

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Oct. 12, 2023

Agenda for Today

- Announcements and reminders
- Research project
- Background on datacenter networking
- eRPC discussion
- PRISM discussion

Next Week

- Talk on Monday (10/16): “Can Great Programmers Be Taught?”
 - 11:00 am – 12:00 pm in CSE 1242
 - John Ousterhout
- No class on Tuesday (10/17)
 - Use the time to work on your project proposal!
- Different office hours next week:
 - Monday (10/16), 12:30-1:30 pm
 - No office hours on Wednesday
 - Thursday (10/19), 2:30-3:30 pm
 - These are a good opportunity to discuss your ideas for your research project

Deadlines

- Course Feedback #FinAid
 - Due Friday 10/13
- Warm-up assignment
 - Due Monday 10/16 at 11:59 pm
- Research project proposal
 - Due Monday 10/23 at 11:59 pm


Research Project

- Original research project
 - Ask a new question or propose a new idea
- Related to OS and/or datacenters
- How to pick a topic?
 - Continue an existing research project
 - Propose your own idea
 - See Canvas for suggestions
- Project proposal
 - 1 page (+ references)
 - Describe the problem, expected solution, and related work
 - List at least 3 milestones and what hardware you need
- See Canvas for details

Datacenter Networking


Incast

- Incast: when a server receives incoming traffic from many senders at once (e.g., 10s-100s)
- Often triggered by one server requesting data from many others
- What problems can this cause?
 - Dropped packets at the switch
 - Queueing delays



Priority Flow Control (PFC)

- RDMA was designed to operate over a lossless network
- How does RoCE prevent packet drops with Ethernet?
 - PAUSE frames
- Drawbacks of PFC
 - Head-of-line blocking
 - Deadlocks
 - Etc.



eRPC Discussion
PRISM Discussion