Complete Sets under Non-Adaptive Reductions are Scarce

Harry Buhrman*
CWI

Dieter van Melkebeek†
CWI & The University of Chicago

Abstract

We investigate the frequency of complete sets for various complexity classes within EXP under non-adaptive reductions in the sense of resource bounded measure. We show that these sets are rare:

- The sets that are complete under $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reductions for NP, the levels of the polynomial-time hierarchy, PSPACE, and EXP have p_2-measure zero for any constant $\alpha < 1$.
- Assuming MA $\neq$ EXP, the $\leqslant^p_{\mathrm{tt}}$-complete sets for PSPACE and the Δ-levels of the polynomial-time hierarchy have p-measure zero.

A key ingredient is the Small Span Theorem, which states that for any set A in EXP at least one of its lower span (i.e., the sets that reduce to A) or its upper span (i.e., the sets that A reduces to) has p_2-measure zero. Previous to our work, the theorem was only known to hold for $\leqslant^p_{k-\mathrm{tt}}$-reductions for any constant k. We establish it for $\leqslant^p_{n^{o(1)}-\mathrm{tt}}$-reductions.

1 Introduction

Lutz introduced resource bounded measure [Lut90] to formalize the notions of scarceness and abundance in complexity theory. His approach makes it possible to express statements like "only a few" or "most" sets in a complexity class $\mathcal{C}$ have property P. Many papers investigate resource bounded measure in relation with complexity theory [JL95, LM94, May94, AS94, LM96, RSC95, Lut96, ASB97].

We can also use resource bounded measure as a tool for separating complexity classes. For example, if we could show that the complete sets in complexity class $\mathcal{C}$ have measure zero and the complete sets in $\mathcal{D}$ do not, we would have separated $\mathcal{C}$ from $\mathcal{D}$.

In this paper we follow that line of research. We investigate complete and hard sets for NP, the levels of the polynomial-time hierarchy, PSPACE and

*Partially supported by the Dutch foundation for scientific research (NWO) through SION project 612-34-002, and by the European Union through NeuroCOLT ESPRIT Working Group Nr. 8556, and HC&M grant ERB-4050-PL-93-0516. E-mail: buhrman@cwi.nl. URL: http://www.cwi.nl/cwi/people/Harry.Buhrman.html.

†Partially supported by the European Union through TMR grant ERB-4001-GT-96-0783 and by the U.S. National Science Foundation through grant CCR 92-53582. Email: dieter@cs.uchicago.edu. URL: http://www.cs.uchicago.edu/~dieter.

EXP, and give some evidence that they have p_2-measure zero. On the other hand, the results of Bennett and Gill [BG81] imply that the $\leqslant^p_{\rm tt}$-hard sets for BPP do not have p_2-measure zero; Allender and Strauss [AS94] even showed they have p_2-measure 1 in EXP.

We use two different approaches to obtain our results. The first strategy yields unhypothesized statements on the border of what is provable by relativizable techniques. Essentially, we significantly improve the Small Span Theorem of Juedes and Lutz [JL95]. The Small Span Theorem for a reducibility $\leqslant^p_r$ states that for any set A in EXP, either the class of sets that $\leqslant^p_r$-reduce to A (called the lower span of A) or the class of sets that A $\leqslant^p_r$-reduces to (the upper span of A) or both have p_2-measure 0. Since the degree of a set is the intersection of its lower and upper span, it implies that every $\leqslant^p_r$-degree has p_2-measure zero, and in particular the $\leqslant^p_r$-complete degree of any complexity class within EXP. The strongest Small Span Theorem previous to our work was due to Ambos-Spies, Neis, and Terwijn [ASNT96], who proved it for $\leqslant^p_{k-\rm tt}$-reductions for any constant k. Even the extension to $\leqslant^p_{\rm btt}$-reductions was a notorious open problem in the area. We establish the Small Span Theorem for $\leqslant^p_{n^{o(1)}-\rm tt}$-reductions, i.e., for non-adaptive reductions that make a subpolynomial number of queries.

Lutz [Lut95] obtained a Small Span Theorem for non-uniform reductions w.r.t. p-space measure. Similar to his proof, our Small Span Theorem follows from the fact that most sets in EXP have a $\leqslant^p_{n^{o(1)}-\rm tt}$-upper span with p_2-measure zero. We actually establish this fact for $\leqslant^p_{n^{\alpha}-\rm tt}$-reductions for any constant $\alpha < 1$. This way, we get stronger results on the scarceness of complete sets than the ones that follow from the Small Span Theorem: Any $\leqslant^p_{n^{\alpha}-\rm tt}$-degree within EXP has p_2-measure zero. Previously, it was only known for $\leqslant^p_{\rm btt}$-reductions that the p_2-measure of the complete sets for EXP have p_2-measure zero [ASNT96, BM95]. We also obtain that the p_2-measure of the $\leqslant^p_{n^{\alpha}-\rm tt}$-hard sets for E and EXP is zero.

Our proofs relativize and are on the edge of the scope of relativizable techniques: Showing the last theorem for unbounded growing exponent α would separate BPP from EXP.

Therefore, we next look at what we can show under a non-relativizing reasonable but yet unproven complexity theoretic hypothesis, namely the assumption that MA $\neq$ EXP. Here we use a different approach to establish the scarceness of complete sets: the connection between completeness and autoreducibility. A set A is autoreducible if on input x, when allowed to make queries to A except for x itself, we can decide in polynomial-time whether $x \in A$. Buhrman, Fortnow and Torenvliet [BFT95] showed that the $\leqslant^p_{\rm T}$-complete sets for EXP are $\leqslant^p_{\rm T}$-autoreducible and that the $\leqslant^p_{\rm tt}$-complete sets for PSPACE and the Δ-levels of the polynomial-time hierarchy are $\leqslant^p_{\rm tt}$-autoreducible. So, measure zero results for the autoreducible sets will imply measure zero results for the complete sets for PSPACE, the Δ-levels of the polynomial-time hierarchy, and EXP.

Babai, Fortnow, Nisan and Wigderson [BFNW93] established the existence of a pseudo-random generator that can be used to simulate BPP in subexponential time for infinitely many input lengths unless MA = EXP. Using their ideas, we are able to prove that the $\leqslant^p_{\rm tt}$-autoreducible sets have p-measure zero unless

MA = EXP. It follows immediately that the $\leqslant^p_{\mathrm{tt}}$-complete sets for PSPACE and the Δ-levels of the polynomial-time hierarchy have p-measure zero.

Summarizing our results:

- We prove a Small Span Theorem for $\leqslant^p_{n^{o(1)}-\mathrm{tt}}$-reductions.
- We show that the complete sets for NP, the levels of the polynomial-time hierarchy, PSPACE, and EXP under $\leqslant^p_{n^{\alpha}-\mathrm{tt}}$-reductions have p_2-measure zero for any $\alpha < 1$.
- We show that the hard sets for E and EXP under $\leqslant^p_{n^{\alpha}-\mathrm{tt}}$-reductions have p_2-measure zero for any $\alpha < 1$.
- We show that the $\leqslant^p_{\mathrm{tt}}$-complete sets for PSPACE and the Δ-levels of the polynomial-time hierarchy have p-measure zero unless MA = EXP (and the polynomial-time hierarchy collapses).

The organization of this paper is as follows: We first give the necessary background on resource bounded measure and on pseudo-random generators. Section 3 describes our results without any complexity theoretic assumption; section 4 those using the hypothesis MA $\neq$ EXP. Finally, we give some comments and mention remaining open problems.

2 Notation and Preliminaries

Most of our complexity theoretic notation is standard. We refer the reader to the textbooks by Balcázar, Díaz and Gabarró [BDG95, BDG90], and by Papadimitriou [Pap94].

A *reduction* of a set A to a set B is a polynomial-time oracle Turing machine M such that $M^B = A$. We say that A reduces to B and write $A \leqslant^p_{\mathrm{T}} B$ ('T' for Turing). The reduction M is *non-adaptive* if the oracle queries M makes on any input are independent of the oracle. In that case we write $A \leqslant^p_{\mathrm{tt}} B$ ('tt' for truth-table). If in addition, the number of queries on an input of length n is bounded by $q(n)$, we write $A \leqslant^p_{q(n)-\mathrm{tt}} B$. For a reducibility $\leqslant^p_r$, we define the *lower span* of a set A as $\mathrm{P}_r(A) = \{B \mid B \leqslant^p_r A\}$, and the *upper span* of A as $\mathrm{P}^{-1}_r(A) = \{B \mid A \leqslant^p_r B\}$. The $\leqslant^p_r$-*degree* of A equals $\mathrm{P}_r(A) \cap \mathrm{P}^{-1}_r(A)$.

A set A is *autoreducible* if there is a reduction M of A to itself such that M^A never queries its input.

2.1 Background on Resource Bounded Measure

For our purposes, we only have to define what it means to have resource bounded measure zero.

Definition 1 *A* supermartingale *is a function* $d : \Sigma^* \to [0, \infty)$ *satisfying*

$$d(w) \geqslant \frac{d(w0) + d(w1)}{2} \tag{1}$$

for every $w \in \Sigma^*$. If equality holds in (1) for all w, d is called a martingale.
A supermartingale succeeds *on a sequence $\omega \in \Sigma^\infty$ if $d(\omega) = \limsup_{w \sqsubseteq \omega,\, w \to \omega} d(w) = \infty$. It* covers *a class $\mathcal{C}$ of sequences if it succeeds on every sequence in $\mathcal{C}$.*

A martingale d describes a strategy for an infinite one-person betting game. At the beginning of the game, an infinite bit sequence ω is fixed but not revealed. The player starts with initial capital $d(\lambda)$, and in each round guesses the next bit of ω and bets some of his capital on that outcome. Then the actual value of the bit is revealed. On a correct guess, the player earns the amount of money he bet, otherwise he looses it. The value of $d(w)$ equals the capital of the player after being revealed the bit sequence w. The player wins on ω if he manages to make his capital arbitrarily high during the game. A supermartingale describes a similar game, but now the player is allowed to throw away some of his capital in every round.

Martingales yield the following characterization:

Theorem 1 *A class $\mathcal{C} \subseteq \Sigma^\infty$ has Lebesgue measure zero iff it can be covered by a martingale iff it can be covered by a supermartingale.*

We obtain a resource bounded variant by putting resource bounds on the martingales.

Definition 2 ([Lut92]) *A (super)martingale d is a p-(super)martingale (resp. p_2-(super)martingale) if we can compute $d(w)$ in time polynomial in $|w|$ (resp. in time $2^{\log^{O(1)}|w|}$).*
A system d_i of (super)martingales is p-uniform (resp. p_2-uniform) if we can compute $d_i(w)$ in time polynomial in $|w| + i$ (resp. in time $2^{\log^{O(1)}(|w|+i)}$).
A class $\mathcal{C} \subseteq \Sigma^\infty$ has p-measure (resp. p_2-measure) zero if it can be covered by a p-supermartingale (resp. p_2-supermartingale). We denote this by $\mu_p(\mathcal{C}) = 0$ (resp. $\mu_{p_2}(\mathcal{C}) = 0$).

As in the unbounded case, the resource bounded measure-zero relations are monotone and closed under union. The following resource bounded version of closure under countable unions holds:

Theorem 2 ([Lut92]) *Let d_i be a p-uniform (resp. p_2-uniform) system of supermartingales such that d_i covers the class $\mathcal{C}_i$. Then $\cup_i \mathcal{C}_i$ has p-measure (resp. p_2-measure) zero.*

Characteristic sequences provide the link between resource bounded measure and complexity theory: We associate with a set $A \subseteq \Sigma^*$ its characteristic sequence $\chi_A = A(s_0)A(s_1)A(s_2)\ldots$, where $s_0, s_1, s_2, \ldots$ is the enumeration of Σ^* in lexicographical order.

The crucial property of the resource bounded measure-zero concepts not shared with the Lebesgue measure-zero concept, is that $\mu_p(\mathrm{E}) \neq 0$ and $\mu_{p_2}(\mathrm{EXP}) \neq 0$ [Lut92].

2.2 Background on Pseudo-Random Generators

Definition 3 ([NW94]) *The* hardness $H_A(n)$ *of a set* A *at length* n *is the largest integer* s *such that for any circuit* C *of size at most* s *with* n *inputs*

$$|\Pr_x[C(x) = A(x)] - \frac{1}{2}| \leqslant \frac{1}{s},$$

where x *is uniformly distributed over* Σ^n.
A pseudo-random generator *is a function* G *that, for each* n, *maps* Σ^n *into* $\Sigma^{r(n)}$ *where* $r(n) > n$. *The* security $S_G(n)$ *of* G *at length* n *is the largest integer* s *such that for any circuit* C *of size at most* s *with* $r(n)$ *inputs*

$$|\Pr_x[C(x) = 1] - \Pr_y[C(G(y)) = 1]| \leqslant \frac{1}{s},$$

where x *is uniformly distributed over* $\Sigma^{r(n)}$ *and* y *over* Σ^n.

For our purposes, we will need a pseudo-random generator computable in E that stretches seeds superpolynomially and has superpolynomial security at infinitely many lengths. We will use the one provided by the following theorem:

Theorem 3 *If* MA $\neq$ EXP, *there is a pseudo-random generator* G *computable in* E *with* $r(n) \in n^{\theta(\log n)}$ *such that for any integer* k, $S_G(n) \geqslant n^k$ *for infinitely many* n.

The proof follows directly from the next results of Babai, Fortnow, Nisan and Wigderson [BFNW93], and Nisan and Wigderson [NW94], combined with some padding.

Theorem 4 ([BFNW93]) *If* MA $\neq$ EXP, *there is a set* $A \in$ EXP *such that for any integer* k, $H_A(n) \geqslant n^k$ *for infinitely many* n.

Theorem 5 ([NW94]) *Given any set* $A \in$ EXP, *there is a pseudo-random generator* G *computable in* EXP *with* $r(n) \in n^{\theta(\log n)}$ *such that* $S_G(n) \geqslant \sqrt{H_A(\sqrt{n})}$.

3 Complete Sets for Non-Adaptive Reductions with n^α Queries and Small Span Theorem

In this section, we establish our results on the measure of complete and hard sets under non-adaptive reductions that hold without any complexity theoretic assumption. The following theorem forms the main ingredient. It states that most sets in EXP have a small upper span under $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reductions for constant $\alpha < 1$, and has a strong connection with the Small Span Theorem we will prove further.

Theorem 6 *For any* $\alpha < 1$,

$$\mu_p(\{A \in \mathrm{EXP} \mid \mu_{p_2}(\mathrm{P}^{-1}_{n^\alpha-\mathrm{tt}}(A)) \neq 0\}) = 0.$$

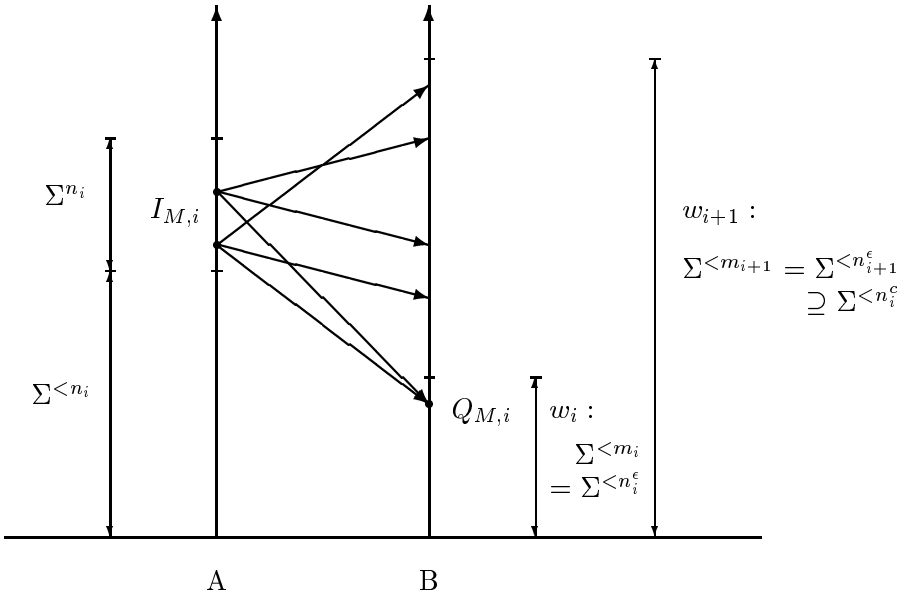

Figure 1: Betting strategies at stage i

We first give an outline of the proof.

Fix a $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reduction M running in time n^c for some constant $c > 0$, and a set $A \in \mathrm{EXP}$. We would like to construct a p_2-martingale that succeeds on any set B for which $M^B = A$. Suppose we are given the initial segment w_i of χ_B corresponding to all strings of length less than m_i. See Figure 1. We can select an input x of length $n_i = m_i^{1/\epsilon}$ for some constant $\epsilon > 0$, and divide the available capital uniformly among the extensions w'_{i+1} of w_i corresponding to all strings of length less than m_{i+1} ($m_{i+1} \geqslant n_i^c$) for which $M^{w'_{i+1}}(x) = A(x)$. This way, our capital at the end of stage i is definitely not smaller than at the beginning, and in case only half or fewer of the extensions pass the consistency test on x, we actually double it or even better. In order to be able to bet on the sets $A \in \mathrm{EXP}$ for which this strategy fails on some set B such that $M^B = A$, we will perform the consistency check not for a single input x of length n_i, but for a certain collection $I_{M,i}$ of $2n_i^\alpha$ inputs x of length n_i: We distribute the available capital uniformly over all extensions w'_{i+1} for which $M^{w'_{i+1}}(x) = A(x)$ for every $x \in I_{M,i}$. If there is an input $x \in I_{M,i}$ for which only half or fewer of the extensions w'_{i+1} satisfy $M^{w'_{i+1}}(x) = A(x)$, we gain a factor of 2 or more in stage i while betting on B. We will try this strategy at every stage i, and we succeed on B if the latter situation occurs for infinitely many of them.

Now, suppose that for some B to which M reduces A, this situation only occurs for finitely many stages. So for almost all stages i, on any input $x \in I_{M,i}$ more than half of the extensions w'_{i+1} of w_i satisfy $M^{w'_{i+1}}(x) = A(x)$. We would like to construct a p-martingale that succeeds on any such $A \in \mathrm{EXP}$ by betting on these x's according to the majority vote of the extensions. We do not know the prefix w_i of χ_B we need for that, but we can guess the values of the bits in this prefix which M queries on inputs $x \in I_{M,i}$, i.e., divide our

capital uniformly over all possible corresponding strategies. In order for this to work, we will make sure that the set $I_{M,i}$ consists of $2n_i^\alpha$ strings of length n_i on which M makes the *same* queries of length less than m_i. This implies we have to distribute our capital among no more than $2^{n_i^\alpha}$ strategies, and at least one of them will realize a relative gain of $2^{|I_{M,i}|} = 2^{2n_i^\alpha} \gg 2^{n_i^\alpha}$, which suffices to succeed on A.

We define the stages as follows:

$$\left\{ \begin{array}{rcl} m_0 & = & 1 \\ m_{i+1} & = & 2^{m_i} \\ n_i & = & m_i^{1/\epsilon}. \end{array} \right. \tag{2}$$

Note that, no matter for what constant c the reduction M runs in time n^c, the stages do not interfere at sufficiently high levels, i.e., $m_{i+1} \leqslant n_i^c$ for i sufficiently large.

Next, we show that for sufficiently large i, the sets $I_{M,i}$ exist for any $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reduction M, and that we can construct them efficiently. Here we need the fact that $\alpha < 1$.

Lemma 7 *Let $\alpha < 1$, $\epsilon \in (0, 1-\alpha)$, and m_i and n_i defined by (2). There is an integer i_0 such that for any $i \geqslant i_0$ and for any $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reduction M, there is a set of strings $Q_{M,i}$ such that*

$$|\{x \in \Sigma^{n_i} \mid Q_M(x) \cap \Sigma^{<m_i} = Q_{M,i}\}| \geqslant 2n_i^\alpha,$$

where $Q_M(x)$ denotes the set of queries M makes on input x. Moreover, we can find the lexicographically first set $Q_{M,i}$ and the lexicographically first subset $I_{M,i}$ of

$$\{x \in \Sigma^{n_i} \mid Q_M(x) \cap \Sigma^{<m_i} = Q_{M,i}\}$$

with $|I_{M,i}| = 2n_i^\alpha$ in time 2^{2n_i}.

Proof (of Lemma 7)
For sufficiently large i, the number of possible values of $Q_M(x) \cap \Sigma^{<m_i}$ for $x \in \Sigma^{n_i}$ is bounded by

$$\sum_{i=0}^{n_i^\alpha} \binom{2^{m_i}-1}{i} \leqslant (2^{m_i})^{n_i^\alpha} = 2^{n_i^{\alpha+\epsilon}} \leqslant \frac{2^{n_i}}{2n_i^\alpha},$$

from which the existence of $Q_{M,i}$ follows. A brute force search does the job. □

We now formalize the above outline.

Proof (of Theorem 6)
We use the notation from Lemma 7. Fix $A \in \mathrm{DTIME}[2^{n^k}]$. Let

$$\pi_{A,M}(w) = \left\{ \begin{array}{ll} 1 & \text{if } |w| < 2^{m_{i_0}} \\ \Pr_{\omega \sqsupseteq w}[\forall\, x \in I_{M,i} : M^\omega(x) = A(x)] & \text{if } 2^{m_{i_0}} \leqslant 2^{m_i} \leqslant |w| < 2^{m_{i+1}}. \end{array} \right.$$

We define the martingale $d_{A,M}$ as follows:

$$\begin{aligned} d_{A,M}(\lambda) &= 1 \\ d_{A,M}(wb) &= \begin{cases} \frac{2\cdot\pi_{A,M}(wb)}{\pi_{A,M}(wb)+\pi_{A,M}(w\overline{b})}\cdot d_{A,M}(w) & \text{if } \pi_{A,M}(wb)+\pi_{A,M}(w\overline{b})\neq 0 \\ d_{A,M}(w) & \text{otherwise.} \end{cases} \end{aligned}$$

This means that for any sufficiently large i (such that $i \geqslant i_0$ and stage $i+1$ does not interfere with stage i) and for any prefix w_i of length $2^{m_i}-1$, the martingale $d_{A,M}$ distributes $2^{2^{m_{i+1}}-2^{m_i}}\cdot d_{A,M}(w_i)$ uniformly over all extensions w'_{i+1} of w_i with $|w'_{i+1}| = 2^{m_{i+1}}-1$ for which $M^{w'_{i+1}}$ and A agree on the membership of every string in $I_{M,i}$.

The defining predicate of $\pi_{A,M}$ depends on at most $|I_{M,i}|\cdot n_i^\alpha \leqslant 2(\log|w|)^{\frac{2\alpha}{\epsilon}}$ positions of ω not fixed by w. It follows that $\pi_{A,M}$ and $d_{A,M}$ can be computed in time $2^{(\log|w|)^{O(\frac{\alpha+k}{\epsilon})}}$.

We distinguish between two cases for the behavior of M and A: Either there are infinitely many stages i such that no matter what the prefix w_i is, there is always an input in $I_{M,i}$ on which only half or fewer of the extensions pass the consistency check between M and A; or else for almost all stages i, there is a prefix w_i such that for any input from $I_{M,i}$, a strict majority of the extensions of w_i make M and A agree on that input.

Case 1 $\exists^\infty i, \forall\, w\in\Sigma^{2^{m_i}-1}, \exists\, x\in I_{M,i} : \Pr_{\omega\sqsupseteq w}[M^\omega(x)=A(x)] \leqslant \frac{1}{2}$.

Then for any $\omega=\chi_B$ such that M reduces A to B, and for any sufficiently large stage i for which the Case 1 condition holds,

$$d_{A,M}(w_{i+1}) \geqslant 2d_{A,M}(w_i),$$

where w_j represents the prefix of ω of length $2^{m_j}-1$. This is because at least half of the extensions w'_{i+1} of w_i with $|w'_{i+1}| = 2^{m_{i+1}}-1$ fail some consistency test. It follows that $d_{A,M}(\omega)=\infty$, and that

$$\mu_{p_2}(\{B \mid M \text{ reduces } A \text{ to } B\}) = 0. \tag{3}$$

Case 2 $\forall^\infty i, \exists\, w\in\Sigma^{2^{m_i}-1}, \forall\, x\in I_{M,i} : \Pr_{\omega\sqsupseteq w}[M^\omega(x)=A(x)] > \frac{1}{2}$.

For any stage i and any $\vec{b}\in\Sigma^{|Q_{M,i}|}$, let $\delta_{M,i,\vec{b}}$ be the martingale with initial capital 1 that only bets on strings of $I_{M,i}$, and for such a string $x\in I_{M,i}$ bets all of its money according to the majority of $M^\omega(x)$ over all sequences $\omega\sqsupseteq v_i$, where v_i is the characteristic string of length $2^{m_i}-1$ in which the bit corresponding to the j-th element of $Q_{M,i}$ equals the j-th bit of $\vec{b}$, and all other bits are say 0. Ties are broken arbitrarily. The martingale

$$\delta_{M,i} = \frac{1}{2^{|Q_{M,i}|}}\Sigma_{\vec{b}}\delta_{M,i,\vec{b}}$$

has initial capital 1 and is computable in time n^2. It has the property

$$\delta_{M,i}(\chi_A) \geqslant \frac{2^{|I_{M,i}|}}{2^{|Q_{M,i}|}} \geqslant \frac{2^{2n_i^\alpha}}{2^{n_i^\alpha}} = 2^{n_i^\alpha},$$

provided i satisfies the Case 2 condition. Since there are infinitely many such i's (actually all of them except for finitely many), the p-martingale $\delta_M = \sum_{i=1}^{\infty} \frac{1}{i^2} \delta_{M,i}$ succeeds on A.

Fix an enumeration M_j of all $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reductions such that we can compute $M_j(x)$ in time polynomial in $2^{|x|} + j$. Then the martingale system δ_{M_j} is p-uniform, so there is a p-martingale δ that succeeds on all sets A for which Case 2 applies for some $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reduction M. Consider any set $A \in \mathrm{EXP}$ not covered by δ. Since the martingale system d_{A,M_j} is p_2-uniform, equation (3) implies that the p_2-measure of $\mathrm{P}^{-1}_{n^\alpha-\mathrm{tt}}(A)$ is zero. □

Our results on the measure of complete sets follow directly from Theorem 6:

Corollary 8 *For any $\alpha < 1$ and $C \in \mathrm{EXP}$, the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-degree of C has p_2-measure zero. In particular, the classes of $\leqslant^p_{n^\alpha-\mathrm{tt}}$-complete sets for* NP*, the levels of the polynomial-time hierarchy,* PSPACE*, and* EXP *all have p_2-measure zero.*

Proof
Suppose not, then for any set A in the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-degree of C, the p_2-measure of $\mathrm{P}^{-1}_{n^\alpha-\mathrm{tt}}(A)$ is not zero, since it contains the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-degree of C. But, by Theorem 6 this would imply that the p-measure of the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-degree of C is zero. □

For the class of $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets, we get:

Corollary 9 *For any $\alpha < 1$ and any complexity class $\mathcal{C}$ such that $\mu_p(\mathcal{C} \cap \mathrm{EXP}) \neq 0$, the class of $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets for $\mathcal{C}$ has p_2-measure zero. In particular, the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets for* E *and* EXP *have p_2-measure zero.*

Proof
By definition, for any set $A \in \mathcal{C}$, the $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets for $\mathcal{C}$ are contained in $\mathrm{P}^{-1}_{n^\alpha-\mathrm{tt}}(A)$. If the class of $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets for $\mathcal{C}$ does not have p_2-measure zero, Theorem 6 yields that $\mu_p(\mathcal{C} \cap \mathrm{EXP}) = 0$. □

The $\leqslant^p_{n^\alpha-\mathrm{tt}}$-hard sets for NP, the levels of the polynomial-time hierarchy, and PSPACE also have p_2-measure zero, provided these classes themselves do not have p-measure zero.

From Theorem 6, we can also deduce a Small Span Theorem. However, we have to settle for a more restrictive reducibility than $\leqslant^p_{n^\alpha-\mathrm{tt}}$, because we need transitivity in the proof, and $\leqslant^p_{n^\alpha-\mathrm{tt}}$ is in general not transitive for any constant $\alpha > 0$. It suffices to keep the number of queries subpolynomial, i.e., asymptotically smaller than n^ϵ for any $\epsilon > 0$.

Theorem 10 (Small Span Theorem) *For any set A, at least one of the following holds:*
$\mu_p(\mathrm{P}_{n^{o(1)}-\mathrm{tt}}(A) \cap \mathrm{EXP}) = 0$ *or* $\mu_{p_2}(\mathrm{P}^{-1}_{n^{o(1)}-\mathrm{tt}}(A)) = 0$.

Proof
We distinguish between two cases:

- Either $\mathrm{P}_{n^{o(1)}-\mathrm{tt}}(A)$ contains a set B such that $\mu_{p_2}(\mathrm{P}^{-1}_{n^{o(1)}-\mathrm{tt}}(B)) = 0$. Then the transitivity of $\leqslant^p_{n^{o(1)}-\mathrm{tt}}$ and the monotonicity of p_2-measure imply that $\mu_{p_2}(\mathrm{P}^{-1}_{n^{o(1)}-\mathrm{tt}}(A)) = 0$.

- Or else $\mathrm{P}_{n^{o(1)}-\mathrm{tt}}(A) \cap \mathrm{EXP}$ is included in $\{B \in \mathrm{EXP} \mid \mu_{p_2}(\mathrm{P}^{-1}_{n^{\alpha}-\mathrm{tt}}(B)) \neq 0\}$ for any $\alpha > 0$. Then Theorem 6 says that $\mu_p(\mathrm{P}_{n^{o(1)}-\mathrm{tt}}(A) \cap \mathrm{EXP}) = 0$.

□

For any set $A \in \mathrm{EXP}$, Theorem 10 states that at least one of its lower span or upper span under $\leqslant^p_{n^{o(1)}-\mathrm{tt}}$-reductions is small.

4 Complete Sets and Autoreducible Sets under Non-Adaptive Reductions

We now show that we can remove the n^α upper bound on the number of queries for some of the results of the previous section, under the likely hypothesis that MA $\neq$ EXP. We exploit the link between completeness and autoreducibility, combined with the following theorem:

Theorem 11 *The class of $\leqslant^p_{\mathrm{tt}}$-autoreducible sets has p-measure zero unless* MA = EXP.

The proof goes as follows: First, for a given $\leqslant^p_{\mathrm{tt}}$-autoreduction M running in time n^c, we construct a martingale d_M computable in time $2^{\log^{O(c)} n}$ that succeeds on all sets which M autoreduces. This implies the known theorem [BL96] that for any c, the class of sets $\leqslant^p_{\mathrm{tt}}$-autoreducible in time n^c, has p_2-measure zero.

Next we show how to mimic the behavior of d_M for any polynomial-time computable M by a (super)martingale $\tilde{d}_M$ computable uniformly in time 2^{an} for some constant a. This boils down to efficiently approximating the probabilities that underlie the construction of d_M. Fairly weak pseudo-random generators suffice for doing so, and their existence follows from the hypothesis MA $\neq$ EXP.

We will use the set system provided by the next lemma.

Lemma 12 *Let M be a $\leqslant^p_{\mathrm{tt}}$-autoreduction computable in time n^c, and m be an integer. Provided $9m^c \log m \leqslant 2^{m-1}$, there is a set $I_{M,m}$ consisting of $3 \log m$ strings of length m such that for no $x \in I_{M,m}$, the autoreduction M queries a string in $I_{M,m}$. Moreover, we can construct such a set $I_{M,m}$ in time 2^{2m}.*

Proof (of Lemma 12)
Consider the following algorithm for constructing $I_{M,m}$. Let the set D originally contain all strings of length m that are queried by the autoreduction for no more than $2m^c$ inputs of length m. Note that there are at least 2^{m-1} such strings. Start with $I_{M,m}$ empty and repeat the following step $3 \log m$ times: Add the lexicographically smallest element x of D to $I_{M,m}$, and delete from D the element x, the queries the autoreduction M makes on input x, and the inputs from D for which the autoreduction M queries x. This construction has

the invariant that for no $x \in I_{M,m}$ the autoreduction M queries a string from $I_{M,m} \cup D$, and for no $x \in D$ it queries a string from $I_{M,m}$. In each step, we remove no more than $3m^c$ elements from D, so we can execute k steps provided $2^{m-1} - 3(k-1)m^c \geqslant 1$. Under the hypothesis given, we can perform $3 \log m$ steps, and the resulting set $I_{M,m}$ satisfies the requirements. We can compute the initial value of D by brute force search, and the whole construction runs in time no more than 2^{2m}. □

Proof (of Theorem 11)
Fix a $\leqslant^p_{\mathrm{tt}}$-autoreduction M computable in time n^c.

We call a sequence ω consistent with M on a set of inputs X if M^ω and ω agree on the membership of every string in X. Let

$$\pi_{M,m}(w) = \Pr_{\omega \sqsupseteq w}[\omega \text{ consistent with } M \text{ on } I_{M,m}],$$

where $I_{M,m}$ is the set constructed in Lemma 12. Consider

$$d_{M,m}(w) = \begin{cases} m^3 \cdot \pi_{M,m}(w) & \text{if } 9m^c \log m \leqslant 2^{m-1} \\ 1 & \text{otherwise.} \end{cases}$$

Note that the consistency of ω with M on $I_{M,m}$ depends on at most $3m^c \log m$ positions in ω, namely those corresponding to the elements of $I_{M,m}$ and the queries the autoreduction makes on $I_{M,m}$. Moreover, for any setting of the queries, there is exactly one consistent setting of the elements of $I_{M,m}$. Therefore, $d_{M,m}(w) = 1$ if $|w| < 2^m$. For sufficiently large m, $d_{M,m}(\omega) = m^3$ if M autoreduces ω. It follows that $d_{M,m}$ is a martingale computable in time $2^{O(\log^{c+1} n)}$, as well as $d_M = \sum_{m=1}^{\infty} \frac{1}{m^2} d_{M,m}$, and that d_M succeeds on any set which M autoreduces.

We now want to construct (super)martingales $\tilde{d}_{M,m}$ and $\tilde{d}_M$ that behave like $d_{M,m}$ and d_M, but are computable uniformly in time 2^{an} for some constant a, i.e., independent of the running time of M. The idea is to efficiently approximate the probability $\pi_{M,m}$ using a pseudo-random generator. The consistency with M on $I_{M,m}$ can be checked by circuits of size $O(m^{c+1})$. So, if we fill in the missing positions using a pseudo-random generator that at length m has security at least say m^{c+4} and produces pseudo-random strings of length at least $3m^c \log m$, we can get an approximation $\tilde{\pi}_{M,m}$ of $\pi_{M,m}$ to within $\epsilon_{M,m} = m^{-(c+4)}$. The pseudo-random generator G provided by Theorem 3 (assuming MA $\neq$ EXP), does this for infinitely many lengths m. We call these m's good. Moreover, we can compute the resulting approximation in time $2^{a_1 m}$ for some constant a_1.

There are still two technical problems we have to solve in order to make sure that $\tilde{d}_{M,m}$ is a supermartingale: First, what to do if m is not good? We will deal with that in a moment. Second, even for a good m, just replacing $\pi_{M,m}$ with $\tilde{\pi}_{M,m}$ in the definition of $d_{M,m}$ might not work. For example, if $\tilde{\pi}_{M,m}$ underestimates $\pi_{M,m}$ on input w, and overestimates it on input $w0$ and $w1$, condition (1) is violated. Note that such a situation can only occur in case the string corresponding to the position right after w is either an element of $I_{M,m}$

or a query the autoreduction M makes on an input from $I_{M,m}$. Since there are no more than $3m^c \log m$ such strings and the limit $\epsilon_{M,m}$ on the estimation error is such that $(3m^c \log m) \cdot \epsilon_{M,m}$ remains bounded, we can remedy this problem by accumulatively subtracting a term $2\epsilon_{M,m}$ from the approximation for $\pi_{M,m}$, and adding a constant to the resulting approximation for $d_{M,m}$. The former modification guarantees that condition (1) is met; the latter is needed after the former in order to keep the values non-negative. More precisely, we define

$$d^*_{M,m}(w) = m^3 \tilde{\pi}_{M,m}(w) + 1 - 2q_{M,m}(w)m^3\epsilon_{M,m},$$

where $q_{M,m}(w)$ denotes the number of positions in w that correspond to elements of $I_{M,m}$ or queries the autoreduction makes on $I_{M,m}$. Note that $0 = q_{M,m}(\lambda) \leqslant q_{M,m}(w) \leqslant q_{M,m}(\omega) \leqslant 3m^c \log m$.

We solve the first problem by explicitly checking for each prefix w that the values $d^*_{M,m}$ proposes for the one bit extensions $w0$ and $w1$, satisfy the defining conditions of a supermartingale. If they do, we accept them; otherwise we enforce the conditions by not betting. So, we define the function $\tilde{d}_{M,m}$ as follows:

$$\begin{aligned} \tilde{d}_{M,m}(\lambda) &= \min(d^*_{M,m}(\lambda), 3) \\ \tilde{d}_{M,m}(wb) &= \begin{cases} d^*_{M,m}(wb) & \text{if } d^*_{M,m}(w0) \geqslant 0 \text{ and } d^*_{M,m}(w1) \geqslant 0 \text{ and} \\ & \quad d^*_{M,m}(w0) + d^*_{M,m}(w1) \leqslant 2\tilde{d}_{M,m}(w) \\ \tilde{d}_{M,m}(w) & \text{otherwise.} \end{cases} \end{aligned} \tag{4}$$

It follows that $\tilde{d}_{M,m}$ is a supermartingale computable in time $2^{a_2 n}$ for some constant a_2 independent of M and m, and that $1 \leqslant \tilde{d}_{M,m}(\lambda) \leqslant 3$.

Claim 1 *If m is good and sufficiently large, $\tilde{d}_{M,m} = d^*_{M,m}$.*

Proof (of Claim 1)
We show that $\tilde{d}_{M,m}(w) = d^*_{M,m}(w)$ for any $w \in \Sigma^*$ by induction on $|w|$. The base case follows from the fact that $d^*_{M,m}(\lambda) = m^3\tilde{\pi}_{M,m}(\lambda) + 1 \leqslant 1 + m^3\epsilon_{M,m} + 1 \leqslant 3$. For the induction step, it suffices to argue that the conditions on the right-hand side of (4) are met, assuming that $\tilde{d}_{M,m}(w) = d^*_{M,m}(w)$. The first two of them are satisfied, since for any string v,

$$d^*_{M,m}(v) \geqslant 1 - 2q_{M,m}(v)m^3\epsilon_{m,m} \geqslant 1 - 6\epsilon_{M,m}m^{c+3}\log m = 1 - \frac{6\log m}{m},$$

which is positive for sufficiently large m. The remaining condition is certainly satisfied in case $d^*_{M,m}(w0) = d^*_{M,m}(w1) = d^*_{M,m}(w)$. In the other case, $q_{M,m}(w0) = q_{M,m}(w1) = q_{M,m}(w) + 1$, and we have that

$$\begin{aligned} & d^*_{M,m}(w0) + d^*_{M,m}(w1) \\ &= m^3(\tilde{\pi}_{M,m}(w0) + \tilde{\pi}_{M,m}(w1)) + 2 - 2(q_{M,m}(w0) + q_{M,m}(w1))m^3\epsilon_{M,m} \\ &\leqslant m^3(\pi_{M,m}(w0) + \pi_{M,m}(w1) + 2\epsilon_{M,m}) + 2 - 4(q_{M,m}(w) + 1)m^3\epsilon_{M,m} \\ &= 2m^3(\pi_{M,m}(w) - \epsilon_{M,m}) + 2 - 4q_{M,m}(w)m^3\epsilon_{M,m} \\ &\leqslant 2m^3\tilde{\pi}_{M,m}(w) + 2 - 4q_{M,m}(w)m^3\epsilon_{M,m} = 2d^*_{M,m}(w) = 2\tilde{d}_{M,m}(w). \end{aligned}$$

□

So, for a good and sufficiently large m we have that

$$\tilde{d}_{M,m}(\omega) = d^*_{M,m}(\omega) \geqslant d_{M,m}(\omega) + 1 - (2q_{M,m}(\omega)+1)m^3\epsilon_{M,m} \geqslant d_{M,m}(\omega).$$

Since there are infinitely many good m's and $d_{M,m}(\omega) = m^3$ if M autoreduces ω, this implies that $\tilde{d}_M = \sum_{m=1}^{\infty} \frac{1}{m^2}\tilde{d}_{M,m}$ is a supermartingale that succeeds on any set that M autoreduces. It is computable in time 2^{an} for some constant a independent of M.

Since for a standard enumeration M_i of all $\leqslant^p_{\mathrm{tt}}$-autoreductions such that $M_i(x)$ is computable in time $(2^{|x|}+i)^{O(1)}$, the supermartingale system $\tilde{d}_{M_i}$ is p-uniform, this proves that the $\leqslant^p_{\mathrm{tt}}$-autoreducible sets have p-measure zero. □

Buhrman, Fortnow and Torenvliet [BFT95] showed that all $\leqslant^p_{\mathrm{tt}}$-complete sets for PSPACE as well as for the Δ-levels of the polynomial-time hierarchy are $\leqslant^p_{\mathrm{tt}}$-autoreducible. The combination with Theorem 11 yields the following result on the measure of $\leqslant^p_{\mathrm{tt}}$-complete sets:

Corollary 13 *The classes of $\leqslant^p_{\mathrm{tt}}$-complete sets for* PSPACE *and for the Δ-levels of the polynomial-time hierarchy have p-measure zero unless* MA = EXP.

5 Discussion and Open Problems

The question whether Theorem 6 holds for constant $\alpha \geqslant 1$, remains open. A positive answer would be the best result provable by relativizable techniques, since Theorem 6 fails for unbounded growing exponent α in a world where BPP = EXP, and such a world exists [Hel86]. This follows from the relativizable result of Bennett and Gill [BG81] that $\mu(\mathrm{P}^{-1}_{\mathrm{tt}}(A)) = 1$ for any $A \in$ BPP. Because of the same reason, relativizable techniques cannot establish the Small Span Theorem for $\leqslant^p_{\mathrm{tt}}$-reductions.

It seems unlikely that our approach allows to establish Theorem 6 for $\alpha \geqslant 1$, because of Lemma 7. For some constant $\epsilon > 0$ and a given $\leqslant^p_{n^\alpha-\mathrm{tt}}$-reduction M, this would require the construction of sets $I_{M,i}$ containing $n_i^{O(1)}$ strings of length n_i and slightly smaller sets $Q_{M,i}$, such that all queries of length less than n_i^ϵ that M makes on inputs from $I_{M,i}$ are in $Q_{M,i}$. However, the following argument shows that for $\alpha \geqslant 1$, it is not even possible for $|I_{M,i}|$ to equal $|Q_{M,i}|$ when for every input $x \in \Sigma^{n_i}$ the queries are chosen from $\Sigma^{<n_i^\epsilon}$ in a Kolmogorov random way. The concatenation σ of all these queries is a Kolmogorov random string of length $2^{n_i}n_i^{\alpha+\epsilon}$. Given a listing of the elements of $Q_{M,i}$, we can describe the queries for elements of $I_{M,i}$ by pointers to that list. Assuming $|I_{M,i}| = |Q_{M,i}| = q$, this leads to a description of σ of length at most $qn_i^\epsilon + q(n_i + n_i^\alpha \log q) + (2^{n_i} - q)n_i^{\alpha+\epsilon} + O(\log q)$, which is asymptotically less than $|\sigma|$ as long as $\log q \leqslant cn_i^\epsilon$ for some constant $c < 1$. Since we have $\log q \in O(\log n_i)$, we get a contradiction to the Kolmogorov randomness of σ.

Regarding our results using the hypothesis MA $\neq$ EXP, we would like to see our technique applied to NP. The $\leqslant^p_{\mathrm{tt}}$-complete sets for NP are only known

to be autoreducible under *non-uniform* $\leqslant^p_{tt}$-autoreductions [BFT95], where the advice consists of random strings fed to a Valiant-Vazirani algorithm. Is there a way to exploit the pseudo-random generators following from the assumption MA $\neq$ EXP to extend Corollary 13 to NP?

For EXP (and E), our approach cannot work, because it relativizes, and in a world where EXP = EXPSPACE, MA $\neq$ EXP but not all $\leqslant^p_{tt}$-complete sets for EXP are $\leqslant^p_{tt}$-autoreducible [BFT95].

Acknowledgments

We would like to thank Jack Lutz and Lance Fortnow for very helpful discussions regarding the Small Span Theorem.

References

[AS94] E. Allender and M. Strauss. Measure on small complexity classes, with applications for BPP. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science*, pages 807–818, New York, 1994. IEEE.

[ASB97] K. Ambos-Spies and L. Bentzien. Separating NP-completeness notions under strong hypothesis. In *Proceedings of the 12th IEEE Conference on Computational Complexity*, New York, 1997. IEEE. To appear.

[ASNT96] K. Ambos-Spies, H.-C. Neis, and S. Terwijn. Genericity and measure for exponential time. *Theoretical Computer Science*, 168(1):3–19, 1996.

[BDG90] J. Balcázar, J. Díaz, and J. Gabarró. *Structural Complexity II*, volume 22 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1990.

[BDG95] J. Balcázar, J. Díaz, and J. Gabarró. *Structural Complexity I*, volume 11 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1995.

[BFNW93] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3:307–318, 1993.

[BFT95] H. Buhrman, L. Fortnow, and L. Torenvliet. Using autoreducibility to separate complexity classes. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 520–528, New York, 1995. IEEE.

[BG81] C. Bennett and J. Gill. Relative to a random oracle, $P^A \neq NP^A \neq co - NP^A$ with probability one. *SIAM Journal on Computing*, 10:96–113, 1981.

[BL96] H. Buhrman and L. Longpré. Compressibility and resource bounded measure. In *Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science*, volume 1046 of *Lecture Notes in Computer Science*, pages 13–24, Berlin, 1996. Springer-Verlag.

[BM95] H. Buhrman and E. Mayordomo. An excursion to the Kolmogorov random strings. In *Proceedings of the 10th IEEE Structure in Complexity Theory Conference*, pages 197–205, New York, 1995. IEEE.

[Hel86] H. Heller. On relativized exponential and probabilistic complexity classes. *Information and Compuation*, 71:231–243, 1986.

[JL95] D. Juedes and J. Lutz. The complexity and distribution of hard problems. *SIAM Journal on Computing*, 24(2):279–295, 1995.

[LM94] J. Lutz and E. Mayordomo. Measure, stochasticity, and the density of hard languages. *SIAM Journal on Computing*, 23:762–779, 1994.

[LM96] J. Lutz and E. Mayordomo. Cook versus Karp–Levin: Separating completeness notions if NP is not small. *Theoretical Computer Science*, 164:141–163, 1996.

[Lut90] J. Lutz. Category and measure in complexity classes. *SIAM Journal on Computing*, 19(6):1100–1131, 1990.

[Lut92] J. Lutz. Almost everywhere high nonuniform complexity. *Journal of Computer and System Sciences*, 44:220–258, 1992.

[Lut95] J. Lutz. A small span theorem for P/poly-Turing reductions. In *Proceedings of the 10th IEEE Structure in Complexity Theory Conference*, pages 324–330, New York, 1995. IEEE.

[Lut96] J. Lutz. Observations on measure and lowness for Δ_2^P. In *Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science*, volume 1046 of *Lecture Notes in Computer Science*, pages 87–97, Berlin, 1996. Springer-Verlag.

[May94] E. Mayordomo. Almost every set in exponential time is P-bi-immune. *Theoretical Computer Science*, 136:487–506, 1994.

[NW94] N. Nisan and A. Wigderson. Hardness vs. randomness. *Journal of Computer and System Sciences*, 49:149–167, 1994.

[Pap94] Ch. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[RSC95] K. Regan, D. Sivakumar, and J. Cai. Pseudorandom generators, measure theory, and natural proofs. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 26–35, New York, 1995. IEEE.