

CSE 291: Operating Systems in Datacenters

Amy Ousterhout

Nov. 16, 2023

Agenda for Today

- Research Tips: Benchmarking
- FPGAs overview
- Coyote 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? another app?





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, latency
 - 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

Make Benchmarking Easy

- Design and experimentation are iterative processes
- Make it easy to iterate
 - Write scripts to run your experiments
 - Automated and reproducible
 - Run smaller experiments first
 - Debug/test in simpler or faster configurations
 - Parallelize your experiments

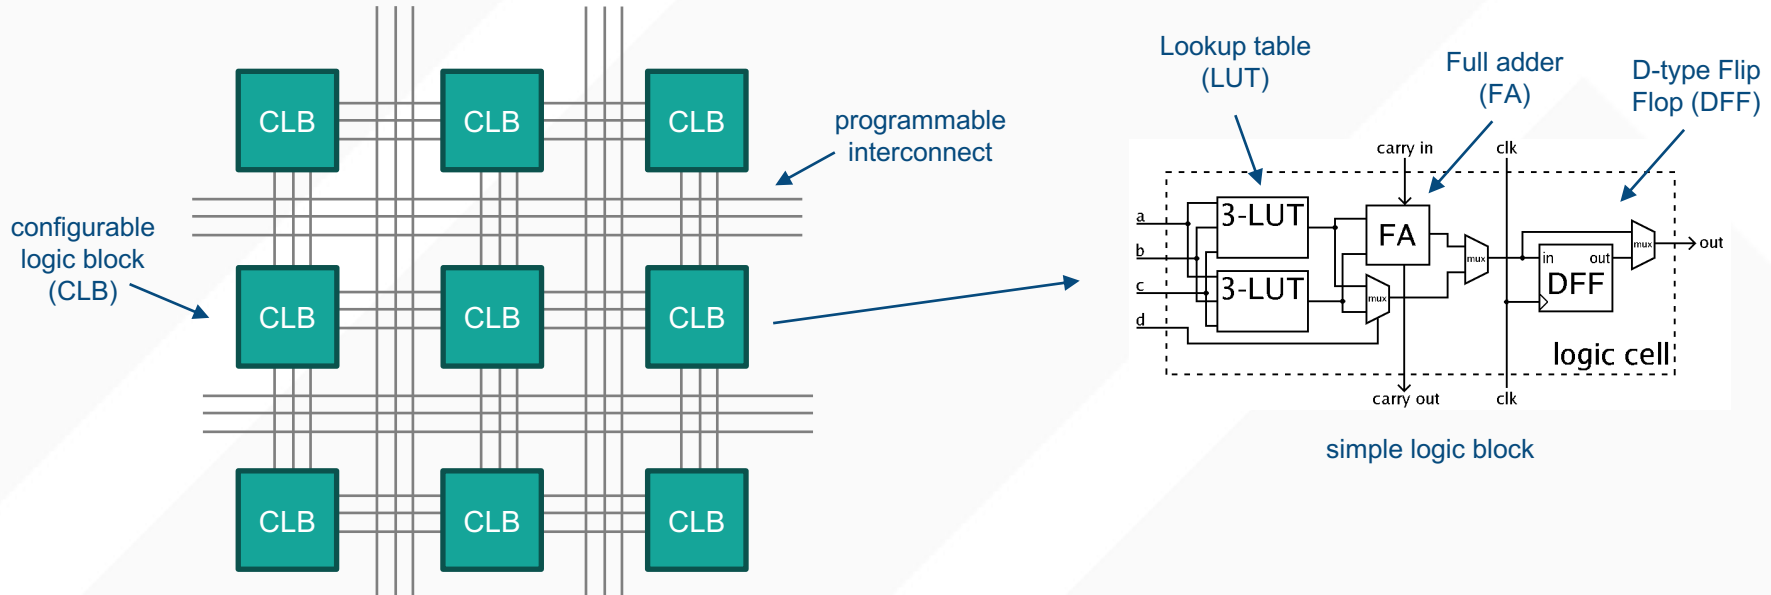
Benchmarking Tips

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

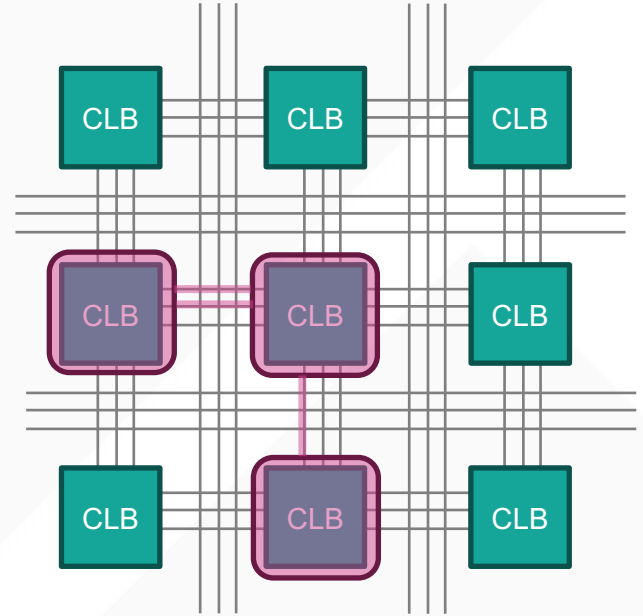
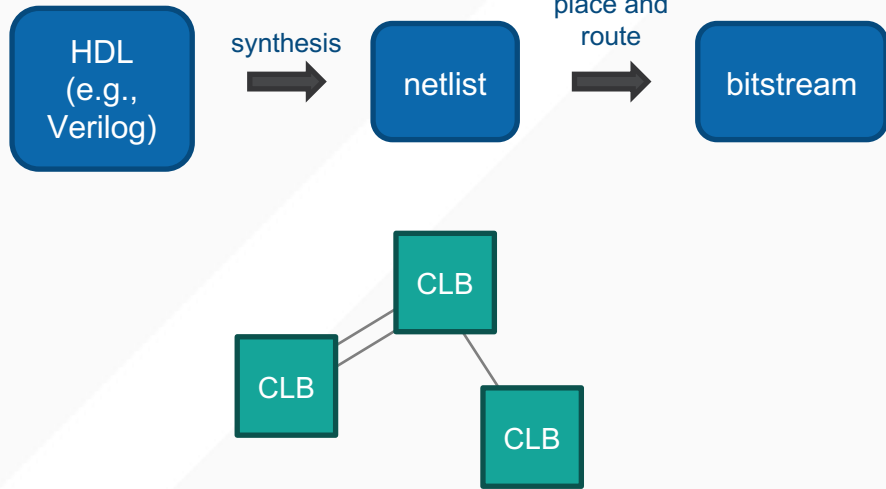
FPGAs

What is an FPGA?

- FPGA: Field Programmable Gate Array
 - Grid of configurable logic blocks (CLBs)
 - Connected by a programmable interconnect
 - RAM, I/O, etc.



How do you program an FPGA?



Coyote Discussion