Is It Possible to Improve Yao's XOR Lemma Using Reductions That Exploit the Efficiency of Their Oracle?

Ronen Shaltiel
University of Haifa, Israel
https://cs.haifa.ac.il/~ronen/
ronen@cs.haifa.ac.il

Abstract

Yao's XOR lemma states that for every function $f : \{0,1\}^k \to \{0,1\}$, if f has hardness $2/3$ for $P/poly$ (meaning that for every circuit C in $P/poly$, $\Pr[C(X) = f(X)] \leq 2/3$ on a uniform input X), then the task of computing $f(X_1) \oplus \ldots \oplus f(X_t)$ for sufficiently large t has hardness $\frac{1}{2} + \epsilon$ for $P/poly$.

Known proofs of this lemma cannot achieve $\epsilon = \frac{1}{k^{\omega(1)}}$, and even for $\epsilon = \frac{1}{k}$, we do not know how to replace $P/poly$ by $\mathrm{AC}^0[\mathrm{PARITY}]$ (the class of constant depth circuits with the gates $\{\mathrm{AND},\mathrm{OR},\mathrm{NOT},\mathrm{PARITY}\}$ of unbounded fan-in).

Recently, Grinberg, Shaltiel and Viola (FOCS 2018) (building on a sequence of earlier works) showed that these limitations cannot be circumvented by *black-box reductions*. Namely, by reductions $\mathrm{Red}^{(\cdot)}$ that given oracle access to a function D that violates the conclusion of Yao's XOR lemma, implement a circuit that violates the assumption of Yao's XOR lemma.

There are a few known reductions in the related literature on worst-case to average case reductions that are *non-black box*. Specifically, the reductions of Gutfreund, Shaltiel and Ta Shma (Computational Complexity 2007) and Hirahara (FOCS 2018)) are "class reductions" that are only guaranteed to succeed when given oracle access to an oracle D from some efficient class of algorithms. These works seem to circumvent some black-box impossibility results.

In this paper we extend the previous limitations of Grinberg, Shaltiel and Viola to class reductions, giving evidence that class reductions cannot yield the desired improvements in Yao's XOR lemma. To the best of our knowledge, this is the first limitation on reductions for hardness amplification that applies to class reductions.

Our technique imitates the previous lower bounds for black-box reductions, replacing the inefficient oracle used in that proof, with an efficient one that is based on limited independence, and developing tools to deal with the technical difficulties that arise following this replacement.

2012 ACM Subject Classification Theory of computation → Circuit complexity

Keywords and phrases Yao's XOR lemma, Hardness amplification, black-box reductions

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2020.10

Category RANDOM

Funding *Ronen Shaltiel*: This research was supported by ISF grant 1628/17.

Acknowledgements We are grateful to Emanuele Viola for very helpful discussions, and to anonymous referees for comments and suggestions.

1 Introduction

Yao's XOR Lemma is a fundamental and celebrated result in complexity theory, that is extensively studied (from various aspects) and has found many applications. See [9] for a survey article.

▶ **Definition 1** (The XOR function). *Given a $f : \{0,1\}^k \to \{0,1\}$, and a number t, we define $f^{\oplus t} : \{0,1\}^{t \cdot k} \to \{0,1\}$, as follows: Given $y \in \{0,1\}^{t \cdot k}$, we view y as $(y_1, \ldots, y_t) \in (\{0,1\}^k)^t$, and define:*

$$f^{\oplus t}(y) = f(y_1) \oplus \ldots \oplus f(y_t)$$

Let U_k denote the uniform distribution on k bit strings. Loosely speaking, Yao's XOR lemma says that if a function f is "mildly hard one average" on input $X \leftarrow U_k$, then as t increases, computing $f^{\oplus t}$ on input $Y \leftarrow U_{tk}$, becomes "very hard on average".

▶ **Lemma 2** (Yao's XOR lemma, for poly-size circuits). *For every $f : \{0,1\}^k \to \{0,1\}$, and $t \leq \text{poly}(k)$ such that $t = \omega(\log k)$:*

If, for every $\text{poly}(k)$ size circuit C, $\Pr_{X \leftarrow U_k}[C(X) = f(X)] < \frac{2}{3}$,

Then, for every constant c, and every $\text{poly}(k)$ size circuit D, $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] < \frac{1}{2} + \frac{1}{k^c}$.[1]

One weakness of Yao's XOR lemma, is that it cannot be used to conclude a statement in which the "hardness on average" $\frac{1}{2} + \frac{1}{k^c}$ is replaced by $\frac{1}{2} + \frac{1}{k^{\omega(1)}}$. This holds, even if the number of repetitions t is increased from slightly larger than $\log k$ (as is the case in Lemma 2) to the maximal choice of $t = \text{poly}(k)$. Specifically, the following question is wide open:

▶ **Open problem 3** (Yao's XOR lemma for subpolynomial error?). *Is it true that for every $f : \{0,1\}^k \to \{0,1\}$, taking $t = \omega(\log k)$ (or even the maximal choice of $t = \text{poly}(k)$) it holds that:*

If, for every $\text{poly}(k)$ size circuit C, $\Pr_{X \leftarrow U_k}[C(X) = f(X)] < \frac{2}{3}$,

Then, for every $\text{poly}(k)$ size circuit D, $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] < \frac{1}{2} + \frac{1}{k^{\omega(1)}}$,

Another weakness of Yao's XOR Lemma is that known proofs fail to prove Yao's XOR lemma when replacing *P/poly* with many interesting constant depth circuit classes. An especially frustrating case is the class $\text{AC}^0[\text{PARITY}]$ of poly-size constant depth circuits over the gates $\{\text{AND,OR,NOT,PARITY}\}$ of unbounded fan-in. There are known lower bounds showing explicit functions that have hardness $\frac{2}{3}$ for $\text{AC}^0[\text{PARITY}]$ (or even $\frac{1}{2} - o(1)$ hardness for circuits of depth d and size $2^{k^{\Omega(1/d)}}$ [25, 28]) but lower bounds with hardness $\frac{1}{2} - \frac{1}{k}$ are unknown. This is a twenty five year old barrier that prevents us from "using the hybrid argument" when constructing pseudorandom generators for $\text{AC}^0[\text{PARITY}]$ (and related classes). This barrier limits the best known pseudorandom generators for $\text{AC}^0[\text{PARITY}]$ (and related classes) [6] to very poor seeds (See [6] for a discussion of this limitation). Specifically, the following question is wide open:

▶ **Open problem 4** (Yao's XOR lemma for constant depth circuits?). *Let G be the set of gates $\{\text{AND,OR,NOT,PARITY}\}$ of unbounded fan-in. Is it true that for every $f : \{0,1\}^k \to \{0,1\}$, taking $t = \omega(\log k)$ (or even the maximal choice of $t = \text{poly}(k)$) it holds that:*

If, for every $\text{poly}(k)$ size, constant-depth circuit C with gates in G, $\Pr_{X \leftarrow U_k}[C(X) = f(X)] < \frac{2}{3}$,

Then, for every $\text{poly}(k)$ size, constant-depth circuit D, with gates in G, $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] < \frac{1}{2} + \frac{1}{k}$,

[1] Naturally, in order to make this asymptotic statement precise, one needs to consider an infinite sequence of functions $\{f_k\}$ with growing input length (so that terms like "poly-size", "$\omega(\log k)$", and "constant" are well defined). We allow ourselves to be imprecise, as a more general, and quantitatively precise statement of Yao's XOR lemma is given below in Lemma 5.

1.1 Proofs of Yao's Lemma as (nonuniform) black-box reductions

Before discussing the best known proofs of Yao's XOR Lemma, let us state the lemma more precisely, in a more general and quantitative form. The next formulation is achieved using Impagliazzo's proof of Yao's XOR lemma [18, 9] together with the quantitative improvement of Klivans and Servedio [21] of Impagliazzo's hard-core lemma [18].

▶ **Lemma 5** (Yao's XOR lemma, General version). *There exist a constant c, and a polynomial p, such that for every $f : \{0,1\}^k \to \{0,1\}$, every $\epsilon, \delta > 0$, and every $t \geq c \cdot \frac{\log(1/\epsilon)}{\delta}$, setting $q = c \cdot \frac{\log(1/\delta)}{\epsilon^2}$, we have that:*

If, for every circuit C of size $s \geq p(t,k,q)$, $\Pr_{X \leftarrow U_k}[C(X) = f(X)] < 1 - \delta$,

Then, for every circuit D of size $s' = \frac{s}{q}$, $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] < \frac{1}{2} + \epsilon$.

The special case of Lemma 2 is obtained by taking s to be a polynomial in k, and $\delta = \frac{1}{3}$. In order to reduce the number of live parameters, we recommend that the reader focuses on these choices on a first reading. We point out that s' (which is the size of D) is smaller by a factor of $q = \Omega(\frac{1}{\epsilon^2})$, than s (which is the size of C). This implies that $s' \leq O(\epsilon^2 \cdot s)$, implying that $\epsilon \geq \Omega(\frac{1}{\sqrt{s}})$, and it is impossible to get $\epsilon < \frac{1}{s}$ with current proofs. (This is a more quantitative way to state the phenomenon in open problem 3).

All known proofs of Yao's XOR lemma work by *reduction*. That is, the proof shows a reduction that transforms a circuit D such that $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$, into a circuit C such that $\Pr_{X \leftarrow U_k}[C(X) = f(X)] \geq 1 - \delta$. All known proofs are "nonuniform black-box reductions", meaning that they provide a reduction (namely an oracle circuit $\mathrm{Red}^{(\cdot)}(x, \alpha)$ where x is an input, and α is an "advice string") and the circuit C is obtained by $C(x) = \mathrm{Red}^D(x, \alpha)$ where α is an "nonuniform advice string" that may depend on f and D.[2] This is made precise in the next definition.

▶ **Definition 6** (Nonuniform black-box reduction for Yao's XOR lemma). *Let $\epsilon, \delta > 0$, and let k, t, a be integers. A $(\frac{1}{2} + \epsilon) \to (1 - \delta)$* black-box reduction for Yao's XOR lemma *(with input length k, t repetitions and advice length a) is an oracle circuit $\mathrm{Red}^{(\cdot)}(x, \alpha)$, where $x \in \{0,1\}^k$ and $\alpha \in \{0,1\}^a$, such that for every $f : \{0,1\}^k \to \{0,1\}$, the following holds:*

For every function $D : \{0,1\}^{tk} \to \{0,1\}$, such that $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$,

there exists $\alpha \in \{0,1\}^a$, such that $\Pr_{X \leftarrow U_k}[\mathrm{Red}^D(X, \alpha) = f(X)] \geq 1 - \delta$.

The version of Yao's XOR lemma stated in Lemma 5, follows by showing the following reduction:

▶ **Lemma 7** (Known black-box reductions for Yao's XOR lemma). *There exist a constant c, and a polynomial p, such that for every integer k, every $\epsilon, \delta > 0$ such that $1 - \delta > \frac{1}{2} + \epsilon$, and every $t \geq c \cdot \frac{\log(1/\epsilon)}{\delta}$, there is a $(\frac{1}{2} + \epsilon) \to (1 - \delta)$ black-box reduction $\mathrm{Red}^{(\cdot)}(x, \alpha)$ for Yao's XOR lemma with input length k, t repetitions, and advice length a such that:*

- *R makes at most $q = c \cdot \frac{\log(1/\delta)}{\epsilon^2}$ queries to its oracle.*
- *R is an oracle circuit of size $r = p(t,k,q)$, (and in particular, $a \leq r$).*
- *R is an oracle circuit of constant depth d over the gates $\{AND, OR, NOT\}$ of unbounded fan-in and also uses one majority gate with fan-in q.*

[2] There is a formal connection between "black box hardness amplification" and list-decodable error correcting code [29], see for example the discussion in [27, 10]. Using this connection, it is known that black-box reductions for Yao's XOR lemma, must be nonuniform and use an advice string if $1 - \delta > \frac{1}{2} + \epsilon$ and $\epsilon < \frac{1}{4}$.

In order to understand the limitations that prevent known proofs from solving the aforementioned open problems, it is instructive to see how Lemma 5 follows from Lemma 7. Specifically, assume (for contradiction) that Lemma 5 does not hold and let D be a circuit of size s' that is violating the conclusion. By Lemma 7, there exists $\alpha \in \{0,1\}^a$, such that the circuit $C(x) = \mathrm{Red}^D(x,\alpha)$ computes $f(X)$ with success $1-\delta$ on $X \leftarrow U_k$. The size of C is bounded by $s = r + a + q \cdot s' \geq q \cdot s'$, and the obtained circuit C has depth at least d, and needs to compute majority on q bits. Summing up:

- The number of queries q made by the reduction is a lower bound on $\frac{s}{s'}$, meaning that $s' \leq \frac{s}{q}$ and as the known reductions have $q \geq \frac{1}{\epsilon^2}$ we cannot expect $\epsilon < \frac{1}{\sqrt{s}}$, and cannot solve open problem 3.
- The fact that the best known reductions requires a majority gate on $q \geq \frac{1}{\epsilon}$ inputs, means that we need to assume hardness against a class that can perform this computation. For $\epsilon = 1/k$, Razborov's lower bound [25] (see also [24]) shows that for every depth d', majority on k bits, cannot be computed by circuits of depth d' and size $2^{k^{\Omega(1/d')}}$ over the gates $\{\text{AND},\text{OR},\text{NOT},\text{PARITY}\}$, explaining why current reductions cannot solve open problem 4.

Limitations on black-box reductions

A sequence of works [32, 27, 12, 3, 2] culminating in [10], shows that known black-box reductions for Yao's XOR lemma must suffer from the limitations above: They require $q = \Omega(\frac{\log(1/\delta)}{\epsilon^2})$ queries, and require computing majority on input length $\Omega(\frac{1}{\epsilon})$.[3]

1.2 Class reductions

On a closer examination, black-box reductions seem to be an overkill for the task of proving Yao's XOR lemma. For proving Yao's XOR lemma, we don't need $\mathrm{Red}^{(\cdot)}$ to succeed given oracle access to *every* function D such that $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$. It is sufficient that Red succeeds only given oracle access to functions D that are *efficiently computable* and belong to the class $\mathcal{D}$ of circuits with size s' (if we're in the setup of open problem 3) and size s' with constant depth (if we're in the setup of open problem 4).

This motivates a notion of *class reduction* (suggested for example in [14]) in which reductions are only required to succeed if given oracle access to a function D that belongs to some class $\mathcal{D}$ of "efficient circuits", and do not need to succeed when given oracle access to a function D that does not belong to $\mathcal{D}$. The definition of *class $\mathcal{D}$ reduction* below is identical to definition 6 with the single exception (that is underlined for emphasis) being that we only require the reduction to succeed when given oracle access to a function D from the class $\mathcal{D}$.

▶ **Definition 8** (Nonuniform class reduction for Yao's XOR lemma)**.** *Let $\epsilon, \delta > 0$ and let k, t, a be integers, and let $\mathcal{D}$ be some class of functions $D : \{0,1\}^{tk} \to \{0,1\}$. A $(\frac{1}{2}+\epsilon) \to (1-\delta)$* class $\mathcal{D}$ reduction for Yao's XOR lemma *(with input length k, t repetitions and advice length a) is an oracle circuit $\mathrm{Red}^{(\cdot)}(x,\alpha)$, where $x \in \{0,1\}^k$ and $\alpha \in \{0,1\}^a$, such that for every $f : \{0,1\}^k \to \{0,1\}$, the following holds:*

For every function $D : \{0,1\}^{tk} \to \{0,1\}$ in the class $\mathcal{D}$, such that $\Pr_{Y \leftarrow U_{tk}}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$,

there exists $\alpha \in \{0,1\}^a$, such that $\Pr_{X \leftarrow U_k}[\mathrm{Red}^D(X,\alpha) = f(X)] \geq 1 - \delta$.

[3] More formally, saying that $\mathrm{Red}^{(\cdot)}$ "requires computing majority on input length $\Omega(1/\epsilon)$" means that every such reduction $\mathrm{Red}^{(\cdot)}$ can be transformed into a circuit (with no oracle) of roughly the same size and depth as $\mathrm{Red}^{(\cdot)}$ for computing the majority function on inputs of length $\Omega(\frac{1}{\epsilon})$.

Note that a black-box reduction is a special case of a class reduction where $\mathcal{D}$ is the class of all boolean functions on tk bits. This raises the following questions:

1. Are there reductions in the literature that are class reductions but not black-box reductions?
2. Can class reductions circumvent the limitations on black-box reductions and solve open problem 3 or open problem 4?

The answer to the first question is affirmative in the sense that there are at least two examples that we are aware of, where a worst-case to average case amplification is proven by a reduction that is not black-box. Furthermore, in both cases, the reduction is a class reduction, and there is strong evidence that it cannot be made black-box.

The first example is a worst-case hardness to average case hardness tradeoff for SAT (with respect to a distribution sampled in quasipolynomial time) by Shaltiel, Gutfreund and Ta-Shma [13] (see also a related work [4, 11, 14]). The correctness of the reduction of [13] relies on the efficiency of the oracle and the term "class reduction" was suggested by Gutfreund and Ta-Shma [14]. It was also argued in [14] that limitations on black-box reductions proven by Bogdanov and Trevisan [5] can be extended to the scenario studied in [13], and show that if the class reduction of [13] (which is non-adaptive) could be made also black-box, then co-NP has nondeterministic circuits of quasipolynomial size.

Another example is Hirahara's recent worst-case to average case reductions for variants of MCSP and MINKT [17]. These reductions are non-black-box, in the sense that their correctness relies on the efficiency of their oracle. The aforementioned work of Bogdanov and Trevisan [5] shows that if these reductions can be made black-box, then these problems are in co-NP/poly, which is not known, and is false, if these problems are NP-complete. See [17] for an elaborate discussion of consequences of the existence of such black-box reductions.

1.3 Our results: limitations on class reductions for Yao's XOR lemma

In this paper we give evidence that the answer to the second question above is negative. We extend the aforementioned limitations of [10] on black-box reductions for Yao's XOR lemma to *class reductions* for any $\mathcal{D}$ of that contains circuits that have polynomial size and constant depth over the gates $\{\text{AND},\text{OR},\text{NOT},\text{PARITY}\}$ with unbounded fan-in. To the best of our knowledge, this is the first example of proving limitations on class reductions in this setup. Our results are stated formally in the next theorem.[4]

▶ **Theorem 9** (Limitations on class reductions for Yao's XOR lemma). *There exist constants $\delta_0 > 0$, $\nu > 0$, $d_0 > 1$ and a polynomial p such that: Let $\mathrm{Red}^{(\cdot)}(x,\alpha)$ be a $(\frac{1}{2}+\epsilon) \to (1-\delta)$ class $\mathcal{D}$ reduction for Yao's XOR lemma, with input length k, t repetitions and advice length a. Assume that:*

- *$\mathrm{Red}^{(\cdot)}$ is a size r oracle circuit, that makes at most q queries.*
- *The class $\mathcal{D}$ contains circuits of size $p(r)$ and depth d_0 over the gates $\{\text{AND},\text{OR},\text{NOT},\text{PARITY}\}$ of unbounded fan-in.*
- *$t, a, \frac{1}{\epsilon}, \frac{1}{\delta} \leq r \leq 2^{\nu \cdot k}$ and $\delta \leq \delta_0$.*

Then the following holds:

[4] We remark that any circuit of size r over the gates $\{\text{AND},\text{OR},\text{NOT},\text{PARITY}\}$ with unbounded fan-in, cannot use fan-in larger than r, and therefore can be simulated by a circuit of size $O(r^2)$ over the standard gates $\{\text{AND},\text{OR},\text{NOT}\}$ with bounded fan-in. This allows us to state our results in a way that captures both circuits of small depth (using gates with unbounded fan-in) and circuits that use the standard gates with bounded fan-in.

- $\text{Red}^{(\cdot)}$ *requires many queries, specifically:* $q = \Omega(\frac{\log(1/\delta)}{\epsilon^2})$.
- $\text{Red}^{(\cdot)}$ *requires majority, specifically: if in addition to the size restriction on* Red*, we also have that* $\text{Red}^{(\cdot)}$ *is an oracle circuit of depth* d *over a set of gates* G *that contains the gates* $\{$AND,OR,NOT,PARITY$\}$ *of unbounded fan-in, then the majority function over* $\Omega(\frac{1}{\epsilon})$ *bits can be computed by a circuit of size* $\text{poly}(r)$ *and depth* $O(d)$ *over the set of gates* G.

What kind of reductions are ruled out by this result?

Theorem 9 achieves exactly the same limitations on *class reductions* for Yao's XOR lemma as the limitations of [10] for *black-box reductions*. This is achieved for any class $\mathcal{D}$ that contains small circuits with constant depth over the gates $\{$AND,OR,NOT,PARITY$\}$ of unbounded fan-in, which is exactly the classes that come up if one wants to use class reductions to solve open problems 3 and 4. This shows that a class reduction cannot circumvent the known limitations on black-box reductions, and additional ideas are needed for solving open problems 3 and 4.

More specifically, for the purpose of solving open problem 3 one wants a $(\frac{1}{2} + \frac{1}{k^{\omega(1)}}) \to \frac{2}{3}$ class $\mathcal{D}$ reduction $\text{Red}^{(\cdot)}$ for Yao's XOR lemma, of size $\text{poly}(k)$, for the class $\mathcal{D}$ of all circuits of size $\text{poly}(k)$. This is ruled out by our lower bound on the number of queries. For the purpose of solving open problem 4 one wants a $(\frac{1}{2} + \frac{1}{k}) \to \frac{2}{3}$ class $\mathcal{D}$ reduction for Yao's XOR lemma, of size $\text{poly}(k)$ and constant depth over the gates $G = \{$AND,OR,NOT,PARITY$\}$, for the class $\mathcal{D}$ of all circuits of size $\text{poly}(k)$ and constant depth over the gates G. This is ruled out by our results that $\text{Red}^{(\cdot)}$ requires majority on inputs of length $\Omega(k)$, and Razborov's lower bound [25] showing that this cannot be done by circuits of depth d' and size $2^{k^{\Omega(1/d)}}$.

A potential weakness of our impossibility results, is that they require that the class $\mathcal{D}$ has circuits of size larger than the reduction (although it is allowed that $\mathcal{D}$ contains only circuits of smaller depth than the reduction). This allows a scenario in which for every polynomial p_1, there exists a larger polynomial p_2 such that there is a $(\frac{1}{2} + \frac{1}{k^{\omega(1)}}) \to \frac{2}{3}$ class $\mathcal{D}_{p_1}$ reduction Red_{p_1} of size $p_2(k)$ for the class $\mathcal{D}_{p_1}$ of circuits of size $p_1(k)$ (but not for the class of circuits of size $p(p_2(k))$ where p is the polynomial in Theorem 9). This is sufficient for solving open problem 3 and is not ruled out by our impossibility results.

An optimistic view is that this may point us to the kind of reductions we need to design, in order to solve the aforementioned open problems. We remark however that the aforementioned reduction by Hirahara [17] *does not* need to assume that the oracle is weaker than the reduction. (The reduction of [13] involves a more complicated scenario where there is also a third entity which is the samplable distribution, and so, it is arguable whether the reduction is more powerful than the oracle).

Theorem 9 is weaker than the results of [10] in the sense that the limitations of [10] apply not only to Yao's XOR lemma, but to *any* hardness amplification technique. More precisely, in the results of [10] one can replace $f^{\oplus t}$ by any other function f' over $n = 2^{o(k)}$ bits, with the same limitations. Our approach cannot give such a general result, but can be extended as follows:

Extension to any efficient hardness amplification construction

Our results immediately extend to any function f' over $n = 2^{o(k)}$ bits such that f' can be efficiently computed given access to f. More precisely, in Theorem 9 one can replace occurrences of the parameter t by n, and the theorem extends to any function f' such that there exists an oracle circuit $\text{Con}^{(\cdot)}$ of size $\text{poly}(r)$, and constant depth over the gates $\{$AND,OR,NOT,PARITY$\}$ of unbounded fan-in, such that $f' = \text{Con}^f$. Moreover, if we omit the restriction that $\text{Con}^{(\cdot)}$ has constant depth, then the theorem holds with respect to any class $\mathcal{D}$ that contains circuits of size $p(r)$.

Extension to hardness amplification based on sufficiently explicit linear codes

Using ideas from [32], our results also extend to the case of $\delta = 2^{-2k}$ (which captures worst-case to average case hardness amplification) for functions f' over $n = 2^{o(k)}$ bits, such that:

$$f'(y) = \sum_{x \in \{0,1\}^k} f(x) \cdot g(x,y),$$

where the sum is taken in the field $\mathbb{F}_2$, and $g : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}$ can be computed by circuits of size $\text{poly}(r)$ and depth d over the set G of gates.

This definition of f' corresponds to "hardness amplification by a linear map". More specifically, we can view g as a matrix A of order $2^k \times 2^n$ over $\mathbb{F}_2$ by $A_{x,y} = g(x,y)$, and view the truth tables of the functions f, f' as vectors over $\mathbb{F}_2$ of dimension 2^k, 2^n, respectively. In this interpretation, the definition of f' above, says that $f' = f \cdot A$, for a matrix A in which the entry $A_{x,y}$ can be efficiently computed given x, y.

Many worst-case to average-case hardness amplification results in the literature choose f' so that the truth table f' is obtained by applying an error correcting code on the truth table of f. (It was observed in [29] that there is a formal connection between black-box reductions for hardness amplification, and list-decodable error correcting codes, see for example [10] for a discussion). Typical choices of this error correcting code are linear codes (most commonly Reed-Muller concatenated with Hadamard) and our results apply to this scenario, with the weaker conclusion that $q = \Omega(\frac{\log r}{\epsilon^2})$, and the same conclusion for the case of majority.

Perspective

Limitations for black-box reductions are extensively studied in various settings in complexity theory and cryptography. In order to prove impossibility results on black-box reduction, it is sufficient to show the existence of an oracle D (*that does not need to be efficient*) on which the reduction cannot succeed.

Many impossibility results and limitations in the literature strongly utilize the ability to choose an oracle D that is not efficient. One notable example is the aforementioned results of Bogdanov and Trevisan [5] (that build on earlier work of Feigenbaum and Fortnow [7]). Indeed, this is why these limitations do not apply to class reductions like the aforementioned results [13, 17].

This work puts an emphasis on whether or not the oracle D that one designs when showing a black-box impossibility result, can be made efficient, and demonstrates that achieving this, has the additional benefit of also ruling out class reductions.

1.4 Some more related work

It is beyond the scope of this paper to survey the vast literature on Yao's XOR lemma and hardness amplification. The reader is referred to [9] for a survey on Yao's XOR lemma, and to [32, 27, 10] for detailed discussions on the more general problem hardness amplification.

A significant advantage of Yao's XOR lemma (over some other suggested methods of hardness amplification) is that the "construction" $f' = f^{\oplus t}$ can be computed very efficiently, when given oracle access to f. A line of work (that is orthogonal to studying the complexity of *reductions* for hardness amplification) is interested in the complexity of *constructions* yielding hardness amplification. This line of work is mostly interested in starting from worst-case hard functions (which correspond to $\delta < 2^{-k}$) and aims to design (or prove impossibility results for) efficient constructions $\text{Con}^{(\cdot)}$ for which one can prove that if f has hardness $1 - \delta$, then $f' = \text{Con}^f$ has hardness $\frac{1}{2} + \epsilon$. (See e.g., [30, 31, 23, 15] for further discussion).

In this orthogonal line of work, there are examples of *constructions* which are non-black-box, and utilize specific properties of the function f (for example that $f \in \text{NP}$ or that f is a low degree polynomial). This is a different form of "non-black-box" than the one studied in this paper, and it is interesting to combine the two orthogonal directions.

There is a large body of work on proving black-box impossibility results in cryptography. This study was initiated by Impagliazzo and Rufich [19] and is concerned both with issues that are related to black-box constructions and to black-box reductions. See for example the discussion in Reingold, Trevisan and Vadhan [26] for a taxonomy of various notions.

2 Technique and a road map for proof

Our results are obtained by carefully examining the argument of the black-box impossibility result of [10], replacing the inefficient oracle with an efficient one, and handling the technical difficulties arising from this modification.

In this section we survey our technique, and give a roadmap of the proof of Theorem 9. We assume the setup of Theorem 9. Specifically, let $\text{Red}^{(\cdot)}(x,\alpha)$ be a $(\frac{1}{2}+\epsilon) \to (1-\delta)$ class $\mathcal{D}$ reduction for Yao's XOR lemma, with input length k, t repetitions and advice length a, which satisfies the requirements of the theorem. Let r be the size of Red and let d (which is not necessarily a constant) be the depth of Red. Our goal is to show that $\text{Red}^{(\cdot)}$ requires many queries, and that $\text{Red}^{(\cdot)}$ requires majority.

Let $f : \{0,1\}^k \to \{0,1\}$ be some function that we choose later, and let $n = tk$ be the input length of $f^{\oplus t}$. We start by surveying the approach of of the previous papers (which only handle *black-box* reductions rather than *class* reductions).

2.1 The approach of [32, 27]

We first introduce the following notation.

▶ **Definition 10** (Random sequences/functions). *For a number $0 \le p \le 1$, and an integer q, we define a distribution Noise^q_p over $\{0,1\}^q$ which consists of q i.i.d. bit variables $\mathsf{Noise}^q_p(1),\ldots,\mathsf{Noise}^p_q(q)$ where each of them has probability p to be one. This notation also allows us to view Noise^q_p as a distribution over functions from $[q]$ to $\{0,1\}$.*

Following [32, 27] (and as done in later works [12, 10]) our plan is to show that a $(\frac{1}{2}+\epsilon) \to (1-\delta)$ reduction $\text{Red}^{(\cdot)}(x,\alpha)$ that makes q queries, can be transformed into a (distribution) over circuits $T : \{0,1\}^q \to \{0,1\}$ with no oracle (that have roughly the same size and depth as Red) that distinguishes $\mathsf{Noise}^q_{1/2-2\epsilon}$ from $\mathsf{Noise}^q_{1/2}$. We will prove the following lemma (which we call the "zoom lemma").

▶ **Lemma 11** (Zoom lemma). *Under the assumption of Theorem 9, for every $x \in \{0,1\}^k$, there exists a circuit T_x over q bits, with size $\text{poly}(r)$ and depth $O(d)$ over the set of gates G, such that:*

- $\Pr_{X \leftarrow U_k}[T_X(\mathsf{Noise}^q_{1/2-2\epsilon}) = 1] \ge 1 - 2\delta$.
- $\Pr_{X \leftarrow U_k}[T_X(\mathsf{Noise}^q_{1/2}) = 1] \le \frac{1}{2} + \frac{1}{200}$.

Shaltiel and Viola [27] (see also [22]) showed that Theorem 9 follows from Lemma 11. This is formally stated and explained in Section A.[5]

[5] On an intuitive level, the connection between the consequence of the zoom lemma and the consequence of Theorem 9 is that the "best way" to distinguish $\mathsf{Noise}^q_{1/2-2\epsilon}$ from $\mathsf{Noise}^q_{1/2}$ is to check whether the

In the remainder of this section, we prove Lemma 11 modulo some other lemmas and claims, that are stated in this section, and proven in later sections of the paper.

2.2 The oracle used for black-box reductions

Lemmas that are similar to the zoom lemma are at the heart of earlier results [27, 12, 10] on *black-box* reductions, and we would like to imitate the argument working with *class* reductions. Let us start by explaining the oracle used in previous works.

Specifically, let us set $N = 2^n$ and identify the set $[N]$ with the set $\{0,1\}^n$ (so that we can think of Noise_p^N as a function $\mathsf{Noise}_p^N : \{0,1\}^n \to \{0,1\}$). We consider the following two (distributions over) oracles $D : \{0,1\}^n \to \{0,1\}$.

- $D_{1/2-2\epsilon}(y) = f^{\oplus t}(y) \oplus \mathsf{Noise}_{1/2-2\epsilon}^N(y)$
- $D_{1/2}(y) = f^{\oplus t}(y) \oplus \mathsf{Noise}_{1/2}^N(y)$.

▶ **Definition 12.** *We say that a function $D : \{0,1\}^n \to \{0,1\}$ is* useful*, if there exists an $\alpha \in \{0,1\}^a$ such that $\Pr_{X \leftarrow U_k}[\mathrm{Red}^D(X,\alpha) = f^{\oplus t}(X)] \geq 1 - \delta$.*

In the oracle $D_{1/2}$, the noise $\mathsf{Noise}_{1/2}^N(y)$ is uniform and completely masks out the information in $f^{\oplus t}(y)$. Intuitively, this means that the oracle $D_{1/2}$ isn't useful for the reduction. On the other hand, a Chernoff bound shows that w.h.p. over choosing $h \leftarrow \mathsf{Noise}_{1/2-2\epsilon}^N$, we have that $|\{y \in \{0,1\}^n : h(y) = 1\}| \leq (\frac{1}{2} - \epsilon) \cdot N$. This gives that w.h.p. over choosing $D \leftarrow D_{1/2-2\epsilon}$, we have that $\Pr_{Y \leftarrow U_n}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$.

If Red is a $(\frac{1}{2} + \epsilon) \to (1 - \delta)$ *black-box* reduction, then by definition, this implies that every such good D is useful. The proof of [32, 27] then proceeds to transform a *black-box* reduction Red into the circuits T_x required from Lemma 11. We will elaborate on this argument shortly.

Our plan is to imitate this argument when Red is not necessarily a *black-box reduction*, and is only guaranteed to be a *class* reduction. Using this weaker assumption, we are not guaranteed that w.h.p. $D \leftarrow D_{1/2-2\epsilon}$ is useful. This is because we are not guaranteed that w.h.p. $D \leftarrow D_{1/2-2\epsilon}$ belongs to the class $\mathcal{D}$, and the reduction does not need to succeed if $D \notin \mathcal{D}$.

2.3 Using limited independence to obtain efficient oracles

We would like to make the oracle $D_{1/2-2\epsilon}$ efficiently computable by small circuits, so that it belongs to $\mathcal{D}$. This presents two difficulties:

1. $f^{\oplus t}$ is harder to compute than f (and f is assumed to be hard).
2. $\mathsf{Noise}_{1/2-2\epsilon}^N$ is a random function, and w.h.p. requires circuits of exponential size.

In order to circumvent the first problem we use an idea from [32] and will choose the function $f : \{0,1\}^k \to \{0,1\}$ in the following way:

▶ **Lemma 13.** *There exist constants c_1 such that for every constant c_2, there exists a function $f : \{0,1\}^k \to \{0,1\}$ such that:*

- *For every circuit $B : \{0,1\}^k \to \{0,1\}$ of size r^{c_2}, $\Pr_{X \leftarrow U_k}[B(X) = f(X)] \leq \frac{1}{2} + \frac{1}{200}$.*
- *f can be computed by a DNF of size $r^{c_1 \cdot c_2}$.*

fraction of ones is below or above $\frac{1}{2} - \epsilon$. This is similar in spirit to majority over inputs of length $\Omega(1/\epsilon)$, and it can be shown that majority on length $\Omega(1/\epsilon)$ can be reduced to this task. A Chernoff bound shows that $q = O(\frac{\log(1/\delta)}{\epsilon^2})$ is sufficient to distinguish between the two distributions with confidence $1 - \delta$, and it can be shown that such a confidence requires $q = \Omega(\frac{\log(1/\delta)}{\epsilon^2})$.

Proof. By a standard counting argument, there exists a constant c_1 such that for every constant c_2, setting $m = c_1 \cdot c_2 \cdot \log r$, there exists a function $g : \{0,1\}^m \to \{0,1\}$ such that for every circuit B of size $2^{m/c_1} = r^{c_2}$, $\Pr_{X \leftarrow U_m}[B(X) = g(X)] \leq \frac{1}{2} + \frac{1}{200}$. By choosing $\nu > 0$ to be sufficiently small as a function of c_2, we can get that $m \leq k$. The function $f : \{0,1\}^k \to \{0,1\}$ is the function that given $x \in \{0,1\}^k$ applies g on the first m bits of x. ◀

We choose f by the lemma, where c_2 is a constant that we choose later. With this choice we have that:

▶ **Corollary 14.** *The function $f^{\oplus t}$ can be computed by circuits of size* $\mathrm{poly}(r)$ *and constant depth over the gates* $\{$AND,OR,NOT,PARITY$\}$ *of unbounded fan-in.*

▶ Remark 15 (Replacing $f^{\oplus t}$ by a different target function f'). Corollary 14 is the only place in the proof where we use specific properties of $f^{\oplus t}$. The corollary holds for every function $f' : \{0,1\}^n \to \{0,1\}$ for which there exists an oracle circuit $\mathrm{Con}^{(\cdot)}$ of size $\mathrm{poly}(r)$ and constant depth over the gates $\{$AND,OR,NOT,PARITY$\}$ of unbounded fan-in, such that $f' = \mathrm{Con}^f$. This means that our results hold for every such function f'. Furthermore, if Con does not have constant depth, then Corollary 14 gives a size bound on f', and this is sufficient to show the lower bound on number of queries with respect to the class $\mathcal{D}$ of circuits of size $p(r)$.

Corollary 14 takes care of the first difficulty above. It says that $f^{\oplus t}$ can be computed by circuits in the class $\mathcal{D}$. We would like to replace $\mathsf{Noise}^N_{1/2-2\epsilon}$ by a (distribution) over efficient circuits in $\mathcal{D}$. Our approach is to replace $\mathsf{Noise}^N_{1/2-2\epsilon}$ (which consists of N independent bits) by a distribution which is ℓ-wise independent, for $\ell = \mathrm{poly}(r)$.

▶ **Definition 16** (ℓ-wise independence with bias p)**.** *A sequence $R_1, \ldots, R_N$ of bit random variables is ℓ-wise independent with bias p, if $R_1, \ldots, R_N$ are ℓ-wise independent, and for every $i \in [N]$, $\Pr[R_i = 1] = p$.*

We will rely on the following theorem by Gutfreund and Viola [16] (which is usually stated for $p = \frac{1}{2}$ but immediately extends to every rational $p = \frac{a}{b}$ as stated below:)

▶ **Theorem 17** ([16])**.** *Let $N = 2^n$. For every integers $\ell \leq N$ and $a \leq b$, setting $p = \frac{a}{b}$, there exists a distribution H^ℓ_p over circuits $h : \{0,1\}^n \to \{0,1\}$ of size* $\mathrm{poly}(n, \ell, b)$ *and depth $O(1)$ (over the gates* $\{$AND,OR,NOT,PARITY$\}$ *of unbounded fan-in) such that the distribution obtained by choosing $h \leftarrow H^\ell_p$ and considering $(h(1), \ldots, h(N))$, is ℓ-wise independent with bias p.*[6]

We will assume w.l.o.g. that $\frac{1}{\epsilon}$ is an integer, and set $\ell = p_0(r)$ for a polynomial p_0 that we will specify later. We define the following two (distributions over) oracles $D : \{0,1\}^n \to \{0,1\}$, in which we replace the independent bits of Noise^N by ℓ-wise independent bits:[7]

---

[6] We remark that the result of Gutfreund and Viola [16] is significantly stronger. More specifically, for our purposes it suffices that there is a family H of ℓ-wise independent hash functions $h : \{0,1\}^n \to \{0,1\}^b$, such that every h can be computed by the type of circuits claimed above. The result of Gutfreund and Viola gives a stronger bound on the size of H, and also shows that there is a *uniform* circuit that given the "index of h" and an input x, computes $h(x)$.

[7] Replacing fully independent oracles by limited independence oracles, and arguing that black-box procedures with few queries cannot tell the difference, is a common approach in proving black-box impossibility results, originating from the work of Goldreich and Krawczyk [8]. It should be noted that even when ignoring the issue of class reductions, and focusing on black-box reductions, we are considering reductions which are *nonuniform*. Nonuniform reductions get an advice string α that depends on the choice of the oracle. Loosely speaking, this may give them information about the "seed" used to generate the limited independence oracle. This creates technical difficulties that do not occur when reductions are uniform.

- $D^{\ell}_{1/2-2\epsilon}(y) = f^{\oplus t}(y) \oplus H^{\ell}_{1/2-2\epsilon}(y)$
- $D^{\ell}_{1/2}(y) = f^{\oplus t}(y) \oplus H^{\ell}_{1/2}(y)$.

We now have that every D in the support of $D^{\ell}_{1/2-2\epsilon}$ has size $r^{c_1 \cdot c_2} + \mathrm{poly}(n, \ell, 1/\epsilon)$ which can be bounded by $p(r)$ for some polynomial p. Furthermore, each such D has constant depth over the set of gates $\{\text{AND},\text{OR},\text{NOT},\text{PARITY}\}$. This gives that every such D is sufficiently efficient, and belongs to the class $\mathcal{D}$.

This will allow us to imitate the argument for black box reductions. Specifically, by Chebyshev's inequality, with probability at least $1 - \frac{1}{\epsilon^2 2^n} \geq \frac{1}{2}$ over choosing $h \leftarrow H^{\ell}_{1/2-2\epsilon}$, we have that $|\{y \in \{0,1\}^n : h(y) = 1\}| \leq (\frac{1}{2} - \epsilon) \cdot N$. This means that with probability at least half over choosing $D \leftarrow D^{\ell}_{1/2-2\epsilon}$, we have that $\Pr_{Y \leftarrow U_n}[D(Y) = f^{\oplus t}(Y)] \geq \frac{1}{2} + \epsilon$. As Red is a $(\frac{1}{2} + \epsilon) \to (1 - \delta)$ class $\mathcal{D}$ reduction, and every such good D belongs to $\mathcal{D}$, we get that:

▷ Claim 18. $\Pr_{h \leftarrow H^{\ell}_{1/2-2\epsilon}}[(f^{\oplus t} \oplus h) \text{ is useful}] \geq \frac{1}{2}$.

2.4 A more general fixed set lemma

We will now proceed in a similar manner to [27, 10]. Specifically, let Advice be a function that given a useful D, produces an advice string α such that $\Pr_{X \leftarrow U_k}[\mathrm{Red}^D(x, \alpha) = f^{\oplus t}(y)] \geq 1 - \delta$ (such an α exists by definition). For every $\alpha \in \{0,1\}^a$, let A_α be the event

$$A_\alpha = \{h : \{0,1\}^n \to \{0,1\} : (f^{\oplus t} \oplus h) \text{ is useful, and } \mathsf{Advice}(f^{\oplus t} \oplus h) = \alpha\}.$$

By averaging over the 2^a advice strings we obtain that:

▷ Claim 19. There exists $\alpha' \in \{0,1\}^a$ s.t. $\Pr_{h \leftarrow H^{\ell}_{1/2-2\epsilon}}[h \in A_{\alpha'}] \geq \frac{1}{2} \cdot 2^{-a} = 2^{-(a+1)}$.

Let $R = H^{\ell}_{1/2-2\epsilon}$ and $Z = (R | R \in A_{\alpha'})$, following [27, 10] we would like to argue that (in some sense to be explained below) for every $x \in \{0,1\}^k$, $\mathrm{Red}^{(\cdot)}(x, \alpha')$ does not distinguish between the oracle $f^{\oplus t} \oplus Z$ (in which bits can be correlated in complicated ways) and the oracle $f^{\oplus t} \oplus R$ (in which bits are ℓ-wise independent). Note that for every $x \in \{0,1\}^k$, and $h : \{0,1\}^n \to \{0,1\}$, we can think of $\mathrm{Red}^{f^{\oplus t} \oplus h}(x, \alpha')$ as a decision tree (that depends on x) that makes q queries to (the truth table of) h.

With this intuition in mind, we will prove the following lemma (which generalizes a "fixed set lemma" proven in [10] for the special case where the random variables $R_1, \ldots, R_N$ are independent).

▶ **Lemma 20** (A more general fixed set lemma)**.** *Let N, a and q be integers, and let $R = (R_1, \ldots, R_N)$ be some distribution, let $A \subseteq \{0,1\}^N$ be an event such that $\Pr[R \in A] \geq 2^{-a}$, and let $Z = (R | R \in A)$. For every $\eta > 0$, there exists a set $B \subseteq [N]$ of size $b \leq O(a \cdot q/\eta)$, and $v \in \{0,1\}^B$ in the support of Z_B, such that for $R' = (R | R_B = v)$ and $Z' = (Z | Z_B = v) = (R | R_B = v, R \in A)$, and every q-query decision tree P, $P(R')$ and $P(Z')$ are η-close.*[8]

Loosely speaking, the proof works by showing that if there exists a q-query decision tree that distinguishes R from Z, then by fixing the variables on some path of that tree, one obtains a distribution R' such that $\Pr[R' \in A] \geq \Pr[R \in A] \cdot (1 + \eta)$. We apply this argument iteratively (using R' as R) until there does not exist a q query decision tree that distinguishes

[8] Two distributions X, Y over the same domain S are η-close if for every $A \subset S$, $|\Pr[X \in A] - \Pr[Y \in A]| \leq \eta$.

R from Z. In each iteration, $\Pr[R \in A]$ increases by a factor of $1+\eta$, and as this probability cannot be larger than one, this process has to stop after $O(a/\eta)$ steps. By then, we have fixed no more than $O(qa/\eta)$ of the variables. The full proof of Lemma 20 appears in Section 4.[9]

We now continue with the proof of Lemma 11. We apply Lemma 20 on $R = H^{\ell}_{1/2-2\epsilon}$ and the event A'_{α}, using $\eta = \delta$, and let Z, R', Z', B, v and b be as in the lemma. It follows that for every $x \in \{0,1\}^k$, the random variables $\mathrm{Red}^{f^{\oplus t} \oplus R'}(x, \alpha')$ and $\mathrm{Red}^{f^{\oplus t} \oplus Z'}(x, \alpha')$ are δ-close. As this holds for every fixed $x \in \{0,1\}^k$, this also holds for an independently chosen $X \leftarrow U_k$, and we obtain that:

$$\Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(X, \alpha') = f(X)] \geq \Pr[\mathrm{Red}^{f^{\oplus t} \oplus Z'}(X, \alpha') = f(X)] - \delta.$$

The support of Z' is contained in $A_{\alpha'}$, and so, for every h in the support of Z', $(f^{\oplus t} \oplus h)$ is useful (with the advice string α') and we get that:

$$\Pr_{X \leftarrow U_k}[\mathrm{Red}^{f^{\oplus t} \oplus Z'}(X, \alpha') = f(X)] \geq 1 - \delta. \tag{1}$$

Combining this with the previous inequality, gives that:

$$\Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(X, \alpha') = f(X)] \geq 1 - 2\delta. \tag{2}$$

The advantage of (2) over (1) is that we have replaced Z' (in which the bits of $Z'([N] \setminus B)$ can be correlated in complicated ways) with R', where $R'([N] \setminus B)$ is $(\ell - b)$-wise independent. This will allow us to relate this oracle to $\mathsf{Noise}^q_{1/2-2\epsilon}$ and prove the zoom lemma.

2.5 Constructing the circuits for the zoom lemma

In order to construct the circuits required for the zoom lemma, we define the following oracle circuit.

▶ **Definition 21.** *We define an oracle circuit $E^{(\cdot)}(x)$ as follows: On input x and oracle h, $E^h(x)$ simulates* $\mathrm{Red}^{(\cdot)}(x, \alpha')$. *Whenever* Red *makes a query y to its oracle, R acts as follows: if $y \notin B$, then R makes the query y to h, and returns $f^{\oplus t}(y) \oplus h(y)$ to* Red. *If $y \in B$, then R returns $f^{\oplus t}(y) \oplus v(y)$ to* Red. *The output of R is the output of* Red *at the end of this simulation.*

With this definition, it is possible to show that:

▶ **Lemma 22.** *By choosing the constant c_2 and the polynomial p_0 to be sufficiently large, we get that:*

- $\Pr[E^{H^q_{1/2-2\epsilon}}(X) = f(X)] \geq 1 - 2\delta$.
- $\Pr[E^{H^q_{1/2}}(X) = f(X)] \leq \frac{1}{2} + \frac{1}{200}$.
- *For every $x \in \{0,1\}^k$, there exists a circuit $T_x : \{0,1\}^q \to \{0,1\}$ of size* $\mathrm{poly}(r)$ *and depth $O(d)$ over the gates G, such that for every $0 \leq p \leq 1$, $T_x(\mathsf{Noise}^q_p) = E^{H^q_p}(x)$.*

We note that this lemma immediately implies Lemma 11. The proof of Lemma 22 appears in Section 3, and is similar in spirit to earlier work [27, 12, 10]. It is in fact significantly simpler, as in this paper, we have the additional advantage that $f^{\oplus t}$ has circuits of size $\mathrm{poly}(r)$ and constant depth.

---

[9] The proof of the fixed set lemma given in [10] also uses an iterative argument: It shows that the existence of a q-query decision tree gives rise to a new distribution Z where the entropy of Z is increased. It is then argued that the iterative process has to stop before (as the entropy of Z is upper bounded by N). This limits the earlier proofs to distributions R where $Z = (R|R \in A)$ has very high entropy, which isn't the case for our choice of oracle.

Organization of the paper

We prove Lemma 22 in Section 3. In Section 4 we prove the more general fixed set lemma (that is Lemma 20). In Appendix A we state and survey the results of [27] showing that the zoom lemma implies the main theorem. In Appendix B we explain how to extend the argument to sufficiently explicit linear codes.

3 Proof of Lemma 22

In this section we prove Lemma 22. We start by proving the first item. Note that $b = O(qa/\delta)$ is bounded by some polynomial in r. We are allowed to choose the polynomial p_0 to be sufficiently large so that $\ell = p_0(r)$ satisfies $(\ell - b) \geq q$. This gives that the $N - b$ coordinates of $R'([N] \setminus B)$ are $(\ell - b)$-wise independent (because R' was obtained by fixing b indices of R which is ℓ-wise independent). The fact that $R'([N] \setminus B)$ are q-wise independent, and that E answers queries in B using v, gives that for every $x \in \{0,1\}^k$, the q queries made by $E^{H^q_{1/2-2\epsilon}}(x)$ are distributed exactly like the queries of $\mathrm{Red}^{f^{\oplus t} \oplus R'}(x, \alpha')$, meaning that:

$$\Pr[E^{H^q_{1/2-2\epsilon}}(x) = f(x)] = \Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(x, \alpha') = f(x)].$$

This immediately means that for an independent $X \leftarrow U_k$:

$$\Pr[E^{H^q_{1/2-2\epsilon}}(X) = f(X)] = \Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(X, \alpha') = f(X)].$$

We have already seen in (2) that:

$$\Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(X, \alpha') = f(X)] \geq 1 - 2\delta,$$

and this gives the first item.

For the second item, we note that if $E^{H^q_{1/2}}$ makes a query $y \notin B$, then it obtains a uniform coin, and the coins obtained on different queries are independent. Recall that on queries $y \in B$, E answers the queries without consulting the oracle. This means that we can simulate $E^{H^q_{1/2}}(x)$ by a randomized circuit $\bar{C}$ that on input x, simulates E and answers queries $y \notin B$ by random coins. It follows that for every $x \in \{0,1\}^k$:

$$\Pr[E^{H^q_{1/2}}(x) = f(x)] = \Pr[\bar{C}(x) = f(x)].$$

This immediately means that for an independent $X \leftarrow U_k$:

$$\Pr[E^{H^q_{1/2}}(X) = f(X)] = \Pr[\bar{C}(X) = f(X)].$$

There exists some fixing for the random coins of $\bar{C}$ such that the obtained (deterministic) circuit C satisfies $Pr[C(X) = f(X)] \geq \Pr[\bar{C}(X) = f(X)]$. The circuit C is hardwired with this choice of random coins, and with α', B, v, and $f^{\oplus t}(B)$. (A crucial observations is that C does not need to compute $f^{\oplus t}$ for $y \notin B$). Overall, this is a circuit of size r^c for some constant c, and by choosing the constant c_2 from to be a larger constant, and using Lemma 13, we have that:

$$\Pr[E^{H^q_{1/2}}(X) = f(X)] \leq \Pr[C(X) = f(X)] \leq \frac{1}{2} + \frac{1}{200}.$$

This proves the second item.

For the third item, we note (once again) that for every p, and for every $x \in \{0,1\}^k$, the distribution of the q answers that $E^{H^q_p}$ obtains from its oracle is distributed like Noise^q_p. This means that for every x, we can construct a circuit T_x that on input Noise^q_p simulates $E^{H^q_p}(x)$,

using its i'th input to answer the i'th query of E. The circuit T_x is hardwired with α', B and v. Unlike the circuit C from the second item, T_x needs to compute $f^{\oplus t}$ on each of the q queries. This can be done using Corollary 14. Overall, T_x is a circuit of size $\text{poly}(r)$ (this time the polynomial is larger than r^{c_2}) and depth $O(d)$ (because on every oracle call of Red, T_x may have to compute $f^{\oplus t}$ (which takes constant depth according to Corollary 14).[10]

4 The fixed set lemma for ℓ-wise independence

In this section we prove Lemma 20. The proof will iteratively applying the following lemma.

▶ **Lemma 23.** *Let $R = (R_1, \ldots, R_N)$ be a distribution, let $A \subseteq \{0,1\}^N$ be an event, and let $Z = (R|R \in A)$. If there exists a q-query decision tree P such that $|\Pr[P(R) = 1] - \Pr[P(Z) = 1]| > \eta$ then there exists $Q \subset [N]$ of size q and $v \in \{0,1\}^Q$ in the support of Z_Q, such that*

$$\Pr[R \in A | R(Q) = v] > (1 + \eta) \cdot \Pr[R \in A].$$

Proof. Let P be a q-query decision tree, and assume w.l.o.g. (by complementing P if necessary) that $\Pr[P(Z) = 1] - \Pr[P(R) = 1] > \eta$. A path in the decision tree corresponds to a subset $Q \subset [N]$ of the q variables queried on the path, and a string $v \in \{0,1\}^q$ of the answers. For every such path, let $path_{Q,v} : \{0,1\}^N \to \{0,1\}$ be the function that evaluates to 1 on input $r = (r_1, \ldots, r_N)$ if $r(Q) = v$ (meaning that the tree P takes the path (Q, v) on input r). Let S be the set of all pairs (Q, v) corresponding to paths of P that answer 1. The path taken by a decision tree is unique, and therefore, for any distribution R on $\{0,1\}^N$, we have that:

$$\Pr[P(R) = 1] = \sum_{(Q,v) \in S} \Pr[path_{Q,v}(R) = 1].$$

▷ Claim 24. There exists a path $(Q, v) \in S$ such that:

$$\Pr[path_{Q,v}(Z) = 1] > (1 + \eta) \cdot \Pr[path_{Q,v}(R) = 1].$$

Proof of claim. This is because otherwise:

$$\begin{aligned}
\Pr[P(Z) = 1] &= \sum_{(Q,v) \in S} \Pr[path_{Q,v}(Z) = 1] \\
&\leq (1 + \eta) \cdot \sum_{(Q,v) \in S} \Pr[path_{Q,v}(R) = 1] \\
&= (1 + \eta) \cdot \Pr[P(R) = 1] \\
&\leq \Pr[P(R) = 1] + \eta.
\end{aligned}$$

◁

[10] Our proof of Lemma 22 relies on the fact that $f^{\oplus t}$ has small constant depth circuits. This allows us to simplify the argument used by some of the previous work [27, 12, 10] which wasn't allowed to assume that the target function $f' = f^{\oplus t}$ can be computed by a small constant depth circuit. The proofs in [27, 12, 10] need to resort to different arguments (and this creates additional difficulties if Red makes adaptive calls to its oracle, meaning that the queries that $\text{Red}^{(\cdot)}(x, \alpha')$ makes are not a function of x and α', and may also depend on previous answers). However, using a clever hybrid argument of [12] and additional ideas explained in [10], it is possible to conclude that T_x has depth $O(d)$ *without* relying on the fact that f' is computable by constant depth circuits. This argument allows choosing f' where $f' = \text{Con}^f$ for an oracle circuit Con that has size $\text{poly}(r)$, but does not necessarily has constant depth, and this gives the aforementioned extension of Theorem 9 to this setup, which now holds for every class $\mathcal{D}$ that contains circuits of size $\text{poly}(r)$.

In particular, the event $\{R(Q) = v\}$ occurs with positive probability, and it follows that:

$$\begin{aligned}\Pr[R \in A | R(Q) = v] &= \frac{\Pr[R \in A \cap R(Q) = v]}{\Pr[R(Q) = v]} \\ &= \frac{\Pr[R \in A] \cdot \Pr[R(Q) = v | R \in A]}{\Pr[path_{Q,v}(R) = 1]} \\ &= \frac{\Pr[R \in A] \cdot \Pr[Z(Q) = v]}{\Pr[path_{Q,v}(R) = 1]} \\ &= \frac{\Pr[R \in A] \cdot \Pr[path_{Q,v}(Z) = 1]}{\Pr[path_{Q,v}(R) = 1]} \\ &> (1+\eta) \cdot \Pr[R \in A].\end{aligned}$$

◀

We are now ready to prove Lemma 20.

Proof of Lemma 20. We consider the following iterative process: At step i, we have:

- A distribution $R^{(i)}$ over $\{0,1\}^N$.
- A set $B^{(i)} \subseteq [N]$.
- $v^{(i)} \in \{0,1\}^{B^{(i)}}$.

We will assume that the following invariant is satisfied:

- $B^{(i)}$ is of size $i \cdot q$.
- $\Pr[R^{(i)} \in A] \geq 2^{-a} \cdot (1+\eta)^i$.
- $R^{(i)}(B^{(i)}) = v^{(i)}$ (with probability one).

Note that the assumption in the lemma fulfills this invariant for $i = 0$ with $R^{(0)} = R.$ and $B^{(0)} = \emptyset$.

At step i, we define $\bar{R} = R^{(i)}([N] \setminus B^{(i)})$. As $R^{(i)}$ is fixed on $B^{(i)}$, we can think of A as an event that only observes the indices in $[N] \setminus B^{(i)}$. More formally, there is an event $\bar{A} \subseteq \{0,1\}^{[N] \setminus B^{(i)}}$ such that $R^{(i)} \in A$ iff $\bar{R} \in \bar{A}$, and

$$\Pr[\bar{R} \in \bar{A}] = \Pr[R^{(i)} \in A] \geq 2^{-a} \cdot (1+\eta)^i.$$

Let $\bar{Z} = (\bar{R} | \bar{R} \in \bar{A})$. If the conclusion of Lemma 20 does not hold with respect to $B^{(i)}, v^{(i)}$, then there exists a q-query decision tree P that distinguishes $R^{(i)}$ from $(R^{(i)} | R^{(i)} \in A)$, with advantage η, and as the two distributions agree on the queries in $B^{(i)}$, we conclude that P distinguishes $\bar{R}$ from $\bar{Z}$ with the same advanatge. We apply Lemma 23 on $\bar{R}$ and $\bar{A}$, and conclude that there exists $Q \subseteq [N] \setminus B^{(i)}$ and $v \in \{0,1\}^Q$ such that

$$\Pr[\bar{R} \in \bar{A} | \bar{R}(Q) = v] > (1+\eta) \cdot \Pr[\bar{R} \in \bar{A}] \geq 2^{-a} \cdot (1+\eta)^{i+1}.$$

We set:

- $B^{(i+1)} = B^{(i)} \cup Q$.
- $v^{(i+1)}$ to be the "concatenation of $v^{(i)}$ and v". More precisely, for $y \in B^{(i)}$, $v_y^{(i+1)} = v_y^{(i)}$ and for $y \in Q$, $v_y^{(i+1)} = v_y$.
- $R^{(i+1)} = (R^{(i)} | R^{(i)}(Q) = v)$. (Note that by definition $B^{(i)} \cap Q = \emptyset$).

We now observe that the invariant is maintained in step $i+1$. Specifically:

- $|B^{(i+1)}| = |B^{(i)}| + q = i \cdot q + q = (i+1) \cdot q$.
- $\Pr[R^{(i+1)} \in A] = \Pr[R^{(i)} \in A | R^{(i)}(Q) = v] = Pr[\bar{R} \in \bar{A} | \bar{R}(Q) = v] \geq 2^{-a} \cdot (1+\eta)^{i+1}$.
- By definition, $R^{(i+1)}(B^{(i+1)}) = v^{(i+1)}$ with probability one.

Therefore, if this process fails to deliver the lemma after i steps, then the invariant is maintained at the end of step i, and in particular, $\Pr[R^{(i)} \in A] \geq 2^{-a} \cdot (1+\eta)^i$. However, this is impossible for $i > \frac{a}{\log(1+\eta)} = \Theta(a/\eta)$, and so, this process has to deliver the lemma within this number of steps. We obtain that the lemma follows with $b = |B| \leq O(\frac{qa}{\eta})$. ◀

5 Conclusion and open problems

Class reductions are known to bypass some limitations on black-box reductions (as explained in Section 1.2). This work demonstrates that it is sometimes possible to extend limitations on black-box reductions to class reductions. Studying the power of class reductions may promote our understanding of how to bypass limitations on black-box reductions. We now mention some more specific open problems:

- Unlike the results of [10], our results do not hold for any construction of target functions f' from f. Is it possible to extend our results to this general setting?
- In Theorem 9, the class $\mathcal{D}$ contains circuits that are polynomially larger than the size of the the reduction. Is it possible to extend our limitations on class reductions with respect to a classes $\mathcal{D}$ of circuits smaller than the circuit size of the reduction?
- Yao's XOR lemma states that *for every* function f, if f is somewhat hard, then $f^{\oplus t}$ is very hard. It makes sense to focus on some specific choice for a somewhat hard function f and prove and improved result for this specific function. If we prove such an assertion by reduction, we can allow the reduction to be tailored to the specific function f, and do no need to require that the reduction performs on any function f, but only on the chosen one. This type of reductions was termed "function specific" by Artemenko and Shaltiel [3], who proved limitations on nonuniform black-box functions specific reductions. It is interesting to understand whether function specific class reductions can circumvent the limitations proven in this paper. We remark that our proof technique indeed relies on the fact that the reduction is not function specific, and must work for *any* function f. This allows us to choose f with specific properties that are useful for our argument.

References

1 Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on np-hardness. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 701–710, 2006. doi:10.1145/1132516.1132614.

2 Benny Applebaum, Sergei Artemenko, Ronen Shaltiel, and Guang Yang. Incompressible functions, relative-error extractors, and the power of nondeterministic reductions. *Computational Complexity*, 25(2):349–418, 2016. doi:10.1007/s00037-016-0128-9.

3 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. doi:10.1007/s00037-012-0056-2.

4 Albert Atserias. Distinguishing SAT from polynomial-size circuits, through black-box queries. In *21st Annual IEEE Conference on Computational Complexity*, pages 88–95, 2006. doi:10.1109/CCC.2006.17.

5 Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. *SIAM J. Comput.*, 36(4):1119–1159, 2006. doi:10.1137/S0097539705446974.

6 Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. On beating the hybrid argument. *Theory of Computing*, 9:809–843, 2013.

7 Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM J. Comput.*, 22(5):994–1005, 1993. doi:10.1137/0222061.

8 Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM J. Comput.*, 25(1):169–192, 1996. doi:10.1137/S0097539791220688.

9 Oded Goldreich, Noam Nisan, and Avi Wigderson. On yao's xor-lemma. 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 273–301. Springer, 2011. doi:10.1007/978-3-642-22670-0_23.

10 Aryeh Grinberg, Ronen Shaltiel, and Emanuele Viola. Indistinguishability by adaptive procedures with advice, and lower bounds on hardness amplification proofs. In *59th IEEE Annual Symposium on Foundations of Computer Science*, pages 956–966, 2018. doi:10.1109/FOCS.2018.00094.

11 Dan Gutfreund. Worst-case vs. algorithmic average-case complexity in the polynomial-time hierarchy. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX and 10th International Workshop on Randomization and Computation, RANDOM*, volume 4110 of *Lecture Notes in Computer Science*, pages 386–397, 2006. doi:10.1007/11830924_36.

12 Dan Gutfreund and Guy N. Rothblum. The complexity of local list decoding. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, 11th International Workshop, APPROX, and 12th International Workshop, RANDOM*, volume 5171 of *Lecture Notes in Computer Science*, pages 455–468, 2008. doi:10.1007/978-3-540-85363-3_36.

13 Dan Gutfreund, Ronen Shaltiel, and Amnon Ta-Shma. If NP languages are hard on the worst-case, then it is easy to find their hard instances. *Computational Complexity*, 16(4):412–441, 2007. doi:10.1007/s00037-007-0235-8.

14 Dan Gutfreund and Amnon Ta-Shma. Worst-case to average-case reductions revisited. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 10th International Workshop, APPROX, and 11th International Workshop, RANDOM*, volume 4627 of *Lecture Notes in Computer Science*, pages 569–583, 2007. doi:10.1007/978-3-540-74208-1_41.

15 Dan Gutfreund and Salil P. Vadhan. Limitations of hardness vs. randomness under uniform reductions. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, 11th International Workshop, APPROX, and 12th International Workshop, RANDOM*, volume 5171 of *Lecture Notes in Computer Science*, pages 469–482, 2008. doi:10.1007/978-3-540-85363-3_37.

16 Dan Gutfreund and Emanuele Viola. Fooling parity tests with parity gates. In *Approximation, Randomization, and Combinatorial Optimization, Algorithms and Techniques, 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX, and 8th International Workshop on Randomization and Computation, RANDOM*, volume 3122 of *Lecture Notes in Computer Science*, pages 381–392, 2004. doi:10.1007/978-3-540-27821-4_34.

17 Shuichi Hirahara. Non-black-box worst-case to average-case reductions within NP. In *59th IEEE Annual Symposium on Foundations of Computer Science*, pages 247–258, 2018. doi:10.1109/FOCS.2018.00032.

18 Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th Annual Symposium on Foundations of Computer Science*, pages 538–545, 1995. doi:10.1109/SFCS.1995.492584.

19 Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 44–61, 1989. doi:10.1145/73007.73012.

20 Russell Impagliazzo and Avi Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing*, pages 220–229, 1997. doi:10.1145/258533.258590.

21 Adam R. Klivans and Rocco A. Servedio. Boosting and hard-core set construction. *Machine Learning*, 51(3):217–238, 2003. doi:10.1023/A:1022949332276.

22 Nutan Limaye, Karteek Sreenivasaiah, Srikanth Srinivasan, Utkarsh Tripathi, and S. Venkitesh. A fixed-depth size-hierarchy theorem for $\mathrm{ac}^0[\oplus]$ via the coin problem. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, pages 442–453, 2019. doi:10.1145/3313276.3316339.

23 Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. On the complexity of hardness amplification. *IEEE Trans. Information Theory*, 54(10):4575–4586, 2008. doi:10.1109/TIT.2008.928988.

24 Igor Carboni Oliveira, Rahul Santhanam, and Srikanth Srinivasan. Parity helps to compute majority. In *34th Computational Complexity Conference*, volume 137, pages 23:1–23:17, 2019. doi:10.4230/LIPIcs.CCC.2019.23.

25 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.

26 Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20, 2004. doi:10.1007/978-3-540-24638-1_1.

27 Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM J. Comput.*, 39(7):3122–3154, 2010. doi:10.1137/080735096.

28 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 77–82, 1987. doi:10.1145/28395.28404.

29 Madhu Sudan, Luca Trevisan, and Salil P. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001. doi:10.1006/jcss.2000.1730.

30 Luca Trevisan and Salil P. Vadhan. Pseudorandomness and average-case complexity via uniform reductions. *Computational Complexity*, 16(4):331–364, 2007. doi:10.1007/s00037-007-0233-x.

31 Emanuele Viola. Hardness vs. randomness within alternating time. In *18th Annual IEEE Conference on Computational Complexity*, page 53, 2003. doi:10.1109/CCC.2003.1214410.

32 Emanuele Viola. *The Complexity of Hardness Amplification and Derandomization*. PhD thesis, Harvard University, 2006.

A Showing that the main theorem follows from the zoom lemma

In this section we show that Theorem 9 follows from Lemma 11. This follows by the earlier work of Shaltiel and Viola [27] which we now explain.

A.1 Consequences distinguishing noise $\frac{1}{2}$ from $(\frac{1}{2} - 2\epsilon)$

The next lemma shows that distinguishing between $\mathsf{Noise}^q_{1/2-2\epsilon}$ and $\mathsf{Noise}^q_{1/2}$ requires many queries.

▶ **Lemma 25** ([32, 27]). *For every $\epsilon, \delta > 0$, such that $\delta < 0.4$, if $T : \{0,1\}^q \to \{0,1\}$ satisfies:*

- $\Pr[T(\mathsf{Noise}^q_{1/2-2\epsilon}) = 1] \geq 1 - \delta$.
- $\Pr[T(\mathsf{Noise}^q_{1/2}) = 1] \leq 0.51$.

Then, $q = \Omega(\frac{\log \frac{1}{\delta}}{\epsilon})$.

The next lemma essentilly shows that distinguishing between $\mathsf{Noise}^q_{1/2-2\epsilon}$ and $\mathsf{Noise}^q_{1/2}$ requires majority on $\Omega(1/\epsilon)$ bits. A technicality is that for this conclusion it is not sufficient to distinguish $\mathsf{Noise}^q_{1/2-2\epsilon}$ from $\mathsf{Noise}^q_{1/2}$, and one needs circuits that distinguish $\mathsf{Noise}^q_{1/2-\frac{2}{j}}$ from $\mathsf{Noise}^q_{1/2}$ for every integer j between 1 and $\log(1/\epsilon)$. This complication is in some sense necessary (see discussion in [27]).

▶ **Lemma 26** ([32, 27]). *For every $\epsilon, \delta > 0$, such that $\delta < 0.4$, and $\frac{1}{\epsilon}$ is an integer. If $T_1, \ldots, T_{\frac{1}{\epsilon}}$ are circuits over q bits, with size $s \geq q$ and depth d (over some set of gates G that includes the standard set $\{\text{AND}, \text{OR}, \text{NOT}\}$ with unbounded fan-in) and for every $j \in [\frac{1}{\epsilon}]$, we have that:*

- $\Pr[T_j(\mathit{Noise}^q_{1/2-\frac{2}{j}}) = 1] \geq 1 - \delta$.
- $\Pr[T_j(\mathit{Noise}^q_{1/2}) = 1] \leq 0.51$.

Then, there exists a circuit A that computes the majority function over $\Omega(\frac{1}{\epsilon})$ bits, and A has size $s \cdot \mathrm{poly}(\frac{1}{\epsilon})$ and depth $d + O(1)$ over the same set of gates G.

A.2 Finishing up

We now have have all the tools to prove that Theorem 9 follows from Lemma 11. We will assume w.l.o.g. that $\frac{1}{\epsilon}$ is an integer. We first observe that a $(\frac{1}{2} + \epsilon) \to (1 - \delta)$ class $\mathcal{D}$ reduction for Yao's XOR lemma, is in particular a $(\frac{1}{2} + \frac{1}{j}) \to (1 - \delta)$ class $\mathcal{D}$ reduction for Yao's XOR lemma, for every $j \in [\frac{1}{\epsilon}]$.

This means that for every $j \in [\frac{1}{\epsilon}]$, we can apply Lemma 11 choosing the parameter ϵ to be $\epsilon = \frac{1}{j}$, and for each such j and $x \in \{0,1\}^k$ we obtain a circuit T^j_x over q bits with size $\mathrm{poly}(r)$ and depth $O(d)$ over the set of gates G such that:

- $\Pr_{X \leftarrow U_k}[T^j_X(\mathsf{Noise}^q_{1/2-\frac{2}{j}}) = 1] \geq 1 - 2\delta$.
- $\Pr_{X \leftarrow U_k}[T^j_X(\mathsf{Noise}^q_{1/2}) = 1] \leq \frac{1}{2} + \frac{1}{200}$.

Applying Markov's inequality to the second item of the lemma, we obtain that there exists a constant $\beta > 0$ such that for every j, for a β fraction of $x \in \{0,1\}^k$,

$$\Pr[T^j_x(\mathsf{Noise}^q_{1/2}) = 1] \leq 0.51.$$

Applying Markov's inequality to the first item of the lemma, we obtain that for every j, for a $1 - \beta/2$ fraction of $x \in \{0,1\}^k$,

$$\Pr[T^j_x(\mathsf{Noise}^q_{1/2-\frac{2}{j}}) = 1] \leq 1 - \frac{4 \cdot \delta}{\beta} \leq 1 - \frac{4 \cdot \delta_0}{\beta}.$$

Together, this gives that for every j, there exists $x \in \{0,1\}^k$ that satisfies both inequalities. Theorem 9 now follows directly from Lemma 25 and Lemma 26, by choosing the constant $\delta_0 > 0$ to be sufficiently small.[11]

B Extending the argument to sufficiently explicit linear codes

We now prove the extension of our results to sufficiently explicit linear codes which is stated in Section 1.3. More specifically, we will prove a version of Theorem 9 that assumes that $\delta = 2^{-2k}$ and replace replace $f^{\oplus t}$ by a function $f' : \{0,1\}^n \to \{0,1\}$ defined by:

$$f'(y) = \sum_{x \in \{0,1\}^k} f(x) \cdot g(x,y),$$

where the sum is taken in the field $\mathbb{F}_2$, and $g : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}$ can be computed by circuits of size $\mathrm{poly}(r)$ and depth d over the set G of gates.

This argument is based on a trick by Viola [32] that we can incorporate into our framework. We will modify the proof of Lemma 11 so that it holds in this setting, the modified version of Theorem 9 will follow from Lemma 11 just as before.

We start by replacing the function f of Lemma 13 with a sightly different function:

[11] The argument above is wasteful, and leads to a rather small constant $\delta_0 > 0$. We remark that with a more careful argument, we could have chosen δ_0 to be any constant smaller than $\frac{1}{2}$, and even allow it to approach $\frac{1}{2}$. More specifically, a more careful analysis can allow $\delta_0 = \frac{1}{2} - O(\log(1/\epsilon))$.

▶ **Lemma 27.** *There exist constants c_1 such that for every constant c_2, there exists a function $f : \{0,1\}^k \to \{0,1\}$ such that there exists $m \leq c_1 \cdot c_2 \cdot \log r$ such that:*

- *For every circuit $B : \{0,1\}^k \to \{0,1\}$ of size r^{c_2}, $\Pr_{X \leftarrow U_m}[B(X \circ 0^{k-m}) = f(X \circ 0^{k-m})] \leq \frac{1}{2} + \frac{1}{200}$.*
- *f can be computed by a DNF of size $r^{c_1 \cdot c_2}$.*

Proof. We repeat the proof of Lemma 13, but this time we take $f(x)$ to be the following function: Let x' denote the first m bits of x and x'' denote the remaining $k-m$ bits. We define $f(x)$ to be $g(x')$ if $x'' = 0^{k-m}$ and zero otherwise. ◀

On a random $X \leftarrow U_k$, $\Pr[f(X) = 0] \geq 1 - 2^{k-m} = 1 - 2^{-\Omega(k)}$. Therefore, it is very easy to compute f with success probability $1 - 2^{-\Omega(k)}$ by simply answering zero. However, by Lemma 27 it is hard for circuits of size r^{c_2} to compute f with success probability 1, or equivalently success probability $1 - \delta$ for $\delta = 2^{-2k}$. This is why this approach can only succeed for very small δ.

With this choice, we can get a corollary that is analogous to Corollary 14.

▶ **Corollary 28.** *The function f' can be computed by circuits of size* $\mathrm{poly}(r)$ *and constant depth over the gates* $\{$AND,OR,NOT,PARITY$\}$ *of unbounded fan-in.*

Proof. The function f' is defined by:

$$f'(y) = \sum_{x \in \{0,1\}^k} f(x) \cdot g(x,y).$$

The sum ranges over 2^k choices of x. However, for our function f, except for $\mathrm{poly}(r)$ of these x (the ones for which the second part of x is $k-m$ zeros) all the remaining x have $f(x) = 0$. This, together with the requirement on g, gives the required result. ◀

The proof proceeds as in Section 2, with the following modifications:

- For $\delta = 2^{-2k} < 2^{-k}$, Red is a $(\frac{1}{2} + \epsilon) \to 1$ reduction. This means in particular that if D is useful, then there exists $\alpha \in \{0,1\}^a$ such that $\Pr_{X \leftarrow U_m}[\mathrm{Red}^D(X \circ 0^{k-m}, \alpha) = f(X \circ 0^{k-m})] = 1$.
- We set $\delta' = \frac{1}{r}$ and will replace some occurrences of δ in the earlier argument by δ'. This is done because the choice of $\delta = 2^{-2k}$ does not satisfy the requirement that $\frac{1}{\delta} \leq r$ made in Theorem 9. Specifically, the requirement that $\frac{1}{\delta} \leq r$ was used to argue that when we apply Lemma 20 with $\eta = \delta$, the size of the set B (which is polynomial in $\frac{1}{\eta}$) is polynomial in r. In order to obtain a set B of size $\mathrm{poly}(r)$ we will now choose $\eta = \delta'$.
- In Section 2.4 we argued that for an independent $X \leftarrow U_k$:

  $$\Pr[\mathrm{Red}^{f^{\oplus t} \oplus R'}(X, \alpha') = f(X)] \geq \Pr[\mathrm{Red}^{f^{\oplus t} \oplus Z'}(X, \alpha') = f(X)] - \delta.$$

  With our modifications we get that for an independent $X \leftarrow U_m$:

  $$\Pr[\mathrm{Red}^{f' \oplus R'}(X \circ 0^{k-m}, \alpha') = f(X \circ 0^{k-m})] \geq \Pr[\mathrm{Red}^{f^{\oplus t} \oplus Z'}(X, \alpha') = f(X)] - \delta'.$$

- This allows us to continue the argument, replacing occurrences of $X \leftarrow U_k$ by $X \circ 0^{k-m}$ for $X \leftarrow U_m$, and occurrences of δ by δ'.
- When we finish the proof and obtain a lower bound of $q = \Omega(\frac{\log(1/\delta')}{\epsilon^2}) = \Omega(\frac{\log r}{\epsilon^2})$ as required. The result on majority is not affected by replacing δ by δ'.