

15-122: Principles of Imperative Computation

Recitation 8

Josh Zimmerman

There is a good visualization of both Selection Sort and Quicksort at <http://www.cs.usfca.edu/~galles/visualization/flash.html>.

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.

Example of Quicksort

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

We split the array so everything on the left is less than or equal to the pivot. After we've done that, we have $[2, 0, 1, 4]$. We then swap the pivot, 2, back to the place where we now know it belongs in the final array – index 2. Therefore we have $[1, 0, 2, 4]$. 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, switching index 0 with index 0 keeps the array exactly the same: Our switch then results in the array $[1, 0]$. Then, we can swap the pivot, 1, to the place where it belongs in the final array – index 1.

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 "sortutil.c0"
2 #use <rand>
3
4 int partition(int[] A, int lower, int pivot_index, int upper)
5 //@requires 0 <= lower && lower <= pivot_index;
6 //@requires pivot_index < upper && upper <= \length(A);
7 //@ensures lower <= \result && \result < upper;
8 //@ensures gt_seg(A[\result], A, lower, \result);
9 //@ensures le_seg(A[\result], A, \result, upper);
10 {
11 // Hold the pivot element off to the left at "lower"
12 int pivot = A[pivot_index];
13 swap(A, lower, pivot_index);
14
15 int left = lower+1; // Inclusive lower bound (lower+1, pivot's at lower)
16 int right = upper; // Exclusive upper bound
17
18 while (left < right)
19 //@loop_invariant lower+1 <= left && left <= right && right <= upper;
20 //@loop_invariant gt_seg(pivot, A, lower+1, left); // Not lower!
21 //@loop_invariant le_seg(pivot, A, right, upper);
22 {
23 if (A[left] < pivot) {
24 left++;
25 } else {
26 //@assert A[left] >= pivot;
27 swap(A, left, right-1); // right-1 because of exclusive upper bound
28 right--;
29 }
30 }
31 //@assert left == right;
32
33 swap(A, lower, left-1);
34 return left-1;
35 }
36
37 void sort(int[] A, int lower, int upper)
38 //@requires 0 <= lower && lower <= upper && upper <= \length(A);
39 //@ensures is_sorted(A, lower, upper);
40 {
41 if (upper - lower <= 1) return;
42 int pivot_index = lower + rand(init_rand(234234)) % (upper - lower); // Pivot at midpoint
43
44 int new_pivot_index = partition(A, lower, pivot_index, upper);
45 sort(A, lower, new_pivot_index);
46 //@assert is_sorted(A, lower, new_pivot_index + 1);
47 sort(A, new_pivot_index + 1, upper);
48 }

```

Checkpoint 0

What would go wrong if the partition function ignored `pivot_index` and picked a new pivot?

Checkpoint 1

Why is swapping the pivot with $\text{left}-1$ the right thing? Why is left wrong? Why is $\text{left}-1$ safe?

Checkpoint 2

Prove that the that partition function is correct.

Loop invariants hold initially

Loop invariants are preserved

Negation of loop guard and loop invariants imply postcondition

Termination

Checkpoint 3

Could you change `partition` (code and/or loop invariants) in order to justify one of the postconditions on line 5 or 6 being `gt_seg` or `lt_seg`?

Checkpoint 4

Using the `rand` library in C0, modify this code to select a random pivot.