

15-122: Principles of Imperative Computation

Recitation 8

Josh Zimmerman


Big-O

The definition of Big-O has a lot of mathematical symbols in it, and so can be very confusing at first. Let's familiarize ourselves with the formal definition and get an intuition behind what it's saying.

First of all, $O(g(n))$ is a set of functions. The formal definition is:

$f(n) \in O(g(n))$ if and only if there is some $c \in \mathbb{R}^+$ and some $n_0 \in \mathbb{R}$ such that for all $n > n_0$, $f(n) \leq c * g(n)$.

That's a bit formal and possibly confusing, so let's look at the intuition. Here's a diagram, courtesy of Wikipedia. (Note that it uses x and x_0 instead of n and n_0 .)



To the left of n_0 , the functions can do anything. To its right, $c * g(n)$ is always greater than or equal to $f(n)$.

In problems 1 and 2, we play with the definition of big-O.

Something that's very important to note about this diagram is that there are infinitely many functions that are in $O(g(n))$: If $f(n) \in O(g(n))$, then $\frac{1}{2}f(n) \in O(g(n))$ and $\frac{1}{4}f(n) \in O(g(n))$ and $2f(n) \in O(g(n))$. In general, for any constant k , $k * f(n) \in O(g(n))$. In problem 3, we prove this.

Something that will come up often with big-O is the idea of a *tight* bound on the runtime of a function. It's technically correct to say that binary search, which takes around $\log(n)$ steps on an array of length n , is $O(n!)$, since $n! > \log(n)$ for all $n > 0$ but it's not very useful. If we ask for a *tight* bound, we want the closest bound you can give. For binary search, $O(\log(n))$ is a tight bound because no function that grows more slowly than $\log(n)$ provides a correct upper bound for binary search.

Unless we specify otherwise, we want a tight bound.

Selection sort

Sorting is a really useful task that's central to many programs (for instance, spreadsheet programs) as well as many algorithms (for instance, Google PageRank needs to sort pages before presenting them).

One such algorithm is selection sort, which we discussed in lecture. The basic idea behind selection sort is that we find the smallest number in the region of the array that we're considering and switch that into

the section of the array that's sorted. Next we shrink the size of the region of the array we're considering by one and repeat until the region we're considering is empty.

Here's a good visualization of selection sort on a series of poles: <http://www.youtube.com/watch?v=LuanFAXgQEw>. In that video, the light blue square points to the smallest element we've found so far, the purple square points to the element we're comparing with the minimum, and the pink rectangle is the region of the array we still need to consider.

This is not a particularly efficient algorithm: if we have i elements in the region of the array we're considering, we need to look at all i of them, every time.

So, the overall amount of work we'll have to do is $n + (n - 1) + (n - 2) + \dots + 2 + 1$, or:

$$\sum_{i=1}^n i = \frac{n(n+1)}{2}$$

Now that you know big-O notation and have started to learn about time complexity, it should be possible for you to figure out what a tight (big-O) upper bound for the complexity of selection sort is.

```

1 #use "sortutil.c0"
2
3 int min_index(int[] A, int lower, int upper)
4 //@requires 0 <= lower && lower < upper && upper <= \length(A);
5 //@ensures lower <= \result && \result < upper;
6 //@ensures le_seg(A[\result], A, lower, upper);
7 {
8     int m = lower;
9     int min = A[lower];
10    for (int i = lower+1; i < upper; i++)
11        //@loop_invariant lower < i && i <= upper;
12        //@loop_invariant le_seg(min, A, lower, i);
13        //@loop_invariant A[m] == min;
14    {
15        if (A[i] < min) {
16            m = i;
17            min = A[i];
18        }
19    }
20    return m;
21 }
22
23 void sort(int[] A, int lower, int upper)
24 //@requires 0 <= lower && lower <= upper && upper <= \length(A);
25 //@ensures is_sorted(A, lower, upper);
26 {
27    for (int i = lower; i < upper; i++)
28        //@loop_invariant lower <= i && i <= upper;
29        //@loop_invariant is_sorted(A, lower, i);
30        //@loop_invariant le_segs(A, lower, i, i, upper);
31    {
32        int m = min_index(A, i, upper);
33        //@assert le_seg(A[m], A, i, upper);
34        swap(A, i, m);
35    }
36    return;
37 }

```

Quicksort

Quicksort is a fast ($O(n \log n)$ on average, $O(n^2)$ worst case) sorting algorithm. The general idea behind quicksort is that we start with an array we need to sort, and then select a *pivot* index. (There are many different methods of picking a pivot index: some strategies always pick index 0 or the largest possible index in the array, some choose a randomly generated index, and some randomly generate three indices, look at the values at those indices, and pick the index of the median value. The different strategies affect the performance of quicksort in different ways.)

After we pick a pivot index, we rearrange the array so that everything to the left of the pivot index in the array is less than or equal to $A[\text{pivot_index}]$ and everything on the right of the pivot index in the array is greater than $A[\text{pivot_index}]$. Then, we call quicksort recursively on the portion of the array that is less than or equal to the pivot and on the portion that is greater than the pivot. If the list is empty or has one element, we just return the list since it's already sorted.

Quicksort can be hard to understand at first, so let's look at an example.

Start with the array [4, 2, 1, 0]. We randomly select 1 as our pivot index. We switch $A[1]$ and $A[n - 1]$, which gives us [4, 0, 1, 2].

We split the array so everything on the left is less than or equal to the pivot. After we've done that, we have [1, 0, 4, 2]. Then, we recursively sort the two portions of the array: [1, 0] and [4].

First, we sort [1, 0]. We must choose 0 as our pivot index, since otherwise the switching won't work correctly — we switch the element at the pivot index with the last element of the array, and if the pivot index is the index of the last element in the array we can't do that switch.

Our switch then results in the array [0, 1]. This already has all of the smallest elements on the left, so we can continue.

Then, we recursively sort the array [0]. It has only one element, so it's sorted already.

Next, we recursively sort the array [] (everything in [0, 1] that's larger than 1). The empty array is already sorted.

Now, we know that [0, 1] is sorted, so we can sort [4].

[4] is already sorted, so we know our whole array, [0, 1, 2, 4] is now sorted.

Now, let's look at the code for quicksort and see what we can say about its correctness:

```
1 #use <rand>
2 #use "sortutil.c0"
3
4 int abs(int x)
5 {
6     if (x < 0)
7         return -x;
8     return x;
9 }
10
11 int partition(int[] A, int lower, int pivot_index, int upper)
12 //@requires 0 <= lower && lower <= pivot_index;
13 //@requires pivot_index < upper && upper <= \length(A);
14 //@ensures lower <= \result && \result < upper;
15 //@ensures ge_seg(A[\result], A, lower, \result);
16 //@ensures le_seg(A[\result], A, \result+1, upper);
17 {
18     int pivot = A[pivot_index];
```

```

19 swap(A, pivot_index, upper-1);
20
21 int left = lower;
22 int right = upper-2;
23 while (left <= right)
24 // @loop_invariant lower <= left && left <= right+1 && right+1 < upper;
25 // @loop_invariant ge_seg(pivot, A, lower, left);
26 // @loop_invariant le_seg(pivot, A, right+1, upper-1);
27 {
28     if (A[left] <= pivot) {
29         left++;
30     } else { // @assert A[left] > pivot;
31         swap(A, left, right);
32         right--;
33     }
34 }
35 // @assert left == right+1;
36 // @assert A[upper-1] == pivot;
37
38 swap(A, left, upper-1);
39 return left;
40 }
41
42 void qsort(int[] A, int lower, int upper, rand_t gen)
43 // @requires 0 <= lower && lower <= upper && upper <= \length(A);
44 // @ensures is_sorted(A, lower, upper);
45 {
46     if (upper - lower <= 1) return; // already sorted
47
48     int pivot_index = lower + abs(rand(gen) % (upper-lower));
49     /* pivot_index = upper-1 or pivot_index = 0 gives O(n^2) for sorted array */
50     /* pivot_index = lower+(upper-lower)/2 efficiency depends on input distribution */
51     /* better yet would be: median of 3 random elements */
52
53     int mid = partition(A, lower, pivot_index, upper);
54     qsort(A, lower, mid, gen);
55     qsort(A, mid+1, upper, gen);
56     return;
57 }

```

Practice!

1. Rank these big-O sets from left to right such that every big-O is a subset of everything to the right of it. (For instance, $O(n)$ goes farther to the left than $O(n!)$ because $O(n) \subset O(n!)$.) If two sets are the same, put them on top of each other.

$O(n!)$ $O(n)$ $O(4)$ $O(n \log(n))$ $O(4n + 3)$ $O(n^2 + 20000n + 3)$ $O(1)$ $O(n^2)$ $O(2^n)$
 $O(\log(n))$ $O(\log^2(n))$ $O(\log(\log(n)))$

2. Using the formal definition of big-O, prove that $n^3 + 300n^2 \in O(n^3)$.

3. Using the formal definition of big-O, prove that if $f(n) \in O(g(n))$, then $k * f(n) \in O(g(n))$ for $k > 0$.

One interesting consequence of this is that $O(\log_i(n)) = O(\log_j(n))$ for all i and j (as long as they're

both greater than 1), because of the change of base formula:

$$\log_i(n) = \frac{\log_j(n)}{\log_j(i)}$$

But $\frac{1}{\log(i)}$ is just a constant! So, it doesn't matter what base we use for logarithms in big-O notation.

4. Let's prove that `partition` is correct. (Note: since `qsort` is a recursive function, it requires a slightly different proof format and we won't discuss it in recitation. Instead of loop invariants, we'd have to base our proof entirely on preconditions and postconditions of `qsort` and `partition`. Recursive programs generally have relatively clean proofs by induction that they're correct. If you don't know what induction is yet, that's fine.)