

COSC 311: ALGORITHMS  
HW1: SORTING AND ASYMPTOTIC ANALYSIS  
Solutions

**Problem 1: Asymptotic analysis.**

(a) Suppose you know that  $f(n) \in O(n^2)$  (and you don't know anything else about  $f$ ). Are each of the following claims **true** for all possible choices of  $f$ , **false** for all possible choices of  $f$ , or **sometimes true** (i.e., true for some specific choices of  $f$  and false for others)? Briefly explain your reasoning. You do not need to give any formal proofs.

(i)  $f(n) \in \Omega(n)$

**Solution:** Sometimes true. For example, the function  $f(n) = n$  is in  $O(n^2)$  and  $\Omega(n)$ , but the function  $f(n) = 3$  is in  $O(n^2)$  but not  $\Omega(n)$ .

(ii)  $f(n) \in \Theta(n^2)$

**Solution:** Sometimes true. For example, the function  $f(n) = 3n^2$  is in  $\Theta(n^2)$ , but the function  $f(n) = 97n$  is not.

(iii)  $f(n) \in O(n^3)$

**Solution:** True. The function  $n^2$  is asymptotically upper bounded by the function  $n^3$ , so any function  $f$  that is in  $O(n^2)$  must also be in  $O(n^3)$ .

(iv)  $f(n) \in O(n \lg n)$

**Solution:** Sometimes true. For example, the function  $f(n) = n$  is in  $O(n^2)$  and  $O(n \ln n)$ , but the function  $f(n) = 4n^2$  is in  $O(n^2)$  but not  $O(n \lg n)$ .

(v)  $f(n) \in \Theta(n^2 \lg n)$

**Solution:** False. For a function to be in  $\Theta(n^2 \lg n)$  it must be, in particular, lower bounded by  $n^2 \lg n$ . But for a function to be in  $O(n^2)$  it must be upper bounded by  $n^2$ . It is not possible to satisfy these two properties simultaneously.

(vi)  $f(n) \in \Omega(n^4)$

**Solution:** False. For a function to be in  $\Omega(n^4)$  it must be lower bounded by  $n^4$ , which is not possible if it is upper bounded by  $n^2$ .

(vii)  $f(n) \in O(14n^2 + 74)$

**Solution:** True. The function  $14n^2 + 74$  is asymptotically equivalent to  $n^2$ ; that is,  $\Theta(14n^2 + 74) = \Theta(n^2)$  (and so their big- $O$  and big- $\Omega$  classes are also equivalent).

(b) Let  $f(n) = 3n^2 + n$ . Prove that  $f(n) \in \Theta(n^2)$ .

**Solution:** Let  $c_1 = 3$ ,  $c_2 = 4$  and  $n_0 = 1$ . We will show that for all  $n \geq n_0$ ,  $c_1 n^2 \leq f(n) \leq c_2 n^2$ . We have:

$$f(n) = 3n^2 + n \geq 3n^2 = c_1 n^2$$

since  $n > 0$ . Thus the first inequality is satisfied. We also have:

$$f(n) = 3n^2 + n \leq 3n^2 + n^2 = 4n^2 = c_2 n^2$$

since, when  $n \geq n_0 = 1$ ,  $n \leq n^2$ . Thus the second inequality is satisfied.

We have shown that there exist a  $c_1$ ,  $c_2$ , and  $n_0$  such that for all  $n \geq n_0$ ,  $c_1 n^2 \leq f(n) \leq c_2 n^2$ . Hence  $f(n) \in \Theta(n^2)$ .

(c) Let  $f(n) = 100n - 7$ . Is  $f(n) \in O(n \lg n)$ ? Is  $f(n) \in \Omega(n \lg n)$ ? Prove or disprove each.

**Solution:**  $f(n) \in O(n \lg n)$ . To see this, let  $c = 100$  and  $n_0 = 2$ . Then for all  $n \geq n_0$  we have:

$$f(n) = 100n - 7 \leq 100n \leq 100n \lg n = cn \lg n.$$

On the other hand,  $f(n) \notin \Omega(n \lg n)$ . To see this, suppose there exists a  $c$  and an  $n_0$  such that for all  $n \geq n_0$   $cn \lg n \leq f(n)$ . That is:

$$\begin{aligned} cn \lg n &\leq 100n - 7 \\ c \lg n &\leq 100 - \frac{7}{n} \\ c \lg n + \frac{7}{n} &\leq 100 \end{aligned}$$

which is a contradiction because the left-hand side approaches  $\infty$  as  $n$  approaches  $\infty$ , but the right-hand side is a constant.

**Problem 2: Analyzing heapsort.** Here is the heapsort algorithm that we wrote together in class:

```
1 Heapsort(A) // A is an unsorted array of length n
2   A.heapify() // turn A into a maxheap
3   for i = A.length-1 to 0
4     swap(A, 0, i)
5     A.siftDown(0)
```

Your job in this problem is to analyze the asymptotic runtime of heapsort, using the steps below to help you. You may find that you need to review the asymptotic analysis of heaps!

(a) In class we wrote the following iterative heapify method:

```
10 A.heapify()
11   for i = (n-1)/2 to 0
12     A.siftDown(i)
```

We'll first analyze the runtime of heapify. Let  $j$  denote the depth of a node in the heap, where the root is at depth 0.

(i) If a node in position  $i$  in the heap array is at depth  $j$ , at most how many swaps can be performed during the call to `siftDown` on line 12 (in terms of  $j$ )?

**Solution:** The total height of the heap is  $\lg n$ , so we can perform at most  $\lg n - j$  swaps when sifting down from depth  $j$ .

(ii) Over the course of the for loop in line 11, how many calls to `siftDown` are made from depth  $j$ ?

**Solution:** There is one node at depth 0 (the root), two nodes at depth 1 (the root's children), four nodes at depth 2, etc. In general, there are  $2^j$  nodes at depth  $j$ , so we make  $2^j$  calls to `siftDown` at depth  $j$ .

(iii) If the heap has  $n$  elements, what values can  $j$  take on (you can assume that  $n$  is one less than a power of 2)?

**Solution:** Again, the total height of the heap is  $\lg n$ , so  $j$  can take on values from 0 to  $\lg n$ .

(iv) Put (i)-(iii) together to write a summation that captures the total number of swaps performed during a call to `heapify`.

**Solution:** The total number of swaps is  $\sum_{j=0}^{\lg n} 2^j (\lg n - j)$ .

(v) Find a closed-form expression (i.e., an expression that does not contain a summation) that gives an upper bound on the number of swaps performed during `heapify`. (Hint: you may find it useful to know that  $\sum_{j=0}^{\infty} \frac{j}{2^j} = 2$ .)

**Solution:** Let  $i = \lg n - j$ . Then we can simplify our summation as follows:

$$\begin{aligned} \sum_{j=0}^{\lg n} 2^j (\lg n - j) &= \sum_{i=0}^{\lg n} 2^{\lg n - i} \cdot i \\ &= 2^{\lg n} \sum_{i=0}^{\lg n} \frac{i}{2^i} \\ &= n \sum_{i=0}^{\lg n} \frac{i}{2^i}. \end{aligned}$$

At this point we use our hint, which says that  $\sum_{j=0}^{\infty} \frac{j}{2^j} = 2$ . Our summation, which only goes up to  $\lg n$ , must be less than the infinite sum. So we have:

$$n \sum_{i=0}^{\lg n} \frac{i}{2^i} \leq 2n.$$

Hence  $2n$  is an upper bound on the runtime for our `heapify` method.

(b) Based on your answer to part (a) and any other analysis you need, what is the asymptotic (big- $O$ ) runtime of heapsort? Give your answer in the simplest form possible, i.e., omitting constant factors and lower-order terms. Briefly explain your reasoning (you might find it helpful to refer to the line numbers in the pseudocode above).

**Solution:** The overall runtime of heapsort is  $O(n \lg n)$ . Here's why:

- Line 2: we've just seen in part (c) that `heapify` takes time  $O(n)$ .

- Lines 3-5: the loop runs  $n$  times. Line 4 is constant time and the `siftDown` in line 7 takes time  $O(\lg n)$ , so in all we have  $O(n \lg n)$  (number of iterations multiplied by time per iteration).

Adding these up, we have  $O(n) + O(n \lg n)$  which is asymptotically equal to  $O(n \lg n)$  since  $O(n) \subset O(n \lg n)$ .

**Problem 3: What do we learn from asymptotic analysis?** In this problem you will explore the value and limitations of asymptotic analysis by experimenting with three sorting algorithms: insertion sort, mergesort, and a new algorithm called bubble sort.

```

1 BubbleSort(A) // A is an unsorted array of length n
2   for i = 0 to n-1
3     for j = 1 to n - 1 - i
4       if A[j-1] > A[j]
5         swap(A, j-1, j)

```

(a) Explain in a short paragraph what bubble sort is doing (i.e., what is the sequence of steps it follows?) and why it works. Your goal is to help another CS undergrad understand what bubble sort does and convince them that this algorithm will in fact result in a sorted array.

**Solution:** Bubble sort “bubbles” larger elements up to the top (right side) of the array. The idea is to walk down the array, comparing each pair of consecutive elements and swapping them if they are out of order. The algorithm walks down the array  $n$  times, at the end of which all  $n$  elements in order.

To understand why bubble sort works, we make the following observation: on the  $i$ th iteration of the outer loop, the  $i$ th largest element ends up in the correct place. Why? Consider the 0th iteration. If the largest element is not already in place, then it’s in some position  $k$  in the middle of the array. Once the inner loop reaches position  $k$ , the largest element gets swapped with the element to the right of it on each remaining iteration of the inner loop. This brings the largest element to the right-most position. On the next iteration, the second-largest element bubbles up to the second-to-last position in the array for the same reason: it is larger than, and gets swapped with, every element it’s compared to in the inner loop until it reaches the end of the unsorted section of the array (the last element is already sorted, so the inner loop only goes up to position  $n-2$ ). Continuing with this line of reasoning, we see that after  $i$  iterations, the  $i$  largest elements are in sorted order at the right of the array, and the first  $n-i$  elements remain unsorted. So after  $n$  iterations the entire array is sorted.

(b) What is the asymptotic runtime of bubble sort? Justify your answer in a sentence or so.

**Solution:** Bubble sort runs in  $\Theta(n^2)$ . Lines 4-5 take constant time. These are inside the loop on line 3, which runs for  $\Theta(n)$  iterations, so the entire loop in lines 3-5 takes  $\Theta(n)$ . This is inside the loop on line 2, which runs for  $n$  iterations, so the entire loop in lines 2-5 (and hence the entire algorithm) takes time  $\Theta(n^2)$ .

(c) Based on your answer to (b), how much wall-clock time do you expect bubble sort to take relative to mergesort? What about insertion sort? (“Wall-clock” time is the amount of time that passes between when you start running your program and when it finishes. This is not the same as the number of operations the computer does while running your code, it instead has to do with the user experience of waiting for your code to finish running.)

**Solution:** Mergesort has asymptotic runtime  $\Theta(n \lg n)$ . Since  $n \lg n$  grows strictly slower than  $n$ , we’d expect mergesort to run faster than bubble sort. Both bubble sort and insertion sort are in  $\Theta(n^2)$  (in the worst case; insertion sort has a  $\Theta(n)$  best-case runtime), so based on the asymptotic analysis we’d expect them to have similar wall-clock runtimes.

(d) Download my `Sort.java` code from:

<https://kgardner.people.amherst.edu/courses/f19/cosc311/hw/hw1/Sort.java>.

The file contains implementations of three sorting algorithms: insertion sort, mergesort, and bubble sort. Your job is to run an experiment to compare their performance. Specifically, do the following:

- Fill an array of length  $n$  with randomly generated ints between 0 and  $10^6$ .
- Sort the array using each of insertion sort, mergesort, and bubble sort. For each algorithm, time how long the algorithm takes by recording the system time before and after running the sorting method. You can do this as follows:

```
long start_time = System.nanoTime();
// whatever code you want to time
long end_time = System.nanoTime();
long time_my_code_took = end_time - start_time;
```

Repeat the above for each of the three sorting algorithms for many values of  $n$ . You will want to try some small values of  $n$  (less than 1000), some large values of  $n$  (greater than 100000), and some in between. Be sure to re-fill your array with a new set of random numbers in between each sorting algorithm—otherwise you will be sorting data that is already in sorted order, which might lead to inaccurate results.

Graph your results: make one graph with all three algorithms, with  $n$  on the  $x$ -axis and time on the  $y$ -axis (you may find that it is easier to see what’s going on if you also make separate graphs for small  $n$  and large  $n$ ).

Discuss your findings. Do your experimental results match up to your predictions in part (b)? Is the asymptotic analysis equally predictive at small  $n$  and large  $n$ ? Are there things you learn about the relative performance of the three algorithms from the experimental results that you don’t learn from the asymptotic analysis? Any other interesting observations?

**Solution:** The asymptotic analysis is a good predictor of the general shape of the runtime curve as a function of  $n$ , particularly as  $n$  gets large (see Figure 1). We can see from the graphs that both bubble sort and insertion sort have a quadratic shape (it helps to zoom in and remove bubble sort from the graph to see that insertion sort is quadratic), whereas the shape for mergesort matches

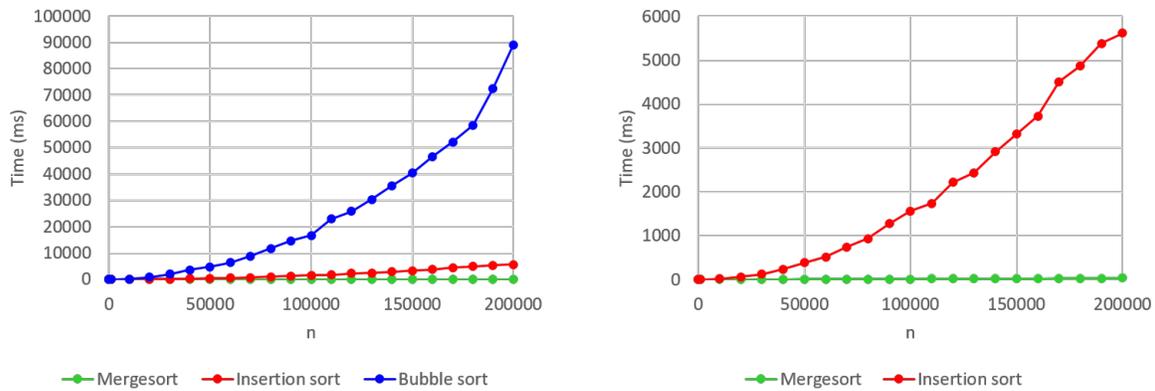


Figure 1: Comparing the runtimes of mergesort, insertion sort, and bubble sort. Left: all three sorting algorithms. Right: zooming in on just mergesort and insertion sort.

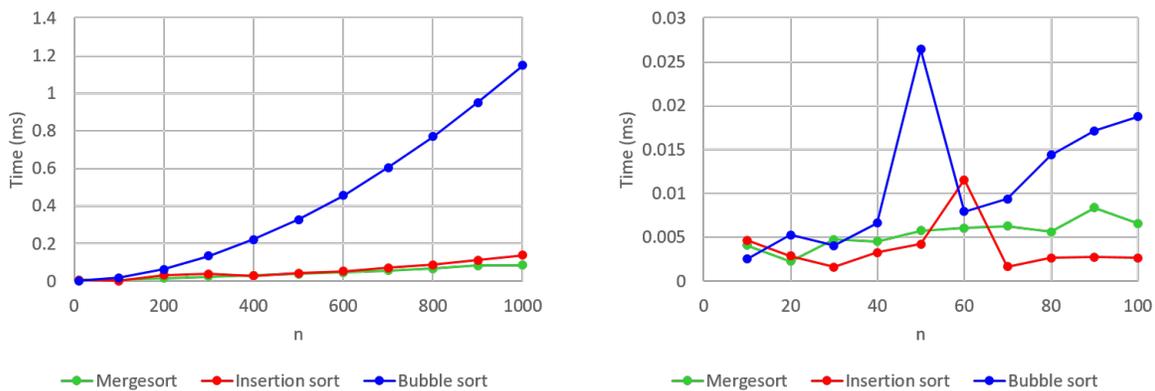


Figure 2: Comparing the runtimes of mergesort, insertion sort, and bubble sort at low  $n$ . Left:  $n \leq 1000$ . Right: zooming in further on  $n \leq 100$ .

the shape of the function  $f(n) = n \lg n$ . Since  $n^2$  and  $n \lg n$  have very different behavior as  $n$  gets large, the asymptotic analysis also does a good job predicting that mergesort will perform better than insertion sort and bubble sort. We can see from the graphs that this is indeed what happens.

When  $n$  is low (see Figure 2), the asymptotic analysis is worse at predicting the relative performance of the three algorithms. We can still see in the graph on the left that bubble sort is worse than mergesort, and even that it looks to be quadratic, but we're no longer able to differentiate between insertion sort and mergesort, which have nearly identical performance when  $n \leq 1000$ . At even lower  $n$  (see the graph on the right for  $n \leq 100$ ), there's too much noise for us to distinguish between the three algorithms at all. This is fine: asymptotic analysis is not really intended to tell us anything useful about what is going on at small values of  $n$ .

What the asymptotic analysis is missing is the ability to differentiate between algorithms that have

the same asymptotic performance. Insertion sort clearly takes much less time to run than bubble sort, even though they both have  $\Theta(n^2)$  worst-case runtime. (This is in part because bubble sort is also  $\Theta(n^2)$  in the best case, whereas insertion sort is  $\Theta(n)$  in the best case. However, it is mostly because insertion sort has much better constant factors.) Asymptotic analysis ignores lower-order terms and constant multiplicative factors, and in this case these factors make a big difference!

What this tells us is that while asymptotic analysis is useful for getting the big picture about how to order algorithms with respect to their runtime, it doesn't tell the full story. When you go out into the world and are asked to select an algorithm to perform a certain task, you can be reasonably confident that, at high enough  $n$ , an  $\Theta(n)$  algorithm will outperform, say, an  $\Theta(n^2)$  algorithm. But not all  $\Theta(n)$  algorithms are created equal, and we need different tools to evaluate the relative performance of algorithms that belong to the same asymptotic class.