Journal of Computer and System Sciences 59, 327–345 (1999)
Article ID jcss.1999.1650, available online at http://www.idealibrary.com on IDEAL®

Hard Sets Are Hard to Find[1]

Harry Buhrman[2]

CWI, 1090 GB, Amsterdam, The Netherlands
E-mail: buhrman@cwi.nl

and

Dieter van Melkebeek[3]

Department of Computer Science, University of Chicago, Chicago, Illinois 60637
E-mail: dieter@cs.uchicago.edu

Received August 13, 1998; revised April 30, 1999

We investigate the frequency of complete sets for various complexity classes within $\mathcal{E}\mathcal{X}\mathcal{P}$ under several polynomial-time reductions in the sense of resource-bounded measure. We show that these sets are scarce: The sets that are $\leqslant^{\mathrm{P}}_{n^{\alpha}-\mathrm{tt}}$-complete for $\mathcal{N}\mathcal{P}$, the levels of the polynomial-time hierarchy, and $\mathcal{PSPACE}$ have p_2-measure zero for any constant $\alpha < 1$; The $\leqslant^{\mathrm{P}}_{n^{c}-\mathrm{T}}$-complete sets for $\mathcal{E}\mathcal{X}\mathcal{P}$ have p_2-measure zero for any constant c; Assuming $\mathcal{MA} \neq \mathcal{E}\mathcal{X}\mathcal{P}$, the $\leqslant^{\mathrm{P}}_{\mathrm{tt}}$-complete sets for $\mathcal{E}\mathcal{X}\mathcal{P}$ have p-measure zero. A key ingredient is the Small Span Theorem, which states that for any set A in $\mathcal{E}\mathcal{X}\mathcal{P}$ 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 best published theorem along these lines held for $\leqslant^{\mathrm{P}}_{\mathrm{btt}}$-reductions. We establish it for $\leqslant^{\mathrm{P}}_{n^{o(1)}-\mathrm{tt}}$-reductions.

1. INTRODUCTION

Lutz introduced resource-bounded measure [16] to formalize the notions of scarceness and abundance in complexity theory. His approach makes it possible to

[1] A preliminary version of this paper was presented at the *13th IEEE Conference on Computational Complexity* [11].

[2] 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 8556, and HC & M Grant ERB-4050-PL-93-0516.

[3] Partially supported by the European Union through Marie Curie Research Training Grant ERB-4001-GT-96-0783, by the U.S. National Science Foundation through Grant CCR 92-53582, and by the Fields Institute. Most of the research was done while visiting CWI and the University of Amsterdam.

0022-0000/99 $30.00

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 [14, 20, 22, 1, 21, 25, 19, 2].

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 $\mathcal{NP}$, the levels of the polynomial-time hierarchy, $\mathcal{PSPACE}$, and $\mathcal{EXP}$, and give some evidence that they have p_2-measure zero. On the other hand, the results of Bennett and Gill [8] imply that the $\leqslant_{\mathrm{tt}}^{\mathrm{P}}$-hard sets for $\mathcal{BPP}$ do not have p_2-measure zero; Allender and Strauss [1] even showed they have p_2-measure 1 in $\mathcal{EXP}$.

We use three different approaches to obtain our results. Two of them yield unhypothesized statements on the border of what is provable by relativizable techniques. First, we significantly improve the Small Span Theorem of Juedes and Lutz [14]. The Small Span Theorem for a reducibility $\leqslant_r^{\mathrm{P}}$ states that for any set A in $\mathcal{EXP}$, either the class of sets that $\leqslant_r^{\mathrm{P}}$-reduce to A (called the lower span of A), or the class of sets that A $\leqslant_r^{\mathrm{P}}$-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 spans, it implies that every $\leqslant_r^{\mathrm{P}}$-degree has p_2-measure zero, and in particular the $\leqslant_r^{\mathrm{P}}$-complete degree of any complexity class within $\mathcal{EXP}$. The strongest Small Span Theorem previous to our work was due to Ambos-Spies, Neis, and Terwijn [4], who proved it for $\leqslant_{\mathrm{btt}}^{\mathrm{P}}$-reductions. The extension to reductions with a non-constant number of queries was a notorious open problem in the area. We establish the Small Span Theorem for $\leqslant_{n^{o(1)}\text{-tt}}^{\mathrm{P}}$-reductions, i.e., for nonadaptive reductions that make a subpolynomial number of queries. Longpré [15] informed us that he obtained a Small Span Theorem for $\leqslant_{\log^{o(1)}n\text{-tt}}^{\mathrm{P}}$-reductions at the end of 1995 using the compressibility method [9].

Lutz [18] obtained a Small Span Theorem for nonuniform reductions w.r.t. *pspace*-measure. Similar to his proof, our Small Span Theorem follows from the fact that most sets in $\mathcal{EXP}$ have a $\leqslant_{n^{o(1)}\text{-tt}}^{\mathrm{P}}$-upper span with p_2-measure zero. We actually establish this fact for $\leqslant_{n^{\alpha}-\mathrm{tt}}^{\mathrm{P}}$-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 $\leqslant_{n^{\alpha}-\mathrm{tt}}^{\mathrm{P}}$-degree within $\mathcal{EXP}$ has p_2-measure zero. Previously, it was only known for $\leqslant_{\mathrm{btt}}^{\mathrm{P}}$-reductions that the p_2-measure of the complete sets for $\mathcal{EXP}$ have p_2-measure zero [4, 10]. We also obtain that the p_2-measure of the $\leqslant_{n^{\alpha}-\mathrm{tt}}^{\mathrm{P}}$-hard sets for $\mathcal{E}$ and $\mathcal{EXP}$ is zero.

Then we take a look at $\mathcal{EXP}$, in particular, and use an ad hoc technique to improve the results of the first approach for this particular case. We show that the $\leqslant_{n^{c}-\mathrm{T}}^{\mathrm{P}}$-complete sets for $\mathcal{EXP}$ have p_2-measure zero for any constant c. Our proofs relativize and are on the edge of the scope of relativizable techniques: Showing the last theorem for unbounded growing exponent c would separate $\mathcal{BPP}$ from $\mathcal{EXP}$.

Therefore, we next look at what we can show under a nonrelativizing reasonable, but yet unproven, complexity theoretic hypothesis, namely the assumption that $\mathcal{MA} \neq \mathcal{EXP}$. Babai, Fortnow, Nisan, and Wigderson [5] established the existence

of a pseudo-random generator that can be used to simulate $\mathcal{BPP}$ in subexponential time for infinitely many input lengths, unless $\mathcal{MA} = \mathcal{EXP}$. Using this pseudo-random generator, Buhrman, Van Melkebeek, Regan, Sivakumar, and Strauss [12] showed that the class of $\leq_{tt}^{p}$-complete sets for each of the Δ-levels of the polynomial-time hierarchy has p-measure zero, unless $\mathcal{EXP} = \mathcal{MA}$. Combining our second approach with theirs and some new ingredients, we are able to prove that the complete sets for $\mathcal{EXP}$ under $\leq_{T}^{p}$-reductions that make their queries in lexicographic order, have p-measure zero unless $\mathcal{EXP} = \mathcal{MA}$. In particular, the $\leq_{tt}^{p}$-complete sets for $\mathcal{EXP}$ have p-measure zero unless $\mathcal{EXP} = \mathcal{MA}$.

Summarizing our results:

- We prove a Small Span Theorem for $\leq_{n^{o(1)}\text{-tt}}^{p}$-reductions.
- We show that the $\leq_{n^{\alpha}-\text{tt}}^{p}$-complete sets for $\Pi_2\mathcal{P}$, the levels of the polynomial-time hierarchy, and $\mathcal{PSPACE}$ have p_2-measure zero for any $\alpha < 1$.
- We show that the $\leq_{n^{\alpha}-\text{tt}}^{p}$-hard sets for $\mathcal{E}$ and $\mathcal{EXP}$ have p_2-measure zero for any $\alpha < 1$.
- We show that the $\leq_{n^{c}-\text{T}}^{p}$-complete sets for $\mathcal{EXP}$ have p_2-measure zero for any constant c.
- We show that the $\leq_{tt}^{p}$-complete sets for $\mathcal{EXP}$ have p-measure zero unless $\mathcal{MA} = \mathcal{EXP}$ (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 for arbitrary subclasses of $\mathcal{EXP}$. Then we discuss our results particular to $\mathcal{EXP}$. Section 4 contains those without any complexity theoretic assumption; Section 5 contains those using the hypothesis $\mathcal{MA} \neq \mathcal{EXP}$. 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 [7, 6], and by Papadimitriou [24].

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 we write $A \leq_{T}^{p} B$ ("T" for Turing). The reduction M is *nonadaptive* if the oracle queries M makes on any input are independent of the oracle. In that case we write $A \leq_{tt}^{p} B$ ("tt" for truth-table). If, in addition, the number of queries on an input of length n is bounded by $q(n)$, we write $A \leq_{q(n)\text{-tt}}^{p} B$. For a reducibility $\leq_{r}^{p}$, we define the *lower span* of a set A as $\mathcal{P}_r(A) = \{B \mid B \leq_{r}^{p} A\}$, and the *upper span* of A as $\mathcal{P}_r^{-1}(A) = \{B \mid A \leq_{r}^{p} B\}$. The $\leq_{r}^{p}$-*degree* of A equals $\mathcal{P}_r(A) \cap \mathcal{P}_r^{-1}(A)$.

An *autoreduction* M is a reduction that never queries its own input; i.e., for any input x and any oracle B, M^B with input x does not query x. A set A is *autoreducible* if there is an autoreduction of A to itself.

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

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

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 $\mathscr{C}$ of sequences if it succeeds on every sequence in $\mathscr{C}$.

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 loses 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 2.2. *A class $\mathscr{C} \subseteq \Sigma^\infty$ 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 2.3 [17]. 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 $\mathscr{C} \subseteq \Sigma^\infty$ 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 $\mu_p(\mathscr{C}) = 0$ (resp. $\mu_{p_2}(\mathscr{C}) = 0$).

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 2.4 [17]. *Let d_i be a p-uniform (resp. p_2-uniform) system of supermartingales such that d_i covers the class $\mathscr{C}_i$. Then $\bigcup_i \mathscr{C}_i$ 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 $\chi_A = A(s_0)\, A(s_1)\, A(s_2) \cdots$, 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 $\mu_p(\mathscr{E}) \neq 0$ and $\mu_{p_2}(\mathscr{E}\mathscr{NP}) \neq 0$ [17].

2.2. Background on Pseudo-Random Generators

DEFINITION 2.5 [23]. 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

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

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 $\mathscr{E}$ that stretches seeds superpolynomially and has superpolynomial security at infinitely many lengths. We will use the one provided by the following theorem.

THEOREM 2.6. *If $\mathscr{NA} \neq \mathscr{EXP}$, there is a pseudo-random generator G computable in $\mathscr{E}$ 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 [5], and Nisan and Wigderson [23], combined with some padding.

THEOREM 2.7 [5]. *If $\mathscr{NA} \neq \mathscr{EXP}$, there is a set $A \in \mathscr{EXP}$ such that for any integer k, $H_A(n) \geqslant n^k$ for infinitely many n.*

THEOREM 2.8 [23]. *Given any set $A \in \mathscr{EXP}$, there is a pseudo-random generator G computable in $\mathscr{EXP}$ with $r(n) \in n^{\theta(\log n)}$ such that $S_G(n) \in \Omega(H_A(\sqrt{n})/n)$.*

3. COMPLETE SETS UNDER NON-ADAPTIVE REDUCTIONS WITH n^α QUERIES AND A SMALL SPAN THEOREM

In this section, we establish our results on the measure of complete and hard sets for complexity classes within $\mathscr{EXP}$. The following theorem forms the main ingredient. It states that most sets in $\mathscr{EXP}$ have a small upper span under $\leqslant^{\mathrm{P}}_{n^\alpha\text{-tt}}$-reductions for constant $\alpha < 1$. Later we also show a strong connection with the Small Span Theorem.

THEOREM 3.1. *For any $\alpha < 1$,*

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

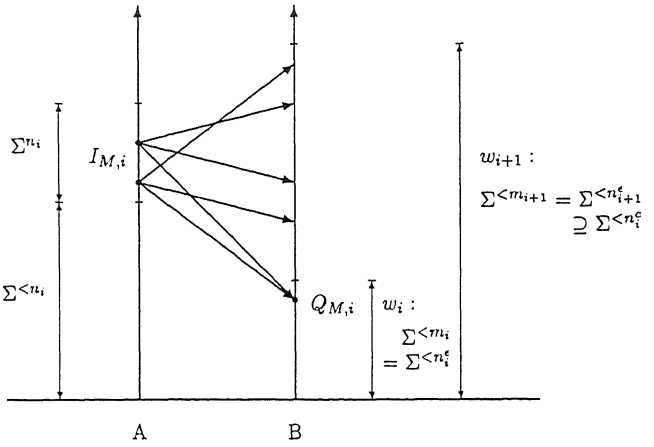

FIG. 1. Betting strategies at stage i.

We first give an outline of the proof. Fix a $\leqslant^{\mathrm{p}}_{n^2-n}$-reduction M running in time n^c for some constant $c>0$, and a set $A \in \mathcal{EXP}$. 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 Fig. 1. We can select an input x of length $n_i = m_i^{1+\varepsilon}$ for some constant $\varepsilon > 0$ 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 better. In order to be able to bet on the sets $A \in \mathcal{EXP}$ 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 $n_i^\chi + 1$ 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 $A \in \mathcal{EXP}$ 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., we 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 $n_i^\chi + 1$ 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 $2^{n_i^\chi}$ strategies, and at least one of them will realize a relative gain of $2^{|I_{M,i}|} = 2^{n_i^\chi + 1} = 2 \cdot 2^{n_i^\chi}$. So, if we do this at every stage with $\frac{2}{3}$ of the capital available at the

beginning of that stage, and leave the other $\frac{1}{3}$ intact, we succeed on A; at almost all stages, we increase our capital with a factor of $\frac{2}{3} \cdot 2 = \frac{4}{3}$, and at the finitely many other stages, we do not lose all of it.

We define the stages as follows:

$$\begin{aligned} m_0 &= 1 \\ m_{i+1} &= 2^{m_i} \\ n_i &= m_i^{1/\varepsilon}. \end{aligned} \tag{2}$$

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 $\leqslant^{\mathrm{P}}_{n^\alpha - \mathrm{tt}}$-reduction M, and that we can construct them efficiently. Here we need the fact that $\alpha < 1$.

LEMMA 3.2. *Let $\alpha < 1$, $\varepsilon \in (0, 1-\alpha)$, 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 $\leqslant^{\mathrm{P}}_{n^\alpha - \mathrm{tt}}$-reduction M, there is a set of strings $Q_{M,i}$ such that*

$$|\{x \in \Sigma^{n_i} \mid Q_M(x) \cap \Sigma^{<m_i} = Q_{M,i}\}| \geqslant n_i^\alpha + 1,$$

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 $|I_{M,i}| = n_i^\alpha + 1$ in time 2^{2n_i}.

Proof of Lemma 3.2. 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

$$\sum_{i=0}^{n_i^\alpha} \binom{2^{m_i}-1}{i} \leqslant (2^{m_i})^{n_i^\alpha} = 2^{n_i^{\alpha+\varepsilon}} \leqslant \frac{2^{n_i}}{n_i^\alpha + 1}, \tag{3}$$

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 3.1. We use the notation from Lemma 3.2. Fix $A \in \mathrm{DTIME}[2^{n^k}]$. Let

$$\pi_{A,M} = \begin{cases} 1, & \text{if } |w| < 2^{m_{i_0}}, \\ \Pr_{\omega \sqsupseteq w}[\forall x \in I_{M,i} : M^\omega(x) = A(x)], & \text{if } 2^{m_{i_0}} \leqslant 2^{m_i} \leqslant |w| < 2^{m_{i+1}}. \end{cases} \tag{4}$$

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

$$d_{A,M}(\lambda) = 1,$$

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

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 $|I_{M,i}| \cdot n_i^{\alpha} \in O((\log |w|)^{2\alpha/\varepsilon})$ positions of ω not fixed by w. It follows that $\pi_{A,M}$ and $d_{A,M}$ can be computed in time $2^{(\log |w|)^{O((\alpha+k)/\varepsilon)}}$.

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,

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

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

$$\mu_{p_2}(\{B \mid M \text{ reduces } A \text{ to } B\}) = 0. \tag{5}$$

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 $b \in \Sigma^{|Q_{M,i}|}$, let $\delta_{M,i,b}$ 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 jth element of $Q_{M,i}$ equals the jth bit of b, and all other bits are, say, 0. Ties are broken arbitrarily. The martingale

$$\delta_{M,i}(w) = \frac{1}{2^{|Q_{M,i}|}} \Sigma_b \delta_{M,i,b}(w)$$

has initial capital 1 and is computable in time $O(|w|^2)$. It has the property that

$$\delta_{M,i}(\chi_A|_{\Sigma^{<n_{i+1}}}) \geq \frac{2^{|I_{M,i}|}}{2^{|Q_{M,i}|}} \geq 2 = 2\delta_{M,i}(\chi_A|_{\Sigma^{<n_i}}),$$

provided i satisfies the Case 2 condition. Since almost all i's do, the following p-martingale δ_M succeeds on A: During stage i, it uses $\delta_{M,i}$ as a strategy on $\frac{2}{3}$ of the capital it has at the beginning of stage i, and does nothing with the other $\frac{1}{3}$.

Fix an enumeration M_j of all $\leq^{\mathrm{P}}_{n^\alpha\text{-tt}}$-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 $\leq^{\mathrm{P}}_{n^\alpha\text{-tt}}$-reduction M. Consider any set $A \in \mathcal{E}.\mathcal{NP}$ not covered by δ. Since the martingale system d_{A,M_j} is p_2-uniform, Eq. (5) implies that the p_2-measure of $\mathscr{P}^{-1}_{n^\alpha\text{-tt}}(A)$ is zero. ∎

Luc Longpré noticed that Theorem 3.1 also holds for $\leq^{\mathrm{P}}_{n^\alpha\text{-T}}$-reductions that make their queries in lexicographical order. It actually suffices that the queries are made in length nondecreasing order.

THEOREM 3.3. *Let $\leq^{\mathrm{P}}_r$ denote the reducibility by polynomial-time Turing machines that query no more than n^α strings on inputs of length n for some constant $\alpha < 1$, and make these queries in length nondecreasing order. Then,*

$$\mu_p(\{A \in \mathcal{E}.\mathcal{NP} \mid \mu_{p_2}(\mathscr{P}_r^{-1}(A)) \neq 0\}) = 0.$$

Proof Sketch. We can extend Lemma 3.2 as follows.

LEMMA 3.4. *Let $\alpha < 1$, $\varepsilon \in (0, 1-\alpha)$, and m_i and n_i be defined by (2). There is an integer i_0 such that for any $i \geq i_0$, for any $\leq^{\mathrm{P}}_r$-reduction M, and for any $b \in \Sigma^{n_i^\alpha}$, there is a set of strings $Q_{M,i,b}$ such that*

$$|\{x \in \Sigma^{n_i} \mid Q^{<m_i}_{M,b}(x) = Q_{M,i,b}\}| \geq n_i^\alpha + 1,$$

where $Q^{<m_i}_{M,b}(x)$ denotes the set of queries of length less than m_i which M makes on input x when the jth-bit of b is given as the answer to the jth query of length less than m_i. Moreover, we can find the lexicographically first set $Q_{M,i,b}$ and the lexicographically first subset $I_{M,i,b}$ of

$$\{x \in \Sigma^{n_i} \mid Q^{<m_i}_{M,b}(x) = Q_{M,i,b}\}$$

with $|I_{M,i,b}| = n_i^\alpha + 1$ in time 2^{2n_i}.

Note that $Q^{<m_i}_{M,b}(x)$ in Lemma 3.4 is well defined, because the queries of length less than m_i which M^ω makes on input x only depend on the prefix of ω of length $2^{m_i}-1$, since M^ω makes its queries in length nondecreasing order. More specifically,

these queries only depend on the part of the prefix that specifies the answers to them, i.e., on b.

The betting strategy for $\mathscr{P}_r^{-1}(A)$ is the same as in Theorem 3.1, except that we use the set $I_{M,i,b}$ of Lemma 3.4, instead of the set $I_{M,b}$ of Lemma 3.2 in formula (4), where b is determined by the prefix of w of length $2^{m_i}-1$.

The martingale $\delta_{M,i}$ is the average over several strategies. Now there is one strategy $\delta_{M,i,b}$ corresponding to every $b \in \Sigma^{n_i^\alpha}$, namely one with initial capital 1 that only bets on strings of $I_{M,i,b}$. On such a string $x \in I_{M,i,b}$, it bets all of its money according to the majority of $M^\omega(x)$ over all sequences $\omega \sqsupseteq v_{i,b}$, where $v_{i,b}$ is the characteristic string of length $2^{m_i}-1$ in which the bit corresponding to the jth query of M on input x equals the jth bit of b, and all other bits are say 0.

The rest of the construction and the analysis are essentially the same as in the proof of Theorem 3.1. ∎

Our results on the measure of complete sets follow directly from Theorem 3.1. By Theorem 3.3, they also hold for the more general reducibility introduced in Theorem 3.3.

COROLLARY 3.5. *For any $\alpha < 1$ and $C \in \mathscr{E}.\mathscr{NP}$, the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-degree of C has p_2-measure zero. In particular, the classes of $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-complete sets for $\mathscr{NP}$, the levels of the polynomial-time hierarchy, $\mathscr{P}/\mathscr{Poly}$, and $\mathscr{E}.\mathscr{NP}$ all have p_2-measure zero.*

Proof. Suppose not, then for any set A in the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-degree of C, the p_2-measure of $\mathscr{P}_{n^\alpha\text{-tt}}^{-1}(A)$ is not zero, since it contains the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-degree of C. But, by Theorem 3.1 this would imply that the p-measure of the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-degree of C is zero. ∎

For the class of $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-hard sets, we get

COROLLARY 3.6. *For any $\alpha < 1$ and any complexity class $\mathscr{C}$ such that $\mu_p(\mathscr{C} \cap \mathscr{E}.\mathscr{NP}) \neq 0$, the class of $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-hard sets for $\mathscr{C}$ has p_2-measure zero. In particular, the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-hard sets for $\mathscr{E}$ and $\mathscr{E}.\mathscr{NP}$ have p_2-measure zero.*

Proof. By definition, for any set $A \in \mathscr{C}$, the $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-hard sets for $\mathscr{C}$ are contained in $\mathscr{P}_{n^\alpha\text{-tt}}^{-1}(A)$. If the class of $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$-hard sets for $\mathscr{C}$ does not have p_2-measure zero, Theorem 3.1 yields that $\mu_p(\mathscr{C} \cap \mathscr{E}.\mathscr{NP}) = 0$. ∎

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

From Theorem 3.1, we can also deduce a Small Span Theorem. However, we have to settle for a more restrictive reducibility than $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$, because we need transitivity in the proof, and $\leqslant_{n^\alpha\text{-tt}}^{\mathrm{P}}$ 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 $\varepsilon > 0$. We write $A \leqslant_{n^{o(1)}\text{-tt}}^{\mathrm{P}} B$ if there exists a subpolynomial function $f(n)$ such that $A \leqslant_{f(n)\text{-tt}}^{\mathrm{P}} B$.

THEOREM 3.7 (Small Span Theorem). *For any set A, at least one of the following holds: $\mu_p(\mathscr{P}_{n^{o(1)}\text{-tt}}(A) \cap \mathscr{E}.\mathscr{NP}) = 0$ or $\mu_{p_2}(\mathscr{P}_{n^{o(1)}\text{-tt}}^{-1}(A)) = 0$.*

Proof. We distinguish between two cases:

- $\mathscr{P}_{n^{o(1)}\text{-tt}}(A)$ contains a set B such that $\mu_{p_2}(\mathscr{P}^{-1}_{n^{o(1)}\text{-tt}}(B)) = 0$. Then the transitivity of $\leqslant^{\mathrm{p}}_{n^{o(1)}\text{-tt}}$ and the monotonicity of p_2-measure imply that $\mu_{p_2}(\mathscr{P}^{-1}_{n^{o(1)}\text{-tt}}(A)) = 0$.
- $\mathscr{P}_{n^{o(1)}\text{-tt}}(A) \cap \mathscr{E}.\mathscr{X}.\mathscr{P}$ is included in $\{B \in \mathscr{E}.\mathscr{X}.\mathscr{P} \mid \mu_{p_2}(\mathscr{P}^{-1}_{n^{\alpha}\text{-tt}}(B)) \neq 0\}$ for any $\alpha > 0$. Then Theorem 3.1 says that $\mu_p(\mathscr{P}_{n^{o(1)}\text{-tt}}(A) \cap \mathscr{E}.\mathscr{X}.\mathscr{P}) = 0$. ∎

For any set $A \in \mathscr{E}.\mathscr{X}.\mathscr{P}$, Theorem 3.7 states that at least one of its lower span or upper span under $\leqslant^{\mathrm{p}}_{n^{o(1)}\text{-tt}}$-reductions is small.

4. COMPLETE SETS FOR $\mathscr{E}.\mathscr{X}.\mathscr{P}$ UNDER ADAPTIVE REDUCTIONS WITH n^c QUERIES

We now show how, in the case of $\mathscr{E}.\mathscr{X}.\mathscr{P}$, we can extend the results of the previous section on the measure of complete sets from $\leqslant^{\mathrm{p}}_{n^{\alpha}\text{-tt}}$-reductions for any $\alpha < 1$ to $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-reductions for any constant c:

THEOREM 4.1. *For any constant c, the class of $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-complete sets for $\mathscr{E}.\mathscr{X}.\mathscr{P}$ has p_2-measure zero.*

The proof technique differs significantly. We exploit the diagonalization power of $\mathscr{E}.\mathscr{X}.\mathscr{P}$ against $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-reductions to show that all $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-complete sets for $\mathscr{E}.\mathscr{X}.\mathscr{P}$ share a structural property that allows the construction of a p_2-martingale succeeding on all of them. We first establish the structural property.

Let $M_1, M_2, \ldots$ be an enumeration of $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-reductions, where M_i runs in time n^i.

LEMMA 4.2. *For any constant c and for any $\leqslant^{\mathrm{p}}_{n^c\text{-T}}$-complete set C for $\mathscr{E}.\mathscr{X}.\mathscr{P}$, there is an index j such that*

$$\forall n, \forall x \in \Sigma^n : M_j^C(\langle 0^j, x\rangle) = \text{minority}_{\omega \supseteq \chi_C|_{\Sigma^{<n}}}[M_j^{\omega}(\langle 0^j, x\rangle)]. \tag{6}$$

The right-hand side of (6) denotes the least probable value of $M_j^{\omega}(\langle 0^j, x\rangle)$ when ω is uniformly distributed over all extensions of the initial segment of χ_C corresponding to all strings of length up to n. Ties are broken in some fixed way, say always 0.

Proof of Lemma 4.2. Let

$$D = \{\langle 0^i, x\rangle \mid \Pr_{\omega \supseteq \chi_C|_{\Sigma^{<|x|}}}[M_i^{\omega}(\langle 0^i, x\rangle) = 1] < \tfrac{1}{2}\}.$$

The above probability is a weighted sum of the accepting leaves of the reduction tree of M_i on input $\langle 0^i, x\rangle$. The weight of a leaf is only nonzero if on its path P all queries of length less than $|x|$ are answered consistent with C, and in that case its weight equals $2^{-q(P)}$, where $q(P)$ denotes the number of other queries made along P. W.l.o.g. we are assuming here that on no path the reduction asks the same query more than once. So, we can decide D on instances $\langle 0^i, x\rangle$ of length n in time

$2^{n^c}(n^c \cdot \text{time}_C(n) + n^j)$. Since $C \in \mathcal{E}.\mathcal{NP}$, this implies $D \in \mathcal{E}.\mathcal{NP}$, and since C is $\leqslant^{\mathrm{P}}_{n^c-\mathrm{T}}$-hard for $\mathcal{E}.\mathcal{NP}$, that there is a $\leqslant^{\mathrm{P}}_{n^c-\mathrm{T}}$-reduction M_j reducing D to C. The index j satisfies (6), because for any $x \in \Sigma^n$,

$$\begin{aligned} M_j^C(\langle 0^j, x\rangle) = 1 &\Leftrightarrow \langle 0^j, x\rangle \in D \\ &\Leftrightarrow \Pr_{\omega \supseteq \chi_C|_{\Sigma^{<n}}}[M_j^\omega(\langle 0^j, x\rangle) = 1] < \tfrac{1}{2} \\ &\Leftrightarrow \text{minority}_{\omega \supseteq \chi_C|_{\Sigma^{<n}}}[M_j^\omega(\langle 0^j, x\rangle)] = 1. \quad \blacksquare \end{aligned}$$

Lemma 4.2 provides a consistency test that eliminates at least half of the remaining possibilities. We now use it in a straightforward way to construct a p_2-martingale covering all $\leqslant^{\mathrm{P}}_{n^c-\mathrm{T}}$-complete sets for $\mathcal{E}.\mathcal{NP}$.

Proof of Theorem 4.1. For any index j, we construct a (uniform) p_2-martingale d_j that succeeds on any set C for which (6) holds. The martingale d_j has initial capital 1, and works in stages defined by

$$n_1 = 1$$

$$n_{i+1} = (n_i + j)^j.$$

The ith stage starts when the martingale has to bet on the string 0^{n_i}. Let w_i denote the prefix seen up to that moment. During stage i, d_j distributes $2^{2^{n_{i+1}}-2^{n_i}} \cdot d_j(w_i)$ uniformly over all extensions w'_{i+1} of w_i with $|w'_{i+1}| = 2^{n_{i+1}} - 1$ for which $M_j^{w'_{i+1}}(\langle 0^j, 0^{n_i}\rangle) = \text{minority}_{\omega \supseteq w_i}[M_j^\omega(\langle 0^j, 0^{n_i}\rangle)]$.

Note that for any set C satisfying (6), d_j at least doubles its capital along C at every stage, so it succeeds on any such C. Therefore, by Lemma 4.2, the martingale system $(d_j)_{j=1}^\infty$ covers the class of $\leqslant^{\mathrm{P}}_{n^c-\mathrm{T}}$-complete sets for $\mathcal{E}.\mathcal{NP}$.

Using the approach of Lemma 4.2, we can compute the minority and the probabilities underlying $d_j(w)$ in time $O(2^{(\log|w|+j)^c}(\log|w|+j)^j)$. So, the martingale system $(d_j)_{j=1}^\infty$ is p_2-uniform. $\blacksquare$

In an analogous way, we get the following theorem for $\mathcal{E}$.

THEOREM 4.3. *For any constant c, the class of $\leqslant^{\mathrm{P}}_{cn-\mathrm{T}}$-complete sets for $\mathcal{E}$ has p-measure zero.*

Ambos-Spies informed us recently that he and Lempp have a new proof of Theorems 4.1 and 4.3 [3].

5. COMPLETE SETS FOR $\mathcal{E}.\mathcal{NP}$ UNDER ADAPTIVE REDUCTIONS

Theorem 4.1 cannot be improved using relativizable techniques, since it fails for unbounded growing exponent c in a world where $\mathcal{BPP} = \mathcal{E}.\mathcal{NP}$ and such a world exists [13]. This follows from the relativizable result of Allender and Strauss [1] that the class of sets that are not $\leqslant^{\mathrm{P}}_{\mathrm{T}}$-hard for BPP has p-measure zero. In this section, we will see what results we can get on the measure of the $\mathcal{E}.\mathcal{NP}$-complete

sets for polynomial-time reductions without an explicit bound on the number of queries, under the likely but unrelativizing hypothesis $\mathcal{NA} \neq \mathcal{EXP}$. We obtain

THEOREM 5.1. *The class of sets complete for $\mathcal{EXP}$ (or $\mathcal{E}$) under $\leq_T^P$-reductions that make their queries in lexicographical order, has p-measure zero unless $\mathcal{EXP} = \mathcal{NA}$. In particular, the class of $\leq_{tt}^P$-complete sets for $\mathcal{EXP}$ (or $\mathcal{E}$) has p-measure zero unless $\mathcal{EXP} = \mathcal{NA}$.*

Buhrman, Van Melkebeek, Regan, Sivakumar, and Strauss [12] used the hypothesis $\mathcal{NA} \neq \mathcal{EXP}$ to show that the class of autoreducible sets under the same type of reductions has p-measure zero. We will use the same idea, namely applying pseudo-random generators to approximate efficiently the probabilities underlying the martingales constructed in the previous section, and that way mimic their behavior by an easier-to-compute martingale. The pseudo-random generators whose existence is known to follow from the assumption $\mathcal{NA} \neq \mathcal{EXP}$ by Theorem 2.6, have superpolynomial security at infinitely many lengths. They will allow us to approximate the underlying probabilities well enough, but only at infinitely many lengths. Therefore, in order for the mimicking martingale to succeed, we will make sure we make a lot of money on these lengths. We will use the following lemma instead of Lemma 4.1 to do so.

LEMMA 5.2. *Fix a pseudo-random generator computable in time 2^{an} for some constant $a > 1$, and with stretching $r(n)$. There is an oracle Turing machine T running in time 2^{2an} with the following property: For any set C complete for $\mathcal{EXP}$ under $\leq_T^P$-reductions that make their queries in lexicographic order, there is an index j of such a reduction M_j such that for any string x,*

$$\Pr_{\omega \sqsupseteq \chi_C|_{\Sigma^{<n}}}[\forall i \in I_n : M_j^{\omega}(\langle 0^j, x, 0^i \rangle) = T^{C \cap \Sigma^{<n}}(\langle 0^j, x, 0^i \rangle)] \leq \frac{2}{n^3}$$
$$\forall i \in I_n : M_j^{C}(\langle 0^j, x, 0^i \rangle) = T^{C \cap \Sigma^{<n}}(\langle 0^j, x, 0^i \rangle), \tag{7}$$

where $n = |x|$ and $I_n = \{1, 2, \ldots, 3 \log n\}$, provided $r(n)$, $S_G(n) \geq n^{j+1}$ and n is sufficiently large.

Lemma 5.2 also holds if we substitute "length nondecreasing" for "lexicographic."

Proof of Lemma 5.2. Consider an input $x \in \Sigma^n$, a prefix $w \in \Sigma^{2^n - 1}$, a string $b \in \Sigma^{3 \log n}$, and an index j such that M_j makes its queries in length nondecreasing order. Recall that M_j runs in time n^j. We can compute the probability

$$\pi_j(x, w, b) = \Pr_{\omega \sqsupseteq w}[\forall i \in I_n : M_j^{\omega}(\langle 0^j, x, 0^i \rangle) = b_i]$$

as the fraction of strings $\beta \in \Sigma^{n^{j+1}}$ such that the predicate underlying π_j holds when the oracle queries of length less than n are answered according to w, and the kth different query of length at least n is answered as β_k. The predicate depends on $o(n^{j+1})$ bits of the prefix w in total, because the queries of length less than n made by M are the same for any β. It follows that the test circuit has size n^{j+1} for sufficiently large n. Therefore, we can approximate $\pi_j(x, w, b)$ to within an additive

term of $1/n^4$ using the pseudo-random generator G at length n, provided $r(n) \geq n^{j+1}$ and $S_G(n) \geq n^{j+1}$.

On input $\langle 0^j, x, 0^i \rangle$, the machine T^w will compute these approximations $\tilde{\pi}_j(x, w, b)$ to $\pi_j(x, w, b)$ for every $b \in \Sigma^{3 \log n}$, select the lexicographically first value $\tilde{b}$ for b that minimizes $\tilde{\pi}_j(x, w, b)$, and output the ith bit of $\tilde{b}$. T can do this in time 2^{2an}.

Note that there is a setting $b^* \in \Sigma^{3 \log n}$ such that $\pi_j(x, w, b^*) \leq 1/n^3$. Inductively set b_i^* such that at least half of the extensions $\omega \sqsupseteq w$ satisfying $M_j^{\omega}(\langle 0^j, x, 0^k \rangle) = b_k^*$, for $1 \leq k < i$, fail the test $M_j^{\omega}(\langle 0^j, x, 0^i \rangle) = b_i^*$. Therefore,

$$\begin{aligned}
\pi_j(x, w, \tilde{b}) &\leq \tilde{\pi}_j(x, w, \tilde{b}) + \frac{1}{n^4} \\
&\leq \tilde{\pi}_j(x, w, b^*) + \frac{1}{n^4} \\
&\leq \tilde{\pi}_j(x, w, b^*) + \frac{2}{n^4} \\
&\leq \frac{1}{n^3} + \frac{2}{n^4} \\
&\leq \frac{2}{n^3},
\end{aligned}$$

which establishes the first part of (7) for any set C.

Now fix a set C complete for $\mathcal{EXP}$ under $\leq_T^p$-reductions that make their queries in lexicographic order, and consider the set

$$D = \{\langle 0^j, x, 0^i \rangle \mid 1 \leq i \leq 3 \log |x| \text{ and } T^{C \cap \Sigma^{<|x|}}(\langle 0^j, x, 0^i \rangle) \text{ accepts}\}.$$

Since $C \in \mathcal{EXP}$, we can also decide D in $\mathcal{EXP}$, and since C is hard for $\mathcal{EXP}$ under $\leq_T^p$-reductions that make their queries in lexicographic order, there is such a reduction M_j reducing D to C. This establishes the second part of (7). ∎

Lemma 5.2 gives a consistency test that eliminates a fraction at least $1 - (2/n^3)$ of the possibilities, and therefore multiplies the capital by a factor of $n^3/2$. For Lemma 4.2 these figures are $\frac{1}{2}$ and 2, respectively. We will now see how we can exploit the larger increase in capital to construct a p-martingale that succeeds on the complete sets for $\mathcal{EXP}$ under $\leq_T^p$-reductions that make their queries in lexicographical order, using the above pseudo-random generator once more.

Proof of Theorem 5.1 *for* $\mathcal{EXP}$. Fix a $\leq_T^p$-reduction M_j running in time n^j that makes its queries in lexicographical order. Let T be the oracle Turing machine given by Lemma 5.2 based on the pseudo-random generator G that follows from the hypothesis $\mathcal{NE} \neq \mathcal{EXP}$ by Theorem 2.6.

Let

$$\pi_{j,m}(w) = \Pr_{\omega \sqsupseteq w}[\forall i \in I_m : M_j^{\omega}(\langle 0^j, 0^m, 0^i \rangle) = T^{w \cap \Sigma^{<m}}(\langle 0^j, 0^m, 0^i \rangle)],$$

and consider

$$d_{j,m}(w) = \begin{cases} m^3 \cdot \pi_{j,m}(w), & \text{if } |w| \geqslant 2^m, \\ 2, & \text{otherwise.} \end{cases}$$

The function $d_{j,m}(w)$ is computable in time $2^{O(\log^{j+1}|w|)}$, and so is $d_j(w) = \sum_{m=1}^{\infty} (1/m^2)\, d_{j,m}(w)$. They are nonnegative and satisfy the supermartingale inequality (1) for all strings w, except possibly for those of length $2^m - 1$. In case of a set C satisfying (7) for $x = 0^m$, the inequality also holds for $w \sqsubseteq \chi_C$ of length $2^m - 1$. Moreover, $d_{j,m}(\chi_C) = m^3$, and $d_j(\chi_C) = \infty$.

We now want to construct (super)martingales $\tilde{d}_{j,m}$ and $\tilde{d}_j$ that behave like $d_{j,m}$ and d_j along χ_C and are computable uniformly in time $|w|^a$ for some constant a, i.e., independent of the running time of M_j. The key idea is to approximate efficiently the probability $\pi_{j,m}$ using the pseudo-random generator G as we did in the proof of Lemma 5.2. Following that approach for some constant a_1, we can compute in time $|w|^{a_1}$ an approximation $\tilde{\pi}_{j,m}(w)$ of $\pi_{j,m}(w)$ to within $\varepsilon_{j,m} = m^{-(j+4)}$, provided $r(m) \geqslant m^{j+1}$ and $S_G(m) \geqslant m^{j+4}$. By Theorem 2.6 (assuming $\mathcal{NP} \neq \mathcal{BPP}$), infinitely many m satisfy the latter conditions; we call such m's good.

There are still two technical problems we have to solve in order to make sure that $\tilde{d}_{j,m}$ is a supermartingale: First, what to do along sets C for which (7) does not hold for $x = 0^m$, and what if m is not good? We will deal with that in a moment. Second, even for a good m along a set C satisfying (7) for $x = 0^m$, just replacing $\pi_{j,m}$ with $\tilde{\pi}_{j,m}$ in the definition of $d_{j,m}$ might not work. For example, if $\tilde{\pi}_{j,m}$ underestimates $\pi_{j,m}$ 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 a query M_j^{ω} makes on some input of the form $\langle 0^j, 0^m, 0^i \rangle$ for some $i \in I_m$ and some $\omega \sqsupseteq w$. As the queries are made in lexicographical order, we can efficiently check the latter condition on w by running M_j^w on every input $\langle 0^j, 0^m, 0^i \rangle$ for $i \in I_m$, and there can be no more than $3m^j \log m$ prefixes w satisfying it along any sequence ω. Since the limit $\varepsilon_{j,m}$ on the estimation error is such that $(3m^j \log m) \cdot \varepsilon_{j,m}$ remains bounded, we can remedy this problem by accumulatively subtracting a term $2\varepsilon_{j,m}$ from the approximation for $\pi_{j,m}$, and adding a constant to the resulting approximation for $d_{j,m}$. The former modification guarantees that condition (1) is met; the latter is needed after the former in order to keep the values nonnegative. More precisely, we define

$$d^*_{j,m}(w) = \begin{cases} m^3 \tilde{\pi}_{j,m}(w) + 1 - 2q_{j,m}(w)\, m^3 \varepsilon_{j,m}, & \text{if } |w| \geqslant 2^m, \\ 4, & \text{otherwise,} \end{cases} \tag{8}$$

where $q_{j,m}(w)$ denotes the number of positions in w that correspond to a query M_j^w makes on an input of the form $\langle 0^j, 0^m, 0^i \rangle$ for some $i \in I_m$. Note that $0 \leqslant q_{j,m}(w) \leqslant q_{j,m}(\omega) \leqslant 3m^j \log m$ and that we can efficiently compute $q_{j,m}(w)$.

We solve the first problem by explicitly checking for each prefix w that the values $d^*_{j,m}$ 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 $\tilde{d}_{j,m}$ as

$$\tilde{d}_{j,m}(\lambda) = 4$$

$$\tilde{d}_{j,m}(wb) = \begin{cases} d^*_{j,m}(wb), & \text{if } d^*_{j,m}(w0) \geqslant 0 \text{ and } d^*_{j,m}(w1) \geqslant 0 \text{ and} \\ & \quad d^*_{j,m}(w0) + d^*_{j,m}(w1) \leqslant 2\tilde{d}_{j,m}(w) \\ \tilde{d}_{j,m}(w), & \text{otherwise.} \end{cases} \tag{9}$$

It follows that $\tilde{d}_{j,m}$ is a supermartingale computable in time $|w|^{a_2}$ for some constant a_2 independent of M_j and m.

CLAIM 5.3. *If m is good and sufficiently large, $\tilde{d}_{j,m}(w) = d^*_{j,m}(w)$ for any $w \sqsubseteq \chi_C$, where C is a set satisfying (7).*

Proof of Claim 5.3. We show that $\tilde{d}_{j,m}(w) = d^*_{j,m}(w)$ for any $w \sqsubseteq \chi_C$ by induction on $|w|$. Clearly, the statement holds for $w = \lambda$. So, it suffices to argue for any string w that the conditions on the right-hand side of (9) are met, assuming that $\tilde{d}_{j,m}(w) = d^*_{j,m}(w)$.

If $|w| < 2^m - 1$, this is true because $d^*_{j,m}(v) = 4$ for $|v| < 2^m$. If $|w| \geqslant 2^m - 1$, the first two conditions on the right-hand side of (9) are satisfied, since for any string v of length $|v| \geqslant 2^m$,

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

which is positive for sufficiently large m. In case $|w| = 2^m - 1$, the remaining condition is met, because

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

In case $|w| \geqslant 2^m$, the remaining condition certainly holds if $d^*_{j,m}(w0) = d^*_{j,m}(w1) = d^*_{j,m}(w)$. Otherwise, $q_{j,m}(w0) = q_{j,m}(w1) = q_{j,m}(w) + 1$, and we have that

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

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

$$\tilde{d}_{j,m}(\chi_C) = d^*_{j,m}(\chi_C) \geqslant d_{j,m}(\chi_C) + 1 - (2q_{j,m}(\omega) + 1)\, m^3 \varepsilon_{j,m} \geqslant d_{j,m}(\chi_C) \qquad (10)$$

for any set C satisfying (7). Since there are infinitely many good m's and $d_{j,m}(\chi_C) = m^3$, this implies that $\tilde{d}_j = \sum_{m=1}^{\infty} (1/m^2)\, \tilde{d}_{j,m}$ is a supermartingale that succeeds on any such set C. It is computable in time $|w|^a$ for some constant a independent of j.

Since for a standard enumeration M_i, including all $\leqslant_T^p$-reductions that make their queries in lexicographical order and such that $M_i(x)$ is computable in time $(2^{|x|} + i)^{O(1)}$, the supermartingale system $\tilde{d}_i$ is p-uniform, Lemma 5.2 finishes the proof of the theorem. ∎

6. DISCUSSION AND OPEN PROBLEMS

The question of whether Theorem 3.1 holds for some constant $\alpha \geqslant 1$, remains open. A positive answer would be the best result provable by relativizable techniques, just as our results in Section 4 are optimal. By the same token, relativizable techniques cannot establish the Small Span Theorem for $\leqslant_{tt}^p$-reductions.

It seems unlikely that our approach allows one to establish Theorem 3.1 for $\alpha \geqslant 1$, because of Lemma 3.2. For some constant $\varepsilon > 0$ and a given $\leqslant_{n^\alpha - tt}^p$-reduction M, this would require the construction of a set $I_{M,i}$ containing $n_i^\alpha + 1$ strings of length n_i and a set $Q_{M,i}$ of size n_i^α, 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 $\Sigma^{<n_i^\varepsilon}$ in a Kolmogorov random way. The concatenation σ of all these queries is a Kolmogorov random string of length $2^{n_i} n_i^{\alpha+\varepsilon}$. 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 $qn_i^\varepsilon + q(n_i + n_i^\alpha \log q) + (2^{n_i} - q)\, n_i^{\alpha+\varepsilon} + O(\log q)$, which is asymptotically less than $|\sigma|$, as long as $\log q \leqslant c n_i^\varepsilon$ 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.

Ambos-Spies, Neis, and Terwijn [4] focused on p-measure, and they established the equivalent of Theorem 3.1 and the Small Span Theorem within $\mathscr{E}$ for $\leqslant_{k\text{-}tt}^p$-reductions for any constant k. A similar Kolmogorov argument as above indicates that our techniques are not powerful enough to extend these results to stronger reductions. Even the $\leqslant_{btt}^p$-case remains open.

ACKNOWLEDGMENTS

We are grateful to Jack Lutz and Lance Fortnow for very helpful discussions regarding the Small Span Theorem. We would like to thank Luc Longpré for the extension of Theorem 3.1 he suggested, Klaus Ambos-Spies for explaining his work in the area, Leen Torenvliet for suggesting the title of the paper, and the anonymous referees for their comments.

REFERENCES

1. 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," pp. 807–818, IEEE Press, New York, 1994.
2. K. Ambos-Spies and L. Bentzien, Separating NP-completeness notions under strong hypothesis, *in* "Proceedings of the 12th IEEE Conference on Computational Complexity," pp. 121–127, IEEE Press, New York, 1997.
3. K. Ambos-Spies, S. Lempp, and G. Mainhardt, Randomness vs. completeness: On the diagonalization strength of resource-bounded random sets, manuscript, March 1998.
4. K. Ambos-Spies, H.-C. Neis, and S. Terwijn, Genericity and measure for exponential time, *Theoret. Comput. Sci.* **168**, No. 1 (1996), 3–19.
5. L. Babai, L. Fortnow, N. Nisan, and A. Wigderson, BPP has subexponential time simulations unless EXPTIME has publishable proofs, *Comput. Complexity* **3** (1993), 307–318.
6. J. Balcázar, J. Díaz, and J. Gabarró, "Structural Complexity II," EATCS Monographs on Theoretical Computer Science, Vol. 22, Springer-Verlag, New York/Berlin, 1990.
7. J. Balcázar, J. Díaz, and J. Gabarró, "Structural Complexity I," EATCS Monographs on Theoretical Computer Science, Vol. 11, Springer-Verlag, New York/Berlin, 1995.
8. C. Bennett and J. Gill, Relative to a random oracle, $P^A \neq NP^A \neq co-NP^A$ with probability one, *SIAM J. Comput.* **10** (1981), 96–113.
9. H. Buhrman and L. Longpré, Compressibility and resource bounded measure, *in* "Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science," pp. 13–24, Lecture Notes in Computer Science, Vol. 1046, Springer-Verlag, New York/Berlin, 1996.
10. H. Buhrman and E. Mayordomo, An excursion to the Kolmogorov random strings, *in* "Proceedings of the 10th IEEE Structure in Complexity Theory Conference," pp. 197–205, IEEE Press, New York, 1995.
11. H. Buhrman and D. van Melkebeek, Hard sets are hard to find, *in* "Proceedings of the 13th IEEE Conference on Computational Complexity," pp. 170–180, IEEE Press, New York, 1998.
12. H. Buhrman, D. van Melkebeek, K. Regan, D. Sivakumar, and M. Strauss, A generalization of resource-bounded measure, with an application, *in* "Proceedings of the 15th Symposium on Theoretical Aspects of Computer Science," pp. 161–171, Lecture Notes in Computer Science, Vol. 1373, Springer-Verlag, New York/Berlin, 1998.
13. H. Heller, On relativized exponential and probabilistic complexity classes, *Inform. and Comput.* **71** (1986), 231–243.
14. D. Juedes and J. Lutz, The complexity and distribution of hard problems, *SIAM J. Comput.* **24**, No. 2 (1995), 279–295.
15. L. Longpré, personal communication, 1997.
16. J. Lutz, Category and measure in complexity classes, *SIAM J. Comput.* **19**, No. 6 (1990), 1100–1131.
17. J. Lutz, Almost everywhere high nonuniform complexity, *J. Comput. System Sci.* **44** (1992), 220–258.
18. J. Lutz, A small span theorem for P/poly-Turing reductions, *in* "Proceedings of the 10th IEEE Structure in Complexity Theory Conference," pp. 324–330, IEEE Press, New York, 1995.
19. J. Lutz, Observations on measure and lowness for Δ_2^P, *in* "Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science," pp. 87–97, Lecture Notes in Computer Science, Vol. 1046, Springer-Verlag, New York/Berlin, 1996.
20. J. Lutz and E. Mayordomo, Measure, stochasticity, and the density of hard languages, *SIAM J. Comput.* **23** (1994), 762–779.
21. J. Lutz and E. Mayordomo, Cook versus Karp–Levin: Separating completeness notions if NP is not small, *Theoret. Comp. Sci.* **164** (1996), 141–163.
22. E. Mayordomo, Almost every set in exponential time is P-bi-immune, *Theoret. Comput. Sci.* **136** (1994), 487–506.

23. N. Nisan and A. Wigderson, Hardness vs. randomness, *J. Comput. System Sci.* **49** (1994), 149–167.

24. C. Papadimitriou, "Computational Complexity," Addison–Wesley, Reading, MA, 1994.

25. 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," pp. 26–35, IEEE Press, New York, 1995.