

CMPSCI 711: More Advanced Algorithms

Section 0: Overview

Andrew McGregor

Last Compiled: April 29, 2012

Goal and Focus of Class

- ▶ **Goal:** Learn general techniques and gain experience designing and analyzing algorithms with randomization and approximation.
- ▶ **Focus:** Syllabus will focus on problems processing large data sets.
 - ▶ **Data Streams:** You process a long sequence $\langle a_1, \dots, a_m \rangle$ of items (numbers, edges of a graph, points in a geometric space etc.) sequentially. What can you compute without storing entire sequence?
 - ▶ **Linear Sketches:** Given a high-dimensional vector $x \in \mathbb{R}^n$, what can you compute about x given a projection $Ax \in \mathbb{R}^k$ where $k \ll n$ and A is a (random) $k \times n$ matrix or your choosing?
 - ▶ **Sampling:** Given a limited number of samples from some data set, what quantities can you accurately estimate with high probability?

Outline

Examples

Class Overview

Example 1: Simple Puzzles

- ▶ **Puzzle 1:** A bag contains n balls which are numbered 1 through n . Balls are removed one-by-one until only one remains. Can you find the number on the remaining ball?
- ▶ **Answer:** Missing value is $(\sum_{i=1}^n i)$ -(sum of removed values).
- ▶ **Puzzle 2:** Suppose bag is initially empty. At each step, a numbered ball is either added or removed. Multiple balls can have same number. Can you determine if all balls in bag have same number?
- ▶ **Answer:** Let f_i be the number of balls with number i in the bag. It's easy to compute $m = \sum f_i$, $U = \sum if_i/m$, and $V = \sum i^2 f_i/m$. Then balls in the bag have the same number iff $U^2 = V$.

Example 2: Randomization and Distinct Elements

- ▶ **Puzzle:** Suppose the bag is initially empty. At each step, a numbered ball is added to the bag. Multiple balls can have same number. Can you estimate how many distinct values are in the bag?
- ▶ **Sorta Answer:** For each $i \in [n]$, pick random $h_i \in_R [0, 1]$. Find

$$X = \min\{h_i : i \text{ added to the bag}\}$$

Can show $\mathbb{E}[X] = \frac{1}{F_0+1}$ where F_0 is the number of distinct values

- ▶ **Problems:**
 - ▶ How can we store all the h_i in limited space?
 - ▶ How can we find F_0 up to a factor $(1 + \epsilon)$ with high probability?
 - ▶ What if balls can be added and removed?

Example 3: Graphs

- ▶ **Puzzle:** Suppose your name is Mark Z. and you can see all the friendships that are forming in the social network that you invented. Can you tell if the network is connected with only $O(n \log n)$ bits of memory where n is the number of people in the network?
- ▶ **Sorta Answer:** Maintain spanning forest of growing network.
- ▶ **Problems:** What can you do if people also de-friend each other?

Example 4: Clustering

- ▶ Suppose you are skimming through holiday photos and you want to find k photos that best summarize your trip. Given two photos p_i and p_j you define a distance metric $d(p_i, p_j)$ between the photos. . .
- ▶ **Puzzle:** Find k photos $\mathcal{C} = \{p_{i_1}, \dots, p_{i_k}\}$ minimizing $\max_j d(p_j, \mathcal{C})$. Suppose we know that optimal solution has value t . Can we find a solution with value at most $2t$ while only remembering k photos.
- ▶ **Answer:** Add a new photo p to \mathcal{C} if $d(p, \mathcal{C}) > 2t$.
- ▶ **Problems:** What if you didn't know value of optimum solution?

Outline

Examples

Class Overview

Approximate Schedule

- ▶ **Numbers (5 lectures):** Basic techniques and estimating numerical statistics such as quantiles, heavy-hitters, frequency moments, etc.
- ▶ **Graphs (4 lectures):** Graph problems including connectivity, cuts, cliques, matchings, sparsification, spanners, etc.
- ▶ **Clustering and Geometry (3 lectures):** Clustering, core-sets, ϵ -nets, convex hulls, minimum enclosing ball, etc.
- ▶ **Sequences (3 lectures):** Longest increasing subsequences, formal language recognition, memory checking, etc.
- ▶ **Lower Bounds (3 lectures):** Basics of communication complexity and applications, information theory, etc.
- ▶ **Special Topics (6 lectures):** Compressed sensing, sliding windows, stochastic streams, distributed monitoring, parallel algorithms, etc.

Material

- ▶ All required material will be handed out in class or posted at:
<https://www.cs.umass.edu/~mcgregor/courses/CS711S12/>
- ▶ No required textbook but you might find the following helpful:
 - ▶ **Data Streams: Algorithms and Applications** by Muthukrishnan
 - ▶ **Randomized Algorithms** by Motwani and Raghavan
 - ▶ **Probability and Computing** by Mitzenmacher and Upfal

Homework, Grading, and Stuff

- ▶ There will be 4 homework which you can do in groups of 3.
- ▶ There will be a couple of lectures for class presentations of a relevant paper of your choosing. I'll post a list of suggestions.
- ▶ Final mark will be based on a combination of homework marks, quality of presentation, and general class participation.