

Structural Average Case Complexity*

Rainer Schuler[†]
Abteilung Theoretische Informatik
Universität Ulm
Oberer Eselsberg
D-89069 Ulm, Germany

Tomoyuki Yamakami[‡]
Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 1A4

Abstract

Levin introduced an average-case complexity measure, based on a notion of “polynomial on μ -average,” and defined “average-case polynomial-time many-one reducibility” among randomized decision problems. We generalize his notions of average-case complexity classes, Random-NP and Average-P. Ben-David et al. use the notation of $\langle \mathcal{C}, \mathcal{F} \rangle$ to denote the set of randomized decision problems (L, μ) such that L is a set in \mathcal{C} and μ is a probability density function in \mathcal{F} . This paper introduces $\text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$ as the class of randomized decision problems (L, μ) such that L is computed by a type- \mathcal{C} machine on μ -average and μ is a density function in \mathcal{F} . These notations capture all known average-case complexity classes as, for example, $\text{Random-NP} = \langle \text{NP}, \text{P-comp} \rangle$ and $\text{Average-P} = \text{Aver}\langle \text{P}, * \rangle$, where P-comp denotes the set of density functions whose distributions are computable in polynomial time, and $*$ denotes the set of all density functions. Mainly studied are polynomial-time reductions between randomized decision problems: many-one, deterministic Turing and nondeterministic Turing reductions and the average-case versions of them. Based on these reducibilities, structural properties of average-case complexity classes are discussed. We give average-case analogues of concepts in worst-case complexity theory; in particular, the polynomial-time hierarchy and Turing self-reducibility, and we show that all known complete sets for Random-NP are Turing self-reducible. A new notion of “real polynomial-time computations” is introduced based on average polynomial-time computations for arbitrary distributions from a fixed set, and it is used to characterize the worst-case complexity classes Δ_k^{P} and Σ_k^{P} of the polynomial-time hierarchy.

1 Introduction

The classical complexity theory of NP-completeness is based on the worst-case analysis of algorithms. A probabilistic analysis has been applied so far only to specific algorithms typically with respect to the uniform distribution for each length of inputs. Levin [22] gave a general framework to perform *average-case analysis* in a way that allows us to discuss many questions of worst-case complexity theory in a more general setting. The average-case analysis considers *randomized problems*, namely, pairs of a decision problem and a probability density (or distribution) function which assigns probabilities to instances.

*A preliminary version of this paper appeared in the Proceedings of the 12th Conference on Foundations of Software Technology and Theoretical Computer Science, New Delhi, December 1992, Lecture Notes in Computer Science, Vol.652 (1992), pp. 128–139.

[†]Supported by the Deutsche Forschungsgemeinschaft, grant No. Scho 302/2-1.

[‡]This work was done while the author was a guest at the Universität Ulm from April to August in 1991.

In [22] Levin implicitly defined the average-case complexity classes “Average-P” and “Random-NP” as analogues of the worst-case complexity classes P and NP, respectively. Roughly speaking, Average-P (called AP in [17, 36]) is the class of randomized problems which are, on the average, solvable in polynomial time for a given distribution on the input, and Random-NP (called DistNP in [14, 5, 30], RNP in [17] and DNP in [36]) is the class of problems in NP together with polynomial-time computable distribution on the inputs. A stimulating question of whether Average-P contains the class Random-NP or not, was raised by Levin in his very terse papers [22, 23]. Up to now, this question has not been solved.

In recent literature [22, 23, 14, 17, 35], many randomized decision problems have been shown to be polynomial-time complete for Random-NP with respect to several different reducibilities. Average-case analysis is very sensitive to the choice of distributions on the instances, since, for example, if a density function μ decreases faster than $2^{-|x|}$ with the length of the instances x , then all NP-complete problems are solvable in time polynomial on μ -average. Even if we stick to “reasonable” distributions, “fast on average” algorithms have been found even for natural NP-complete problems. For example, the *satisfiability problem*, SAT with a natural distribution [10], the *graph 3-colorability problem* with a natural distribution [39], and the *Hamiltonian circuit problem* with a natural distribution and edge probability 1/2 [8] are all in Average-P. In the decade since Levin’s fundamental papers, several results have revealed the significant role of distributions on the average-case complexity (see, e.g., [36]). However, it is still open whether every NP-complete problem is complete for Random-NP for some reasonable distribution. In this paper, our main interest lies in the analysis of structural properties of average-case complexity classes. We extend well-known notions of worst-case complexity theory such as time- and space-bounded computation, nondeterministic Turing reducibility, self-reducibility and relativization to average-case analysis. Furthermore, we consider sets of distribution functions and discuss the (maximal) average-case complexity of problems under every distribution chosen from the given set.

Ben-David et al. [5] introduced the notation $\langle \mathcal{C}, \mathcal{F} \rangle$ to discuss the average case complexity of decision problems in a general setting. The class $\langle \mathcal{C}, \mathcal{F} \rangle$ contains all randomized problems (L, μ) (recall that a randomized problem here is always a pair of a decision problem and a density function on the instances) such that L is in the complexity class \mathcal{C} and μ is in a class of density functions \mathcal{F} . Random-NP can simply be denoted as $\langle \text{NP}, pcomp \rangle$, where *pcomp* (polynomial-time computable) denotes the set of density functions whose distribution functions can be approximated by a deterministic Turing machine in polynomial time.

Schapire [30] has given a different characterization of Levin’s notion of “polynomial on μ -average”, which will be used in this paper. He has shown that a function g from strings to non-negative real numbers is *polynomial on μ -average*, in Levin’s sense, if and only if there exists a polynomial p such that $\text{Prob}_\mu[\{x \mid g(x) > p(|x| \cdot r)\}] < 1/r$ for any positive real number r . Schapire’s characterization is intuitive and can be easily generalized to the notion of “ f on μ -average” by replacing the polynomial p by a function f which is defined on non-negative real numbers.

This paper introduces the notation $\text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$ to denote the set of all problems (L, μ) such that μ is in \mathcal{F} , and L is recognized by a “type- \mathcal{C} on μ -average” machine. Here, \mathcal{C} denotes a time- or space-bounded complexity class, and a *type- \mathcal{C} on μ -average machine* is a Turing machine which respects, on μ -average, the time- or space-bound, respectively. By using this notation, for example, Average-P, the set of problems solvable in average polynomial time, can be denoted as $\text{Aver}\langle \text{P}, * \rangle$, where $*$ is the set of all density functions. Levin’s most intriguing open question is rephrased as whether $\langle \text{NP}, \text{P-comp} \rangle \subseteq \text{Aver}\langle \text{P}, * \rangle$ holds or not.

A central concept in this paper are different types of *polynomial-time reductions* which are all generalizations of the polynomial time many-one reducibility used by Levin. Roughly speaking,

a reduction function between two randomized problems reduces a set of strings to another set of strings and also satisfies the so-called *domination condition* between density functions, which ensures that likely instances are mapped to likely instances. Since Levin’s work, random many-one reducibility [5] and polynomial-time deterministic Turing reducibility [5, 19] have been defined and studied. All those reducibilities can be extended by allowing their many-one reduction functions or oracle reduction machines to be polynomial-time bounded on μ -average. We redefine many-one reduction, deterministic Turing reduction and their average-case extensions and study their properties.

In this paper, a *polynomial-time nondeterministic Turing reduction* between randomized decision problems is defined in a way that captures both deterministic Turing reducibility and random many-one reducibility. Let M be a nondeterministic oracle Turing machine, E a set and p a polynomial. By $Q(M, E, x, y)$, we denote the set of strings which are queried by M with oracle E on input x on computation path y , and $\text{Acc}(M, E, x)$ (resp. $\text{Rej}(M, E, x)$) denotes the set of (codes of) accepting (resp. rejecting) computation paths given by M with E on input x . We introduce the density function μ' that is induced from μ , M and E as: $\mu'(x, y) = \mu(x)/|\text{Acc}(M, E, x)|$ if y is an accepting path of M^E on input x ; $\mu'(x, y) = \mu(x)/|\text{Rej}(M, E, x)|$ if there is no accepting paths and y is a rejecting path; otherwise, $\mu'(x, y) = 0$. The machine M polynomial-time nondeterministic (on μ -average) Turing reduces (D, μ) to (E, ν) if $D = L(M^E)$, M^E is polynomial (on μ -average) time bounded, and there exists a density function ν , which polynomially (on μ -average) dominates μ' , such that $\nu(z) \geq \text{Prob}_\nu[\{(x, y) | z \in Q(M, E, x, y)\}]$ for all strings z .

A randomized decision problem (L, μ) is said to be $\langle \mathcal{C}, \mathcal{F} \rangle$ -complete if (L, μ) is in $\langle \mathcal{C}, \mathcal{F} \rangle$, and every problem in $\langle \mathcal{C}, \mathcal{F} \rangle$ is polynomial-time many-one reducible to (L, μ) . Levin has first proven that the *randomized tiling problem* is $\langle \text{NP}, \text{P-comp} \rangle$ -complete [22]. Another typical $\langle \text{NP}, \text{P-comp} \rangle$ -complete problem is the *randomized bounded halting problem* [17]. Both problems are NP-complete problems together with some natural distribution. However, if (L, μ) is $\langle \text{NP}, \text{P-comp} \rangle$ -complete, then L must be NP-complete. As we noted before, it is not known whether every NP-complete set L has an appropriate density function μ which forces (L, μ) to be $\langle \text{NP}, \text{P-comp} \rangle$ -complete.

This paper extends the notion of average-case complexity further to discuss questions raised in structural complexity theory. A significant property of sets is the (*Turing*) *self-reducibility* [25]. In worst-case complexity theory, many known NP-complete sets are self-reducible. Analogously, we define a *self-reducibility* in average-case analysis. A randomized decision problem is self-reducible if it is polynomial-time Turing reducible to itself, while querying only strings of length smaller than the input. We prove the existence of $\langle \Sigma_k^{\text{P}}, \text{P-comp} \rangle$ -complete sets which are self-reducible by demonstrating that the k -th level of the randomized bounded halting problem $(\text{RH}^k, \mu_{\text{RH}})$ is $\langle \Sigma_k^{\text{P}}, \text{P-comp} \rangle$ -complete and also self-reducible. From the fact that most of the known $\langle \text{NP}, \text{P-comp} \rangle$ -complete problems are p-isomorphic [3], they turn out to be self-reducible.

Similar to the Meyer-Stockmeyer polynomial-time hierarchy [26] in worst-case complexity theory, we give a precise definition of its average-case analogue, the *average polynomial-time hierarchy*, founded on a relativization of $\text{Aver}\langle \text{P}, \mathcal{F} \rangle$ and $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle$ which is done in a similar fashion as we obtain relativized classes P^B and NP^B from P and NP , respectively. To be more precise, recall that P^B (resp. NP^B) is equivalent to the collection of sets which are polynomial-time deterministic (resp. nondeterministic) Turing reducible to B . Analogously, we define a class $\text{Aver}\langle \text{P}, \mathcal{F} \rangle^{(E, \nu)}$ (resp. $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}$) as a collection of randomized decision problems (D, μ) , with $\mu \in \mathcal{F}$, which are deterministic (resp. nondeterministic) polynomial-time on μ -average Turing reducible to (E, ν) .

We show that $\langle \text{NP}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{P}, \mathcal{F} \rangle$ and $\langle \text{NP}, \mathcal{F} \rangle \not\subseteq \text{Aver}\langle \text{P}, \mathcal{F} \rangle$ in some relativized worlds, using the technique of Baker et al. [1]. This implies that any proof technique which can be relativized will not solve Levin’s question.

Let $\text{Aver}\langle\Delta_k^{\text{P}}, \mathcal{F}\rangle$ denote the k -th level of the average polynomial-time hierarchy. We consider the general question of whether or not $\langle\Sigma_k^{\text{P}}, \text{P-comp}\rangle \subseteq \text{Aver}\langle\Delta_k^{\text{P}}, *\rangle$. This paper shows that this is not the case unless every tally set in Σ_k^{P} is in Δ_k^{P} . Hence, we believe that $\langle\Sigma_k^{\text{P}}, \text{P-comp}\rangle \not\subseteq \text{Aver}\langle\Delta_k^{\text{P}}, *\rangle$, since it seems very likely that there exist tally sets in $\Sigma_k^{\text{P}} - \Delta_k^{\text{P}}$.

Finally we discuss a tie between worst-case and average-case complexity classes. Note that every polynomial-time computation is polynomial on μ -average even if μ is chosen by some (powerful) adversary, that is, every set in P is computable by a deterministic Turing machine which is polynomial-time bounded on μ -average for every density function μ . In other words, if A is in P then $(A, \mu) \in \text{Aver}\langle\text{P}, *\rangle$ for every μ . On the other hand, P can be expected to be the largest class which satisfies this property. This is indeed true [24]. This observation can be generalized as follows. For a class \mathcal{C} , we define the *real \mathcal{C} over \mathcal{F}* , denoted by $\mathcal{C}_{\mathcal{F}}$, to be the class of all sets A such that $(A, \mu) \in \text{Aver}\langle\mathcal{C}, *\rangle$ for every density function μ in \mathcal{F} . Clearly $\text{P} \subseteq \text{P}_{\text{P-comp}}$, and Levin's question of whether $\langle\text{NP}, \text{P-comp}\rangle \subseteq \text{Aver}\langle\text{P}, *\rangle$ can now be stated as whether $\text{NP} \subseteq \text{P}_{\text{P-comp}}$. We show that $\Delta_k^{\text{P}} = \Delta_k^{\text{P}}_{\text{REC-comp}}$ and $\Sigma_k^{\text{P}} = \Sigma_k^{\text{P}}_{\text{REC-comp}}$ for all levels $k > 0$, where REC-comp denotes the collection of all recursive density functions. As a particular case, we can show that $\text{P} = \text{P}_{\text{E-comp}}$.

We show that, relative to random oracle, $\text{NP}_{\text{P-comp}}$ is different from $\text{P}_{\text{P-comp}}$ with probability 1.

2 Preliminaries

In this paper, we follow the standard definitions and notations of complexity theory; see, e.g., [11, 2]. Here we briefly present only necessary notations and notions.

Let $\Sigma = \{0, 1\}$ be fixed. By Σ^* we denote the set of all strings over Σ , and Σ^n denotes the set of all strings of length n . The set Σ^* admits the standard lexicographic order: $\epsilon < 0 < 1 < 00 < 01 < \dots$, where ϵ denotes the empty string. The successor and predecessor of x in this order are denoted by x^+ and x^- , respectively. A subset A of Σ^* is often identified with its characteristic function, i.e., $A(x) = 1$ if $x \in A$, and $A(x) = 0$ otherwise. The cardinality of a set A and the length of a string $x \in \Sigma^*$ are respectively denoted by $|A|$ and $|x|$. The *complement* of a set A is denoted by \overline{A} . For two sets A and B , let $A \triangle B$ denote the *symmetric difference* of A and B and let $A \oplus B$ be the *disjoint union* of A and B .

Let \mathbb{N} be the set of all *non-negative integers* and let \mathbb{R}^+ be the set of all *non-negative real numbers*. A *non-negative integer* can be identified with its binary representation, and we often refer to strings as non-negative integers. Especially, Σ^* is identified with the set $\mathbb{D} = \{m/2^n \mid m, n \in \mathbb{N}\}$ of *non-negative dyadic rational numbers*, i.e., a string $d_1d_2\dots d_n$ in Σ^* is identified with the number $d_12^{-1} + d_22^{-2} + \dots + d_n2^{-n}$ in \mathbb{D} .

A set is *tally* if it is a subset of $\{0\}^*$, and a set A is (*polynomially*) *sparse* if there is a polynomial p such that $|A \cap \Sigma^n| \leq p(n)$ for all $n \in \mathbb{N}$. Denote by TALLY the class of all tally sets and by SPARSE the class of all sparse sets.

A formal definition of *Turing machines* with semi-infinite tapes are given in, for example, [12, 11, 18, 2], and we assume the reader's familiarity of it. In this paper, we are interested in only *resource bounded* algorithms and assume that all Turing machines are designed in such a way that all computation paths have the same length. Therefore, for every oracle Turing machine M and a set A , the running *time* of the machine M with oracle A on input x , denoted by $\text{Time}_M^A(x)$, is simply defined to be the length of some possible computation, and the *space* complexity, $\text{Space}_M^A(x)$, is defined to be the maximum, over all configurations of M with oracle A on input x , of the number of tape-squares in use.

A (non-)deterministic oracle Turing machine *accepts* an input x , if there is an accepting computation of M with oracle A on input x ; otherwise, M *rejects* x . If M is probabilistic, then M

accepts x if $\text{Prob}_M[M \text{ on input } x \text{ halts in an accepting state}] > \frac{1}{2}$; otherwise, M rejects x , where $\text{Prob}_M[Q(M)]$ denotes the probability that $Q(M)$ holds.

Let $\text{Acc}(M, A, x)$ denote the set of (codes of) *accepting* computations of M on input x with oracle A , and similarly $\text{Rej}(M, A, x)$ denotes that of *rejecting* computation paths. Let $Q(M, A, x, y)$ be the set of strings queried by M with A on input x on computation path y . If M is deterministic, then we simply denote by $Q(M, A, x)$ the set of all strings queried by M on input x with oracle A .

As usual, $L(M, A)$ denotes the set of strings accepted by M with oracle A , and we simply say that M^A computes a set B if $B = L(M, A)$. For a machine M , $M^A(x)$ denotes the output of a computation of M on input x . For a deterministic Turing machine M (also called a *transducer*), M computes a function f if $f(x) = M^A(x)$ for all $x \in \Sigma^*$.

For any function t on \mathbb{N} , a Turing machine M with oracle A is *t-time bounded* (resp. *t-space bounded*) if $\text{Time}_M^A(x) \leq t(|x|)$ (resp. $\text{Space}_M^A(x) \leq t(|x|)$) for all x . Let $\text{DTIME}(t)$, $\text{NTIME}(t)$, and $\text{DSPACE}(t)$ denote the class of all sets computable by deterministic t -time bounded, non-deterministic t -time bounded, and deterministic t -space bounded Turing machines, respectively. Finally, let $\text{BPTIME}(t)$ denote the class of sets B computable by t -time bounded probabilistic Turing machines with bounded error probability, i.e., there is a constant ϵ , $0 \leq \epsilon < \frac{1}{2}$, such that $\text{Prob}_M[M(x) = B(x)] < 1 - \epsilon$ holds for all x .

We assume that the reader is familiar with standard notations of complexity classes, such as P , NP , Δ_k^P , Σ_k^P , PH (the polynomial-time hierarchy), BPP , PSPACE , E (linear exponential time), NE , EXP (exponential time), NEXP and FP (polynomial-time computable functions). Moreover, we let $\text{ESPACE} = \bigcup_{k>0} \text{DSPACE}(2^{kn+k})$ and $\text{BPE} = \bigcup_{k>0} \text{BPTIME}(2^{kn+k})$.

A function f on Σ^* is *p-honest* if there exists a polynomial p such that $p(|f(x)|) \geq |x|$ for any string x , and f is *p-invertible* if there is a function g in FP such that $g \circ f(x) = x$ for all x . A function f from \mathbb{N} to \mathbb{R}^+ is called *unbounded* if $\lim_{k \rightarrow \infty} \sup_{n>k} f(n) = \infty$.

A set D is *polynomial-time many-one reducible* to a set E if there exists a function f (in FP) such that, for all x , $x \in D$ if and only if $f(x) \in E$. A set D is *polynomial-time many-one complete* for a class \mathcal{C} if $D \in \mathcal{C}$ and every set in \mathcal{C} is polynomial-time many-one reducible to D . We simply say that D is *\mathcal{C} -complete* for a class \mathcal{C} if it is polynomial-time many-one complete for \mathcal{C} .

In the following, we use the following pairing function $\langle \cdot, \cdot \rangle$ [28], from $\Sigma^* \times \Sigma^*$ onto Σ^* , which is defined as follows: $\langle x, y \rangle = d(x)y$ if $|y| \leq 1$; otherwise $\langle x, y \rangle = d(x)i_2[(y^-)^-]$, where $d(\epsilon) = \epsilon$, $d(0x) = 00d(x)$, $d(1x) = 11d(x)$, $i_2[0x] = 01x$, and $i_2[1x] = 10x$ for all x . This pairing function is monotone, i.e., $x \leq x'$ and $y \leq y'$ imply $\langle x, y \rangle \leq \langle x', y' \rangle$, and it is computable in linear-time in the lengths of x and y . Moreover, it holds that $2|x| + |y| \leq |\langle x, y \rangle| \leq 2|x| + |y| + 1$ for all x and y . This pairing function is generalized recursively to a bijection from $(\Sigma^*)^k$ onto Σ^* as $\langle x_1, x_2, \dots, x_k \rangle = \langle x_1, \langle x_2, \dots, x_k \rangle \rangle$.

A function $\mu : \Sigma^* \rightarrow [0, 1]$ is a *density function* if $\sum_{x \in \Sigma^*} \mu(x) = 1$, and its corresponding *distribution* μ^* is given by $\mu^*(x) = \sum_{z \leq x} \mu(z)$. To avoid confusion, we remark here that density functions are also called “probability distributions” in [7, 3, 36] or “probability functions” in [17], and distributions are called “probability distributions” in [14, 17]. Let μ_n denote the *conditional density function* for strings of length n , i.e., $\mu_n(x) = \mu(x) / \sum_{y:|y|=n} \mu(y)$ whenever $\sum_{y:|y|=n} \mu(y) \neq 0$; otherwise 0. We use the notation $\mu(x, y)$ and $\mu(x, y, z)$ to denote $\mu(\langle x, y \rangle)$ and $\mu(\langle x, y, z \rangle)$, respectively.

For a finite domain D , the *uniform density function* on D is defined as $1/|D|$ for every x in D . The *standard density function* ν_{st} on Σ^* is defined as $\nu_{\text{st}}(x) = \frac{6}{\pi^2} \cdot (|x|+1)^{-2} \cdot 2^{-|x|}$ for all x . Although the standard distribution is called “uniform” in, e.g., [5], actually only its conditional distribution is uniform for all lengths n . We note that there are other ways to define a “standard” density function. See Gurevich [17] for a discussion. A density function μ is called *flat* if $\mu(x) \leq 2^{-|x|^c}$ for

some constant $\epsilon > 0$ [17]. A density function μ is *positive* if $\mu(x) > 0$ for all x .

For a density function μ , we use the notation $\text{Prob}_\mu[\{x \mid Q(x)\}]$ to denote the probability that property $Q(x)$ holds, where x is chosen randomly according to μ . For set A , let $\mu(A)$ denote $\sum_{x \in A} \mu(x)$.

In his papers [22, 23], Levin considers “polynomial-time computable” distributions as reasonable to discuss the average time-complexity of NP problems. Later a more generalized notion, i.e., “polynomial-time samplable” distributions, has been proposed [5]. This paper follows Gurevich [17] to define the notion of “polynomial-time computability” of distributions.

Definition 2.1 ([21, 17]) Let g be a distribution from Σ^* to $[0, 1]$ and f be a function on \mathbb{N} .

1. g is *f-time computable* (resp. *f-space computable*), if there exists a deterministic f -time (resp. f -space) bounded transducer T such that, for all $x \in \Sigma^*$ and all $k \geq 0$,

$$|g(x) - T(x01^k)| \leq 2^{-k}.$$

2. g is *F-time computable* (resp. *F-space computable*) for a class \mathcal{F} , if there exists a function $f \in \mathcal{F}$ such that g is f -time (resp. f -space) computable.
3. Let L-comp, P-comp, PSPACE-comp and EXP-comp denote the class of density functions whose *distribution* functions are logarithmic-space, polynomial-time, polynomial-space and exponential-time computable, respectively. Let REC-comp denote the set of all *recursive* density functions.

We remark that if a distribution is \mathcal{F} -time computable, then the density function is also \mathcal{F} -time computable. The converse may not always hold since it is shown, for example, in [17] that if $P \neq NP$, then there is a polynomial-time computable density function such that its associated distribution can not be computed in polynomial time.

Note also that Ben-David et al. [5] use a stronger definition of polynomial-time computability, i.e., for all x , the value of $\mu^*(x)$ is exactly computed by some polynomial-time bounded transducer. Let SP-comp denote the class of these density functions. Naturally, if $g \in \text{SP-comp}$, then $g(x)$ is either 0 or greater than $2^{-p(|x|)}$ for some polynomial p . However, for every $\mu \in \text{P-comp}$, there exists a positive $\nu \in \text{SP-comp}$ such that $\nu(x)$ has at most $4 + 2|x|$ binary digits and $4\nu(x) > \mu(x)$ for all x [17]. Also, if $\mu(x) > 2^{-p(|x|)}$ for some polynomial p , then there exists a total, one-one, p-invertible function $f \in \text{FP}$ such that, for all x , $4 \cdot 2^{-|f(x)|} \leq \mu(x) \leq 20 \cdot 2^{-|f(x)|}$ [3, 37].

A central concept in average-case complexity is “a computation being time (space) bounded on the average for some distribution.” For a discussion on the definition of “polynomial on μ -average,” see [14, 17]. This paper uses a characterization of polynomial on μ -average given by Schapire [30], since it can easily be extended to define the notion of “ f on μ -average” for an arbitrary function f .

Definition 2.2 (cf. [30]) Let f be a function on \mathbb{R}^+ and let μ be a density function. A function $g : \Sigma^* \rightarrow \mathbb{R}^+$ is *f on μ -average* if $\text{Prob}_\mu[\{x \mid g(x) > f(|x| \cdot r)\}] < 1/r$ for any real number $r > 0$. For a class \mathcal{C} of functions, g is *F on μ -average* if there exists a function $f \in \mathcal{C}$, and g is f on μ -average.

It follows immediately from this definition that increasing the value of r also increases the probability weight of the set of strings x with the property that $g(x) \leq f(|x| \cdot r)$. It will be $1 - 1/r$. One significant consequence of this fact is that if g is f on μ -average, then $g(x) \leq f(|x|/\mu(x))$ for all x with $\mu(x) > 0$. This fact can be seen as follows. Suppose that there exists an x_0 such that

$g(x_0) > f(|x_0|/\mu(x_0))$ and $\mu(x_0) > 0$. Choose $r = 1/\mu(x_0)$. Then $\mu(x_0) \leq \text{Prob}_\mu[\{x \mid g(x) > f(|x| \cdot r)\}] < 1/r = \mu(x_0)$, a contradiction.

Definition 2.2 allows us to consider different sets of functions which are bounded on average. For instance, if $\mathcal{F} = \{\lambda x.(x^k + k) \mid k > 0\}$ (the set of polynomials), then we obtain the notion of *polynomial on μ -average* as defined in [22, 23] and used in [14, 30, 17, 5, 36]. Similarly, if $\mathcal{F} = \{\lambda x.(c \log x + d) \mid c, d \geq 0\}$ (the set of logarithmic functions), then the above definition yields the notion of *logarithmic on μ -average* as defined in [5] and also used in [15]. The following lemma gives the justification of the above definitions.

Lemma 2.3 ([30]) *Let g be a function from Σ^* to \mathbb{R}^+ .*

1. *The function g is polynomial on μ -average if and only if it is polynomial on μ -average in the sense of Levin, i.e., for some $\delta > 0$,*

$$\sum_{x:|x|>0} \frac{g(x)^\delta}{|x|} \mu(x) < \infty.$$

2. *The function g is logarithmic on μ -average if and only if, for some $\delta > 0$,*

$$\sum_{x:|x|>0} \frac{(2^{g(x)})^\delta}{|x|} \mu(x) < \infty.$$

Proof. The proof of the claim (1) follows [30]. Without loss of generality, assume that $g(\epsilon) = 0$ and that g is $\lambda n.cn^k$ on μ -average for some k . By definition, for any real number $r > 0$, it holds that $\text{Prob}_\mu[\{x \mid g(x) > c(r|x|)^k\}] < 1/r$. This implies that $\text{Prob}_\mu[\{x \mid g(x) > c(r|x|^2)^k\}] < 1/r$. In other words, $\text{Prob}_\mu[\{x \mid g(x)^{1/2k}|x|^{-1} > c^{1/2k}r^{1/2}\}] < 1/r$. Hence, for every integer $t > 0$ and $r = t^2/c^{1/k}$ it follows that $\text{Prob}_\mu[\{x \mid g(x)^{1/2k}|x|^{-1} > t\}] < c^{1/k}/t^2$. Then,

$$\sum_{x:|x|>0} \frac{g(x)^{1/2k}}{|x|} \mu(x) \leq \sum_{t=1}^{\infty} \text{Prob}_\mu[\{x \mid t-1 < \frac{g(x)^{1/2k}}{|x|} \leq t\}] \cdot t = \sum_{t=0}^{\infty} \text{Prob}_\mu[\{x \mid \frac{g(x)^{1/2k}}{|x|} > t\}] < \infty.$$

Conversely, assume that $\sum_{x:|x|>0} g(x)^\delta |x|^{-1} \mu(x) \leq N$ for some number $N > 0$ and choose a positive integer k such that $1/k < \delta$. Markov's inequality shows us that $\text{Prob}_\mu[\{x \mid g(x)^{1/k}|x|^{-1} > r \cdot N\}] < 1/r$ for any real number $r > 0$. Hence, it yields that $\text{Prob}_\mu[\{x \mid g(x) > (rN|x|)^k\}] < 1/r$.

To see the claim (2), note that a function g is logarithmic on μ -average if and only if $\lambda x.2^g(x)$ is polynomial on μ -average. ■

The next definition follows the notion of “polynomial domination” introduced by Levin [22, 23]. The domination condition between density functions is crucial in the definition of reducibilities among randomized problems in Section 4. Intuitively, it ensures that if an algorithm is fast (e.g., polynomial) on average for a distribution μ , then it is also fast on average for all distributions which are dominated by μ .

Definition 2.4 Let μ_1, μ_2 and ν be density functions.

1. Let t and \mathcal{T} be a function and a set of functions from Σ^* to \mathbb{R}^+ , respectively. The density function μ_2 *t-dominates* μ_1 if $\mu_2(x) \cdot t(x) \geq \mu_1(x)$ for all $x \in \Sigma^*$, and μ_2 *T-dominates* μ_1 if there exists a function $t' \in \mathcal{T}$ such that μ_2 *t'-dominates* μ_1 .

2. Let t and \mathcal{T} be a function and a set of functions on \mathbb{R}^+ , respectively. The density function μ_2 t on ν -average dominates μ_1 if there exists a function t' , from Σ^* to \mathbb{R}^+ , such that t' is t on ν -average and μ_2 t' -dominates μ_1 , and μ_2 \mathcal{T} on ν -average dominates μ_1 if there exists a function t in \mathcal{T} such that μ_2 t on ν -average dominates μ_1 .

This definition enables us to consider logarithmic, polynomial and exponential domination and domination on μ -average. For example, if \mathcal{T} is the set of polynomials, then μ_2 *polynomially dominates* μ_1 , denoted by $\mu_1 \preceq^P \mu_2$, and μ_2 *polynomial on ν -average dominates* μ_1 , denoted by $\mu_1 \preceq^{P,\nu} \mu_2$, respectively. In [17], polynomial domination and polynomial on μ -average domination are called “domination” and “weakly domination”, respectively.

We will now give a general definition of “time- and space-bounded on average” for Turing machines.

Definition 2.5 Let M be an oracle Turing machine, A a set, μ a density function, and let t and \mathcal{T} be a function and a set of functions on \mathbb{R}^+ , respectively. The machine M^A is *t -time bounded on μ -average* if the function Time_M^A is t on μ -average, and M^A is *\mathcal{T} -time bounded on μ -average* if there exists a function $t \in \mathcal{T}$ such that M^A is t -time bounded on μ -average. The notions of *t -space bounded on μ -average* and *\mathcal{T} -space bounded on μ -average* are defined similarly using Space_M^A instead of Time_M^A .

For instance, if \mathcal{T} is the set of polynomials, then we say that M^A is *polynomial-time (or polynomial-space) bounded on μ -average* as in [22, 23]. If a function f is computed by a deterministic transducer which is polynomial-time bounded on μ -average, we say that f is *computable in time polynomial on μ -average*.

We observe that the quantifier characterization of nondeterministic and probabilistic Turing machines holds in average-case setting. Recall that all sets in NP can be characterized by an existential quantifier and deterministic Turing machines as follows: a set D is in NP if and only if there exist a polynomial p and a polynomial-time deterministic Turing machine M such that $D = \{x \mid \exists y[|y| = p(|x|) \wedge \langle x, y \rangle \in L(M)]\}$ [40].

Proposition 2.6 *For every set D and every density function μ , the following statements are equivalent.*

1. *There exists a nondeterministic Turing machine M such that $D = L(M)$ and M is polynomial-time bounded on μ -average.*
2. *There exists a function p from Σ^* to \mathbb{N} and a polynomial-time bounded deterministic Turing machine M such that p is computable in time polynomial on μ -average and $D = \{x \mid \exists y[|y| = p(|x|) \wedge \langle x, y \rangle \in L(M)]\}$.*

Proposition 2.7 *For every set D , every density function μ , the following statements are equivalent.*

1. *There exists a bounded-error probabilistic Turing machine M such that $D = L(M)$ and M is polynomial-time bounded on μ -average.*
2. *For every function q that is computable in time polynomial on μ -average, there exists a probabilistic Turing machine M such that M is polynomial-time bounded on μ -average, $D = L(M)$ and $\text{Prob}_M[M(x) = D(x)] \geq 1 - 2^{-q(x)}$ for all x .*

3. For every function q that is computable in time polynomial on μ -average, there exists a function f from Σ^* to \mathbb{N} and a polynomial-time bounded deterministic Turing machine M such that f is computable in time polynomial on μ -average, $D = L(M)$, and $\text{Prob}_\nu[\{y \mid x \in D \text{ iff } \langle x, y \rangle \in L(M)\}] \geq 1 - 2^{-q(x)}$ for all x , where ν is the uniform density function on $\Sigma^{f(x)}$.

Proof. The proof of Proposition 2.6 is straightforward and follows from the standard technique of encoding nondeterministic computation paths into strings and the fact that Time_M is computable by a deterministic transducer that is polynomial-time bounded on μ -average.

The same argument shows that (3) implies (2) in Proposition 2.7. Clearly (2) implies (1). Thus, it remains to show that (1) implies (3). Assume that (1) holds. Then, Time_M is polynomial on μ -average. Now we perform the usual probability amplification (see, e.g., [2], p. 139). We simulate the machine M $p(n)$ ($= O(q(n))$) times and accept x as soon as more than $p(n)/2$ simulations accept x , and reject x as soon as more than $p(n)/2$ simulations reject x . Hence, if we choose $f(x) = p(|x|) \cdot k \cdot \text{Time}_M(x)$, then (1) implies (3). ■

3 Randomized Decision Problems

The basic objects of average-case complexity theory are (decision or search) problems together with distributions on instances i.e., a density function assigns probabilities to instances of those problems. Then the time and space complexity of an algorithm for that problem is measured under the assumption that the inputs occur according to the given distribution. The hope is to show that even for (some) intractable problems, hard instances occur only with small probability. Hence, some algorithm should run efficiently on average.

Some NP-complete problems, such as the *satisfiability problem*, the *graph 3-colorability problem* and the *Hamiltonian circuit problem*, can be solved by deterministic algorithms in time polynomial on average with respect to reasonably chosen density functions [10, 8, 39].

This paper will focus only on *decision problems* (readers interested in *search problems* are referred to [5, 7]). For a decision problem D and a density function μ , the pair (D, μ) is called a *randomized (decision) problem*. Here, (D, μ) means that instances of the decision problem D are given randomly according to μ ; in other words, a string s occurs as input to some algorithm deciding D with probability $\mu(s)$. *Average-case complexity classes* are sets of randomized problems. We note that, in [22, 23], Levin has considered pairs of a decision problem and a distribution function (also called *distribution problems*). See [5] for more details.

We will consider two different types of average-case complexity classes. The first type is defined by a worst-case complexity class and a class of density functions. In the second type of classes, the resource bounds of the complexity class are taken with respect to the given density functions.

Definition 3.1 ([5]) Let \mathcal{C} be a complexity class and \mathcal{F} be a class of density functions. The *randomized class* $\langle \mathcal{C}, \mathcal{F} \rangle$ is the set $\{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D \in \mathcal{C}\}$.

Definition 3.2 Let t be a function on \mathbb{N} and \mathcal{F} be a class of density functions. Time- and space-bounded *average classes* are defined as follows.

1. $\text{Aver}\langle \text{DTIME}(t), \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D = L(M) \text{ for a deterministic Turing machine } M \text{ which is } t\text{-time bounded on } \mu\text{-average}\}$.
2. $\text{Aver}\langle \text{NTIME}(t), \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D = L(M) \text{ for a nondeterministic Turing machine } M \text{ which is } t\text{-time bounded on } \mu\text{-average}\}$.

3. $\text{Aver}\langle \text{DSpace}(t), \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D = L(M) \text{ for a deterministic Turing machine } M \text{ which is } t\text{-space bounded on } \mu\text{-average}\}$.
4. $\text{Aver}\langle \text{NSpace}(t), \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D = L(M) \text{ for a nondeterministic Turing machine } M \text{ which is } t\text{-space bounded on } \mu\text{-average}\}$.
5. $\text{Aver}\langle \text{BPTIME}(t), \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F} \text{ and } D = L(M) \text{ for a bounded-error probabilistic Turing machine } M \text{ which is } t\text{-time bounded on } \mu\text{-average}\}$.

Using the above definitions, one can consider average-case analogues of many known time- or space-bounded complexity classes. For example, NP with polynomial-time computable distributions on the inputs, as defined in [5], is expressed as $\langle \text{NP}, \text{SP-comp} \rangle$. The set of randomized problems solvable in polynomial-time on average (AverageP or AP in [17, 36]) is denoted by $\text{Aver}\langle \text{P}, * \rangle$. Here, $*$ denotes the set of *all* density functions. $\text{Aver}\langle \text{P}, \text{P-comp} \rangle$ and $\text{Aver}\langle \text{NP}, \text{P-comp} \rangle$ (denoted by AP_{P} and ANP_{P} in [36]) is the set of problems (D, μ) such that $\mu \in \text{P-comp}$, and D is respectively solvable deterministically and nondeterministically in polynomial-time on μ -average. The class of randomized problems which are solvable in logarithmic-space on average (Average-logspace [5] and averageL in [15]) is denoted by $\text{Aver}\langle \text{L}, * \rangle$.

Note that Ben-David et al. [5] use the notation $\text{AverDTime}(t(n))$ to denote $\text{Aver}\langle \text{DTIME}(t), * \rangle$ (also denoted by $\text{AvDTime}(t(n))$ in [29]). Several important randomized decision problems which belong to $\langle \text{NP}, \text{P-comp} \rangle$ can be found in [17].

The next propositions follow immediately from the definitions of the average-case complexity classes.

Proposition 3.3 *Let $\mathcal{C} \in \{\text{DTIME}(t), \text{NTIME}(t), \text{DSpace}(t), \text{NSpace}(t), \text{BPTIME}(t)\}$ for some function t on \mathbb{N} , and let \mathcal{F} be a class of density functions. Then, $\langle \mathcal{C}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$.*

Proposition 3.4 *Let \mathcal{F} be a class of density functions and let t be a function on \mathbb{N} .*

1. $\text{Aver}\langle \text{DTIME}(t), \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{NTIME}(t), \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{NSpace}(t), \mathcal{F} \rangle$.
2. $\text{Aver}\langle \text{DTIME}(t), \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{DSpace}(t), \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{DTIME}(2^{t(n)}), \mathcal{F} \rangle$.

It is natural to ask whether $(D, \nu) \in \text{Aver}\langle \text{DTIME}(t), \mathcal{F} \rangle$ implies that $(D, \mu) \in \text{Aver}\langle \text{DTIME}(to h), \mathcal{F} \rangle$ if $\mu \in \mathcal{F}$ and ν h -dominates μ for some function h . An affirmative answer for a special case is given by the following lemma. A class \mathcal{T} of functions is said to be *closed under composition with polynomials* if, for any function t and any polynomial p , $t \in \mathcal{T}$ implies $\lambda x.t(p(x)) \in \mathcal{T}$.

Lemma 3.5 *Let μ and ν be density functions, \mathcal{T} a class of functions on \mathbb{R}^+ which is closed under composition with polynomials, and let h be a function from Σ^* to \mathbb{R}^+ . If ν polynomially on μ -average dominates μ and h is \mathcal{T} on ν -average, then h is also \mathcal{T} on μ -average.*

Proof. Assume that $\mu \preceq^{\text{P}, \mu} \nu$, and h is t on ν -average for some function $t \in \mathcal{T}$. Choose a function q which is polynomial on μ -average such that, for all $x \in \Sigma^*$, $\nu(x) \cdot q(x) \geq \mu(x)$. By assumption, $\text{Prob}_{\nu}[\{x \mid h(x) > t(|x| \cdot r)\}] < 1/r$ for all $r > 0$. Since \mathcal{T} is closed under composition with polynomials, without loss of generality, we assume that $h(\epsilon) < t(0)$ and therefore $\text{Prob}_{\mu}[\{x \mid h(\epsilon) > t(0)\}] = 0$. Since q is polynomial on μ -average, there exists a polynomial p such that $\text{Prob}_{\mu}[\{x \mid q(x) > p(|x| \cdot r)\}] < 1/r$ for all $r > 0$.

Let μ_n and ν_n denote the conditional probability of strings of length n of μ and ν , respectively. Note that, if $q(x) \leq p(|x| \cdot r)$ for strings x of length n , then $\mu(\Sigma^n) \cdot \mu_n(x) \leq \nu(\Sigma^n) \cdot p(n \cdot r) \cdot \nu_n(x)$

and that $\nu(\Sigma^n) \cdot \text{Prob}_{\nu_n}[\{x \mid h(x) > t(|x| \cdot r)\}] < 1/r$ for all $n \in \mathbb{N}$ and $r > 0$. Now define g as $g(x) = t(4x^3 \cdot p(2x))$. Since the class \mathcal{T} is closed under composition with polynomials, g is in \mathcal{T} . Then, for all $n \in \mathbb{N}$ and all $r \in \mathbb{R}^+$, it follows that $g(n \cdot r) \geq t(n \cdot r^3 \cdot 4n^2 \cdot p(n \cdot 2r))$. It remains to show that h is g on μ -average. For every real number $r > 0$,

$$\begin{aligned}
& \text{Prob}_\mu[\{x \mid h(x) > g(|x| \cdot r)\}] \\
& \leq \text{Prob}_\mu[\{x \mid q(x) > p(|x| \cdot 2r)\}] + \text{Prob}_\mu\{x \mid q(x) \leq p(|x| \cdot 2r) \wedge h(x) > g(|x| \cdot r)\} \\
& < \frac{1}{2r} + \sum_{n=1}^{\infty} \nu(\Sigma^n) \cdot p(2nr) \cdot \text{Prob}_{\nu_n}[\{x \mid q(x) \leq p(2nr) \wedge h(x) > g(nr)\}] \\
& \leq \frac{1}{2r} + \sum_{n=1}^{\infty} \nu(\Sigma^n) \cdot p(2nr) \cdot \text{Prob}_{\nu_n}[\{x \mid h(x) > t(|x| \cdot 4n^2 r^3 \cdot p(2nr))\}] \\
& < \frac{1}{2r} + \sum_{n=1}^{\infty} \frac{p(2nr)}{4n^2 r^3 \cdot p(2nr)} = \frac{1}{2r} + \frac{6}{2\pi^2 r^3} < \frac{1}{r}.
\end{aligned}$$

■

Theorem 3.6 *Let $\mathcal{C} \in \{\text{L}, \text{P}, \text{NP}, \text{PSPACE}, \text{NPSpace}, \text{EXP}\}$, and assume $(D, \nu) \in \text{Aver}(\mathcal{C}, \mathcal{F})$, for some class of density functions \mathcal{F} . For all $\mu \in \mathcal{F}$, if ν polynomially on μ -average dominates μ , then $(D, \mu) \in \text{Aver}(\mathcal{C}, \mathcal{F})$.*

Proof. Since the class of polynomials, logarithms, and exponentials are closed under composition with polynomials, the theorem immediately follows from Lemma 3.5. ■

From the definition, it is obvious that if \mathcal{C}_1 is a proper subset of \mathcal{C}_2 , then $\langle \mathcal{C}_1, \mathcal{F} \rangle$ is also a proper subset of $\langle \mathcal{C}_2, \mathcal{F} \rangle$. A similar hierarchy result can be shown for the second type of average-case complexity classes. See also [15, 13].

Theorem 3.7 *Let s and t be space- and time-constructible functions, respectively, and $s'(n) \in \omega(s(n \cdot f(n)))$ and $t'(n) \in \omega(t(n \cdot f(n)))$ for some nondecreasing unbounded function f from \mathbb{N} to \mathbb{R}^+ . Assume that \mathcal{F} contains a density function μ_f of the form*

$$\mu_f(x) = \begin{cases} \frac{1}{c \cdot f(|x|)} \cdot 2^{-|x|} & \text{if } f(|x| - 1) < k^2 \leq f(|x|) \text{ for some integer } k > 2, \\ 0 & \text{otherwise,} \end{cases}$$

where c is a constant with $\frac{1}{4} > c > 0$.

1. $\text{Aver}\langle \text{DSPACE}(s(n)), \mathcal{F} \rangle \subsetneq \text{Aver}\langle \text{DSPACE}(s'(n)), \mathcal{F} \rangle$.
2. $\text{Aver}\langle \text{DTIME}(t(n)), \mathcal{F} \rangle \subsetneq \text{Aver}\langle \text{DTIME}(\log t(n) \cdot t'(n)), \mathcal{F} \rangle$.

Proof. (1) Choose a nondecreasing, unbounded function f such that $\mu_f \in \mathcal{F}$ and $s'(n) \in \omega(s(n \cdot f(n)))$. It is shown in [38, 27] that there exists a set A in $\text{DSPACE}(s'(n))$ which are *random* for $\text{DSPACE}(s(n \cdot f(n)))$, i.e., for any set $B \in \text{DSPACE}(s(n \cdot f(n)))$ and for any $\epsilon > 0$, there exists an integer n_0 such that, for all $n > n_0$,

$$\left| \frac{|A^n \Delta B^n|}{2^n} - \frac{1}{2} \right| < \epsilon.$$

Clearly $(A, \mu_f) \in \langle \text{DSPACE}(s'(n)), \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{DSPACE}(s'(n)), \mathcal{F} \rangle$.

We show that $(A, \mu_f) \notin \text{Aver}\langle \text{DSPACE}(S(n)), \mathcal{F} \rangle$. Assume on the contrary that $(A, \mu) \in \text{Aver}\langle \text{DSPACE}(s(n)), \mathcal{F} \rangle$ and let M be a deterministic Turing machine which accepts A and is s -space bounded on μ_f -average, i.e., $\text{Prob}_{\mu_f}[\{x \mid \text{Space}_M(x) > s(|x| \cdot r)\}] < 1/r$ for any real number $r > 0$. Then, a set $D = \{x \mid \text{Space}_M(x) < s(n \cdot f(n)) \wedge M(x) = 1\}$ is in $\text{DSPACE}(s(n \cdot f(n)))$, and therefore A is random for $\{D\}$. Choose $\epsilon = 1/8$. The randomness of A yields that there exists an integer $n_0 > 0$ such that, for all $n > n_0$,

$$\left| \frac{|A^n \Delta D^n|}{2^n} - \frac{1}{2} \right| < 1/8.$$

Note that, for all n , $\mu_f(\Sigma^n) = 1/(c \cdot f(n))$ if $f(n-1) < k^2 \leq f(n)$ some constant k . It follows that, for some $n > n_0$,

$$\text{Prob}_{\mu_f}[\{x \in \Sigma^n \mid \text{Space}_M(x) > s(n \cdot f(n))\}] \leq \frac{1}{f(n)} \cdot c \cdot f(n) < \frac{1}{4}.$$

Therefore, D is identical to A on at least $3/4$ of all strings of length n . This contradicts the randomness of A .

(2) The proof for the time-bounded classes is similar to (1) and follows from the fact that, for every time-constructible function $t(n) > n$ and $t'(n) \in \omega(t(n))$, there exists a set in $\text{DTIME}(\log t(n) \cdot t'(n))$ which is random for $\text{DTIME}(t(n))$ [38, 27]. ■

Corollary 3.8 *Let p and p' be polynomials such that $p'(n) \in \omega(p(n))$.*

1. $\text{Aver}\langle \text{DTIME}(p(n)), \text{P-comp} \rangle \subsetneq \text{Aver}\langle \text{DTIME}(\log p(n) \cdot p'(n) + n), \text{P-comp} \rangle$.
2. $\text{Aver}\langle \text{DSPACE}(p(n)), \text{P-comp} \rangle \subsetneq \text{Aver}\langle \text{DSPACE}(p'(n) + \log n), \text{P-comp} \rangle$.

We show a basic relationship between worst-case complexity and average-case complexity on strings with high probability. To show this, we first introduce a property of an average-case complexity class $\text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$.

Definition 3.9 For a sparse set S and a polynomial q , let $\mu_{S,q}$ denote a density function such that $\mu_{S,q}(x) \geq 1/q(|x|)$ for all $x \in S$. A class $\text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$ has the *sparse interpolation property* if, for any set A , any infinite sparse set S and any polynomial q such that $(A, \mu_{S,q}) \in \text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$, there exists a set $B \in \mathcal{C}$ such that $A \cap S' \subseteq B \subseteq A$. The set B is called an *interpolant* of A and S .

Lemma 3.10 *For a class $\mathcal{C} \in \{\text{P}, \text{NP}, \text{BPP}, \text{PSPACE}\}$, $\text{Aver}\langle \mathcal{C}, * \rangle$ has the sparse interpolation property.*

Proof. We show the case $\mathcal{C} = \text{NP}$. Take any sparse set S and a polynomial q and assume that $(A, \mu_{S,q}) \in \text{Aver}\langle \text{NP}, * \rangle$. There exists a Turing machine M which computes A such that Time_M is p on $\mu_{S,q}$ -average for some polynomial p . Note that $\text{Time}_M(x) \leq p(|x|/\mu_{S,q}(x))$ for all x with $\mu_{S,q}(x) > 0$. Let N simulate M on input x in $p(|x| \cdot q(|x|))$ steps. If the simulation of M does not terminate within $p(|x| \cdot q(|x|))$ steps, then N rejects x . Let $B = L(N)$. Clearly $B \subseteq A$. Since $q(|x|) \geq 1/\mu_{S,q}(x)$ for all $x \in S$, N completely simulates M on all inputs x in S , and thus, $A \cap S = B \cap S$. Clearly N is polynomial-time bounded. Therefore, $B \in \text{NP}$. ■

In Section 6, we will extend Lemma 3.10 to the k -th level of an “average polynomial-time hierarchy” which is an average-case version of the Meyer-Stockmeyer polynomial-time hierarchy.

One of the most interesting open questions is whether NP sets can be solved in average polynomial time for every polynomial time computable distribution function, i.e., if $\langle \text{NP}, \text{P-comp} \rangle \subseteq \text{Aver}\langle \text{P}, * \rangle$ holds or not. Clearly if $\text{P} = \text{NP}$, then $\langle \text{NP}, \mathcal{F} \rangle$ is included in $\text{Aver}\langle \text{P}, \mathcal{F} \rangle$. Ben-David et al. [5] first gave a partial answer to this question by showing that $\langle \text{NP}, \text{P-comp} \rangle \not\subseteq \text{Aver}\langle \text{P}, * \rangle$ if $\text{E} \neq \text{NE}$. Under the same assumption, it also holds that $\langle \text{NP}, \text{L-comp} \rangle \not\subseteq \text{Aver}\langle \text{P}, * \rangle$ [5]. Ben-David et al. actually show that $\langle \text{NP} \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \text{P}, * \rangle$ if and only if $\text{NP} \cap \text{TALLY} \subseteq \text{P}$.

Here we consider the bounded-error probabilistic class $\text{Aver}\langle \text{BPP}, * \rangle$ and extend the above result by Ben-David et al. to $\text{Aver}\langle \text{BPP}, * \rangle$.

Theorem 3.11 $\langle \text{NP} \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \text{BPP}, * \rangle$ if and only if $\text{NP} \cap \text{TALLY} \subseteq \text{BPP}$.

Proof. From the assumption that $\text{NP} \cap \text{TALLY} \subseteq \text{BPP}$, it immediately follows that $\langle \text{NP} \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \text{BPP}, * \rangle$. Assume conversely that $\langle \text{NP} \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \text{BPP}, * \rangle$. Choose a density function μ satisfying $\mu(x) \propto (|x| + 1)^{-2}$ if $x \in \{0\}^*$, and 0 otherwise. Note that $\mu \in \text{L-comp}$. By our assumption, for every set $A \in \text{NP} \cap \text{TALLY}$, $(A, \mu) \in \text{Aver}\langle \text{BPP}, * \rangle$. That is, there is a probabilistic Turing machine M with exponential small error probability computing A which runs in time q on μ -average for some polynomial q . This yields that $\text{Time}_M(x) \leq q(|x|/\mu(x)) \leq q(3|x|(|x| + 1)^2)$ for all $x \in \{0\}^*$. Therefore, A is in BPP. ■

A standard padding argument shows that $\text{NP} \cap \text{TALLY} \subseteq \text{BPP}$ if and only if $\text{NE} \subseteq \text{BPE}$. Hence, we have the following conclusion.

Corollary 3.12 $\text{NE} \not\subseteq \text{BPE}$ implies $\langle \text{NP}, \text{P-comp} \rangle \not\subseteq \text{Aver}\langle \text{BPP}, * \rangle$.

4 Many-One and Turing Reductions

A theory of average NP-completeness was initiated by Levin in his terse papers [22, 23]. Levin has introduced a polynomial-time many-one reducibility between randomized decision problems and has shown that the *randomized tiling problem*, the “tiling problem” with a natural distribution, is complete for Random-NP or, in our notation, complete for $\langle \text{NP}, \text{P-comp} \rangle$. Intuitively, a reduction from (D_1, μ_1) to (D_2, μ_2) should reduce a set D_1 to a set D_2 and ensure a *domination condition* between μ_1 and μ_2 that guarantees that instances in D_1 that occur with high probability are reduced to instances in D_2 that occur with high probability. The notions of deterministic Turing reducibility and random many-one reducibility have been introduced in [22, 5]. The latter is especially suitable for randomized algorithms (see [5, 19, 7]). Also considered so far were logspace many-one reductions [5] and logspace many-one reductions that are polynomially honest [15].

We first recall from [22, 5] the definition of polynomial time many-one and Turing reducibility.

Definition 4.1 ([22]) Let (D_1, μ_1) and (D_2, μ_2) be randomized decision problems.

1. (D_1, μ_1) is *polynomial-time many-one reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_m^{\text{P}} (D_2, \mu_2)$, if there exists a density function ν and a function f such that
 - i. $f \in \text{FP}$;
 - ii. for all x , $x \in D_1$ if and only if $f(x) \in D_2$; and
 - iii. $\mu_1 \leq^{\text{P}} \nu$, and $\mu_2(y) \geq \text{Prob}_\nu[\{x \mid f(x) = y\}]$ for all y .

2. (D_1, μ_1) is *average polynomial-time many-one reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_m^{\text{p,av}} (D_2, \mu_2)$, if f is polynomial on μ_1 -average and $\mu_1 \leq^{\text{p}, \mu_1} \nu$.

The condition (iii) on the density functions in the above definition is simply called the *domination condition* for the reduction. There are several weaker definitions of the domination conditions; however, the one used here simplifies the reductions and guarantees that the reducibilities are reflexive and transitive. See [17] for a discussion.

Definition 4.2 ([5]) Let (D_1, μ_1) and (D_2, μ_2) be randomized decision problems.

1. (D_1, μ_1) is *polynomial-time Turing reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_T^{\text{p}} (D_2, \mu_2)$, if there exists a density function ν and an oracle Turing machine M such that
 - i. M^{D_2} is polynomial-time bounded;
 - ii. $D_1 = L(M, D_2)$; and
 - iii. $\mu_1 \leq^{\text{p}} \nu$, and $\mu_2(y) \geq \text{Prob}_\nu[\{x \mid y \in Q(M, D_2, x)\}]$ for all y .

Here $Q(M, D_2, x)$ is the set of strings queried by M with oracle D_2 on input x .

2. (D_1, μ_1) is *average polynomial-time Turing reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_T^{\text{p,av}} (D_2, \mu_2)$, if M is polynomial-time bounded on μ -average and $\mu_1 \leq^{\text{p}, \mu_1} \nu$.

Let α be a reducibility. For a class \mathcal{C} , a problem (D, μ) is α -hard for \mathcal{C} if every problem (E, ν) in \mathcal{C} is α -reducible to (D, μ) , and (D, μ) is α -complete for \mathcal{C} if it is in \mathcal{C} and is α -hard for \mathcal{C} . If \mathcal{C} is of the form $\langle \mathcal{C}', \mathcal{F} \rangle$ (resp. $\text{Aver}(\mathcal{C}', \mathcal{F})$), then let \mathcal{C} -complete abbreviate “many-one complete” (resp. “many-one complete on average”).

It is shown in [17] that if (D, μ) is \leq_m^{p} -hard for $\langle \text{NP}, \text{P-comp} \rangle$ for a density function D in EXP and a flat density function μ , then $\text{EXP} = \text{NEXP}$. Recall that the standard density function is flat. Hence, under the condition of $\text{EXP} \neq \text{NEXP}$, no decision problem with the standard density function is $\langle \text{NP}, \text{P-comp} \rangle$ -complete. By the result of Wang and Belanger [36], every $\langle \text{NP}, \text{P-comp} \rangle$ -complete problem is also $\text{Aver}(\langle \text{NP}, \text{P-comp} \rangle)$ -complete, i.e., $\leq_m^{\text{p,av}}$ -complete for $\text{Aver}(\langle \text{NP}, \text{P-comp} \rangle)$.

Now we introduce an average-case version of the nondeterministic Turing reducibility. This will be used for the definition of an “average polynomial-time hierarchy” in Section 6. Recall that $\text{Acc}(M, D, x)$ (resp. $\text{Rej}(M, D, x)$) is the set of (codes of) all accepting (resp. rejecting) computation paths of M with oracle D on input x , and $Q(M, D, x, y)$ is the set of strings queried by M with oracle D on input x on computation path y .

Definition 4.3 Let (D_1, μ_1) and (D_2, μ_2) be randomized decision problems.

1. (D_1, μ_1) is *polynomial-time nondeterministic Turing reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_T^{\text{np}} (D_2, \mu_2)$, if there exist a density function ν and a nondeterministic Turing machine M such that
 - i. M^{D_2} is polynomial-time bounded;
 - ii. $D_1 = L(M, D_2)$; and
 - iii. $\mu'_1 \leq^{\text{p}} \nu$, and $\mu_2(z) \geq \text{Prob}_\nu[\{(x, y) \mid z \in Q(M, D_2, x, y)\}]$ for all z ,

where μ'_1 is the density function induced from μ , M and D_2 as follows:

$$\mu'_1(x, y) = \begin{cases} \mu_1(x)/|\text{Acc}(M, D_2, x)| & \text{if } y \in \text{Acc}(M, D_2, x), \\ \mu_1(x)/|\text{Rej}(M, D_2, x)| & \text{if } \text{Acc}(M, D_2, x) = \emptyset \text{ and } y \in \text{Rej}(M, D_2, x), \\ 0 & \text{otherwise.} \end{cases}$$

2. (D_1, μ_1) is *average polynomial-time nondeterministic Turing reducible* to (D_2, μ_2) , denoted by $(D_1, \mu_1) \leq_T^{\text{np,av}} (D_2, \mu_2)$, if M^{D_2} is polynomial-time bounded on μ_1 -average and $\mu'_1 \preceq^{\text{p}, \mu'_1} \nu$.

We note that, in the case that the reduction machine always has one computation path, the nondeterministic Turing reduction coincides with deterministic one in Definition 4.2.

In the following, we state basic properties of the reducibilities (cf. [17, 5]).

Proposition 4.4 *Let $A_i = (D_i, \mu_i)$, $i = 1, 2, 3$, be randomized decision problems.*

1. *Polynomial-time reducibility implies average polynomial-time reducibility, i.e., for every $\alpha \in \{\text{m}, \text{T}\}$ and $\beta \in \{\text{p}, \text{np}\}$, if $A_1 \leq_\alpha^\beta A_2$ then $A_1 \leq_\alpha^{\beta, \text{av}} A_2$.*
2. *Many-one reducibility implies deterministic Turing reducibility, i.e., if $A_1 \leq_m^{\text{p}} A_2$ then $A_1 \leq_T^{\text{p}} A_2$, and if $A_1 \leq_m^{\text{p,av}} A_2$ then $A_1 \leq_T^{\text{p,av}} A_2$.*
3. *Deterministic Turing reducibility implies nondeterministic Turing reducibility, i.e., if $A_1 \leq_T^{\text{p}} A_2$ then $A_1 \leq_T^{\text{np}} A_2$, and if $A_1 \leq_T^{\text{p,av}} A_2$ then $A_1 \leq_T^{\text{np,av}} A_2$.*
4. *Nondeterministic Turing reducibility is reflexive, i.e., $A_1 \leq_T^{\text{np}} A_1$ and $A_1 \leq_T^{\text{np,av}} A_1$.*
5. *Many-one reducibility and deterministic Turing reducibility are reflexive and transitive, i.e., for every $\leq_\alpha \in \{\leq_m^{\text{p}}, \leq_m^{\text{p,av}}, \leq_T^{\text{p}}, \leq_T^{\text{p,av}}\}$, $A_1 \leq_\alpha A_1$, and if $A_1 \leq_\alpha A_2$ and $A_2 \leq_\alpha A_3$ then $A_1 \leq_\alpha A_3$.*

Proof. The claims (1)-(4) follow immediately from the definitions. The claim for reflexivity in (5) is obvious. Here we show that $\leq_T^{\text{p,av}}$ is transitive. The other proofs for transitivity are analogous. Assume that $(D_1, \mu_1) \leq_T^{\text{p,av}} (D_2, \mu_2)$ via a deterministic Turing machine M_1 and a density function ν_1 and assume that $(D_2, \mu_2) \leq_T^{\text{p,av}} (D_3, \mu_3)$ via a deterministic machine M_2 and a density function ν_2 . We show that $(D_1, \mu_1) \leq_T^{\text{p,av}} (D_3, \mu_3)$.

By definition, there exist two functions f_1 and f_2 which are polynomial on μ_1 -average and on μ_2 -average, respectively, such that $f_1(x) \cdot \nu_1(x) \geq \mu_1(x)$ and $f_2(x) \cdot \nu_2(x) \geq \mu_2(x)$ for all x . Assume that $f_1(x) > 0$ and $f_2(x) > 0$ for all strings x .

We define a new machine M as follows: on input x , M deterministically simulates M_1 on x , and whenever M_1 queries a string y , M deterministically simulates M_2 on input y . Especially, in the case that x is the empty string ϵ , M is designed not to query any strings, even if M_2 queries some strings to oracle D_3 , but their oracle answers are encoded in a program of M . Clearly $D_1 = L(M, D_3)$. Note that

$$\text{Time}_M^{D_3}(x) \leq c \cdot \left(\text{Time}_{M_1}^{D_2}(x) + \sum_{y \in Q(M_1, D_2, x)} \text{Time}_{M_2}^{D_3}(y) \right)$$

for some constant $c > 0$.

Let $f(x) = f_1(x) \cdot (\sum_{y \in Q(M_1, D_2, x)} f_2(y) + 1)$ and choose a density function ν such that $\nu(\epsilon) = 1 - \sum_x \mu_1(x)/f(x)$ and $f(x) \cdot \nu(x) = \mu_1(x)$ for all strings x different from ϵ . For each string z , we define a set A_z such that that $(x, y) \in A_z$ if and only if $z \in Q(M_2, D_3, y)$, $y \in Q(M_1, D_2, x)$ and $(x', y) \notin A_z$ for any string $x' < x$. Then, for all $z \in \bigcup_x Q(M, D_3, x)$,

$$\mu_3(z) \geq \text{Prob}_{\nu_2}[\{y \mid z \in Q(M_2, D_3, y)\}] \geq \sum_{y: z \in Q(M_2, D_3, y)} \frac{\mu_2(y)}{f_2(y)}$$

$$\begin{aligned}
&\geq \sum_{x:z \in Q(M, D_3, x)} \sum_{y:(x,y) \in A_z} \frac{1}{f_2(y)} \cdot \mu_2(y) \geq \sum_{x:z \in Q(M, D_3, x)} \sum_{y:(x,y) \in A_z} \frac{1}{f_2(y)} \cdot \frac{\mu_1(x)}{f_1(x)} \\
&\geq \sum_{x:z \in Q(M, D_3, x)} \frac{\mu_1(x)}{f_1(x) \cdot \sum_{y \in Q(M_1, D_2, x)} f_2(y)} \geq \text{Prob}_\nu[\{x \mid z \in Q(M, D_3, x)\}]
\end{aligned}$$

since $\sum_{i=1}^n \frac{1}{a_i} \geq 1 / \sum_{i=1}^n a_i$.

It remains to show that M and f are polynomial on μ_1 -average. Let p_1 and p_2 be polynomials such that, for any real number $r > 0$,

$$\begin{aligned}
\text{Prob}_{\mu_1}[\{x \mid \text{Time}_{M_1}^{D_2}(x) > p_1(|x| \cdot r)\}] &< 1/r, \text{ and} \\
\text{Prob}_{\mu_2}[\{y \mid \text{Time}_{M_2}^{D_3}(y) > p_2(|y| \cdot r)\}] &< 1/r.
\end{aligned}$$

Now we define a polynomial s as $s(z) = c \cdot p_1(2z) \cdot (1 + p_2(2z \cdot p_1(2z))) + c_0$, where $c_0 = \text{Time}_M^{D_3}(\epsilon)$. Note that $|Q(M_1, D_2, x)| \leq \text{Time}_{M_1}^{D_2}(x)$ for all x . Then, for all x and $r > 0$,

$$\begin{aligned}
&\text{Prob}_{\mu_1}[\{x \mid \text{Time}_M^{D_3}(x) > s(|x| \cdot r)\}] \\
&\leq \text{Prob}_{\mu_1}[\{x \mid \text{Time}_{M_1}^{D_2}(x) > p_1(|x| \cdot 2r)\}] + \text{Prob}_{\mu_1}[\{x \mid \text{Time}_{M_1}^{D_2}(x) \leq p_1(|x| \cdot 2r) \\
&\quad \wedge \sum_{y \in Q(M_1, D_2, x)} \text{Time}_{M_2}^{D_3}(y) > p_2(p_1(|x| \cdot 2r) \cdot 2r) \cdot p_1(|x| \cdot 2r)\}] \\
&< 1/2r + \text{Prob}_{\mu_1}[\{x \mid \exists y \in Q(M_1, D_2, x)[\text{Time}_{M_2}^{D_3}(y) > p_2(|y| \cdot 2r)\}] \\
&\leq 1/2r + \text{Prob}_{\mu_2}[\{y \mid \text{Time}_{M_2}^{D_3}(y) > p_2(|y| \cdot 2r)\}] < 1/r.
\end{aligned}$$

The proof that f is polynomial on μ_1 -average is similar, and thus, the claim is established. \blacksquare

The following lemma shows how the domination condition for Turing reducibilities works.

Lemma 4.5 *Assume that (E, ν) is computable by a nondeterministic Turing machine N in time polynomial on ν -average.*

1. *Assume that $(D, \mu) \leq_T^{\text{p,av}} (E, \nu)$ via a machine M and let $h(x) = \sum_{z \in Q(M, D, x)} \text{Time}_N(z)$. Then, h is polynomial on μ -average.*
2. *Assume that $(D, \mu) \leq_T^{\text{np,av}} (E, \nu)$ via a machine M and let $h(x) = \min_{y \in \text{Acc}(M, D, x)} \sum_{z \in Q(M, D, x, y)} \text{Time}_N(z)$ if $x \in D$; otherwise, $h(x) = \min_{y \in \text{Rej}(M, D, x)} \sum_{z \in Q(M, D, x, y)} \text{Time}_N(z)$. Then, h is polynomial on μ -average.*

Proof. We prove the claim (2). Since $(D, \mu) \leq_T^{\text{np,av}} (E, \nu)$ via nondeterministic oracle Turing machine M , there exist a density function ν and a polynomial p_D such that $D = L(M, E)$, $\mu' \preceq^{\text{p}, \mu'} \nu$, Time_M^E is p_D on μ -average, and $\nu(z) \geq \text{Prob}_\nu[\{(x, y) \mid z \in Q(M, E, x, y)\}]$ for all z , where μ' is the density function induced from μ , M and E as in Definition 4.3. Assume that $\text{Time}_M^E(x) > |x|$ for all x .

Choose a nondeterministic Turing machine N and a polynomial p_E such that $E = L(N)$ and Time_N is p_E on ν -average. Furthermore, let p be a polynomial and q a function such that q is p on μ -average and $q(x) \cdot \nu(x) \geq \mu'(x)$ for all x .

Now define a polynomial s as

$$s(z) = p_D(6z) \cdot p_E \left(p_D(6z) \cdot 6z^2 \cdot (p_D(6z) + 2)^2 \cdot p(6z \cdot p_D(6z)) \right) + c_0,$$

where $c_0 = \text{Time}_M^E(\epsilon)$.

We show that h is s on μ -average. For simplicity, let A_x denote $\text{Acc}(M, E, x)$ and let $Q_{x,y}$ denote $Q(M, E, x, y)$. Note that $|Q_{x,y}|, |z| \leq \text{Time}_M^E(x)$ and $|\langle x, y \rangle| \leq 2|x|(\text{Time}_M^E(x) + 2)$ if $x \neq \epsilon$. First we show that $\text{Prob}_\mu[\{x \in D \mid h(x) > s(r \cdot |x|)\}] < 1/2r$. For any real number $r > 0$,

$$\begin{aligned} & \text{Prob}_\mu[\{x \in D \mid h(x) > s(r \cdot |x|)\}] \\ & \leq \text{Prob}_\mu[\{x \in D \mid \text{Time}_M^E(x) > p_D(6r \cdot |x|)\}] \\ & \quad + \text{Prob}_\mu[\{x \in D \mid \text{Time}_M^E(x) \leq p_D(6r \cdot |x|) \wedge \forall y \in A_x[\sum_{z \in Q_{x,y}} \text{Time}_N(z) > s(r \cdot |x|)]\}] \\ & \leq \frac{1}{6r} + \text{Prob}_\mu[\{x \in D \mid \forall y \in A_x \exists z \in Q_{x,y}[\text{Time}_N(z) > p_E(|z| \cdot 6r \cdot |\langle x, y \rangle|^2 \cdot p(6r \cdot |\langle x, y \rangle|))]\}] \end{aligned}$$

The latter term can be calculated as follows:

$$\begin{aligned} & \leq \text{Prob}_{\mu'}[\{(x, y) \mid x \in D \wedge \exists z \in Q_{x,y}[\text{Time}_N(z) > p_E(|z| \cdot 6r \cdot |\langle x, y \rangle|^2 \cdot p(6r \cdot |\langle x, y \rangle|))]\}] \\ & = \sum_{n=1}^{\infty} \mu'(\Sigma^n) \cdot \text{Prob}_{\mu'_n}[\{(x, y) \mid x \in D \wedge \exists z \in Q_{x,y}[\text{Time}_N(z) > p_E(|z| \cdot 6rn^2 \cdot p(6rn))]\}] \\ & \leq \text{Prob}_{\mu'}[\{(x, y) \mid q(\langle x, y \rangle) > p(6r \cdot |\langle x, y \rangle|)\}] + \\ & \quad + \sum_{n=1}^{\infty} v(\Sigma^n) \cdot p(6rn) \cdot \text{Prob}_{\nu_n}[\{(x, y) \mid \exists z \in Q_{x,y}[\text{Time}_N(z) > p_E(|z| \cdot 6rn^2 \cdot p(6rn))]\}] \\ & \leq \frac{1}{6r} + \sum_{n=1}^{\infty} p(6rn) \cdot \text{Prob}_{\nu}[\{z \mid \text{Time}_N(z) > p_E(|z| \cdot 6rn^2 \cdot p(6rn))\}] \\ & \leq \frac{1}{6r} + \sum_{n=1}^{\infty} p(6rn) \cdot \frac{1}{6rn^2 \cdot p(6rn)} = \frac{1}{6r} + \frac{1}{\pi^2 r} < \frac{1}{3r}. \end{aligned}$$

A similar argument shows that $\text{Prob}_\mu[\{x \in \overline{D} \mid h(x) > s(r \cdot |x|)\}] < 1/2r$. Thus, h is s on μ -average. \blacksquare

Lemma 4.6 *Let \mathcal{F} be a class of density functions.*

1. $\text{Aver}\langle \text{NP}, * \rangle$ is closed under $\leq_n^{\text{p,av}}$ -reductions [17].
2. $\text{Aver}\langle \text{P}, * \rangle$ is closed under $\leq_T^{\text{p,av}}$ -reductions [5]. Moreover, $\text{Aver}\langle \text{P}, \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^{\text{p,av}} (E, \nu) \text{ and } (E, \nu) \in \text{Aver}\langle \text{P}, * \rangle\}$.
3. $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^{\text{np,av}} (E, \nu) \text{ and } (E, \nu) \in \langle \text{P}, * \rangle\}$.

Proof. The claims (1) and (2) follow from Lemma 4.5. Here we show the claim (3). Clearly if $(D, \mu) \in \text{Aver}\langle \text{NP}, \mathcal{F} \rangle$, then $(D, \mu) \leq_T^{\text{np,av}} (\emptyset, \nu_{\text{st}})$. To see the other direction, we assume that $\mu \in \mathcal{F}$, $(D, \mu) \leq_T^{\text{np,av}} (E, \nu)$ and $(E, \nu) \in \langle \text{P}, * \rangle$. We will show that (D, μ) belongs to $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle$.

Since $(D, \mu) \leq_T^{\text{np,av}} (E, \nu)$, there exist a polynomial p_D and a nondeterministic oracle Turing machine M_D such that $D = L(M_D, E)$ and $\text{Time}_{M_D}^E$ is p_D on μ -average. Let M_E be a deterministic Turing machine and p_E a polynomial such that $E = L(M_E)$ and Time_{M_E} is p_E -time bounded.

Note that $\text{Time}_{M_E}(z) \leq p_E(\text{Time}_{M_D}(x))$ for all y and z in $Q(M_D, E, x, y)$. Now we consider a machine M which nondeterministically simulates the computation of M_D , and whenever M_D makes

a query z , M deterministically simulates M_E on input z . By definition, D is computed by M , and on each computation path y of M on x , the number of steps that M actually takes is bounded by

$$\begin{aligned} &\leq c \cdot \left(\text{Time}_{M_D}^E(x) + \sum_{z \in Q(M_D, E, x, y)} \text{Time}_{M_E}(z) \right) \\ &\leq c \cdot \left(\text{Time}_{M_D}^E(x) + \text{Time}_{M_D}^E(x) \cdot \max_{z \in Q(M_D, E, x, y)} \text{Time}_{M_E}(z) \right) \\ &\leq c \cdot \text{Time}_{M_D}^E(x) \cdot (1 + p_E(\text{Time}_{M_D}^E(x))). \end{aligned}$$

Hence, we can redefine M such that all computation paths of M have the same length. The function Time_M is polynomial on μ -average, since $\text{Time}_{M_D}^E$ is polynomial on μ -average, and the class of functions which are polynomial on μ -average is closed under composition with polynomials [17]. Therefore, $(D, \mu) \in \text{Aver}\langle \text{NP}, \mathcal{F} \rangle$. \blacksquare

Although $\text{Aver}\langle \text{P}, * \rangle$ is closed under polynomial-time many-one reducibility, it is shown in [36] that there exist two problems (A, μ) and (B, ν) such that $(A, \mu) \leq_m^{\text{P}} (B, \nu)$, and $(B, \nu) \in \text{Aver}\langle \text{P}, \text{P-comp} \rangle$, but $(A, \mu) \notin \text{Aver}\langle \text{P}, \text{P-comp} \rangle$.

The following theorem shows that the domination condition is necessary in the definition of reducibilities.

Theorem 4.7 *For every recursive decision problem D not in P , there exist density functions μ_1 and μ_2 in P-comp such that (D, μ_1) and (D, μ_2) are incomparable, i.e., $(D, \mu_1) \not\leq_T^{\text{P}} (D, \mu_2)$ and $(D, \mu_2) \not\leq_T^{\text{P}} (D, \mu_1)$.*

Proof. The proof proceeds by a slow diagonalization technique. Let M_1, M_2, \dots denote a standard enumeration of all the deterministic polynomial-time oracle Turing machines. We identify every non-negative integer i with the $(i+1)$ -th string on the lexicographic order: $\epsilon < 0 < 1 < 00 < 01 < \dots$. We recursively define density functions μ_1, μ_2 and an auxiliary function r on \mathbb{N} .

At step 0, let $r(0) = 0$, $\mu_1(0) \propto 1$, and $\mu_2(0) \propto 1$.

At step n , $n > 0$, the values $r(n), \mu_1(n)$ and $\mu_2(n)$ are defined as follows. On input n , an initial segment of the sequence

$$\begin{aligned} &\langle r(0), D(0), \mu_1(0), \mu_2(0) \rangle, \langle r(1), D(1), \mu_1(1), \mu_2(1) \rangle, \dots \\ &\dots, \langle r(n-1), D(n-1), \mu_1(n-1), \mu_2(n-1) \rangle, \end{aligned}$$

is computed by repeating the same construction procedure for $|n|$ steps. Let m be the largest integer, if any, for which $\langle r(m), D(m), \mu_1(m), \mu_2(m) \rangle$ is completely computed. If none exists, then let $r(n) = 1$, $\mu_1(n) \propto 1$, and $\mu_2(n) \propto 1$. Otherwise, each step depends on whether $r(n)$ is even or odd.

First suppose that $r(m)$ is even and let $i = r(m)/2$. The sequence $M_i^D(0), M_i^D(1), \dots$ is computed by simulating M_i until either more than $|n|$ steps are done, or on some input y , $M_i^D(y)$ queries a string larger than m to oracle D . Let k be the last integer for which the simulation of $M_i^D(k)$ can be completed. If there exists a $y \leq k$ such that either (i) $M_i^D(y) \neq D(y)$, or (ii) M_i^D on input y queries some w satisfying that $\mu_1(y) > \mu_2(w) = 0$, then let $r(n) = r(m) + 1$, $\mu_1(n) = 0$, and $\mu_2(n) \propto 2^{-(|n|+1)}$. Clearly (D, μ_1) is not Turing reducible to (D, μ_2) via M_i . Otherwise, let $r(n) = r(m)$, $\mu_1(n) \propto 2^{-(|n|+1)}$, and $\mu_2(n) = 0$.

If $r(m)$ is odd, then let $i = (r(m) - 1)/2$ and change the roles of μ_1 and μ_2 . It is not hard to see that $\mu_1, \mu_2 \in \text{P-comp}$ since, in each step, we quit the simulations after $|n|$ steps are done.

Claim 1 $\text{range}(r) = \mathbb{N}$, where $\text{range}(r) = \{r(z) | z \in \mathbb{N}\}$.

Proof of Claim. Assume $\text{range}(r) \neq \mathbb{N}$. Take a minimal integer n_0 such that $r(n_0)$ is the maximum in $\text{range}(r)$.

First assume $r(n_0)$ is even. Let $i = r(n_0)/2$ and n be large enough such that $\langle r(n_0), D(n_0), \mu_1(n_0), \mu_2(n_0) \rangle$ is constructed. Note that for every $y > n$ $\mu_1(y) > 0$ and $\mu_2(y) = 0$. For any $y > n$, we have $M_i^D(y) = D(y)$, and $M_i^D(y)$ does not query any string w , where $\mu_1(y) > \mu_2(w) = 0$. Hence, M_i^D computes D on all inputs and it queries only strings shorter than n_0 . This implies that D is in P and this contradicts our assumption. The same argument holds if $r(n_0)$ is odd. ■

Therefore, μ_1 and μ_2 are well-defined. and thus, we complete the proof. ■

5 Complete Problems and Self-Reducibility

The *randomized tiling problem* is the first problem that was shown to be $\langle \text{NP}, \text{P-comp} \rangle$ -complete [22, 23]. In the past decade, several other randomized decision problems have been proven to be $\langle \text{NP}, \text{P-comp} \rangle$ -complete [3, 16, 17, 14, 30]. One of the most useful randomized problem is the *randomized bounded halting problem* $(\text{RH}, \mu_{\text{RH}})$ that is defined as follows: $\text{RH} = \{ \langle i, x, 1^n \rangle \mid M_i \text{ accepts } x \text{ within } n \text{ steps} \}$ and $\mu_{\text{RH}}(\langle i, x, 1^n \rangle) \propto (|i| + 1)^{-2}(|x| + 1)^{-2}(n + 1)^{-2}2^{-(|i|+|x|)}$, where M_0, M_1, \dots is a fixed enumeration of all nondeterministic Turing machines. A proof that $(\text{RH}, \mu_{\text{RH}})$ is $\langle \text{NP}, \text{P-comp} \rangle$ -complete can be found in [14, 17, 5]. We note that each of the complete problems is a pair of an NP-complete set and a *natural* density function. However, a randomized satisfiability problem and a randomized graph 3-colorability problem are not $\langle \text{NP}, \text{P-comp} \rangle$ -complete for reasonable natural density functions [10, 39]. In [36], Wang and Belanger show that for every set D , if D is \leq_m^{P} -hard for NP, then there exists a density function μ such that (D, μ) is \leq_m^{P} -hard for $\langle \text{NP}, \text{P-comp} \rangle$. However, it is not known whether every NP-complete set D has a density function μ such that (D, μ) is $\langle \text{NP}, \text{P-comp} \rangle$ -complete.

The (*Turing*) *self-reducibility* has been introduced into worst-case complexity theory by Meyer and Paterson [25]. All known NP-complete problems are self-reducible and every self-reducible set belongs to PSPACE. It is natural to ask whether the notion of self-reducibility is applicable to randomized decision problems.

Definition 5.1 A polynomial-time computable partial order $<$ is *OK* if there exists a polynomial p such that

1. every strictly descending chain is finite and is polynomial in the length of its maximum element, i.e., if $x_k < x_{k-1} < \dots < x_2 < x_1$ is a descending chain starting from x_1 , then $k \leq p(|x_1|)$, and
2. for every x and y , $x < y$ implies $|x| \leq p(|y|)$.

Definition 5.2 A randomized decision problem (D, μ) is (*Turing*) *self-reducible* if there exists an OK partial order and a deterministic oracle Turing machine M such that $(D, \mu) \leq_T^{\text{P}} (D, \mu)$ via M , and for every input x , all query strings in the computation of M on input x are smaller than x with respect to the partial order.

Clearly every randomized problem in $\langle \text{P}, * \rangle$ is self-reducible, and every self-reducible randomized problem is in $\langle \text{PSPACE}, * \rangle$. Moreover, the set of all self-reducible problems is closed under

polynomial isomorphism, i.e., if $(D, \mu) \cong^p (E, \nu)$ and (E, ν) is self-reducible, then (D, μ) is self-reducible.

One of the classical self-reducible NP-complete problems is the *satisfiability problem*, SAT. However, we do not know a simple density function μ such that (SAT, μ) is $\langle \text{NP}, \text{P-comp} \rangle$ -complete. Franco and Paull [10] show that SAT with a natural probability distribution on formulas is in $\text{Aver}(\text{P}, *)$. So, we consider, as a canonical example, the randomized bounded halting problem again. More precisely, we consider the k -th level of the randomized bounded halting problem $(\text{RH}^k, \mu_{\text{RH}})$ that is defined as follows. Assume that M_0, M_1, \dots is an effective enumeration of all nondeterministic oracle Turing machines, and define $\text{RH}(A) = \{ \langle i, x, 1^n \rangle \mid M_i^A \text{ accepts } x \text{ in less than } n \text{ steps} \}$. Let $\text{RH}^1 = \text{RH}(\emptyset)$ and $\text{RH}^{k+1} = \text{RH}(\text{RH}^k)$ for $k \geq 1$. We note that RH^k is Σ_k^{P} -complete.

Lemma 5.3 *For any $k > 0$, $(\text{RH}^k, \mu_{\text{RH}})$ is $\langle \Sigma_k^{\text{P}}, \text{P-comp} \rangle$ -complete.*

Proof. The case $k = 1$ is shown in [14, 17, 5]. Now let $k > 1$. The proof follows [36]. For any set $D \in \Sigma_k^{\text{P}}$ and any density function $\mu \in \text{P-comp}$, we will show that $(D, \mu) \leq_m^{\text{P}} (\text{RH}^k, \mu_{\text{RH}})$. Without loss of generality, we assume that $|\mu(x)| \leq 2|x| + 4$ (see Lemma 1.6 in [17]). Since RH^{k-1} is Σ_{k-1}^{P} -complete, there exists a polynomial-time nondeterministic oracle Turing machine M such that $D = L(M, \text{RH}^{k-1})$. Let a function g be defined as follows: on input x , g deterministically computes a minimal string y such that $\mu^*(x^-) < 0.y1 \leq \mu^*(x)$. Now consider a machine N : on input y , N first computes a string x that $y = g(x)$, by a binary search in time polynomial in $|x|$, and if x exists, then N nondeterministically simulates M on input x ; otherwise, N simply rejects x . Note that $|g(x)| \leq q(|x|)$ for some absolute polynomial q , and that $\mu(x) < 2^{-|g(x)|}$. Now let i be an index such that $L(M_i) = L(N)$. Let p be a polynomial time bound of M_i . The desired reduction f is now defined as $f(x) = \langle i, g(x), 1^{p(|x|)} \rangle$. Note that f is one-one and reduces D to RH^k . It suffices to check that f satisfies the domination condition. Since i is a constant in the reduction, it follows that

$$\mu_{\text{RH}}(\langle i, g(x), 1^{p(|x|)} \rangle) = \frac{c}{(|i| + 1)^2 (|g(x)| + 1)^2 (p(n) + 1)^2 \cdot 2^{|i|}} 2^{-|g(x)|} \geq \frac{1}{s(|x|)} \cdot \mu(x),$$

where s is a polynomial such that $c \cdot s(n) \geq (|i| + 1)^2 (q(n) + 1)^2 (p(n) + 1)^2 \cdot 2^{|i|}$. Hence, $(D, \mu) \leq_m^{\text{P}} (\text{RH}^k, \mu_{\text{RH}})$. \blacksquare

Theorem 5.4 *For each $k > 0$, $(\text{RH}^k, \mu_{\text{RH}})$ is self-reducible.*

Proof. Without loss of generality, we assume that nondeterministic Turing machines have transition functions with exactly two nondeterministic choices. Consider the following encoding of nondeterministic oracle Turing machines M : let $\langle M \rangle$ be the code of a set of finite states (each state q_i is simply encoded by $\langle 0, i \rangle$), the initial state, the final states and the transition function, where the transition function is given by a table in which each row consists of 4 quintuples; $(q_i, 0, q_j^1, t^1, s^1)$, $(q_i, 0, q_j^2, t^2, s^2)$, $(q_i, 1, q_j^3, t^3, s^3)$ and $(q_i, 1, q_j^4, t^4, s^4)$.

Let $g(M, x, t)$ be 0 if $t \notin \{0, 1\}$; otherwise, let $g(M, x, t)$ be a code of a nondeterministic Turing machine M' that simulates M on input x , but the first nondeterministic step of M is done (deterministically) depending on the value of t . The code of M' is obtained from the code of M with a new initial state and one additional row in the transition table, which describes the first step. We assume that g is one-one, computable in polynomial time, and $|g(M, x, t)| \leq |M| + c \log |M| + c$ for some constant c .

We define $<$ as OK partial order on strings of the form $\langle M, \langle x, s \rangle, 1^n \rangle$ as follows: $\langle M', \langle x', s' \rangle, 1^m \rangle < \langle M, \langle x, s \rangle, 1^n \rangle$ if $x' = x$, $s' = st$, where $t \in \{0, 1\}$, $M' = g(M, x, t)$, and $|s'| + m = |s| + n$, i.e., $m = n - 1$. Every string of the form $\langle M, \langle x, \epsilon \rangle, 1^n \rangle$ is the largest string in the order $<$. Each chain starting from $\langle M, \langle x, \epsilon \rangle, 1^n \rangle$ has at most length n .

Now let N be an oracle Turing machine which, on input $F = \langle M, \langle x, s \rangle, 1^n \rangle$ works as follows: if $n = 0$, then N accepts F exactly when M is in an accepting configuration after $|s|$ deterministic steps; otherwise, N computes two strings $F^0 = \langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle$ and $F^1 = \langle g(M, s, 1), \langle x, s1 \rangle, 1^{n-1} \rangle$ and accepts exactly when either F^0 or F^1 (or both) appears to be in the oracle.

Clearly N reduces RH^k to RH^k by querying only strings which are smaller than input with respect to $<$. It remains to show that, for some polynomial q , $\mu_{\text{RH}}(F') \geq \sum_F \mu_{\text{RH}}(F)/q(|F|)$, where F ranges over all strings which are reduced to F' by N . From the construction of the strings F^0 and F^1 , it follows immediately that they are only asked on input F . Therefore, it suffices to show that $\mu_{\text{RH}}(F^0) \geq \mu_{\text{RH}}(F)/q(|F^0|)$. Let $p(\langle M, \langle x, s \rangle, 1^n \rangle) = (|M| + 1)^2(|\langle x, s \rangle| + 1)^2(n + 1)^2$. Thus,

$$\begin{aligned} \mu_{\text{RH}}(F^0) &= \mu_{\text{RH}}(\langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle) \\ &\geq \frac{c}{p(\langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle)} 2^{-(|g(M, s, 0)| + |\langle x, s0 \rangle|)} \\ &\geq \frac{c}{p(\langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle)} 2^{-(|M| + c \log |M| + |\langle x, s \rangle| + c + 1)} \\ &\geq \frac{c}{2^{c+1} \cdot p(\langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle) \cdot |M|^c} 2^{-(|M| + |\langle x, s \rangle|)} \\ &\geq \frac{\mu_{\text{RH}}(\langle M, \langle x, s \rangle, 1^n \rangle)}{q(\langle g(M, s, 0), \langle x, s0 \rangle, 1^{n-1} \rangle)} = \frac{\mu_{\text{RH}}(F)}{q(|F^0|)} \end{aligned}$$

for some polynomial q . ■

Belanger and Wang [3] show that the following $\langle \text{NP}, \text{P-comp} \rangle$ -complete problems are polynomially isomorphic to each other: the *randomized bounded halting problem* [14, 17, 5], the *randomized tiling problem* [22, 23, 17], the *randomized Post correspondence problem* [17], and the *randomized word problem for Thue systems* [3]. From Theorem 5.4, we immediately conclude:

Corollary 5.5 *The following randomized decision problems are all self-reducible: the randomized tiling problem, the randomized Post correspondence problem, and the randomized word problem for Thue systems.*

We note that if every pair of $\langle \text{NP}, \text{P-comp} \rangle$ -complete problems is polynomially isomorphic, then every $\langle \text{NP}, \text{P-comp} \rangle$ -complete problem is self-reducible.

6 Average Polynomial Time Hierarchy

The Meyer-Stockmeyer polynomial-time hierarchy is introduced in [26] based on polynomial-time deterministic and nondeterministic Turing reducibilities and is a central concept in worst-case complexity theory. Here, Turing reductions are used to define new classes over P and NP in an analogous way to the Kleene arithmetical hierarchy.

The theory of average NP -completeness can be similarly generalized to an arbitrary level of the Meyer-Stockmeyer polynomial-time hierarchy by using Turing reducibilities among randomized

decision problems. We have already seen in Section 5 that all the classes $\langle \Sigma_k^P, \text{P-comp} \rangle$, $k > 0$, have \leq_m^P -complete sets. It is natural to ask whether, e.g., $\langle \Sigma_k^P, \text{P-comp} \rangle$ is contained in an average-case version of Δ_k^P or not. To answer this question, we introduce a notion of an *average polynomial-time hierarchy*, which is based on average polynomial-time Turing reducibilities discussed in Section 4, in analogy with the Meyer-Stockmeyer polynomial-time hierarchy.

First we define a relativization of the average-case complexity classes, $\text{Aver}\langle \text{P}, \mathcal{F} \rangle$ and $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle$, to an oracle (E, ν) .

Definition 6.1 Let (E, ν) be a randomized decision problem and let \mathcal{F} be a set of density functions.

1. $\langle \text{P}, \mathcal{F} \rangle^{(E, \nu)} = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^P (E, \nu)\}$.
2. $\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)} = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^{\text{NP}} (E, \nu)\}$.
3. $\text{Aver}\langle \text{P}, \mathcal{F} \rangle^{(E, \nu)} = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^{\text{P, av}} (E, \nu)\}$.
4. $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)} = \{(D, \mu) \mid \mu \in \mathcal{F}, (D, \mu) \leq_T^{\text{NP, av}} (E, \nu)\}$.

From the definitions of Turing reducibilities, it immediately follows that, for any randomized problem (E, ν) , $\text{Aver}\langle \text{P}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{P}, \mathcal{F} \rangle^{(E, \nu)}$, $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}$, $\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)} \subseteq \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}$, and $\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)} \subseteq \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}$.

Definition 6.2 Let \mathcal{C} be a set of randomized decision problems and let \mathcal{F} be a set of density functions.

1. $\langle \text{P}, \mathcal{F} \rangle^{\mathcal{C}} = \{(D, \mu) \mid \exists (E, \nu) \in \mathcal{C} [(D, \mu) \in \langle \text{P}, \mathcal{F} \rangle^{(E, \nu)}]\}$.
2. $\langle \text{NP}, \mathcal{F} \rangle^{\mathcal{C}} = \{(D, \mu) \mid \exists (E, \nu) \in \mathcal{C} [(D, \mu) \in \langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}]\}$.
3. $\text{Aver}\langle \text{P}, \mathcal{F} \rangle^{\mathcal{C}} = \{(D, \mu) \mid \exists (E, \nu) \in \mathcal{C} [(D, \mu) \in \text{Aver}\langle \text{P}, \mathcal{F} \rangle^{(E, \nu)}]\}$.
4. $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{\mathcal{C}} = \{(D, \mu) \mid \exists (E, \nu) \in \mathcal{C} [(D, \mu) \in \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)}]\}$.

Lemma 4.6 immediately yields the following closure properties.

Lemma 6.3 Let \mathcal{F} be a class of density functions.

1. $\text{Aver}\langle \text{P}, \mathcal{F} \rangle^{\text{Aver}\langle \text{P}, * \rangle} = \text{Aver}\langle \text{P}, \mathcal{F} \rangle^{(\text{P}, *)} = \text{Aver}\langle \text{P}, \mathcal{F} \rangle$.
2. $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(\text{P}, *)} = \text{Aver}\langle \text{NP}, \mathcal{F} \rangle$.

Up to now, it is unknown whether “in the unrelativized world” $\langle \text{NP}, \text{P-comp} \rangle \subseteq \text{Aver}\langle \text{P}, \text{P-comp} \rangle$ or not. Ben-David et al. [5] show that if $\langle \text{NP}, \text{P-comp} \rangle \subseteq \text{Aver}\langle \text{P}, \text{P-comp} \rangle$, then $\text{E} = \text{NE}$. Here we give two, contradicting relativized results, an inclusion and a separation.

Theorem 6.4 *There exist randomized decision problems (A, μ) and (B, ν) such that*

$$\begin{aligned} \langle \text{NP}, \text{P-comp} \rangle^{(A, \mu)} &\subseteq \text{Aver}\langle \text{P}, \text{P-comp} \rangle^{(A, \mu)}, \text{ and} \\ \langle \text{NP}, \text{P-comp} \rangle^{(B, \nu)} &\not\subseteq \text{Aver}\langle \text{P}, \text{P-comp} \rangle^{(B, \nu)}. \end{aligned}$$

Proof. The desired oracle (A, μ) is a special version of the randomized halting problem. Let $\langle N, x, 1^t \rangle$ be in A if the nondeterministic Turing machine N with oracle A accepts x in less than t steps. Note that this is a valid definition since N can make only queries smaller than $1^t < \langle N, x, 1^t \rangle$. Define $\mu(\langle N, x, 1^t \rangle) \propto (|N| + 1)^{-2}(|x| + 1)^{-2}(t + 1)^{-2} \cdot 2^{-(|N|+|x|)}$.

Recall that the randomized halting problem is \leq_m^P -hard for $\langle \text{NP}, \text{P-comp} \rangle$. By exactly the same argument, it follows that (A, μ) is \leq_m^P -hard for $\langle \text{NP}, \text{P-comp} \rangle^{(A, \mu)}$. That is, it actually holds that $\langle \text{NP}, \text{P-comp} \rangle^{(A, \mu)} = \langle \text{P}, \text{P-comp} \rangle^{(A, \mu)}$.

Let B be the oracle set used by Baker et al. [1] to separate P from NP using the tally set $L(B) = \{0^n \mid \exists y [|y| = n \wedge y \in B]\}$ in $\text{NP}^B - \text{P}^B$ and let $\nu(x) \propto (|x| + 1)^{-2}$ if $x \in B$; otherwise, $(|x| + 1)^{-2} \cdot 2^{-|x|}$. Consider the randomized problem $(L(B), \eta)$, where $\eta(x) \propto (|x| + 1)^{-2}$, if $x \in \{0\}^*$, and 0 otherwise. Clearly $(L(B), \eta)$ is in $\langle \text{NP}, \text{P-comp} \rangle^{(B, \nu)}$, and thus, it is in $\text{Aver} \langle \text{NP}, \text{P-comp} \rangle^{(B, \nu)}$. Now assume $(L(B), \eta)$ is in $\text{Aver} \langle \text{P}, \text{P-comp} \rangle^{(B, \nu)}$. There exist a deterministic Turing machine M witnesses that $(L(B), \eta) \leq_T^{\text{NP, av}} (B, \nu)$. Since M is polynomial-time bounded on η -average, it yields that $L(B) \in \text{P}^B$. This is a contradiction against the fact that $L(B) \notin \text{P}^B$. ■

We now give a definition of an *average polynomial-time hierarchy*, which is an average-case analogue of the Meyer-Stockmeyer polynomial-time hierarchy in worst-case complexity theory.

Definition 6.5 Let $k > 1$ and let \mathcal{F} be a class of density functions.

1. $\text{Aver} \langle \Delta_0^P, \mathcal{F} \rangle = \text{Aver} \langle \Sigma_0^P, \mathcal{F} \rangle = \text{Aver} \langle \text{P}, \mathcal{F} \rangle$.
2. $\text{Aver} \langle \Delta_k^P, \mathcal{F} \rangle = \text{Aver} \langle \text{P}, \mathcal{F} \rangle^{\text{Aver} \langle \Sigma_{k-1}^P, * \rangle}$.
3. $\text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle = \text{Aver} \langle \text{NP}, \mathcal{F} \rangle^{\text{Aver} \langle \Sigma_{k-1}^P, * \rangle}$.
4. $\text{Aver} \langle \text{PH}, \mathcal{F} \rangle = \bigcup_{k \geq 0} \text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle$.

We here remark that oracle sets used in the definition are not restricted to the class $\text{Aver} \langle \Sigma_{k-1}^P, \mathcal{F} \rangle$, since the domination condition of Turing reducibility already puts a constraint on the complexity of the density function of the oracle.

Note also that Lemma 4.6(1) can be easily extended to the class $\text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle$, namely, $\text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle$ is closed under $\leq_m^{\text{P, av}}$ -reductions.

The following two propositions give reasonable evidence that the above defined average polynomial-time hierarchy has a structure similar to that of the worst-case polynomial-time hierarchy.

Proposition 6.6 Let $k \geq 1$ and let \mathcal{F} be any class of density functions.

1. $\text{Aver} \langle \Delta_k^P, \mathcal{F} \rangle \subseteq \text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle$.
2. $\text{Aver} \langle \Sigma_k^P, \mathcal{F} \rangle \subseteq \text{Aver} \langle \Delta_{k+1}^P, \mathcal{F} \rangle$.

Proof. The proposition follows immediately from Definition 6.5. ■

Proposition 6.7 Let \mathcal{F} be any class of density functions, then $\text{Aver} \langle \text{BPP}, \mathcal{F} \rangle \subseteq \text{Aver} \langle \Sigma_2^P, \mathcal{F} \rangle$.

Proof. Let (D, μ) be an arbitrary problem in $\text{Aver} \langle \text{BPP}, \mathcal{F} \rangle$. By Proposition 2.7, there exists a polynomial p , a p -time bounded deterministic Turing machine M , and a function f such that f is computable in time polynomial on μ -average, and for all x , there are more than $2^{f(x)-|x|}$ strings w of length $f(x)$ such that $x \in D$ if and only if M accepts $\langle x, w \rangle$.

We define a nondeterministic oracle Turing machine M_0 with an oracle set E as follows. On input x , M_0 first computes the value $f(x)$, then it guesses a non-negative integer m , and distinct strings $u_1, \dots, u_{f(x)}$ and w of length $f(x)$, and queries the string $\langle x, u_1 \cdots u_{f(x)}, w \rangle$ to the oracle E . If the string is in E , then M_0 simulates M on input $\langle x, w \rangle$; otherwise, it rejects the input.

The desired oracle E is defined by the nondeterministic machine M_1 that works as follows: on input $\langle x, u_1 \cdots u_m, w \rangle$, M_1 guesses v of length m and x_1, \dots, x_m of length $|x|$, and checks if $M(\langle x_i, u_i \oplus v \rangle) \neq M(\langle x_i, w \rangle)$, for all $i \leq m$, where $u \oplus v$ denotes the bitwise addition of u and v . If this is true, then M_0 accepts the input; otherwise, M_0 rejects the input. Then it follows that $D = L(M_0, E)$ (for a proof, see [2], pp. 170-173).

Note that M_1 is polynomial-time bounded and therefore, $E \in \text{NP}$. Furthermore, M_0 is polynomial-time bounded on μ -average f -time bounded, since f is computable in time polynomial on μ -average.

To show that $(D, \mu) \leq_T^{\text{np,av}} (E, \nu)$, we define a density function ν on E as follows. Take the density function μ' induced from μ , M_0 and E and define $\nu(z) = \text{Prob}_{\mu'}[\{(x, y) \mid z \in Q(M_0, E, x, y)\}]$. Hence, (D, μ) is in $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{(E, \nu)} \subseteq \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{\text{Aver}\langle \text{NP}, * \rangle} = \text{Aver}\langle \Sigma_2^{\text{P}}, \mathcal{F} \rangle$. ■

The definition of the average polynomial-time hierarchy implies that the average-case complexity classes of the hierarchy contain the associated worst-case complexity classes with arbitrary classes of density functions.

Lemma 6.8 *Let $k \geq 1$ and let \mathcal{F} be any class of density functions.*

1. $\langle \Delta_k^{\text{P}}, \mathcal{F} \rangle \subseteq \langle \text{P}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle} \subseteq \text{Aver}\langle \text{P}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle} \subseteq \text{Aver}\langle \Delta_k^{\text{P}}, \mathcal{F} \rangle$.
2. $\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle \subseteq \langle \text{NP}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle} \subseteq \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle} \subseteq \text{Aver}\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$.

Proof. We show the claim (2) since the proof of (1) is analogous. The proof proceeds by induction on k .

The case $k = 1$ is obvious. Let $k > 1$ and assume that $(A, \mu) \in \langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$. There is a set $B \in \Sigma_{k-1}^{\text{P}}$, a polynomial p , and a nondeterministic oracle Turing machine M , which is p -time bounded, such that M computes A with oracle B .

Take μ' induced from μ , M and B as in Definition 4.3. A density function ν on the oracle E is given by $\nu(z) = \text{Prob}_{\mu'}[\{(x, y) \mid z \in Q(M, B, x, y)\}]$. Therefore, $(A, \mu) \leq_T^{\text{np}} (B, \nu)$, where $(B, \nu) \in \langle \Sigma_{k-1}^{\text{P}}, * \rangle$. This implies $(A, \mu) \in \langle \text{NP}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle}$. Clearly it holds that $(A, \mu) \in \text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle}$. Since, by induction hypothesis, $\langle \Sigma_{k-1}^{\text{P}}, * \rangle$ is included in $\text{Aver}\langle \Sigma_{k-1}^{\text{P}}, * \rangle$, it yields that $(A, \mu) \in \text{Aver}\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$. ■

Lemma 3.10 can be extended into an arbitrary level of the average polynomial-time hierarchy. We first see a key lemma.

Lemma 6.9 *Let $k \geq 0$ and assume that $(A, \mu) \leq_T^{\text{np,av}} (B, \nu)$ and $(B, \nu) \in \text{Aver}\langle \Sigma_k^{\text{P}}, * \rangle$. For any set S and any polynomial q , there exist sets $C_0 \in \Sigma_k^{\text{P}}$, $C_1 \in \Pi_k^{\text{P}}$ and S' such that $A \cap S' \subseteq C_0 \subseteq A$, $\overline{A} \cap S' \subseteq C_1 \subseteq \overline{A}$ and $\mu(S^n) - \mu(S'^n) \leq 1/q(n)$ for all $n \in \mathbb{N}$.*

Proof. The proof proceeds by induction on k . The base case $k = 0$ follows from Lemma 3.10.

Let $k \geq 1$ and assume that $(A, \mu) \in \text{Aver}\langle \Sigma_{k+1}^{\text{P}}, * \rangle$. By definition, there is a nondeterministic Turing machine M and a randomized decision problem $(B, \nu) \in \text{Aver}\langle \Sigma_k^{\text{P}}, * \rangle$ such that $(A, \mu) \leq_T^{\text{np,av}} (B, \nu)$ via M . Let p be a polynomial such that $\text{Prob}_{\mu}[\{x \mid \text{Time}_M^B(x) > p(|x| \cdot m)\}] < 1/m$ for any positive real number m .

Let μ' be the density function induced from μ , M and B . Furthermore, let v be a density function, such that $\mu' \preceq^{p, \mu'} v$ and $\nu(z) \geq \text{Prob}_v[\{(x, y) \mid z \in Q(M, B, x, y)\}]$. Since $\mu' \preceq^{p, \mu'} v$, there exists a function f , which is r on μ' -average, such that $v(x, y) \cdot f(x, y) \geq \mu'(x, y)$, where r is a polynomial.

Claim 2 *There exists a nondeterministic Turing machine M' and a randomized problem (B', ν') in $\text{Aver}\langle \Sigma_k^p, \mathcal{F} \rangle$ such that $(A, \mu) \leq_T^{\text{np, av}} (B', \nu')$ via M' , and all strings queried by M' with oracle B' on input x are of length greater than $|x|$.*

Proof. Define $B' = \{z01^n \mid z \in B\}$ and

$$\nu'(w) = \begin{cases} \nu(z) \cdot \frac{\text{Prob}_{\mu'}[\{(x, y) \mid |x|=n \wedge z \in Q(M, B, x, y)\}]}{\text{Prob}_{\mu'}[\{(x, y) \mid z \in Q(M, B, x, y)\}]} & \text{if } w = z01^n \text{ for some } z, n, \\ 0 & \text{otherwise.} \end{cases}$$

Define a new oracle Turing machine M' which works as follows: on input x , simulate M on input x , and whenever M queries a string z , M' queries $z01^{|x|}$ to the oracle. It is easy to see that $(A, \mu) \leq_T^{\text{np, av}} (B', \nu')$ via M' . We next show that (B', ν') is in $\text{Aver}\langle \Sigma_k^p, * \rangle$. Note that (B', ν') is average polynomial-time many-one reducible to (B, ν) . Since $\text{Aver}\langle \Sigma_k^p, * \rangle$ is closed under \leq_m^p -reduction, (B', ν') is in $\text{Aver}\langle \Sigma_k^p, * \rangle$. \blacksquare

Therefore, without loss of generality, we assume that M queries only strings of length greater than length of inputs. For each $n \in \mathbb{N}$, it yields that

$$\text{Prob}_\mu[\{x \mid \text{Time}_M^B(x) > p(|x| \cdot 3q(n))\}] < 1/3q(n), \text{ and}$$

$$\text{Prob}_{\mu'}[\{(x, y) \mid f(x, y) > r(|\langle x, y \rangle| \cdot 3q(n))\}] < 1/3q(n).$$

Let $\text{Flip}(x)$ be $\text{Acc}(M, B, x)$ if $x \in D$, or else let it be $\text{Rej}(M, B, x)$. We define a set T_n by

$$T_n = \{(x, y) \mid x \in S^n \text{ and } \text{Time}_M^B(x) \leq p(|x| \cdot 3q(n)) \text{ and} \\ y \in \text{Flip}(x) \text{ and } \mu'(x, y) \leq r(|\langle x, y \rangle| \cdot 3q(n)) \cdot v(x, y)\},$$

and let $T = \bigcup_{n>0} T_n$. It is easy to see that $\mu(S^n) - \mu'(T_n) \leq 2/3q(n)$. For any pair (x, y) in T ,

$$\begin{aligned} \mu'(x, y) &\leq r(|\langle x, y \rangle| \cdot 3q(n)) \cdot v(x, y) \\ &\leq r((2n + p(n \cdot 3q(n) + 1)) \cdot 3q(n)) \cdot v(x, y) \end{aligned}$$

since $|y| \leq p(|x| \cdot 3q(n))$, and thus, it holds that $\mu'(x, y) \leq s(n) \cdot v(x, y)$ for some polynomial s .

Now let $Z = \{z \mid \exists(x, y) \in T[z \in Q(M, B, x, y)]\}$. Recall that (B, ν) is in $\text{Aver}\langle \Sigma_k^p, * \rangle$. Hence, it follows, by induction hypothesis, that, for any polynomial l , there exist a subset Z' of Z and sets $C'_0 \in \Sigma_k^p$ and $C'_1 \in \Pi_k^p$ such that $B \cap Z' \subseteq C'_0 \subseteq B$, $\overline{B} \cap Z' \subseteq C'_1 \subseteq \overline{B}$ and $\nu(Z^n) - \nu(Z'^n) \leq 1/l(n)$. Denote by Z_n (resp. Z'_n) the set of strings in Z (resp. Z') whose length is between n and $p(n \cdot 3q(n))$. Now choose $l(n) = 3q(n) \cdot s(n) \cdot p(n \cdot 3q(n))$ for all $n \in \mathbb{N}$. Then, it yields that $\nu(Z_n) - \nu(Z'_n) \leq 1/3q(n)s(n)$.

Let $T' = \{(x, y) \in T \mid Q(M, B, x, y) \subseteq Z'\}$. Note that, for all $n \in \mathbb{N}$,

$$T_n - T'_n \subseteq \{(x, y) \in T_n \mid Q(M, B, x, y) \cap (Z_n - Z'_n) \neq \emptyset\}.$$

Hence, it holds that $\nu(Z_n - Z'_n) \geq \nu(T_n - T'_n)$ for all $n \in \mathbb{N}$. Then, it follows that

$$1/3q(n) \geq (\nu(Z_n) - \nu(Z'_n)) \cdot s(n) \geq (\nu(T_n) - \nu(T'_n)) \cdot s(n) \geq \mu'(T_n) - \mu'(T'_n).$$

The desired set S' is defined as $S' = \{x \mid \exists y[(x, y) \in T']\}$. Note that $\mu(S'_n) \geq \mu'(T'_n)$ holds. It immediately follows from our definition that $S' \subseteq S$ and $\mu(S^n) - \mu(S'^n) \leq 1/q(n)$ for all $n \in \mathbb{N}$.

Let M_0 be an oracle Turing machine with oracle X defined as follows. On input x , simulate M on x in time $p(|x| \cdot 3q(|x|))$, and whenever M queries a string z , M_0 queries $\langle 0, z \rangle$ and $\langle 1, z \rangle$ to its oracle X . If $\langle 0, z \rangle \in X$ and $\langle 1, z \rangle \notin X$, continue the simulation with the oracle answer “yes”, and if $\langle 0, z \rangle \notin X$ and $\langle 1, z \rangle \in X$, then continue the simulation with the oracle answer “no”; otherwise, immediately reject the input x . The machine M_0 accepts x only when M halts and accepts it. Similarly we define a machine M_1 by interchanging the oracle answers and accepts the input if M halts in time $p(|x| \cdot 3q(|x|))$ and rejects it. Now let $C_0 = L(M_0, C'_0 \oplus C'_1)$ and $C_1 = L(M_1, C'_0 \oplus C'_1)$. By definition of the oracle machines M_0 and M_1 , it yields that $A \cap S' \subseteq C_0 \subseteq A$ and $\overline{A} \cap S' \subseteq C_1 \subseteq \overline{A}$. \blacksquare

Proposition 6.10 *For $k \geq 1$, $\text{Aver}\langle \Delta_k^P, * \rangle$ and $\text{Aver}\langle \Sigma_k^P, * \rangle$ have the sparse interpolation property.*

Proof. We show the case $\text{Aver}\langle \Sigma_k^P, * \rangle$ here. The case $k = 1$ follows from Lemma 3.10. Let $k \geq 2$ and assume that $(A, \mu_{S,q}) \in \text{Aver}\langle \Sigma_k^P, * \rangle$ for a sparse set S and a polynomial q . It follows from Lemma 6.9 that there exists a set $C \in \Sigma_k^P$ and a subset S' of S such that $A \cap S' \subseteq C \subseteq A$ and $\mu_{S,q}(S^n) - \mu_{S,q}(S'^n) \leq 1/2q(n)$ for all $n \in \mathbb{N}$. It suffices to show that $S' = S$. Assume that there exists a string $x \in S - S'$. Let $n = |x|$. Since $\mu_{S,q}(x) \geq 1/q(|x|)$,

$$\frac{1}{q(n)} \leq \mu_{S,q}(S^n) - \mu_{S,q}(S'^n) \leq \frac{1}{2q(n)}.$$

This is a contradiction. \blacksquare

The average polynomial-time hierarchy allows to construct an average-version of the *high and low hierarchy in NP* [31] to refine the structure within NP. It may be possible that some NP-complete problems with natural distributions which are unknown to be either in $\text{Aver}\langle P, * \rangle$ or in $\langle \text{NP}, \text{P-comp} \rangle$ -complete fall into a “low (or high) hierarchy in $\text{Aver}\langle \text{NP}, \text{P-comp} \rangle$.”

If we return to Levin’s fundamental question of whether or not $\langle \text{NP}, \text{P-comp} \rangle \subseteq \text{Aver}\langle P, * \rangle$, we can now raise a more general question of whether $\langle \Sigma_k^P, \text{P-comp} \rangle \subseteq \text{Aver}\langle \Delta_k^P, * \rangle$ holds or not. However, to answer this question turns out to be very hard, since the following claim closes the gap between average-case and worst-case.

Theorem 6.11 *For any $k \geq 1$, $\langle \Sigma_k^P \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \Delta_k^P, * \rangle$ if and only if $\Sigma_k^P \cap \text{TALLY} \subseteq \Delta_k^P$.*

Proof. It suffices to prove the “only if” part of the theorem. Let $\mu(x) = \frac{6}{\pi^2(|x|+1)^2}$ if $x \in \{0\}^*$, and 0 otherwise. Clearly μ is in L-comp. Suppose that $\langle \Sigma_k^P \cap \text{TALLY}, \text{L-comp} \rangle \subseteq \text{Aver}\langle \Delta_k^P, * \rangle$, and $A \in \Sigma_k^P \cap \text{TALLY}$. Since $(A, \mu) \in \text{Aver}\langle \Delta_k^P, * \rangle$, Proposition 6.10 shows the existence of a set $B \in \Delta_k^P$ such that $A \cap \{0\}^* = B \cap \{0\}^*$. Now define $B' = B \cap \{0\}^*$. Since $A \subseteq \{0\}^*$, it follows that $A = B' \in \Delta_k^P$. \blacksquare

Thus it seems likely that $\langle \Sigma_k^P, \text{P-comp} \rangle$ is not a subset of $\text{Aver}\langle \Delta_k^P, * \rangle$, since otherwise every tally Σ_k^P set is in Δ_k^P .

Corollary 6.12 *Let $k \geq 1$. If $\text{Aver}\langle \Delta_k^P, \text{P-comp} \rangle = \text{Aver}\langle \Sigma_k^P, \text{P-comp} \rangle$, then $\Sigma_k^P \cap \text{TALLY} \subseteq \Delta_k^P$.*

7 Real Polynomial Time Computability

This section establishes a direct link to the classical framework of worst-case complexity theory. This link casts a light on the essential role of average-case analysis in the study of worst-case complexity.

As we have seen, average-case analysis is very sensitive to the selection of distributions. For example, fast decreasing density functions help average polynomial-time bounded machines to solve hard sets; however, there exist sets in EXP that are not solvable in average polynomial-time if we choose, e.g., a density function ν defined as $\nu(x) \propto (|x| + 1)^{-2}$ for $x \in \{0\}^*$, and 0 otherwise. This approach toward average-case analysis does not capture an important feature of average-case analysis. To see this feature, we try to abstract a notion of “rare instances” under any “reasonable” distribution and to make it independent from individual distributions only to rely on a general property.

Observe that the class P is the largest class of sets which are computable in average polynomial-time with respect to all density functions [24]. We will consider the class of sets which are deterministically computable in average polynomial-time for every density function in P-comp, and refer to it as a “real P over P-comp.” First we introduce a more general notion of “real \mathcal{C} over \mathcal{F} .”

Definition 7.1 Let \mathcal{C} be a complexity class and let \mathcal{F} be a class of density functions. The *real \mathcal{C} over \mathcal{F}* , denoted by $\mathcal{C}_{\mathcal{F}}$, is the class of languages D such that $(D, \mu) \in \text{Aver}\langle \mathcal{C}, * \rangle$ for every $\mu \in \mathcal{F}$.

This new definition formalizes a significant property of the associated average-case complexity classes. The next proposition indicates the importance of those new classes.

Proposition 7.2 Let $\langle \mathcal{C}, \mathcal{F} \rangle$ and $\text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$ be any randomized and average-case complexity classes, respectively. Then, $\mathcal{C} \subseteq \mathcal{D}_{\mathcal{F}}$ if and only if $\langle \mathcal{C}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \mathcal{D}, \mathcal{F} \rangle$.

Proof. Assume that $\mathcal{C} \subseteq \mathcal{D}_{\mathcal{F}}$ and (A, μ) is in $\langle \mathcal{C}, \mathcal{F} \rangle$. Since A belongs to $\mathcal{D}_{\mathcal{F}}$, it follows that $(A, \mu) \in \text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$. Conversely, assume that $\langle \mathcal{C}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \mathcal{D}, * \rangle$. Let D be any a set in \mathcal{D} . For every $\mu \in \mathcal{F}$, since $(D, \mu) \in \text{Aver}\langle \mathcal{D}, \mathcal{F} \rangle$, we obtain $(D, \mu) \in \text{Aver}\langle \mathcal{C}, \mathcal{F} \rangle$. Hence, D belongs to $\mathcal{D}_{\mathcal{F}}$. ■

We call $\{\Delta_k^{\mathcal{P}_{\mathcal{F}}}, \Sigma_k^{\mathcal{P}_{\mathcal{F}}} | k > 0\}$ the *real polynomial-time hierarchy* with respect to a class \mathcal{F} of density functions. By Proposition 7.2, Levin’s original question of whether $\langle \text{NP}, \text{P-comp} \rangle$ is contained in $\text{Aver}\langle \text{P}, * \rangle$ is simply rephrased as whether $\text{NP} \subseteq \text{P}_{\text{P-comp}}$ holds or not.

Lemma 7.3 Let \mathcal{F} be any class of density functions which contains the standard density function.

1. $\text{P} \subseteq \text{P}_{\mathcal{F}} \subseteq \text{E}$.
2. $\text{NP} \subseteq \text{NP}_{\mathcal{F}} \subseteq \text{NE}$.
3. $\text{BPP} \subseteq \text{BPP}_{\mathcal{F}} \subseteq \text{BPE}$.
4. $\text{PSPACE} \subseteq \text{PSPACE}_{\mathcal{F}} \subseteq \text{ESPACE}$.

Proof. Here we give only the proof of (1) since the rest of the claims follow by a similar argument. Since $\langle \text{P}, \mathcal{F} \rangle \subseteq \text{Aver}\langle \text{P}, \mathcal{F} \rangle$, it follows that $\text{P} \subseteq \text{P}_{\mathcal{F}}$. Now we show that $\text{P}_{\mathcal{F}} \subseteq \text{E}$. Let A be any set in $\text{P}_{\mathcal{F}}$. Since $(A, \nu_{\text{st}}) \in \text{Aver}\langle \text{P}, \mathcal{F} \rangle$, there exists a polynomial p and a deterministic Turing machine

M which is p -time bounded on ν_{st} -average such that M computes A . It clearly yields that, for all x ,

$$\text{Time}_M(x) \leq p(|x|/\nu_{\text{st}}(x)) = p(\pi^2|x|(|x|+1)^2 \cdot 2^{|x|}/6) \leq 2^{c|x|}$$

for some adequate constant $c > 0$. Therefore, we have $A \in \text{DTIME}(2^{cn})$. \blacksquare

The next result follows immediately from Lemma 6.8.

Lemma 7.4 *Let $k > 0$ and let \mathcal{F} be any class of density functions.*

1. $\Delta_k^{\text{P}} \subseteq \Delta_k^{\text{P}\mathcal{F}}$.
2. $\Sigma_k^{\text{P}} \subseteq \Sigma_k^{\text{P}\mathcal{F}}$.

Tally sets play a significant role in average-case analysis. From Lemma 3.10 and proposition 6.10, the average-case complexity of tally sets turns out to equal the worst-case complexity of them.

Proposition 7.5 *For every $\mathcal{C} \in \{\Delta_k^{\text{P}}, \Sigma_k^{\text{P}}, \text{BPP}, \text{PSPACE}\}$, $\mathcal{C}_{\text{P-comp}} \cap \text{TALLY} \subseteq \mathcal{C}$.*

Proof. The proof is similar to Theorem 6.11. \blacksquare

Recall that REC-comp denotes the set of recursive density functions, i.e., all “computable” density functions (under Church’s Thesis). If we take REC-comp as a set of density functions \mathcal{F} , then the real computable classes collapse to their worst-case counterparts.

To prove this, we show that if $\text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$ has the sparse interpolation property, then $\mathcal{C}_{\text{REC-comp}} \subseteq \mathcal{C}$. In the proof of the following lemma, we use the notion of infinite, recursive, proper hard cores [9]. A set H is called a *proper hard core for A with respect to \mathcal{C}* if $H \subseteq A$, and for all $D \in \mathcal{C}$, if $D \subseteq A$, then $|D \cap H|$ is finite.

Lemma 7.6 *Let $\text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$ be an average-case complexity class. If $\text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$ has the sparse interpolation property, then $\mathcal{C}_{\text{REC-comp}} \subseteq \mathcal{C}$.*

Proof. Suppose that $\text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$ has the sparse interpolation property. We show that $\mathcal{C}_{\text{REC-comp}} \subseteq \mathcal{C}$ by leading to a contradiction. Now assume that there exists a set $A \in \mathcal{C}_{\text{REC-comp}} - \mathcal{C}$. By [9], there exists an infinite, recursive, proper hard core H for A with respect to \mathcal{C} . We note that if $\mathcal{C} = \text{P}$, then H is in the class E (see, e.g., [2]). Thus, for any set $B \in \mathcal{C}$, if $B \subseteq A$ then $B \cap H$ is finite. Now let S be a recursive, infinite, sparse subset of H . Let $q(n) = |S \cap \Sigma^n|$. Consider the density function $\mu_{S,q}$ such that $\mu_{S,q}(x) \propto (|x|+1)^{-2} \cdot q(|x|)^{-1}$ for all $x \in S$, and $\mu_{S,q}(x) = 0$ otherwise. Clearly $\mu_{S,q} \in \text{REC-comp}$. Since $(A, \mu_{S,q}) \in \text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$, there exists an interpolant $B' \in \mathcal{C}$ of A and S . We then have $B' \cap H \supseteq S$, and thus $B' \cap H$ is infinite. This contradicts the fact that H is a proper hard core for A . \blacksquare

Theorem 7.7 *Let $k > 0$.*

1. $\Delta_k^{\text{P}}_{\text{REC-comp}} = \Delta_k^{\text{P}}$.
2. $\Sigma_k^{\text{P}}_{\text{REC-comp}} = \Sigma_k^{\text{P}}$.
3. $\text{BPP}_{\text{REC-comp}} = \text{BPP}$.
4. $\text{PSPACE}_{\text{REC-comp}} = \text{PSPACE}$.

Proof. By Lemma 7.6, it suffices to show that, for $\mathcal{C} \in \{\Delta_k^P, \Sigma_k^P, \text{BPP}, \text{PSPACE}\}$, $\text{Aver}\langle \mathcal{C}, \text{REC-comp} \rangle$ has the sparse interpolation property. The claim for $\mathcal{C} \in \{\text{BPP}, \text{PSPACE}\}$ follows From Lemma 3.10. The claim for $\mathcal{C} \in \{\Delta_k^P, \Sigma_k^P\}$ follows from Proposition 6.10. ■

Theorem 7.7 shows that the definition of the average polynomial-time hierarchy in Section 6 is a reasonable generalization of the worst-case polynomial-time hierarchy.

Note that, in the proof of Lemma 7.6, the complexity of the distribution $\mu_{S,q}$ depends only on the complexity of the complexity core. Since all sets not in P have complexity cores in E , we get the following corollary.

Corollary 7.8 $\text{P}_{\text{E-comp}} = \text{P}$.

Very recently, $\text{P}_{\text{P-comp}}$ is shown to be different from P and NP [32].

At the end of this section, we show that, relative to random oracle, $\text{NP}_{\text{P-comp}}$ is different from $\text{P}_{\text{P-comp}}$ with probability 1, i.e., Lebesgue measure of the set $\{X | \text{P}_{\text{P-comp}}^X \neq \text{NP}_{\text{P-comp}}^X\}$ is 1.

Definition 7.9 Let X be a set of strings.

1. Let $\text{P}_{\mathcal{F}}^X$ be the collection of all sets A such that, for any density function μ in \mathcal{F} , $(A, \mu) \leq_T^{\text{P,av}} (X, \nu)$ for some density function ν .
2. Let $\text{NP}_{\mathcal{F}}^X$ be the collection of all sets A such that, for any density function μ in \mathcal{F} , $(A, \mu) \leq_T^{\text{NP,av}} (X, \nu)$ for some density function ν .

Proposition 7.10 With probability 1, $\text{P}_{\text{P-comp}}^X \neq \text{NP}_{\text{P-comp}}^X$ relative to a random oracle X .

Proof. For any oracle A , $\text{P}_{\text{P-comp}}^A = \text{NP}_{\text{P-comp}}^A$ implies $\text{NP}^A \cap \text{TALLY} \subseteq \text{P}^A$. Bennett and Gill [4] have proven that, relative to a random oracle X , $\text{NP}^X \cap \text{TALLY} \not\subseteq \text{P}^X$ with probability 1. Hence, we get the desired result. ■

8 Conclusions and Open Problems

We have discussed structural properties of average-case complexity classes. Especially reducibilities have played a central role in our study of structural properties of average-case complexity classes. This paper introduces an average-case counterpart of the Meyer-Stockmeyer polynomial-time hierarchy based on the deterministic and nondeterministic Turing reducibilities between randomized decision problems, and we have seen that this hierarchy has a structure similar to its counterpart in worst-case complexity. Below, we give some problems mentioned in this paper.

1. Let D be any Σ_k^P -complete set. Does there exist a “natural” density function μ such that (D, μ) is \leq_m^{P} -complete for $\langle \Sigma_k^P, \text{P-comp} \rangle$? For example if A is NP -complete with p-honest reductions, then there exists a density function μ whose distribution is polynomial-time computable relative to $\#\text{P}$ such that (A, μ) is \leq_m^{P} -hard for $\langle \text{NP}, \text{P-comp} \rangle$.
2. Are all $\langle \Sigma_k^P, \text{P-comp} \rangle$ -complete problems Turing self-reducible? Can we extend the notion of polynomial-time Turing self-reducibility for randomized problems by allowing the reduction to be polynomial-time bounded on average?

3. In Lemma 2.6, $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle$ is characterized in terms of both deterministic machine models and logical formulas with the existential quantifier. Can this result be extended to characterize the class $\text{Aver}\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$?
4. Is $\text{Aver}\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$ contained in $\text{Aver}\langle \Delta_k^{\text{P}}, \mathcal{F} \rangle$ for some reasonable class \mathcal{F} ? Recall that if $\text{Aver}\langle \Delta_k^{\text{P}}, \text{P-comp} \rangle$ is equal to $\text{Aver}\langle \Sigma_k^{\text{P}}, \text{P-comp} \rangle$, then all tally Σ_k^{P} -sets are in Δ_k^{P} .
5. Is $\text{Aver}\langle \Sigma_k^{\text{P}}, \mathcal{F} \rangle$ different from $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle^{\langle \Sigma_{k-1}^{\text{P}}, * \rangle}$, $k > 1$, for a reasonable class \mathcal{F} , such as P-comp ?
6. In worst-case complexity theory, many important complexity classes, such as UP, CP and $\oplus\text{P}$, are used to classify intractable problems and to investigate their structural properties. Is it reasonable to consider those classes in average-case complexity theory ?
7. Very recently, it is shown that $\text{P}_{\text{P-comp}} \neq \text{P}$ [32]. Can we extend this result to show that $\Sigma_k^{\text{P}}_{\text{P-comp}} \neq \Sigma_k^{\text{P}}$ or $\Delta_k^{\text{P}}_{\text{P-comp}} \neq \Delta_k^{\text{P}}$?
8. Is it possible to show that an NP-complete problem with a natural density function is in the “low hierarchy in $\text{Aver}\langle \text{NP}, \mathcal{F} \rangle$ ” ?
9. What is a reasonable relativization of classes, such as $\text{Aver}\langle \text{BPP}, \mathcal{F} \rangle$ and $\text{Aver}\langle \text{PSPACE}, \mathcal{F} \rangle$?
10. Recall from [18] the definition of the time-complexity of a nondeterministic Turing machine M . Here we define Time_M to be the minimal length of accepting computations of M on input x if one exists; otherwise, $\text{Time}_M(x)$ is always set to 1. Can we develop a theory founded on this type of nondeterministic Turing machines which are polynomial-time (or polynomial-space) bounded on μ -average ? In this setting, for example, we can prove that $\text{Aver}\langle \text{P}, * \rangle \neq \text{Aver}\langle \text{NP}, * \rangle$ by choosing a non-recursive, recursively enumerable set $A = \{M(0), M(1), \dots\}$ by a deterministic machine M and defining a density function μ by: $\mu(x) \propto (|x| + 1)^{-2} \cdot 2^{-\sum_{i=0}^n \text{Time}_M(i)}$ if $x \in A$ and $n = \min\{k \mid M(k) = x\}$, or else 0.
11. An alternative definition of real polynomial-time hierarchy is given by: $\Delta_k^{\text{P}}_{\mathcal{F}} = \text{P}_{\mathcal{F}}^{\Sigma_{k-1}^{\text{P}}}$ and $\Sigma_k^{\text{P}}_{\mathcal{F}} = \text{NP}_{\mathcal{F}}^{\Sigma_{k-1}^{\text{P}}}$. Is it possible to develop a theory based on these $\Delta_k^{\text{P}}_{\mathcal{F}}$ and $\Sigma_k^{\text{P}}_{\mathcal{F}}$? For example, it is not hard to see that $\Delta_k^{\text{P}}_{\text{REC-comp}} = \Delta_k^{\text{P}}$ and $\Sigma_k^{\text{P}}_{\text{REC-comp}} = \Sigma_k^{\text{P}}$ also in this setting.

Acknowledgments

The authors would like to thank Professor José Balcázar and Professor Uwe Schöning for organizing workshops in Barcelona and Ulm in 1991 where many discussions on average-case complexity theory took place and thank anonymous referees who gave them useful comments and also criticism for an early draft. The second author is specially grateful to Professor Uwe Schöning for allowing him to stay at the Universität Ulm from April to August in 1991.

References

- [1] T. Baker, G. Gill, and R. Solovay, Relativizations of the $P \stackrel{?}{=} NP$ question, *SIAM J. Comput.* **4** (1975), 431–442.
- [2] J.L. Balcazar, J. Díaz, and J. Gabarró, *Structural Complexity I, II*, Springer-Verlag, Berlin, Heidelberg, (I) 1988, (II) 1990.
- [3] J. Belanger and J. Wang, Isomorphisms of NP complete problems on random instances, in “Proceedings, 8th IEEE Conference on Structure in Complexity Theory”, 65–74, 1993.
- [4] C.H. Bennett and J. Gill, Relative to a random oracle A , $P^A \neq NP^A \neq co-NP^A$ with probability 1, *SIAM J. Comput.* **10** (1981), 96–113.
- [5] S. Ben-David, B. Chor, O. Goldreich, and M. Luby, On the theory of average case complexity, *J. Comput. System Sci.* **44** (1992), 193–219.
- [6] L. Berman and J. Hartmanis, On isomorphism and density of NP and other complete sets, *SIAM J. Comput.* **6** (1977), 305–322.
- [7] A. Blass and Y. Gurevich, Randomized reductions of search problems, *SIAM J. Comput.* **22** (1993), 949–975.
- [8] B. Bollobas, T.I. Fenner and A.M. Frieze, An algorithm for finding Hamiltonian cycle in a random graph, in “Proceedings, 17th ACM Symposium on Theory of Computing”, 430–439, 1985.
- [9] R. Book, D.Z. Du, and D. Russo, On polynomial and generalized complexity cores, in “Proceedings, 3rd Conference on Structure in Complexity Theory”, 236–250, 1988.
- [10] J. Franco and M. Paull, Probabilistic analysis of the Davis Putnum procedure for solving the satisfiability problem, *Discrete Applied Math.* **5** (1983), 77–87.
- [11] M.R. Garey and D.J. Johnson, *Computers and Intractability, A Guide to the Theory of NP-Completeness*, W.H.Freeman and Company, New York, 1979.
- [12] J. Gill, Computational complexity of probabilistic Turing machines, *SIAM J. Comput.* **6** (1977), 675–695.
- [13] M. Goldmann, P. Grape, and J.Håstad, On average time hierarchy, *Inform. Processing Letters* **49** (1994), 15–20.
- [14] O. Goldreich, Towards a theory of average case complexity (a survey), Technical Report No.507, Israel Institute of Technology, 1988.
- [15] P. Grape, Two results in average case complexity, Technical Report TRITA-NA-9105, Royal Institute of Technology, Stockholm, Sweden, 1991.
- [16] Y. Gurevich, Complete and incomplete randomized NP problems, in “Proceedings, 28th IEEE Symposium on Foundations of Computer Science”, 111–117, 1987.
- [17] Y. Gurevich, Average case complexity, *J. Comput. System Sci.* **42** (1991), 346–398.

- [18] J.E. Hopcroft and J.D. Ullman, *Introduction to Automata Theory, Languages, and Computation*, Addison-Wesley, Massachusetts, 1979.
- [19] R. Impagliazzo and L.A. Levin, No better ways to generate hard NP-instances than picking uniformly at random, in “Proceedings, 31th IEEE Symposium on Foundations of Computer Science”, 812–821, 1990.
- [20] D. Joseph and P. Young, Some remarks on witness functions for nonpolynomial and noncomplete sets in NP, *Theoret. Comput. Sci.* **39** (1985), 225–237.
- [21] K.I. Ko and H. Friedman, Computational complexity of real functions, *Theoret. Comput. Sci.* **20** (1982), 323–352.
- [22] L. Levin, Problems, complete in “average” instance, in “Proceedings, 16th ACM Symposium on Theory of Computing”, p.465, 1984.
- [23] L. Levin, Average case complete problems, *SIAM J. Comput.* **15** (1986), 285–286.
- [24] M. Li and P.M.B. Vitányi, Average case complexity under the universal distribution equals worst-case complexity, *Information Processing Letters* **42** (1992), 145–149.
- [25] A. Meyer and M. Paterson, With what frequency are apparently intractable problems difficult ?, Technical Report MIT/LCS/TM-126, MIT, 1979.
- [26] A.R. Meyer and L.J. Stockmeyer, The equivalence problem for regular expressions with squaring requires exponential time, in “Proceedings, 13th IEEE Symposium on Switching and Automata Theory”, 125–129, 1972.
- [27] M. Mundhenk and R. Schuler, Random languages for non-uniform complexity classes, *Journal of Complexity* **7** (1991), 296–310.
- [28] K.W. Regan, Minimum-complexity pairing functions, *J. Comput. System Sci.* **45** (1992), 285–295.
- [29] R. Reischuk and C. Schindelbauer, Precise average case complexity, in “Proceedings, 10th Symposium on Theoretical Aspect of Computer Science”, Lecture Notes in Computer Science, Vol.665, Springer-Verlag, 650–661, 1993.
- [30] R.E. Schapire, The emerging theory of average-case complexity, Technical Report, MIT/LCS/TM-431, Massachusetts Institute of Technology, 1990.
- [31] U. Schöning, A low and a high hierarchy within NP, *J. Comput. System Sci.* **27** (1983), 14–28.
- [32] R. Schuler, On average polynomial time, Technical Report, Ulmer Informatik Berichte Nr. 94-12, Universität Ulm, 1994.
- [33] L.J. Stockmeyer, The polynomial-time hierarchy, *Theoret. Comput. Sci.* **3** (1977), 1–22.
- [34] L. Valiant, The complexity of computing the permanent, *Theoret. Comput. Sci.* **5** (1979), 189–201.
- [35] R. Venkatesan and L. Levin, Random instances of a graph coloring problem are hard, in “Proceedings, 20th ACM Symposium on Theory of Computing”, 217–222, 1988.

- [36] J. Wang and J. Belanger, On average P vs. average NP, *in* Complexity Theory – Current Research, editors K. Ambos-Spies, S. Homer and U. Schöning, Cambridge University press, 47–67, 1993.
- [37] J. Wang and J. Belanger, On the NP-isomorphism problem with respect to random instances, to appear in *J. Comput. System Sci.*, 1994.
- [38] R.E. Wilber, Randomness and the density of hard problems, *in* “Proceedings, 24th IEEE Symposium on Foundations of Computer Science”, 335–342, 1983.
- [39] H. Wilf, Some examples of combinatorial averaging, *American Math. Monthly* **92** (1985), 250–261.
- [40] C. Wrathall, Complete sets and the polynomial-time hierarchy, *Theoret. Comput. Sci.* **3** (1977), 23–33.