

Topics in Reliable Distributed Systems

049011



WINTER 2017-2018
IDIT KEIDAR

Time & Place

2

- Mondays 14:30-16:10 (no break)
- Meyer 353
- Schedule on course web page (check for updates)
<http://webee.technion.ac.il/~idish/049011/>
- To join the mailing list, please send me e-mail **today**
 - to `idish@ee`
 - with **049011 signup** in the subject
 - and an empty body (will not be read)
- Include the course number **049011** in the subject in all emails you send me about the course

General Course Info

3

- Graduate level
 - Open to excellent undergraduates with permission from me
 - Can help you find a thesis topic
- Prerequisite
 - Some basic course in distributed computing
- Format: reading group & seminar
 - Discussion of **recent** research papers
- Please select a paper by November 2
 - Papers allocated on first-come-first-served basis
 - Need approval from me

Selected SOSP'17 Papers 1

4

Bug hunting

- Pensieve: Non-Intrusive Failure Reproduction for Distributed Systems using the Event Chaining Approach
- Canopy: An End-to-End Performance Tracing And Analysis System

Selected SOSP'17 Papers 2

5

Scalability

- Algorand: Scaling Byzantine Agreements for Cryptocurrencies
- Scaling a file system to many cores using an operation log
- SVE: Distributed Video Processing at Facebook Scale

Selected [SOSP'17](#) Papers 3

6

In-network computing

- Eris: Coordination-Free Consistent Transactions Using In-Network Concurrency Control
- NetCache: Balancing Key-Value Stores with Fast In-Network Caching
- KV-Direct: High-Performance In-Memory Key-Value Store with Programmable NIC

Adaptation and repair

- Rocksteady: Fast Migration for Low-latency In-memory Storage

Selected [SOSP'17](#) Papers 4

7

Privacy

- Atom: Horizontally Scaling Strong Anonymity
- Stadium: A Distributed Metadata-Private Messaging System
- Prochlo: Strong Privacy for Analytics in the Crowd
- NOVA-Fortis: A Fault-Tolerant Non-Volatile Main Memory File System

Storage systems

- PebblesDB: Building Key-Value Stores using Fragmented Log-Structured Merge Trees

Selected SOSP'17 Papers 5

8

Data analytics

- Sub-millisecond Stateful Stream Querying over Fast-evolving Linked Data
- Optimizing Big-Data Queries Using Program Synthesis
- Low-Latency Analytics on Colossal Data Streams with SummaryStore

Papers, slides, and talks available [online](#)

Selected OSDI'16 Papers

9

- Just Say NO to Paxos Overhead: Replacing Consensus with Network Ordering
- Realizing the Fault-Tolerance Promise of Cloud Storage Using Locks with Intent
- Consolidating Concurrency Control and Consensus for Commits under Conflicts
- Correlated Crash Vulnerabilities

Papers, slides, and talks available [online](#)

Selected SOSP'15 Papers

10

- No compromises: distributed transactions with consistency, availability, and performance
- Implementing Linearizability at Large Scale and Low Latency
- High-Performance ACID via Modular Concurrency Control
- Paxos Made Transparent
- Holistic Configuration Management at Facebook

Papers, slides, and talks available [online](#)

Other conferences

11

- Seeing is Believing: A client-centric specification of database isolation (PODC'17) ([link](#))
- Recovering Shared Objects Without Stable Storage (DISC'17) ([link](#))
- Black-box Concurrent Data Structures for NUMA Architectures (ASPLOS'17) ([link](#))
- I Can't Believe It's Not Causal! Scalable Causal Consistency with No Slowdown Cascades (NSDI'17) ([link](#))

Requirements and Grading

12

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

Reading the Papers

13

- This is a reading group
 - You have to read each paper before it is discussed
 - 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

Weekly Submission Requirement

14

- A short review of the presented paper
 - (unless you are the one presenting it)
 - 1/2 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
 - Include the course number **049011** in the subject

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

Paper Review Structure & Content

15

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 & weaknesses (2-3 bullets each, see next slide)
4. Interesting **question** (or at most 2) for further research (see slide after next)

Evaluating Strengths & Weaknesses (Examples)

16

- 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

17

- 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

When It's Your Turn to Present a Paper

18

- **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 1/2 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

19

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
 - Recycling publicly available slides is strongly encouraged!
4. **Discussion: evaluation**
 - Strengths and weaknesses

Contact Me

20

- **Idit Keidar <idish@ee>**
 - Please send me e-mail **today** with **049011 signup** in the subject, and an **empty body**
 - Warning: Technion spam filter may block email from company addresses
- **Let me know by November 2**
 - What you would like to present (first-come-first-served)
 - Any scheduling constraints
- **Please include the course number **049011** in the subject in all correspondences about the course**
- **Schedule will be posted on the course web page**