

# CS167: Readings in Algorithms

Tim Roughgarden

March 31, 2014

# Outline for today

1. Course goals
2. Plan for course
3. Deliverables
4. Example topics
5. Technical topic: counting triangles via Map-Reduce. (see separate course notes)

# Course goal #1

Message: algorithms rule!

- As vibrant and important a research field as ever.
  - New algorithms/data structures didn't stop in the 70s!
  - New progress on old problems; new problems from new applications; old problems on new architectures
  - All papers from 21st century, if not the last few years
- Example: Symposium on Discrete Algorithms (SODA – top algorithms conference),  $\approx 150$  papers/year
  - Also lots of cool algorithms in more applied conferences (machine learning, data mining, etc.)

# Course goal #2

Message: research is different/humbling, but cool.

- Not like problem sets, at all.
  - Formulating the right problem often the hardest part.
  - No idea what's true or false, trivial or impossible.
  - Requires some serious creativity.
- Reading research papers is hard.
  - Often written only for experts, and even then, badly.
  - Will develop “paper reading survival kit”.
  - Even understanding 10% of a paper can be a victory.

# Course goal #3

Message: small classes are cool.

- Can easily graduate with a Stanford CS degree without any professor knowing your name.
- I don't get to know any undergrads.
- This class will be different.

# Course goal #4

Message: presentation skills are important.

- In research and otherwise.
- In this class, more means to an end:
  - Best way to make sure you understand this stuff is to explain it (to yourself through writing it up, and to others by telling them about it).

# Plan for course

- First 2-3 weeks: instructor lectures on topics of broad interest.
  - Current trends in algorithms research; survival guide for reading papers; tips for speaking and writing; landscape of relevant conferences; etc.
- Next 1-2 weeks: paper reading, group discussions.
- Last 5-6 weeks: student presentations
  - 25-30 minute per student (solo or in pairs)

# Deliverables

- Overall grade out of 100 points.
  1. 20 points: attendance.
  2. 20 points: paper responses. (.5-1 page per paper; paper-specific guidelines will be given)
  3. 20 points: paper presentation (25-30 min)
  4. 20 points: practice presentation (1 wk before)
  5. 20 points: write-up to accompany presentation (due 1 week after)



# Example Topics

- Bread-and-butter algorithms
  - Traveling Salesman Problem, bipartite matching
- Bread-and-butter data structures
  - E.g., latest development in hashing, heaps
- Algorithms for social networks
  - Triangle counting, de-anonymization, etc.

# Example Topics (con'd)

- Algorithms for massive data sets
  - Map-reduce, dimensionality reduction, etc.
- Algorithms in machine learning
  - E.g., topic modeling, feature selection
- Miscellaneous cool stuff
  - Algorithmic LLL, pseudorandom data, etc.