

Analysis of Shellsort and Related Algorithms

Robert Sedgewick*
Princeton University

This is an abstract of a survey talk on the theoretical and empirical studies that have been done over the past four decades on the Shellsort algorithm and its variants. The discussion includes: upper bounds, including linkages to number-theoretic properties of the algorithm; lower bounds on Shellsort and Shellsort-based networks; average-case results; proposed probabilistic sorting networks based on the algorithm; and a list of open problems.

1 Shellsort

The basic Shellsort algorithm is among the earliest sorting methods to be discovered (by D. L. Shell in 1959 [36]) and is among the easiest to implement, as exhibited by the following C code for sorting an array $a[1], \dots, a[r]$:

```
shellsort(itemType a[], int l, int r)
{ int i, j, h; itemType v;
  int incs[16] = { 1391376, 463792, 198768, 86961, 33936,
                  13776, 4592, 1968, 861, 336,
                  112, 48, 21, 7, 3, 1 };
  for ( k = 0; k < 16; k++)
    for (h = incs[k], i = l+h; i <= r; i++)
      {
        v = a[i]; j = i;
        while (j > h && a[j-h] > v)
          { a[j] = a[j-h]; j -= h; }
        a[j] = v;
      }
}
```

The algorithm is based on *insertion sort*: proceed left to right through a file, inserting each element into position among the elements to its left (which are in sorted order) by moving the larger ones one position to the right. Shellsort is a sequence of interleaved insertion sorts based on an *increment sequence*: for each increment h , insertion-sort the h subfiles consisting of the i th, $(i + h)$ th,

*Supported by the National Science Foundation, Grant No. DCR-8605962.

$(i + 2h)$ th, \dots items, for $i = 0, 1, \dots, h - 1$. This process is called *h-sorting* the file. Basic information about the algorithm and many of its properties may be found in algorithms textbooks, for example [9], [17], and [30]. Some of this information is summarized in this paper.

The performance of Shellsort depends on the increment sequence. The specific values used in the code above are explained below. The essential point is that it is trivial to substitute another set of increments for those used in this code. We seek better increments—no matter how sophisticated the mathematical analysis or how extensive the empirical studies we use, they can be quickly verified and put to immediate practical use. Moreover, despite extensive research on learning properties of different families of increment sequences and on lower bounds on performance, we still cannot deny the possibility that some other set of increments might make this program run two or three times faster, which would make it as fast or faster than the best known sorting methods. This possibility makes Shellsort worthy of close study.

The operation of moving larger elements one position to the right is equivalent to doing a sequence of *compare-exchange* operations where we compare $a[j-h]$ with $a[j]$ and exchange them if necessary to put the smaller of the two items on the left. If we make this change, and remove the test $a[j-h] > v$ from the implementation above, we get a specification of a *sorting network*, where the compare-exchange operations and the order in which they are performed does not depend on the input data. Another reason to study Shellsort and its variants is the possibility of progress on the longstanding open problem of developing optimal sorting networks that are useful in practice.

2 Upper Bounds

A file that is 2-sorted and 3-sorted can be 1-sorted in one insertion-sort pass, using just N comparisons. Similarly, a file that is 4-sorted and 6-sorted can be 2-sorted in one insertion-sort pass, using just N comparisons; and a file that is 6-sorted and 9-sorted can be 3-sorted in one insertion-sort pass, using just N comparisons. This generalizes to give a simple increment sequence with provable performance within a factor of $\log N$ of optimal [28].

Theorem (Pratt, 1971) *The running time of Shellsort is $\Theta(N(\log N)^2)$ for the increments 1 2 3 4 6 9 8 12 18 27 16 24 36 54 81 \dots*

The increment sequence is formed by starting at 1, then proceeding through the sequence appending 2 and 3 times each increment encountered. This sequence has the best known asymptotic performance of any increment sequence for Shellsort, and it also defines a sorting network with $O(N(\log N)^2)$ compare-exchange operations. However, it is generally not competitive in practice because of the constant factors involved. There are simply too many increments.

Instead, increment sequences that are approximately geometric, and therefore use $O(\log N)$ increments to sort N items, are normally used in Shellsort. The number of passes is a fixed cost, but the actual time taken depends on the “sortedness” of the file, one pass to the next. For many years, the focus of research on Shellsort was on proving upper bounds on the running time for sequences that use $O(\log N)$ passes. The question of whether such a sequence exists for which the total running time is $O(N \log N)$ was not resolved until the early 1990s (see the next section).

Shell proposed using powers of two for the increment sequence. It is easy to see that this choice can lead to quadratic running times, because elements in odd positions are not compared to elements in even positions until the 1-sort at the end. Early practitioners noticed this problem and the difficulty of precisely characterizing the running time of the method. The early studies of Shellsort are intriguing [11][19], but we begin with theoretical results that were developed quite some time after Shellsort was invented [25][28]:

Theorem (Papernov-Stasevich, 1965; Pratt, 1971) *The running time of Shellsort is $\Theta(N^{3/2})$ for the increments 1 3 7 15 31 63 127 255 511 ...*

The upper bound involves a tradeoff between two easy bounds: the first is for large increments (small subfiles), and the second is for small increments (large subfiles).

Lemma *The number of comparisons required to h_j -sort a file of size N is $O(N^2/h_j)$.*

Lemma *The number of comparisons required to h_j -sort a file that is h_{j+1} -sorted and h_{j+2} -sorted is $O(Nh_{j+1}h_{j+2}/h_j)$, if h_{j+1} and h_{j+2} are relatively prime.*

See, for example, [17] or [30] for proofs of these lemmas and other details. These bounds are both $O(N^{3/2})$ when the increments are $O(N^{1/2})$, and it follows that the total cost is $O(N^{3/2})$ because of the geometric growth of the increments.

Pratt proved that the upper bound is tight by giving a construction of inputs for which $O(N^{3/2})$ comparisons are required.

Pratt also showed that the bound is tight under rather general assumptions: if increments are within a constant additive factor of a geometric sequence (and a few other technical conditions hold), then the running time of Shellsort is $O(N^{3/2})$. It turns out that the constant factors involved give the best total running times when the growth ratio of the geometric sequence is roughly between 2 and 4. Knuth recommended using the sequence 1 4 13 40 121 364 1093 3280 9841 ..., because it is easy to compute, uses relatively few (about $\log_3 N$) increments, and does well in empirical studies. Later empirical studies have shown a ratio of about 11/5 to be the best for geometric increment sequences [38][9].

Pratt's proof seemed to apply to all the increment sequences of practical interest, and Knuth's sequence works well in practice for moderate-sized sorting problems (and is still widely used today), so Shellsort received little attention for another decade. However, better increment sequences exist and are useful [31].

Theorem (Sedgewick, 1982) *The running time of Shellsort is $O(N^{4/3})$ for the increments 1 8 23 77 281 1073 4193 16577 ... $4^{j+1} + 3 \cdot 2^j + 1$.*

The proof technique involves a relationship to the *Frobenius problem* from number theory [4][5][10][12][16][24][34]. The *Frobenius number* for a set of distinct integers is defined to be the number of positive integers that cannot be expressed as a linear combination of the integers in the set. This function arises in improving the bound in the second Lemma given above.

Once the $O(N^{3/2})$ barrier was broken, further improvements soon followed [14][35]:

Theorem (Incerpi-Sedgewick, 1985; Selmer, 1987) *There are increment sequences for which the running time of Shellsort is $O(N^{1+1/k})$, using $O(\log N)$ increments.*

Selmer's proof involves generalizing the Sedgewick construction using general bounds from the literature on the Frobenius problem for increment sequences formed by taking the product of k consecutive terms from the sequence a_1, a_2, \dots , with $a_i = 2^{i+5} - 7$ for $k = 2, 3$ and $a_i = 2^{i+5} - 45$ for $k = 4, 5, \dots, 9$. The Incerpi-Sedgewick construction also involves increments with large common multiples. In both cases, the increments (and the implied constant in the O -notation) are quite large and therefore not practical. A modification to include smaller increments gives practical sequences and better asymptotic results [14].

Theorem (Incerpi-Sedgewick, 1985) *For any $\epsilon > 0$ there is an increment sequence for which the running time of Shellsort is $O(N^{1+\epsilon/\log N})$, using $(8/\epsilon^2) \log N$ passes.*

A simple proof of the same asymptotic result, due to Chazelle, comes from generalizing Pratt's $O(N \log^2 N)$ method: instead of using 2 and 3 to build the increment sequence, use α and $(\alpha + 1)$ for fixed α [6]. The asymptotic worst-case running time works out to be

$$N(\log N)^2 \frac{\alpha^2}{(\lg \alpha)^2} \quad \text{which is} \quad O(N^{1+\epsilon/\log N}) \quad \text{for} \quad (\lg \alpha)^2 = O(\log N).$$

This method also has too few small increments to be competitive in practice, and it implies the use of a different increment sequence for each N .

The Incerpi-Sedgewick construction corrects these problems by building up a triangle of increments from a base sequence a_1, a_2, a_3, \dots , as follows:

a_1	$a_1 a_2$	$a_1 a_2 a_3$	$a_1 a_2 a_3 a_4$	$a_1 a_2 a_3 a_4 a_5$.
	$a_1 a_3$	$a_1 a_2 a_4$	$a_1 a_2 a_3 a_5$	$a_1 a_2 a_3 a_4 a_6$.
		$a_1 a_3 a_4$	$a_1 a_2 a_4 a_5$	$a_1 a_2 a_3 a_5 a_6$.
			$a_1 a_3 a_4 a_5$	$a_1 a_2 a_4 a_5 a_6$.
				$a_1 a_3 a_4 a_5 a_6$.
					.

The increment sequence used in the program above is derived from this table, using the base sequence 1, 3, 7, 16, 41, 101, 247. These values (including the use of $a_1 = 1$, which changes the construction slightly) was suggested by Knuth[18]. This increment sequence does quite well in empirical studies.

3 Lower Bounds

None of the increment sequences above achieve the goal of yielding an optimal sort. Weiss showed the Incerpi-Sedgewick upper bounds to be tight, indicating that new constructions would be required for better upper bounds [38][41]. The next advances were towards improved lower bounds. Weiss gave strong evidence that the Incerpi-Sedgewick sequences are the best possible for $O(\log N)$ increments, showing this fact to be implied by a certain conjecture about inversions related to Frobenius patterns [38][41]. This result is implied by general bounds for Shellsort, which were proven soon thereafter by Poonen [27].

Theorem (Poonen) *Shellsort requires at least $N^{1+c/\sqrt{M}}$ comparisons in the worst case to sort a file of size N in M passes, for some $c > 0$.*

A simpler proof of this same result is given by Plaxton and Suel [26]. Taking $M = \Omega(\log N)$ shows that the Incerpi-Sedgewick or Chazelle constructions are optimal for short increment sequences, and taking M to be slightly larger gives the lower bound $\Theta(N \log N)^2 / (\log \log N)^2$ for the worst-case complexity of Shellsort.

Ajtai, Komlós, and Szemerédi showed in 1983 that optimal sorting networks have $O(N \log N)$ compare-exchange operations [1][2], but the constants involved in the upper bounds are quite large, and no optimal networks are known that are competitive with the network associated with Pratt’s increment sequence for Shellsort given above or the classical Batcher’s network [3], which also has $O(N(\log N)^2)$ compare-exchanges. Poonen’s bound is relevant to this problem, but Cypher proved a stronger result for networks: Shellsort networks with decreasing increments must have $\Theta(N \log N)^2 / \log \log N$ compare-exchange operations [7].

These theoretical results not only tell us that the Shellsort is nonoptimal, but also that any Shellsort using $O(\log N)$ increments must, for example, use $\Omega(N(\log N)^k)$ comparisons for all k . There does remain a gap in the complexity

results: is the worst-case complexity of Shellsort $\Theta(N \log N)^2 / (\log \log N)^2$ or $\Theta(N \log N)^2 / \log \log N$ or $\Theta(N(\log N)^2)$ or something in between? Poonen conjectures that Pratt's networks are optimal; on the other hand, not all of the increments in the Pratt-like sequences seem necessary.

A significant gap also exists because of the nature of the functions involved and the effect of unspecified constants. Consider the following table:

N	$\lg N$	$(\lg N)^2$	$\frac{(\lg N)^2}{(\lg \lg N)^2}$	$\frac{1}{N^{2\sqrt{\lg N}}}$	$\frac{1}{N^{\sqrt{\lg N}}}$	$\frac{2}{N^{\sqrt{\lg N}}}$
10^3	10	100	9	3	9	80
10^6	20	400	21	5	22	487
10^9	30	900	56	9	80	6327

Small constants, particularly those in the exponents, are quite significant for file sizes of practical interest, and differences in asymptotic performance may not be discernible for any realistic file sizes. Furthermore, the constants involved in theoretical work such as the Ajtai, Komlós, and Szemerédi optimal sorting network or the Poonen and Cypher lower bounds are not small. Thus, despite the lower bounds, there may be nonoptimal Shellsorts that are more efficient than all other sorts (even optimal ones) for practical sorting problems.

Moreover, all the results are worst-case results, and proving the bounds tight has needed intricate constructions. Does Shellsort run significantly faster for random permutations than it does in the worst case for practical increment sequences?

4 Average Case

Two-pass Shellsort, using just the increments h and 1, is closely related to the classical ballot problem, and the average running time for random permutations can be explicitly determined [17][33]:

Theorem (Knuth) *Two-pass $(h, 1)$ Shellsort uses $2N^2/h + \sqrt{\pi N^3 h}$ comparisons to sort a random file of size N .*

Taking $h = O(N^{1/3})$ gives the result that Shellsort uses $O(N^{5/3})$ comparisons, on the average, to sort a random file of size N , asymptotically less than the quadratic worst-case time. This extends to give average-case results for the case where all the increments divide, but that case is not of practical interest.

For three increments, an intricate argument by Yao gives an exact expression for the average running time [43]:

Theorem (Yao) *Three-pass $(h, k, 1)$ Shellsort uses $2N^2/h + (\sqrt{\pi N^3 h}/8 - \sqrt{\pi N^3/8h})/k + \psi(h, k)N$ comparisons to sort a random file of size N .*

The definition for $\psi(h, k)$ in this expression is extremely complex and is omitted. Unfortunately, it does not give enough information to find the values of h and k that minimize the cost, and it is not known whether the best asymptotic result for three increments is lower than the $O(N^{5/3})$ worst case.

Beyond these results, the analysis of the average-case of Shellsort for any increment sequence of practical interest is completely open.

5 Variants of Shellsort

Dobosiewicz [8] was among the first to notice that using the Shellsort increments and basic structure, but substituting another operations for the full h -sort, may give an efficient way to sort random files. He proposed replacing the h -sort with what might be called an *h-bubble* pass: move from left to right through the file, compare-exchanging each element with the one h to its right.

Empirical Result (Dobosiewicz, 1980) *Replacing the h -sort in Shellsort by an h -bubble pass gives an algorithm that nearly always sorts when the increment sequence is geometric, with ratio less than 1.33 .*

The imprecise phrase “nearly always sorts” indicates a probabilistic sorting method. That is, the method might leave some items unsorted. Either we announce that we sort all but a few files, which occur with small probability, or we can run the algorithm with increments of 1 until the sort completes. No specific results on performance have been proven.

A more symmetric version of bubble sort is to alternate left-right and right-left passes through the file. We define a *h-shake* pass to be an h -bubble pass followed by a similar pass in the opposite direction, from right to left through the file, compare-exchanging each element with the one h to its left. This leads to a faster method [13][15]:

Empirical Result (Incerpi-Sedgewick, 1985) *Replacing the h -sort in Shellsort by an h -shake pass gives an algorithm that nearly always sorts when the increment sequence is geometric, with ratio less than 1.7 .*

Robbins [29] found this method to be among the fastest for a certain range of file sizes on a vector supercomputer; Weiss [38] also validated this result with extensive simulations.

Both of these methods also define probabilistic sorting networks, but they (and other networks directly based on Shellsort) suffer from the basic problem that the depth for each pass is linear. For example, the last pass involves compare-exchanging the first and second elements, then the second and third, then the third and fourth, and so forth, which requires linear time, even on a parallel machine. This property severely limits the utility of the networks.

Remarkably, there is a simple alternative which seems to lead to low-depth networks: use an *h-brick* pass, where items in positions $i, i + 2h, i + 4h, i + 6h,$

... are compare-exchanged with items in positions $i + h, i + 3h, i + 5h, i + 7h, \dots$ respectively, for i from 0 to $h - 1$, then items in positions $i + h, i + 3h, i + 5h, i + 7h, \dots$ are compare-exchanged with those in positions $i + 2h, i + 4h, i + 6h, i + 8h, \dots$ respectively. This can be done in two parallel steps, and seems sufficient to build useful networks [23][32]:

Empirical Result (Lemke, Sedgewick, 1994) *Replacing the h -sort in Shellsort by an h -brick pass gives an algorithm that nearly always sorts when the increment sequence is geometric, with ratio less than 1.22 .*

This seems to lead to simple probabilistic sorting networks with depth about $4 \lg N$. These networks are far simpler than other probabilistic networks than have been defined [21][22], but analytic verification that the networks sort with high probability has not been completed. Lemke developed a probabilistic model that explains some of the behavior, but the question of the existence of a probabilistic depth $\log N$ network based on a variant of Shellsort remains open.

6 Open Problems

A number of questions remain about the performance of Shellsort and its variants, some of which may lead to results of direct practical utility.

Are there increment sequences that perform better than known ones in practice? Given the large number of potential sequences, we certainly can answer this question in the affirmative. Finding a sequence that leads to running times 25% lower than the best known certainly would be of practical interest, and reducing the time by a factor of two would give a sorting algorithm competitive with the fastest known. Specific results such as the best sequence for $N = 10^3$ or $N = 10^6$ also would be interesting to know. Experimenting with increment sequences can be expensive, and many approaches for searching for better ones have been tried (see, for example, [38]). Promising approaches include: doing minor perturbations of known increment sequences; adding random increments; intermixing increment sequences designed for different reasons; and adding more smaller increments to sequences that lack them (for example, Chazelle's or Selmer's increments)

What is the complexity of Shellsort? Lowering the upper bound by a factor of $\log \log N$ or $(\log \log N)^2$ may have practical impact; denying this possibility would also be of interest.

What is the average running time of Shellsort, as a function of N , for the best choice of t increments, with $t > 2$? No asymptotic results on the average running time are known for the types of increments used in practice. Is the average running time of Shellsort $O(N \log N)$ for some increment sequence? Do increment sequences designed to minimize the average running time differ signif-

icantly from increment sequences designed to minimize the worst-case running time?

Do different types of increment sequences improve the performance of Shellsort variants where the h -sort is replaced by an h -bubble, h -shake, or h -brick pass? Substantial improvements in the performance of Shellsort itself were found by exploiting Frobenius effects, large common factors among increments, and so forth. Are similar improvements available for the variants?

Is there an increment sequence for which replacing the h -sort in Shellsort by an h -brick pass yields a probabilistic sorting network of depth $\alpha \lg N$? For theoretical purposes, a proof for any fixed value of α would be of interest; for practical purposes, empirical evidence substantiating some $\alpha < 2$ would be important and useful. The possibility of the existence of low-depth networks based on Shellsort variants remains the most compelling reason to continue research on the algorithm.

The basic operations on which Shellsort and its variants are based are quite amenable to fast implementation in many hardware environments, and improvements to these methods have the potential to provide the fastest way to sort for many applications. Analysis of the algorithm also has brought many interesting theoretical challenges, and the final chapters on this topic remain to be written.

7 References

This paper summarizes the work of many authors, and certainly cannot do justice to the detailed information given in their original works, listed here.

1. M. Ajtai, J. Komlós, and E. Szemerédi. “An $O(n \log n)$ sorting network,” in *Proc. 15th Ann. ACM Symp. on Theory of Computing*, 1983.
2. M. Ajtai, J. Komlós, and E. Szemerédi. “Sorting in $c \log n$ parallel steps,” *Combinatorica* 3, 1983, 1–19.
3. K. Batcher. Sorting networks and their applications in *Proceedings of the AFIPS Spring Joint Computer Conference* 32, 1968, 307–314.
4. A. Brauer. “On a problem of partitions,” *Amer. J. Math.* 64, 1942, 299–312.
5. W. Curran-Sharp. “Solution to Problem 7382 (Mathematics),” *Ed. times (London)* 1, 1884.
6. B. Chazelle. Private communication, 1983.
7. R. Cypher. “A lower bound on the size of Shellsort sorting networks,” *SIAM J. Computing* 22, 1993, 62–71.
8. W. Dobosiewicz. “An efficient variation of bubble sort.” *Information Processing Letters* 11, 1980, 5–6.
9. G. Gonnet and R. Baeza-Yates. *Handbook of Algorithms and Data Structures*, 2nd edition, Addison-Wesley, Reading, MA, 1991.

10. H. Greenberg. "An algorithm for a linear diophantine equation and a problem of Frobenius," *Numer. Math.* **34**, 1980, 349–352.
11. T. Hibbard. "An empirical study of minimal storage sorting," *Communications of the ACM* **6**, 1963, 206–213.
12. G. Hofmeister. "Zu einem Problem von Frobenius," *Norske Vid. Selsk. Skr.* **5**, 1966, 1–37.
13. J. Incerpi. *A Study of the Worst Case of Shellsort*, Ph.D. thesis, Brown University, Department of Computer Science, 1985.
14. J. Incerpi and R. Sedgewick. "Improved Upper Bounds on Shellsort," *J. of Computer and System Sciences* **31**, 1985, 210–224.
15. J. Incerpi and R. Sedgewick. "Practical variations of Shellsort," *Information Processing Letters* **26**, 1987, 37–43.
16. S. Johnson. "A linear diophantine problem," *Canad. J. Math.* **12**, 1960, 390–398.
17. D. Knuth. *The Art of Computer Programming. Volume 3: Sorting and Searching*, Addison-Wesley, Reading, MA, 1973.
18. D. Knuth. Private communication, 1995.
19. R. Lazarus and R. Frank. "A high-speed sorting procedure," *Communications of the ACM* **3**, 1960, 20–22.
20. T. Leighton. "Tight bounds on the complexity of parallel sorting," in *Proc. 16th Ann. ACM Symp. on Theory of Computing*, 1984.
21. T. Leighton. *Introduction to parallel algorithms and architectures*, Morgan Kaufmann, San Mateo, CA, 1992, 672.
22. T. Leighton and G. Plaxton. "A (fairly) simple circuit that (usually) sorts," in *Proc. 31st IEEE Symposium on Foundations of Computer Science*, 1990, 264–274.
23. P. Lemke. "The performance of randomized Shellsort-like network sorting algorithms," SCAMP working paper no. P18/94, Institute for Defense Analyses, Princeton, NJ, 1994.
24. M. Lewin. "On a linear diophantine problem," *Bull. London Math. Soc.* **5**, 1973, 75–78.
25. A. Papernov and G. Stasevich. "A method of information sorting in computer memories," *Problems of Information Transmission* **1**, 1965, 63–75.
26. C. Plaxton and T. Suel. "Improved lower bounds for Shellsort," *J. of Algorithms*, to appear. Preliminary version in *Proc. 33rd IEEE Symposium on Foundations of Computer Science*, 1992, 226–235.
27. B. Poonen. "The worst case in Shellsort and related algorithms," *J. of Algorithms* **15**, 1993, 101–124.
28. V. Pratt. *Shellsort and Sorting Networks*, Garland, New York, 1979; Ph.D. thesis, Stanford University, 1971.

29. D. Robbins. "Experiments with shaker sort on the CRAY-2," SCAMP working paper no. 22/89, Institute for Defense Analyses, Princeton, NJ, 1989.
30. R. Sedgewick. *Algorithms*, 2nd edition, Addison-Wesley, Reading, Mass, 1988.
31. R. Sedgewick. "A new upper bound for Shellsort," *J. Algorithms* 7, 1986, 159–173.
32. R. Sedgewick. "Bricksort networks," in preparation.
33. R. Sedgewick and P. Flajolet. *An Introduction to the Analysis of Algorithms*, Addison-Wesley, Reading, Mass., 1996.
34. E. Selmer. "On the linear diophantine problem of Frobenius," *J. Reine Angew. Math.* 294, 1977, 1–17.
35. E. Selmer. "On Shellsort and the Frobenius problem," TR 87-27, Department of Mathematics, University of Bergen, Norway, 1987.
36. D. Shell. "A high-speed sorting procedure," *Communications of the ACM* 2, 1959, 30–32.
37. J. Vitter and P. Flajolet, "Analysis of algorithms and data structures," in *Handbook of Theoretical Computer Science A: Algorithms and Complexity*, J. van Leeuwen, ed., Elsevier, Amsterdam, 1990, 431–524.
38. M. Weiss. *Lower Bounds for Shellsort*, Ph.D. thesis, Princeton University, Department of Computer Science, June 1987.
39. M. Weiss. "Empirical study of the expected running time of Shellsort," *Computer Journal* 34, 1991, 88–91.
40. M. Weiss and R. Sedgewick. "Bad cases for shakersort," *Information Processing Letters* 28, 1988, 133–136.
41. M. Weiss and R. Sedgewick. "Tight lower bounds for Shellsort," *J. of Algorithms* 11, 1990, 242–251.
42. M. Weiss and R. Sedgewick. "More on Shellsort increment sequences," *Information Processing Letters* 34, 1990, 267–270.
43. A. Yao. "An analysis of $(h, k, 1)$ Shellsort," *Journal of Algorithms* 1, 1980, 14–50.