

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Oct. 4, 2022

Agenda for Today

- Reminders
- Introduction to RDMA and RPCs
- FaRM discussion
- Where do research ideas come from?

Reminders

- Leading a discussion
 - See notes on Canvas
- First Day Survey #FinAid
 - Due Friday 10/7
- Warm-up assignment
 - Due Tuesday 10/11 at 11:59 pm


RDMA and RPCs

The Shift Towards Storing Data in Memory

- Disk is a poor fit for modern datacenter applications
 - Disk is much slower to access than memory (10 ms vs. 100 ns)
 - Datacenter workloads require random access
- RAM is becoming much cheaper
- Feasible to store a significant fraction (or all of) your app's data in memory, distributed across a cluster



- 500 GB disk
- 800 Mbps
- 10 ms latency



Lower latency,
higher throughput

- 20 machines, each with:
 - 128 GB DRAM
 - 40 Gbps
 - 10 μ s latency

How Should Programs Access Remote Memory?

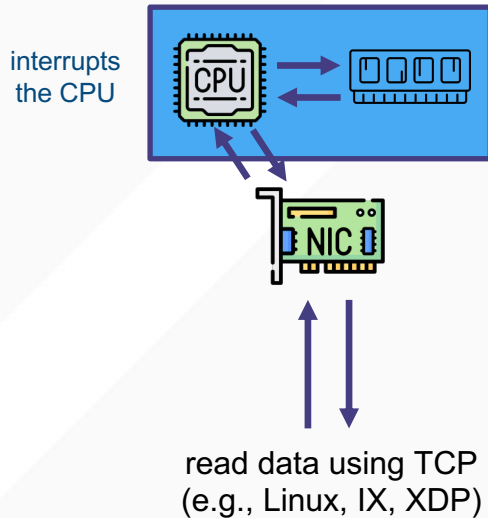
- Access data one word at a time, similar to local memory
 - `movl remote_addr %eax`
- Access a chunk of data at once (e.g., 64 bytes, 1 KB)
 - FaRM
- Access multiple dependent chunks of data at once
 - PRISM
- Execute a function on the remote server via RPC
 - eRPC

← today


← Thursday

CPU-Based Memory Access vs. RDMA

- Remote Direct Memory Access
 - Access memory directly from the NIC



“RPC” or
“two-sided RDMA”



“one-sided RDMA”

RDMA – An Old Technology

- First proposed in 1993
- Used in super computers (HPC) for many years
- Relied on Infiniband in the 2000s
 - Lossless network
 - Expensive
- RoCE (~2010)
 - RDMA over Converged Ethernet, pronounced “rocky”
 - Provides a reliable network and enables RDMA over regular Ethernet
 - Cheaper than Infiniband
 - Made it easier to adopt RDMA in datacenters

FaRM Discussion

Where Do Research Ideas Come From?

- How do researchers come up with the problems, solutions, or observations in their papers?
- Why does this matter?
 - Helps to understand the perspective of authors
 - Inspiration for coming up with your own research ideas