The Isomorphism Conjecture Fails Relative to a Random Oracle

Stuart A. Kurtz[*]
University of Chicago

Stephen R. Mahaney[†]
University of Arizona

James S. Royer[‡]
Syracuse University

December 16, 1996

Abstract

Berman and Hartmanis [BH77] conjectured that there is a polynomial-time computable isomorphism between any two languages complete for NP with respect to polynomial-time computable many-one (Karp) reductions. Joseph and Young [JY85] gave a structural definition of a class of NP-complete sets—the k-creative sets—and defined a class of sets (the K_f^k's) that are necessarily k-creative. They went on to conjecture that certain of these K_f^k's are not isomorphic to the standard NP-complete sets. Clearly, the Berman–Hartmanis and Joseph–Young conjectures cannot both be correct.

We introduce a family of strong one-way functions, the *scrambling* functions. If f is a scrambling function, then K_f^k is not isomorphic to the standard NP-complete sets, as Joseph and Young conjectured, and the Berman-Hartmanis conjecture fails. Indeed, if scrambling functions exist, then the isomorphism also fails at higher complexity classes such as EXP and NEXP. As evidence for the existence of scrambling functions, we show that much more powerful one-way functions—the *annihilating* functions—exist relative to a random oracle.

[*] The first author was supported in part by NSF Grant DCR–8602562

[†] The second author was supported in part by NSF Grant CCR-91-10745.

[‡] The third author was supported in part by NSF Grants DCR–8602991 and CCR-89-011154.

Random oracles are the first examples of oracles relative to which the isomorphism conjecture fails with respect to higher classes such as EXP and NEXP.

1 Introduction

The relationship between the Berman-Hartmanis isomorphism conjecture and existence of one-way functions has been the subject of considerable research and conjecture in recent years [JY85, KLD86, HH91, FKR89].

We prove that the isomorphism conjecture is incompatible with the existence of *scrambling functions*, a type of powerful one-way function. To provide plausibility for the hypothesis that scrambling functions exist, we show that they exist relative to a random oracle, i.e., the set of oracles relative to which scrambling functions exist has measure one in the standard Lebesgue measure on languages. As a corollary, we obtain that the isomorphism conjecture fails with respect to a random oracle.

The remainder of Section 1 consists of three parts: a historical survey, a precise statement of our results, and some possible directions for future research. Section 2 describes our notation and nomenclature. Section 3 considers the structural consequences of the existence of scrambling functions, and Section 4 establishes the existence of scrambling functions (and still more powerful one-way functions called *annihilating functions*) relative to a random oracle.

This paper has been written so that it might be read in two different ways. General readers who want to know what we did and what it means, but not necessarily how we did it, will want to read the remainder of this section, using Section 2 as a general reference. Those readers who are interested in our proofs may find it more economical to skim the preliminaries and continue their reading with Sections 3 and 4.

1.1 A Brief Survey

In this section, we briefly survey the research that led to this work. The reader may wish to consult Young's excellent survey [You90] of structural research on isomorphisms, as well as the surveys by Mahaney [Mah86] and and Kurtz, Mahaney, and Royer [KMR90].

1.1.1 The Structural Approach

In [BH77] Berman and Hartmanis made the following conjecture:

The Isomorphism Conjecture *All* NP*-complete languages are polynomial-time isomorphic to one another.*

That is, in the terminology of Section 2, the NP-complete m-degree collapses.

As evidence for this conjecture, they adapted the proof of the Cantor-Bernstein Theorem to show that the paddable NP-complete languages are isomorphic to one another. As all languages polynomial-time isomorphic to a paddable language must themselves be paddable, it follows that the isomorphism conjecture is equivalent to the assertion that all NP-complete languages are paddable. By surveying the literature of the time on NP-complete languages, Berman and Hartmanis established that all of the then-known NP-complete languages were paddable, thereby providing empirical evidence for their conjecture.

In the years immediately following the isomorphism conjecture, research centered not on the conjecture itself, but rather on structural predictions of the conjecture. For example, the isomorphism conjecture predicts that there are no sparse NP-complete languages. Mahaney, building on work of P. Berman and Fortune, (cf. [Mah82]) verified this prediction under the hypothesis P $\neq$ NP.

Another direction pursued in the years immediately following the conjecture was to "relocate" it to other natural degrees. In their original article, Berman and Hartmanis conjectured not only that the NP-complete degree collapsed, but also that the PSPACE-complete degree collapsed. While they could not prove their conjectures, Berman [Ber77] was able to obtain a number of important partial results, e.g., that the complete m-degree for EXP consists of a single 1-li degree.

A re-examination of the isomorphism conjecture began with Joseph and Young's [JY85] definition of a new class of NP-complete languages—the k-creative languages. Joseph and Young considered the following specific k-creative languages:

$$K_f^k \quad = \quad \{\, f(i) : \Phi_i(f(i)) \leq |i| \cdot |f(i)|^k + |i| \,\},$$

where $k \geq 1$, f is a polynomial-time computable, honest function, and Φ_i is the run-time function of the i-th nondeterministic Turing Machine in some reasonable indexing of TMs. At present, it is only known how to pad a K_f^k when f is polynomial-time invertible. Joseph and Young went on to assert the following conjecture.

The Joseph-Young Conjecture *There exists a one-way function f such that for some k, K_f^k is nonpaddable.*

This conjecture goes beyond asserting that the Berman-Hartmanis conjecture fails: it asserts that a language with a specific form witnesses the failure. The K_f^k languages are not merely m-complete for NP, they are 1-li complete. Thus, the Joseph-Young conjecture predicts that the 1-li complete degree for NP fails to collapse. Selman [Sel92] observed that the the Joseph-Young conjecture entails the following, simpler, conjecture:

The Encrypted Complete Set Conjecture *There exists a one-way function f such that $f(\mathrm{SAT})$ is not is isomorphic to* SAT.

Intuitively, $f(\mathrm{SAT})$ is an encrypted version of SAT. It is easy to see that $f(\mathrm{SAT})$ is in NP: $x \in f(\mathrm{SAT})$ if and only if there exists an instance φ of SAT and a truth assignment ν such that $f(\varphi) = x$ and $\nu(\varphi) = \mathrm{true}$. As f is itself 1-1 and length increasing, we must have $f : \mathrm{SAT} \leq^{\mathrm{P}}_{\text{1-li}} f(\mathrm{SAT})$, and therefore $f(\mathrm{SAT})$ is NP-complete. Without further knowledge about f, there is essentially only one clear choice for a padding function:

$$p_{f(\mathrm{SAT})}(x, y) = f(p_{\mathrm{SAT}}(g(x), y))$$

where p_{SAT} is a padding function for satisfiability, and $g : f(\mathrm{SAT}) \leq^{\mathrm{P}}_{\text{1-li}} \mathrm{SAT}$. This candidate padding function is easily seen to be 1-1 (as g and p_{SAT} and f are 1-1), and length-increasing in both arguments (again, because g and p_{SAT} and f are). Unfortunately, the only way to recover y from $f(p_{\mathrm{SAT}}(g(x), y))$ appears to require the ability to invert f, and this we cannot do.

To summarize the foregoing discussion, we have the following four statements:

JYC: The Joseph-Young Conjecture some K_f^k is nonpaddable.

ECSC: The Encrypted Complete Set Conjecture there is a 1-way function f such that $f(\mathrm{SAT})$ is nonpaddable.

IC$_{1-li}$: The 1-li Isomorphism Conjecture The complete 1-li degree of NP consists of a single polynomial time isomorphism type, and

IC: The Isomorphism Conjecture The complete m-degree of NP consists of a single polynomial time isomorphism type.

with the following known structure:

$$\mathbf{JYC} \Rightarrow \mathbf{ECSC} \Rightarrow \neg\mathbf{IC}_{1-li} \Rightarrow \neg\mathbf{IC}.$$

It is not known which, if any, of these implications can be reversed at NP. Not surprisingly, more is known about higher complexity classes. A well-known theorem of Berman [Ber77] states that the complete m-degree of EXP consists of a single 1-li degree, i.e., $\neg\mathbf{IC}^{\mathrm{EXP}} \Rightarrow \neg\mathbf{IC}^{\mathrm{EXP}}_{1-li}$. A more recent theorem of Watanabe [Wat91, Theorem 4] states that $\neg\mathbf{IC}^{\mathrm{C}} \Rightarrow \mathbf{ECSC}^{\mathrm{C}}$ for deterministic complexity classes C that contain coNEXP.

The Berman-Hartmanis and Joseph-Young conjectures both predicted properties of the NP-complete degrees that were not known to hold anywhere: Berman and Hartmanis predicted that the complete m-degree for NP collapses, and yet no nontrivial collapsing m-degree was known to exist; while Joseph and Young predicted that the complete 1-li degree for NP does not collapse, and yet no noncollapsing 1-li degree was known to exist.

If one-way functions fail to exist, then every one-one length-increasing polynomial-time computable function is necessarily invertible, and so by results of Berman and Hartmanis, any two 1-li equivalent sets must be polynomial-time isomorphic, and thus every 1-li degree collapses. Therefore, a minimal hypothesis for the construction of a noncollapsing 1-li degree is the existence of a one-way function. Watanabe [Wat85] conjectured that the existence of a one-way function *is* an adequate hypothesis for the construction of a noncollapsing 1-li degree, and he was proven correct by Ko, Long, and Du [KLD86]. This validation of a prediction of the Joseph-Young conjecture is an important piece of evidence in its favor.

Somewhat later, we [KMR88] showed that there are nontrivial collapsing m-degrees, providing analogous evidence in favor of the Berman-Hartmanis conjecture. (See [KMR90] for a more complete discussion of this.)

1.1.2 Relativizations

Relativizations have long been used in complexity theory to probe the limitations of our proof techniques. Indeed, the use of relativizations has been so successful that at times it seems that *any* reasonable complexity theoretic statement holds relative to *some* oracle. The various isomorphism conjectures are a notable exception to this trend: it has proven very difficult to produce oracles relative to which one can decide the various conjectures.

The one simple relativization is an oracle relative to which the complete m-degree for EXP collapses. By Berman's theorem that the m-complete for degree EXP consists of a 1-li degree, it suffices to take an oracle relative to which one-way functions fail to exist, for if one-way functions don't exist, then *all* 1-li degrees must collapse. The original Baker-Gill-Solovay oracle relative to which P = NP suffices.

In contrast, and in spite of widely perceived similarities between NP and EXP, progress has not been made in obtaining an oracle relative to which the complete degree for NP collapses.

Kurtz [Kur83b] provided the first example of an oracle relative to which P $\neq$ NP and yet the isomorphism conjecture fails. The failure of the isomorphism conjecture relative to Kurtz's oracle is different from that predicted by Joseph and Young, as it is obtained by splitting the m-complete degree for NP into several 1-degrees. Hartmanis and Hemachandra [HH91], by combining Kurtz's construction with Rackoff's [Rac82] construction of an oracle relative to which P = UP $\neq$ NP, construct an oracle relative to which both conjectures fail. Thus, while the Berman-Hartmanis and the Joseph-Young conjectures cannot both be true, they can both be false.

After publication of the Hartmanis-Hemachandra paper, this was the state of knowledge on oracles and isomorphisms: we knew of oracles relative to which complete degrees at or above EXP collapsed, and we knew of oracles relative to which NP did not collapse, but we knew of no natural complexity class for which oracles of both sorts existed. Thus, while we knew of oracles that made the Berman-Hartmanis and Joseph-Young conjectures fail, we didn't know of any oracles relative to which either conjecture succeeded.

After [KLD86] and [KMR88] appeared, we hoped to break the impasse. We expected that the techniques of [KLD86] could be exploited in an oracle construction relative to which the Joseph-Young conjecture holds; and we expected that the techniques of [KMR88] could be used to construct an oracle relative to which the Berman-Hartmanis conjecture holds.

We achieved limited success [KMR87] by constructing a sparse oracle relative to which there is a collapsing m-degree in NP. As sparse oracles seem less likely to distort structural relationships, we take this as evidence for the proposition that some m-degree in NP collapses.

Homer and Selman [HS88] achieved the first breakthrough, producing an oracle relative to which the m-complete degree for Σ_2^p collapses, as well as an oracle relative to which it fails to collapse. By their efforts, the complete m-degree for Σ_2^p became the first natural degree to have both collapsing and noncollapsing relativizations. Recently, Fenner, Fortnow, and Kurtz [FFK92] demonstrated the existence of an oracle relative to which the isomorphism conjecture is true. Although the construction of these two oracles were substantial technical advances, neither is so natural as to provide a plausible source of intuition about the unrelativized case.

In this paper, we show that the encrypted complete set conjecture holds relative to a random oracle, i.e., the set of oracles relative to which the encrypted complete set conjecture holds has measure one in the usual Lebesgue measure on languages. Relative to a random oracle, higher complexity classes such as PSPACE and EXP fail to collapse, and so we provide numerous examples of natural complexity classes that can be relativized in both directions.

1.1.3 Random Oracles

We think of relativizations as providing alternative computational universes. One often stated complaint about relativizations, however, is that these alternative universes are not consistent with one another. In other words, if $\mathcal{T}$ is a statement about complexity theory, e.g., P $\neq$ NP, there may be oracles A and B such that $\mathcal{T}^A$ is true, and yet $\mathcal{T}^B$ is false. Because of this, the existence of an oracle A relative to which $\mathcal{T}^A$ holds can only be taken as weak evidence for (unrelativized) $\mathcal{T}$.

In 1981, Bennett and Gill made the following dramatic conjecture [BG81].

The Random Oracle Conjecture *If the set of oracles A relative to which a complexity theoretic statement $\mathcal{T}^A$ holds has measure one, i.e., if $\mathcal{T}$ holds relative to a* random *oracle, then the unrelativized $\mathcal{T}$ is true.*

They based their conjecture on the following intuition: We know that there are (qualitatively) good pseudorandom languages A in P. We expect that pseudorandom and random languages will have essentially the same properties. Thus, it is reasonable to expect that $\mathcal{T}^A$ will equal $\mathcal{T}^R$ for random R. But since A is in P, $\mathcal{T}^A = \mathcal{T}$, and so we expect that $\mathcal{T}$ will equal $\mathcal{T}^R$ as desired. Alternatively, either the Random Oracle Conjecture is true, or there are essential structural properties of random sets that cannot be captured by any pseudorandom set.

Since 1981, the random oracle conjecture has been refuted twice. Kurtz [Kur83a] pointed out that $\mathrm{coNP} \subseteq \mathrm{P}^{\mathrm{SAT}}$, but $\mathrm{coNP}^R \not\subseteq \mathrm{P}^{\mathrm{SAT},R}$ for random R. More recently, Chor, Goldreich, and Haståd [CGH90] showed that $\mathrm{coNP}^R \not\subseteq \mathrm{IP}^R$ for random R, which together with the Fortnow, Karloff, Lund, Nisan [FKLN92] proof that $\mathrm{PH} \subseteq \mathrm{IP}$ gives another counterexample.

Both counterexamples have a common flavor. Imagine that there are exponentially many boxes, one of which contains a prize. If the prize is placed at random, then a computational agent that can only examine polynomially many boxes has essentially no chance of finding the prize, no matter how powerful he may be. If, on the other hand, the prize is placed only pseudorandomly, then a sufficiently powerful computational agent will find it every time. Both counterexamples rely on finding computational agents ($\mathrm{P}^{\mathrm{SAT}}$ or IP) that are strong enough to defeat any polynomial time pseudorandom prize-hiding strategy; but these agents must themselves be defeated when presented with a truly random prize-hiding strategy. Stated somewhat differently, both $\mathrm{P}^{\mathrm{SAT}}$ and IP are sufficiently powerful to search any pseudorandom language; but for random R, neither $\mathrm{P}^{\mathrm{SAT},R}$ nor IP^R is powerful enough to search R.

The main result of this paper is that the isomorphism conjecture fails relative to a random oracle. In our case, the computational agents we will need to consider will be in P, and so, in some sense, we end up relying on the fact that P^R isn't powerful enough to search R. But now we're back to the observation upon which the random oracle conjecture was based: there seem to be good pseudorandom languages in P, and we only need them to be good enough to defeat P, not $\mathrm{P}^{\mathrm{SAT}}$ or IP. We believe that such pseudorandom languages exist, and that therefore the unrelativized isomorphism conjecture fails.

1.2 Overview of New Results

This section surveys the technical contributions of this paper. A one-way function (cf. Definition 3.1) is a polynomial-time computable, one-one, honest function that is not polynomial-time invertible. We have not been able to make progress on the Joseph-Young conjecture under the hypothesis that "vanilla" one-way functions exist. We have, however, been able to make considerable progress under a stronger hypothesis: the existence of scrambling functions.

Definition 3.2 A function f is a *scrambling function* if and only if f is a one-one, honest, polynomial-time computable function such that range(f) does not contain a nonempty paddable language.

The existence of scrambling functions implies the Joseph-Young conjecture:

Theorem 3.3 *If scrambling functions exist, the complete 1-li degree for* NP *fails to collapse.*

Theorem 3.4 *If f is a scrambling function, then K_f^k is a nonpaddable 1-li complete language for* NP*.*

Theorem 3.7 *If scrambling functions exist, then the complete 1-li degrees for* NP*,* PSPACE*,* EXP*,* NEXP*, and* RE *all fail to collapse.*

In as much as a direct proof of the existence of scrambling functions seems to be well beyond our immediate ability, as a surrogate, we looked for an oracle relative to which such functions exist. It was intuitively obvious that scrambling functions must exist relative to a random oracle. In fact, much more powerful one-way functions exist relative to a random oracle:

Definition 3.8 A function f is an *annihilating function* if and only if f is a one-way function such that all polynomial-time decidable subsets of range(f) are sparse.

It is not difficult to see that an annihilating function is necessarily a scrambling function.

Theorem 4.9 *Annihilating functions exist relative to a random oracle.*

Combining Theorems 3.7 and 4.9 yields

Theorem 4.10 *Relative to a random oracle, the complete 1-li degrees for* NP*,* PSPACE*,* EXP*,* NEXP*, and* RE *do not collapse. In particular, the isomorphism conjecture fails relative to a random oracle.*

1.3 Further Questions

We see a number of opportunities for improving on these results.

A first opportunity is to weaken the structural hypotheses that suffice to prove the encrypted complete set conjecture. We speculate that the encrypted complete set conjecture is equivalent to the existence of some sort of one-way function, more powerful than the "vanilla" one-way functions, and weaker than our scrambling functions.

A second opportunity is to explore additional structural consequences of the existence of scrambling and/or annihilating functions. It seems that the existence of annihilating functions ought to have profound structural consequences, and yet none of our structural theorems requires this power. In particular, we would like to see a proof that the existence of annihilating functions implies the complete m-degree for NP consists of a single 1-li degree, or perhaps that the existence of annihilating functions implies that the polynomial-time hierarchy separates. In view of Theorem 4.9, these structural consequences would immediately hold relative to a random oracle.

We would like to see structural hypotheses that are equivalent to the existence of these strong one-way functions, much as P $\neq$ UP is equivalent to the existence of one-way functions. This sort of structural taxonomy of one-way functions seems to have a great deal of promise.

A final opportunity is to look for more powerful structural properties that hold relative to random or generic oracles. We have found random oracles, in particular, to be a valuable "laboratory" for exploring the plausibility of various structural hypotheses. In particular, random oracles tend to be very good at separating deterministic and nondeterministic complexity classes, and at producing languages with very strong immunity properties.

2 Background Notation and Terminology

2.1 Numbers, Strings, and Languages

The set of natural numbers, $\{0, 1, 2, \ldots\}$, is denoted by ω. We identify each $x \in \omega$ with the x-th string over $\Sigma = \{0, 1\}$ in the lexicographic ordering on Σ^* and use natural numbers and strings over Σ^* interchangeably.

Languages are subsets of Σ^*. The characteristic functions are the total functions from ω to $\{0, 1\}$, and are denoted collectively by 2^ω. We identify each language $L \subseteq \Sigma^*$ with its characteristic function: $L(w) = 1$ means $w \in L$ and $L(w) = 0$ means $w \notin L$. The complement of L in Σ^*, i.e., $\Sigma^* - L$, is denoted by $\overline{L}$.

The cardinality of a language L is denoted by $||L||$. A language L is *sparse* if and only if there is a polynomial p such that for every n there are at most $p(n)$ elements of L of length at most n, i.e., $||L \bigcap \Sigma^n|| < p(n)$ for all $n \in \omega$.

2.2 Reducibility, Equivalence, and Isomorphism

If A and B are languages, and if $f\colon \Sigma^* \to \Sigma^*$ is computable in polynomial-time, then A is *polynomial-time many-one reducible to B via f* if and only if

$$x \in A \iff f(x) \in B$$

for all x in Σ^*. This relation is denoted by $f\colon A \leq_{\mathrm{m}}^{\mathrm{p}} B$. Often the specific reducibility f will not be mentioned, leaving us to simply write $A \leq_{\mathrm{m}}^{\mathrm{p}} B$. If $f\colon A \leq_{\mathrm{m}}^{\mathrm{p}} B$, and f is also one-one, then we say A is *polynomial-time one-one reducible* to B, and write $f\colon A \leq_{1}^{\mathrm{p}} B$. If f is length-increasing as well as one-one, then we say A is *polynomial-time 1-li reducible* to B, and write $f\colon A \leq_{\text{1-li}}^{\mathrm{p}} B$. As a general rule, we are only interested in (possibly relativized) polynomial-time reducibilities, and so we abbreviate *polynomial-time many-one reducible* by *m-reducible*, *polynomial-time one-one reducible* by *1-reducible*, and *polynomial-time 1-li reducible* by *1-li reducible*.

A language L is *complete* for a class C with respect to a reducibility $\leq_{\mathrm{r}}$ if and only if L is in C, and for all $L' \in$ C, $L' \leq_{\mathrm{r}} L$. In the literature, the term NP-*complete* is often used without specifying the intended reducibility. In the early literature, NP-*complete* usually meant with respect to *log-space* reductions. In more recent literature, NP-*complete* has come to mean with respect to *m-reductions*. We use the term in this latter sense.

If $A \leq_{\mathrm{m}}^{\mathrm{p}} B$ and $B \leq_{\mathrm{m}}^{\mathrm{p}} A$, then we say A and B are *polynomial-time many-one equivalent*, and write $A \equiv_{\mathrm{m}}^{\mathrm{p}} B$. The notions of *1-equivalent* and *1-li equivalent* are defined analogously. The collection of languages equivalent to a language A is called the *degree* of A. Thus, the m-degree of A is $\{B : A \equiv_{\mathrm{m}}^{\mathrm{p}} B\}$. The set of NP-complete languages is an important example of an m-degree.

If $f\colon A \leq_{\mathrm{m}}^{\mathrm{p}} B$, where f is one-one, onto, and polynomial-time invertible, then we say that f is a *polynomial-time isomorphism* between A and B, and write $f\colon A \cong^{\mathrm{p}} B$. We abbreviate *polynomial-time isomorphism* by *isomorphism*. We say that a degree *collapses* if and only if all of its members are isomorphic to one another. Notice that m-degrees and 1-degrees are always unions of isomorphism classes, but 1-li degrees need not be [KLD86].

A function f is *honest* if and only if there is a polynomial p such that for every $x \in \Sigma^*$, $|x| \leq p(|f(x)|)$. A function f is *one-way* if and only if f is a one-one, honest, polynomial-time computable function that has no polynomial-time computable inverse, i.e., there is no polynomial-time computable g such that, for all x, $g(f(x)) = x$. One-way functions are not known, but are widely believed, to exist.

2.3 Padding and Pairing

A *padding function* $\langle\langle\cdot,\cdot\rangle\rangle$ is a polynomial-time computable, one-one function from pairs of strings to strings that is polynomial-time invertible in both arguments [MY85]. A language A is *paddable* [BH77] if and only if for all x and y,

$$x \in A \iff \langle\langle x, y\rangle\rangle \in A.$$

If a padding function $\langle\langle\cdot,\cdot\rangle\rangle$ is also *onto*, then we say $\langle\langle\cdot,\cdot\rangle\rangle$ is a *polynomial-time pairing function.*

Let $\langle\cdot,\cdot\rangle$ be the standard Rogers' pairing function [Rog67, Page 64], where $\langle x, y\rangle = \frac{1}{2}((x+y)^2 + x + 3y)$. It is easy to see that $\langle x, y\rangle$ is a polynomial-time pairing function according to our definitions; moreover, $\langle\cdot,\cdot\rangle$ is length-nondecreasing in both arguments. Let $B \otimes C$ denote $\{\langle b, c\rangle : b \in B \wedge c \in C\}$. If B is m-complete for NP, then $A = B \otimes \Sigma^*$ is 1-li complete for NP.

Most complexity classes are closed under $\otimes$, i.e., if $\mathcal{C}$ is a complexity class, and $A, B \in \mathcal{C}$, then $A \otimes B \in \mathcal{C}$. Moreover, most complexity classes contain Σ^*. For such complexity classes, the construction above yields a simple but important result: If B is m-complete for $\mathcal{C}$, then $B \otimes \Sigma^*$ is paddable and 1-li complete for $\mathcal{C}$.

3 Structural Theorems

In this section, we consider various strengthenings of the definition of a one-way function. We show that if one-way functions of a particular type—scrambling functions—exist, then the complete 1-li degrees of several complexity classes all fail to collapse and, in particular, the isomorphism conjecture fails.

Definition 3.1. A function f is a *one-way function* if and only if f is honest, one-one, polynomial-time computable, and *not* polynomial-time invertible.

Our definition of one-way function requires totality, which is not the case in all presentations, e.g., [GS88]. Berman [Ber77], Grollmann and Selman [GS84, GS88], and Ko [Ko85] show that the existence of one-way functions is equivalent to P $\neq$ UP. Ko, Long, and Du [KLD86] show by a simple padding construction that if one-way functions exist, then length-increasing one-way functions exist.

We introduce two more powerful variants of the notion of a one-way function, and show that if these functions exist, then the complete 1-li degree for NP (and for many other natural complexity classes) does not collapse.

Definition 3.2. A function f is a *scrambling function* if and only if f is a one-way function and range(f) does not contain a nonempty paddable language.

As with "vanilla" one-way functions [KLD86, Proposition 2.1], if scrambling functions exist, then length-increasing scrambling functions exist.

The existence of scrambling functions implies that the encrypted complete set conjecture is valid.

Theorem 3.3 *If scrambling functions exist, then the complete 1-li degree for* NP *fails to collapse, and so the isomorphism conjecture fails.*

Proof: Let f be a length-increasing scrambling function and let A be paddable 1-li complete for NP. Consider $B = f(A)$. Since f is honest, it follows that B is in NP, and, moreover, since $f\colon A \leq_{\text{1-li}}^{\text{p}} B$, B is 1-li complete for NP.

As $B \subset$ range(f), B cannot be paddable. As paddability is an isomorphism invariant, A and B are not isomorphic.

□

It is natural to ask whether the existence of scrambling functions implies the Joseph-Young conjecture.

To this end, let $\Phi_i(x)$ denote the the running time of the i-th nondeterministic Turing machine on input x. Recall that Joseph and Young define

$$K_f^k \quad = \quad \{f(i) : \Phi_i(f(i)) < |i| \cdot |f(i)|^k + |i|\}$$

for $k > 0$ and one-one, honest, polynomial-time computable f.

By definition each K_f^k is a subset of range(f), and so is not paddable if f is a scrambling function. By the analysis in [JY85], the K_f^k's will be 1-li complete for NP whenever f is an honest polynomial-time computable 1-1 function. We have

Theorem 3.4 *If f is a scrambling function, then each K_f^k is a nonpaddable 1-li complete language for* NP*.*

The proof of Theorem 3.3 is far more general than it might initially appear. In particular, the hypothesis that A was 1-li complete for NP was only used to ensure that $B \leq^{\mathrm{P}}_{\text{1-li}} A$. By isolating this hypothesis, we can extend the proof of Theorem 3.3 to obtain the noncollapse of many other 1-li degrees.

Definition 3.5. A language A is *image complete* if and only if for every polynomial-time computable 1-li function f, $f(A) \leq^{\mathrm{P}}_{\text{1-li}} A$. A 1-li degree is *image complete* if and only if all of its members are image complete.

Theorem 3.6 *If scrambling functions exist, then every image complete 1-li degree with a paddable element fails to collapse.*

Image completeness is a property shared by the complete languages for most natural complexity classes containing NP. In particular, the 1-li complete languages for NP, PSPACE, EXP, NEXP, and RE are all image complete.

Theorem 3.7 *If scrambling functions exist, then the complete 1-li degrees for* NP*,* PSPACE*,* EXP*,* NEXP*, and* RE *all fail to collapse.*

In Section 4, we will show that there are oracles relative to which scrambling functions exist, indeed, that much more powerful sorts of one-way functions exist relative to random oracles.

Definition 3.8. A function f is an *annihilating* function if and only if f is a one-way function such that all polynomial-time decidable subsets of range(f) are sparse.

As before, if annihilating functions exist, then length-increasing annihilating functions exist.

Annihilating functions are, in one sense, the most powerful sort of one-way function possible, for the range of every polynomial-time computable one-one function must contain sparse sets in P with arbitrarily large polynomial census functions.

It is easy to see that every annihilating function is a scrambling function. The main result of the Section 4 is that annihilating functions exist relative to a random oracle.

4 Randomness

In this section we show that annihilating functions exist relative to a random oracle. Although it is easy to give heuristic arguments for this result, its proof requires a measure theoretic argument. Section 4.1 sketches most of the background needed for the proof which appears in Section 4.2. Readers should not be intimidated by measure theory. As we shall see below, the standard measure on 2^ω is a direct generalization of probability over finite spaces. We refer readers who want a systematic introduction to measure theory to any of the excellent texts of Dudley [Dud89], Oxtoby [Oxt80], or Rudin [Rud87].

4.1 Measure Theory Background

Here we briefly develop the standard Lebesgue measure on 2^ω and discuss two results of general measure theory: countable subadditivity and Kolmogorov's zero-one law.

Recall that 2^ω is the collection of all total functions from ω to $\{0,1\}$, or, equivalently, the collection of all infinite sequences of 0's and 1's. Let's view 2^ω as the collection of all possible infinite sequences of independent tosses of a fair coin. For $\mathcal{A} \subseteq 2^\omega$, the measure of $\mathcal{A}$ (written $\mu(\mathcal{A})$) is simply the probability that an element of 2^ω is in $\mathcal{A}$. One can formally define this measure as follows.

Let σ range over finite sequences of 0's and 1's and let $\langle\!\langle \sigma \rangle\!\rangle$ denote the collection of all infinite sequences that begin with σ, i.e., that have σ as an initial subsequence. Fix an arbitrary σ of length m. The probability that m independent tosses of a fair coin will produce σ is 2^{-m}. So, the probability that a randomly chosen *infinite* sequence begins with σ should be 2^{-m}. Thus we define $\mu(\langle\!\langle \sigma \rangle\!\rangle) = 2^{-m}$.

To extend μ beyond measuring the $\langle\!\langle \sigma \rangle\!\rangle$'s, the idea is to take $\mu(\mathcal{A})$ as the limit of measures of approximations to $\mathcal{A}$. We say that a countable collection of σ's, $\sigma_0, \sigma_1, \ldots,$ *covers* $\mathcal{A}$ if and only if $\mathcal{A} \subseteq \cup_{i=0}^{\infty} \langle\!\langle \sigma_i \rangle\!\rangle$ and we define the *size* of this cover to be $\sum_{i=0}^{\infty} \mu(\langle\!\langle \sigma_i \rangle\!\rangle)$. The *outer measure of* $\mathcal{A}$ (written $\mu^*(\mathcal{A})$) is the greatest lower bound of the sizes of covers of $\mathcal{A}$. One would like to define $\mu(\mathcal{A}) = \mu^*(\mathcal{A})$ for arbitrary $\mathcal{A}$, but there is a problem. Using the axiom of choice one can construct an $\mathcal{A}_0$ such that $\mu^*(\mathcal{A}_0) + \mu^*(\overline{\mathcal{A}_0}) > 1$, where $\overline{\mathcal{A}_0}$ denotes the complement of $\mathcal{A}_0$ in 2^ω [Oxt80]. On sets such as $\mathcal{A}_0$, μ^* fails to make sense as a probability measure. We thus say that $\mathcal{A}$ is *measurable* if and only if $\mu^*(\mathcal{A}) + \mu^*(\overline{\mathcal{A}}) = 1$ and then define $\mu(\mathcal{A}) = \mu^*(\mathcal{A})$ for measurable $\mathcal{A}$ and leave $\mu(\mathcal{A})$ undefined otherwise. It can be shown (cf. [Dud89, Oxt80, Rud87]) that all Borel sets are measurable. All of the $\mathcal{A}$ considered below will be first order definable, therefore Borel, and therefore measurable. We will use the term *probability* as a synonym for *measure*. E.g., if we say that a random oracle R is in $\mathcal{A}$ with probability ρ, this means that $\mathcal{A}$ has measure ρ.

Countable subadditivity is the property of μ that, if $\langle \mathcal{A}_i \rangle_{i \in N}$ is a sequence of measurable sets, then

$$\mu(\bigcup_{i \in N} \mathcal{A}_i) \quad \leq \quad \sum_{i \in N} \mu(\mathcal{A}_i).$$

In probabilistic language this says that the probability of a countable union of events is bounded above by the sum of the probabilities of the individual events. Countable subadditivity, simple reasoning about limits, and finitary probability theory are the primary mathematical tools in the next section.

A *tail set* is a subset $\mathcal{P}$ of 2^ω that is closed under finite variants, i.e., if X and Y are subsets of ω such that $X \bigtriangleup Y$ is finite, then $X \in \mathcal{P} \iff Y \in \mathcal{P}$. *Kolmogorov's zero-one law* [Oxt80, Theorem 21.3] states that a measurable tail set must have measure 0 or 1. The zero-one law thus gives us a means to convert bounds of measures of sets to exact measures, e.g., to show that a tail set $\mathcal{A}$ has measure 1, it is enough to show that $\mathcal{A}$ has positive measure.

If P is a predicate with $\mu(\{R : P^R\}) = 1$, then we say P *holds relative to a random oracle*. In essence, this defines our use of the word *random*. Structural properties such as $\{ X : \mathrm{P}^X \neq \mathrm{NP}^X \}$ are definable tail sets, and so have measure 0 or 1 by Kolmogorov's zero-one law. Informally, this means that there is a well-defined "measure 1" theory.

4.2 Annihilating functions exist relative to a random oracle

We focus our attention on the following function:

$$\xi^R(x) \quad = \quad R(x1)R(x10)\ldots R(x10^{3|x|}).$$

Intuitively, ξ^R maps x to a string of length $3|x|+1$ by copying and concatenating $3|x|+1$ independent "bits" from the oracle R. Note that for distinct x and x', the parts of the oracle that determine $\xi^R(x)$ and $\xi^R(x')$ are disjoint, i.e., the values $\xi^R(x)$ and $\xi^R(x')$ are independent. As ξ^R maps strings of length n to strings of length $3n+1$, it is honest.

Proposition 4.1 *ξ^R is an annihilating function with probability at least* $1/2$.

Lemmas 4.2 through 4.8 establish this proposition which, together with an application of Kolmogorov's zero-one law, will yield the existence of annihilating functions relative to a random oracle (Theorem 4.9).

Lemma 4.2 *ξ^R is one-one with probability at least* $1/2$.

Proof: If $\xi^R(a) = \xi^R(b)$, then a and b must have the same length. We show

$$\text{Prob}\,[\,(\exists a, b \in \Sigma^n)[\,a \neq b \;\;\&\;\; \xi^R(a) = \xi^R(b)\,]\,] \quad \leq \quad 2^{-n-2}. \tag{1}$$

Since $\sum_{n\in\omega} 2^{-n-2} = 1/2$, (1) implies that the probability that ξ^R *fails* to be one-one is no more than $1/2$ and so the lemma follows.

If a and b are distinct elements of length n, then the probability that they have the same image under ξ^R is exactly $1/2^{3n+1}$. There are $\binom{2^n}{2}$ distinct pairs of elements of length n, and so the probability that there exist two strings of length n having the same image under ξ^R can be bounded above by

$$\binom{2^n}{2}\frac{1}{2^{3n+1}} \quad = \quad \frac{2^{n-1}(2^n-1)}{2^{3n+1}} \quad < \quad \frac{2^{n-1}2^n}{2^{3n+1}} \quad = \quad \frac{1}{2^{n+2}}.$$

Therefore, (1) and the lemma follow.

□

The next lemma helps reduce the problem of showing Proposition 4.1 to the problem of reasoning about the behavior of individual machines that try to decide subsets of the range of ξ^R. Let $\widehat{M}$ range over relativized, polynomially-clocked, deterministic TMs in which the clocks do *not* depend on the oracle. For each such $\widehat{M}$, define

$$\mathcal{S}_{\widehat{M}} \quad = \quad \left\{\, R : L(\widehat{M}^R) \subseteq \text{range}(\xi^R) \implies L(\widehat{M}^R) \text{ is sparse} \,\right\}.$$

Lemma 4.3 *If, for each $\widehat{M}$, $\mu(\mathcal{S}_{\widehat{M}}) = 1$, then ξ^R is an annihilating function with probability at least* $1/2$.

Proof: The argument is organized into a series of three claims. First define

$$\mathcal{S} \quad = \quad \left\{\, R : (\forall M) \begin{bmatrix} \text{if } M^R \text{ is polynomial-time and } L(M^R) \\ \subseteq \text{range}(\xi^R), \text{ then } L(M^R) \text{ is sparse} \end{bmatrix} \right\},$$

where M ranges over deterministic, relativized TMs.

Claim 1 *If $\mu(\mathcal{S}) = 1$, then ξ^R is an annihilating function with probability at least* $1/2$.

The probability that ξ^R is an annihilating function is easily seen to be $\mu(\mathcal{S} \cap \{\, R : \xi^R \text{ is honest \& 1-1} \,\})$. Suppose $\mu(\mathcal{S}) = 1$. Then,

$$\begin{aligned} \mu(\mathcal{S} \cap \{\, R : \xi^R \text{ is honest \& 1-1} \,\}) &= \mu(\{\, R : \xi^R \text{ is honest \& 1-1} \,\}) \\ &= \mu(\{\, R : \xi^R \text{ is 1-1} \,\}) \\ &\geq 1/2, \quad \text{by Lemma 4.2.} \end{aligned}$$

Therefore, Claim 1 follows.

Claim 2 $\mathcal{S} = \cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}$.

Suppose that $R \in \mathcal{S}$. Then, by the definitions of $\mathcal{S}$ and the $\mathcal{S}_{\widehat{M}}$'s, $R \in \cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}$. Hence, $\mathcal{S} \subseteq \cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}$.

Suppose $R \in \overline{\mathcal{S}}$. Then, for some M, M^R is polynomial-time, $L(M^R) \subseteq$ range(ξ^R), and $L(M^R)$ is not sparse. Clearly, then, there is an $\widehat{M}$ such that $L(\widehat{M}^R) \subseteq$ range(ξ^R) and $L(\widehat{M}^R)$ is not sparse. So, $R \in \cup_{\widehat{M}} \overline{\mathcal{S}_{\widehat{M}}} = \overline{\cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}}$. Hence, $\mathcal{S} \supseteq \cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}$.

Therefore, Claim 2 follows.

Claim 3 *If, for each $\widehat{M}$, $\mu(\mathcal{S}_{\widehat{M}}) = 1$, then $\mu(\cap_{\widehat{M}} \mathcal{S}_{\widehat{M}}) = 1$.*

By countable subadditivity, the union of countably many sets of measure 0 is itself a set of measure 0. So, the intersection of countably many measure 1 sets is also measure 1. Therefore, Claim 3 follows.

Putting the three claims together, we obtain the lemma.

□

For the rest of this section fix an arbitrary $\widehat{M}$ and let $\widehat{L}^R = L(\widehat{M}^R)$. The aim of Lemmas 4.5 through 4.8 is to show that:

Relative to a random oracle R, if $\widehat{L}^R \subseteq$ range(ξ^R), then $\widehat{L}^R$ is sparse.

As our choice of $\widehat{M}$ was arbitrary, by the previous lemma it follows that this suffices to establish Proposition 4.1.

We introduce some more terminology. We say that $\widehat{M}^R$ *on argument y examines x* if and only if in the course of the computation of $\widehat{M}^R(y)$ the machine queries its oracle about some string of the form $x10^k$ for $k \leq 3|x|$. Intuitively, this means that the computation learns some information about the value of $\xi^R(x)$.

We decompose $\widehat{L}^R$ into two disjoint languages:

$$Q^R = \left\{ y \in \widehat{L}^R : \begin{array}{l} \widehat{M}^R \text{ on argument } y \text{ examines} \\ \text{some } x \text{ such that } \xi^R(x) = y \end{array} \right\}.$$

$$U^R = \left\{ y \in \widehat{L}^R : \begin{array}{l} \widehat{M}^R \text{ on argument } y \text{ fails to exam-} \\ \text{ine any } x \text{ such that } \xi^R(x) = y \end{array} \right\}.$$

Let's view $\widehat{M}^R$ as trying to accept a subset of range(ξ^R). One can then think of Q^R as being the "responsible" subset of $\widehat{L}^R$, i.e., the subset of $\widehat{L}^R$ consisting of those $y \in \widehat{L}^R$ for which $\widehat{M}^R$ has successfully obtained a preimage of y under ξ^R. In contrast, U^R is the "irresponsible" subset of $\widehat{L}^R$ as it contains $y \in \widehat{L}^R$ for which $\widehat{M}^R$ was not able to obtain a preimage under ξ^R. By a reasonably straightforward argument, the proof of Lemma 4.7 establishes that Q^R is sparse relative to a random oracle R. By a rather more involved argument, the proof of Lemma 4.6 shows that, relative to a random oracle R, if $U^R \subseteq$ range(ξ^R), then U^R is finite. At an intuitive level, this last assertion seems quite plausible: If you say that $y \in \Sigma^{3n+1}$ is in the range of ξ^R without querying R about any preimage x, you have a probability of $2^n/2^{3n+1} = 1/2^{2n+1}$ of being correct. You cannot expect to have an infinite run of wins against such odds. However, the formalization of this heuristic argument leads us into an analysis of a topological structure that ξ^R imposes on 2^ω.

Definition 4.4. We say that R and S are *x-variants* (written $R \sim_x S$) if and only if $R \triangle S \subseteq \{ x10^k : k \leq 3n \}$, i.e., R and S are identical except perhaps on the strings that determine the value of ξ on input x.

Clearly $\sim_x$ is an equivalence relation for every string x and every R has exactly $2^{3|x|+1}$ many x-variants (including R itself). The next lemma states the fundamental fact we use about x-variants.

Lemma 4.5 (The x-Variant Density Lemma) *Suppose $\epsilon \geq 0$, $x \in \omega$, and $\mathcal{A}$ is a measurable subset of 2^ω such that for every oracle R,*

$$\frac{\|\{ S : S \sim_x R \} \cap \mathcal{A}\|}{\|\{ S : S \sim_x R \}\|} \leq \epsilon.$$

Then, $\mu(\mathcal{A}) \leq \epsilon$.

Proof: Here, and only here, do we make use of more tools from measure theory than those introduced in Section 4.1. See [Dud89] or [Rud87] for background on product measures.

Let $n = |x|$ and let ω' be a distinct copy of ω. We factor 2^ω into a product of $\{0,1\}^{3n+1}$ and $2^{\omega'}$ in which, for each $R \in 2^\omega$, $R = (r_0, R_1)$ where $r_0 = R(x1)R(x10)\ldots R(x10^{3n})$ $(= \xi^R(x))$ and R_1 is just the sequence R with each of the elements at positions $x1$, $x10$, $\ldots$, and $x10^{3n}$ removed. Let $1_{\mathcal{A}}$ be the characteristic function of $\mathcal{A}$ over $\{0,1\}^{3n+1} \times 2^{\omega'}$, that is, for $R = (r_0, R_1)$, $1_{\mathcal{A}}(r_0, R_1) = 1$ when $R \in \mathcal{A}$ and $1_{\mathcal{A}}(r_0, R_1) = 0$ when $R \notin \mathcal{A}$.

We view $\{0,1\}^{3n+1}$ as a measure space under the uniform, normalized counting measure, i.e., each element of $\{0,1\}^{3n+1}$ has weight 2^{-3n-1}. Let μ_0 be this measure. Also let μ_1 be the standard Lebesgue measure on $2^{\omega'}$. One can show that μ is the product of measures μ_0 and μ_1, and so, by Fubini's Theorem [Dud89, Rud87],[1] we have that

$$\mu(\mathcal{A}) \quad = \quad \int\int 1_{\mathcal{A}}(r_0, R_1)\, d\mu_0(r_0)\, d\mu_1(R_1).$$

Now, using the terminology of the prior two paragraphs, the inequality of the hypothesis can be restated as: for each R_1 in $2^{\omega'}$,

$$\int 1_{\mathcal{A}}(r_0, R_1)\, d\mu_0(r_0) \quad \leq \quad \epsilon.$$

Therefore,

$$\mu(\mathcal{A}) \quad = \quad \int\int 1_{\mathcal{A}}(r_0, R_1)\, d\mu_0(r_0)\, d\mu_1(R_1) \quad \leq \quad \int \epsilon\, d\mu_1(R_1) \quad = \quad \epsilon$$

as required.

□

We introduce one more bit of terminology for the proof of the next lemma. If L is an oracle-dependent language, then we say L^R *on argument* y *depends on* x if and only if there is an S, $S \sim_x R$, such that $y \in L^R \bigtriangleup L^S$. Clearly, if our $\widehat{L}^R$ on argument y depends on x, then $\widehat{M}^R$ on argument y must examine x, but the converse is not necessarily true. Note that by the definition of U, if $y = \xi^R(x) \in U^R$, then $\widehat{M}^R$ on argument y does *not* depend on x.

Lemma 4.6 *With probability 0,* U^R *is an infinite subset of* $\mathrm{range}(\xi^R)$*.*

[1] Actually, we are using a special case of Fubini's Theorem, the *Product Measure Existence Theorem* [Dud89, Theorem 4.4.4].

Proof: For each k, define

$$\mathcal{R}(k) \quad = \quad \left\{ R \ : \ \begin{array}{l} U^R \subseteq \text{range}(\xi^R) \text{ and for some } x \\ \text{with } |x| \geq k \text{ we have } \xi^R(x) \in U^R \end{array} \right\}.$$

A little playing with quantifiers shows that $\cap_{k\geq 0}\mathcal{R}(k)$ is the collection of all oracles R such that U^R is an infinite subset of range(ξ^R). Thus, to show the lemma it suffices to prove that $\mu(\mathcal{R}(k)) \to 0$ as $k \to \infty$. To help establish this convergence, we define, for each $x \in \omega$:

$$\begin{aligned} \mathcal{C}(x) &= \left\{ R \ : \ U^R \subseteq \text{range}(\xi^R) \text{ and } \xi^R(x) \in U^R \right\}. \\ \mathcal{M}(x) &= \left\{ R \ : \ \text{for some } x' \neq x, \ \xi^R(x') = \xi^R(x) \right\}. \\ \mathcal{C}'(x) &= \left\{ R \ : \ \begin{array}{l} \text{(i) } U^R \subseteq \text{range}(\xi^R), \text{ (ii) } \xi^R(x) \in U^R, \text{ and} \\ \text{(iii) } x \text{ is the only string that } \xi^R \text{ maps to } \xi^R(x) \end{array} \right\} \\ &= \mathcal{C}(x) - \mathcal{M}(x). \end{aligned}$$

Note that $\mathcal{R}(k) = \cup_{n\geq k} \cup_{x\in\Sigma^n} \mathcal{C}(x)$ and $\mathcal{C}(x) \subseteq \mathcal{C}'(x) \cup \mathcal{M}(x)$. So, to bound $\mu(\mathcal{R}(k))$, we bound the $\mu(\mathcal{C}(x))$'s and, to bound $\mu(\mathcal{C}(x))$, we bound $\mu(\mathcal{C}'(x))$ and $\mu(\mathcal{M}(x))$.

Fix x and let $n = |x|$.

We first bound $\mu(\mathcal{M}(x))$. Since $\xi^R(x) = \xi^R(y)$ implies that $|y| = n$, the only strings different from x that ξ^R could map to $\xi^R(x)$ are the $y \in (\Sigma^n - \{ x \})$. There are $2^n - 1$ many such y and each has a 1 in 2^{3n+1} chance to map to $\xi^R(x)$. Since these events are pairwise independent, we therefore have

$$\mu(\mathcal{M}(x)) \quad = \quad \frac{2^n - 1}{2^{3n+1}}. \tag{2}$$

Next we bound $\mu(\mathcal{C}'(x))$. We show

$$\mu(\mathcal{C}'(x)) \quad \leq \quad \frac{1}{2^{3n+1}}. \tag{3}$$

To establish this, it suffices by Lemma 4.5 to argue that, for each R, at most one of R's 2^{3n+1} many x-variants can be in $\mathcal{C}'(x)$.

Pick an arbitrary R. If none of R's x-variants is in $\mathcal{C}'(x)$, we are done. So, suppose that at least one of R's x-variants is in $\mathcal{C}'(x)$. Without loss of generality suppose R itself is in $\mathcal{C}'(x)$. Let $y = \xi^R(x)$ and let S be an arbitrary x-variant of R distinct from R.

Claim 1 $y \in U^S$.

Since $R \in \mathcal{C}'(x)$, by clause (ii) in the definition of $\mathcal{C}'(x)$, $y \in U^R$. As U^R on argument y does not depend on x and since $y \in U^R$, it follows that $y \in U^{S'}$ for each $S' \sim_x R$. In particular, $y \in U^S$. Hence, Claim 1 follows.

Claim 2 $y \notin \text{range}(\xi^S)$.

Since $S \sim_x R$, but $S \neq R$, we have $\xi^S(x) \neq \xi^R(x) = y$. By clause (iii) in the definition of $\mathcal{C}'(x)$, ξ^R maps each string in $(\Sigma^* - \{x\})$ to someplace other than y. But, since $S \sim_x R$, ξ^S and ξ^R act identically on $(\Sigma^* - \{x\})$. Hence, Claim 2 follows.

Thus, by Claims 1 and 2, $y \in (U^S - \text{range}(\xi^S))$. So, by clause (i) in the definition of $\mathcal{C}'(x)$, $S \notin \mathcal{C}'(x)$. Since S was an arbitrary x-variant of R distinct from S, we therefore have that R is the only one of its x-variants in $\mathcal{C}'(x)$—as required. We thus obtain (3).

Now, since $\mathcal{C}(x) \subseteq \mathcal{C}'(x) \cup \mathcal{M}(x)$,

$$\begin{aligned} \mu(\mathcal{C}(x)) &\leq \mu(\mathcal{C}'(x)) + \mu(\mathcal{M}(x)) \\ &\leq \frac{1}{2^{3n+1}} + \frac{2^n - 1}{2^{3n+1}} \qquad \text{(by (2) and (3))} \\ &= \frac{1}{2^{2n+1}}. \end{aligned}$$

Since for each k we have $\mathcal{R}(k) = \cup_{n \geq k} \cup_{x \in \Sigma^n} \mathcal{C}(x)$, it follows by countable subadditivity that

$$\mu(\mathcal{R}(k)) \leq \sum_{n \geq k} \sum_{x \in \Sigma^n} \mu(\mathcal{C}(x)) \leq \sum_{n \geq k} \sum_{x \in \Sigma^n} \frac{1}{2^{2n+1}} = \sum_{n \geq k} \frac{1}{2^{n+1}} = \frac{1}{2^k}.$$

Therefore, $\lim_{k \to \infty} \mu(\mathcal{R}(k)) = 0$ as was to be shown.

□

We note that in the above argument one needs to use only two facts about U: that (a) for each x and y, if $y = \xi^R(x) \in U^R$, then U^R on argument y does *not* depend x and that (b) for each y, the set $\{R : y \in U^R\}$ is measurable. (Condition (b) guarantees that the $\mathcal{C}(x)$'s are measurable which in turn allows application of Lemma 4.5.) Thus, so long as U satisfies these two conditions, it does not have to be given by a polynomial-time deterministic TM, in fact U^R does not even have to be computable relative to R!

Lemma 4.7 *Q^R is sparse with probability* 1.

Proof: Let p be a polynomial such that, for each oracle R and input y, $p(|y|)$ bounds the run time of $\widehat{M}$. Recall that by the definitions of ξ^R and Q^R, $Q^R \subseteq \text{range}(\xi^R) \subseteq \cup_{n \geq 0} \Sigma^{3n+1}$.

For the moment fix a $y \in \Sigma^{3n+1}$. For any given $x \in \Sigma^n$, the probability that $\xi^R(x) = y$ is $1/2^{3n+1}$. Recall that the value of ξ^R on x depends solely on $R(x1), \ldots, R(x10^{3n})$ and is thus independent of the rest of the oracle R. Therefore, for each x, $1/2^{3n+1}$ is an upper bound on the probability that $\xi^R(x) = y$ and $\widehat{M}^R$ examines x on argument y. Since for any R, $\widehat{M}^R$ on argument y can examine no more than $p(3n+1)$ many x's of length n, the probability that $\widehat{M}^R$ on argument y examines a preimage of y is at most $p(3n+1)/2^{3n+1}$.

As the bound of the prior paragraph was for an arbitrary $y \in \Sigma^{3n+1}$, it follows that the expected number of elements of length $3n+1$ accepted by $\widehat{M}^R$ is bounded above by $2^{3n+1} \cdot (p(3n+1)/2^{3n+1}) = p(3n+1)$. By Markov's Inequality we know that if X is a nonnegative random variable and $a > 0$, then $\text{Prob}[X > a \cdot EX] < 1/a$. Thus, the probability that Q^R can contain more than $n^2 \cdot p(3n+1)$ many elements of length $3n+1$ is less than n^{-2}.

Therefore, for each k: the probability that there exists some $n > k$ such that $\|Q^R \cap \Sigma^{3n+1}\| > n^2 \cdot p(3n+1)$ is bounded above by

$$\sum_{n>k} \frac{1}{n^2} \quad < \quad \frac{1}{k-1}.$$

Thus, it follows that Q^R is sparse with probability 1.

□

Lemma 4.8 *With probability* 1, *if* $\widehat{L}^R \subseteq \text{range}(\xi^R)$, *then* $\widehat{L}^R$ *is sparse.*

Proof: Recall that $\widehat{L}^R$ is the disjoint union of Q^R and U^R. By Lemma 4.6, with probability 1, if U^R is a subset of range(ξ^R), then U^R is finite. By Q^R's definition, it is a subset of range(ξ^R) and, by Lemma 4.7, Q^R is sparse with probability 1. Thus, the lemma follows.

□

Therefore, by Lemmas 4.3 and 4.8, Proposition 4.1 follows. We can now prove

Theorem 4.9 *Annihilating functions exist relative to a random oracle.*

Proof: If an annihilating function exists with respect to an oracle R, then clearly an annihilating function exists with respect to all its finite variants. By Kolmogorov's zero-one law, the measure of the set of oracles R such that there is an annihilating function relative to R has measure 0 or 1. By Proposition 4.1, there is a set of positive measure on which ξ^R is an annihilating function. The theorem follows.

□

The next theorem is an immediate consequence of Theorem 4.9, Theorem 3.7, and the fact that all annihilating functions are scrambling functions.

Theorem 4.10 *Relative to a random oracle, the complete 1-li degrees for* NP*,* PSPACE*,* EXP*,* NEXP*, and* RE *do not collapse. In particular, the isomorphism conjecture fails relative to a random oracle.*

We also observe that relative to a random oracle, the NP-complete languages require exponential time to compute deterministically:

Corollary 4.11 *Relative to a random oracle, the smallest deterministic class that*

- *is closed under precomposition with the polynomial-time computable functions; and*
- *contains* NP

is EXP*.*

Proof: The proofs of Lemmas 4.3 through 4.8, *mutatis mutandis*, show that if T^R is a deterministic oracle Turing machine that makes fewer than $2^n/n^2$ queries on strings of length $3n+1$ for infinitely many n, then T^R cannot accept range(ξ^R) relative to a random oracle. The corollary follows immediately.

□

Independently of and simultaneously with our work, Rudich [Rud88] proved that, relative to a random oracle, there is a one way function f such that no BPP-machine can invert f on a nonsparse set. His proof is essentially the proof of Lemma 4.7, together with the observation that P = BPP relative to a random oracle. By combining our Theorem 3.9 with Rudich's observations, we obtain the following purely complexity theoretic result:

Corollary 4.12 *Relative to a random oracle, there is a* UP *set whose only* BPP *subsets are sparse.*

5 Acknowledgments

We would like to acknowledge the contributions of a number of our colleagues. The following people commented on an earlier version of this paper: Joan Feigenbaum, Stephen Fenner, Lane Hemachandra, Stephen Homer, Neil Immerman, Jeffrey Legarias, Tim Long, Alan Selman, Janos Simon, and Osamu Watanabe. We thank them for their time and effort. We thank the anonymous referees, who favored us with extensive and helpful comments. This paper was significantly improved by their efforts. We thank László Babai for several helpful discussions. We thank Alan Selman for sharing with us his yet-unpublished research on the Joseph-Young conjecture. We thank Charles Bennett and John Gill for several discussions about the intuitive foundations of the Random Oracle Hypothesis. Finally, we would like to thank the two gracious antagonists—Juris Hartmanis and Paul Young—whose differing insights have provided us so with many fine problems.

References

[Ber77] Leonard Berman, *Polynomial reducibilities and complete sets*, Ph.D. thesis, Cornell University, 1977.

[BG81] Charles H. Bennett and John Gill, *Relative to a random oracle A,* $\mathrm{P}^A \neq \mathrm{NP}^A \neq \mathrm{coNP}^A$*, with probability 1*, SIAM Journal on Computing **10** (1981), 96–113.

[BH77] Leonard Berman and Juris Hartmanis, *On isomorphism and density of* NP *and other complete sets*, SIAM Journal on Computing **6** (1977), 305–322.

[CGH90] Benny Chor, Oded Goldreich, and Johan Håstad, *The random oracle hypothesis is false*, Tech. Report 631, Dept. of Computer Science, Technion Institute, 1990.

[Dud89] Richard M. Dudley, *Real analysis*, Wadsworth & Brooks/Cole, 1989.

[FFK92] Stephen A. Fenner, Lance Fortnow, and Stuart A. Kurtz, *The isomorphism conjecture holds relative to an oracle*, Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science (New York), IEEE, New York, 1992, pp. 30–39.

[FKLN92] Lance Fortnow, Howard Karloff, Carsten Lund, and Noam Nisan, *Algebraic methods for interactive proof systems*, Journal of the ACM **39** (1992), 859–868.

[FKR89] Stephen A. Fenner, Stuart A. Kurtz, and James S. Royer, *Every polynomial-time 1-degree collapses iff P = PSPACE*, Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science, 1989, pp. 624–629.

[GS84] Joachim Grollman and Alan L. Selman, *Complexity measures for public-key cryptosystems*, Proceedings of the 25th Annual IEEE Symposium on Foundations of Computer Science, 1984, pp. 495–503.

[GS88] Joachim Grollman and Alan L. Selman, *Complexity measures for public-key cryptosystems*, SIAM Journal on Computing **17** (1988), 309–335.

[HH91] Juris Hartmanis and Lane A. Hemachandra, *One-way functions and the nonisomorphism of* NP*-complete sets*, Theoretical Computer Science **81** (1991), 155–163.

[HS88] Steven Homer and Alan L. Selman, *Oracles for structural properties: The isomorphism problem and public-key cryptography*, Proceedings of the 4th Annual Structure in Complexity Theory, 1988, pp. 3–14.

[JY85] Deborah Joseph and Paul Young, *Some remarks on witness functions for nonpolynomial and noncomplete sets in* NP, Theoretical Computer Science **39** (1985), 225–237.

[KLD86] Ker-I Ko, Timothy J. Long, and Ding-Zhu Du, *On one-way functions and polynomial-time isomorphisms*, Theoretical Computer Science **47** (1986), 263–276.

[KMR87] Stuart A. Kurtz, Stephen R. Mahaney, and James S. Royer, *Progress on collapsing degrees (extended abstract)*, Proceedings of the 2nd Annual Structure in Complexity Theory, 1987, pp. 126–131.

[KMR88] Stuart A. Kurtz, Stephen R. Mahaney, and James S. Royer, *Collapsing degrees*, Journal of Computer System Science **37** (1988), 247–268.

[KMR90] Stuart A. Kurtz, Stephen R. Mahaney, and James S. Royer, *The structure of complete degrees*, Complexity Theory Retrospective (A. Selman, ed.), Springer-Verlag, 1990, pp. 108–146.

[Ko85] Ker-I Ko, *On some natural complete operators*, Theoretical Computer Science **37** (1985), 1–30.

[Kur83a] Stuart A. Kurtz, *On the random oracle hypothesis*, Information and Control **57** (1983), 40–47.

[Kur83b] Stuart A. Kurtz, *A relativized failure of the Berman-Hartmanis conjecture*, Tech. Report 83–001, University of Chicago, 1983.

[Mah82] Stephen R. Mahaney, *Sparse complete sets for* NP*: solution of a conjecture of Berman and Hartmanis*, Journal of Computer System Science **25** (1982), 130–143.

[Mah86] Stephen R. Mahaney, *Sparse sets and reducibilities*, Studies in Complexity Theory (Ronald V. Book, ed.), John Wiley & Sons, Inc., 1986, pp. 63–118.

[MY85] Stephen R. Mahaney and Paul Young, *Reductions among polynomial isomorphism types*, Theoretical Computer Science **39** (1985), 207–224.

[Oxt80] John C. Oxtoby, *Measure and category*, 2nd ed., Springer-Verlag, New York, 1980.

[Rac82] Charles Rackoff, *Relativized questions involving probabilistic algorithms*, Journal of the ACM **29** (1982), 261–268.

[Rog67] Hartley Rogers, Jr., *Theory of recursive functions and effective computability*, McGraw-Hill, New York, 1967.

[Rud87] Walter Rudin, *Real and complex analysis*, 3rd ed., McGraw-Hill, 1987.

[Rud88] Steven Rudich, *Limits on the provable consequences of one-way functions*, Ph.D. thesis, University of California at Berkeley, 1988.

[Sel92] Alan Selman, *A survey of one way functions in complexity theory*, Mathematical Systems Theory **25** (1992), 203–221.

[Wat85] Osamu Watanabe, *On one-one polynomial time equivalence relations*, Theoretical Computer Science **38** (1985), 157–165.

[Wat91] Osamu Watanabe, *On the p-isomorphism conjecture*, Theoretical Computer Science **83** (1991), 337–343.

[You90] Paul Young, *Juris Hartmanis: Fundamental contributions to isomorphism problems*, Complexity Theory Retrospective (A. Selman, ed.), Springer-Verlag, 1990, pp. 28–58.