like tight hierachies can be obtained this way, both for the time complexity and the complexity of the distributions. Based on these notions, starting with distributional complexity classes we have presented meaningful definitons of average case complexity classes the elements of which are languages in the standard sense. They are directly comparable to worst case classes.

These ideas can also be applied to other cases like the analysis of average space complexity.

References

[BCGL 92] S. Ben-David, B. Chor, O. Goldreich, M. Luby, *On the Theory of Average Case Complexity*, J. CSS, 1992, 44; see also Proc. 29. FoCS, 1989, 204-261.

[Gure 91] Y. Gurevich *Average Case Completeness*, J. CSS 42, 1991, 346-398.

[ImLe 90] R. Impagliazzo, L. Levin, *No Better Ways to Generate Hard $\mathcal{NP}$ Instances than Picking Uniformly at Random*, Proc. 31. FoCS, 1990, 812-821.

[John 84] D. Johnson, *The $\mathcal{NP}$ -Completeness Column*, J. of Algorithms 5, 1984, 284-299.

[Levi 86] L. Levin, *Average Case Complete Problems*, SIAM J. Computing 15, 1986, 285-286.

[Milt 91] P. Miltersen, *The Complexity of Malign Ensembles,* Proc. 6. StruC, 1991, 164-171, see also SIAM J. Computing 22, 1993, 147-156.

[MiSeLe92] D. Mitchell, B. Selman, H. Levesque *Hard and Easy Distributions of SAT Problems*, Proc. 10. Nat. Conf. on Artificial Intelligence, 1992, 459-465.

[ReSc93] R. Reischuk, C. Schindelhauer, *Precise Average Case Complexity*, Proc. 10. STACS, 1993, 650-661.

[Schi91] C. Schindelhauer, *Neue Average Case Komplexitätsklassen,* Diplomarbeit, Technische Hochschule Darmstadt, 1991.

[VeLe 88] R. Venkatesan, L. Levin, *Random Instances of Graph Coloring Problems are Hard*, Proc. 20. SToC, 1988, 217-222.

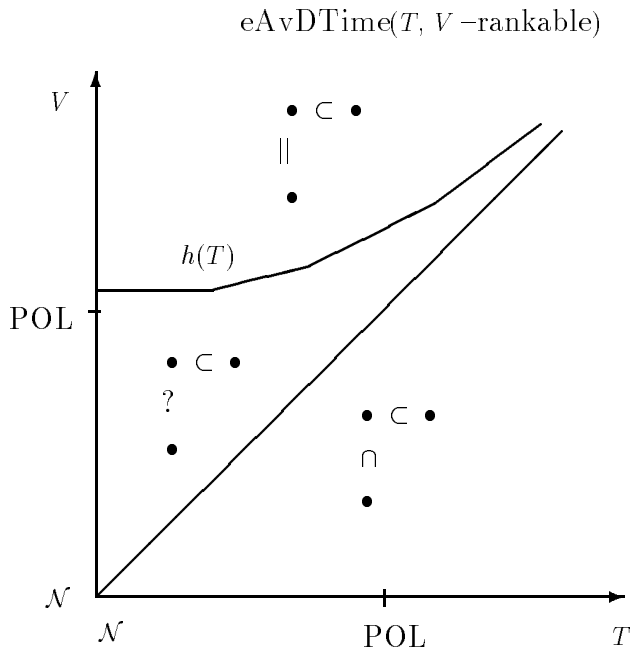

Figure 3: Hierarchies between expected average case complexity classes eAvDTime(T)

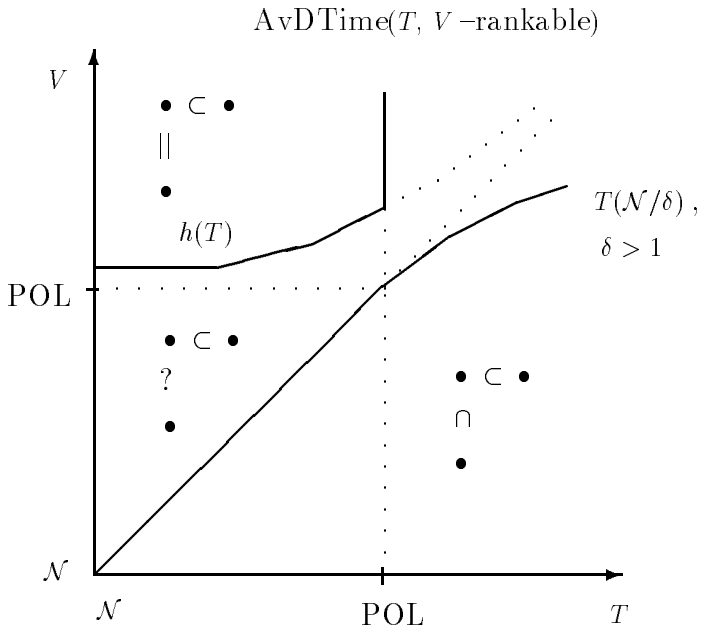

Figure 4: Hierarchies between average case complexity classes AvDTime(T)

Remember that $T_1 \leq T_2(o(\mathcal{N}))$. So f is well-defined, and such simulations of M_i on input x can be performed in time $T_2(|x|)$.

$P \in \mathrm{DTime}(T_2)$ follows as above.

$P \cap L \notin \mathrm{AvDTime}(T_1, \{\mathrm{rank}_L\})$:

Assume $P \cap L \in \mathrm{AvDTime}(T_1, \mathrm{rank}_L)$ and let M_i be a DTM for this language with $(\mathrm{time}_{M_i}, \mathrm{rank}_L) \in \mathrm{Av}(T_1)$. Again letting $l := R_L(n_i) + 2^{2n_i+1}$ we get the contradiction

$$\begin{aligned}\sum_{\mathrm{rank}_L(x) \leq l} \frac{{T_1}^{-1}\mathrm{time}_{M_i}(x)}{|x|} &\geq \sum_{R_L(n_i) \leq \mathrm{rank}_L(x) \leq R_L(n_i)+2^{2n_i+1}} \frac{{T_1}^{-1}(T_1(2 \cdot |x|)}{|x|} \\ &= \left(R_L(n_i) + 2^{2n_i+1} - R_L(n_i) + 1\right) \cdot 2 \;\geq\; 2^{2n_i+2} + 1 \;\geq\; l + 1 \,.\end{aligned}$$

∎

Combining the last four lemmata we get

Theorem 10 *Let $\delta > 1$. Then for all bounds T, T', V_1, V_2 with $\mathcal{N} \leq V_1 \leq o(V_2)$, $V_2 \leq O(T)$ and $V_2(\delta \cdot \mathcal{N}) \leq O(T')$ holds*

$$\begin{aligned}eAvDTime(T, V_2\text{-}rankable) &\subset eAvDTime(T, V_1\text{-}rankable)\,, \\ AvDTime(T', V_2\text{-}rankable) &\subset AvDTime(T', V_1\text{-}rankable)\,.\end{aligned}$$

Proof: Choose $T = T_1$ where T_1 is the complexity bound considered in the previous Lemmata. The language $P \cap L$ constructed above in Lemmata 9 and 11 is contained in the average class eAvDTime(T_1, V_1 –rankable) according to Lemma 9, but not in eAvDTime(T, V_2 –rankable) according to Lemma 9 since $\mathrm{rank}_L \in V_2$–rankable (Lemma 9). The analogous property for the Av-measure holds for the language obtained from Lemmata 10 and 12 and $T' = T_1$. ∎

Corollary 9 *For $\mathcal{N} \leq V_1 \leq o(V_2)$, $V_1 \leq O(T)$ and $T \in POL$*

$$AvDTime(T,\ V_2\text{-}rankable) \subset AvDTime(T,\ V_1\text{-}rankable)\,.$$

Figures 3 and 4 give a pictorial description of the hierarchies implied by the last two theorems. Each point • in the diagram represents a complexity class eAvDTime(T, V –rankable), respectively AvDTime(T, V –rankable). $\bullet \subset \bullet$ means strict inclusion between the two classes, $\bullet = \bullet$ equality and $\bullet \,?\, \bullet$ that the relation cannot be deduced from the results obtained above.

8 Conclusions

We have shown that the average case time complexity of an algorithm can be estimated as precisely as in the worst case. Ranking the input space and measuring the complexity of a distribution with respect to its rankability has turned out to be an appropriate and natural concept. Classical results

and a diagonal language like in the proof of theorem 9 by

$$P \;:=\; \{x \mid \text{for } i \text{ with } f(i) \le |x| < f(i+1) \text{ holds:} \\ i \cdot T_1(|x|) \le T_2(|x|) \quad \text{and} \quad [\text{time}_{M_i}(x) > 2 \cdot T_1(|x|) \quad \text{or} \quad x \notin L(M_i)] \;\} \;.$$

Since $T_1 \le o(T_2)$ the sequence n_i is well defined (always $< \infty$).

1. $P \in \text{DTime}(T_2)$:

   On input x, first compute all $f(j)$ and n_j smaller than $\log |x|$.
   Because of the strong growth of f and the time constructibility of all time bounds this can easily be done in time $T_2(|x|)$.

   Compute i with $f(i) \le |x| < f(i+1)$ in time $O(|x|)$ using the last computed n_i.

   Test whether $i \cdot T_1(|x|) \le T_2(|x|)$.
   If this does not hold reject, otherwise simulate the machine M_i for $2 \cdot T_1(|x|)$ steps.
   Accept iff M_i does not accept within that period.

   This simulation costs at most $2 \cdot T_1(|x|) \cdot i \le 2 \cdot T_2(|x|)$ many steps.

2. $P \cap L \notin \text{eAvDTime}(T_1, \{\text{rank}_L\})$:

   Assume that M_i computes $P \cap L$ in expected average time T_1 with respect to the ranking rank_L. By definition of n_i and R_L it holds for all i

   $$|\{x \in L |\; f(i) \le |x| < f(i+1) \text{ and } T_1(|x|) \cdot i \le T_2(|x|)\}| \ge R_L(n_i) + 2^{2n_i+1}$$

   Furthermore, there exists an infinitive set of machine indices $i_1, i_2 \ldots$ such that $L(M_{i_k}) = P$. Hence, the set $P \cap L$ is infinite.

   Observe that for inputs $x \in P \cap L$ it must hold $x \in L(M_i)$ and hence $\text{time}_{M_i} > 2T_1(|x|)$. Let $l := R_L(n_i) + 2^{2n_i+1}$. Then $l \le 2^{2n_i+2}$ since $R_L(n_i) \le n_i$.

   The time complexity of M_i can be estimated by

   $$\begin{aligned} \sum_{\text{rank}_L(x) \le l} \frac{\text{time}_{M_i}(x)}{T_1(|x|)} \;&\ge\; \sum_{R_L(n_i) \le \text{rank}_L(x) \le R_L(n_i) + 2^{2n_i+1}} \frac{2 \cdot T_1(|x|)}{T_1(|x|)} \\ &=\; \left(R_L(n_i) + 2^{2n_i+1} - R_L(n_i) + 1\right) \cdot 2 \;\ge\; 2^{2n_i+2} + 1 \;\ge\; l + 1 \;, \end{aligned}$$

   contradicting the assumption that M_i were T_1–time bounded on the average.

∎

Lemma 12 *Let $\delta > 1$, $V_2 \ge \mathcal{N}$, $V_2(\delta \cdot \mathcal{N}) \le T_1$ and $T_1 \le T_2(o(\mathcal{N}))$. Then for all infinite $L \in DTime(V_2)$ there exists a language $P \in DTime(T_2)$ such that*

$$P \;\cap\; L \;\notin\; AvDTime(T_1, \{rank_L\}) \;.$$

Proof: Define P and the sequence n_i as above, but replace the condition $\frac{T_2(j)}{T_1(j)} \ge i$ by $\frac{T_2(j)}{T_1(2 \cdot j)} \ge i$. Modify the simulation in the proof of lemma 11 to testing whether a machine M_i does not accept an input x in time $T_1(2 \cdot |x|)$.

- $rank_L \in V_2\text{-rankable} \setminus V_1\text{-rankable}$ *and*
- $P \cap L \in AvDTime(T_1, V_1\text{-rankable})$ *for all* $P \in DTime(T_2)$.

Proof: Use the same L as in the previous lemma. The time bound T_2 is now defined as:

$$T_2(n) := T_1\left(\left\lfloor n \cdot R_L(n) \cdot \left(1 - \delta^{-1}\right)\right\rfloor\right) .$$

We already know that $\mathrm{rank}_L \in V_2\text{-rankable} \setminus V_1\text{-rankable}$. It remains to show that for all $P \in \mathrm{DTime}(T_2)$:

$$P \cap L \in \mathrm{AvDTime}(T_1(2 \cdot \mathcal{N}),\ V_1\text{-rankable}) .$$

For a given $P \in \mathrm{DTime}(T_2)$ and for a given rank function $r \in V_1$-rankable construct a machine M as above.

- For the finitely many x with $r(x) \leq (\mathrm{rank}_L(x)+1)^2$ the answer is encoded into the machine and obtained in linear time.
- M tests if $x \in L$ and rejects if not.
  This now takes time at most $O(T_1(|x|/\delta))$.
- In the remaining case, $\mathrm{rank}_L(x) < \infty$ and $r(x) > (\mathrm{rank}_L(x)+1)^2$, decide $x \in P$ within at most $T_2(|x|)$ additional steps.

Then for all $l \in \mathbb{N}$:

$$\begin{aligned}
& \sum_{r(x)\leq l} \frac{{T_1}^{-1}(\mathrm{time}_M(x))}{|x|} \\
\leq\ & \sum_{\substack{r(x)\leq l \text{ and} \\ r(x)\leq(\mathrm{rank}_L(x)+1)^2}} \frac{{T_1}^{-1}(T_1(|x|/\delta))}{|x|} \;+\; \sum_{\substack{r(x)\leq l \text{ and} \\ r(x)>(\mathrm{rank}_L(x)+1)^2}} \frac{{T_1}^{-1}(T_2(|x|))}{|x|} \\
\leq\ & \frac{l}{\delta} + \sum_{\mathrm{rank}_L(x)\leq\sqrt{l}-1} \frac{R_L(|x|)\cdot|x|\cdot(1-\frac{1}{\delta})}{|x|} \leq \frac{l}{\delta} + \left(1-\frac{1}{\delta}\right)\cdot\sum_{i=1}^{\lfloor\sqrt{l}\rfloor-1} i \leq l .
\end{aligned}$$

Therefore $(\mathrm{time}_M, r) \in \mathrm{Av}(T_1)$. ∎

7.2 The Separation

Lemma 11 *Let* $\mathcal{N} \leq V_2 \leq o(T_1)$ *and* $T_1 \leq o(T_2)$ *. Then for all* $\delta > 1$ *and for all infinite* $L \in DTime(V_2)$*, there exists a* $P \in DTime(T_2)$ *such that*

$$P \cap L \notin eAvDTime(T_1, \{rank_L\}) .$$

Proof: Define

$$\begin{aligned}
f(0) &:= 1 , \\
n_i &:= \min\left\{ b \geq f(i) \;\middle|\; \forall\, j \in [b \ldots R_L^{-1}(R_L(b) + 2^{2b+1})] :\ \frac{T_2(j)}{T_1(j)} \geq i \right\} , \\
f(i+1) &:= 2^{2n_i+1} ,
\end{aligned}$$

For a language L with rank function $\mathrm{rank}_L(x)$ define

$$R_L(n) \;:=\; \max\{\mathrm{rank}_L(x) \;\big|\; |x| \le n,\ \mathrm{rank}_L(x) < \infty\}\ .$$

Lemma 9 *Let T_1, V_1, V_2 be complexity bounds with the properties $\mathcal{N} \le V_1 \le o(V_2)$ and $V_2 \le O(T_1)$. Then there exists a language $L \in DTime(V_2)$ and a complexity bound $T_2 \in \omega(T_1)$ such that*

- *$rank_L \in V_2$ –rankable $\setminus$ V_1 –rankable and*
- *$P \cap L \in eAvDTime(T_1,\ V_1$ –rankable$)$ for all $P \in DTime(T_2)$.*

Proof: Let L be a language with the properties as described in lemma 8, and define $T_2(n) := T_1(n) \cdot R_L(n)$.

1. $\mathrm{rank}_L \in V_2$ –rankable$\setminus$ V_1 –rankable follows immediately from the construction of L since any rank function of complexity at most V_1 will infinitely often generate rank values larger than rank_L.

2. $\forall P \in \mathrm{DTime}(T_2)\quad P \cap L \in \mathrm{eAvDTime}(T_1,\ V_1$ –rankable$)$:

   Let r be the rank function of a distribution that is V_1 –rankable, and choose $P \in \mathrm{DTime}(T_2)$. To decide whether $x \in P \cap L$ do the following:

   - For the finitely many inputs x with $r(x) \le (\mathrm{rank}_L(x)+1)^2$ prepare a decision table beforehand to get the result in linear time.
   - Otherwise, test whether $x \in L$. Reject, if not.
     Since $L \in \mathrm{DTime}(V_2)$ and, by assumption, $V_2 \le O(T_1)$ this question can be decided in time $T_1(|x|)$ (using a linear speedup if necessary).
   - In the remaining case it holds $x \in L$ and thus $\mathrm{rank}_L(x) < \infty$, and furthermore $r(x) > (\mathrm{rank}_L(x)+1)^2$.
     Then, use a T_2 –time bounded algorithm to decide whether $x \in P$.

   To estimate the average time complexity of a machine M performing this computation we get for any $l \in \mathbb{N}$:

$$\sum_{r(x)\le l} \frac{\mathrm{time}_M(x)}{2\cdot T_1(|x|)} \;\le\; \sum_{\substack{r(x)\le l \text{ and}\\ r(x)\le(\mathrm{rank}_L(x)+1)^2}} \frac{T_1(|x|)}{2\cdot T_1(|x|)} \;+\; \sum_{\substack{r(x)\le l \text{ and}\\ r(x)>(\mathrm{rank}_L(x)+1)^2}} \frac{T_2(|x|)}{2\cdot T_1(|x|)}$$

$$\le\; \frac{l}{2} \;+\; \sum_{\mathrm{rank}_L(x)<\sqrt{l}-1} \frac{T_1(|x|)\cdot R_L(|x|)}{2\cdot T_1(|x|)} \;\le\; \frac{l}{2} \;+\; \frac{1}{2}\cdot \sum_{i=1}^{\lfloor\sqrt{l}\rfloor-1} i \;\le\; l\ .$$

   Therefore $(\mathrm{time}_M, r) \in \mathrm{eAv}\,(2\cdot T_1)$. By a linear speedup, one can find a machine M' with $L(M') = L(M) = P \cap L$ and $(\mathrm{time}_{M'}, r) \in \mathrm{eAv}\,(T_1)$.

∎

Lemma 10 *Let $\delta > 1$ and T_1, V_1, V_2 be complexity bounds with $\mathcal{N} \le V_1 \le o(V_2)$ and $V_2(\delta\mathcal{N}) \le O(T_1)$. Then, there exists a language $L \in DTime(V_2)$ and a complexity bound $T_2 \in \omega(T_1)$ such that*

RANKL(x)
for input $x = 1^m$ find l such that $m = n_l$
if l does not exist then return $(-,-,\infty)$
if $l = 1$
    then return $(1,1,1)$
    else $(j,k,r) := \mathtt{RANKL}(x_{l-1})$
let I_k be the output of a V_1-rank accumulator on input 1^k
$I := I_k \cap [1..b]$, where $b := \min\left\{j, \frac{V_2(m)}{6 \cdot V_1(m)}\right\}$
for all $i \in I$
    simulate machine M_i on input x for at most $V_1(|x|)$ steps
    if by that time M_i has output a rank less than $(r+2)^2$
        then return $(j, k+1, \infty)$
return $(j+1, j, r+1)$

We have to show that the language L constructed this way possesses the properties stated above.

1. $\mathrm{rank}_L \in V_2$-rankable:
   Assume inductively that on input x_{l-1} the time complexity of **RANKL** is bounded by $V_2(n_{l-1})$. Now consider $x = 1^m$ with $m > n_{l-1}$. We may assume that determining n_{l-1}, resp. n_l takes time at most $\frac{1}{6}V_2(m)$.

   If $m = n_l$ then due to the properties of this sequence the time for the recursive call of $\mathtt{RANKL}(x_{l-1})$ is bounded by $\frac{1}{2}V_2(n_l)$. Also, we may assume that the output of the rank accumulator can be obtained in time $\frac{1}{6}V_2(n_l)$. The time needed for the simulation of all machines in I is bounded by
   $$b \cdot V_1(m) \;\le\; \frac{V_2(m)}{6 \cdot V_1(m)} \cdot V_1(m) \;=\; \frac{1}{6}V_2(n_l)\ .$$
   This proves that rank_L can be computed in time V_2.

2. L is infinite:
   Assume that for some input x_l a finite rank r is computed by **RANKL**. As long as the next candidates $x = x_{l+1}, x_{l+2}, \ldots$ do not get a finite rank the variable j stays fixed and k is increased each time. By definition of a V_1-rank-accumulator, eventually for some k the set $I = I_k \cap [1..b]$ with $b = \min\{j, V_2(|x|)/6V_1(|x|)\}$ only contains indices of machines that compute a pseudo rank function. Any one of them can hurt the condition "output" $< (r+1)^2$ only a finite number of times. Thus, eventually the rank $r+1$ will be assigned to some successor of x_l.

3. For all $R \in V_1$-rankable and for almost all x:
   $$rank_L(x) < \infty \qquad \Longrightarrow \qquad R(x) \ge (rank_L(x)+1)^2\ .$$
   The index of a V_1-time bounded machine M computing R will eventually be considered in the *for*-loop on some input x_l. Then for all inputs x_s with $s \ge l$ holds:
   if **RANKL** generates a finite rank r it will not be larger than $\sqrt{R(x_s)} - 1$.

∎

Lemma 8 *For all complexity bounds V_1, V_2 with $\mathcal{N} \leq V_1 \leq o(V_2)$ there exists an infinite language L with the following properties:*

- *$rank_L \in V_2$ -rankable,*
- *for all $R \in V_1$ -rankable: $rank_L(x) < \infty \implies R(x) \geq rank_L(x)^2$ for almost all x.*

Proof: The language L to be constructed will be a subset of 1^*. Remember that according to our general assumption on complexity bounds V_2 is monotone and time-constructible. We start with a sequence $n_1 = 1 < n_2 < n_3 < \ldots \subseteq \mathbb{N}$ with the following properties:

- $V_2(n_{l+1}) \geq 2 \cdot V_2(n_l)$
- for any $m \in \mathbb{N}$ the element n_l with $n_l \leq m < n_{l+1}$ can be computed in time $O(V_2(m))$.

For example, such a sequence can be obtained as follows: Assume n_l has been defined. Then, for $k = 1, 2, 3, \ldots$ consider the sequence $V_2(2^k \cdot n_l)$ till the first k such that

$$V_2(2^k \cdot n_l) \;\geq\; 2^k \cdot V_2(n_l) \;.$$

It is easy to see that such a value, let us denote it by $\bar{k}$, must exist. Otherwise, the partial sums of V_2 were bounded quadratically, which contradicts the condition that V_2 grows faster than linear. Considering n_l fixed and m growing, one would get

$$\begin{aligned}\sum_{a=1}^{n_l+2^m} V_2(a) \;&=\; \sum_{a=1}^{n_l-1} V_2(a) \;+\; \sum_{a=n_l}^{n_l+2^m} V_2(a) \;\leq\; O(1) \;+\; \sum_{a'=1}^{m} 2^{a'-1} \cdot V_2(2^{a'} \cdot n_l) \\ &\leq\; O(1) \;+\; \sum_{a'=1}^{m} 2^{a'-1} \cdot 2^{a'} \cdot V_2(n_l) \;\leq\; O(1) \;+\; 2^{2m} \cdot V_2(n_l) \;\leq\; O((n_l + 2^m)^2) \;.\end{aligned}$$

Define $n_{l+1} := 2^{\bar{k}} \cdot n_l$. Then this sequence obviously fulfills the first condition. Since V_2 is time constructible, given n_l the next value n_{l+1} can be computed in time

$$\sum_{k=1}^{\bar{k}} V_2(2^k \cdot n_l) \;\leq\; \sum_{k=1}^{\bar{k}-1} 2^k \cdot V_2(n_l) + V_2(2^{\bar{k}} \cdot n_l) \;\leq\; 2 \cdot V_2(n_{l+1}) \;.$$

Thus, the whole sequence $n_1, \ldots, n_{l+1}$ can be computed in time $O(V_2(n_{l+1}))$, and for $m = n_{l+1}$ the second condition holds, too. For arbitrary m first compute $V_2(m)$ and then set a time limit of $O(V_2(m))$. Start computing the sequence $n_1, \ldots$ until either n_{l+1} has been found or the time limit is over. In the second case, the last value computed is the desired n_l. Define

$$x_l := 1^{n_l} \;.$$

The following algorithm computes rank_L, and thus implicitely defines L. On input x_l a triple (j, k, r) is generated where r denotes the rank of x_l in L. Thus $x_l \in L$ iff $r < \infty$. j and k are additional parameters that are needed in this recursive procedure.

Definition 12 *A* **pseudo rank function** *is function* $r : \Sigma^* \to \mathbb{N} \cup \{\infty\}$ *with the property that for all* $l \in \mathbb{N}$ *holds:* $|\{x | r(x) \leq l\}| \leq l$.

An example of a pseudo rank function that is not a rank function is the function $1 \mapsto 4$, $2 \mapsto 5$, $3 \mapsto 6$, ... Some rank values never appear, but it does not happen that too many small ranks are generated. One cannot enumerate all machines computing pseudo rank functions either, but at least one can construct a dynamic set of candidates such that each machine fulfilling this property will eventually be contained in this set, and each machine not fulfilling this property will eventually be thrown out.

Definition 13 *Let* V *be a time bound. A* **V-rank-accumulator** *is a DTM* M *that on input* 1^n *computes a set of machine indices* $I_n \subset \mathbb{N}$ *with the following property:*

- $\forall\, i \in \mathbb{N}$ *holds:* M_i *computes a pseudo rank function and is* V*-time bounded* $\iff$
  $\exists\, n_0 \quad \forall n \geq n_0 \qquad i \in I_n$.

Lemma 7 *For every complexity bound* V *there exists a* V*-rank-accumulator* M *. Furthermore,* M *can be made to run in linear time.*

Proof: Let $c : \mathbb{N} \to \mathbb{N}$ be a monotone, unbounded function that can be computed in linear time. The following algorithm defines a V-rank-accumulator.

on input 1^n compute $c(n)$ and define $I := [1 \ldots c(n)]$
for all $i \in I$
    $R_i := \emptyset$
    for all $x \in \Sigma^{\leq c(n)}$
        simulate machine M_i on input x for at most $V(|x|)$ steps
        if M_i has not terminated by that time
            then remove i from I
            else add its output to the multiset R_i
    for all $l \in R_i$
        if $|\{r \in R_i \mid r \leq l\}| > l$ then remove i from I
    return I

The correctness of this algorithm can be seen easily. Its time complexity on input 1^n is given by

$$T_{V,c}(n) \leq O\Big(c(n) \cdot |\Sigma|^{c(n)+1} \cdot V(c(n))\Big)$$

(remember that complexity bounds like V were assumed to be monotone increasing). For any V one can find a monotone and unbounded function c such that $T_{V,c}$ grows at most linearly, and $c(n)$ can be computed in time $O(n)$. For example, the following function has these properties

$$c(n) := \min\left\{V^{-1}(\sqrt{n}), \frac{\log n}{3 \log |\Sigma|}\right\} .$$

∎

$$
\begin{aligned}
\sum_{2^{n_i} \le \mathrm{rank}_{\mathrm{uni}}(x) < 2^{2n_i+1}} \frac{|x|+1}{|x|} &\le 2^{2n_i+1} \quad \Rightarrow \\
\sum_{2^{n_i} \le \mathrm{rank}_{\mathrm{uni}}(x) < 2^{2n_i+1}} 1 + \frac{1}{|x|} &\le 2^{2n_i+1} \quad \Rightarrow \\
(2^{2n_i+1} - 2^{n_i}) \cdot \left(1 + \frac{1}{2^{n_i}}\right) &\le 2^{2n_i+1} \quad \Rightarrow \\
2 \cdot 2^{n_i} &\le 2^{n_i} + 1
\end{aligned}
$$

∎

Now observe that if $V_1 \le V_2$ then V_1–rankable $\subseteq$ V_2–rankable and hence

$$\mathrm{eAvDTime}(T,\ V_1\text{–rankable}) \supseteq \mathrm{eAvDTime}(T,\ V_2\text{–rankable})\ .$$

Therefore, for $V_1 \ge \mathcal{N}$

$$
\begin{aligned}
\mathrm{DTime}(T) &\subseteq \mathrm{eAvDTime}(T, V_1\text{–rankable}) \\
&\subseteq \mathrm{eAvDTime}(T, \{\mathrm{uni}\})
\end{aligned}
$$

and similarly for the other measure. This implies

Corollary 8 *Let $T_1, V \ge (1+\epsilon)\mathcal{N}$ and $T_2 \in \omega(T_1)$. Then*

$$
\begin{aligned}
\mathit{eAvDTime}(T_1, V\text{–}\mathit{rankable}) &\subset \mathit{eAvDTime}(T_2, V\text{–}\mathit{rankable})\ , \\
\mathit{AvDTime}(T_1, V\text{–}\mathit{rankable}) &\subset \mathit{AvDTime}(T_2, V\text{–}\mathit{rankable})\ .
\end{aligned}
$$

∎

7 Distribution Hierarchies of Average Case Classes

In the previous section we have seen that like in the worst case any slight increase in the average time bound gives more computational power. Now we want to investigate the same question with respect to the complexity of the distributions. Will a more severe restriction on the set of distributions, for which one requires a good average case behaviour, allow to solve more algorithmic problems within a given average time bound?

7.1 Diagonalization with respect to Sets of Rank Functions

Again a proof for this property is likely to use a special kind of diagonalization, but there is a slight technical problem arises. In general, it is not possibile to enumerate exactly those machines that compute rank functions of a certain time complexity because this task is equivalent to the halting problem. We circumvent this difficulty by considering a broader class of rank functions.

- To exceed the average time bound T_1 we let the index i grow very slowly. So, we get for each index i at least as many finite ranks (relevant for the diagonalization) as for all indices $1 \ldots i-1$ together. This is achieved by the rapid growth of the n_i.
- To be able to compute the inverse of the resulting function f in time $T_2(|x|)$ we choose $f(i+1)$ as 2^{2n_i+1}.

1. $L \in \mathrm{DTime}(T_2)$:

   To decide L by a DTM M, first for an input x compute the corresponding i such that $f(i) \leq |x| < f(i+1)$. For this aim the machine succesively computes all $f(j)$ for $j \leq |x|$. This is easy in the given time bound since f grows very fast ($f \in \omega(\mathrm{itexp})$). For a given $f(j)$, $f(j+1)$ is computed by incrementing two counters b, j until b fulfills the condition for n_j. If $j > \log|x|$, the algorithm halts: the last computed value of f is $f(i)$. This costs at most $(\log f(j+1) - f(j)) \cdot T_2(\log f(j+1)) \leq T_2(j+1)$ many steps for almost all j. It is not necessary to compute $f(i+1)$. Hence, i is computed in time $O(T_2)$.

   If $i \cdot T_1(|x|) > T_2(|x|)$ the machine rejects x.

   Otherwise, M simulates $2 \cdot T_1(|x|)$ steps of the machine M_i on input x. By assumption this can be done in time $i \cdot T_1(|x|) \leq T_2(|x|)$. If M_i accepts x within that many steps M rejects x, else it accepts.

2. $L \notin \mathrm{eAvDTime}(T_1, \{\mathrm{uni}\})$:

   Assume that $L \in \mathrm{eAvDTime}(T_1, \{\mathrm{uni}\})$. Let M_i be a machine that accepts L. By construction of the sequence n_i we have that for all $x \in [n_i, 2n_i+1]$: $i \cdot T_1(|x|) \leq T_2(|x|)$. Hence, the first condition in the definition of L is always fulfilled for such x. Now $L = L(M_i)$ implies that for such x the case $x \notin L(M_i)$ in the definition cannot occur, hence $\mathrm{time}_{M_i}(x) > 2 \cdot T_1(|x|)$. We claim that the average time complexity of M_i with respect to the uniform distribution is larger than T_1 both for the eAv and the Av–measure. Otherwise, for the first measure we would get the following contradiction. Let $l = 2^{2n_i+1} - 1$, then

$$\begin{aligned}
\sum_{\mathrm{rank}_{\mathrm{uni}}(x) \leq l} \frac{\mathrm{time}_{M_i}(x)}{T_1(|x|)} &\leq l && \Rightarrow \\
\sum_{2^{n_i} \leq \mathrm{rank}_{\mathrm{uni}}(x) < 2^{2n_i+1}} \frac{\mathrm{time}_{M_i}(x)}{T_1(|x|)} &\leq l && \Rightarrow \\
\sum_{2^{n_i} \leq \mathrm{rank}_{\mathrm{uni}}(x) < 2^{2n_i+1}} 2 &\leq 2^{2n_i+1} && \Rightarrow \\
2 \cdot (2^{2n_i+1} - 2^{n_i}) &\leq 2^{2n_i+1} && \Rightarrow \\
2^{2n_i+1} &\leq 2 \cdot 2^{n_i}
\end{aligned}$$

3. $L \notin \mathrm{AvDTime}(T_1, \{\mathrm{uni}\})$:

   For the other measure a contradiction is derived similarly.

$$\begin{aligned}
\sum_{\mathrm{rank}_{\mathrm{uni}}(x) \leq l} \frac{T_1^{-1}(\mathrm{time}_{M_i}(x))}{|x|} &\leq l && \Rightarrow \\
\sum_{2^{n_i} \leq \mathrm{rank}_{\mathrm{uni}}(x) < 2^{n_i+1}} \frac{T_1^{-1}(\mathrm{time}_{M_i}(x))}{|x|} &\leq l && \Rightarrow
\end{aligned}$$

$$
\begin{aligned}
&\leq \quad 1 + \sum_{\substack{\mathrm{rank}_\mu(z) \leq l \\ \text{and } z \notin L_E}} 1 + \sum_{\substack{\mathrm{rank}_\mu(1^{2^j}) \leq l \\ \text{and } 1^{2^j} \notin \{x,y\}}} \frac{\log\log(2^{2 \cdot 2^{2^j}})}{2^j} \\
&= \quad 1 + l - e_l + \sum_{\substack{\mathrm{rank}_\mu(1^{2^j}) \leq l \\ \text{and } 1^{2^j} \notin \{x,y\}}} 1 + \frac{1}{2^j} \\
&\leq \quad l - e_l + 1 + (e_l - 2) + 1 \quad = \quad l \; .
\end{aligned}
$$

Thus, $L \in \mathrm{AvDTime}(T, \mu)$ for all distributions μ. ∎

6 Time Hierarchies of Average Case Classes

For average case complexity classes with a fixed bound on the rankability of the distribution we get tight hierarchy results. First we show that even under uniform distribution we cannot always solve a problem of a higher worst case complexity class. This substantially improves theorem 4.2 in [Milt 91], which gives a weaker separation for the expected measure.

Theorem 9 *For time-constructible and monotone increasing time bounds T_1, T_2 with $T_1 \leq o(T_2)$ holds*

$$
\begin{aligned}
DTime(T_2) \setminus eAvDTime(T_1, \{uni\}) &\neq \emptyset \; , \\
DTime(T_2) \setminus AvDTime(T_1, \{uni\}) &\neq \emptyset \; .
\end{aligned}
$$

Proof: Define a function $f : \mathbb{N} \longrightarrow \mathbb{N}$ as follows.

$$
\begin{aligned}
f(0) &:= 1 \\
n_i &:= \min\left\{ b \geq f(i) \;\middle|\; \forall j \in [b; 2b+1] \quad \frac{T_2(j)}{T_1(j)} \geq i \right\} \\
f(i+1) &:= 2^{2n_i+1}
\end{aligned}
$$

With the help of f we define a diagonal language L by

$$
\begin{aligned}
\boldsymbol{L} \;:=\; \{x \mid &\text{for } i \text{ with } f(i) \leq |x| < f(i+1) \text{ holds:} \\
&i \cdot T_1(|x|) \leq T_2(|x|) \quad \text{and} \left[\; \mathrm{time}_{M_i}(x) > 2 \cdot T_1(|x|) \quad \text{or} \quad x \notin L(M_i) \; \right] \Big\} \; .
\end{aligned}
$$

The motivation for these definitions is as follows:

- Like in a normal diagonalization a string x belongs to L if a corresponding machine M_i on input x computes longer than $2 \cdot T_1(|x|)$ steps or rejects.
- To guarantee $L \in \mathrm{DTime}(T_2)$ such a simulation of M_i should only be performed if there is enough time, that means $i \cdot T_1(|x|) \leq T_2(|x|)$.

Corollary 6

$$eAvDTime(T, C) \;=\; DTime(T)$$

for a set C of distributions that are $O(T \cdot \log^2 T)$–rankable with an oracle for $H_{T(\delta\mathcal{N})}$ $(\delta > 1)$.

Using the fact that the time bound h_T grows at most like $T(n) \cdot 2^n$ and that for polynomial bounds the constant δ can be omitted one obtains

Corollary 7 *For complexity bounds T, T' with $T' \in POL$ holds:*

$$\begin{aligned} eAvDTime(T, (T \cdot EXL)\text{–}rankable) &= DTime(T), \\ AvDTime(T', (T' \cdot EXL)\text{–}rankable) &= DTime(T'). \end{aligned}$$

Notice that for the stricter average case measure the equality is only claimed for polynomial bounds T'. The motivation for considering average case measures other than the expectation was to keep the concept of polynomial time reductions. This aim seems to imply that the precision of the measure has to be weekened. Levin's approach is very coarse grain, while the one presented here will turn out to be as precise as the worst case measure in the most interesting range of polynomial bounds. But for time bounds much larger than polynomials the situation becomes somewhat different.

Theorem 8 *For $T \geq EEXL$ and the set of all distributions U holds:*

$$DTime(T) \subset AvDTime(T, U) \;.$$

Proof:
The relation " $\subseteq$ " follows from the definition.
That the deterministic class is strictly contained in the average time class can be seen as follows. Let $T = \exp\exp$ be the double exponential time bound and $L_E := \{1^{2^i} \mid i \in \mathbb{N}\}$.
Let $L \subseteq L_E$ be a language in $\mathrm{DTime}(T^2) \setminus \mathrm{DTime}(T)$ and μ be an arbitrary distribution. Let x, y be the elements in L_E with the two smallest ranks according to μ. One can construct a machine M for L that takes linear time for all inputs in $\overline{L_E} \cup \{x, y\}$ and time T^2 else.

This means for input x that

$$\frac{T^{-1}(\mathrm{time}_M(x))}{|x|} \quad \leq \quad \frac{\log\log|x|}{|x|} \quad \leq \quad 1/2$$

and similarly for y.

Then, for each l let $e_l := |\{z | \mathrm{rank}_\mu(z) \leq l \text{ and } z \in L_E\}|$:

$$\sum_{\mathrm{rank}_\mu(z) \leq l} \frac{T^{-1}(\mathrm{time}_M(z))}{|z|} \quad \leq \quad 1 + \sum_{\substack{\mathrm{rank}_\mu(z) \leq l \\ \text{and } z \notin L_E}} \frac{T^{-1}(|z|)}{|z|} + \sum_{\substack{\mathrm{rank}_\mu(1^{2^j}) \leq l \\ \text{and } 1^{2^j} \notin \{x,y\}}} \frac{T^{-1}(\mathrm{time}_M(1^{2^j}))}{|1^{2^j}|}$$

long-time(x)

$\underline{\text{if}}$ $|x| = 1$ $\underline{\text{then}}$ $\underline{\text{return}}$ $(0, 0, 0, \infty)$;

$(\mathbf{index}, j, s, r)$:= **long-time**$(z_{\text{itlog}(|x|)-1})$

case 1: error $(\mathbf{index}, x) = \mathbf{true}$ $\underline{\text{then}}$
$\underline{\text{if}}$ $(c(\mathbf{index}, \text{itlog}(|x|-1)), x-1) \in H_{\delta T}$
$\underline{\text{then}}$ $\underline{\text{return}}$ $(\mathbf{index}+1, 0, s+1, \infty)$;
$\underline{\text{then}}$ $\underline{\text{return}}$ $(\mathbf{index}+1, 0, s, \infty)$;

case 2: $(c(\mathbf{index}, \text{itlog}(|x|-1)), x-1) \in H_{\delta T}$ $\underline{\text{then}}$
$\underline{\text{return}}$ $(\mathbf{index}, j+1, s+1, \infty)$;

case 3: $\text{time}_{M_{\text{index}}}(x) \leq \delta T(|x|)$ $\underline{\text{then}}$
$\underline{\text{return}}$ $(\mathbf{index}, j, s, \infty)$;

case 4: $j > \frac{s-j}{\delta-1}$ $\underline{\text{then}}$
include x into X_i and $\underline{\text{return}}$ $(\mathbf{index}+1, 0, s+1, s+1)$;

case 5: $\underline{\text{else}}$
include x into X_i and $\underline{\text{return}}$ $(\mathbf{index}, j+1, s+1, s+1)$;

In case 1 the index can be increased since $L(M_i) \neq L$ and one does not have to worry about machine M_i anymore. If case 2 applies there is already one string from the interval $]z_{\text{itlog}(|x|)-1}, z_{\text{itlog}(|x|)}]$ included into $X_{\mathbf{index}}$ and we will not choose anymore. In the last two cases M_i spends too much time on input x and hence this input will be put into $X_{\mathbf{index}}$. If we now have enough such strings (case 4) the index will be increased, otherwise we still have to find more strings for this set.

This procedure will not get stuck at some index i because this would imply that only for finitely many x holds $\text{time}_{M_i}(x) > \delta T(|x|)$. But then another machine would accept L in worst case time $\delta \cdot T$. By a linear speedup L could also be accepted in worst case time T, a contradiction to the choice of L.

By construction, the sets X_i fulfill the desired conditions (a), (b) and (c).

Lemma 6 *$\text{rank}_D \in O(h_{\delta T})$ –rankable*

Proof: Let $d(n)$ be an upper bound on the run time of $\mathbf{long-time}$ on input of length n.

1. The computation of **long-time**$(z_{\text{itlog}(|x|)-1})$ requires at most $d(\log|x|)$ steps.
2. **test**$(\mathbf{index}, x)$ can be computed in time $|x| \cdot \mathbf{index}^2 / \log|x|$, where $\mathbf{index} \leq \text{itlog}|x|$.
3. The simulation of M_{index} for $\delta \cdot T(|x|)$ steps costs time $\delta \cdot T(|x|) \cdot \text{itlog}(|x|)^2 \leq h_{\delta T}(|x|)$.
4. To simulate the oracle, time $h_{\delta T}(|x|)$ is sufficient by definition of $h_{\delta T}$.

Therefore, $\text{rank}_D \in O(h_{\delta T})$ –rankable. ∎

The proof for the Av-measure is almost identical replacing δT by $T(\delta \cdot \mathcal{N})$ in the construction above. ∎

Miltersen has shown that there exists a distribution μ malign for $\text{DTime}(\mathcal{N}^k)$ with respect to the expected time measure, which can be computed in polynomial time with an $\Sigma_2^{\mathcal{P}}$ –oracle ([Milt 91]). The theorem above shows that an $\mathcal{NP}$ –oracle suffices.

3. AvDTime(T, $h_{T(\delta\cdot\mathcal{N})}$–rankable) $\subseteq$ DTime($T(\delta\mathcal{N})$) :

   As above, consider any language $L \notin \mathrm{DTime}(T(\delta\mathcal{N}))$. Now replace the first property of X_i by

   (a') $\forall x \in X_i \quad \mathrm{time}_{M_i} > T(\delta|x|)$.

   Assume that there exists a DTM M_i with $L(M_i) = L$ and $(\mathrm{time}_{M_i}, \mathrm{rank}_D) \in \mathrm{Av}(T)$, that means

$$\forall l \qquad \sum_{\mathrm{rank}_D(x)\le l} \frac{T^{-1}(\mathrm{time}_{M_j}(x))}{|x|} \;\le\; l$$

   Again, for $l = \sum_{j=1}^{i} |X_j|$ we derive a contradiction

$$\sum_{\mathrm{rank}_D(x)\le l} \frac{T^{-1}(\mathrm{time}_{M_i}(x))}{|x|} \;\ge\; \sum_{x\in X_i} \frac{T^{-1}(T(\delta|x|))}{|x|} \;\ge\; \sum_{x\in X_i} \frac{\delta|x|}{|x|} \;=\; \delta|X_i| \;>\; l \;.$$

The sequence X_i is constructed as follows. Define the iterated exponential function and its inverse by $\mathrm{itexp}(n) := \underbrace{\exp((\cdots\exp(}_{n \text{ times}} 1)\cdots))$, $\mathrm{itlog} := \mathrm{itexp}^{-1}$ and strings $z_i := 1^{\mathrm{itexp}(i)}$. Note that for any string x with respect to the lexicografical ordering $z_{\mathrm{itlog}(|x|)-1} < x \le z_{\mathrm{itlog}(|x|)}$.

Let M^* be a DTM for L .

To make the computation of the ranks efficient for each machine M_i we we will use another machine with identical time behaviour for large inputs, but a linear time bound for small inputs. It is irrelevant whether this machine accepts the same language as M_i . More formally, let $c : \mathbb{N}\times\mathbb{N} \to \mathbb{N}$ be a function such that for all i

$$\mathrm{time}_{M_{c(i,k)}}(x) := \begin{cases} \mathrm{time}_{M_i}(x) & \text{if } x \ge z_k\,, \\ |x| & \text{if } x < z_k\,. \end{cases}$$

It is easy to see that such indices $c(i,k)$ of size bounded by $O(i\cdot k)$ can be computed time $O(i\cdot k)$.

Below we describe a recursive algorithm **long-time**, which on input x outputs a tupel $(\mathbf{index}, s, j, r)$. The meaning of these components are as follows. **index** denotes the index of the set X_{index} into which x might be inserted. j counts the number of strings lexicographically smaller than or equal to x that are contained in X_{index} , while s counts those in D . r equals $\mathrm{rank}_D(x)$. Thus in particular, the last component gives the rank function. The other information will be needed in the recursive calls to achieve the porperties discussed above.

As a subroutine we will check whether a machine M_i behaves different from M^* for any inputs z much smaller than the current input x . This is done by the following procedure.

```
error(i, x)
for all z with |z| ≤ log log |x|
    simulate machines M* and M_i on input z for |x|/(log |x|)^2 steps;
    if either (both have accepted) or (both have rejected or have not terminated)
        then return true
return false
```

The main algorithm looks as follows.

Theorem 7 *For all $\delta > 1$ and $T \geq 2\delta\mathcal{N}$ holds*

$$\begin{aligned} eAvDTime(T, h_{\delta T}\text{-}rankable) &= DTime(T), \\ AvDTime(T, h_{T(\delta\mathcal{N})}\text{-}rankable) &\subseteq DTime(T(\delta\mathcal{N})). \end{aligned}$$

Proof:

1. eAvDTime(T, $h_{\delta T}$ –rankable) $\supseteq$ DTime(T)
   follows from the definition.

2. eAvDTime(T, $h_{\delta T}$ –rankable) $\subseteq$ DTime(T) :
   We will prove this inclusion indirectly. Let L be a language that is not in DTime(T) . It suffices to show that their exists a distribution in $h_{\delta T}$ –rankable with respect to which L is not average T-time bounded. To obtain the rank function of this distribution we will construct a sequence of finite sets X_i containing all inputs of finite rank and hence positive weight. Then

   $$D := \bigcup_{i \in \mathbb{N}} X_i$$

   is the support of this distribution and the ranking will be the lexicografical one

   $$\text{rank}_D(x) \;:=\; |\{z \in D \mid z \leq x\}| \quad \text{for } x \in D$$

   according to Definition 2. The X_i and the ranks of their elements will be obtained by the algorithm `long-time` defined below. We will show in Lemma 6 that a DTM can compute the rank function rank_D in time $h_{\delta T}$.

   The idea behind the construction of the X_i is as follows. X_i will be used to diagonalize against machine M_i . If $L(M_i) = L$ then the following properties will be achieved:

   (a) $\forall x \in X_i \qquad \text{time}_{M_i}(x) > \delta\, T(|x|)$,
   that means each set X_i contains only inputs on which M_i spends much time (sufficiently more than the average).

   (b) $|X_i| \;>\; (\delta - 1)^{-1} \sum_{j=1}^{i-1} |X_j|$:
   these sets are chosen large enough such that there are enough long computations to prevent a good average case behavior.

   (c) For all $j < l$ holds: $\qquad |x| < |y| \qquad \forall\, x \in X_j \quad \forall\, y \in X_l$,
   this way, the rank function is strictly monotone increasing and the problems with too many small ranks or holes in the ranking do not occur.

   Assume that for each i we have constructed such an X_i and that there exists a machine M_i that accepts L in average time T with respect to the distribution defined by rank_D . This would requrire

   $$\forall l \qquad \sum_{\text{rank}_D(x) \leq l} \frac{\text{time}_{M_i}(x)}{T(|x|)} \;\leq\; l \;.$$

   Let $l := \sum_{j=1}^{i} |X_j|$. Then using property (b), we get the contradiction

   $$\begin{aligned} \sum_{\text{rank}_D(x) \leq l} \frac{\text{time}_{M_i}(x)}{T(|x|)} &\geq \sum_{x \in X_i} \frac{\delta T(|x|)}{T(|x|)} \\ &= \delta|X_i| \;=\; |X_i|(\delta - 1) + |X_i| \;>\; (l - |X_i|) + |X_i| \;=\; l \;. \end{aligned}$$

5.2 More Complex Distributions

Next, we will show that allowing arbitrary distributions there is no difference between the average case and the worst case.

To this aim we want to construct a rank function that for any DTM M with $L(M) \notin \mathrm{DTime}(T)$ gives small ranks to inputs with long computations exceeding the time bound T. So, there is no difference between DTime and eAvDTime (or AvDTime) with respect to this distribution. The main difficulty is to control the number of small ranks. If, for example, a machine outputs the number 1 twice this machine does not compute a rank function. On the other hand this machine must not leave any gaps in the ranking it outputs. For example, in the sequence of ranks $2, 2, 5, 5, 6, 7, 8, 9 \ldots$ rank 3 is missing. One solution to this problem is first to compute all ranks the machine outputs for smaller strings. But this requires exponential time.

Another solution is to use the following language as oracle.

Definition 11

$$\mathbf{H_T} := \{(x, 1^i) \mid \exists z \leq x : \ time_{M_i}(z) > T(|z|)\} \ .$$

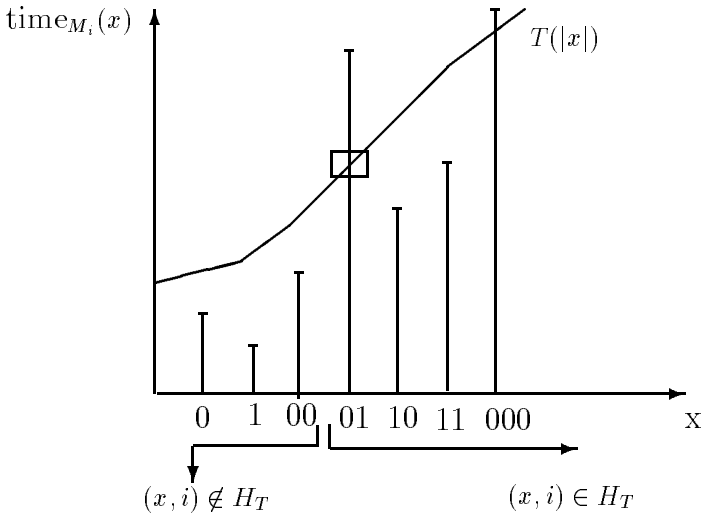

Figure 2: The definition of H_T.

Figure 2 visualize this set H_T. Note that H_T is in $\mathcal{NP}$ if T is polynomially bounded and constructible. It is also complete for $\mathcal{NP}$ since the **bounded halting problem**

$$\texttt{NBH} \ := \ \{(x01^t0^i) \mid time_{M_i}(x) \leq t\}$$

can be reduced to it.

Let $\Omega(T^2) \ \leq \ \boldsymbol{h}_T \ \leq \ O(exp \cdot T)$ be a time bound such that $H_T \in \mathrm{DTime}(h_T)$.

Proof: For $T \in \omega(\mathcal{N})$ holds: Solving the equation

$$T \circ \exp \frac{\mathcal{N}}{2} \;=\; T' \cdot \exp$$

with T' as unknown one gets a complexity bound $T' \in \omega(T)$. Construct a language L as in the last two proofs: For the AvDTime classes this is done with respect to the complexity bound T, for the eAvDTime classes with respect to T'. In both cases we get the same language and it holds

$$\begin{aligned} L &\in \mathrm{AvDTime}(T, \{\mathrm{uni}\}) \;, \\ L &\notin \mathrm{eAvDTime}(o(T'), \{\mathrm{uni}\}) \supseteq \mathrm{eAvDTime}(T, \{\mathrm{uni}\}) \;. \end{aligned}$$

∎

If we restrict complexity classes to unary languages defined over a one letter alphabet then measuring the average complexity by the expectation obviously does not make sense because it will be exactly equal to the worst case measure. However, both of our new measures are able to provide a meaningful tool as the following result shows.

Theorem 6 *There exists a unary languages L in $DTime(\mathcal{N}^2)$ whose average case complexity cannot be bounded by any function in $o(\mathcal{N}^2)$ when taking the expectation. However,*

$$L \;\in\; eAvDTime(2\mathcal{N}, \{uni\}) \;\cap\; AvDTime(2\mathcal{N}, \{uni\}) \;.$$

Proof: Let $L_1 \subseteq 1^*$ be a unary language in $DTime(\exp \mathcal{N}^2) \setminus DTime(o(\exp \mathcal{N}^2))$. Define the unary language L by

$$L \;:=\; \{1^n \mid n = 2^m \text{ for some } m \in \mathbb{N} \text{ and } 1^m \in L_1\} \;.$$

Then $L \in DTime(\mathcal{N}^2) \setminus DTime(o(\mathcal{N}^2))$, which also implies that the time complexity measured by the expectation cannot be bounded by $o(\mathcal{N}^2)$. A simple calculation on the other hand yields

$$\begin{aligned} \sum_{\mathrm{rank}_{\mathrm{uni}}(x) \le l} \frac{\mathrm{time}_M(x)}{2^{|x|}} \;&\le\; \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \le l \\ |x| \notin \{2^k \mid k \in \mathbb{N}\}}} \frac{|x|}{2^{|x|}} \;+\; \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \le l \\ |x| \in \{2^k \mid k \in \mathbb{N}\}}} \frac{|x|^2}{2^{|x|}} \\ &\le\; \frac{l}{2} + \frac{1}{2} \sum_{n=1}^{\log l - 1} 2^n \;\le\; l \;. \end{aligned}$$

∎

This example shows that in the case of unary languages these measures tolerate larger peaks in the runtime of a machine if they do not happen too often. Thus, for unary languages we also get an averaging effect, but now over different input lengths. On the other hand, it is not hard to see that restricted to unary languages there can be at most a linear factor difference between these average case and the worst case measure.

Note that this result does not contradict Theorem 1, where we have shown that for the uniform distribution the standard average time complexity based on $Time_M^{\mu}(n)$ defines the same complexity classes as $\mathrm{eAv}(T)$. In that case the average was taken with respect to a ground set of strings defined over an alphabet of size at least 2.

More precisely, there exists a machine M for L that rejects all strings not of the form $1^k0^{2^k-k}$ in linear time, and takes time $\frac{1}{2}T \cdot \exp$ for the remaining inputs. This yields $L \in \mathrm{eAvDTime}(T, \{\mathrm{uni}\})$ since for all l

$$\begin{aligned}
& \sum_{\mathrm{rank}_{\mathrm{uni}}(x) \leq l} \frac{\mathrm{time}_M(x)}{T(|x|)} \\
\leq & \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \leq l \\ x \notin \{1^k0^{2^k-k} \mid k \in \mathbb{N}\}}} \frac{|x|}{T(|x|)} + \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \leq l \\ x \in \{1^k0^{2^k-k} \mid k \in \mathbb{N}\}}} \frac{1}{2} \cdot \frac{2^{|x|} \cdot T(|x|)}{T(|x|)} \\
\leq & \frac{l}{2} + \frac{1}{2} \sum_{n=1}^{\log l - 1} 2^n \leq l \; .
\end{aligned}$$

Note that the previous theorem implies $L \notin \mathrm{eAvDTime}(o(T), \{\mathrm{uni}\})$. ∎

The AvDTime-classes behave very differently in this respect. For any bound T the blowup is not necligible, it is fully exponential, as it is the case, for example, when comparing nondeterministic or probabilistic time classes with deterministic ones. This may be another indication that this concept is the more natural one for complexity investigations.

Theorem 5 *For* $T \geq 2\mathcal{N}$

$$eAvDTime\Big(T, \{uni\}\Big) \quad \setminus \quad DTime\left(o(T \circ \exp \frac{\mathcal{N}}{2})\right) \quad \neq \quad \emptyset \; .$$

Proof: Again, let L_1 be again an unary language in

$$L_1 \in \mathrm{DTime}(\exp(T \circ \exp \frac{\mathcal{N}}{2}))) \setminus \mathrm{DTime}(o(\exp(T \circ \exp \frac{\mathcal{N}}{2})))$$

and define $L := \{xy \mid x \in L_1 \text{ and } y = 0^{2^{|x|}-|x|}\}$.
$L \in \mathrm{DTime}(T \circ \exp \frac{\mathcal{N}}{2}) \setminus \mathrm{DTime}(o(T \circ \exp \frac{\mathcal{N}}{2}))$, and all strings not of the form $1^k0^{2^k-k}$ can be rejected in linear time by an appropriate machine M . Then $L \in \mathrm{AvDTime}(T, \{\mathrm{uni}\}$ follows from

$$\begin{aligned}
& \sum_{\mathrm{rank}_{\mathrm{uni}}(x) \leq l} \frac{T^{-1}(\mathrm{time}_M(x))}{|x|} \\
\leq & \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \leq l \\ x \notin \{1^k0^{2^k-k} \mid k \in \mathbb{N}\}}} \frac{T^{-1}(|x|)}{|x|} + \sum_{\substack{\mathrm{rank}_{\mathrm{uni}}(x) \leq l \\ x \in \{1^k0^{2^k-k} \mid k \in \mathbb{N}\}}} \frac{T^{-1}(T(\exp(|x|/2)))}{|x|} \\
\leq & \frac{l}{2} + \frac{1}{2} \sum_{n=1}^{\log l - 1} 2^n \leq l \; .
\end{aligned}$$

Note that $L \notin \mathrm{AvDTime}(T(o(\mathcal{N})), \{\mathrm{uni}\})$ due to Theorem 3. ∎

Corollary 5 *Let* $T \in \omega(\mathcal{N})$ *and the function* $T/\mathcal{N}$ *be monotone increasing, then*

$$eAvDTime(T, \{uni\}) \subset AvDTime(T(2 \cdot \mathcal{N}, \{uni\}) \; .$$

5.1 Uniform Distributions

Theorem 3

$$\begin{aligned} eAvDTime(T,\{uni\}) &\subseteq DTime(T\cdot EXL)\ , \\ AvDTime(T,\{uni\}) &\subseteq DTime(T\circ EXL)\ . \end{aligned}$$

Proof: For $L \in \text{eAvDTime}(T,\{\text{uni}\})$ let M be a expected T-bounded DTM with $L(M) = L$. Then for all strings z it must hold

$$\frac{\text{time}_M(z)}{T(|z|)} \quad \leq \quad \sum_{x\leq z} \frac{\text{time}_M(x)}{T(|x|)} \quad \leq \quad \text{rank}_{\text{uni}}(z)\ .$$

Thus, for every input z we get the worst case estimate

$$\text{time}_M(z) \ \leq \ \text{rank}_{\text{uni}}(z)\cdot T(|z|) \ \leq \ \exp O(|z|)|\ \cdot T(|z|)\ .$$

Now, assume $L \in \text{AvDTime}(T,\{\text{uni}\})$ and let M be DTM for L such that for all strings z,

$$\frac{T^{-1}(\text{time}_M(z))}{|z|} \quad \leq \quad \sum_{x\leq z} \frac{T^{-1}(\text{time}_M(x))}{|x|} \quad \leq \quad \text{rank}_{\text{uni}}(z)\ .$$

Then,

$$\text{time}_M(z) \ \leq \ T(\ \text{rank}_{\text{uni}}(z)\cdot|z|) \ \leq \ T(\exp O(|z|)\ )\ .$$

∎

Note that for the eAvDTime-classes there is an exponential increase in the worst case bound if the time bound T is a polynomial. On the other hand, if T itself is exponential the additional factor EXL has only little influence. It is not hard to show that for small time bounds the exponential increase cannot be avoided.

Theorem 4 *For any complexity bound $T \geq 2^{\mathcal{N}}$ holds*

$$eAvDTime(T,\{uni\}) \quad \setminus \quad DTime(o(T\cdot\exp)) \quad \neq \quad \emptyset\ .$$

Proof: Let $L_1 \subseteq 1^*$ be a unary language in

$$\text{DTime}\left(\exp(\frac{1}{2}T\cdot\exp)\right) \ \setminus \text{DTime}\Big(\exp(o(T\cdot\exp))\Big)$$

and define

$$L \ := \ \{xy \mid x \in L_1 \text{ and } y = 0^{2^{|x|}-|x|}\}\ .$$

Then,

$$L \ \in \ \text{DTime}\left(\frac{1}{2}T\cdot\exp\right) \setminus \text{DTime}(o(T\cdot\exp))\ .$$

Theorem 2 *For arbitrary* $\delta, \epsilon > 0$ *and* $T \geq (1/\delta + \epsilon)\ \mathcal{N}$ *with the additional property that the function* $T/\mathcal{N}$ *is monotone increasing and all classes of distributions* C *holds*

$$eAvDTime(T, C) \subseteq AvDTime(T((1+\delta)\mathcal{N}), C) .$$

Proof: For $L \in$ eAvDTime(T,C) and $\mu \in C$, let M be a DTM with $L(M) = L$ and $(\text{time}_M, \mu) \in$ eAv(T). Then, by a linear speedup there also exists a machine M' for L with $\text{time}_{M'}(x) \leq \delta \cdot \text{time}_M(x)$ for all inputs on which M spends time at least $T(|x|)$. Again, divide the inputs into the two subsets

$$\begin{aligned} I_1 &:= \{x \mid \text{time}_{M'}(x) > T(|x|)\} , \\ I_2 &:= \{x \mid \text{time}_{M'}(x) \leq T(|x|)\} . \end{aligned}$$

For $x \in I_1$ one can conclude as in Lemma 5

$$\frac{\text{time}_{M'}(x)}{T(|x|)} \geq \frac{T^{-1}(\text{time}_{M'}(x))}{|x|} .$$

Thus,

$$\begin{aligned} \delta \cdot l &\geq \sum_{x \in I_1, \text{rank}_\mu(x) \leq l} \frac{\text{time}_{M'}(x)}{T(|x|)} \geq \sum_{x \in I_1,\ \text{rank}_\mu(x) \leq l} \frac{\text{time}_{M'}(x)}{T(|x|)} \\ &\geq \sum_{x \in I_1,\ \text{rank}_\mu(x) \leq l} \frac{T^{-1}(\text{time}_{M'}(x))}{|x|} . \end{aligned}$$

For $x \in I_2$ it holds

$$\frac{T^{-1}(\text{time}_{M'}(x))}{|x|} \leq 1 .$$

Therefore,

$$\begin{aligned} \sum_{x \in I_1,\ \text{rank}_\mu(x) \leq l} \frac{T^{-1}(\text{time}_{M'}(x))}{|x|} &\leq l \quad \text{and} \\ \sum_{\text{rank}_\mu(x) \leq l} \frac{T^{-1}(\text{time}_{M'}(x))}{|x|} &\leq l(1+\delta) . \end{aligned}$$

∎

The investigations in the next section will show that in general the second measure defines broader complexity classes. Already for the uniform ranking and a polynomial bound T it thus holds

$$\text{eAvDTime}(T,\ \text{rank}_{\text{uni}}) \subset \text{AvDTime}(T,\ \text{rank}_{\text{uni}}) .$$

5 The Relation between Average Case and Worst Case Complexity Classes

From a complexity point of view sticking to a fixed distribution also does not make much sense for the following reason. If one restricts the average case analysis to a simple distribution like the uniform one the best relation between average and worst case classes that seems to be obtainable is an exponential gap.

Lemma 5 *Let T be a time bound and the function $T/\mathcal{N}$ $(n \mapsto \frac{T(n)}{n})$ be monotone increasing, then*

$$eAv(T) \subseteq Av(T(2 \cdot \mathcal{N})) \ .$$

Proof: Let $(f, \rho) \in \mathrm{eAv}\,(T)$. Thus,

$$\forall l \qquad \sum_{\rho(x) \leq l} \frac{f(x)}{T(|x|)} \ \leq \ l \ .$$

To prove the claim one has to show

$$\forall l \qquad \sum_{\rho(x) \leq l} \frac{T^{-1}(f(x))}{|x|} \ \leq \ 2 \cdot l \ .$$

Divide the inputs into the two subsets

$$\begin{aligned} I_1 &:= \{x \mid f(x) > T(|x|)\} \ , \\ I_2 &:= \{x \mid f(x) \leq T(|x|)\} \ . \end{aligned}$$

If $f(z) > T(|z|)$ then $T^{-1}(f(z)) \geq T^{-1}(T(|z|) + 1) \geq |z|$. Using the assumption that $g(n) := T(n)/n$ is monotone increasing a simple calculation yields for such x

$$\begin{aligned} g(T^{-1}(f(z))) &\geq g(|z|) && \Longleftrightarrow \\ \frac{T(T^{-1}(f(z)))}{T^{-1}(f(z)))} &\geq \frac{T(|z|)}{|z|} && \Longrightarrow \\ \frac{f(z)}{T^{-1}(f(z))} &\geq \frac{T(|z|)}{|z|} && \Longleftrightarrow \\ \frac{f(z)}{T(|z|)} &\geq \frac{T^{-1}(f(z))}{|z|} \ . \end{aligned}$$

Then

$$\sum_{x \in I_1, \rho(x) \leq l} \frac{T^{-1}(f(x))}{|x|} \ \leq \ \sum_{x \in I_1, \rho(x) \leq l} \frac{f(z)}{T(|z|)} \ \leq \ l \ .$$

For $x \in I_2$ it holds

$$\frac{T^{-1}(f(x))}{|x|} \ \leq \ 1 \ .$$

Therefore

$$\sum_{x \in I_2, \rho(x) \leq l} \frac{T^{-1}(f(x))}{|x|} \ \leq \ l$$

∎

Theorem 1 *For all* $T \geq \alpha\mathcal{N}$ *holds*

$$eAvDTime(T, \{\mu_{\text{uni}}\}) \;=\; \{L \mid \exists \text{ a DTM } M \text{ with } L(M) = L \text{ and } Time_M^{\mu_{\text{uni}}} \leq T\} \;.$$

Proof: "$\subseteq$":
Let $L \in \text{eAvDTime}(T, \{\mu_{\text{uni}}\})$. Thus, there exists a machine M with $L(M) = L$ and $(\text{time}_M, \mu_{\text{uni}}) \in \text{eAv}(T)$. Thus,

$$\begin{aligned}
\forall l \sum_{\text{rank}_{\text{uni}}(x) \leq l} \frac{\text{time}_M(x)}{T(|x|)} \;\leq\; l \quad &\Longleftrightarrow \quad \forall n \quad \sum_{|x| \leq n} \frac{\text{time}_M(x)}{T(|x|)} \;\leq\; |\Sigma^{\leq n}| \\
&\Longrightarrow \quad \forall n \quad \sum_{|x|=n} \frac{\text{time}_M(x)}{|\Sigma|^n} \;\leq\; |\Sigma| \cdot T(n) \;.
\end{aligned}$$

If M is sped up by the factor $2 \cdot |\Sigma|$ we get a machine M' with $L(M') = L$ and

$$\text{time}_{M'}(x) \;\leq\; \max\left\{\frac{\text{time}_M(x)}{2 \cdot |\Sigma|}, (1+\epsilon)|x|\right\}$$

for some $\epsilon > 0$. If T grows at least as $\alpha\mathcal{N}$ for an appropriate α (for example $\alpha = 8|\Sigma|^2$) then one can show

$$\forall n \qquad \sum_{|x|=n} \frac{\text{time}_{M'}(x)}{|\Sigma|^n} \;\leq\; T(n) \;,$$

which means that M' is expected T-time bounded.

"$\supseteq$":
Let M be a DTM such that $L(M) = L$ and $Time_M^{\mu_{\text{uni}}} \leq T$, which means

$$\begin{aligned}
\forall n \quad \sum_{|x|=n} \frac{\text{time}_M(x)}{|\Sigma|^n} \;\leq\; T(n) \quad &\Longleftrightarrow \quad \forall n \quad \sum_{|x|=n} \frac{\text{time}_M(x)}{T(n)} \;\leq\; |\Sigma|^n \\
&\Longrightarrow \quad \forall n \quad \sum_{|x| \leq n} \frac{\text{time}_M(x)}{T(n)} \;\leq\; |\Sigma^{\leq n}| \\
&\Longleftrightarrow \quad \forall l \sum_{\text{rank}_{\text{uni}}(x) \leq l} \frac{\text{time}_M(x)}{T(n)} \;\leq\; l \\
&\Longleftrightarrow \quad (\text{time}_M, \mu_{\text{uni}}) \;\in\; \text{eAv}(T) \;.
\end{aligned}$$

∎

4.2 eAvDTime versus AvDTime

Let us now compare these two concepts of bounding the average run time. It will be shown that the generalized expected time measure is at least as strict as the average measure. Thus, algorithms with known expected running time can be used to get upper bounds for the average time complexity with respect to the $\text{Av}(T)$ measure.

4 Modification of the Expected Measure

4.1 Bounds for the Expectation Using Rank Functions

The generalization by monotone transformations can also be applied to the classical expected measure.

Definition 9 *The set* $\mathbf{eAv}(\boldsymbol{T})$ *(expected average* T*) contains all pairs* (f,μ) *such that for all monotone transformations* m *of* μ

$$\sum_x m(\mu(x))\ \frac{f(x)}{T(|x|)}\ \leq\ 1\ .$$

Similar to above this condition can be simplified to

Proposition 3

$$(f,\mu)\in eAv(T) \qquad \Longleftrightarrow \qquad \forall l \qquad \sum_{\mathrm{rank}_\mu(x)\leq l}\ \frac{f(x)}{T(|x|)}\ \leq\ l\ .$$

Proof: "$\Rightarrow$" by definition.
"$\Leftarrow$": Identical to the proof of proposition 1 and corollary 1 when substituting $\frac{T^{-1}(f(x))}{|x|}$ by $\frac{f(x)}{T(|x|)}$.
∎

As for the classical expected measure one can easily find examples showing that this measure is not closed under polynomial transformations, too. The complexity classes $\mathrm{eAvDTime}(T)$ are then defined similarly to $\mathrm{AvDTime}(T)$.

Definition 10 *Let* C *be a set of distributions, then*

$$\mathbf{eAvDTime}(\boldsymbol{T},\boldsymbol{C})\ :=\ \{L\mid \forall\mu\in C\ \ \exists DTM\ \ M\ \ \textit{with}\ L(M)=L\ \ \text{ and }\ \ (time_M,\mu)\in eAv(T)\}\ .$$

It is not hard to see that the average run time being bounded by T for every input length, that means $\forall n\ \ Time_M^\mu\leq T(n)$, is a stricter condition than membership in $\mathrm{eAv}(T)$. A machine contradicting the property $Time_M^\mu\leq T(n)$ for some n may still be average bounded with respect to the eAv-measure if for smaller input length n' the average is sufficiently smaller than $T(n')$.

There is at least one special case, where both measures define the same complexity class. Consider uniform distributions μ_{uni} with the rank function

$$\mathrm{rank}_{\mathrm{uni}}(x)\ :=\ |\Sigma^{\leq|x|}|\ .$$

Thus as ranks only values of the form $l=|\Sigma^{\leq|x|}|$ appear. In this case one can prove for an appropriate constant α that only depends on Σ the following equality.

We then define the following complexity classes.

Definition 5

$$\begin{aligned} \mathbf{DistDTime}(\boldsymbol{T}) &:= \{(L,\mu) \mid \exists M \in DTM:\ L(M) = L \text{ and } (time_M, \mu) \in Av(T)\} \\ \mathbf{DistP} &:= \bigcup_{T \in \mathrm{POL}} DistDTime(T)\ . \end{aligned}$$

These sets are called **distributional complexity classes**. Their elements are pairs of an algorithmic problem (a language) and a probability distribution. They are formally different from the standard worst case classes. But the ultimate goal is to compare the worst case and the average case complexity of problems. Therefore the following definition turns out to be suitable for this purpose. We consider the distributional complexity of languages L with respect to a set C of distributions and require that the average complexity is bounded for all these distributions.

Definition 6 *Let T be a complexity bound and C a set of distributions, resp. rankings. Then*

$$\mathbf{AvDTime}(\boldsymbol{T}, \boldsymbol{C}) := \{L \mid \forall \mu \in C \quad (L,\mu) \in DistDTime(T)\}\ .$$

In order to restrict the complexity of the input distributions a natural approach is to put a time limit for computing the corresponding rank functions.

Definition 7 *Let T* **−rankable** *be the set of all distributions μ, resp. rankings ρ for which there exists a DTM M that on input x computes $bin(\rho(x))$ in time $T(|x|)$.*

Most average complexity classes we will be studying in the following will be of the form

$$\mathrm{AvDTime}(T,\ V\text{−rankable})\ ,$$

where T and V are complexity bounds.

Alternative notions for restricting distributions have been defined in [Levi 86], [Gure 91] and [BCGL 92].

Definition 8 *A distribution μ belongs to the class T* **-computable** *if there is a deterministic TM that on input $(x, 1^i)$ outputs the first i bits of the binary expansion of $\mu^*(x) := \sum_{z \le x} \mu(z)$ in time $T(|(x,1^i)|)$.*
μ is T **-sampleable** *if one can find a probabilistic TM that outputs each string $x \in \Sigma^*$ with probability $\mu(x)$, and this within $T(|x|)$ steps.*

These concepts are not directly comparable because each rank function represents a whole equivalence class of distributions. In a subsequent paper we will discuss the relation between these conditions and the rankability property in more detail.

b) follows similarly by applying g^{-1} , which is monotone, to $g \circ f$ and $g \circ T$ and noticing that for strictly increasing functions $g^{-1} \circ g = \mathcal{N}$. ∎

In particular, we get the result that this measure is closed under linear or polynomial transformations.

Corollary 3 *For any constants* c, c_1, c_2 *holds:*

$$\begin{aligned}
(f,\mu) \in Av(T) \quad &\Longrightarrow \quad (c_1 f + c_2, \mu) \in Av(c_1 T + c_2) \qquad \text{and} \\
&\qquad\quad (f^c, \mu) \in Av(T^c) \\
g \in Pol(f) \quad \text{and} \quad (f,\mu) \in Av(T) \quad &\Longrightarrow \quad (g,\mu) \in Av(Pol(T)) \\
g \in Pol(f) \quad \text{and} \quad (f,\mu) \in Av(Pol) \quad &\Longrightarrow \quad (g,\mu) \in Av(POL)
\end{aligned}$$

To estimate sums and products of complexity bounds is a little more complicated. For this purpose, we first will consider the maximum operator. For two functions f, g let $\max[f,g]$ denote the function defined by $n \mapsto \max\{f(n), g(n)\}$.

Lemma 4 *For* $0 < \beta < 1$ *holds:*

$$(f,\mu) \in Av(T_f) \quad \text{and} \quad (g,\mu) \in Av(T_g) \qquad \Longrightarrow \qquad (\max[f,g],\mu) \;\in\; Av\left(\max[T_f(\tfrac{\mathcal{N}}{\beta}), T_g(\tfrac{\mathcal{N}}{1-\beta})]\right)$$

Proof: Let $D \;:=\; \max[T_f(\frac{\mathcal{N}}{\beta}), T_g(\frac{\mathcal{N}}{1-\beta})]$. Then,

$$\begin{aligned}
(f,\mu) \in \mathrm{Av}(T_f) \quad &\Longrightarrow \quad \forall l \sum_{\mathrm{rank}_\mu(x) \le l} \min\left\{\alpha_x \mid f(x) \le T_f\left(\frac{\alpha_x}{\beta} \cdot |x|\right)\right\} \;\le\; l \cdot \beta \\
&\Longrightarrow \quad \forall l \sum_{\mathrm{rank}_\mu(x) \le l} \min\{\alpha_x | f(x) \le D(\alpha_x \cdot |x|)\} \;\le\; l \cdot \beta \;.
\end{aligned}$$

Similarly,

$$(g,\mu) \in \mathrm{Av}\,(T_g) \qquad \Longrightarrow \qquad \forall l \sum_{\mathrm{rank}_\mu(x) \le l} \min\{\alpha_x | g(x) \le D(\alpha_x \cdot |x|)\} \;\le\; l \cdot (1-\beta) \;.$$

Therefore,

$$\forall l \sum_{\mathrm{rank}_\mu(x) \le l} \min\{\alpha_x \mid g(x) \le D(\alpha_x \cdot |x|) \;\text{ and }\; f(x) \le D(\alpha_x \cdot |x|)\} \;\le\; l \;,$$

which implies by Proposition 2 that $(\max[f,g],\mu) \in \mathrm{Av}\,(D)$. ∎

Hence, for sum and product using that $f + g \le 2 \cdot \max\{f,g\}$, resp. $f \cdot g \le \max\{f^2, g^2\}$ we obtain

Corollary 4

$$\begin{aligned}
(f,\mu) \in Av(T_f) \quad \text{and} \quad (g,\mu) \in Av(T_g) \quad &\Longrightarrow \quad (f+g,\mu) \in Av\Big(2 \cdot \max[T_f(2 \cdot \mathcal{N}), T_g(2 \cdot \mathcal{N})]\Big) \\
&\text{and} \quad (f \cdot g,\mu) \in Av\Big(\max[T_f(2 \cdot \mathcal{N})^2, T_g(2 \cdot \mathcal{N})^2]\Big)|.
\end{aligned}$$

Corollary 2 *If μ_1, μ_2 are distributions with $rank_{\mu_1} = rank_{\mu_2}$ then*

$$(f, \mu_1) \in Av(T) \qquad \Longleftrightarrow \qquad (f, \mu_2) \in Av(T) \ .$$

A rank function ρ represents a whole equivalence class of distributions, namely those μ with $\mathrm{rank}_\mu = \rho$. Therefore, in the following we do not differentiate between pairs (f, μ) containing distributions and pairs (f, ρ) referring to the corresponding rank functions.

The folowing gives another useful characterization of these sets which helps to simplify estimations.

Proposition 2

$$(f, \mu) \in Av(T) \qquad \Longleftrightarrow \qquad \forall l \sum_{rank_\mu(x) \leq l} \min\{\alpha_x \mid f(x) \leq T(\alpha_x \cdot |x|)\} \ \leq \ l \ .$$

Proof: By Corollary 1

$$(f, \mu) \in \mathrm{Av}\,(T) \qquad \Longleftrightarrow \qquad \forall l \sum_{\mathrm{rank}_\mu(x) \leq l} \frac{T^{-1}(f(x))}{|x|} \ \leq \ l \ .$$

By the definition of the inverse of T

$$\Longleftrightarrow \qquad \forall l \sum_{\mathrm{rank}_\mu(x) \leq l} \frac{\min\{m \mid f(x) \leq T(m)\}}{|x|} \ \leq \ l \ .$$

Now we substitute m by $\alpha_x \cdot |x|$

$$\Longleftrightarrow \qquad \forall l \sum_{\mathrm{rank}_\mu(x) \leq l} \min\{\alpha_x \mid f(x) \leq T(\alpha_x \cdot |x|)\} \ \leq \ l \ .$$

∎

Before we will use this notion to define complexity classes let us investigate closure and transformation properties of these sets of average case bounds.

Lemma 3
a) For a monotone function $g : \mathcal{N} \to \mathcal{N}$ holds:

$$(f, \mu) \in Av(T) \qquad \Longrightarrow \qquad (g \circ f, \mu) \in Av(g \circ T) \ .$$

b) If, in addition, $g > 0$ is strictly increasing also the inverse implications holds:

$$(g \circ f, \mu) \in Av(g \circ T) \qquad \Longrightarrow \qquad (f, \mu) \in Av(T) \ .$$

Proof: a) By proposition 2,

$$\begin{aligned}
(f, \mu) \in \mathrm{Av}(T) \quad &\Longrightarrow \quad \forall l \sum_{\mathrm{rank}_\mu(x) \leq l} \min\{\alpha_x | f(x) \leq T(\alpha_x \cdot |x|)\} \ \leq \ l \\
&\Longrightarrow \quad \forall l \sum_{\mathrm{rank}_\mu(x) \leq l} \min\{\alpha_x | g(f(x)) \leq g(T(\alpha_x \cdot |x|))\} \ \leq \ l \\
&\Longrightarrow \quad (g \circ f, \mu) \in \mathrm{Av}(g \circ T) \ .
\end{aligned}$$

We will show in lemma 3, that this generalization does not destroy polynomial bounds that are increased by a polynomial.

Because of the universal quantifier over all monotone transformations the above condition for $\mathrm{Av}(T)$ is hard to verify. But there exists an equivalent, more practical characterization of $\mathrm{Av}(T)$. Consider the special case of *threshold functions* $\mathtt{thr}_l : [0,1] \to [0,1]$ as monotone transformations, where for $l = \mathrm{rank}_\mu(x)$ we define $\mathtt{thr}_l(z) := 1/l$ if $z \geq \mu(x)$ and 0 else.

Proposition 1

$$(f,\mu) \in Av(T) \qquad \Longleftrightarrow \qquad \forall l \quad \sum_x \mathtt{thr}_l(\mu(x)) \, \frac{T^{-1}(f(x))}{|x|} \;\leq\; 1 \; .$$

Proof: "$\Rightarrow$" follows from the definition.
"$\Leftarrow$": Let $x_1, \; x_2, \; \ldots$ be an enumeration of all strings s with $\mu(x) \neq 0$ such that $\mu(x_i) \geq \mu(x_{i+1})$. Let $a^n := (a_1, \ldots, a_n)$ where $a_i := T^{-1}(f(x_i))/|x_i|$. Assume $(f,\mu) \notin \mathrm{Av}(T)$. Then there exists an integer n with $\mu(x_n) > \mu(x_{n+1})$ such that for some transformation m

$$\sum_{i=1}^{n} m(\mu(x_i)) a_i > 1 \; .$$

To obtain the maximum value of this sum with respect to m we must solve a linear optimization problem. The solution consists of a set of vectors which describe an $(n-r)$–dimensional simplex (r is the number of indices $i < n$ where $\mu(x_i) = \mu(x_{i+1})$). The extremal points are $(\underbrace{1/i, \ldots, 1/i}_{i \text{ times}}, 0, \ldots, 0)$.
The solution is the set of points of largest distance from the $(n-1)$-dimensional hyperplane with the normal vector a that contains the origin $(0, \ldots, 0)$. This set contains at least one extremal point of the simplex. Hence, there exists an integer l with $a_1 = a_2 = \cdots = a_l = 1/l$ such that the sum above achieves its maximum. ∎

As an immediate consequence of this proposition we obtain the following fundamental result, which shows that an average bound can be computed without considering all possible transformations.

Corollary 1

$$(f,\mu) \in Av(T) \qquad \Longleftrightarrow \qquad \forall l \quad \sum_{\mathrm{rank}_\mu(x) \leq l} \frac{T^{-1}(f(x))}{|x|} \;\leq\; l \; .$$

Proof: Observe that for a threshold function $\mathtt{thr}_l$

$$\sum_x \mathtt{thr}_l(\mu(x)) \, \frac{T^{-1}(f(x))}{|x|} \;=\; \sum_{\mathrm{rank}_\mu(x) \leq l} \frac{T^{-1}(f(x))}{l \cdot |x|} \; .$$

So, Corollary 1 follows from proposition 1. ∎

In the following we will use this simpler characterization to decide membership in $\mathrm{Av}(T)$. Since the condition only involves the ranking of the input space derived from the distribution one also gets

Definition 2

$$\mathbf{rank}_\mu(\boldsymbol{x}) \quad := \quad \begin{cases} |\{z \in \Sigma^* \mid \mu(z) \geq \mu(x)\}| & \text{if } \mu(x) > 0, \\ \infty & \text{if } \mu(x) = 0, \end{cases}$$

where ∞ plays the role of a number that is greater than any natural number. A machine computing the rank function is assumed to output a special symbol in the second case. **Uniform distributions** *give identical weights to inputs of the same length. Shorter inputs receive strictly larger probabilities than longer ones. Corresponding to such distributions we have the* **uniform ranking** *of the input space Σ^* defined by*

$$\mathbf{rank_{uni}}(\boldsymbol{x}) \quad := \quad 1 + |\Sigma| + |\Sigma|^2 + \ldots + |\Sigma|^{|x|} \ .$$

Later, we will also consider special rankings derived from subsets $L \subseteq \Sigma^*$. For this purpose, let us introduce the notion

$$\mathbf{rank}_L(\boldsymbol{x}) \quad := \quad \begin{cases} |\{y \in L | y \leq x\}| & \text{if } x \in L, \\ \infty & \text{else.} \end{cases}$$

The set of distributions $\tilde{\mu}$ equivalent to μ can be generated by *monotone transformations* of μ. Figure 1 shows a sample of a probability distribution and its rank function.

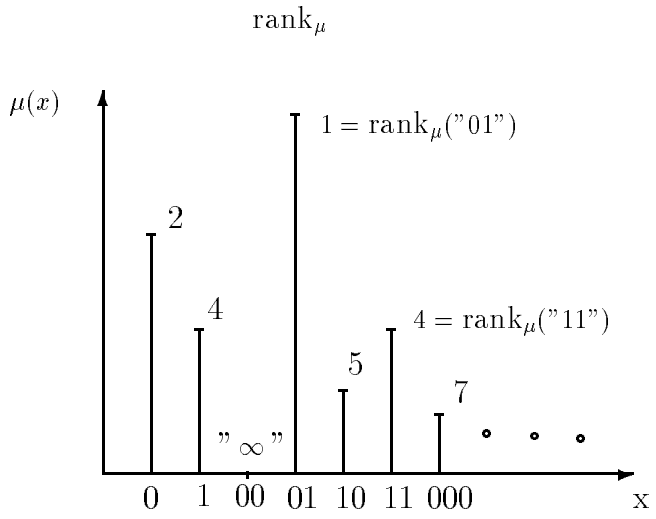

Figure 1: The rank function derived from a distribution μ.

Definition 3 *A real-valued monotone function $m : [0,1] \to [0,1]$ is called a* **monotone transformation** *of the distribution μ if $\sum_x m(\mu(x)) \leq 1$.*

Definition 4 *The set* $\mathbf{Av}(\boldsymbol{T})$ *(average T) contains all pairs (f, μ) consisting of a function $f : \Sigma^* \to \mathbb{N}$ and a distribution μ such that for all monotone transformations m of μ*

$$\sum_x \ m(\mu(x)) \ \frac{T^{-1}(f(x))}{|x|} \ \leq \ 1 \ .$$

Lemma 1 *For a monotone decreasing strictly positive function* $u : \Sigma^* \to [0;1]$ *and a real-valued function* $h : [0,1] \to [0,1]$ *with a monotone increasing first deriative* h' *in* $(0,1)$ *it holds for* $u(x) \neq u(x-1)$

$$h'(u(x-1)) \;\leq\; \frac{h(u(x)) - h(u(x-1))}{u(x) - u(x-1)} \;\leq\; h'(u(x)) \;.$$

Proof: Applying the mean value theorem there exists $z_0 \in (0;1)$ such that

$$h'(z_0) \;=\; \frac{h(u(x)) - h(u(x-1))}{u(x) - u(x-1)} \;.$$

Since h' is monotone increasing, it holds $h'(u(x-1)) \leq h'(z_0) \leq h'(u(x))$. ∎

Lemma 2 *For every monotone decreasing strictly positive probability distribution* $\mu : \Sigma^* \to [0,1]$ *there exists a distribution* $\tilde{\mu}$ *such that*

$$\lim_{x \to \infty} \frac{\tilde{\mu}(x)}{\mu(x)} \;=\; \infty \;.$$

Proof: Let $\mu^*(x) := \sum_{z \leq x} \mu(z)$ be the distribution function of μ according to the lexicographical ordering of strings in Σ. Then we define a new probability distribution $\tilde{\mu}$ by the distribution function

$$\tilde{\mu}^*(x) \;:=\; 1 - \sqrt{1 - \mu^*(x)} \;.$$

Thus the probability distribution is given by $\tilde{\mu}(x) \;:=\; \tilde{\mu}^*(x) - \tilde{\mu}^*(x-1)$. Then by lemma 1 for $h(x) := 1 - \sqrt{1-x}$ and $u(x) := \mu^*(x)$ we get

$$\frac{\tilde{\mu}(x)}{\mu(x)} \;\geq\; \frac{\mu^*(x-1)}{\sqrt{1 - \mu^*(x-1)}} \;.$$

Note that $\frac{\mu^*(x)}{\sqrt{1-\mu^*(x)}}$ tends to infinity. ∎

Hence, Levin defined only the term *polynomial on the average* as a bound for the time complexity:

Definition 1 *A function* $f : \Sigma^* \to \mathbb{N}$ *belongs to the class* **lAv(POL)** *(Levin average polynomial) with respect to a distribution* μ *iff for some number* k

$$\sum_x \; \mu(x) \; \frac{f(x)^{1/k}}{|x|} \;<\; \infty \;.$$

We will now present a stricter average case measure. The idea is to look simultaneously at all distributions $\tilde{\mu}$ which generate the same ordering of the inputs by their probabilities as μ, for example if $\mu(10001) < \mu(11)$ then $\tilde{\mu}(10001) \leq \tilde{\mu}(11)$. μ–average bounded by T is then defined to be $\tilde{\mu}$–average bounded by T in the sense above for all such $\tilde{\mu}$. Thus, only the ranking of the inputs by decreasing probabilities matters. We therefore define for a distribution μ its corresponding rank function rank_μ by

The following sets of complexity bounds will be of special interest:

$$\begin{aligned} \mathrm{POL} &:= \bigcup_{k \in \mathbb{N}} O(\mathcal{N}^k) \ , \\ \mathrm{EXL} &:= \exp(\Theta(\mathcal{N})) \qquad \text{and} \\ \mathrm{EEXL} &:= \exp\exp(\Theta(\mathcal{N})) \ . \end{aligned}$$

If $\mathcal{T}$ is a set of complexity bounds $T \leq \mathcal{T}$ means that for some $T' \in \mathcal{T}$ holds $T \leq T'$. For a complexity bound T, which does not necessarily have to be injective, we define the inverse T^{-1} by

$$T^{-1}(n) := \min\{m \mid T(m) \geq n\} \ .$$

Useful properties of this definition are

$$\begin{aligned} T(T^{-1}(n)) &\geq n \ , \\ T^{-1}(T(n)) &\leq n \ , \\ T^{-1}(T(n)+1) &> n \ . \end{aligned}$$

Let $M_1, M_2, \ldots$ be an enumeration of all deterministic Turing machines (DTM). In some cases we will also consider nondeterministic machines (NTM). We may assume that all machines have only two work tapes, implying that one can use a universal machine with only a constant factor slowdown. More explicitely, we assume that there is a machine U that can simulate each M_i on any input x in time at most $i \cdot \mathrm{time}_{M_i}(x)$, where $\mathrm{time}_{M_i}(x)$ denotes the number of steps of M_i on input x. $M(x)$ denotes the output that M generates on an input $x \in \Sigma^*$ and $L(M)$ (in case of acceptors) the language accepted by M. Unless otherwise stated, we will always assume that the alphabet Σ has size at least 2. At least for the standard approach using the expectation unary alphabets obviously do not make much sense.

For an ordering of binary strings, $x \leq y$, we refer to the lexicographical ordering. We consider probability measures (density functions) $\mu : \Sigma^* \to [0,1]$ over the input space. μ has to satisfy $\sum_x \mu(x) \leq 1$. $\mathrm{bin} : \mathbb{N} \to \{0,1\}^*$ denotes the standard correspondence between natural numbers and binary strings.

3 Refinement of Levin's Average Case Measure

In the introduction we already discussed the problem how to measure the average time complexity of machines with respect to different probability distributions precisely. Levin's solution essentially can only distinguish between polynomial and superpolynomial growth. The problem with the uniform distribution mentioned above can somehow be diminished, by giving Σ^n instead of a total weight n^{-2} weight proportional to $n^{-1} \cdot \log^{-2} n$ or even better proportional to

$$n^{-1} \cdot \log^{-1} n \cdot \ldots \cdot (\log^{[k-1]} n)^{-1} \cdot (\log^{[k]} n)^{-2}$$

for some k, where $\log^{[k]}$ denotes the k-th iteration of the logarithm function. Still, it can never be completely resolved which can be seen as follows.

In practice, one often does not know the values of the distribution exactly, but for each pair of inputs at least one can decide which input is more likely. This way, the whole analysis is greatly simplified. Each ranking of the input space describes a whole equivalence class of distributions, and we get rid of the influence of the asymptotic growth of the probability measure.

A **distributional problem** is a pair (L, μ), consisting of a language defined over an alphabet Σ and a probability distribution μ on Σ^*. We define *distributional complexity classes* **DistDTime**($\boldsymbol{T}$) containing all pairs (L, μ), for which there exists a DTM accepting L that is μ–average T–time bounded in this generalized sense.

Given a language $L \in \mathrm{DTime}(T)$ and a DTM M for L, it is easy to see that by cycling through all inputs of length n one can find an x, on which M spends the maximal time for inputs of length n. If a probability measure μ gives all its weight for inputs of length n to this x then the average time of M (in the expected sense) with respect to this μ equals the worst case complexity. $\mu(x)$ can be computed in time $O(2^{|x|} \cdot T(|x|))$. Using this idea, Miltersen has shown that allowing exponential time overhead a measure μ can be constructed that is *malign* for all expected T–time bounded machines ([Milt 91]). That means their expected time complexity with respect to this μ is no more than a constant factor smaller than their worst case complexity.

On the other hand, restricting an average case analysis to some simple distributions may yield results with little practical value. The satisfiability problem, for example, has been shown quickly solvable for certain symmetric distributions, but the input space generated this way seems to be of not interest for applications in AI (see for example the discussion in [MiSeLe92]). These observation motivate to consider *average case complexity classes* **AvDTime**($\boldsymbol{T}$, $\boldsymbol{C}$) consisting of all languages L that can be recognized in μ–average time T for distributions μ of complexity at most C, for certain bounds C. That way, average case complexity classes are directly comparable to the standard worst case classes since both contain only languages.

In this paper a notation different from the one in previous research on average case complexity will be used because we feel that this new one is more appropriate and natural. There should be a clear distinction between distributional classes, where distributions appear explicitly, and average case classes the elements of which are languages in the usual sense. From a complexity theoretic point of view one is more interested in the second kind of classes.

The complexity of a distribution is measured by its **rankability**, that is the effort to compute the ranks of the elements in the input space. Previous approaches have bounded the complexity of distributions using the notion of **computable** and **sampleable** ([Levi 86],[Gure 91],[BCGL 92]). A preliminary version of some of these results was presented at the 10 th STACS–Conference [ReSc93].

2 Notations

A complexity bound is a function $T : \mathbb{N} \to \mathbb{N}$. Let $\mathcal{N}$ denote the identity function on the natural numbers, i.e. the linear complexity bound. All complexity bounds T considered in this paper are assumed to be monotone increasing, time-constructible and at least as large as $\mathcal{N}$. To simplify the notation, for a constant $\alpha > 0$

$$T(\alpha\mathcal{N})$$

means the function defined by $n \mapsto T(\lfloor \alpha \cdot n \rfloor)$.

1 Introduction and Overview

Levin observed that a sound definition of average case complexity and complexity classes is not at all obvious ([Levi 86]). The classical notion of average-case time complexity of a machine M with respect to given probability distributions μ_n on inputs x of length n takes the expectation

$$\boldsymbol{Time}_M^{\mu}(\boldsymbol{n}) \;:=\; \sum_{|x|=n} \mu_n(x) \cdot time_M(x) \;,$$

where $\mathrm{time}_M(x)$ denotes the running time of M on x and $\mu := \mu_1, \mu_2, \ldots$. The machine M is **μ–average T–time bounded** (in the expected sense) for a resource bound $T : \mathbb{N} \longrightarrow \mathbb{N}$, if $Time_M^{\mu} \leq T$, that means for all n

$$\sum_{|x|=n} \mu_n(x) \cdot \frac{time_M(x)}{T(|x|)} \;\leq\; 1 \;.$$

The problem with this definition is that polynomial time simulations of polynomial average time machines can result in superpolynomial average time complexity. It was resolved by Levin by applying the inverse of T to the fraction, thus requiring

$$\sum_{|x|=n} \mu_n(x) \cdot \frac{T^{-1}(\mathrm{time}_M(x))}{|x|} \;\leq\; 1 \;.$$

This definition does not take into account that the weights of different input length may be very unequal. Thus one considers only distributions μ defined over the whole set of inputs and requires

$$\sum_{x} \mu(x) \cdot \frac{T^{-1}(time_M(x))}{|x|} \;\leq\; 1 \;.$$

M is then called (Levin)– **μ–average T–time bounded**. For a discussion of this approach see the detailed exposition in [Gure 91].

Still there remains an unpleasant property, the influence of the functional growth of $\mu(x)$ on the time bound T. If, for example, one takes the "standard" *uniform probability distribution,* which assigns probability

$$\mu_{\mathrm{uni}}(x) \;:=\; 6/\pi^2 \cdot |x|^{-2} \cdot 2^{-|x|}$$

to a string $x \in \{0,1\}^*$ a machine making n^2 steps on every input of length n would already be average $O(\mathcal{N}^{1+\epsilon})$–time bounded for arbitrary $\epsilon > 0$, where $\mathcal{N}$ denotes the identity function on $\mathbb{N}$. This problem can be resolved to a certain extent (see [Gure 91]), but not completely.

Our first contribution to the average case analysis will be a new definition of average T–time bounded, which gets rid of this problem. It will allow us to differenciate between bounds T_1 and T_2 for any $T_1 \leq o(T_2)$. The idea is to bound the complexity of a machine not only with respect to the probability distribution μ, but with respect to all monotone transformations of μ. At first glance, it seems that this complicates the analysis even more. But we will show that this larger set of conditions is equivalent to a very simple property of the distribution μ, which does not involve probabilities explicitly anymore. The only thing that matters is the ranking of the inputs by μ, that is the sequence of inputs ordered by decreasing probabilities.

Precise Average Case Complexity Measures

Rüdiger Reischuk*
Christian Schindelhauer†

TR-93-049
August 1993

Abstract

A new definition is given for the average growth of a function $f : \Sigma^* \to \mathbb{N}$ with respect to a probability measure μ on Σ^*. This allows us to define meaningful average case distributional complexity classes for arbitrary time bounds (previously, one could not guarantee arbitrary good precision). It is shown that basically only the ranking of the inputs by decreasing probabilities are of importance.

To compare the average and worst case complexity of problems we study average case complexity classes defined by a time bound and a bound on the complexity of possible distributions. Here, the complexity is measured by the time to compute the rank functions of the distributions. We obtain tight and optimal separation results between these average case classes. Also the worst case classes can be embedded into this hierarchy. They are shown to be identical to average case classes with respect to distributions of exponential complexity.

Key words. worst case complexity, expectation, average case complexity, distributions, distributional complexity classes, time hierarchies, rank functions, rankability hierarchies

AMS(MOS) subject classifications. 68Q10, 68Q15, 68Q25, 60E05

*Institut für Theoretische Informatik, Technical University Darmstadt, Alexanderstraße 10, 64283 Darmstadt, Germany, e-mail: reischuk@iti.informatik.th-darmstadt.de
Part of this work was done during a visit at the ICSI in Berkeley
†Institut für Theoretische Informatik, Technical University Darmstadt, Alexanderstraße 10, 64283 Darmstadt, Germany, e-mail: schindel@iti.informatik.th-darmstadt.de