

Constant-error pseudorandomness proofs from hardness require majority

Emanuele Viola*

July 26, 2018

Abstract

Research in the 80's and 90's showed how to construct a pseudorandom generator from a function that is hard to compute on more than 99% of the inputs. A more recent line of works showed however that if the generator has small error, then the proof of correctness cannot be implemented in subclasses of TC^0 , and hence the construction cannot be applied to the known hardness results. This paper considers a typical class of pseudorandom generator constructions, and proves an analogous result for the case of large error.

The construction of *pseudorandom generators* from *hard functions* is a fundamental paradigm in complexity theory. Call a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ δ -hard if every “small circuit” fails to compute f on at least a δ fraction of the inputs. Perhaps the simplest example of the paradigm is this. Given f which is $(1/2 - \epsilon)$ -hard the *repetition* pseudorandom generator

$$PRG : \{0, 1\}^{\ell \cdot s} \rightarrow \{0, 1\}^{\ell \cdot s + s}$$

$$PRG(x_1, x_2, \dots, x_s) := (x_1, x_2, \dots, x_s) \circ f(x_1) \circ f(x_2) \circ \dots \circ f(x_s)$$

has error $O(\epsilon s)$, by which we mean that small circuits cannot distinguish its output distribution from uniform with advantage more than $O(\epsilon s)$. The proof relies on the *hybrid argument*, see for example [FSUV13] and the discussion there. Hence, from hardness $1/2 - \epsilon$ we can obtain $\Omega(1/\epsilon)$ bits of pseudorandomness, and this is tight for black-box reductions [FSUV13].

This repetition approach has poor output/input ratio less than two. The landmark paper by Nisan [Nis91] gave a better construction using a *design* $S_1, S_2, \dots, S_s \subseteq \{1, 2, \dots, m\}$. Here the generator gets an input σ and computes from it s inputs x_1, x_2, \dots, x_s for f where x_i is the bits of σ indexed by S_i .

$$N : \{0, 1\}^m \rightarrow \{0, 1\}^{m+s}$$

$$N(\sigma) := \sigma \circ f(x_1) \circ f(x_2) \circ \dots \circ f(x_s).$$

*Supported by NSF CCF award 1813930.

The repetition generator can be seen as a trivial design where the S_i are disjoint. Nisan allows us to take s much larger than m , and so the stretch is essentially unaffected if we drop σ from the output. Jumping ahead, our result only applies if σ is present in the output: the case where σ is not revealed is an interesting open problem.

Starting from less hard f . If the hardness of f is less, say if it is a constant bounded away from $1/2$, then the above approach gives only a constant number of bits of pseudorandomness. Since at least [NW94] the way around this has been to *amplify the hardness* of f to obtain another function f' with hardness $1/2 - \epsilon$ where $\epsilon \leq 1/s$, and then apply a generator construction to f' . The simplest and most well-known hardness amplification is the XOR lemma originating from oral presentations by Yao in the 80's, see [GNW95]. Here f' is simply the XOR of k independent copies of f :

$$f' : \{0, 1\}^{\ell \cdot k} \rightarrow \{0, 1\}$$

$$f'(x_1, x_2, \dots, x_k) := \bigoplus_{j \leq k} f(x_j).$$

For a concrete example, if we combine this hardness amplification with the repetition generator we obtain the following construction where $m = \ell ks$:

$$G : \{0, 1\}^m \rightarrow \{0, 1\}^{m+s}$$

$$G(x_{ij})_{i \leq s, j \leq k} := (x_{ij})_{i \leq s, j \leq k} \circ \bigoplus_{j \leq k} f(x_{1j}) \circ \bigoplus_{j \leq k} f(x_{2j}) \circ \dots \circ \bigoplus_{j \leq k} f(x_{sj}). \quad (1)$$

If we start with f that has constant hardness then in Yao's XOR lemma for any given s we can take $k = O(\log s)$ to obtain hardness say $1/2 - 1/s^2$ which is sufficient to show that G has error $O(1/s)$.

The extension of Nisan's generator where we first apply Yao's XOR lemma is known as the Nisan-Wigderson generator [NW94]. It is as follows

$$NW : \{0, 1\}^m \rightarrow \{0, 1\}^{m+s}$$

$$NW(\sigma) := \sigma \circ \bigoplus_{j \leq k} f(x_{1j}) \circ \bigoplus_{j \leq k} f(x_{2j}) \circ \dots \circ \bigoplus_{j \leq k} f(x_{sj}), \quad (2)$$

where now all the x_{ij} are computed from σ .

Hardness amplification proofs require majority. This hardness amplification step is problematic for *restricted computational models*. This is because most or all reductions that prove hardness amplification are *black-box*, and a line of research [Vio04, Vio06, LTW11, SV10, GR08, AS14, GSV18] has shown that these reductions cannot be implemented in a class less powerful than TC^0 . The sense in which this holds is made precise below, but basically the reduction circuits that amplify hardness to $1/2 - \epsilon$ can be used to compute majority on $\Omega(1/\epsilon)$ bits. This means that when ϵ is small enough these reductions can only be used when starting from a lower bound against at least TC^0 , a long-standing open problem which is believed to require substantially new techniques [RR97, NR04]. In particular, they cannot be used for several classes for which we do have δ -hard functions for $\delta = 0.1$ but not

$\delta = 1/2 - 1/\sqrt{\ell}$ for functions on ℓ bits. Such classes include AC^0 circuits with parity gates [Raz87, Smo87] or with one majority gate [ABFR94]. We refer to [GSV18] for additional discussion and pointers.

Moreover, most proofs of pseudorandom generators $PRG : \{0, 1\}^m \rightarrow \{0, 1\}^{m+s}$ also yield error about $1/s$. For example in the above example of G if we amplify hardness to $1/s^2$ we obtain error $O(1/s)$. And the same limitations apply to a pseudorandomness proof with error ϵ as to a hardness amplification proof to hardness $1/2 - \epsilon$ [GSV18].

Constant-error pseudorandomness. However such limitations were not known for pseudorandomness proofs with *constant error*. Pseudorandom generators with constant error are important in their own right. For example they suffice for derandomization, and are not known for several frontier classes such as AC^0 with parity gates. In this paper we show that even the proofs of the constructions above *with constant error* require TC^0 . For s bits of stretch we obtain the same limitations established in [GSV18] for amplifying hardness to $1/2 - 1/s$. We state a more general and abstract result, where the generator computes inputs x_{ij} from the seed σ via a function L , then evaluates f at those inputs and combines the output via a function H . Following [Vio06, SV10] we show that the proof of correctness can be used to distinguish uniform bits from slightly biased bits. We denote by $N_{1/2-\epsilon}^q$ a q -tuple of i.i.d. bits coming up 1 with probability $1/2 - \epsilon$.

Theorem 1. *There is a constant c such that if $\max\{a, q, 1/\delta, g, k\} \leq 2^{\ell/c}$ and $\delta \leq 1/3$ the following holds.*

Hypothesis: Let $G^f(\sigma) = \sigma \circ G_1^f(\sigma) \circ G_1^f(\sigma) \circ \dots \circ G_s^f(\sigma)$ be a pseudorandom-generator construction mapping a seed $\sigma \in \{0, 1\}^m$ and a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ to $m + s$ output bits defined by

$$G_i^f(\sigma) = H(f(x_{i1}), f(x_{i2}), \dots, f(x_{ik}))$$

where $H : \{0, 1\}^k \rightarrow \{0, 1\}$ is a function, and the $s \cdot k$ values $x_{ij}, i \leq s, j \leq k$ are obtained from σ by the map $L : \{0, 1\}^m \rightarrow \{0, 1\}^{\ell \cdot s \cdot k}$. Suppose that with probability $1 - 1/1000$ over σ , the x_{ij} values $L(\sigma)$ are distinct.

Let $\{C(\cdot, \alpha)\}_{\alpha \in \{0, 1\}^a}$ be a family of 2^a oracle circuits such that for every α the circuit $C(\cdot, \alpha)$ has size g , depth κ , makes q oracle queries, and uses gates $\Xi \supseteq \{\text{And}, \text{Or}, \text{Not}\}$.

Suppose that for every $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and $d : \{0, 1\}^{m+s} \rightarrow \{0, 1\}$ such that

$$|\mathbb{P}[d(G^f(\sigma)) = 1] - \mathbb{P}[d(U_{m+s}) = 1]| \geq \frac{1}{100},$$

where σ and U_{m+s} are uniform in $\{0, 1\}^m$ and $\{0, 1\}^{m+s}$ respectively, there exists $\alpha \in \{0, 1\}^a$ such that

$$\mathbb{P}_{x \in \{0, 1\}^\ell}[C^d(x, \alpha) = f(x)] \geq 1 - \delta.$$

Conclusion: Then for every $\epsilon = \Omega(1/s)$ there is a distribution T on circuits $t : \{0, 1\}^q \rightarrow \{0, 1\}$ that distinguish $N_{1/2}^q$ from $N_{1/2-\epsilon}^q$ with advantage $1 - O(\delta)$, and have depth $O(\kappa)$ and size $\text{poly}(g)$. The circuits t have gates Ξ , and in addition have gates for H and L .

The rationale behind the hypothesis is that if f is hard for small circuits then d can't be a small circuit, and so the output of G is pseudorandom.

By the arguments in sections 5 and 6 of [SV10], the conclusion of the theorem implies

(A) $q \geq \Omega(\log(1/\delta)/s^2)$, and

(B) there is a circuit with depth $O(\kappa)$, size $\text{poly}(g)$, and the same gates as T that computes majority on inputs of length $\Omega(s)$. In particular if H, L, C have depth $O(1)$ then the size of C must be exponential in $s^{\Omega(1)}$. By contrast, for unbounded-depth circuits the size can be polynomial.

The function L is the identity map in the repetition PRG, and is a projection in NW. In both cases it is trivially computable in AC^0 . The function H is simply parity on k bits in both cases.

The theorem assumes that with probability $1 - 1/1000$ over σ the x_{ij} are distinct. We now remark that for typical settings of both generators (1) and (2) this assumption holds. When each x_{ij} is the projection of σ on a set of ℓ bits, as long as the pairwise intersection of the sets is $\leq \ell/2$ then the probability that two x_{ij} will be equal is $\leq \binom{ks}{2} 2^{-\ell/2}$ by a union bound. In particular if $k, s = \ell^{O(1)}$ the assumption is satisfied for large enough ℓ . (In a typical Nisan-Wigderson design, the bound on the intersection is in fact much smaller.)

The corollaries (A) and (B) are the same obtained in [GSV18] for amplification to hardness $1/2 - 1/s$. As explained there, the bounds are tight for hardness amplification. Specifically, the bound on the number of queries matches [KS03], and the result about majority is tight because there exist proofs of Yao's XOR lemma which can be implemented in AC^0 with one majority gate [Kli01] (for a simplification of [Kli01], due to Klivans and Vadhan, see [Vio09]). We note that such tightness extends to the generators (1) and (2), because the pseudorandomness proof from a $1/2 - 1/s$ hard function can be implemented in AC^0 with only one query [Vio07].

Note that for constructions with small error ϵ [GSV18] rules out AC^0 reductions even for one bit of stretch, that is $s = 1$. In our setting of constant error ϵ we can't do that because it is possible to get $s = \text{poly log } g$ bits of stretch via an AC^0 reduction of size g , and our result shows that this is tight.

Theorem 1 would immediately follow from [GSV18] if there was a way to turn a constant-error generator into a $1/2 - 1/s$ hard function without using majority. No such way is known, and in fact we suspect one can derive a black-box separation for this task along the lines of [GSV18].

Open problems and discussion

Problem 2. Prove Theorem 1 without any restriction on k . Prove Theorem 1 when the seed σ is not part of the output (for $s \gg m$).

Fixed F . Our proof is black-box in both the use of the distinguisher d and of the hard function f . The result is actually false for specific functions f ! [FSUV13] showed with an AC^0 reduction that the repetition generator is pseudorandom *with no error loss* when

starting with any resamplable function, such as Parity. Using this they constructed the best-known generators for classes such as AC^0 with modular gates. However it is not clear how to push their techniques to Nisan-style pseudorandom generators with much better stretch. A natural proof strategy would require sampling the output distribution of the generator, which is impossible [LV12].

So it would be very interesting to understand if the techniques in this paper can be extended to handle a fixed f like Parity. One difficulty is that our techniques do not work unless one further exploits that the reduction's queries are a low-complexity function of the input x . In a nutshell, this is because a reduction on input x could make for example the query $(x_1 0 x_2 0 \dots x_t 0)$ where each x_i has the same parity as x . Then a distinguisher such as those we define later would allow to compute the parity of x . Interestingly, this is exactly how the proof in [FSUV13] works. Intuitively we cannot do this for a uniform function f , and our proof formalizes the intuition.

Because Nisan's generator applied to the parity function is just the uniform distribution over a vector space, we can ask:

Problem 3. Understand the pseudorandomness properties of vector spaces $V \subseteq \{0, 1\}^n$ with bit-wise addition modulo 2. In particular, for how small m there exists a space of dimension m that fools AC^0 with mod 3 gates? What about a single mod 3 gate? What can be proved using reductions that are black-box in the use of the distinguisher?

1 Proof discussion

We explain the basic idea of the proof in the case of the repetition generator (1). Take F to be a random function $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and the distinguisher d as a distribution D defined as follows. Recall that the input to D is a string z of length $m + s$ which consists of sk inputs x_{ij} to F and s additional bits b_i . On input z the distinguisher D selects a uniform pointer $P(z) \in \{1, 2, \dots, s\}$ and then answers one if and only if the bit b_i matches the output of the pseudorandom generator, that is, if $b_i = \bigoplus_{j \leq k} F(x_{ij})$. Such a D breaks the pseudorandom generator, as it always outputs 1 when the input is pseudorandom, but if b_i is uniform then the equation $b_i = \bigoplus_{j \leq k} F(x_{ij})$ is satisfied with probability $1/2$.

Now imagine, for some fixed α , a reduction circuit $C^D(x^*, \alpha)$ that is trying to compute $F(x^*)$ by querying D . Equivalently, the circuit is trying to distinguish $F_0 := F|F(x^*) = 0$ from $F_1 := F|F(x^*) = 1$. We can assume that in any query z to D the inputs x_{ij} are all different (by instructing D to answer 0 otherwise). Then the answers to each query in F_0 and F_1 are close: they are random bits whose statistical distance is $O(1/s)$. When s is large enough, this difference is too small to be detected by a constant-depth circuit.

Formalizing this idea presents several difficulties. First, the above definition of D actually does not work. This is because, using advice, the reduction could ask queries where all the values $F(x_{ij})$ are 0 except possibly $F(x^*)$. In such a situation the answers of D would be always 0 under F_0 , and non-zero with probability $\geq 1/s$ under F_1 . This *can* be detected in constant-depth (by a simple Or). The fix is to pad the bits b_i with a balanced string, so that the answer bits always come up 1 with probability bounded away from 0 and 1.

To handle advice, we use the *fixed-set lemma* from [GSV18], which basically says that conditioning on a specific α can be thought of as fixing some values of F and D . As in [GSV18], this lemma appears to greatly simplify the argument.

We note that although we use this lemma from [GSV18], and other results from [SV10], our proof presents several differences. The main one is in the choice of the oracle. In previous works this choice is straightforward: the oracle is simply obtained by perturbing the encoding of f with suitable noise, and the main focus is on the analysis. In our case the choice is less obvious. This difference then propagates to several other parts of the proof, for example in how we handle adaptive queries.

2 Formal proof

Let $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ be a uniform function. We define the distinguisher D as follows. Let $P : \{0, 1\}^{m+s} \rightarrow [3s]$ be a uniform function. On input $z = \sigma \circ b_1 \circ \dots \circ b_s \in \{0, 1\}^{m+s}$, compute $L(\sigma) = \{x_{ij}\}_{i \leq s, j \leq k}$. For simplicity we think of the output of $L(\sigma)$ as a multiset. If the x_{ij} are not all distinct then output zero. Otherwise compute the string $v(z) \in \{0, 1\}^s$ where the bit i is the indicator of whether b_i is the correct output of the hard function, i.e., it is the indicator of $H(F(x_{i1}), F(x_{i2}), \dots, F(x_{ik})) = b_i$. Return the bit $P(z)$ of the string $v(z)0^s1^s \in \{0, 1\}^{3s}$.

Claim 4. For every F , with probability $\Omega(1)$ over the choice of P , D distinguishes the PRG from uniform with advantage $\Omega(1)$.

Proof. We first prove the claim ignoring the fact that D outputs 0 if the x_{ij} are not all distinct.

On any input z from the PRG, D outputs 1 unless $P(z)$ lands on 0^s . The latter happens with probability $1/3$ independently. Hence with probability say $1 - 1/100$ over P , assuming s is large enough, D outputs 1 with probability $\geq 2/3 - 1/100$ over a uniform output of the PRG. Note we can assume that s is large enough for else the conclusion of the theorem is trivial.

We now analyze the behavior for a uniform input u . First note that at least say a 0.99 fraction of the strings $u \in \{0, 1\}^{m+s}$ are such that $v(u)$ has Hamming weight $\leq 0.51s$, again using that s is large enough. (Note that this statement only depends on F , which is fixed.) For any such string u , the distinguisher outputs 1 with probability $\leq (1 + 0.51)/3$ over the choice of $P(u)$. Hence with probability $1 - 1/100$ over P the distinguisher outputs 1 on at most a $(1+0.52)/3$ fraction of those strings, and hence overall on at most a $1/100 + (1+0.52)/3$ fraction of the strings in $\{0, 1\}^{m+s}$.

Hence with probability $1 - 1/100$ over P , D has an advantage of

$$\geq (2/3 - 1/100) - (1/100 + (1 + 0.52)/3) = 0.14.$$

Finally we take into account the probability that the x_{ij} are not all distinct. This happens with probability $\leq 1/1000$ by assumption. Hence we still have advantage $\geq 1/100$. \square

By the above claim, averaging over F , and our assumption, there is an advice $\alpha \in \{0, 1\}^a$ such that with probability $\geq \Omega(2^{-a})$ over F and P the event $A := “\mathbb{P}_{x \in \{0, 1\}^\ell}[C^D(x, \alpha) = F(x)] \geq 1 - \delta”$ happens. We now use the *fixed-set lemma* from [GSV18], restated next.

Lemma 5. *Let N, a, q' be integers. Let $Y = (Y_1, \dots, Y_N)$ be independent random variables, each uniform over some finite set Σ . Let $A \subseteq \Sigma^N$ be an event such that $\mathbb{P}[Y \in A] \geq 2^{-a}$, and let $X = (Y|Y \in A)$. For every $\eta > 0$, there exists a set $B \subseteq [N]$ of size $\text{poly}(a, q', \eta^{-1})$, and $v \in \{0, 1\}^B$ such that for $Y' := (Y|Y_B = v)$ and $X' := (X|X_B = v) = (Y|Y_B = v, Y \in A)$, and every q' -query decision tree t , $|\mathbb{P}[t(Y') = 1] - \mathbb{P}[t(X') = 1]| \leq \eta$.*

We apply the lemma to the $2^\ell + 2^{m+s}$ random variables $F(x), P(z)$. We set $q' = O(qk)$ and $\eta = 1/2^{\ell/c}$.

Remark 6. The fixed-set lemma is only stated for variables with the same range Σ , whereas $F(x)$ and $P(z)$ have different range. However we can for example rewrite $F(x)$ and $P(z)$ as fixed functions of random variables with the same range Σ (for example $|\Sigma| = 2 \cdot 3s$ will do). So we can apply it in our setting as well.

The lemma gives a set B of size $\text{poly}(a, q', 1/\eta) = \text{poly}(2^{\ell/c})$ and a string v such that for every x

$$|\mathbb{P}_{F,P}[C^D(x, \alpha) = F(x)|A, (FP)_B = v] - \mathbb{P}_{F,P}[C^D(x, \alpha) = F(x)|(FP)_B = v]| \leq \eta, \quad (3)$$

where note in the second probability there is no conditioning on A . Here we are using that checking if $C^D(x, \alpha) = F(x)$ can be computed by a decision tree with access to the variables $F(x)$ and $P(z)$ making $\leq 1 + q + qk = O(qk)$ queries. The tree first queries $F(x)$, then it simulates $C^D(x, \alpha)$, answering each oracle query $D(z)$ by querying $P(z)$ and then the k corresponding values of F .

Recall that for every F, P such that A holds we have by definition of A that

$$\mathbb{P}_x[C^D(x, \alpha) = F(x)] \geq 1 - \delta.$$

Hence the same holds if we average over F, P and further condition:

$$\mathbb{P}_{x,F,P}[C^D(x, \alpha) = F(x)|A, (FP)_B = v] \geq 1 - \delta.$$

By Equation 3 we can drop the conditioning over A and have

$$\mathbb{P}_{x,F,P}[C^D(x, \alpha) = F(x)|(FP)_B = v] \geq 1 - 2\delta,$$

using that $\delta \geq 2^{-\ell/c}$.

We now want to fix a “good” x such that $F(x)$ is uniform and C does not get too much “information” about $F(x)$ from the oracle. The first condition simply means $x \notin B$. The second is more complicated. We need to avoid that C makes a query to a z such that $P(z)$ is fixed to a bit that depends on $F(x)$; as that could give away the value $F(x)$ (for example if the answer is $F(x) \oplus 0 \oplus 0 \oplus \dots \oplus 0$). For simplicity, we shall fix F at all values $x \in \{0, 1\}^\ell$ that appear in any $z \in B$. We can enlarge B to include such problematic x :

$$B' := B \cup \{x \in \{0, 1\}^\ell : x \in L(\sigma) \text{ for some } \sigma \circ b \in B\}.$$

Note that $|B'| \leq |B| + |B| \cdot k \cdot s \leq \text{poly}(2^{\ell/c})$. By averaging there exists a corresponding fixing v' such that again

$$\mathbb{P}_{x,F,P}[C^D(x, \alpha) = F(x) | (FP)_{B'} = v'] \geq 1 - 2\delta.$$

By the bound on the size of B' we have that there exists a fixed $x^* \notin B'$ such that

$$\mathbb{P}_{F,P}[C^D(x^*, \alpha) = F(x^*) | (FP)_{B'} = v'] \geq 1 - 3\delta,$$

using that $\delta \geq \text{poly}(2^{\ell/c})/2^\ell$ which is implied by $2^{-\ell/c} \geq \text{poly}(2^{\ell/c})/2^\ell$ and hence true for a large enough c .

Now let us hardwire x^* and α and write C_0 for $C(x^*, \alpha)$. So we have

$$\mathbb{P}_{F,P}[C_0^D = F(x^*) | (FP)_{B'} = v'] \geq 1 - 3\delta.$$

Next we produce a series of circuits C_i to arrive to the desired T . Each circuit will have depth $O(\kappa)$, size $\text{poly}(g)$, will not increase the number of oracle queries, and will use the same gates Ξ used by C_0 , and also gates for H and L .

Our first task is to get rid of the queries $z \in B$. If $z \in B$ then $P(z)$ is a fixed value v_z , corresponding to an evaluation $H = H(F(x_1), F(x_2), \dots, F(x_k))$. All the x_i belong to B' , hence the corresponding values of F are fixed and so $D(z)$ is fixed if $z \in B$. Construct a circuit C_1 which has all the values $D(z)$ for $z \in B$ stored in a table, and answers a query $z \in B$ with that fixed value. So we have

$$\mathbb{P}_{F,P}[C_1^D = F(x^*) | (FP)_{B'} = v'] \geq 1 - 3\delta,$$

and we now know that C_1 only makes queries z where $P(z)$ is uniform.

At this point the conditioning on $P_{B'} = v'$ is immaterial, hence we drop it. So we have

$$\mathbb{P}_{F,P}[C_1^D = F(x^*) | F_{B'} = v'] \geq 1 - 3\delta.$$

Recall that $F(x^*)$ is still uniform conditioned on $F_{B'} = v'$ because $x \notin B'$. Hence the previous equation can be rewritten as

$$\frac{1}{2}(\mathbb{P}_{F,P}[C_1^D = 1 | F_{B'} = v', F(x^*) = 1] + \mathbb{P}_{F,P}[C_1^D = 0 | F_{B'} = v', F(x^*) = 0]) \geq 1 - 3\delta.$$

We rewrite this as

$$\mathbb{P}_{F,P}[C_1^D = 1 | F_{B'} = v', F(x^*) = 1] - \mathbb{P}_{F,P}[C_1^D = 1 | F_{B'} = v', F(x^*) = 0] \geq 1 - 6\delta. \quad (4)$$

Now for $b \in \{0, 1\}$ define the oracle E_b as follows. The input is a string $w \in \{0, 1, ?, 1-?\}^s$ with exactly one occurrence of $?$ (possibly as $1-?$). The output is obtained by replacing $?$ with b (and $1-?$ with $1-b$) to obtain $w_b \in \{0, 1\}^s$ and then outputting a uniformly selected bit of $w_b 0^s 1^s$.

Lemma 7. *There is a distribution on circuits C_2 such that for every $b \in \{0, 1\}$ we have*

$$\mathbb{P}[C_2^{E_b} = 1] = \mathbb{P}_{F,P}[C_1^D = 1 | F_{B'} = v', F(x^*) = b].$$

Proof. The high-level idea is that C_2 fills the unfixed values of $F(x)$ on the fly, for every $x \neq x^*$, while keeping a table of its choices. More precisely, initialize the table with the values $F_{B'} = v'$. Then arrange the oracle gates of C_1 in levels $1, 2, \dots$ with level 1 being closest to the input. To perform the queries $z_1 = \sigma_1 b_1, z_2 = \sigma_2 b_2, \dots, z_t = \sigma_t b_t, t \leq q$, at some level i , C_2 first computes $L' := L(\sigma_1) \cup L(\sigma_2) \cup \dots \cup L(\sigma_t)$. For every $x \in L'$ that is in the table, it fetches the corresponding value of F . For the others, except for x^* , it tosses a coin, and stores the value in the table. From these values it computes t strings $w'_1, w'_2, \dots, w'_t \in \{0, 1, ?, 1-\?\}^s$ by computing H . If an evaluation of H depends on the unknown value $F(x^*)$ then C_2 writes $?$ or $1-?$ depending on what this dependency is. (If H syntactically depends on $F(x^*)$ but actually the other input bits of H already determine the value of H , C_2 outputs a value in $\{0, 1\}$. This obviously never happens when $H = \text{Parity}$.)

Recall from above that there is at most one occurrence of x^* in each $L(\sigma_i)$ (for else the oracle always outputs 0 and the query isn't made). Hence each of the strings w'_i has ≤ 1 occurrence of $?$. Then the strings w_1, w_2, \dots, w_t are constructed by performing bit-by-bit equality with b_1, b_2, \dots, b_t . More formally $w_i := 1^s \oplus w'_i \oplus b_i$, where \oplus is bit-wise xor and the expression is the indicator of the equality function (here $(1-?) \oplus 1 = ?$ etc.). For those strings $w_i \in \{0, 1\}^s$, the oracle query is answered simply by the circuit by returning a uniform bit of $w_i 0^s 1^s$. For the others C_2 queries E_b . \square

Now we further simplify the oracle to zoom in on strings which are nearly balanced. Let E'_b be the oracle that takes as input $w \in \{?, 1-\?\}$, substitutes b for $?$ to obtain $w_b \in \{0, 1\}$ and then outputs a uniform bit of $w_b 0^s 1^s$.

Lemma 8. *There is a distribution on circuits C_3 of size polynomial in that of C_2 and constant depth such that for every $b \in \{0, 1\}$ we have*

$$\mathbb{P}[C_3^{E'_b} = 1] = \mathbb{P}[C_2^{E_b} = 1].$$

Proof. Think of answering an oracle query to E_b as follows. On input $w \in \{0, 1, ?, 1-\?\}^s$, write $w = w'w''$ where $w'' \in \{?, 1-\?\}$. First toss a *selection coin* to decide if the answer will be a uniform bit from w' or from $w'' 0^s 1^s$. In the former case we can answer the query without invoking the oracle. In the latter case the oracle query is answered using E'_b . \square

Now we want to replace the oracles E'_0, E'_1 with inputs $N_{1/2}^q, N_{1/2-\epsilon}^q$. Note E'_0 on input $?$ outputs $N_{s/(2s+1)}$ and on input $1-?$ outputs $N_{(s+1)/(2s+1)}$. E'_1 does the same but with $?$ and $1-?$ swapped. We use the following lemma to map pairs of distributions.

Lemma 9. *Let p, γ, ϵ be such that the following quantities are in $[0, 1]$: $1/2 - \epsilon, p + \gamma, p - \gamma, p + \gamma + \gamma/2\epsilon, p + \gamma - \gamma/2\epsilon$. There is a distribution on functions $M : \{0, 1\} \rightarrow \{0, 1\}$ such that $M(N_{1/2}) \equiv N_{p+\gamma}$ and $M(N_{1/2-\epsilon}) \equiv N_{p-\gamma}$, where \equiv denotes equivalence as distributions.*

Proof. $M(0)$ outputs 1 with probability $\alpha := p + \gamma - \gamma/2\epsilon$, $M(1)$ outputs 1 with probability $\beta := p + \gamma + \gamma/2\epsilon$. Then

$$\begin{aligned}\mathbb{P}[M(N_{1/2}) = 1] &= (1/2)\beta + (1/2)\alpha = p + \gamma, \\ \mathbb{P}[M(N_{1/2-\epsilon}) = 1] &= (1/2 - \epsilon)\beta + (1/2 + \epsilon)\alpha = p + \gamma + \epsilon(-\beta + \alpha) = p - \gamma.\end{aligned}$$

□

Consider the circuit C_4 that on input a string $v \in \{0, 1\}^q$ simulates C_3 . Oracle query at gate i on input $?$ is answered applying Lemma 9 to bit v_i of v with $p = 1/2$ and $\gamma = -1/2(2s + 1)$. If v_i is sampled according to $N_{1/2}$ then we get $N_{1/2-1/2(2s+1)} = N_{s/(2s+1)}$, and if v_i is sampled according to $N_{1/2-\epsilon}$ then we get $N_{1/2+1/2(2s+1)} = N_{(s+1)/(2s+1)}$, as desired. On input $1-?$ we instead pick $\gamma = +1/2(2s + 1)$. In both cases the hypotheses of the lemma hold for $\epsilon = \Omega(1/s)$. This shows that

$$\begin{aligned}\mathbb{P}[C_4(N_{1/2}^q) = 1] &= \mathbb{P}[C_3^{E'_0} = 1], \\ \mathbb{P}[C_4(N_{1/2-\epsilon}^q) = 1] &= \mathbb{P}[C_3^{E'_1} = 1].\end{aligned}$$

Combining this with Equation 4 concludes the proof.

References

- [ABFR94] James Aspnes, Richard Beigel, Merrick Furst, and Steven Rudich. The expressive power of voting polynomials. *Combinatorica. An Journal on Combinatorics and the Theory of Computing*, 14(2):135–148, 1994.
- [AS14] Sergei Artemenko and Ronen Shaltiel. Lower bounds on the query complexity of non-uniform and adaptive reductions showing hardness amplification. *Computational Complexity*, 23(1):43–83, 2014.
- [FSUV13] Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. On beating the hybrid argument. *Theory of Computing*, 9:809–843, 2013.
- [GNW95] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao’s XOR lemma. Technical Report TR95–050, *Electronic Colloquium on Computational Complexity*, March 1995. www.eccc.uni-trier.de/.
- [GR08] Dan Gutfreund and Guy Rothblum. The complexity of local list decoding. In *12th Intl. Workshop on Randomization and Computation (RANDOM)*, 2008.
- [GSV18] Aryeh Grinberg, Ronen Shaltiel, and Emanuele Viola. Indistinguishability by adaptive procedures with advice, and lower bounds on hardness amplification proofs. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2018. Available at <http://www.ccs.neu.edu/home/viola/>.
- [Kli01] Adam R. Klivans. On the derandomization of constant depth circuits. In *Workshop on Randomization and Computation (RANDOM)*. Springer, 2001.
- [KS03] Adam Klivans and Rocco A. Servedio. Boosting and hard-core sets. *Machine Learning*, 53(3):217–238, 2003.

- [LTW11] Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. Complexity of hard-core set proofs. *Computational Complexity*, 20(1):145–171, 2011.
- [LV12] Shachar Lovett and Emanuele Viola. Bounded-depth circuits cannot sample good codes. *Computational Complexity*, 21(2):245–266, 2012.
- [Nis91] Noam Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica. An Journal on Combinatorics and the Theory of Computing*, 11(1):63–70, 1991.
- [NR04] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. of the ACM*, 51(2):231–262, 2004.
- [NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. of Computer and System Sciences*, 49(2):149–167, 1994.
- [Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in *Mathematical Notes of the Academy of Sci. of the USSR*, 41(4):333–338, 1987.
- [RR97] Alexander Razborov and Steven Rudich. Natural proofs. *J. of Computer and System Sciences*, 55(1):24–35, August 1997.
- [Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *19th ACM Symp. on the Theory of Computing (STOC)*, pages 77–82. ACM, 1987.
- [SV10] Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM J. on Computing*, 39(7):3122–3154, 2010.
- [Vio04] Emanuele Viola. The complexity of constructing pseudorandom generators from hard functions. *Computational Complexity*, 13(3-4):147–188, 2004.
- [Vio06] Emanuele Viola. *The Complexity of Hardness Amplification and Derandomization*. PhD thesis, Harvard University, 2006.
- [Vio07] Emanuele Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM J. on Computing*, 36(5):1387–1403, 2007.
- [Vio09] Emanuele Viola. On the power of small-depth computation. *Foundations and Trends in Theoretical Computer Science*, 5(1):1–72, 2009.