Scalability and the Isomorphism Problem

Judy Goldsmith * and Steven Homer †

November 2, 1995

*Department of Computer Science, University of Kentucky, Lexington, KY . This research was supported in part by NSF grants RII-9003056 and CCR-9315354

†Department of Computer Science, Boston University, Boston, MA 02215. This research was supported in part by NSF grants CCR-9103055 and CCR-9400229.

Abstract

Scalable sets are defined and their properties studied. It is shown that the set of scalable sets is the isomorphism closure of the set of rankable sets and that every scalable set is P-isomorphic to some rankable set. Scalable sets coincide with P-printable sets when sparse, and with P-paddable sets when thick. Using scalability as a tool, the P-isomorphism question for polynomial-time computable sets of similar densities is examined.

1 Introduction

This paper defines and investigates the new concept of scalability for polynomial-time sets. A set is scalable if there is an efficient method for computing the number of elements in the set (or its complement) which are less than a given element, relative to some polynomial-time computable and invertible order on Σ^*. All scalable sets are polynomial-time computable. Scalability generalizes the previously studied concept of ranking which was based on the same property with respect to a fixed ordering, the lexicographic ordering, on strings. Scalability better captures the intuition of an efficiently computable order on elements of a set. This notion is naturally complexity-theoretic in that it is invariant under polynomial-time isomorphisms. Furthermore, as shown below, it unifies several different important and previously unrelated complexity-theoretic concepts.

There are several reasons for studying this generalization of ranking. If polynomial-time computable functions are taken as the fundamental class of feasibly computable functions, then any reasonable complexity-theoretic notion, such as those capturing the structural properties of a problem, should be P-isomorphism invariant[1]. Scalability is the smallest complexity-theoretic generalization of rankability in the sense that the scalable sets are the smallest class containing the rankable sets and closed under P-isomorphism. This notion of closure can be taken as a touchstone to determine the possible usefulness and naturalness of new concepts; in this sense scalability is a fundamental notion for complexity theory.

Scalability can be viewed as a notion of efficient enumeration of a set, similar to several other actively studied concepts. Central among them is rankability, defined by Goldberg and Sipser [GS85] and studied by Huynh [Huy88], and Hemachandra and Rudich [HR90]. A set is rankable if, with respect to the lexicographic ordering on strings, one can compute in polynomial time the number of elements in the set which are less than a given element. Rankability is a useful concept, but as we will see it has the weakness that it is not closed under P-isomorphism unless

[1] We use the notation *P-isomorophism, P-rankable,* etc. to indicate that the function in question is computable in polynomial time; we use *P-set* to indicate a set recognizable in polynomial time.

every set in P is rankable. Other notions of enumerability of P-sets include minimal perfect hash functions (studied in the infinite case by Goldsmith, Hemachandra, and Kunen [GHK90]), P-enumerability (Hemachandra, Hoene, Siefkes and Young [HHSY]), monotone P-enumerability (Hoene [Hoe91]), P-paddability (introduced by Berman and Hartmanis [BH77], and studied by Mahaney and Young [MY85]), and, for sparse sets, P-printability (first studied by Allender [A86], and Allender and Rubinstein [AR88]). One of the first definitions of an enumeration property was Young's definition of enumerable sets [You69]: a set is enumerable in time $t(n)$ if there is a program that enumerates at least n elements of the set in the first $t(n)$ steps, for all n. That definition excludes sufficiently sparse sets, and includes (for t a polynomial) all sets that contain P-printable subsets. Selman [S78] considered a quite different notion of enumerability. His notion is defined on NP sets (in fact it characterizes NP) and is not directly comparable to the enumerations of P-sets. Most of these later concepts were introduced either to give a complexity-theoretic version of a recursion-theoretic enumeration property, or to generalize P-printability to sets of other densities. Scalability coincides with P-printability on the sparse sets, and with P-paddability on "thick" sets, as well as generalizing rankability. It implies both compressibility (the existence of minimal perfect hash functions), and P-enumerability. Thus it is a more general notion of P-time enumerability.

The density of P-sets plays a central role both in relation to enumeration properties of scalable sets, as discussed above, and in other questions related to scalability. One such question is whether and under what conditions sets in P are P-isomorphic. This isomorphism question is almost always asked of complete sets for hard classes, and is difficult to answer. For P-sets several of the difficulties with the isomorphism problems for less efficiently computable classes disappear. We are left with a "core" isomorphism problem, which is more elementary than the general problem and whose solution may yield a better understanding of the isomorphism problems for complete sets. In particular, since we are dealing with sets in P, polynomial-time reductions between them always exist and it is always possible to efficiently decide membership in the sets. Since sets in P can have widely differing densities, obviating any chance they have of being P-isomorphic, we restrict our isomorphism question to pairs of polynomial-time sets whose densities are equivalent (that is, between which there are 1-1, polynomial honest, polynomial-bounded, but not necessarily efficiently computable, functions; a function is polynomially honest, or just honest, if its inverse is also polynomially bounded). It is for such sets that we consider the isomorphism problem and gain some understanding using the concept of scalability and the methods which arise in their study: two scalable sets of equivalent densities are necessarily P-isomorphic.

We begin the next section by formally defining a scaled well-ordering and a scalable set.

The relationship between scalability and rankability is examined, as well as the identification of thick scalable sets and thick P-paddable P-sets and of sparse scalable sets and P-printable sets. We consider the isomorphism problem for polynomial-time sets. Finally we study the strength of the assumption that non-scalable P-sets exist, considering the consequence of their existence with different densities.

2 Scalability and Its Properties

In this section we define scalability and relate it to other, previously considered concepts. The idea is that a set is scalable if there is a polynomial-time computable, invertible and honest ordering of strings relative to which we can determine how many elements are in the set that are less than a given element.

Definition 1 *A well-ordering W of Σ^* is* scaled *if there is a polynomial-time computable and polynomial-time invertible bijection $f : \mathcal{N} \to \Sigma^*$ which maps a number n to the n^{th} string in W. We call f the* scale *for W.*

Note that the lexicographic ordering is trivially scaled.

Definition 2 *A set A is* scalable *if there is a scaled well-ordering W and a polynomial time function $g : \Sigma^* \to \mathcal{N}$ such that for any x, $g(x)$ is the number of elements in A which are less than or equal to x with respect to the ordering W. We say that A is scaled with respect to W.*

It is immediate that all scalable sets are in P. (To test whether $x \in A$, compute $g(x)$ and $g(x')$, where x' is the predecessor of x with respect to W. Then $x \in A$ iff $g(x) = g(x') + 1$.) Thus we are concerned here with the structure and properties of polynomial-time computable sets.

Definition 3 (Goldberg and Sipser [GS85]) *A set B is* rankable *if it is scalable with respect to the lexicographic ordering.*

We show that scalability is preserved under polynomial isomorphism. We later show this is not the case for rankability, assuming that $P \neq P^{\#P}$. We begin by proving that scalable sets are exactly those which are P-isomorphic to rankable sets.

Definition 4 *A* polynomial time isomorphism *is a bijection $h : \Sigma^* \to \Sigma^*$ which is polynomial time computable and polynomial-time invertible.*
A set A is polynomial time isomorphic *to a set B (written $A \cong_p B$) if there is a polynomial time isomorphism h with $h(A) = B$.*

Proposition 1 *Let A be scalable and $A \cong_p B$. Then B is scalable.*

Proof: Let h be the polynomial time isomorphism from A to B. Since A is scalable there is a well-ordering W of Σ^* which is scaled by a function $f : \mathcal{N} \to \Sigma^*$. And there is a function $g : \Sigma^* \to \mathcal{N}$ witnessing the scalability of A.

Now h induces a well-ordering V of Σ^* defined by $x <_V y$ iff $h^{-1}(x) <_W h^{-1}(y)$. V is scaled, as the function $h \circ f$ is a scale for V. That V is a well-ordering follows from the fact that W is, and that h is a bijection from Σ^* to itself. Moreover $h \circ f$ is a scale for V since f is a scale for W and h is a polynomial time isomorphism.

We claim that B is scalable using V as the well-ordering and $g \circ h^{-1} : \Sigma^* \to \mathcal{N}$ as the function proving scalability. It is straightforward to check that, for any $x \in \Sigma^*$, the number of elements in B less than or equal to x with respect to V is equal to the number of elements in A less than or equal to $h^{-1}(x)$ with respect to W, and this last quantity equals $g(h^{-1}(x))$. □

We next clarify the relationship between scalable and rankable sets.

Proposition 2 *A set A is scalable if and only if there is a rankable set B with $B \cong_p A$.*

Proof:

⇒) Let $\nu : \Sigma^* \to \mathcal{N}$ map a binary string x to its associated natural number and let $b(n)$ be the inverse function, from natural numbers to their binary representations. Let $f : \mathcal{N} \to \Sigma^*$ be the scale for an order W and let $g : \Sigma^* \to \mathcal{N}$ be a function witnessing that A is scalable with respect to f. Thus, given $x \in \Sigma^*$, x is the $f^{-1}(x)^{th}$ element in order W, and $b \circ f^{-1}(x)$ is the $f^{-1}(x)^{th}$ string in Σ^* under the lexicographic order.

Define $B = \{y : \ f \circ \nu(y) \in A\}$. Since $f \circ \nu$ is a P-isomorphism, $B \cong_p A$. Furthermore, $r_B(y) = (g \circ f \circ \nu)(y)$ is the rank of y in B.

⇐) This follows immediately from Proposition 1 since every rankable set is also scalable and Proposition 1 says that scalability is preserved under polynomial isomorphism. □

In the introduction, we mentioned several enumeration properties that are generalized by scalability. We have just seen that scalability is the P-isomorphism closure of ranking. We now give the definitions of two more of these properties, and immediately observe that scalability generalizes them.

Definition 5 (Goldsmith, Hemachandra, and Kunen [GHK90]) *A set A has a minimal perfect hash function, or is* compressible, *if there is a p-time function h that maps A one-one onto Σ^*.*

Note that an infinite scalable set is compressible: if A is infinite and scalable, with scale f, then $\nu \circ f$ maps A one-one onto Σ^*. The converse, even for sets in P, may fail. For instance, if there is a non-P-printable set in P with exactly one element of each length, the set will be compressible but not scalable.

Definition 6 (Hemachandra, Hoene, Siefkes, and Young [HHSY]) *A set A is* P-enumerable *if there is a p-time function f and string x such that $A = \{x, f(x), f \circ f(x), f \circ f \circ f(x)...\}$.*

Note that P-enumerable sets cannot have superpolynomial gaps in the lengths of their elements. Scalable sets that do not have such gaps are P-enumerable. For instance, the set $\{1^k \mid k \text{ a stack of 2's}\}$ is scalable, but is too sparse to be P-enumerable.

We now consider scalable sets with very low and very high densities, namely sparse and thick sets.

Definition 7 *1. A set A is* P-sparse *if there is a polynomial p such that for all n, $||A^n|| \leq p(n)$.*

2. A set $B \subseteq \Sigma^*$ is thick *if there is a polynomial p, a one-to-one mapping r from Sigma* to B such that $|r(x)| \leq p(|x|)$, and a one-to-one mapping r' from Sigma* to $\overline{B}$ such that $|r'(x)| \leq p(|x|)$.*

Note that the mappings r and r' need not be efficient. Intuitively a thick set has the same density (up to a polynomial factor) as does SAT (see [G88]).

We next want to consider P-printable sets and relate them to sets which are scalable and sparse.

Definition 8 *A set S is* P-printable *if there is a polynomial time function which on input 1^n outputs the elements of S which have length $\leq n$.*

We first consider P-sparse sets and prove that scalability coincides with P-printability. Observe that all infinite compressible sets have infinite P-printable subsets [GHK90]. Thus, all infinite scalable sets have P-printable subsets, since scalability implies compressibility.

Theorem 9 *If A is sparse, then A is P-printable if and only if A is scalable.*

Proof: Allender showed that rankability and P-printability are equivalent for sparse sets [A86]. This proves one direction. Suppose, therefore, that S is scalable and sparse. Let $S \cong_P S'$, where S' is rankable. Further, let $q(n)$ bound the P-isomorphism, h, between S and S'. To P-print S up to length n, first P-print S' up to length $q(n)$. Apply h to each string thus produced. If the resulting string is of length less than n, put it into a list. Then sort the (polynomial size) list, and print it. $\square$

The study of isomorphisms of NP-complete sets, and the isomorphism conjecture that all NP-complete sets are P-isomorphic, led to the consideration of *P-paddable* sets [BH77]. There are two reasons that the isomorphism conjecture for NP-complete sets might fail. Either, there is a NP-complete set that is not of the same density as SAT, i.e., not thick, or there is a thick complete set that is not P-paddable.

We write $A \equiv B$ if A and B are of equivalent density, that is, if there is a P-length bounded, P-honest bijection $f : \Sigma^* \to \Sigma^*$ with $f(A) = B$. (Note that we make no assumption concerning the complexity of computing f.)

Definition 10 *A set A is P-paddable if there are P-time computable functions p and d such that*

- *for all $x, y \in \Sigma^*$, $x \in A \leftrightarrow p(x, y) \in A$, and*
- *$d(p(x, y)) = y$, for all $x, y \in \Sigma^*$.*

Theorem 11 *If $A \in P$ is thick, then A is P-paddable if and only if A is scalable.*

Proof: Remember that A is thick if and only if the lexicographic isomorphism between A and $E = \{x0 : x \in \Sigma^*\}$ is polynomially bounded. Observe that E is both scalable (in fact, rankable) and P-paddable.

Let A be a thick, P-paddable set in P. Notice that $A \leq_m^p E$ via the 1-1, length-increasing, polynomial-time computable and invertible function

$$f(x) = \begin{cases} x0 & x \in A \\ x1 & x \notin A. \end{cases}$$

If we can show that $E \leq^P A$ by a one-one, length-increasing, polynomial-time invertible function, then by a theorem of Berman and Hartmanis [BH77], the two sets are P-isomorphic. Let $p(x, n)$ be a padding function for A. Let $g(x0) = p(a, x)$ for some fixed element a of A, and $g(x1) = p(b, x)$ for some fixed element b of $\bar{A}$. Then g has the desired properties, so $A \cong_P E$, and thus is scalable.

Conversely, let B be a thick, scalable set. We will prove that B is P-paddable. As B is thick, B and E have the same density.

We can then define a padding function h' for B. Assume $x \in B$ and x is the t^{th} element of B with respect to B's scale. We define $h'(x, y)$ to be the $\langle t, y\rangle^{th}$ element of B. We can find this element using binary search; since B is thick, the upper bound on the binary search is bounded by some polynomial in $|x| + |y|$, so it can be found in polynomial time. We define h' similarly for elements of $\bar{B}$.

Note that E has a simple P-padding function h derived from a polynomial-time pairing function $< \, , \, >$, defined by $h(x0, y) = \langle x, y \rangle 0$ and $h(x1, y) = \langle x, y \rangle 1$. Clearly h is also polynomial-time invertible. $\Box$

We have seen that scalability is preserved under P-isomorphism. We now turn to showing that rankability is not so preserved. Hemachandra and Rudich have shown that there are non-rankable sets in P if and only if $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$. We show that, if $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$, then rankability and scalability differ. Thus, if $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$, the property of being rankable is not preserved under P-isomorphism.

Theorem 12 *There is a P-set that is scalable but not rankable if and only if* $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$.

Proof: To prove this, we use the set, SAT_W, the so-called "witness set" for SAT. Formally,

$$SAT_W = \{B\#a | \ B \text{ is a Boolean formula and } a \text{ a satisfying assignment for } B\}.$$

Clearly, SAT_W is in P for any reasonable interpretation of $x\#y$; if the pairing function is interpreted appropriately, it is easily shown that SAT_W is not rankable if and only if $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$. (This is the heart of the proof that P is rankable iff $P = P^{\#P}$.) Furthermore, since SAT is thick, and it and its complement polynomially embed in SAT_W and *its* complement, respectively, SAT_W is also thick. Finally, SAT_W is P-paddable, using the standard padding argument for SAT: $\mathrm{pad}(B\#a, k) = B\&(z_0 \vee \neg z_0)\&...\&(z_k \vee \neg z_k)\#a10^k$.

If $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$ then SAT_W is not rankable. Since SAT_W is thick and P-paddable, by Theorem 11 it must be scalable.

If there is a P-set that is not rankable, by the aforementioned result of Hemachandra and Rudich, $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$. $\Box$

Corollary 13 *Rankability is P-isomorphism invariant if and only if every P-set is rankable.*

Proof: By Theorem 11, SAT_W is scalable, and thus by Proposition 2, P-isomorphic to a rankable set B. If there is a non-rankable P-set, then $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$ in which case SAT_W is not rankable, although it is P-isomorphic to B. $\Box$

We want to understand if and when P-isomorphisms exist between P-sets. By Theorem 11, the isomorphism of thick P-sets can be reduced to previous results concerning the isomorphism of P-paddable sets. Namely, it is proved in Mahaney and Young [MY85] that any two polynomial time many-one equivalent P-paddable sets are P-isomorphic. Hence any two thick, scalable sets are P-isomorphic. P-isomorphisms between P-sets are more difficult to analyze for sets with intermediate densities, or gaps. This situation leads to complications and considerations not present when considering P-isomorphisms of sets of similar densities.

The isomorphism problem for P-sets is: if two P-sets A and B are of equivalent densities, are they P-isomorphic? The strongest possible situation regarding this hypothesis would be that any two P-sets of equivalent density are P-isomorphic. We call this situation the *isomorphism hypothesis (for P)* and consider the strength of this hypothesis and how it relates to scalable sets. The result of Mahaney and Young, together with that of Theorem 11, tells us that the isomorphism hypothesis is true for thick sets if and only if they are all scalable. Our first result says that this isomorphism hypothesis holds for P-sets of all densities if all P-sets are scalable. That is, two scalable sets of the same density are P-isomorphic.

Theorem 14 *If $A \equiv B$ and A and B are both scalable then $A \cong_p B$.*

Proof: This follows from the fact that if $C \equiv D$ and both C and D are rankable, then $C \cong_P D$. If A and B are scalable, then, by Proposition 2, there are rankable sets C and D with $A \cong_P C$ and $B \cong_P D$. Thus, $C \cong_P A \equiv B \cong_P D$, so $C \equiv D$. □

Finally we consider the question of whether all P-sets are scalable. We expect that, at least under some reasonable hypothesis, there are non-scalable P-sets. We would like to find a condition which is equivalent to every set in P being scalable, in the same way that $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$ is such a condition for rankability of all P-sets. We mention here some partial results. Observe that, by Theorem 12, if the isomorphism hypothesis fails then there is a non-scalable set. Allender and Rubinstein [AR88] have already considered the question of when there are sparse sets in P that are not P-printable. They proved this is the case if and only if there are sparse sets in FewP$-$P. Since sparse sets are scalable iff they are P-printable, this answers the question of when all sparse sets in P are scalable.

There are some well-known connections between the isomorphism conjecture and the existence of one-way functions. For example, if there are no one-way functions then the isomorphism conjecture holds for exponential time. There have been several sets proposed as non-P-paddable, $\leq_m^P$-complete sets for NP. One such set, discussed, for example, in Joseph and Young [JY85] is $f(SAT)$, where f is a one-way function. What effect would the existence of a one-way function have on the scalability of sets in P? The following results of Allender give a partial answer, again for P-sparse sets.

Theorem 15 (Allender [A86]) *There is a sparse set in P that is not P-printable if and only if there is an honest, polynomial-to-one, polynomial-time computable function that cannot be* *polynomially inverted on* $\{0\}^*$.

3 Conclusions and Open Questions

The class of scalable sets, the P-isomorphism closure of the rankable sets, has been defined and the relationships between scalable, rankable, P-printable and thick P-sets have largely been determined here. Scalability seems the most universal of the generalizations of rankability that have been studied in a P-time setting. However, it is not clear whether scalability is a weaker hypothesis than rankability; we do not know whether, for instance, $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$ implies the existence of a non-scalable set in P.

An oracle result, for example an oracle relative to which all P-sets are scalable and $\mathrm{P} \neq \mathrm{P}^{\#\mathrm{P}}$, would be of interest, since the Hemachandra-Rudich result relativizes. It is not too difficult to construct oracles relative to which there exist non-scalable sets. A simple variant of the original Baker-Gill-Solovay construction yields an oracle relative to which there is a sparse non-scalable set. A slightly more complicated, though still straightforward, construction yields an oracle relative to which there is a thick set which is not scalable. That is, relative to this oracle $O \in$ EXP we have $P^O \neq NP^O$ and O is not P^O-scalable. The proof is a standard diagonalization argument and is omitted here.

Several other open questions remain. What is implied by the existence of a non-scalable P-set? Is there a characterization for the failure of the isomorphism hypothesis for P in terms of reducibilities? What does the isomorphism hypothesis for P-sets imply? Does it imply the isomorphism conjecture for NP-complete sets, or E-complete sets?

Acknowledgements

The authors would like to thank Andrew Klapper for his careful proofreading.

References

[AR88] E. Allender and R. Rubinstein. P-printable sets. *SIAM Journal on Computing*, 17(6):1193–1202, 1988.

[A86] E. Allender. The complexity of sparse sets in P. *Proceedings 1st Conf. on Structural Complexity Theory*, pages 1-11, Springer-Verlag, 1986.

[BH77] L. Berman and J. Hartmanis. On isomorphisms and density of NP and other complete sets. *SIAM Journal on Computing*, 6(2):305–322, 1977.

[GS85] A. Goldberg and M. Sipser. Compression and ranking. In *Proceedings of the 17th ACM Symposium on Theory of Computing*, pages 440–448, 1985.

[G88] J. Goldsmith. Polynomial isomorphisms and near-testable sets. Ph.D. Thesis, University of Wisconsin, 1988. Also available as University of Wisconsin Technical Report # 816 (1989).

[GHK90] J. Goldsmith, L. Hemachandra and K. Kunen. Polynomial-time compressibility *computational complexity* 2:18–39, 1992.

[HHSY] L. Hemachandra, A. Hoene, D. Siefkes, and P. Young. On sets polynomially enumerable by iteration. *Theoretical Computer Science*, 80(2):203–225, 1991.

[HR90] L. Hemachandra and S. Rudich. On the complexity of ranking. *Journal of Computer and System Sciences*, 41(2):251–271, 1990.

[Hoe91] A. Hoene. Monotonic enumerations of complete sets. University of Washington, Technical Report 91-02-02, 1991.

[Huy88] D. Huynh. The complexity of ranking. In *Proceedings of the 3rd Structure in Complexity Theory Conference*, pages 204–212. IEEE Computer Society Press, June 1988.

[JY85] D. Joseph and P. Young. Some remarks on witness functions for non-polynomial and non-complete sets in NP. *Theoretical Computer Science*, 39:225–237, 1985.

[MY85] S. Mahaney and P. Young. Reductions among polynomial isomorphism types. *Theoretical Computer Science*, 39:207–224, 1985.

[S78] A. Selman. Polynomial time enumeration reducibility. *SIAM J. Comput.* 7(4):440-457, 1978.

[You69] P. Young. Toward a theory of enumerations. *Journal of the ACM,* 16:328–348, 1969.