Electronic Colloquium on Computational Complexity, Report No. 187 (2017)

A Note on the Limitations of Two Black-Box Techniques in Quantified Derandomization

Roei Tell *

December 3, 2017

Abstract

The quantified derandomization problem of a circuit class $\mathcal{C}$ with a function $B : \mathbb{N} \to \mathbb{N}$ is the following: Given an input circuit $C \in \mathcal{C}$ over n bits, deterministically distinguish between the case that C accepts all but $B(n)$ of its inputs and the case that C rejects all but $B(n)$ of its inputs. This generalizes the standard derandomization problem, which is the special case where $B(n) = 2^n/3$.

A major goal of this framework is to serve as a stepping-stone on the way to standard derandomization. Specifically, we want to construct reductions of the standard derandomization problem of a class $\mathcal{C}$ to the quantified derandomization problem (e.g., using strong error-reduction), and to then construct an algorithm that solves the latter.

In this note we show that if both the reduction (from standard derandomization to quantified derandomization) and the algorithm (for quantified derandomization) are constructed using two specific natural techniques *that only rely on "black-box" access to the input circuit* C, then a naive combination of the two algorithms does not suffice to yield a standard derandomization of $\mathcal{C}$. That is, when using these two techniques, the parameter value $B(n)$ to which standard derandomization is reduced is necessarily larger than the value $B(n)$ that the quantified derandomization algorithm can handle.

1 Introduction

For a circuit class $\mathcal{C}$ and a function $B : \mathbb{N} \to \mathbb{N}$, the $(\mathcal{C}, B)$-quantified derandomization problem is the following: Given a circuit $C \in \mathcal{C}$ over n input bits, deterministically distinguish between the case that C accepts all but at most $B(n)$ of its inputs and the case that C rejects all but at most $B(n)$ of its inputs. This problem was introduced by Goldreich and Wigderson [GW14], and constitutes a generalization of the classical

*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel. Email: roei.tell@weizmann.ac.il

ISSN 1433-8092

problem of standard derandomization, which is obtained by setting $B(n) = 2^n/3$. We call $B(n)$ the number of exceptional inputs for the circuit C.

A major goal in the study of quantified derandomization is to use this framework in order to solve the standard derandomization problem. Specifically, fix a circuit class $\mathcal{C}$, and assume that we constructed an algorithm that reduces the standard derandomization problem of $\mathcal{C}$ to quantified derandomization of $\mathcal{C}$ with $B^{\mathtt{red}}$ exceptional inputs (e.g., such an algorithm can get as input a circuit $C \in \mathcal{C}$ over m bits with at most $2^m/3$ exceptional inputs, and construct a circuit $C' \in \mathcal{C}$ over $n = \mathrm{poly}(m)$ bits with at most $B^{\mathtt{red}}(n)$ exceptional inputs, such that the most frequent output of C' equals the most frequent output of C). Also assume that we constructed an algorithm for quantified derandomization of $\mathcal{C}$ that can handle $B^{\mathtt{alg}}$ exceptional inputs. Then, if $B^{\mathtt{alg}}(n) \geq B^{\mathtt{red}}(n)$, we can combine both algorithms in the straightforward way to obtain an algorithm for standard derandomization of $\mathcal{C}$.

In this note we show *a limitation of two specific natural techniques* in quantified derandomization. One technique is used to reduce standard derandomization to quantified derandomization, and is based on error-reduction using seeded extractors that are computable by circuits in $\mathcal{C}$. The other technique is used to construct quantified derandomization algorithms, and is based on pseudorandom distributions of restrictions that simplify every circuit $C \in \mathcal{C}$, with high probability. The two foregoing techniques, which are described in detail in Section 2, only rely on "black-box" access to the input circuit C; that is, the algorithm (either for the reduction or for quantified derandomization) does not use the explicit description of C, beyond the guarantee that $C \in \mathcal{C}$ and the ability to evaluate C at arbitrarily-chosen inputs.

Informally, our main theorem asserts that the straightforward combination of the two foregoing techniques cannot suffice to yield a standard derandomization algorithm. This is the case since the function $B^{\mathtt{red}}$ for quantified derandomization to which standard derandomization is reduced is *necessarily larger* than the function $B^{\mathtt{alg}}$ that the quantified derandomization algorithm can handle. That is,

Theorem 1 *(a limitation of two "black-box" techniques in quantified derandomization; informal). Assume that there exists a reduction of standard derandomization of a class $\mathcal{C}$ to quantified derandomization of $\mathcal{C}$ with $B^{\mathtt{red}}$ exceptional inputs that is based on seeded extractors that are computable in $\mathcal{C}$. Also assume that there exists a quantified derandomization algorithm for $\mathcal{C}$ with $B^{\mathtt{alg}}$ exceptional inputs that is based on a distribution over restrictions that simplifies every $C \in \mathcal{C}$ to a constant, with high probability. Then, $B^{\mathtt{alg}}(n) < B^{\mathtt{red}}(n)$.*

This result is particularly meaningful for derandomization of $\mathcal{AC}^0$. In this setting, we know of a reduction to quantified derandomization with $B^{\mathtt{red}}$ exceptional inputs and of a quantified derandomization algorithm with $B^{\mathtt{alg}}$ inputs such that $B^{\mathtt{red}}$ and $B^{\mathtt{alg}}$ are very close (i.e., both are of the form $2^{n/\mathrm{poly}\log(n)}$; see [Tel17a, Thms. 1 & 3]). However, the algorithms for both these results rely on the aforementioned "black-box" techniques, and therefore Theorem 1 implies that one cannot hope to obtain a standard derandomization algorithm for $\mathcal{AC}^0$ by merely improving the parameters of the underlying technical results. Nevertheless, for other circuit classes, results in

quantified derandomization were obtained using different techniques, *which are not "black-box"*; in particular, such results were obtained for various subclasses of $\mathcal{AC}^0[\oplus]$ and for sparse $\mathcal{TC}^0$ circuits (see Section 2 for more details).

Organization. In Section 2 we describe the two "black-box" techniques that are the focus of the current text. In Section 3 we prove the main theorem (i.e., Theorem 1). In Section 4 we discuss several relaxations of the hypotheses of the main theorem (i.e., several natural modifications of the two "black-box" techniques) that do not seem sufficient to bypass the conclusion of the theorem.

2 Two black-box techniques for quantified derandomization

Throughout the text, whenever we refer to a class $\mathcal{C}$ of circuits, we will always think of $\mathcal{C}$ as a *set of circuits*, rather than a collection of such sets. That is, we consider $\mathcal{C} = \bigcup_{n\in\mathbb{N}} \mathcal{C}_n$, where $\mathcal{C}_n$ is some fixed set of circuits over n input bits. This is done merely for simplicity of presentation, and all results extend to the standard setting (i.e., when considering circuit families) in a natural way.

2.1 Error-reduction using a $\mathcal{C}$-computable sampler

The first technique that we discuss is used to reduce standard derandomization of a circuit class $\mathcal{C}$ to quantified derandomization of $\mathcal{C}$. This technique is based on *randomness-efficient error reduction*, using seeded extractors (equivalently, averaging samplers; see, e.g., [Vad12, Cor. 6.24], [Gol08, Apdx. D.4.1.2]) that are computable by $\mathcal{C}$-circuits. That is,

Definition 2 *($\mathcal{C}$-computable sampler). For $B^{\mathtt{red}} : \mathbb{N} \to \mathbb{N}$, we say that a function $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ is a* Boolean sampler with $B^{\mathtt{red}}$ bad inputs *if it satisfies the following: For every $T \subseteq \{0,1\}^m$, for all but at most $B^{\mathtt{red}}$ of the inputs $x \in \{0,1\}^n$ it holds that $\Pr_{i\in\{0,1\}^s}[Samp(x,i) \in T] \in |T|/2^m \pm 1/10$. [1] For a circuit class $\mathcal{C}$, we say that $Samp$ is* computable in $\mathcal{C}$ *if for every fixed $i \in \{0,1\}^s$, each output bit of the function $Samp^{(i)}(x) = Samp(x,i)$ is computable by a circuit in $\mathcal{C}$.*

Samplers that are computable in a circuit class $\mathcal{C}$ can be used for error-reduction of $\mathcal{C}$-circuits. Specifically, assume that for every $m \in \mathbb{N}$ we can efficiently construct circuits for the output bits of a $\mathcal{C}$-computable sampler $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ with $B^{\mathtt{red}}$ bad inputs. Then, given a circuit $C \in \mathcal{C}$ over m input bits, we can efficiently construct a circuit $C' : \{0,1\}^n \to \{0,1\}$ that gets input $x \in \{0,1\}^n$, computes the 2^s outputs of the sampler on x (each output is an m-bit string), evaluates C on each of these outputs, and outputs the majority of the evaluations of C; that is, $C'(x) = MAJ\left(\{C(Samp(x,i))\}_{i\in\{0,1\}^s}\right)$. If C accepts (resp., rejects) all but at most $2^m/3$ of its inputs, then C' accepts (resp., rejects) all but at most $B^{\mathtt{red}}(n)$ of its inputs.

[1] Indeed, for simplicity, we fixed the accuracy parameter of the sampler to $1/10$.

We typically want to minimize the overhead of C' with respect to the original circuit C, since we reduce the standard derandomization problem of C to quantified derandomization of C'. Hence, typical settings of the parameters are $n = \text{poly}(m)$ and $s = O(\log(n))$, such that the size of C' is polynomial in the size of C. Also observe that the majority function in the definition of C' can be replaced by an "approximate majority" (i.e., a function that distinguishes between strings with relative Hamming weight ≤ 0.49 and strings with relative Hamming weight ≥ 0.51, as in [Ajt83, Vio09]). For further discussion of the effect of the "overhead" when constructing C', see Section 3.

Observe that the algorithm above (that constructs the circuit C') uses the same sampler $Samp$, regardless of the input circuit $C \in \mathcal{C}$. Thus, in order to evaluate C' at any point, the algorithm does not need the explicit description of C, but rather only "black-box" access to C (i.e., the ability to evaluate C at arbitrarily-chosen points).

This approach has been used to reduce standard derandomization to quantified derandomization in the contexts of $\mathcal{AC}^0$ circuits (see [Tel17a, Thm. 1], which uses the extractor of [CL16]), of $\mathcal{AC}^0[\oplus]$ circuits (see [GW14, Thm. 1.4] and [Tel17a, Thm. 5]), and of $\mathcal{TC}^0$ circuits (see [Tel17b, Thm. 1.2]). A somewhat different approach for error-reduction was taken in [GW14, Thm. 3.4 in the full version]; their approach is also insufficient to bypass the limitation in our main theorem, but is interesting in its own right (see further discussion in Section 4.2).

2.2 A quantified derandomization algorithm that uses pseudorandom restrictions

The second technique that we discuss can be used to construct a quantified derandomization algorithm. This technique relies on the existence of a distribution $\mathbf{S}_n$ over subsets of $\{0,1\}^n$ (i.e., over "restrictions") such that for any $C \in \mathcal{C}$ over n input bits, with high probability over $S \sim \mathbf{S}_n$ it holds that $C\restriction_S$ is constant. That is,

Definition 3 *(simplifier sets). For $B^{\mathtt{alg}} : \mathbb{N} \to \mathbb{N}$, we say that a distribution $\mathbf{S}_n$ over subsets of $\{0,1\}^n$ is* a distribution of simplifier sets of size more than $B^{\mathtt{alg}}$ for $\mathcal{C}$ *if the following conditions hold:*

1. *Every subset S in the support of $\mathbf{S}_n$ is of size $|S| > B^{\mathtt{alg}}(n)$.*
2. *For every $C \in \mathcal{C}$ over n input bits, $\Pr_{S \sim \mathbf{S}_n}[C\restriction_S \text{ is constant}] > 1/2$.*

To see why simplifier sets are useful for quantified derandomization, let $C \in \mathcal{C}$ be a circuit over n bits with $B^{\mathtt{alg}}(n)$ exceptional inputs. Then, with probability more than $1/2$ over $S \sim \mathbf{S}_n$ it holds that $C\restriction_S$ is a constant function; and since $B^{\mathtt{alg}}(n) < |S|$, whenever $C\restriction_S$ is a constant function, this constant equals the most frequent output of C. Now, assume that we can efficiently sample a succint representation of a set $S \sim \mathbf{S}_n$ using only a few (say, $O(\log(n))$) random bits, and that this representation allows to efficiently find some input $x \in S$. In this case, we can solve the quantified derandomization problem as follows: We enumerate the choices of $S \sim \mathbf{S}_n$, evaluate C on an (arbitrary) input in each choice of S, and output the majority value among

the evaluations of C. Note that the only information about C that the algorithm in this approach uses, other than the fact that $C \in \mathcal{C}$, is the ability to evaluate C on arbitrarily-chosen inputs.

This ("black-box") approach has been used to construct quantified derandomization algorithms for $\mathcal{AC}^0$ (see [GW14, Thm. 1.3] and [Tel17a, Thm. 3]) and for $\mathcal{TC}^0$ circuits of depth two (see [Tel17b, Thm. 7.2]). Nevertheless, for other circuit classes, *pseudorandom restriction algorithms that are "non-black-box"* have been constructed and used for quantified derandomization; these classes include various subclasses of $\mathcal{AC}^0[\oplus]$ (see [Tel17a, Thm. 6]) and sparse $\mathcal{TC}^0$ circuits (see [Tel17b, Thm. 1]).

3 Proof of the main theorem

Towards formally stating and proving Theorem 1, let us now carefully track how the combination of the two techniques from Section 2 works. We are interested in standard derandomization of a circuit class $\mathcal{C}$. That is, we are given a circuit $C \in \mathcal{C}$ over m input bits, and want to distinguish between the case that C accepts all but at most $2^m/3$ of its inputs and the case that C rejects all but at most $2^m/3$ of its inputs.

To do so, we fix some sampler $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ with $B^{\mathtt{red}}$ bad inputs, and consider the "error-reduced" circuit $C' : \{0,1\}^n \to \{0,1\}$ such that $C'(x) = MAJ\left(\{C(Samp(x,i))\}_{i \in \{0,1\}^s}\right)$. [2] We think of the circuit C' as belonging to some new circuit class, which we denote by $\widehat{\mathcal{C}}$; for example, we can define $\widehat{\mathcal{C}}$ to be the class of circuits of the form $C''(x) = \Phi\left(\{C_0(Samp(x,i)\}_{i \in \{0,1\}^s}\right)$, where $C_0 \in \mathcal{C}$ and Φ is one of several "simple composition" functions (the majority function being one of them). Indeed, it might be the case that $\widehat{\mathcal{C}} \subseteq \mathcal{C}$, if $\mathcal{C}$ is closed to the overhead involved in constructing circuits such as C', but we do not assume so. However, we assume that $Samp$ is computable in $\widehat{\mathcal{C}}$, which holds under reasonable definitions of $\widehat{\mathcal{C}}$. [3]

Now we are interested in quantified derandomization of the class $\widehat{\mathcal{C}}$ with $B^{\mathtt{red}}$ bad inputs, using simplifier sets. However, the following lemma asserts that in any distribution of simplifier sets of size more than $B^{\mathtt{alg}}$ for a class that can compute $Samp$ (and in particular for $\widehat{\mathcal{C}}$), the size of the sets satisfies $B^{\mathtt{alg}}(n) < B^{\mathtt{red}}(n)$. When reading the lemma's statement, we encourage the reader to think of a sampler $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ such that $s = O(\log(n))$ and $m = n^{\Omega(1)}$.

Lemma 4 *(simplifier sets and samplers). Let $\widehat{\mathcal{C}}$ be a circuit class, and let $B^{\mathtt{red}} : \mathbb{N} \to \mathbb{N}$ and $B^{\mathtt{alg}} : \mathbb{N} \to \mathbb{N}$. For $n \in \mathbb{N}$ and $s, m \in \mathbb{N}$ such that $2^{s+3m/4} \leq 2^m/5$, assume that there exists a Boolean sampler $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ with $B^{\mathtt{red}}$ bad inputs that is computable in $\widehat{\mathcal{C}}$. Also assume that there exists a distribution $\mathbf{S}_n$ of simplifier sets of size more than $B^{\mathtt{alg}}$ for $\widehat{\mathcal{C}}$. Then $B^{\mathtt{alg}}(n) < B^{\mathtt{red}}(n)$.*

[2] Recall that the majority function can also be replaced by "approximate majority".

[3] For example, if $\widehat{\mathcal{C}}$ is indeed defined as all circuits of the form $\Phi\left(\{C_0(Samp(x,i)\}_{i \in \{0,1\}^s}\right)$, where in particular C_0 and Φ can be any pair of dictator functions, then $Samp$ is computable in $\widehat{\mathcal{C}}$.

Proof. Assuming that a distribution $\mathbf{S}_n$ as in the hypothesis exists, we will construct a set $T \subseteq \{0,1\}^m$ such that for more than $B^{\mathtt{alg}}(n)$ inputs $x \in \{0,1\}^n$, the set T is "over-sampled" by the sampler $Samp$ with input x; that is, for every such x it holds that $\Pr_{i \in \{0,1\}^s}[Samp(x,i) \in T] > |T|/2^m + 1/10$. Thus, the number $B^{\mathtt{red}}(n)$ of "bad" inputs for the sampler is larger than $B^{\mathtt{alg}}(n)$.

In order to construct $T \subseteq \{0,1\}^m$, we will first fix a single set $S \subseteq \{0,1\}^n$ of size $|S| > B^{\mathtt{alg}}(n)$ of inputs for the sampler; these inputs will later on be the ones that will cause the sampler to "over-sample" T. To do so, for any fixed choice of $S \sim \mathbf{S}_n$, let us call an index $i \in \{0,1\}^s$ $\mathtt{bad\ under}$ S if at least a quarter of the output bits of $Samp^{(i)}|_S$ are constant functions (i.e., when restricting the function $Samp^{(i)}(x) = Samp(x,i)$ to the set S, at least $m/4$ of the m output bits of $Samp^{(i)}|_S$ become constant functions). For any $i \in \{0,1\}^s$, recall that each output bit of $Samp^{(i)}$ can be computed by a $\widehat{\mathcal{C}}$-circuit, and therefore (since $\mathbf{S}_n$ simplifies $\widehat{\mathcal{C}}$) the expected number of output bits of $Samp^{(i)}$ that become constant under $S \sim \mathbf{S}_n$ is more than half. Thus, there exists some set $S \sim \mathbf{S}_n$ such that at least one third of the indices $i \in \{0,1\}^s$ are bad under S. [4] We now fix any such set S, and note that $|S| > B^{\mathtt{alg}}(n)$ (since $\mathbf{S}_n$ is of size more than $B^{\mathtt{alg}}$).

Let us now construct $T \subseteq \{0,1\}^m$ that is "over-sampled" by the sampler when given any input $x \in S$. To do so, let $\mathcal{B} \subseteq \{0,1\}^s$ be the set of bad indices under S. For any $i \in \mathcal{B}$, let $\Phi_i \subseteq [m]$ be the set of output bits of $Samp^{(i)}|_S$ that are constant; that is, $\Phi_i = \left\{ j \in [m] : \exists \sigma_j \in \{0,1\}, (Samp^{(i)}|_S)_j \equiv \sigma_j \right\}$ and $|\Phi_i| \geq m/4$ (since i is bad under S). Also, let Q_i be the subcube of $\{0,1\}^m$ corresponding to the non-constant output bits of $Samp^{(i)}|_S$; that is, $Q_i = \left\{ z \in \{0,1\}^m : \forall j \in \Phi_i, z_j = (Samp^{(i)}|_S)_j \right\}$. We define T to be the union of the subcubes Q_i for all $i \in \mathcal{B}$; that is, $T = \bigcup_{i \in \mathcal{B}} Q_i$.

Note that $|T| < 2^m/5$, since for every $i \in \mathcal{B}$ it holds that $|Q_i| \leq 2^{3m/4}$ (and since we assumed that $2^s \cdot 2^{3m/4} < 2^m/5$). On the other hand, for every $x \in S$, the probability over $i \in \{0,1\}^s$ that $Samp(x,i) \in T$ is lower bounded by the probability that $i \in \mathcal{B}$, which is at least $1/3$. Therefore, for any $x \in S$ it holds that $\Pr_{i \in \{0,1\}^s}[Samp(x,i) \in T] \geq 1/3 > |T|/2^m + 1/10$. ∎

We mention that lemmas similar to Lemma 4 were proved for the specific setting of $\mathcal{AC}^0$ circuits (i.e., when $\mathcal{C} = \mathcal{AC}^0$) in [Vio05, Thm. 6.4] and [GVW15, Thm. 5.4], using a different proof strategy. [5] Relying on Lemma 4 and on the discussion that preceded the lemma's statement, we can now formally state our main theorem:

Theorem 5 *(a limitation of two "black-box" techniques in quantified derandomization; Theorem 1, restated). Let $B^{\mathtt{red}} : \mathbb{N} \to \mathbb{N}$ and $B^{\mathtt{alg}} : \mathbb{N} \to \mathbb{N}$. For $n, s, m \in \mathbb{N}$ such that*

[4] For any $i \in \{0,1\}^s$, denote by $\alpha_i(S)$ the random variable that is the number of output bits of $Samp^{(i)}|_S$ that are constant functions. Then, we have that $\Pr_{S \sim \mathbf{S}_n}[\alpha_i(S) \geq m/4] \geq 1/3$ (otherwise $\mathbb{E}_{S \sim \mathbf{S}_n}[\alpha_i(S)] < \frac{1}{3} \cdot m + \frac{2}{3} \cdot (m/4) = m/2$). Thus, the expected number of bad indices under $S \sim \mathbf{S}_n$ is at least $2^s/3$.

[5] Instead of relying directly on the existence of simplifier sets, they relied on the *low average sensitivity* of $\mathcal{AC}^0$ circuits; this property follows from a stronger notion of simplifier sets, which in particular are subcubes (i.e., it follows from Håstad's switching lemma; see [LMN93, Bop97, Tal17]). In comparison, our proof is simpler, and also applies to classes of functions with *high sensitivity* that have distributions of simplifier sets (recall that the simplifier sets in Definition 3 are not necessarily subcubes).

$2^{s+3m/4} \leq 2^m/5$, *let* $C : \{0,1\}^m \to \{0,1\}$, *and let* $Samp : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ *be a Boolean sampler with* $B^{\mathtt{red}}$ *bad inputs. Let* $\widehat{\mathcal{C}}$ *be any circuit class that can compute the function* $C'(x) = MAJ\left(\{C(Samp(x,i))\}_{i\in\{0,1\}^s}\right)$ *and such that* $Samp$ *is computable in* $\widehat{\mathcal{C}}$. *Then, for any distribution* $\mathbf{S}_n$ *of simplifier sets of size more than* $B^{\mathtt{alg}}$ *for* $\widehat{\mathcal{C}}$ *it holds that* $B^{\mathtt{alg}}(n) < B^{\mathtt{red}}(n)$.

We stress that Theorem 5 holds *regardless of the class* $\mathcal{C}$ for which we wanted to solve the standard derandomization problem. This is the case since the limitation in Theorem 5 only relies on the hypothesis that $Samp$ is computable in $\widehat{\mathcal{C}}$, and not on the fact that the circuit C' (which depends on C) is computable in $\widehat{\mathcal{C}}$.

Indeed, we did not use the fact that $\widehat{\mathcal{C}}$ contains a circuit (i.e., C') that computes a function that is "more complicated" than (the output bits of) $Samp$. Intuitively, we expect that "complicated" circuits will require simplifier sets that are smaller than simplifier sets for "simpler" circuits (e.g., Håstad's switching lemma [Hås87] yields simplifier sets of size $2^{\Omega(n/\log^{d-1}(n))}$ for circuits of depth d). In particular, in typical situations we expect that simplifier sets for $\widehat{\mathcal{C}}$ will need to be even smaller than simplifier sets for $Samp$. In such situations, the upper bound on $B^{\mathtt{alg}}(n)$ in Theorem 5 is not tight; that is, in such situations $B^{\mathtt{red}}$ and $B^{\mathtt{alg}}$ are far apart by significantly more than just a single bit (as is asserted in the theorem).

4 Strengthenings of the main theorem: Natural relaxations that do not suffice to bypass the limitation in Theorem 5

Following Theorem 5, the main question we are faced with is *which relaxations of the hypotheses* are sufficient to avoid the conclusion of the theorem, and thus to bypass the limitation arising from it. We now mention a few natural relaxations that *do not* seem sufficient to bypass this limitation.

4.1 Simplifier sets that simplify $\mathcal{C}$-circuits to non-constant functions

In Definition 3 we assumed that for every $C \in \mathcal{C}$, with high probability over $S \sim \mathbf{S}_n$ it holds that $C|_S$ is *constant*. We now note that if we assume that $C|_S$ simplifies to a "simple" (non-constant) function, then we can still obtain a corresponding quantified derandomization algorithm (with a mild loss in the parameters); but that in *some* natural settings, this relaxation does not suffice to bypass the limitations in Theorem 5.

Let us first see why this relaxation still suffices for quantified derandomization. Fix a circuit $C : \{0,1\}^n \to \{0,1\}$ with at most $B^{\mathtt{alg}}(n)$ exceptional inputs. Assume that there exists a distribution $\mathbf{S}_n$ over subsets $S \subseteq \{0,1\}^n$ of size at least $3 \cdot B^{\mathtt{alg}}(n)$ such that $\Pr_{S\sim\mathbf{S}_n}[C|_S \in \mathcal{C}_{Simple}] > 1/2$, where $\mathcal{C}_{Simple}$ is some class of "simple" functions. Further assume that an algorithm can efficiently sample a succint representation of $S \sim \mathbf{S}_n$ using few random bits, and that given a representation of S such that $C|_S \in \mathcal{C}_{Simple}$, the algorithm can efficiently approximate the acceptance probability of $C|_S$, up to error $1/10$. Note that for a strict majority of the choices of $S \sim \mathbf{S}_n$, the acceptance probability

of $C|_S$ is either at most $1/3$ or at least $2/3$ (since $|S| > 3 \cdot B^{\mathtt{alg}}(n)$ for any $S \sim \mathbf{S}_n$), and we can distinguish between the two cases by estimating the acceptance probability of $C|_S \in \mathcal{C}_{Simple}$. Thus, a quantified derandomization algorithm can enumerate the choices of $S \sim \mathbf{S}_n$, decide for each choice whether the acceptance probability of $C|_S$ is at most $1/3$ or at least $2/3$, and rule according to a majority vote. [6]

Nevertheless, in some cases, this relaxation does not seem sufficient to bypass the limitation in Theorem 5. This is since for *some* natural classes $\mathcal{C}_{Simple}$ of "very simple" functions, a random restriction simplifies every $C \in \mathcal{C}_{Simple}$ to a constant function, with high probability; for example, this holds for constant-depth circuits [Hås87] and for linear threshold functions [KW16]. In these cases, the existence of $\mathbf{S}_n$ as above implies the existence of $\mathbf{S}'_n$ that meets the stronger definition (i.e., Definition 3), with a quantitative loss in the parameter $B^{\mathtt{alg}}$ that depends on $\mathcal{C}_{Simple}$ (i.e., the loss is induced by the random restriction that turns functions in $\mathcal{C}_{Simple}$ to constant functions).

4.2 A sampler that only samples $\mathcal{C}$-events

Our requirement from the sampler in Definition 2 was information-theoretic: For *any set* $T \subseteq \{0,1\}^m$, we required that for all but $B^{\mathtt{red}}$ inputs, the sampler will hit T with an approximately correct probability (i.e., $\Pr_{i \in \{0,1\}^s}[Samp(x,i) \in T] \in |T|/2^m \pm 1/10$). However, since we only want to use the sampler to approximate the acceptance probability of a circuit $C \in \mathcal{C}$, one may consider a relaxation in wish we only require that the sampler "appropriately samples" sets T that are decidable by $\mathcal{C}$ circuits. That is,

Definition 6 *(sampler for $\mathcal{C}$-events). We say that a function $Samp_{\mathcal{C}} : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}^m$ is a* Boolean sampler for $\mathcal{C}$-events with $B^{\mathtt{red}}$ bad inputs *if it satisfies the following: For every $T \subseteq \{0,1\}^m$ such that $T = C^{-1}(1)$ for some $C \in \mathcal{C}$, for all but at most $B^{\mathtt{red}}$ of the inputs $x \in \{0,1\}^n$ it holds that $\Pr_{i \in \{0,1\}^s}[Samp(x,i) \in T] \in |T|/2^m \pm 1/10$.*

The point that we wish to make is that in many natural settings, the relaxation in Definition 6 does not suffice to bypass the limitation in Theorem 5. This is the case because the "over-sampled" set $T \subseteq \{0,1\}^m$ that was constructed in the proof of Theorem 5 can be decided by a circuit that is a *DNF of size at most* 2^s. [7] (Using the notation of the proof, T is the union of $|\mathcal{B}| \leq 2^s$ subcubes $Q_1, ..., Q_{|\mathcal{B}|} \subseteq \{0,1\}^m$.) In particular, if the initial circuit C belongs to a class $\mathcal{C}$ that contains DNFs of size 2^s, then the relaxation in Definition 6 does not suffice to bypass the limitation in Theorem 5.

Detour. The notion of a sampler for $\mathcal{C}$-events might be of independent interest. The main point is that potentially, in some settings, the relaxation embodied in the definition of samplers for $\mathcal{C}$-events might allow to construct such samplers with better

[6] The algorithm may not be able to correctly decide whether the acceptance probability of $C|_S$ is at most $1/3$ or at least $2/3$ when $C|_S \notin \mathcal{C}_{Simple}$, but the latter event only happens in the minority of the choices $S \sim \mathbf{S}_n$.

[7] Recall that a typical setting of the parameters is $s = O(\log(n))$, and therefore the size of such a DNF is $2^s = \mathrm{poly}(n)$.

parameters than the parameters of information-theoretic samplers (i.e., as in Definition 2). For an example of a construction of a sampler for $\mathcal{AC}^0$-events, see [GW14, Thm. 3.4 in the full version]: They constructed an $\mathcal{AC}^0$-computable sampler for $\mathcal{AC}^0$-events with $2^{n/\text{poly}\log(n)}$ bad inputs and $m = n^{\Omega(1)}$.

We note, however, that this construction from [GW14] was later superseded by a construction of an $\mathcal{AC}^0$-computable sampler in the *information-theoretic* sense (i.e., as in Definition 2) that also has $2^{n/\text{poly}\log(n)}$ bad inputs, and that has $m = n/\text{poly}\log(n)$; see [CL16, Thms. 1.5 & 1.7]. Moreover, this construction, coupled with Lemma 4 and with Håstad's switching lemma [Hås87], implies that $\mathcal{AC}^0$-computable samplers for $\mathcal{AC}^0$-events *cannot* have significantly better parameters than $\mathcal{AC}^0$-computable samplers in the information-theoretic sense (i.e., as in Definition 2). This is the case since Håstad's switching lemma yields a distribution of simplifier sets of size $2^{\Omega(n/\log^{d-2}(n))}$ for depth-d circuits, and thus Lemma 4 implies that any sampler computable by depth-d circuits (even if it is a sampler only for DNF-events) must have at least $2^{\Omega(n/\log^{d-2}(n))}$ bad inputs; [8] and the number of bad inputs in the information-theoretic constructions in [CL16] is already $2^{\Omega(n/\log^{d-10}(n))}$ (or $2^{n/\log^{\Omega(d)}(n)}$, if one wishes to maximize the output length m), which nearly matches this lower-bound.

4.3 Samplers that are not efficiently computable

To combine the two techniques from Section 2 into a single algorithm, we need an efficient *uniform* algorithm that can compute $Samp(x, i)$ for arbitrarily-chosen $x \in \{0,1\}^n$ and $i \in \{0,1\}^s$; that is, we need the sampler not only to be computable in the class $\mathcal{C}$, but also to be efficiently computable by a uniform algorithm.

We note, however, that the limitation in Theorem 5 holds even if we use a sampler that is not necessarily efficiently computable by a uniform algorithm. This is the case since the argument in Theorem 5 did not rely on such a hypothesis regarding the sampler, and thus holds also for "non-uniform" samplers.

Acknowledgements

The author thanks his advisor, Oded Goldreich, for his encouragement to write the note and to publish it, and for his close guidance in the processes of conceptualizing the result and of writing the note.

References

[Ajt83] M. Ajtai. Σ_1^1-formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.

[8] The lower bound on the number of bad inputs for $\mathcal{AC}^0$-computable samplers for $\mathcal{AC}^0$-events can also be derived using the proofs of [Vio05, GVW15]: In their proofs, the "over-sampled" set can also be decided by a DNF of size that is at most exponential in the seed length of the extractor.

[Bop97] Ravi B. Boppana. The average sensitivity of bounded-depth circuits. *Information Processing Letters*, 63(5):257–261, 1997.

[CL16] Kuan Cheng and Xin Li. Randomness extraction in AC0 and with small locality. *Electronic Colloquium on Computational Complexity: ECCC*, 23:18, 2016.

[Gol08] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, New York, NY, USA, 2008.

[GVW15] Oded Goldreich, Emanuele Viola, and Avi Wigderson. On randomness extraction in AC0. In *Proc. 30th Annual IEEE Conference on Computational Complexity (CCC)*, pages 601–668, 2015.

[GW14] Oded Goldreich and Avi Widgerson. On derandomizing algorithms that err extremely rarely. In *Proc. 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 109–118. 2014. Full version available online at *Electronic Colloquium on Computational Complexity: ECCC*, 20:152 (Rev. 2), 2013.

[Hås87] Johan Håstad. *Computational Limitations of Small-depth Circuits*. MIT Press, 1987.

[KW16] Daniel M. Kane and Ryan Williams. Super-linear gate and super-quadratic wire lower bounds for depth-two and depth-three threshold circuits. In *Proc. 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 633–643, 2016.

[LMN93] Nathan Linial, Yishay Mansour, and Noam Nisan. Constant depth circuits, Fourier transform, and learnability. *Journal of the Association for Computing Machinery*, 40(3):607–620, 1993.

[Tal17] Avishay Tal. Tight bounds on the fourier spectrum of AC0. In *Proc. 32nd Annual IEEE Conference on Computational Complexity (CCC)*, pages 15:1–15:31, 2017.

[Tel17a] Roei Tell. Improved bounds for quantified derandomization of constant-depth circuits and polynomials. In *Proc. 32nd Annual IEEE Conference on Computational Complexity (CCC)*, pages 18:1 – 18:49, 2017.

[Tel17b] Roei Tell. Quantified derandomization of linear threshold circuits. *Electronic Colloquium on Computational Complexity: ECCC*, 24:145, 2017.

[Vad12] Salil P. Vadhan. *Pseudorandomness*. Foundations and Trends in Theoretical Computer Science. Now Publishers, 2012.

[Vio05] Emanuele Viola. The complexity of constructing pseudorandom generators from hard functions. *Computational Complexity*, 13(3-4):147–188, 2005.

[Vio09] Emanuele Viola. On approximate majority and probabilistic time. *Computational Complexity*, 18(3):337–375, 2009.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il