

Query Complexity in Errorless Hardness Amplification

Thomas Watson*

June 23, 2013

Abstract

An errorless circuit for a boolean function is one that outputs the correct answer or “don’t know” on each input (and never outputs the wrong answer). The goal of errorless hardness amplification is to show that if f has no size s errorless circuit that outputs “don’t know” on at most a δ fraction of inputs, then some f' related to f has no size s' errorless circuit that outputs “don’t know” on at most a $1 - \epsilon$ fraction of inputs. Thus the hardness is “amplified” from δ to $1 - \epsilon$. Unfortunately, this amplification comes at the cost of a loss in circuit size. This is because such results are proven by reductions which show that any size s' errorless circuit for f' that outputs “don’t know” on at most a $1 - \epsilon$ fraction of inputs could be used to construct a size s errorless circuit for f that outputs “don’t know” on at most a δ fraction of inputs. If the reduction makes q queries to the hypothesized errorless circuit for f' , then plugging in a size s' circuit yields a circuit of size $\geq qs'$, and thus we must have $s' \leq s/q$. Hence it is desirable to keep the query complexity to a minimum.

The first results on errorless hardness amplification were obtained by Bogdanov and Safra. They achieved query complexity $O((\frac{1}{\delta} \log \frac{1}{\epsilon})^2 \cdot \frac{1}{\epsilon} \log \frac{1}{\delta})$ when f' is the XOR of several independent copies of f . We improve the query complexity (and hence the loss in circuit size) to $O(\frac{1}{\epsilon} \log \frac{1}{\delta})$, which is optimal up to constant factors for nonadaptive black-box errorless hardness amplification.

Bogdanov and Safra also proved a result that allows for errorless hardness amplification within NP. They achieved query complexity $O(k^3 \cdot \frac{1}{\epsilon} \log \frac{1}{\delta})$ when f' consists of any monotone function applied to the outputs of k independent copies of f , provided the monotone function satisfies a certain combinatorial property parameterized by δ and ϵ . We improve the query complexity to $O(\frac{k}{t} \cdot \frac{1}{\epsilon} \log \frac{1}{\delta})$, where $t \geq 1$ is a certain parameter of the monotone function.

As a side result, we prove a lower bound on the advice complexity of black-box reductions for errorless hardness amplification.

1 Introduction

Traditionally, an algorithm for solving a computational problem is required to be correct on all inputs and is judged in terms of its efficiency (the amount of computational resources it uses). One criticism of this model is that it is too strict: In practice, an algorithm only needs to be correct on “real-world” inputs and not on contrived worst-case inputs. To address this issue within the framework of complexity theory, researchers developed the theory of average-case complexity

*Computer Science Division, University of California, Berkeley. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-0946797 and by the National Science Foundation under Grant No. CCF-1017403.

(starting with the work of Levin [15]). In this theory, an algorithm is judged in terms of both its efficiency and the fraction of inputs on which it fails to solve the problem correctly. The topic of this paper is the relationship between these two measures of the quality of an algorithm.

There are two standard settings for average-case complexity. In the original setting proposed by Levin [15], one only considers *errorless algorithms*, which are required to output the correct answer or “don’t know” on each input.¹ An errorless algorithm is judged in terms of both its efficiency and the fraction of inputs on which it outputs “don’t know”. We refer to this setting as *errorless average-case complexity*. In the other setting, one considers arbitrary algorithms which may output the wrong answer rather than just “don’t know” on an input. We refer to this setting as *non-errorless average-case complexity*. Errorless average-case complexity is an intermediate setting between worst-case complexity and non-errorless average-case complexity.

We first discuss non-errorless average-case complexity. A boolean function is said to be *mildly average-case hard* if no efficient algorithm can compute it on almost all inputs. Applications such as derandomization and cryptography require functions that are *strongly average-case hard*, meaning that no efficient algorithm can compute the function on noticeably more than half the inputs. This motivates hardness amplification, which is the problem of transforming a mildly average-case hard function into a strongly average-case hard function. A classic result in this area is the XOR Lemma [16, 11, 6, 13], which states that the XOR of sufficiently many independent copies of a mildly average-case hard function is strongly average-case hard, for suitable models of computation such as boolean circuits.

However, the XOR Lemma (as well as the numerous subsequent results on hardness amplification) incurs an unfortunate loss in circuit size. Suppose the original function f is mildly average-case hard in the sense that no size s circuit succeeds on at least a $1 - \delta$ fraction of inputs, and we wish for the new function f' to be strongly average-case hard in the sense that no size s' circuit succeeds on at least a $1/2 + \epsilon$ fraction of inputs. Then we would like s' to be as large as possible, but the XOR Lemma only gives results with s' smaller than s . This is because such results are proven by reductions which show that if f' is not strongly average-case hard, then a circuit witnessing this could be used to construct a circuit witnessing that f is not mildly average-case hard. If the reduction makes q queries to the hypothesized circuit, then plugging in a size s' circuit yields a circuit of size $\geq qs'$, and thus we must have $s' \leq s/q$. Hence the query complexity q governs the loss in circuit size. For the XOR Lemma, the query complexity is well-understood (at least in the case of nonadaptive black-box reductions). The proof due to Impagliazzo [11] and Klivans and Servedio [14] shows that $q = O(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$ queries are sufficient, and Shaltiel and Viola [20] showed that $q = \Omega(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$ queries are necessary for nonadaptive black-box reductions such as the ones used to prove the upper bound.

Bogdanov and Safra [3] initiated the study of hardness amplification in Levin’s original setting of errorless average-case complexity. A boolean function is said to be *mildly errorless average-case hard* if no efficient errorless algorithm (say, size s circuit) can compute it on almost all inputs (say, a $1 - \delta$ fraction). A function is said to be *strongly errorless average-case hard* if no efficient errorless algorithm (say, size s' circuit) can compute it on a noticeable fraction of inputs (say, an ϵ fraction). Note that in the non-errorless setting, computing a boolean function on half the inputs is trivial (using constant 0 or constant 1), but in the errorless setting, computing a boolean function

¹Actually, Levin proposed considering algorithms that are correct on all inputs but which are efficient “on average” with respect to a random input. Under a suitable formalization, such algorithms are equivalent to errorless algorithms that may fail on a small fraction of inputs but are efficient on all inputs.

on even a small fraction of inputs is nontrivial. The goal of errorless hardness amplification is to transform a mildly errorless average-case hard function f into a strongly errorless average-case hard function f' . Such results suffer from a loss in circuit size for the same reason as in the non-errorless setting. Bogdanov and Safra [3] showed that $q = O\left(\left(\frac{1}{\delta} \log \frac{1}{\epsilon}\right)^2 \cdot \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ queries are sufficient when f' is the XOR of several independent copies of f . The result of Shaltiel and Viola [20] can be modified without difficulty to show that $q = \Omega\left(\frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ queries are necessary for nonadaptive black-box reductions. We close the gap by showing that $q = O\left(\frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ queries are sufficient (via a nonadaptive black-box reduction).

Another natural goal for hardness amplification is to guarantee that if f represents an NP language at some input length, then f' also represents an NP language at some input length. In the non-errorless setting this goal has been studied in numerous works [19, 22, 10, 23, 4, 17, 7], and in the errorless setting this goal has been studied by Bogdanov and Safra [3]. We significantly improve the query complexity of the Bogdanov-Safra result.

1.1 The Errorless XOR Lemma

Given $f : \{0,1\}^n \rightarrow \{0,1\}$ we define $f^{\oplus k} : \{0,1\}^{n \times k} \rightarrow \{0,1\}$ as follows: $f^{\oplus k}(x_1, \dots, x_k) = f(x_1) \oplus \dots \oplus f(x_k)$.

Definition 1 (Errorless Average-Case Hardness). *We say a circuit $A : \{0,1\}^n \rightarrow \{0,1,\perp\}$ is a δ -errorless circuit for $f : \{0,1\}^n \rightarrow \{0,1\}$ if*

- (i) $A(x) \in \{f(x), \perp\}$ for all $x \in \{0,1\}^n$, and
- (ii) $\Pr_x[A(x) = \perp] \leq \delta$ where $x \in \{0,1\}^n$ is chosen uniformly at random.

We say f is (s, δ) -hard if it has no δ -errorless circuit of size $\leq s$.

Theorem 1 (Query-Optimal Errorless XOR Lemma). *If f is (s, δ) -hard then $f' = f^{\oplus k}$ is $(s', 1 - \epsilon)$ -hard where $s' = s / \left(\frac{4}{\epsilon} \ln \frac{2}{\delta}\right)$, provided $k \geq \frac{16}{\delta} \ln \frac{2}{\epsilon}$.*

We prove Theorem 1 in Section 2.² The query complexity of the reduction underlying Theorem 1 corresponds to the ratio s/s' . Bogdanov and Safra [3] proved a version of Theorem 1 where $s' = s / \left(k^2 \cdot \frac{2}{\epsilon} \ln \frac{2}{\delta}\right)$, provided $k \geq \frac{2}{\delta} \ln \frac{2}{\epsilon}$.³ For the best value of k , they only achieve $O\left(\left(\frac{1}{\delta} \log \frac{1}{\epsilon}\right)^2 \cdot \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ query complexity.

We prove Theorem 1 by a reduction similar to the one used in [3]. Our contribution is a new, tight analysis of the reduction. The crux of the reduction is a randomized procedure that solves f errorlessly (meaning that for each input x it may output $f(x)$ with some probability and \perp with some probability, but it never outputs $\overline{f(x)}$) while making one query to a hypothesized $(1 - \epsilon)$ -errorless circuit A' for f' . Suppose for some $\beta > 0$ we knew that $\leq \delta/2$ fraction of inputs x are bad in the sense that the probability the procedure outputs $f(x)$ is $< \beta$. Then by amplifying the success probability on the good inputs and hard-wiring the randomness appropriately, we obtain a δ -errorless circuit A for f , via a reduction with query complexity $O\left(\frac{1}{\beta} \log \frac{1}{\delta}\right)$. The heart of our

²In the statement of Theorem 1, it would be more accurate to say $s' = s / \left(\frac{4}{\epsilon} \ln \frac{2}{\delta}\right) - O(1)$ to account for the trivial circuitry needed to combine the results of the $\frac{4}{\epsilon} \ln \frac{2}{\delta}$ queries the reduction makes. Throughout this paper, we ignore such details. We also ignore details arising from the fact that numbers such as $\frac{4}{\epsilon} \ln \frac{2}{\delta}$ might not be integers.

³Actually, their proof gives $s' = s / \left(k^2 \cdot \frac{2}{\epsilon} \ln \frac{2}{\delta}\right)$, provided $k \geq \frac{1}{\delta} \ln \frac{2}{\epsilon}$, but a minor tweak to their proof yields the stated result.

improvement over the Bogdanov-Safra proof is in arguing that we can take $\beta = \epsilon/4$. To prove this, we suppose the fraction of bad inputs is $> \delta/2$ and prove that then A' must compute f' on $< \epsilon$ fraction of inputs. The procedure outputs $f(x)$ if and only if the query is an input on which A' computes f' ; furthermore the distribution of this query (x_1, \dots, x_k) is obtained by setting $x_i = x$ for a uniformly random i and picking $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k$ uniformly at random. Consider the following two distributions on queries to A' : the uniform distribution, and the distribution obtained by picking a random bad x and running the procedure on input x . We wish to show that A' computes f' with probability $< \epsilon$ under the former distribution, and we know A' computes f' with probability $< \beta = \epsilon/4$ under the latter. We show the two distributions are “close” in the sense that the probability of any event under the former is less than twice the probability under the latter plus $\epsilon/2$: For every query (x_1, \dots, x_k) with, say, a ν fraction of bad coordinates, its probability under the former distribution is μ/ν times larger than under the latter distribution, where $\mu > \delta/2$ is the fraction of bad inputs. By a Chernoff bound, $\nu \geq \mu/2$ for all but an $\epsilon/2$ fraction of queries. Our formal proof of Theorem 1 can be viewed as implicitly following the above intuition, though we rearrange the argument somewhat to highlight the similarity with the (more involved) proof of our result on errorless hardness amplification within NP.

It can be shown that $\Omega(\frac{1}{\epsilon} \log \frac{1}{\delta})$ queries are needed by any nonadaptive black-box reduction achieving errorless hardness amplification, regardless of how f' is constructed from f (see Section 1.3 for the precise statement). Since our proof of Theorem 1 (and the Bogdanov-Safra proof) is by a nonadaptive black-box reduction, this shows that Theorem 1 is optimal in a sense. Shaltiel and Viola [20] gave a general technique for lower bounding the query complexity of nonadaptive black-box reductions, and they applied it to non-errorless hardness amplification (including the XOR Lemma and the Direct Product Lemma) and to constructions of pseudorandom generators from average-case hard functions. Similarly, we observe that their technique applies to errorless hardness amplification. Artemenko and Shaltiel [1] have proven a $\Omega(\frac{1}{\epsilon})$ query lower bound even for *adaptive* black-box reductions. The optimal $\Omega(\frac{1}{\epsilon} \log \frac{1}{\delta})$ lower bound for adaptive reductions remains open.

1.2 Monotone Errorless Amplification

Consider the problem of errorless hardness amplification within NP. That is, if f is computable in nondeterministic polynomial time, then we want f' to also be computable in nondeterministic polynomial time. Taking $f' = f^{\oplus k}$ does not guarantee this. We instead consider more general constructions of the form $f' = C \circ f^k$ where $C : \{0, 1\}^k \rightarrow \{0, 1\}$, and $f^k : \{0, 1\}^{n \times k} \rightarrow \{0, 1\}^k$ is defined as $f^k(x_1, \dots, x_k) = (f(x_1), \dots, f(x_k))$. In the setting of the XOR Lemma, the combiner function C is the k -bit parity function. If C is monotone (that is, $C(y_1, \dots, y_k) \leq C(z_1, \dots, z_k)$ whenever $y_i \leq z_i$ for all $i \in [k]$) and f and C are both computable in nondeterministic polynomial time, then f' is guaranteed to be computable in nondeterministic polynomial time. This approach dates back to [19, 10].

Bogdanov and Safra [3] showed that this construction yields errorless hardness amplification provided the monotone combiner function C satisfies a certain combinatorial property. To describe this property, we need some definitions from [3] (though we use somewhat different notation). Fix $b \in \{0, 1\}$. Given a monotone function $C : \{0, 1\}^k \rightarrow \{0, 1\}$ and a string $y \in \{0, 1\}^k$, we say that coordinate $i \in [k]$ is *b-sensitive* if flipping the i th bit of y causes the value of C to flip from b to \bar{b} , and we let $\sigma(C, y, b)$ denote the set of b -sensitive coordinates. That is,

$$\sigma(C, y, b) = \{i \in [k] : C(y) = b \text{ and } C(y \oplus e_i) = \bar{b}\}.$$

Note that if $C(y) = \bar{b}$ then $\sigma(C, y, b) = \emptyset$ and if $C(y) = b$ then by the monotonicity of C , $\sigma(C, y, b)$ only contains coordinates i such that $y_i = b$. For $p \in [0, 1]$, we use $y \sim_p \{0, 1\}^k$ to denote that y is sampled from the p -biased distribution, that is, each bit is independently set to 1 with probability p .

Definition 2. For $b \in \{0, 1\}$, a function $C : \{0, 1\}^k \rightarrow \{0, 1\}$ is a (t, ρ, p, b) -amplifier if C is monotone and

$$\Pr_{y \sim_p \{0, 1\}^k} \left[|\sigma(C, y, b)| \geq t \right] \geq 1 - \rho.$$

Note that a monotone function $C : \{0, 1\}^k \rightarrow \{0, 1\}$ is a (t, ρ, p, b) -amplifier if and only if its monotone complement $C^\dagger : \{0, 1\}^k \rightarrow \{0, 1\}$ is a $(t, \rho, 1 - p, \bar{b})$ -amplifier, where C^\dagger is defined as

$$C^\dagger(y_1, \dots, y_k) = \overline{C(\bar{y}_1, \dots, \bar{y}_k)}.$$

For reasons discussed in [3], we consider the following one-sided version of Definition 1.

Definition 3 (One-Sided Errorless Average-Case Hardness). For $b \in \{0, 1\}$, we say a circuit $A : \{0, 1\}^n \rightarrow \{0, 1, \perp\}$ is a (δ, b) -errorless circuit for $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if

- (i) $A(x) \in \{f(x), \perp\}$ for all $x \in \{0, 1\}^n$, and
- (ii) $\Pr_x[A(x) = \perp] \leq \delta$ where $x \in f^{-1}(b)$ is chosen uniformly at random.

We say f is (s, δ, b) -hard if it has no (δ, b) -errorless circuit of size $\leq s$.

Note that if f is (s, δ) -hard then f is either $(s/2, \delta, 0)$ -hard or $(s/2, \delta, 1)$ -hard.

Theorem 2 (Monotone Errorless Amplification Lemma). For $b \in \{0, 1\}$, if f is (s, δ, b) -hard and $C : \{0, 1\}^k \rightarrow \{0, 1\}$ is a (t, ρ, p, b) -amplifier then $f' = C \circ f^k$ is $(s', 1 - \epsilon)$ -hard where $s' = s / \left(\frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta} \right)$, provided $t \geq \frac{16}{\delta} \ln \frac{4}{\epsilon}$, $\rho \leq \epsilon/4$, and $p = \Pr_x[f(x) = 1]$.

We prove Theorem 2 in Section 3. The query complexity of the reduction underlying Theorem 2 corresponds to the ratio s/s' . Bogdanov and Safra [3] proved a version of Theorem 2 where $s' = s / \left(k^3 \cdot \frac{64}{\epsilon^2} \ln \frac{2}{\delta} \right)$, provided $t \geq \frac{4}{\delta} \ln \frac{8}{\epsilon}$ and $\rho \leq \epsilon/2$. Their argument involves considering the subcubes of $\{0, 1\}^{n \times k}$ given by $f^k(x_1, \dots, x_k) = y$ for each y individually and then combining the results for the different subcubes using a nontrivial probabilistic argument. We show how to give a direct argument that handles all the subcubes simultaneously. This idea alone actually simplifies the proof and reduces the query complexity to $O(k^2 \cdot \frac{1}{\epsilon} \log \frac{1}{\delta})$. Combining this idea with the ideas from our analysis in the proof of Theorem 1 allows us to further reduce the query complexity to $O\left(\frac{k}{t} \cdot \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$. We believe this bound on the query complexity cannot be improved without exploiting some non-obvious structural property of (t, ρ, p, b) -amplifiers; however, we could not come up with a compelling way to formalize this.

Bogdanov and Safra [3] showed how to construct good amplifiers (with large t and small ρ) and how to use Theorem 2 with these amplifiers to do uniform and nonuniform errorless hardness amplification within NP. For simplicity in describing their amplifiers, we restrict to the case $p = 1/2$ (so only balanced functions f are allowed). Bogdanov and Safra showed that for all constants $\alpha, \beta \geq 0$ and sufficiently large k , every $(t, \rho, 1/2, b)$ -amplifier with $t = k^\alpha$ and $\rho = k^{-\beta}$ satisfies $\alpha + \beta \leq 1/2$. They constructed an amplifier they call OR-DDNF_α (which they view as a

derandomization of Talagrand’s random DNF [21]) that achieves near-optimal parameters: $t = k^\alpha$ and $\rho = O(k^{-(1/2-\alpha)} \cdot \sqrt{\log k})$, but only for $b = 0$. Of course, $\text{OR-DDNF}_\alpha^\dagger$ achieves the same parameters for $b = 1$, but OR-DDNF_α is in NP whereas $\text{OR-DDNF}_\alpha^\dagger$ is not known to be in NP. Hence $\text{OR-DDNF}_\alpha^\dagger$ cannot be used for errorless hardness amplification within NP unless $k \leq O(\log n)$. Bogdanov and Safra constructed another amplifier they call HOL-AMP_α (based on the “holographic function” of Benamini, Schramm, and Wilson [2]). For $b = 0$ this amplifier achieves $t = k^\alpha$ and $\rho = O(k^{-(1/3-\alpha)} \cdot \text{polylog } k)$, and of course $\text{HOL-AMP}_\alpha^\dagger$ achieves the same parameters for $b = 1$. These parameters are worse than for OR-DDNF_α , but the advantage is that HOL-AMP_α and $\text{HOL-AMP}_\alpha^\dagger$ are computable in polynomial time, so they can both be used for errorless hardness amplification within NP.

To plug one of these amplifiers into Theorem 2 and minimize the query complexity, one should pick an $0 \leq \alpha < 1/2$ and then choose the smallest k allowed by the provisions $k^\alpha = t \geq \frac{16}{\delta} \ln \frac{4}{\epsilon}$ and $\rho \leq \epsilon/4$. For small α , the former provision is more restrictive, and for large α the latter is more restrictive. The provisions in the original Bogdanov-Safra version of Theorem 2 are the same as ours (up to constant factors), but the query complexity is a factor $k^2 t / \epsilon$ larger than ours.

Bogdanov and Safra [3] used the HOL-AMP_α and $\text{HOL-AMP}_\alpha^\dagger$ amplifiers to show that if every problem in NP has polynomial-size errorless circuits that succeed on at least $n^{-2/9+o(1)}$ fraction of inputs, then every problem in NP has polynomial-size errorless circuits that succeed on at least $1 - n^{-c}$ fraction of inputs (for every constant c). This particular result is only concerned with the loss in circuit size at the granularity of polynomial versus not polynomial, and our result only gives a polynomial improvement in the loss in circuit size. Hence our result does not help to improve the $2/9$ exponent, which arises from the parameters of the amplifier.

1.3 Black-Box Lower Bounds

We give lower bounds on the query complexity and advice complexity of black-box errorless hardness amplification proofs. We allow ourselves to identify strings with functions; for example, we identify $\{0, 1\}^{2^n}$ with the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}$.

Definition 4. *An $(n, n', \delta, \epsilon, \alpha)$ -black-box errorless hardness amplification is a pair (Enc, Dec) with $\text{Enc} : \{0, 1\}^{2^n} \rightarrow \{0, 1\}^{2^{n'}}$ and $\text{Dec} : \{0, 1, \perp\}^{2^{n'}} \times \{0, 1\}^\alpha \rightarrow \{0, 1, \perp\}^{2^n}$, such that for all $f \in \{0, 1\}^{2^n}$ and $A' \in \{0, 1, \perp\}^{2^{n'}}$ there exists an $a \in \{0, 1\}^\alpha$ such that the following holds, where $f' = \text{Enc}(f)$ and $A = \text{Dec}(A', a)$. If*

- (i) $A'(x') \in \{f'(x'), \perp\}$ for all $x' \in \{0, 1\}^{n'}$, and
- (ii) $\Pr_{x'}[A'(x') = \perp] \leq 1 - \epsilon$,

then

- (i) $A(x) \in \{f(x), \perp\}$ for all $x \in \{0, 1\}^n$, and
- (ii) $\Pr_x[A(x) = \perp] \leq \delta$.

We say it is q -query if there is an algorithm that takes (x, a) as input, makes q (possibly adaptive) queries to A' , and outputs $A(x)$. We further say it is nonadaptive if the algorithm makes its queries to A' nonadaptively.

In fact, our lower bounds hold even for the following weaker type of reduction.

Definition 5. An $(n, n', \delta, \epsilon, \alpha)$ -black-box non-errorless to errorless hardness amplification is a pair (Enc, Dec) with $\text{Enc} : \{0, 1\}^{2^n} \rightarrow \{0, 1\}^{2^{n'}}$ and $\text{Dec} : \{0, 1, \perp\}^{2^{n'}} \times \{0, 1\}^\alpha \rightarrow \{0, 1\}^{2^n}$, such that for all $f \in \{0, 1\}^{2^n}$ and $A' \in \{0, 1, \perp\}^{2^{n'}}$ there exists an $a \in \{0, 1\}^\alpha$ such that the following holds, where $f' = \text{Enc}(f)$ and $A = \text{Dec}(A', a)$. If

- (i) $A'(x') \in \{f'(x'), \perp\}$ for all $x' \in \{0, 1\}^{n'}$, and
- (ii) $\Pr_{x'}[A'(x') = \perp] \leq 1 - \epsilon$,

then $\Pr_x[A(x) \neq f(x)] \leq \delta$. We say it is q -query if there is an algorithm that takes (x, a) as input, makes q (possibly adaptive) queries to A' , and outputs $A(x)$. We further say it is nonadaptive if the algorithm makes its queries to A' nonadaptively.

In the proof of Theorem 1 we show that there exists a q -query nonadaptive $(n, n', \delta, \epsilon, \alpha)$ -black-box errorless hardness amplification where $q = \frac{4}{\epsilon} \ln \frac{2}{\delta}$, $n' = kn$, and $\alpha = (\log_2 k + (k - 1)n + 1) \cdot q$, provided $k \geq \frac{16}{\delta} \ln \frac{2}{\epsilon}$.

Theorem 3 ([20]). *There exists a universal constant $c > 1$ such that the following holds. If there exists a q -query nonadaptive $(n, n', \delta, \epsilon, \alpha)$ -black-box non-errorless to errorless hardness amplification then $q \geq \frac{1}{c} \cdot \frac{1}{\epsilon} \ln \frac{1}{\delta}$, provided $n \geq c$, $n' \geq c$, $2^{-n/c} \leq \delta \leq 1/3$, $2^{-n/c} \leq \epsilon \leq 1/3$, and $\alpha \leq 2^{n/c}$.*

Shaltiel and Viola [20] actually proved a similar result for the fully non-errorless setting (with the conclusion $q \geq \frac{1}{c} \cdot \frac{1}{\epsilon^2} \ln \frac{1}{\delta}$). Their proof involves using noise that flips bits (representing the answers to queries) with probability either $1/2 - \epsilon$ or $1/2$, and they argue that a nonadaptive reduction can be used to distinguish between the two cases, and this requires $\Omega(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$ bits (queries). We adapt their proof to our errorless setting as follows. The two cases instead involve using noise that masks the bit with \perp with probability $1 - \epsilon$ and reveals the correct bit with probability ϵ , or using “noise” that masks the bit with \perp with probability 1. A nonadaptive reduction in our setting can be used to distinguish between the two cases, and this requires $\Omega(\frac{1}{\epsilon} \log \frac{1}{\delta})$ bits (queries). The rest of the Shaltiel-Viola proof goes through with some minor changes but without major changes, so we omit the argument.

Building on the techniques of [20], Artemenko and Shaltiel [1] proved the following lower bound, which applies to *adaptive* reductions but which falls short of the tight $\Omega(\frac{1}{\epsilon} \log \frac{1}{\delta})$ lower bound by a factor of $\log \frac{1}{\delta}$. Like Theorem 3, the result of [1] holds for “non-errorless to errorless” hardness amplification, and thus simultaneously applies to both the fully non-errorless setting and the fully errorless setting.

Theorem 4 ([1]). *There exists a universal constant $c > 1$ such that the following holds. If there exists a q -query $(n, n', \delta, \epsilon, \alpha)$ -black-box non-errorless to errorless hardness amplification then $q \geq \frac{1}{c} \cdot \frac{1}{\epsilon}$, provided $n \geq c$, $n' \geq c$, $\delta \leq 1/3$, $2^{-n/c} \leq \epsilon \leq 1/3$, and $\alpha \leq 2^{n/c}$.*

We now turn our attention to lower bounds on advice complexity.

Theorem 5. *If there exists an $(n, n', \delta, \epsilon, \alpha)$ -black-box non-errorless to errorless hardness amplification then $2^\alpha \geq \log_3 \frac{1}{\epsilon}$, provided $\delta < 1/8$ and $\epsilon \geq 1/1.01^{1.01^{2^n}}$.*

We prove Theorem 5 in Section 4. In the fully non-errorless setting, lower bounds on advice complexity correspond to lower bounds on list size for approximately list-decoding error-correcting codes from flipped bits. Such a lower bound was given in [18] (see also [9]). In the non-errorless to errorless setting, lower bounds on advice complexity correspond to lower bounds on list size for approximately list-decoding error-correcting codes from erasures. For exact decoding, such lower bounds were given in [8, 5]. Our proof of Theorem 5 is simpler and cleaner than the proofs of the latter results (at the cost of achieving worse constants), and it handles approximate decoding. Also, the presentation in [8], which is geared toward coding theorists, views the rate, list size, and fraction of erasures as constants for an infinite family of codes. Our presentation is geared toward complexity theorists, who are interested more generally in the asymptotic relationships among all the parameters.

1.4 Preliminaries

We use the following standard Chernoff bound several times.

Theorem 6. *If X_1, \dots, X_τ are fully independent indicator random variables each with expectation μ , then $\Pr [\sum_{j=1}^\tau X_j < \mu\tau/2] < e^{-\mu\tau/8}$.*

2 Proof of Theorem 1

We prove the contrapositive. Suppose f' is not $(s', 1 - \epsilon)$ -hard and thus there is a circuit A' of size $\leq s'$ such that

- (i) $A'(x_1, \dots, x_k) \in \{f'(x_1, \dots, x_k), \perp\}$ for all x_1, \dots, x_k , and
- (ii) $\Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) = \perp] \leq 1 - \epsilon$.

We give a nonuniform reduction that makes $\frac{4}{\epsilon} \ln \frac{2}{\delta}$ nonadaptive queries to A' and combines the results with some trivial computation, yielding a circuit A that witnesses that f is not (s, δ) -hard. To start out, we give a randomized algorithm (Algorithm 1) that solves f errorlessly using oracle access to A' and oracle access to f . The oracle queries to f only depend on the randomness (and not on the input), and later we will hard-wire a particular choice of randomness to get a circuit without oracle access to f .

Input: $x \in \{0, 1\}^n$
Output: $f(x)$ or \perp

- 1 **repeat** $\frac{4}{\epsilon} \ln \frac{2}{\delta}$ times
- 2 pick $i \in [k]$ and $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k \in \{0, 1\}^n$ uniformly at random
- 3 if $A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp$ then halt and output

$f(x_1) \oplus \dots \oplus f(x_{i-1}) \oplus A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \oplus f(x_{i+1}) \oplus \dots \oplus f(x_k)$
- 4 **end**
- 5 halt and output \perp

Algorithm 1: Reduction for Theorem 1

Define the good set

$$G = \left\{ x \in \{0, 1\}^n : \Pr_{i, x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k} [A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp] \geq \epsilon/4 \right\},$$

and define the bad set $B = \{0, 1\}^n \setminus G$. That is, G is the set of inputs for which each iteration of the loop has at least an $\epsilon/4$ probability of producing output.

Claim 1. $|B| \leq (\delta/2) \cdot 2^n$.

Proof. Suppose for contradiction that $|B| > (\delta/2) \cdot 2^n$. Let $\gamma = |B|/2^{n+1}$. We define the event

$$W = \left\{ (x_1, \dots, x_k) \in \{0, 1\}^{n \times k} : |\{i : x_i \in B\}| \geq \gamma k \right\}.$$

That is, W is the event that at least a γ fraction of coordinates are bad. We have

$$\begin{aligned} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp] &\leq \Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W] + \\ &\quad \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W]. \end{aligned}$$

We show that both terms on the right side are $< \epsilon/2$, thus contradicting property (ii) of A' .

Bounding the first term. Applying Theorem 6 with X_i as the indicator variable for $x_i \in B$, and with $\tau = k$ and $\mu = |B|/2^n$, we have

$$\Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W] < e^{-k \cdot |B|/2^{n+3}} < e^{-k\delta/16} \leq \epsilon/2$$

where the middle inequality follows by our assumption on $|B|$ and the last inequality follows by $k \geq \frac{16}{\delta} \ln \frac{2}{\epsilon}$.

Bounding the second term. For each $S \subseteq [k]$ we define the event

$$W_S = \left\{ (x_1, \dots, x_k) \in \{0, 1\}^{n \times k} : \forall i \ x_i \in B \Leftrightarrow i \in S \right\}.$$

Note that the W_S 's are disjoint and

$$W = \bigcup_{S : |S| \geq \gamma k} W_S.$$

We have

$$\begin{aligned} &\Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W] \\ &= \sum_{S : |S| \geq \gamma k} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S] \\ &\leq \frac{1}{\gamma k} \sum_{S \subseteq [k]} |S| \cdot \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S] \\ &= \frac{1}{\gamma k} \sum_{i \in [k]} \sum_{S \ni i} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S] \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\gamma k} \sum_{i \in [k]} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } x_i \in B] \\
&= \frac{1}{\gamma k} \sum_{i \in [k]} \sum_{x \in B} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \mid x_i = x] \cdot \Pr_{x_1, \dots, x_k} [x_i = x] \\
&= \frac{1}{\gamma k 2^n} \sum_{x \in B} \sum_{i \in [k]} \Pr_{x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k} [A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp] \\
&= \frac{1}{\gamma k 2^n} \sum_{x \in B} k \cdot \Pr_{i, x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k} [A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp] \\
&< \frac{1}{\gamma k 2^n} \sum_{x \in B} k \cdot \epsilon/4 \\
&= \frac{\epsilon/4}{\gamma 2^n} \cdot |B| \\
&= \epsilon/2
\end{aligned}$$

where the second and fifth lines follow by the disjointness of the W_S 's, and the remaining lines follow by simple rearrangements. \square

The rest of the proof of Theorem 1 is similar to the argument from [3]. First we note that for all $x \in \{0, 1\}^n$ and all choices of randomness, Algorithm 1 does indeed output either $f(x)$ or \perp . This follows trivially from the fact that if $A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp$ then

$$A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) = f(x_1) \oplus \dots \oplus f(x_{i-1}) \oplus f(x) \oplus f(x_{i+1}) \oplus \dots \oplus f(x_k)$$

by property (i) of A' . Next we observe that for each $x \in G$, we have

$$\Pr_{\text{randomness}} [\text{Algorithm 1 outputs } \perp] \leq (1 - \epsilon/4)^{\frac{4}{\epsilon} \ln \frac{2}{\delta}} \leq \delta/2.$$

Therefore

$$\begin{aligned}
\Pr_{x, \text{randomness}} [\text{Algorithm 1 outputs } \perp] &\leq \Pr_x [x \in B] + \\
&\quad \mathbb{E}_x \left[\Pr_{\text{randomness}} [\text{Algorithm 1 outputs } \perp] \mid x \in G \right] \\
&\leq \delta/2 + \mathbb{E}_x [\delta/2 \mid x \in G] \\
&= \delta
\end{aligned}$$

where the second inequality follows by Claim 1 and by the above observation. It follows that there exists a setting of the randomness such that

- (i) Algorithm 1 outputs $f(x)$ or \perp for all x , and
- (ii) $\Pr_x [\text{Algorithm 1 outputs } \perp] \leq \delta$.

To get a circuit A that witnesses that f is not (s, δ) -hard, just hard-wire the randomness and the values of $f(x_1) \oplus \dots \oplus f(x_{i-1}) \oplus f(x_{i+1}) \oplus \dots \oplus f(x_k)$ needed for this choice of randomness, and plug in the hypothesized circuit A' . Since A' has size $\leq s'$ and Algorithm 1 makes $\frac{4}{\epsilon} \ln \frac{2}{\delta}$ queries to A' , A has size $\leq s' \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta} = s$. Note that Algorithm 1 can trivially be implemented with *nonadaptive* access to A' .

3 Proof of Theorem 2

Fix an arbitrary bit b . We prove the contrapositive. Suppose f' is not $(s', 1 - \epsilon)$ -hard and thus there is a circuit A' of size $\leq s'$ such that

- (i) $A'(x_1, \dots, x_k) \in \{f'(x_1, \dots, x_k), \perp\}$ for all x_1, \dots, x_k , and
- (ii) $\Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) = \perp] \leq 1 - \epsilon$.

We give a nonuniform reduction that makes $\frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta}$ nonadaptive queries to A' and combines the results with some trivial computation, yielding a circuit A that witnesses that f is not (s, δ, b) -hard. To start out, we give a randomized algorithm (Algorithm 2) that solves f errorlessly using oracle access to A' and oracle access to f and $\sigma(C, \cdot, b)$. The oracle queries to f and $\sigma(C, \cdot, b)$ only depend on the randomness (and not on the input), and later we will hard-wire a particular choice of randomness to get a circuit without oracle access to f or $\sigma(C, \cdot, b)$.

Input: $x \in \{0, 1\}^n$
Output: $f(x)$ or \perp

- 1 **repeat** $\frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta}$ times
- 2 pick $i \in [k]$ and $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k \in \{0, 1\}^n$ uniformly at random
- 3 if $A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp$ and $i \in \sigma(C, y, b)$ where

$$y = (f(x_1), \dots, f(x_{i-1}), b, f(x_{i+1}), \dots, f(x_k))$$
 then halt and output $A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k)$
- 4 **end**
- 5 halt and output \perp

Algorithm 2: Reduction for Theorem 2

Define the good set

$$G = \left\{ x \in f^{-1}(b) : \Pr_{i, x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k} [A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp \text{ and } i \in \sigma(C, y, b)] \geq \epsilon t / 4k \right\}$$

where y is as in line 3 of Algorithm 2, and define the bad set $B = f^{-1}(b) \setminus G$. That is, G is the set of inputs in $f^{-1}(b)$ for which each iteration of the loop has at least an $\epsilon t / 4k$ probability of producing output.

Claim 2. $|B| \leq (\delta/2) \cdot |f^{-1}(b)|$.

Proof. Suppose for contradiction that $|B| > (\delta/2) \cdot |f^{-1}(b)|$. Let $\gamma = |B|/2|f^{-1}(b)|$. We define the event

$$W = \left\{ (x_1, \dots, x_k) \in \{0, 1\}^{n \times k} : \left| \left\{ i : x_i \in B \text{ and } i \in \sigma(C, f^k(x_1, \dots, x_k), b) \right\} \right| \geq \gamma t \right\}.$$

That is, W is the event that at least a $\gamma t/k$ fraction of coordinates are both bad and b -sensitive. We have

$$\begin{aligned} \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp] &\leq \Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W] + \\ &\quad \Pr_{x_1, \dots, x_k} [A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W]. \end{aligned}$$

We show that both terms on the right side are $< \epsilon/2$, thus contradicting property (ii) of A' .

Bounding the first term. We have

$$\begin{aligned} \Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W] &\leq \Pr_{x_1, \dots, x_k} [|\sigma(C, f^k(x_1, \dots, x_k), b)| < t] + \\ &\quad \Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W \mid |\sigma(C, f^k(x_1, \dots, x_k), b)| \geq t]. \end{aligned}$$

To show that this is $< \epsilon/2$, we show that the first of the two terms on the right side is $\leq \epsilon/4$ and the second is $< \epsilon/4$. Since C is a (t, ρ, p, b) -amplifier, we have

$$\Pr_{x_1, \dots, x_k} [|\sigma(C, f^k(x_1, \dots, x_k), b)| < t] = \Pr_{y \sim_p \{0,1\}^k} [|\sigma(C, y, b)| < t] \leq \rho \leq \epsilon/4.$$

We have

$$\begin{aligned} &\Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W \mid |\sigma(C, f^k(x_1, \dots, x_k), b)| \geq t] \\ &= \mathbb{E}_{y \sim_p \{0,1\}^k} \left[\Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W \mid f^k(x_1, \dots, x_k) = y] \mid |\sigma(C, y, b)| \geq t \right]. \end{aligned}$$

Fix any $y \in \{0,1\}^k$ such that $|\sigma(C, y, b)| \geq t$, and for now let us abbreviate $\sigma(C, y, b)$ as σ . Then using $(x_i)_{i \notin \sigma}$ to denote the sequence of inputs x_i (chosen uniformly at random) for indices $i \notin \sigma$ (and similarly $(x_i)_{i \in \sigma}$), we have

$$\begin{aligned} &\Pr_{x_1, \dots, x_k} [(x_1, \dots, x_k) \notin W \mid f^k(x_1, \dots, x_k) = y] \\ &= \mathbb{E}_{(x_i)_{i \notin \sigma}} \left[\Pr_{(x_i)_{i \in \sigma}} [(x_1, \dots, x_k) \notin W \mid f(x_i) = b \ \forall i \in \sigma] \mid f(x_i) = y_i \ \forall i \notin \sigma \right] \end{aligned}$$

since $\sigma \subseteq \{i : y_i = b\}$. Now fix any $(x_i)_{i \notin \sigma}$ such that $f(x_i) = y_i$ for all $i \notin \sigma$. Then we have

$$\begin{aligned} &\Pr_{(x_i)_{i \in \sigma}} [(x_1, \dots, x_k) \notin W \mid f(x_i) = b \ \forall i \in \sigma] \\ &= \Pr_{(x_i)_{i \in \sigma}} \left[|\{i \in \sigma : x_i \in B\}| < \gamma t \mid f(x_i) = b \ \forall i \in \sigma \right] \\ &\leq \Pr_{(x_i)_{i \in \sigma}} \left[|\{i \in \sigma : x_i \in B\}| < \gamma \cdot |\sigma| \mid f(x_i) = b \ \forall i \in \sigma \right] \end{aligned}$$

where the inequality follows by $t \leq |\sigma|$. Applying Theorem 6 with X_j as the indicator variable for $x_i \in B$ where i is the j th value in σ and x_i is chosen uniformly from $f^{-1}(b)$, and with $\tau = |\sigma|$ and $\mu = |B|/|f^{-1}(b)|$, we have that the latter quantity is less than

$$e^{-|\sigma| \cdot |B|/8|f^{-1}(b)|} < e^{-|\sigma| \cdot \delta/16} \leq e^{-t\delta/16} \leq \epsilon/4$$

where the first inequality follows by our assumption on $|B|$, the middle inequality follows by $|\sigma| \geq t$, and the last inequality follows by $t \geq \frac{16}{\delta} \ln \frac{4}{\epsilon}$. This establishes that

$$\Pr_{x_1, \dots, x_k} \left[(x_1, \dots, x_k) \notin W \mid |\sigma(C, f^k(x_1, \dots, x_k), b)| \geq t \right] < \epsilon/4.$$

Bounding the second term. This is similar to the corresponding part of the analysis in the proof of Theorem 1. For each $S \subseteq [k]$ we define the event

$$W_S = \left\{ (x_1, \dots, x_k) \in \{0, 1\}^{n \times k} : \forall i \left(x_i \in B \text{ and } i \in \sigma(C, f^k(x_1, \dots, x_k), b) \right) \Leftrightarrow i \in S \right\}.$$

Note that the W_S 's are disjoint and

$$W = \bigcup_{S : |S| \geq \gamma t} W_S.$$

Using the shorthand y as in line 3 of Algorithm 2, we have

$$\begin{aligned} & \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W \right] \\ &= \sum_{S : |S| \geq \gamma t} \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S \right] \\ &\leq \frac{1}{\gamma t} \sum_{S \subseteq [k]} |S| \cdot \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S \right] \\ &= \frac{1}{\gamma t} \sum_{i \in [k]} \sum_{S \ni i} \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } (x_1, \dots, x_k) \in W_S \right] \\ &= \frac{1}{\gamma t} \sum_{i \in [k]} \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } x_i \in B \text{ and } i \in \sigma(C, f^k(x_1, \dots, x_k), b) \right] \\ &= \frac{1}{\gamma t} \sum_{i \in [k]} \sum_{x \in B} \Pr_{x_1, \dots, x_k} \left[A'(x_1, \dots, x_k) \neq \perp \text{ and } i \in \sigma(C, f^k(x_1, \dots, x_k), b) \mid x_i = x \right] \cdot \Pr_{x_1, \dots, x_k} [x_i = x] \\ &= \frac{1}{\gamma t 2^n} \sum_{x \in B} \sum_{i \in [k]} \Pr_{x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k} \left[A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp \text{ and } i \in \sigma(C, y, b) \right] \\ &= \frac{1}{\gamma t 2^n} \sum_{x \in B} k \cdot \Pr_{i, x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k} \left[A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp \text{ and } i \in \sigma(C, y, b) \right] \\ &< \frac{1}{\gamma t 2^n} \sum_{x \in B} k \cdot \epsilon t / 4k \\ &= \frac{\epsilon/4}{\gamma 2^n} \cdot |B| \\ &\leq \epsilon/2 \end{aligned}$$

where the second and fifth lines follow by the disjointness of the W_S 's, the last line follows by $|f^{-1}(b)| \leq 2^n$, and the remaining lines follow by simple rearrangements. For the seventh line, we used the fact that $x \in B$ implies $f(x) = b$ and thus $y = f^k(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k)$. \square

We now finish the proof of Theorem 2. First we note that for all $x \in \{0, 1\}^n$ and all choices of randomness, Algorithm 2 does indeed output either $f(x)$ or \perp . This follows trivially from the facts that $A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) \neq \perp$ implies

$$A'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) = f'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k)$$

by property (i) of A' , and $i \in \sigma(C, y, b)$ implies $f'(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) = f(x)$, where y is as in line 3 of Algorithm 2. Next we observe that for each $x \in G$, we have

$$\Pr_{\text{randomness}} [\text{Algorithm 2 outputs } \perp] \leq (1 - \epsilon t / 4k)^{\frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta}} \leq \delta / 2.$$

Therefore, picking $x \in f^{-1}(b)$ uniformly at random, we have

$$\begin{aligned} \Pr_{x, \text{randomness}} [\text{Algorithm 2 outputs } \perp] &\leq \Pr_x [x \in B] + \\ &\quad \mathbb{E}_x \left[\Pr_{\text{randomness}} [\text{Algorithm 2 outputs } \perp] \mid x \in G \right] \\ &\leq \delta / 2 + \mathbb{E}_x [\delta / 2 \mid x \in G] \\ &= \delta \end{aligned}$$

where the second inequality follows by Claim 2 and by the above observation. It follows that there exists a setting of the randomness such that

- (i) Algorithm 2 outputs $f(x)$ or \perp for all $x \in \{0, 1\}^n$, and
- (ii) $\Pr_x [\text{Algorithm 2 outputs } \perp] \leq \delta$ where $x \in f^{-1}(b)$ is chosen uniformly at random.

To get a circuit A that witnesses that f is not (s, δ, b) -hard, just hard-wire the randomness and the correct responses to the $\sigma(C, \cdot, b)$ queries (which only depend on the randomness and not on x), and plug in the hypothesized circuit A' . In fact, the iterations for which $i \notin \sigma(C, y, b)$ for this particular choice of randomness can simply be eliminated. Since A' has size $\leq s'$ and Algorithm 2 makes $\leq \frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta}$ queries to A' , A has size $\leq s' \cdot \frac{k}{t} \cdot \frac{4}{\epsilon} \ln \frac{2}{\delta} = s$. Note that Algorithm 2 can trivially be implemented with *nonadaptive* access to A' .

4 Proof of Theorem 5

Let (Enc, Dec) be an $(n, n', \delta, \epsilon, \alpha)$ -black-box non-errorless to errorless hardness amplification with $\delta < 1/8$ and $\epsilon \geq 1/1.01^{1.01^{2^n}}$. We use the notation $\Delta(f_1, f_2)$ for the relative Hamming distance between bit strings f_1 and f_2 . We begin with a completely standard claim that asserts the existence of a good error-correcting code. We include the proof for completeness.

Claim 3. *There exists an $F \subseteq \{0, 1\}^{2^n}$ such that*

- (1) $\forall f_1, f_2 \in F : \text{if } f_1 \neq f_2 \text{ then } \Delta(f_1, f_2) > 2\delta$, and
- (2) $|F| = 4 \log_3 \frac{1}{\epsilon}$.

Proof. Pick $F \subseteq \{0, 1\}^{2^n}$ randomly by choosing $f_1, \dots, f_{4 \log_3 \frac{1}{\epsilon}} \in \{0, 1\}^{2^n}$ independently uniformly at random and setting $F = \{f_1, \dots, f_{4 \log_3 \frac{1}{\epsilon}}\}$. To prove (1) and (2), it suffices to show that

$$\Pr [\exists i_1, i_2 \in \{1, \dots, 4 \log_3 \frac{1}{\epsilon}\} : i_1 \neq i_2 \text{ and } \Delta(f_{i_1}, f_{i_2}) \leq 2\delta] < 1.$$

For each pair $i_1 \neq i_2$, since $\delta < 1/8$ we have

$$\Pr [\Delta(f_{i_1}, f_{i_2}) \leq 2\delta] \leq \Pr [\Delta(f_{i_1}, f_{i_2}) < 1/4] < e^{-2^n/16}$$

by applying Theorem 6 with X_j as the indicator variable for $f_{i_1}(x) \neq f_{i_2}(x)$ where x is the j th string in $\{0, 1\}^n$, and with $\tau = 2^n$ and $\mu = 1/2$. By a union bound, the probability in question is at most $(4 \log_3 \frac{1}{\epsilon})^2 \cdot e^{-2^n/16} < 1$ since $\epsilon \geq 1/1.01^{1.01^{2^n}}$. \square

For the rest of the proof of Theorem 5 we fix a set F as in Claim 3.

Claim 4. *There exists an $E \subseteq F$ and an $A' \in \{0, 1, \perp\}^{2^{n'}}$ such that*

- (i) $\forall f \in E : A'(x') \in \{f'(x'), \perp\}$ for all $x' \in \{0, 1\}^{n'}$, where $f' = \text{Enc}(f)$, and
- (ii) $\Pr_{x'}[A'(x') = \perp] \leq 1 - \epsilon$, and
- (iii) $|E| = \log_3 \frac{1}{\epsilon}$.

Proof. For each $x' \in \{0, 1\}^{n'}$ let $m_{x'} = \text{majority}_{f \in F} \text{Enc}(f)(x')$, breaking a tie arbitrarily, and define

$$M_{x'} = \{f \in F : \text{Enc}(f)(x') = m_{x'}\}.$$

For any set $E \subseteq F$, define

$$M_E^{-1} = \{x' \in \{0, 1\}^{n'} : E \subseteq M_{x'}\}.$$

We construct a sequence of sets $E_0 \subseteq E_1 \subseteq E_2 \subseteq \dots \subseteq E_{\log_3 \frac{1}{\epsilon}} \subseteq F$ such that for all i , $|E_i| = i$ and $|M_{E_i}^{-1}| \geq 2^{n'}/3^i$. Then we can take $E = E_{\log_3 \frac{1}{\epsilon}}$ and

$$A'(x') = \begin{cases} m_{x'} & \text{if } x' \in M_E^{-1} \\ \perp & \text{otherwise} \end{cases}$$

and (i), (ii), and (iii) all follow immediately. We do the construction inductively. The base case $i = 0$ is trivial. Now assume $i \in \{0, 1, \dots, \log_3 \frac{1}{\epsilon} - 1\}$ and we have a set $E_i \subseteq F$ with $|E_i| = i$ and $|M_{E_i}^{-1}| \geq 2^{n'}/3^i$. For each $x' \in M_{E_i}^{-1}$, since $i \leq |F|/4$ we have

$$\Pr_{f \in F \setminus E_i} [f \in M_{x'}] = \frac{|M_{x'} \setminus E_i|}{|F \setminus E_i|} \geq \frac{\frac{1}{2}|F| - i}{|F| - i} \geq 1/3$$

where f is chosen uniformly at random. Thus for some $f \in F \setminus E_i$ we have $\Pr_{x' \in M_{E_i}^{-1}} [f \in M_{x'}] \geq 1/3$ where x' is chosen uniformly at random. For this fixed f , setting $E_{i+1} = E_i \cup \{f\}$ we have $|E_{i+1}| = |E_i| + 1 = i + 1$ and

$$|M_{E_{i+1}}^{-1}| = |\{x' \in M_{E_i}^{-1} : f \in M_{x'}\}| = |M_{E_i}^{-1}| \cdot \Pr_{x' \in M_{E_i}^{-1}} [f \in M_{x'}] \geq |M_{E_i}^{-1}|/3 \geq 2^{n'}/3^{i+1}.$$

This finishes the induction step. \square

Now to prove Theorem 5, suppose for contradiction that $2^\alpha < \log_3 \frac{1}{\epsilon}$. Then by the pigeonhole principle there must exist $f_1, f_2 \in E$ such that $f_1 \neq f_2$ and the advice string corresponding to f_1 and A' equals the advice string corresponding to f_2 and A' . Call this advice string a , and let $A = \text{Dec}(A', a)$. By Definition 5, we must have $\Delta(A, f_1) \leq \delta$ and $\Delta(A, f_2) \leq \delta$. But this is impossible because $\Delta(f_1, f_2) > 2\delta$ by property (1) in Claim 3.

5 Open Problems

There remains a multiplicative $\log \frac{1}{\delta}$ gap between the known upper and lower bounds on query complexity for adaptive black-box reductions for errorless hardness amplification.

It is an open problem to prove some sort of uniform version of the Errorless XOR Lemma. Impagliazzo et al. [12] proved a sort of uniform version of the (non-errorless) XOR Lemma, but their techniques do not seem to apply to the errorless setting.

Another open problem concerns a version of the Errorless XOR Lemma where k different mildly errorless average-case hard functions are XORed together. The original Bogdanov-Safra proof can be used to show that this construction does yield a strongly errorless average-case hard function, but our proof does not seem to apply, so it is not known if this more general hardness amplification can be done with query complexity $O(\frac{1}{\epsilon} \log \frac{1}{\delta})$.

It remains open to construct polynomial-time computable $(t, \rho, 1/2, b)$ -amplifiers that achieve near-optimal tradeoff between t and ρ . (The optimal tradeoff would be $t/\rho = k^{1/2}$.) The functions HOL-AMP_α (for $b = 0$) and $\text{HOL-AMP}_\alpha^\dagger$ (for $b = 1$) are suboptimal (with $t/\rho \approx k^{1/3}$). The function OR-DDNF_α (for $b = 0$) is near-optimal, but not known to be polynomial-time computable (though it is in NP, which is good enough for hardness amplification within NP). For $b = 1$, it is not even known how to achieve near-optimal tradeoff with a function in NP.

Acknowledgments

I thank Salil Vadhan and anonymous reviewers for helpful comments.

References

- [1] S. Artemenko and R. Shaltiel. Lower bounds on the query complexity of non-uniform and adaptive reductions showing hardness amplification. In *Proceedings of the 15th International Workshop on Randomization and Computation*, pages 377–388, 2011.
- [2] I. Benjamini, O. Schramm, and D. Wilson. Balanced boolean functions that can be evaluated so that every input bit is unlikely to be read. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 244–250, 2005.
- [3] A. Bogdanov and M. Safra. Hardness amplification for errorless heuristics. In *Proceedings of the 48th IEEE Symposium on Foundations of Computer Science*, pages 418–426, 2007.
- [4] J. Buresh-Oppenheim, V. Kabanets, and R. Santhanam. Uniform hardness amplification in NP via monotone codes. Technical Report TR06-154, Electronic Colloquium on Computational Complexity, 2006.

- [5] G. Cohen, S. Litsyn, and G. Zémor. Upper bounds on generalized distances. *IEEE Transactions on Information Theory*, 40(6):2090–2092, 1994.
- [6] O. Goldreich, N. Nisan, and A. Wigderson. On Yao’s XOR-Lemma. *Studies in Complexity and Cryptography*, pages 273–301, 2011.
- [7] P. Gopalan and V. Guruswami. Hardness amplification within NP against deterministic algorithms. *Journal of Computer and System Sciences*, 77(1):107–121, 2011.
- [8] V. Guruswami. Chapter 10: List decoding from erasures. In *List Decoding of Error-Correcting Codes (Lecture Notes in Computer Science)*, volume 3282. Springer-Verlag, 2005.
- [9] V. Guruswami and S. Vadhan. A lower bound on list size for list decoding. *IEEE Transactions on Information Theory*, 56(11):5681–5688, 2010.
- [10] A. Healy, S. Vadhan, and E. Viola. Using nondeterminism to amplify hardness. *SIAM Journal on Computing*, 35(4):903–931, 2006.
- [11] R. Impagliazzo. Hard-core distributions for somewhat hard problems. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 538–545, 1995.
- [12] R. Impagliazzo, R. Jaiswal, V. Kabanets, and A. Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. *SIAM Journal on Computing*, 39(4):1637–1665, 2010.
- [13] R. Impagliazzo and A. Wigderson. P=BPP unless E has subexponential circuits: Derandomizing the XOR Lemma. In *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 220–229, 1997.
- [14] A. Klivans and R. Servedio. Boosting and hard-core sets. *Machine Learning*, 53(3):217–238, 2003.
- [15] L. Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.
- [16] L. Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.
- [17] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. Improved hardness amplification in NP. *Theoretical Computer Science*, 370(1-3):293–298, 2007.
- [18] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. On the complexity of hardness amplification. *IEEE Transactions on Information Theory*, 54(10):4575–4586, 2008.
- [19] R. O’Donnell. Hardness amplification within NP. *Journal of Computer and System Sciences*, 69(1):68–94, 2004.
- [20] R. Shaltiel and E. Viola. Hardness amplification proofs require majority. *SIAM Journal on Computing*, 39(7):3122–3154, 2010.
- [21] M. Talagrand. How much are increasing sets positively correlated? *Combinatorica*, 16(2):243–258, 1996.

- [22] L. Trevisan. List decoding using the XOR Lemma. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science*, pages 126–135, 2003.
- [23] L. Trevisan. On uniform amplification of hardness in NP. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 31–38, 2005.