

CMPSCI 311: Introduction to Algorithms

Akshay Krishnamurthy and Andrew McGregor

University of Massachusetts

Last Compiled: September 6, 2016

CS 311: Intro to Algorithms

- ▶ **Instructor:** Akshay Krishnamurthy and Andrew McGregor
- ▶ **Where:** Hasbrouck Laboratory 124
- ▶ **When:** Tu/Th 11:30–12:45
- ▶ **Discussion Sections:** Wednesdays 10:10, 11:15, 12:05 (Please stick to assigned section)
- ▶ **TAs:** Walter Brown and Archan Ray
- ▶ **Office hours:**
 - ▶ Akshay: 1pm Tuesday in CS 258.
 - ▶ Andrew: 3pm Thursday in CS 334.
 - ▶ Archan: TBA
 - ▶ Walter: TBA

What is Algorithm Design?

How do you write a computer program to solve a complex problem?

- ▶ Computing similarity between DNA sequences
- ▶ Routing packets on the Internet
- ▶ Scheduling final exams at a college
- ▶ Assign medical residents to hospitals
- ▶ Find all occurrences of a phrase in a large collection of documents
- ▶ Finding the smallest number of coffee shops that can be built in the US such that everyone is within 20 minutes of a coffee shop.

DNA sequence similarity

- ▶ **Input:** two n -bit strings s_1 and s_2
 - ▶ $s_1 = \text{AGGCTACC}$
 - ▶ $s_2 = \text{CAGGCTAC}$
- ▶ **Output:** minimum number of insertions/deletions to transform s_1 into s_2
- ▶ **Algorithm:** ????
- ▶ Even if the objective is precisely defined, we are often not ready to start coding right away!

What is Algorithm Design?

- ▶ **Step 1:** Formulate the problem precisely
- ▶ **Step 2:** Design an algorithm
- ▶ **Step 3:** Prove the algorithm is correct
- ▶ **Step 4:** Analyze its running time

Important: this is an iterative process, e.g., sometimes you'll even want to redesign the algorithm to make it easier to prove that it is correct.

Course Goals

- ▶ Learn how to apply the algorithm design process... by practice!
- ▶ Learn specific algorithm design techniques
 - ▶ Greedy
 - ▶ Divide-and-conquer
 - ▶ Dynamic Programming
 - ▶ Network Flows
- ▶ Learn to communicate precisely about algorithms
 - ▶ Proofs, reading, writing, discussion
- ▶ Prove when no exact efficient algorithm is possible
 - ▶ Intractability and NP-completeness

Grading Breakdown

- ▶ **Participation (10%):** Discussion assignments and other contributions.
- ▶ **Homework (25%):** Homework (every two weeks due Friday) and online quiz (every weekend due Monday).
- ▶ **Midterm 1 (20%):** Focus on first third of lectures. 7pm Thursday 13th October.
- ▶ **Midterm 2 (20%):** Focus on second third of lectures. 7pm Thursday 17th November.
- ▶ **Final (25%):** Covers all lectures. 1pm Friday 16th December.

Course Information

Course websites: Slides, homework, course information:

people.cs.umass.edu/~mcgregor/courses/CS311F16/

Quizzes, solutions, grades will be posted at:

moodle.umass.edu

Forums for discussion and contacting instructors and TA's:

piazza.com

Announcements: Check your UMass email daily and log into Piazza regularly for course announcements.

Policies

- ▶ **Online Quizzes:** Quizzes must be submitted before 8pm Monday. No late quizzes allowed but we'll ignore your lowest scoring quiz.
- ▶ **Homework:** To be submitted via Moodle. Marks for any homework that is late by 5 minutes up to 24 hours, will be halved. However, you are allowed to submit *one* homework up to 24 hours late without penalty; don't use this lifeline prematurely! Homework later than 24 hours receives no credit.

Collaboration and Academic Honesty

- ▶ **Homework:** You may discuss homework with other students but the writeup and code must be your own. There'll be formal action if cheating is suspected. You must list your collaborators and any printed or online sources at the top of each assignment.
 - ▶ **Online Quizzes:** Should be done entirely on your own although it's fine to consult the book and slides as you do the quiz. Again, there'll be formal action if cheating is suspected.
 - ▶ **Discussions:** Groups for the discussion section exercises will be assigned randomly at the start of each session. You must complete the discussion session exercise with your assigned group.
 - ▶ **Exams:** Closed book and no electronics. Cheating will result in an F in the course.
- * If it doubt whether something is allowed, ask...

Stable Matching and College Admissions

- ▶ Suppose there are n colleges c_1, c_2, \dots, c_n and n students s_1, s_2, \dots, s_n .
- ▶ Each college has a ranking of all the students that they could admit and each student has a ranking of all the colleges. For simplicity, suppose each college can only admit one student.
- ▶ Can we match students to colleges such that everyone is *happy*?
 - ▶ Not necessarily, e.g., if UMass was everyone's top choice.
- ▶ Can we match students to colleges such that matching is *stable*?
 - ▶ **Stability:** Don't want to match c with s and c' with s' if c and s' would prefer to switch to being matched with each other.
 - ▶ Yes! And there's an efficient algorithm to find that matching.

Propose-and-Reject (Gale-Shapley) Algorithm

Initially all colleges and students are free

```
while some college is free and hasn't proposed to every student
do
  Choose such a college  $c$ 
  Let  $s$  be the highest ranked student to whom  $c$  has not
  proposed
  if  $s$  is free then
     $c$  and  $s$  become matched
  else if  $s$  is matched to  $c'$  but prefers  $c$  to  $c'$  then
     $c'$  becomes unmatched
     $c$  and  $s$  become matched
  else
     $s$  rejects  $c$  and  $c$  remains free
  end if
end while
```

▷ s prefers c'

Analyzing the Algorithm

- ▶ Some natural questions:
 - ▶ Can we guarantee the algorithm terminates?
 - ▶ Can we guarantee the every college and student gets a match?
 - ▶ Can we guarantee the resulting allocation is stable?
- ▶ Some initial observations:
 - ▶ College propose to students in order of college's preferences.
 - ▶ Every student only "upgrades" during the algorithm.

Can we guarantee the algorithm terminates?

- ▶ Yes! Proof...
 - ▶ Note that in every round, some college proposes to some student that they haven't already proposed to.
 - ▶ Any college can propose to at most n different students and there are only n colleges.
 - ▶ There are at most n^2 different rounds of the algorithm.

Can we guarantee all colleges and students get a match?

- ▶ Yes! Proof by contraction...
 - ▶ Suppose not all colleges and students have matches. Then there exists unmatched college c and unmatched student s .
 - ▶ When algorithm terminates c has proposed to every student.
 - ▶ Student s was never matched during the algorithm.
 - ▶ When c proposed to student s , s was unmatched and yet rejected c . Contradiction!

Can we guarantee the resulting allocation is stable?

- ▶ Yes! Proof by contraction with a case analysis...
 - ▶ Suppose there exists colleges c and c' and students s and s' such that c is matched to s' and c' is matched to s but c and s would prefer to be matched to each other.
 - ▶ Case 1: s got an offer from c
 - ▶ Since s isn't matched to c at the end of the algorithm, they must have rejected c 's offer at some point and are therefore must be matched to a college they prefer.
 - ▶ Case 2: s never got an offer from c
 - ▶ c must have been matched to a student that they prefer.
 - ▶ Either way, at least one of c or s would not want to be matched to the other more than their current assignment.

For Thursday

- ▶ Think about:
 - ▶ Would it be better or worse for the students if we ran the algorithm with the students proposing.
 - ▶ Can a student get an advantage by lying about their preferences?
- ▶ Read: Chapter 1, course policies, handout about Nobel prize
- ▶ Enroll in Piazza, log into Moodle, and visit the course webpage.