Improved upper bounds on shellsort

Janet Incerpi, Robert Sedgewick

IIRIA

CENTRE DE ROCQUENCOURT

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
BP105
78153 Le Chesnay Cedex
France
Tél.: (3) 954 90 20

IMPROVED UPPER BOUNDS ON SHELLSORT

Janet INCERPI
Robert SEDGEWICK

Résumé: *Des résultats de Pratt avaient conduit à la conjecture que le temps d'exécution du tri Shell ("Shellsort") etait $\Theta(N^{3/2})$ lorsque le nombre de passes est en $O(\log N)$. Un résultat récent de Sedgewick a conduit à une borne en $O(N^{4/3})$ mais ce dernier résultat ne peut être amélioré sans une avancée significative sur un problème classique de théorie des nombres. Nous utilisons ici une méthode nouvelle permettant d'atteindre une borne en $O(N^{1+4\sqrt{2\log n}})$.*

Improved Upper Bounds on Shellsort

by

Janet Incerpi
Dept. of Computer Science
Brown University
Providence, RI

Robert Sedgewick
INRIA
78150 Rocquencourt
France

Abstract: *The running time of Shellsort, with the number of passes restricted to $O(\log N)$, was thought for some time to be $\Theta(N^{3/2})$, due to general results of Pratt. Sedgewick recently gave an $O(N^{4/3})$ bound, but extensions of his method to provide better bounds seem to require new results on a classical problem in number theory. In this paper, we use a different approach to achieve $O(N^{1+4/\sqrt{2 \lg N}})$.*

Introduction

Shellsort is a fundamental but little-understood sorting algorithm. A brief description of the algorithm is given below. It is based on a table h_1, h_2, ... of integers called an *increment sequence*. In practice, increment sequences are chosen heuristically based on partial analytic results which have been derived for some specific increment sequences. This algorithm is an attractive candidate for detailed study because it is closely related to classical problems in number theory and because theoretical results translate directly to practice. (A practioner can make immediate use of a good increment sequence, no matter how intricate the analysis.) It is difficult to deny the existence of increment sequences that would make Shellsort the sorting method of choice, for most situations. Moreover, relatively few types of increment sequences have been tried. Some references for Shellsort and some of the analysis that has been done are [6], [8], [9], and [10]; some of this information is summarized below.

The Shellsort algorithm works as follows: given an increment sequence h_1, h_2, ..., a file is sorted by successively h_j-*sorting* it, for j from some integer t down to 1. An array $a[1]$, ..., $a[N]$ is defined to be h_j-sorted if $a[i - h_j] \leq a[i]$ for i from $h_j + 1$ to N. The method used for h_j-sorting is *insertion sort*: for i from $h_j + 1$ to N, we sort the sequence ..., $a[i - 2h_j]$, $a[i - h_j]$, $a[i]$ by taking advantage of the fact that the sequence ..., $a[i - 2h_j]$, $a[i - h_j]$ is already sorted, so $a[i]$ can be inserted by moving larger elements one position to the right in the sequence, then putting $a[i]$ in the place vacated.

A fundamental property of this process is that, if we h-sort a file which is k-sorted, then the file remains k-sorted. Thus, when we come to h_j-sort the file during Shellsort, we know that it is h_{j+1}-, h_{j+2}-, ..., h_t-sorted. This ordering makes the h_j-sort less expensive than if we were to h_j-sort a randomly ordered file.

Shellsort sorts properly whenever the increment sequence ends with $h_1 = 1$, but the running time of the algorithm clearly is quite dependent on the specific increment sequence used. Unfortunately, we have little guidance on how to pick the "best" increment sequences. All the results that we have relate to specific sequences (from a quite large universe) and leave open the possibility of an undiscovered increment sequence with far better performance characteristics than those that have been tried to date.

From a practical standpoint, Shellsort leads to a simple and compact sorting program which works well for small files and for files which are already partially ordered. It is the practical method of choice for files with less than several hundred elements, and each

new increment sequence that we discover raises this bound. Empirical tests by several researchers indicate that there might exist increment sequences for which the average running time is $O(N \log N)$ (for example, see [4]).

From a theoretical standpoint, the study of increment sequences for Shellsort is important because of the potential for a simple constructive proof of the existence of an $O(N \log N)$ sorting network. (An increment sequence of length $O(\log N)$ for which each insertion requires a constant number of steps would imply this.) This was an open problem in the theory of sorting for some time; the existence of such a network was recently presented by Ajtai, Komlos, and Szemerdi [1] but their construction is hardly practical. This result makes the search for a short proof based on Shellsort even more appealing. Weaker results (for example, an $O(N \log N)$ average case) are also worth pursuing because of the practical implications.

In this paper, we are interested in worst-case bounds for the total running time of Shellsort for particular increment sequences. Specifically, we're most interested in increment sequences of length $O(\log N)$: this would be required for an optimal sorting network, and such sequences are the most viable from a practical standpoint. Even with this restriction, the space of possible increment sequences is quite large. For simplicity, in this paper we assume that the sequence increases (although there is no particular requirement for this). Further, we make the following distinction:

Definition: A Shellsort implementation is said to be *uniform* if the increments used to sort N items are all the numbers less than N (taken in decreasing order) from a *fixed* infinite increasing sequence $h_1, h_2, \ldots$.

A non-uniform Shellsort might use a different increment sequence for each file size. Both types are used in practice, though uniform implementations have been studied more heavily. For example, Knuth [6] recommends using a uniform implementa-

tion based on the sequence 1, 4, 13, 40, ..., $\frac{1}{2}(3^k - 1)$, ... On the other hand, in order to use a uniform sequence one must calculate an appropriate starting place and/or save the sequence, so some practitioners find it more convenient to use non-uniform sequences such as $\lfloor N/2 \rfloor$, $\lfloor \lfloor N/2 \rfloor /2 \rfloor$, etc. Unless designed with care, non-uniform sequences are susceptible to bad worst-case performance for some file sizes. Consequently, uniform implementations are more widely used and studied. We use the terminology "uniform $f(N)$-sequence" to refer to an infinite sequence for which the number of integers less than N is $f(N)$.

Shellsort and the Frobenius Problem

To prove upper bounds on the number of steps required for Shellsort, we're interested specifically in the following function:

Definition: $n_d(a_1, a_2, \ldots, a_k) \equiv$ the number of multiples of d which cannot be represented as linear combinations (with non-negative integer coefficients) of $a_1, a_2, \ldots, a_k$.

We assume that $a_1, a_2, \ldots, a_k$ are > 1 (otherwise all integers could be represented) and that $a_1, a_2, \ldots, a_k$ are *independent*: that none can represented as a linear combination with non-negative integer coefficients of the others (otherwise it could be deleted from the list without affecting the result). More important, for $n_d(a_1, a_2, \ldots, a_k)$ to be defined, it must be the case that $a_1, a_2, \ldots, a_k$ do not have a common factor which is not shared by d (otherwise, only those multiples of d which share that common factor could be represented as linear combinations of $a_1, a_2, \ldots, a_k$, and there are an infinite number of multiples of d which do not).

This function is related to Shellsort by the following lemma:

Lemma 1. The number of steps required to h_j-sort a file which is h_{j+1}-, h_{j+2}-, ..., h_t-sorted is

$$O(N n_{h_j}(h_{j+1}, h_{j+2}, \ldots, h_t)).$$

Proof: The number of steps required to insert element $a[i]$ is the number of elements among $a[i-h_j]$, $a[i-2h_j]$, ... which are greater than $a[i]$. Any element $a[i-x]$ with x a linear combination of h_{j+1}, ..., h_t must be less than $a[i]$ since the file is h_{j+1}-, h_{j+2}-, ..., h_t-sorted. Thus, an upper bound on the number of steps to insert $a[i]$, for $1 \leq i \leq N$, is the number of multiples of h_j which are not expressible as linear combinations of $h_{j+1}, h_{j+2}, \ldots, h_t$ or $n_{h_j}(h_{j+1}, h_{j+2}, \ldots, h_t)$. ∎

When $d = 1$, we have $n_1(a_1, a_2, \ldots, a_k)$ (or just $n(a_1, a_2, \ldots, a_k)$) which is the number of positive integers which cannot be represented as linear combinations with nonnegative coefficients of a_1, a_2, ..., a_k. A closely related function is $g(a_1, a_2, \ldots, a_k)$, the largest integer which cannot be so represented. These functions are well-studied in number theory [5,10,11]: to find $g(a_1, \ldots, a_k)$ is the so-called *Frobenius problem*.

The function which arises in Shellsort is related to the standard Frobenius function by the following lemma:

Lemma 2. For $a_1, a_2, \ldots, a_k$ relatively prime,

$$n_d(a_1, a_2, \ldots, a_k) < \frac{g(a_1, a_2, \ldots, a_k)}{d}.$$

Proof: (Note that $g(a_1, a_2, \ldots, a_k)$ is undefined unless $a_1, a_2, \ldots, a_k$ are relatively prime.) Every number greater than $g(a_1, a_2, \ldots, a_k)$ can be represented as a linear combination of $a_1, a_2, \ldots, a_k$; in the worst case all multiples of d less than $g(a_1, a_2, \ldots, a_k)$ cannot. ∎

Previous upper bound proofs for Shellsort have used a combined version of these lemmas:

Lemma 3. The number of steps required to h_j-sort a file which is h_{j+1}-, h_{j+2}-, ..., h_t-sorted is

$$O(N g(h_{j+1}, h_{j+2}, \ldots, h_t)/h_j).$$

Proof: Immediate from Lemmas 1 and 2. ∎

Specific bounds are obtained by solving the Frobenius problem for specific increment sequences. For $k = 2$, we have the original Frobenius problem whose solution dates at least to 1884:

Lemma 4. If a_1 and a_2 are relatively prime, then $g(a_1, a_2) = (a_1 - 1)(a_2 - 1)$.

Proof: See Knuth[6], Ex. 5.2.1–21, or [2]. ∎

For example, this leads directly to an upper bound for h_j-sorting of $O(N h_j)$ when $h_j = 2^j + 1$, since

$$N \frac{g(h_{j+1}, \ldots, h_t)}{h_j} \leq N \frac{g(h_{j+1}, h_{j+2})}{h_j}$$
$$= O(N \frac{h_{j+1}, h_{j+2}}{h_j})$$
$$= O(N h_j).$$

This is the bound of Papernov and Stasevich [7], which was generalized by Pratt [8] to cover a large family of "almost geometric" increment sequences.

Upper bounds on h_j-sorting for sequences with geometric growth translate to upper bounds on the total number of steps required by Shellsort as follows:

Lemma 5. Suppose that an increment sequence h_1, h_2, ... is used to Shellsort a file of size N, with $h_j = \Theta(\alpha^j)$ for some constant α. If the number of steps for h_j-sorting is $O(N h_j^{1/c})$, then the total running time for Shellsort is $O(N^{1+1/(c+1)})$.

Proof: The increments used are h_1, ..., h_t, where t is the largest integer such that h_t is less than N. We use the bound $O(N^2/h_j)$ for large h_j (this comes from considering h_j independent subfiles of size N/h_j, each of which could require $O((N/h_j)^2)$ steps) and the bound $O(N h_j^{1/c})$ for small h_j, switching at $h_j = \Theta(N^{c/(c+1)})$, when both bounds are $O(N^{1+1/(c+1)})$. The total number of steps for Shellsort is

$$\sum_{1 \leq j \leq t_0} O(N h_j^{1/c}) + \sum_{t_0 \leq j \leq t} O\left(\frac{N^2}{h_j}\right).$$

where t_0 is such that $h_{t_0} = \Theta(N^{c/(c+1)})$. Both sums are geometric and are bounded by their largest term $O(N^{1+1/(c+1)})$ in both cases. ∎

For example, Lemma 5, with $c = 1$, gives the $O(N^{3/2})$ upper bound for Shellsort of Papernov and Stasevich [7] and Pratt [8]. In fact, Pratt showed this bound to be tight for a large family of increment sequences (encompassing most of those that have been proposed), where the increments are within an additive constant of a geometric progression.

Sedgewick [10] used general results of Selmer [11] and Johnson [5] for the Frobenius problem for $k = 3$ to develop increment sequences that grow geometrically and the upper bound for h_j-sorting is $O(Nh_j^{1/2})$. This leads to an $O(N^{4/3})$ upper bound for Shellsort (Lemma 5 with $c = 2$). (These sequences are of the form $\alpha 4^j + \beta 2^j + \nu$, not within an additive constant of a geometric progression.) Unfortunately, there are few available results on the Frobenius problem for $k > 3$, and such results would seem to be required to get better upper bounds using this approach.

Furthermore, we can show that a bound of the type $O(N^{1+1/(c+1)})$ is the best that can be achieved with this approach because we have a lower bound on the Frobenius function:

Lemma 6. For $a_1, a_2, \ldots, a_k$ increasing,

$$n(a_1, a_2, \ldots, a_k) = \Omega(a_1^{1+1/(k-1)}).$$

Sketch of Proof: Define $L(m) = \{x \mid x = c_1 a_1 + \cdots + c_k a_k \text{ with } c_1, c_2, \ldots, c_k \geq 0 \text{ and } c_1 + \cdots + c_c = m\}$. Then if $x \in L(M)$ we know that $x \geq ma_1$. Furthermore, $|L(M)| \leq \binom{m+k-1}{k-1}$. Now, for any constant m_0, we have the inequality

$$n(a_1, a_2, \ldots, a_k) \geq (m_0 + 1)a_1 - \sum_{1 \leq m \leq m_0} |L(m)|$$
$$= (m_0 + 1)a_1 - \binom{m_0 + k}{k}.$$

But this function is maximized at $m_0 \approx a^{1/k-1}$ when both terms are about $a_1^{1+1/k-1}$. ∎

If we only consider the effects of c increments $h_{j+1}, \ldots, h_{j+c}$ when h_j-sorting and we use the approach of Lemma 3 (the standard approach), then

Lemma 6 says that the best bound that we can hope for for h_j-sorting is $O(h_j^{1/(c-1)})$, which translates to an $O(N^{1+1/c})$ Shellsort bound by Lemma 5. Thus, the bounds of Papernov-Stasevich ($c = 2$) and Sedgewick ($c = 3$) are best possible in this sense. Below, we show how to achieve $O(N^{1+1/c})$ with $(c \log n)$-uniform sequences for any c, though we do so by circumventing the standard approach of Lemma 3 and using Lemma 2 directly, not by developing new results on the Frobenius problem. Furthermore, we show how this method extends to provide even better bounds. We do so by turning attention to increments which have large common divisors, then by exploiting specific properties of the generalized Frobenius function with two arguments, $n_d(r, s)$.

Generalized Frobenius Problem

It is possible to completely characterize the generalized Frobenius function with two arguments. We have:

Theorem 1. For any positive integers r, s, and d with $\gcd(d, r, s) = \gcd(r, s)$:

$$n_d(r, s) = n_{\frac{d \gcd(d,r,s)}{\gcd(d,r)\gcd(d,s)}}\left(\frac{r}{\gcd(d,r)}, \frac{s}{\gcd(d,s)}\right).$$

If d, r, and s are pairwise relatively prime, then

$$n_d(r, s) = n\left(r, s, \frac{r + b_1 s}{d}, \ldots, \frac{(d-1)r + b_{d-1}s}{d}\right)$$

where b_i is the unique integer between 0 and $d - 1$ such that $ir + b_i s \equiv 0 \pmod{d}$. Note that $n_d(r, s)$ is undefined if $\gcd(d, r, s) \neq \gcd(r, s)$.

Proof: Omitted.

For the constructions of the next section, we actually use a special case of this theorem which can be proved directly from the definition.

Corollary: For integer $z > 1$, $n_{dz}(rz, sz) = n_d(r, s)$. This property holds for more than two arguments: we have $n_{dz}(a_1 z, a_2 z, \ldots, a_k z) = n_d(a_1, a_2, \ldots, a_k)$, but

4

a full characterization such as Theorem 1 for more than two arguments seems complicated.

Applying Lemmas 2 and 4 with the corollary to Theorem 1, we have $n_{dz}(rz, sz) = n_d(r, s) < rs/d$ (if r and s are relatively prime), which is less by a factor of z than the bound $rzsz/dz$ which derives from direct application of Lemmas 2 and 4.

Increment Sequences

Our increment sequences represent a compromise between two classical increment sequences that have been proposed for Shellsort. The first, proposed by Shell, is the geometric sequence $1, 2, 4, 8, 16, \ldots$. The problem with this sequence is that the generalized Frobenius function is always undefined, since even after the application of Theorem 1 $h_{j+1}, h_{j+2}, \ldots$ have a common factor (2) which is not shared by h_j. The practical effect of this is that the worst case is $\Theta(N^2)$, for example for a shuffled file with the $N/2$ smallest elements in the odd positions and the $N/2$ largest elements in the even positions. Because of this effect, Shellsort increment sequences are normally designed to have successive increments relatively prime.

A notable exception is the sequence $1, 2, 3, 4, 6, 9, \ldots$ given by Pratt, which is defined by appending $2z$ and $3z$ to the sequence for every element z in the sequence. Thus, by the corollary to Theorem 1, the running time for each increment is $n_1(2,3) = 1$. Unfortunately, there are $\Theta(\log^2 N)$ increments less than N, and even after applying tradeoffs as in Lemma 5, the running time is always $\Theta(N \log^2 N)$. Increment sequences with $O(\log N)$ increments are of more interest because in principle, the running time for such sequences could be $O(N \log N)$ on the average (or even in the worst case), and in practice, the large number of passes required for Pratt's sequence makes it slower than typical $O(\log N)$-pass Shellsorts.

Thus, our goal is to design a geometrically increasing sequence in which successive increments have both large common factors and small relatively prime factors. Our method for doing so is to build up increments by multiplying together selected terms of a "base" sequence $a_1, a_2, \ldots$.

Given a constant c, we associate c increments with each term of the base sequence, each increment formed by multiplying together c terms of the base sequence. To simplify the discussion, we'll first consider explicitly the increment sequence formed for $c = 3$; the extension to larger c follows directly.

Specifically, for $c = 3$, we form an increment sequence by interleaving the three sequences

$$a_1 a_2 a_3, \ a_2 a_3 a_4, \ \ldots, \ a_i a_{i+1} a_{i+2}, \ \ldots$$
$$a_1 a_2 a_4, \ a_2 a_3 a_5, \ \ldots, \ a_i a_{i+1} a_{i+3}, \ \ldots$$
$$a_1 a_3 a_4, \ a_2 a_4 a_5, \ \ldots, \ a_i a_{i+2} a_{i+3}, \ \ldots$$

(and, of course, prepending 1). Now, each increment has exactly two "a" factors in common with two increments that appear later in the sequence, which leads directly to an application of the corollary to Theorem 1. We have

$$n_{a_i a_{i+1} a_{i+2}}(a_{i+1} a_{i+2} a_{i+3}, a_{i+1} a_{i+2} a_{i+4})$$
$$= n_{a_i}(a_{i+3}, a_{i+4})$$
$$n_{a_i a_{i+1} a_{i+3}}(a_{i+1} a_{i+2} a_{i+3}, a_{i+1} a_{i+3} a_{i+4})$$
$$= n_{a_i}(a_{i+2}, a_{i+4})$$
$$n_{a_i a_{i+2} a_{i+3}}(a_{i+1} a_{i+2} a_{i+3}, a_{i+2} a_{i+3} a_{i+4})$$
$$= n_{a_i}(a_{i+1}, a_{i+4})$$

If the elements $a_{i+1}, a_{i+2}, a_{i+3}$, and a_{i+4} are all relatively prime, and if each term is within a constant factor of the previous, then these are all $O(a_i)$, by Lemmas 2 and 4. Therefore, by Lemma 3, the number of steps to h-sort is $O(h^{1/3})$ for each increment h in this sequence. (For 1-sorting, we must argue separately that the running time is $O(1)$ if $a_1, a_2, \ldots, a_6$ are all $O(1)$: the running time for 1-sorting is $O(n(a_1 a_2 a_3, a_4 a_5 a_6))$ since those two increments are relatively prime.) Now, by Lemma 5, we get an $O(N^{5/4})$ bound for this sequence.

The extension of this argument to general c is straightforward:

Theorem 2. Given a constant c, there exists a uniform ($\log N$)-sequence of increments for which the running time of Shellsort is $O(N^{1+1/(c+1)})$.

Sketch of Proof: As before, the increment sequence is 1 followed by an interleaving of the c sequences

$$\{(\prod_{0 \le k \le c} a_{i+k})/a_{i+c_0}\}_{i \ge 1}$$

where c_0 ranges from c down to 1. For example, for $c = 5$ we have

$$a_1 a_2 a_3 a_4 a_5, \ldots, a_i a_{i+1} a_{i+2} a_{i+3} a_{i+4}, \ldots$$
$$a_1 a_2 a_3 a_4 a_6, \ldots, a_i a_{i+1} a_{i+2} a_{i+3} a_{i+5}, \ldots$$
$$a_1 a_2 a_3 a_5 a_6, \ldots, a_i a_{i+1} a_{i+2} a_{i+4} a_{i+5}, \ldots$$
$$a_1 a_2 a_4 a_5 a_6, \ldots, a_i a_{i+1} a_{i+3} a_{i+4} a_{i+5}, \ldots$$
$$a_1 a_3 a_4 a_5 a_6, \ldots, a_i a_{i+2} a_{i+3} a_{i+4} a_{i+5}, \ldots$$

Now, we note that each increment has exactly $c-1$ factors in common with two increments that appear later in the sequence, which allows application of the corollary to Theorem 1. When

$$d = \frac{1}{a_{i+c_0}} \prod_{0 \le k \le c} a_{i+k},$$

$$r = \prod_{0 \le k < c} a_{i+1+k},$$

$$s = \frac{1}{a_{(i+1)+(c_0-1)}} \prod_{0 \le k \le c} a_{i+1+k},$$

then

$$n_d(r, s) = n_{a_i}(a_{i+c_0}, a_{i+c+1}).$$

(For example, when $c = 5$, $n_{a_1 a_2 a_3 a_4 a_6}(a_2 a_3 a_4 a_5 a_6, a_2 a_3 a_4 a_6 a_7) = n_{a_1}(a_5, a_7)$.) This works except for $c_0 = 1$, when we take $s = (\prod_{0 \le k < c} a_{i+2+k})$ which still gives $n_d(r, s) = n_{a_i}(a_{i+1}, a_{i+c+1})$. Again, if all $a_{i+1}, \ldots, a_{i+c+1}$ are relatively prime and related by a constant factor then these are all $O(a_i)$ which leads to a bound of $O(Nh^{1/c})$ for each increment in this sequence. This gives a Shellsort bound of $O(N^{1+1/(c+1)})$ by Lemma 5, using the same argument as before for 1-sorting.

The proof is completed by exhibiting a sequence $a_1, a_2, \ldots$ which satisfies the conditions above which

is easy because of the density of primes. For example, we can take a_i to be the smallest prime greater than or equal to 2^i, to get a geometrically increasing sequence of primes $1, 2, 5, 11, 17, \ldots$ which satisfies the conditions. ∎

Note that the constant implied by the O-notation in Theorem 2 is exponential in c. This makes the increment sequences hardly of practical use. Next, we'll examine sequences built according to the same principle as those above but which have good practical performance and even better asymptotic bounds:

Theorem 3. There exists a uniform ($\log N$)-sequence for which the running time of Shellsort is

$$O(N^{1+4/\sqrt{2\lg N}}).$$

Proof: As above, we start with a base sequence a_1, a_2, a_3, ... of relatively prime integers. In this case, we construct the sequence as follows:

a_1	$a_1 a_2$	$a_1 a_2 a_3$	$a_1 a_2 a_3 a_4$	$\ldots$
	$a_1 a_3$	$a_1 a_2 a_4$	$a_1 a_2 a_3 a_5$	$\ldots$
		$a_1 a_3 a_4$	$a_1 a_2 a_4 a_5$	$\ldots$
			$a_1 a_3 a_4 a_5$	$\ldots$
				$\ldots$

The cth column in the table is formed by starting with $\prod_{1 \le i \le c} a_i$, then multiplying each element in the previous column by a_{c+1}. This ensures that each increment exactly divides two increments which appear later in the sequence. The following table gives the upper bound for the increment appearing in the corresponding position in the above sequence:

$n(a_2, a_3)$	$n(a_3, a_4)$	$n(a_4, a_5)$	$n(a_5, a_6)$	$\ldots$
	$n(a_2, a_4)$	$n(a_3, a_5)$	$n(a_4, a_6)$	$\ldots$
		$n(a_2, a_5)$	$n(a_3, a_6)$	$\ldots$
			$n(a_2, a_6)$	$\ldots$
				$\ldots$

If we use c columns in the table, then we use $\frac{1}{2}(c^2 - c)$ increments, all less than $\prod_{1 \le i \le c} a_i$ with a total cost of less than $N(\sum_{1 \le i \le c} a_i)^2$. (This bound follows quickly from the fact that $n(r,s) < rs$.) Once again, we achieve good asymptotics by proper choice of the base sequence. Specifically, we take a_i to be the smallest prime greater than or equal to 2^i, so that

$$a_i = O(2^i)$$
$$\prod_{1 \le i \le c} a_i = O(2^{\frac{1}{2}(c^2 - c)})$$
$$\sum_{1 \le i \le c} a_i = O(2^c)$$

Using all the increments less than N corresponds to taking $c = \sqrt{2 \lg N}$, for a total cost of

$$O(N 2^{2\sqrt{2 \lg N}}) = O(N^{1 + 4/\sqrt{2 \lg N}}),$$

as desired. ∎

There is a quite simple proof of the same asymptotic result for non-uniform sequences, due to B. Chazelle [3]. This result actually motivated the search for the sequence of Theorem 3.

Proof of Theorem 3: non-uniform case (Chazelle)
Simply use Pratt's method, starting with $(\alpha - 1)$ and α for an appropriately chosen α (instead of 2 and 3). The running time is bounded by $N\alpha^2$ for each of the $O((\log_\alpha N)^2)$ increments, for a total of

$$N (\lg N)^2 \frac{\alpha^2}{(\lg \alpha)^2}.$$

Now, take α such that $(\lg \alpha)^2 = \lg N$, or $\alpha = 2^{\sqrt{\lg N}}$, for a total of $N \lg N 2^{2\sqrt{\lg N}} = O(N^{1 + (2 + \epsilon)/\sqrt{\lg N}})$ for any $\epsilon > 0$. ∎

Chazelle's proof tends to favor larger increments (close to N), while the given proof of Theorem 3 tends to favor smaller increments (closer to 1).

The table below shows the number of exchanges required by Knuth's sequence, the uniform sequence suggested after Theorem 3, and the non-uniform sequence, averaged over a few random files for various file sizes.

	10000	40000
Knuth	242110	1317825
Thm 3 (uniform)	219536	1054873
Thm 3 (non-uniform)	242248	1153723

Conclusions

Despite the substantial improvements that we have been able to make in upper bounds for Shellsort, the results still pertain to particular increment sequences of somewhat artificial construction and there seems to be room for improvement. Furthermore, even the bounds derived for the given sequences are not tight. For example, they only derive from the effects of a few of the previous passes and they don't take into account obvious correlations in insertion costs of successive elements.

It seems likely that better bounds can be obtained by taking such effects into account, and these are worth exploring because of the direct practical benefits that accrue. The question of whether there exists an increment sequence of $O(\log N)$ numbers which produces an $O(N \log^2 N)$ or $O(N \log N)$ Shellsort still remains open.

References

[1]   M. Ajtai, J. Komlos, and E. Szemeredi, "An $O(n \log n)$ Sorting Network" *Proceedings 15th Annual ACM Symposium of Theory of Computing*, April 1983, Boston, MA.

[2]   W. J. Curran-Sharp, Solution to Problem 7382 (Mathematics), *Educational Times*, London 1 (1884).

[3]   B. M. Chazelle, private communication, 1983.

[4]   W. Dobosiewicz, "Sorting by distributive partitioning", *Information Processing Letters* 7 (1) (1978), 1-7.

[5]   S. M. Johnson, "A linear diophantine problem", *Canadian J. Math.* 12 (1960), 390-398.

[6] D. E. Knuth, *The Art of Computer Programming. Volume 3: Sorting and Searching*, Addison-Wesley, Reading, Mass. (1973).

[7] A. A. Papernov and G. V. Stasevich, "A method of information sorting in computer memories", *Problems of Information Transmission* **1**, (3) (1965), 63–75.

[8] V. Pratt, *Shellsort and Sorting Networks*, Garland Publishing, New York (1979). (Originally presented as the author's Ph.D. thesis, Stanford University, 1971.)

[9] R.Sedgewick, *Algorithms*, Addison-Wesley, 1983.

[10] R.Sedgewick, "A New Upper Bound for Shellsort", *J. of Algorithms* (to appear).

[11] E. S. Selmer, "On the linear diophantine problem of Frobenius", *J. reine angew. Math.* **294** (1977), 1–17.

[12] D. L. Shell, "A high-speed sorting procedure", *Communications of the ACM* **2**, 7(1959), 30–32.