CSE 4/586 Distributed Systems

Murat Demirbas

[2020-08-31 Mon]

Distributed Systems

Distributed systems need radically different software than centralized systems do.
A. Tannenbaum
`http://www.youtube.com/watch?v=eakKfY5aHmY`

Distributed Systems

A collection of autonomous nodes

- communicating over a network to address a problem collectively
- with no shared memory
- with no common physical clock

Distribution goals

- Scalability
- Reliability
- Accommodating geographic separation

Distribution challenges

- Coordination despite lack of global view (lack of global state/time)
- Handling failures of the nodes and communication
- Reasoning about nondeterminism due to concurrency
- Scalability

Distributed systems examples

- Cloud computing
- Grid computing
- Client-server systems
- Peer-to-peer systems, decentralized systems
- Network protocols
- Wireless sensor networks
- Ad hoc and mobile networks

Our course

Our course covers principles for design & analysis of distributed systems

- reasoning about distributed programs
- handling the lack of global state and global time
- achieving distributed consensus in the presence of faults
- designing fault-tolerance for distributed systems

The course also reviews state-of-the-art distributed systems, particularly cloud computing systems

Topics

- Introduction, 2 phase-commit
- Safety/progress properties and proving program properties
- Consensus, Paxos
- Failure detectors, Faults and fault-tolerance
- Time: logical clocks, State: distributed snapshots

- Datacenter computing, Cloud computing
- NoSQL databases, CAP theorem, Distributed databases
- Big data, Big data analytics
- Decentralized ledgers and blockchains

Course approach

Learning about Distributed Systems:
`http://muratbuffalo.blogspot.com/2020/06/learning-about-distributed-systems.html`

Non-topics

- Internet protocols
- Security
- Parallel computing
- Web services

Logistics

- Class location: Zoom
- Mon/Wed 6:30pm-7:50pm
- Office hours: Mon/Wed 8pm-9pm

We will use Piazza as a social Q&A web service for classrooms. Piazza is a mixture between a wiki and a forum.

No required textbook

If you like to use a reference book, you can consider:

- Maarten van Steen Andrew S. Tanenbaum. Distributed Systems. Freely available at `https://www.distributed-systems.net/index.php/books/distributed-systems-3rd-edition-2017/`

- Paolo Sivilotti, Introduction to Distributed Systems, 2005. Freely available at `http://www.cse.buffalo.edu/~demirbas/CSE586/book.pdf`

Grading

- %20 Assignments (4-5 will be assigned)
- %20 TLA+ project
- %10 Quizzes (via Blackboard; one each week, best 10 quizzes will be used)
- %25 Midterm1 (open book)
- %25 Midterm2 (open book, inclusive of Midterm1 content)

My expectations

- Review the assigned reading/slides
- Work independently and ethically on the assignments
- Ask questions

All assignments are due on the day and time posted

Extra work will NOT be given to improve your grade

Why do I do this?

From student evaluations in Fall 2019

Murat Demirbas has a clear passion for this field and his work, and he brings that passion to class. There was a clear shift in student sentiment from being almost confused by his more direct teaching style to embracing it, and I think the professor did an excellent job adapting to the class's needs. Professor Demirbas treats his students as intellectuals capable of understanding big picture concepts rather than grunts who need to grind away at theory until they've been molded to fit the course needs. If more professors taught this way, I think more students would be encouraged to work with formal research and would learn the skills to be able to understand the important concepts of a class rather than regurgitating lecture slides.

Academic Integrity

Zero tolerance on cheating!

- 0 in the particular assignment/exam for first attempt
- Fail the course on the second attempt
- Students who do share their work with others are as responsible for academic dishonesty as the student receiving the material
- Consult the University Statements on Academic Integrity: `http://www.cse.buffalo.edu/shared/policies/academic.php`.

About Me

- background in distributed systems (OSU, MIT)
- worked on wireless sensor networks for 8 years (OSU, UB)
- cloud computing for the past 10 years
- sabbatical at Microsoft Azure Cosmos DB in 2018/19
- `http://muratbuffalo.blogspot.com/`

Questions ?