Average-Case Complexity of Shellsort

(*Preliminary Version*)

Tao Jiang[1], Ming Li[2], and Paul Vitányi[3]

[1] Dept of Computing and Software, McMaster University,
Hamilton, Ontario L8S 4K1, Canada, jiang@cas.mcmaster.ca
Supported in part by NSERC and CITO grants
[2] Dept of Computer Science, University of Waterloo,
Waterloo, Ontario N2L 3G1, Canada, mli@math.uwaterloo.ca
Supported in part by NSERC and CITO grants and Steacie Fellowship
[3] CWI, Kruislaan 413, 1098 SJ Amsterdam,
The Netherlands, paulv@cwi.nl
Supported in part via NeuroCOLT II ESPRIT Working Group

Abstract. We prove a general lower bound on the average-case complexity of Shellsort: the average number of data-movements (and comparisons) made by a p-pass Shellsort for any incremental sequence is $\Omega(pn^{1+1/p})$ for every p. The proof method is an incompressibility argument based on Kolmogorov complexity. Using similar techniques, the average-case complexity of several other sorting algorithms is analyzed.

1 Introduction

The question of a nontrivial general lower bound (or upper bound) on the average complexity of Shellsort (due to D.L. Shell [15]) has been open for about four decades [7, 14]. We present such a lower bound for p-pass Shellsort for every p.

Shellsort sorts a list of n elements in p passes using a sequence of increments $h_1, \ldots, h_p$. In the kth pass the main list is divided in h_k separate sublists of length $\lceil n/h_k \rceil$, where the ith sublist consists of the elements at positions j, where $j \bmod h_k = i - 1$, of the main list ($i = 1, \ldots, h_k$). Every sublist is sorted using a straightforward insertion sort. The efficiency of the method is governed by the number of passes p and the selected increment sequence $h_1, \ldots, h_p$ with $h_p = 1$ to ensure sortedness of the final list. The original $\log n$-pass [1] increment sequence $\lfloor n/2 \rfloor, \lfloor n/4 \rfloor, \ldots, 1$ of Shell [15] uses worst case $\Theta(n^2)$ time, but Papernov and Stasevitch [10] showed that another related sequence uses $O(n^{3/2})$ and Pratt [12] extended this to a class of all nearly geometric increment sequences and proved this bound was tight. The currently best asymptotic method was found by Pratt [12]. It uses all $\log^2 n$ increments of the form $2^i 3^j < \lfloor n/2 \rfloor$ to obtain time $O(n \log^2 n)$ in the worst case. Moreover, since every pass takes at least n steps, the average complexity using Pratt's increment sequence is $\Theta(n \log^2 n)$. Incerpi and Sedgewick [2] constructed a family of increment sequences for which

[1] "log" denotes the binary logarithm and "ln" denotes the natural logarithm.

Shellsort runs in $O(n^{1+\epsilon/\sqrt{\log n}})$ time using $(8/\epsilon^2)\log n$ passes, for every $\epsilon > 0$. B. Chazelle (attribution in [13]) obtained the same result by generalizing Pratt's method: instead of using 2 and 3 to construct the increment sequence use a and $(a+1)$ for fixed a which yields a worst-case running time of $n\log^2 n(a^2/\ln^2 a)$ which is $O(n^{1+\epsilon/\sqrt{\log n}})$ for $\ln^2 a = O(\log n)$. Plaxton, Poonen and Suel [11] proved an $\Omega(n^{1+\epsilon/\sqrt{p}})$ lower bound for p passes of Shellsort using any increment sequence, for some $\epsilon > 0$; taking $p = \Omega(\log n)$ shows that the Incerpi-Sedgewick / Chazelle bounds are optimal for small p and taking p slightly larger shows a $\Theta(n\log^2 n/(\log\log n)^2)$ lower bound on the worst case complexity of Shellsort. Since every pass takes at least n steps this shows an $\Omega(n\log^2 n/(\log\log n)^2)$ lower bound on the worst-case of every Shellsort increment sequence. For the *average-case* running time Knuth [7] showed $\Theta(n^{5/3})$ for the best choice of increments in $p = 2$ passes; Yao [17] analyzed the average case for $p = 3$ but did not obtain a simple analytic form; Yao's analysis was improved by Janson and Knuth [3] who showed $O(n^{23/15})$ average-case running time for a particular choice of increments in $p = 3$ passes. Apart from this no nontrivial results are known for the average case; see [7, 13, 14].

Results: We show a general $\Omega(pn^{1+1/p})$ lower bound on the average-case running time of p-pass Shellsort under uniform distribution of input permutations for every p. [2] This is the first advance on the problem of determining general nontrivial bounds on the *average-case* running time of Shellsort [12, 7, 17, 2, 11, 13, 14]. Using the same simple method, we also obtain 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 [7] and Tarjan [16] for the worst case.

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 [8] Chapter 6, and recent work is [1]. The most spectacular successes of the method occur in the computational complexity analysis of algorithms.

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. A string x is incompressible relative to y if $C(x|y)$ is large in the appropriate sense. For details see [8]. Here we use that, both absolutely and relative to any fixed string y, there are incompressible strings of

[2] The trivial lower bound is $\Omega(pn)$ comparisons since every element needs to be compared at least once in every pass.

every length, and that *most* strings are nearly incompressible, by *any* standard. [3] 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. 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 set, then for every y every element $x \in A$ has complexity $C(x|A,y) \leq \log |A| + O(1)$.*

The first lemma is proved by simple counting. The second lemma 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 Shellsort

A Shellsort computation consists of a sequence comparison and inversion (swapping) operations. In this analysis of the average-case lower bound we count just the total number of data movements (here inversions) executed. The same bound holds for number of comparisons automatically. The average is taken over the uniform distribution of all lists of n items.

The proof is based on the following intuitive idea: There are $n!$ different permutations. Given the sorting process (the insertion paths in the right order) one can recover the correct permutation from the sorted list. Hence one requires $n!$ pairwise different sorting processes. This gives a lower bound on the minimum of the maximal length of a process. We formulate the proof in the crisp format of incompressibility.

Theorem 1. *The average number of inversions in p-pass Shellsort on lists of n keys is at least $\Omega\left(pn^{1+1/p}\right)$ for every increment sequence.*

Proof. Let the list to be sorted consist of a permutation π of the elements $1,\ldots,n$. Consider a $(h_1,\ldots,h_p)$ Shellsort algorithm A where h_k is the increment in the kth pass and $h_p = 1$. For any $1 \leq i \leq n$ and $1 \leq k \leq p$, let $m_{i,k}$ be the number of elements in the h_k-*chain* containing element i that are to the left of i at the beginning of pass k and are larger than i. Observe that $\sum_{i=1}^{n} m_{i,k}$ is the number of inversions in the initial permutation of pass k, and that the

[3] By a simple counting argument one can show that whereas some strings can be enormously compressed, like strings of the form $11\ldots1$, 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 $C(x|y) \geq n$.

insertion sort in pass k requires precisely $\sum_{i=1}^{n}(m_{i,k}+1)$ comparisons. Let M denote the total number of inversions:

$$M := \sum_{k=1}^{p}\sum_{i=1}^{n} m_{i,k}. \tag{1}$$

Claim. Given all the $m_{i,k}$'s in an appropriate fixed order, we can reconstruct the original permutation π.

Proof. The $m_{i,p}$'s trivially specify the initial permutation of pass p. In general, given the $m_{i,k}$'s and the final permutation of pass k, we can easily reconstruct the initial permutation of pass k. □

Let M as in (1) be a fixed number. Let permutation π be an incompressible permutation having Kolmogorov complexity

$$C(\pi|n,A,P) \geq \log n! - \log n. \tag{2}$$

where P is the encoding program in the following discussion. The description in Claim 2 is effective and therefore its minimum length must exceed the complexity of π:

$$C(m_{1,1},\ldots,m_{n,p}|n,A,P) \geq C(\pi|n,A,P). \tag{3}$$

Any M as defined by (1) such that every division of M in $m_{i,k}$'s contradicts (3) would be a lower bound on the number of inversions performed. There are

$$D(M) := \sum_{i=1}^{np-1} \binom{M}{np-i} = \binom{M+np-1}{np-1} \tag{4}$$

possible divisions of M into np nonnegative integral summands $m_{i,k}$'s. Every division can be indicated by its index j in an enumeration of these divisions. Therefore, a self-delimiting description of M followed by a description of j effectively describes the $m_{i,k}$'s. The length of this description must by definition exceed its Kolmogorov complexity. That is,

$$\log D(M) + \log M + 2\log\log M \geq C(m_{1,1},\ldots,m_{n,p}|n,A,P) + O(1).$$

We know that $M \leq pn^2$ since every $m_{i,k} \leq n$. We can assume $p < n$.[4] Together with (2) and (3), we have

$$\log D(M) \geq \log n! - 4\log n + O(1). \tag{5}$$

By (4) $\log D(M)$ is bounded above by [5]

$$\log \binom{M+np-1}{np-1} = (np-1)\log\frac{M+np-1}{np-1} + M\log\frac{M+np-1}{M}$$

[4] Otherwise we require at least n^2 comparisons.

[5] Use the following formula ([8], p. 10),

$$\log\binom{a}{b} = b\log\frac{a}{b} + (a-b)\log\frac{a}{a-b} + \frac{1}{2}\log\frac{a}{b(a-b)} + O(1).$$

$$+\frac{1}{2}\log\frac{M+np-1}{(np-1)M}+O(1).$$

By (5) we have $M \to \infty$ for $n \to \infty$. Therefore, the second term in the right-hand side equals

$$\log\left(1+\frac{np-1}{M}\right)^{M} \to \log e^{np-1}$$

for $n \to \infty$. Since $0 < p < n$ and $n \le M \le n^3$,

$$\frac{1}{2(np-1)}\log\frac{M+np-1}{(np-1)M} \to 0$$

for $n \to \infty$. Therefore, the total right-hand side goes to

$$(np-1)\left(\log\left(\frac{M}{np-1}+1\right)+\log e\right)$$

for $n \to \infty$. Together with (5) this yields

$$M = \Omega(pn^{1+1/p}).$$

Therefore, the running time of the algorithm is as stated in the theorem for every permutation π satisfying (2). By lemma 1 at least a $(1-1/n)$-fraction of all permutations π require that high complexity. Therefore, the following is a lower bound on the expected number of inversions of the sorting procedure:

$$(1-\frac{1}{n})\Omega(pn^{1+1/p})+\frac{1}{n}\Omega(0)=\Omega(pn^{1+1/p})$$

This gives us the theorem. □

Compare our lower bound on the average-case with the Plaxton-Poonen-Suel $\Omega(n^{1+\epsilon/\sqrt{p}})$ worst case lower bound [11]. Some special cases of the lower bound on the average-case complexity are:

1. When $p = 1$, this gives asymptotically tight bound for the average number of inversions for Insertion Sort.
2. When $p = 2$, Shellsort requires $\Omega(n^{3/2})$ inversions (the tight bound is known to be $\Theta(n^{5/3})$ [7]);
3. When $p = 3$, Shellsort requires $\Omega(n^{4/3})$ inversions (the best known upper bound is $O(n^{23/15})$ in [3]);
4. When $p = \log n/\log\log n$, Shellsort requires $\Omega(n\log^2 n/\log\log n)$ inversions;
5. When $p = \log n$, Shellsort requires $\Omega(n\log n)$ inversions. When we consider comparisons, this is of course the lower bound of average number of comparisons for every sorting algorithm.
6. In general, when $p = p(n) > \log n$, Shellsort requires $\Omega(n \cdot p(n))$ inversions (it requires that many comparisons anyway since every pass trivially makes n comparisons).

In [14] it is mentioned that the existence of an increment sequence yielding an average $O(n\log n)$ Shellsort has been open for 30 years. The above lower bound on the average shows that the number p of passes of such an increment sequence (if it exists) is precisely $p = \Theta(\log n)$; all the other possibilities are ruled out.

3 Sorting with Queues and Stacks

Knuth [7] and Tarjan [16] have studied the problem of sorting using a network of queues or stacks. In particular, the main variants of the problem are: assuming the stacks or queues are arranged sequentially or in parallel, how many stacks or queues are needed to sort n numbers. Here, the input sequence is scanned from left to right. 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 [16].

3.1 Sorting with Sequential Stacks

The sequential stack sorting problem is in [7] exercise 5.2.4-20. We have k stacks numbered $S_0, \ldots, S_{k-1}$ arranged sequentially from right to left. The input is a permutation π of the elements $1, \ldots, n$. Initially we feed the elements of π to 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 to the left) or push an incoming element on a stack. The question is how many stack 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 [7, 16]. Here we prove that the same lower bound also holds on the average with a very simple incompressibility argument.

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

Proof. Fix an incompressible permutation π such that

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

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

Assume that k stacks is 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. [6] 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 incompressible permutations π.

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

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

□

[6] In fact since each stack corresponds to precisely n pushes and n pops where the pushes and pops form a "balanced" string, the Kolmogorov complexity of such a sequence is at most $g(n) := 2n - \frac{3}{2}\log n + O(1)$ bits. So $2kg(n)$ bits would suffice to specifiy the input permutation. But this does not yield a nontrivial improvement.

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 is $n-1$. However, most sequences do not need these many stacks to sort in the 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 [16] and the bounds on the length of longest increasing subsequences of random permutations given in, [5, 9, 4]. However, the proofs in [5, 9, 4] 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 3. *On the average, the number of parallel stacks needed to sort n elements is $O(\sqrt{n})$.*

Proof. Consider an incompressible permutation π satisfying

$$C(\pi|n) \geq \log n! - \log n. \tag{6}$$

We use the following trivial algorithm (which is described in [16]) to sort π with stacks in the parallel arrangement . Assume that the stacks are named $S_0, S_1, \ldots$ and the input sequence is denoted as $x_1, \ldots, x_n$.

Algorithm Parallel-Stack-Sort

1. For $i = 1$ to n do
   Scan the stacks from left to right, and push x_i on the 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 [16]. 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 the 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. Using Stirling approximation and the fact that $\sqrt{n} \leq m = o(n)$, we the above expression is upper bounded by:

$$\begin{aligned}
&\log n! + \log\frac{(n/e)^n}{(m/e)^{2m}((n-m)/e)^{n-m}} + O(\log n) \\
\approx\ &\log n! + m\log\frac{n}{m^2} + (n-m)\log\frac{n}{n-m} + m\log e + O(\log n) \\
\approx\ &\log n! + m\log\frac{n}{m^2} + 2m\log e + O(\log n)
\end{aligned}$$

This description length must exceed the complexity of the permutation which is lower-bounded in (6). This requires that (approximately) $m \leq e\sqrt{n} = O(\sqrt{n})$.

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

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

□

Theorem 4. *On the average, the number of parallel stacks required to sort a permutation is* $\Omega(\sqrt{n})$.

Proof. Let A be any sorting algorithm using parallel stacks. Fix an incompressible 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 as a sequence of the following items: "push to stack i" and "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 term requires $\log T$ bits. In total, we use $2n$ terms precisely 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 with length $2n\log T$ 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:

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

□

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 [16], 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 [16] and the bounds in [5, 9, 4] (with sophisticated proofs). Our proofs are almost trivial given the proofs in the previous subsection.

Theorem 5. *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 3. We use a slightly modified greedy algorithm as described in [16]:

Algorithm Parallel-Queue-Sort

1. For $i = 1$ to n do
   Scan the queues from left to right, and append x_i on the 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 can 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 3. □

Theorem 6. *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 4 except that we should replace "push" with "enqueue" and "pop" with "dequeue". □

4 Conclusion

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 [1]. This methodology and applications can be easily taught to undergraduate students.

The average-case performance of Shellsort has been one of the most fundamental and interesting open problems in the area of algorithm analysis. The simple average-case analysis of Insertion Sort (1-pass Shellsort), stack-sort and queue-sort are further examples to demonstrate the generality and simplicity of our technique in analyzing sorting algorithms in general. Some open questions are:

1. Tighten the average-case lower bound for Shellsort. Our bound 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?

5 Acknowledgements

We thank Don Knuth, Ian Munro, and Vaughan Pratt for discussions and references on Shellsort.

References

1. H. Buhrman, T. Jiang, M. Li, and P. Vitányi, New applications of the incompressibility method, in *the Proceedings of ICALP'99*.
2. J. Incerpi and R. Sedgewick, Improved upper bounds on Shellsort, *Journal of Computer and System Sciences*, 31(1985), 210–224.
3. S. Janson and D.E. Knuth, Shellsort with three increments, *Random Struct. Alg.*, 10(1997), 125-142.
4. S.V. Kerov and A.M. Versik, Asymptotics of the Plancherel measure on symmetric group and the limiting form of the Young tableaux, *Soviet Math. Dokl.* 18 (1977), 527-531.
5. J.F.C. Kingman, The ergodic theory of subadditive stochastic processes, *Ann. Probab.* 1 (1973), 883-909.
6. A.N. Kolmogorov, Three approaches to the quantitative definition of information. *Problems Inform. Transmission*, 1:1(1965), 1–7.
7. D.E. Knuth, *The Art of Computer Programming, Vol.3: Sorting and Searching*, Addison-Wesley, 1973 (1st Edition), 1998 (2nd Edition).
8. M. Li and P.M.B. Vitányi, *An Introduction to Kolmogorov Complexity and its Applications*, Springer-Verlag, New York, 2nd Edition, 1997.
9. B.F. Logan and L.A. Shepp, A variational problem for random Young tableaux, *Advances in Math.* 26 (1977), 206-222.
10. A. Papernov and G. Stasevich, A method for information sorting in computer memories, *Problems Inform. Transmission*, 1:3(1965), 63–75.
11. C.G. Plaxton, B. Poonen and T. Suel, Improved lower bounds for Shellsort, *Proc. 33rd IEEE Symp. Foundat. Comput. Sci.*, pp. 226–235, 1992.
12. V.R. Pratt, *Shellsort and Sorting Networks*, Ph.D. Thesis, Stanford Univ., 1972.
13. R. Sedgewick, Analysis of Shellsort and related algorithms, presented at the *Fourth Annual European Symposium on Algorithms*, Barcelona, September, 1996.
14. R. Sedgewick, Open problems in the analysis of sorting and searching algorithms, Presented at *Workshop on Prob. Analysis of Algorithms*, Princeton, 1997.
15. D.L. Shell, A high-speed sorting procedure, *Commun. ACM*, 2:7(1959), 30–32.
16. R.E. Tarjan, Sorting using networks of queues and stacks, *Journal of the ACM*, 19(1972), 341–346.
17. A.C.C. Yao, An analysis of $(h, k, 1)$-Shellsort, *J. of Algorithms*, 1(1980), 14–50.