

Relativized Worlds Without Worst-Case to Average-Case Reductions for NP

Thomas Watson*

June 10, 2012

Abstract

We prove that relative to an oracle, there is no worst-case to average-case reduction for NP. We also handle classes that are somewhat larger than NP, as well as worst-case to *errorless*-average-case reductions. In fact, we prove that relative to an oracle, there is no worst-case to errorless-average-case reduction from NP to BPP_{\parallel}^{NP} . We also handle reductions from NP to the polynomial-time hierarchy and beyond, under restrictions on the number of queries the reductions can make.

1 Introduction

The study of average-case complexity concerns the power of algorithms that are allowed to make mistakes on a small fraction of inputs. Of particular importance is the relationship between worst-case complexity and average-case complexity. For example, cryptographic applications require average-case hard problems, and it would be desirable to base the existence of such problems on minimal, worst-case complexity assumptions.

For the class PSPACE, it is known that worst-case hardness and average-case hardness are equivalent [BFNW93]. That is, if PSPACE is worst-case hard then it is also average-case hard. For the class NP, the situation is not well-understood. A central open problem in average-case complexity is to prove that if NP is worst-case hard then it is also average-case hard. Considering the lack of progress toward proving this proposition, a natural goal is to exhibit barriers to proving it, by ruling out certain general proof techniques. Bogdanov and Trevisan [BT06b] considered the possibility of a *proof by reduction*. Building on [FF93], they showed that the proposition cannot be proven by a nonadaptive randomized reduction unless the polynomial-time hierarchy collapses; it remains open to provide evidence against the existence of adaptive reductions.¹ Another possibility that has been considered is a *relativizing proof*. In 1995, Impagliazzo and Rudich claimed [Imp95] that they had constructed a relativized heuristica, which is a world in which NP is worst-case hard but average-case easy, thus ruling out this possibility. However, they have since retracted their claim. We make progress toward obtaining relativized heuristica, by ruling out the possibility of a

*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.

¹It can be shown without difficulty that there is no *deterministic* adaptive worst-case to average-case reduction for NP unless $P = NP$.

relativizing proof by reduction. Our barrier holds even for adaptive reductions. More formally, we prove that there exists an oracle relative to which there is no reduction of type

$$(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP} \Rightarrow \text{NP} \subseteq \text{BPP}$$

where $(\text{NP}, \text{PSAMP})$ is the class of distributional NP problems under polynomial-time samplable distributions, and HeurBPP is the class of distributional problems with polynomial-time average-case randomized algorithms.

We also generalize this result in various ways. The proposition that if NP is worst-case hard then it is also average-case hard concerns average-case algorithms that may output the wrong answer on a small fraction of inputs. In light of the aforementioned barriers, it is natural to consider the following proposition, which is potentially easier to prove: If NP is worst-case hard then it is also hard for *errorless* average-case algorithms, which may output “don’t know” on a small fraction of inputs but must never output the wrong answer.² Our result generalizes to rule out relativizing proofs by reduction of this proposition. Further, we show how to rule out relativizing proofs by reduction that if NP is worst-case hard then certain classes larger than NP are errorless-average-case hard.

Independently of our work, Impagliazzo [Imp11] has succeeded in constructing a relativized heuristica; we discuss his result in Section 1.3 below.

1.1 Notions of Reductions and Relationship to Previous Work

Various models of worst-case to average-case reductions for NP have been considered in the literature, and they can be informally classified as follows.

For the moment let us gloss over the issue of which distribution on inputs an average-case algorithm is judged with respect to. A worst-case to average-case reduction for NP must show that for every $L_1 \in \text{NP}$ there exists an $L_2 \in \text{NP}$ such that if L_2 has a polynomial-time average-case algorithm then L_1 has a polynomial-time worst-case algorithm. The worst-case algorithm for L_1 depends on the hypothesized average-case algorithm for L_2 in some way, which we call the *decoding*. There are the following four natural types of dependence, in decreasing order of strength.

- (1) Black-box dependence means that the worst-case algorithm for L_1 has oracle access to the average-case algorithm for L_2 , and it must solve L_1 on all inputs for *every* oracle that solves L_2 on most inputs, regardless of whether the oracle represents an efficient algorithm.
- (2) The worst-case algorithm for L_1 might have oracle access to the average-case algorithm for L_2 but only be guaranteed to solve L_1 when the oracle is, in fact, an efficient average-case algorithm for L_2 .
- (3) The worst-case algorithm for L_1 might require the code of an efficient average-case algorithm for L_2 .
- (4) The dependence can be arbitrary, meaning that if L_2 has an efficient average-case algorithm then L_1 has an efficient worst-case algorithm. This type of dependence allows for arbitrary proofs that if NP is worst-case hard then it is also average-case hard.

²An equivalent notion of an errorless average-case algorithm is one that always outputs the correct answer but whose running time is only “polynomial-on-average” [Lev86].

For the first three types, the algorithm that solves L_1 with the aid of a hypothesized average-case algorithm for L_2 is called the reduction itself. In this paper we consider type (1) decoding. Note that since our results are about relativization, the reductions we consider have access to two oracles: the reduction oracle (representing the hypothesized average-case algorithm) and the relativization oracle.

Bogdanov and Trevisan [BT06b] also considered type (1) decoding. They showed that such a reduction cannot exist unless the polynomial-time hierarchy collapses, provided the reduction is nonadaptive in its oracle access to the hypothesized average-case algorithm. Compared to the Bogdanov-Trevisan barrier, our barrier has the advantages that it is unconditional and it applies to adaptive reductions, but has the disadvantage that it only applies to reductions that relativize.

Gutfreund et al. [GSTS07] showed a *positive* result, namely that there is a worst-case to average-case reduction for NP with type (2) decoding, under a distribution on inputs that is samplable in slightly-superpolynomial time. Building on this result, Gutfreund and Ta-Shma [GTS07] showed that under a certain weak derandomization hypothesis, there is a worst-case to average-case reduction from NP to nondeterministic slightly-superpolynomial time with type (2) decoding, under the uniform distribution on inputs. Moreover, the results of [GSTS07, GTS07] relativize.

A natural goal is to extend our results to handle type (2) decoding. However, this turns out to be as hard as extending our results to handle type (4) decoding (which was independently accomplished by Impagliazzo [Imp11], at least for NP). For example, we claim that relative to every oracle, the following are equivalent.

(A) There is no reduction of type

$$(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP} \Rightarrow \text{NP} \subseteq \text{BPP}$$

with type (2) decoding.

(B) $(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP}$ and $\text{NP} \not\subseteq \text{BPP}$.

Clearly (B) implies (A). To see that (A) implies (B), consider two cases. If $\text{NP} \subseteq \text{BPP}$, then there is a trivial reduction that ignores the hypothesized HeurBPP algorithm for $(\text{NP}, \text{PSAMP})$. If $(\text{NP}, \text{PSAMP}) \not\subseteq \text{HeurBPP}$, then there is some problem in $(\text{NP}, \text{PSAMP})$ for which every algorithm is vacuously an appropriate type (2) decoder, because the universal quantification over HeurBPP algorithms for that problem is over an empty set.

The classes P^{PP} and PSPACE both have an $O(n)$ -query worst-case to average-case reduction under the uniform distribution, with type (1) decoding, using multilinear extensions [BFNW93]. Moreover, these results relativize.

For certain promise problems regarding lattices, e.g., certain versions of the shortest vector problem, worst-case to average-case reductions to problems in $(\text{NP}, \text{PSAMP})$ are known, with type (1) decoding. However, these lattice problems are not known or believed to be NP-hard. We refer to [BT06a, Reg06] for surveys of these results.

Another aspect of worst-case to average-case reductions is the *encoding*, which refers to the way in which L_2 depends on L_1 . Black-box encoding means that the algorithm that defines L_2 has oracle access to L_1 , and for *every* language L_1 (not just those in NP), if the corresponding L_2 has an efficient average-case algorithm then L_1 has an efficient worst-case algorithm (via one of the above four types of decoding).

Viola [Vio05a, Vio05b] proved two results about worst-case to average-case reductions with black-box encoders implementable in the polynomial-time hierarchy. In [Vio05a] he proved unconditionally that such a reduction with type (1) decoding does not exist. In [Vio05b] he proved that if such a reduction with type (4) decoding exists then PH is average-case hard, and thus basing the average-case hardness of PH on the worst-case hardness of PH in this way is no easier than unconditionally proving the average-case hardness of PH.

1.2 Results

Our first result concerns the class $\text{BPP}_{\parallel}^{\text{NP}}$. Recall that AvgZPP denotes the class of distributional problems with polynomial-time errorless average-case randomized algorithms.

Theorem 1. *There exists an oracle relative to which there is no reduction of type*

$$(\text{BPP}_{\parallel}^{\text{NP}}, \text{PSAMP}) \subseteq \text{AvgZPP} \Rightarrow \text{UP} \subseteq \text{BPP}.$$

Note that the type of reduction considered in Theorem 1 is weaker than a worst-case to average-case reduction for NP, because $\text{BPP}_{\parallel}^{\text{NP}}$ is larger than NP, AvgZPP is smaller than HeurBPP , and UP is smaller than NP. Ruling out weaker reductions yields a stronger result.

If we restrict our attention to reductions that use a limited number of queries, then we can handle classes even larger than $\text{BPP}_{\parallel}^{\text{NP}}$.

Theorem 2. *For every polynomial q there exists an oracle relative to which there is no q -query reduction of type*

$$(\text{PH}, \text{PSAMP}) \subseteq \text{AvgZPP} \Rightarrow \text{UP} \subseteq \text{BPP}.$$

Since $\text{BPP}_{\parallel}^{\text{NP}} \subseteq \text{PH}$ holds relative to every oracle, it may appear at first glance that Theorem 2 subsumes Theorem 1. The reason it does not is because of the order of the quantifiers. In Theorem 2, the reduction may not make as many queries as it likes; it may only make a fixed polynomial q number of queries even though its running time may be an arbitrarily high degree polynomial.

If we are willing to sacrifice all but two queries, then we can go quite a bit further than PH.

Theorem 3. *For every uniform complexity class of languages \mathcal{C} there exists an oracle relative to which there is no 2-query reduction of type*

$$(\mathcal{C}, \text{PSAMP}) \subseteq \text{AvgZPP} \Rightarrow \text{UP} \subseteq \text{BPP}.$$

The term “uniform complexity class of languages” has a somewhat technical meaning, which is explained in Section 2, but it encompasses all “ordinary” complexity classes such as PSPACE and EXP^{EXP} .

Our theorems can be generalized in various ways. For example, Theorem 1 and Theorem 2 both hold with AvgZPP replaced by the deterministic version AvgP , by essentially the same proofs.³ We have chosen to state the results using AvgZPP because we feel it is more natural to allow randomized algorithms in average-case complexity. As another example, Theorem 1 holds with BPP replaced by BQP, by inserting a quantum query lower bound for the OR function [BBBV97] at the appropriate point in the argument, instead of a randomized lower bound. We have chosen

³For Theorem 1 exactly the same proof works; for Theorem 2 a minor tweak is needed.

the particular statements of our three theorems so as to highlight the interesting aspects and make the relationships among them clear.

In the original ECCC version of this paper [Wat10], we also proved two results similar to Theorem 1. One is a generalization of Theorem 1 where $\text{BPP}_{\parallel}^{\text{NP}}$ is generalized to allow multiple rounds of adaptivity in the NP oracle access (up to $o(n/\log n)$ rounds). In the other result, $\text{BPP}_{\parallel}^{\text{NP}}$ is replaced with the class BPP_{path} , which was introduced in [HHT97] and which captures the power of polynomial-time randomized computations conditioned on efficiently testable events. Relative to every oracle, $\text{P}_{\parallel}^{\text{NP}} \subseteq \text{BPP}_{\text{path}} \subseteq \text{BPP}_{\parallel}^{\text{NP}}$, and thus this result is subsumed by Theorem 1. However, our proof of the BPP_{path} result is still interesting because the heart of the proof is genuinely different from the heart of our proof of Theorem 1, and it exploits the definition of BPP_{path} in a particularly intuitive way, without going through approximate counting. We refer the reader to [Wat10] for details about these results.

1.3 Independent Work

Independently of our work, Impagliazzo [Imp11] has succeeded in constructing a relativized heuristic, even for errorless average-case algorithms. In fact, he constructs an oracle relative to which $(\text{NP}, \text{PSAMP}) \subseteq \text{AvgP}$ but $\text{UP} \not\subseteq \text{P/poly}$. Thus relative to his oracle, there is no worst-case to average-case reduction for NP with any of the four types of decoding discussed in Section 1.1. This subsumes our result for NP (which only applies to black-box decoding).

The results of [Imp11] do not subsume our results for classes higher than NP, although Impagliazzo conjectures that this may be possible using his techniques. Furthermore, our techniques can be adapted without difficulty (though we do not argue this here) to show that there exists an oracle relative to which there is no reduction of type

$$(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP} \Rightarrow (\text{NP}, \text{PSAMP}) \subseteq \text{AvgZPP}$$

while it remains open to prove that there exists an oracle relative to which $(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP}$ but $(\text{NP}, \text{PSAMP}) \not\subseteq \text{AvgZPP}$.

Another benefit of our paper is that our techniques are genuinely different (and more elementary) than Impagliazzo’s. The outline of his argument is that he puts a random permutation into the oracle so that inverting the permutation is a worst-case hard NP problem, and to make NP average-case easy he puts the answers to all NP problems into the oracle but “censors” a certain small fraction of the answers in a way that preserves the worst-case hardness of inverting the permutation. Ensuring consistency between the conflicting requirements is a delicate business, and involves results on random restrictions of so-called matching DNFs. In contrast, our proof of our result for NP does not use such machinery. While much of the work in Impagliazzo’s proof is in the analysis rather than the construction of the oracle, our argument is more directly adversarial. We use a potential function technique to guide the construction to converge to an oracle with the desired properties. We employ elementary counting arguments to achieve this.

Also, our proofs of Theorem 2 and Theorem 3 illustrate a new connection between lower bounds for error-correcting codes and relativized lower bounds.

1.4 Organization

In Section 2 we provide preliminaries, which clarify the precise meanings of our theorems. In Section 3 we give the intuition for our proofs. In Section 4 we describe the basic setup that is common to

the formal proofs of all three theorems. Section 5 contains the formal proof of Theorem 1. Section 6 contains the formal proof of Theorem 2. Section 7 contains the formal proof of Theorem 3. In Section 8 we conclude the paper with a list of open problems regarding oracles in average-case complexity.

2 Preliminaries

We refer the reader to the textbooks [AB09, Gol08] for background on complexity theory and definitions of standard complexity classes. We refer the reader to the survey paper [BT06a] for background on average-case complexity. In this section we provide preliminaries that are not completely standard.

2.1 Complexity Classes

For any randomized algorithm M , we let M_r denote M using internal randomness r .

We now define the average-case complexity classes we need. Recall that in average-case complexity, we study distributional problems (L, D) where L is a language and $D = (D_1, D_2, \dots)$ is an ensemble of probability distributions, where D_n is distributed over $\{0, 1\}^n$. Recall that PSAMP denotes the class of polynomial-time samplable ensembles, and \mathcal{U} denotes the class consisting of only the uniform ensemble U . If \mathcal{C} is a class of languages and \mathcal{D} is a class of ensembles then $(\mathcal{C}, \mathcal{D}) = \{(L, D) : L \in \mathcal{C} \text{ and } D \in \mathcal{D}\}$.

Definition 1. *HeurBPP denotes the class of distributional problems (L, D) that have a polynomial-time heuristic scheme, that is, a randomized algorithm M that takes as input x and $\delta > 0$, runs in time polynomial in $|x|$ and $1/\delta$, and for all n and all $\delta > 0$ satisfies*

$$\Pr_{x \sim D_{n,r}} [M_r(x, \delta) \neq L(x)] \leq \delta.$$

Definition 2. *AvgZPP denotes the class of distributional problems (L, D) that have a polynomial-time errorless heuristic scheme, that is, a randomized algorithm M that takes as input x and $\delta > 0$, runs in time polynomial in $|x|$ and $1/\delta$, always outputs $L(x)$ or \perp , and for all n and all $\delta > 0$ satisfies*

$$\Pr_{x \sim D_{n,r}} [M_r(x, \delta) = \perp] \leq \delta.$$

2.2 Reductions

In this section we informally explain what we mean when we say there exists a reduction of type

$$\mathcal{C}'_2 \subseteq \mathcal{C}_2 \Rightarrow \mathcal{C}'_1 \subseteq \mathcal{C}_1$$

where $\mathcal{C}'_2, \mathcal{C}_2, \mathcal{C}'_1, \mathcal{C}_1$ are four complexity classes. In Section 2.3 below we give formal definitions for the specific classes to which our theorems apply.

A complexity class is a set of computational problems, such as languages or distributional problems. We assume for concreteness that each of \mathcal{C}_1 and \mathcal{C}_2 is defined in the following way. By an *input-output relationship* we mean a randomized function. There is a set of algorithms, each of which induces an input-output relationship. That is, each algorithm takes an input and

produces an output sampled from some distribution depending on the input. There is a predicate that indicates for each input-output relationship and each computational problem whether the input-output relationship *solves* the problem. There is a notion of computational resources used by the algorithms, and an algorithm is said to be *efficient* if it satisfies certain resource constraints. The class is defined as the set of problems solved by efficient algorithms. This type of definition encompasses classes defined in terms of (uniform or nonuniform) deterministic, randomized, or quantum algorithms, but it could be generalized to handle other models as well.

We also assume that for \mathcal{C}_1 there is an analogous set of algorithms that can make queries to a *reduction oracle*, which represents an input-output relationship.⁴ We assume that plugging any algorithm from \mathcal{C}_2 's set into the reduction oracle yields an algorithm from \mathcal{C}_1 's set.

Now suppose P_1 is a computational problem of the appropriate kind for \mathcal{C}_1 and P_2 is a computational problem of the appropriate kind for \mathcal{C}_2 .

Definition 3. *A reduction of type*

$$P_2 \in \mathcal{C}_2 \Rightarrow P_1 \in \mathcal{C}_1$$

is an algorithm from \mathcal{C}_1 's set of reduction oracle algorithms, such that (i) for every reduction oracle that solves P_2 according to \mathcal{C}_2 , the reduction solves P_1 according to \mathcal{C}_1 , and (ii) for every reduction oracle, the reduction satisfies \mathcal{C}_1 's resource constraints if we pretend each query to the reduction oracle uses any amount of resources allowed by \mathcal{C}_2 's resource constraints.

In other words, if the reduction oracle is correct then the reduction is correct, and if we pretend the reduction oracle is efficient then the reduction is efficient.

Note that if we plug an actual, efficient algorithm for P_2 (according to \mathcal{C}_2) into the reduction oracle of such a reduction, then the reduction becomes an efficient algorithm for P_1 (according to \mathcal{C}_1). Thus if there exists a reduction satisfying Definition 3 then $P_2 \in \mathcal{C}_2$ implies $P_1 \in \mathcal{C}_1$. But the reduction must work even when the reduction oracle is an input-output relationship that is not efficiently implementable.

As an example, suppose $\mathcal{C}_2 = \text{BPTIME}(2^{n^\epsilon})$. Then the reduction must solve P_1 according to \mathcal{C}_1 when the reduction oracle is any randomized function from $\{0, 1\}^*$ to $\{0, 1\}$ that, on input w , returns $P_2(w)$ with probability $\geq 2/3$.⁵ Further, the reduction must satisfy the resource constraints of \mathcal{C}_1 when we pretend each query of length n to the reduction oracle takes time $O(2^{n^\epsilon})$.

Definition 4. *We say there exists a reduction of type*

$$\mathcal{C}'_2 \subseteq \mathcal{C}_2 \Rightarrow \mathcal{C}'_1 \subseteq \mathcal{C}_1$$

if for every $P_1 \in \mathcal{C}'_1$ there exists a $P_2 \in \mathcal{C}'_2$ and a reduction of type

$$P_2 \in \mathcal{C}_2 \Rightarrow P_1 \in \mathcal{C}_1.$$

We make a few remarks about Definition 4.

⁴In particular, the reduction oracle is not like a relativization oracle, which just answers queries to a language.

⁵One might wonder about reductions that can also choose the randomness used by the reduction oracle. While this would be more general in one sense, it would be more restrictive in the sense that it would limit the randomness complexity of the reduction oracle. In this paper, queries are always just inputs to an input-output relationship as defined above.

- When \mathcal{C}'_1 has an appropriately complete problem P_1 , this is equivalent to saying there exists a $P_2 \in \mathcal{C}'_2$ and a reduction of the above type, for the fixed problem P_1 .
- Note that we do not require that the reduction is uniform in the sense of there being a fixed algorithm R that computes the reduction for every $P_1 \in \mathcal{C}'_1$ given the code for a \mathcal{C}'_1 -type algorithm for P_1 .
- Note that when we say there is a reduction of the above type, this assertion gets weaker as \mathcal{C}'_2 and \mathcal{C}_1 get larger and \mathcal{C}_2 and \mathcal{C}'_1 get smaller.

2.3 Relativization

When we relativize to an oracle language A , every computation gets unrestricted oracle access to A . This includes samplers and reductions. Thus reductions have access to two oracles: the reduction oracle and the relativization oracle. When we write $R^{B,A}$ we mean B is the reduction oracle and A is the relativization oracle for reduction R .

To illustrate the formal framework set up so far, we give the precise statement of Theorem 1. There exists a language A and a language $L_1 \in \text{UP}^A$ such that for all languages $L_2 \in (\text{BPP}_{\parallel}^{\text{NP}})^A$, all ensembles $D \in \text{PSAMP}^A$, and all polynomial-time randomized reductions $R^{\circ,\circ}$, $R^{\circ,A}$ is not of type

$$(L_2, D) \in \text{AvgZPP}^A \Rightarrow L_1 \in \text{BPP}^A.$$

The latter means that there exists an $x \in \{0, 1\}^*$ and a randomized function $B : \{0, 1\}^* \times \mathbb{R}_{>0} \rightarrow \{0, 1, \perp\}$ which is a valid AvgZPP oracle for (L_2, D) , such that

$$\Pr_{r,B} \left[R_r^{B,A}(x) = L_1(x) \right] < 2/3$$

where the probability is over both the internal randomness of R and the randomness of B (each query is answered with fresh independent randomness). When we say B is a valid AvgZPP oracle for (L_2, D) we mean that $B(w, \delta)$ always returns $L_2(w)$ or \perp , and for all n and all $\delta > 0$,

$$\Pr_{w \sim D_n, B} [B(w, \delta) = \perp] \leq \delta.$$

When we say $R^{\circ,\circ}$ runs in polynomial time, this includes the fact that each query $B(w, \delta)$ to the reduction oracle is charged time polynomial in $|w|$ and $1/\delta$. In other words, δ must always be at least inverse polynomial. Throughout the paper we tacitly assume that “polynomial-time reductions” have this restriction, since \mathcal{C}_2 is always AvgZPP. We clarify that $D \in \text{PSAMP}^A$ means that for some randomized algorithm S° , $S^A(n)$ runs in time polynomial in n and outputs a sample distributed according to D_n . Finally, we clarify that $(\text{BPP}_{\parallel}^{\text{NP}})^A$ is the class of languages L_2 for which there exists a language $L_3 \in \text{NP}^A$ and a polynomial-time randomized algorithm $M^{\circ,\circ}$ that accesses its first oracle nonadaptively, such that for all $x \in \{0, 1\}^*$,

$$\Pr_r \left[M_r^{L_3,A}(x) = L_2(x) \right] \geq 2/3.$$

Regarding Theorem 2 and Theorem 3, there is one further issue to consider. For reductions that are allowed an unlimited number of queries (like in Theorem 1), the error probability of $1/3$ in the definition of BPP is unimportant since it can be amplified from $1/2 - 1/\text{poly}(n)$ to $1/2^{\text{poly}(n)}$.

However, amplification increases the number of queries, so the error probability is not arbitrary for Theorem 2 and Theorem 3. For example, the existence of a q -query $(1/2 - 1/\text{poly}(n))$ -error reduction of type

$$(\text{PH}, \text{PSAMP}) \subseteq \text{AvgZPP} \Rightarrow \text{UP} \subseteq \text{BPP}$$

does not seem to imply the existence of a q -query $1/3$ -error reduction of the same type, but it still does imply that if $(\text{PH}, \text{PSAMP}) \subseteq \text{AvgZPP}$ then $\text{UP} \subseteq \text{BPP}$. For this reason, we allow an error probability of $1/2 - 1/\text{poly}(n)$ (for arbitrarily high degree polynomials) in Theorem 2 and in Theorem 3.

2.4 Uniform Complexity Classes

We now precisely define the restriction on \mathcal{C} in Theorem 3.

Definition 5. *We say that \mathcal{C} is a uniform complexity class of languages if there is a countable collection of functions $\{M_1, M_2, \dots\}$ mapping oracle languages A to languages M_i^A , such that the following three conditions all hold.*

- *For every i and every x , $M_i^A(x)$ only depends on a finite number of bits of A .*
- *For every i and every x there exists a property $P_{i,x}(A)$ that only depends on the bits of A that $M_i^A(x)$ depends on, such that $\mathcal{C}^A = \{M_i^A : \forall x P_{i,x}(A)\}$.*
- *For every i and every linear-time computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ there exists a j such that for all A the following two conditions hold: $M_j^A = M_i^A \circ f$, and if $M_i^A \in \mathcal{C}^A$ then $M_j^A \in \mathcal{C}^A$.*

The second condition says the class is defined by a property of the computation (for example, bounded error) holding for all inputs. The third condition says the class is closed under linear-time deterministic mapping reductions. Observe that $\text{BPP}_{\parallel}^{\text{NP}}$, PH , PSPACE , and EXP^{EXP} are all examples of uniform complexity classes under this definition.

3 Intuition

In Section 3.1 we describe the intuition behind the proof of Theorem 1. Then in Section 3.2 we describe the intuition behind the proofs of Theorem 2 and Theorem 3.

3.1 Intuition for Theorem 1

We start by informally describing how to construct an oracle relative to which there is no reduction of type

$$(\text{NP}, \mathcal{U}) \subseteq \text{HeurBPP} \Rightarrow \text{UP} \subseteq \text{BPP}.$$

To obtain Theorem 1 we must strengthen HeurBPP to AvgZPP,⁶ strengthen \mathcal{U} to PSAMP, and strengthen NP to $\text{BPP}_{\parallel}^{\text{NP}}$. We describe how to do these things below.

⁶Usually AvgZPP is thought of as being a weaker class than HeurBPP (since $\text{AvgZPP} \subseteq \text{HeurBPP}$), but it is stronger in our situation.

Fix an arbitrary NP-type algorithm M and an arbitrary polynomial-time randomized reduction R , and fix a sufficiently large n . For simplicity we assume that on inputs of length n , R only queries the reduction oracle on inputs of length n^d (for some positive integer d) and only with some fixed polynomially small δ ; thus we can omit the δ from the queries. We consider relativization oracles of the form $A : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, which we think of as $2^n \times 2^n$ tables. Let $L_1^A : \{0, 1\}^n \rightarrow \{0, 1\}$ be defined by $L_1^A(x) = \bigvee_y A(xy)$. That is, L_1^A is the language of strings x such that there exists a 1 in the x th row of A . Let $L_2^A : \{0, 1\}^{n^d} \rightarrow \{0, 1\}$ denote the language computed by M^A . We only consider A, L_1^A, L_2^A at these input lengths since all other input lengths are irrelevant.

We explain how to diagonalize against the pair M, R . We wish to construct an A such that for some $x \in \{0, 1\}^n$ and some deterministic⁷ reduction oracle $B : \{0, 1\}^{n^d} \rightarrow \{0, 1\}$, B agrees with L_2^A on at least a $1 - \delta$ fraction of inputs and $R^{B,A}(x)$ outputs $L_1^A(x)$ with probability $< 2/3$. This will show that R fails to be a reduction of type

$$(L_2^A, U) \in \text{HeurBPP}^A \Rightarrow L_1^A \in \text{BPP}^A.$$

We also need to ensure that there is at most one 1 in each row of A so that $L_1^A \in \text{UP}^A$, but this will fall right out of the construction. We construct A through an iterative process, and we use a potential function argument to show that this process makes steady progress toward our goal. The process iteratively modifies the relativization oracle, and we use A to denote the relativization oracle throughout the whole process.⁸ Thus the table denoted by A changes many times throughout our argument, and the languages L_1^A and L_2^A change accordingly. Initially A is all 0's.

Let us consider the computation of R on some input x . It is trying to figure out whether there is a 1 in the x th row of A , in other words, compute $L_1^A(x)$. It has two sources of information about $L_1^A(x)$: the relativization oracle A itself, and the reduction oracle B . If R did not have access to B , then we could diagonalize in a standard way: Observe how R behaves given that the x th row of A is all 0's. If R outputs 1 with high probability, then we are done. If R outputs 1 with low probability, then we find a bit in the x th row that R queries with only tiny probability and flip that bit (such a bit must exist because R does not have enough time to keep an eye on the entire row); then R still outputs 1 with low probability, but now $x \in L_1^A$. Thus R must rely on the reduction oracle B for help.

Our construction has two stages. The goal of stage 1 is to gain the upper hand by rendering B useless to R . Then in stage 2 we deliver the coup de grâce with the standard diagonalization argument. We cannot guarantee that B is useless for every x , but we only need it to be useless for some x . Specifically, suppose we could set up A in such a way that there exists an x such that

- (1) the x th row of A is all 0's, and
- (2) for all y , flipping $A(xy)$ would cause $L_2^A(w)$ to change for at most a δ fraction of w 's.

Then declaring B to be L_2^A for the particular A we have set up, we know that we can leave A alone or we can flip any bit in the x th row, and for all these possibilities B is a valid HeurBPP oracle for the new L_2^A . Then we can observe the behavior of R on input x , using this fixed B for the

⁷ B will be deterministic here even though randomness is allowed; this makes the result stronger.

⁸More formally, we could say we define a sequence of relativization oracles A_0, A_1, A_2, \dots that leads to some final version $A_k = A$. We omit the subscripts throughout the argument and simply refer to A with the understanding that this means the "current" version.

reduction oracle, and diagonalize against R in the standard way with the assurance that whatever happens to A during this second stage, B will remain valid.

How do we set up A so that such an x exists? We do this iteratively. In each iteration, we find a certain x whose row is currently all 0's, which is our "best guess" for the good x . If condition (2) is satisfied for this x , then we are done. Otherwise, there is some column y that violates condition (2). Then we flip the bit $A(xy)$ to 1 and continue with the next iteration. We just need to show that there are $< 2^n$ iterations before we succeed. For this, we define a potential function Φ^A that assigns an energy value to A . The key is to show that if y violates condition (2) for our best guess x , then flipping $A(xy)$ must cause a significant decrease in potential. Since Φ^A must remain bounded, there cannot be too many iterations before M is beaten into submission and our best guess x works.

Let us hold off on the definition of Φ^A and focus on finding a best guess x . Our ultimate goal is to ensure that if we flip any bit in the x th row, most of the inputs to L_2^A "don't notice". There is an asymmetry between inputs that are accepted by M^A and those that are rejected. If $w \in \{0, 1\}^{n^d}$ is such that $M^A(w)$ rejects, then if *any* of the exponentially many computation paths "notices" a change in A , the whole computation could become accepting. However, if $M^A(w)$ accepts, then we can pick an arbitrary accepting computation path of $M^A(w)$ to be the "designated" one. Only polynomially many bits of A are queried by M on this path, and as long as none of these bits is flipped, w "won't notice" any change to A because $M^A(w)$ will still accept. In particular, there are only polynomially many x 's such that $M^A(w)$ queries some bit in the x th row on the designated path. Thus for every w with $L_2^A(w) = 1$, the vast majority of x have the property that flipping any bit in the x th row does not cause $L_2^A(w)$ to change to 0. By an averaging argument, most x have the property that for most $w \in \{0, 1\}^{n^d}$, flipping any bit in the x th row does not cause $L_2^A(w)$ to change from 1 to 0. For the current A , there must exist an x with the latter property and such that the x th row is all 0's, since (by induction) we know there are not very many x 's with a 1 in their row currently. This is our best guess x .

We know that flipping any bit in the x th row causes only a small fraction of all $w \in \{0, 1\}^{n^d}$ to change from 1 to 0 under L_2^A . This is good, but it is only half the story. We would also like that flipping any bit in the x th row causes only a small fraction of w 's to change from 0 to 1. Suppose we budget a $\delta/2$ fraction of w 's to change from 1 to 0, and a $\delta/2$ fraction to change from 0 to 1. Now if some y violates condition (2), then it must be the case that flipping $A(xy)$ causes at least a $\delta/2$ fraction of w 's to change from 0 to 1. We want to define the potential function so that having w 's change from 0 to 1 under L_2^A causes a decrease in potential. A natural choice is

$$\Phi^A = \Pr_{w \sim U_{n^d}} [L_2^A(w) = 0].$$

Flipping $A(xy)$ causes at least a $\delta/2$ probability mass to leave the event $L_2^A(w) = 0$. However, as much as a $\delta/2$ probability mass could enter the event due to w 's that change from 1 to 0, which could essentially cancel out the drop in potential from the w 's that changed from 0 to 1! The solution is to change our budgeting. If we budget a $\delta/3$ fraction of w 's to change from 1 to 0 and a $2\delta/3$ fraction to change from 0 to 1, then flipping $A(xy)$, where y violates condition (2), causes at least a $2\delta/3$ probability mass to leave the event, while at most a $\delta/3$ probability mass enters the event. Thus Φ^A goes down by at least $\delta/3$, and there are at most $3/\delta < 2^n$ iterations before our best guess x works. This concludes the argument.

Very roughly, the big picture is as follows. For an input that is accepted by M^A , it is easy to ensure that the answer under L_2^A does not change when we make modifications to A . For an input

that is rejected by M^A , we cannot ensure that the answer does not change, but the point is that if it does change, then *we can ensure that it does not change again*, since the input is now accepted.

3.1.1 Intuition for Strengthening HeurBPP to AvgZPP

Let x denote our best guess at the end of stage 1. Suppose we knew that there exists a set $W \subseteq \{0,1\}^{n^d}$ of density at most δ such that for all $w \notin W$ and all y , flipping $A(xy)$ does not change $L_2^A(w)$. Then setting

$$B(w) = \begin{cases} L_2^A(w) & \text{if } w \notin W \\ \perp & \text{if } w \in W \end{cases} \quad (1)$$

where A is the relativization oracle at the end of stage 1, we would have that B is a valid AvgZPP oracle for L_2^A no matter whether we leave A alone or flip any bit in the x th row. Then we could diagonalize in the standard way, by observing how R behaves on input x using this fixed B and the current A , and either leaving A alone or flipping some bit in the x th row to make R output the wrong answer with high probability.

The existence of such a W is too much to ask for. However, this is only because we were trying to find a B that would remain a valid AvgZPP oracle for *all* of the $2^n + 1$ diagonalization options. We do not really need all these options. Let Y be an arbitrary fixed set of columns of size $|Y| = 4t$, where t is the running time of R on inputs of length n . Then running R on input x with any fixed B and the current A , there must be a $y \in Y$ such that $A(xy)$ gets queried with probability $\leq 1/4$. If R outputs 1 with probability $\leq 1/3$ then after flipping this $A(xy)$, R outputs 1 with probability $< 2/3$ and hence errs. Thus it suffices to have $4t + 1$ diagonalization options, namely leaving A alone or flipping some $A(xy)$ with $y \in Y$. Suppose we knew that there exists a set $W \subseteq \{0,1\}^{n^d}$ of density at most δ such that for all $w \notin W$ and all $y \in Y$, flipping $A(xy)$ does not change $L_2^A(w)$. Then defining B as in Equation (1), we could diagonalize by either leaving A alone or flipping $A(xy)$ for some $y \in Y$ with the assurance that whatever happens, B will remain valid.

Now the existence of such a W is *not* too much to ask for. Using the argument for the HeurBPP case with a small adjustment of parameters, we can ensure that flipping any bit in the x th row causes $L_2^A(w)$ to change for at most a $\delta/4t$ fraction of w 's. Then we can take W to be the set of all w such that there exists a $y \in Y$ such that flipping $A(xy)$ changes $L_2^A(w)$.

3.1.2 Intuition for Strengthening \mathcal{U} to PSAMP

There are two approaches: one that is direct, and one that uses a result of Impagliazzo and Levin [IL90]. Neither is difficult. We first describe the direct approach.

First, observe that if U_{n^d} were replaced by some other distribution on $\{0,1\}^{n^d}$ that is independent of A , then the whole argument above would carry through, just by replacing “fraction of w 's” with “probability mass of w 's” under this distribution. Now in addition to M and R , we need to worry about an arbitrary polynomial-time sampler S , and we need to ensure that B is a valid AvgZPP oracle for (L_2^A, D^A) , where D^A denotes the distribution sampled by $S^A(n^d)$. If S did not query A at all, then D^A would be independent of A and thus we could use the same argument, by the above observation. Two issues arise because S is allowed to query A . First, when we flip a bit during stage 1, this affects

$$\Phi^A = \Pr_{w \sim D^A} [L_2^A(w) = 0]$$

in terms of not only the event but also the distribution. Second, when we flip a bit during stage 2, this affects the distributional problem (L_2^A, D^A) for which B needs to be a valid AvgZPP oracle, in terms of not only the language but also the distribution.

Handling these issues is just a matter of tweaking the argument to ensure that our modifications to A cause only small statistical deviations in D^A . Specifically, consider the beginning of an iteration of stage 1, and let D denote D^A for the current A (thus D is fixed and will not react to changes in A). Now suppose we choose our best guess x as before, but based on this distribution D . Then by the above argument we know that for every y , flipping $A(xy)$ would either cause

$$\Pr_{w \sim D} [L_2^A(w) = 0]$$

to go down by a significant amount, or cause $L_2^A(w)$ to change with only small probability over $w \sim D$. It can be shown that this is good enough for our purpose provided that for all y , flipping $A(xy)$ results in a D^A that is statistically very close to D . To ensure the latter, we choose our best guess x not only so that the x th row is all 0's and flipping any bit in the x th row only causes a small probability mass of $w \sim D$ to change from 1 to 0 under L_2^A , but also so that the probability $S^A(n^d)$ queries any bit in the x th row is small. This is possible because the vast majority of x 's satisfy the latter condition since S runs in polynomial time.

An alternative approach to handling PSAMP uses a result due to Impagliazzo and Levin [IL90]. They proved that if \mathcal{C} is a class of languages containing NP and satisfying certain simple closure properties, then relative to every oracle, there exists a reduction of type

$$(\mathcal{C}, \mathcal{U}) \subseteq \text{AvgZPP} \Rightarrow (\mathcal{C}, \text{PSAMP}) \subseteq \text{AvgZPP}.$$

The proof of this result appears in Section 5.2 of [BT06a] and is based on a result of Impagliazzo and Luby on distributionally inverting one-way functions [IL89]. By composing this reduction with the hypothesized reduction, we can assume without loss of generality that the distributional problem we are reducing to uses the uniform ensemble. In the formal proof of Theorem 1, rather than use the Impagliazzo-Levin result we opt to directly handle the samplable ensembles because doing so makes the argument self-contained at only a slight cost in complicatedness.

3.1.3 Intuition for Strengthening NP to $\text{BPP}_{\parallel}^{\text{NP}}$

There exists a simple reduction of type

$$(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP} \Rightarrow (\text{BPP}_{\parallel}^{\text{NP}}, \text{PSAMP}) \subseteq \text{HeurBPP}$$

(and this fact relativizes). Thus if we consider HeurBPP rather than AvgZPP, then the result for $\text{BPP}_{\parallel}^{\text{NP}}$ follows from the result for NP by composing reductions. To handle AvgZPP, we directly adapt the NP argument to work for $\text{BPP}_{\parallel}^{\text{NP}}$. Let us revert from PSAMP to \mathcal{U} .

Instead of a single algorithm M we have a pair M, N where N is an NP-type algorithm and M is a polynomial-time randomized algorithm that accesses its first oracle nonadaptively. We let L_3^A denote the language computed by N^A , and we let L_2^A denote the language computed by $M^{L_3^A, A}$ (assuming bounded error is satisfied for every input).⁹ Let us make the simplifying assumption

⁹We again only deal with L_2^A on inputs of length n^d , but we consider L_3^A on all input lengths. We could assume all queries M makes to its first oracle have the same length, but it turns out this would not make the proof any simpler.

that M has oracle access only to L_3^A and not to A . (Extending the argument to the general case is not difficult; it just involves taking an extra precaution when picking our best guess x to ensure that hardly any w 's "notice" changes to A via the second oracle.)

The differences from the above proof are in the definition of the potential function Φ^A , the choice of our best guess x , and the argument that if some y violates condition (2) for our best guess x , then flipping $A(xy)$ causes a significant decrease in potential. Let $M_r^{L_3^A}(w)_j \in \{0, 1\}^*$ denote the j th query to L_3^A made by the computation $M_r^{L_3^A}(w)$ (in our simplified setting, this query does not depend on A), and consider the bits

$$L_3^A\left(M_r^{L_3^A}(w)_j\right)$$

over the choice of w, r, j . We define Φ^A to be the fraction of these bits that are 0.

Consider an arbitrary iteration of stage 1, and let A denote the current relativization oracle. By choosing our best guess x appropriately, we can ensure that the x th row of A is all 0's, and no matter what y is, only a tiny fraction of the w, r, j bits go from 1 to 0 when we flip $A(xy)$ (and thus bits going from 1 to 0 can only contribute a tiny increase in potential). Suppose there is a y such that flipping $A(xy)$ causes $L_2^A(w)$ to change for a significant fraction of w 's. We want it to be the case that flipping $A(xy)$ also causes a significant decrease in potential, and for this it suffices to show that a significant fraction of w, r, j bits go from 0 to 1. Let A' denote A with $A(xy)$ flipped to 1. For each w such that $L_2^{A'}(w) \neq L_2^A(w)$, it must be the case that

$$M_r^{L_3^{A'}}(w) \neq M_r^{L_3^A}(w) \tag{2}$$

for at least $1/3$ of the r 's. Thus we know that Inequality (2) holds for a significant fraction of pairs w, r . If Inequality (2) holds for w, r then there must exist a j such that the w, r, j bit changes when we go from A to A' . But the fraction of pairs w, r such that the w, r, j bit goes from 1 to 0 for some j is tiny (at most polynomially larger than the fraction of triples w, r, j that go from 1 to 0). Thus a significant fraction of pairs w, r are such that the w, r, j bit goes from 0 to 1 for some j , and hence a significant fraction (possibly a polynomially smaller fraction) of triples w, r, j go from 0 to 1. Thus we have a significant decrease in potential when we flip $A(xy)$.

3.2 Intuition for Theorem 2 and Theorem 3

It is well-known that error-correcting codes can be used to construct worst-case to average-case reductions, at least for large complexity classes such as PSPACE [BFNW93, STV01]. To be applicable, the codes must have very efficient encoders (since this dictates the complexity of the language being reduced to) and very efficient decoders (since this dictates the complexity of the reduction itself). Our strategy for proving Theorem 2 and Theorem 3 is to set up the relativization oracle in such a way that error-correcting codes are in some sense the *only* way to construct worst-case to average-case reductions of the appropriate types, and then argue that the efficiency of the resulting encoders and decoders is too good to be true. That is, we would like to be able to extract a good error-correcting code from any purported reduction and then apply known lower bounds on the efficiency of encoders and decoders for such codes. For Theorem 2, we use a result due to Viola [Vio05a] which states that good error-correcting codes¹⁰ cannot be encoded by small constant-depth

¹⁰His result even applies to list-decodable codes, but we do not need this stronger result.

circuits. For Theorem 3, we use a lower bound due to Kerenidis and de Wolf [KdW04] on the length of 2-query locally decodable codes.

Our approach for Theorem 2 and Theorem 3 is in some sense a *dual* approach to the one we used for Theorem 1. As before, we have a reduction R that is trying to solve a problem with the aid of a relativization oracle A and a reduction oracle B . Before, our goal was to *render B useless to R* so we could focus on how R interacted with A . Now, our goal is to *render A useless to R* so we can focus on how R interacts with B . Before, we found a good *row* of A and filled in that row adversarially. Now, we find a good *column* of A and fill in that column adversarially.

Unlike in the proof of Theorem 1, we cannot use the Impagliazzo-Levin result to reduce PSAMP to \mathcal{U} since it uses too many queries. But again, directly handling the samplable ensembles presents no major difficulties. Thus, for the rest of this section we assume PSAMP is replaced by \mathcal{U} .

The basic setup is the same as before. We have an algorithm M (PH-type for Theorem 2 or arbitrary complexity for Theorem 3). We have a polynomial-time randomized reduction R that uses a limited number of queries to the reduction oracle. For simplicity we assume that on inputs of length n , R only queries the reduction oracle on inputs of length n^d (for some positive integer d) and only with some fixed polynomially small δ . We construct a sequence of relativization oracles $A : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, and we define $L_1^A : \{0, 1\}^n \rightarrow \{0, 1\}$ by $L_1^A(x) = \bigvee_y A(xy)$, and we let $L_2^A : \{0, 1\}^{n^d} \rightarrow \{0, 1\}$ denote the language computed by M^A . For the final version of A , we want $R^{B,A}(x)$ to output $L_1^A(x)$ with probability $< 1/2 + 1/n^{\log n}$ for some $x \in \{0, 1\}^n$ and some $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ that agrees with L_2^A on at least a $1 - \delta$ fraction of inputs and returns \perp on the rest. We have $1/2 + 1/n^{\log n}$ instead of $2/3$ for the reason discussed at the end of Section 2.3.

Let us start by pretending that R never queries A . Then it is completely straightforward to extract a good binary error-correcting code from M, R : Pick an arbitrary column y and define

$$C : \{0, 1\}^{2^n} \rightarrow \{0, 1\}^{2^{n^d}}$$

by viewing the input as a function $Z : \{0, 1\}^n \rightarrow \{0, 1\}$ and the output as a function $C(Z) : \{0, 1\}^{n^d} \rightarrow \{0, 1\}$ given by $C(Z) = L_2^{A_Z}$ where A_Z denotes the relativization oracle with Z as the y th column and 0's everywhere else. If R really is of the hypothesized type no matter which Z we use, then it immediately follows that R is a decoder that recovers any bit $Z(x) = L_1^{A_Z}(x)$ of the information word from any corrupted code word B that has at most a δ fraction of erasures (and no flipped bits).

For Theorem 2, note that C has relative minimum distance $> \delta$ and each bit of C is encodable by a small constant-depth circuit since M is a PH-type algorithm with oracle access to Z [FSS84]. This contradicts a result of Viola [Vio05a] which says that such a code cannot exist. Thus there must be some Z for which R is not of the hypothesized type.

For Theorem 3, note that C is a 2-query locally decodable code in the sense that each bit of the information word can be recovered with probability at least $1/2 + 1/n^{\log n}$ assuming there are at most a δ fraction of erasures.¹¹ Since the code word length is only *quasipolynomial* in the information word length, this contradicts a result of Kerenidis and de Wolf [KdW04] which says that the length of such a code must be *nearly exponential*.¹² Thus there must be some Z for which

¹¹Usually, locally decodable codes are defined in terms of flipped bits rather than erasures, but they are equivalent up to small differences in parameters.

¹²The lower bound is only *nearly exponential* since the relative minimum distance and the advantage over $1/2$ in correct decoding probability are subconstant in our case.

R is not of the hypothesized type. Since the lower bound holds regardless of the complexity of encoding, we can handle *any* uniform complexity class of languages.

Now we return to the “real world” where R may query A . Then the above argument, with an arbitrary fixed y , does not work because R might know y in which case R can easily go look up the answers to L_1^A in the y th column. We must choose y so as to “hide” the answers from R . Restricting the number of queries R can make to B is essential for this: If R can make n queries then M can easily let R know what y is by explicitly writing y over and over again in the truth table L_2^A , and R would have no trouble retrieving this information from any B that has sufficient agreement with L_2^A . (Of course in Theorem 2, R can use n , or any fixed polynomial, number of queries. But this is easily remedied by just adding $2^{\text{poly}(n)}$ columns to the table A , with a high enough degree polynomial, so that we can hide the answers from R . Henceforth we assume R only uses $n^{o(1)}$ queries, so that we can stick with 2^n columns.)

Suppose we could choose y so that for every x and every $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$, the probability that $R^{B,0}(x)$ (where 0 denotes the all 0’s relativization oracle) queries a bit in the y th column is at most $1/2n^{\log n}$. Then we would know that for every Z , every x , and every B that is valid for L_2^{AZ} , the probability $R^{B,0}(x)$ outputs $L_1^{AZ}(x)$ is within $1/2n^{\log n}$ of the probability $R^{B,AZ}(x)$ outputs $L_1^{AZ}(x)$ and is hence at least $1/2 + 1/2n^{\log n}$. This would suffice for a contradiction, because we could use $R^{B,0}$ for the decoder. Actually this property of y is more than we really need. If we replace “every x ” with “most x ” then we could just remove the bad x ’s from consideration, at a small loss in the information word length, and we would still get a contradiction. Now to find such a y , we use the fact that quantifying over all B is the same as quantifying over all paths of adaptivity in R ’s access to B , and there are a limited number of such paths. Specifically, for every x and every r there are only a small number of columns of the relativization oracle that get queried by $R_r^{o,0}(x)$ over all possible reduction oracles (namely, at most the running time of R times 3 to the number of reduction oracle queries). By an averaging argument, there is some y such that for most x ’s, all but a $1/2n^{\log n}$ fraction of r ’s are such that $R_r^{B,0}(x)$ does not query any bit in the y th column, for any B . This is good enough for our purpose.

The bottom line is that there are basically only two ways M could help R solve L_1^A : by telling R the answers, or by telling R where to find the answers in A . The former is impossible because then we would have an error-correcting code that is too good to be true, and the latter is impossible because R cannot make enough queries to B to retrieve the identity of y .

4 Generic Setup for the Formal Proofs

We first need the following complicated-looking lemma, which just says that in all three of our theorems, we can assume without loss of generality that on inputs of length n , any candidate reduction only queries the reduction oracle on inputs of length n^d and only with $\delta = 1/n^d$ for some positive integer d .

Lemma 1. *For every polynomial-time randomized reduction $R^{\circ,\circ}$ (where the reduction oracle is of the form $\{0, 1\}^* \times \mathbb{R}_{>0} \rightarrow \{0, 1, \perp\}$) there exists a polynomial-time randomized reduction $R_{\text{clean}}^{\circ,\circ}$ and a positive integer d such that the following holds. For every polynomial-time sampler S° there exists a polynomial-time sampler S_{clean}° , and for every uniform complexity class of languages \mathcal{C} and every i there exists an i_{clean} , such that for every relativization oracle A , the following properties all hold.*

- If $R^{\circ, A}$ is of type

$$(M_i^A, D^A) \in \text{AvgZPP}^A \Rightarrow L \in \text{BPP}^A$$

for some language L , where D^A is the ensemble sampled by S^A , then $R_{\text{clean}}^{\circ, A}$ is of type

$$(M_{i_{\text{clean}}}^A, D_{\text{clean}}^A) \in \text{AvgZPP}^A \Rightarrow L \in \text{BPP}^A$$

where D_{clean}^A is the ensemble sampled by S_{clean}^A .

- On inputs of length n , R_{clean} only queries the reduction oracle on inputs of length n^d and only with $\delta = 1/n^d$.
- R_{clean} always makes the same number of queries to the reduction oracle as R does.
- If $M_i^A \in \mathcal{C}^A$ then $M_{i_{\text{clean}}}^A \in \mathcal{C}^A$.

Proof sketch. The basic idea is to take the answers to all the inputs to M_i^A up to the longest length R on inputs of length n could possibly query the reduction oracle, and put them in some larger input length n^d . Here d needs to be large enough that $1/n^d$ times the longest length R could query is less than the smallest value of δ that R could possibly query (which is at least inverse polynomial). The reason for multiplying by the longest length is that an error of $1/n^d$ in the AvgZPP oracle could get amplified by this amount when restricted to any particular input length that is stored “within” n^d . The index i_{clean} is just the j guaranteed by Definition 5 for index i and the mapping reduction we just informally described. \square

We now describe the basic setup that is common to the proofs of all three theorems. However, this setup will need to be customized a bit for each of the three proofs.

We have a uniform complexity class of languages \mathcal{C} with enumeration $\{M_1, M_2, \dots\}$. Consider an arbitrary triple i, S, R where $i \in \mathbb{N}$, S is a polynomial-time sampler, and R is a polynomial-time randomized reduction. Using Lemma 1 we can assume without loss of generality that on inputs of length n , R only queries the reduction oracle on inputs of length n^d and only with $\delta = 1/n^d$ for some positive integer d . For an arbitrary relativization oracle $A \subseteq \{0, 1\}^*$ we make the following definitions. Let L_1^A denote the NP^A language defined by

$$L_1^A = \{x : \exists y \text{ such that } |y| = |x| \text{ and } xy \in A\}.$$

If M_i^A defines a language in \mathcal{C}^A then let L_2^A denote this language.¹³ Let D^A denote the PSAMP^A ensemble defined by S^A .

We wish to construct a relativization oracle A^* so that $L_1^{A^*} \in \text{UP}^{A^*}$ (by ensuring that in the definition of $L_1^{A^*}$, y is always unique if it exists) and so that for all i, S, R , either $M_i^{A^*}$ fails to define a language in \mathcal{C}^{A^*} , or otherwise

$$\Pr_{r, B} \left[R_{r, R}^{B, A^*}(x) = L_1^{A^*}(x) \right] < 2/3$$

for some $x \in \{0, 1\}^*$ and some randomized function $B : \{0, 1\}^* \times \mathbb{R}_{>0} \rightarrow \{0, 1, \perp\}$ which is a valid AvgZPP oracle for $(L_2^{A^*}, D^{A^*})$, thereby ensuring that the reduction R°, A^*} fails to be of type

$$(L_2^{A^*}, D^{A^*}) \subseteq \text{AvgZPP}^{A^*} \Rightarrow L_1^{A^*} \subseteq \text{BPP}^{A^*}.$$

¹³Technically M_i^A equals the language L_2^A according to Definition 5, but the notation L_2^A is more convenient for the proofs.

We construct a sequence of relativization oracles by starting with \emptyset and adding strings and never taking them back out. We take A^* to be the limit of this sequence. Throughout the proofs, we simply refer to the “current” A with the understanding that this is the set of strings that have been included so far. We diagonalize against each triple i, S, R in sequence. After each round of diagonalization, we have the requirement that A^* matches the current A up through a certain input length, and we know that the current A contains no strings longer than that length. Now consider an arbitrary round, and suppose i, S, R is the triple to diagonalize against.

If there exists an A' consistent with the requirements of previous rounds and such that $M_i^{A'}$ fails to define a language in $\mathcal{C}^{A'}$, say with x as the violating input, then we update A to match A' up through the largest input length $M_i^{A'}(x)$ can query, and we require that A^* matches the new A up through this input length. This ensures that $M_i^{A^*}$ fails to define a language in \mathcal{C}^{A^*} , and we can move on to the next round.

Otherwise, we know that whatever we do to A , L_2^A will always be defined. Choose n large enough so that the following three things hold.

- The relativization oracle is fresh for all input lengths $\geq n$.
- The asymptotic constraints throughout the arguments are satisfied.
- The “relevant computations” all run in time $n^{\log n}$ without a big O.

The “relevant computations” include S on input n^d , R on inputs of length n , and (depending on the theorem) possibly the underlying computations of M_i on inputs of length n^d . We construct A at input length $2n$ to ensure that at the end of this round,

$$\Pr_{r_R, B} \left[R_{r_R}^{B, A}(x) = L_1^A(x) \right] < 2/3$$

for some $x \in \{0, 1\}^n$ and some randomized function $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ which is a valid AvgZPP oracle for (L_2^A, D^A) at input length n^d with respect to $\delta = 1/n^d$. Note that it makes sense to run $R^{B, A}(x)$ since this computation only queries B on inputs of length n^d and only with $\delta = 1/n^d$ (so we are justified in omitting the δ). This suffices to diagonalize against i, S, R because we can require that A^* matches the new A up through input length $n^{\log n}$ and up through the longest input length M_i can query on inputs of length n^d , thus ensuring the following three things.

- $L_1^{A^*}(x) = L_1^A(x)$.
- $R^{B, A^*}(x)$ behaves the same as $R^{B, A}(x)$.
- $L_2^{A^*}|_{n^d} = L_2^A|_{n^d}$ and $D_{n^d}^{A^*} = D_{n^d}^A$, which implies that B is a valid AvgZPP oracle for $(L_2^{A^*}, D^{A^*})$ at input length n^d with respect to $\delta = 1/n^d$ and can thus be extended to a full valid AvgZPP oracle for $(L_2^{A^*}, D^{A^*})$ without changing the behavior of $R^{B, A^*}(x)$.

5 Proof of Theorem 1

We use the setup from Section 4, customized as follows. We have $\mathcal{C} = \text{BPP}_{\parallel}^{\text{NP}}$, and M_i corresponds to a pair M, N where M is a $\text{BPP}_{\parallel}^{\circ}$ -type algorithm and N is an NP-type algorithm. Thus L_2^A is the $(\text{BPP}_{\parallel}^{\text{NP}})^A$ language computed by $M^{L_3^A, A}$ where L_3^A denotes the NP^A language computed by N^A .

Also, M on inputs of length n^d , as well as N on all inputs that could be queried by M on inputs of length n^d , count as “relevant computations” and thus all run in time $n^{\log n}$ without a big O .

Assume without loss of generality that for some positive integer e , M on inputs of length n^d always makes exactly n^e queries to its first oracle, and let $M_{r_M}^{L_3^A, A}(w)_j \in \{0, 1\}^*$ denote the j th of these queries when the input is w and the randomness is r_M .

5.1 Main Construction

Recall that M, N, S, R, n are fixed. For all relativization oracles A (not just the one we have constructed so far) we define the potential

$$\Phi^A = \mathbb{E}_{r_S, r_M, j} \left[1 - L_3^A \left(M_{r_M}^{L_3^A, A} (S_{r_S}^A(n^d))_j \right) \right]$$

where $j \in \{1, \dots, n^e\}$ is chosen uniformly at random. The construction has two stages.

Stage 1. This stage proceeds in iterations. For a given iteration, let A denote the current relativization oracle after the previous iteration. If there exist $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$ such that $x \notin L_1^A$ and $\Phi^{A \cup \{xy\}} \leq \Phi^A - 1/n^{3 \log n}$ then update $A := A \cup \{xy\}$ and continue with the next iteration. Otherwise, halt stage 1 and proceed to stage 2.

The following lemma is the technical heart of the proof of Theorem 1. We first finish the proof of Theorem 1 assuming the lemma, and then we prove the lemma in Section 5.2.

Lemma 2. *At the end of stage 1, there exists an $x \in \{0, 1\}^n$ such that $x \notin L_1^A$ and for all $y \in \{0, 1\}^n$,*

$$\Pr_{r_S} \left[L_2^{A \cup \{xy\}}(S_{r_S}^A(n^d)) \neq L_2^A(S_{r_S}^A(n^d)) \right] \leq 1/8n^{d+\log n} \quad (3)$$

and

$$\Pr_{r_S} \left[S_{r_S}^{A \cup \{xy\}}(n^d) \neq S_{r_S}^A(n^d) \right] \leq 1/2n^d. \quad (4)$$

Stage 2. Let A denote the current relativization oracle at the end of stage 1, and let x be as guaranteed by Lemma 2. Let $Y \subseteq \{0, 1\}^n$ be an arbitrary set of size $4n^{\log n}$. Define a deterministic reduction oracle $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ by

$$B(w) = \begin{cases} L_2^A(w) & \text{if } L_2^{A \cup \{xy\}}(w) = L_2^A(w) \text{ for all } y \in Y \\ \perp & \text{otherwise} \end{cases}.$$

There are two cases.

Case 1. If

$$\Pr_{r_R} \left[R_{r_R}^{B, A}(x) = 1 \right] > 1/3$$

then we will use A for the relativization oracle at the beginning of the next round of diagonalization, without changing it. Since $x \notin L_1^A$, we have

$$\Pr_{r_R} \left[R_{r_R}^{B, A}(x) = L_1^A(x) \right] < 2/3.$$

We just need to verify that B is a valid AvgZPP oracle for (L_2^A, D^A) at input length n^d with respect to $\delta = 1/n^d$. Obviously, $B(w)$ always returns $L_2^A(w)$ or \perp , by our definition of B . We have

$$\begin{aligned}
\Pr_{w \sim D_{n^d}^A} [B(w) = \perp] &= \Pr_{r_S} [B(S_{r_S}^A(n^d)) = \perp] \\
&= \Pr_{r_S} [\exists y \in Y \text{ such that } L_2^{A \cup \{xy\}}(S_{r_S}^A(n^d)) \neq L_2^A(S_{r_S}^A(n^d))] \\
&\leq \sum_{y \in Y} \Pr_{r_S} [L_2^{A \cup \{xy\}}(S_{r_S}^A(n^d)) \neq L_2^A(S_{r_S}^A(n^d))] \\
&\leq \sum_{y \in Y} 1/8n^{d+\log n} \\
&= |Y| \cdot 1/8n^{d+\log n} \\
&= 1/2n^d \\
&\leq 1/n^d = \delta
\end{aligned}$$

where the fourth line follows by Lemma 2. Thus we have succeeded in diagonalizing against M, N, S, R as described at the end of Section 4.

Case 2. If

$$\Pr_{r_R} [R_{r_R}^{B,A}(x) = 1] \leq 1/3$$

then for each $y \in Y$ we define

$$\pi_y = \Pr_{r_R} [R_{r_R}^{B,A}(x) \text{ queries } A(xy)].$$

Since $R^{B,A}(x)$ runs in time $n^{\log n}$, we have $\sum_{y \in Y} \pi_y \leq n^{\log n}$. Thus there exists a $y \in Y$ such that $\pi_y \leq n^{\log n}/|Y| = 1/4$. Fix this y . We will update the relativization oracle to be $A \cup \{xy\}$ for the end of this round of diagonalization. Since $x \in L_1^{A \cup \{xy\}}$, we have

$$\begin{aligned}
\Pr_{r_R} [R_{r_R}^{B, A \cup \{xy\}}(x) = L_1^{A \cup \{xy\}}(x)] &\leq \Pr_{r_R} [R_{r_R}^{B,A}(x) = 1 \text{ or } R_{r_R}^{B, A \cup \{xy\}}(x) \neq R_{r_R}^{B,A}(x)] \\
&\leq \Pr_{r_R} [R_{r_R}^{B,A}(x) = 1 \text{ or } R_{r_R}^{B,A}(x) \text{ queries } A(xy)] \\
&\leq \Pr_{r_R} [R_{r_R}^{B,A}(x) = 1] + \pi_y \\
&\leq 1/3 + 1/4 \\
&< 2/3.
\end{aligned}$$

We just need to verify that B is a valid AvgZPP oracle for $(L_2^{A \cup \{xy\}}, D^{A \cup \{xy\}})$ at input length n^d with respect to $\delta = 1/n^d$. Since $y \in Y$, we have that for all w , if $B(w) \neq \perp$ then $B(w) = L_2^A(w) = L_2^{A \cup \{xy\}}(w)$, by our definition of B . We also have

$$\begin{aligned}
\Pr_{w \sim D_{n^d}^{A \cup \{xy\}}} [B(w) = \perp] &= \Pr_{r_S} [B(S_{r_S}^{A \cup \{xy\}}(n^d)) = \perp] \\
&\leq \Pr_{r_S} [B(S_{r_S}^A(n^d)) = \perp \text{ or } S_{r_S}^{A \cup \{xy\}}(n^d) \neq S_{r_S}^A(n^d)]
\end{aligned}$$

$$\begin{aligned}
&\leq \Pr_{r_S} \left[B(S_{r_S}^A(n^d)) = \perp \right] + \Pr_{r_S} \left[S_{r_S}^{A \cup \{xy\}}(n^d) \neq S_{r_S}^A(n^d) \right] \\
&\leq 1/2n^d + 1/2n^d \\
&= 1/n^d = \delta
\end{aligned}$$

where the fourth line follows by the calculation from case 1 and by Lemma 2. Thus we have succeeded in diagonalizing against M, N, S, R as described at the end of Section 4.

5.2 Proof of Lemma 2

For all A (not just the one we have constructed so far) and all r_S, r_M, j , let us define

$$\Phi_{r_S, r_M, j}^A = 1 - L_3^A \left(M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \right)$$

and

$$\Phi_{r_S, r_M}^A = \mathbb{E}_j \left[\Phi_{r_S, r_M, j}^A \right]$$

so that $\Phi^A = \mathbb{E}_{r_S, r_M} \left[\Phi_{r_S, r_M}^A \right]$. We always have $0 \leq \Phi^A \leq 1$.

From here on out, A denotes the current relativization oracle at the end of stage 1. Since there are at most $n^{3 \log n}$ iterations before stage 1 terminates, we have

$$\Pr_{x \in \{0,1\}^n} [x \in L_1^A] \leq n^{3 \log n} / 2^n$$

where x is chosen uniformly at random. For $x \in \{0,1\}^n$ define

$$p_x = \Pr_{r_S, r_M} \left[\exists y \in \{0,1\}^n \text{ such that } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d)) \text{ queries } A(xy) \right].$$

Since $M_{r_M}^{L_3^A, A}(w)$ runs in time $n^{\log n}$ for all $w \in \{0,1\}^{n^d}$, we have $\sum_x p_x \leq n^{\log n}$ and thus

$$\Pr_{x \in \{0,1\}^n} [p_x > 1/n^{3 \log n}] < n^{4 \log n} / 2^n.$$

For every $v \in L_3^A$ pick an arbitrary accepting computation path of $N^A(v)$ to be the “designated” path. For $x \in \{0,1\}^n$ define

$$\begin{aligned}
q_x = &\Pr_{r_S, r_M, j} \left[M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \in L_3^A \text{ and } \exists y \in \{0,1\}^n \text{ such that} \right. \\
&\left. N^A \left(M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \right) \text{ queries } A(xy) \text{ on the designated path} \right]
\end{aligned}$$

where j is chosen uniformly at random. Since $N^A(v)$ runs in time $n^{\log n}$ for every v of interest, we have $\sum_x q_x \leq n^{\log n}$ and thus

$$\Pr_{x \in \{0,1\}^n} [q_x > 1/n^{4 \log n}] < n^{5 \log n} / 2^n.$$

For $x \in \{0,1\}^n$ define

$$s_x = \Pr_{r_S} \left[\exists y \in \{0,1\}^n \text{ such that } S_{r_S}^A(n^d) \text{ queries } A(xy) \right].$$

Since $S^A(n^d)$ runs in time $n^{\log n}$, we have $\sum_x s_x \leq n^{\log n}$ and thus

$$\Pr_{x \in \{0,1\}^n} [s_x > 1/n^{3 \log n}] < n^{4 \log n} / 2^n.$$

By a union bound we find that

$$\begin{aligned} & \Pr_{x \in \{0,1\}^n} [x \notin L_1^A \text{ and } p_x \leq 1/n^{3 \log n} \text{ and } q_x \leq 1/n^{4 \log n} \text{ and } s_x \leq 1/n^{3 \log n}] \\ & > 1 - (n^{3 \log n} / 2^n) - (n^{4 \log n} / 2^n) - (n^{5 \log n} / 2^n) - (n^{4 \log n} / 2^n) \\ & > 0. \end{aligned}$$

Thus there exists an $x \in \{0,1\}^n$ such that $x \notin L_1^A$ and $p_x \leq 1/n^{3 \log n}$ and $q_x \leq 1/n^{4 \log n}$ and $s_x \leq 1/n^{3 \log n}$. Fix this x . We claim that this x satisfies the conditions of Lemma 2. Suppose for contradiction that there exists a $y \in \{0,1\}^n$ such that either Inequality (3) does not hold or Inequality (4) does not hold. Fix this y . We claim that $\Phi^{A \cup \{xy\}} \leq \Phi^A - 1/n^{3 \log n}$, thus contradicting the fact that stage 1 halted. Henceforth we let A' denote $A \cup \{xy\}$. We partition the joint sample space of S 's internal randomness and M 's internal randomness into five events.

$$\begin{aligned} E_1 &= \{(r_S, r_M) : S_{r_S}^{A'}(n^d) \neq S_{r_S}^A(n^d)\} \\ E_2 &= \{(r_S, r_M) : (r_S, r_M) \notin E_1 \text{ and } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d)) \text{ queries } A(xy)\} \\ E_3 &= \{(r_S, r_M) : (r_S, r_M) \notin E_1 \cup E_2 \text{ and } \exists j \text{ such that } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \in L_3^A \setminus L_3^{A'}\} \\ E_4 &= \{(r_S, r_M) : (r_S, r_M) \notin E_1 \cup E_2 \cup E_3 \text{ and } \exists j \text{ such that } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \in L_3^{A'} \setminus L_3^A\} \\ E_5 &= \{(r_S, r_M) : (r_S, r_M) \notin E_1 \cup E_2 \cup E_3 \cup E_4\} \end{aligned}$$

Claim 1. $\Pr_{r_S, r_M} [(r_S, r_M) \in E_1] \leq 1/n^{3 \log n}$ and for all $(r_S, r_M) \in E_1$, $\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \leq 1$.

Claim 2. $\Pr_{r_S, r_M} [(r_S, r_M) \in E_2] \leq 1/n^{3 \log n}$ and for all $(r_S, r_M) \in E_2$, $\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \leq 1$.

Claim 3. $\Pr_{r_S, r_M} [(r_S, r_M) \in E_3] \leq 1/n^{3 \log n}$ and for all $(r_S, r_M) \in E_3$, $\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \leq 1$.

Claim 4. $\Pr_{r_S, r_M} [(r_S, r_M) \in E_4] \geq 1/n^{2 \log n}$ and for all $(r_S, r_M) \in E_4$, $\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \leq -1/n^e$.

Claim 5. $\Pr_{r_S, r_M} [(r_S, r_M) \in E_5] \leq 1$ and for all $(r_S, r_M) \in E_5$, $\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \leq 0$.

From these five claims it follows that

$$\begin{aligned} \Phi^{A'} - \Phi^A &= \mathbb{E}_{r_S, r_M} [\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A] \\ &= \sum_{k=1}^5 \mathbb{E}_{r_S, r_M} [\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A \mid (r_S, r_M) \in E_k] \cdot \Pr_{r_S, r_M} [(r_S, r_M) \in E_k] \\ &\leq 1/n^{3 \log n} + 1/n^{3 \log n} + 1/n^{3 \log n} - 1/n^{e+2 \log n} \\ &\leq -1/n^{3 \log n} \end{aligned}$$

which is what we wanted to show.

Proof of Claim 1. The first assertion follows because

$$\Pr_{r_S, r_M} [(r_S, r_M) \in E_1] \leq \Pr_{r_S} [S_{r_S}^A(n^d) \text{ queries } A(xy)] \leq s_x \leq 1/n^{3 \log n}.$$

The second assertion follows trivially from the fact that $\Phi_{r_S, r_M}^{A'} \leq 1$ and $\Phi_{r_S, r_M}^A \geq 0$. \square

Proof of Claim 2. The first assertion follows because

$$\Pr_{r_S, r_M} [(r_S, r_M) \in E_2] \leq \Pr_{r_S, r_M} [M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d)) \text{ queries } A(xy)] \leq p_x \leq 1/n^{3 \log n}.$$

The second assertion follows trivially from the fact that $\Phi_{r_S, r_M}^{A'} \leq 1$ and $\Phi_{r_S, r_M}^A \geq 0$. \square

Proof of Claim 3. The first assertion follows because

$$\begin{aligned} \Pr_{r_S, r_M} [(r_S, r_M) \in E_3] &\leq \Pr_{r_S, r_M} [\exists j \text{ such that } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \in L_3^A \setminus L_3^{A'}] \\ &\leq \Pr_{r_S, r_M} \left[\exists j \text{ such that } M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j \in L_3^A \text{ and} \right. \\ &\quad \left. N^A(M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))_j) \text{ queries } A(xy) \text{ on the designated path} \right] \\ &\leq n^e \cdot q_x \\ &\leq 1/n^{3 \log n} \end{aligned}$$

where the second-to-last line follows using a union bound and the last line follows because $q_x \leq 1/n^{4 \log n}$. The second assertion follows trivially from the fact that $\Phi_{r_S, r_M}^{A'} \leq 1$ and $\Phi_{r_S, r_M}^A \geq 0$. \square

Proof of Claim 4. This claim is in some sense the crux of the whole proof. Since $1/n^{3 \log n} \leq 1/2n^d$, Claim 1 implies that Inequality (4) holds and therefore Inequality (3) does not hold. We claim that if $(r_S, r_M) \notin E_1 \cup E_2 \cup E_3 \cup E_4$ then

$$M_{r_M}^{L_3^{A'}, A'}(S_{r_S}^A(n^d)) = M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d)).$$

This is because every query $M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))$ makes to its second oracle has the same answer under A' and A , and every query it makes to its first oracle has the same answer under $L_3^{A'}$ and L_3^A . Thus the computations $M_{r_M}^{L_3^{A'}, A'}(S_{r_S}^A(n^d))$ and $M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))$ proceed identically, making the same queries and receiving the same answers, and hence they produce the same output. The first assertion now follows because

$$\begin{aligned} &\Pr_{r_S, r_M} [(r_S, r_M) \in E_4] \\ &\geq \Pr_{r_S, r_M} [(r_S, r_M) \notin E_1 \cup E_2 \cup E_3 \text{ and } M_{r_M}^{L_3^{A'}, A'}(S_{r_S}^A(n^d)) \neq M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))] \\ &\geq \Pr_{r_S, r_M} [M_{r_M}^{L_3^{A'}, A'}(S_{r_S}^A(n^d)) \neq M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))] - \sum_{k=1}^3 \Pr_{r_S, r_M} [(r_S, r_M) \in E_k] \\ &\geq \Pr_{r_S} [L_2^{A'}(S_{r_S}^A(n^d)) \neq L_2^A(S_{r_S}^A(n^d))] / 3 - \sum_{k=1}^3 \Pr_{r_S, r_M} [(r_S, r_M) \in E_k] \end{aligned}$$

$$\begin{aligned}
&> 1/24n^{d+\log n} - 1/n^{3\log n} - 1/n^{3\log n} - 1/n^{3\log n} \\
&\geq 1/n^{2\log n}
\end{aligned}$$

where the third line follows by a union bound and the second-to-last line follows by the negation of Inequality (3) and by Claim 1, Claim 2, and Claim 3.

We now argue the second assertion. Since $(r_S, r_M) \notin E_1$, we have $S_{r_S}^{A'}(n^d) = S_{r_S}^A(n^d)$. Let w denote this string. Since $(r_S, r_M) \notin E_1 \cup E_2$, we have

$$M_{r_M}^{L_3^{A'}, A'}(w)_j = M_{r_M}^{L_3^A, A}(w)_j$$

for all j . Let v_j denote these strings, and note that $\Phi_{r_S, r_M, j}^{A'} = 1 - L_3^{A'}(v_j)$ and $\Phi_{r_S, r_M, j}^A = 1 - L_3^A(v_j)$. Since $(r_S, r_M) \notin E_3$, we have $\Phi_{r_S, r_M, j}^{A'} \leq \Phi_{r_S, r_M, j}^A$ for all j . By the definition of E_4 , we have $\Phi_{r_S, r_M, j}^{A'} = 0$ and $\Phi_{r_S, r_M, j}^A = 1$ for some j . Therefore,

$$\Phi_{r_S, r_M}^{A'} - \Phi_{r_S, r_M}^A = \frac{1}{n^e} \sum_j (\Phi_{r_S, r_M, j}^{A'} - \Phi_{r_S, r_M, j}^A) \leq -1/n^e. \quad \square$$

Proof of Claim 5. The first assertion is trivial. We now argue the second assertion. In the proof of Claim 4 we argued that if $(r_S, r_M) \notin E_1 \cup E_2 \cup E_3 \cup E_4$ then the computations $M_{r_M}^{L_3^{A'}, A'}(S_{r_S}^{A'}(n^d))$ and $M_{r_M}^{L_3^A, A}(S_{r_S}^A(n^d))$ proceed identically, making the same queries and receiving the same answers. In particular, $\Phi_{r_S, r_M, j}^{A'} = \Phi_{r_S, r_M, j}^A$ for all j , which implies that $\Phi_{r_S, r_M}^{A'} = \Phi_{r_S, r_M}^A$. \square

6 Proof of Theorem 2

Fix a polynomial q . We use the setup from Section 4, customized as follows. We have $\mathcal{C} = \text{PH}$, and M_i corresponds to a PH-type algorithm M . We redefine

$$L_1^A = \{x : \exists y \text{ such that } |y| = |x| + 2q(|x|) \text{ and } xy \in A\}$$

using $|y| = |x| + 2q(|x|)$ instead of $|y| = |x|$, and thus we need to construct A at input length $2n + 2q(n)$ rather than $2n$. We only diagonalize against reductions R that use at most q queries to the reduction oracle. Also, M on inputs of length n^d counts as “relevant computations” and thus runs in time $n^{\log n}$ without a big O . For the reason discussed at the end of Section 2.3, we have the stronger requirement that at the end of this round,

$$\Pr_{r_R, B} \left[R_{r_R}^{B, A}(x) = L_1^A(x) \right] < 1/2 + 1/n^{\log n}$$

with $1/2 + 1/n^{\log n}$ instead of $2/3$. Finally, note that it can never be the case that M^A fails to define a language in PH^A , since PH is a syntactically defined class.

We generalize the notion of a reduction oracle: If $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}^{\mathbb{N}}$ is a deterministic function then running $R_{r_R}^{B, A}(x)$ means that for each w , the i th time the computation queries $B(w)$ it gets $B(w)(i)$ as a response. Thus a randomized function $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ is a distribution over such deterministic functions, where each $B(w)(i)$ is independent and the distribution of $B(w)(i)$ depends only on w and not on i .

6.1 Main Construction

Recall that M, S, R, n are fixed. Let A denote the current relativization oracle at the beginning of this round. For $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^{n+2q(n)}$ define

$$p_{x,y} = \Pr_{r_R} \left[\exists B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}^{\mathbb{N}} \text{ and } \exists x' \in \{0, 1\}^n \text{ such that } R_{r_R}^{B,A}(x) \text{ queries } A(x'y) \right]$$

and

$$p_y = \mathbb{E}_{x \in \{0,1\}^n} [p_{x,y}]$$

where x is chosen uniformly at random. For each $x \in \{0, 1\}^n$ and r_R , the computation $R_{r_R}^{B,A}(x)$ has at most $3^{q(n)}$ computation paths over the possible responses it could get from B (recall that A is fixed). On each of these computation paths, $R_{r_R}^{B,A}(x)$ can query at most $n^{\log n}$ bits of A since it runs in time $n^{\log n}$. Thus there are at most $n^{\log n} 3^{q(n)}$ pairs $(x', y) \in \{0, 1\}^n \times \{0, 1\}^{n+2q(n)}$ for which there exists a $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}^{\mathbb{N}}$ such that $R_{r_R}^{B,A}(x)$ queries $A(x'y)$. It follows that $\sum_y p_y \leq n^{\log n} 3^{q(n)}$ and thus

$$\Pr_{y \in \{0,1\}^{n+2q(n)}} [p_y > 1/2n^{\log n}] < 2n^{2\log n} 3^{q(n)} / 2^{n+2q(n)}$$

where y is chosen uniformly at random. For $y \in \{0, 1\}^{n+2q(n)}$ define

$$s_y = \Pr_{r_S} \left[\exists x' \in \{0, 1\}^n \text{ such that } S_{r_S}^A(n^d) \text{ queries } A(x'y) \right].$$

Since $S^A(n^d)$ runs in time $n^{\log n}$, we have $\sum_y s_y \leq n^{\log n}$ and thus

$$\Pr_{y \in \{0,1\}^{n+2q(n)}} [s_y > 1/2n^d] < 2n^{d+\log n} / 2^{n+2q(n)}.$$

By a union bound we find that

$$\begin{aligned} & \Pr_{y \in \{0,1\}^{n+2q(n)}} [p_y \leq 1/2n^{\log n} \text{ and } s_y \leq 1/2n^d] \\ & > 1 - (2n^{2\log n} 3^{q(n)} / 2^{n+2q(n)}) - (2n^{d+\log n} / 2^{n+2q(n)}) \\ & > 0. \end{aligned}$$

Thus there exists a $y \in \{0, 1\}^{n+2q(n)}$ such that $p_y \leq 1/2n^{\log n}$ and $s_y \leq 1/2n^d$. Fix this y . Now

$$\Pr_{x \in \{0,1\}^n} [p_{x,y} \geq 1/n^{\log n}] \leq 1/2$$

and thus there exists a set $X \subseteq \{0, 1\}^n$ of size $|X| = 2^{n-1}$ such that for all $x \in X$, $p_{x,y} < 1/n^{\log n}$. To prove the theorem, it suffices to show that there exists a $Z \subseteq \{xy : x \in X\}$, an $x \in X$, and a randomized function $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ which is a valid AvgZPP oracle for $(L_2^{\text{AUZ}}, D^{\text{AUZ}})$ at input length n^d with respect to $\delta = 1/n^d$, such that

$$\Pr_{r_R, B} [R_{r_R}^{B, \text{AUZ}}(x) = L_1^{\text{AUZ}}(x)] < 1/2 + 1/n^{\log n}$$

because we can then update the relativization oracle to be $A \cup Z$ for the end of this round.

Suppose for contradiction that this does not hold. We can assume that r_S is sampled uniformly at random from $\{0, 1\}^{n^{\log n}}$ when S is run on input n^d . Define an error-correcting code

$$C : \{0, 1\}^{2^{n-1}} \rightarrow \{0, 1\}^{2^{n^{\log n}}}$$

as follows, where the information word is viewed as a subset $Z \subseteq \{xy : x \in X\}$ and the code word is viewed as a function $C(Z) : \{0, 1\}^{n^{\log n}} \rightarrow \{0, 1\}$.

$$C(Z)(r_S) = L_2^{A \cup Z}(S_{r_S}^A(n^d))$$

Claim 6. *The relative minimum distance of C is $> 1/2n^d$.*

We prove Claim 6 shortly. Let k denote the number of quantifiers M uses, and recall that M runs in time $n^{\log n}$ on inputs of length n^d . Since each bit of $C(Z)$ corresponds to running $M^{A \cup Z}$ on a fixed input of length n^d , each bit of $C(Z)$ is computable by a circuit of depth k and size $2^{n^{\log n}}$ where each input to the circuit is the output of a deterministic computation running in time $n^{\log n}$ with oracle access to $A \cup Z$. Since A is fixed, each of the inputs to this circuit is computable by a DNF with top fan-in $2^{n^{\log n}}$ and bottom fan-in $n^{\log n}$ whose inputs correspond to strings in $\{xy : x \in X\}$, that is, coordinates of the information word.

The bottom line is that there exists a binary error-correcting code with information word length 2^{n-1} and relative minimum distance $> 1/2n^d$ such that each bit of the code word is computable by a circuit of depth $k + 2$ and size $2^{2n^{\log n}} n^{\log n}$. This contradicts the following result.

Theorem 4 (Viola [Vio05a]). *If there exists a binary error-correcting code with information word length ν and relative minimum distance γ such that each bit of the code word is computable by a circuit of depth κ and size σ , then $\nu\gamma \leq O(\log^{\kappa-1} \sigma)$.*

Theorem 4 holds regardless of the rate of the code.

Proof of Claim 6. In Figure 1 we exhibit a decoder that can handle up to a $1/2n^d$ fraction of erasures. For an arbitrary $Z \subseteq \{xy : x \in X\}$, assume that C' agrees with $C(Z)$ on at least a $1 - 1/2n^d$ fraction of r_S 's and outputs \perp on the rest. Then we just need to show that $Z' = Z$. We do this by showing that for an arbitrary $x \in X$,

$$\Pr_{r_S, B} \left[R_{r_S}^{B, A}(x) = L_1^{A \cup Z}(x) \right] > 1/2$$

which implies that $xy \in Z'$ if and only if $x \in L_1^{A \cup Z}$ if and only if $xy \in Z$.

We start by showing that B is a valid AvgZPP oracle for $(L_2^{A \cup Z}, D^{A \cup Z})$ at input length n^d with respect to $\delta = 1/n^d$. We have that $B(w)$ always equals $L_2^{A \cup Z}(w)$ or \perp , since if r_S is such that $S_{r_S}^A(n^d) = w$ and $C'(r_S) \neq \perp$ then

$$C'(r_S) = C(Z)(r_S) = L_2^{A \cup Z}(S_{r_S}^A(n^d)) = L_2^{A \cup Z}(w).$$

We have

$$\Pr_{r_S, B} \left[B(S_{r_S}^A(n^d)) = \perp \right] = \sum_{w \in \{0, 1\}^{n^d}} \Pr_{r_S, B} \left[B(S_{r_S}^A(n^d)) = \perp \mid S_{r_S}^A(n^d) = w \right] \cdot \Pr_{r_S, B} \left[S_{r_S}^A(n^d) = w \right]$$

• **Input:** $C' : \{0, 1\}^{n^{\log n}} \rightarrow \{0, 1, \perp\}$

• **Output:** $Z' \subseteq \{xy : x \in X\}$ given by

$$Z' = \left\{ xy : \Pr_{r_{R,B}} \left[R_{r_R}^{B,A}(x) = 1 \right] > 1/2 \right\}$$

where the randomized function $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ is defined by

$$\Pr_B [B(w) = b] = \Pr_{r_S} \left[C'(r_S) = b \mid S_{r_S}^A(n^d) = w \right]$$

if

$$\Pr_{r_S} \left[S_{r_S}^A(n^d) = w \right] > 0$$

and otherwise

$$\Pr_B [B(w) = \perp] = 1$$

Figure 1: Decoder for Claim 6

$$\begin{aligned} &= \sum_{w \in \{0,1\}^{n^d}} \Pr_B [B(w) = \perp] \cdot \Pr_{r_S} \left[S_{r_S}^A(n^d) = w \right] \\ &= \sum_{w \in \{0,1\}^{n^d}} \Pr_{r_S} \left[C'(r_S) = \perp \mid S_{r_S}^A(n^d) = w \right] \cdot \Pr_{r_S} \left[S_{r_S}^A(n^d) = w \right] \\ &= \Pr_{r_S} [C'(r_S) = \perp] \\ &\leq 1/2n^d \end{aligned}$$

and

$$\Pr_{r_S} \left[S_{r_S}^{A \cup Z}(n^d) \neq S_{r_S}^A(n^d) \right] \leq \Pr_{r_S} \left[\exists z \in Z \text{ such that } S_{r_S}^A(n^d) \text{ queries } A(z) \right] \leq s_y \leq 1/2n^d$$

and thus

$$\begin{aligned} \Pr_{w \sim D^{A \cup Z, B}} [B(w) = \perp] &= \Pr_{r_{S,B}} \left[B(S_{r_S}^{A \cup Z}(n^d)) = \perp \right] \\ &\leq \Pr_{r_{S,B}} \left[B(S_{r_S}^A(n^d)) = \perp \text{ or } S_{r_S}^{A \cup Z}(n^d) \neq S_{r_S}^A(n^d) \right] \\ &\leq \Pr_{r_{S,B}} \left[B(S_{r_S}^A(n^d)) = \perp \right] + \Pr_{r_S} \left[S_{r_S}^{A \cup Z}(n^d) \neq S_{r_S}^A(n^d) \right] \\ &\leq 1/2n^d + 1/2n^d \\ &= 1/n^d = \delta. \end{aligned}$$

Now we have

$$\Pr_{r_{R,B}} \left[R_{r_R}^{B, A \cup Z}(x) \neq R_{r_R}^{B,A}(x) \right] \leq \mathbb{E}_B \left[\Pr_{r_R} \left[\exists z \in Z \text{ such that } R_{r_R}^{B,A}(x) \text{ queries } A(z) \right] \right]$$

$$\begin{aligned}
&\leq \mathbb{E}_B [p_{x,y}] \\
&= p_{x,y} \\
&< 1/n^{\log n}
\end{aligned}$$

and thus

$$\begin{aligned}
\Pr_{r_{R,B}} \left[R_{r_R}^{B,A}(x) = L_1^{A \cup Z}(x) \right] &\geq \Pr_{r_{R,B}} \left[R_{r_R}^{B,A \cup Z}(x) = L_1^{A \cup Z}(x) \text{ and } R_{r_R}^{B,A \cup Z}(x) = R_{r_R}^{B,A}(x) \right] \\
&\geq \Pr_{r_{R,B}} \left[R_{r_R}^{B,A \cup Z}(x) = L_1^{A \cup Z}(x) \right] - \Pr_{r_{R,B}} \left[R_{r_R}^{B,A \cup Z}(x) \neq R_{r_R}^{B,A}(x) \right] \\
&> (1/2 + 1/n^{\log n}) - 1/n^{\log n} \\
&= 1/2
\end{aligned}$$

where the third line follows by our contradiction assumption. \square

7 Proof of Theorem 3

We use the setup from Section 4, customized as follows. We only diagonalize against reductions R that use at most 2 queries to the reduction oracle. For the reason discussed at the end of Section 2.3, we have the stronger requirement that at the end of this round,

$$\Pr_{r_{R,B}} \left[R_{r_R}^{B,A}(x) = L_1^A(x) \right] < 1/2 + 1/n^{\log n}$$

with $1/2 + 1/n^{\log n}$ instead of $2/3$. The proof is so similar to the proof of Theorem 2 that we just sketch how it plays out. We can work with $|y| = n$ (rather than $|y| = n + 2q(n)$ as in the proof of Theorem 2).

7.1 Main Construction

Recall that M_i, S, R, n are fixed. Let A denote the current relativization oracle at the beginning of this round. There exists a $y \in \{0, 1\}^n$ such that $p_y \leq 1/4n^{\log n}$ and $s_y \leq 1/2n^d$, and there exists a set $X \subseteq \{0, 1\}^n$ of size $|X| = 2^{n-1}$ such that for all $x \in X$, $p_{x,y} \leq 1/2n^{\log n}$. Then there exists a $Z \subseteq \{xy : x \in X\}$, an $x \in X$, and a randomized function $B : \{0, 1\}^{n^d} \rightarrow \{0, 1, \perp\}$ which is a valid AvgZPP oracle for $(L_2^{A \cup Z}, D^{A \cup Z})$ at input length n^d with respect to $\delta = 1/n^d$, such that

$$\Pr_{r_{R,B}} \left[R_{r_R}^{B,A \cup Z}(x) = L_1^{A \cup Z}(x) \right] < 1/2 + 1/n^{\log n}$$

since otherwise we can extract an error-correcting code

$$C : \{0, 1\}^{2^{n-1}} \rightarrow \{0, 1\}^{2^{n \log n}}$$

with the following properties. There is a randomized decoder that can handle up to a $1/2n^d$ fraction of erasures, and it recovers any bit of the information word with probability at least

$$(1/2 + 1/n^{\log n}) - 1/2n^{\log n} = 1/2 + 1/2n^{\log n}.$$

To recover any bit, the decoder runs $R^{B,A}(x)$ for some $x \in \{0,1\}^n$ and some randomized function B . Since R makes at most 2 queries to B , and since each query to B can be answered with at most 1 query to the corrupted code word C' , the decoder makes at most 2 queries to C' .

The bottom line is that there exists a binary error-correcting code with information word length 2^{n-1} and code word length $2^{n^{\log n}}$ and a decoder that uses 2 queries to recover any bit of the information word with probability at least $1/2 + 1/2n^{\log n}$ when at most a $1/2n^d$ fraction of the code word bits are erased. This contradicts the following result.

Theorem 5 (Kerenidis and de Wolf [KdW04]). *If there exists a binary error-correcting code with information word length ν and code word length μ and a decoder that uses 2 queries to recover any bit of the information word with probability at least $1/2 + \epsilon$ when at most a γ fraction of the code word bits are erased, then $\mu \geq 2^{\Omega(\gamma\epsilon^3\nu)}$.*

Remarkably, the proof of Theorem 5 is based on quantum information theory. Kerenidis and de Wolf proved the stronger bound $\mu \geq 2^{\Omega(\gamma\epsilon^2\nu)}$ assuming that the decoder is guaranteed to work even if a γ fraction of the code word bits are flipped rather than just erased. The extra ϵ in the exponent in Theorem 5 grossly accounts for the generalization from flips to erasures. It may be possible to prove the stronger bound for erasure decoders, but Theorem 5 as stated is already good enough for our purpose.

The complexity of M_i is immaterial because Theorem 5 holds without any constraints on the efficiency of the encoder.

8 Open Problems

Impagliazzo [Imp11] gave an oracle relative to which $(\text{NP}, \text{PSAMP}) \subseteq \text{AvgZPP}$ but $\text{NP} \not\subseteq \text{BPP}$, but it is open to give an oracle relative to which $(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP}$ but $(\text{NP}, \text{PSAMP}) \not\subseteq \text{AvgZPP}$.

Impagliazzo and Levin [IL90] proved that relative to every oracle, there exists a nonadaptive reduction of type

$$(\text{NP}, \mathcal{U}) \subseteq \text{HeurBPP} \Rightarrow (\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP}.$$

This reduction uses polynomially many queries. It is open to construct such a reduction using a smaller number of queries, ideally a mapping reduction. It would be interesting to prove that there exists an oracle relative to which no such mapping reduction exists.

Bogdanov and Trevisan [BT06b] proved that relative to every oracle, if there exists a nonadaptive reduction of type

$$(\text{NP}, \text{PSAMP}) \subseteq \text{HeurBPP} \Rightarrow \text{NP} \subseteq \text{BPP}$$

then the polynomial-time hierarchy collapses to the third level. It is open to extend this result to adaptive reductions. It would be interesting to prove that there exists an oracle relative to which such an adaptive reduction exists and yet the polynomial-time hierarchy is infinite. Can the ‘‘Book trick’’ [For99] be used? Less generally, it would be interesting to prove that there exists an oracle relative to which an adaptive reduction of the above type exists but no nonadaptive reduction of the above type exists.

Acknowledgments

First and foremost, I thank Luca Trevisan for suggesting the research topic. I thank Luca Trevisan, Dieter van Melkebeek, Ronen Shaltiel, and anonymous reviewers for helpful comments.

References

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
- [BBBV97] Charles Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997.
- [BFNW93] László Babai, Lance Fortnow, Noam Nisan, and Avi Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3:307–318, 1993.
- [BT06a] Andrej Bogdanov and Luca Trevisan. Average-case complexity. *Foundations and Trends in Theoretical Computer Science*, 2(1), 2006.
- [BT06b] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. *SIAM Journal on Computing*, 36(4):1119–1159, 2006.
- [FF93] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 22(5):994–1005, 1993.
- [For99] Lance Fortnow. Relativized worlds with an infinite hierarchy. *Information Processing Letters*, 69(6):309–313, 1999.
- [FSS84] Merrick Furst, James Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
- [Gol08] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [GSTS07] 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.
- [GTS07] Dan Gutfreund and Amnon Ta-Shma. Worst-case to average-case reductions revisited. In *Proceedings of the 11th International Workshop on Randomization and Computation*, pages 569–583, 2007.
- [HHT97] Yenjo Han, Lane Hemaspaandra, and Thomas Thierauf. Threshold computation and cryptographic security. *SIAM Journal on Computing*, 26(1):59–78, 1997.
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography. In *Proceedings of the 30th IEEE Symposium on Foundations of Computer Science*, pages 230–235, 1989.

- [IL90] Russell Impagliazzo and Leonid Levin. No better ways to generate hard NP instances than picking uniformly at random. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science*, pages 812–821, 1990.
- [Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th IEEE Conference on Structure in Complexity Theory*, pages 134–147, 1995.
- [Imp11] Russell Impagliazzo. Relativized separations of worst-case and average-case complexities for NP. In *Proceedings of the 26th IEEE Conference on Computational Complexity*, pages 104–114, 2011.
- [KdW04] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004.
- [Lev86] Leonid Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.
- [Reg06] Oded Regev. Lattice-based cryptography. In *Proceedings of the 26th International Cryptology Conference*, pages 131–141, 2006.
- [STV01] Madhu Sudan, Luca Trevisan, and Salil Vadhan. Pseudorandom generators without the XOR Lemma. *Journal of Computer and System Sciences*, 62(2):236–266, 2001.
- [Vio05a] Emanuele Viola. The complexity of constructing pseudorandom generators from hard functions. *Computational Complexity*, 13(3-4):147–188, 2005.
- [Vio05b] Emanuele Viola. On constructing parallel pseudorandom generators from one-way functions. In *Proceedings of the 20th IEEE Conference on Computational Complexity*, pages 183–197, 2005.
- [Wat10] Thomas Watson. Relativized worlds without worst-case to average-case reductions for NP. Technical Report TR10-042, Electronic Colloquium on Computational Complexity, 2010.