


CS/COE 1501

www.cs.pitt.edu/~lipschultz/cs1501/

Introduction



Meta-notes

- These notes are intended for use by students in CS1501 at the University of Pittsburgh. They are provided free of charge and may not be sold in any shape or form.
- These notes are NOT a substitute for material covered during course lectures. If you miss a lecture, you should definitely obtain both these notes and notes written by a student who attended the lecture.
- Material from these notes is obtained from various sources, including, but not limited to, the following:
 - Algorithms in C++ by Robert Sedgwick
 - Algorithms, 4th Edition by Robert Sedgwick and Kevin Wayne
 - Introduction to Algorithms, by Cormen, Leiserson and Rivest
 - Various Java and C++ textbooks
 - Various online resources (see notes for specifics)

Instructor Info

- Michael Lipschultz (mil28@pitt.edu)
Office: 6148 Sennott Square
- TAs:
 - TBA

NO RECITATIONS THIS WEEK

A note about email

- Address all emails to both the instructor and the TA
 - I may not be able to respond during working hours

Course Info

- Website:
 - www.cs.pitt.edu/~lipschultz/cs1501/
- Review the Course Information and Policies
- **Assignments will not be accepted after the deadline**
 - No late assignment submissions
 - If you do not submit an assignment by the deadline, you will receive a **0** for that assignment

Up until now, your classes have focused on how you *could* solve a problem. Here, we will start to look at how you *should* solve a problem.

Alright, then, down to business...

First some definitions:

- Offline problem
 - We provide the computer with some input and after some time receive some acceptable output
- Algorithm
 - A step-by-step procedure for solving a problem or accomplishing some end
- Program
 - An algorithm expressed in a language the computer can understand

An algorithm solves a problem if it produces an acceptable output on *every* input

Goals of the course (1)

- To learn to convert non-trivial *algorithms* into *programs*
 - Many seemingly simple algorithms can become much more complicated as they are converted into programs
 - Algorithms can also be very complex to begin with, and their implementation must be considered carefully
 - Various issues will always pop up during *implementation*
 - Such as?...

Example

```
Input: SPJ query  $q$  on relations  $R_1, \dots, R_n$   
Output: A query plan for  $q$   
  
1: for  $i = 1$  to  $n$  do {  
2:    $\text{optPlan}(\{R_i\}) = \text{accessPlans}(R_i)$   
3:    $\text{prunePlans}(\text{optPlan}(\{R_i\}))$   
4: }  
5: for  $i = 2$  to  $n$  do {  
6:   for all  $S \subset \{R_1, \dots, R_n\}$  such that  $|S| = i$  do {  
7:      $\text{optPlan}(S) = \emptyset$   
8:     for all  $O \subset S$  do {  
9:        $\text{optPlan}(S) = \text{optPlan}(S) \cup \text{joinPlans}(\text{optPlan}(O), \text{optPlan}(S \setminus O))$   
10:       $\text{prunePlans}(\text{optPlan}(S))$   
11:     }  
12:   }  
13: }  
14:  $\text{finalizePlans}(\text{optPlan}(\{R_1, \dots, R_n\}))$   
15:  $\text{prunePlans}(\text{optPlan}(\{R_1, \dots, R_n\}))$   
16: return  $\text{optPlan}(\{R_1, \dots, R_n\})$ 
```

Fig. 1. "Classic" dynamic programming algorithm.

- Pseudocode for dynamic programming algorithm for relational query optimization
- The optimizer portion of the PostgreSQL codebase is over 28,000 lines of code (i.e., not counting blank/comment lines)

Goals of the course (2)

- To see and understand differences in algorithms and how they affect the run-times of the associated programs
 - Different algorithms can be used to solve the same problem
 - Different solutions can be compared using many metrics
 - Run-time is a big one
 - Better run-times can make an algorithm more desirable
 - Better run-times can sometimes make a problem solution feasible where it was not feasible before
 - There are other metrics, though...

How to determine an algorithm's performance

- Implement it and measure performance
 - Any problems with this approach?
- Algorithm Analysis
 - Determine *resource usage* as a function of *input size*
 - Measure *asymptotic* performance
 - Performance as input size increases to infinity

Let's consider ThreeSum example from text

- Problem:
 - Given a set of arbitrary integers (could be negative), find out how many distinct triples sum to exactly zero
- Simple solution: triple for loops!

```
public static int count(int[] a) {
    int n = a.length;
    int cnt = 0;
    for (int i = 0; i < n; i++) {
        for (int j = i+1; j < n; j++) {
            for (int k = j+1; k < n; k++) {
                if (a[i] + a[j] + a[k] == 0) {
                    cnt++;
                }
            }
        }
    }
    return cnt;
}
```

Definition of Big O?

- Big O
 - Upper bound on asymptotic performance
 - As we go to infinity, function representing resource consumption will not exceed specified function
 - E.g., Saying runtime is $O(n^3)$ means that as input size (n) approaches infinity, actual runtime will not exceed n^3

Wait...

- Assuming that definition...
 - Is ThreeSum $O(n^4)$?
 - What about $O(n^5)$?
 - What about $O(3^n)$??
- If all of these are true, why was $O(n^3)$ what we jumped to to start?

Big O isn't the whole story

- Big Omega
 - Lower bound on asymptotic performance
- Theta
 - Upper and Lower bound on asymptotic performance
 - Exact bound

Formal definitions

- $f(x)$ is $O(g(x))$ if constants c and x_0 exist such that:
 - $|f(x)| \leq c * |g(x)| \quad \forall x > x_0$
- $f(x)$ is $\Omega(g(x))$ if constants c and x_0 exist such that:
 - $|f(x)| \geq c * |g(x)| \quad \forall x > x_0$
- if $f(x)$ is $O(g(x))$ and $\Omega(g(x))$, then $f(x)$ is $\Theta(g(x))$
 - $c_1, c_2,$ and x_0 exist such that:
 - $c_1 * |g(x)| \leq f(x) \leq c_2 * |g(x)| \quad \forall x > x_0$
- May also see $f(x) \in O(g(x))$ or $f(x) = O(g(x))$ used to mean that $f(x)$ is $O(g(x))$
 - Same for Ω and Θ

Mathematically modelling runtime

- Runtime primarily determined by two factors:
 - Cost of executing each statement
 - Determined by machine used, environment running on the machine
 - Frequency of execution of each statement
 - Determined by program and input

Let's consider ThreeSum example from text

```
public static int count(int[] a) {  
    int n = a.length;  
    int cnt = 0;  
    for (int i = 0; i < n; i++) {  
        for (int j = i+1; j < n; j++) {  
            for (int k = j+1; k < n; k++) {  
                if (a[i] + a[j] + a[k] == 0) {  
                    cnt++;  
                }  
            }  
        }  
    }  
    return cnt;  
}
```


Quick algorithm analysis in detail

- Ignore multiplicative constants and lower terms
- Use standard measures for comparison

Common orders of growth

- Constant - 1
- Logarithmic - $\log n$
- Linear - n
- Linearithmic - $n \log n$
- Quadratic - n^2
- Cubic - n^3
- Exponential - 2^n
- Factorial - $n!$

Graphical orders of growth



Tilde approximations and Order of Growth

- ThreeSum order of growth:
 - Upper bound: $O(n^3)$
 - Lower bound: $\Omega(n^3)$
 - And hence: $\Theta(n^3)$

- Tilde approximations?
 - Introduced in section 1.4 of the text
 - In this case: $\sim n^3/6$

How can we ignore lower order terms and multiplicative constants???

- Remember, this is asymptotic analysis

f(n) / n =	10	100	1,000	10,000
$n^3/6 - n^2/2 + n/3$	120	161,700	166,167,000	166,616,670,000
$n^3/6$	167	166,667	16,6666,667	166,666,666,667
n^3	1,000	1,000,000	1,000,000,000	1,000,000,000,000

Easy to get Theta for ThreeSum

- Why do we need to bother with Big O and Big Omega?

Further thoughts on ThreeSum

- Is there a better way to solve the problem?
- What if we sorted the array first?
 - Pick two numbers, then binary search for the third one that will make a sum of zero
 - $a[i] = 10$, $a[j] = -7$, binary search for -3
 - Still have two for loops, but we replace the third with a binary search
 - Runtime now?
 - What if the input data isn't sorted?
- See `ThreeSumFast.java`

Topics for the term

- Sorting
- Searching
- Hashing
- Compression
- Heaps and Priority Queues
- Graph Algorithms
- Cryptography
- P vs NP
- Heuristic Approximation
- Dynamic Programming