

CSE 486/586: Distributed Systems

Mid-Semester Review

Ethan Blanton

Department of Computer Science and Engineering
University at Buffalo

Topics to date

So far, we have covered:

- Internet architecture
- Android development (will not directly appear on exam)
- Failure detectors
- Time
 - Physical clocks, logical clocks, causal ordering
- Global states (snapshots)
- Reliable multicast
- Gossiping
- Peer-to-Peer architectures
 - Napster, Gnutella, KaZaA, BitTorrent

The Big Picture

In addition to individual topics, you should know:

- **“It depends on the application.”**
 - What protocol?
 - How should it be configured?
 - *etc.*
- **The Internet is indifferent to your suffering.**
 - Messages can be lost, arbitrarily delayed, duplicated, reordered, ...
 - *Many of the big problems in DS come from asking: What happens if this message is lost? Delayed? Reordered?*

Programming Assignments

I'm not going to ask you (directly) about your PAs.

Your PAs evaluate your understanding of specific implementation details.

I *might* ask you about concepts used in your PAs.

Internet Architecture

You should understand:

- The **model of the Internet** as we have discussed it
- Roughly where the various responsibilities lie
 - (Naming, routing, reliability, *etc.*)
- The Internet design goals
- The Internet protocols' **approach to reliability**

Failure Detectors

You should understand:

- Our model (**crash-stop**) of failures
- The advantages and disadvantages of:
 - **Heartbeating** detectors
 - **Ping** detectors
 - **Gossip** detectors
- The limitations of **asynchronous systems**

Physical Time

You should understand:

- What we mean by **physical clocks**
- The notion of an ideal time, approximated by **UTC**
- Time **synchronization** protocols
- Clock **skew** vs. **drift**

Logical Time

You should understand:

- Logical clocks (**Lamport** and **vector**, *etc.*)
- **Causal ordering**
- Advantages and disadvantages over physical clocks

Global States

Closely related to logical time.

You should understand:

- The snapshot algorithm
- **Consistent** states/cuts/snapshots
- The implications of **FIFO ordering**

Reliable Multicast

You should understand:

- B-multicast
- R-multicast
- The distinction between **receive** and **deliver**
- Ordered multicast (**causal**, **FIFO**, **total**)
- ISIS

Gossiping

You should understand:

- The advantages and disadvantages of:
 - Gossiping vs. B-multicast
 - Gossiping vs. R-multicast
- The **cost/latency/reliability tradeoff** of b and c
- Gossiping as a failure detector

Peer-to-Peer Architectures

You should understand:

- The spectrum of architectures from **centralized** to **fully distributed**
- Where the protocols that we discussed fit in
- What each protocol **does and does not solve**
 - In particular, the differences between **BitTorrent** and the others
- The free-rider problem

Procedures

The midterm is **closed book**.

You **may not** consult notes during the exam.

Use blue or black ink.

Do not bring:

- Cell phones
- Laptops
- Wearables

(If you have a phone/laptop in your bag, that's fine ...
but I don't want to see or hear it!)