Electronic Colloquium on Computational Complexity, Report No. 86 (2004)

Pseudorandomness for Approximate Counting and Sampling

Ronen Shaltiel* Christopher Umans†

October 11, 2004

Abstract

We study computational procedures that use both randomness and nondeterminism. Examples are Arthur-Merlin games and approximate counting and sampling of NP-witnesses. The goal of this paper is to derandomize such procedures under the weakest possible assumptions.

Our main technical contribution allows one to "boost" a given hardness assumption. One special case is a proof that

$$\mathrm{EXP} \not\subseteq \mathrm{NP}/\mathrm{poly} \Rightarrow \mathrm{EXP} \not\subseteq \mathrm{P}_{||}^{\mathrm{NP}}/\mathrm{poly}.$$

In words, if there is a problem in EXP that cannot be computed by poly-size nondeterministic circuits then there is one which cannot be computed by poly-size circuits which make non-adaptive NP oracle queries. This in particular shows that the various assumptions used over the last few years by several authors to derandomize Arthur-Merlin games (i.e., show AM $=$ NP) are in fact all *equivalent*. In addition to simplifying the framework of AM derandomization, we show that this "unified assumption" suffices to derandomize several other probabilistic procedures.

For these results we define two new primitives that we regard as the natural pseudorandom objects associated with *approximate counting* and *sampling* of NP-witnesses. We use the "boosting" theorem (as well as some hashing techniques) to construct these primitives using an assumption that is no stronger than that used to derandomize Arthur-Merlin games. As a consequence, under this assumption, there are *deterministic* polynomial time algorithms that use *non-adaptive* NP-queries and perform the following tasks:

- approximate counting of NP-witnesses: given a Boolean circuit A, output r such that $(1-\epsilon)|A^{-1}(1)| \leq r \leq |A^{-1}(1)|$.
- pseudorandom sampling of NP-witnesses: given a Boolean circuit A, produce a polynomial-size sample space that is computationally indistinguishable from the uniform distribution over $A^{-1}(1)$.

We also present applications. For example, we observe that Cai's proof that $S_2^p \subseteq \mathrm{ZPP}^{\mathrm{NP}}$ and the learning algorithm of Bshouty et al. can be seen as non-randomized reductions to sampling. As a consequence they can be derandomized under the assumption stated above, which is weaker than the assumption that was previously known to suffice.

*Department of Computer Science, University of Haifa, Mount Carlel, Haifa 31905, Israel. Some of this work was done while at the Weizmann Institute and supported by the Koshland Scholarship. Email: ronen@haifa.ac.il.

†Department of Computer Science, California Institute of Technology, Pasadena, CA 91125. Email: umans@cs.caltech.edu. This research was supported by NSF grant CCF-0346991.

ISSN 1433-8092

1 Introduction

One of the major areas in complexity is the study of the power of randomness in various computational settings. In certain contexts randomness affords additional power. But for broad classes of problems it has been demonstrated over the last decade that randomness can be simulated deterministically, under widely accepted complexity assumptions.

The central object used in these derandomization results is a *pseudorandom generator* (PRG), which is an efficient deterministic procedure that generates a *discrepancy set* – a set of strings with the property that no test (from a pre-specified class of tests) can distinguish a random string in the discrepancy set from a uniformly random string. We say that a PRG *fools* this class of tests. A probabilistic procedure is derandomized by replacing its random bits with strings from the discrepancy set; the procedure cannot behave noticeably differently than it would with truly random bits, as then it would constitute a distinguishing test. As a consequence derandomizing stronger probabilistic algorithms typically requires pseudorandom generators that produce discrepancy sets for stronger classes of tests.

An efficient pseudorandom generator for some class of tests immediately implies an efficiently computable function which is hard for these tests. More specifically, an efficient pseudorandom generator that fools small circuits implies the existence of a language in a uniform complexity class (e.g., $E = DTIME(2^{O(n)})$) that lies outside a non-uniform complexity class (e.g. $P/poly$). Thus constructing such pseudorandom generators amounts to proving circuit lower bounds for explicit functions, a task that is currently beyond our reach. Consequently, this line of research focuses on constructing pseudorandom generators under a *hardness assumption*[1]. In this context the goal is to achieve derandomization results under the weakest possible hardness assumptions.

One of the main efforts in derandomization over the last decade has focused on the class BPP which can be derandomized given access to pseudorandom generators that fool small circuits. Here the appropriate hardness assumption is that there exists a language in E that requires exponential size circuits (i.e., the language cannot be computed by size $2^{\epsilon n}$ circuits, for some $\epsilon > 0$).[2] A sequence of results [NW94, BFNW93, Imp95, IW97] showed that under this hardness assumption $BPP = P$. A further sequence of papers achieved a *quantitatively optimal* hardness vs. randomness tradeoff [ISW99, ISW00, SU01, Uma02].

An analogous line of work [AK01, KvM02, MV99, SU01] derandomized Arthur-Merlin games [Bab85, BM88, GMR89]. (Recall that the class AM contains important and natural problems like graph non-isomorphism that are not known to be in NP). These works achieved $AM = NP$ under progressively *qualitatively* weaker hardness assumptions. The first results required average-case hardness for circuits that make non-adaptive queries to an NP oracle, while the latest results require only hardness for SV-nondeterministic circuits[3] In this paper we show that the various different assumptions used to derandomize AM are in fact equivalent.

Another line of research [Sto83, JVV86, BGP00] addresses procedures which approximately count and sample NP-witnesses. More precisely, given a Boolean circuit A the first task is to approximately count the number of accepting inputs of A, and the second is to sample a random accepting input. Note that both problems are NP-hard and thus any such procedure must use nondeterminism unless $NP = P$. The current procedures for these tasks also use randomization: they are probabilistic algorithms which use an NP-oracle. We remark that procedures for counting can use non-adaptive NP-queries whereas the known procedures for sampling use adaptive NP-queries. In this paper we show how to derandomize these procedures and show that under a hardness assumption that is no stronger that used to derandomize AM, both these tasks can be performed by polynomial time *deterministic* algorithms which make *non-adaptive* NP-queries.

In order to achieve these results we make a technical contribution and a conceptual contribution. Our main technical result is a "downward collapse theorem" that implies (as a special case): [4]

$$E \subseteq P_{||}^{NP}/\text{poly} \Rightarrow E \subseteq NP/\text{poly}.$$

[1] This "hardness vs. randomness paradigm" was initiated by [BM84, Yao82, Sha81]. It should be noted that the notion of pseudorandom generators in these papers is different than the one we use here. In particular, in this paper we follow a paradigm initiated by [NW94] which allow pseudorandom generators which given a size bound s, run in time polynomial in s and output a discrepancy set for size s circuit. The reader is referred to [Gol98] for a survey on pseudorandomness and its applications and to [Kab02] for a recent survey which focuses on derandomization.

[2] One of the confusing aspects of all the results in this area is that the assumptions involve "exponential time" classes. In actual applications these assumptions are "scaled down" to say that there exists a function on $O(\log n)$ bits which is computable in polynomial time and cannot be computed by size n^c circuits (for some constant c).

[3] SV-nondeterministic circuits are the nonuniform analog of the class $NP \cap coNP$ (see definition 2.2).

[4] The notation $A_{||}^{B}$ says that A uses *non-adaptive* queries to the oracle B.

This downward collapse shows that all of the various flavors of nondeterministic hardness assumptions considered in the literature are equivalent. This unifies and simplifies a number of past results. This result is also helpful when derandomizing other probabilistic procedures which involve randomness and nondeterminism. It allows us to start from a weak hardness assumption, boost it to a stronger hardness assumption, and then use pseudorandom generators for stronger classes of tests, namely circuits which make non-adaptive NP-queries.

Our conceptual contribution lies in defining what we regard as the natural "derandomization objects" associated with approximate counting and sampling. These are *relative-error approximators* (for approximate counting) and *conditional discrepancy sets* (for sampling). The first is a strengthening of additive-error approximators (which derandomize BPP), and the second is a generalization of discrepancy sets (which "sample" from the uniform distribution). We show how to obtain relative-error approximators and conditional discrepancy sets in polynomial time with non-adaptive NP oracle access, under a hardness assumption no stronger than that used for derandomizing AM. Note that in particular this suggests that the "true complexity" of these problems is $FP_{||}^{NP}$ and in particular that adaptive NP queries are not necessary for sampling. Loosely speaking, our technique uses the strong pseudorandom generators obtained by boosting the initial hardness assumption to derandomize the probabilistic procedures for approximate counting and sampling. We remark that this derandomization relies on the specific imp lementation of this procedures, and that some additional tricks are needed to obtain procedures that make *nonadaptive* queries to an NP-oracle.

We also give several applications of relative error approximators and conditional discrepancy sets. We obtain the following collapses under a hardness assumption similar to that used for derandomizing AM: $S_2^P = P^{NP}$ and $BPP_{path} = P_{||}^{NP}$. The first collapse comes from viewing Cai's result [Cai01] (that places S_2^P in ZPP^{NP}) as a *non-randomized reduction* of S_2^P to sampling, which in turn is derandomized via conditional discrepancy sets. Similarly, we view a fundamental result by Bshouty et al. [BCG+96] (concerning the learning of circuits using equivalence queries) as a reduction to sampling, and derandomize it in the same way.

Outline

In Section 2 we present definitions of the various types of nondeterministic circuits and hardness assumptions. In Section 3 we describe our main results and relation to prior work. In Section 4 we describe the major ideas and techniques used in the proofs; Sections 5 and 6 contain the full proofs. Finally in Section 7 we conclude with some open problems.

2 Nondeterministic circuits and hardness

We assume that the reader is familiar with (deterministic) Boolean circuits. We use the convention that the size of a circuit is the total number of wires and gates. Nondeterministic circuits come in several flavors, which we define below. We remark that a main contribution of this paper lies in showing that the multitude of hardness assumptions defined below are all equivalent – unfortunately, in order to show that, we need to be able to discuss all of the various assumptions below.

Definition 2.1 (nondeterministic and co-nondeterministic circuits). *A nondeterministic (co-nondeterministic) circuit is a Boolean circuit C with a set of n inputs x, and a second set of inputs y. The* function computed by C*, denoted $f_C : \{0,1\}^n \to \{0,1\}$ is defined by $f_C(x) = 1$ iff $\exists y\ C(x,y) = 1$ ($\forall y\ C(x,y) = 0$).*

The uniform analogue of nondeterministic circuit is the class NP. The uniform analogue of co-nondeterministic circuits is $coNP$. Single-valued nondeterministic circuits have $NP \cap coNP$ as their uniform analogue.

Definition 2.2 (single-valued nondeterministic circuits). *A single-valued nondeterministic (or SV-nondeterministic) circuit is a Boolean circuit C with a set of n inputs x, a second set of inputs y, and two outputs* **value** *and* **flag**. *Circuit C computes the function $f : \{0,1\}^n \to \{0,1\}$ if the following hold:*

- *for every x, y, if $C(x,y)$ has 1 at its* **flag** *gate then $C(x,y)$ has $f(x)$ at its* **value** *gate, and*
- *for every x, there exists some y for which $C(x,y)$ has 1 at its* **flag** *gate.*

Note that a circuit C may meet the syntactic demands of this definition, and yet not compute any function. When we refer to a SV-nondeterministic circuit, we always mean a circuit C that in fact computes a function according to this definition, and we refer to that unique function as the *function computed by* C.

Definition 2.3 (NP-circuits and NP$||$-circuits). *An NP-circuit is a Boolean circuit C that is also permitted to use SAT-oracle gates. SAT-oracle gates are gates with many inputs and a single output that is 1 iff the input is in SAT.*

An NP$||$-circuit is a pair of Boolean circuits C_{pre} and C_{post}. On input x, C_{pre} outputs a number of queries $q_1, q_2, \ldots, q_m$. Circuit C_{post} receives x together with m bits $a_1, a_2, \ldots, a_m$, where $a_i = 1$ iff q_i is in SAT, and outputs a single answer bit.

We could also have defined NP$||$-circuits to be NP-circuits in which no path from the output gate to an input gate encounters more than one SAT-oracle gate; the above definition makes explicit the pre- and post- processing phase. For NP$||$-circuits so defined, their size is the sum of the sizes of C_{pre} and C_{post}.

We will frequently speak of a language L that is "hard for" a class of circuits. Of course this hardness can be quantified by the size of the circuit. For clarity, we have chosen only to present the "high-end" results that follow when this hardness is exponential, even though more general results are true using our methods. Consequently, we only need the following definitions:

Definition 2.4 (worst-case hardness for exponential-size circuits). *A language L is worst-case hard for exponential-size (deterministic, nondeterministic, co-nondeterministic, SV-nondeterministic, NP-, or NP$||$-) circuits if there exists a constant $\epsilon > 0$ for which every circuit of the prescribed type and size at most $2^{\epsilon n}$, fails to compute L restricted to inputs of length n, for all sufficiently large n.*

Definition 2.5 (average-case hardness for exponential-size circuits). *A language L is α-hard for exponential-size (deterministic, nondeterministic, co-nondeterministic, SV-nondeterministic, NP-, or NP$||$-) circuits if there exists a constant $\epsilon > 0$ for which every circuit of the prescribed type and size at most $2^{\epsilon n}$, fails to compute L restricted to inputs of length n on at least $(1-\alpha)2^n$ such inputs, for all sufficiently large n.*

Note that the definition of $(1 - 2^{-n})$-hard coincides with the definition of worst-case hard.

Definition 2.6 (worst-case and average-case hardness of complexity classes). *A complexity class $\mathcal{C}$ is worst-case hard (resp. α-hard) for exponential-size circuits of a given type if there exists a language $L \in \mathcal{C}$ that is worst-case hard (resp. α-hard) for exponential-size circuits of that type.*

We also sometimes say "$\mathcal{C}$ *requires* exponential-size circuits" of a given type to mean $\mathcal{C}$ is worst-case hard for exponential-size circuits of that type.

3 Main results

Several of our results apply to any complexity class for which one can compute the low-degree extension within that class. To make these results easier to state we introduce the following definition:

Definition 3.1. *We say that a complexity class $\mathcal{C}$* allows low-degree extension *if $E^{\mathcal{C}^{\leq O(n)}} \subseteq \mathcal{C}$, where the notation $\mathcal{C}^{\leq O(n)}$ means that the E oracle machine makes only linear-length queries.*

Examples of complexity classes $\mathcal{C}$ that support low-degree extension are: E, $NE \cap coNE$, E^{NP}, $E_{||}^{NP}$.

3.1 Unifying hardness assumptions

Several authors [AK01, KvM02] have observed that the PRG constructions intended to derandomize BPP can be adapted to construct discrepancy sets that fool efficient *non-deterministic* tests under stronger hardness assumptions. Just as PRGs that fool efficient deterministic tests imply $BPP = P$, PRGs that fool efficient non-deterministic tests imply $AM = NP$.

Several hardness assumptions sufficient to achieve $AM = NP$ have been considered in the literature. All of these hardness assumptions (and the others we will consider in this paper) have the following form: there exists a language

L in some "high" uniform class (examples are E, $NE \cap coNE$, $E_{||}^{NP}$ and E^{NP}) which requires exponential size circuits from some non-uniform circuit class[5]. Three non-uniform circuit classes have been discussed in the literature in relation to AM. These are

- SV-nondeterministic circuits, used by Milersen and Vinodchandran [MV99] and later Shaltiel and Umans [SU01],
- non-deterministic (and co-nondeterministic) circuits, used by Arvind and Kobler [AK01], and
- $NP||$-circuits, used by Klivans and van Melkebeek [KvM02][6],

listed in order from weaker to stronger. Perhaps the best way to understand these circuit classes is to think of them as nonuniform analogs of $NP \cap coNP$, NP (and $coNP$), and $P_{||}^{NP}$, respectively. Figure 1 summarizes the various hardness assumptions and pseudorandom generators implying $AM = NP$ and known relationships between them.

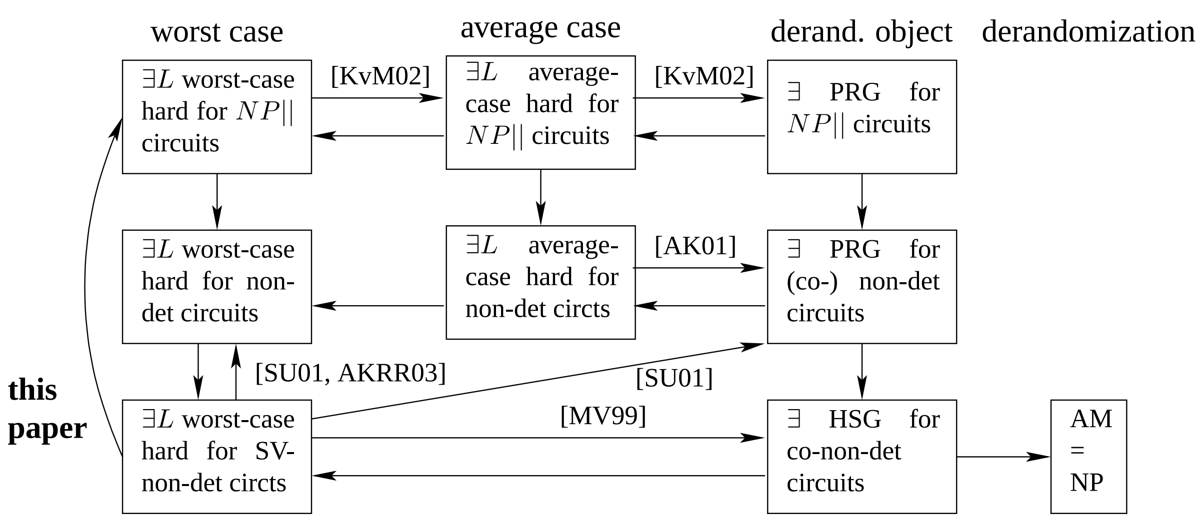

Figure 1: Assumptions implying $AM = NP$. In all cases L is a language in $NE \cap coNE$. The phrase "L worst-case (resp., average-case) hard for" means "L cannot be computed exactly by (resp., approximated by) size $2^{\epsilon n}$ for some $\epsilon > 0$." Arrows indicate implications; unlabelled arrows correspond to implications that follow from standard arguments.

Notice that with the exception of the $AM = NP$ box, there are two strongly connected components, consisting of the top row and the bottom two rows. In this paper we show that *all of the hardness assumptions considered in the literature are equivalent*. In addition to clarifying the situation, this result somewhat simplifies the task of building a PRG sufficient to derandomize AM. One can replace previous constructions [MV99, SU01] that are specialized for derandomizing AM under an SV-nondeterministic hardness assumption by *any* relativizing construction of ordinary pseudorandom generators (designed to derandomize BPP). Furthermore, by "climbing up" using the result of this paper, we get a construction of PRGs against NP||-circuits using only hardness for SV-nondeterministic circuits.

[5]We stress that it is the choice of the nonuniform circuit class that typically plays an important role in the argument. Loosely speaking, this choice determines the class of tests to be fooled by the generator. The choice of the uniform class determines the efficiency of the generator. For example, choosing this class to be E gives a generator which runs in P, whereas $NE \cap coNE$ (or E^{NP}) give a generator which runs in $NP \cap coNP$ (or P^{NP}). We encourage the reader to ignore the precise choice of the uniform class at a first reading and focus on the choice of the nonuniform class.

[6]Actually, the paper in question refers to SAT oracle circuits, but their argument works just as well for $NP||$-circuits, giving a stronger result.

3.1.1 A downward collapse theorem

The equivalence of the various hardness assumptions is implied by the following downward collapse theorem, which may be of independent interest:

Theorem 3.2 (downward collapse). *Let $\mathcal{C}$ be any complexity class that allows low-degree extension. If every language in $\mathcal{C}$ has $NP||$ circuits of size $s(n)$ then every language in $\mathcal{C}$ has SV-nondeterministic circuits of size $s(n)^{O(1)}$.*

A special case of Theorem 3.2 is:

$$E \subseteq P_{||}^{NP}/\text{poly} \Rightarrow E \subseteq NP/\text{poly} \cap coNP/\text{poly}.$$

We remark that it is widely believed that $P_{||}^{NP}$ is stronger than $NP \cap coNP$ and that $NP||$-circuits are stronger than SV-nondeterministic circuits. Nevertheless, a collapse of E to the stronger class implies a further collapse to the weaker class.

The following Corollary is the contrapositive version of Theorem 3.2 which states that a "weak" hardness assumption implies a stronger one:

Corollary 3.3. *For every class $\mathcal{C}$ that allows low-degree extension, if $\mathcal{C}$ is worst-case hard for exponential-size SV-nondeterministic circuits then $\mathcal{C}$ is worst-case hard for exponential-size NP$||$-circuits.*

Corollary 3.3 will allow us to derandomize many probabilistic algorithms and classes using hardness for SV-nondeterministic circuits by first "boosting" this assumption to hardness for $NP||$-circuits, then working with the pseudorandom generators obtained by the latter assumption.

3.2 Derandomization objects for approximate counting and sampling

In this section we define two generic computational objects – relative-error approximators, and conditional discrepancy sets. These objects are natural and make no reference to nondeterminism. They are intended to capture approximate counting and sampling, and they generalize and strengthen two existing and widely used objects: additive-error approximators and (ordinary) discrepancy sets.

3.2.1 Relative-error approximators

Ordinary pseudo-random generators allow one to obtain an *additive* approximation of the acceptance probability of circuits:

Definition 3.4. *An (additive-error)* approximator *is a procedure that takes as input a Boolean circuit A, and $\epsilon > 0$, and outputs a real number ρ for which*

$$\left|\Pr_x[A(x) = 1] - \rho\right| \leq \epsilon.$$

Indeed additive approximation is in some sense the *raison d'etre* of ordinary PRGs: additive approximation of the acceptance probability of circuits allows one to derandomize BPP. Relative error approximation allows *approximate counting*, and is much more difficult (it is NP-hard). We will be concerned with *one-sided relative-error* approximations of the acceptance probability of circuits:

Definition 3.5. *A* one-sided relative-error approximator *is a procedure that takes as input a Boolean circuit A, and $\epsilon > 0$, and outputs a real number ρ for which*

$$(1 - \epsilon) \Pr_x[A(x) = 1] \leq \rho \leq \Pr_x[A(x) = 1].$$

We give a construction of deterministic relative-error approximators under a hardness assumption for SV-nondeterministic circuits.

Theorem 3.6 (construction of relative-error approximators). *If $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits, then there is a deterministic* one-sided relative-error approximator *that runs in time polynomial in the length of its input and $1/\epsilon$, with* non-adaptive *access to an NP oracle.*

As an immediate corollary, we obtain

Corollary 3.7. *If $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits, then for every function $f : \{0,1\}^n \to \mathbb{N}$ in $\#P$, and every $\epsilon > 0$, there is a deterministic procedure P running in poly(n, ϵ^{-1}) time with non-adaptive access to an NP-oracle for which (on input x):*

$$(1 - \epsilon)f(x) \leq P(x) \leq f(x);$$

in other words, every problem in $\#P$ can be approximated in $FP_{||}^{NP}$.

Note that it was shown in [Sto83, JVV86, BGP00] that using randomness and an NP-oracle, it is possible to uniformly sample NP-witnesses. This implies a *fully polynomial-time randomized approximation scheme* (FPRAS) with access to an NP-oracle, for every problem in $\#P$. However, no deterministic *fully polynomial-time approximation schemes* (FPAS's) with access to an NP-oracle are known for any $\#P$-complete problem; the above corollary gives FPAS's that make non-adaptive NP oracle queries for all problems in $\#P$, albeit under a complexity assumption.

3.2.2 Conditional discrepancy sets

Ordinary pseudo-random generators are sometimes called "discrepancy set generators," since they produce the following object:

Definition 3.8. *An (n, s, ϵ)-*discrepancy set *is a subset $T \subseteq \{0,1\}^n$ with the property that for all Boolean circuits C of size at most s:*

$$\left|\Pr_x[C(x) = 1] - \Pr_{t \in T}[C(t) = 1]\right| \leq \epsilon.$$

A discrepancy set is a "good sample" of *strings* $x \in \{0,1\}^n$, with respect to any property $\mathcal{P}$ that is decidable by small Boolean circuits. Of course one particularly useful such property is the property that a BPP machine with a fixed input accepts when given string x as its random coins.

Frequently one wishes to obtain a "good sample" of *strings* $x \in S$ for some subset $S \subseteq \{0,1\}^n$. Again, the sample should be good with respect to any property $\mathcal{P}$ that is recognizable by small Boolean circuits. For example S may be the set of 3-colorings of a given graph; a property of interest might be the property of having two specified nodes colored with the same color. A large body of literature is devoted to sampling various structures (e.g., colorings, matchings, contingency tables, etc...), often employing Markov Chain Monte Carlo methods.

We define *conditional discrepancy sets* as the derandomization object associated with such sampling in its full generality. We will allow the set S to be any set recognizable by a small Boolean circuit; that is, $S = A^{-1}(1)$ for some small circuit A. Conditional discrepancy sets capture "pseudorandomly sampling an accepting input of A" and can be seen to be a natural generalization of ordinary discrepancy sets.

Definition 3.9. *Let $S \subseteq \{0,1\}^n$ be some subset. An S-*conditional (n, s, ϵ)-discrepancy set *is a subset $T \subseteq \{0,1\}^n$ with the property that for all Boolean circuits C of size at most s:*

$$\left|\Pr_x[C(x) = 1 | x \in S] - \Pr_{t \in T}[C(t) = 1 | t \in S]\right| \leq \epsilon.$$

Our main result here is a procedure to efficiently generate conditional discrepancy sets under a hardness assumption (which is the same hardness assumption used to derandomize AM):

Theorem 3.10 (construction of conditional discrepancy sets). *If $E_{||}^{NP}$ (resp. E^{NP}) requires exponential size SV-nondeterministic circuits, then there is a deterministic procedure that takes as input a Boolean circuit A that accepts a subset $S \subseteq \{0,1\}^n$, an integer s, and $\epsilon > 0$, and outputs an S-*conditional (n, s, ϵ)-discrepancy set *T. The procedure runs in poly$(|A|, n, s, 1/\epsilon)$ time with* non-adaptive *(resp. adaptive) access to an NP oracle.*

3.3 Applications

3.3.1 Applications of Theorem 3.2

We use Theorem 3.2 to prove our two other main theorems, regarding the deterministic construction of relative-error approximators and conditional discrepancy sets. An additional application is given in the next theorem.

Theorem 3.11. *If $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits, then $BPP_{||}^{NP} = P_{||}^{NP}$.*

Klivans and van Melkebeek [KvM02] formalized the notion of a *relativizing* PRG construction, and observed that such constructions can be used to fool circuit classes that are stronger than deterministic circuits, if one is willing to make a similarly stronger hardness assumption. For example, this observation allows the construction of PRGs that fool NP||-circuits, assuming there exist languages that are hard for NP||-circuits. Our Corollary 3.3 states that hardness for SV-nondeterministic circuits implies hardness for NP||-circuits. As a consequence, existing relativizing PRG constructions (e.g. [IW97, STV01, SU01, Uma02]) may be used directly to fool NP|| circuits, assuming only hardness for SV-nondeterministic circuits. As stated in the above theorem, this in turn derandomizes the class $BPP_{||}^{NP}$ using a weaker assumption than previously known.

We next present an additional application of Corollary 3.3 that turns out to be critical in the proof of Theorem 3.10. The following remarkable observation is found in [KvM02]: if $E_{||}^{NP}$ requires exponential size NP||-circuits, then there is a polynomial-time procedure to *produce* a satisfying assignment of a given circuit C that uses *non-adaptive* access to an NP-oracle. Note that the standard (unconditional) method uses adaptive access. The non-adaptive procedure comes from noting that there is a polynomial time algorithm that makes non-adaptive NP queries to test whether the outcome of applying the Valiant-Vazirani reduction to a satisfiable circuit C (for a specific choice of random bits) succeeds in producing a circuit that has a unique satisfying assignment. Using a PRG for NP||-circuits, it is then possible to deterministically produce a list of candidate circuits from C, one of which is guaranteed to have a unique satisfying assignment. For this circuit C', we can find the satisfying assignment by making the following queries in parallel: "Does C' have a satisfying assignment that assigns x_i true?" and "Does C' have a satisfying assignment that assigns x_i false?" for each i. The overall procedure requires only non-adaptive NP-oracle access. Corollary 3.3 gives us the same consequence from a weaker hardness assumption:

Theorem 3.12. *If $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits, then there is a procedure that, given a circuit C, outputs a satisfying assignment for C if one exists, and runs in polynomial time with* non-adaptive *NP-oracle access.*

Finally, using Corollary 3.3 together with the "hardness amplification" results of [STV01] gives a hardness amplification result for *nondeterministic* circuits: it states that worst-case hardness implies average-case hardness for nondeterministic circuits. This problem is extensively studied for *deterministic* circuits [BFNW93, Imp95, IW97, STV01]; a hardness amplification for nondeterministic circuits was first given in [SU01]. The present route (using Corollary 3.3 and [STV01]) gives a simpler and more modular proof of:

Theorem 3.13 ([SU01]). *Let $\mathcal{C}$ be a complexity class that allows low-degree extension. For every $\epsilon > 0$, if $\mathcal{C}$ is hard for size s nondeterministic circuits then $\mathcal{C}$ is $(1/2 + \epsilon)$-hard for size $s' = (s\epsilon/n)^{\Omega(1)}$ nondeterministic circuits.*

3.3.2 Applications of relative error approximators

In addition to giving a (conditional) derandomization of approximate counting, we obtain the following further application of Theorem 3.6:

Theorem 3.14. *If $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits, then $BPP_{path} = P_{||}^{NP}$.*

The class BPP_{path} was defined by Han, Hemaspaandra and Theirauf [HHT97]. It is the class of languages L for which there exists a non-deterministic polynomial-time Turing Machine M for which

$$x \in L \quad \Rightarrow \quad \text{at least } 2/3 \text{ of the computation paths of } M \text{ accept} \tag{1}$$

$$x \notin L \quad \Rightarrow \quad \text{at least } 2/3 \text{ of the computation paths of } M \text{ reject.} \tag{2}$$

Notice that the computation paths need not all make the same number of non-deterministic choices; if they are required to, we just get BPP. In contrast to BPP, BPP_{path} is quite powerful: it is known to contain $P^{NP}_{||}$ [HHT97]. The above theorem suggests it is probably *equal* to $P^{NP}_{||}$.

Proof of Theorem 3.14. Let L be a language in BPP_{path} with associated non-deterministic Turing Machine M. Let $p(n)$ be an upper bound on the running time of M on an input of length n.

Fix an input x. Given a string $y \in \{0,1\}^{p(|x|)}$ consider the following deterministic procedure: simulate M using successive bits of y as M's non-deterministic choices. When M halts, if the remainder of y is all-zeros, then accept, otherwise reject. Let $D_x(y)$ be the circuit outputting 1 iff this procedure accepts, and let S be the set of strings accepted by D_x.

We consider one other deterministic procedure: given $y \in \{0,1\}^{p(|x|)}$, simulate M using successive bits of y as M's non-deterministic choices and accept if and only if M accepts. Let $C_x(y)$ be the circuit outputting 1 iff this procedure accepts. Observe that the probability over computation paths of M that M accepts input x is exactly:

$$\Pr_y[C_x(y) = 1 | D_x(y) = 1],$$

since each 1 of D_x corresponds to a unique computation path.

We use the one-sided relative-error approximator of Theorem 3.6 twice (in parallel), once with input C_x, and once with input $D_x \wedge C_x$, and $\epsilon = 1/10$. Let ρ_1 and ρ_2 be the two approximations. Notice that

$$(1-\epsilon)\Pr_y[C_x(y) = 1 | D_x(y) = 1] \le (\rho_2/\rho_1) \le (1-\epsilon)^{-1}\Pr_y[C_x(y) = 1 | D_x(y) = 1].$$

We accept iff $\rho_2/\rho_1 > 1/2$, which is guaranteed to happen iff $\Pr_y[C_x(y) = 1 | D_x(y) = 1] \ge 2/3$. The entire procedure runs in time poly$(|x|)$ with non-adaptive NP oracle access. ☐

3.3.3 Applications of conditional discrepancy sets

The class S_2^p was defined by [Can96] and [RS98]. It is the class of languages L for which there is a polynomial-time predicate R for which:

$$x \in L \;\Rightarrow\; \exists y\, \forall z\; R(x,y,z) = 1 \quad (3)$$
$$x \notin L \;\Rightarrow\; \exists z\, \forall y\; R(x,y,z) = 0. \quad (4)$$

Cai [Cai01] recently showed that the class S_2^p (which contains P^{NP} and MA) is contained in ZPP^{NP}. One consequence of this result is that under a hardness assumption sufficient to derandomize ZPP^{NP}, the class S_2^p collapses to P^{NP}. This is remarkable because S_2^p is defined by alternating quantifiers and has more of the flavor of the Polynomial-Time Hierarchy than any randomized complexity class; yet derandomization techniques yield a surprising collapse.

We view Cai's result as a reduction of S_2^p to sampling, and thus obtain the following collapse as an application of Theorem 3.10. Note that this result does *not* follow directly from $S_2^P \subseteq ZPP^{NP}$ using straightforward derandomization techniques – that would require a stronger hardness assumption (given current technology) in order to cope with the adaptive NP-queries used in Cai's algorithm.

Theorem 3.15. *If E^{NP} requires exponential size SV-nondeterministic circuits, then $S_2^p = P^{NP}$.*

Proof. Let L be a language in S_2^p, and let R be the associated polynomial-time predicate for which Eqs. (3) and (4) hold. By padding if necessary we may assume that $|x| = |y| = |z| = n$. Let s be the running time of R.

The procedure to decide if $x \in L$ operates in rounds. Initially, we set $i = 0$, and $S_0 = \{0,1\}^n$, and observe that S_0 is clearly recognized by a trivial circuit C_0. We now begin round 0.

In round i we do the following:

1. In P^{NP}, generate the S_i-conditional $(n, s^3, 1/2)$-discrepancy set $T_i \subseteq \{0,1\}^n$.

2. If $\forall z\; \vee_{t \in T_i} R(x,t,z) = 1$ then accept.

3. Otherwise, find some z_i for which $\vee_{t\in T_i} R(x,t,z_i) = 0$.

4. Define $S_{i+1} = \{y : y \in S_i \wedge R(x,y,z_i) = 1\}$, and observe that S_{i+1} is recognized by a circuit C_{i+1} of size $O(s^2 + |C_i|)$.

5. If $S_{i+1} = \emptyset$, then reject; otherwise, begin round $i+1$.

Notice that step 2 requires a single NP-oracle query, as does step 5, and that step 3 involves finding an NP-witness in the usual way with multiple NP-oracle queries.

The main claim is that the number of rounds before this procedure either accepts or rejects is at most $n+1$. Notice that at step 3, we must have that

$$\Pr_y[R(x,y,z_i) = 1 | y \in S_i] \leq 1/2,$$

since $\Pr_{t\in T_i}[R(x,t,z_i) = 1 | y \in S_i] = 0$ and the circuit computing R with x and z_i hard-wired has size at most $O(s^2) < s^3$, and T_i is an S_i-conditional $(n, s^3, 1/2)$-discrepancy set. Thus $|S_{i+1}| \leq |S_i|/2$ for all i. Since we start with $|S_0| = 2^n$, we have $|S_{n+1}| \leq 1/2$ which implies $|S_{n+1}| = 0$, so we halt after at most $n+1$ rounds.

For correctness, observe that if we accept, we have found that the complement of Eq. (4) holds; if we reject, then $\forall y\ \exists z_i\ R(x,y,z_i) = 0$, and thus the complement of Eq. (3) holds. □

In a similar manner, the result by Bshouty et al. [BCG+96] on learning of circuits using equivalence queries may be regarded as a reduction to sampling. We thus obtain, using Theorem 3.10:

Theorem 3.16. *If E^{NP} requires exponential size SV-nondeterministic circuits, then there is a deterministic procedure with access to an NP-oracle that learns an unknown Boolean circuit C of size s on n inputs in time poly(s,n) using equivalence queries.*

Proof. We use the notation $[y]$ to indicate function computed by the Boolean circuit described by string y. Define the function $R : \{0,1\}^s \times \{0,1\}^n \to \{0,1\}$ by $R(y,z) = [y](z)$.

The learning procedure is very similar to the algorithm in the proof of Theorem 3.15. The procedure operates in rounds. Initially, we set $i = 0$, and $S_0 = \{0,1\}^n$, and observe that S_0 is clearly recognized by a trivial circuit C_0. We now begin round 0.

In round i we do the following:

1. In P^{NP}, generate the S_i-conditional $(s, s^3, 1/4)$-discrepancy set $T_i \subseteq \{0,1\}^s$.

2. Make the equivalence query: "maj $_{t\in T_i} R(t,z) \equiv C(z)$?" If the answer is YES, then we are done.

3. If the answer is NO, then we are given a counterexample z_i for which $v_i = \text{maj}_{t\in T_i} R(t,z_i) \neq C(z_i)$.

4. Define $S_{i+1} = \{y : y \in S_i \wedge R(y,z_i) \neq v_i\}$, and observe that S_{i+1} is recognized by a circuit C_{i+1} of size $O(s^2 + |C_i|)$.

5. Begin round $i+1$.

Note that the procedure must terminate because $|S_{i+1}| < |S_i|$ for all i, and there is some $y \in S_i$ such that $[y] = C$, for all i. As in the proof of Theorem 3.15 then main claim is that the number of rounds before completion is at most $O(s)$. At step 3, we must have

$$\Pr_y[R(y,z_i) = C(z_i) | y \in S_i] \leq 3/4.$$

This is true because $\Pr_{t\in T_i}[R(t,z_i) = v_i | y \in S_i] \geq 1/2$, which implies $\Pr_{t\in T_i}[R(t,z_i) = C(z_i) | y \in S_i] \leq 1/2$. Furthermore, the circuit computing R with z_i hard-wired has size at most $O(s^2) < s^3$, and T_i is an S_i-conditional $(s, s^3, 1/4)$-discrepancy set, which implies $\Pr_y[R(y,z_i) = C(z_i) | y \in S_i] \leq 1/2 + \epsilon = 3/4$, as claimed.

Thus $|S_{i+1}| \leq (3/4)|S_i|$ for all i, and we start with $|S_0| = 2^s$, so we must halt after at most $O(s)$ rounds with a positively answered equivalence query. □

We remark that Theorem 3.15 and Theorem 3.16 are just two examples where a ZPP^{NP} algorithm for sampling is used as a critical subroutine (see, e.g., the discussion in [BGP00] regarding applications in interactive proofs). Often this is the *only* randomness used in these procedures, and so conditional discrepancy sets suffice for derandomization in a variety of settings.

4 Overview of the techniques

In this section we present the main technical ideas in this paper in an informal manner; the full proofs are in the following two sections.

4.1 Proof of the downward collapse theorem

We show in Theorem 3.2 that for every sufficiently strong complexity class $\mathcal{C}$, if $\mathcal{C}$ is computable by small NP||-circuits then $\mathcal{C}$ is computable by small SV-nondeterministic circuits. This certainly does not mean that one can always transform small NP||-circuits into small SV-nondeterministic circuits. In particular, there are small NP||-circuits for Satisfiability and we do not expect Satisfiability to have small SV-nondeterministic circuits, as this would mean that $NP \subseteq coNP/poly$ and collapse the polynomial hierarchy. Indeed, this observation demonstrates the main problem we need to overcome. Whenever an NP||-circuit calls its NP-oracle, it gets a result no matter whether the query asked is answered positively or negatively. An SV-nondeterministic circuit can attempt to simulate an NP||-circuit by guessing which queries are answered positively, together with witnesses for those queries – in this way it can "verify" some queries that are answered positively. But it can not be sure that it has correctly guessed *all* of the positively answered queries, precisely because it is incapable of verifying negative answers (assuming $NP \not\subseteq coNP/poly$).

The main idea in the proof is that when the function to be computed is a low degree multivariate polynomial, a small SV-nondeterministic circuit *can* in fact verify negative answers, in an indirect way. Every function in a sufficiently strong class $\mathcal{C}$ has a multivariate polynomial "low-degree extension" [BF90] that lies in the same class. Thus the trick that allows SV-nondeterministic circuits to simulate NP||-circuits on low-degree polynomials implies the existence of small SV-nondeterministic circuits for all functions in class $\mathcal{C}$ if the class has small NP||-circuits.

We now describe the trick that exploits the low-degree extension[7]. We're given a small NP||-circuit which computes some low degree multivariate polynomial $f : F^d \to F$ (for some field F of size q). For simplicity, let's assume that this circuit makes a single NP-query. We now describe how we construct a small SV-nondeterministic circuit for f. For every input x in the domain of f, let $A(x)$ denote the answer to the NP-query asked on x. Let p denote the fraction of x's for which the query is answered positively. We hardwire p to our SV-nondeterministic circuit. Now, on input x the new circuit passes a random low degree curve through x (we denote the degree of this curve by r). Except for x, the other q points on this curve are r-wise independent and therefore with high probability the fraction of points y on the curve for which $A(y) = 1$ is in the range $(p - \delta, p + \delta)$ for some small δ.[8] The circuit now guesses $(p - \delta)q$ points on the curve along with witnesses showing that the queries corresponding to these points are answered positively. It assumes that these queries are answered positively and the queries for the remaining points on the curve are answered negatively. The critical observation is that this assumption can be incorrect on at most a 2δ fraction of the points on the curve. The circuit now simulates the NP||-circuit (which makes no further NP queries) on all q points on the curve, and the final evaluations it receives differ from the correct evaluations on at most $2\delta q$ points. Finally, because the function f restricted to the curve is a low-degree polynomial, the circuit can run a decoding algorithm for Reed-Solomon codes [WB86] to correct the errors and obtain the correct answers for all points on the curve, and in particular the circuit obtains $f(x)$.

4.2 Building relative-error approximators

Our relative-error approximators build on a line of work which gives probabilistic algorithms that use an NP-oracle to approximately count NP-witnesses [Sto83, JVV86, BGP00] (for more information see the discussion in [BGP00]). Such algorithms are given a deterministic circuit A on n bits and wish to produce a relative approximation of the size of the set $S = \{x|A(x) = 1\}$. The algorithms work by finding a hash function $h : \{0,1\}^n \to \{0,1\}^k$ with the property that for every image $y \in \{0,1\}^k$ the size of the preimage $S_y = \{x \in S|h(x) = y\}$ is roughly n^2, which implies that $|S|$ is approximately $n^2 2^k$.

[7] A similar idea was used in [SU01] to build PRGs for nondeterministic circuits. It may also be viewed as a non-trivial "scaling down" of $EXP^{NP}_{||} \subseteq NEXP/poly \cap coNEXP/poly$ – a containment credited to Harry Buhrman on Lance Fortnow's weblog.

[8] By choosing the degree r large enough we can show that there exist *fixed* points $v_1, \cdots, v_r \in F^d$ such that for every x the fraction of points y such that $A(y) = 1$ on the degree r curve that passes through $x; v_1, \cdots, v_r$ is in the range $(p - \delta, p + \delta)$. In the final construction we also hardwire the points $v_1, \cdots, v_r$ to the circuit.

To find such a hash function, we choose a random hash function $h : \{0,1\}^n \to \{0,1\}^k$ from an n-wise independent hash family, and use the NP oracle to check whether there exists a $y \in \{0,1\}^k$ whose preimage has size greater than n^2. We do this for $k = 1, 2, 3, \ldots$, stopping with the first h which is good in the sense that there does not exist such a y whose preimage is "too large". By the pigeonhole principle, a good h does not exist for k such that $n^2 2^k < |S|$; for slightly larger k a random h from the n-wise independent hash family is good with high probability. Thus, the algorithm stops with the "correct" value of k, with high probability.

We would like to derandomize this procedure. Since it is not a decision problem we cannot use PRGs directly[9]. Instead we derandomize this procedure by using the particular way it operates (a general method that has been suggested by [KvM02] for such circumstances). Rather than choosing the hash functions randomly, we try all of the hash functions that are described by outputs of a PRG for nondeterministic circuits. For the "correct" k, one of the hash functions we try is good, because the generator fools the nondeterministic circuit which, given h, checks whether it is good. In addition, some care must be taken to obtain less-coarse approximations, and to ensure that the overall procedure runs in $FP_{||}^{NP}$, rather than FP^{NP}.

4.3 Constructing conditional discrepancy sets

An S-conditional discrepancy set for small circuits is a set $T \subseteq S$ such that no small (deterministic) circuit can distinguish a random element from T from a random element in S. This generalizes "regular" discrepancy sets for small circuits (for which the set S is simply $\{0,1\}^n$). Given a set S, encoded by a circuit A such that $S = \{x|A(x) = 1\}$, our goal is to output an S-conditional discrepancy set T.

As with relative-error approximation, our approach is based on algorithms which uses an NP-oracle to sample (or count) accepting inputs of A [Sto83, JVV86, BGP00]. Fix a hash function $h : \{0,1\}^n \to \{0,1\}^k$ which is good in the sense defined above. To sample a random element from S, one can choose a random image y, use the NP oracle to find all the preimages of y (there are approximately n^2 of them), and choose a random one.

Our procedure for producing conditional discrepancy sets is a derandomization of this algorithm, that uses a PRG for NP||-circuits. We first deterministically find a good hash function h as explained above. Then, we include in the conditional discrepancy set T the preimages of *only* those y that are outputs of a PRG G for NP||-circuits; here we make critical use of Theorem 3.12 to perform this step using only non-adaptive NP oracle access.

The proof that T is in fact an S-conditional discrepancy set is somewhat subtle. Given a (deterministic) circuit that distinguishes a random element in T from a random element in S, we need to construct a NP||-circuit D that is a distinguisher for the PRG G, thus leading to a contradiction. Care is needed to ensure that the distinguisher D makes only non-adaptive NP oracle queries – and this is especially crucial here because a distinguisher that makes adaptive queries is not guaranteed to be fooled by the PRG G that is based on only an SV-nondeterministic hardness assumption.

5 Proof of Theorem 3.2 and its applications

We begin with some definitions and preliminaries.

5.1 Preliminaries

Given a function $f : X \to Y$ and $S \subseteq X$ we use $f(S)$ to denote the (multi-)set $\{f(x)|x \in S\}$.

5.1.1 Discrepancy sets and pseudorandom generators

In this paper we define pseudorandom generators in terms of discrepancy sets.

[9]For the case of decision problems every probabilistic algorithm can be derandomized if one has a sufficiently strong pseudorandom generator. However, there are tasks (which are not decision problems) that can be easily solved by a probabilistic algorithm and cannot be solved by a deterministic algorithm. For example, a probabilistic algorithm can easily produce a string with high Kolmogorov complexity whereas no deterministic algorithm can output such a string.

Definition 5.1 (discrepancy set). *Let $\mathcal{D}$ be a subset of all functions from $\{0,1\}^n$ to $\{0,1\}$. A set $T \subseteq \{0,1\}^n$ is an* (n, ϵ)-discrepancy set *for $\mathcal{D}$ if for every $D \in \mathcal{D}$,*

$$\left| \Pr_{x \in \{0,1\}^n}[D(x) = 1] - \Pr_{t \in T}[D(t) = 1] \right| \leq \epsilon.$$

Commonly $\mathcal{D}$ is the set of functions with size s deterministic circuits; in this case we use the shorthand (n, s, ϵ)-*discrepancy set* (as in Subsection 3.2.2). A pseudorandom generator is a function whose output is a discrepancy set[10].

Definition 5.2 (pseudorandom generator). *Let $\mathcal{C}$ be a complexity class. A pseudorandom generator (PRG) for $\mathcal{C}$ is a procedure which on input 1^n outputs a $(n, 1/n)$-discrepancy set for the set $\mathcal{D}$ of all characteristic functions of languages in $\mathcal{C}$ restricted to length n.*

In this paper $\mathcal{C}$ will typically be the class of those languages with nondeterministic circuits of a given type, and whose size is a fixed polynomial.

5.1.2 Low-degree polynomials

The low-degree extension of a function embeds the function in a low-degree polynomial.

Definition 5.3 (low-degree extension). *Let $f : \{0,1\}^n \to \{0,1\}$ be a function, $\mathbb{F}_q$ the field with q elements, and h and d integers for which $h^d \geq 2^n$. Let H be a subset of $\mathbb{F}_q$ of size h, and let I be an efficiently computable injective mapping from $\{0,1\}^n$ to H^d.*

The low-degree extension of f with respect to q, h, d *is the (unique) d-variate polynomial $\hat{f} : \mathbb{F}_q^d \to \mathbb{F}$ with degree $h - 1$ in each variable, for which $\hat{f}(I(x)) = f(x)$ for all $x \in \{0,1\}^n$ and $\hat{f}(v) = 0$ for $v \in (H^d \setminus Im(I))$.*

It is often helpful to think of field elements as binary strings of length $\log q$. From this viewpoint, $\hat{f}$ is a function from $d \log q$ bits to $\log q$ bits. We will often consider a version of the low degree extension which outputs a single bit. This *boolean version* of the low-degree extension is denoted $\hat{f}_{bool} : \{0,1\}^{d \log q + \log \log q} \to \{0,1\}$ and is defined by $\hat{f}_{bool}(x, i) = \hat{f}(x)_i$.

The following properties of low-degree extensions are trivial and standard:

Proposition 5.4 (properties of the low-degree extension). *For $\hat{f}$ and $\hat{f}_{bool}$ as defined above, the following hold:*

- *$\hat{f}$ has total degree hd, and*
- *$\hat{f}_{bool}$ is computable in time poly$(2^n, \log q, d)$ given oracle access to f.*

Complexity classes that allow low-degree extension (see Definition 3.1) contain the (boolean) low-degree extensions of every function in that class; Theorem 3.2 applies to all such classes.

Definition 5.5 (parametric curves). *Let F_q be the field with q elements, and let $f_1, f_2, \ldots f_q$ be an enumeration of the elements of F_q. Given $v_1, v_2, \ldots, v_r \in \mathbb{F}_q^d$, for $r \leq q$, we define the* curve passing through *$v_1, v_2, \cdots, v_r$ to be the unique degree $r - 1$ polynomial function $c : \mathbb{F}_q \to \mathbb{F}_q^d$ for which $c(f_i) = v_i$ for all i. A curve c is* one to one *if $i \neq j$ implies $c(f_i) \neq c(f_j)$.*

The function $\hat{f} \circ c$ is the *restriction of $\hat{f}$ to the curve c*. It is a low-degree univariate polynomial; in coding terms, it is a Reed-Solomon codeword.

Theorem 5.6 (decoding of Reed-Solomon codes [WB86]). *Let $\mathbb{F}_q$ be the field with q elements. Given t pairs (x_i, y_i) of elements of $\mathbb{F}_q$, there is a unique polynomial $g : \mathbb{F}_q \to \mathbb{F}_q$ of degree at most u for which $p(x_i) = y_i$ for at least a pairs, provided $a > (t + u)/2$. Furthermore, there is a polynomial time algorithm that finds g.*

---

[10] A more standard formulation is that a pseudorandom generator "stretches" a short seed into a long pseudorandom string, with the property that the set of all pseudorandom strings is a discrepancy set. Our definition asks the pseudorandom generator to output all pseudorandom strings at once. This difference is immaterial in in this paper as we will be concentrating on discrepancy sets with polynomial size, and thus the entire set can be output in polynomial time if each individual string can be generated in polynomial time.

5.2 Random curves that pass through a fixed point

In this subsection we prepare some technical machinery needed for the proof of Theorem 3.2. We will repeatedly use the following tail-inequality for r-wise independent random variables:

Lemma 5.7 ([BR94]). *Let $r > 4$ be an even integer. Suppose $X_1, X_2, \ldots, X_q$ are r-wise independent random variables taking values in $[0,1]$. Let $X = \sum X_i$, and $A > 0$. Then:*

$$\Pr[|X - E[X]| \geq A] \leq 8 \cdot \left(\frac{r \cdot E[X] + r^2}{A^2}\right)^{r/2}.$$

We prove a technical lemma regarding the sampling properties of low-degree parametric curves. The points on a random degree r parametric curve are r-wise independent; a well-known consequence of this fact (using, e.g., Lemma 5.7) is that the points on such a curve are a good "oblivious sampler" (see the survey [Gol97]). This means that for any function $h : F^d \to [0,1]$ the average of $h(x)$ over the points on a random curve is with high probability close to the average over the whole space. We show below that this holds even if an adversary gets to choose the first point on the curve. Because the remaining points on the curve are still r-wise independent it remains a good sampler.

We need the following notation:

Definition 5.8. *Let $W \subseteq Z$ be finite sets and let $h : Z \to [0,1]$ be an arbitrary function. The* average of h over W *is defined by:*

$$\mu_W(h) = \frac{1}{|W|}\sum_{i \in W} h(i)$$

We will use $c_{(x,v_1,v_2,\ldots,v_r)}$ to denote the curve passing through $x, v_1, v_2, \ldots v_r$ (see Definition 5.5). We require that $c_{(x,v_1,v_2,\ldots,v_r)}(0) = x$; i.e., the enumeration of the field elements in Definition 5.5) starts with 0. Also, below $\mathbb{F}_q$ is the field of size q, and $\mathbb{F}_q^* = \mathbb{F}_q \setminus \{0\}$.

Lemma 5.9. *Let r be an integer for which $2 \leq r < q$. For every point $x \in \mathbb{F}_q^d$, function $h : \mathbb{F}_q^d \to [0,1]$, and $\delta > 0$, the following hold:*

1. $\Pr_{v_1,\ldots,v_r \in \mathbb{F}_q^d}\left[\left|\mu_{c_{(x,v_1,\ldots,v_r)}(\mathbb{F}_q^*)}(h) - \mu_{\mathbb{F}_q^d}(h)\right| \geq \delta\right] \leq 8 \cdot \left(\frac{2r}{t\delta^2}\right)^{r/2}$*, and*
2. $\Pr_{v_1,\ldots,v_r \in \mathbb{F}_q^d}\left[c_{(x,v_1,\ldots,v_r)}\ \text{isn't one to one}\right] \leq \frac{1}{q^{d-2}}$.

Proof. Fix x and h, and let $v_1, \ldots, v_r$ be chosen uniformly and independently from $\mathbb{F}_q^d$. Define random variables Y_a by $Y_a = c_{(x,v_1,\ldots,v_r)}(a)$. It is standard that for every $a \in F_q^*$, Y_a is uniformly distributed over $\mathbb{F}_q^d$, and that the random variables $\{Y_a\}_{a \in F_q^*}$ are r-wise independent. Now we define the random variables $R_a = h(Y_a)$. It follows that for every $a \in F_q^*$, $E[R_i] = \mu_{\mathbb{F}_q^d}(h)$, and that $\{R_a\}_{a \in F_q^*}$ are r-wise independent. Let $R = \sum_{a \in F_q^*} R_a$. We apply Lemma 5.7 with $A = |F_q^*|\delta = (q-1)\delta$ to conclude:

$$\Pr_{v_1,\cdots,v_r \in \mathbb{F}_q^d}\left[\left|\mu_{c_{(x,v_1,\ldots,v_r)}(\mathbb{F}_q^*)}(h) - \mu_{\mathbb{F}_q^d}(h)\right| \geq \delta\right] = \Pr[|R - E[R]| \geq A] \leq 8 \cdot \left(\frac{A + r^2}{A^2}\right)^{r/2} \leq 8 \cdot \left(\frac{2r}{(q-1)\delta^2}\right)^{r/2}.$$

This proves (1). For (2), we observe that for every $a \neq a' \in \mathbb{F}_q$,

$$\Pr_{v_1,\ldots,v_r \in \mathbb{F}_q^d}[c_{(x,v_1,\ldots,v_r)}(f) = c_{(x,v_1,\ldots,v_r)}(f')] = \frac{1}{q^d},$$

and taking a union bound over all (at most q^2) such pairs yields the desired result. □

We will be interested in curves that are good samplers for k functions simultaneously. The following is a corollary of the above lemma; it is an easy application of a union bound:

Corollary 5.10. *Let r be an integer for which $2 \le r < q$. Let $h_1, h_2, \ldots, h_k$ be functions from $\mathbb{F}_q^d$ to $[0,1]$. For every point $x \in \mathbb{F}_q^d$ and $\delta > 0$, the probability over a random choice of points $v_1, \ldots, v_r \in \mathbb{F}_q^d$ that $c_{(x,v_1,\ldots,v_r)}$ is one-to-one and*

$$\left|\mu_{c_{(x,v_1,\ldots,v_r)}(\mathbb{F}_q^*)}(h_i) - \mu_{\mathbb{F}_q^d}(h_i)\right| < \delta$$

for all $1 \le i \le k$, is at least

$$1 - \left(8k\left(\frac{2r}{(q-1)\delta^2}\right)^{r/2} + \frac{1}{q^{d-2}}\right).$$

5.3 Proof of the downward collapse theorem

In this subsection we prove Theorem 3.2. We refer the reader to the informal description of the technique in the introduction (Section 4.1).

Let L be an arbitrary language in $\mathcal{C}$, and let $f : \{0,1\}^n \to \{0,1\}$ be the restriction of (the characteristic function of) L to inputs of length n. Throughout the proof we assume that n is sufficiently large, $n \le s(n) \le 2^n$, and that $s(O(n)) \le s(n)^{O(1)}$.

Let $\hat{f}$ be the low-degree extension of f with respect to parameters q, h, d chosen as follows (they are expressed in terms of a fourth parameter r):

- $r = 2(n + \log(32s(n)^5))$
- $h = (4r)^2(9s(n))^4$
- $d = \lceil n/\log h\rceil + 3$
- q smallest prime power larger than $9hdr$.

Note that $\mathcal{C}$ allows low-degree extension, and so by Proposition 5.4, the function family consisting of (boolean versions of) the low-degree extensions of L for each input length, with parameters as defined above, lies in $\mathcal{C}$.

Thus, by the hypothesis of the theorem, $\hat{f}_{bool}$ has an NP$||$-circuit of size $s(n')$, where $n' = \log(q^d) + \log(q) = O(n)$ is the input length of $\hat{f}_{bool}$. We will construct a *probabilistic* SV-nondeterministic circuit C' computing $\hat{f}_{bool}$ of size $s' = s(n')^c$, for a constant c (it will be clear in the exposition below what is meant by a "probabilistic SV-nondeterministic circuit"). We will then transform C' into an SV-nondeterministic circuit C'' computing f by fixing a "good" random string, and using the function I that accompanies the low-degree extension (recall Definition 5.3). The resulting circuit C'' will have size $s(n')^c + poly(n)$. Since $s(n')^c = s(O(n))^c = s(n)^{O(1)}$, we will conclude that L has circuits of size $s(n)^{O(1)}$. As L was arbitrary, this will prove the theorem.

Let C_{pre}, C_{post} be the Boolean circuits that describe the NP$||$-circuit of size $s(n')$ that computes $\hat{f}_{bool}$ (recall Definition 2.3). With $\log q$ parallel copies of C_{pre} and C_{post}, we can construct an NP$||$-circuit with $\log q$ outputs that computes $\hat{f}$. Let $Q_1(x), \ldots, Q_k(x)$ and $A_1(x), \ldots, A_k(x)$ be the queries and answers associated with this circuit, respectively, on input $x \in \mathbb{F}_q^d$. Without loss of generality we assume that exactly k queries are made on every input x. We define $p_i = \mu_{\mathbb{F}_q^d}(A_i)$.

We focus first on constructing C', the probabilistic SV-nondeterministic circuit. Circuit C' makes use of C_{pre} and C_{post}, as well as $p_1, p_2, \ldots, p_k$ as non-uniform advice. We set $\delta = 1/(9k)$. On input (x, b), circuit C' wants to compute $\hat{f}_{bool}(x, b)$; it performs the following steps:

- Pick $v_1, v_2, \ldots, v_r \in \mathbb{F}_q^d$ uniformly at random, and set $x_a = c_{(x,v_1,v_2,\ldots,v_r)}(a)$, so the x_a are the q points along a random curve passing through $x, v_1, v_2, \ldots v_r$. Simulate C_{pre} to compute queries $Q_i(x_a)$ for $1 \le i \le k$ and $a \in \mathbb{F}_q^*$.
- Set $n_i = \lfloor(p_i - \delta)(q-1)\rfloor$. For $1 \le i \le k$, guess $z_i \in \{0,1\}^{\mathbb{F}_q^*}$ with exactly n_i ones, and strings $\{w_{i,a}\}_{a \in \mathbb{F}_q^*}$.
- For $1 \le i \le k$ and $a \in \mathbb{F}_q^*$, check that $(z_i)_a = 1$ implies $w_{i,a}$ is a witness that query $Q_i(x_a)$ is answered positively; otherwise, set the **flag** output to 0 and halt.

- Compute $y_a = C_{post}(x_a, (z_1)_a, (z_2)_a, \ldots, (z_k)_a)$ for $a \in F_q^*$.
- Run the algorithm of Theorem 5.6 on the $q-1$ pairs (f_a, y_a) with $u = hdr$ to obtain a polynomial $g : \mathbb{F}_q \to \mathbb{F}_q$ of degree u. Set the **value** output to the b-th bit of $g(0)$, and set the **flag** output to 1.

The following claim will allow us to fi x the coin-flips of circuit C', described above, to get an SV-nondeterministic circuit computing f.

Claim 5.10.1. *For every $x \in \mathbb{F}_q^d$ and $b \in [\log q]$, with probability at least $1 - \frac{2^{-n}}{2\log q}$ over the choice of $v_1, \ldots, v_r$, the following two conditions hold:*

1. *For all guesses $z_i, w_{i,a}$ for which the* **flag** *output is set to one, the* **value** *output is $\hat{f}_{bool}(x, b)$.*
2. *There exist guesses $z_i, w_{i,a}$ such that the* **flag** *output is set to one.*

Proof. Fix an $x \in \mathbb{F}_q^d$. We apply Corollary 5.10 to conclude that the probability over a random choice of points $v_1, \ldots, v_r \in \mathbb{F}_q^d$ that

$$c_{(x,v_1,\ldots,v_r)} \text{ is one-to-one and } \left|\mu_{c_{(x,v_1,\ldots,v_r)}(\mathbb{F}_q^*)}(A_i) - \mu_{\mathbb{F}_q^d}(A_i)\right| < \delta \text{ for all } 1 \le i \le k \tag{5}$$

is at least

$$1 - \left(8k\left(\frac{2r}{(q-1)\delta^2}\right)^{r/2} + \frac{1}{q^{d-2}}\right).$$

By our choice of parameters:

$$\left(8k\left(\frac{2r}{(q-1)\delta^2}\right)^{r/2} + \frac{1}{q^{d-2}}\right) \le 8s(n)\log q\left(\frac{1}{2}\right)^{r/2} + \frac{1}{q^{d-2}} \le \frac{2^{-n}}{4\log q} + \frac{2^{-n}}{4\log q} \le \frac{2^{-n}}{2\log q}.$$

The fi rst inequality it true because $k \le s(n)\log q$, $\delta^{-2} = (9k)^2 \le (9s(n)\log q)^2$ and

$$(q-1)/\log^2 q \ge \sqrt{q} \ge \sqrt{h} \ge (4r)(9s(n))^2$$

(for suffi ciently large q). The second inequality follows from our choice of r and d, and the fact that $\log q = O(n) \le s(n)^2$ (for suffi ciently large q).

We will show that whenever (5) holds, the two items in the claim hold. We begin with the second item. Since (5) holds, for each i we know that there are at least n_i distinct indices for which $A_i(x_a) = 1$; we choose z_i to be a string with ones in exactly n_i of these indices. For each index a for which $(z_i)_a = 1$, there is a witness $w_{i,a}$ showing that query $Q_i(x_a)$ is answered positively (since $A_i(x_a) = 1$). Thus there exists a choice of of the $z_i, w_{i,a}$ for which the **flag** output is set to one.

Now, we turn to the fi rst item. Once the verifi cation in the third bullet above is complete, we know that for all i, and all $a \in F_q^*$, $(z_i)_a = 1$ implies $A_i(x_a) = 1$, and that there are at least n_i such a for which $(z_i)_a = 1$. We also know, by (5), that the number of a for which $A_i(x_a) = 1$ is at most $\lceil (p_i + \delta)(q-1)\rceil$. Thus we can bound the number of "errors attributable to query i" as follows:

$$|\left\{a : a \in \mathbb{F}_q^*, A_i(x_a) \neq (z_i)_a\right\}| \le \lceil (p_i+\delta)(q-1)\rceil - \lfloor (p_i-\delta)(q-1)\rfloor \le 2\delta q,$$

and the number of "errors" overall as follows:

$$|\left\{a : a \in \mathbb{F}_q^* \text{ for which } \exists i\ A_i(x_a) \neq (z_i)_a\right\}| \le 2\delta qk.$$

For every a that is not an "error," $y_a = \hat{f}(x_a)$. We conclude that for at least $(q-1) - 2\delta qk = (1-2\delta k)q - 1$ of the pairs (a, y_a), we have $y_a = p(a)$, where $p(w)$ is the degree hdr "restriction to the curve" $p(w) = \hat{f} \circ c_{(x_1,v_1,v_2,\ldots,v_r)}(w)$.

If the number of pairs that agree with $p(w)$ is greater than $(q-1+hdr)/2$, then the algorithm of Theorem 5.6 returns $p(w)$, and our circuit outputs the b-th bit of $p(0) = \hat{f}(x)$ as desired. Thus to conclude the proof we verify that

$$(1-2\delta k)q - 1 = \frac{7/9}{q} - 1 > \frac{q-1+hdr}{2},$$

which holds by our choice of q. ∎

Now, recall that the low-degree extension is accompanied by a polynomial-time computable function I from $\{0,1\}^n$ into $\mathbb{F}_q^d$. Consider the set of inputs to C' given by

$$S = \{(x, b) : x \in I(\{0,1\}^n), b \in [\log q]\}$$

and note that $|S| = (\log q)2^n$. Thus there must be a fixing of the coin-flips of C' so that the two statements in the above claim hold for all inputs in S.

Our SV-nondeterministic circuit C'' computing f is built as follows:

- on input $y \in \{0,1\}^n$, compute $x = I(y)$
- use circuit C' with the "good" random coin-flips hardwired to compute $\hat{f}_{bool}(x, b)$ for every $b \in [\log q]$.
- these $\log q$ bits give us $\hat{f}(x) = \hat{f}(I(y)) = f(y)$. Output $f(y)$.

Because non-adaptive queries to an SV-nondeterministic circuit may be simulated by an SV-nondeterministic circuit, the resulting circuit C'' is an SV-nondeterministic circuit. Finally, we can verify that its size is poly(n) + $s(n')^c$ for some constant c. This concludes the proof of Theorem 3.2.

5.4 Application of Theorem 3.2: Derandomizing $BPP_{||}^{NP}$

The following is a slight refinement of a theorem in [KvM02] (we use the additional fact that the NP oracle access in their argument is always non-adaptive):

Theorem 5.11. *[KvM02] If $E_{||}^{NP}$ (resp. E) requires exponential size NP$||$-circuits then there is a PRG for linear-size NP$||$-circuits that runs in polynomial time with non-adaptive access to an NP oracle (resp. polynomial time).*

We obtain the following improvement:

Theorem 5.12. *If $E_{||}^{NP}$ (resp. E) requires exponential size SV-nondeterministic circuits then there is a PRG for linear-size NP$||$-circuits that runs in polynomial time with non-adaptive access to an NP oracle (resp. polynomial time).*

Proof. Combine Theorem 5.11 with Corollary 3.3. □

We use this to prove Theorem 3.11.

Proof of Theorem 3.11. Given a $BPP_{||}^{NP}$ algorithm $A(x, y)$ for language L and an input x let $C_x(Y) = A(x, y)$. By padding if necessary, C_x can be computed by a linear-size NP$||$-circuit. We run the PRG of Theorem 5.12 on input $1^{|x|}$ to produce a discrepancy set T that fools circuit C_x. We compute $A(x, t)$ for each $t \in T$, and output the majority. This constitutes a deterministic algorithm that decides language L in polynomial time with non-adaptive access to an NP oracle. □

Also, as explained in the introduction, Theorem 3.11 gives an alternative way of constructing PRGs for nondeterministic circuits from an SV-nondeterministic hardness assumption. This permits the use of "standard constructions" in this setting, whereas previous constructions [MV99, SU01] were specialized to the nondeterministic case.

5.5 Application of Theorem 3.2: hardness amplification

Hardness amplification results transform functions which are hard on the worst case into functions which are hard on the average. In a sequence of works [BFNW93, Imp95, IW97, STV01] it was shown that for every class which allows low degree extension if the class is hard on the worst case for small *deterministic* circuits then the class is hard on average for small *deterministic* circuits. We restate the best such results (which is by [STV01]) in the following way:[11]

[11] We remark that [STV01] is concerned with list-decoding the Reed-Solomon code. The statement given here is less general.

Theorem 5.13. *[STV01] Let $\mathcal{C}$ be a class which allows low degree extension. There exists a constant c such that for every function $f : \{0,1\}^n \to \{0,1\}$ such that $f \in \mathcal{C}$ and $\epsilon > 2^{-n}$ there is a function $\hat{f} : \{0,1\}^{n'=O(n)} \to \{0,1\}$ such that $\hat{f} \in \mathcal{C}$ and for every function $D : \{0,1\}^{n'} \to \{0,1\}$ such that*

$$\Pr_{x \in_R \{0,1\}^{n'}}[D(x) = \hat{f}(x)] \geq 1/2 + \epsilon$$

there is an oracle circuit C such that C^D computes f, and the size of C is $(n/\epsilon)^c$.

Indeed, this is a complicated way to say that if f is hard for size s deterministic circuits then $\hat{f}$ is hard on average for slightly smaller deterministic circuits. We chose to state Theorem 5.13 this way, because in this form it also gives a hardness amplification result for other classes of circuits. The following corollary is an example.

Corollary 5.14. *Let $\mathcal{C}$ be a class which allows low degree extension. If $\mathcal{C}$ is hard for size s NP$||$-circuits then $\mathcal{C}$ is $1/2+\epsilon$-hard for size $s' = (s\epsilon/n)^{\Omega(1)}$ NP$||$-circuits.*

Proof. One only has to notice that if D is a size s' NP$||$-circuits then C^D (from Theorem 5.13) is a size $s' \cdot poly(n, 1/\epsilon)$ NP$||$-circuit. □

It is important to note that this argument does not work directly for nondeterministic circuits. The reason is that it does not follow that if D is a nondeterministic circuit and C is a deterministic circuit then C^D is a nondeterministic circuit. (Consider for example the case where D computes SAT and C flips the result. The circuit C^D computes $coSAT$ which is not believed to be computable by a nondeterministic circuit.) A hardness amplification result for nondeterministic circuits was proven in [SU01]. Corollary 3.3 gives an alternative proof of this result (Theorem 3.13).

Proof of Theorem 3.13. By Corollary 3.3 we have that there exists a function $\hat{f} : \{0,1\}^{O(n)} \to \{0,1\}$ in $\mathcal{C}$ that is hard for size $s' = s^{\Omega(1)}$, and the Theorem follows from Corollary 5.14 which gives that $\mathcal{C}$ is hard on average even for NP$||$-circuits. □

The proof above makes two consecutive steps of the low degree extension (one in Corollary 3.3 and the other in Theorem 5.13). This is not necessary. We mention that the arguments of the two steps above can be combined to give a more direct proof. We defer this to a later version of the paper.

6 Proofs of Theorem 3.6 and Theorem 3.10

We will need the following lemma:

Lemma 6.1. *Let $H_{n,k}$ be an n-wise independent family of hash functions mapping n bits to k bits, and let $S \subseteq \{0,1\}^n$. Then for every $1 \geq \delta > 0$, and sufficiently large n:*

$$\Pr_{h \in H_{n,k}}\left[\exists y \text{ for which } |\{x : h(x) = y \wedge x \in S\}| > (1+\delta)\frac{|S|}{2^k}\right] \leq 1/2,$$

provided $2^k \leq \delta^2 n^{-3}|S|$.

Proof. Fix $y \in \{0,1\}^k$, and let I_x be the indicator random variable for the event $h(x) = y$. Notice that $\mathrm{E}[I_x] = 2^{-k}$ and that the I_x are n-wise independent. Define $I = \sum_{x \in S} I_x$; we have $\mathrm{E}[I] = |S|2^{-k}$ by linearity of expectation. Applying Lemma 5.7, we get:

$$\begin{aligned} \Pr\left[|\{x : h(x) = y \wedge x \in S\}| > (1+\delta)\frac{|S|}{2^k}\right] &\leq \Pr\left[I - \mathrm{E}[I] \geq \delta\mathrm{E}[I]\right] \leq 8 \cdot \left(\frac{n\mathrm{E}[I] + n^2}{(\delta \mathrm{E[I]})^2}\right)^{n/2} \\ &\leq 8 \cdot \left(\frac{2n}{\delta^2 \mathrm{E[I]}}\right)^{n/2} \leq 8 \cdot \left(\frac{2}{n^2}\right)^{n/2} < 2^{-(n+1)}. \end{aligned}$$

Applying a union bound over all $2^k < 2^n$ different y, we obtain the stated result. □

The main procedure that is used in the proofs of Theorem 3.10 and Theorem 3.6 makes use of a PRG that fools non-deterministic circuits. It takes a circuit C that accepts a subset S of $\{0,1\}^n$, and outputs a hash function from n bits to k bits whose preimages partition S nearly evenly. Additionally, it outputs an S-conditional discrepancy set.

Lemma 6.2. *There is a function family that takes as input:*

- *a circuit C on n bits, and*
- *$\delta > 0$, and*
- *an integer s, and*
- *the truth table T of a function on $t = O(\log |C|, \log n, 1/\delta, \log s)$ bits that cannot be computed by NP$||$-circuits of size $2^{\gamma t}$ for some constant $\gamma > 0$,*

and outputs:

- *an integer k, and*
- *a hash function $h : \{0,1\}^n \to \{0,1\}^k$, and*
- *an integer B with $B =$ poly$(n, 1/\delta)$, and*
- *a multiset R*

for which the following hold:

- $\forall y \; |\{x : h(x) = y \wedge C(x) = 1\}| \le (1+\delta)B$*, and*
- $2^k B \le |C^{-1}(1)|$*, and*
- *the multiset $S = \{x : h(x) \in R \wedge C(x) = 1\}$ is a $C^{-1}(1)$-conditional $(n, s, 2\delta)$-discrepancy set.*

This function family is in $FTIME(2^{O(t)})^{NP}_{||}$.

Proof. Set $N = \delta^{-2} n^3$. We observe that for all $1 \le k \le n$ and all $0 \le e < N$, given a description of some $h \in H_{n,k}$ we can test if

$$\exists y \text{ for which } |\{x : h(x) = y \wedge C(x) = 1\}| > (1+\delta)(N+e) \qquad (6)$$

in nondeterministic time $m = \text{poly}(N, |C|)$. Using known constructions (e.g., [SU01]) we can produce from T a $(m, 1/4)$-discrepancy set $U \subseteq \{0,1\}^m$ for the set of nondeterministic circuits of size m. With known constructions, an enumeration of U can be computed efficiently and deterministically from T. Let M_k be an efficiently computable mapping from strings of length m to $H_{n,k}$ such that M_k is uniform on $H_{n,k}$ when its input is chosen uniformly.

For all triples (k, e, u) with $k = 1, 2, \ldots n$, $e = 0, 1, \ldots N-1$, and $u \in U$, we test whether Eq. (6) holds for hash function $h = M_k(u)$. This entails poly$(N, |C|)$ parallel NP queries altogether. We label each of these queries with a triple (k, e, u). Order the queries lexicographically (with k changing the slowest), and let (k^*, e^*, u^*) be the first triple for which Eq. (6) *does not* hold for $h = M_{k^*}(u^*)$. We claim that

$$(N + e^*) \le \frac{|C^{-1}(1)|}{2^{k^*}} \le (1+\delta)(N+e^*). \qquad (7)$$

This is true because: by the pigeonhole principle, if $|C^{-1}(1)|/2^{k^*} > (1+\delta)(N+e^*)$, then Eq. (6) holds; and if $(N+e^*) > |C^{-1}(1)|/2^{k^*}$ then $k^* > \lfloor \log_2(|C^{-1}(1)|/N) \rfloor$ and so it must be the case that Eq. (6) holds for $k = \lfloor \log_2(|C^{-1}(1)|/N) \rfloor$ and all e, for *all* $h = M_k(u)$, $u \in U$ (by our choice of the lexicographically *first* triple). However, Lemma 6.1 implies that for this k, and e for which $(N+e) = \lceil |C^{-1}(1)|/2^k \rceil$, Eq. (6) holds for at most $1/2$ of the $h \in H_{n,k}$, which contradicts the fact that U is a discrepancy set.

At this point we have k^* and e^* satisfying Eq. 7 and a hash function $h^* : \{0,1\}^n \to \{0,1\}^{k^*}$ (namely $h^* = M_{k^*}(u^*)$) for which

$$\forall y \; |\{x : h^*(x) = y \wedge C(x) = 1\}| \le (1+\delta)(N+e^*). \qquad (8)$$

At this point we observe that the integer k^*, the hash function h^*, and the integer $B = (N + e^*)$ satisfy the properties stated in the lemma.

Now, let s' be some fixed polynomial in $B, |C|, s$ to be determined later. Using known constructions (e.g., [KvM02]) we can produce from T a (k^*, δ)-discrepancy set $R \subseteq \{0,1\}^{k^*}$ for the set of NP||-circuits of size m. With known constructions, an enumeration of R can be computed efficiently and deterministically from T. The remainder of the proof is devoted to proving that

$$S = \{x : h^*(x) \in R \wedge C(x) = 1\},$$

is a $C^{-1}(1)$-conditional $(n, s, 2\delta)$-discrepancy set as required.

Suppose for the purpose of contradiction that there is a distinguisher $f : \{0,1\}^n \to \{0,1\}$ computable by a size s circuit for which

$$|\Pr_x[f(x) = 1|C(x) = 1] - \Pr_{t \in S}[f(t) = 1|C(t) = 1]| > 2\delta. \tag{9}$$

We use f to describe a distinguisher $g : \{0,1\}^{k^*} \to \{0,1\}$ computable by a size s' NP||-circuit that "catches" the discrepancy set R. On input $y \in \{0,1\}^{k^*}$, g uses $(1+\delta)B$ *non-adaptive* NP queries to determine $\ell_y = |\{x : h^*(x) = y \wedge C(x) = 1 \wedge f(x) = 1\}|$, and g then outputs 1 with probability $\ell_y/((1+\delta)B)$.

We know that $2^{k^*}B \leq |C^{-1}(1)| = 2^{k^*}(1+\delta)B$. Thus

$$\Pr_x[f(x) = 1|C(x) = 1] = \frac{\sum_y \ell_y}{|C^{-1}(1)|} \leq \frac{\sum_y \ell_y}{2^{k^*}B} = \frac{1}{2^{k^*}} \sum_y \frac{\ell_y}{B} = (1+\delta) \Pr_y[g(y) = 1]$$

Similarly, we known that $|S| \leq |R|(1+\delta)B$ and so:

$$\Pr_{r \in R}[g(r) = 1] = \frac{1}{|R|} \sum_{r \in R} \frac{\ell_r}{(1+\delta)B} \leq \frac{\sum_{r \in R} \ell_r}{|S|} = \Pr_{t \in S}[f(t) = 1|C(t) = 1].$$

We may assume that Eq. 9 holds without the absolute value, by inverting f if necessary. Then we get:

$$\Pr_{r \in R}[g(r) = 1] \leq \Pr_{t \in S}[f(t) = 1|C(x) = 1] \quad < \quad \Pr_x[f(x) = 1|C(x) = 1] - 2\delta \leq (1+\delta) \Pr_y[g(y) = 1] - 2\delta$$

and so g distinguishes a random element from R from a truly random element with advantage greater than δ. We may fix g's random coins to preserve this advantage, and notice that g is computable by a size $s' = \text{poly}(B, |C|, s)$ circuit that makes non-adaptive NP oracle queries. This contradicts the fact that R is a discrepancy set against size s' NP||-circuits and so S must indeed be an $C^{-1}(1)$-conditional $(n, s, 2\delta)$-discrepancy set, as desired.

We output k^*, the hash function h^*, the integer $B = (N + e^*)$, and the multiset R, which satisfy the properties stated in the lemma. □

We are now in a position to prove Theorems 3.6 and 3.10. We will need the following fact about composing functions computable with *non-adaptive* NP oracle access:

Lemma 6.3. *Let $f = \{f_n\}$ and $g = \{g_n\}$ be length-preserving function families in $FTIME(t(n))_{||}^{NP}$ and $FTIME(s(n))_{||}^{NP}$ respectively. Then the function family $(f \circ g)$ defined by $(f \circ g)(x) = f(g(x))$ is in $FTIME(\text{poly}(t(n)s(n)n))_{||}^{NP}$.*

Proof. We are given an input x of length n, and we wish to compute $f(g(x))$. Let M_f and M_g be the deterministic oracle Turing Machines associated with f and g.

For $i = 0, 1, 2, \ldots, s(n)$; $j = 0, 1, 2, \ldots t(n)$; $k = 1, 2, \ldots n$; $b \in \{0,1\}$; $z = 0, 1, 2$ we guess:

- a transcript T_g for a computation of M_g on input x in which exactly j queries are answered positively. The transcript includes witnesses $w_\ell^{(g)}$ for each positively answered query $q_\ell^{(g)}$, and an output y_g.
- a transcript T_f for a computation of M_f on input y_g in which exactly i queries are answered positively. The transcript includes witnesses $w_\ell^{(f)}$ for each positively answered query $q_\ell^{(f)}$, and an output y_f.

and check if *all* of the following hold:

- Each $w_\ell^{(g)}$ is a valid witness for query $q_\ell^{(g)}$, and T_g is a valid transcript for the operation of M_g on input x *with the queries $q_\ell^{(g)}$ answered positively and all other queries answered negatively*, with output y_g.
- Each $w_\ell^{(f)}$ is a valid witness for query $q_\ell^{(f)}$, and T_f is a valid transcript for the operation of M_f on input y_g *with the queries $q_\ell^{(f)}$ answered positively and all other queries answered negatively*, with output y_f; OR $z < 1$.
- The k-th bit of y_f is b; OR $z < 2$.

This describes the $3n(s(n)+1)(t(n)+1)$ non-adaptive NP oracle queries. We label each of these queries with a tuple (i, j, k, b, z). Equipped with the answers to these oracle queries, we will (deterministically) compute $f(g(x))$.

Let i^* be the largest value of i for which query $(i, 0, 1, 0, 0)$ is answered positively. We claim that every valid witness for query (i^*, j, k, b, z) for $z \geq 0$ must have $y_g = g(x)$. First, observe that if M_g has exactly i oracle queries answered positively on input x, then query $(i, 0, 1, 0, 0)$ will be answered positively: a witness is obtained by taking T_g to be the *correct* transcript for the operation of M_g on input x together with valid witnesses $w_\ell^{(g)}$ for every positively answered query $q_\ell^{(g)}$. Moreover, no query $(i, 0, 1, 0, 0)$ with i larger than the true number of queries answered positively on input x will be answered positively, because it is impossible to have valid witnesses for that many queries. Therefore the only witnesses for a query of the form (i^*, j, k, b, z) must have a correct transcript T_g for the operation of M_g on input x, and therefore they must end with the correct output $y_g = g(x)$.

Similarly, let j^* be the largest value of j for which query $(i^*, j, 1, 0, 1)$ is answered positively. An identical argument as above shows that every valid witness for query (i^*, j^*, k, b, z) for $z \geq 1$ must have $y_f = f(g(x))$.

Finally, we see that $(i^*, j^*, k, b, 2)$ can be answered positively if and only if b is the value of the k-th bit of y_f. Thus we can determine the string y_f by examining the answers to these $2n$ queries. We output y_f. □

Proof of Theorem 3.6. We are given a circuit A on n bits, and $\epsilon > 0$. Set $\delta = \epsilon/(1-\epsilon)$ and set t as in the statement of Lemma 6.2. We describe our procedure in several steps, and then apply Lemma 6.3 to assemble them into a single procedure that uses non-adaptive NP-oracle access.

- We are assuming that $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits. By Corollary 3.3, $E_{||}^{NP}$ also contains languages that require exponential size NP||-circuits. Let L be such a language in $E_{||}^{NP}$. We first produce the truth table T of L restricted to length t inputs. Since $L \in E_{||}^{NP}$ this procedure is in $FTIME(2^{O(t)})_{||}^{NP}$.
- Apply the function family of Lemma 6.2, with inputs A, δ, n, T. This produces output k, h, B and R, and runs in time $FTIME(2^{O(t)})_{||}^{NP}$.
- The resulting output has integers k and B for which

  $$2^k B \leq |A^{-1}(1)| \leq (1+\delta)2^k B.$$

  We can then output $\rho = (2^k B)/(2^n)$, and the above equation implies:

  $$(1-\epsilon)\Pr_x[A(x)=1] \leq \rho \leq \Pr_x[A(x)=1]$$

  as required.

After applying Lemma 6.3, the overall running time of the procedure is polynomial in $|A|, n$ and $1/\epsilon$ and it uses only non-adaptive NP oracle access. □

Proof of Theorem 3.10. We are given a circuit A on n bits, an integer s, and $\epsilon > 0$, and we want to produce a $A^{-1}(1)$-conditional (n, s, ϵ)-discrepancy set. Set $\delta = \epsilon/2$ and set t as in the statement of Lemma 6.2. We describe our procedure in several steps, and then apply Lemma 6.3 to assemble them into a single procedure that uses non-adaptive NP-oracle access.

- We are assuming that $E_{||}^{NP}$ requires exponential size SV-nondeterministic circuits. By Theorem 3.2, $E_{||}^{NP}$ also contains languages that require exponential size NP||-circuits. Let L be such a language in $E_{||}^{NP}$. We first produce the truth table T of L restricted to length t inputs. Since $L \in E_{||}^{NP}$ this procedure is in $FTIME(2^{O(t)})_{||}^{NP}$.
- Apply the function family of Lemma 6.2, with inputs A, δ, s, T. This produces output k, h, B and R, and runs in time $FTIME(2^{O(t)})_{||}^{NP}$.
- Finally, we produce from R an enumeration of the multiset $S = \{x : h(x) \in R \wedge A(x) = 1\}$. This can be accomplished by making queries of the form "Is S_i a multiset of size i for which $S_i \subseteq \{x : h(x) \in R \wedge A(x) = 1\}$?" for each i up to $2^{O(t)}$ (which is an upper bound on $|R|$). By Theorem 3.12, we can actually produce such multisets S_i using non-adaptive NP oracle queries, and we find all of the S_i in parallel. Finally we output the largest S_i that is found, which must equal S, which is the desired $A^{-1}(1)$-conditional (n, s, ϵ)-discrepancy set.

After applying Lemma 6.3, the overall running time of the procedure is polynomial in $|A|, n, s$ and $1/\epsilon$ and it uses only non-adaptive NP oracle access.

If we assume instead that E^{NP} requires exponential-size SV-nondeterministic circuits then step 1 runs in $FTIME(2^{O(t)})^{NP}$ and the last step can use an NP oracle adaptively to find S in the usual way. In this case the procedure has the same overall running time but uses adaptive NP oracle access. □

7 Conclusions and open problems

Our "downward collapse theorem" states that for every sufficiently strong class $\mathcal{C}$ if $\mathcal{C}$ has small NP||-circuits then C has small SV-nondeterministic circuits. A very natural open problem is try to extend the "downward collapse" theorem to handle *adaptive* NP queries. That is, show that if E is computable by small NP-circuits then E is computable by small NP||-circuits. We remark that all the proofs in this paper relativize. Thus, we find it interesting to check whether the statement above relativizes.

The proof of the "downward collapse" theorem gives a *nonuniform* reduction from an NP||-circuit that computes $\hat{f}$ to an SV-nondeterministic circuit that computes f. We observe that achieving uniform versions of this theorem would give new *unconditional* results. For example, replacing the nonuniform reduction with a uniform probabilistic reduction (which succeeds with probability say $2/3$) would show that $EXP \subseteq P_{||}^{NP}$/poly $\Rightarrow EXP = AM$. This in turn gives the following unconditional results: $AMEXP \not\subseteq P_{||}^{NP}$/poly (by an argument similar to that of [BFT98]) and $AM \subseteq \cap_{\delta>0}[io]NTIME(2^{n^{\delta}})^{NP}$ (see the discussion section in [GSTS03]). A somewhat easier task may be to try and show that $EXP \subseteq P_{||}^{NP} \Rightarrow EXP = AM$. Obtaining this would give the unconditional result that $AMEXP \not\subseteq P_{||}^{NP}$. We note that our results already give $EXP \subseteq P_{||}^{NP} \Rightarrow EXP \subseteq AM/\log$. To see this, observe that with some minor modifications to the parameters used in the proof of Theorem 3.2, it is sufficient to supply $\sum_i p_i$ as nonuniform advice, rather than $p_1, p_2, \ldots, p_k$.

We have shown how to construct relative error approximators and conditional discrepancy sets using hardness for nondeterministic circuits. However, we do not know whether the existence of any of these objects entails such a hard function. In particular, the "standard arguments" which show that "pseudorandomness entails hardness" only give hardness for *deterministic* circuits. Is it possible to construct the objects above using a weaker hardness assumption? Is it possible to derandomize AM using these objects?

Acknowledgements

We thank Russell Impagliazzo for pointing out the class BPP_{path}. We thank Salil Vadhan for helpful comments.

References

[AK01] V. Arvind and J. Kobler. On pseudorandomness and resource-bounded measure. *TCS: Theoretical Computer Science*, 255, 2001.

[AKRR03] E. Allender, M. Koucky, D. Ronneburger, and S. Roy. Derandomization and distinguishing complexity. In *Proc. 18th Annual IEEE Conference on Computational Complexity*, pages 209–220, 2003.

[Bab85] L. Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, 1985.

[BCG$^+$96] N. H. Bshouty, R. Cleve, R. Gavalda, S. Kannan, and C. Tamon. Oracles and queries that are sufficient for exact learning. *Journal of Computer and System Sciences*, 52(3):421–433, 1996.

[BF90] D. Beaver and J. Feigenbaum. Hiding instances in multioracle queries. In *7th Annual Symposium on Theoretical Aspects of Computer Science*, volume 415 of *lncs*, pages 37–48, Rouen, France, 22–24 February 1990. Springer.

[BFNW93] 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.

[BFT98] H. Buhrman, L. Fortnow, and T. Thierauf. Nonrelativizing separations. In *Proc. 13th Annual IEEE Conference on Computational Complexity*, pages 8–12, 1998.

[BGP00] M. Bellare, O. Goldreich, and E. Petrank. Uniform generation of NP-witnesses using an NP-oracle. *INFCTRL: Information and Computation (formerly Information and Control)*, 163, 2000.

[BM84] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13(4):850–864, November 1984.

[BM88] L. Babai and S. Moran. Arthur-merlin games: A randomized proof system and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36:254–276, 1988.

[BR94] M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. In Shafi Goldwasser, editor, *Proceedings: 35th Annual Symposium on Foundations of Computer Science, November 20–22, 1994, Santa Fe, New Mexico*, pages 276–287, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1994. IEEE Computer Society Press.

[Cai01] J.-Y. Cai. $S_2^P \subseteq ZPP^{NP}$. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS-01)*, pages 620–629, 2001.

[Can96] R. Canetti. On BPP and the polynomial-time hierarchy. *Information Processing Letters*, 57:237–241, 1996.

[GMR89] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal of Computing*, 18(1):186–208, 1989.

[Gol97] O. Goldreich. A sample of samplers – A computational perspective on sampling (survey). In *ECCCTR: Electronic Colloquium on Computational Complexity, technical reports*, 1997.

[Gol98] O. Goldreich. *Modern Cryptography, Probabilistic Proofs and Pseudorandomness*. Springer-Verlag, Algorithms and Combinatorics, 1998.

[GSTS03] D. Gutfreund, R. Shaltiel, and A. Ta-Shma. Uniform hardness versus randomness tradeoffs for Arthur-Merlin games. In *Proc. 18th Annual IEEE Conference on Computational Complexity*, 2003.

[HHT97] Y. Han, L.A. Hemaspaandra, and T. Thierauf. Threshold computation and cryptographic security. *SIAM J. Comput.*, 26(1):59–78, 1997.

[Imp95] 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 October 1995. IEEE.

[ISW99] R. Impagliazzo, R. Shaltiel, and A. Wigderson. Near-optimal conversion of hardness into pseudo-randomness. In IEEE, editor, *40th Annual Symposium on Foundations of Computer Science: October 17–19, 1999, New York City, New York,*, pages 181–190, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1999. IEEE Computer Society Press.

[ISW00] R. Impagliazzo, R. Shaltiel, and A. Wigderson. Extractors and pseudo-random generators with optimal seed length. In ACM, editor, *Proceedings of the thirty second annual ACM Symposium on Theory of Computing: Portland, Oregon, May 21–23, [2000]*, pages 1–10, New York, NY, USA, 2000. ACM Press.

[IW97] R. Impagliazzo and A. Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 220–229, El Paso, Texas, 4–6 May 1997.

[JVV86] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43(2-3):169–188, ???? 1986.

[Kab02] V. Kabanets. Derandomization: A brief overview. In *ECCCTR: Electronic Colloquium on Computational Complexity, technical reports*, 2002.

[KvM02] A. R. Klivans and D. van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM Journal on Computing*, 31(5):1501–1526, October 2002.

[MV99] P. B. Miltersen and N. V. Vinodchandran. Derandomizing Arthur-Merlin games using hitting sets. In *40th Annual Symposium on Foundations of Computer Science (FOCS '99)*, pages 71–80, 1999.

[NW94] N. Nisan and A. Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49(2):149–167, October 1994.

[RS98] A. Russell and R. Sundaram. Symmetric alternation captures BPP. *Computational Complexity*, 7(2):152–162, 1998.

[Sha81] A. Shamir. The generation of cryptographically strong pseudo-random sequences. In Allen Gersho, editor, *Advances in Cryptology: A Report on CRYPTO 81*, page 1. Department of Electrical and Computer Engineering, U. C. Santa Barbara, 24–26 August 1981.

[Sto83] L. Stockmeyer. The complexity of approximate counting. In ACM, editor, *Proceedings of the fifteenth annual ACM Symposium on Theory of Computing, Boston, Massachusetts, April 25–27, 1983*, pages 118–126, New York, NY, USA, 1983. ACM Press.

[STV01] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom generators without the XOR lemma. *Journal of Computer and System Sciences*, 62:236–266, 2001.

[SU01] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudo-random generator. In IEEE, editor, *42nd IEEE Symposium on Foundations of Computer Science: proceedings: October 14–17, 2001, Las Vegas, Nevada, USA*, pages 648–657, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 2001. IEEE Computer Society Press.

[Uma02] C. Umans. Pseudo-random generators for all hardnesses. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC-02)*, pages 627–634, New York, May 19–21 2002. ACM Press.

[WB86] L.R. Welch and E.R. Berlekamp. Error correction for algebraic block codes. U.S. Patent No. 4,633,470, issued December 30, 1986.

[Yao82] A. C. Yao. Theory and applications of trapdoor functions (extended abstract). In *23th Annual Symposium on Foundations of Computer Science (FOCS '82)*, pages 80–91, Los Alamitos, Ca., USA, November 1982. IEEE Computer Society Press.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'