

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 15, 2022

Agenda for Today

- Research Tips: Benchmarking
- Disaggregation overview
- LegoOS discussion

Research Tips: Benchmarking

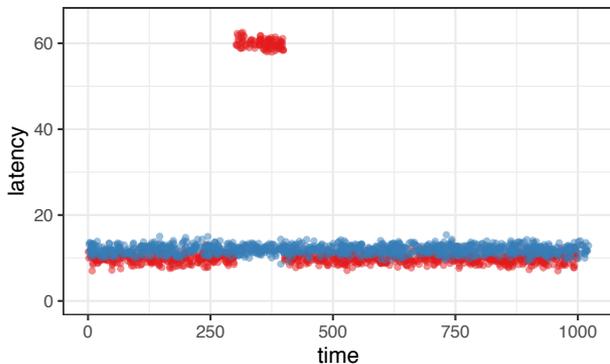
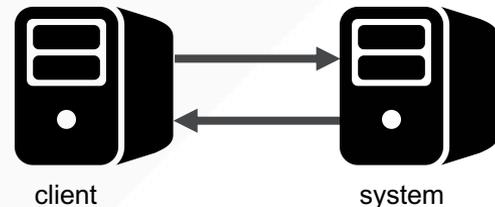
Research Tips: Benchmarking

- Strategies for effectively measuring systems performance
- Why is this important?
 - Reading papers:
 - Do the experiments measure/show what they intend to?
 - Do the results make sense?
 - Benchmarking in your own research projects:
 - Conduct accurate and meaningful measurements

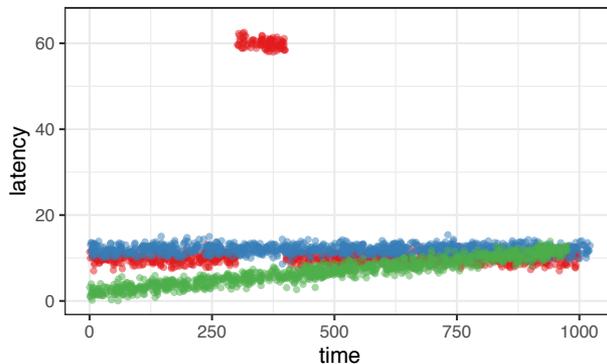
Benchmarking (Fake) Example

- Comparing the latency of different systems, A and B
 - Which system has better latency?
- Suppose we measure them and find:

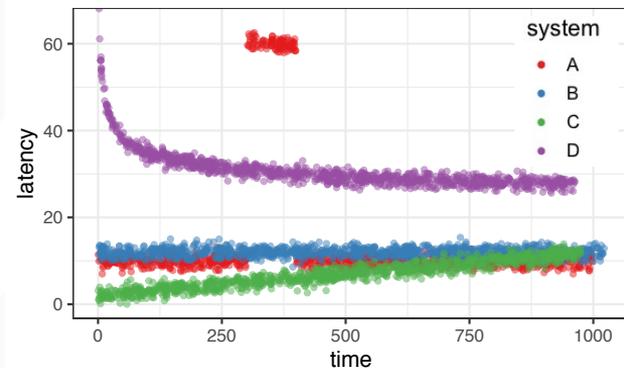
System	A	B	C	D
Average latency	14.7	12.0	6.9	31.0



latency spike – maybe a bug?



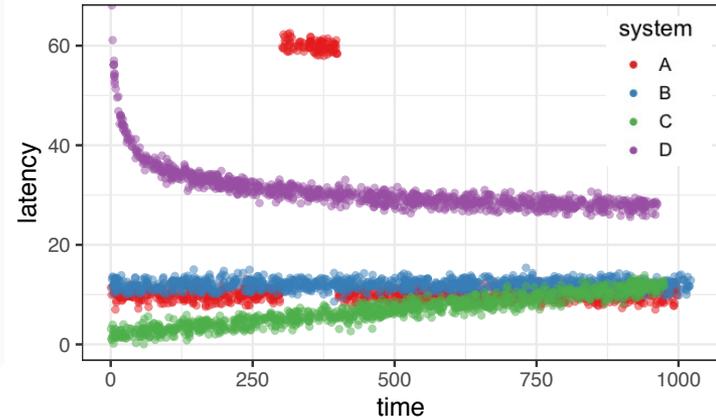
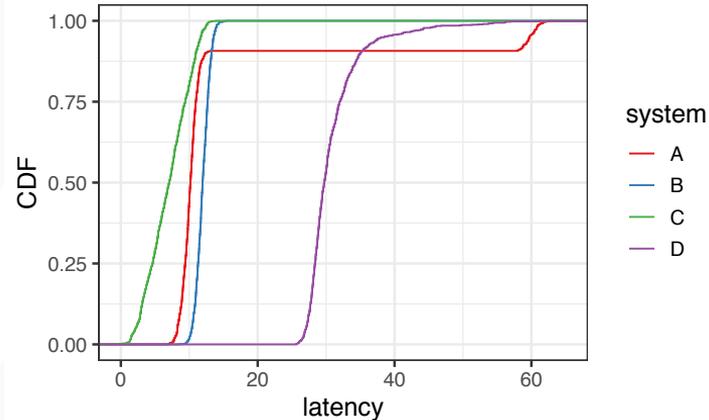
increasing latency – queue building up?



decreasing latency – warm-up period?

Always Start with the Raw Data

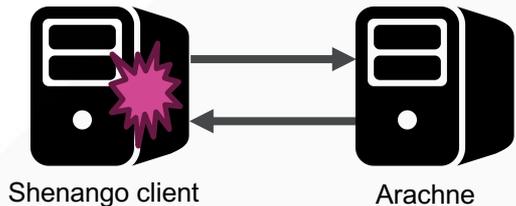
- Start by trying to understand the raw data
 - What is the distribution of it?
 - How does it change over time?
 - Does the data behave as you expected?
- Once you understand the raw data, then you can summarize it
 - Averages, medians, 99th percentile, etc.



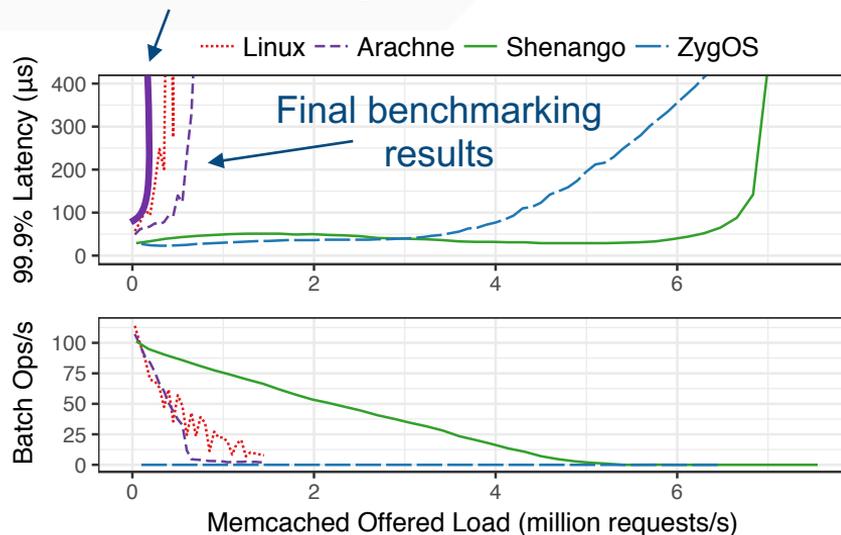
Benchmarking (Real) Example

- Comparing Shenango to Arachne
- Arachne performed poorly! All done?
- But... performance was worse than reported in their paper 🤔
- More measurements: Arachne had a lot more TCP retransmissions
- More measurements: bug in Shenango's TCP stack!

bug in Shenango's
TCP stack!



First benchmarking attempt



Measure One Level Deeper

- Application-level metrics
 - Job completion time, requests per second, etc.
- But measure lower-level statistics too!
 - Network: bandwidth, packets per second
 - CPU: CPU utilization over time, per app, context switches
 - Transport protocol: retransmissions, timeouts, dropped packets
 - Statistics that are specific to your system
 - E.g., in Shenango: core reallocations, work stealing

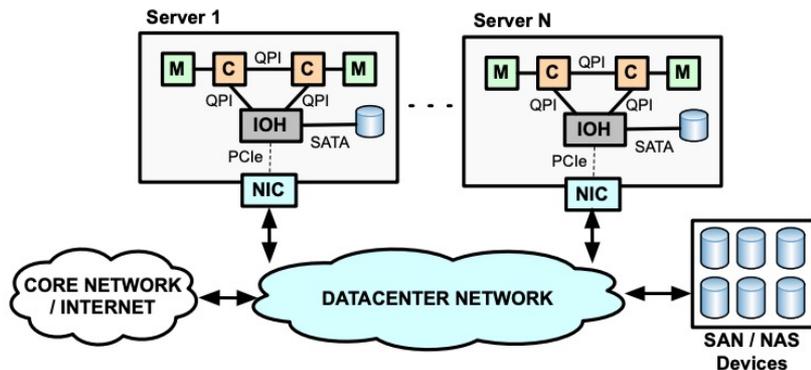
Benchmarking Tips

- Start with the raw data
- Always measure one level deeper
- Want more tips?
 - “Systems Benchmarking Crimes”
 - <https://gernot-heiser.org/benchmarking-crimes.html>
 - “Always Measure One Level Deeper” [CACM 2018]

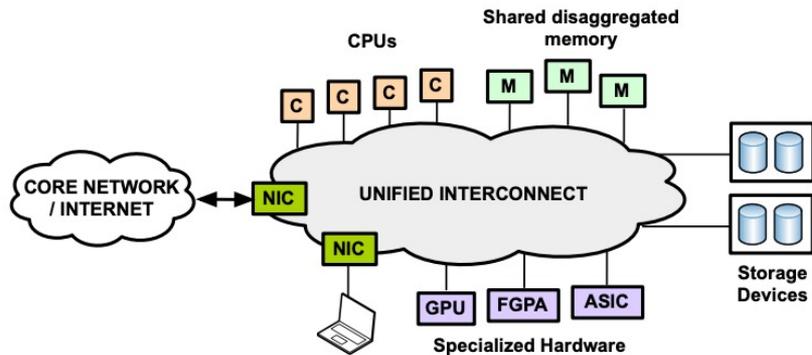
Disaggregation

What is disaggregation?

- Disaggregated datacenter
 - Resources directly connect to the network in “resource blades”
 - No notion of a “server” anymore



(a) Current datacenter

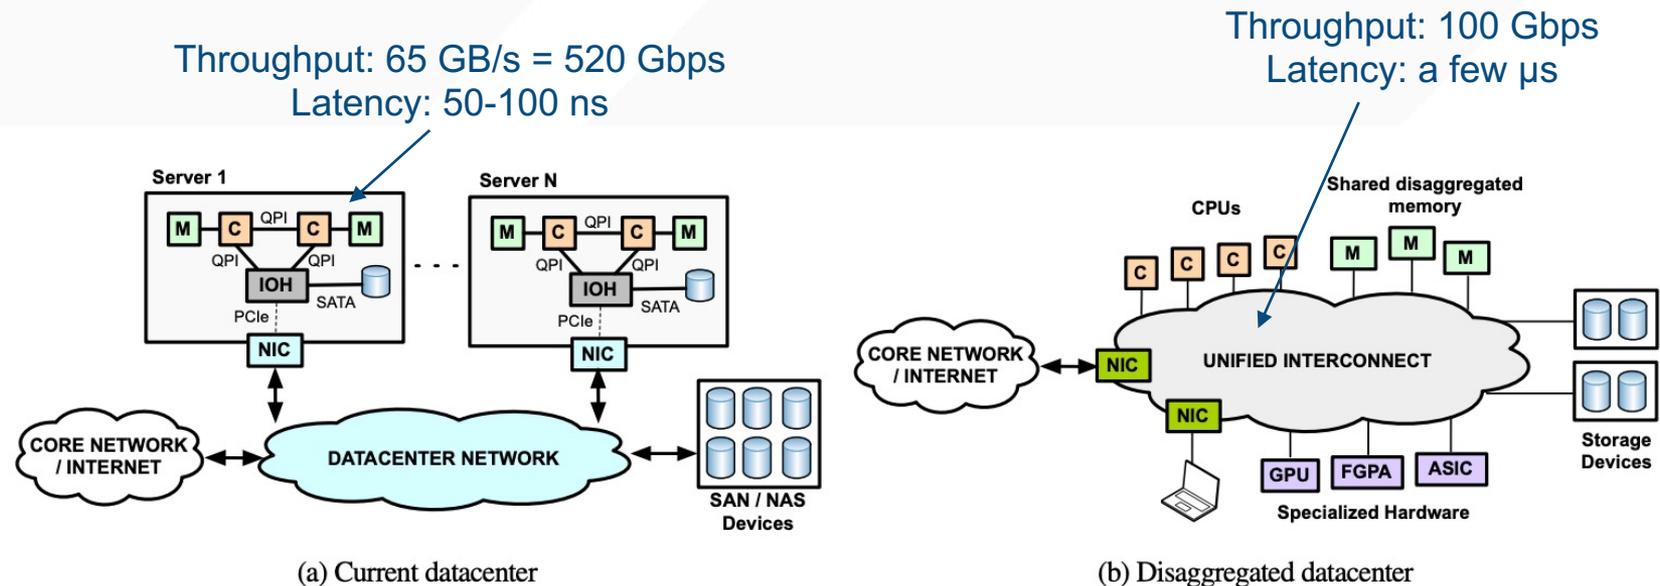


(b) Disaggregated datacenter

“Network Requirements for Resource Disaggregation” [OSDI ‘16]

What makes disaggregation challenging?

- Traditional datacenter networks have much lower throughput and higher latency than server interconnects



“Network Requirements for Resource Disaggregation” [OSDI ‘16]

Hardware for Disaggregation

- Many announcements of disaggregated hardware:
 - SeaMicro
 - Intel Rack Scale Architecture, 2013
 - HP The Machine, announced in 2014
 - UC Berkeley Firebox, 2014

RIP HPE's The Machine product, 2014-2016: We hardly knew ye

Remains of lab experiment, including ReRAM, will be scattered into future gear

 [Chris Mellor](#)

Tue 29 Nov 2016 / 07:28 UTC

[HOME](#) > [NEWS](#) > [IT HARDWARE](#)

AMD kills off SeaMicro server business

Dense server subsidiary shut down after AMD posts a bigger-than-expected loss

April 17, 2015 By: [Max Smolaks](#)  [Comment](#)

But... no widespread adoption yet

Disaggregation-Related Research Questions

- What kind of network is required?
 - “Network Requirements for Resource Disaggregation” [OSDI ‘16]
 - “Pond: CXL-Based Memory Pooling Systems for Cloud Platforms” [ASPLOS ‘23]
- Transparent or not?
 - “AIFM: High-Performance, Application-Integrated Far Memory” [OSDI ‘20]
- How to share disaggregated resources?
 - “MIND: In-Network Memory Management for Disaggregated Data Centers” [SOSP ‘21]

Want to learn more about disaggregation?

- WORDS 2022

- Workshop on Resource Disaggregation and Serverless Computing
- November 17th, 2022
- UCSD and virtual
- Registration is free
- <https://www.wordsworkshop.org/>

Keynote and Invited Industry Talk 9:30-11:00 PT (session chair - Ryan Kosta)

- Keynote: Marcos Aguilera, VMware - "Memory disaggregation: Why it will happen now and bring a surprise impact"

Abstract: Disaggregation refers to moving hardware resources outside the box. Looking at history, we learn that disaggregation succeeds based two key factors: burning issue and technical feasibility. We also learn that, once successful, disaggregation brings an impact that is much broader than originally envisioned. We look at memory disaggregation from this perspective; we make the case why memory disaggregation will finally happen now and discuss some of its potential surprise impact: fluid memory, instantaneous VM migration, and cheap Byzantine Fault Tolerance. The wildly different nature of these applications suggest that memory disaggregation will be a research topic of wide interest in the years to come.

LegoOS Discussion