Algorithms - Basic Concepts

Algorithms – so what is an algorithm, anyway?

The dictionary definition: An algorithm is a well-defined computational procedure that takes input and produces output. This class will deal with how to design algorithms and understand their complexity in terms of runtime, space, and correctness.

Examples: data compression, path-finding, game-playing, scheduling, bin packing

Example algorithm: Insertion Sort

Insertion sort is how people often sort a hand of cards. Starting from the left, go towards the right and make sure that everything on the left hand side is correctly sorted. For example:

Original:	5	3	8	2	10	7
Pass 1:	3	5|	8	2	10	7
Pass 2:	3	5	8|	2	10	7
Pass 3:	2	3	5	8|	10	7
Pass 4:	2	3	5	8	10|	7
Pass 5:	2	3	5	7	8	10|

Here is the pseudocode for the Insertion sort algorithm.
Typically we will use pseudocode and not actual implementation code.

Insertion-Sort(A)   ; Assume our arrays start at index 1 and not at 0
        For j← 2 to length(A)
        Do
                key←A[j]
                i← j-1
                while i>0 and A[i] > key
                        do
                                A[i+1]← A[i]
                                i← i-1
                A[i+1]←key

Run through algorithm on our input, 5 3 8 2 10 7
J=2
        A=[5 3 8 2 10 7] Key=3
        I=1
        A[1]>3 so A[2]=A[1]
        A[1]=3

J=3
    A=[3 5 8 2 10 7] Key=8
    I=2
    A[2] < 8
    A[3]=8
J=4
    A=[3 5 8 2 10 7] Key=2
    I=3
    A[3]>2 so A[4]=A[3]
    A[2]>2 so A[3]=A[2]
    A[1]>2 so A[2]=A[1]
    A[1]=2
    J=5
    A=[2 3 5 8 10 7]
    …
    The final two passes are left as an exercise for you to verify.

It looks like this algorithm works to sort the input. But how do we analyze it? There are also certainly many other ways to sort the data (and we'll look at lots of them!) What makes one algorithm better than another?

Let's make an assumption that we will use throughout most of this class.

RAM model. This is a uniprocessor, random access machine with no parallelism (PRAM).

Run time and space will generally depend on the size of the input.

Input Size 'n' : Definition of n depends on the problem. It may be bits, bytes, but is usually the number of items in the input. For sorting, n=# of items to sort. For a graph, n could be the number of nodes or the number of links.

Running time: defined as the number of steps that are executed in the program.

Let's count up the runtime for this algorithm. This can be a little tricky since the loop doesn't always run the same number of times, but varies depending on the input.

Let t_j = the # of times the inner while loop is examined for the current value of j. This includes checking for the exit condition of the loop.

Code	Cost	Times
For j← 2 to length(A)	C1	n
Do		
key←A[j]	C2	n-1
i← j-1	C3	n-1
while i>0 and A[i] > key	C4	$\sum_{j=2}^{n} t_j$
do		
A[i+1]← A[i]	C5	$\sum_{j=2}^{n} (t_j - 1)$
i← i-1	C6	$\sum_{j=2}^{n} (t_j - 1)$
A[i+1]←key	C7	n-1

The statement for the outer j loop will run n times (once for each element, once to check the exit condition). The inner statements run n-1 times. Since we've defined t_j as the number of times the inner loop is executed for a value of j, to get the total times it is run we just add up the sum or all values of j.

Our total runtime is then:

$$C1(n)+C2(n-1)+C3(n-1)+C4(\sum_{j=2}^{n} t_j)+C5(\sum_{j=2}^{n} t_j - 1)+C6(\sum_{j=2}^{n} t_j - 1)+C7(n-1)$$

Doing a little math:

$$C1(n)+C2(n)-C2+C3(n)-C3+C4(\sum_{j=2}^{n} t_j)+(C5+C6)(\sum_{j=2}^{n} t_j - 1)+C7(n)-C7$$

$$(C1+C2+C3+C7)(n) - (C2+C3+C7) + (C4)(\sum_{j=2}^{n} t_j) + (C5+C6)(\sum_{j=2}^{n} (t_j - 1))$$

What's the best case? If the input is already sorted. In this case, $\forall j, t_j = 1$. That is the inner loop only gets executed once since we'll immediately find the right spot for the last item.

Runtime for this case:

$$(C1+C2+C3+C7)(n) - (C2+C3+C7) + (C4)(\sum_{j=2}^{n} 1) + (C5+C6)(0)$$
$$(C1+C2+C3+C7)(n) - (C2+C3+C7) + (C4)(n-1)$$

$$(C1+C2+C3+C4+C7)(n)-(C2+C3+C7+C4)$$

$$= (CONSTANT1)n - CONSTANT2$$

This is a linear function of n ; y=ax+b. The runtime will increase linearly as the size of the input increases
.

What's the worst case? Lets say that the input conspires to provide the worst runtime. This would happen if the array was in reverse order, because then we have to compare with every other number in the inner loop.

	5	4	3	2	1

j=2 t_j = 2 compares
j=3 t_j = 3 compares
j=4 t_j = 4 compares
…

So we know that $t_j = j$

$$\sum_{j=2}^{n} t_j = \sum_{j=2}^{n} j = \frac{n(n+1)}{2} - 1$$

$$\sum_{j=2}^{n} t_j - 1 = \sum_{j=2}^{n} j - \sum_{j=2}^{n} 1$$

$$= \frac{n(n+1)}{2} - 1 - (n-1)$$

$$= \frac{n^2 + n}{2} - \frac{2n}{2}$$

$$= \frac{n(n-1)}{2}$$

Plug these into our formula:

$$(C1+C2+C3+C7)(n) - (C2+C3+C7) + (C4)(\sum_{j=2}^{n} t_j) + (C5+C6)(\sum_{j=2}^{n} t_j - 1)$$

$$(C1+C2+C3+C7)(n) - (C2+C3+C7) + (C4)(\frac{n(n+1)}{2} - 1) + (C5+C6)(\frac{n(n-1)}{2})$$

Of the form $(Constant1)(n^2) + (Constant2)(n) + Constant3$

This is a quadratic; an^2+bn+c; the runtime will increase based on the square of the input.


Average case is also quadratic.

How much space does this take? Sort is done in place, so just a few variables. This requires constant space.

Function Growth

θ-notation : Gives asymptotically tight bound on a functions growth

Definition:

$$\theta\big(g(n)\big) = \{f(n):\exists c_1,c_2,n_0:0 \le c_1 g(n) \le f(n) \le c_2 g(n)\} \text{ for all } n \ge n_0$$



$F(n)=3n$ is $\theta(n)$ since we can find $g(n)=n$, $c1=1$, $c2=4$.

Insertion sort is not $\theta(n^2)$ since it is linear in the best case.
However, average and worst case insertion sort is $\theta(n^2)$.

O-notation: Big-O notation. This gives an asymptotic upper bound.
Definition:

$$O\big(g(n)\big) = \big\{ f(n) : \exists c_1, n_0 : 0 \le f(n) \le c_1 g(n) \big\} \text{ for all } n \ge n_0$$

Insertion sort is $O(n^2)$. Tight upper bound in the worst case, loose in the best case.
Looser: $O(n^3)$, $O(2^n)$.
Use little-o notation if you know that it is a loose bound.
$o(n^3)$ if $f = n^2$.

Ω - Omega notation. Omega provides an asymptotic lower bound.
Definition:

$$\Omega\big(g(n)\big) = \big\{ f(n) : \exists c_1, n_0 : 0 \le c_1 g(n) \le f(n) \big\} \text{ for all } n \ge n_0$$
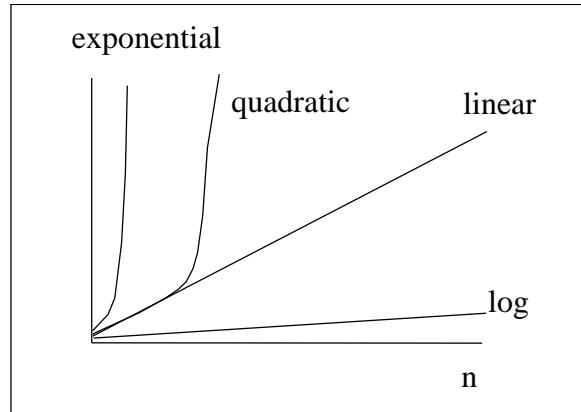

$\Omega(1)$ – Constant time, this is a trivial lower bound for most cases.
For insertion sort, $\Omega(n)$: linear time is the best possible

Running Time Comparison:

θ	n=10	n=100	n=1000	n=10000
lg(n)	2	5	7	9
n	10	100	1000	10000
n(lg(n))	20	500	7000	90000
n^2	100	10000	1,000,000	100,000,000
2^n	1024	1.3×10^{30}	huge	very huge

For comparison: estimate of the number of atoms in the universe is about 10^{80} (i.e. about 2^{270}).

One of the tools we will use for analysis and solutions is Divide and Conquer

Recursive approach to solve problems. The idea is to break the problem up into sub-problems with the same solution, but smaller size. Solve these problems recursively, then combine the sub-solutions to solve the original problem.

Divide – into sub-problems
Conquer – Sub-problems recursively
Combine – Sub-problem solutions to original problem

Let's look at Merge Sort:
        Divide n elements into two subsequences to be sorted of size n/2
        Conquer – sort subsequences recursively with merge sort
        Combine – merge sorted subsequences into big sorted answer
Need termination criteria for recursion –
        Quit if sequence to sort is length 1

Pseudocode:

    Merge_Sort(A,p,r)

        If p<r then

$$q \leftarrow \left\lfloor \frac{p+r}{2} \right\rfloor$$

           Merge_Sort(A,p,q)
           Merge_Sort(A,q+1,r)
           Merge(A,p,q,r)   ; Merges A[p..q] with A[q+1..r] into A[p..q]

Call with Merge_Sort(A,1,length(A))
q gets the median, merge left, right

A[1…5…10] becomes Merge_Sort(A,1,5), Merge_Sort(A,6,10)
How do we merge?

A1:	1	5	8	30	31	50
A2:	2	10	11	12	15

Combined: 1 2 5 8 10 11 12 15 30 31 50

Easy to do; just start at the front of each array, compare the pointers, and put the smallest into a new array and then increment the pointer that was the smallest.   We can't do this merge in-place though (using the original array storage location only) so we need to make a copy of the merged array and then copy it back to the original array.

If n is the number of elements to merge, then it takes $\theta(n)$ time to merge.

Example: A= 1 5 3 7 11 8 2 4
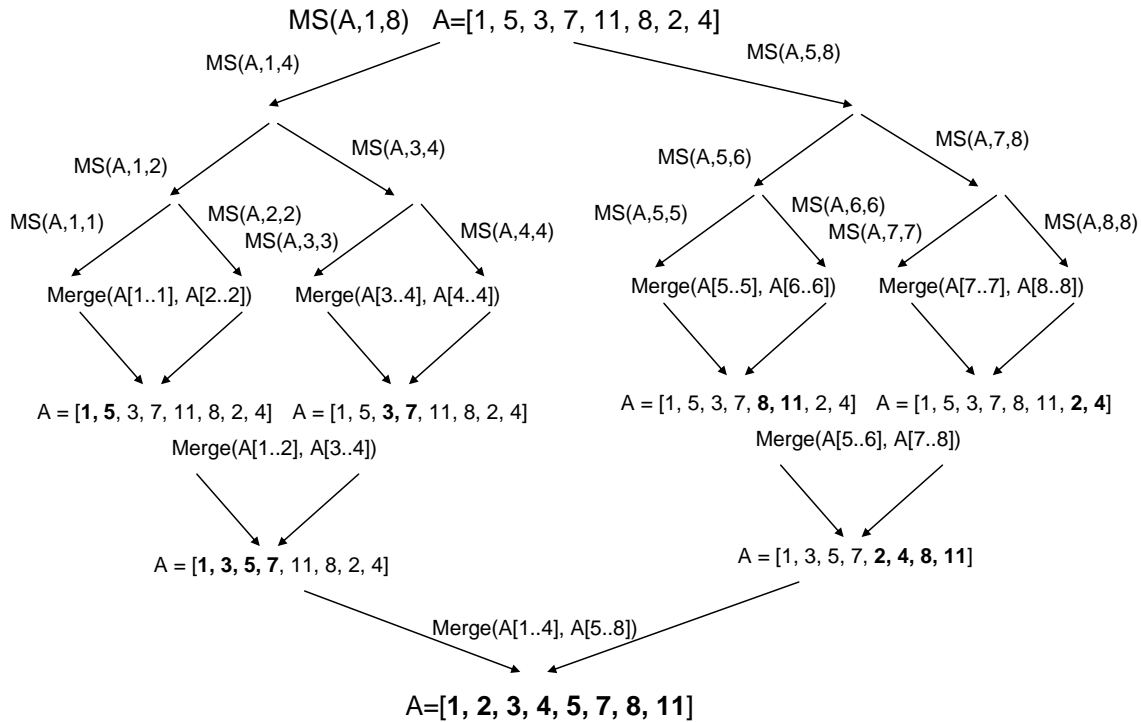
How much work is done? It's the cost to split + the cost to merge.
Let's define T(n) to be the runtime for a problem of size(n). This is called a **recurrence relation**, since it is recursively defined.

$$\text{Recurrence:} \quad T(n) = \begin{cases} \Theta(1), n = 1 \\ 2T\left(\dfrac{n}{2}\right) + \Theta(n), n > 1 \end{cases}$$

The 2T(n/2) comes from the divide, and the θ(n) comes from the merge.
Later we will show that Merge Sort is θ(nlgn) in runtime by illustrating a number of techniques to solve these recurrence relations.

Math Review

$\lfloor x \rfloor$ - floor of x, round down
$\lceil x \rceil$ - ceiling of x, round up

lg(x) = log$_2$ x
ln x = log$_e$ x

$a = b^{\log_b a}$
log$_c$(ab) = log$_c$(a) + log$_c$(b)
$$\log_b a = \frac{\log_c a}{\log_c b}$$
log$_b$ (1/a) = -log$_b$a
log$_b$a = 1 / (log$_a$b)
$a^{\log_b n} = n^{\log_b a}$

iterated log : log(log n) = log$^{(2)}$n
This grows very slowly! Almost the same as constant time.

$$\sum_{k=1}^{n} k = 1 + 2 + 3 + 4 + ... n = \frac{1}{2}n(n+1) = \Theta(n^2)$$

$$\sum_{k=1}^{n} \frac{1}{k} = 1 + \frac{1}{2} + \frac{1}{3} + ... \frac{1}{n} = \ln(n) + O(1) = \Theta(\ln n) \qquad \text{Harmonic series}$$

$$\sum_{k=0}^{n} a_k - a_{k-1} = a_n - a_0 \qquad \text{Telescoping}$$

There are some other summations we will use, but we'll discuss them as we get to that part of the course.

Monte Carlo Methods

Throughout this class we'll focus primarily on analytic methods to characterize algorithms. However, there is another class of statistical/probabilistic methods that are commonly referred to as Monte Carlo algorithms. From http://mathworld.wolfram.com/MonteCarloMethod.html : Any method which solves a problem by generating suitable random numbers and observing that fraction of the numbers obeying some property or properties. The method is useful for obtaining numerical solutions to problems which are too complicated to solve analytically. It was named by S. Ulam Eric Weisstein's World of Biography, who in 1946 became the first mathematician to dignify this approach with a name, in honor of a relative having a propensity to gamble (Hoffman 1998, p. 239).

Numerical methods that are known as Monte Carlo methods can be loosely described as statistical simulation methods, where statistical simulation is defined in quite general terms to be any method that utilizes sequences of random numbers to perform the simulation. Monte Carlo methods have been used for centuries, but only in the past several decades has the technique gained the status of a full-fledged numerical method capable of addressing the most complex applications. The name "Monte Carlo" was coined during the Manhattan Project of World War II, because of the similarity of statistical simulation to games of chance, and because the capital of Monaco was a center for gambling and similar pursuits.

Monte Carlo is now used routinely in many diverse fields, but has gained some of the greatest popularity in computer science and physics. For example, consider a complex physical system (solar system, nuclear reaction, quantum chromodynamics, traffic patterns, etc.) that has been modeled through a series of pdf's (probability density functions).   The behavior and interaction of the model with the input is often too complex to determine deterministically. Instead, the system is simulated by running many trials using random inputs selected from a proper range.   The resulting output helps us learn about the outputs of the system and the complexity of the system.

As another example, the merge sort example we previously covered could be analyzed using a monte carlo method by randomly selecting inputs, running the algorithm, and measuring the runtime (or the number of comparisons). By averaging out the average runtime, we could determine that merge sort is a θ(nlgn) algorithm.