

CPSC 320 Sample Solution, Asymptotic Analysis

January 18, 2018

1 Comparing Orders of Growth for Functions

For each of the functions below, give the best Θ bound you can find and then arrange these functions by increasing order of growth.

SOLUTION:

Original function	Good Θ bound	Notes
$\ln n$	$\Theta(\lg n)$	$\lg n$, $\lg n$, and $\ln n$ differ by a constant factor
$\frac{n}{\log n}$	$\Theta\left(\frac{n}{\lg n}\right)$	But that does not make $\log n$ a constant factor!
$55n + 4$	$\Theta(n)$	
$1.5n \lg n$	$\Theta(n \lg n)$	
$2n \log(n^2) = 4n \log n$	$\Theta(n \lg n)$	(same "rank" as previous line)
$n + n^2$	$\Theta(n^2)$	
$(n \lg n)(n + 1) = n^2 \lg n + n \lg n$	$\Theta(n^2 \lg n)$	
$\sqrt{n} \sqrt{n} = (n^{\frac{1}{2}}) \sqrt{n}$	$\Theta(n^{\frac{\sqrt{n}}{2}})$	(challenge problem ; what's the "best" form?)
2^n	$\Theta(2^n)$	
$1.6^{2n} = (1.6^2)^n = 2.56^n$	$\Theta(2.56^n)$	
$n!$	$\Theta(n!)$	
$(n + 1)!$	$\Theta((n + 1)!) = \Theta(n \times n!)$	

Note: $\lg n = \log_2 n$, $\ln n = \log_e n$, and $\log n = \log_{10} n$.

For the (challenge-problem-only) comparison of $n^{\frac{\sqrt{n}}{2}}$ and 2^n , one approach is to make them look more similar. $f(x) = 2^{\lg f(x)}$. So, $n^{\frac{\sqrt{n}}{2}} = 2^{\lg n^{\frac{\sqrt{n}}{2}}} = 2^{\frac{\sqrt{n}}{2} \lg n}$. Now, if we take the ratio of these two functions: $\frac{2^n}{2^{\frac{\sqrt{n}}{2} \lg n}} = 2^{n - \frac{\sqrt{n}}{2} \lg n}$. $n - \frac{\sqrt{n}}{2} \lg n \in \Theta(n)$; so, $2^{n - \frac{\sqrt{n}}{2} \lg n}$ goes to infinity, and 2^n dominates our big nasty function. Just for fun, note that that means our big nasty function is **sub-exponential** yet **super-polynomial**.

2 Functions/Orders of Growth for Code

Give and briefly **justify** good Θ bounds on the worst-case running time of each of these pseudocode snippets dealing with an array A of length n . Note: we use 1-based indexing; so, the legal indexing of A is: $A[1], A[2], \dots, A[n]$.

SOLUTION NOTES: I've annotated the pseudocode below with notes and put my conclusions on worst-case running time below. Note: any line of code that is executed takes $\Omega(1)$ time (i.e., at least some constant amount of time to process); so, anything I annotate as $O(1)$ is in $\Theta(1)$.

Finding the maximum in a list:

Let `max = -infinity` $O(1)$

```

For each element a in A:      n iterations
  If max < a:                  0(1)  could make this true often, but with..
    Set max to a              0(1)  no asymp diff to loop body runtime anyway
Return max                    0(1)

```

SOLUTION: n iterations of a loop whose body takes constant time to run—regardless of the conditional contained inside—plus some constant-time setup and teardown. Overall, this is $\Theta(n)$.

"Median-of-three" computation:

```

Let first = A[1]              0(1)
Let last = A[n]               0(1)
Let middle = A[floor(n/2)]    0(1)

If first <= middle And middle <= last:  0(1)
  return middle                    0(1)
Else If middle <= first And first <= last: 0(1)
  return first                      0(1)
Else:                               0(1)
  return last                        0(1)

```

SOLUTION: Nothing exciting happening here. Every single operation is $O(1)$; so, it doesn't matter what happens in the conditional, the whole thing takes $\Theta(1)$ time.

Counting inversions:

```

Let inversions = 0           0(1)
For each index i from 1 to n:  n iterations
  For each index j from (i+1) to n:  n-i iterations
    If a[i] > a[j]:                0(1)
      Increment inversions          0(1)
Return inversions            0(1)

```

SOLUTION: The inner loop takes constant time per iteration. I've seen many times the pattern of $n - i$ iterations of an inner loop within an outer loop that lets i range over n values. So, I might jump to saying this is $\Theta(n^2)$, but let's get there step by step.

We can express the runtime as a sum:

$$\begin{aligned}
 \sum_{i=1}^n \sum_{j=i+1}^n 1 &= \sum_{i=1}^n (n - i) \\
 &= \left(\sum_{i=1}^n n \right) - \left(\sum_{i=1}^n i \right) \\
 &= n^2 - \frac{n(n+1)}{2} \\
 &= n^2 - \frac{n^2 + n}{2} \\
 &= n^2/2 - n/2 \\
 &\in \Theta(n^2)
 \end{aligned}$$

It's interesting to note also that in the worst case, where the array is in reverse-sorted order, we increment **inversions** $\Theta(n^2)$ times as well, which means there can be $\Theta(n^2)$ inversions in an array of length n . We'll end up coming back to that.

Repeated division:

```

Let count = 0           0(1)
While n > 0:           some number of iterations
    count = count + 1   0(1)
    n = floor(n/2)      0(1)
Return count           0(1)

```

SOLUTION: This loop isn't just a counting loop. It divides n by two each time. How many times can we divide n by 2 before it reaches 0? Well... an infinite number. Fortunately, we're actually taking the floor of $n/2$, and so we will reach 0.

If you don't already see where this is going, you might make a table of values of n and number of iterations to guide you toward it:

Initial value of n	Number of loop iterations
0	0
1	1
2	2
3	2
4	3
5	3
6	3
7	3
8	4

At some point, you'll notice that powers of 2 matter and realize that the right side is roughly the log of the left. Specifically, the number of iterations is $\lfloor \lg n \rfloor + 1$ (but just 0 when $n = 0$).

So, we have a logarithmic number of iterations; each iteration takes constant time; plus constant setup and finish time, for: $\Theta(\lg n)$ runtime.

3 Progress Measures for While Loops

Assume that `FindNeighboringInversion(A)` consumes an array A and returns an index i such that $A[i] > A[i+1]$ or returns -1 if no such inversion exists. Let's work out a bound on the number of iterations of the loop below in terms of n , the length of the array A .

```

Let i = FindNeighboringInversion(A)
While i >= 0:
    Swap A[i] and A[i+1]
    Set i to FindNeighboringInversion(A)

```

1. **Give and work through two small inputs** that will be useful for studying the algorithm. (What is "useful"? Try to find one that is simply common/representative and one that really stresses the algorithm.)

SOLUTION: Let's try this as a "representative" input [8, 2, 3, 9, 7]. And, how about this to stress the algorithm, because it has the maximum possible number of inversions (each pair of elements is inverted with respect to each other) [4, 3, 2, 1]?

The algorithm is "non-deterministic" in the sense that it doesn't specify which inversion to select at each step. I'll try the leftmost and think back to this decision as I solve future parts, in case it's important.

Then, for the first example:

```

[8, 2, 3, 9, 7]
[2, 8, 3, 9, 7]

```

[2, 3, 8, 9, 7]
[2, 3, 8, 7, 9]
[2, 3, 7, 8, 9]

That happens to be the same number of steps as the number of inversions (4) in the input.

For the second:

[4, 3, 2, 1]
[3, 4, 2, 1]
[3, 2, 4, 1]
[2, 3, 4, 1]
[2, 3, 1, 4]
[2, 1, 3, 4]
[1, 2, 3, 4]

Again, that happens to be the same number of steps as the number of inversions (6)! (Cool, but perhaps totally irrelevant.)

2. **Define an inversion** (not just a neighboring one), and **prove that if an inversion exists at all, a neighboring inversion exists**.

SOLUTION: Well, here's the definition of a neighboring inversion "an index i such that $A[i] > A[i+1]$ ". The neighboring part is because we compare index i to index $i+1$. So, let's let it be any two indexes, but insist their elements are out of order with respect to the indexes: "indexes i and j where $i < j$ but $A[i] > A[j]$ ".

Let's double-check that that makes sense. It means one element comes before another in the array and yet is larger, which seems reasonable. It applies to the 8 and 3 in our first small example but not to the 8 and 9. It applies to every pair of elements in our second example.

We can sketch a quick proof that if there is an inversion, there's a neighboring inversion. Assume an array A contains an inversion. Then, for some indices $i < j$, $A[i] > A[j]$. There has to be a neighboring inversion in this range or each element will be at least as large as the last from i up to j and so $A[j]$ will be at least as large as $A[i]$, which isn't true.

Here's that argument laid out in more detail as an inductive proof. If $A[i] > A[i+1]$ (including if $j = i+1$), then this is a neighboring inversion (base case). Otherwise, assume that if there is an inversion in the shorter array $A[i+1..j]$ (which must still have at least two elements), then there is a neighboring inversion in that portion of the array. Since we're not in the base case, we know that $A[i] \leq A[i+1]$ and $A[j] < A[i]$. Therefore, $A[j] < A[i+1]$, which means that there is an inversion in $A[i+1..j]$ between j and $i+1$, and so we know there is a neighboring inversion in that range. QED

3. Give **upper- and lower-bounds on the number of inversions** in A .

SOLUTION: By our definition, an inversion requires two distinct elements of A . There are $\binom{n}{2} = \frac{n(n-1)}{2} = \frac{n^2}{2} - \frac{n}{2}$ pairs. Can they all be inverted? In our second example, we do achieve $\frac{4^2}{2} - \frac{4}{2} = 8 - 2 = 6$ inverted pairs. If we extend our second example to an array of length n in reverse-sorted order, then any pair of elements $A[i]$ and $A[j]$ with $i < j$ will indeed be out-of-order with respect to each other.

So, the upper-bound is $\binom{n}{2} \in O(n^2)$.

An already-sorted array will have no inversions. So, the lower bound is 0.

4. Give a "measure of progress" for each iteration of the loop in terms of inversions. (I.e., how can we measure that we're making progress toward terminating the loop?)

SOLUTION: In our examples above, we took exactly as many steps as there were inversions. That's because each swap patched an inversion without "disturbing" any others (because all elements besides the two swapped were still in the same order with respect to each other).¹ So, one measure of progress would be the number of remaining inversions, which goes down by 1 at each iteration.

5. Give an upper-bound on the number of iterations the loop could take.

SOLUTION: Since we reduce the number of inversions by 1 in each iteration and the number of inversions is at most $\binom{n}{2} \in O(n^2)$, then that is also an upper-bound on the number of iterations of the loop.

6. Prove that this algorithm sorts the array A (i.e., removes all inversions from the array).

SOLUTION: Note that if there's an inversion, then there's a neighboring inversion. In $O(n^2)$ steps, there will be no inversions left, which means no neighboring inversions, either. (Every neighboring inversion is an inversion.) Thus, the loop ends in $O(n^2)$ steps. With no neighboring inversions, there must be no inversions at all (the contrapositive of our theorem above), and so the array is sorted.

So, in $O(n^2)$ steps, the loop has to end (because there are no more inversions and so no more neighboring inversions, which triggers the loops termination condition).

¹Note that this is true regardless of the arbitrary choice we made to address the leftmost neighboring inversion first.