



UNIVERSITY OF GOTHENBURG

Data structures

Sorting

Dr. Alex Gerdes DIT961 - VT 2018

Searching



 Suppose I give you an array, and ask you to find if a particular value is in it, say 4 (or 2)

- The only way is to look at each element in turn
- Quiz: what is the complexity?
 - O(n)
- This is called *linear search*
- You might have to look at every element before you find the right one



• But what if the array is *sorted*?

• Then we can use *binary search*!



- Suppose we want to look for 4
- We start by looking at the element half way along the array, which happens to be 3





- 3 is less than 4
- Since the array is sorted, we know that 4 must come after 3
- We can ignore everything before 3





- Now we repeat the process
- We look at the element half way along what's left of the array, this happens to be 7





- 7 is greater than 4
- Since the array is sorted, we know that 4 must come before 7
- We can safely ignore everything after 7





- We repeat the process
- We look half way along the array again
- We find 4!



BINÄRY SEARCH

idea-instructions.com/binary-search/ v1.0, CC by-nc-sa 4.0









Performance of binary search



- Binary search repeatedly chops the array in half
 - If we double the size of the array, we need to look at one more array element
 - With an array of size 2^n , after *n* tries, we are down to 1 element
 - On an array of size *n* takes **O(log n)** time!
- On an array of a billion elements, need to search 30 elements (compared to a billion tries for linear search!)



- Keep two indices lo and hi, they represent the part of the array to be searched
- Let mid = (lo + hi) / 2 and look at a[mid], then either set lo = mid + 1 or hi = mid - 1, depending on the value of a[mid]



CHALMERS

- Keep two indices lo and hi, they represent the part of the array to be searched
- Let mid = (lo + hi) / 2 and look at a[mid], then either set lo = mid + 1 or hi = mid 1.
 depending on the value of a[mid]



Sorting



Zillions of sorting algorithms (bubblesort, insertion sort, selection sort, quicksort, heapsort, mergesort, shell sort, counting sort, ...)





- Why is sorting important? Because sorted data is much easier to deal with!
 - Searching use binary instead of linear search

Sorting

- Finding duplicates takes linear instead of quadratic time etc.
- Most sorting algorithms are based on *comparisons*
 - Compare elements is one bigger than the other? If not, do something about it!
 - Advantage: they can work on all sorts of data
 - Disadvantage: specialised algorithms for e.g. sorting lists of integers can be faster

Bubblesort



- Go through the array, comparing adjacent elements
 - If we find two that are in the wrong order, swap them
- Once we reach the end of the array, go back and start again!



Compare a[0] and a[1]:





Compare a[1] and a[2]:





Compare a[2] and a[3]:





Compare a[3] and a[4]:





Back to the beginning!





Compare a[1] and a[2]:





Compare a[2] and a[3]:





Compare a[3] and a[4]:





Back to the beginning!



Bubblesort



- How do we know when to stop going back to the beginning?
 - When the array is sorted
- How many loops until that happens?
 - Each time we loop through the array, at least one more element ends up in the right place: the biggest element that was in the wrong place before
- So repeat as many times as there are elements in the input array

Insertion sort



Imagine someone is dealing you cards. Whenever you get a new card you put it into the right place in your hand:



This is the idea of *insertion* sort.





Start by "picking up" the 5:

5





Then insert 3 into the right place:







Then the 9:







Then the 2:







Finally the 8:





- Insertion sort does *n* insertions for an array of size *n*
- Does this mean it is O(n)?
 - No! An insertion is not constant time.
- To insert into a sorted array, you must move all the elements up one, which is O(n)
- Thus total is $O(n^2)$



- This version of insertion sort needs to make a new array to hold the result
- An *in-place* sorting algorithm is one that doesn't need to make temporary arrays
 - Has the potential to be more efficient
- Let's make an in-place insertion sort!
- Basic idea: loop through the array, and insert each element into the part which is already sorted





• The first element of the array is sorted:







• Insert the 3 into the correct place:





• Insert the 9 into the correct place:





• Insert the 2 into the correct place:







• Insert the 8 into the correct place:



 One way to do it: repeatedly swap the element with its neighbour on the left, until it's in the right position



while n > 0 and array[n] > array[n-1] swap array[n] and array[n-1] n = n-1







• An improvement: instead of swapping, move elements upwards to make a "hole" where we put the new value



In-place insertion sort







- An aside: we have the invariant that array[0..i) is sorted
 - An invariant is something that holds whenever the loop body starts to run
 - Initially, i = 1 and array[0..1) is sorted
 - As the loop runs, more and more of the array becomes sorted
 - When the loop finishes, i = n, so array[0...) is sorted
 the whole array!

Selection sort



- Find the smallest element of the array, and delete it
- Find the smallest remaining element, and delete it
- And so on...
- Quiz: complexity?
- Finding the smallest element is O(n), so total complexity is $O(n^2)$





The smallest element is 2:

2

We also delete 2 from the input array





Now smallest element is 3:



We delete 3 from the input array





Now smallest element is 5:

We delete from the input array (... and so on)



- Instead of deleting the smallest element, swap it with the first element!
- The next time round, ignore the first element of the array: we know it's the smallest one
- Instead, find the smallest element of the rest of the array, and swap it with the second element





The smallest element is 2:







The smallest element in the rest of the array is 3:





The smallest element in the rest of the array is 5:





The smallest element in the rest of the array is 8:



for i = 0 to a.length-1 find the smallest element in a[i .. a.length) swap it with a[i]



- All the algorithms so far are $O(n^2)$ in the worst case
- One of them is O(n) in the best case (a sorted array) which?
 - Answer: insertion sort
 - This makes insertion sort the best of our three algorithms it's actually a fast sorting algorithm in general for small lists
 - The other two are bad, but selection sort is the basis for a better algorithm, heapsort