The Zero-One Law Holds for BPP

Dieter van Melkebeek*
DIMACS
Rutgers University
96 Frelinghuysen Road
Piscataway, NJ 08854-8018
USA

Abstract

We show that BPP has p-measure zero if and only if BPP differs from EXP. The same holds when we replace BPP by any complexity class $\mathcal{C}$ that is contained in BPP and is closed under tt-reductions. The zero-one law for each of these classes $\mathcal{C}$ follows: Within EXP, $\mathcal{C}$ has either measure zero or else measure one.

Keywords: theory of computation, computational complexity, resource-bounded measure, zero-one law, derandomization.

1 Introduction

The zero-one law of classical Lebesgue measure states that any reasonable class of languages that is closed under finite variation has measure zero or one. It is an open question whether the zero-one law carries over to resource-bounded measure as developed by Lutz [6]. In particular, we do not know whether every complexity class within exponential time has measure zero or one.

On one hand, Regan, Sivakumar and Cai [10] showed that if strong pseudo-random generators exist then the zero-one law fails for the class of languages with polynomial-size circuits. No unconditional counterexamples are known. On the other hand, no nontrivial positive examples were known either. We establish the first one, namely for the class BPP. It remains open whether BPP has measure zero or one within exponential time but we show that one of the two must hold.

By the measure conservation property [6], BPP having p-measure zero implies that BPP $\neq$ EXP. Conversely, Buhrman, Fenner and Fortnow [4] showed that if MA $\neq$ EXP, then BPP has p-measure zero. This state of affairs left open the possibility that BPP differs from EXP, but does not have p-measure zero either. We exclude that possibility.

Using Impagliazzo and Wigderson's work on pseudo-randomness secure against uniform adversaries [5], we are able to bridge the remaining gap between BPP and MA, and show that the weaker hypothesis BPP $\neq$ EXP already implies that BPP has p-measure zero. So, establishing the latter is equivalent to separating BPP from EXP:

Theorem 1.1 BPP *has p-measure zero iff* BPP $\neq$ EXP.

*Research done while supported by the Fields Institute, Toronto. Email: dieter@dimacs.rutgers.edu. Web: http://dimacs.rutgers.edu/~dieter.

In the terminology of Lutz and Mayordomo [8], the main result of Impagliazzo and Wigderson [5] implies that if $\mathsf{BPP} \neq \mathsf{EXP}$ then no language in BPP is weakly $(2^n, 0, 2^n)$-stochastic. Theorem 1.1 thus follows from Lutz and Mayordomo's theorem that the class of languages that are not weakly $(2^n, 0, 2^n)$-stochastic has p-measure zero [7, Theorem 10.1]. We will give an elementary proof of Theorem 1.1 that does not use the general theorem of Lutz and Mayordomo on weak stochasticity.

The zero-one law for BPP follows from Theorem 1.1.

Corollary 1.2 BPP *either has p-measure zero or else has p-measure one.*

Note that Corollary 1.2 is in fact as strong as Theorem 1.1, since Regan et al. [10] showed that BPP having p-measure one implies that $\mathsf{BPP} \cap \mathsf{E} = \mathsf{E}$ and therefore that $\mathsf{BPP} = \mathsf{EXP}$.

Allender and Strauss observed that their work [1] allows us to generalize Theorem 1.1 and Corollary 1.2 to any class (instead of BPP) contained in BPP and closed under tt-reductions.

Theorem 1.3 *Let $\mathcal{C}$ be any class contained in* BPP *and closed under tt-reductions. Then $\mathcal{C}$ has p-measure zero iff $\mathcal{C} \neq \mathsf{EXP}$.*

Corollary 1.4 *Let $\mathcal{C}$ be any class contained in* BPP *and closed under tt-reductions. Then $\mathcal{C}$ either has p-measure zero or else has p-measure one.*

For the same reason as above, Corollary 1.4 is as powerful as Theorem 1.3.

2 Notation and Preliminaries

Most of our complexity theoretic notation is standard. We refer the reader to the textbooks by Balcázar, Díaz and Gabarró [3, 2], and by Papadimitriou [9].

BPP denotes the complexity class of languages that can be decided by probabilistic polynomial-time Turing machines with bounded two-sided error. EXP denotes $\cup_{c>0}\mathsf{DTIME}[2^{n^c}]$ and E denotes $\cup_{c>0}\mathsf{DTIME}[2^{cn}]$. EXP contains both BPP and E.

A *reduction* of a language A to a language B is a polynomial-time oracle Turing machine M such that $M^B = A$. The reduction M is *non-adaptive* if the oracle queries M makes on any input are independent of the oracle. A non-adaptive reduction is called a *truth-table* reduction or tt-reduction for short. A language B is *hard* for a complexity class $\mathcal{C}$ if every language $A \in \mathcal{C}$ reduces to B. If the same holds for tt-reductions we call B *tt-hard* for $\mathcal{C}$.

The rest of this section will give the necessary background on resource-bounded measure. For our purposes, we only need to specify what it means to have p-measure zero or one. Lutz [6] defined having p-measure zero using the concept of p-martingales.

Definition 2.1 *A* martingale *is a function $d : \{0,1\}^* \to [0, \infty)$ satisfying*

$$d(w) = \frac{d(w0) + d(w1)}{2}$$

for every $w \in \{0,1\}^*$. If $d(w)$ is computable in time polynomial in $|w|$, d is called a p-martingale. *A martingale d* succeeds *on a sequence $\omega \in \{0,1\}^\infty$ if $d(\omega) \doteq \limsup_{w \sqsubseteq \omega, w \to \omega} d(w) = \infty$.*

We will view a martingale as 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 part 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 large during the game.

Definition 2.2 *A complexity class $\mathcal{C}$ has* p-measure zero *if there is a single p-martingale d that succeeds on the characteristic sequence of every language in $\mathcal{C}$. In that case, we say that d* covers *$\mathcal{C}$.*

The property of having p-measure zero is monotone, i.e., every subclass of a p-measure zero class also has p-measure zero. Lutz [6] showed that E does not have p-measure zero, which he called the *measure conservation property*. It follows that EXP does not have p-measure zero either. The property of having p-measure zero is also closed under finite union, i.e., if two classes have p-measure zero then so has their union.

These properties guarantee that the following definition of having p-measure one is consistent in the sense that no class can have both p-measure zero and p-measure one.

Definition 2.3 *A complexity class $\mathcal{C}$ has* p-measure one *if the complement of $\mathcal{C}$ has p-measure zero.*

We will use the next theorem by Lutz [6] about countable unions of p-measure zero classes:

Theorem 2.4 *Suppose d_i, $i \in \mathbb{N}$, is a martingale covering class $\mathcal{C}_i$. If $d(i,w) \doteq d_i(w)$ is computable in time polynomial in $i + |w|$, then $\cup_i \mathcal{C}_i$ has p-measure zero.*

3 Main Result

In this section we will prove Theorem 1.1.

The left-to-right implication follows immediately from the measure conservation property. To establish the right-to-left implication, we will use the next result by Impagliazzo and Wigderson [5].

Theorem 3.1 *If $\mathsf{BPP} \neq \mathsf{EXP}$, then for every language $A \in \mathsf{BPP}$ and all constants $\epsilon > 0$ and $d > 0$, there is a language $B \in \mathsf{DTIME}[2^{n^\epsilon}]$ such that the following holds: For any length-preserving probabilistic polynomial-time Turing machine M,*

$$\Pr[A(M(1^m)) = B(M(1^m))] > 1 - m^{-d} \tag{1}$$

for infinitely many lengths m, where the probability is over the internal coin tosses of M.

We will apply Theorem 3.1 with $\epsilon = 1$, d any constant, and $M(1^m)$ the uniform distribution on strings of length m. In that case, for sufficiently large m, (1) implies

$$\Pr_{|x|=m}[A(x) = B(x)] > \frac{2}{3}, \tag{2}$$

where the probability is with respect to the uniform distribution over $\{0,1\}^m$.

Fix a set $B \in \mathsf{DTIME}[2^n]$. For any integer $m > 0$, consider the strategy for the game of Section 2 that only bets on the membership bits of strings of length m, and for each of these strings puts $\frac{1}{3}$ of its capital on the outcome that the membership is the same as for B. The corresponding martingale $d_{B,m}(w)$, say with initial capital $d_{B,m}(\lambda) = 1$, is identically 1 for $|w| < 2^{m+1}$, and constant $d_{B,m}(\omega)$ for $w \sqsubset \omega$ with $|w| \geqslant 2^{m+2}$. Computing $d_{B,m}(w)$ for $2^{m+1} \leqslant |w| < 2^{m+2}$ essentially amounts to deciding whether $x \in B$ for strings x of length m in lexicographic order. Since there are 2^m strings of length m, and for each of them we can decide membership to B in time 2^m, it follows that we can compute $d_{B,m}(w)$ in time $|w|^2 \log^{O(1)} |w|$.

Note that any correct bet on the membership of a string of length m, increases the capital by a factor of $1 + \frac{1}{3} = \frac{4}{3}$, whereas an incorrect one reduces it by a factor of $1 - \frac{1}{3} = \frac{2}{3}$. If ω is the characteristic sequence of a language A for which (2) holds, more than $\frac{2}{3}$ of the bets on strings of length m are correct. Therefore,

$$d_{B,m}(\omega) > \left(\frac{4}{3}\right)^{\frac{2}{3}2^m} \left(\frac{2}{3}\right)^{\frac{1}{3}2^m} d_{B,m}(\lambda) = \left(\frac{2^{\frac{5}{3}}}{3}\right)^{2^m} \in 2^{\Omega(2^m)}.$$

Now consider

$$d_B(w) \doteq \sum_{m=1}^{\infty} \frac{1}{m^2} d_{B,m}(w). \tag{3}$$

Since $\sum_{m=1}^{\infty} \frac{1}{m^2}$ converges, d_B is a well-defined martingale. Because $d_{B,m}(w) = 1$ for $m \geqslant \log |w|$, computing $d_B(w)$ really only requires evaluating the first $\log |w|$ terms of the right-hand side of (3), which we can do in time $|w|^2 \log^{O(1)} |w|$. Moreover, the above argument shows that $d_B(\omega) = \infty$ if ω is the characteristic sequence of any language A for which (2) holds for infinitely many m. Let $\mathcal{C}_B$ denote the class of such languages A.

Let $B_1, B_2, \ldots$ be a standard enumeration of $\mathsf{DTIME}[2^n]$ obtained by clocking Turing machines. Then there is a fixed Turing machine deciding $x \in B_i$ in time polynomial in $i + 2^{|x|}$. By the previous analysis, this implies that $d(i, w) \doteq d_{B_i}(w)$ is computable in time polynomial in $i + |w|$. Therefore, by Theorem 2.4, $\mathcal{C} \doteq \cup_i \mathcal{C}_{B_i}$ has p-measure zero. By Theorem 3.1, the assumption $\mathsf{BPP} \neq \mathsf{EXP}$ implies that $\mathcal{C}$ contains BPP. So, BPP has p-measure zero if $\mathsf{BPP} \neq \mathsf{EXP}$.

4 Generalization

In this section, we establish Theorem 1.3.

Assume that $\mathcal{C}$ satisfies the conditions of the theorem and that $\mathcal{C}$ does not have measure zero. Allender and Strauss [1] showed that the p-measure of the class of languages that are not hard for BPP is zero. Their proof also works for tt-reductions, i.e., even the class of languages that are not tt-hard for BPP has p-measure zero. It follows that $\mathcal{C}$ contains a language that is tt-hard for BPP. Since $\mathcal{C}$ is closed under tt-reductions, this implies that $\mathsf{BPP} \subseteq \mathcal{C}$. On the other hand, by hypothesis $\mathcal{C} \subseteq \mathsf{BPP}$. Therefore, $\mathcal{C} = \mathsf{BPP}$, and Theorem 1.1 yields that $\mathcal{C} = \mathsf{EXP}$, since $\mathcal{C}$ does not have measure zero. This finishes the proof.

Acknowledgments

I am grateful to Eric Allender and Martin Strauss for allowing me to include the generalizations given by Theorem 1.3 and Corollary 1.4 in this paper. I would like to thank Jack Lutz for comments on an earlier version, and Lance Fortnow for suggesting the title of the paper. Thanks also to an anonymous referee for pointing out an additional link to the literature, and to Wolfgang Merkle for pointing out a typo in the abstract of an earlier version.

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*, pages 807–818. IEEE, 1994.

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

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

[4] H. Buhrman, S. Fenner, and L. Fortnow. Results on resource-bounded measure. In *Proceedings of the 24th International Colloquium On Automata, Languages and Programming*, volume 1256 of *Lecture Notes in Computer Science*, pages 188–194. Springer-Verlag, 1997.

[5] R. Impagliazzo and A. Wigderson. Randomness vs. time: De-randomization under a uniform assumption. In *Proceedings of the 39th IEEE Symposium on Foundations of Computer Science*, pages 734–743. IEEE, 1998.

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

[7] J. Lutz. The quantitative structure of exponential time. In L. Hemaspaandra and A. Selman, editors, *Complexity Theory Retrospective II*, pages 225–260. Springer-Verlag, 1997.

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

[9] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[10] 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. IEEE, 1995.