Average-Case Analysis of Algorithms Using Kolmogorov Complexity*

JIANG Tao (姜 涛)[1], LI Ming (李 明)[2] and Paul M.B. Vitányi[3]

[1]*Department of Computing Science, University of California, Riverside, CA 92521, USA*

[2]*Department of Computer Science, University of California, Santa Barbara, CA 93106, USA*

[3]*CWI and University of Amsterdam, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands*

E-mail: jiang@cs.ucr.edu; mli@cs.ucsb.edu; paulv@cwi.nl

Abstract   Analyzing the average-case complexity of algorithms is a very practical but very difficult problem in computer science. In the past few years, we have demonstrated that Kolmogorov complexity is an important tool for analyzing the average-case complexity of algorithms. We have developed the incompressibility method. In this paper, several simple examples are used to further demonstrate the power and simplicity of such method. We prove bounds on the average-case number of stacks (queues) required for sorting sequential or parallel Queuesort or Stacksort.

Keywords   Kolmogorov complexity, algorithm, average-case analysis, sorting

1   Introduction

We obtain some results on the average number of stacks or queues (sequential or parallel) required for sorting under the uniform distribution on input permutations. These problems have been studied before by Knuth[1] and Tarjan[2] for the worst case. This is a further application of the incompressibility method on sorting algorithms following [3].

The goal of the paper is not to obtain these bounds, but to demonstrate a new methodology of analyzing the average-case complexity of algorithms via these problems.

Kolmogorov Complexity and the Incompressibility Method: The technical tool to obtain our results is the incompressibility method. This method is especially suited for the average case analysis of algorithms and machine models, whereas average-case analysis is usually more difficult than worst-case analysis using more traditional methods. A survey of the use of the incompressibility method is in [4] Chapter 6, and recent work is [5]. A recent application is a general nontrivial lower bound on the average-case complexity of multipass Shellsort[3], a problem that had seen remarkably little progress in the last forty years[1].

Informally, the Kolmogorov complexity $C(x)$ of a binary string x is the length of the shortest binary program (for a fixed reference universal machine) that prints x as its only output and then halts[6]. A string x is *incompressible* if $C(x)$ is at least $|x|$, the approximate length of a program that simply includes all of x literally. Similarly, the *conditional* Kolmogorov complexity of x with respect to y, denoted by $C(x|y)$, is the length of the shortest program that, with extra information y, prints x. And a string x is incompressible relative to y if $C(x|y)$ is large in the appropriate sense. For details see [4]. Here we use that, both absolutely and relative to any fixed string y, there are incompressible strings of every length,

and that *most* strings are nearly incompressible, by *any* standard[1]. Another easy one is that significantly long subwords of an incompressible string are themselves nearly incompressible by *any* standard, even relative to the rest of the string[2]. In the sequel we use the following easy facts (sometimes only implicitly).

Lemma 1. *Let c be a positive integer. For every fixed y, every finite set A contains at least $(1 - 2^{-c})|A| + 1$ elements x with $C(x|A, y) \geq \lfloor \log |A| \rfloor - c$.*

Lemma 2. *If A is a finite set, then for every y, every element $x \in A$ has complexity $C(x|A, y) \leq \log |A| + O(1)$.*

Lemma 1 is proved by simple counting. Lemma 2 holds since x can be described by first describing A in $O(1)$ bits and then giving the index of x in the enumeration order of A.

2 Average Case of Bubble Sort: An Example

Generally speaking, the difficulty of analyzing the average-case complexity comes from the fact that one has to analyze the time complexity for all inputs of each length and then compute the average. This is a formidable task.

Using the incompressibility method, we choose just one input — a representative input. Via Kolmogorov complexity, we can show that the time complexity of this input is in fact the average-case complexity of all inputs of the algorithm on this length. Finding such a "representative input" is impossible, but we know it exists and this is sufficient.

The mothodology is best described via examples. We present the following average-case analysis of Bubble Sort.

Example 1. A description and average-case analysis of Bubble Sort can be found in [1]. It is well-known that Bubble Sort uses $\Theta(n^2)$ comparisons/exchanges on the average. We present a very simple proof of this fact. The number of exchanges is obviously at most n^2, so we only have to consider the lower bound. In Bubble Sort we make at most $n - 1$ passes from left to right over the permutation to be sorted and move the largest element we have currently found right by exchanges. For a permutation π of the elements $1, \ldots, n$, we can describe the total number of exchanges by $M := \sum_{i=1}^{n-1} m_i$ where m_i is the initial distance of element $n - i$ to its final position. Note that in every pass more than one element may "bubble" right but that means simply and in the future passes of the sorting process an equal number of exchanges will be saved for the element to reach its final position. That is, every element executes a number of exchanges going right that equals precisely the initial distance between its start position and its final position. A simple analysis as in [3] shows that $\log(M/n) \geq \log n + O(1)$ for a random permutation. As before this holds for an overwhelming fraction of all permutations, and hence gives us an $\Omega(n^2)$ lower bound on the expected number of comparisons/exchanges.

3 Sorting with Queues and Stacks

Knuth[1] and Tarjan[2] have studied the problem of sorting using a network of queues or stacks. In particular, the main variant of the problem is: assuming the stacks or queues are arranged sequentially as shown in Fig.1 or in parallel as shown in Fig.2, then how many

---

[1] By a simple counting argument one can show that whereas some strings can be enormously compressed, like strings of the form $11 \ldots 1$, the majority of strings can hardly be compressed at all. For every n there are 2^n binary strings of length n, but only $\sum_{i=0}^{n-1} 2^i = 2^n - 1$ possible shorter descriptions. Therefore, there is at least one binary string x of length n such that $C(x) \geq n$. Similarly, for every length n and any binary string y, there is a binary string x of length n such that $C(x|y) \geq n$.

[2] Strings that are incompressible are patternless, since a pattern could be used to reduce the description length. Intuitively, we think of such patternless sequences as being random, and we use "random sequence" synonymously with "incompressible sequence". It is possible to give a rigorous formalization of the intuitive notion of a random sequence as a sequence that passes all effective tests for randomness, see for example [4].

stacks or queues are needed to sort n elements with comparisons only. The input sequence is scanned from left to right and the elements follow the arrows to go to the next stack or queue or output. We will concentrate on the average-case analyses of the above two main variants, although our technique in general apply to arbitrary acyclic networks of stacks and queues as studied in [2].


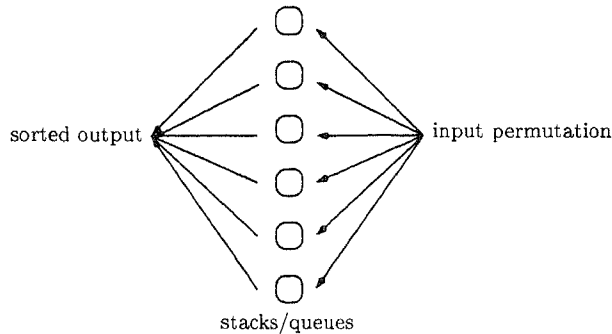Fig.1. Six stacks/queues arranged in sequential order.

Fig.2. Six stacks/queues arranged in parallel order.

3.1 Sorting with Sequential Stacks

The sequential stack sorting problem is in [1] exercise 5.2.4-20. We have k stacks numbered $S_0, \ldots, S_{k-1}$. The input is a permutation π of the elements $1, \ldots, n$. Initially we push the elements of π on S_0 at most one at a time in the order in which they appear in π. At every step we can pop a stack (the popped elements will move left in Fig.1) or push an incoming element on a stack. The question is how many stacks are needed for sorting π. It is known that $k = \log n$ stacks suffice, and $\frac{1}{2} \log n$ stacks are necessary in the worst-case[1,2]. Here we prove that the same lower bound also holds on the average with a very simple incompressibility argument.

Theorem 1. *On the average, at least $\frac{1}{2} \log n$ stacks are needed for sequential stack sort.*

Proof. Fix a random permutation π such that

$$C(\pi|n, P) \geq \log n! - \log n = n \log n - O(\log n),$$

where P is an encoding program to be specified in the following.

Assume that k stacks are sufficient to sort π. We now encode such a sorting process. For every stack, exactly n elements pass through it. Hence we need perform precisely n pushes and n pops on every stack. Encode a push as 0 and a pop as 1. It is easy to prove that different permutations must have different push/pop sequences on at least one stack. Thus with $2kn$ bits, we can completely specify the input permutation π. Then, as before,

$$2kn \geq \log n! - \log n = n \log n - O(\log n).$$

Hence, approximately $k \geq \frac{1}{2} \log n$ for the random permutation π.

Since most (a $(1 - 1/n)$th fraction) permutations are incompressible, we calculate the average-case lower bound as:

$$\frac{1}{2} \log n \cdot \frac{n-1}{n} + 1 \cdot \frac{1}{n} \approx \frac{1}{2} \log n. \qquad \square$$

3.2 Sorting with Parallel Stacks

Clearly, the input sequence $2, 3, 4, \ldots, n, 1$ requires $n-1$ parallel stacks to sort. Hence the worst-case complexity of sorting with parallel stacks, as shown in Fig.2, is $n - 1$. However, most sequences do not need this many stacks to sort in parallel arrangement. The next two theorems show that, on the average, $\Theta(\sqrt{n})$ stacks are both necessary and sufficient. Observe that the result is actually implied by the connection between sorting with parallel stacks and *longest increasing subsequences* given in [2] and the bounds on the length of longest increasing subsequences of random permutations given in [7–9]. However, the proofs in [7–9] use deep results from probability theory (such as Kingman's ergodic theorem) and are quite sophisticated. Here we give simple proofs using incompressibility arguments.

Theorem 2. *On the average, the number of parallel stacks needed to sort n elements is* $O(\sqrt{n})$.

Proof. Consider an incompressible permutation π satisfying

$$C(\pi|n) \geq \log n! - \log n. \tag{1}$$

We use the following trivial algorithm (which is described in [2]) to sort π with stacks in the parallel arrangement as shown in Fig.2. Assume that the stacks are named $S_0, S_1, \ldots$ and the input sequence is denoted as $x_1, \ldots, x_n$.

Algorithm 1. Parallel-Stack-Sort

1. For $i = 1$ to n do
   Scan the stacks from left to right, and push x_i on the first stack S_j whose top element is larger than x_i. If such a stack doesn't exist, put x_i on the first empty stack.
2. Pop the stacks in the ascending order of their top elements.

We claim that algorithm Parallel-Stack-Sort uses $O(\sqrt{n})$ stacks on the permutation π. First, we observe that if the algorithm uses m stacks on π then we can identify an increasing subsequence of π of length m as in [2]. This can be done by a trivial backtracing starting from the top element of the last stack. Then we argue that π cannot have an increasing subsequence of length longer than $e\sqrt{n}$, where e is a natural constant, since it is compressible by at most $\log n$ bits.

Suppose that σ is a longest increasing subsequence of π and $m = |\sigma|$ is the length of σ. Then we can encode π by specifying:

1) a description of this encoding scheme in $O(1)$ bits;
2) the number m in $\log m$ bits;
3) the combination σ in $\log \binom{n}{m}$ bits;
4) the locations of the elements of σ in π in at most $\log \binom{n}{m}$ bits; and
5) the remaining π with the elements of σ deleted in $\log(n - m)!$ bits.

This takes a total of

$$\log(n - m)! + 2\log \frac{n!}{m!(n - m)!} + \log m + O(1) + 2\log \log m$$

bits, where the last item $\log \log m$ is needed for self-delimiting encoding of m. Using Stirling approximation and the fact that $\sqrt{n} \leq m = o(n)$, the above expression is upper bounded by:

$$\log n! + \log \frac{(n/e)^n}{(m/e)^{2m}((n - m)/e)^{n-m}} + O(\log n)$$

$$\approx \log n! + m \log \frac{n}{m^2} + (n - m) \log \frac{n}{n - m} + m \log e + O(\log n)$$

$$\approx \log n! + m \log \frac{n}{m^2} + 2m \log e + O(\log n)$$

This description length must exceed the complexity of the permutation which is lower-bounded in 1). This requires that (approximately) $m \leq e\sqrt{n} = O(\sqrt{n})$.

This yields an average complexity of Parallel-Stack-Sort of:

$$O(\sqrt{n}) \cdot \frac{n-1}{n} + n \cdot \frac{1}{n} = O(\sqrt{n}). \qquad \square$$

Theorem 3. *On the average, the number of parallel stacks required to sort a permutation is $\Omega(\sqrt{n})$.*

Proof. Let A be a sorting algorithm using parallel stacks. Fix a random permutation π with $C(\pi|n, P) \geq \log n! - \log n$, where P is the program to do the encoding discussed in the following. Suppose that A uses T parallel stacks to sort π. This sorting process involves a sequence of moves, and we can encode this sequence of moves by a sequence of the following terms:

- push to stack i,
- pop stack j,

where the element to be pushed is the next unprocessed element from the input sequence and the popped element is written as the next output element. Each of these terms requires $\log T$ bits. In total, we use precisely $2n$ terms since every element has to be pushed once and popped once. Such a sequence is unique for every permutation.

Thus we have a description of an input sequence in $2n \log T$ bits, which must exceed $C(\pi|n, P) \geq n \log n - O(\log n)$. It follows that $T \geq \sqrt{n} = \Omega(\sqrt{n})$.

This yields the average-case complexity of A:

$$\Omega(\sqrt{n}) \cdot \frac{n-1}{n} + 1 \cdot \frac{1}{n} = \Omega(\sqrt{n}). \qquad \square$$

3.3 Sorting with Parallel Queues

It is easy to see that sorting cannot be done with a sequence of queues. So we consider the complexity of sorting with parallel queues. It turns out that all the result in the previous subsection also hold for queues.

As noticed in [2], the worst-case complexity of sorting with parallel queues is n since the input sequence $n, n-1, \ldots, 1$ requires n queues to sort. We show in the next two theorems that, on the average, $\Theta(\sqrt{n})$ queues are both necessary and sufficient. Again, the result is implied by the connection between sorting with parallel queues and longest *decreasing* subsequences given in [2] and the bounds in [7–9] (with sophisticated proofs). Our proofs are almost trivial given the proofs in the previous subsection.

Theorem 4. *On the average, the number of parallel queues needed to sort n elements is upper bounded by $O(\sqrt{n})$.*

Proof. The proof is very similar to the proof of Theorem 2. We use a slightly modified greedy algorithm as described in [2]:

Algorithm 2. Parallel-Queue-Sort

1. For $i = 1$ to n do
   Scan the queues from left to right, and append x_i to the first queue whose rear element is smaller than x_i. If such a queue doesn't exist, put x_i on the first empty queue.
2. Delete the front elements of the queues in the ascending order.

Again, we claim that algorithm Parallel-Queue-Sort uses $O(\sqrt{n})$ queues on any permutation π that cannot be compressed by more than $\log n$ bits. We first observe that if the algorithm uses m queues on π then a decreasing subsequence of π of length m can be identified, and we then argue that π cannot have a decreasing subsequence of length longer than $e\sqrt{n}$, in a way analogous to the argument in the proof of Theorem 2. $\square$

Theorem 5. *On the average, the number of parallel queues required to sort a permutation is* $\Omega(\sqrt{n})$.

Proof. The proof is the same as the one for Theorem 3 except that we should replace "push" with "enqueue" and "pop" with "dequeue". □

4   Open Questions

The incompressibility method is a good tool to analyzing the average-case complexity of sorting algorithms. Simplicity has been our goal. Examples of such average-case analyses of some other algorithms are given in [5]. This methodology and applications can be easily taught to undergraduate students.

The average-case performance of Shellsort[3] has been one of the most fundamental and interesting open problems in the area of algorithm analysis. The simple average-case analysis of Bubble Sort, stack-sort and queue-sort are further examples to demonstrate the generality and simplicity of our technique in analyzing sorting algorithms in general. We believe the method can be widely applied for analyzing the average case complexity of many other problems. Some specific open questions are:

1) Tighten the average-case lower bound for Shellsort. The bound in [3] is not tight for $p = 2$ passes.

2) For sorting with sequential stacks, can we close the gap between $\log n$ upper bound and the $\frac{1}{2} \log n$ lower bound?

References

[1] Knuth D E. The Art of Computer Programming. Vol.3: Sorting and Searching, Addison-Wesley, 1973 (1st Edition), 1998 (2nd Edition).
[2] Tarjan R E. Sorting using networks of queues and stacks. *Journal of the ACM,* 1972, 19: 341–346.
[3] Jiang T, Li M, Vitányi P. A lower bound on the average-case complexity of Shellsort. *J. Assoc. Comp. Mach.,* to appear. Also in *ICALP'99,* July 11–15, 1999, Prague.
[4] Li M, Vitányi P M B. An Introduction to Kolmogorov Complexity and Its Applications. Springer-Verlag, New York, 2nd Edition, 1997.
[5] Buhrman H, Jiang T, Li M, Vitányi P. New applications of the incompressibility method. In *ICALP'99,* July 11–15, 1999, Prague. Also in *Theoretical Computer Science,* 1999, 235.
[6] Kolmogorov A N. Three approaches to the quantitative definition of information. *Problems Inform. Transmission,* 1965, 1(1): 1–7.
[7] Kerov S V, Versik A M. Asymptotics of the Plancherel measure on symmetric group and the limiting form of the Young tableaux. *Soviet Math. Dokl.,* 1977, 18: 527–531.
[8] Kingman J F C. The ergodic theory of subadditive stochastic processes. *Ann. Probab.,* 1973, 1: 883–909.
[9] Logan B F, Shepp L A. A variational problem for random Young tableaux. *Advances in Math.,* 1977, 26: 206–222.
[10] Shell D L. A high-speed sorting procedure. *Commun. ACM,* 1959, 2(7): 30–32.

JIANG Tao received the B.S. degree in computer science and technology from the University of Science and Technology of China in 1984 and the Ph.D. degree in computer science from the University of Minnesota in 1988. From Jan. 1989 to Sept. 1999, he was a faculty member at Department of Computing and Software, McMaster University, Hamilton, Ontario, Canada. During 1995–1996, he took a research leave at University of Washington, USA, and at Gunma University, Japan. He joined University of California-Riverside as a professor of computer science in Sept. 1999, while taking a leave from McMaster University.

His research interests include computational molecular biology, design and analysis of algorithms, computational complexity, and information retrieval. He has published actively in many theoretical computer science journals and served on program committees for many international conferences. He is presently serving on the editorial board of International Journal of Foundations of Computer Science (IJFCS).

LI Ming is a professor of computer science at the University of Waterloo, and a guest professor at Peking University and University of Science and Technology of China. He received his Ph.D. degree from Cornell University in 1985. He is a recipient of Canada's prestigious E.W.R. Steacie Fellowship Award in 1996, and the 1997 Award of Merit from the Federation of Chinese Canadian Professionals. He is a coauthor of the book "An Introduction to Kolmogorov Complexity and Its Applications" (Springer-Verlag, 1993, 2nd Edition, 1997), and "A Course on Java Programming Language" (in Chinese, Science Press, 1997). He currently serves on the editorial boards of Journal of Computer and System Sciences, Journal of Computer Science and Technology, Information and Computation, and Journal of Combinatorial Optimization, International Journal of Foundation of Computer Science. His main research interests is bioinformatics. His recent book together with Paul Vitányi "Description Complexity" (Chinese, Science Press, 1999) won first prize for China's National Science and Technology Book Award.

Paul M.B. Vitányi obtained his Ph.D. degree from the Free University of Amsterdam (the Netherlands) in 1978. He holds positions at the CWI (Center for Mathematics and Computer Science) National Research Institute in Amsterdam where he heads the Algorithms and Complexity Research Group, and at the University of Amsterdam where he heads a similar group and is a professor of computer science. He has worked in the areas of cellular automata, computational complexity, distributed and parallel computing, machine learning and prediction, physics of computation, quantum computing, and description complexity. He serves on the editorial boards of Distributed Computing, Information Processing Letters, Theory of Computing Systems, Parallel Processing Letters, and has co-authored about 150 research papers and books, among others with Li Ming "An Introduction to Kolmogorov Complexity and Its Applications", Springer-Verlag, New York, 1993 (2nd Edition 1997), parts of which have been translated into Russian, Japanese and Chinese. (Web page: http://www.cwi.nl/$\sim$paulv/)