

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Oct. 3, 2023

Agenda for Today

- Announcements
- Introduction to multicore and heterogeneity
- Multikernel discussion

Announcements

- Enrollment
 - The class is close to full
 - If you hope to join, continue doing the course work
 - Give me your email after class
- Office hours
 - Wednesday 1-2 pm or by appointment in CSE 3130
- Piazza
 - Linked from course website
 - Intended for students to assist each other


Announcements, continued

- Warm-up assignment
 - Posted on the course website
 - Individual assignment
 - Due Monday 10/16 at 11:59 pm (after the add deadline)
 - Don't leave it until the last minute

Introduction to Multicore and Heterogeneity

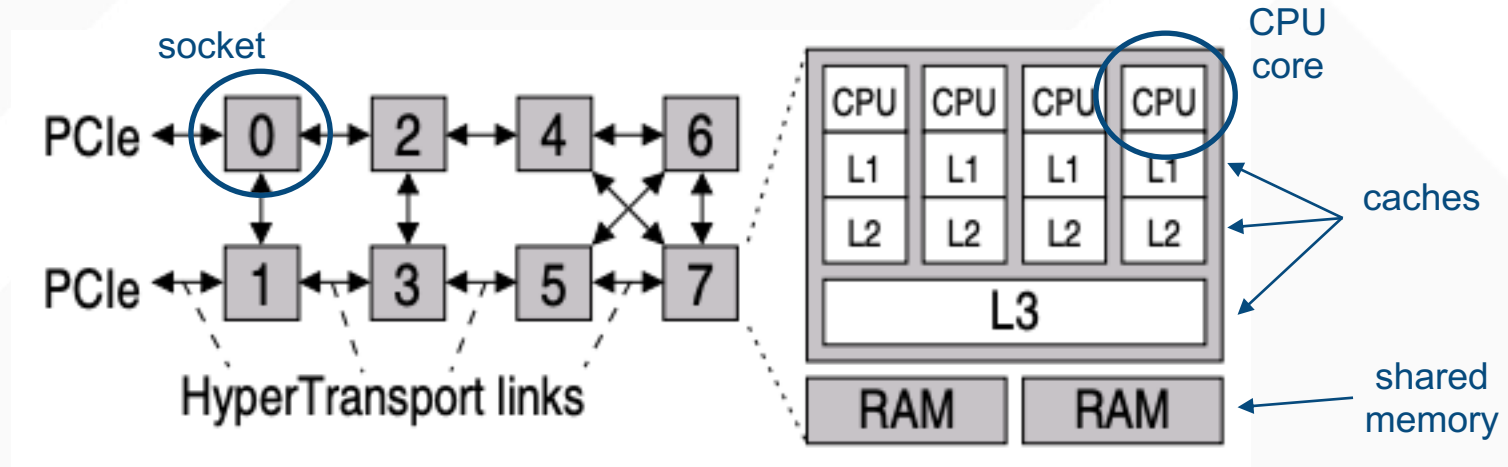
The Rise of Multicore

- First multicore CPUs?
 - IBM's Power4 in 2001 had 2 cores
- Driven by the end of Moore's Law



Multicore Architectures

- Cores grouped into sockets
 - Also referred to as “multi-core processors” or “NUMA nodes”
 - Interconnect in between them
- Cache coherence
 - Keeping data in separate caches consistent




Challenges of Multicore

- How do you design **applications** that scale well across many cores?
- How do you build an **OS** that scales well? ← today
 - How do you design the application-OS **interface** to enable scalability?
“The Scalable Commutativity Rule” [SOSP ‘13]
 - How do you build a scalable **network stack**? ← Thursday
 - How do you make **CPU scheduling** scalable? ← later this quarter
 - How do you scale **memory management**? ← later this quarter

Heterogeneity

- Heterogeneous processors
 - Different power consumption ← common in mobile and tablets
 - E.g., ARM's big.LITTLE announced in 2011
 - Different ISAs
- Other types of heterogeneity:
 - GPUs
 - SmartNICs
 - FPGAs
 - Accelerators

common in
datacenters



Challenges of Heterogeneity

- May not be cache coherent
- Different memory layouts
- Different models of computation (e.g., GPU vs. CPU)

Multikernel Discussion