Electronic Colloquium on Computational Complexity, Report No. 302 (2026)

Tight Depth Lower Bounds for Stretch Amplification

Oliver Korten*

Institute for Advanced Study

Abstract

We investigate the *depth overhead* of stretch-amplification for pseudorandom generators in the black-box model. We show that any black box construction of a PRG $H : \{0,1\}^m \to \{0,1\}^M$ which makes $\mathrm{poly}(n)$ queries to a PRG $G : \{0,1\}^n \to \{0,1\}^N$ in d adaptive phases during each evaluation (and with a black-box reduction to the security of G) must satisfy:

$$M \leq (m + O(\log n))\Big(\frac{N}{n - O(\log n)}\Big)^d$$

This implies that the standard amplification procedure [GM84, GGM86] based on repeated composition of G is essentially optimal. In particular, amplifying $n \mapsto n+1$ to $m \mapsto 2m$ requires depth overhead $d = \Omega(\frac{n}{\log n})$, and amplifying $n \mapsto 2n$ to $m \mapsto m^2$ requires depth overhead $d = \Omega(\log n)$[1]; each of these bounds is tight up to a constant factor. As an immediate corollary we obtain tight lower bounds on the depth complexity of black-box constructions of PRFs from PRGs. Both of these problems (pertaining to amplification and PRF construction) have been studied in various prior works [V05, MV11, BJP11, MV15, BMM24, ADV26]. Despite progress on special cases in which the structure of the reductions is constrained, it was previously unknown whether completely parallel constructions ($d = 1$) existed in every stretch regime.

The kinds of amplification procedures we investigate have more recently found use outside of cryptography in the study of *range avoidance* ("Avoid") [K21]. We show an analogous result in this setting: any black-box $\mathbf{FP}^{\mathbf{NP}}$ reduction from an Avoid instance $G : \{0,1\}^n \to \{0,1\}^N$ to a second Avoid instance $H : \{0,1\}^m \to \{0,1\}^M$ which has "depth d" in the above sense must satisfy the same bound relating n, N, m, M, d. Using this impossibility result for Avoid, we construct an oracle relative to which $\mathbf{E} \not\subseteq \mathbf{NC}/\mathrm{poly}$ but $\mathbf{E}^{\mathbf{NP}} \subseteq \mathbf{P}/\mathrm{poly}$. Our oracle separation implies that in the relativized setting, the bounded arithmetic theory S^1_2 cannot prove the existence of functions outside $\mathbf{P}/\mathrm{poly}$, even given as an axiom the existence of functions outside $\mathbf{NC}/\mathrm{poly}$.

AI Usage: The proofs of the main results were found through an extended interaction with GPT-6 Astra. Initially the author presented Astra with two of his own proofs for restricted cases of the main theorems, and asked for these arguments to be extended to the full $d = 1$ case; Astra completed this immediately in the case of Avoid. The author then suggested how to adapt the proof to larger d, and after further discussion the general proof was obtained. The author had sketched a proof of the oracle separation assuming the amplification result for Avoid several years ago; machine assistance was used to work out some missing details in this argument. The contents of this document were written entirely by the named author, who takes full responsibility for their correctness.

*Email: korten@ias.edu. This material is based upon work supported by the National Science Foundation under Grant No. DMS-2424441.

[1] The above stated formula gives $d = \Omega(\log m)$, but a trivial separate argument handles the case $m < \frac{n}{4}$.

ISSN 1433-8092

1 Introduction

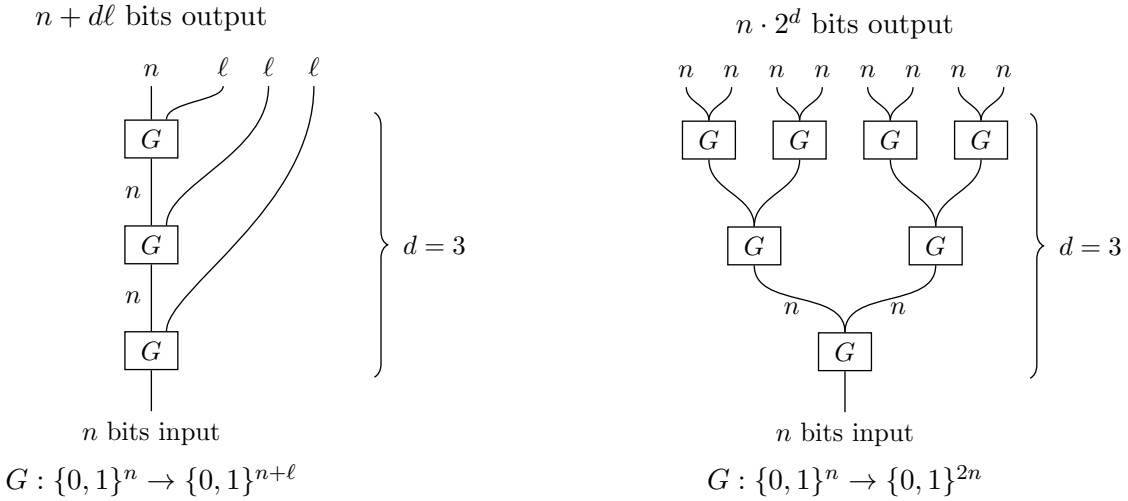
We start by introducing separately the two main subjects of study (PRG stretch amplification and AVOID stretch amplification) and giving a high-level description of our results and their context. The depth complexity of black-box PRG amplification is an established problem studied in various prior works, while the analogous question for AVOID has up until now received no formal treatment (aside from being mentioned as an open problem in a survey previously written by the author [K25a]).

1.1 Stretch Amplification for PRGs

A *pseudorandom generator* $G : \{0,1\}^n \to \{0,1\}^N$, $N > n$, is an efficiently computable length-extending function with the property that a random output of G is indistinguishable from a random N-bit string to any polynomial-time observer. First introduced in [BM82, Y82], PRGs have since come to be understood as the most fundamental primitive in cryptography, their existence being equivalent to that of one-way functions [HILL99].

An important parameter of a pseudorandom generator is its *stretch*, which refers to the quantitative relationship between its input length n and output length N; we will say that a generator with these parameters has "stretch $n \mapsto N$." In many important cases, a generator with very small stretch $n \mapsto n+1$ is considerably easier to build from a given hardness assumption (see e.g. [GL89] for a famous early example), whereas a generator with larger stretch $n \mapsto 2n$ or $n \mapsto n^2$ is necessary for key applications (e.g. pseudorandom one-time pads). This motivates the question of whether a PRG with small stretch can be generically *amplified* into a PRG with large stretch. This question was answered affirmatively in early work by Goldreich and Micali [GM84]: for any constant k, we can take a PRG with the minimal possible stretch $n \mapsto n+1$, and amplify it to a generator with $n \mapsto n^k$ stretch while incurring only a polynomial loss in security.

A well-known drawback of the amplification procedure in [GM84] is that it incurs a significant overhead in the *circuit depth* of the generator. In particular, starting with a generator G with stretch $n \mapsto n+\ell$, it produces a generator H with stretch $n \mapsto n+m$, where on each input the generator H calls the original generator G exactly $d = \frac{m}{\ell}$ many times in series. If we implement H by a Boolean circuit in the obvious way, the depth of H will increase relative to the circuit depth of G by a factor d; we hence refer to d as the *depth overhead* or *depth complexity* of the construction. If our starting generator G has slightly larger stretch $n \mapsto 2n$, a more depth-efficient construction was provided in [GGM86], which allows us to achieve stretch $n \mapsto m$ in depth $d = \log \frac{m}{n}$. These standard constructions are displayed below:



Aside from these composition-based constructions, there is one more construction of relevance due to Goldreich and Levin [GL89], which allows us to amplify stretch $n \mapsto n+1$ to $m \mapsto m + \log m$ for some $m = \Theta(n)$ while using only depth $d = 1$ (i.e., a completely parallel construction). In all regimes, the most depth-efficient amplification constructions known follow from composing some subset of these three reductions (first, [GL89], then [GM84], then [GGM86]). Perhaps the most important "critical regimes" are:

Construction	Source stretch	Target stretch	Depth	Target seed length
[GL89, GM84]	$n \mapsto n+1$	$m \mapsto 2m$	$d = O(n/\log n)$	$m = \Theta(n)$
[GGM86]	$n \mapsto 2n$	$m \mapsto m^2$	$d = O(\log n)$	$m = n$

The question of improving the depth overhead of PRG amplification has been studied in a few works [V05, L06, MV11, BJP11, MV15]. The primary motivation for finding better constructions arises from the study of low-depth cryptography, particularly the study of cryptography in $\mathbf{NC^0}$. In seminal work of [AIK06b], it was shown that under standard cryptographic assumptions, there exist PRGs $G : \{0,1\}^n \to \{0,1\}^{n+1}$ which are computable in $\mathbf{NC^0}$: each output bit is a function of only $O(1)$ input bits (in particular, only 4). Unfortunately, such PRGs have found limited use due to their minimal stretch. On the other hand, a wide variety of applications [AIK08, IKOS08, DV21, JLS21, JLS22] have been found for $\mathbf{NC^0}$ PRGs which have larger stretch $n \mapsto 2n$ or $n \mapsto n^{1+\epsilon}$. Perhaps most notably among these, $\mathbf{NC^0}$ PRGs with polynomial stretch $n \mapsto n^{1+\epsilon}$ are a key element in the celebrated constructions of *indistinguishability obfuscation* by Jain, Lin and Sahai [JLS21, JLS22]. It is thus natural to ask whether we can amplify a low-stretch $\mathbf{NC^0}$ generator into a high-stretch $\mathbf{NC^0}$ generator:

Question 1: Can we amplify $n \mapsto n+1$ stretch to $m \mapsto 2m$ (for some $m = \text{poly}(n)$), or $n \mapsto 2n$ to $m \mapsto m^2$, with constant depth overhead?

The classical amplification procedures highlighted in the above figures/table are clearly insufficient to answer this question. We note that Question 1 is of interest in the case of $\mathbf{NC^1}$ PRGs as well; any amplification construction with $d = \omega(1)$, when applied to an $\mathbf{NC^1}$ PRG G, will in general give a generator H^G outside $\mathbf{NC^1}$ [2]. Some cryptographic constructions proposed in [AIK06a, GKW17] specifically rely on $\mathbf{NC^1}$ PRGs with linear/polynomial stretch.

A positive answer to this question would take the form of a direct construction. To give a negative answer, we can only hope to prove the impossibility of such a construction in the *black-box model* (unless we aim to also resolve $\mathbf{P}$ vs. $\mathbf{NP}$ along the way). Such a negative result was first pursued in a pair of concurrent works [MV11, BJP11], which asked whether nontrivial amplification was possible with $d = 1$, i.e. *with no depth overhead whatsoever*. To make this precise, they viewed a $d = 1$ amplification procedure $H^G : \{0,1\}^m \to \{0,1\}^M$ of a high stretch generator $m \mapsto M$ from a low stretch generator $G : \{0,1\}^n \to \{0,1\}^N$ as a bounded-query oracle procedure, where on each input x, $H^G(x)$ is computed by making $\text{poly}(n)$ queries to G *in parallel*: the set of queries $w_1, \ldots, w_{\text{poly}(n)}$ to G is chosen first purely as a function of x, then the value of G is queried on all of these inputs, and finally the output $H^G(x)$ is chosen as a function of x and all the query outcomes. They also mandate that the *security reduction* is black-box: there is a query efficient algorithm which, given query access to G and any distinguisher T for H^G, is able to distinguish the output of G from random (the security reduction is not required to be parallel in any sense).

[2] We are referring here to whether the direct implementation of H^G obtained by plugging in G is an $\mathbf{NC^1}$ circuit; it is another matter whether the function computed by H^G can be implemented in lower depth by a different circuit.

Miles and Viola [MV11] were able to establish the impossibility of such a reduction under substantial but well-motivated restrictions on how $H^G(x)$ is computed as a function of $x, G(w_1), \ldots, G(w_{\text{poly}(n)})$ (they call this "postprocessing"). They first show impossibility in the case that $H^G(x)$ is obtained from projecting the concatenated string $x, G(w_1), \ldots, G(w_{\text{poly}(n)})$ onto some subset of coordinates. They then extend this result to handle post-processing by $\mathbf{AC^0}$ circuits of subexponential size.[3] This constraint on the postprocessing is relevant for the application to $\mathbf{NC^0}$ amplification; any black box construction H^G which has the property that H^G is $\mathbf{NC^0}$ whenever G is $\mathbf{NC^0}$ must have all its intermediate computations between calls to G lying in $\mathbf{NC^0}$.[4] Nonetheless the result falls short of answering Question 1 in two respects: first even for $d = 1$ the postprocessing restriction rules out applicability to amplification of $\mathbf{NC^1}$ PRGs, and second (more importantly) it says nothing about the case $d > 1$. In the case of [BJP11], impossibility for $d = 1$ reductions is shown in a setting with unrestricted postprocessing, but where severe limitations are placed on the number of queries to G and the structural dependence of these queries on the seed of H^G.

In this work we completely resolve the question studied in [MV11, BJP11] for parallel constructions, handling arbitrary postprocessing, and then extend this result to the full range of source/target stretches and possible depth parameters d. For $d > 1$, we can define the depth complexity of a black-box reduction by the number of *adaptive phases* in which G is queried in the evaluation of H^G. A depth-d construction makes $\text{poly}(n)$ queries to G, spread across d rounds of adaptivity; the queries in round i are chosen as a function of the outcomes in previous rounds, and are made simultaneously. This precise model was first suggested in [J11, Section 4.7, p. 168]. Our lower bound on depth is tight up to constant factors across the entire standard parameter regime, and hence resolves the depth complexity of black-box stretch amplification for PRGs.

As an immediate corollary, we are also able to derive essentially tight lower bounds on the depth complexity of constructing a *pseudorandom function (PRF)* from a PRG. The classical construction of [GGM86], combined with a hashing trick due to Levin [L87, Section 5.4], has remained the state-of-the-art here. A wide range of works have been dedicated to finding more depth-efficient constructions of PRFs [NR99, NR04, NRR02, LW09, BPR12, AR16, YS16, DJK25, DJK26a, DJK26b], and a few works have explicitly investigated lower bounds for black-box constructions which build PRFs from PRGs [MV15, BMM24, ADV26]. As in the case of PRG amplification, it was previously unknown whether completely parallel constructions were possible here ($d = 1$).

1.2 Range Avoidance and Hardness Amplification

In recent years, a search problem called *Range Avoidance* ("Avoid" for short) has generated significant attention due to its connections to circuit lower bounds, explicit construction problems and derandomization. We refer the reader to the survey [K25a] and the references therein for a more comprehensive introduction to the area. The problem Avoid is defined as follows: given a circuit $G : \{0,1\}^n \to \{0,1\}^N$, $N > n$, find $z \in \{0,1\}^N \setminus \text{range}(G)$.

Perhaps the most important property of the search problem Avoid is the following observation in [K21]: the problem of constructing the *truth table* of a Boolean function $f : \{0,1\}^n \to \{0,1\}$ which requires $2^{\Omega(n)}$ circuit complexity can be reduced in $2^{O(n)}$ time to Avoid (this is time polynomial in the description size of f). This follows by considering the *truth table generator* $G : \{0,1\}^{O(s \log s)} \to \{0,1\}^{2^n}$ which maps descriptions of circuits of size s to the truth tables they compute, and has the following implication:

[3] These two results in [MV11] are actually incomparable; the result for more powerful postprocessing requires imposing a limitation on the "uniformity" of the security reduction. We discuss the distinction between uniform and nonuniform security reductions in Section 1.3.

[4] This requires the same qualification as in footnote 2.

If Avoid *is solvable in polynomial time, then* $\mathbf{E}$ *requires* $2^{\Omega(n)}$ *circuit size. If* Avoid *is solvable in polynomial time with an* $\mathbf{NP}$ *oracle, then* $\mathbf{E}^{\mathbf{NP}}$ *requires* $2^{\Omega(n)}$ *circuit size.*

Whether or not $\mathbf{E}, \mathbf{E}^{\mathbf{NP}}$ have polynomial size circuits is among the most important open problems in complexity theory. Pioneering work of Nisan, Wigderson and Impagliazzo [NW94, IW97] showed that this question is intimately connected to *derandomization*: if $\mathbf{E}$ requires $2^{\Omega(n)}$ size circuits then $\mathbf{BPP} = \mathbf{P}$, and if $\mathbf{E}^{\mathbf{NP}}$ requires $2^{\Omega(n)}$ size circuits then $\mathbf{BPP} \subseteq \mathbf{P}^{\mathbf{NP}}$. In the case of $\mathbf{NP}$-oracle algorithms the connection between circuit lower bounds and Avoid algorithms becomes an exact equivalence: there is a constant $\epsilon > 0$ and a language in $\mathbf{E}^{\mathbf{NP}}$ which requires circuits of size $2^{\epsilon n}$ for sufficiently large n *if and only if* Avoid is solvable in polynomial time with an $\mathbf{NP}$ oracle. This was shown in [K21] to follow from a result of Jeřábek [J04] in Bounded Arithmetic.

Ren, Santhanam and Wang [RSW22] extended the connection between range avoidance and circuit lower bounds to the case of circuit classes smaller than $\mathbf{P}/\text{poly}$. Generally speaking, if we want to prove a lower bound against a circuit class $\mathcal{C}$, it suffices to solve "$\mathcal{C}$-Avoid," which denotes the special case of Avoid in which the input $G : \{0,1\}^n \to \{0,1\}^N$ is a circuit from the class $\mathcal{C}$. For example, if there is a polynomial time algorithm for $\mathbf{NC^1}$-Avoid, meaning instances $G : \{0,1\}^n \to \{0,1\}^{n^2}$ in which each output bit of G is specified by a polynomial size formula, then $\mathbf{E} \not\subseteq \mathbf{NC}/\text{poly}$. Overall we end up with the following precise picture:

To prove that $\mathbf{E}$ *(resp.* $\mathbf{E}^{\mathbf{NP}}$*) does not have circuits of size* $s(n)$ *and depth* $d(n)$*, it suffices to give a polynomial time (resp. with an* $\mathbf{NP}$ *oracle) algorithm which solves* Avoid *on instances* $G : \{0,1\}^{O(s(n)\log s(n))} \to \{0,1\}^{2^n}$ *in which* G *has circuit depth* $\approx d(n)$*.*

This connection holds regardless of whether we measure the depth for bounded or unbounded fan-in circuits (provided we use the same measure for the circuits in the lower bound and in the Avoid instances). In light of this, an appealing recipe (roughly suggested in [RSW22]) for proving new circuit lower bounds might go as follows:

1. Design an Avoid algorithm $\mathcal{A}$ that solves instances of a certain depth/stretch, with parameters that are apparently insufficient to directly yield a circuit lower bound.

2. Design a depth-efficient *stretch amplification reduction* which reduces Avoid instances sufficient for new circuit lower bounds to Avoid instances solvable by the algorithm $\mathcal{A}$.

3. Combine (1) and (2) to obtain a new circuit lower bound.

There are two particularly appealing candidates for instantiating this recipe. The first candidate is related to $\mathbf{AC^0}$-Avoid. On the one hand, [CHLR23] give a polynomial time $\mathbf{NP}$ oracle algorithm for $\mathbf{AC^0}$-Avoid for instances with quasipolynomial stretch $n \mapsto 2^{(\log n)^{c_d}}$ where c_d depends on the depth d of the $\mathbf{AC^0}$ circuits in question.[5] On the other hand, [RSW22] show that if there is a polynomial time $\mathbf{NP}$-oracle algorithm for stretch $n \mapsto 2^{(\log n)^c}$ for arbitrary depth $\mathbf{AC^0}$-Avoid where c is a universal constant, then $\mathbf{E}^{\mathbf{NP}} \not\subseteq \mathbf{NC^1}$. The second important candidate arises in the context of $\mathbf{NC^0}$-Avoid. On the one hand, a sequence of works [GLW22, GGNS23, KPI25, GLY26] have developed polynomial time algorithms for $\mathbf{NC^0}$-Avoid in the case that the stretch is $n \mapsto n^c$ for some constant c depending on the locality of the circuit (state-of-the-art is $c \approx \frac{k}{2}$ for locality k [KPI25, GLY26]). On the other hand, it is shown in [RSW22] that an efficient algorithm for $\mathbf{NC^0}$-Avoid with stretch $n \mapsto n+1$ (even when the locality equals 4) yields an efficient algorithm for general $\mathbf{NC^1}$-Avoid and hence implies circuit lower bounds for $\mathbf{NC^1}$.

[5] The runtime of the algorithm is polynomial in the size of the Avoid instance, which in this case is quasipolynomial in the parameter n.

The kinds of stretch-amplification reductions we seek here are of a very similar form to those sought for PRGs: given some $G : \{0,1\}^n \to \{0,1\}^N$, we want to construct $H : \{0,1\}^m \to \{0,1\}^M$ with larger stretch and comparable depth, such that given any AVOID solution for H we can find an AVOID solution for G (possibly using an **NP** oracle). Moreover, the known constructions for AVOID stretch amplification [K21] are exactly the same iterated-composition constructions used in the PRG case (recall the figures above), and achieve the same parameters.[6]

In this work we show that the composition-based amplification reductions for AVOID are essentially optimal in the black box model: to obtain the desired stretch, we cannot have a significantly more depth-efficient construction than what is known in the black-box setting. It is therefore not possible to realize the above lower bound recipes, unless the amplification step (step 2) involves *non-black-box* techniques. We then utilize this impossibility result for AVOID to show that there are strong limitations on our ability to *amplify* lower bounds against depth-restricted circuit classes to lower bounds against general circuits by black-box methods: we give an oracle relative to which **E** does not have circuits of depth $o(\sqrt{n})$ and size $2^{\sqrt{n}}$, but nonetheless every language in the larger class $\mathbf{E}^{\mathbf{NP}}$ has circuits of fixed polynomial size (with no depth restriction).

We emphasize that this result is much stronger than merely giving an oracle relative to which **P**/poly and **NC**/poly are separated (which has been known since [M92]). In particular, say that we aim to prove the "hardness amplification" statement $\mathbf{E}^{\mathbf{NP}} \not\subseteq \mathbf{NC}/\text{poly} \to \mathbf{E}^{\mathbf{NP}} \not\subseteq \mathbf{P}/\text{poly}$. The simplest way to prove this would be to show that $\mathbf{P}/\text{poly} = \mathbf{NC}/\text{poly}$, but this is widely believed to be false. Instead, one might hope to prove the amplification statement by taking some purported language L in $\mathbf{E}^{\mathbf{NP}}$ which is hard for depth-restricted circuits, and constructing from it some other language $L' \in \mathbf{E}^{\mathbf{NP}}$ which is hard for general circuits; L' might perform a computation on inputs of length n by querying L on inputs of length $O(n)$ and performing $2^{O(n)}$ time **NP**-oracle postprocessing. In fact, this is exactly the kind of hardness amplification procedure presented in [K21]:

Theorem ([K21])**.** $\mathbf{E}^{\mathbf{NP}}$ *contains a language of circuit complexity* $2^{\Omega(n)}$ *for sufficiently large* n *if and only if it contains a language of circuit complexity* $\Omega(2^n/n)$ *for sufficiently large* n*. This holds relative to every oracle.*

In this case, the amplification result is not obtained by collapsing the non-uniform class $\mathbf{Size}[o(2^n/n)]$ to $\mathbf{Size}[2^{o(n)}]$, and in fact such a collapse is unconditionally impossible by a counting argument. Instead the proof proceeds by a reduction which starts with some mildly hard function in $\mathbf{E}^{\mathbf{NP}}$ and uses some additional $2^{O(n)}$-time **NP** oracle computations to transform it into a different function of higher circuit complexity. Our oracle separation shows that when it comes to amplifying depth lower bounds to size lower bounds, no such reduction is possible in the relativizing setting.

Finally, we will apply this oracle separation to obtain an independence result in relativized Bounded Arithmetic. In an influential sequence of works by Jeřábek [J04, J05, J07], a rich subtheory of bounded arithmetic (now known as APC_1) was formed around the logical strength of the *dual weak pigeonhole principle* (the principle stating every AVOID instance has a solution). This work predated and heavily influenced the later study of AVOID as a search problem [K21]. The theory APC_1 is formed by augmenting the standard "polynomial time reasoning" theory PV_1 with the dual weak pigeonhole principle for polynomial time functions. Jeřábek demonstrates that APC_1 is capable of powerful probabilistic reasoning, and in particular can consistently assign approximate probabilities to every polynomial-time recognizable event. It is also shown in [J04, Corollary 3.6] that if we expand our "base theory" from PV_1 to S^1_2, the dual weak pigeonhole principle is equivalent to an axiom asserting the existence of functions with circuit complexity $2^{\Omega(n)}$. Our final main result

[6] There is one exception to this correspondence, namely the Goldreich-Levin step [GL89] which amplifies $n \mapsto n+1$ to $m \mapsto m + \log m$, $m = O(n)$ with no depth overhead. We know of no analogous amplification procedure for AVOID.

is that, in the relativized setting, over S_2^1 a strong lower bound for depth-restricted circuits *does not imply* a superpolynomial lower bound for general circuit size. In fact we prove that it does not even imply a sufficiently large fixed polynomial lower bound.

1.3 Our Results

We now give more formal statements of our main results.

PRG Stretch Amplification and PRF Constructions: We show that the known PRG stretch amplification procedures are essentially optimal in terms of depth complexity. The following statement is only "informal" insofar as we have yet to precisely define a black-box reduction (our definition is completely standard):

Theorem 1 (See (1) for the formal statement)**.** *Every fully black-box stretch amplification reduction with the following parameters:*

1. *The reduction amplifies stretch $n \mapsto N$ to stretch $m \mapsto M$*
2. *The security reduction maps any 0.9-advantage distinguisher for the high-stretch generator to an ϵ-advantage distinguisher for the low-stretch generator.*
3. *The number of queries made in the construction and the security reduction is bounded by q, and the adaptivity depth of the construction is bounded by d.*

must satisfy the bound:

$$M \leq \big(m + O(\log(qd))\big) \left(\frac{N}{n - O(\log(q/\epsilon))}\right)^d,$$

provided $q, \epsilon^{-1} < 2^{cn}$ for a suitable universal constant $c > 0$.

In the most standard setting we have $q, \epsilon^{-1} = \mathrm{poly}(n)$ (i.e. the construction and reduction run in polynomial time and achieve polynomial advantage) in which case we obtain the following lower bounds on depth which are tight up to a constant factor:

Source stretch	Target stretch	Our lower bound	Standard Construction
$n \mapsto n+1$	$m \mapsto 2m$	$d = \Omega(n/\log n)$	$d = O(n/\log n)$ [GM84, GL89]
$n \mapsto 2n$	$m \mapsto m^2$	$d = \Omega(\log n)$	$d = O(\log n)$ [GGM86]

We note that the second row requires an additional observation on top of Theorem 1: when $q, \epsilon^{-1} \leq 2^{cn}$ for a sufficiently small constant c, we can assume that $m \geq n/4$ (the argument is easy, see Appendix B).

As an immediate corollary, we obtain an essentially tight depth lower bound for fully black-box constructions of PRFs from PRGs. Our lower bound holds even for PRFs with nonadaptive security:

Theorem 2 (See (2) for the formal statement)**.** *Suppose a fully black-box construction takes a PRG $G : \{0,1\}^n \to \{0,1\}^N$ and constructs from it a PRF $(F_x^G : \{0,1\}^\ell \to \{0,1\})_{x \in \{0,1\}^k}$, where $x \in \{0,1\}^k$ is its key. Suppose that each evaluation of F makes at most q queries to G across d adaptive rounds, and that the security reduction maps any advantage-0.9, nonadaptive t-query*

distinguisher for F with $t \leq 2^\ell$ into an advantage-ϵ distinguisher for G, with the reduction making at most q queries in total to G and the F distinguisher. Then

$$t \leq \big(k + O(\log(ntqd))\big) \left(\frac{N}{n - O(\log(ntq))}\right)^d,$$

provided $k, \ell, N \leq \mathrm{poly}(n)$, $q, \epsilon^{-1} \leq (nt)^{O(1)}$, and $\log(ntq) = o(n)$.

Known constructions are able to achieve, for any t and some $q = \mathrm{poly}(tn)$, $\epsilon \geq \frac{1}{\mathrm{poly}(tn)}$, $k = \Theta(n)$, a construction satisfying:

$$t \geq \frac{1}{n}\left(\frac{N}{n - \log(nt)}\right)^{\Omega(d)}$$

(see Proposition 1 for precise parameters). Hence we have the following comparison where $0 < \alpha < 1$, $C > 1$ are fixed constants:

Source stretch	**PRF query budget**	**Our lower bound**	**Standard Construction**
$n \mapsto 2n$	$t = n^{\omega(1)}, q, \epsilon^{-1} = \mathrm{poly}(t)$	$d = \omega(\log n)$	any $d = \omega(\log n)$
$n \mapsto n^C$	$t = 2^{n^\alpha}, q, \epsilon^{-1} = \mathrm{poly}(t)$	$d = \Omega(n^\alpha / \log n)$	$d = O(n^\alpha / \log n)$

The intended interpretation of the table is that, in order to achieve a construction with t at least as large as the given bound, and q, ϵ^{-1} no larger than $\mathrm{poly}(t)$, we need the depth at least that in the lower bound column, and we can achieve the depth in the upper bound column. The top right entry means that for any $d(n) = \omega(\log n)$, there is a construction of the given form with depth $d(n)$, achieving some $t = n^{\omega(1)}$ and $q, \epsilon^{-1} = \mathrm{poly}(t)$.

Note on uniform vs. non-uniform security reductions:

In the main body of this document, our results for PRG amplification and PRF construction are established in the more restrictive model of security reductions studied in [BJP11] rather than the more general model of [MV11]. In [BJP11] and the majority of this paper, the security reduction is a single randomized $\mathrm{poly}(n)$-query algorithm $R^{G,T}$ which gets access to G and the distinguisher T, and distinguishes G whenever T distinguishes H^G. In contrast, [MV11] allows the security reduction R to receive some $q = \mathrm{poly}(n)$ bits of advice depending on G, T; the security reduction must satisfy that for every G, T such that T distinguishes H^G, there exists some advice string $a \in \{0,1\}^q$ such that $R_a^{G,T}$ distinguishes G.[7] Our results extend to this setting as well (with no loss in parameters) after some additional work. For clarity of exposition we relegate this extension to Appendix A since it is notably more complicated without involving any significant new idea.

Amplification for Range Avoidance: We now turn to the case of Avoid, where we obtain the same quantitative bounds. Again, the theorem statement here is somewhat ill-defined as we have yet to formally introduce the notion of reduction:

Theorem 3 (See (3) for the formal statement)**.** *Any black-box $\mathbf{FP}^{\mathbf{NP}}$ stretch amplification reduction for* Avoid *with the following parameters:*

1. *The reduction amplifies stretch $n \mapsto N$ to stretch $m \mapsto M$*

[7] [MV11] phrases this equivalently in terms of oracle circuits.

2. *The solution transfer procedure is an* $\mathbf{FP}^{\mathbf{NP}}$*-query algorithm of complexity* q *which produces an* Avoid *solution for the* $n \mapsto N$ *instance when given as input any* Avoid *solution for the constructed* $m \mapsto M$ *instance.*

3. *The construction makes at most* q *queries and has* adaptivity depth d*.*

must satisfy the bound:

$$M \leq \big(m + O(\log(qd))\big) \left(\frac{N}{n - O(\log(q))}\right)^d,$$

provided $q < 2^{cn}$ *for a suitable universal constant* $c > 0$*.*

For $q = \text{poly}(n)$, we obtain the following comparison with the standard composition reductions for Avoid [K21]. Both upper bounds use $m = n$.

Source stretch	**Target stretch**	**Our lower bound**	**Standard Construction**
$n \mapsto n+1$	$m \mapsto 2m$	$d = \Omega(n/\log n)$	$d = O(n)$
$n \mapsto 2n$	$m \mapsto m^2$	$d = \Omega(\log n)$	$d = O(\log n)$

As in the case of PRGs, these bounds follow from combining the stated theorem with an easy argument that handles the case $m < \frac{n}{4}$ (see Appendix B). The reader may observe that in contrast to PRG amplification, our lower bound for Avoid does not match the known upper bound in the first row up to constant factors; this is due to the lack of an analogue of the Goldreich-Levin construction which allowed us to amplify $n \mapsto n+1$ to $m \mapsto m + \log n$ for some $m = \Theta(n)$ with no depth overhead in the PRG setting. It is easy to see that the same construction does not directly work for Avoid.

An Oracle Separation: In the last section of this paper, we apply our main results to give the following oracle separation:

Theorem 4. *There exists an oracle* O *such that* $\mathbf{E}^O \not\subseteq \mathbf{NC}^O/\text{poly}$ *and* $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\text{poly}$*. In fact, for this oracle* O *there exists a language* $L \in \mathbf{E}^O$ *such that* L *does not have* O*-oracle circuits of size* $2^{\sqrt{n}}$ *and depth* $\sqrt{n}/8$ *for all sufficiently large* n*.*

Using standard witnessing theorems, this oracle yields the following relativized separation in bounded arithmetic:

Theorem 5 (See (5) for the formal statement)**.** *There is a fixed universal constant* k *such that the following holds:* $S^1_2(\alpha) + \text{DepthLB}(\alpha) \nvdash \text{CircuitLB}(\alpha)$*, where* DepthLB *states that there exists a function which does not have* $2^{\sqrt{n}}$*-size* α*-oracle circuits of depth* $\sqrt{n}/8$*, and* $\text{CircuitLB}(\alpha)$ *asserts that there exists a function which does not have* n^k*-size* α*-oracle circuits (of unrestricted depth).*

1.4 Proof Overviews

The technical lemma underlying all of our main results (Lemma 1) concerns the combinatorial structure of what we call (q, d)*-constructions*. A (q, d)-construction is any oracle construction

$$H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M \quad \text{taking an oracle } G : \{0,1\}^n \to \{0,1\}^N$$

which has the following property: on any input $x \in \{0,1\}^m$, $H^G(x)$ is computed by making q queries to the oracle G in d adaptive phases.

It is easiest to think about our approach in the context of range avoidance. To simplify matters we'll work with **FP** reductions rather than $\mathbf{FP}^{\mathbf{NP}}$ reductions for now. In this case, what we want to show is that if d is small and the stretch $m \mapsto M$ is much larger than $n \mapsto N$, then no $\text{poly}(n)$ query algorithm can find a string outside the range of G, *even if it is given as advice some* $y \notin \text{range}(H^G)$. Fix some $q = \text{poly}(n)$ which bounds the number of queries made both in the construction H^G and in the solver (the algorithm which solves Avoid for G given a solution for Avoid on H^G). Our approach will be to show that there is some way to reveal the values of G on all but $q+1$ inputs in $\{0,1\}^n$, such that afterwards there is a *fixed string* $y \in \{0,1\}^M$ which lies outside the range of H^G for *any completion of the oracle* G obtained by assigning values to the remaining $q+1$ inputs. If we can prove this, the result follows immediately: we supply this fixed "outer solution" y, and argue that it still requires $> q$ queries to solve Avoid for G subject to what has been revealed so far. This part is easy: if at least $q+1$ inputs to G are left free, then after any q-query algorithm terminates there will be some $w \in \{0,1\}^n$ on which the value of G remains unknown; this precludes any Avoid solution for G from being determined, since $G(w)$ can still cover any element of the codomain. Our key lemma (Lemma 1) will supply exactly this. In particular, we will show that for a random partial assignment to G leaving an expected $2q$ inputs free, if $Y_x \subseteq \{0,1\}^M$ are the possible outputs $H^G(x)$ can map to according to some completion G of the partial assignment, then $\sum_x |Y_x| < 2^M$ with high probability. We thus determine that there is some $y \notin \bigcup_{x \in \{0,1\}^m} Y_x$, which is exactly the outer solution we seek.

We will describe the argument in full in the simplest case $d = 1$ (non-adaptive constructions). Set $s = 2q$. We will choose a random set $S \subseteq \{0,1\}^n$ by including each $w \in \{0,1\}^n$ independently with probability $s/2^n$, and then form a partial assignment $\rho : \{0,1\}^n \to \{0,1\}^N \cup \{*\}$ ($*$ means free/unassigned) by randomly fixing the value of $\rho(w)$ for all $w \notin S$, and leaving $\rho(w) = *$ for all $w \in S$. With overwhelming probability $|S| \geq q+1$. For each x, let $Y_x^\rho = \{H^G(x) \mid G \supseteq \rho\}$, where $G \supseteq \rho$ means that G is a completion of the partial assignment ρ i.e. it agrees with ρ at every input assigned by ρ. We aim to show

$$\mathbb{E}_\rho \sum_{x \in \{0,1\}^m} |Y_x^\rho| \leq 2^{M-1}$$

from which it follows that some ρ with $\geq q+1$ free inputs satisfies $\bigcup_x Y_x^\rho \neq \{0,1\}^M$. For each x, let $Q_x \subseteq \{0,1\}^n$ be the set of q queries made to G in the execution $H^G(x)$ (since the construction is nonadaptive, Q_x does not depend on G). Let $t_x(\rho) = |S \cap Q_x|$, $\ell = n - \log qs$. Our argument is based around the following key observation: $t_x(\rho)$ is typically very small, and $|Y_x^\rho|$ can be directly upper bounded in terms of $t_x(\rho)$. Specifically:

1. For any t, $\Pr_\rho[t_x(\rho) = t] \leq \binom{q}{t}(s/2^n)^t \leq 2^{-\ell t}$. This follows by taking a union bound over the $\binom{q}{t}$ possible values of $S \cap Q_x$; for each, the probability that every one of the t inputs in this set lies in S is $(s/2^n)^t$.

2. Whenever $t_x(\rho) = t$, $|Y_x^\rho| \leq 2^{Nt}$. This follows since each element of Y_x^ρ can be specified by listing the values of G on the t critical queries of $H^G(x)$ in order (all other query answers are predetermined by ρ).

Now, let τ be some threshold to be determined shortly. We break our analysis into cases based on whether $t_x(\rho) \geq \tau$:

$$\mathbb{E}_\rho \sum_{x\in\{0,1\}^m} |Y_x^\rho| \le \sum_{x\in\{0,1\}^m} \sum_{t=0}^{q} \Pr[t_x(\rho)=t]\cdot \mathbb{E}_\rho[|Y_x^\rho| \mid t_x(\rho)=t] \le \sum_{x\in\{0,1\}^m} \sum_{t=0}^{q} 2^{-\ell t}\min\{2^M, 2^{Nt}\}$$
$$\le \sum_{x\in\{0,1\}^m} \Big(\sum_{\substack{0\le t\le q\\ t\ge\tau}} 2^{M-\ell t} + \sum_{\substack{0\le t\le q\\ t<\tau}} 2^{t(N-\ell)} \Big) \le 2^m(q+1)\cdot\max\{2^{M-\ell\tau}, 2^{\tau(N-\ell)}\}$$

To optimize this bound we set τ so that $M-\ell\tau = \tau(N-\ell)$, which yields $\tau = \frac{M}{N}$. The above is then bounded by $(q+1)2^{m+M-\ell\frac{M}{N}}$; this is at most 2^{M-1} provided $M \ge (m+\log(q+1)+1)\frac{N}{\ell}$ and the proof is complete. Contrapositively, we have shown that any valid reduction must satisfy:

$$M \le (m+\log(q+1)+1)\frac{N}{\ell} = (m+O(\log n))\frac{N}{n-O(\log n)}$$

The crux of this argument relied on a tradeoff between the unlikely event in which x sees many critical queries, and the upper bound on the amount of "entropy" $H^G(x)$ can generate as a function of this number of critical queries (where this "entropy" quantity is conditioned on the partial assignment). For larger d, the same principle applies. We will show that, except with small probability, we can bound the number of critical queries made in round j in terms of the number made during the preceding rounds. This will in turn yield a recurrence for the overall number of critical queries which, when unrolled (and subjected to the same case analysis above using an appropriate threshold τ) gives the quantitative bound needed for Theorem 3.

We now turn to the case of PRG amplification; here we require some additional nontrivial arguments on top of our key random restriction Lemma (Lemma 1). The high-level structure of the argument will proceed as follows:

1. Generate a random partial assignment ρ as above, and then a random completion of this partial assignment to some $G : \{0,1\}^n \to \{0,1\}^N$. Instead of $s = 2q$, we will require some larger setting $s = \text{poly}(q\epsilon^{-1})$ (recall that s represents the expected number of free inputs in the random restriction).

2. Construct a distinguisher $T : \{0,1\}^M \to \{0,1\}$, $T(y) = 1 \leftrightarrow \{\exists G' \supseteq \rho, y \in \text{range}(H^{G'})\}$

3. Show with high probability over G, T we have that T distinguishes H^G, while the security reduction fails to distinguish G when given oracle access to both G and T.

Now, Lemma 1 will take care of the first condition in (3) directly: after the random restriction ρ, with high probability the vast majority of strings[8] $y \in \{0,1\}^M$ will lie outside the range of $H^{G'}$ *for all* $G' \supseteq \rho$ and hence $T(y) = 0$, while every output $y = H^G(x)$ will satisfy $T(y) = 1$. The remaining difficulty will be in showing that the security reduction fails to distinguish G when given oracle access to T; for this remaining part we will no longer need to make use of the fact that H^G has small adaptivity depth.

We proceed by considering a correlated experiment in which we sample $w \sim \{0,1\}^n, z \sim \{0,1\}^N$ uniformly along with ρ, G, and consider the "planted generator" $G' = G[w \mapsto z]$ which modifies the value of G on w to equal z; we then redefine ρ at w to obtain some ρ' with $\rho' \subseteq G'$ (if $\rho(w) = *$ originally then no modification is necessary, else we set $\rho'(w) = z$) and define T' in terms of ρ' in exactly the same way that T was defined in terms of ρ. Because G is distributed as a

[8]For Avoid we only needed one such string, but our argument above supplied $\approx \frac{1}{2}$ of all M-bit strings. A trivial adjustment yields 99% of strings.

uniformly random function, in this correlated experiment in which we sampled the entire sequence (G, G', T, T', w, z), we see that (G', T', z) has the same marginal distribution as $(G, T, G(w))$. This will allow us to show that the security reduction fails by arguing that no efficient query procedure can distinguish (G', T', z) from (G, T, z) in this correlated experiment. This last step will be carried out by showing that in some sense the query procedure can only distinguish (G', T', z) from (G, T, z) by finding the "hidden planting location" w, which is easily shown to require many queries. The high-level structure of this correlated planting argument is very similar to prior work establishing that random oracles can be used to construct PRGs [BMG07].

There is an additional issue we must address; a naive application of this argument would involve averaging the signed distinguishing advantage of our reduction over the random choice of (G, G', T, T', w, z), which has the potential to destroy the advantage in the case that it oscillates between positive and negative values. To get around this issue, we employ a generic argument due to [BG11] which lets us assume that the security reduction always achieves a nonnegative signed advantage.

Oracle Construction: We next turn to our oracle construction in Section 5. The first key ingredient is the observation in [K21] mentioned earlier in the introduction that proving a circuit lower bound for the exponential time class $\mathbf{E}$ (resp. $\mathbf{E}^{\mathbf{NP}}$) is equivalent to giving an efficient algorithm (resp. $\mathbf{NP}$-oracle algorithm) for a certain instance of Avoid often called the "truth table generator." This connection is relativizing and was shown in [RSW22] to extend to depth-restricted circuit classes, wherein the relevant truth table generator will also be computable in restricted depth. We additionally observe that the truth table generator itself has quite large *stretch* (the particular stretch will correspond to the quantitative strength of the circuit lower bound). The second main ingredient is a result in [K21] (also relativizing, and based on the aforementioned theorem of Jeřábek in Bounded Arithmetic) which shows that for $\mathbf{E}^{\mathbf{NP}}$, a lower bound against general Boolean circuits is in fact equivalent to the existence of an $\mathbf{FP}^{\mathbf{NP}}$ algorithm for *all instances* of Avoid (not just some particular truth table generator). Hence it will suffice to construct an oracle O with the following properties:

1. Avoid is solvable in polynomial time *specifically on* the O-oracle low-depth truth table generator (which has very large stretch); this will establish $\mathbf{E}^O \not\subseteq \mathbf{NC}^O/\text{poly}$.

2. There are still some O-oracle instances of Avoid (perhaps of smaller stretch) which are hard for $\mathbf{FP}^{\mathbf{NP}^O}$ algorithms; this will establish that $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\text{poly}$.

Now, the idea is that to ensure (1) and (2) simultaneously, we need to find some fixed solution to a large-stretch instance of Avoid which is generated with small O-oracle depth, such that this fixed solution *does not help* $\mathbf{FP}^{\mathbf{NP}}$ algorithms solve Avoid on smaller stretch instances. The standard definition of O-oracle circuit depth (studied previously in [W87, M92]) is precisely the "adaptivity depth" we analyzed in the previous sections. We will interpret (a part of) our oracle as defining a hard small-stretch instance of Avoid, and treat the large-stretch, low O-depth Avoid instance as a low-depth stretch amplifying reduction for Avoid, and apply our impossibility results.

This description is an oversimplification in a few respects. First, as is usually the case in such arguments, we need to deal with the effects of our choices in the construction of the oracle across various input lengths, and perform an explicit diagonalization. Second, for (2) we will actually need to construct an explicit O-oracle Avoid instance which requires exponential time for $\mathbf{NP}^O$-oracle algorithms to solve on all sufficiently large input lengths. Finally, we will not actually derive our oracle as a black-box from our impossibility results for stretch-amplifying reductions (Theorem 3), but instead use directly the random restriction lemma underlying it (Lemma 1).

2 Preliminaries

We start with some general terminology. For an integer n, we use $\mathcal{U}_n$ to denote the uniform distribution on $\{0,1\}^n$. For a distribution $\mathcal{D}$ over a finite set X, we use $x \sim \mathcal{D}$ to denote a sample from $\mathcal{D}$, and $x \sim X$ to denote a uniformly random sample from X.

We now introduce notation for partial assignments to functions $\{0,1\}^n \to \{0,1\}^N$. A partial assignment $\rho : \{0,1\}^n \to \{0,1\}^N \cup \{*\}$ assigns each input in $\{0,1\}^n$ either a value in $\{0,1\}^N$ or the symbol $*$, indicating that the input is free (unassigned). For a partial assignment ρ, define

$$\mathrm{dom}(\rho) := \{w : \rho(w) \neq *\},\ \mathrm{free}(\rho) := \{w : \rho(w) = *\},\ |\rho| := |\mathrm{dom}(\rho)|.$$

For a pair ρ, τ of partial assignments, we say that τ extends ρ, written $\tau \supseteq \rho$, if $\mathrm{dom}(\tau) \supseteq \mathrm{dom}(\rho)$ and $\tau(w) = \rho(w)$ for every $w \in \mathrm{dom}(\rho)$. A total assignment $G : \{0,1\}^n \to \{0,1\}^N$ is a special case of a partial assignment; in this case when $G \supseteq \rho$ we say G is a "completion of ρ." We say that two partial assignments ρ, τ are *compatible* if $\rho(w) = \tau(w)$ for every $w \in \mathrm{dom}(\rho) \cap \mathrm{dom}(\tau)$. For a function $G : \{0,1\}^n \to \{0,1\}^N$ and a set $X \subseteq \{0,1\}^n$, we use $G|_X$ for the partial assignment which takes the values of G inside X and is undefined outside X; observe that $G \supseteq G|_X$ for any X. For $p \in [0,1]$ and $n, N \in \mathbb{N}$ we define a distribution $\mathcal{R}_p^{n,N}$ over partial assignments $\rho : \{0,1\}^n \to \{0,1\}^N \cup \{*\}$ defined as follows: to sample $\rho \sim \mathcal{R}_p^{n,N}$, for each $w \in \{0,1\}^n$, we independently set $\rho(w) = *$ with probability p, and $\rho(w) = z \in \{0,1\}^N$ with probability $(1-p)2^{-N}$ for each $z \in \{0,1\}^N$. Equivalently, we sample $G : \{0,1\}^n \to \{0,1\}^N$ uniformly at random, then sample $S \subseteq \{0,1\}^n$ by including each $w \in \{0,1\}^n$ with probability p, and finally set $\rho = G|_{\overline{S}}$.

2.1 Amplification Procedures

We start with a definition that will be common to both the PRG and Avoid settings. A "construction" merely describes how a generator H with larger stretch can be built from a generator G with smaller stretch; the relevant parameters are the number of times it queries G and the *adaptivity depth* of these queries.

Definition 1 ((q,d)-Construction)**.** *A (q,d)-construction $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ from inner function $G : \{0,1\}^n \to \{0,1\}^N$ is a query procedure which, on every input $x \in \{0,1\}^m$ and for every oracle G, outputs a value $H^G(x) \in \{0,1\}^M$ which is determined by making at most q queries to G in at most d adaptive rounds. More formally the computation of $H^G(x)$ proceeds as follows: first a sequence of queries $w^1, \ldots, w^\ell \in \{0,1\}^n$ is determined purely as a function of the input x; the procedure then evaluates $G(w^1), \ldots, G(w^\ell) \in \{0,1\}^N$, and as a function of these returned values and the original input x, a second batch of queries is determined. This proceeds for d rounds, and the total number of queries to G across all rounds is bounded by q. At the end, as a function of every query outcome and the original input x, a value $y \in \{0,1\}^M$ is output. We call d the* adaptivity depth *of the construction. Without loss of generality, no input is queried more than once in an execution of $H^G(x)$.*

The common basis of our two main results will be a statement purely about the combinatorial structure of (q,d) constructions (Lemma 1). We now turn to the two objects of interest which utilize these constructions. We will start with PRG stretch amplification. We first introduce some general definitions relevant to the study of PRGs:

Definition 2 (Distinguishers and Advantage)**.** *A* distinguisher *on $\{0,1\}^L$ is a probability distribution D over Boolean functions $\{0,1\}^L \to \{0,1\}$. For a generator $F : \{0,1\}^\ell \to \{0,1\}^L$, the* signed

advantage *and* advantage *of D are, respectively,*

$$\mathrm{SAdv}_F(D) := \Pr[D(F(\mathcal{U}_\ell)) = 1] - \Pr[D(\mathcal{U}_L) = 1],$$
$$\mathrm{Adv}_F(D) := \left|\mathrm{SAdv}_F(D)\right|,$$

where the probabilities are taken over $\mathcal{U}_\ell$ (resp. $\mathcal{U}_L$) and the internal randomness of D simultaneously.

Definition 3 (PRG Stretch Amplification Reduction)**.** *Let $N > n$, $M > m$, and $0 < \gamma, \epsilon \leq 1$. A (q, γ, ϵ, d) amplification reduction from $n \mapsto N$ stretch to $m \mapsto M$ stretch is a pair (H, R) such that $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ is a (q, d)-construction from inner function $G : \{0,1\}^n \to \{0,1\}^N$, and $R^{G,T}$ is a randomized query procedure outputting a value in $\{0,1\}$ and querying oracles $G : \{0,1\}^n \to \{0,1\}^N$, $T : \{0,1\}^M \to \{0,1\}$, such that for every possible G, T we have:*

$$\mathrm{Adv}_{H^G}(T) \geq \gamma \quad \Longrightarrow \quad \mathrm{Adv}_G(R^{G,T}) \geq \epsilon$$

On every input and for every choice of its randomness, R makes at most q queries in total to G and T.

We emphasize that there is no adaptivity-depth restriction on R. It would be fitting to allow also the test T to be a fully general "distinguisher" in the sense of Definition 2 (a distribution over Boolean-valued functions) just as R is, rather than a fixed deterministic function. However, since our results here will be negative, we lose nothing by allowing the reduction to work only on deterministic T.

We next define black-box constructions of PRFs from PRGs. We will restrict our attention here to PRFs with *nonadaptive security* and which output a single bit. As our results are negative, this only makes them stronger.

Definition 4 (PRG-to-PRF Construction)**.** *Let $N > n$, $k, \ell, q, d \geq 1$, $1 \leq t \leq 2^\ell$, and $0 < \gamma, \epsilon \leq 1$. A* fully black-box PRG-to-PRF construction with nonadaptive security *and parameters $(q, t, \gamma, \epsilon, d)$ is defined by a pair (F, R). The function F is a (q, d) construction of the form:*

$$F^{(\cdot)} : \{0,1\}^k \times \{0,1\}^\ell \to \{0,1\}, \qquad F_x^G(u) := F^G(x, u),$$

taking an oracle $G : \{0,1\}^n \to \{0,1\}^N$, so q bounds the number of queries to G per evaluation $F_x^G(u)$.

A deterministic nonadaptive t-query distinguisher consists of a tuple $\bar{u} = (u^1, \ldots, u^t)$ of distinct inputs in $\{0,1\}^\ell$ and a Boolean test $T : \{0,1\}^t \to \{0,1\}$. Its advantage is

$$\mathrm{Adv}^{\mathrm{prf}}_{F^G}(\bar{u}, T) := \left| \Pr_{x \sim \mathcal{U}_k}[T(F_x^G(u^1) \cdots F_x^G(u^t)) = 1] - \Pr[T(\mathcal{U}_t) = 1] \right|.$$

The uniform t-bit string on the right side represents the values of a uniformly random function $f : \{0,1\}^\ell \to \{0,1\}$ on the distinct inputs $u^1, \ldots, u^t$.

The reduction $R_{\bar{u}}^{G,T}$ is a randomized query procedure making at most q queries in total to oracles $G : \{0,1\}^n \to \{0,1\}^N$, $T : \{0,1\}^t \to \{0,1\}$, on every input and for every choice of its randomness. Overall the construction must satisfy the following for all $\bar{u}, G, T$:

$$\mathrm{Adv}^{\mathrm{prf}}_{F^G}(\bar{u}, T) \geq \gamma \quad \Longrightarrow \quad \mathrm{Adv}_G(R_{\bar{u}}^{G,T}) \geq \epsilon.$$

Note that we defined a (q, d) construction for multioutput functions $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$, but the definition extends in the obvious way to single output functions $F^{(\cdot)} : \{0,1\}^k \times \{0,1\}^\ell \to \{0,1\}$ (the length of the output had no special role in the definition).

We now turn to amplification reductions for Avoid. We start with the definition of $\mathbf{FP}^{\mathbf{NP}}$ query procedures:

Definition 5 ($\mathbf{FP^{NP}}$ Query Procedures)**.** *An $\mathbf{FP^{NP}}$ query procedure Q of complexity q querying a function $G : \{0,1\}^n \to \{0,1\}^N$ is defined as follows. Q^G proceeds in q phases. At each phase, a single "$\mathbf{NP}$-query" to G of complexity q is made, which is defined as follows: some DNF ϕ of width q over the propositional variables $(\mathsf{G}(w)_i)_{w\in\{0,1\}^n, i\leq N}$ is presented, and at unit cost it is determined whether the function G satisfies ϕ under the natural interpretation, e.g. if the proposition is $\neg\mathsf{G}(w)_i$, then G will satisfy this literal if and only if $G(w)_i = 0$. Based on this determination (and all previous), we choose the next DNF to query, and so on. At the end the query procedure outputs a value in $\{0,1\}^*$.*

Definition 6 (Avoid amplification reduction)**.** *Let $N > n$ and $M > m$. A (q,d)* amplification reduction for Avoid *from $n \mapsto N$ stretch to $m \mapsto M$ stretch is a pair (H, Solve), where $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ is a (q,d)-construction from an oracle $G : \{0,1\}^n \to \{0,1\}^N$, $\mathrm{Solve}^{(\cdot)} : \{0,1\}^M \to \{0,1\}^N$ is an $\mathbf{FP^{NP}}$ query procedure of complexity q, and jointly they satisfy the following for every $G : \{0,1\}^n \to \{0,1\}^N$, $y \in \{0,1\}^M$:*

$$y \notin \mathrm{range}(H^G) \quad \Longrightarrow \quad \mathrm{Solve}^G(y) \notin \mathrm{range}(G).$$

We emphasize that no adaptivity-depth restriction is placed on the query procedure $\mathrm{Solve}^{(\cdot)}$.

2.2 Relativized Complexity Classes

We define here the relativized complexity classes needed for our main oracle separation; our definitions, including for oracle circuit depth, are standard, see [W87, M92][9].

We start with relativized circuit complexity. An "oracle circuit" $C^{(\cdot)} : \{0,1\}^n \to \{0,1\}$ is a circuit in which each gate is either one of $\{\neg, \vee, \wedge\}$ of fan-in 1,2,2 respectively, or else an "oracle gate" which can have arbitrary fan-in. The total size of C, denoted $|C|$, is the number of wires in the circuit C. The depth of C, denoted $\mathrm{depth}(C)$, is the length of the longest path of wires from the output to some input. For an oracle $O : \{0,1\}^* \to \{0,1\}^*$, $C^O : \{0,1\}^n \to \{0,1\}$ is given by substituting O into the oracle gates in the natural way. For functions $s, d : \mathbb{N} \to \mathbb{N}$, we define $\mathbf{SizeDepth}^O[s(n), d(n)]$ to be the set of languages L such that there exists a sequence of O-oracle circuits $(C_n)_{n\in\mathbb{N}}$, with $|C_n| \leq s(n)$, $\mathrm{depth}(C_n) \leq d(n)$, such that $C^O_{|x|}(x) = 1 \leftrightarrow x \in L$ for all $x \in \{0,1\}^*$. We define $\mathbf{Size}^O[s(n)] := \mathbf{SizeDepth}^O[s(n), s(n)]$. We will study the classes:

$$\mathbf{P}^O/\mathrm{poly} = \bigcup_{k\in\mathbb{N}} \mathbf{Size}^O[n^k], \qquad \mathbf{NC}^O/\mathrm{poly} = \bigcup_{k\in\mathbb{N}} \mathbf{SizeDepth}^O[n^k, \log^k(n)]$$

We next define the classes $\mathbf{E}^O, \mathbf{E}^{\mathbf{NP}^O}$. We start by defining the language $\mathbf{SAT}^O$ for an oracle O. The language consists of all descriptions of oracle circuits $C^{(\cdot)} : \{0,1\}^n \to \{0,1\}$, $n \in \mathbb{N}$, such that C^O is satisfiable (using a standard string encoding of circuits). The class $\mathbf{E}^O$ (resp. $\mathbf{E}^{\mathbf{NP}^O}$) is the set of languages decidable in $2^{O(n)}$ time with access to an oracle for O (resp. for both $\mathbf{SAT}^O$ and O).

3 The Key Lemma, and Amplification Lower Bounds for Avoid

Here we prove the key lemma underlying all of the subsequent results in this paper. We then apply this to show our main impossibility result for stretch-amplification reductions for Avoid, which will follow almost immediately. The case of stretch-amplification for PRGs will require a substantial additional argument, and so we leave it for the following section.

[9] We use a slightly more restrictive metric for depth, i.e. for a given circuit our measure will never be larger than that in [W87, M92]. This will only make our results stronger since we show a depth lower bound in our oracle construction.

Lemma 1. *Let $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ be a (q,d) construction from an oracle $G : \{0,1\}^n \to \{0,1\}^N$, $N > n$, $q, d \geq 1$. Let $1 \leq s < \frac{2^n}{q}$ be any parameter and let $p = s/2^n$. Then*

$$\Pr_{\rho \sim \mathcal{R}_p^{n,N}, y \sim \mathcal{U}_M}[\exists G \supseteq \rho, y \in \text{range}(H^G)] \leq \text{poly}(qd) \cdot 2^{m - M(\frac{n - \log qs}{N})^d}$$

In particular, the probability on the left-hand-side of the above inequality is at most $2^{-\lambda}$ provided $M \geq (m + O(\log qd) + \lambda)\left(\frac{N}{n - \log qs}\right)^d$.

Proof. For a partial assignment ρ, let $Y^\rho = \bigcup_{G \supseteq \rho} \text{range}(H^G)$. For $x \in \{0,1\}^m$, a partial assignment ρ and $G \supseteq \rho$, let $t_x^G(\rho)$ be the number of queries made to G on inputs in $\text{free}(\rho)$ during the execution $H^G(x)$; we call such queries "critical queries." For each $j \leq d$ let $t_{x,j}^G(\rho)$ be the number of critical queries made during the first j adaptivity phases (so $t_{x,d}^G(\rho) = t_x^G(\rho)$). For $x \in \{0,1\}^m$, a partial assignment ρ and some $t \geq 0$, let $Y_{x,t}^\rho = \{H^G(x) \mid G \supseteq \rho, t_x^G(\rho) = t\}$. We will show that for all $x \in \{0,1\}^m$ we have:

$$\mathbb{E}_\rho \sum_{t=0}^{q} |Y_{x,t}^\rho| \leq \text{poly}(qd) \cdot 2^{M - M(\frac{n - \log qs}{N})^d} \tag{$\dagger$}$$

Since $Y^\rho = \bigcup_{x \in \{0,1\}^m, 0 \leq t \leq q} Y_{x,t}^\rho$, this will directly yield the lemma:

$$\Pr_{\rho,y}[\exists G \text{ extending } \rho, y \in \text{range}(H^G)] = \mathbb{E}_\rho 2^{-M}|Y^\rho| \leq \mathbb{E}_\rho \sum_{x \in \{0,1\}^m} \sum_{t=0}^{q} 2^{-M}|Y_{x,t}^\rho|$$

$$\leq 2^{m-M} \max_{x \in \{0,1\}^m} \left(\mathbb{E}_\rho \sum_{t=0}^{q} |Y_{x,t}^\rho|\right) \leq \text{poly}(qd) \cdot 2^{m - M(\frac{n - \log qs}{N})^d}$$

So it remains only to prove ($\dagger$). Fix some particular $x \in \{0,1\}^m$ for the remainder of the proof. Let $\ell = n - \log qs$, $\alpha = \frac{N - \ell}{(N/\ell)^d - 1}$. We will show the following key claims:

1. For any threshold $\tau \geq 0$, $\Pr_\rho[\exists t \geq \tau, Y_{x,t}^\rho \neq \emptyset] \leq \text{poly}(qd) \cdot 2^{-\tau\alpha}$ (where the $\text{poly}(qd)$ term does not depend on the threshold τ).
2. $\mathbb{E}_\rho |Y_{x,t}^\rho| \leq 2^{(N-\ell)t}$ for all $0 \leq t \leq q$.

Assuming these, for any chosen threshold $\tau \geq 0$ we will have

$$\mathbb{E}_\rho \sum_{t=0}^{q} |Y_{x,t}^\rho| \leq (q+1)2^M \Pr_\rho[\exists t \geq \tau, Y_{x,t}^\rho \neq \emptyset] + \sum_{0 \leq t \leq q, t < \tau} \mathbb{E}_\rho |Y_{x,t}^\rho|$$
$$\leq \text{poly}(qd) \cdot 2^{M - \tau\alpha} + (q+1)2^{(N-\ell)\tau} \leq \text{poly}(qd) \cdot \max\{2^{M-\tau\alpha}, 2^{(N-\ell)\tau}\}$$

We can then choose τ such that $M - \tau\alpha = (N - \ell)\tau$; this choice yields $\tau = \frac{M}{N - \ell + \alpha}$ and hence $M - \tau\alpha = (N-\ell)\tau = M - M(\ell/N)^d$, giving the overall bound $\mathbb{E}_\rho \sum_{t=0}^{q} |Y_{x,t}^\rho| \leq \text{poly}(qd) \cdot 2^{M - M(\ell/N)^d}$ which is exactly ($\dagger$).

We now establish (1) and (2). We suppress x in the notation $t_{x,j}^G(\rho)$ since we have fixed it. To prove (1), we establish the following: for any $j \leq d$, $t_{j-1}, t_j \in \mathbb{N}$, we have:

$$\Pr_\rho[\exists G \supseteq \rho, t_{j-1}^G(\rho) = t_{j-1} \wedge t_j^G(\rho) = t_j] \leq 2^{N t_{j-1} - \ell t_j} \qquad \text{(Critical Query bound)}$$

To show this, define a "critical transcript of type (j, t_{j-1}, t_j)" to be a pair $(P, \bar{z})$ where $P \subseteq [q], |P| = t_j$ and $\bar{z} = (z^1, \ldots, z^{t_{j-1}}), z^i \in \{0,1\}^N$. We say that some ρ is consistent with this critical transcript if for some $G \supseteq \rho$, during the execution of $H^G(x)$ through the first j adaptivity rounds, the query indices $P \subseteq [q]$ represent exactly the queries (in chronological order) on which $H^G(x)$ queried G on some w with $\rho(w) = *$, and through the first $j-1$ rounds if w^i was the i^{th} such critical query then $G(w^i) = z^i$. Now, for any such critical transcript $(P, \bar{z})$ of type (j, t_{j-1}, t_j), the probability that ρ is consistent with it is at most p^{t_j}. To see this, note that for any $G \supseteq \rho$ realizing this critical transcript, the entire transcript of the simulation of $H^G(x)$ up until just after the queries for round j are chosen (but not yet queried) depends only on $\rho, P, \bar{z}$; moreover, before each of the t_j critical queries $w^1, \ldots, w^{t_j}$ is made, we can imagine that $\rho(w^i)$ is chosen independently with probability p to equal $*$ (recall from Section 2 that we can assume no repeated queries are ever made). We then take a union bound over all critical transcripts of type (j, t_{j-1}, t_j), of which there are at most $q^{t_j} 2^{Nt_{j-1}}$, and get an overall bound of $q^{t_j} 2^{Nt_{j-1}} p^{t_j} = 2^{Nt_{j-1} - \ell t_j}$.

Now, for some G, ρ and threshold $\tau \geq 0$, we claim that if $t_d^G(\rho) \geq \tau$, then we must have that $\ell t_j^G(\rho) - N t_{j-1}^G(\rho) \geq \alpha\tau$ for some j. To see this, note that if $\ell t_j^G(\rho) - N t_{j-1}^G(\rho) < \alpha\tau$ for all $j \leq d$ then we can strictly upper bound each term in the sequence $(t_1^G(\rho), \ldots, t_d^G(\rho))$ by its corresponding term in the sequence $(t_1, \ldots, t_d)$ where

$$t_0 := 0,\ t_j := \frac{N}{\ell} t_{j-1} + \frac{\alpha\tau}{\ell}, \qquad \text{which yields } t_d = \frac{\alpha\tau}{\ell} \sum_{i=0}^{d-1} \left(\frac{N}{\ell}\right)^i = \tau$$

It then follows that:

$$\Pr_\rho[\exists G \supseteq \rho, t_d^G(\rho) \geq \tau] \leq \sum_{j=1}^{d} \sum_{\substack{0 \leq t_{j-1}, t_j \leq q \\ \ell t_j - N t_{j-1} \geq \alpha\tau}} \Pr[\exists G \supseteq \rho, t_{j-1}^G(\rho) = t_{j-1} \wedge t_j^G(\rho) = t_j]$$

$$\leq \sum_{j=1}^{d} \sum_{\substack{0 \leq t_{j-1}, t_j \leq q \\ \ell t_j - N t_{j-1} \geq \alpha\tau}} 2^{N t_{j-1} - \ell t_j} \leq d(q+1)^2 2^{-\alpha\tau}$$

where we used the Critical Query bound in the second step. Hence we have established (1). For (2), we use a similar argument to that used in the Critical Query bound. Define a terminal critical transcript of type t to be a pair $(P, \bar{z})$ with $P \subseteq [q], |P| = t$ and $\bar{z} \in (\{0,1\}^N)^t$; we say that ρ is consistent with this transcript if, for some $G \supseteq \rho$, P represents the query steps in chronological order on which $H^G(x)$ queried some w with $\rho(w) = *$, and $\bar{z}$ represents the values that G took on each of the critical queries. As before, we can bound the probability that ρ is consistent with a given transcript $(P, \bar{z})$ by p^t, and the number of such transcripts is $\leq q^t 2^{Nt}$. Finally observe that $\rho, P, \bar{z}$ completely determine the value $H^G(x)$ (for any $G \supseteq \rho$ realizing this transcript); in other words, for a fixed ρ, each transcript contributes at most one element to $Y_{x,t}^\rho$; on the other hand every element of $Y_{x,t}^\rho$ arises from some critical transcript. Hence overall:

$$\mathbb{E}_\rho |Y_{x,t}^\rho| \leq \sum_{P, \bar{z}} \Pr_\rho[\rho \text{ is consistent with } (P, \bar{z})] \leq p^t q^t 2^{Nt} = 2^{(N-\ell)t}$$

□

We now apply this to stretch amplification reductions for Avoid. We need the following result, an elementary proof of which is given in [K25a]; the core of the argument first appeared in [W83] but is stated in a different terminology.

Lemma 2 ([W83, K25a]). *Any* $\mathbf{FP}^{\mathbf{NP}}$ *query algorithm solving* Avoid *on all instances* $G : \{0,1\}^n \to \{0,1\}^N$ *must have complexity* $\geq \sqrt{2^n}$.

Theorem 3. *Let* (H, Solve) *be a* (q, d) *stretch amplification reduction for Avoid from* $n \mapsto N$ *stretch to* $m \mapsto M$ *stretch, where* $q \geq 2$, $d \geq 1$, *and* $8q^3 < 2^n$. *Then*

$$M \leq \big(m + O(\log(qd))\big) \left(\frac{N}{n - \log(8q^3)}\right)^d.$$

Proof. Set $s = 8q^2$ and $p = s/2^n$. Say that $M \geq (m + C\log(qd))\left(\frac{N}{n-\log qs}\right)^d$ for a suitably large constant C. Then by Lemma 1, with probability $\geq \frac{1}{2}$ over $\rho \sim \mathcal{R}_p^{n,N}$ and $y \sim \{0,1\}^M$, we have that $y \notin \mathrm{range}(H^G)$ for all $G \supseteq \rho$. The probability that $|\mathrm{free}(\rho)| \leq s/2$ is at most $e^{-s/8}$, hence there must exist ρ, y such that $y \notin \mathrm{range}(H^G)$ for all $G \supseteq \rho$ and $|\mathrm{free}(\rho)| > s/2 = 4q^2$.

Now, the procedure $\mathrm{Solve}^G(y)$ must successfully find an Avoid solution for G whenever y is a solution for H^G; in particular, it solves Avoid on all instances $G \supseteq \rho$. Let r be such that $q^2 < 2^r \leq |\mathrm{free}(\rho)|$. Then, identifying $\{0,1\}^r$ with a subset of $\mathrm{free}(\rho)$, we obtain a $\mathbf{FP}^{\mathbf{NP}}$ query procedure with complexity q which solves Avoid on all instances $F : \{0,1\}^r \to \{0,1\}^N$, $q^2 < 2^r$. This contradicts Lemma 2. □

We conclude this section with a brief observation used later in our oracle construction (Section 5): the distribution of the values that ρ took whenever $\rho(w) \neq *$ had no relevance to the proof of Lemma 1. Hence:

Observation 1. *Let* $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$, q, d, n, N, s, p *be exactly as in Lemma 1. Let* $\mathcal{G}$ *be any distribution on functions* $G : \{0,1\}^n \to \{0,1\}^N$ *and let* $\mathcal{R}_p(\mathcal{G})$ *be the distribution on partial assignments* $\rho : \{0,1\}^n \to \{0,1\}^N \cup \{*\}$ *given by sampling* $G \sim \mathcal{G}$, *then sampling* $S \subseteq \{0,1\}^n$ *by including each input independently with probability* p, *and outputting* $\rho = G|_{\overline{S}}$. *Then the exact same bound in Lemma 1 holds with the distribution* $\mathcal{R}_p(\mathcal{G})$ *in place of* $\mathcal{R}_p^{n,N}$.

4 PRG Stretch Amplification and PRG to PRF Constructions

In this section we establish our main impossibility results for PRG stretch amplification and constructions of PRFs from PRGs. The results pertaining to PRFs will follow immediately from our stretch amplification results by standard reductions.

We start with the following lemma which is a special case of a very general transformation presented in [BG11, Proposition 1.1]; it allows us to focus on reductions which always achieve nonnegative signed advantage. This will allow us to argue in our main result ruling out PRG amplification reductions that the distinguisher $R^{G,T}$ in the security reduction has small advantage by an averaging argument (this kind of averaging is unavailable when the signed advantages can be both positive and negative). In Appendix A where we prove a strengthening of Theorem 1 involving reductions with advice, we will side-step this argument and reason directly about the absolute advantage; this has a mild cost in terms of the complexity of the argument and so we prefer using the reduction to nonnegative advantage here.

Lemma 3 ([BG11]). *Suppose* (H, R) *is an amplification reduction such that, for every* G, T *it holds that* $\mathrm{Adv}_{H^G}(T) \geq \gamma \Longrightarrow \mathrm{Adv}_G(R^{G,T}) \geq \epsilon$. *Then there is an amplification reduction* $(H, \tilde{R})$ *such that* $\mathrm{SAdv}_G(\tilde{R}^{G,T}) \geq 0$ *for all* G, T, *and* $\mathrm{Adv}_{H^G}(T) \geq \gamma \Longrightarrow \mathrm{SAdv}_G(\tilde{R}^{G,T}) \geq \epsilon^2$. *If* R *makes at most* q *queries in total (to* G *and* T*), then* $\tilde{R}$ *makes at most* $3q + 1$ *queries in total.*

Proof. On input $z \in \{0,1\}^N$, our randomized distinguisher $\tilde{R}$ will sample independent $w \sim \mathcal{U}_n$, $z' \sim \mathcal{U}_N$, and compute $a = R^{G,T}(G(w))$, $b = R^{G,T}(z')$, $c = R^{G,T}(z)$ using independent randomness for each invocation. $\tilde{R}$ will then output c if $(a,b) = (1,0)$, $\neg c$ if $(a,b) = (0,1)$, and 0 otherwise. Fix G, T, and let $\alpha = \Pr[R^{G,T}(G(\mathcal{U}_n)) = 1]$, $\beta = \Pr[R^{G,T}(\mathcal{U}_N) = 1]$, and $\Delta = \alpha - \beta$. Since a, b are independent of the challenge and its evaluation,

$$\mathrm{SAdv}_G(\tilde{R}^{G,T}) = \big(\alpha(1-\beta) - (1-\alpha)\beta\big)\Delta = \Delta^2.$$

Thus the signed advantage is always nonnegative, and $|\Delta| \geq \epsilon$ implies that the signed advantage is at least ϵ^2. The construction uses three invocations of R and one additional query to G, so at most $3q + 1$ queries in total. □

It remains only to rule out an amplification reduction with nonnegative signed advantage. As mentioned in the introduction, the main correlated experiment we study in this proof is an established method used previously to prove the existence of PRGs (among other primitives) in the random oracle model [BMG07].

Theorem 1. *There is a universal constant $c > 0$ such that the following holds. Let (H, R) be a $(q, 0.9, \epsilon, d)$ PRG amplification reduction from $n \mapsto N$ stretch to $m \mapsto M$ stretch, where $q \geq 2$, $d \geq 1$, $0 < \epsilon \leq 1$, and $q, \epsilon^{-1} \leq 2^{cn}$. Then*

$$M \leq \big(m + O(\log(qd))\big)\left(\frac{N}{n - O(\log(q/\epsilon))}\right)^d.$$

Proof. By assumption, $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ is a (q,d) construction from an oracle $G : \{0,1\}^n \to \{0,1\}^N$, and $R^{G,T}$ makes $\leq q$ queries to T, G. Using Lemma 3, we may assume up to negligible loss in parameters that $R^{G,T}$ has nonnegative signed advantage for all T, G, and that this advantage is $\geq \epsilon$ for every T, G such that $\mathrm{Adv}_{H^G}(T) \geq 0.9$. Let $s = \mathrm{poly}(q\epsilon^{-1})$ be a parameter to be set later and let $p = s/2^n$. Our proof proceeds as follows:

1. Sample $\rho \sim \mathcal{R}_p^{n,N}$, then sample a uniformly random completion $G \supseteq \rho$.
2. Set $T = T_\rho$, where $T_\rho(y) = 1$ if and only if there exists $G' \supseteq \rho$ such that $y \in \mathrm{range}(H^{G'})$.
3. Evaluate the correctness of our stretch reduction for this particular choice of G, T.

Assuming towards a contradiction that $M \geq (m + C\log qd)(\frac{N}{n - \log qs})^d$ for a sufficiently large universal constant C, we will show that when G, T are sampled in this way we have:

$$\text{(a)} \quad \Pr_{G,T}[\mathrm{Adv}_{H^G}(T) \geq 0.9] > \frac{1}{2} \qquad \text{(b)} \quad \Pr_{G,T}[\mathrm{SAdv}_G(R^{G,T}) < \epsilon] > \frac{1}{2}$$

Taking a union bound we conclude that for some G, T both conditions hold, contradicting the correctness of the stretch reduction. The first condition will follow directly from Lemma 1. Observe that for our choice of G, T, we have that $T(H^G(x)) = 1$ for all $x \in \{0,1\}^m$. On the other hand,

$$\begin{aligned}\Pr_{G,T}[\Pr_y[T(y) = 1] \geq 0.1] &\leq \Pr_\rho[\Pr_y[\exists G' \supseteq \rho, y \in \mathrm{range}(H^{G'})] \geq 0.1] \\ &\leq 10 \Pr_{\rho,y}[\exists G' \supseteq \rho, y \in \mathrm{range}(H^{G'})] \leq \frac{1}{3}\end{aligned}$$

where we use the bound from Lemma 1 and our choice of parameters in the inequality $\Pr_{\rho,y}[\exists G' \supseteq \rho, y \in \mathrm{range}(H^{G'})] \leq \frac{1}{30}$ (and a Markov argument in the preceding step).

Now, to prove (b) we show that $\mathbb{E}_{G,T}[\mathrm{SAdv}_G(R^{G,T})] < \frac{\epsilon}{2}$. Since we assumed the signed advantage is always nonnegative this will suffice by another Markov argument. We need to bound

$$\mathbb{E}_{G,T}\big(\mathbb{E}_{R,w}R^{G,T}(G(w)) - \mathbb{E}_{R,z}R^{G,T}(z)\big) = \mathbb{E}_R\Big(\big(\mathbb{E}_{G,T,w}R^{G,T}(G(w))\big) - \big(\mathbb{E}_{G,T,z}R^{G,T}(z)\big)\Big)$$

where $\mathbb{E}_R$ denotes an expectation over the internal randomness of R. It in turn suffices to prove $\big(\mathbb{E}_{G,T,w}R^{G,T}(G(w))\big) - \big(\mathbb{E}_{G,T,z}R^{G,T}(z)\big) < \frac{\epsilon}{2}$ for every fixing of the internal randomness of R and then average. Hence for the remainder we assume R is deterministic.

Let $\mathcal{D}$ name the distribution over (G,T) we have been concerned with thus far. Now, consider the following experiment: sample $G : \{0,1\}^n \to \{0,1\}^N$ at random, sample $S \subseteq \{0,1\}^n$ by choosing each input to lie in S with probability p, sample $w \sim \{0,1\}^n$, and $z \sim \{0,1\}^N$ (all of these independently). Now set $G' = G[w \mapsto z]$, where this notation means

$$G'(w') = \begin{cases} z \text{ if } w' = w \\ G(w') \text{ otherwise} \end{cases}$$

Define $T = T_{G|_{\overline{S}}}$, $T' = T_{G'|_{\overline{S}}}$. Overall we have a joint distribution $\mathcal{D}'$ over tuples (G, G', T, T', w, z). We now observe:

1. The distribution $(G,T,z)_{(G,G',T,T',w,z)\sim\mathcal{D}'}$ equals the distribution $(G,T,z)_{(G,T)\sim\mathcal{D},z\sim\{0,1\}^N}$.
2. The distribution $(G',T',z)_{(G,G',T,T',w,z)\sim\mathcal{D}'}$ equals the distribution $(G,T,G(w))_{(G,T)\sim\mathcal{D},w\sim\{0,1\}^n}$.

Hence, by linearity of expectation, in order to prove $\big(\mathbb{E}_{G,T,w}R^{G,T}(G(w))\big) - \big(\mathbb{E}_{G,T,z}R^{G,T}(z)\big) < \frac{\epsilon}{2}$ it suffices to show:

$$\mathbb{E}_{G,G',T,T',w,z}\big[R^{G',T'}(z) - R^{G,T}(z)\big] < \frac{\epsilon}{2}$$

We show that $\Pr[R^{G',T'}(z) \neq R^{G,T}(z)] < \frac{\epsilon}{2}$ which will conclude the proof. We need to compare the executions $R^{G',T'}(z), R^{G,T}(z)$. To do this we introduce a third "reference execution" and argue that both of these will equal the reference execution with high probability. Recall that our sampling process for generating (G, G', T, T', w, z) involved another random variable S which we have since dropped from our notation. We reintroduce it here, so that now the relevant joint distribution is (G, G', T, T', S, w, z). Define $S^+ = S \cup \{w\}$, $\rho^+ = G|_{\overline{S^+}} = G'|_{\overline{S^+}}$, $T^+ = T_{\rho^+}$. Now, our "reference execution" will be $R^{G,T^+}(z)$; we refer to $R^{G',T'}(z)$ as the "planted execution" and $R^{G,T}(z)$ as the "unplanted execution." We aim to show that both the planted and unplanted executions match the reference execution except with probability $\frac{\epsilon}{2}$. In particular, we will define a set $Q \subseteq \{0,1\}^n$ jointly distributed with $(G, G', T, T', T^+, S, S^+, w, z)$. We show that whenever the planted/unplanted execution differs from the reference, we must have $w \in Q$. We then show that $\Pr[w \in Q] < \frac{\epsilon}{2}$.

To define Q, we first include every input at which G is queried in the execution $R^{G,T^+}(z)$. In addition, for each y on which this execution queries T^+ and receives the answer $T^+(y) = 1$, we choose some $x \in \{0,1\}^m$ and some $F \supseteq \rho^+$ such that $H^F(x) = y$ and include in Q all inputs at which F is queried in the execution $H^F(x)$; choose these x, F canonically depending only on ρ^+, y. The total size of Q is $O(q^2)$. We need to show that whenever one of the planted/unplanted executions differs from the reference we must have $w \in Q$. In order for such a difference to occur, one of the planted/unplanted executions must either query G' on an input on which it differs from G (this only applies for the planted case), or else query T/T' on an input on which it differs from T^+. Consider the first occurrence of such an event (first in chronological order over the steps of the execution when running all three executions in parallel). Before the relevant query is made, we have

reached the same internal state in the reference and the planted (resp. unplanted) execution. If the divergence of the execution occurs because the planted execution queries G' on an input on which it differs from G, then this input must be w; but then the reference must query G on w as well (since they are in the same state before this query), so $w \in Q$. Otherwise, we have that the next query is $y \in \{0,1\}^M$, and either $T'(y) \neq T^+(y)$ or $T(y) \neq T^+(y)$. Note that for any pair of partial assignments $\tau \supseteq \sigma$, we must have $T_\tau(y) \leq T_\sigma(y)$ for every y, so in particular $T(y), T'(y) \leq T^+(y)$. The only case to consider is therefore when $T^+(y) = 1$, while one of $T(y), T'(y)$ equals 0. In this case, let (x, F) be the canonical witness that $T^+(y) = 1$ (depending only on ρ^+, y). We claim that w must have been queried in the execution $H^F(x)$ and hence $w \in Q$; otherwise we would have that, for $F^0 = F[w \mapsto G(w)]$, $F^1 = F[w \mapsto G'(w)]$, (x, F^0) and (x, F^1) each form a witness that $T(y) = 1$ and $T'(y) = 1$ respectively.

Now, observe that Q is completely determined by G, z, S^+. Conditioned on G, z, S^+ the distribution of w is uniform on S^+, so we have

$$\Pr[w \in Q] \leq \mathbb{E}_{G,z,S^+} \frac{|Q \cap S^+|}{|S^+|} \leq \Pr[|S^+| < \frac{s}{2}] + O(\frac{q^2}{s}) \leq e^{-s/8} + O(\frac{q^2}{s})$$

We can strictly upper bound the final value by $\frac{\epsilon}{2}$ provided $s \geq \frac{Cq^2}{\epsilon}$ for a sufficiently large universal constant C.

□

4.1 Constructions of PRFs from PRGs

We now apply Theorem 1 to give depth lower bounds for constructing PRFs from PRGs. We refer the reader to Definition 4; we will use here the same parameters $q, t, \epsilon, d, k, \ell, n, N$. In the following we say that a reduction has *polynomial loss* if $q, \epsilon^{-1} \leq (nt)^{O(1)}$.

Theorem 2 (Necessary depth)**.** *Any polynomial-loss PRG-to-PRF construction with parameters $(q, t, 0.9, \epsilon, d)$, $q, d \geq 1$, and $\log(ntq) = o(n)$ satisfies*

$$t \leq \big(k + O(\log(ntqd))\big) \left(\frac{N}{n - O(\log(ntq))}\right)^d.$$

Proof. Fix distinct $u^1, \ldots, u^t \in \{0,1\}^\ell$ and set $H^G(x) = F_x^G(u^1) \cdots F_x^G(u^t)$. Evaluating these in parallel gives a (tq, d)-construction with seed length k and output length t, and $\mathrm{Adv}_{H^G}(T) = \mathrm{Adv}^{\mathrm{prf}}_{F^G}(\bar{u}, T)$. Since $R_{\bar{u}}$ makes at most q queries, $(H, R_{\bar{u}})$ is a $(tq, 0.9, \epsilon, d)$ stretch amplification reduction when $t > k$; we can then apply Theorem 1. If $t \leq k$, the bound is immediate. □

We compare this to what is achievable from known constructions. In the standard regime $t = n^{\omega(1)}$, our lower bound is essentially tight:

Proposition 1 (Achievability [GM84, GGM86, GL89, L87])**.** *There is a constant $0 < c < 1/2$ such that, for all sufficiently large n, $\ell \leq \mathrm{poly}(n)$, and $1 \leq t \leq \min\{2^\ell, 2^{cn}\}$, known constructions are able to achieve a PRG-to-PRF construction with $\gamma = 0.9$, $k = \Theta(n)$, $q, \epsilon^{-1} \leq (nt)^{O(1)}$, only polynomially many queries to G in each PRF evaluation, and adaptivity depth d satisfying*

$$t \geq \frac{1}{n} \left(\frac{N}{n - \log(nt)}\right)^{\Omega(d)}.$$

This proposition follows directly from the stated prior work [GM84, GGM86, GL89, L87]; we have not found a source which precisely states a general bound of this form so we spell out the construction briefly. We first apply the amplification methods in [GM84, GGM86, GL89] to obtain a generator $H^G : \{0,1\}^m \to \{0,1\}^{m^2}$ for some $m = \Theta(n)$ in depth $d_0 = O(\frac{\log n}{\log(N/(n-\log tn))})$ (apply [GL89] to obtain additive $\log tn$ stretch, [GM84] to obtain doubling stretch, then [GGM86] to obtain quadratic stretch, skipping the first/second steps if the initial stretch is already large enough). We then use the quadratic stretch PRG to build the PRF as follows: first use Levin's domain extension trick [L87] to hash the PRF input to length $O(\log nt)$ (using a $1/(100t^2)$-almost-universal hash with an independent $O(\log(nt))$-bit seed included in the key), and then apply [GGM86] with branching factor m for this compressed input length using H^G as the PRG. The depth of this PRF construction in terms of the calls it makes to H^G is $d_1 = O(1 + \frac{\log nt}{\log n})$, so its overall depth in terms of calls to G is $d = d_0 d_1 = O(\frac{\log nt}{\log(N/(n-\log(nt)))})$ which gives the bound in Proposition 1.

5 Oracle Separations

In this section we use our techniques to give the following oracle separation:

Theorem 4. *There exists an oracle O such that $\mathbf{E}^O \not\subseteq \mathbf{NC}^O/\mathrm{poly}$ and $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\mathrm{poly}$. In fact, for this oracle O there exists a language $L \in \mathbf{E}^O$ such that L does not have O-oracle circuits of size $2^{\sqrt{n}}$ and depth $\sqrt{n}/8$ for all sufficiently large n.*

We then apply this to give the relativized separation in bounded arithmetic mentioned in the introduction; this will follow directly from Theorem 4 by standard witnessing theorems. To construct our oracle we use the following consequence of the relativizing reduction from Avoid to hard truth table construction in [K21] (as mentioned in the introduction, this lemma follows from previous work of [J04] in bounded arithmetic):

Lemma 4 ([K21])**.** *Let O be any oracle and let $(G_n : \{0,1\}^n \to \{0,1\}^{2n})_{n\in\mathbb{N}}$ be uniformly computable in polynomial time with access to the oracle O. Suppose that for every deterministic machine M with access to O and $\mathbf{SAT}^O$, running in time $\lfloor 2^{n/32} \rfloor$ on input 1^n and outputting a string in $\{0,1\}^{2n}$, we have that M outputs a string outside the range of G_n on input 1^n for at most finitely many n. Then $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\mathrm{poly}$.*

We will not peer into the proof of this lemma formally, but it is worth noting that the polynomial-size circuits computing every language in $\mathbf{E}^{\mathbf{NP}^O}$ here are precisely the GGM iterations [GGM86] of the function G_n, with the key fixed to an appropriate advice string depending on the language being computed.

Proof of Theorem 4. Our proof proceeds in two parts; we first show that we may fix some partial information about the oracle to guarantee that $\mathbf{E}^O \not\subseteq \mathbf{NC}^O/\mathrm{poly}$; we then show that after this partial fixing of information, enough freedom remains to allow the collapse $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\mathrm{poly}$.

1. Fixing an NC Lower Bound after Random Restriction: We will consider oracles $O(G,\theta)$ indexed by a pair G,θ, where $G = (G_n : \{0,1\}^n \to \{0,1\}^{2n})_{n\in\mathbb{N}}$, and $\theta = (\theta_k : \{0,1\}^k \to \{0,1\})_{k\in\mathbb{N}}$; we will treat $\theta_k : \{0,1\}^k \to \{0,1\}$ also as a string $\theta_k \in \{0,1\}^{2^k}$ at various points. The oracle $O(G,\theta)$ is defined by

$$O(G,\theta)\big(0\,1^n 0\,w\ \mathrm{bin}_{\lceil\log(2n)\rceil}(i)\big) = (G_n(w))_{i+1},$$
$$O(G,\theta)\big(1u0^{2^{k+2}-k-1}\big) = \theta_k(u),$$

for $n, k \geq 1$, $w \in \{0,1\}^n$, $u \in \{0,1\}^k$, and $0 \leq i < 2n$; the value of $O(G,\theta)$ is zero everywhere else. Here $\mathrm{bin}_r(i)$ denotes the r-bit binary encoding of i.

Let $p_n = 2^{-7n/8}$. Independently sample $\rho_n \sim \mathcal{R}_{p_n}^{n,2n}$ for each $n \in \mathbb{N}$. For such a choice of $\rho = (\rho_n)_{n\in\mathbb{N}}$, we will consider a family of oracles $\mathcal{O}^\rho := \{O(G,\theta) \mid G_n \supseteq \rho_n \text{ for all } n\}$. Now, let $r_k = \lfloor 2^{\sqrt{k}} \rfloor$, $d_k = \lfloor \sqrt{k}/8 \rfloor$. For an oracle O, let $H_k^O : \{0,1\}^{m_k} \to \{0,1\}^{2^k}$ be the truth table generator which maps descriptions of O-oracle circuits with k inputs, with size at most r_k and depth at most d_k, to the truth tables they compute under the oracle O. Under a standard encoding we will have:

1. $m_k = O(r_k \log r_k)$
2. For each circuit description $x \in \{0,1\}^{m_k}$, we can evaluate $H_k^O(x)$ by making at most 2^{2k} queries to O in at most d_k adaptive phases (for sufficiently large k).
3. If $O = O(G,\theta)$, for each circuit description $x \in \{0,1\}^{m_k}$, $H_k^O(x)$ will only query bits of G_n with $n \leq r_k$ and values $\theta_\ell(u)$ with $\ell \leq \lfloor \log r_k \rfloor \leq \sqrt{k}$; queries outside the encoding have answer zero and hence are predetermined. In particular, there are at most $2r_k$ potentially accessible θ-bits. These bounds follow from our encoding of the oracle $O(G,\theta)$.

We will now claim that, with probability 1 over ρ, there exists $(y_k : \{0,1\}^k \to \{0,1\})_{k\in\mathbb{N}}$ such that for all $O \in \mathcal{O}^\rho$ we have:

$$y_k \notin \mathrm{range}(H_k^O) \text{ for all } k \text{ sufficiently large}$$

This will follow from Lemma 1. To prove it, fix some k and condition on $\rho_{<4k} := (\rho_n)_{n<4k}$. Let $t_k = \sum_{n<4k} 2n|\mathrm{free}(\rho_n)|$ and define a map

$$\widetilde{H}_k^{O,\rho_{<4k}} : \{0,1\}^{m_k+t_k+2r_k} \to \{0,1\}^{2^k}$$

Given (x,a,b), this map simulates $H_k^O(x)$, using the t_k bits of a to fill in the values of the free inputs of $\rho_{<4k}$ and the $2r_k$ bits of b to specify all accessible θ-bits. Thus when $O = O(G,\theta)$, only queries to G_n with $4k \leq n \leq r_k$ require querying the actual oracle O, and at most $q := 2^{2k}$ such queries occur across at most d_k phases of adaptivity. Moreover,

$$\bigcup_{O\in\mathcal{O}^\rho} \mathrm{range}(H_k^O) = \bigcup_{O\in\mathcal{O}^\rho} \mathrm{range}(\widetilde{H}_k^{O,\rho_{<4k}}).$$

Now, we'd like to treat $\tilde{H}_k^{O,\rho_{<4k}}$ as a stretch-amplifying construction querying some G_n and apply Lemma 1 to argue that for any conditioning of $\rho_{<4k}$ such that t_k is suitably small, with high probability over ρ there will be some fixed y_k outside the union of ranges $\bigcup_{O\in\mathcal{O}^\rho} \mathrm{range}(\tilde{H}_k^{O,\rho_{<4k}})$. A slight subtlety is that $\tilde{H}_k^{O,\rho_{<4k}}$ is not querying a single map G_n, but rather a sequence $G_{4k}, \ldots, G_{r_k}$, and the random restrictions are being applied to these maps at different free-probability rates $p_{4k} \geq \cdots \geq p_{r_k}$. To get around this, we will combine the information $G_{4k}, \ldots, G_{r_k}$ into a single map F_k, and consider a restriction being applied to F_k at the highest free-probability rate p_{4k}. In particular, let $F_k : \{0,1\}^{r_k+1} \to \{0,1\}^{2r_k}$; interpret the input of F_k as a pair (n,w) with $n \in \{4k, \ldots, r_k\}$ and $w \in \{0,1\}^n$ (according to any suitable fixed encoding), set $F_k(n,w) = G_n(w)$, where $G_n(w) \in \{0,1\}^{2^n}$ is padded with zeroes if necessary so that it lies in $\{0,1\}^{2r_k}$; hence we interpret F_k as a function:

$$F_k : \bigcup_{n\in\{4k,\ldots,r_k\}} \{0,1\}^n \to \bigcup_{n\in\{4k,\ldots,r_k\}} \{0,1\}^{2n}$$

where the "input region" $\{0,1\}^n$ is mapped into the "output region" $\{0,1\}^{2n}$ for each $n \in \{4k, \ldots, r_k\}$. We consider the distribution $\mathcal{R}_k$ on partial assignments, where $\tau \sim \mathcal{R}_k$ is sampled

by setting $\tau(w) = *$ independently with probability p_n and setting $\tau(w) = z$ with probability $(1-p_n)2^{-2n}$ for each $z \in \{0,1\}^{2n}$, where $n \in \{4k, \ldots, r_k\}$ is the value such that w lies in the input region $\{0,1\}^n$. Define $\mathcal{R}'_k$ by replacing p_n with p_{4k} throughout this sampling rule in every region. Let $\tilde{H}_k^{F,\rho_{<4k}}$ be the function defined by using the values of F to answer the queries to O in $\tilde{H}_k^{O,\rho_{<4k}}$. We claim that:

$$\Pr_{\tau\sim\mathcal{R}_k}\left[\bigcup_{F'\supseteq\tau}\operatorname{range}\big(\widetilde{H}_k^{F',\rho_{<4k}}\big)=\{0,1\}^{2^k}\right] \le \Pr_{\tau\sim\mathcal{R}'_k}\left[\bigcup_{F'\supseteq\tau}\operatorname{range}\big(\widetilde{H}_k^{F',\rho_{<4k}}\big)=\{0,1\}^{2^k}\right]$$
$$\le \operatorname{poly}(qd_k)\,2^{m_k+t_k+2r_k-2^k\left(\frac{\log(1/(qp_{4k}))}{2r_k}\right)^{d_k}} = \operatorname{poly}(qd_k)\,2^{m_k+t_k+2r_k-2^k\left(\frac{3k}{4r_k}\right)^{d_k}}$$
$$\le 2^{m_k+t_k+2r_k+O(k)-2^{7k/8}}.$$

The first inequality follows from the fact that freeing additional inputs in a partial assignment can only enlarge the union of ranges over its completions. The second inequality applies Lemma 1; specifically, we use the modified version of the lemma discussed in Observation 1. At this point it remains to remove the conditioning on $\rho_{<4k}$. Since each input to ρ_n is free with probability p_n, we have

$$\mathbb{E}_\rho t_k = \sum_{n<4k} 2n\,2^n p_n = \sum_{n<4k} 2n\,2^{n/8} \le \operatorname{poly}(k)2^{k/2}.$$

By Markov,

$$\Pr_\rho[t_k > 2^{3k/4}] \le \operatorname{poly}(k)2^{-k/4}.$$

Splitting according to this event and applying the conditional bound above gives

$$\Pr_\rho\left[\bigcup_{O\in\mathcal{O}^\rho}\operatorname{range}(H_k^O)=\{0,1\}^{2^k}\right] \le \operatorname{poly}(k)2^{-k/4} + 2^{m_k+2r_k+2^{3k/4}+O(k)-2^{7k/8}}.$$

Since $m_k + 2r_k = 2^{O(\sqrt{k})}$, these failure probabilities are summable over k. By Borel–Cantelli, with probability one, $\bigcup_{O\in\mathcal{O}^\rho}\operatorname{range}(H_k^O)$ is a proper subset of $\{0,1\}^{2^k}$ for every sufficiently large k. Choose y_k outside this union for each such k, and choose the finitely many remaining y_k arbitrarily.

We now make an additional easy claim: with probability 1 over ρ, $|\operatorname{free}(\rho_n)| \ge \frac{1}{2}2^{n/8}$ for all sufficiently large n. Indeed, $|\operatorname{free}(\rho_n)|$ is binomial with mean $2^n p_n = 2^{n/8}$, so a Chernoff bound gives $\Pr_\rho\big[|\operatorname{free}(\rho_n)| < \frac{1}{2}2^{n/8}\big] \le \exp\big(-\frac{2^{n/8}}{8}\big)$. These probabilities are summable over n, so again Borel–Cantelli proves the claim.

At this point we may now fix a restriction sequence ρ satisfying both key properties: there exists some $y = (y_k)_{k\in\mathbb{N}}$ as above, and $|\operatorname{free}(\rho_n)| \ge \frac{1}{2}2^{n/8}$ for all sufficiently large n. We will now fix some ρ, one of its corresponding y's, and consider from now on only oracles in the set $\mathcal{O}^{\rho,y} = \{O(G,y) \mid G_n \supseteq \rho_n \text{ for all } n\}$. We claim that for any $O \in \mathcal{O}^{\rho,y}$ we have that $\mathbf{E}^O \not\subseteq \mathbf{NC}^O/\operatorname{poly}$; in fact, there is a language in $\mathbf{E}^O$ which does not have O-oracle circuits of size $2^{\sqrt{n}}$ and depth $\sqrt{n}/8$ for sufficiently large n. In particular the witnessing language will be:

$$L^O := \bigcup_{n\ge1}\left\{x \in \{0,1\}^n : O\big(1x0^{2^{n+2}-n-1}\big) = 1\right\}.$$

Clearly $L^O \in \mathbf{E}^O$. We observe that if $O = O(G,y)$ then for $x \in \{0,1\}^n$, $x \in L^O \leftrightarrow y_n(x) = 1$. By construction we have that y_n avoids the truth table generator for O-oracle circuits of size $\lfloor 2^{\sqrt{n}}\rfloor$ and depth $\lfloor\sqrt{n}/8\rfloor$.

2. Making Range Avoidance hard: Let $T(n) = \lfloor 2^{n/32} \rfloor$. We will choose $O = O(G, y) \in \mathcal{O}^{\rho,y}$ so that every deterministic machine with access to O and $\mathbf{SAT}^O$, running in time $T(n)$ on input 1^n, fails to find a string outside range(G_n) for all sufficiently large n. Applying Lemma 4 this will immediately yield $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\text{poly}$ and finish the proof, since G_n is computable in poly(n) time with an $O(G, y)$ oracle. The argument here is essentially identical to Wilson's construction of an oracle relative to which $\mathbf{E}^{\mathbf{NP}} \subseteq \mathbf{P}/\text{poly}$ [W83] (recall that Wilson's oracle separation is also roughly equivalent to Lemma 2; it requires only some additional diagonalization bookkeeping).

Enumerate all candidate machines as $M_1, M_2, \ldots$. Our construction will proceed in stages indexed by $\mathbb{N}$. In a given stage, we maintain some $\tau \supseteq \rho$ (this notation is shorthand for the statement "$\tau_n \supseteq \rho_n$ for all n"). We will extend τ in each stage to some $\tau' \supseteq \tau$; we don't introduce a new variable name for the extension of τ, we will just write "extend τ as follows" in each stage. For each n, after stage n we will have $\tau_n(w) \neq *$ for all $w \in \{0,1\}^n$; in this way the construction will uniquely determine some $G \supseteq \rho$, and hence some $O \in \mathcal{O}^{\rho,y}$. Choose n_0 large enough that for all $n \geq n_0$,

$$n^2\big(T(n)^2 + T(n) + 1\big) = 2^{n/16 + O(\log n)} < \tfrac{1}{2}2^{n/8} \leq |\text{free}(\rho_n)| \tag{1}$$

Initialize $\tau = \rho$. For each $n < n_0$, complete τ_n arbitrarily. Now, for each stage $n \geq n_0$ we will simulate each of $M_1, \ldots, M_n$ on 1^n for $T(n)$ steps as follows. Say that τ is our current partial assignment. If some M_i queries a bit of some $G_k(w)$, we extend τ to fix its value $\tau_k(w) \in \{0,1\}^{2k}$ arbitrarily, and then respond with the requested bit of $\tau_k(w)$. All other direct queries to the oracle are already determined by the condition $O \in \mathcal{O}^{\tau,y}$. For a $\mathbf{SAT}^O$ query, in the form of some oracle circuit $C^{(\cdot)}$, $|C^{(\cdot)}| \leq T(n)$, we do the following. If there exists some $O' \in \mathcal{O}^{\tau,y}$ and some x such that $C^{O'}(x) = 1$, we immediately fix all queries $G_k(w)$ made to G in the computation $C^{O'}(x)$ and extend τ to agree with O' on these, and then proceed in the simulation of M_i assuming this $\mathbf{SAT}^O$ query output "SATISFIABLE." If no such O' existed, we make no adjustment to τ, and continue the simulation of M_i assuming this query had the output "UNSATISFIABLE." It is easy to see that by this rule, we can guarantee that we have provided the correct $\mathbf{SAT}^O$ oracle responses for all oracles O satisfying $O \in \mathcal{O}^{\tau,y}$. Observe that each machine simulation extends τ on at most $T(n)^2 + T(n)$ additional inputs. If, at the end of its computation, some machine M_i outputs a string $z \in \{0,1\}^{2n}$, we choose some remaining free input $w \in \text{free}(\tau_n)$ and then extend τ to set $\tau_n(w) = z$. Such an input w always exists by (1). At the end of stage n, if any $\tau_n(w)$ remains undefined, extend τ by setting it to 0^{2n}.

In this staged process, we are able to guarantee that for all machines M_n and all $n' \geq \max\{n, n_0\}$, if $\tau^{(n')}$ is the state of the maintained partial assignment τ after stage n', then $M_n^{O,\mathbf{SAT}^O}(1^{n'})$ fails to find a string in $\{0,1\}^{2n'} \setminus \text{range}(G_{n'})$ in $T(n')$ steps for all $O \in \mathcal{O}^{\tau^{(n')},y}$. For the unique oracle $O \in \bigcap_{n \geq n_0} \mathcal{O}^{\tau^{(n)},y}$ we hence have that every machine fails to solve Avoid on G_n^O for all but finitely many n and we are done. □

We now move on to the relativized separation in Bounded Arithmetic. We start with the following observation, that the above theorem can be strengthened so that the circuit upper bound is $\mathbf{Size}^O[n^k]$ for a fixed universal constant k:

Observation 2. *There exists a universal constant k and an oracle O such that (1) there is a language $L \in \mathbf{E}^O$ such that L does not have O-oracle circuits of size $2^{\sqrt{n}}$ and depth $\sqrt{n}/8$ for all sufficiently large n, and (2) every language $L \in \mathbf{E}^{\mathbf{NP}^O}$ has O-oracle circuits of size n^k for all $n \geq n_L$, where n_L depends on L.*

In fact if we were more careful in the proof of Theorem 4 we could establish this directly for a very reasonable constant k, but this would require peering into the proof of Lemma 4. Instead we

observe that this follows as a black box from Theorem 4, using the fact that for every oracle O, $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{P}^O/\text{poly}$ if and only if there is a universal constant k such that $\mathbf{E}^{\mathbf{NP}^O} \subseteq \mathbf{Size}^O[O(n^k)]$. This follows from the relativizing fact that $\mathbf{E}^{\mathbf{NP}^O}$ contains complete problems under linear time reductions.

Our results concern the bounded arithmetic theory $S^1_2(\alpha)$, which is a relativized theory of bounded arithmetic in which an uninterpreted unary predicate symbol α is added to the language. The BASIC axioms remain unchanged (and do not involve α), and the polynomial induction schema is extended to all $\Sigma^b_1(\alpha)$ formulas over the expanded language $\mathcal{L}_2(\alpha) := \mathcal{L}_2 \cup \{\alpha\}$, where $\mathcal{L}_2$ is the standard language of bounded arithmetic. See [K95] for a comprehensive introduction to bounded arithmetic.

Now, we consider the following $\forall\Sigma^b_2(\alpha)$ sentences in the language $\mathcal{L}_2(\alpha)$, where k is the universal constant from Observation 2. Each is parameterized by a fixed standard number n_0:

1. $\text{DepthLB}_{n_0}(\alpha)$: For all a, $n := ||a|| \geq n_0$, there exists $f : \{0,1\}^n \to \{0,1\}$ such that for all α-oracle circuits C, $|C| \leq 2^{\sqrt{n}}$, $\text{depth}(C) \leq \sqrt{n}/8$, there is $x \in \{0,1\}^n$ such that $C^\alpha(x) \neq f(x)$.
2. $\text{CircuitLB}_{n_0}(\alpha)$: For all a, $n := ||a|| \geq n_0$, there exists $f : \{0,1\}^n \to \{0,1\}$ such that for all α-oracle circuits C, $|C| \leq n^k$, there is $x \in \{0,1\}^n$ such that $C^\alpha(x) \neq f(x)$.

Note that the final quantifier over $x \in \{0,1\}^n$ in each sentence is sharply bounded. We formalize these sentences in $\mathcal{L}_2(\alpha)$ in the standard way, following Jeřábek. Note that the restriction to n with $n = ||a||$ for some a is necessary if we wish to quantify over truth tables of functions $f : \{0,1\}^n \to \{0,1\}$ in the natural way. We prove the following relativized separation:

Theorem 5. *For all sufficiently large $n_0 \in \mathbb{N}$ and every $n_0' \in \mathbb{N}$,*

$$S^1_2(\alpha) + \text{DepthLB}_{n_0}(\alpha) \nvdash \text{CircuitLB}_{n_0'}(\alpha)$$

Let Φ be a $\forall\Sigma^b_2(\alpha)$ sentence of the form $\Phi = \forall a \exists x \leq t_1(a) \forall y \leq t_2(a,x) \phi(a,x,y)$ where $\phi(a,x,y)$ is a sharply bounded $\mathcal{L}_2(\alpha)$ formula and t_1, t_2 are $\mathcal{L}_2$ terms. For an oracle $O : \{0,1\}^* \to \{0,1\}$, we can define a $(\mathbf{F\Sigma^P_2})^O$ search problem $\text{Search}^O(\Phi)$ as follows: given $a \in \mathbb{N}$ in binary, find $x \leq t_1(a)$ such that for all $y \leq t_2(a,x)$, $\phi^O(a,x,y)$ holds, where the $\mathcal{L}_2$ symbols are interpreted via the standard model of arithmetic and $\phi^O(a,x,y)$ is defined by substituting O for α (coding inputs to α, which are elements of $\mathbb{N}$, by their binary encodings in $\{0,1\}^*$ in order to feed them into O). If the sentence Φ is true in the standard model of arithmetic when O is substituted for α, then this will be a total search problem, i.e. it will lie in $(\mathbf{TF\Sigma^P_2})^O$. We require the following standard consequence of the polynomial round KPT-style witnessing theorem for S^1_2 [P92, K92, K25b] [10]:

Theorem 6 ([P92, K92, K25b])**.** *Let Φ_0, Φ_1 be $\forall\Sigma^b_2(\alpha)$ formulas such that $S^1_2(\alpha) + \Phi_0(\alpha) \vdash \Phi_1(\alpha)$. Then for every oracle O, if $\text{Search}^O(\Phi_0) \in \mathbf{FP}^O$ then $\text{Search}^O(\Phi_1) \in \mathbf{FP}^{\mathbf{NP}^O}$.*

We can now prove Theorem 5:

Proof of Theorem 5. For an oracle O, the search problems associated with $\text{DepthLB}_{n_0}(\alpha)$ and $\text{CircuitLB}_{n_0'}(\alpha)$ are given as follows:

1. Given 1^{2^n}, $n \geq n_0$ output the truth table of a function $f : \{0,1\}^n \to \{0,1\}$ which does not have O-oracle circuits of size $2^{\sqrt{n}}$ and depth $\sqrt{n}/8$.

[10] [P92] proves this in the case Φ_0 is omitted (and its search problem is removed from the statement of the theorem); [K92] provides an alternate proof. The form including Φ_0 follows from the relativization of the witnessing theorem for $S^1_2(\text{PV}) + T_{\text{PV}}$ in [K25b, Theorem 8.1.1(2)], where T_{PV} is the true universal theory in the language of PV.

2. Given 1^{2^n}, $n \geq n_0'$ output the truth table of a function $f : \{0,1\}^n \to \{0,1\}$ which does not have O-oracle circuits of size n^k.

For the specific oracle we have constructed above (Observation 2) and some n_0 sufficiently large, the first problem is in $\mathbf{FP}^O$ but the second is not in $\mathbf{FP}^{\mathbf{NP}^O}$ for any n_0'. Applying Theorem 6 we derive Theorem 5. □

Acknowledgments

The author thanks Toniann Pitassi and Tianqi Yang for numerous valuable discussions about the subject of this work.

References

[ADV26] Bar Alon, Itai Dinur, and Muthuramakrishnan Venkitasubramaniam. Lower bounds on black-box constructions of pseudorandom functions. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, volume 16800 of *Lecture Notes in Computer Science*, pages 159–189. Springer, 2026. Full version: https://arxiv.org/abs/2608.14501. doi:10.1007/978-3-032-35367-2_6.

[AIK06a] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006. doi:10.1007/s00037-006-0211-8.

[AIK06b] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. *SIAM Journal on Computing*, 36(4):845–888, 2006. doi:10.1137/S0097539705446950.

[AIK08] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. On pseudorandom generators with linear stretch in NC^0. *Computational Complexity*, 17(1):38–69, 2008. doi:10.1007/s00037-007-0237-6.

[AR16] Benny Applebaum and Pavel Raykov. Fast pseudorandom functions based on expander graphs. In *Theory of Cryptography – TCC 2016-B, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 27–56. Springer, 2016. doi:10.1007/978-3-662-53641-4_2.

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737. Springer, 2012. doi:10.1007/978-3-642-29011-4_42.

[BMG07] Boaz Barak and Mohammad Mahmoody-Ghidary. Lower bounds on signatures from symmetric primitives. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*, pages 680–688. IEEE Computer Society, 2007. URL: https://www.boazbarak.org/Papers/siglb.pdf.

[BMM24] Amos Beimel, Tal Malkin, and Noam Mazor. Structural lower bounds on black-box constructions of pseudorandom functions. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part V*, volume 14924 of *Lecture Notes in Computer Science*, pages 459–488. Springer, 2024. Full version: https://eprint.iacr.org/2024/1104. doi:10.1007/978-3-031-68388-6_16.

[BM82] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *23rd Annual Symposium on Foundations of Computer Science (FOCS 1982)*, pages 112–117. IEEE Computer Society, 1982. doi:10.1109/SFCS.1982.72.

[BG11] Zvika Brakerski and Oded Goldreich. From absolute distinguishability to positive distinguishability. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 141–155. Springer, 2011. First appeared as ECCC Report TR09-031, 2009; https://eccc.weizmann.ac.il/report/2009/031/. doi:10.1007/978-3-642-22670-0_17.

[BJP11] Josh Bronson, Ali Juma, and Periklis A. Papakonstantinou. Limits on the stretch of non-adaptive constructions of pseudo-random generators. In Yuval Ishai, editor, *Theory of Cryptography – TCC 2011*, volume 6597 of *Lecture Notes in Computer Science*, pages 504–521. Springer, 2011. doi:10.1007/978-3-642-19571-6_30.

[CHLR23] Yeyuan Chen, Yizhi Huang, Jiatu Li, and Hanlin Ren. Range avoidance, remote point, and hard partial truth table via satisfying-pairs algorithms. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1058–1066, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585147.

[DV21] Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. In *Proceedings of the 34th Conference on Learning Theory (COLT 2021)*, volume 134 of *Proceedings of Machine Learning Research*, pages 1358–1394. PMLR, 2021. URL: https://proceedings.mlr.press/v134/daniely21a.html.

[DJK25] Youlong Ding, Aayush Jain, and Ilan Komargodski. A new approach for LPN-based pseudorandom functions: Low-depth and key-homomorphic. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pages 1898–1909. ACM, 2025. Full version: https://eprint.iacr.org/2025/898. doi:10.1145/3717823.3718210.

[DJK26a] Youlong Ding, Aayush Jain, and Ilan Komargodski. Logarithmic-depth pseudorandom functions from well-founded code-based assumptions. Cryptology ePrint Archive, Paper 2026/261, 2026. TCC 2026. https://eprint.iacr.org/2026/261.

[DJK26b] Youlong Ding, Aayush Jain, and Ilan Komargodski. Pseudorandom functions in NC^1 from LWE/LPN/CDH (Or: How to build PRFs in NC^1, generically). arXiv:2608.25213, 2026. https://arxiv.org/abs/2608.25213. doi:10.48550/arXiv.2608.25213.

[GGNS23] Karthik Gajulapalli, Alexander Golovnev, Satyajeet Nagargoje, and Sidhant Saraogi. Range Avoidance for Constant Depth Circuits: Hardness and Algorithms. In Nicole Megow and Adam Smith, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*, volume 275 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 65:1–65:18, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops-dev.dagstuhl.de/entities/document/10.4230/LIPIcs.APPROX/RANDOM.2023.65, doi:10.4230/LIPIcs.APPROX/RANDOM.2023.65.

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986. doi:10.1145/6490.6503.

[GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 25–32. ACM, 1989. Full version: https://www.cs.bu.edu/fac/lnd/pdf/hard.pdf. doi:10.1145/73007.73010.

[GM84] Oded Goldreich and Silvio Micali. The weakest pseudo-random generator implies the strongest one. Unpublished manuscript, October 1984. URL: https://www.wisdom.weizmann.ac.il/~oded/annot/node15.html.

[GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2017)*, pages 612–621. IEEE Computer Society, 2017. Full version: https://eprint.iacr.org/2017/274. doi:10.1109/FOCS.2017.62.

[GLW22] Venkatesan Guruswami, Xin Lyu, and Xiuhan Wang. Range Avoidance for Low-Depth Circuits and Connections to Pseudorandomness. In Amit Chakrabarti and Chaitanya Swamy, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, volume 245 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:21, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: https://drops-dev.dagstuhl.de/entities/document/10.4230/LIPIcs.APPROX/RANDOM.2022.20, doi:10.4230/LIPIcs.APPROX/RANDOM.2022.20.

[GLY26] Venkatesan Guruswami, Xin Lyu, and Weiqiang Yuan. Cell-probe lower bounds via semi-random CSP refutation: Simplified and the odd-locality case. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 372–387. SIAM, 2026. URL: https://epubs.siam.org/doi/10.1137/1.9781611978971.17, doi:10.1137/1.9781611978971.17.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.

[IW97] Russell Impagliazzo and Avi Wigderson. P = BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC 1997)*, pages 220–229. ACM, 1997. doi:10.1145/258533.258590.

[IKOS08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008)*, pages 433–442. ACM, 2008. doi:10.1145/1374376.1374438.

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pages 60–73. ACM, 2021. Full version: https://arxiv.org/abs/2008.09317. doi:10.1145/3406325.3451093.

[JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_p$, DLIN, and PRGs in NC^0. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EURO-*

CRYPT 2022, Part I, volume 13275 of *Lecture Notes in Computer Science*, pages 670–699. Springer, 2022. Full version: https://eprint.iacr.org/2021/1334. doi:10.1007/978-3-031-06944-4_23.

[J04] Emil Jeřábek. Dual weak pigeonhole principle, Boolean complexity, and derandomization. *Annals of Pure and Applied Logic*, 129(1–3):1–37, 2004. URL: https://users.math.cas.cz/~jerabek/papers/wphp.pdf, doi:10.1016/j.apal.2003.12.003.

[J05] Emil Jeřábek. *Weak Pigeonhole Principle, and Randomized Computation*. PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, 2005. URL: https://users.math.cas.cz/~jerabek/papers/phd.pdf.

[J07] Emil Jeřábek. Approximate counting in bounded arithmetic. *Journal of Symbolic Logic*, 72(3):959–993, 2007. URL: https://users.math.cas.cz/~jerabek/papers/apx.pdf.

[J11] Ali Juma. *Leakage Resilience and Black-Box Impossibility Results in Cryptography*. PhD thesis, University of Toronto, 2011. URL: https://www.cs.toronto.edu/~ajuma/phdthesis.pdf.

[K21] Oliver Korten. The hardest explicit construction. In *62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2021)*, pages 433–444. IEEE, 2021. Full version: arXiv:2106.00875v3. doi:10.1109/FOCS52979.2021.00051.

[K25a] Oliver Korten. Range avoidance and the complexity of explicit constructions. *Bulletin of the EATCS*, 145:94–134, 2025. URL: https://bulletin.eatcs.org/index.php/beatcs/article/view/825.

[KPI25] Oliver Korten, Toniann Pitassi, and Russell Impagliazzo. Stronger cell probe lower bounds via local PRGs. *Electron. Colloquium Comput. Complex.*, TR25-030, 2025. To appear in FOCS 2025. URL: https://eccc.weizmann.ac.il/report/2025/030, arXiv:TR25-030.

[K92] Jan Krajíček. No counter-example interpretation and interactive computation. In Yiannis N. Moschovakis, editor, *Logic from Computer Science*, volume 21 of *Mathematical Sciences Research Institute Publications*, pages 287–293. Springer-Verlag, 1992. URL: https://www2.karlin.mff.cuni.cz/~krajicek/scan21.pdf.

[K95] Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*, volume 60 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1995. URL: https://www.karlin.mff.cuni.cz/~krajicek/kniha.pdf.

[K25b] Jan Krajíček. *Proof Complexity Generators*, volume 497 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 2025. URL: https://www.karlin.mff.cuni.cz/~krajicek/k4.pdf, doi:10.1017/9781009611664.

[L87] Leonid A. Levin. One way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987. Full version: https://www.cs.bu.edu/fac/lnd/pdf/onew.pdf. doi:10.1007/BF02579323.

[LW09] Allison B. Lewko and Brent Waters. Efficient pseudorandom functions from the decisional linear assumption and weaker variants. In *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS 2009)*, pages 112–120. ACM, 2009. doi:10.1145/1653662.1653677.

[L06] Chi-Jen Lu. On the complexity of parallel hardness amplification for one-way functions. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography – TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 462–481. Springer, 2006. doi:10.1007/11681878_24.

[MV11] Eric Miles and Emanuele Viola. On the complexity of non-adaptively increasing the stretch of pseudorandom generators. In Yuval Ishai, editor, *Theory of Cryptography – TCC 2011*, volume 6597 of *Lecture Notes in Computer Science*, pages 522–539. Springer, 2011. doi:10.1007/978-3-642-19571-6_31.

[MV15] Eric Miles and Emanuele Viola. On the complexity of constructing pseudorandom functions (especially when they don't exist). *Journal of Cryptology*, 28(3):509–532, 2015. Full version: https://eprint.iacr.org/2016/343. doi:10.1007/s00145-013-9161-x.

[M92] Peter Bro Miltersen. Circuit depth relative to a random oracle. *Information Processing Letters*, 42(6):295–298, 1992. Preprint: https://tidsskrift.dk/daimipb/article/download/6610/5732/22107. doi:10.1016/0020-0190(92)90225-K.

[NR99] Moni Naor and Omer Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. *Journal of Computer and System Sciences*, 58(2):336–375, 1999. doi:10.1006/jcss.1998.1618.

[NR04] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *Journal of the ACM*, 51(2):231–262, 2004. doi:10.1145/972639.972643.

[NRR02] Moni Naor, Omer Reingold, and Alon Rosen. Pseudorandom functions and factoring. *SIAM Journal on Computing*, 31(5):1383–1404, 2002. doi:10.1137/S0097539701389257.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994. doi:10.1016/S0022-0000(05)80043-1.

[P92] Pavel Pudlák. Some relations between subsystems of arithmetic and complexity of computations. In Yiannis N. Moschovakis, editor, *Logic from Computer Science*, volume 21 of *Mathematical Sciences Research Institute Publications*, pages 499–519. Springer-Verlag, 1992. URL: https://users.math.cas.cz/~pudlak/some-relations.pdf.

[RSW22] Hanlin Ren, Rahul Santhanam, and Zhikun Wang. On the range avoidance problem for circuits. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 640–650, 2022. doi:10.1109/FOCS54457.2022.00067.

[V05] Emanuele Viola. On constructing parallel pseudorandom generators from one-way functions. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity (CCC)*, pages 183–197. IEEE Computer Society, 2005. Full version: https://eprint.iacr.org/2005/159. doi:10.1109/CCC.2005.16.

[W83] Christopher B. Wilson. Relativized circuit complexity. In *24th Annual Symposium on Foundations of Computer Science (FOCS 1983)*, pages 329–334. IEEE Computer Society, 1983.

[W87] Christopher B. Wilson. Relativized NC. *Mathematical Systems Theory*, 20(1):13–29, 1987.

[Y82] Andrew C. Yao. Theory and applications of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (FOCS 1982)*, pages 80–91. IEEE Computer Society, 1982. doi:10.1109/SFCS.1982.95.

[YS16] Yu Yu and John Steinberger. Pseudorandom functions in almost constant depth from low-noise LPN. In *Advances in Cryptology – EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 154–183. Springer, 2016. doi:10.1007/978-3-662-49896-5_6.

A Security Reductions with Advice

In this section we show that our depth lower bound for PRG amplification can be strengthened to hold for nonuniform security reductions (this distinction was mentioned in the introduction). Since the PRF depth lower bound (Theorem 2) was derived in a completely generic way from Theorem 1, we can apply this argument to strengthen Theorem 2 to hold for nonuniform reductions as well.

Definition 7 (Amplification with nonuniform security reductions)**.** *Let $N > n$, $M > m$, and $0 < \gamma, \epsilon \leq 1$. A (q, γ, ϵ, d) amplification reduction with nonuniform security from $n \mapsto N$ stretch to $m \mapsto M$ stretch is a pair (H, R) such that $H^{(\cdot)} : \{0,1\}^m \to \{0,1\}^M$ is a (q, d)-construction from inner function $G : \{0,1\}^n \to \{0,1\}^N$, and for each advice string $a \in \{0,1\}^q$, $R_a^{G,T}$ is a randomized query procedure outputting a value in $\{0,1\}$ and querying oracles $G : \{0,1\}^n \to \{0,1\}^N$, $T : \{0,1\}^M \to \{0,1\}$, such that for every possible G, T we have:*

$$\mathrm{Adv}_{H^G}(T) \geq \gamma \quad \Longrightarrow \quad \exists a \in \{0,1\}^q,\ \mathrm{Adv}_G(R_a^{G,T}) \geq \epsilon$$

On every input and for every choice of its randomness and advice a, R_a makes at most q queries in total to G and T.

Theorem 7. *There is a universal constant $c > 0$ such that the following holds. Let (H, R) be a $(q, 0.9, \epsilon, d)$ amplification reduction with nonuniform security from $n \mapsto N$ stretch to $m \mapsto M$ stretch, where $q \geq 2$, $d \geq 1$, $0 < \epsilon \leq 1$, and $q, \epsilon^{-1} \leq 2^{cn}$. Then*

$$M \leq \big(m + O(\log(qd))\big) \left(\frac{N}{n - O(\log(q/\epsilon))}\right)^d.$$

Proof. We proceed as in the proof of Theorem 1: for some parameters $h = \mathrm{poly}(q\epsilon^{-1})$, $s = \mathrm{poly}(hq)$ to be chosen later, set $p = s/2^n$, sample G, S in the same way as in Theorem 1, and set $\rho = G|_{\overline{S}}$, $T = T_\rho$. The application of Lemma 1 establishing that T distinguishes H^G with high probability is unchanged; in particular, if $M \geq (m + C \log qd)(\frac{N}{n - \log qs})^d$ for a sufficiently large universal constant C, then with probability at least $\frac{1}{2}$ T will distinguish H^G with advantage ≥ 0.9. We will not apply the reduction to nonnegative signed advantage (Lemma 3) used in the proof of Theorem 1, and instead bound the absolute advantage directly. We will show that for every choice of $a \in \{0,1\}^q$, we have

$$\Pr_{G,T}[\mathrm{Adv}_G(R_a^{G,T_\rho}) \geq \epsilon] \leq O(e^{-\Omega(h\epsilon^2/q)}) \tag{$\dagger$}$$

Taking a union bound over $a \in \{0,1\}^q$:

$$\Pr_{G,T}[\exists a \in \{0,1\}^q,\ \mathrm{Adv}_G(R_a^{G,T_\rho}) \geq \epsilon] \leq 2^q O(e^{-\Omega(h\epsilon^2/q)})$$

For a sufficiently large choice of $h = O(q^2\epsilon^{-2})$ this will give the desired result. We hence suppress the advice a for the remainder of the proof and establish ($\dagger$) for any fixed R (with the understanding that $R = R_a$ for some fixing of a).

To prove ($\dagger$) we will modify the planted experiment from the proof of Theorem 1 to plant h values instead of 1. We introduce two experiments which we call the "subsample advantage experiment" and the "planted experiment" and show that various relevant parameters in these experiments have an identical distribution. We start with the subsample experiment. Sample G, S as above,

sample $w^1, \ldots, w^h \sim \{0,1\}^n$ uniformly *without replacement*, $u^1, \ldots, u^h \sim \{0,1\}^N$ uniformly with replacement, and a uniform random sign vector $b \sim \{\pm 1\}^h$. We then set $T = T_{G|_{\overline{S}}}$ and define:

$$z^i = \begin{cases} G(w^i) \text{ if } b_i = 1 \\ u^i \text{ if } b_i = -1 \end{cases}, \qquad r_i = 2 \cdot R^{G,T}(z^i) - 1, \qquad \nu = \frac{1}{h}\sum_{i=1}^{h} b_i \cdot r_i$$

where in the definition of $r_1, \ldots, r_h$, R uses fresh independent internal randomness across these h executions. Let $\mathcal{D}_1$ be the distribution of $(G, S, T, \bar{w}, \bar{z}, b, r, \nu)$ generated in the process.

We next describe a distribution $\mathcal{D}_2$ defining the "planted experiment." Sample G, S as before, sample $w^1, \ldots, w^h \sim \{0,1\}^n$ without replacement, $z^1, \ldots, z^h \sim \{0,1\}^N$ with replacement, and $b \sim \{\pm 1\}^h$ uniformly. Define

$$G' = G[w^i \mapsto z^i \mid \; b_i = 1], \qquad T' = T_{G'|_{\overline{S}}} \qquad r_i = 2 \cdot R^{G',T'}(z^i) - 1, \qquad \nu = \frac{1}{h}\sum_{i=1}^{h} b_i \cdot r_i$$

where as above we use fresh internal R-randomness in the definition of each r_i. We take $\mathcal{D}_2$ to be the distribution of $(G', S, T', \bar{w}, \bar{z}, b, r, \nu)$ generated in this experiment. We claim that in fact $\mathcal{D}_1 = \mathcal{D}_2$. This follows by observing that the distributions of $\bar{w}, b$ are identical in both, and conditioned on $\bar{w}, b$, G' is a uniformly random function and z^i equals $G'(w^i)$ when $b_i = 1$, and is otherwise an independent uniform element of $\{0,1\}^N$.

Now, we want to argue that $\mathrm{Adv}_G(R^{G,T})$ is small with high probability over G, T. We will first argue, using the $\mathcal{D}_1$ experiment, that with overwhelming probability $|\mathrm{SAdv}_G(R^{G,T}) - \nu|$ is very small, so that it suffices to bound $|\nu|$. We then argue using the $\mathcal{D}_2$ experiment that with overwhelming probability $|\nu|$ is small, which will complete the proof.

For the first step we will argue that $|\mathrm{SAdv}_G(R^{G,T}) - \nu|$ is small for every conditioning of G, T. Fix some G, T for the remainder of this argument, and define

$$a(w) = \Pr_R[R^{G,T}(G(w)) = 1], \quad \alpha = \mathbb{E}_{w \sim \{0,1\}^n} a(w), \quad \beta = \Pr_{R, z \sim \{0,1\}^N}[R^{G,T}(z) = 1], \quad \Delta = \alpha - \beta$$

Observe that $\Delta = \mathrm{SAdv}_G(R^{G,T})$. Now, if we condition as well on $\bar{w}$, we see that the summands $b_i r_i$ are independent with means $a(w^i) - \beta$. By Hoeffding their average ν concentrates around $\widehat{\Delta}(\bar{w}) := \frac{1}{h}\sum_{i \le h}(a(w^i) - \beta)$; in particular $\Pr\left[|\nu - \widehat{\Delta}(\bar{w})| > \frac{\epsilon}{8} \mid G, T, \bar{w}\right] \le 2e^{-h\epsilon^2/128}$. Next we observe that:

$$\widehat{\Delta}(\bar{w}) - \Delta = \frac{1}{h}\Big(\sum_{i \le h}(a(w^i) - \beta)\Big) - (\alpha - \beta) = \Big(\frac{1}{h}\sum_{i \le h} a(w^i)\Big) - \mathbb{E}_{w \sim \{0,1\}^n} a(w)$$

We can hence control the deviation of $\widehat{\Delta}(\bar{w})$ from Δ using the sampling-without-replacement Hoeffding bound: we are sampling h uniform inputs in $\{0,1\}^n$ without replacement, and asking for the deviation of the sample mean of $a(\cdot) \in [0,1]$ from its population mean. Applying this bound gives $\Pr_{\bar{w}}\left[|\widehat{\Delta}(\bar{w}) - \Delta| > \frac{\epsilon}{8} \mid G, T\right] \le 2e^{-h\epsilon^2/32}$. The triangle inequality and a union bound then give $\Pr[|\nu - \Delta| > \frac{\epsilon}{4} \mid G, T] \le 4e^{-h\epsilon^2/128}$.

We now show, using the $\mathcal{D}_2$ experiment, that $\Pr[|\nu| \ge \epsilon/4] \le 2^{-\Omega(h\epsilon^2/q)}$ which will complete the proof. In particular we will show that this holds for an appropriate setting of s bounded by $s = O(h^3 q^2)$. In the planted experiment, we interpret R as attempting to "guess" the value of the signs b_i by making queries to G', T'. We measure R's success by evaluating it at each z^i and subtracting the fraction of incorrect answers from the fraction of correct answers to compute its overall "score" ν. To argue the score is small, we define a set $D \subseteq [h]$, jointly distributed with the

other parameters of $\mathcal{D}_2$, indicating which of the w^i will get "discovered" by one of the executions of R defining $r_1, \ldots, r_h$ if we were to run the executions in series (the definition of "discovered" will be more general than strictly requiring G' to be queried on w^i). We will argue that D is small with high probability, and conditioned on D being small the score ν must also be small with high probability.

Initialize $D = \emptyset$; for the current value of D, define $S_D = S \cup \{w^j \mid j \notin D\}$, $\rho_D = G'|_{\overline{S_D}}$, $T_D = T_{\rho_D}$. Let $\rho' = G'|_{\overline{S}}$; so $T' = T_{\rho'}$. Observe that for any D, $\rho_D \subseteq \rho'$, hence for all y $T'(y) \leq T_D(y)$. Now, iterate over each $i \leq h$ in order and add elements to D by considering the execution of $R^{G',T'}(z^i)$. If the next query is to $G'(w^j)$ for some j, we perform the query and add j to the discovered set D if it was not already present. If the next query is to $T'(y)$ for some y we perform this query and then initiate the following loop to expand D before moving on to R's next query. If $T_D(y) = 0$ then we do not modify D and continue to R's next query immediately. If $T_D(y) = 1$, we fix some canonical τ (defined in the next sentence) and for each j such that $\tau(w^j) \neq *$, add j into D. By "canonical witness" we mean: τ is compatible with ρ_D, $|\tau| \leq q$, there exists $x \in \{0,1\}^m$ such that $H^F(x) = y$ for all $F \supseteq \tau$, and amongst all such partial assignments, τ is the lexicographically first in some standard ordering. Now, if D did not grow at all during this step, we proceed to R's next query. Otherwise we repeat this process: for the expanded set D, if $T_D(y) = 0$ then we halt and proceed in R's execution, otherwise we expand D again using the canonical witness for $T_D(y) = 1$. Each repeated step here strictly increases $|D|$ so eventually this process terminates and we move to R's next query.

We make two key claims about this process for generating D. We first argue that if D_y is the state of D after $T'(y)$ was queried and the repeated loop terminated thereafter, then we must have $T_{D_y}(y) = T'(y)$. We already know that $T'(y) \leq T_{D_y}(y)$, so we only need to prove that $T_{D_y}(y) \leq T'(y)$. This follows since the canonical witness τ for $T_{D_y}(y) = 1$ satisfies $\tau(w^i) = *$ for every $i \notin D_y$, and therefore is compatible with ρ'; hence there exists some $F \supseteq \rho', F \supseteq \tau$ and some x such that $H^F(x) = y$, so $T'(y) = 1$. Using this we next claim that, conditioned on the entire transcript of the execution defining D (including the internal randomness that R uses in each of the h executions), we have that $(b_i)_{i \notin D}$ are independent uniform signs. It suffices to show that after fixing the R-randomness, we can exactly determine the entire execution transcript from $(G, S, \bar{w}, \bar{z}, (b_i)_{i \in D})$. To do this we modify the execution. When we query $G'(w)$, if $w \neq w^i$ for any i then we can answer according to $G(w)$ which is known; otherwise if $G'(w^i)$ is queried, we must have $i \in D$, so we can determine $G'(w^i)$ from G, w^i, z^i, b_i. If on the other hand we are about to query T' on y we instead run the loop extending the previous D_{prev} to D_y directly until it terminates, and then proceed as if the query to T' took the value $T_{D_y}(y)$. We know $T_{D_y}(y) = T'(y)$, so by induction we have reached the same state in the R simulation. To perform the loop extending D_{prev} to D_y, we need only inspect D_{prev} (which by induction we can obtain from $G, S, \bar{w}, \bar{z}, (b_i)_{i \in D}$ and the R-randomness), together with $G, S, \bar{w}, \bar{z}$, the R-randomness, and the bits b_i for each $i \in D_y \setminus D_{\text{prev}}$.

We now show that, for an appropriate choice of s satisfying $s = O(h^3 q^2)$, (1) for all $t \geq 0$, $\Pr[|D| > t] \leq 2^{-t/q} + e^{-s/8}$, and (2) for all $0 < \delta \leq 1$, $\Pr[|\nu| \geq \delta \mid |D| \leq \frac{\delta}{2} h] \leq 2e^{-h\delta^2/8}$. Together these give us $\Pr[|\nu| \geq \frac{\epsilon}{4}] \leq 2^{-h\epsilon/(8q)} + e^{-s/8} + 2e^{-h\epsilon^2/128} = O(e^{-\Omega(h\epsilon^2/q)})$ and we are done. We start with (2) which is easy. Condition on any complete transcript of the executions performed above in the construction of D with $|D| \leq \frac{\delta}{2} h$ (including the internal randomness of each R execution and the discovered indices and signs); for each $i \notin D$ r_i is fixed, and $(b_i)_{i \notin D}$ remain uniform independent signs, hence $(b_i r_i)_{i \notin D}$ remain uniform independent signs. Overall

$$|\nu| = \left| \frac{1}{h} \sum_{i \in D} b_i r_i + \frac{1}{h} \sum_{i \notin D} b_i r_i \right| \leq \frac{\delta}{2} + \left| \frac{1}{h} \sum_{i \notin D} b_i r_i \right|$$

By Hoeffding, $|\frac{1}{h} \sum_{i \notin D} b_i r_i| < \delta/2$ except with probability $2e^{-h\delta^2/8}$.

We now prove (1). By another application of Hoeffding we know that $\Pr[|S| < \frac{s}{2}] \leq e^{-s/8}$; we prove that $\Pr[|D| > tq \mid |S| \geq \frac{s}{2}] \leq 2^{-t}$ which gives (1). For some set $U \subseteq \{0,1\}^n$, $|U| \geq \frac{s}{2}$, condition on the event $S \cup \{w^1, \ldots, w^h\} = U$, as well as on a particular value of $G, \bar{z}, b$, the internal R-randomness of each of the h executions, and values $\eta_i = \mathbf{1}\{w^i \in S\}, i \leq h$. After such conditioning the only remaining randomness in the execution sequence lies in the values of $\bar{w}$, whose conditional distribution is that of a uniform without-replacement sample from U (the set S is completely determined by $\bar{w}$ at this point, and equals $U \setminus \{w^i \mid \eta_i = 0\}$).

We now consider running the execution of all the $R(z^i)$ in the second form described above, where we only expose the information about G', T' necessary to make the next query; this was the formulation we used to prove the claim that $(b_i)_{i \notin D}$ are uniform conditioned on the partial execution. During this execution, define a "basic step" to be either one direct query to G', or else one iteration of the loop associated with a query to T'. We argue that conditioned on any partial execution with current set D_0 (pausing the execution immediately before some basic step), the probability that $|D_{\text{final}}| > |D_0|$ is at most $\frac{1}{2}$ (where D_{final} is the final state of D after all executions are complete). This gives the recurrence $\Pr[|D| > (t+1)q] \leq \frac{1}{2}\Pr[|D| > tq]$ and hence $\Pr[|D| > tq] \leq 2^{-t}$, since each time D grows in a given basic step, it grows by at most q. It suffices to argue that conditioned on any partial execution and any D_0 for that stage, the probability that any particular basic step in the remaining execution discovers a new input is at most $\frac{1}{2h(h+1)q}$; we then take a union bound over all $\leq h(h+1)q$ basic steps.

Let such a partial execution and D_0 be fixed, and consider the remaining execution. We will maintain a set $W \subseteq \{0,1\}^n$ of inspected inputs of size at most $h(h+1)q^2$, so that if $D \supseteq D_0$ is the current set of discovered indices, then $\{w^i \mid i \notin D\}$ is distributed as a uniform subset of $U \setminus W$ of size $h - |D|$. If the basic step involves a query to $G'(w)$ we add w into the set W. Before we add it, the probability that $w \in \{w^i \mid i \notin D\}$ is at most $\frac{h}{|U \setminus W|}$, and otherwise D does not grow. Since $|U \setminus W| \geq \frac{s}{2} - h(h+1)q^2$, this probability is bounded by $\frac{1}{2h(h+1)q}$ for a sufficiently large choice of $s = O(h^3q^2)$. The relevant property on W is maintained. If the next basic step is during a loop involving a query $T'(y)$, the probability that D strictly grows is bounded by the probability that some set of size $\leq q$ (the inputs in $\{0,1\}^n$ which are assigned values by the canonical witness) intersects $\{w^i \mid i \notin D\}$. Since the latter is a random subset of $U \setminus W$ of size $h - |D| \leq h$ and $|U| - |W| \geq \frac{s}{2} - h(h+1)q^2$ this probability is at most $\frac{1}{2h(h+1)q}$ for a sufficiently large choice of $s = O(h^3q^2)$. We then add all inputs assigned by the canonical witness in each step of the loop to W, which maintains the stated property on W. Over the course of this entire process we add at most $(h+1)q$ inputs to W whenever we simulate a single query of some $R^{G',T'}(z^i)$ (since the loop can run for at most $(h+1)$ rounds and we add at most q inputs to W at each step of the loop), so over the course of all $\leq hq$ query steps simulated the size of W is never more than $h(h+1)q^2$ as promised.

□

B Seed-Length Bounds

Lemma 5 (Seed-length bounds)**.** *Let $q \geq 1$. The following bounds hold for any amplification from stretch $n \mapsto N$ to stretch $m \mapsto M$, with $M \geq m + 4$:*

(i) If (H, R) is a $(q, 0.9, \epsilon, d)$ PRG amplification reduction and $q, \epsilon^{-1} < 2^{n/100}$ then $m \geq \frac{n}{4}$.

(ii) If (H, Solve) is a (q, d) amplification reduction for Avoid *and $q < 2^{n/4}$ then $m \geq \frac{n}{4}$*

Proof sketch. We sketch the argument for PRGs. If $m < \frac{n}{4}$, then the reduction implies the existence of a $q' = q^2 2^{\frac{n}{4}}$ query algorithm $\widetilde{R}^{(\cdot)}$ querying a single oracle, such that for all G $\text{Adv}_G(\widetilde{R}^G) \geq \epsilon$.

This follows since we can plug in the particular choice of T to the amplification reduction which tests each $y \in \{0,1\}^M$ for membership in range(H^G); this test T can be simulated with $q2^m < q2^{\frac{n}{4}}$ queries to G directly. Ruling out the existence of such $\widetilde{R}$ is now a special case of the argument in the latter half of the proof of Theorem 1; our task is strictly easier since we no longer have to deal with the test T. The case of Avoid is handled similarly, applying a brute force search to find a string outside the range of H^G and then applying Lemma 2. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il