Searching & Sorting in Java – Shell Sort

Although our previous sorting algorithms (insertion, selection, bubble) work well enough for short lists, they tend to slow down drastically with long lists. This is due to the large number of comparisons that are needed to find the correct positions for each data element.

This is not always true. If, for example, the data is almost in order, an algorithm like insertion sort will be very fast, because only a few comparisons are needed for each value.

Shell sort adapts insertion sort to produce a sorting technique that works well even with long lists that are randomly ordered. The key is that data can be moved long distances in the list with only a few comparisons.

The sort is based upon the following idea: Rather than sorting the entire list at once, we sort every k^{th} element. Such a list is said to be *k-sorted*. A k-sorted list is made up of *k* sublists, each of which is sorted, interleaved together.

Suppose we are given an unsorted list of integers.

34	21	40	12	27	18	29	13	25	17	11	38

This can be decomposed into three sublists, each with 4 elements.

34			12			29			17		
	21			27			13			11	
		40			18			25			38

If we now sort each sublist independently, we obtain

12			17			29			34		
	11			13			21			27	
		18			25			38			40

which yields the following 3-sorted list.

12	11	18	17	13	25	29	21	38	34	27	40

If we compare the unsorted list, to the 3-sorted list, to the sorted list, we see that the 3-sorted list has moved significantly toward the sorted goal by sorting small sublists (which is comparatively fast).

Unsorted	34	21	40	12	27	18	29	13	25	17	11	38
3-sorted	12	11	18	17	13	25	29	21	38	34	27	40
Sorted	11	12	13	17	18	21	25	27	29	34	38	40

Shell sort repeatedly used insertion sort to create sorted sublists. Initially a *k*-sorted list is created using a large value of *k* so that values can be moved long distances. On subsequent passes, smaller values of *k* are used, until, on the final pass, the value of *k* is set to one. A 1-sorted list is completely sorted.

Consider the following example which creates lists that are 7-sorted, 3-sorted, and then 1-sorted.

Unsorted	64	31	10	40	22	49	82	20	29	56	40	18	19	27	26

First Pass

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	20							26							64
		29							31						
			10							56					
				40							40				
					18							22			
						19							49		
							27							82	

7-sorted	20	29	10	40	18	19	27	26	31	56	40	22	49	82	64

Second Pass

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	20			27			40			49			56		
		18			26			29			40			82	
			10			19			22			31			64

3-sorted	20	18	10	27	26	19	40	29	22	49	40	31	56	82	64

Third Pass 1-sorted	10	18	19	20	22	26	27	29	31	40	40	49	56	64	82

In using Shell Sort, we noted that the first pass should used a large value of k so that values could move long distances quickly. We also noted that the last pass must use $k = 1$ to completely sort the list. This discussion begs the question: What is the best sequence of values for k? The answer to that question is not well defined, despite a considerable amount of work done on the problem.

One sequence that has been found to give good results follows the pattern 1, 4, 13, 40, ...

Each term in the sequence is three times the previous term, plus one.

Exercises

1. Given the following data, show how they would appear after they have been 5-sorted.

   26   37   21   41   63   19   61   72   55   29   47   18   26   22

2. Starting with the same set of data from question 1, show how they would appear 4-sorted.

3. How would you answer the following argument against using Shell Sort? "The last step of Shell Sort, using *k* = 1, is simply a normal insertion sort. Since Shell Sort performs many preliminary steps before this final one, it must be slower than a single insertion sort."

4. Suppose you were going to write a version of Shell Sort using the sequence of k-sorts suggested previously. For a list containing *n* elements, the first value of *k* that should be used is the largest value in the sequence that is less than *n*. For example, in a list of 50 elements, the largest *k* value would be 40, so the first pass should be 40-sorted.

   (a) Write a sequence of statements that will initialize *k* correctly for a given value of *n*.

   (b) Write a statement that will, for any value of *k* in the sequence, produce the next smaller value of *k*.

5.
   (a) Write a method `shellSort` to sort an array of `int` values in ascending order. In performing the *k*-sorts, use the sequence of values of *k* suggested in the text. Be sure to use insertion sort at each stage of the sort.

   (b) Test your method by writing a complete program that first generates an array of 500 random `int` values in the range [0, 999], prints this array (ten values per line), sorts the array using your `shellSort` method, and then prints the resulting array (ten values per line).

6. Experiment with using Shell Sort with sequences where *k* is other than the suggested values, testing your sequences on large arrays of integers and noting the time required by the sort in each case. To measure the time taken by a sort, you can use the method `currentTimeMillis()` in the `System` class. The method has the signature

   ```
   public static long currentTimeMillis()
   ```

   and it returns a long value, the number of milliseconds since midnight, January 1, 1970.