CSE373: Data Structure \& Algorithms Lecture 18: Comparison Sorting

Kevin Quinn

Fall 2015

## Introduction to Sorting

- Stacks, queues, priority queues, and dictionaries all focused on providing one element at a time
- But often we know we want "all the things" in some order
- Humans can sort, but computers can sort fast
- Very common to need data sorted somehow
- Alphabetical list of people
- List of countries ordered by population
- Search engine results by relevance
- ...
- Algorithms have different asymptotic and constant-factor tradeoffs
- No single "best" sort for all scenarios
- Knowing one way to sort just isn't enough


## More Reasons to Sort

General technique in computing:
Preprocess data to make subsequent operations faster

Example: Sort the data so that you can

- Find the $\mathbf{k}^{\text {th }}$ largest in constant time for any $\mathbf{k}$
- Perform binary search to find elements in logarithmic time

Whether the performance of the preprocessing matters depends on

- How often the data will change (and how much it will change)
- How much data there is


## The main problem, stated carefully

For now, assume we have $n$ comparable elements in an array and we want to rearrange them to be in increasing order

## Input:

- An array A of data records
- A key value in each data record
- A comparison function (consistent and total)

Effect:

- Reorganize the elements of $\mathbf{A}$ such that for any $i$ and $j$, if $\mathbf{i}<\boldsymbol{j}$ then $\mathrm{A}[\mathrm{i}] \leq \mathrm{A}[\mathrm{j}]$
- (Also, A must have exactly the same data it started with)
- Could also sort in reverse order, of course

An algorithm doing this is a comparison sort

## Variations on the Basic Problem

1. Maybe elements are in a linked list (could convert to array and back in linear time, but some algorithms needn't do so)
2. Maybe ties need to be resolved by "original array position"

- Sorts that do this naturally are called stable sorts
- Others could tag each item with its original position and adjust comparisons accordingly (non-trivial constant factors)

3. Maybe we must not use more than $O(1)$ "auxiliary space"

- Sorts meeting this requirement are called in-place sorts

4. Maybe we can do more with elements than just compare

- Sometimes leads to faster algorithms

5. Maybe we have too much data to fit in memory

- Use an "external sorting" algorithm


## Sorting: The Big Picture

Surprising amount of neat stuff to say about sorting:


## Insertion Sort

- Idea: At step $\mathbf{k}$, put the $\mathbf{k}^{\text {th }}$ element in the correct position among the first $\mathbf{k}$ elements
- Alternate way of saying this:
- Sort first two elements
- Now insert 3rd element in order
- Now insert $4^{\text {th }}$ element in order
- ...
- "Loop invariant": when loop index is $i$, first $i$ elements are sorted
- Time?

Best-case $\qquad$ Worst-case $\qquad$ "Average" case $\qquad$

## Insertion Sort

- Idea: At step $\mathbf{k}$, put the $\mathbf{k}^{\text {th }}$ element in the correct position among the first $\mathbf{k}$ elements
- Alternate way of saying this:
- Sort first two elements
- Now insert 3rd element in order
- Now insert $4^{\text {th }}$ element in order
- ...
- "Loop invariant": when loop index is i, first i elements are sorted
- Time?

| Best-case $O(n)$ | Worst-case $O\left(n^{2}\right)$ | "Average" case $O\left(n^{2}\right)$ |
| :---: | :---: | :---: | :---: | :---: |
| start sorted | start reverse sorted | (see text) |

## Selection sort

- Idea: At step $\mathbf{k}$, find the smallest element among the not-yetsorted elements and put it at position $k$
- Alternate way of saying this:
- Find smallest element, put it $1^{\text {st }}$
- Find next smallest element, put it $2^{\text {nd }}$
- Find next smallest element, put it $3^{\text {rd }}$
- ...
- "Loop invariant": when loop index is $i$, first $i$ elements are the $i$ smallest elements in sorted order
- Time?

Best-case $\qquad$ Worst-case $\qquad$ "Average" case $\qquad$

## Selection sort

- Idea: At step $\mathbf{k}$, find the smallest element among the not-yetsorted elements and put it at position $k$
- Alternate way of saying this:
- Find smallest element, put it $1^{\text {st }}$
- Find next smallest element, put it $2^{\text {nd }}$
- Find next smallest element, put it $3^{\text {rd }}$
- ...
- "Loop invariant": when loop index is $i$, first $i$ elements are the $i$ smallest elements in sorted order
- Time?

$$
\begin{aligned}
& \text { Best-case } O\left(n^{2}\right) \text { Worst-case } O\left(n^{2}\right) \text { "Average" case } O\left(n^{2}\right) \\
& \text { Always } T(1)=1 \text { and } T(n)=n+T(n-1)
\end{aligned}
$$

## Mystery

This is one implementation of which sorting algorithm (for ints)?

```
void mystery(int[] arr) {
    for(int i = 1; i < arr.length; i++) {
        int tmp = arr[i];
        int j;
        for(j=i; j > 0 && tmp < arr[j-1]; j--)
                arr[j] = arr[j-1];
        arr[j] = tmp;
    }
}
```

Note: Like with heaps, "moving the hole" is faster than unnecessary swapping (constant-factor issue)

## Insertion Sort vs. Selection Sort

- Different algorithms
- Solve the same problem
- Have the same worst-case and average-case asymptotic complexity
- Insertion-sort has better best-case complexity; preferable when input is "mostly sorted"
- Other algorithms are more efficient for non-small arrays that are not already almost sorted
- Insertion sort may do well on small arrays


## Aside: We Will Not Cover Bubble Sort

- It is not, in my opinion, what a "normal person" would think of
- It doesn't have good asymptotic complexity: $O\left(n^{2}\right)$
- It's not particularly efficient with respect to common factors

Basically, almost everything it is good at some other algorithm is at least as good at

- Perhaps people teach it just because someone taught it to them?

Fun, short, optional read:
Bubble Sort: An Archaeological Algorithmic Analysis, Owen Astrachan, SIGCSE 2003
http://www.cs.duke.edu/~ola/bubble/bubble.pdf

## The Big Picture

Surprising amount of juicy computer science: 2-3 lectures...


## Heap sort

- Sorting with a heap is easy:
- insert each arr[i], or better yet use buildHeap
- for (i=0; i < arr.length; i++) arr[i] = deleteMin();
- Worst-case running time: $O(n \log n)$
- We have the array-to-sort and the heap
- So this is not an in-place sort
- There's a trick to make it in-place...


## In-place heap sort

## But this reverse sorts how would you fix that?

- Treat the initial array as a heap (via buildHeap)
- When you delete the $i^{\text {th }}$ element, put it at arr[n-i]
- That array location isn't needed for the heap anymore!



## "AVL sort"

- We can also use a balanced tree to:
- insert each element: total time $O(n \log n)$
- Repeatedly deleteMin: total time $O(n \log n)$
- Better: in-order traversal $O(n)$, but still $O(n \log n)$ overall
- But this cannot be made in-place and has worse constant factors than heap sort
- both are $O(n \log n)$ in worst, best, and average case
- neither parallelizes well
- heap sort is better


## "Hash sort"???

- Don't even think about trying to sort with a hash table!
- Finding min item in a hashtable is $O(\mathrm{n})$, so this would be a slower, more complicated selection sort


## Divide and conquer

Very important technique in algorithm design

1. Divide problem into smaller parts
2. Independently solve the simpler parts

- Think recursion
- Or potential parallelism

3. Combine solution of parts to produce overall solution
(The name "divide and conquer" is rather clever.)

## Divide-and-Conquer Sorting

Two great sorting methods are fundamentally divide-and-conquer

1. Mergesort: Sort the left half of the elements (recursively) Sort the right half of the elements (recursively)
Merge the two sorted halves into a sorted whole
2. Quicksort: Pick a "pivot" element Divide elements into less-than pivot and greater-than pivot
Sort the two divisions (recursively on each) Answer is sorted-less-than then pivot then sorted-greater-than

## Mergesort



- To sort array from position lo to position hi:
- If range is 1 element long, it is already sorted! (Base case)
- Else:
- Sort from lo to (hi+lo)/2
- Sort from (hi+lo) /2 to hi
- Merge the two halves together
- Merging takes two sorted parts and sorts everything
- $O(n)$ but requires auxiliary space...


## Example, Focus on Merging

 Start with:| 8 | 2 | 9 | 4 | 5 | 3 | 1 | 6 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

 Start with:| 8 | 2 | 9 | 4 | 5 | 3 | 1 | 6 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge, copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge,
copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge, copy back to original array)

## Example, focus on merging

Start with:


After recursion: (not magic ©)


Merge:
Use 3 "fingers" and 1 more array

(After merge, copy back to original array)


## Example, Showing Recursion



## Some details: saving a little time

- What if the final steps of our merge looked like this:

- Wasteful to copy to the auxiliary array just to copy back...


## Some details: saving a little time

- If left-side finishes first, just stop the merge and copy back:

- If right-side finishes first, copy dregs into right then copy back



## Some details: Saving Space and Copying

Simplest / Worst:
Use a new auxiliary array of size (hi-lo) for every merge

Better:
Use a new auxiliary array of size n for every merging stage

Better:
Reuse same auxiliary array of size n for every merging stage

Best (but a little tricky):
Don't copy back - at $2^{\text {nd }}, 4^{\text {th }}, 6^{\text {th }}, \ldots$ merging stages, use the original array as the auxiliary array and vice-versa

- Need one copy at end if number of stages is odd


## Swapping Original / Auxiliary Array ("best")

- First recurse down to lists of size 1
- As we return from the recursion, swap between arrays

(Arguably easier to code up without recursion at all)


## Linked lists and big data

We defined sorting over an array, but sometimes you want to sort linked lists

One approach:

- Convert to array: $O(n)$
- Sort: O(n log $n$ )
- Convert back to list: $O(n)$

Or: merge sort works very nicely on linked lists directly

- Heapsort and quicksort do not
- Insertion sort and selection sort do but they're slower

Merge sort is also the sort of choice for external sorting

- Linear merges minimize disk accesses
- And can leverage multiple disks to get streaming accesses


## Analysis

Having defined an algorithm and argued it is correct, we should analyze its running time and space:

To sort $n$ elements, we:

- Return immediately if $n=1$
- Else do 2 subproblems of size $n / 2$ and then an $O(n)$ merge

Recurrence relation:

$$
\begin{aligned}
& \mathrm{T}(1)=\mathrm{c}_{1} \\
& \mathrm{~T}(n)=2 \mathrm{~T}(n / 2)+\mathrm{c}_{2} n
\end{aligned}
$$

## One of the recurrence classics...

For simplicity let constants be 1 (no effect on asymptotic answer)

$$
\begin{aligned}
T(1) & =1 \\
T(n) & =2 T(n / 2)+n \\
& =2(2 T(n / 4)+n / 2)+n \\
& =4 T(n / 4)+2 n \\
& =4(2 T(n / 8)+n / 4)+2 n \\
& =8 T(n / 8)+3 n \\
& \cdots \\
& =2^{k} T\left(n / 2^{k}\right)+k n
\end{aligned}
$$

So total is $2^{k} T\left(n / 2^{k}\right)+k n$ where

$$
n / 2^{k}=1 \text {, i.e., } \log n=k
$$

That is, $2^{\log n} T(1)+n \log n$

$$
=n+n \log n
$$

$$
=O(n \log n)
$$

## Or more intuitively...

This recurrence is common you just "know" it's $O(n \log n)$

Merge sort is relatively easy to intuit (best, worst, and average):

- The recursion "tree" will have $\log n$ height
- At each level we do a total amount of merging equal to $n$



## Quicksort

- Also uses divide-and-conquer
- Recursively chop into two pieces
- Instead of doing all the work as we merge together, we will do all the work as we recursively split into halves
- Unlike merge sort, does not need auxiliary space
- $O(n \log n)$ on average $\odot$, but $O\left(n^{2}\right)$ worst-case $)^{\circ}$
- Faster than merge sort in practice?
- Often believed so
- Does fewer copies and more comparisons, so it depends on the relative cost of these two operations!


## Quicksort Overview

1. Pick a pivot element
2. Partition all the data into:
A. The elements less than the pivot
B. The pivot
C. The elements greater than the pivot
3. Recursively sort A and C
4. The answer is, "as simple as $A, B, C$ "
(Alas, there are some details lurking in this algorithm)

## Think in Terms of Sets



## Example, Showing Recursion



## Details

Have not yet explained:

- How to pick the pivot element
- Any choice is correct: data will end up sorted
- But as analysis will show, want the two partitions to be about equal in size
- How to implement partitioning
- In linear time
- In place


## Pivots

- Best pivot?
- Median
- Halve each time

- Worst pivot?
- Greatest/least element
- Problem of size n-1
- O( $n^{2}$ )



## Potential pivot rules

While sorting arr from 10 (inclusive) to hi (exclusive)...

- Pick arr[lo] or arr[hi-1]
- Fast, but worst-case occurs with mostly sorted input
- Pick random element in the range
- Does as well as any technique, but (pseudo)random number generation can be slow
- Still probably the most elegant approach
- Median of 3, e.g., arr[lo], arr[hi-1], arr[(hi+lo) /2]
- Common heuristic that tends to work well


## Partitioning

- Conceptually simple, but hardest part to code up correctly
- After picking pivot, need to partition in linear time in place
- One approach (there are slightly fancier ones):

1. Swap pivot with $\operatorname{arr}[10]$
2. Use two fingers i and $\mathbf{j}$, starting at lo+1 and hi-1
3. while (i < j)
if (arr[j] > pivot) j--
else if (arr[i] < pivot) i++
else swap arr[i] with arr[j]
4. Swap pivot with arr [i] *
*skip step 4 if pivot ends up being least element

## Example

- Step one: pick pivot as median of 3
- lo = 0, hi = 10

| 0 |  |  |  | 3 | 4 | 5 | 6 | 7 |  | 8 | 9 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 8 | 1 |  |  | 9 | 0 | 3 | 5 | 2 |  | 7 | 6 |  |

- Step two: move pivot to the lo position



## Often have more than

## Example

 one swap during partition this is a short exampleNow partition in place

$$
\begin{aligned}
& \begin{array}{|l|l|l|l|l|l|l|l|l|l|}
\hline 6 & 1 & 4 & 9 & 0 & 3 & 5 & 2 & 7 & 8 \\
\hline
\end{array} \\
& \begin{array}{|l|l|l|l|l|l|l|l|l|l|}
\hline 6 & 1 & 4 & 9 & 0 & 3 & 5 & 2 & 7 & 8 \\
\hline
\end{array} \\
& \begin{array}{|l|l|l|l|l|l|l|l|l|l|}
\hline 6 & 1 & 4 & 2 & 0 & 3 & 5 & 9 & 7 & 8 \\
\hline
\end{array}
\end{aligned}
$$

Move fingers

Swap

Move fingers


Move pivot

$$
\begin{array}{|l|l|l|l|l|l|l|l|l|l|}
\hline 5 & 1 & 4 & 2 & 0 & 3 & 6 & 9 & 7 & 8 \\
\hline
\end{array}
$$

## Analysis

- Best-case: Pivot is always the median
$T(0)=T(1)=1$
$\mathrm{T}(n)=2 \mathrm{~T}(n / 2)+n \quad$-- linear-time partition
Same recurrence as mergesort: $O(n \log n)$
- Worst-case: Pivot is always smallest or largest element

$$
\begin{aligned}
& \mathrm{T}(0)=\mathrm{T}(1)=1 \\
& \mathrm{~T}(n)=1 \mathrm{~T}(n-1)+n
\end{aligned}
$$

Basically same recurrence as selection sort: $O\left(n^{2}\right)$

- Average-case (e.g., with random pivot)
- O( $n \log n$ ), not responsible for proof (in text)


## Cutoffs

- For small $n$, all that recursion tends to cost more than doing a quadratic sort
- Remember asymptotic complexity is for large $n$
- Common engineering technique: switch algorithm below a cutoff
- Reasonable rule of thumb: use insertion sort for $n<10$
- Notes:
- Could also use a cutoff for merge sort
- Cutoffs are also the norm with parallel algorithms
- Switch to sequential algorithm
- None of this affects asymptotic complexity


## Cutoff skeleton

```
void quicksort(int[] arr, int lo, int hi) {
    if(hi - lo < CUTOFF)
        insertionSort(arr,lo,hi);
    else
}
```

Notice how this cuts out the vast majority of the recursive calls

- Think of the recursive calls to quicksort as a tree
- Trims out the bottom layers of the tree


## Cool Resources

- http://www.sorting-algorithms.com/
- https://www.youtube.com/watch?v=t8g-iYGHpEA

