct product lemma does not. Specifically, our results show that if a class $\mathcal{D}$ proves a $(\delta = 1/3) \to (1/2 - \epsilon)$ black-box hardness amplification, such as Yao's XOR lemma, then "$\mathcal{D}$ can compute majority," and in particular $\mathcal{D}$ requires either large depth or exponential size in $1/\epsilon$ (Theorem 1.3). On the other hand, there are black-box proofs of the $\delta \to (1 - \epsilon)$ direct-product lemma that can be implemented by small constant-depth circuits for arbitrary $\epsilon > 0$. For example, this is achieved by the proof of Goldreich et al. [20].[3]

Another difference can be seen in the number of queries. The proof of the direct-product lemma in [20] uses $q = O\left(\log(1/\delta)/\epsilon\right)$ queries, and note that for small ϵ this beats our $\Omega\left(\log(1/\delta)/\epsilon^2\right)$ lower bound that applies to XOR lemmas (Theorem 1.4).

Finally, we point out that the techniques in this paper show that $q = \Omega(\log(1/\delta)/\epsilon)$ queries are necessary for black-box proofs of the direct-product lemma (details omitted), which again matches the upper bound in [20].

Our main result: The general form.

We now state our main result that hardness amplification requires majority in its full generality. Previously, we had stated a corollary of it that was tailored to circuit depth (Theorem 1.3). The general form of our results shows that the circuits $\mathcal{D}$ in a black-box hardness amplification $(Amp, \mathcal{D})$ can be used to compute the majority function by a small constant-depth circuit. The way in which we are going to use a circuit $D \in \mathcal{D}$ is simple and explained next. First, let us remark that since the circuit makes non-adaptive oracle queries, for a fixed $x \in \{0,1\}^k$ the output of $D^h(x)$ is a function $D_x : \{0,1\}^q \to \{0,1\}$ of q evaluations of h at fixed points $y_1, y_2, \ldots, y_q \in \{0,1\}^n$ (again, the y_i's depend on x only): $D^h(x) = D_x(h(y_1), \ldots, h(y_q))$. Let us formally state this key definition.

DEFINITION 1.5. *Let $D^h(x)$ be an oracle circuit that makes q non-adaptive queries to its oracle. For a fixed input x we denote by $D_x : \{0,1\}^q \to \{0,1\}$ the function that maps the q oracle answers to the output $D^h(x) \in \{0,1\}$.*

We are going to show that having access to the above functions $D_x : \{0,1\}^q \to \{0,1\}$ for a few distinct $D \in \mathcal{D}$ and $x \in \{0,1\}^k$ is sufficient to compute majority.

[1] See Theorem 5.20 in the full version of [21].

[2] The proof in Impagliazzo's paper gives $q = O\left(\log(1/\epsilon\delta)/\epsilon^2\right)$ (when using the min-max proof for the hard-core theorem). However, a more efficient version (in terms of queries) of the hard-core theorem is given in [37], and using it one can push the number of queries to $q = O(\log(1/\delta)/\epsilon^2)$.

[3] We remark that the proof that appears in [20] does not directly achieve this. However, several researchers have independently observed that this is possible via a simple modification. We also mention that an unpublished manuscript [53] gives an alternative proof of the direct-product lemma that is also implementable by constant-depth circuits.

THEOREM 1.6 (DECODING REQUIRES MAJORITY). *There is a universal constant $C > 1$ such that the following holds. Let $(Amp, \mathcal{D})$ be a q-query non-adaptive $(1/2-\gamma) \to (1/2-\epsilon)$ black-box hardness amplification. Suppose that $q, \log|\mathcal{D}|, 1/\gamma \leq 2^{k/C}$, and $n, k \geq C^2$, and that $\gamma \geq 1/\log(1/\epsilon)$.*

Then there is a circuit of depth C and size $(q/\epsilon)^C$ with oracle access to (at most $(q/\epsilon)^C$ of) the functions $\{D_x : \{0,1\}^q \to \{0,1\}\}_{D\in\mathcal{D}, x\in\{0,1\}^k}$ that computes majority on inputs of length $1/\epsilon$.

To understand the above theorem, let us briefly see how to obtain Theorem 1.3 from it. Suppose that $\mathcal{D}$ consists of circuits of size s and depth d, that $1/2-\gamma = 1/3$, and that $s \leq 2^{\gamma\cdot k}$ for a suitable universal constant γ. First, we verify that the hypothesis of Theorem 1.6 is satisfied. This is because the circuits in $\mathcal{D}$ make at most $q \leq s \leq 2^{\gamma\cdot k} \leq 2^{k/C}$ queries – where the last inequality holds for a small enough γ – and $|\mathcal{D}| \leq 2^{s^{O(1)}}$ which implies $\log|\mathcal{D}| \leq s^{O(1)} \leq 2^{k/C}$ – where again the last inequality holds for a small enough γ. At this point, observe that the functions $D_x : \{0,1\}^q \to \{0,1\}$ are also computable by circuits of size s and depth d. Substituting these circuits for the oracle gates in the circuit of depth C and size $(q/\epsilon)^C$ given by the above theorem, we obtain a circuit of depth $C \cdot d = O(d)$ and size $(q/\epsilon)^C \cdot s = \mathrm{poly}(s/\epsilon)$ that computes the majority function on inputs of length $1/\epsilon$. As we mentioned earlier, by known lower bounds for the majority function [27, 49, 54] we obtain Theorem 1.3: $s \geq \exp\left((1/\epsilon)^{\Omega(1/d)}\right)$.

1.4 Related work

The inapplicability of hardness amplification techniques against low-complexity classes seems to have been observed independently by several researchers, and is also pointed out in [1] and in [60, Section 10]. The latter paper informally conjectures the main result of this work that proving hardness amplification requires computing majority (Theorem 1.6). A preliminary version of this work [62, Chapter 6] proves the conjecture in the special case where the class $\mathcal{D}$ in Definition 1.2 is small. The main result in this paper addresses for the first time the general case when there is no bound on the size of $\mathcal{D}$. The same preliminary version [62, Chapter 6] also proved a qualitatively weaker lower bound on the number of queries. We note that a recent work by Lu et al. [44] addresses the necessity of both majority and many queries in proofs of Impagliazzo's hard-core set theorem [32]. Specifically, [44] introduces two notions of black-box proof of the hard-core set theorem, and shows that one proof cannot be implemented by small constant-depth circuits, and that the other requires many oracle queries. Their arguments only apply to proofs of the hard-core set theorem, whereas our work addresses arbitrary black-box hardness amplification.

We remark that there is a variety of features that it is interesting to study and optimize of q-query $\delta \to (1/2-\epsilon)$ hardness amplification $(Amp, \mathcal{D})$ with input lengths k, n. We discuss the most relevant ones next.

Optimizing the ratio between k and n: E.g. [7, 32, 35, 55]. This is in particular relevant to obtain conclusions such as $P = BPP$ under the assumption that E requires exponential-size circuits [35].

Optimizing $|\mathcal{D}|$ = advice = list size: [36, 55, 59, 56, 58, 33, 34]. This is in particular relevant when $\mathcal{D}$ is a class of uniform machines (as opposed to circuits).

Optimizing the number of queries q: [32, 37], as well as the literature on locally-decodable codes (see, e.g., [57]). As discussed in Section 1.1, this is particularly relevant to the loss in circuit size incurred by hardness amplification.

The complexity of Amp: [47, 56, 60, 61, 58, 29, 39, 43, 41, 42] This line of research is orthogonal to this paper which studies the complexity of $\mathcal{D}$ and does not place any restriction on *Amp*. For context, we mention that the complexity of *Amp* is a key issue when we want to guarantee that the amplified function $Amp(f)$ lies in a specific class whenever the starting function f does. An example of this is the line of work on hardness amplification within *NP* [47, 56, 29, 58, 43] which started with the remarkable result by O'Donnell [47].

Relaxed definitions of hardness amplification: There are other works that study different, less demanding models of hardness amplification which are tailored to important questions such as worst-case to average-case connections within *NP* [10, 9, 61]. These works are incomparable with ours, one key difference being that they impose computational restrictions on the starting function f and the amplified function $Amp(f)$, whereas our results do not.

Finally, we would like to mention that there is a long line of research that is devoted to proving average-case hardness results for circuit classes below TC^0, e.g. [27, 25, 38, 4, 11, 23, 64]. With a few exceptions (discussed below) this research has been independent of hardness amplification, and our results may be interpreted as a partial explanation for this independence. The work by Klivans [38] stands out. Exploiting precisely the fact that computing majority is all that is needed for hardness amplification, Klivans uses a lower bound for constant-depth circuits with *one* majority gate [5] to give an alternative proof of the strong average-case hardness of parity for constant-depth circuits *without* majority gates. We remark that [38] does not contradict the results in this paper, but rather matches them by showing that a lower bound for a class with majority gates is sufficient for hardness amplification; see the full version of this paper for more on the status of lower bounds for constant-depth circuits with few (e.g. one) majority gates.

2. OVERVIEW OF THE PROOF

In this section we give a high level overview of the ideas that come into the proofs of our main results (Theorems 1.6,1.4). Within this section we allow ourselves to oversimplify and ignore some technicalities; the reader is referred to the formal proofs for precise details.

The Zoom Theorem.

Both the result about the necessity of majority (Theorem 1.6) and our lower bound on the number of queries (Theorem 1.4) rely on a theorem that we call "the Zoom Theorem." Let us first recall the setup. We are given a non-adaptive q-query $\delta \to (1/2-\epsilon)$ black-box hardness amplification $(Amp, \mathcal{D})$ where *Amp* maps functions $f : \{0,1\}^k \to \{0,1\}$ into functions $Amp(f) : \{0,1\}^n \to \{0,1\}$ (think $n = k^{O(1)}$). Recall that $\mathcal{D}$ is a class of oracle circuits and that for any circuit $D \in \mathcal{D}$ and input $x \in \{0,1\}^k$, Definition 1.5 defines a function $D_x : \{0,1\}^q \to \{0,1\}$ which captures the way D uses the answer to its q oracle queries to compute its output.

An informal statement of the Zoom Theorem follows (see Theorem 3.2 for a precise statement).

INFORMAL THEOREM 2.1 (ZOOM THEOREM). *There exists a circuit $D \in \mathcal{D}$ and an input $x \in \{0,1\}^k$ such that there is a function $T : \{0,1\}^q \to \{0,1\}$ of roughly the same complexity as D_x that satisfies:*

1. $\Pr[T(N^1_{1/2}, \ldots, N^q_{1/2}) = 1] \geq 0.49$, *where $(N^1_{1/2}, \ldots, N^q_{1/2})$ is a vector of q independent bits with probability of being 1 equal to $1/2$ (i.e., the vector is uniform in $\{0,1\}^q$).*
2. $\Pr[T(N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon}) = 1] \leq 2\delta$, *where $(N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon})$ is a vector of q independent bits with probability of being 1 equal to $1/2 - \epsilon$.*

We refer to the distributions $(N^1_{1/2}, \ldots, N^q_{1/2})$ and $(N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon})$ above as "uniform noise" and "bounded noise," respectively. Loosely speaking, the theorem says that T (which has the same complexity as circuits in $\mathcal{D}$) distinguishes between uniform noise and bounded noise.

Usefulness of the Zoom Theorem.

Our two main results follow from the Zoom Theorem. On an intuitive level, it seems that the natural way to decide whether a string $w \in \{0,1\}^q$ was chosen according to uniform noise or according to bounded noise is to compute the Hamming weight of w (which we denote by $weight(w)$) and decide according to whether $weight(w) \leq (1/2 - \epsilon/2)q$. Note that if T implements this strategy then it can indeed be used to compute majority. Furthermore note that when implementing this strategy, a Chernoff-style bound shows that $q = O(\log(1/\delta)/\epsilon^2)$ independent variables are sufficient in order to distinguish uniform noise from bounded noise (at rate $1/2-\epsilon$) with confidence $1-\delta$. Our bound on the number of queries essentially follows from the fact that this bound on q is tight.

Let us be more precise in explaining how the "necessity of majority" Theorem 1.6 follows from the Zoom Theorem. We would like to argue that T can be used to compute majority on inputs z of length $\ell := 1/\epsilon$. For simplicity, we explain how to use T to accomplish a slightly easier task, namely distinguishing between inputs z with $weight(z) = \ell/2$ and inputs z with $weight(z) = \ell/2 - 1$ (in the formal proof we essentially show that computing majority can be reduced to this simpler task). Given an input $z \in \{0,1\}^\ell$ we generate a string $w \in \{0,1\}^q$ where w_i is obtained by picking a random index $j \in [\ell]$ and setting $w_i = z_j$. In words, each bit in w is filled with a bit from a random position in z. Note that if $weight(z) = \ell/2$ then w is distributed like uniform noise, whereas if $weight(z) = \ell/2 - 1$ then w is distributed like bounded noise, because $weight(z)/\ell = 1/2 - 1/\ell = 1/2 - \epsilon$. It follows that we can use T to distinguish between the two cases (and recall that T has roughly the same complexity as circuits in $\mathcal{D}$).

This key idea was communicated to us by Madhu Sudan.

Finally, we point out that although the above reduction is randomized, at the end we obtain a *deterministic* circuit that computes majority. For this we also exploit that the relevant probabilities in the above reduction are sufficiently bounded away that they can be amplified using circuits of constant-depth by the result [2] (see also [3, 63]).

2.1 Proving the Zoom Theorem when $\mathcal{D}$ contains a single circuit

The proof of the Zoom Theorem is the main technical contribution of this paper. What makes this problem challenging is that the class $\mathcal{D}$ can be very large (e.g. $|\mathcal{D}| = \exp(k)$). We explain how we handle such large $\mathcal{D}$ later on. As a warm-up, we outline of the argument in the case that $\mathcal{D}$ contains only one circuit D. We consider a probability space with four independent random variables:

- A uniformly chosen function $F : \{0,1\}^k \to \{0,1\}$. We think of F as the original hard function.
- An input $X \in \{0,1\}^k$ that is uniformly distributed. We think of X as a random input to F.
- A uniformly chosen function $UN : \{0,1\}^n \to \{0,1\}$. We refer to UN as "uniform noise function."
- A function $BN : \{0,1\}^n \to \{0,1\}$ where for every $y \in \{0,1\}^n$, $BN(y)$ is an independent bit with probability of being 1 equal to $1/2-\epsilon$. We refer to BN as "bounded noise function."

We first consider the setting in which D is run with oracle $Amp(F) \oplus UN$. (This is an oracle that on input $y \in \{0,1\}^n$ returns $Amp(F)(y) \oplus UN(y)$). Note that the uniform noise function UN "masks out" the values of $Amp(F)$ and therefore the circuit D receives no information about F. Thus, D cannot possibly compute a function that is correlated with F:

$$\Pr\left[D^{Amp(F)\oplus UN}(X) \neq F(X)\right] = \Pr\left[D^{UN}(X) \neq F(X)\right] \geq 0.49. \quad (1)$$

We also consider the setting in which D is run with oracle $Amp(F) \oplus BN$. (This is an oracle that on input $y \in \{0,1\}^n$ returns $Amp(F)(y) \oplus BN(y)$). Since BN corresponds to bounded noise at rate $1/2-\epsilon$, we have that this oracle agrees with $Amp(F)$ on a $(1/2+\epsilon)$ fraction of inputs and therefore, by the definition of black-box hardness amplification:

$$\Pr\left[D^{Amp(F)\oplus BN}(X) \neq F(X)\right] \leq \delta. \quad (2)$$

Intuitively, the inequalities (1), (2) are going to translate into the two items of the Zoom Theorem. We now explain this part of the argument. Let us examine the computation of D on an input $x \in \{0,1\}^k$ with the two different oracles: In both cases D prepares the same q queries $y_1, \ldots, y_q \in \{0,1\}^n$ to the oracle and receives answers $a_1, \ldots, a_q$ from the oracle. It then outputs $D_x(a_1, \ldots, a_q)$. The high level idea is that when run on random $X \in \{0,1\}^k$, D_X distinguishes between the two oracles and therefore distinguishes between uniform noise and bounded noise. More precisely, by an averaging argument we can fix the random variables F and X and obtain a fixed function T that essentially equals D_X and distinguishes between bounded noise and uniform noise.

2.2 Extending the argument to the case when $\mathcal{D}$ is large

We would like to imitate the proof above when the class $\mathcal{D}$ contains many circuits. For concreteness let us assume that $\mathcal{D}$ contains 2^{k^2} circuits, i.e. $|\mathcal{D}| = \exp\left(k^2\right)$. In this general case the definition of black-box hardness amplification only says that for any choice of f, h where h agrees with $Amp(f)$

on a $(1/2+\epsilon)$ fraction of inputs *there exists* a circuit $D \in \mathcal{D}$ such that D^h agrees with f on a $1-\delta$ fraction of inputs. Note that the circuit D is a function of both f and h, and let us denote this function by $circuit(f,h)$.

We would like to imitate the previous argument. However, when we use oracle $Amp(F) \oplus BN$, we do not know which circuit $D \in \mathcal{D}$ is the "correct circuit", i.e. $circuit(F, Amp(F) \oplus BN)$. More formally, we have that $circuit(F, Amp(F) \oplus BN)$ is a random variable that in particular depends on BN. In the previous argument we applied a *fixed* function D_x on the answers $a_1, \ldots, a_q$ that were returned by the oracle. However, the function D_x for $D = circuit(F, Amp(F) \oplus BN)$ that we want to apply on $a_1, \ldots, a_q$ is now a random variable that *depends* on $a_1, \ldots, a_q$ and we cannot use the argument above.

Going back to the case of a single circuit.

To avoid the aforementioned problem we start by fixing the random variable $circuit(F, Amp(F) \oplus BN)$ to its most likely value. That is, let D be the most likely value of $circuit(F, Amp(F) \oplus BN)$ and let $E = E(F, BN)$ be the event

$$E := \{circuit(F, Amp(F) \oplus BN) = D\}.$$

Note that the probability of E is at least $1/|\mathcal{D}| = 2^{-k^2}$ (which is small but not *too* small). We have that $circuit(F, Amp(F) \oplus BN)$ is *fixed* in E (which means that in E we only need to consider *one fixed circuit* D). From now on we restrict our attention to E. That is, let F', BN' denote the distribution of F, BN when conditioned on the event E. Note that this conditioning can skew the distribution of F', BN' and that these variables are no longer distributed like the original variables F, BN and in particular may become dependent. For the purpose of explaining the argument let us assume the unjustified assumption that F' and BN' are independent. (In the actual argument we bypass this problem by fixing F to some fixed function f before conditioning on the event E).

We would like to imitate the argument of the previous section in this new probability space. Indeed, we are back to dealing with one fixed circuit D. However, the previous argument critically relies on properties of BN: Most notably that for any $y_1, \ldots, y_q \in \{0,1\}^n$, the random variable $(BN(y_1), \ldots, BN(y_q))$ is distributed like bounded noise. This may not necessarily hold for BN'.

An information-theoretic lemma.

In order to handle this problem, we use the following Lemma (stated informally; cf. the full version of this work for a precise statement).

INFORMAL LEMMA 2.2. *Let $V_1, \ldots, V_t$ be independent and identically distributed random variables. Let E be an event whose probability is "not too small." Then for any integer q there exists a "large" set $G \subseteq [t]$ such that for every $i_1, \ldots, i_q \in G$, the distribution $(V_{i_1}, \ldots, V_{i_q})$ "does not change significantly" when conditioning on E.*

This lemma can be viewed as a generalization of a Lemma by Raz (in which $q = 1$) that is used in his parallel repetition theorem [48]. We have recently found out that this lemma follows easily from the results in [14, Section 4].

We apply the lemma to the random variables $\{BN(y)\}_{y \in \{0,1\}^n}$. We conclude that there exists a large set $G \subseteq \{0,1\}^n$ such that for any $y_1, \ldots, y_q \in G$ the variable

$$(BN'(y_1), \ldots, BN'(y_q))$$

is statistically close to

$$(BN(y_1), \ldots, BN(y_q)).$$

This lemma intuitively helps us recover the previous argument in the new probability space: We consider the operation of $D^{Amp(F') \oplus BN'}$ on an input $x \in \{0,1\}^k$. If the queries $y_1, \ldots, y_q \in \{0,1\}^n$ that D makes are all in the "good set" G, then the rest of the proof essentially goes through. This is because on these q queries the distribution of the bounded noise function is statistically close to its initial distribution and we can continue with the previous argument.

However, even though the set G of "good queries" is large, it may be the case that on every input $x \in \{0,1\}^k$, D makes a "bad query" $y' \notin G$. We have no control on the distribution $BN'(y')$ when $y' \notin G$; for example, it may be correlated with the value of BN' on another query y, and so we cannot relate this distribution to that of bounded noise (in which different coordinates are independent and distributed in the same way).

Fixing bad queries.

In order to address this issue we further refine the probability space by fixing the value of BN' at some bad queries. The high level idea is that by fixing the bounded noise function on these queries we "remove dependencies" between the answers that the circuit D sees when making its queries. This part of the argument is more technical and we will not describe it in detail. However, we point out that fixing bad queries is a tricky business as whenever we fix a bad query we change the probability space, which in turn skews the distribution of the bounded noise function and may result in introducing new bad queries (and it seems that we make no progress as we can never fully get rid of bad queries). In the actual argument we fix the bounded noise function only on those queries that are *heavy* in the sense that they are "asked frequently" by D. The rationale is that even if fixing the bounded noise function on these queries skews the distribution and introduces new bad queries we do make progress as the new bad queries are queries that are not asked frequently by D. Finally, we argue that bad queries that are not asked frequently by D do not hurt us too much when implementing the initial argument (because on an intuitive level, this means that D asks good queries "most of the time").

One technical point that we want to make is that for implementing the approach above we must make sure that the number of bad queries that are introduced after fixing the frequent queries does not depend on the number of frequent queries that we fix. This is because in the actual argument we do a union bound over all bad queries and argue that the probability that a random input queries *any* bad query (that is not already fixed) is low. This allows us to ignore bad queries as the weight of inputs which query bad queries is small.

3. THE ZOOM THEOREM

Both our result about the necessity of majority (Theorem 1.6) and our lower bound on the number of queries (Theorem

1.4) rely on a theorem which we call the "Zoom Theorem" and is our main technical contribution. In this section we state this theorem formally; we refer the reader to the full version of this paper for a proof of the zoom theorem and for the proofs of our main theorems from the zoom theorem. The Zoom Theorem shows that given a non-adaptive q-query black-box $\delta \to (1/2-\epsilon)$ hardness amplification $(Amp, \mathcal{D})$ we can "zoom in" on a particular function $D_x : \{0,1\}^q \to \{0,1\}$, where $D \in \mathcal{D}, x \in \{0,1\}^k$ (cf. Definition 1.5 for the definition of D_x) that is distinguishing noise rate $1/2$ from noise rate $1/2 - \epsilon$. The distinguisher will not quite be a function D_x but rather (a distribution on) *projections* of such functions, which are simply functions that can be obtained from D_x by fixing some input variables to constants and complementing others. We give the formal definition of a projection and then we state the zoom theorem.

DEFINITION 3.1. *Let $d = d(y_1, \ldots, y_q) : \{0,1\}^q \to \{0,1\}$ be a function. A* projection *of d is a function $d' : \{0,1\}^q \to \{0,1\}$ that can be obtained from d by fixing some input variables to constants and complementing others, and possibly complementing the output. Formally, there are $a_1, \ldots, a_q, b_1, \ldots, b_q, c \in \{0,1\}$ such that for any $y_1, \ldots, y_q \in \{0,1\}$, $d'(y_1, \ldots, y_q) = d((y_1 \cdot a_1) \oplus b_1, \ldots, (y_q \cdot a_q) \oplus b_q) \oplus c$.*

THEOREM 3.2 (ZOOM THEOREM). *There is a universal constant $C > 1$ such that the following holds. Let $(Amp, \mathcal{D})$ be a non-adaptive q-query $\delta \to (1/2 - \epsilon)$ hardness amplification scheme. Suppose that $q, \log|\mathcal{D}| \leq 2^{k/C}$, and $n, k \geq C^2$.*

Then there is a distribution T on functions $t : \{0,1\}^q \to \{0,1\}$ such that

1. $\Pr_{T, N^1_{1/2}, \ldots, N^q_{1/2}}[T(N^1_{1/2}, \ldots, N^q_{1/2}) = 1] \geq 1/2 - 2^{-k/C}$, *where $(N^1_{1/2}, \ldots, N^q_{1/2})$ is a vector of q independent bits with probability of being 1 equal to $1/2$ (i.e., the vector is uniform in $\{0,1\}^q$),*

2. $\Pr_{T, N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon}}[T(N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon}) = 1] \leq \delta + 2^{-k/C}$, *where $(N^1_{1/2-\epsilon}, \ldots, N^q_{1/2-\epsilon})$ is a vector of q independent bits with probability of being 1 equal to $1/2 - \epsilon$, and*

3. *each $t \in T$ is a projection of a function D_x for some $D \in \mathcal{D}$ and $x \in \{0,1\}^k$. I.e., every $t \in T$ can be obtained from D_x for some $D \in \mathcal{D}, x \in \{0,1\}^k$ by fixing some input variables to constants and complementing others, and possibly complementing the output.*

4. OPEN PROBLEMS

One weakness of our result is that we can only handle black-box hardness amplification which use *nonadaptive* circuits. While to the best of our knowledge most known black-box hardness amplification results use nonadaptive circuits, it is an interesting open problem to extend the results in this work to the case of *adaptive* circuits. We remark that, for some specific functions Amp, the techniques in this work already give some results on adaptive circuits when the amount of non-uniformity $|\mathcal{D}|$ of the black-box hardness amplification is small (e.g., $|\mathcal{D}| = \text{poly}(1/\epsilon)$). In particular, one can show that achieving the parameters of the hardness amplification in [19] (based on the Hadamard code) or the parameters of the hardness amplification in [55] (based on Reed-Muller codes), requires computing majority. The details of these results appear in [62, Chapter 6].

Another problem that deserves more investigation is whether something similar to our results can be said about pseudorandom generator constructions. For example, is computing majority necessary for a black-box construction of a pseudorandom generator with constant error from a (1/3)-hard function?

Acknowledgements.

We would like to thank Madhu Sudan, Salil Vadhan, and Avi Wigderson for many helpful conversations and comments. Ronen would like to thank Avi Wigderson for introducing him to the subject of hardness amplification and for many discussions. Emanuele would like to thank Salil Vadhan for many helpful discussions during the preliminary stages of this work [62, Chapter 6], and Madhu Sudan also for helpful discussions and especially for a key suggestion (cf. Section 2 and [62, Section 6.2]). We thank Russell Impagliazzo for pointing out [14] to us. We also thank Hoeteck Wee for pointing out [13] to us.

5. REFERENCES

[1] M. Agrawal. Hard sets and pseudo-random generators for constant depth circuits. In *Twenty First Foundations of Software Technology and Theoretical Computer Science, December 13-15, Bangalore, India*, pages 58–69. Springer-Verlag, 2001.

[2] M. Ajtai. Σ^1_1-formulae on finite structures. *Ann. Pure Appl. Logic*, 24(1):1–48, 1983.

[3] M. Ajtai. Approximate counting with uniform constant-depth circuits. In *Advances in computational complexity theory (New Brunswick, NJ, 1990)*, pages 1–20. Amer. Math. Soc., Providence, RI, 1993.

[4] N. Alon and R. Beigel. Lower bounds for approximations by low degree polynomials over z_m. In *Proceedings of the Sixteenth Annual Conference on Computational Complexity*, pages 184–187. IEEE, June 18–21 2001.

[5] J. Aspnes, R. Beigel, M. Furst, and S. Rudich. The expressive power of voting polynomials. *Combinatorica*, 14(2):135–148, 1994.

[6] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Comput. Complexity*, 1(1):3–40, 1991.

[7] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3(4):307–318, 1993.

[8] D. Beaver and J. Feigenbaum. Hiding instances in multioracle queries. In *7th Annual Symposium on Theoretical Aspects of Computer Science*, volume 415 of *Lecture Notes in Computer Science*, pages 37–48, Rouen, France, 22–24 Feb. 1990. Springer.

[9] A. Bogdanov and L. Trevisan. Average-case complexity. *Foundations and Trendső in Theoretical Computer Science*, 2(1), 2006.

[10] A. Bogdanov and L. Trevisan. On worst-case to average-case reductions for np problems. *SIAM J. Comput.*, 36(4):1119–1159, 2006.

[11] J. Bourgain. Estimation of certain exponential sums arising in complexity theory. *C. R. Math. Acad. Sci.*

Paris, 340(9):627–631, 2005.

[12] J.-Y. Cai, A. Pavan, and D. Sivakumar. On the hardness of the permanent. In *16th International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Volume 1563, pages 90–99, Trier, Germany, 1999. Springer-Verlag.

[13] R. Canetti, G. Even, and O. Goldreich. Lower bounds for sampling algorithms for estimating the average. *Information Processing Letters*, 53(1):17–25, 1995.

[14] J. Edmonds, R. Impagliazzo, S. Rudich, and J. Sgall. Communication complexity towards lower bounds on circuit depth. *Computational Complexity*, 10(3):210–246, 2001.

[15] U. Feige and C. Lund. On the hardness of computing the permanent of random matrices. *Computational Complexity*, 6(2):101–132, 1996.

[16] M. L. Furst, J. B. Saxe, and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, April 1984.

[17] O. Goldreich. *Modern cryptography, probabilistic proofs and pseudorandomness*, volume 17 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999.

[18] O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, Cambridge, 2001.

[19] O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, Washington, 15–17 May 1989.

[20] O. Goldreich, N. Nisan, and A. Wigderson. On Yao's XOR lemma. Technical Report TR95–050, Electronic Colloquium on Computational Complexity, March 1995. `http://www.eccc.uni-trier.de/eccc`.

[21] S. Goldwasser, D. Gutfreund, A. Healy, T. Kaufman, and G. N. Rothblum. Verifying and decoding in constant depth. In *STOC*, pages 440–449, 2007.

[22] P. Gopalan and V. Guruswami. Hardness amplification within np against deterministic algorithms. In *Proceedings of the 23nd Annual Conference on Computational Complexity*. IEEE, June 23–26 2008.

[23] F. Green, A. Roy, and H. Straubing. Bounds on an exponential sum arising in Boolean circuit complexity. *C. R. Math. Acad. Sci. Paris*, 341(5):279–282, 2005.

[24] V. Guruswami and V. Kabanets. Hardness amplification via space-efficient direct products. In J. R. Correa, A. Hevia, and M. A. Kiwi, editors, *LATIN*, volume 3887 of *Lecture Notes in Computer Science*, pages 556–568. Springer, 2006.

[25] A. Hajnal, W. Maass, P. Pudlák, M. Szegedy, and G. Turán. Threshold circuits of bounded depth. *J. Comput. System Sci.*, 46(2):129–154, 1993.

[26] K. A. Hansen and P. B. Miltersen. Some meet-in-the-middle circuit lower bounds. In *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, Lecture Notes in Computer Science, Volume 3153, pages 334 – 345, August 22–27 2004.

[27] J. Håstad. *Computational limitations of small-depth circuits*. MIT Press, 1987.

[28] J. Håstad and M. Goldmann. On the power of small-depth threshold circuits. *Comput. Complexity*, 1(2):113–129, 1991.

[29] A. Healy, S. P. Vadhan, and E. Viola. Using nondeterminism to amplify hardness. *SIAM J. Comput.*, 35(4):903–931, 2006.

[30] IEEE. *Proceedings of the 20th Annual Conference on Computational Complexity*, June 12–15 2005.

[31] IEEE. *Proceedings of the 22nd Annual Conference on Computational Complexity*, June 13–16 2007.

[32] R. Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th Annual Symposium on Foundations of Computer Science*, pages 538–545, Milwaukee, Wisconsin, 23–25 Oct. 1995. IEEE.

[33] R. Impagliazzo, R. Jaiswal, and V. Kabanets. Approximately list-decoding direct product codes and uniform hardness amplification. In *FOCS*, pages 187–196. IEEE Computer Society, 2006.

[34] R. Impagliazzo, R. Jaiswal, V. Kabanets, and A. Wigderson. Uniform direct-product theorems: Simplified, optimized, and derandomized. In *Proceedings of the 40th Annual ACM Symposium on the Theory of Computing (STOC)*, Victoria, Canada, 17–20 May 2008.

[35] R. Impagliazzo and A. Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 220–229, El Paso, Texas, 4–6 May 1997.

[36] R. Impagliazzo and A. Wigderson. Randomness vs time: derandomization under a uniform assumption. *J. Comput. System Sci.*, 63(4):672–688, 2001. Special issue on FOCS 98.

[37] A. Klivans and R. A. Servedio. Boosting and hard-core sets. *Machine Learning*, 53(3):217–238, 2003.

[38] A. R. Klivans. On the derandomization of constant depth circuits. In *Proceedings of the Fifth International Workshop on Randomization and Approximation Techniques in Computer Science*, August 18–20 2001.

[39] H. Lin, L. Trevisan, and H. Wee. On hardness amplification of one-way functions. In J. Kilian, editor, *TCC*, volume 3378 of *Lecture Notes in Computer Science*, pages 34–49. Springer, 2005.

[40] R. Lipton. New directions in testing. In *Proceedings of DIMACS Workshop on Distributed Computing and Cryptography*, volume 2, pages 191–202. ACM/AMS, 1991.

[41] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. On the complexity of hardness amplification. In *Proceedings of the 20th Annual Conference on Computational Complexity* [30], pages 170–182.

[42] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. Impossibility results on weakly black-box hardness amplification. In E. Csuhaj-Varjú and Z. Ésik, editors, *FCT*, volume 4639 of *Lecture Notes in Computer Science*, pages 400–411. Springer, 2007.

[43] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. Improved hardness amplification in np. *Theor. Comput. Sci.*, 370(1-3):293–298, 2007.

[44] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. On the complexity

of hard-core set constructions. In L. Arge, C. Cachin, T. Jurdzinski, and A. Tarlecki, editors, *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 183–194. Springer, 2007.

[45] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. ACM*, 51(2):231–262, 2004.

[46] N. Nisan and A. Wigderson. Hardness vs randomness. *J. Computer & Systems Sciences*, 49(2):149–167, Oct. 1994.

[47] R. O'Donnell. Hardness amplification within NP. *J. Comput. Syst. Sci.*, 69(1):68–94, Aug. 2004.

[48] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803 (electronic), 1998.

[49] A. A. Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Mat. Zametki*, 41(4):598–607, 623, 1987.

[50] A. A. Razborov and S. Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, Aug. 1997.

[51] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudorandom generator. *J. ACM*, 52(2):172–216, 2005.

[52] R. Shaltiel and C. Umans. Pseudorandomness for approximate counting and sampling. *Computational Complexity*, 15(4):298–341, 2006.

[53] R. Shaltiel, E. Viola, and A. Wigderson. Unpublished manuscript, 2005.

[54] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, pages 77–82, New York City, 25–27 May 1987.

[55] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. System Sci.*, 62(2):236–266, 2001. Special issue on the Fourteenth Annual IEEE Conference on Computational Complexity (Atlanta, GA, 1999).

[56] L. Trevisan. List decoding using the XOR lemma. In *44th Annual Symposium on Foundations of Computer Science*, pages 126–135, Cambridge, Massachusetts, 11–14 Oct. 2003. IEEE.

[57] L. Trevisan. Some applications of coding theory in computational complexity. In *Complexity of computations and proofs*, volume 13 of *Quad. Mat.*, pages 347–424. Dept. Math., Seconda Univ. Napoli, Caserta, 2004.

[58] L. Trevisan. On uniform amplification of hardness in np. In H. N. Gabow and R. Fagin, editors, *STOC*, pages 31–38. ACM, 2005.

[59] L. Trevisan and S. Vadhan. Pseudorandomness and average-case complexity via uniform reductions. *Comput. Complex.*, 16(4):331–364, 2007.

[60] E. Viola. The complexity of constructing pseudorandom generators from hard functions. *Comput. Complexity*, 13(3-4):147–188, 2004.

[61] E. Viola. On constructing parallel pseudorandom generators from one-way functions. In *Proceedings of the 20th Annual Conference on Computational Complexity* [30], pages 183–197.

[62] E. Viola. *The Complexity of Hardness Amplification and Derandomization*. PhD thesis, Harvard University, 2006. `http://www.eccc.uni-trier.de/eccc`.

[63] E. Viola. On approximate majority and probabilistic time. In *Proceedings of the 22nd Annual Conference on Computational Complexity* [31], pages 155–168.

[64] E. Viola and A. Wigderson. Norms, xor lemmas, and lower bounds for GF(2) polynomials and multiparty protocols. In *Proceedings of the 22nd Annual Conference on Computational Complexity* [31].

[65] A. C. Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, 3–5 Nov. 1982. IEEE.

[66] A. C.-C. Yao. Separating the polynomial-time hierarchy by oracles. In *Proc. 26th annual symposium on Foundations of computer science*, pages 1–10. IEEE Press, 1985.