

CSE 486/586 Distributed Systems Wrap-up

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586

CSE 486/586 Administrivia

- PA4 due this Friday (5/8) @ 11:59am
- Final: 5/15, Friday, 11:45am – 2:45pm
 - NSC 201
 - Everything
 - **No restroom use** (this quickly becomes chaotic)
 - Multiple choices
- Important things about the final week
 - PA4 scores will be posted before the final.
 - No office hours next week
- Final grading
 - ~ 1 week before the posting deadline (Sat, 5/23)
 - I'm shooting for Wednesday (5/20) for posting, Thursday (5/21) for reviewing, and Friday for finalizing.
 - This will change if there's any delay in grading at the scoring center.

CSE 486/586

2

Building a Distributed System

- *"The number of people who know how to build really solid distributed systems...is about ten"*
 - Scott Shenker, Professor at UC Berkeley
- Are you confident now?
- What were the most interesting topic to you?

CSE 486/586

3

Things We Discussed (Midterm)

- Networking basics (feat. the Internet)
- Failure detection
- Time synchronization
- Logical time & global states
- P2P & DHT
- Reliable multicast
- Consensus basics
- Mutual exclusion & leader election
- RPC

CSE 486/586

4

Things We Discussed

- Transactions & concurrency control
- Replication
- Gossiping
- Distributed file systems
- Paxos
- BFT
- Security

CSE 486/586

5

The Way I See It

- We've learned some of the building blocks & fundamental results...
 - Networking basics, failure detection, logical time, reliable multicast, mutual exclusion, leader election, transactions, concurrency control, replication, gossiping, Paxos, BFT, ...
- ...and how real systems get built using those...
 - P2P, DHT, Dynamo, ...
- ...and also got some experience in building/using the fundamental building blocks...
 - Ordered multicast for messaging, a DHT, and a replicated key-value storage

CSE 486/586

6


Distributed Systems 10 Questions

- Course goal: answering **10 questions on distributed systems**
 - At the end of the semester, if you can answer only 10 questions about distributed systems, you'll probably get an A.
 - Easy enough!
- What are those questions?
 - Organized in 6 themes
 - 1~2 questions in each theme
 - A few (or several) lectures to answer each question

CSE 486/586

7

Theme 1: Hint



CSE 486/586

8

Theme 1: Communications

- Q1: **how do you talk to another machine?**
 - A: Networking basics
 - Know how to use socket now?
- Q2: **how do you talk to multiple machines at once?**
 - A: Multicast
 - What is "reliable multicast"?
 - What orderings are there for ordered multicast?
- Q3: **can you call a function/method/procedure running in another machine?**
 - A: RPC
 - What is a stub compiler (generator)?

CSE 486/586

9

Theme 2: Hint



CSE 486/586

10

Theme 2: Concurrency

- Q4: **how do you control access to shared resources?**
 - A: Distributed mutual exclusion, leader election, etc.
 - Ring election? Modified ring election? Bully algorithm?

CSE 486/586

11

Theme 3: Hint



CSE 486/586

12


Theme 3: Consensus

- Q5: **how do multiple machines reach an agreement?**
 - A: it's impossible! (the FLT result), but algorithms do exist that get around the impossibility (Paxos, BFT, etc.)
 - What are the phases for Paxos?

CSE 486/586

13

Theme 4: Hint



CSE 486/586

14


Theme 4: Storage Management

- Q6: **how do you locate where things are and access them?**
 - A: DHT, distributed file systems, etc.
 - Consistent hashing?

CSE 486/586

15

Theme 5: Hint



CSE 486/586

16

Theme 5: Non-Byzantine Failures

- Q7: **how do you know if a machine has failed?**
 - A: Failure detection
 - What is the fundamental limit of a failure detector?
- Q8: **how do you program your system to operate continually even under failures?**
 - A: Replication, gossiping
 - Linearizability? Sequential consistency? One-copy serializability?

CSE 486/586

17

Theme 6: Hint



CSE 486/586

18

Theme 6: Byzantine Failures

- Q9: **how do you deal with attackers?**
 - A: Security
 - What is a digital certificate?
- Q10: **what if some machines malfunction?**
 - A: Byzantine fault tolerance
 - To tolerate f faulty nodes, how many nodes do we need in total?

CSE 486/586

19

Acknowledgements

- These slides contain material developed and copyrighted by Indranil Gupta (UIUC).

CSE 486/586

20