

# CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 1, 2022

## Agenda for Today

- Announcements
- SmartNICs overview
- iPipe discussion
- nanoPU discussion

# Announcements


- Discussions on Canvas
  - A place to ask questions about common hardware or software
  - Help out your peers
  - Notes about how to save state on CloudLab between experiments

# SmartNICs

# Incentives to Offload

- End of Moore's Law
  - Can we increase compute capacity with accelerators?
- Short tasks (e.g., 1-2  $\mu$ s GET in a key-value store)
  - Can we reduce software overheads?
- Increasing network speeds
  - How can we support high bandwidth links?


SmartNICs!




# Different Varieties of SmartNICs

- ASIC-based SmartNICs
  - Limited flexibility (e.g., match-action tables)
  - Best performance
- FPGA-based SmartNICs
- Multicore SoC-based SmartNICs → iPipe
- Codesign the NIC and CPU → nanoPU

AccelNet  
→



(b) Azure SmartNIC Gen2, 50GbE w/ on-board NIC  
FPGA-based SmartNIC



SoC-based SmartNIC

# Research Questions Raised by SmartNICs

- What tasks should we offload to SmartNICs?
  - OS-level functionality: checksums, virtualization, transport layer, scheduling
  - App-level functionality: serialization, encryption, compression, part of an app, complete app
- How dynamic should this be?
- How should the SmartNIC and CPU communicate?
  - DMA
  - RDMA
  - Directly connected

**iPipe Discussion**  
**nanoPU Discussion**