

1 Sorted Runtimes

We want to sort an array of N **unique** numbers in ascending order. Determine the best case and worst case runtimes of the following sorts:

- (a) Once the runs in merge sort are of $size \leq N/100$, we perform insertion sort on them.

Best Case: $\Theta(\quad)$, Worst Case: $\Theta(\quad)$

Solution:

Best Case: $\Theta(N)$, Worst Case: $\Theta(N^2)$

Once we have 100 runs of size $N/100$, insertion sort will take best case $\Theta(N)$ and worst case $\Theta(N^2)$ time. The constant number of linear time merging operations don't add to the runtime.

- (b) We can only swap adjacent elements in selection sort.

Best Case: $\Theta(\quad)$, Worst Case: $\Theta(\quad)$

Solution:

Best Case: $\Theta(N^2)$, Worst Case: $\Theta(N^2)$

The best case and worst case don't change since swapping at most doubles the work each iteration, which produces the same asymptotic runtime as normal selection sort.

- (c) We use a linear time median finding algorithm to select the pivot in quicksort.

Best Case: $\Theta(\quad)$, Worst Case: $\Theta(\quad)$

Solution:

Best Case: $\Theta(N \log(N))$, Worst Case: $\Theta(N \log(N))$

Doing an extra N work each iteration of quicksort doesn't asymptotically change the best case runtime, but it improves the worst case runtime.

- (d) We implement heapsort with a min-heap instead of a max-heap. You may modify heapsort but must maintain constant space complexity.

Best Case: $\Theta(\quad)$, Worst Case: $\Theta(\quad)$

Solution:

Best Case: $\Theta(N \log(N))$, Worst Case: $\Theta(N \log(N))$

While a max-heap is better, we can make do with a min-heap by placing the smallest element at the right end of the list until the list is sorted in **descending order**. Once the list is in descending order, it can be sorted in ascending order with a simple linear time pass.

(e) We run an optimal sorting algorithm of our choosing knowing:

- There are at most N inversions

Best Case: $\Theta(N)$, Worst Case: $\Theta(N)$

Solution:

Best Case: $\Theta(N)$, Worst Case: $\Theta(N)$

Recall that insertion sort takes $\Theta(N + K)$ time, where K is the number of inversions. If K is at most N , then, insertion sort has the best and worst case runtime of $\Theta(N)$. Here is an explanation for why no sorting algorithm can surpass this. Notice for our algorithm to terminate we *either* need to address every inversion or look at every element. Since there are at most N inversions, knowing that we have addressed every inversion would take us at least $\Theta(N)$ time. Looking at every element in the list would also take us $\Theta(N)$ time. In either case, we see the runtime of any sorting algorithm cannot be faster than $\Theta(N)$.

- There is exactly 1 inversion

Best Case: $\Theta(1)$, Worst Case: $\Theta(N)$

Solution:

Best Case: $\Theta(1)$, Worst Case: $\Theta(N)$

The inversion may be the first two elements, in which case constant time is needed. Or, it may involve elements at the end, in which case N time is needed. It can be proven quite simply that no sorting algorithm can achieve a better runtime than above for the best and worst case.

- There are exactly $(N^2 - N)/2$ inversions

Best Case: $\Theta(N)$, Worst Case: $\Theta(N)$

Solution:

Best Case: $\Theta(N)$, Worst Case: $\Theta(N)$

If a list has $N(N - 1)/2$ inversions, it means it is sorted in descending order! So, it can be sorted in ascending order with a simple linear time pass. We know that reversing any array is a linear time operation, so the optimal runtime of any sorting algorithm is $\Theta(N)$.

2 Shuffled Exams

For this problem, we will be working with `Exam` and `Student` objects, both of which have only one attribute: `sid`, which is a number like any student ID.

Gradescope thought it was ready for the midterm. It had meticulously created two arrays, one of `Exams` and the other of `Students`, and ordered both on `sid` such that the i th `Exam` in the `Exams` array has the same `sid` as the i th `Student` in the `Students` array. Note the arrays are not necessarily sorted by `sid`. However, Gradescope crashed, and the `Students` array was shuffled, but the `Exams` array somehow remained untouched. Time is precious, so you must design a $O(N)$ time algorithm to reorder the `Students` array appropriately **without** changing the `Exams` array! For partial credit, you may reorder **both** the `Students` and `Exams` arrays such that i th `Exam` in the `Exams` array has the same `sid` as the i th `Student` in the `Students` array.

Hint: Use radix sort.

Solution:

Let's begin by creating an `ExamWrapper` class that contains two attributes — an `Exam` instance and the `index` of the corresponding `Exam` in the `Exams` array. Next, for each `Exam`, create the corresponding `ExamWrapper` instance.

Run radix sort on the `ExamWrappers`, sorting them on the `sid` of the `Exam` instances. Similarly run radix sort on the list of `Students`, sorting them on `sid` as well. Note that both iterations of radix sort take linear time since the `sid` is of fixed length and of base 10.

At this point in the algorithm, we have satisfied the partial credit criteria, but we still need move the i th `Student` to its proper place relative to the original `Exams` array. To achieve this, for the i th `Student`, we will access the i th `ExamWrapper`, and set the `index` of the i th `Student` as the `ExamWrapper`'s `index` attribute.

3 Bears and Beds

The hot new Cal startup AirBearsnBeds has hired you to create an algorithm to help them place their customers in the best possible homes to improve their experience. They are currently in their alpha stage so their only customers (for now) are bears. Now, a little known fact about bears is that they are very, very picky about their bed sizes: they do not like their beds too big or too little - they like them just right. Bears are also sensitive creatures who don't like being compared to other bears, but they are perfectly fine with trying out beds.

The Problem:

Given a list of Bears with unique but unknown sizes and a list of Beds with corresponding but also unknown sizes (not necessarily in the same order), return a list of Bears and a list of Beds such that that the i th Bear in your returned list of Bears is the same size as the i th Bed in your returned list of Beds. Bears can only be compared to Beds and we can get feedback on if the Bed is too large, too small, or just right. In addition, Beds can only be compared to Bears and we can get feedback if the Bear is too large for it, too small for it, or just right for it.

The Constraints:

Your algorithm should run in $O(N \log N)$ time on average. It may be helpful to figure out the naive $O(N^2)$ solution first and then work from there.

Solution:

Our solution will modify quicksort. Let's begin by choosing a pivot from the Bears list. To avoid quicksort's worst case behavior on a sorted array, we will choose a random Bear as the pivot. Next we will partition the Beds into three groups — those less than, equal to, and greater than the pivot Bear. Next, we will select a pivot from the Bears list. This is very important — our pivot Bed will be the Bed that is equal to the pivot Bear. Given that the Beds and Bears have unique sizes, we know that **exactly** one Bed will be equal to the pivot Bear. Next we will partition the Beds into three groups — those less than, equal to, and greater than the pivot Bed.

Next, we will "match" the pivot Bear with the pivot Bed by adding them to the Bears and Beds lists at the same index, which is as easy as just adding to the end. Finally, in the same fashion as quicksort, we will have two recursive calls. The first recursive call will contain the Beds and Bears that are **less** than their respective pivots. The second recursive call will contain the Beds and Bears that are **greater** than their respective pivots.