

Topics in Reliable Distributed Systems 048961

1

WINTER 2011-2012
IDIT KEIDAR

Time & Place

2

- Fischbach 403
- Mondays 14:30-16:10 (no break)
- Updates on course web page
<http://webee.technion.ac.il/~idish/048961/>
- Please send me e-mail **today**
 - to idish@ee
 - with 048961 in the subject
 - and an empty body (will not be read)

Course Overview

3

- **Graduate level**
 - Open to excellent undergraduates with permission from me
 - Can help you find a thesis topic
- **Format: reading group & seminar**
 - Discussion of **recent** research papers
- **Please select a paper by Nov 7 (in 2 weeks)**
 - Need approval from me

This Year's Foci

4

- **Scale-out in data centers and clouds**
 - Systems research papers on scalable storage systems
 - Systems deployed in industry
- **Race detection and deterministic execution**

Where to Find Exciting Recent Papers

5

- Presented at this very moment in SOSP
<http://sigops.org/sosp/sosp11/current/index.html>

SOSP: Key-Value, Storage, Geo-Replication

6

- SILT: A Memory-Efficient, High-Performance Key-Value Store
- Scalable Consistency in Scatter
- Fast Crash Recovery in RAMCloud
- Design Implications for Enterprise Storage Systems via Multi-Dimensional Trace Analysis
- Transactional storage for geo-replicated systems
- Don't Settle for Eventual: Scalable Causal Consistency for Wide-Area Storage with COPS

SOSP “Reality” Session

7

Real-world Industrial Distributed Systems

- Google's [Thialfi: A Client Notification Service for Internet-Scale Applications](#)
- Microsoft's [Windows Azure Storage: A Highly Available Cloud Storage Service with Strong Consistency](#)

More Research On Scalable Cloud Systems

8

- Centrifuge: integrated lease management and partitioning for cloud services
Adya et al. NSDI'10
- Volley: Automated data placement for geo-distributed cloud services
Agarwal et al. NSDI'10
- The Case for Determinism in Database Systems
Thomson and Abadi VLDB'10
- Boom analytics: exploring data-centric, declarative programming for the cloud
Alvaro et al. EuroSys'10

Older Scalable Cloud Systems In Industry

9

- Google Bigtable
Bigtable: A Distributed Storage System for Structured Data
- Apache Hadoop
HadoopDB: an architectural hybrid of MapReduce and DBMS technologies for analytical workloads
- Yahoo! PNUTS
PNUTS: Yahoo!'s hosted data serving platform

SOSP: Races and Determinism

10

- Dthreads: Efficient Deterministic Multithreading
- Efficient Deterministic Multithreading through Schedule Relaxation
- Pervasive Detection of Process Races in Deployed Systems
- Detecting and Surviving Data Races using Complementary Schedules

Requirements and Grading

11

- **Reading the papers (one a week)**
- **10 short paper reviews – 20%**
- **Participating in class discussions – 10%**
- **Presenting one of the papers – 70%**

Reading The Papers

12

- **This is a reading group**
 - You have to read each paper before it is discussed
 - Most will be conference papers
- **Read the entire paper**
 - Be familiar with all of its content
 - No need to understand everything, check previous work, or memorize details
- **We will focus on critical reading**

Submission Requirement

13

- **A short review of each paper**
 - (unless you are the one presenting it)
 - ½ to 1 page long (not more!)
 - See next slide for required structure/content
- **By e-mail to me and to the paper's presenter**
 - By the night* before the lecture

***Any time before 8:00am the morning of the lecture is considered part of the night before**

Paper Review Structure & Content

14

- 1. One paragraph **overview****
 - What question is the paper is trying to answer?
 - What are the main (most important) results?
- 2. One subjective paragraph on **your experience****
 - What did you learn?
 - What didn't you understand?
 - Were you surprised? Disappointed?
- 3. Opinion and **evaluation****
 - How did you like the paper?
 - Main strengths and weaknesses (2-3 bullets each)
- 4. Interesting **question** (or at most 2) for further research**

Evaluating Strengths & Weaknesses (Examples)

15

- Is the paper answering the “right” question?
 - Does it make reasonable assumptions?
- How novel is the solution?
- Is the solution technically sound?
- Is the solution difficult or easy?
- How well is the solution evaluated?
- Expected impact (hard to guess)
- Writing level: is the paper clearly written? Is it self-contained?

Thinking of Interesting Research Questions

16

- Ask why they do what they do the way they do it
- Examples
 - Can you think of an interesting related problem?
 - ✦ Similar– a variant of this one
 - ✦ Complementary
 - Is the problem interesting under different assumptions?
 - How would you tackle the problem if posed to you?
 - ✦ Would you try another approach instead of the one in the paper?
 - Is the solution optimal under some metric?
 - ✦ Can it be proven optimal?
 - ✦ Can it be improved?
 - ✦ Are there other interesting metrics to optimize for?
- Pick one question, develop it in a couple of sentences

Presenting a Paper

17

- **Fully understand the paper**
 - Be able to explain it
- **Put the paper in context**
 - Compare with similar work
 - Which ideas are new and which existed before?
- **Plan an hour-long presentation (time yourself!)**
 - Leave ½ hour for Q&A, discussion
- **Prepare a list of topics to discuss in class**
 - Include questions in your slides
- **Discuss the presentation with me beforehand**

Presentation Structure

18

- 1. Background**
 - Historical perspective, old work in the area
 - Recent related work
- 2. Discussion: the paper's contribution**
 - Main contribution(s)
 - Novelty relative to previous work
- 3. Explanation of technically difficult parts**
 - Use of animation is encouraged
- 4. Discussion: evaluation**
 - Strengths and weaknesses

Contact Me

19

- **Idit Keidar <idish@ee>**
 - Please send me e-mail **today** with **048961** in the subject, and an **empty body**
 - Warning: Technion spam filter may block email from company addresses
- **Let me know by Nov 7**
 - What you would like to present (first-come-first-served)
 - Any scheduling constraints
- **Schedule will be posted on the course web page**