UC San Diego

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 7, 2023

UC San Diego

Agenda for Today

- Introduction to the Datacenter Tax
- Accelerometer discussion

UC San Diego

Datacenter Tax

Profiling a warehouse-scale computer

- “Profiling a warehouse-scale computer”, ISCA 2015
  - Analyzed jobs on 20,000+ Google machines over 3 years
  - Measured how CPU cycles are spent
  - Conducted a microarchitectural analysis
- Coined the term “datacenter tax”

How CPU Cycles were Spent (in 2015)

- 22-27% of CPU cycles spent in the “datacenter tax”
  - “functions unique to performing computation that transcends a single machine”
  - “prime candidates for hardware acceleration”
- Almost 20% of CPU cycles spent in the kernel

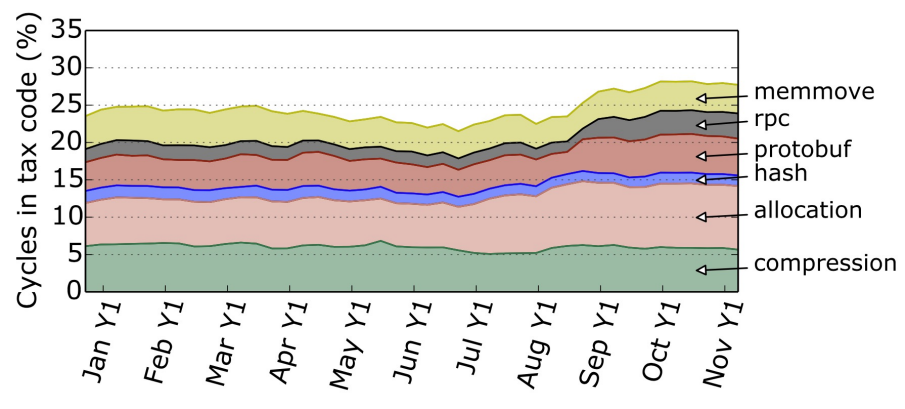

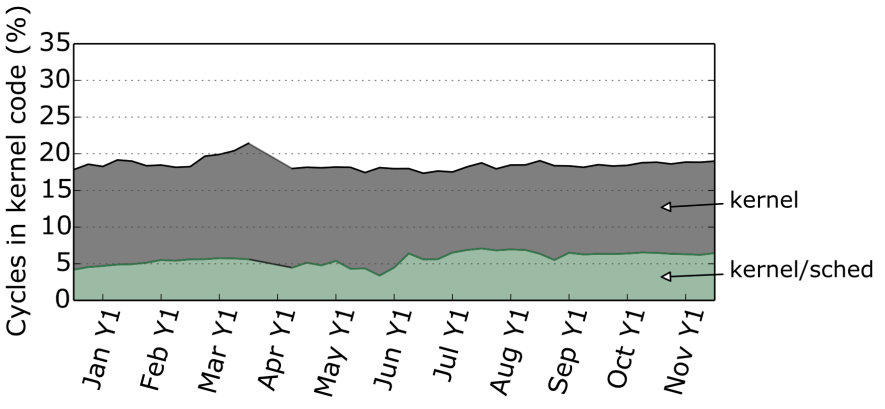

“Profiling a warehouse-scale computer”, ISCA 2015

Microarchitectural Analysis

- Which parts of the CPU microarchitecture are typically the bottleneck?
  - Use performance counters to measure this
- Findings:
  - Instruction cache problems
  - Low IPC (instructions per cycle)
  - …

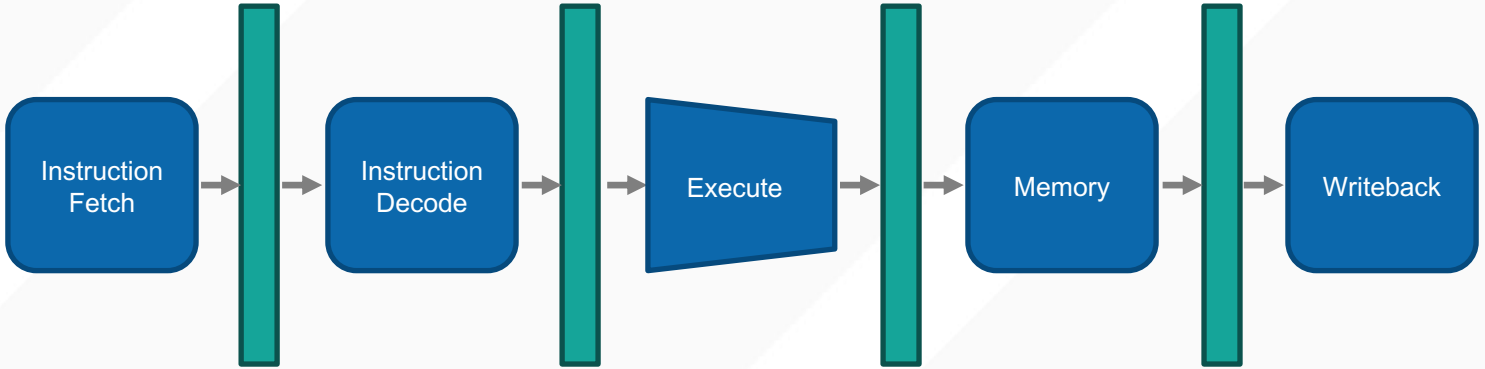

Hardware Offloads

- How much will performance improve with offloads?
  - Accelerometer [ASPLOS ‘20]
- SmartNICs
  - AccelNet [NSDI ‘18]
  - iPipe [SIGCOMM ‘19]
- GPUs and TPUs
  - TensorFlow [OSDI ‘16]
- FPGAs
  - Coyote [OSDI ‘20]

UC San Diego

Accelerometer Discussion