Seminar Algorithms for Database Systems

Michael Böhlen Przemysław Uznański

February 24, 2017



- Goals of the Seminar
- Logistics
- Topics
- Next Steps
- Grading

Goals of the Seminar

- Get a deep understanding of an advanced topic/paper in algorithms for database systems
- Focus on one technical paper/topic (not an overview of ten papers!)
- Give a talk on this paper, write a manuscript, lead the discussion
- Act as a buddy to another student (read paper, read report, give feedback)

Prerequisites

• Good understanding of database systems and of algorithms

Attendance Limitation

- Total seminar attendance will be limited to 18
- If more than 18 students want to take the seminar, we will choose randomly among all students

Topics Preference

- You will get the chance to submit your preferences on topics
- We will use RSD to assign 1) topics and 2) buddies

Topics

- Cache oblivious: matrix transpose, fft, sorting
- Cache-conscious data placement techniques
- Cache oblivious B-trees
- Cache conscious B-trees
- Cache oblivious priority queue
- Cost functions for database operations
- Cache obliviousness: lower bounds
- Cache performance in online transaction processing
- Mesh Layouts for block-based caches
- Cache-oblivious dynamic programming
- Cache-oblivious databases
- Cache-oblivious bloom filters
- Cache + CPU/GPU architectures
- Interval join + cache
- GPU-based pipelined query processing engine

Next steps (1/2): topics

- Until Sunday (26.2.2016), 23:59, send an email to: przemyslaw.uznanski@inf.ethz.ch containing the following information:
 - Name
 - Student number (ETH/UZH)
 - preferences for topics (cf. seminar page)
- We use RSD to assign topics to students
- We also assign buddies (using common sense)
- We assign advisors (topic-based) and dates (random/common sense)
- We publish/send you the list of all assigned topics, buddies and advisors

Next steps (2/2): preparing manuscript + talk

- Read your paper (and related papers)
- Write manuscript (4 8 pages) on how would you present it: motivation, formal model, interesting proofs, main results
- Send your manuscript to buddy + the advisor 3 weeks before your talk
- $\bullet\,$ Meet with the advisor and buddy ~ 2 weeks before the talk
- submit final version of manuscript 1 week before the talk (STRICT)
- prepare talk, rehearse with buddy
- talks: blackboard talk (30min) + lead discussion (10min) on:
 - Saturday April 8 at UZH
 - Saturday May 13 at ETH
- NO SLIDES

Grading

In order of importance:

- Presentation
- Manuscript
- Buddy (you are responsible for quality of your partner talk)
- Active Participation

Useful link on how to read papers and give talks: http://cs.stanford.edu/~rishig/courses/s14.html