



On Constructing Parallel Pseudorandom Generators from One-Way Functions

Emanuele Viola*

Division of Engineering and Applied Sciences

Harvard University

Cambridge, MA 02138

viola@eecs.harvard.edu

April 2, 2005

Abstract

We study pseudorandom generator (PRG) constructions $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ from one-way functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. We consider PRG constructions of the form $G^f(x) = C(f(q_1) \dots f(q_{\text{poly}(n)}))$ where C is a polynomial-size constant depth circuit (i.e. AC^0) and C and the q 's are generated from x arbitrarily. We show that every black-box PRG construction of this form must have stretch s bounded as $s \leq l \cdot (\log^{O(1)} n)/m + O(1) = o(l)$. This holds even if the PRG construction starts from a one-to-one function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ where $m \geq 5n$. This shows that either adaptive queries or sequential computation are necessary for black-box PRG constructions with constant factor stretch (i.e. $s = \Omega(l)$) from one-way functions, even if the functions are one-to-one.

On the positive side we show that if there is a one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ that is regular (i.e. the number of preimages of $f(x)$ depends on $|x|$ but not on x) and computable by polynomial-size constant depth circuits then there is a PRG $: \{0, 1\}^l \rightarrow \{0, 1\}^{l+1}$ computable by polynomial-size constant depth circuits. This complements our negative result above because one-to-one functions are regular.

We also study constructions of average-case hard functions starting from worst-case hard ones, i.e. hardness amplifications. We show that if there is an oracle procedure Amp^f in the polynomial time hierarchy (PH) such that Amp^f is average-case hard for every worst-case hard f , then there is an average-case hard function in PH *unconditionally*. Bogdanov and Trevisan (FOCS '03) and Viola (CCC'03) show related but incomparable negative results.

Keywords: Pseudorandom generator construction, one-way function, black-box, constant-depth circuit, hardness amplification, restriction, noise sensitivity.

*Research supported by NSF grant CCR-0133096, US-Israel BSF grant 2002246, ONR grant N-00014-04-1-0478.

1 Introduction

A rigorous notion of *pseudorandom generators* (PRGs) was introduced in the seminal works of Blum and Micali [BM] and Yao [Yao] and have since found a striking variety of applications in Cryptography and Complexity Theory. A PRG $G : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is an efficient procedure that stretches l input bits into $l + s$ output bits such that the output of the PRG is indistinguishable from random to efficient adversaries. That is, for every *probabilistic polynomial time machine* (PPT) A we have

$$\left| \Pr[A(G(U_l)) = 1] - \Pr[A(U_{l+s}) = 1] \right| \leq \epsilon$$

where U_n denotes a uniform random variable in $\{0, 1\}^n$ and ϵ is negligible in $l + s$.

While the existence of PRGs is a major open problem, there has been a series of fascinating works constructing PRGs from weaker and weaker assumptions. Most of these works construct PRGs starting from *one-way functions* [BM, Yao, Lev, GL, GKL, HILL]. Informally, a function is one-way if it is easy to compute but hard to invert on average. (The existence of one-way functions implies that $P \neq NP$, but the converse is not known to hold.) For a discussion of pseudorandom generators we refer to the reader to the excellent book by Goldreich [Gol].

1.1 PRG constructions with polynomial stretch

A crucial parameter of every PRG $G : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is its *stretch* s , that one wants as big as possible. Note that s is only relevant in relation with the input length l , since from a PRG $G' : \{0, 1\}^l \rightarrow \{0, 1\}^{l+1}$ one can trivially construct, for every polynomial p , a PRG $G'' : \{0, 1\}^{pl} \rightarrow \{0, 1\}^{pl+p}$ with stretch p . G'' is the concatenation of the output of p copies of G' on p independent seeds. However, in this way we will never get stretch equal to the seed length (pl).

But in many applications one needs the stretch to be linear in the input length. Two important such applications are Naor's bit-commitment [Nao] and private-key encryption, where starting from a small key of length l , one wants to generate many bits $l + s \gg l$ that can be used to encrypt messages in a "stream cipher".

So suppose we want to build a PRG $G : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ where $s = l$. To achieve this, using any of the known constructions [BM, Yao, Lev, GL, GKL, HILL], one works in two steps. First, starting from a one-way function f one builds a PRG G_1^f with small stretch, say one bit:

$$G_1^f(x) = H^f(x) \circ b^f(x)$$

where $G_1^f(x) : \{0, 1\}^l \rightarrow \{0, 1\}^{l+1}$, $H^f : \{0, 1\}^l \rightarrow \{0, 1\}^l$ and $b^f : \{0, 1\}^l \rightarrow \{0, 1\}$. Then, to get arbitrary polynomial stretch, one uses the following construction due to Goldreich and Micali (see [Gol], Section 3.3.2): from G_1^f construct $G_2^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ defined as

$$G_2^f(x) := b^f(x) \circ b^f(H^f(x)) \circ b^f(H^f(H^f(x))) \circ \dots \circ b^f(\underbrace{H^f(\dots(H^f(x)\dots))}_{l+s-1}). \quad (1)$$

Construction (1) is very *sequential* in the following sense: the i -th evaluation of H depends on the output of the $(i - 1)$ -th evaluation of H , and hence the straightforward circuit for G_2^f has depth at least s .

Notice that, once we have a PRG with linear stretch, say $G : \{0, 1\}^l \rightarrow \{0, 1\}^{2l}$, to improve the stretch one can use the ‘very efficient’ construction given by Goldreich, Goldwasser and Micali in [GGM]. However, there remains the problem of constructing a PRG with linear stretch.

1.2 The main problem we study

The main question addressed in this paper is the following: *Are PRG constructions with arbitrary stretch inherently sequential?* This problem is motivated by the question, both practical and philosophical, of how much cryptography can be done in low complexity classes.

Of course, we must be more precise about what we mean by ‘PRG construction’ and ‘sequential’.

We now discuss PRG constructions. We consider *black-box* PRG constructions, as in many other works starting with the seminal paper by Impagliazzo and Rudich [IR]. Roughly speaking, $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is a *black-box PRG construction from one-way function f* if there is a fixed PPT M such that, for *every* (computationally unbounded) oracle function f and adversary A , if A *breaks* the PRG then M *inverts* f . I.e., if A distinguishes the output of PRG from truly random, then M , when given oracle access to both f and A , can find a preimage of $f(U_n)$ with noticeable probability. The idea is that if $f(U_n)$ cannot be inverted with noticeable probability by a PPT (i.e., if f is one-way) then no PPT can break G^f , and so G^f is a PRG. Most results in Cryptography, and in particular most PRG constructions (for example [BM, Yao, Lev, GL, GKL, HILL]) are proved via black-box constructions.

We now define *parallel* PRG constructions. The notion of a parallel PRG construction we look at in this paper is the following, where $l = l(n), m = m(n)$ and $s = s(n)$:

PRG construction $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$

$$G^f(x) = C_x (f(q_{x,1}) \dots f(q_{x,\text{poly}(n)}))$$

where $C_x : \{0, 1\}^{\text{poly}(n)} \rightarrow \{0, 1\}^{l+s}$ is a constant depth circuit of size $\text{poly}(n)$,
and $C_x, q_{x,1}, \dots, q_{x,\text{poly}(n)}$ are generated from x arbitrarily.

Table 1: Parallel PRG construction.

By a *constant depth circuit* we mean AC^0 , i.e. a constant depth circuit with OR,AND gates, where AND, OR gates have unbounded fan-in (see e.g., [Hås]).

PRG constructions in the form in Table 1 are intuitively parallel in the following sense: (1) The queries made to f are *non-adaptive* (i.e. they do not depend on f but only on x), and (2) C_x is a *constant depth* circuit. It seems interesting to study black-box PRG constructions relaxing either (1) or (2), but we can prove our main negative results only when both apply.

Finally, note we do not make any assumption on how $C_x, q_{x,1}, \dots, q_{x,\text{poly}(n)}$ are generated from x : This makes our negative result stronger, while in our positive results all the computation is done by a constant depth circuit.

1.3 Our Results on PRG Constructions

Next we discuss our results for PRG constructions. We have both positive and negative results. We start with the latter.

Theorem 1.1 (This Paper). *Let $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ be a black-box PRG construction (Def. 2.1) in the form in Table 1. Then the following hold:*

1. *If G^f starts from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with $m = \log^{\omega(1)} n$ and $m = n^{O(1)}$, then $s \leq l \cdot \frac{\log^{O(1)} n}{m} + O(1)$.*
2. *If G^f starts from one-to-one one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{5n}$ then $s \leq l \cdot \frac{\log^{O(1)} n}{5n} + O(1)$.*

Theorem 1.1-(1) shows that black-box PRG constructions from one-way functions can only be parallel (i.e., in the form in Table 1) if $s = o(l)$. And Theorem 1.1-(2) shows this holds even if the PRG construction starts with one-to-one functions, as long as the range of the one-to-one one-way function is sufficiently bigger than its domain.

In this work we also sketch a proof of the fact that (essentially) Theorem 1.1-(1) holds even for a less restrictive kind of black-box constructions, i.e. *mildly* black-box constructions.

It seems natural at this point to ask: is there *any* PRG construction in the form in Table 1, even with stretch $s = 1$? The problems we must solve (to answer this question) depend on what kind of one-way function the PRG construction starts from. We now explain these problems and our contributions.

From (generic) one-way function: The only PRG construction that works in this case is [HILL]. This construction, even to produce a PRG with stretch $s = 1$, uses as a component the Goldreich-Micali Construction (1) discussed earlier in Section 1.1. As already noted, Construction (1) is not parallel. In this case we do not know if there is any parallel PRG construction.¹

From one-to-one one-way function: The main problem here is that existing PRG constructions (e.g. [HILL]) apply *randomness extractors* [NZ] to the evaluations of the one-way function. While it is known that constant-depth circuits cannot compute extractors with good parameters [MNT, Vio], in this work we show that constant-depth circuits can compute extractors for the setting of parameters that arises in some of these PRG constructions. Using this fact we obtain a parallel PRG construction (with some stretch $s \geq 1$). We actually show parallel PRG constructions starting from the more general class of *regular* one-way functions, where a function f is regular if the number of preimages of $f(x)$ depends on $|x|$ but not on x .

From one-way permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$: In this case we can use the PRG construction $GL^\pi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n+1}$ by Goldreich and Levin and defined as $GL^\pi(y, r) := (\pi(y), r, \langle y, r \rangle)$, where $\langle y, r \rangle$ is a general hard-core predicate (see [GL, Gol]). This is of the form in Table 1, because for every input $x = (y, r)$ the circuit C_x that has $\langle y, r \rangle$ and r hardwired and is defined as $C_x(\pi(y)) := (\pi(y), r, \langle y, r \rangle)$ is trivially constant depth. But what happens if we require that all the computation be done in constant depth? We cannot use

¹However, Construction (1) can be dispensed with if one allows for $O(\log n)$ bits of nonuniformity in the PRG construction. One can then obtain a parallel PRG construction with our techniques. Details omitted.

directly the above construction since it requires computing a general hard-core predicate (i.e., $\langle y, r \rangle$), which cannot be done in constant depth [GNR]. We bypass this problem by showing that it is sufficient to compute the output distribution of a general hard-core predicate over a random input, and this in fact can be done by a constant-depth circuit.

We now state our positive results for PRG constructions.

Theorem 1.2 (This Paper). *There is a black-box PRG construction $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+1}$ from regular one-way functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that G is of the form in Table 1. Moreover, all the computation is done by a uniform constant depth circuit of size $\text{poly}(n)$. The input length of G is $l = 2n$ starting from one-way permutations $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $l = \text{poly}(n)$ starting from regular one-way functions. In particular:*

1. *If there is a one-way permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ computable by uniform constant depth circuits of size $\text{poly}(n)$ then there is a PRG $G : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n+1}$ computable by uniform constant depth circuits of size $\text{poly}(n)$.*
2. *If there is a regular one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ computable by uniform constant depth circuits of size $\text{poly}(n)$ then there is a PRG $G : \{0, 1\}^l \rightarrow \{0, 1\}^{l+1}$ computable by uniform constant depth circuits of size $\text{poly}(n)$.*

Note Theorem 1.2-(2) matches Theorem 1.1-(2) because one-to-one functions are regular.

1.4 Related Work

A concurrent and beautiful work related to ours is the one by Applebaum, Ishai and Kushilevitz [AIK]. They show that the existence of a ‘moderately easy’ PRG, say in NC^1 (i.e. computable by circuits of polynomial-size and logarithmic-depth), implies the existence of a PRG in NC^0 (i.e. computable by circuits of polynomial-size and constant-depth with *bounded fan-in*). Note that in this work we consider the strictly larger class AC^0 (the class of functions computable by circuits of polynomial-size and constant depth with *unbounded fan-in*.) However, the NC^0 PRG of [AIK] has *sublinear stretch even if the original NC^1 PRG has polynomial stretch*. This is interesting in relation with our negative results that only rule out parallel black-box PRG constructions with linear stretch. (They also prove analogous connections for other cryptographic primitives, such as one-way functions.) Moreover, they improve on our Theorem 1.2 obtaining the same results for constant-depth circuits with bounded fan-in (whereas our Theorem 1.2 refers to constant-depth circuits with unbounded fan-in, specifically our result uses fan-in $\log^{1+\epsilon} n$). (This result of theirs uses techniques similar to ours.)

Another beautiful work related to ours is the one by Gennaro and Trevisan [GT]. They show that there is no black-box PRG construction $G^\pi : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ that makes less than $s/\omega(\log n)$ queries to π . This holds even if G^f is a PRG construction from one-way permutations. The work of Gennaro and Trevisan is thus a tradeoff between the stretch of the PRG construction and the number of queries it makes to f . The difference with our work is the following: We are not concerned with *how many* queries G^f makes to f , rather we are concerned with *how these queries are made and processed*. Note that our bounds

in Theorem 1.1 essentially do *not* depend on the number of queries made to f (except for the hidden constant in the negligible factor $\log^{O(1)} n$). Rather, they depend on the parallel structure of G .

There exist several other works addressing the complexity of PRGs. Kharitonov, Goldberg and Yung [KGY] and Yu and Yung [YY] prove strong negative results about the ability of various automata and other space-restricted devices to compute PRGs. Linial, Mansour and Nisan [LMN] prove that constant depth circuits cannot compute pseudorandom functions (an object related to PRGs). Impagliazzo and Naor [IN] show how to construct PRGs based on the assumed intractability of the subset sum problem. In particular, they show how to construct a PRG $\{0, 1\}^n \rightarrow \{0, 1\}^{n+\log n}$ computable by constant depth circuits. Reif and Tygar [RT] and Naor and Reingold [NR] construct PRGs under specific number-theoretic complexity assumptions.

In [Vio] the author studies the complexity of a different kind of PRG constructions, namely those based on the Nisan-Wigderson paradigm [NW]. The results in [Vio] do not seem to be comparable to the ones in this paper, as they apply to this different kind of PRG constructions. However, both works use results on the noise sensitivity of constant-depth circuits in the proof of their negative results.

1.5 Worst-case Hardness Amplification

Another problem we study in this paper is the problem of *worst-case hardness amplification*, which is the problem of producing an average-case hard function starting from a worst-case hard function. A motivation for studying this problem is to establish connections between average-case complexity and worst-case complexity, as it has been accomplished for high complexity classes such *PSPACE* and *EXP* (e.g. [Lip, BF, BFL, FL, CPS, STV, TV, Vio]). Most constructions in these works are black-box both in the use of the worst-case hard function f and in the ‘proof of correctness’. Namely they exhibit efficient algorithms Amp and R such that for *every* function f and *every* adversary A , if A computes Amp^f well on average then R^A computes f everywhere. Note that if f is worst-case hard then R^A cannot be a small circuit. Since R is efficient this means that A cannot be a small circuit, and hence Amp^f is average-case hard.

There are results showing that such connections (i.e., between worst-case and average-case hardness) for classes within the polynomial-time hierarchy (PH) are unlikely to be provable using these kind of black-box techniques: Bogdanov and Trevisan [BT], building on [FF], show that every hardness amplification within NP such that its proof of correctness is black-box and R is non-adaptive implies that the PH collapses, and therefore such a hardness amplification is unlikely to exist. In a previous work [Vio] we showed (unconditionally) that there is no hardness amplification within PH where both the use of f and the proof of correctness are black-box.

In this paper we obtain the first negative result on hardness amplifications within PH that are black-box only in the use of f . Specifically, we show that exhibiting such hardness amplification procedures is equivalent to exhibiting an average-case hard function in PH , in which case no hardness amplification is needed. We give one necessary definition and then our result.

Definition 1.3. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is worst-case-hard (resp., ϵ -hard) for size S if every circuit of size S fails to compute f on some input (resp., on at least ϵ fraction of inputs).

Theorem 1.4 (This Paper). Let $S(n)$ be such that for every c and sufficiently large n , $S(n) \leq 2^n/n^c$. Suppose there is a constant a and an oracle machine Amp in PH such that for every $f : \{0, 1\}^n \rightarrow \{0, 1\}$ that is worst-case hard for size $S = S(n)$, $\text{Amp}^f : \{0, 1\}^{n^a} \rightarrow \{0, 1\}$ is .3-hard for size $S' = S'(n)$. Then there is a constant b and a function f' in PH such that $f' : \{0, 1\}^{n^b} \rightarrow \{0, 1\}$ is .1-hard for size $S'(n)$.

1.6 Techniques

We now sketch the main ideas in the proof of Theorem 1.1-(1). Similarly to other works [IR, GT], the idea is to choose the oracle function f at random from a certain distribution \tilde{F} , then show that (1) \tilde{F} is one-way w.h.p. but also (2) there is an adversary that breaks $G^{\tilde{F}}$ w.h.p., thus contradicting the fact that G is a black-box PRG construction. The main new ingredient in this work is that some of the bits in the truth-table of \tilde{F} are *fixed*, and we will give them for free to the adversary. We will then show that for this fixing of bits $G^{\tilde{F}}$ is easy to break. One of the challenges is of course showing that \tilde{F} is still one-way after these bits have been fixed. More specifically, the bits to be fixed are chosen by applying a *random restriction* [FSS] to the truth table of the oracle f . Since $G^f(x)$ is a constant depth function of evaluations of f , and because after applying a random restriction to a constant depth circuit the circuit ‘tremendously simplifies’ (see e.g. [FSS, Hås]), it is possible to exhibit an adversary that breaks the output of G . More specifically, we will use the fact that constant depth circuits have *low noise sensitivity* [LMN, Bop, Vio], which means that after fixing most of its input bits, a constant depth circuit becomes very biased, i.e. its output does not change much when the few unfixed input bits are filled at random.

For Theorem 1.1-(2) we define a particular distribution on restrictions that also ensures that the oracle f is one-to-one.

We now sketch the main ideas of our negative result about Hardness Amplification, Theorem 1.4. As before, we will choose the oracle function f at random from a certain distribution \tilde{F} where some of the bits are fixed in such a way that, (1) \tilde{F} is still worst-case hard w.h.p., but (2) $\text{Amp}^{\tilde{F}}$ is trivialized. Again, the bits to be fixed are chosen applying a random restriction to the truth table of the oracle f . The idea now is that since $\text{Amp}^{\tilde{F}}$ is trivialized, we can dispense with the oracle and construct an average-case hard function h from scratch, thus proving the theorem. But the problem is that we don’t know what is the fixing of the bits that satisfies (1) and (2). An idea would be to include the fixing of the bits in the input to the function h , but the problem is that the size of this fixing of bits is of the order of the truth table of the oracle f , i.e. 2^n , while we need the input length of h to be polynomial in n (since the circuit size S' in Theorem 1.4 is relative to the input length of f). To overcome this problem, we *derandomize* the random restriction. I.e., we create a pseudorandom distribution on restrictions that can be generated using only $\text{poly}(n)$ random bits, yet still w.h.p. satisfies (1) and (2). Now, the function h takes σ as part of the input, where σ is of size $\text{poly}(n)$ and is used to generate a pseudorandom restriction. Now the input length of h is polynomial in n and the theorem can be proved. This pseudorandom

distribution on restrictions is obtained using Nisan’s unconditional PRG against constant depth circuits [Nis]. The challenges of course are showing that after this derandomization (1) and (2) still hold. In particular, for (2) we show that a constant depth circuit becomes very biased even after applying a pseudorandom restriction. The idea of using Nisan’s generator to derandomize restrictions already appeared in [CSS].

1.7 Organization

This paper is organized as follows. In Section 2 we discuss notation. In Section 3 we prove our negative result for black-box PRG constructions in the form in Table 1 from one-way function, i.e. we prove Theorem 1.1-(1). The analogous result for one-to-one one-way functions is proved in Section 6. In Appendix B we sketch a proof of the fact that (essentially) Theorem 1.1-(1) also holds for “mildly black-box” PRG constructions. In Section 4 we prove our positive result about PRG constructions in constant-depth circuits from one-way permutations, i.e. we prove Theorem 1.2-(1). We omit the details of the proof of the analogous result for regular one-way function, i.e. Theorem 1.2-(2) (but the key idea is discussed at the end of Section 4). Our negative result for hardness amplification, i.e. Theorem 1.4, is proved in Section 5. We discuss some open problems in Section 7. Appendix A contains some proof details.

2 Preliminaries

We denote by U_n the uniform random variable over $\{0, 1\}^n$ and by $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ a uniform random function. $\Delta(x, y)$ is the relative Hamming distance between vectors $x, y \in \{0, 1\}^n$, i.e. $\Pr_i[x_i \neq y_i]$. Throughout the paper $\epsilon(n)$ denotes a quantity negligible in n , i.e. $1/n^{\omega(1)}$. We write ‘w.h.p.’ for ‘with high probability’, i.e. with probability $1 - o(1)$.

Restrictions: A restriction ρ on t bits is an element of $\{0, 1, *\}^t$, where we think of the *’s as values yet to be chosen. For $x \in \{0, 1\}^t$ we denote by $x^\rho \in \{0, 1\}^t$ the string obtained from ρ by substituting the *’s with the corresponding bits of x . Note x^ρ only depends on the bits of x corresponding to * in ρ . We often consider restrictions on $b \cdot m$ bits, where b can be as large as 2^n , and it will be convenient to view such restrictions as functions $\rho : [b] \rightarrow \{0, 1, *\}^m$, where $[b] := \{1, \dots, b\}$. When $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}^m$ we think of $\rho(i)$ as a partial assignment to the output $f(i)$ of some function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. The following is a key definition: for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ we denote by $f^\rho : \{0, 1\}^n \rightarrow \{0, 1\}^m$ the function defined by

$$f^\rho(x) := f(x)^{\rho(x)}.$$

For a fixed restriction ρ , a key random variable we will look at is F^ρ (recall $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a uniform random function). Note that F^ρ can be seen as the distribution on functions whose truth table is obtained starting from the truth table of ρ and replacing each * with a uniform and independent random bit. The standard [FSS] distribution on restrictions R_δ is the one where each symbol in the restriction is independently * with probability δ and otherwise it is a uniform and independent random bit. When we say that $\rho : [b] \rightarrow \{0, 1, *\}^m$ is random in R_δ we mean that each of the $b \cdot m$ symbols in the truth table

of ρ is independently $*$ with probability δ and otherwise it is a uniform and independent random bit.

We would like to point out some differences between our notation for restrictions (above) and the notation more commonly used in literature (e.g. [FSS, Hås]). The notation commonly used in literature is the following: for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and a restriction $\rho : \{1, 2, \dots, n\} \rightarrow \{0, 1, *\}$ with s stars (i.e. exactly s distinct indexes i such that $\rho(i) = *$) one denotes by $f|_\rho : \{0, 1\}^s \rightarrow \{0, 1\}^m$ the function obtained by ‘restricting’ f on the s bits mapped to $*$ by ρ , where the other bits are fixed as prescribed by ρ . Note that we have $f^\rho(x) = f|_{\rho(x)}(y)$, where $y \in \{0, 1\}^s$ is the projections of x on the s bits mapped to $*$ by ρ . Our notation is more convenient in our setting where the restriction applied to f depends on the input x . To avoid confusion we will never use the notation $f|_\rho$ in the rest of the paper.

Black-box PRG constructions: Now we formally define black-box PRG constructions.

Definition 2.1 (Black-box PRG Construction). *An oracle machine $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is a black-box PRG construction from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ if there exists an oracle PPT M such that for sufficiently large n , for every $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and every $A : \{0, 1\}^{l+s} \rightarrow \{0, 1\}$, if A breaks G^f , i.e.*

$$|\Pr[A(G^f(U_l)) = 1] - \Pr[A(U_{l+s}) = 1]| \geq 1/4$$

then $M^{f,A}$ inverts f , i.e.

$$\Pr[f(M^{f,A}(f(U_n))) = f(U_n)] \geq 1/n.$$

We say G is from one-to-one one-way function (resp., from regular one-way function) if the above is only required to hold when f is one-to-one (resp., regular).

We think of l, s, m as functions of n . Recall a function $f : \{0, 1\}^* \rightarrow \{0, 1\}$ is regular if the number of preimages of $f(x)$ depends on $|x|$ but not on x . The values $1/4$ and $1/n$ in Definition 2.1 can be substituted by $1/p(n)$ for any polynomial $p(n)$. We fix them for concreteness. In Definition 2.1, and throughout the paper, probabilities are (implicitly) taken also over the internal coin tosses of the PPTs. For more on black-box constructions we refer the reader to the survey in the paper by Reingold, Trevisan and Vadhan [RTV] (in their taxonomy, Definition 2.1 defines a ‘fully black-box’ PRG construction).

3 Proof of Theorem 1.1-(1)

In this section we prove our negative result about black-box PRG constructions that start from (generic) one-way function, i.e Theorem 1.1-(1). We now proceed to sketch the main ideas in the proof. Suppose $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is a PRG construction from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and let M be the inverting machine required by Definition 2.1. The high level idea is to come up with, for sufficiently large n , a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and a computationally unbounded adversary $A : \{0, 1\}^{l+s} \rightarrow \{0, 1\}$ such that A breaks G^f but $M^{f,A}$ does *not* invert f , thus contradicting the fact that G^f is a PRG construction by Definition 2.1. The construction of f and A will be probabilistic, i.e. we will show a

distribution on functions and adversaries that w.h.p. satisfies the above. This certainly ensures that there exist some f, A satisfying the above. Our final distribution on functions will be F^ρ for a suitable restriction $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}^m$. Recall from Section 2 that F^ρ can be seen as the random function obtained from the truth table of ρ replacing the $*$'s with (uniform and independent) random bits.

The main idea in the proof is to find a *fixed* restriction $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}^m$ that satisfies the following two properties:

- I. For every oracle A , with high probability over F , $M^{F^\rho, A}$ does not invert F^ρ , i.e.:

$$\Pr_{F, U_n} [F^\rho(M^{F^\rho, A}(F^\rho(U_n))) = F^\rho(U_n)] \leq \epsilon(n).$$

- II. There is a fixed function $g : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that

$$E_{F, U_l} [\Delta(G^{F^\rho}(U_l), G^g(U_l))] \leq \frac{\text{poly log}(n)}{m}.$$

(Where E denotes expectation and recall that Δ denotes relative Hamming distance.)

Intuitively, (I) says that F^ρ is hard to invert just because of the randomness left in F even after fixing some of the bits in its truth table according to ρ . (II) says that ρ trivializes $G^{F^\rho}(U_l)$ because, on average, the output of $G^{F^\rho}(U_l)$ is close in Hamming distance to a vector that does not depend on the oracle. Before discussing how to construct the restriction ρ , let us show how to prove the theorem once we have such a ρ .

Proof of Theorem 1.1-(1), assuming ρ satisfies (I) and (II). Let g be the function given by Property (II), and let d be a sufficiently large constant. Consider the following (computationally unbounded) adversary

$$A_g(z) := 1 \text{ iff } \exists x \in \{0, 1\}^l : \Delta(G^g(x), z) \leq \frac{\log^d n}{m}.$$

Claim 3.1. *There is a constant d' such that if $s \geq l \cdot \frac{\log^{d'} n}{m} + d'$ then w.h.p. over F , we have that A_g breaks $G^{F^\rho}(U_l)$, i.e.*

$$\left| \Pr_{U_l} [A_g(G^{F^\rho}(U_l)) = 1] - \Pr_{U_{l+s}} [A_g(U_{l+s}) = 1] \right| \geq 1/4.$$

The proof of the theorem follows from Claim 3.1 as follows: By the claim, w.h.p. over F , A_g breaks $G^{F^\rho}(U_l)$, so by definition of black-box PRG construction, M^{F^ρ, A_g} inverts F^ρ w.p. at least $1/n$. But this contradicts Property (I).

Proof of Claim 3.1. By Property (II) and Markov's inequality, w.h.p. over F , we have

$$\Pr_{U_l} [A_g(G^{F^\rho}(U_l)) = 1] \geq 1/2.$$

On the other hand, we also have:

$$\begin{aligned} \Pr[A_g(U_{l+s}) = 1] &\leq \frac{|\{z : A_g(z) = 1\}|}{2^{l+s}} \\ &\leq \frac{\sum_{x \in \{0,1\}^l} |\{z : \Delta(G^g(x), z) \leq (\log^d n)/m\}|}{2^{l+s}} \leq \frac{2^l \cdot 2^{H(\frac{\log^d n}{m}) \cdot (l+s)}}{2^{l+s}}, \end{aligned}$$

where for every x we bound

$$|\{z : \Delta(G^g(x), z) \leq (\log^d n)/m\}| \leq 2^{H(\frac{\log^d n}{m})(l+s)},$$

where $H(p) = p \log(1/p) + (1-p) \log(1/(1-p))$ is the binary entropy function (see any book on Coding Theory). Since $H(\log^d n/m) \leq \log^d n/m$ for some constant d' , we have $\Pr[U_{l+s} \in A_g] \leq 1/4$ as

$$s \geq (l+s) \frac{\log^{d'} n}{m} + 2.$$

To conclude, recall from the statement of Theorem 1.1-(1) that $m = \log^{\omega(1)} n$. □

□

3.1 Constructing ρ

We now turn to the problem of constructing ρ that satisfies Properties (I) and (II). Again, our construction of ρ will be probabilistic. That is, we will show a distribution on restrictions that satisfies both Properties (I) and (II) w.h.p.. This certainly guarantees the existence of one fixed ρ that satisfies both Properties (I) and (II). We start with some intuition and then we give the actual construction.

Noise Sensitivity of Constant Depth Circuits: For property (II) we use the low noise sensitivity of constant depth circuits. Recall from Section 2 that the standard distribution on restrictions R_δ is the distribution on restrictions where each symbol is independently * with probability δ and otherwise it is a uniform independent random bit.

Lemma 3.2 ([LMN, Bop] Low Noise Sensitivity of Constant Depth Circuits). *Let $C : \{0, 1\}^t \rightarrow \{0, 1\}^{t'}$ be a circuit of size S and depth d . Let $\rho \in R_\delta$ then*

$$E_{\rho \in R_\delta, U_t, U_t'} \left[\Delta \left(C(U_t^\rho), C(U_t'^\rho) \right) \right] \leq O(\delta \cdot \log^{d-1} S).$$

For completeness, we show in Appendix A a simple derivation of Lemma 3.2 from known results [Vio]. Then, assuming G makes r queries to f , we have the following, taking expectations over random choice of uniform random functions $F, F' : \{0, 1\}^n \rightarrow \{0, 1\}^m$, random input $x \in \{0, 1\}^n$ and random $\rho \in R_\delta$.

$$\begin{aligned} &E \left[\Delta \left(G^{F^\rho}(x), G^{F'^\rho}(x) \right) \right] \\ &= E \left[\Delta \left(C_x(F^\rho(q_1), \dots, F^\rho(q_r)), C_x(F'^\rho(q_1), \dots, F'^\rho(q_r)) \right) \right] \\ &= E \left[\Delta \left(C_x(U_{rm}^\rho), C_x(U'_{rm}{}^\rho) \right) \right] \tag{2} \\ &\leq O(\delta \log^{d-1} S) \quad (\text{By Lemma 3.2}) \end{aligned}$$

Where Equation (2) follows from the definition of ρ and F^ρ , assuming without loss of generality that $G^{F^\rho}(x)$ never queries the same input twice. So by fixing $F' = g$ and then applying Markov's inequality, we have that most ρ satisfy Property (II) with the expectation at most $O(\delta \log^{d-1} S)$, which is at most $(\text{poly log } n)/m$, as required by Property (II), when $\delta \leq (\text{poly log } n)/m$ (recall the size of C_x is $S = \text{poly } n$). The conclusion is that for property (II) the standard distribution R_δ suffices when $\delta \leq (\text{poly log } n)/m$.

F^ρ one-way: Recall for Property (I) we want F^ρ to still be one-way after we fix ρ , even relative to an oracle that depends on ρ (i.e. A). The main problem here is that the restriction could conceivably leak information about the input. For example, if $m = 2n$, i.e. the range of F is $\{0, 1\}^{2n}$, then one could consider the pathological restriction $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}^{2n}$ such that for every x the first n symbols of $\rho(x)$ are x . Now the inversion of $F^\rho(x)$ is trivial, because the output completely reveals the input. A similar problem arises w.h.p. when ρ is selected according to the distribution R_δ for m sufficiently large. (In fact, we take advantage of this in Section 6.) To overcome this problem we need another idea. The idea is to ensure that for most x there is a superpolynomial number of y 's such that $\rho(x) = \rho(y)$. This implies that after we are given $F^\rho(U_n)$ we have little information about U_n , since information-theoretically U_n is uniform on a set of superpolynomial size. To achieve this we compose a random restriction $\rho : [b] \rightarrow \{0, 1, *\}^n$ in R_δ with a random function $h : \{0, 1\}^n \rightarrow [b]$ for $b = n^{\omega(1)}$. We now give the formal definition of this distribution and then show that it satisfies properties (I) and (II) w.h.p..

Definition 3.3. *The distribution \tilde{R} on restrictions $\rho_0 \circ h : \{0, 1\}^n \rightarrow \{0, 1, *\}^m$ is defined as follows. Set $\delta := \log^4 n/n$ and $b := n^{\log n}$. Let $\rho_0 : [b] \rightarrow \{0, 1, *\}^n$ be a random restriction in R_δ . Let $h : \{0, 1\}^n \rightarrow [b]$ be a random function, then*

$$(\rho_0 \circ h)(x) := \rho_0(h(x))$$

We now show that w.h.p. $\rho_0 \circ h$ satisfies both Properties (I) and (II). Thus Theorem 1.1-(1) follows from the next two lemmas.

Lemma 3.4. *A random $\rho = \rho_0 \circ h \in \tilde{R}$ satisfies Property (I) w.h.p.*

Proof. We can assume without loss of generality that $(\rho_0 \circ h)(y)$ contains at least $\log^2 n$ $*$'s for every y . This is because the probability that this does not happen is at most, using union bounds (recall $\delta = \log^4 n/m$, $m = \log^{\omega(1)} n$, $m = n^{O(1)}$, and $b = n^{\log n}$):

$$\begin{aligned} b \binom{m}{m - \log^2 n} (1 - \delta)^{m - \log^2 n} &= b \binom{m}{\log^2 n} (1 - \delta)^{m - \log^2 n} \\ &\leq n^{\log n} (e \cdot m)^{\log^2 n} (1 - \delta)^{\delta^{-1} \delta (m - \log^2 n)} \leq n^{\log n} e^{O(\log^3 n)} (1/e)^{\log^4 n - o(1)} \leq \epsilon(n). \end{aligned}$$

Now we fix any such $\rho_0 \circ h$ and we analyze the inversion probability over random F and random input $X \equiv U_n$.

By the pigeon hole principle, there are at most $b2^{n/2}$ inputs x such that there are fewer than $2^{n/2}$ y 's such that $h(x) = h(y)$, i.e. $|h^{-1}(h(x))| \leq 2^{n/2}$. Since $b = n^{\log n}$ there is only an exponentially small fraction of such x 's. So let us assume without loss of generality that X is such that there are at least $2^{n/2}$ y 's such that $h(X) = h(y)$.

In the following we restrict our attention to the queries M makes to the oracle function $F^{\rho_0 \circ h}$, and we make no assumption on the queries it makes to the adversary A . First, without loss of generality assume that M queries its output (clearly it does not hurt for M to check its answer before outputting it) and that M never queries any input twice (this can be accomplished by keeping a simple list of the inputs queried and of the oracle answers).

Since M queries its output, the probability (over F, X) that M inverts $F^{\rho_0 \circ h}(X)$ is the probability that (1) M queries X or (2) M queries $y \neq X$ such that $F^{\rho_0 \circ h}(y) = F^{\rho_0 \circ h}(X)$. We bound these two events separately.

(1): By our assumption, there are at least $2^{n/2}$ y 's such that $h(X) = h(y)$. Therefore, given $F^{\rho_0 \circ h}(X)$, X is uniform on a set of inputs of size at least $2^{n/2}$. Hence the probability that M queries X is negligible because M only makes $\text{poly}(n)$ queries.

A more formal argument goes as follows: Suppose M queries X with nonnegligible probability. We construct another (computationally unbounded, depending on h and ρ_0) machine M' *without oracle access to $F^{\rho_0 \circ h}$* that outputs a polynomial size list containing X with nonnegligible probability. But this is impossible because, as we said before, given $F^{\rho_0 \circ h}(X)$, X is uniform on a set of inputs of size at least $2^{n/2}$. M' simply simulates M and whenever M queries $F^{\rho_0 \circ h}$ at q , it adds q to the list and answers the query with $U'_m{}^{\rho_0 \circ h}(q)$ where U'_m is uniform and independent from all the previous query answers. It is easy to see that the probability that M queries X is the same as the probability that X is in the list that M' outputs.

(2): On the other hand, whenever M queries $y \neq X$ then by definition $F^{\rho_0 \circ h}(y) = U'_m{}^{\rho_0 \circ h}(y)$, where U'_m is uniform and independent from X and the state of M , since M never queries the same input twice. Since by our assumption $\rho_0 \circ h(y)$ has at least $\log^2 n$ *'s, it follows that with probability at most $1/n^{\log n}$ we have $F^{\rho_0 \circ h}(y) = F^{\rho_0 \circ h}(X)$. Again, since M only makes $\text{poly}(n)$ queries, the probability that M ever queries $y \neq X$ such that $F^{\rho_0 \circ h}(y) = F^{\rho_0 \circ h}(X)$ is negligible.

Therefore, the total probability that M inverts $F^{\rho_0 \circ h}(X)$ is negligible. \square

Lemma 3.5. *A random $\rho_0 \circ h \in \tilde{R}$ satisfies Property (II) w.h.p.*

Proof. In order to show Property (II) all we need is Equation (2) in Page 11 to go through *approximately*. Let $F, F' : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be random uniform random functions, $x \in \{0, 1\}^n$ a random input, $\rho_0 \circ h$ a random restriction in \tilde{R} and ρ a random restriction in R_δ for $\delta := \log^4 n/m$ (i.e. the same δ in Definition 3.3.) Consider the random variable

$$V := \Delta\left(C_x(F^{\rho_0 \circ h}(q_1), \dots, F^{\rho_0 \circ h}(q_r)), C_x(F'^{\rho_0 \circ h}(q_1), \dots, F'^{\rho_0 \circ h}(q_r))\right).$$

We show

$$E[V] \leq E\left[\Delta\left(C_x(U_{rm}^\rho), C_x(U'_{rm}{}^\rho)\right)\right] + \epsilon(n). \quad (3)$$

Note Lemma 3.5 follows from Inequality 3 as explained before in Page 11: we use Lemma 3.2 to bound $E[\Delta(C_x(U_{rm}^\rho), C_x(U'_{rm}{}^\rho))]$ by $O(\delta \text{poly log } n)$, then we fix $F' = g$ and use Markov's inequality. So all we need to show is that Inequality 3 holds:

$$\begin{aligned} E[V] &\leq E\left[V \mid \forall i \neq j, h(q_i) \neq h(q_j)\right] + \Pr[\exists i \neq j : h(q_i) = h(q_j)] \\ &\leq E\left[\Delta\left(C(U_{rm}^\rho), C(U'_{rm}{}^\rho)\right)\right] + r^2/b \end{aligned}$$

Where we use the fact that conditioned on the event “ $h(q_i) \neq h(q_j)$ for every $i \neq j$ ” the induced distribution of $\rho_0 \circ h$ is exactly R_δ . And then we use the fact that a random function mapping in $[b]$ has collision probability $1/b$, i.e., for every $a \neq b$ we have $\Pr_h[h(a) = h(b)] \leq 1/b$. Since G makes only r queries to f , the probability that there are $i \neq j$, such that $h(q_i) = h(q_j)$ is at most r^2/b . Noticing that r^2/b is negligible because $b = n^{\log n}$ and $r = \text{poly}(n)$ concludes the proof. \square

4 PRG constructions in constant depth circuits

In this section we show our PRG construction in constant-depth circuits from one-way permutations, i.e., we prove the following theorem which is a restatement of Theorem 1.2-(1).

Theorem 4.1 (Theorem 1.2-(1)). *If there is a one-way permutation $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ in constant depth circuits then there is PRG $G : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n+1}$ in constant depth circuits.*

We use the same pseudorandom distribution of Goldreich and Levin [GL], and our only difficulty is showing how it can be generated in constant depth circuits. We denote by $\langle x, y \rangle$ the Goldreich-Levin general hard-core predicate [GL], i.e. $\sum_i x_i y_i \pmod{2}$.

Theorem 4.2 ([GL]). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a one-way permutation. Then*

$$GL^f(x, y) := (f(x), y, \langle x, y \rangle)$$

is a PRG.

At first glance GL^f does not seem to be computable in constant depth circuits, because parity is not [FSS, Hås]. In the following lemma we show how to circumvent this problem.

Lemma 4.3. *There is a constant depth circuit $C : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n+1}$ such that for every $x \in \{0, 1\}^n, C(x, U_n)$ is distributed as $(U'_n, \langle x, U'_n \rangle)$.*

Theorem 4.1 follows from Lemma 4.3 simply defining $G^f(x, y) := (f(x), C(x, y))$.

The key observation to prove Lemma 4.3 is that while constant depth circuits cannot compute the parity function, constant depth circuits *can* generate a random x together with its parity. To see this, consider the constant depth circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$ such that $C(r_1, \dots, r_n) := (r_1, r_2 \oplus r_1, r_3 \oplus r_2, \dots, r_n \oplus r_{n-1}, r_n)$. It is easy to see that $C(U_n)$ outputs a random value in $\{0, 1\}^n$ and its parity, and moreover C is constant depth. This observation is from [BL].

To prove Lemma 4.3 we use the same approach, but only on the bits of x that are 1.

Proof of Lemma 4.3. Let the input be $x = x_1, \dots, x_n$ and $r = r_1, \dots, r_n$, and consider the circuit $C : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{n+1}, C(x, r) = r'b$, where $r' = r'_1, \dots, r'_n$ and $b \in \{0, 1\}$, defined as follows:

$$r'_i := \begin{cases} r_i & \text{if } x_i = 0 \\ r_i & \text{if } x_i = 1 \text{ and } \nexists j < i : x_j = 1 \\ r_i \oplus r_j & \text{if } x_i = 1 \text{ and } j \\ & \text{is the biggest index } j < i : x_j = 1 \end{cases}$$

and

$$b := r_i \text{ where } i \text{ is the biggest index such that } x_i = 1 \quad (b = 0 \text{ if there is no such } i).$$

It is not too hard to see that $C(x, U_n)$ is distributed like $(U_n, \langle x, U_n \rangle)$. It is also easy to check that C can be implemented in constant depth. (Indeed, this follows from the fact that we defined it using first-order logic.) \square

For constructions from one-to-one and regular one-way functions the only additional thing we need are extractors. While constant-depth circuits cannot in general compute extractors with good parameters [Vio], it can be shown that they can compute extractors (specifically, one based on the hash function due to Carter and Wegman [CW]) for the parameters of interest here (i.e. seed length polynomial in the source length) (details omitted).

5 Worst-case Hardness Amplification

In this section we prove Theorem 1.4. Let us first recall the definition of a hard function.

Definition 5.1. *A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is worst-case-hard (resp., ϵ -hard) for size S if every circuit of size S fails to compute f on some input (resp., on at least ϵ fraction of inputs).*

Before proving Theorem 1.4 we make some remarks.

Remarks on Theorem 1.4: (1) We focus on amplification up to constant (i.e. .3). Notice that by Yao's XOR lemma (cf. [GNW]) if PH has a $1/\text{poly}(n)$ -hard function for size $S'(n)$ then PH has a .3-hard function for size $\text{poly}(S'(n^\Omega(1)))$.

(2) We required (in the statement of the theorem) that for every c and for sufficiently large n , $S(n) \leq 2^n/n^c$. This is because for $S \geq 2^n/n^c$ for some fixed constant c , the oracle is already $1/\text{poly}(n)$ -hard by a counting argument (given in [Vio]), and therefore by the previous item PH has a .3-hard function.

(3) Similar results hold for hardness amplifications $\text{Amp}^f : \{0, 1\}^{l(n)} \rightarrow \{0, 1\}$ running in time $t(n)$ with a constant number of alternations, for a wide range of parameters $l(n), t(n)$. Here we set $l(n) = \text{poly}(n)$ and $t(n) = \text{poly}(n)$ for simplicity of exposition.

In this section functions are boolean. In particular F will denote a uniform random function $F : \{0, 1\}^n \rightarrow \{0, 1\}$. Accordingly, we take restrictions ρ on 2^n bits, $\rho : \{0, 1\}^n \rightarrow \{0, 1\}$, which we see as a partial assignment to the truth table of $f : \{0, 1\}^n \rightarrow \{0, 1\}$.

To prove Theorem 1.4 we build a certain pseudorandom distribution on restrictions. This is the main technical lemma of this section.

Lemma 5.2. *For every constant c there is a distribution \tilde{R}^c on restrictions $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}$ such that:*

(1) *Every $\rho \in \tilde{R}^c$ is described by $\text{poly}(n)$ bits σ . We denote by ρ_σ the restriction described by σ . For random σ , we have that ρ_σ is random in \tilde{R}^c . There is a polynomial time algorithm such that given σ and $x \in \{0, 1\}^n$, computes $\rho_\sigma(x)$.*

(2) For every circuit C of size 2^{n^c} and depth c on $t := 2^n$ bits: w.p. $1 - o(1)$ over $\rho_\sigma \in \tilde{R}^c$,

$$\text{Bias}_{U_t}[C(U_t^{\rho_\sigma})] \geq 1 - o(1).$$

(Where the Bias of a 0 – 1 random variable X is $|\Pr[X = 0] - \Pr[X = 1]|$.)

(3) There is a constant d such that w.h.p. over $\rho_\sigma \in \tilde{R}^c$, ρ_σ has at least $2^n/n^d$ *'s.

We now assume the above Lemma and prove Theorem 1.4.

Proof of Theorem 1.4. By standard techniques (see e.g., [FSS, Hås]), the oracle algorithm in PH can be turned into an exponential size constant-depth circuit whose input is the truth table of the oracle f . In particular, let c be such that $\text{Amp}^f(x)$ has depth c and size 2^{n^c} when turned into a constant depth circuit whose only input is the truth table of f . Consider the distribution on restrictions \tilde{R}^c whose existence is guaranteed by Lemma 5.2. Consider $\text{Amp}^{F^{\rho_\sigma}}$, where $\rho_\sigma \in \tilde{R}^c$. We need a couple of lemmas.

Lemma 5.3. *W.h.p. over $\rho_\sigma \in \tilde{R}^c$ and F , $F^{\rho_\sigma} : \{0, 1\}^n \rightarrow \{0, 1\}$ is worst-case hard for size $S(n)$. In particular, w.h.p. over $\rho_\sigma \in \tilde{R}^c$ and F , $\text{Amp}^{F^{\rho_\sigma}} : \{0, 1\}^{n^a} \rightarrow \{0, 1\}$ is .3-hard for size $S'(n)$.*

Lemma 5.4. *There is a PH machine A' such that given σ and an input x rounds $\text{Amp}^{F^{\rho_\sigma}}(x)$ to its most likely value, over the choice of F , whenever $\text{Bias}_F[\text{Amp}^{F^{\rho_\sigma}}(x)] \geq .2$. I.e., if $\Pr_F[\text{Amp}^{F^{\rho_\sigma}}(x) = 1] \geq .6$ then $A'(\sigma, x) = 1$, and if $\Pr_F[\text{Amp}^{F^{\rho_\sigma}}(x) = 0] \geq .6$ then $A'(\sigma, x) = 0$.*

Now for the proof of Theorem 1.4. By Lemma 5.2, for every x , w.p. $1 - o(1)$ over σ , $\text{Bias}_F[\text{Amp}^{F^{\rho_\sigma}}(x)] \geq 1 - o(1)$. This holds because, for fixed x , $\text{Amp}^f(x)$ is a constant depth function of the truth table of f . Let A' be the oracle PH machine in Lemma 5.4. By Lemma 5.4 we have:

$$\Pr_{x, \sigma, F}[A'(\sigma, x) \neq \text{Amp}^{F^{\rho_\sigma}}(x)] \leq o(1) + o(1) = o(1). \quad (4)$$

Thus there is a function $\eta(n) = o(1)$ such that

$$\Pr_{\sigma, F}[\Delta(A'(\sigma, \cdot), \text{Amp}^{F^{\rho_\sigma}}) \geq \eta(n)] \leq \eta(n). \quad (5)$$

Where $\Delta(f, f')$ denotes the relative Hamming distance of the truth tables of $f, f' : \{0, 1\}^n \rightarrow \{0, 1\}$. Thus:

$$\begin{aligned} & \Pr_{\sigma, F}[A'(\sigma, \cdot) \text{ is not } .2\text{-hard for size } S'(n)] \leq \\ & \leq \Pr_{\sigma, F}[\Delta(A'(\sigma, \cdot), \text{Amp}^{F^{\rho_\sigma}}) > .1 \text{ or } \text{Amp}^{F^{\rho_\sigma}} \text{ is not } .3\text{-hard for size } S'(n)] \\ & \leq o(1) + o(1) \quad (\text{By Inequality (5) and Lemma 5.3}) \\ & \leq o(1). \end{aligned}$$

(where we use the fact that if $A'(\sigma, \cdot)$ is at relative Hamming distance at most .1 from $\text{Amp}^{F^{\rho_\sigma}}$ that is .3-hard, then $A'(\sigma, \cdot)$ must be .2-hard. For else the same circuit computing $A'(\sigma, \cdot)$ with error less than .2 would compute $\text{Amp}^{F^{\rho_\sigma}}$ with error less than $.2 + .1 = .3$).

Now, since w.h.p. over σ we have that $A'(\sigma, \cdot)$ is $.2$ -hard for size $S'(n)$, we have that every circuit of size $S'(n)$ fails to compute A' on at least a $.2(1 - o(1)) > .1$ fraction of inputs, and thus A' is $.1$ -hard for size $S'(n)$.

To finish the proof note that A' is in PH and that it has input length n^b for some b . \square

We now discuss Lemmas 5.3 and 5.4.

Proof of Lemma 5.3. This is a simple counting argument. By Lemma 5.2 there is a constant d such that ρ has at least $2^n/n^d$ $*$'s w.h.p.. Whenever this happens, F^ρ is uniform (over the choice of F) on a set of $2^{2^n/n^d}$ functions. There are at most $2^{O(S \cdot \log S)}$ circuits of size S . By assumption, for sufficiently large n , $S(n) \leq 2^n/n^{d+2}$. This means in particular that for sufficiently large n , $O(S \cdot \log S) \leq (2^n/n^d)/2$. Therefore, for sufficiently large n , the probability (over F) that F^ρ is not worst-case hard for size S is at most $2^{O(S(n) \cdot \log S(n)) - 2^n/n^d} \leq 2^{2^n/(2 \cdot n^d)} = o(1)$. \square

Lemma 5.4 was essentially proved by Nisan in [Nis] (see also [NW]) (In [Nis, NW] the lemma is stated as “almost- $PH=PH$ ”). We omit the details here. (Note Lemma 5.4 does not immediately follow from the more well-known fact that $BPP \subseteq PH$ (even though the latter result is used in Nisan’s proof). This is because in 5.4 the machine Amp has access (through the oracle) to an exponential (as opposed to polynomial) number of random bits.)

5.1 Pseudorandom restrictions

In this section we prove Lemma 5.2. A key tool is Nisan’s pseudorandom generator against constant depth circuits.

Theorem 5.5 ([Nis]). *For every constant c and every n there is a generator $Nis : \{0, 1\}^{\text{poly log } n} \rightarrow \{0, 1\}^n$ such that (1) given x and $i \leq n$ we can compute the i -th bit of $Nis(x)$ in time $\text{poly log}(n)$ and (2) N is $1/n$ -pseudorandom for circuits of size n and depth c . That is, for every circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size n and depth c :*

$$\left| \Pr[C(Nis(U_{\text{poly log } n})) = 1] - \Pr[C(U_n) = 1] \right| \leq 1/n.$$

Proof of Lemma 5.2. Let $\delta := 1/n^{c^2}$. We know by Lemma 3.2 that for every circuit $C : \{0, 1\}^t \rightarrow \{0, 1\}$ of size 2^{n^c} and depth c on t bits:

$$\Pr_{\rho \in R_\delta, U_t, U'_t} [C(U_t^\rho) \neq C(U'_t^\rho)] \leq O(\delta \log^{c-1} 2^{n^c}) = o(1). \quad (6)$$

The above equation in turn implies that w.p. $1 - o(1)$ over R_δ , $\text{Bias}_{U_t}[C(U_t^\rho)] \geq 1 - o(1)$ (see below). Moreover by a Chernoff bound the fraction of $*$'s in $\rho \in R_\delta$ will be concentrated around δ .

So R_δ satisfies Items (2) and (3) in Lemma 5.2. But the problem is that R_δ requires at least 2^n bits to be generated, while we aim to a distribution on restrictions which can be generated with $\text{poly } n$ bits. To this aim we *derandomize* R_δ .

Let W be a canonical circuit that given $I := O(2^n \log(1/\delta))$ random bits generates R_δ (say we use blocks of $O(\log(1/\delta))$ bits to put a $*$ with probability δ). It is easy to see that there is such a circuit W of size $\text{poly}(2^n)$ and depth $O(1)$.

We now define \tilde{R}^c . Consider $Nis : \{0, 1\}^{\log^d I} \rightarrow \{0, 1\}^I$ which is $2^{n^{c'}}$ -pseudorandom against depth c' , for a constant c' to be determined later. For every c' there is a constant d such that such a Nis exists according to theorem 5.5. Then

$$\tilde{R}^c := W(Nis(U_{\log^d I})).$$

We now prove that \tilde{R}^c has the required properties.

(1) By construction \tilde{R}^c can be generated with $\text{poly}(n)$ random bits. As shown in [Nis], given i and the random bits σ we can compute the i -th symbol of $\rho_\sigma \in \tilde{R}^c$ in polynomial time.

(2) Let $t := 2^n$. First we show that, even under such restrictions, circuits of size 2^{n^c} and depth c still have low noise sensitivity, the claim about the bias then easily follows. To show this we use an approach similar to one used in [HVV]. As noticed there, the noise sensitivity of a constant depth circuit C equals the acceptance probability of another (slightly bigger) constant depth circuit C' defined as follows: Given a restriction ρ , C' tosses coins for U_t and U'_t and answers 1 if and only if $C(U_t^\rho) \neq C(U'_t{}^\rho)$. It is easy to see that such a C' can be implemented in constant depth. Combining C' with our constant depth circuit W that given random bits generates a random restriction in R_δ we obtain another constant depth circuit $C'' := C' \circ W$. Now, the acceptance probability of C'' over a truly random input is the noise sensitivity of C with respect to R_δ , while the acceptance probability of C'' over a pseudorandom input generated using Nisan's PRG Nis is the noise sensitivity of C with respect to \tilde{R}^c . Therefore, since C'' cannot distinguish the output of Nisan's PRG from truly random, we deduce that the noise sensitivity of C with respect to \tilde{R}^c is close to the noise sensitivity of C with respect to R_δ .

Therefore, choosing a sufficiently large constant c' in the definition of \tilde{R}^c we have that, for every $C : \{0, 1\}^t \rightarrow \{0, 1\}$ of size 2^{n^c} and depth c :

$$\begin{aligned} & \Pr_{\rho_\sigma \in \tilde{R}^c, U_t, U'_t} [C(U_t^{\rho_\sigma}) \neq C(U'_t{}^{\rho_\sigma})] \\ & \leq \Pr_{\rho \in R_\delta, U_t, U'_t} [C(U_t^\rho) \neq C(U'_t{}^\rho)] + o(1) \quad (\text{by pseudorandomness}) \\ & \leq o(1) \quad (\text{by Equation (6)}) \end{aligned}$$

We now deduce a claim about the bias. By above there is a function $\eta(n) = o(1)$ such that

$$\Pr_{\rho_\sigma \in \tilde{R}^c} \left[\Pr_{U_t, U'_t} [C(U_t^{\rho_\sigma}) \neq C(U'_t{}^{\rho_\sigma})] \leq \eta(n) \right] \geq 1 - \eta(n).$$

Noticing that $\Pr_{U_t, U'_t} [C(U_t^{\rho_\sigma}) \neq C(U'_t{}^{\rho_\sigma})] \leq o(1)$ implies that $\text{Bias}_{U_t}[C(U_t^{\rho_\sigma})] = 1 - o(1)$ concludes the proof of this item.

(3) Ajtai [Ajt] shows the following:

Lemma 5.6 ([Ajt]). *For every i there is a circuit C of size $\text{poly}(2^n)$ and depth $O(1)$ such that, given u and a bit string of length 2^n :*

If the bit string has more than $u + 2^n/n^i$ occurrences of '1' then C outputs 1.

If the bit string has fewer than $u - 2^n/n^i$ occurrences of '1' then C outputs 0.

We expect a random $\rho \in R_\delta$ to have $\delta 2^n$ $*$'s. By a concentration bound the probability that it has less than $(\delta 2^n)/2$ is $o(1)$. Since $\delta = 1/n^{c^2} = 1/\text{poly}(n)$, by Lemma 5.6 there is a constant depth circuit of size $\text{poly}(2^n)$ that can distinguish the cases, say, “more than $\delta/2$ fraction of $*$'s” and “less than $\delta/3$ fraction of $*$'s” (setting $u := (\delta 2^n)/2.5$ in Lemma 5.6). Therefore, choosing a sufficiently large constant c' in the definition of \tilde{R}^c , by pseudorandomness, we have that $\rho_\sigma \in \tilde{R}^c$ has at least $(\delta 2^n)/3$ $*$'s w.h.p.. \square

6 PRG constructions from one-to-one one-way functions

In this section we prove our negative result for black-box parallel PRG constructions from one-to-one one-way functions, i.e. the proof of Theorem 1.1-(2).

The problem is that the functions $F^{\rho_0 \circ h}$ defined in Section 3 are not one-to-one. To ensure this property, we define another distribution on restrictions and from this a new distribution on one-to-one functions. This definition is slightly elaborate because injectivity is in tension with the fact that we need $\rho(x)$ to *not* uniquely identify x (to preserve the one-wayness of the oracle, see Section 3).

We denote by $\rho(x)_k \in \{0, 1, *\}$ the k -th symbol of $\rho(x)$. We say that a restriction $\rho : [b] \rightarrow \{0, 1, *\}^{cn}$ *splits* if for every $i \neq j$ there is $k > \log^2 n$ such that $\rho(i)_k = 1$ and $\rho(j)_k = 0$, or $\rho(i)_k = 0$ and $\rho(j)_k = 1$. The idea is that if ρ splits then for every function $f : [b] \rightarrow \{0, 1\}^{cn}$ we have that the function f^ρ is injective. For technical reasons we require $k > \log^2 n$.

Definition 6.1. *Let $c := 5$. The distribution \overline{R} on restrictions $\overline{\rho_0 \circ h} : \{0, 1\}^n \rightarrow \{0, 1, *\}^{cn}$ is defined in stages as follows. Set $\delta := (\log^4 n)/n$ and $b := 2^{n - \log^2 n}$.*

Let $\overline{h} : \{0, 1\}^n \rightarrow [b]$ be a random function such that for every $i \in [b]$ there are exactly $n^{\log n}$ inputs $x \in \{0, 1\}^n$ such that $\overline{h}(x) = i$.

*Let $\rho' : [b] \rightarrow \{0, 1, *\}^{cn}$ be a random restriction in R_δ such that ρ' splits. (We can think of ρ' as being generated by repeated sampling from R_δ until found one that splits.)*

Then let the random restriction $\overline{\rho_0}$ be equal to ρ' except for every i if $\rho'(i)$ contains less than $\log^2 n$ $$'s, then set $\overline{\rho_0}(i)_j := *$ for every $j \leq \log^2 n$. (I.e. this forces $\overline{\rho_0}(i)$ to have at least $\log^2 n$ $*$'s for every i .) Then define*

$$\overline{\rho_0 \circ h}(x) := \overline{\rho_0}(\overline{h}(x))$$

By $\overline{F^{\rho_0 \circ h}} : \{0, 1\}^n \rightarrow \{0, 1\}^{cn}$ we denote a random one-to-one function F consistent with $\overline{\rho_0 \circ h}$, i.e. such that $F(x)^{\overline{\rho_0 \circ h}(x)} = F(x)$ for every x . In other words, $\overline{F^{\rho_0 \circ h}}$ is a random function obtained from the truth table of $\overline{\rho_0 \circ h}$ replacing the $*$'s with random bits, *conditioned on the event that $\overline{F^{\rho_0 \circ h}}$ is one-to-one.*

It is easy to check that the space of restrictions \overline{R} is not empty, i.e. there exist restrictions that satisfy Definition 6.1. It is also easy to see that this guarantees that the space of functions $\overline{F^{\rho_0 \circ h}}$ is not empty, because in Definition 6.1 ρ splits and, for every $i \leq b$, $\rho(i)$ has at least $\log^2 n$ $*$'s and finally there are only $n^{\log n}$ inputs mapping to the same i through h .

All that is left to do is to show that $\overline{\rho_0 \circ h}$ satisfies Properties (I) and (II) from Page 10 w.h.p.. Of course, these properties must now be satisfied for our new space of random functions, namely $\overline{F^{\rho_0 \circ h}}$. Since this slightly changes the properties, we now repeat them.

- i. For every oracle A , with high probability over $\overline{F^{\rho_0 \circ h}}$, $M^{\overline{F^{\rho_0 \circ h}}, A}$ does not invert $\overline{F^{\rho_0 \circ h}}$, i.e.:

$$\Pr_{\overline{F^{\rho_0 \circ h}}, U_n} \left[M^{\overline{F^{\rho_0 \circ h}}, A}(\overline{F^{\rho_0 \circ h}}(U_n)) = U_n \right] \leq \epsilon(n).$$

- ii. There is a fixed function $g : \{0, 1\}^n \rightarrow \{0, 1\}^{c \cdot n}$ such that

$$E_{\overline{F^{\rho_0 \circ h}}, U_l} \left[\Delta \left(G^{\overline{F^{\rho_0 \circ h}}}(U_l), G^g(U_l) \right) \right] \leq \frac{\text{poly log}(n)}{m}.$$

Note Property (i) is simpler than Property (I) on Page 10 because now $\overline{F^{\rho_0 \circ h}}$ is always one-to-one.

Lemma 6.2. *A random $\overline{\rho_0 \circ h} \in \overline{R}$ satisfies Property (i) w.h.p..*

Proof. The proof is similar to the proof of Lemma 3.4. Let $X \equiv U_n$ be a random input. As in Lemma 3.4 assume that M queries its output and that never queries the same input twice. Given $\overline{F^{\rho_0 \circ h}}(X)$, by construction X is uniform over a set superpolynomial size. Since M only makes $\text{poly}(n)$ query, M queries X with negligible probability.

A more formal argument goes as follows: Suppose M queries X with nonnegligible probability. We construct another (computationally unbounded, depending on \overline{h} and $\overline{\rho_0}$) machine M' without oracle access to $\overline{F^{\rho_0 \circ h}}$ that outputs a polynomial size list containing X with nonnegligible probability. But this is impossible because, as we said before, given $\overline{F^{\rho_0 \circ h}}(X)$, X is uniform on a set of inputs of superpolynomial size. M' simply simulates M and whenever M queries $\overline{F^{\rho_0 \circ h}}$ at q , it adds q to the list and answers the query with $U'_m{}^{\overline{\rho_0 \circ h}(q)}$, where U'_m is a uniform and independent random variable such that $U'_m{}^{\overline{\rho_0 \circ h}(q)}$ is different from all the previous query answers and from $\overline{F^{\rho_0 \circ h}}(X)$ (this is always possible since M only makes $\text{poly}(n)$ queries and $\overline{\rho_0}(i)$ contains at least $\log^2 n$ *'s for every i). It is easy to see that the probability that M queries X is the same as the probability that X is in the list that M' outputs. \square

Lemma 6.3. *A random $\overline{\rho_0 \circ h} \in \overline{R}$ satisfies Property (ii) w.h.p..*

Proof. Again, in order to show Property (II) all we need is Inequality (2) in Page 11 to go through approximately. Let $x \in \{0, 1\}^l$ be a random input, q_1, \dots, q_r the $r \leq \text{poly}(n)$ queries made by C_x , $\overline{\rho_0 \circ h}$ be random in \overline{R} , let $\rho' : [b] \rightarrow \{0, 1, *\}^m$ and $\rho \in \{0, 1, *\}^{rm}$ be truly random restrictions in R_δ (for $\delta := (\log^4 n)/n$ as in Definition 6.1). Let $c := 5$ (again as in Definition 6.1) and $m := c \cdot n$.

Consider the random variable

$$V := \Delta \left(C_x(\overline{F^{\rho_0 \circ h}}(q_1), \dots, \overline{F^{\rho_0 \circ h}}(q_r)), C_x(\overline{F^{\rho_0 \circ h}}(q_1), \dots, \overline{F^{\rho_0 \circ h}}(q_r)) \right).$$

As in Lemma 3.5, Lemma 6.3 follows from the following inequality:

$$E[V] \leq E \left[\Delta \left(C_x(U_{rm}^\rho), C_x(U'_{rm}{}^\rho) \right) \right] + \epsilon(n). \quad (7)$$

To prove Inequality 7, notice that $E[V]$ is at most

$$\begin{aligned}
& E \left[V \mid \forall i \neq j : \bar{h}(q_i) \neq \bar{h}(q_j) \right] + \Pr \left[\forall i \neq j : \bar{h}(q_i) \neq \bar{h}(q_j) \right] \\
\leq & E \left[V \mid \forall i \neq j : \bar{h}(q_i) \neq \bar{h}(q_j) \right] + \epsilon(n) \\
= & E \left[\Delta \left(C_x(U_m^1 \overline{\rho_0 \circ h}(q_1)), \dots, U_m^r \overline{\rho_0 \circ h}(q_r), C_x(U_m^1 \overline{\rho_0 \circ h}(q_1)), \dots, U_m^r \overline{\rho_0 \circ h}(q_r)) \right) \mid \right. \\
& \left. \forall i \neq j : \bar{h}(q_i) \neq \bar{h}(q_j) \right] + \epsilon(n) \\
\leq & E \left[\Delta \left(C_x(U_m^1 \rho'(\bar{h}(q_1))), \dots, U_m^r \rho'(\bar{h}(q_r)), C_x(U_m^1 \rho'(\bar{h}(q_1))), \dots, U_m^r \rho'(\bar{h}(q_r))) \right) \mid \right. \\
& \left. \forall i \neq j : \bar{h}(q_i) \neq \bar{h}(q_j) \right] + \epsilon(n) \\
\leq & E \left[\Delta \left(C_x(U_{rm}^\rho), C_x(U_{rm}'^\rho) \right) \right] + \epsilon(n)
\end{aligned}$$

Above, the second inequality follows from the fact that $b = n^{\omega(1)}$ (this is not too hard to check). In the next equality we use the fact that the distribution $(\overline{F}^{\rho_0 \circ h}(q_1), \dots, \overline{F}^{\rho_0 \circ h}(q_r))$ equals the distribution of $(U_m^1 \overline{\rho_0 \circ h}(q_1), \dots, U_m^r \overline{\rho_0 \circ h}(q_r))$ whenever for every $i \neq j$ we have $\bar{h}(q_i) \neq \bar{h}(q_j)$ (here we use that $\overline{\rho_0}$ splits).

In the next inequality (i.e. in the second to last), we replace $\overline{\rho_0}$ with ρ' (recall $\rho' : [b] \rightarrow \{0, 1\}^m$ is a truly random restriction). The inequality follows from the fact that, with high probability, the distribution of any $r = \text{poly}(n)$ fixed values of $\overline{\rho_0}$ looks like the distribution of the corresponding values of ρ' . To show this latter claim we need to bound two probabilities.

First, the probability that $\rho' \in R_\delta$ does not split can be bound as follows. Fix $i \neq j$. The probability that does not exist $k > \log^2 n$ such that $\rho'(i)_k = 1$ and $\rho'(j)_k = 0$, or $\rho'(i)_k = 0$ and $\rho'(j)_k = 1$ is at most

$$(\delta + \delta + 1/2)^{cn - \log^2 n} \leq (2/3)^{4n}.$$

So by a union bound the probability that there are $i \neq j$ such that does not exist $k > \log^2 n$ such that $\rho'(i)_k = 1$ and $\rho'(j)_k = 0$, or $\rho'(i)_k = 0$ and $\rho'(j)_k = 1$ is at most $(2^n)^2 \cdot (2/3)^{4n}$ which is negligible.

Second, the probability that there exists $i \leq r$ such that $\rho'(\bar{h}(q_i))$ has less than $\log^2 n$ *'s is at most

$$r \binom{cn}{n - \log^2 n} (1 - \delta)^{cn - \log^2 n}$$

which is negligible (recall $r = \text{poly}(n)$ and cf. the similar bound inside the proof of Lemma 3.4). \square

7 Open Problems

(1) Is there a parallel black-box PRG construction with linear stretch from one-way permutations?

(2) Is there a *uniform* parallel black-box PRG construction with any stretch from one-way functions? Our techniques only give a nonuniform one.

(3) In Theorem 1.2-(2), can we achieve $l = (n+m)$ poly log n ? This would match Theorem 1.1-(2). It would be enough to show the existence of an extractor (that extracts almost all the min-entropy) with linear seed length and computable by constant depth circuits. Such an extractor is not ruled out by [Vio] nor given by our positive results, as we only get polynomial seed length.

8 Acknowledgements

Many thanks to Salil Vadhan for helpful discussions, pointing out a mistake in an early version of this work and for his very helpful reading of this paper. We thank Oded Goldreich for helpful discussions and comments. We thank Alex Healy for helpful conversations, rediscovering with us some of the results in [BL] that then led to Lemma 4.3, and for his helpful reading of this paper. We thank Minh Nguyen for her helpful reading of the introduction. We thank Yuval Ishai for sending us an early version of [AIK] and for helpful discussions about PRGs. We thank Ronen Shaltiel for helpful discussions about PRGs.

References

- [Ajt] M. Ajtai. Σ_1^1 -formulae on finite structures. *Ann. Pure Appl. Logic*, 24(1):1–48, 1983.
- [ABO] M. Ajtai and M. Ben-Or. A Theorem on Probabilistic Constant Depth Computation. In ACM, editor, *Proceedings of the sixteenth annual ACM Symposium on Theory of Computing, Washington, DC, April 30–May 2, 1984*, pages 471–474, 1984. ACM order no. 508840.
- [AIK] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in NC^0 . In *45th Annual Symposium on Foundations of Computer Science*, Rome, ITALY, 17–19 Oct. 2004. IEEE.
- [BFL] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Comput. Complexity*, 1(1):3–40, 1991.
- [BF] D. Beaver and J. Feigenbaum. Hiding Instances in Multioracle Queries. In *7th Annual Symposium on Theoretical Aspects of Computer Science*, volume 415 of *Lecture Notes in Computer Science*, pages 37–48, Rouen, France, 22–24 Feb. 1990. Springer.
- [BM] M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. *SIAM J. on Computing*, 13(4):850–864, Nov. 1984.
- [BT] A. Bogdanov and L. Trevisan. On Worst-Case to Average-Case Reductions for NP Problems. In *44th Annual Symposium on Foundations of Computer Science*, Cambridge, Massachusetts, 11–14 Oct. 2003. IEEE.
- [Bop] R. B. Boppana. The average sensitivity of bounded-depth circuits. *Inform. Process. Lett.*, 63(5):257–261, 1997.

- [BL] R. B. Boppana and J. C. Lagarias. One-way functions and circuit complexity. *Inform. and Comput.*, 74(3):226–240, 1987.
- [CPS] J.-Y. Cai, A. Pavan, and D. Sivakumar. On the Hardness of the Permanent. In *16th International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Trier, Germany, March 4–6 1999. Springer-Verlag. To appear.
- [CSS] J.-Y. Cai, D. Sivakumar, and M. Strauss. Constant Depth Circuits and the Lutz Hypothesis. In *38th Annual Symposium on Foundations of Computer Science*, pages 595–604, Miami Beach, Florida, 20–22 Oct. 1997. IEEE.
- [CW] J. L. Carter and M. N. Wegman. Universal classes of hash functions. *J. Comput. System Sci.*, 18(2):143–154, 1979.
- [FL] U. Feige and C. Lund. On the Hardness of Computing the Permanent of Random Matrices. *Computational Complexity*, 6(2):101–132, 1996.
- [FF] J. Feigenbaum and L. Fortnow. Random-Self-Reducibility of Complete Sets. *SIAM J. on Computing*, 22(5):994–1005, Oct. 1993.
- [FSS] M. L. Furst, J. B. Saxe, and M. Sipser. Parity, Circuits, and the Polynomial-Time Hierarchy. *Mathematical Systems Theory*, 17(1):13–27, April 1984.
- [GT] R. Gennaro and L. Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 305–313. IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.
- [GNR] M. Goldmann, M. Näslund, and A. Russell. Complexity bounds on general hard-core predicates. *J. Cryptology*, 14(3):177–195, 2001.
- [Gol] O. Goldreich. *Foundations of Cryptography. Volume 1 - Basic Techniques*. Cambridge University Press, Cambridge, 2001.
- [GGM] O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. *J. ACM*, 33(4):792–807, Oct. 1986.
- [GKL] O. Goldreich, H. Krawczyk, and M. Luby. On the existence of pseudorandom generators. *SIAM J. Comput.*, 22(6):1163–1175, 1993.
- [GL] O. Goldreich and L. A. Levin. A Hard-Core Predicate for all One-Way Functions. In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, Washington, 15–17 May 1989.
- [GNW] O. Goldreich, N. Nisan, and A. Wigderson. On Yao’s XOR lemma. Technical Report TR95–050, Electronic Colloquium on Computational Complexity, March 1995. <http://www.eccc.uni-trier.de/eccc>.

- [Hås] J. Håstad. *Computational limitations of small-depth circuits*. MIT Press, 1987.
- [HILL] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396 (electronic), 1999.
- [HVV] A. Healy, S. Vadhan, and E. Viola. Using nondeterminism to amplify hardness. In *Proceedings of the Thirty-Six Annual ACM Symposium on the Theory of Computing*, pages 192–201, Chicago, IL, 13–15 June 2004. Invited to *SIAM Journal of Computing*, STOC Special Issue.
- [IN] R. Impagliazzo and M. Naor. Efficient Cryptographic Schemes Provably as Secure as Subset Sum. *Journal of Cryptology: the journal of the International Association for Cryptologic Research*, 9(4):199–216, Fall 1996.
- [IR] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, Washington, 15–17 May 1989.
- [KGY] M. Kharitonov, A. V. Goldberg, and M. Yung. Lower Bounds for Pseudorandom Number Generators. In *30th Annual Symposium on Foundations of Computer Science*, pages 242–247, Research Triangle Park, North Carolina, 30 Oct.–1 Nov. 1989. IEEE.
- [Lev] L. A. Levin. One way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.
- [LMN] N. Linial, Y. Mansour, and N. Nisan. Constant depth circuits, Fourier transform, and learnability. *J. Assoc. Comput. Mach.*, 40(3):607–620, 1993.
- [Lip] R. Lipton. New Directions in Testing. In *Proceedings of DIMACS Workshop on Distributed Computing and Cryptography*, 1989.
- [MNT] Y. Mansour, N. Nisan, and P. Tiwari. The Computational Complexity of Universal Hashing. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (May 14–16 1990: Baltimore, MD, USA)*, pages 235–243, New York, NY 10036, USA, 1990. ACM Press.
- [Nao] M. Naor. Bit Commitment Using Pseudorandomness. *Journal of Cryptology: the journal of the International Association for Cryptologic Research*, 4(2):151–158, 1991.
- [NR] M. Naor and O. Reingold. Number-Theoretic Constructions of Efficient Pseudo-Random Functions. In *38th Annual Symposium on Foundations of Computer Science*, pages 458–467, Miami Beach, Florida, 20–22 Oct. 1997. IEEE.
- [Nis] N. Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica*, 11(1):63–70, 1991.

- [NW] N. Nisan and A. Wigderson. Hardness vs Randomness. *J. Computer & Systems Sciences*, 49(2):149–167, Oct. 1994.
- [NZ] N. Nisan and D. Zuckerman. Randomness is Linear in Space. *J. Comput. Syst. Sci.*, 52(1):43–52, Feb. 1996.
- [RT] J. H. Reif and J. D. Tygar. Efficient Parallel Pseudo-Random Number Generation. In H. C. Williams, editor, *Advances in Cryptology—CRYPTO '85*, volume 218 of *Lecture Notes in Computer Science*, pages 433–446. Springer-Verlag, 1986, 18–22 Aug. 1985.
- [RTV] O. Reingold, L. Trevisan, and S. Vadhan. Notions of Reducibility between Cryptographic Primitives. In *Proceedings of the 1st Theory of Cryptography Conference (Feb 19-21, 2004: Cambridge, MA, USA)*. Springer-Verlag, 2004.
- [STV] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. System Sci.*, 62(2):236–266, 2001. Special issue on the Fourteenth Annual IEEE Conference on Computational Complexity (Atlanta, GA, 1999).
- [TV] L. Trevisan and S. Vadhan. Pseudorandomness and Average-Case Complexity via Uniform Reductions. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity*, pages 129–138, Montréal, CA, May 2002. IEEE.
- [Vio] E. Viola. The Complexity of Constructing Pseudorandom Generators from Hard Functions. *Comput. Complexity*, 13(3-4):147–188, 2004.
- [Yao] A. C. Yao. Theory and Applications of Trapdoor Functions (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, 3–5 Nov. 1982. IEEE.
- [YY] X. Yu and M. Yung. Space Lower-Bounds for Pseudorandom-Generators. In *Ninth Annual Structure in Complexity Theory Conference*, pages 186–197. IEEE Computer Soc., Los Alamitos, CA, 1994.

A Proof of Lemma 3.2

The next lemma, from [Vio], is based on the result in [Bop].

Lemma A.1 ([Vio]). *Let $D : \{0, 1\}^n \rightarrow \{0, 1\}$ be a circuit of size S and depth d . Let $X \in \{0, 1\}^n$ be a random input and let \tilde{X} be obtained from X by flipping each bit independently with probability $\delta < 1/2$. Then:*

$$\Pr_{X, \tilde{X}} [D(X) \neq D(\tilde{X})] \leq O(\delta \log^{d-1} S).$$

We obtain Lemma 3.2 as follows:

$$\begin{aligned}
E_{\rho \in R_\delta, U_t, U'_t} \left[\Delta \left(C(U_t^\rho), C(U'_t{}^\rho) \right) \right] &= \Pr_{\rho \in R_\delta, U_t, U'_t, i} [C(U_t^\rho)_i \neq C(U'_t{}^\rho)_i] \\
&= \Pr_{X, \tilde{X}, i} [C(X)_i \neq C(\tilde{X})_i] \\
&\quad (\text{where } \tilde{X} \text{ is obtained from } X \text{ flipping each bit independently with probability } \delta/2) \\
&\leq O(\delta \log^{d-1} S). \quad (\text{by Lemma A.1})
\end{aligned}$$

B Mildly black-box PRG construction

In this section we elaborate on our claim that (essentially) the negative result in Theorem 1.1-(1) also holds for *mildly* black-box PRG constructions, a less restrictive notion of black-box construction (cf. [RTV]). We give one necessary definition and then we state our result.

Definition B.1 (Mildly black-box PRG construction). *An oracle machine $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ is a mildly black-box PRG construction from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ if for every PPT A there exists an oracle PPT M such that for sufficiently large n , for every $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ if*

$$|\Pr[A(G^f(U_l))] - \Pr[A(U_{l+s})]| \geq 1/4$$

then

$$\Pr[f(M^f(f(U_n))) = U_n] \geq 1/n.$$

Theorem B.2 (This Paper). *Let $G^f : \{0, 1\}^l \rightarrow \{0, 1\}^{l+s}$ be a mildly black-box PRG construction (Def. B.1) in the form in Table 1, but with the additional requirement that the circuit C_x and the queries $q_{x,1}, \dots, q_{x,\text{poly}(n)}$ are generated from x in polynomial time (as opposed to arbitrarily).*

Suppose G^f starts from one-way function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with $m = \log^{\omega(1)} n$ and $m = n^{O(1)}$. Then there is a function $\eta(n) = o(l)$ such that if $s \geq \eta(n)$ then $P \neq NP$.

We now give some intuition about the proof. Consider the proof of our negative result for (fully) black-box PRG constructions (Section 3). Loosely speaking, we now want the adversary A to be a PPT. One of the problems is that our adversary depends on the function g in Property (II). Intuitively, however, A can invert G^g only knowing the restriction $\rho := \rho_0 \circ h$ (recall g is obtained probabilistically as $g = f^\rho$ for some f). This is because of the low noise sensitivity of constant depth circuits: more formally, we know that the output of the generator $G^{F^\rho}(x)$ will be very biased (over F). Therefore, if A' knew the restriction ρ and the input x to the generator, it could compute $\text{Round}G^g(x)$ which equals to $G^{F^\rho}(x)$ except each bit is rounded according to the bias (over F) of $G^{F^\rho}(x)$. Since G is biased, the output of $\text{Round}G^g(x)$ will be close in Hamming distance to $G^{F^\rho}(x)$. Thus A' could tell if an input comes from the generator by *guessing* an input x and a restriction ρ as follows:

$$A'(z) := 1 \text{ iff } \exists x \in \{0, 1\}^l, \exists \rho : \Delta(\text{Round}G(x, \rho), z) \leq \eta$$

We need to ensure two things. First A' must be efficient, and second A' should be a good distinguisher. Both problems will be solved by showing a certain pseudorandom distribution on restrictions (for the efficiency of A' we will also use $P = NP$). (Intuitively for A' to be a good distinguisher we need a description of restrictions which is shorter than the stretch of the generator.) More formally we obtain the following distribution on restrictions:

Lemma B.3. *For every constant c there is a distribution \tilde{R}^c on restrictions $\rho : \{0, 1\}^n \rightarrow \{0, 1, *\}^m$ such that:*

(1) *Every $\rho \in \tilde{R}^c$ is described by $\text{poly log } n$ bits σ . We denote by ρ_σ the restriction described by σ . For random σ , we have that ρ_σ is random in \tilde{R}^c . There is a polynomial time algorithm such that given σ and $i \in \{0, 1\}^n$, computes $\rho_\sigma(i)$.*

(2) *There is a function $\eta'(n) = o(1)$ such that for every circuit C of size n^c and depth c , and fixed $q_1, \dots, q_{n^c} \in \{0, 1\}^n$: w.p. $1 - \eta'(n)$ over $\rho \leftarrow \tilde{R}^c$, $\text{Bias}_F[C(F^\rho(q_1) \dots F^\rho(q_{n^c}))] \geq 1 - \eta'(n)$, where the bias is only over the choice of F (ρ is fixed).*

(3) *The probability over $\rho \in \tilde{R}^c$ that there is $i \in \{0, 1\}^n$ such that $\rho(i)$ has less than $\log^2 n$ $*$'s is negligible. For every $\rho \in \tilde{R}^c$, the probability over $x \in \{0, 1\}^n$ that there are fewer than $2^{n/2}$ y 's such that $\rho(x) = \rho(y)$ is negligible.*

Then, similarly as we did in Section 3, we can prove the following two properties:

- i. For every PPT M , with high probability over F and $\rho \in \tilde{R}^c$, M does not invert F^ρ , i.e.:

$$\Pr_{F, U_n} [F^\rho(M(F^\rho(U_n))) = F^\rho(U_n)] \leq \epsilon(n).$$

- ii. $E_{F, U_l, \rho \in \tilde{R}^c} [\Delta(G^{F^\rho}(U_l), \text{RoundG}(U_l, \rho))] \leq o(1)$.

Property (i) can be shown as in Lemma 3.4 using Lemma B.3-(3). Property (ii) follows from Lemma B.3-(2). The rest of the proof is similar to the proof on Page 10: A counting argument shows that A' is a good distinguisher when the stretch of the PRG is too big (i.e. $s \geq l \cdot \eta(n)$). Here we use Property (ii). Then one shows that A' is efficient (i.e. a PPT) under the assumption that $P = NP$. This is because if $P = NP$ then guessing x, ρ and computing RoundG can be done in polynomial time. To compute RoundG first note that we required (in the statement of the theorem) that the circuit C_x and the queries $q_{x,1}, \dots, q_{x, \text{poly}(n)}$ are generated from x in polynomial time. Then we use the fact that approximating the acceptance probability of a circuit is in PH , and that $PH = P$ if $P = NP$. So by definition of mildly black-box PRG construction there is a PPT M that inverts F^ρ . But this contradicts Property (i).

We conclude this section giving some intuition of how to obtain the distribution on restrictions in Lemma B.3. Recall our distribution on restrictions in Section 3 consisted of a hash function $h : \{0, 1\}^n \rightarrow [b]$, for $b = n^{\log n}$, and a restriction $\rho_0 : [b] \rightarrow \{0, 1, *\}^m$. We obtain the distribution in Lemma B.3 by derandomizing both h and ρ_0 .

Derandomization of h : For h we use any family of hash functions which can be generated using $\text{poly log}(n)$ random bits and that has low collision probability, for example the following standard construction (which we use with $m := \log n$.)

Lemma B.4. *Let n be a prime power. For every t there is a family of hash functions $\tilde{h}_s : \{0, 1\}^n \rightarrow \{0, 1\}^{t \log n}$ with seed length $|s| = t \log n$ such that for every $x \neq y$, $\Pr_s[\tilde{h}_s(x) = \tilde{h}_s(y)] \leq 1/n^{t-1}$.*

Proof. Construction Fix a field F of size n^t . The seed s represents an element in F . We see x as a univariate polynomial $p_x : F \rightarrow F$ where p_x has degree $n/t \log n \leq n$. Then $\tilde{h}_s(x) := p_x(s)$.

Analysis: Fix $x \neq y$. Then clearly $p_x \neq p_y$. So

$$\Pr_s[p_x(s) = p_y(s)] \leq n/|F| = 1/n^{t-1}$$

where we use the well known fact that two distinct polynomial of degree $\leq n$ over F agree in at most $n/|F|$ fraction of points. \square

Derandomization of ρ_0 : The derandomization of ρ_0 is similar to what we did in Section 5.1. It again uses Nisan's *unconditional* PRG against constant depth circuits, Theorem 5.5. Our new distribution on restrictions $\tilde{\rho}_0 : [b] \rightarrow \{0, 1, *\}^m$ is obtained by plugging a random seed into Nisan's PRG and interpreting its output bits as choices for $\{0, 1, *\}$.

We must ensure that Items (2) and (3) in Lemma B.3 hold. (Item (1) is easy to check.)

Item (3): First we need to ensure that w.h.p. for every $i \in \{0, 1\}^n$ $\tilde{\rho}_0(i)$ has $\log^2 n$ *'s. Since this holds for a truly random restriction $\rho_0 : [b] \rightarrow \{0, 1, *\}^m$ (by a concentration bound), it would be enough to show that we can check with a constant depth circuit if a block has $\log^2 n$ *'s. Then the result follows from pseudorandomness of Nisan's generator N . But this seems problematic, because it is known that constant depth circuits cannot count! (See e.g. [Hås].) However, it was shown by Ajtai and Ben-Or [ABO] that constant depth circuits of size $\text{poly}(n)$ can count up to $\log^2 n$. So using their result we can guarantee that $\tilde{\rho}_0$ is such that w.h.p. for every $i \in [b]$, $\tilde{\rho}_0(i)$ has at least $\log^2 n$ *'s. The second statement in Item (3) holds by a pigeon hole principle.

Item (2): First we argue that under $\tilde{\rho}_0$ constant depth circuits still have high bias (over the choice of the random bits for the *'s). This is obtained by showing (as in Section 5.1) that the bias of a constant depth circuit C is essentially the acceptance probability of another (slightly bigger) constant depth circuit C' , then arguing by pseudorandomness of N . To obtain Item (3) we argue as in the proof of Property (II) in Section 3. By an appropriate choice of parameters for the hash function in Lemma B.4 (i.e. $t := \log n$) we can ensure that for every fixed q_1, \dots, q_r , their images through \tilde{h} will be pairwise different w.p. $1 - \epsilon(n)$. Whenever this happens we have that the circuit will have high bias because what we said above about $\tilde{\rho}_0$.