And, what about debugging for multi-FPGA systems?

Paul Chow

High-Performance Reconfigurable Computing Group
University of Toronto

Outline

- Context
- What we've been thinking
- The Debug Governor
- FFShark
- Pharos
- Simulation
- Final thoughts

CONTEXT

How do you debug something running on this?

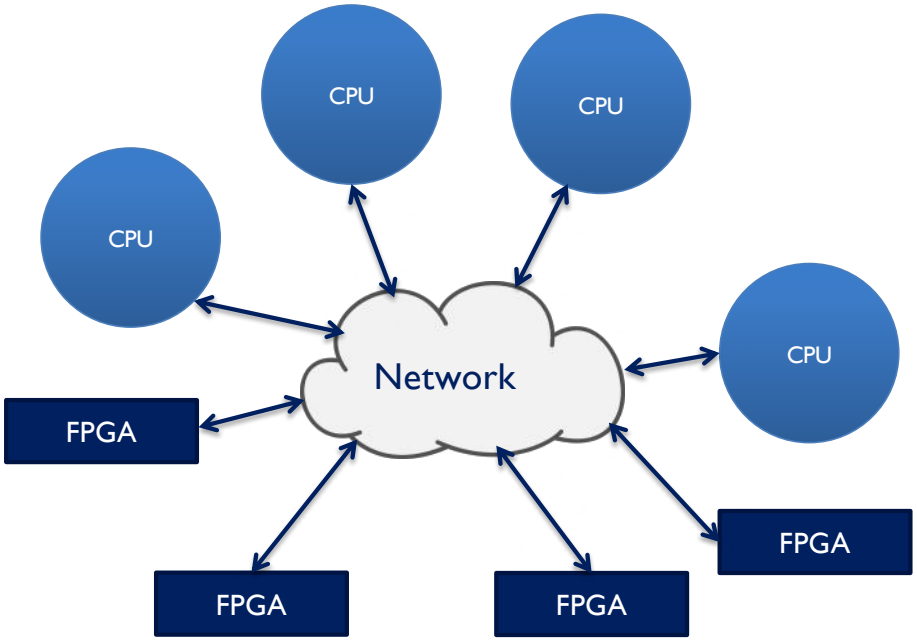

Network-connected heterogeneous platform

- Scalable and elastic

Examples include: Microsoft, cloudFPGA, UofT Galapagos

A bit (actually, a lot) harder than this

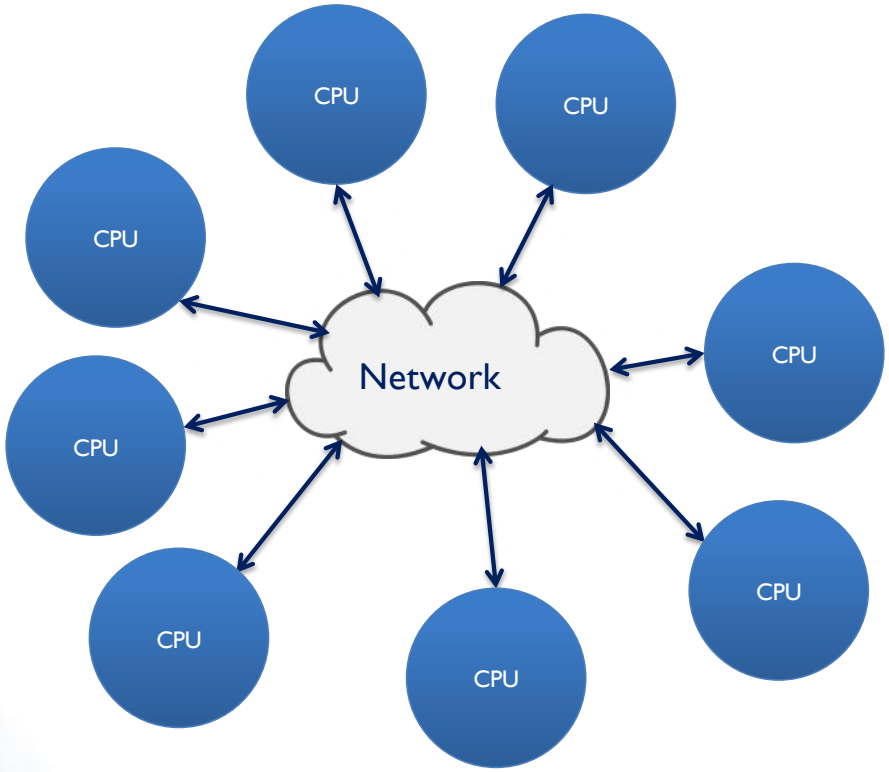

Typical multiprocessor platform

- Often MPI-based, many tools
- And it's software

Back to this

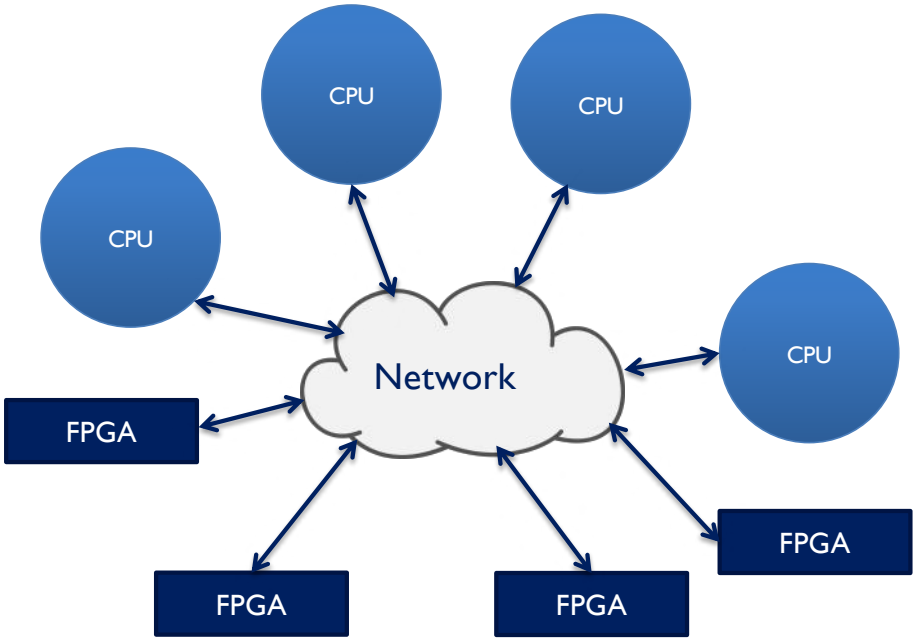

- There are no (standard) tools
- If anything exists, they are custom, one-off, and work on only one platform
- Catapult Flight Data Recorder records interesting state

We need to build an ecosystem of tools that can evolve and be used across multiple heterogeneous platforms

WHAT WE'VE BEEN THINKING

Ideas

- Compare software and hardware environments
- What does software do?
- High and low-level debugging
- Wireshark
- Logging
- Performance debugging

Distributed System Debugging

Software	Hardware
• Highly observable	• Difficult to observe
• Logging "in series" with program code	• Logging "in parallel" (zero runtime cost)
• Underlying OS	• Underlying "shell"
• Breakpoints (maybe)	• No breakpoints
• Some consistent structure	• No consistent structure

Software

- Many new techniques being developed and applied to large software systems
  - Dynamic invariant detection, model inference, declarative specifications
- Hardware doesn't seem ready for that

HIGH AND LOW-LEVEL DEBUGGING

Definition

- Have multiple nodes
- Strategy
  - High-level – Which node exhibits the problem?
  - Low-level – Debug on that node

High-Level Debugging

- Log "API calls" instead of single instructions
- Goal: find location to apply low-level debugger

Visualizations

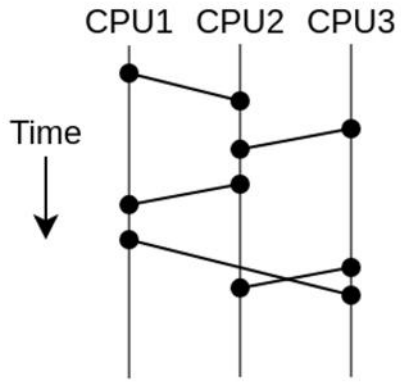

Performance Debugging

	# Calls	Avg. Lat.
`fn1()`	102	10 µs
`fn2()`	4	5 ms
`fn3()`	760	70 µs

Failure Reproduction

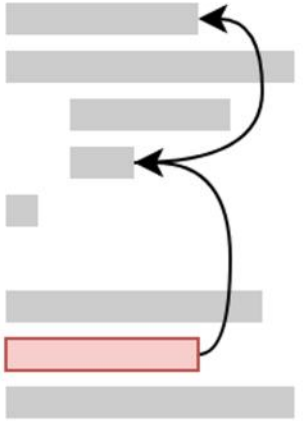

Low-Level Debugging

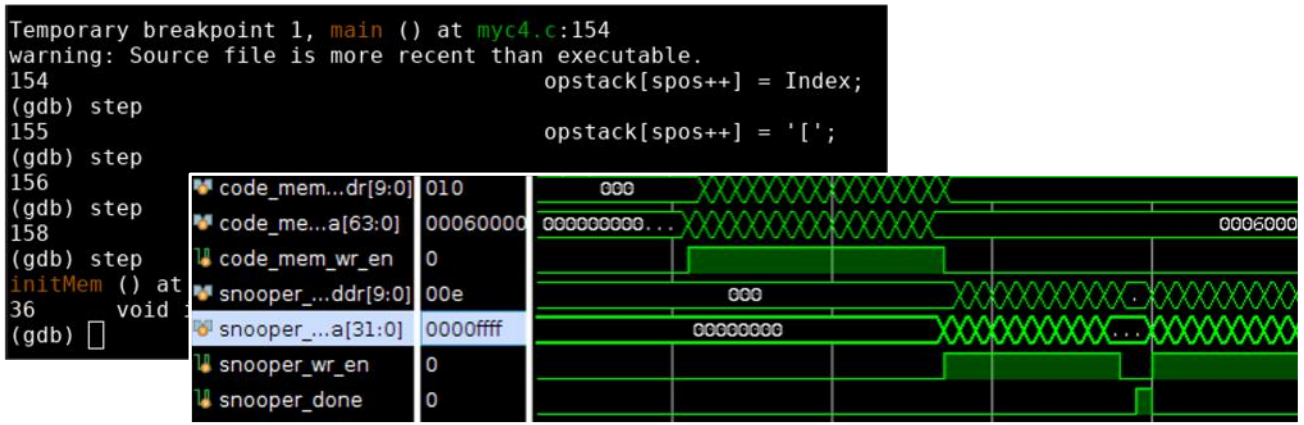

- Software single-stepping, watching state, etc.
- Hardware – SignalTap, ILA
- In a large project, "scrubbing a building with a toothbrush"

ASSUMPTIONS AND OBSERVATIONS

Observation

Many designs are dataflow graphs:

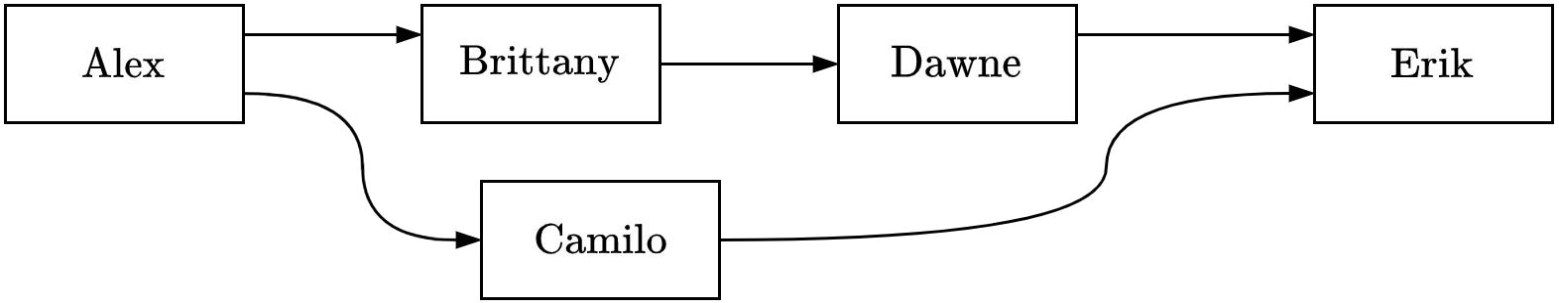

Observation

Many designs are dataflow graphs:

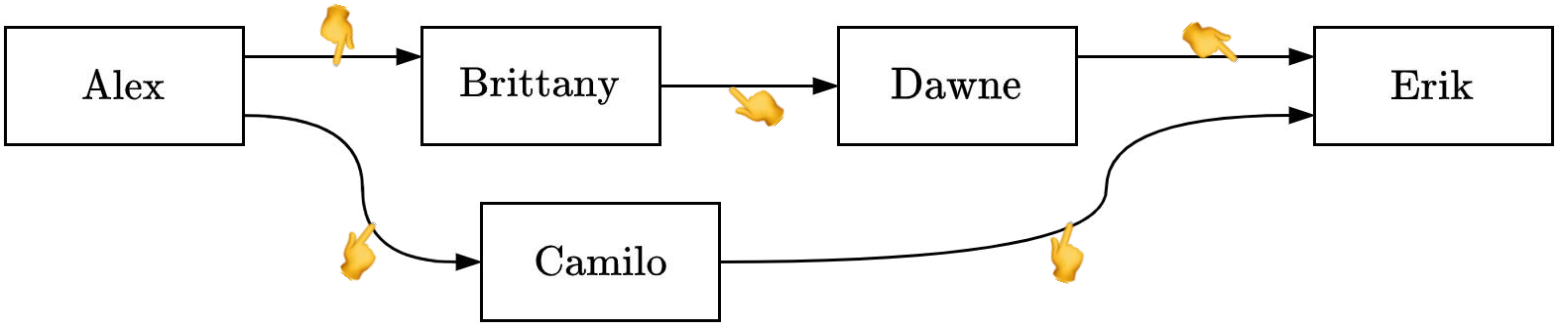

Multiple FPGAs

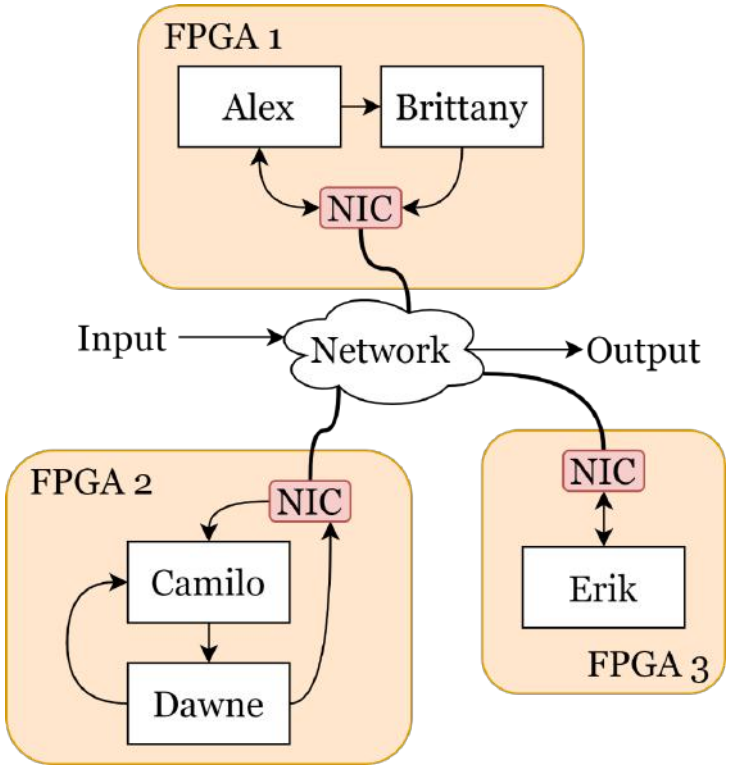

Our Debuggers

- Debug Governors
- FFShark

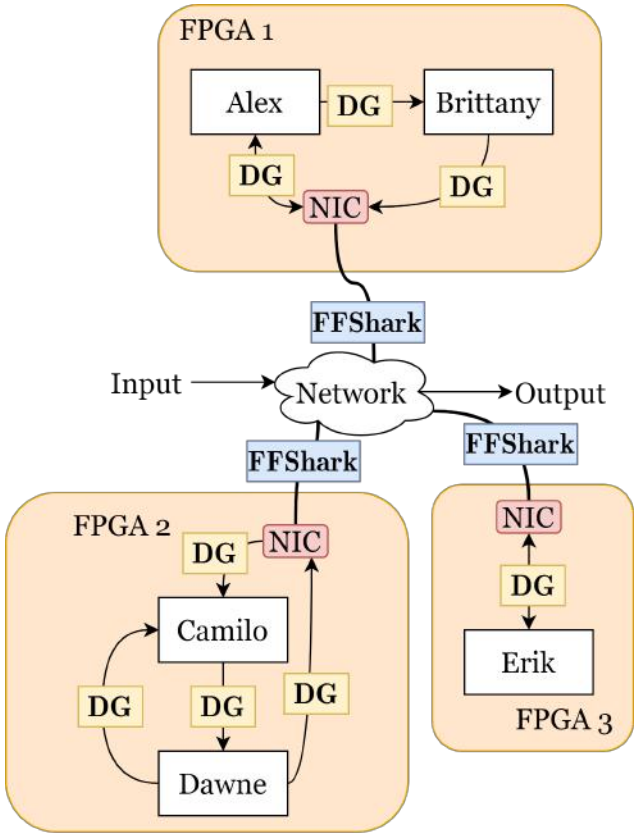

Marco Merlini, Isamu Poy, Paul Chow, FPGA 2021

INTERACTIVE DEBUGGING AT IP BLOCK INTERFACES IN FPGAS

DEBUG GOVERNOR - OVERVIEW

Overview

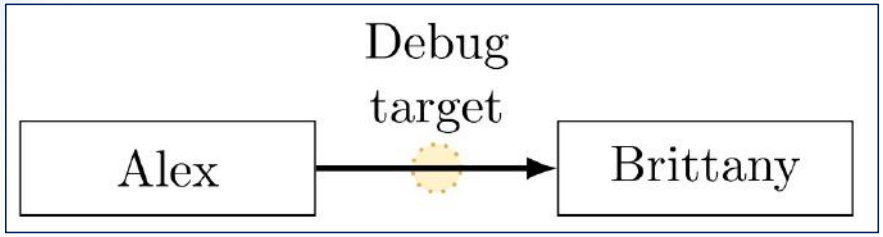

Overview

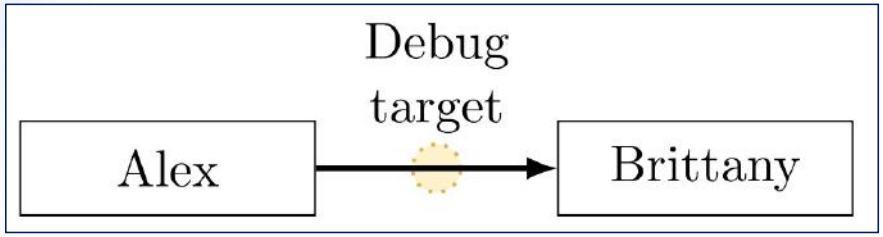
Debug
target
Alex
Brittany

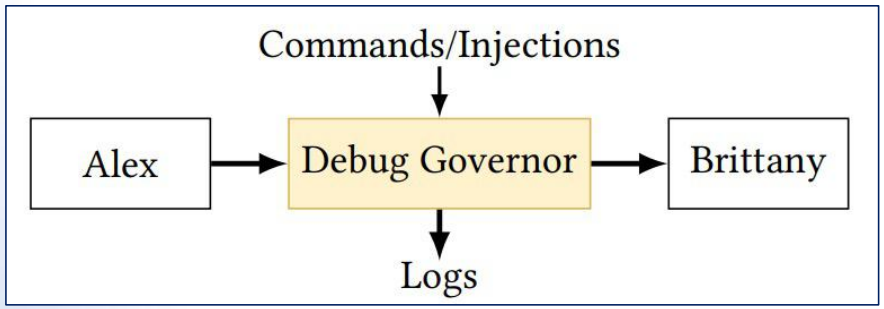
Commands/Injections
Alex
Debug Governor
Brittany
Logs

Overview

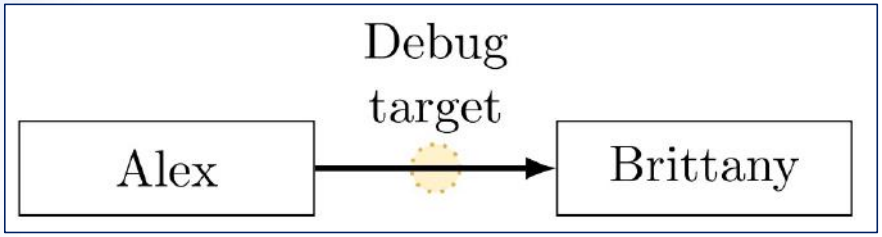

- Pause
  - Alex is prevented from sending data

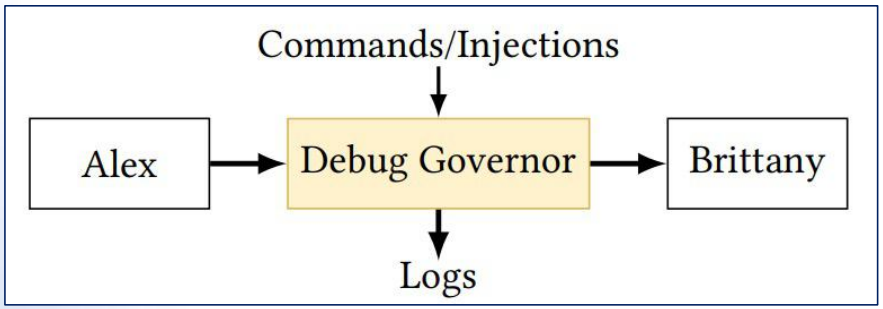

Overview

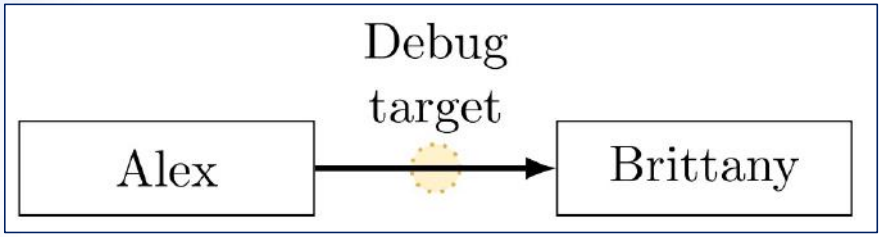

- Pause
  - Alex is prevented from sending data

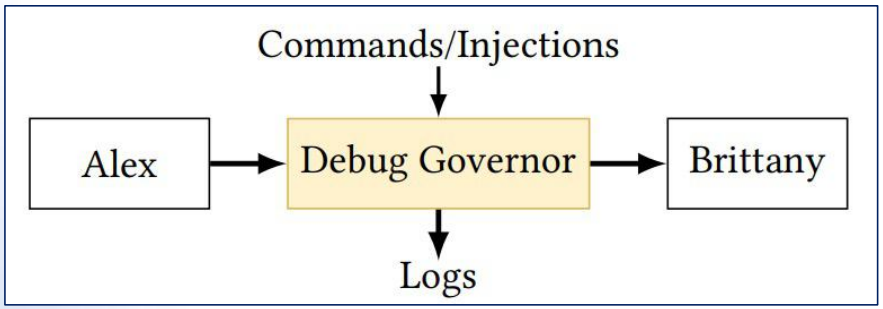

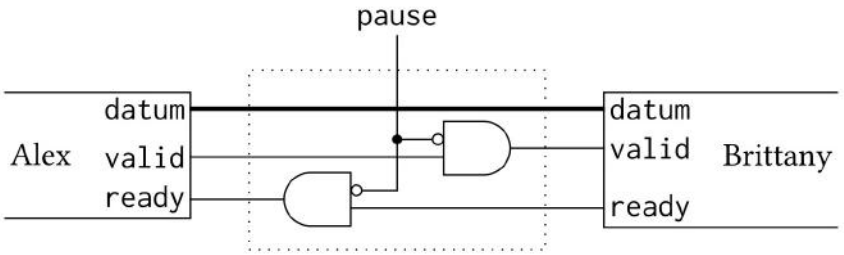

Overview

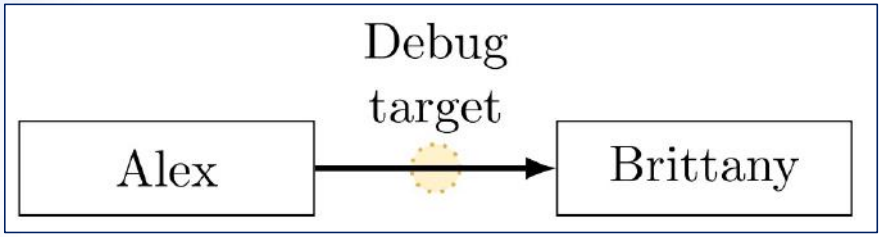

- Pause
- Log
  - Data from Alex are duplicated and sent to the developer.

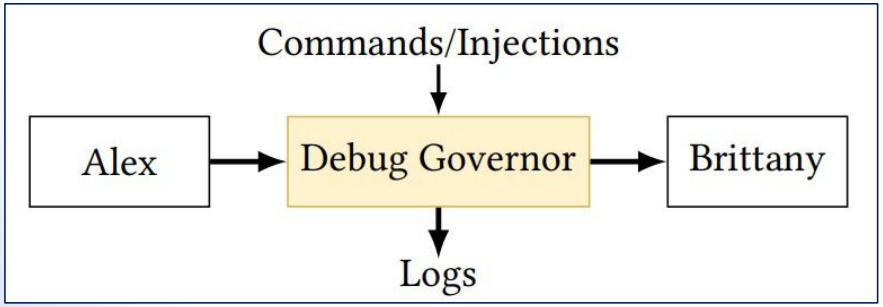

Overview

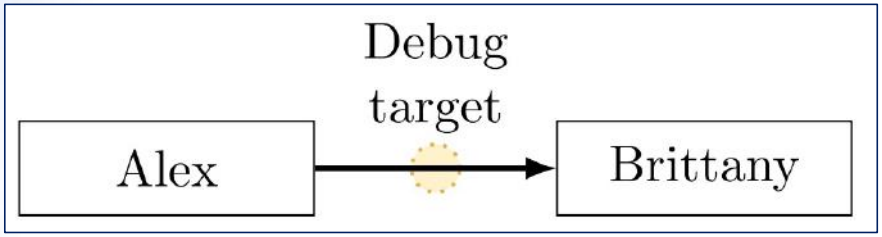

- Pause
- Log
- Drop
  - Alex can send data, but they are not sent to Brittany

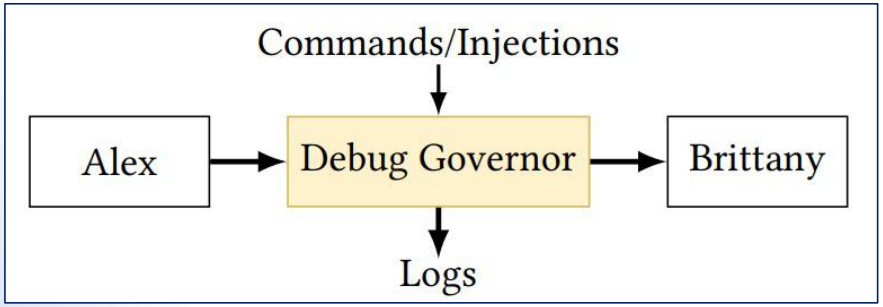

Overview

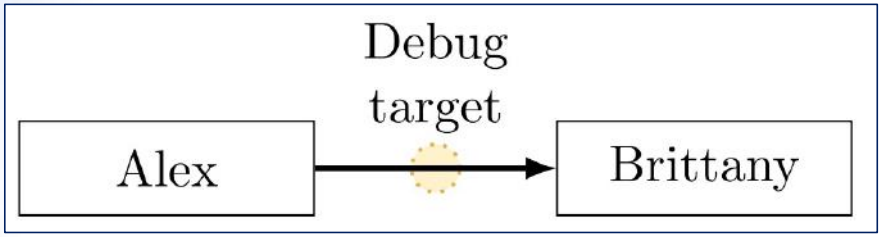

- Pause
- Log
- Drop
- Inject
  - Developer can send data to Brittany.

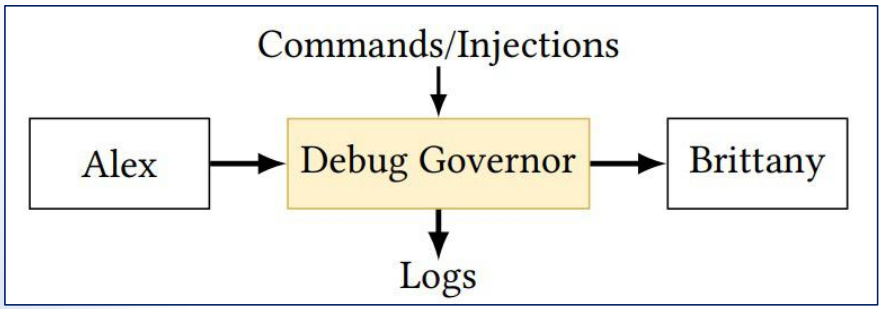

Overview

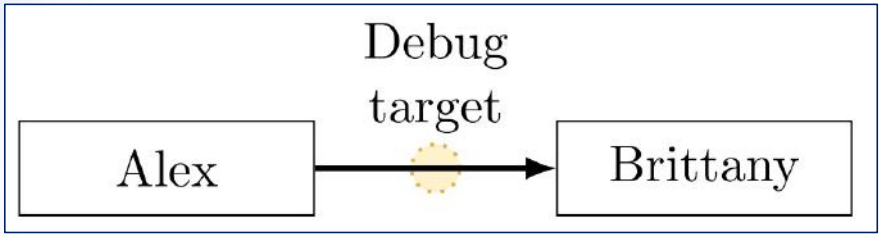

- Pause
- Log
- Drop
- Inject

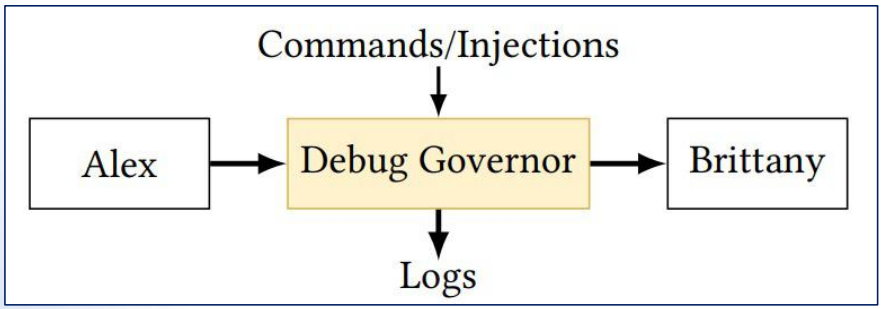

+ Any subset thereof

Single-Stepping

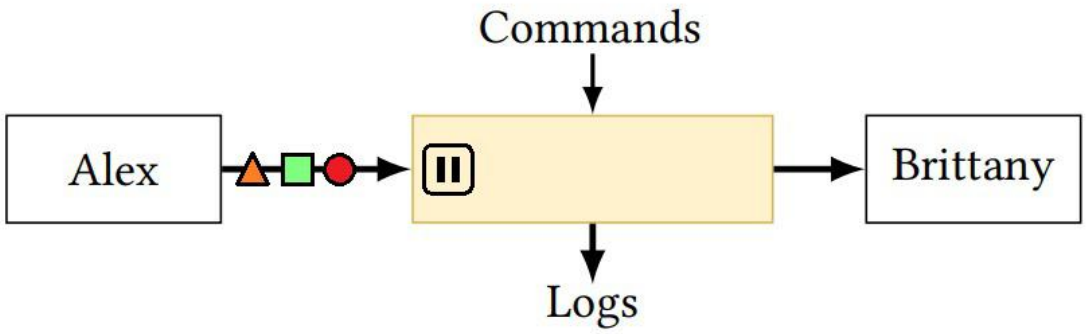

Single-Stepping

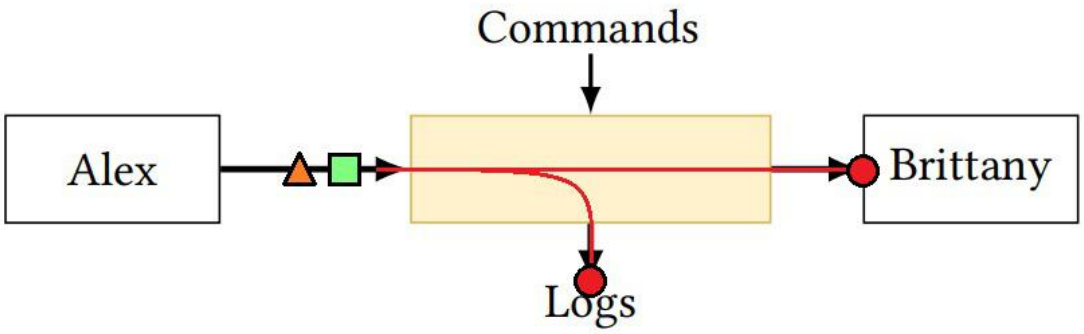

Single-Stepping

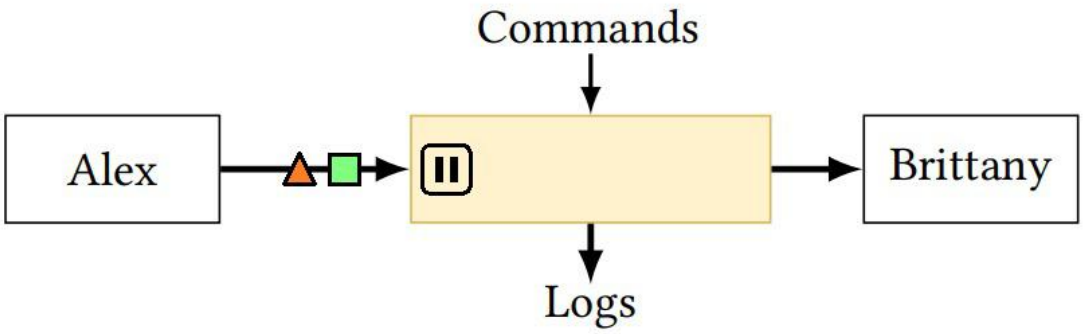

Juan Camilo Vega, Marco Merlini and Paul Chow, FCCM 20

FFSHARK: A 100G FPGA IMPLEMENTATION OF BPF FILTERING FOR WIRESHARK

Network Debugging with Wireshark

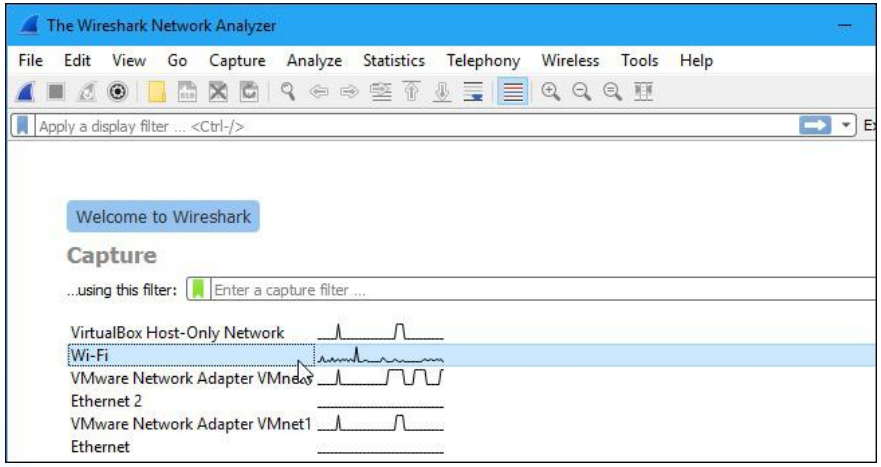

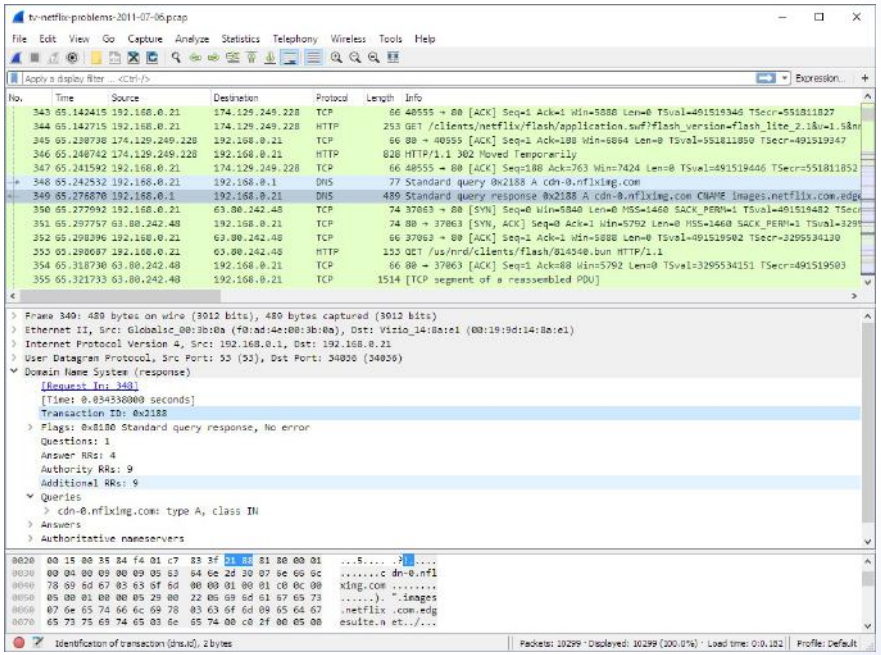

Debugging 10G and 100G networks?

On a CPU, Wireshark is difficult at 10G, and *extremely* difficult at 100G

Example: Intel i9 processor

- Sixteen PCIe 3.0 lanes (8 Gbps each)
- 5 GHz
- 64-bit

Must transfer, filter, and copy accepted packets in 3 cycles/word!

Improving Packet Capture Performance

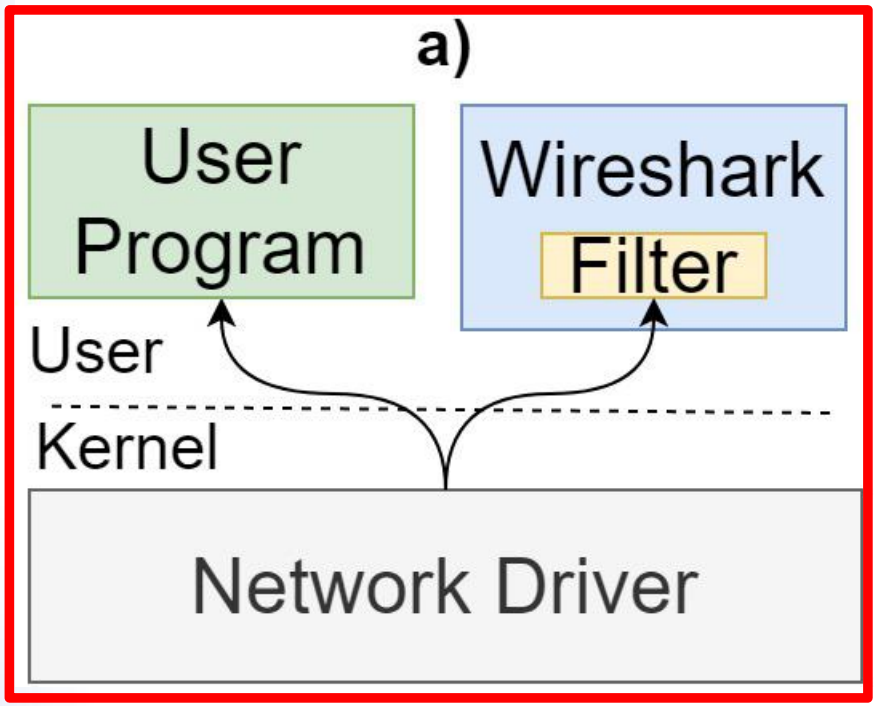

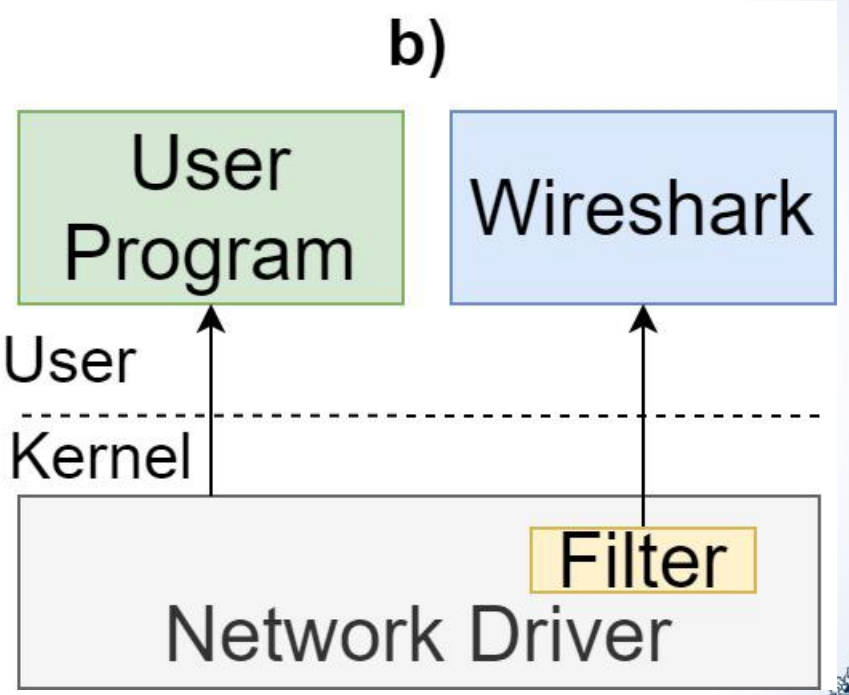

Improving Packet Capture Performance

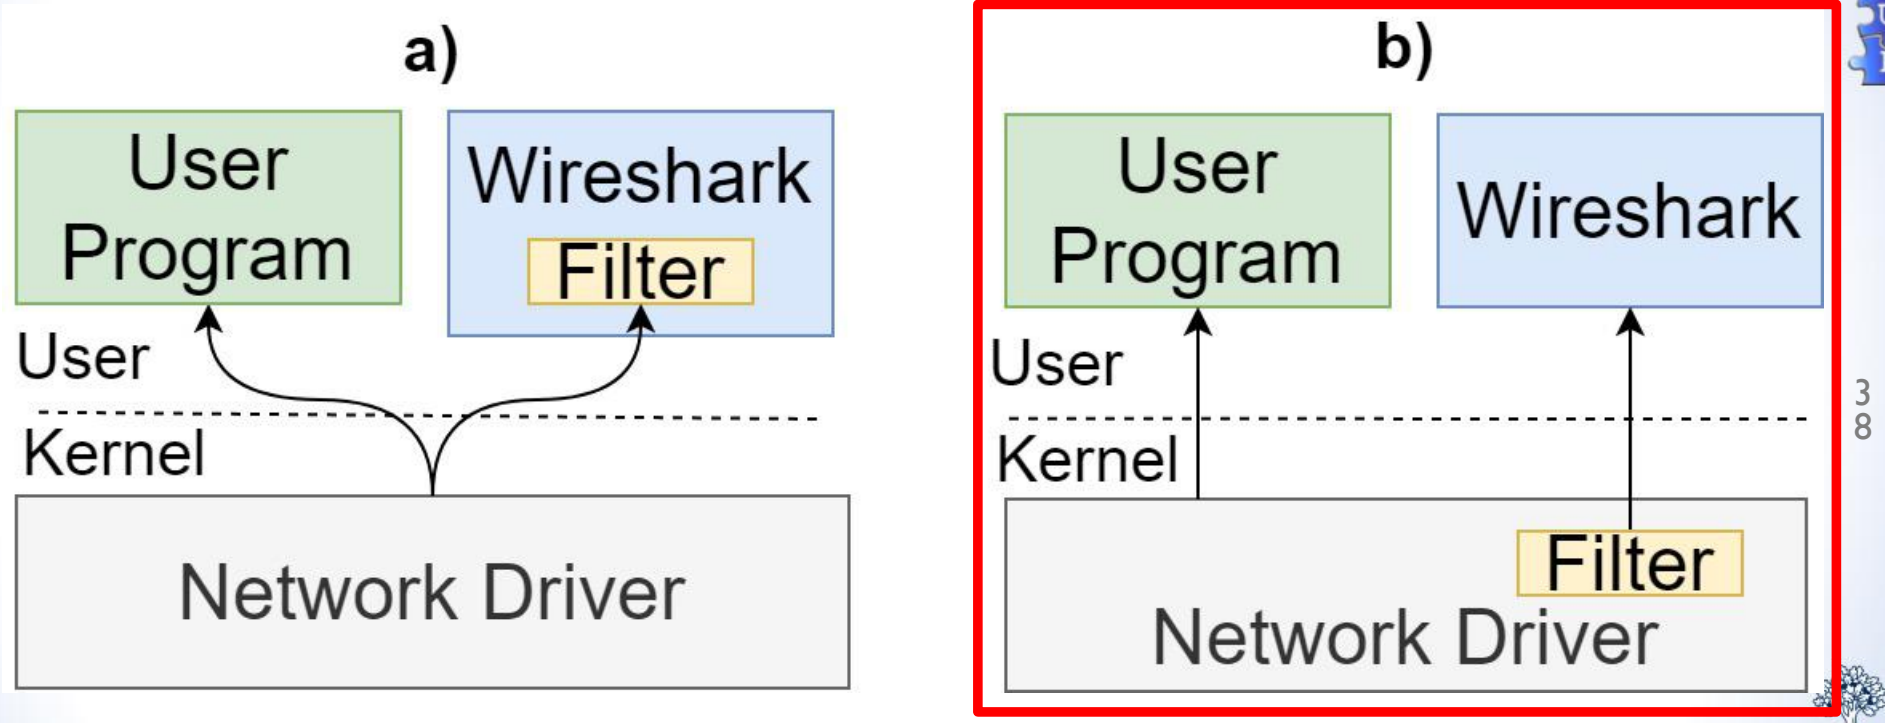

Improving Packet Capture Performance

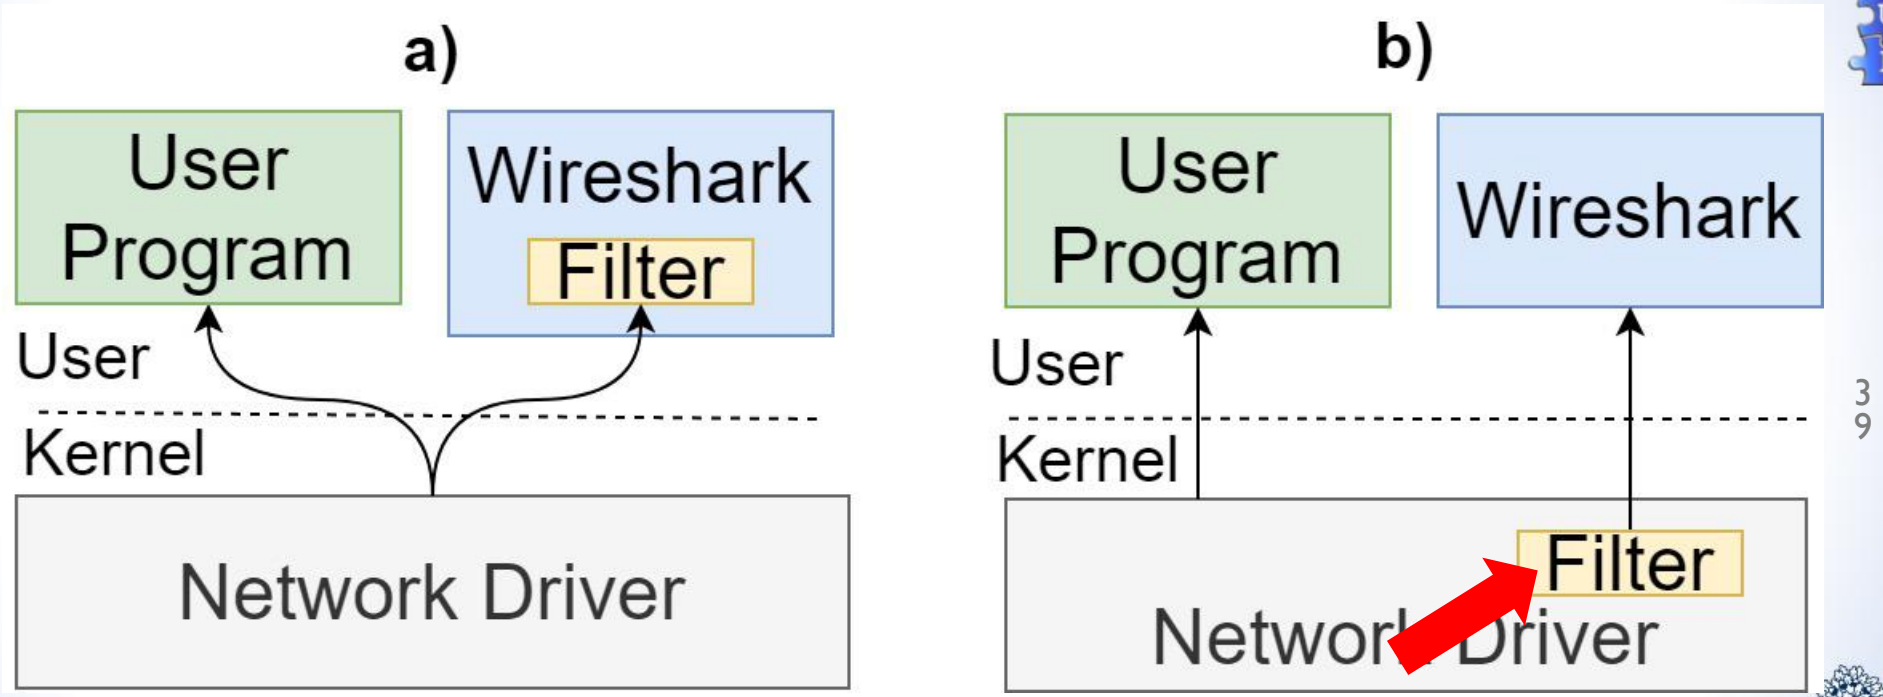

FFShark

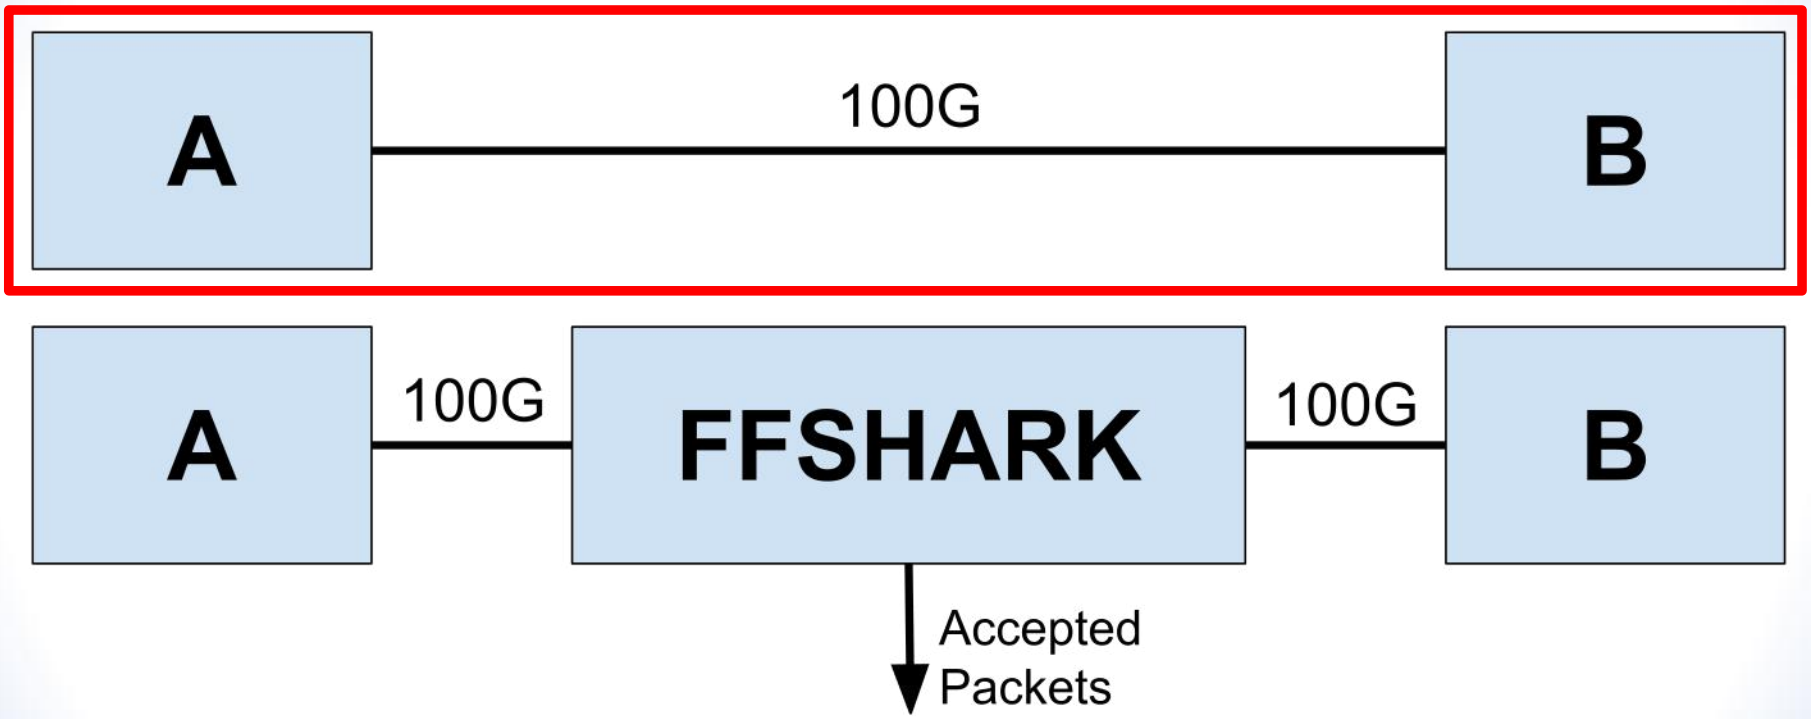

FFShark

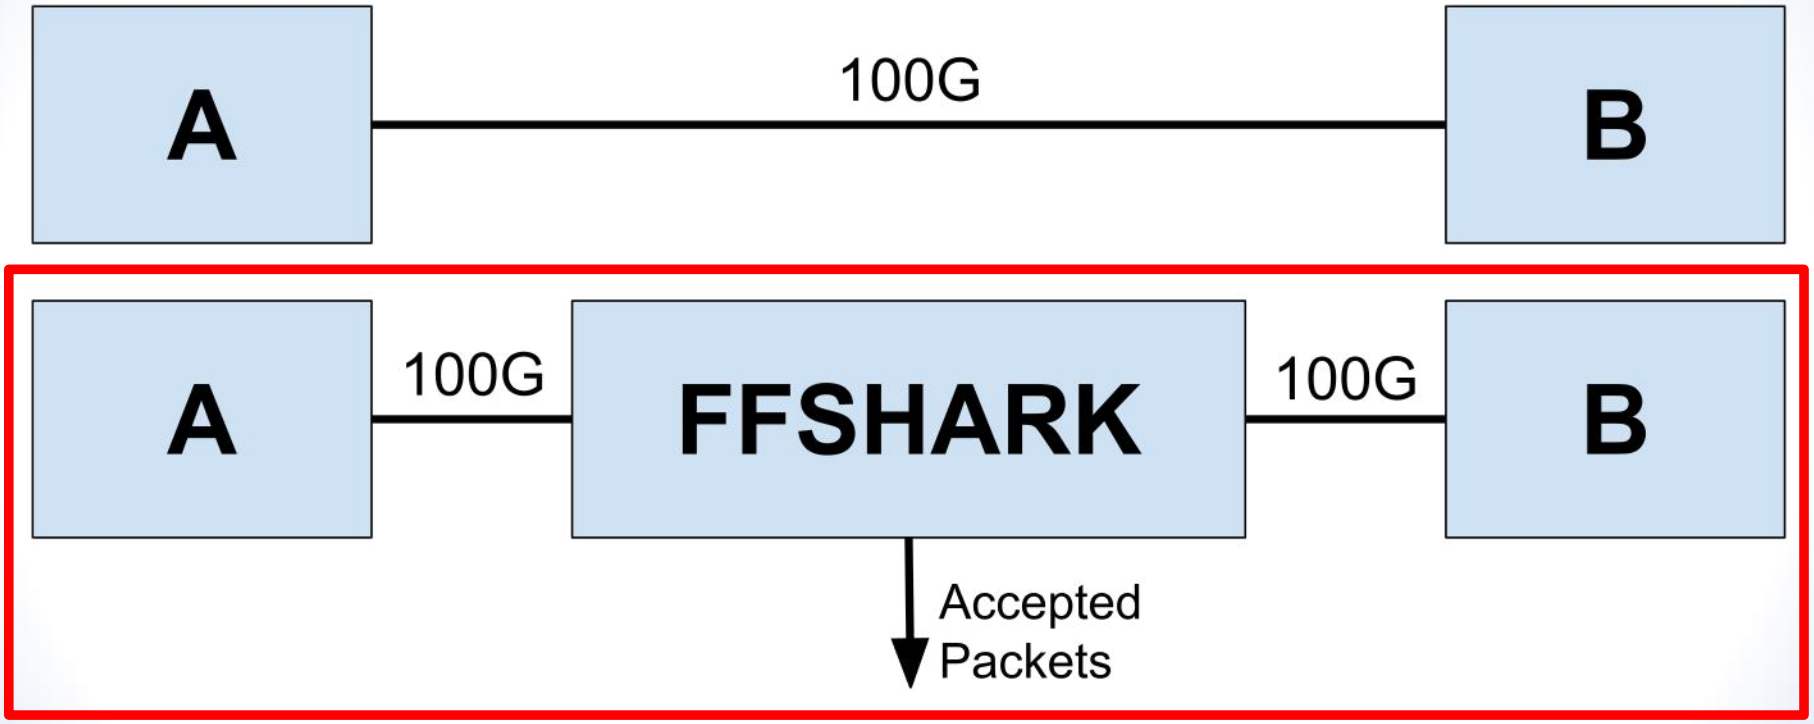

Architecture of the FFShark System

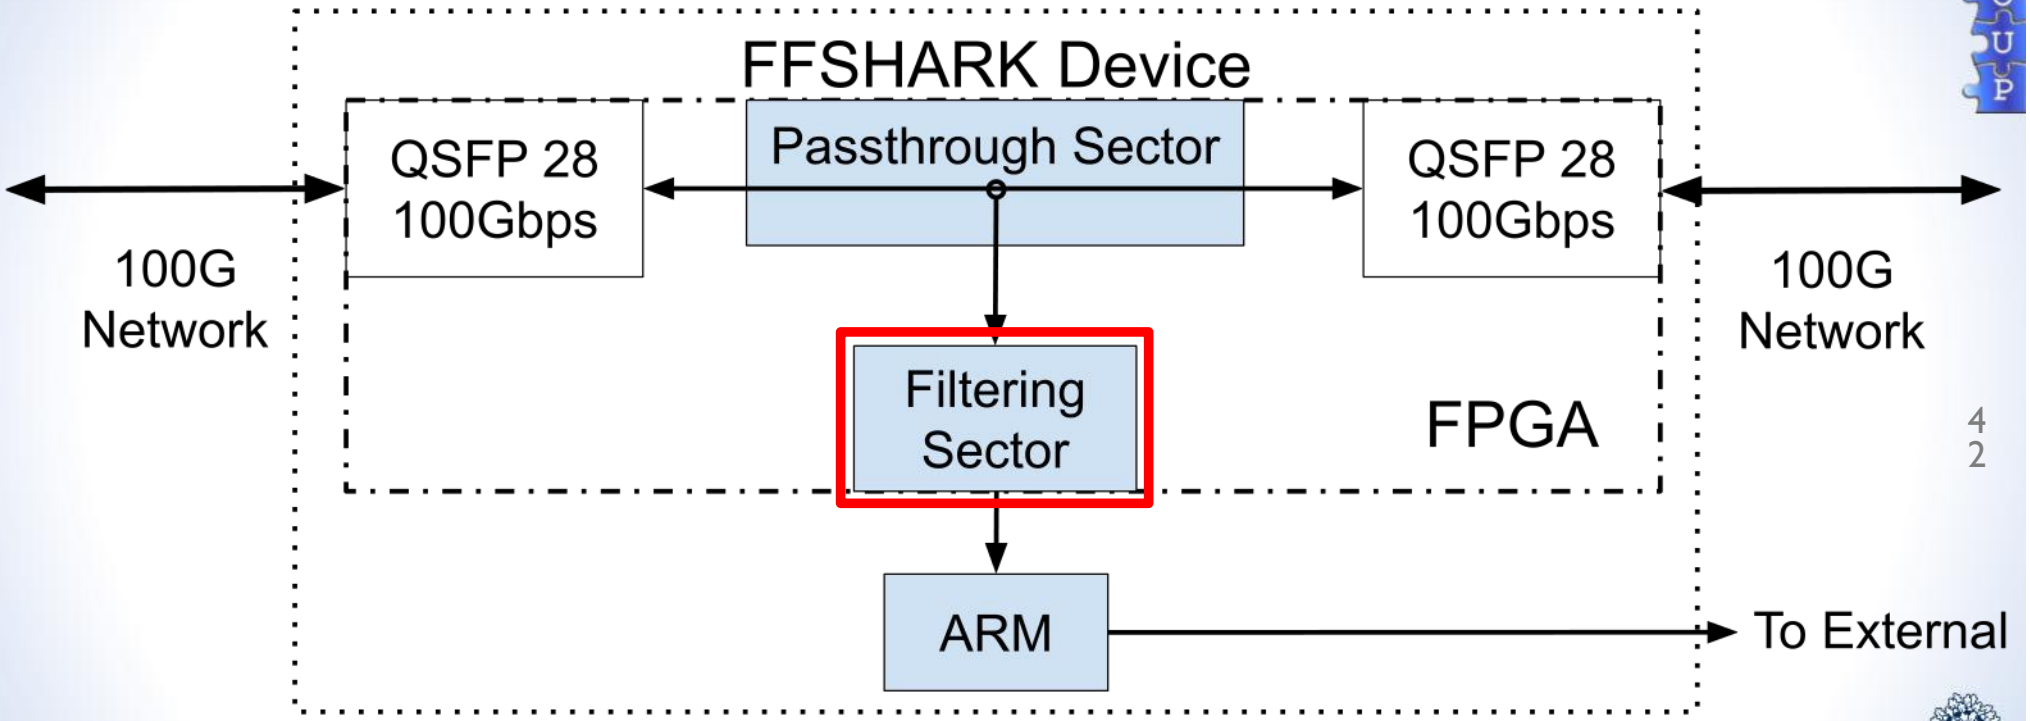

Arzhang Rafii and Paul Chow, FPL 2021, Thursday Session 4A

PHAROS: A MULTI-FPGA PERFORMANCE MONITOR

Motivation

- Monitoring events in a single FPGA
  - Which event happened first?
  - How long did each event take?

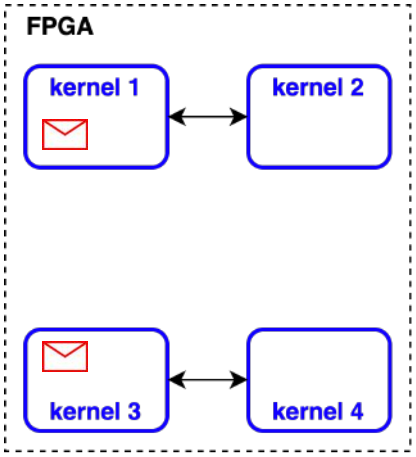

Motivation

- Monitoring events in a single FPGA
  - Which event happened first?
  - How long did each event take?
- Use a single timer to generate timestamps

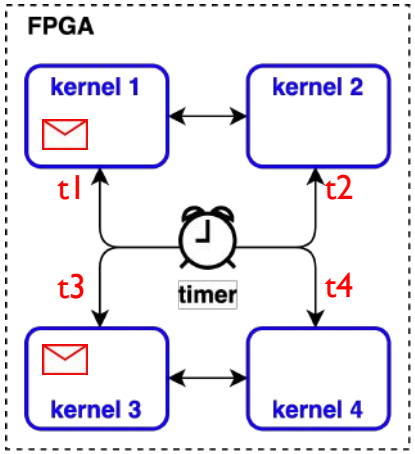

Motivation

- Monitoring events in a single FPGA
  - Use a single timer to generate timestamps
- Monitoring events in a multi-FPGA system
  - How to order events in the right sequence?

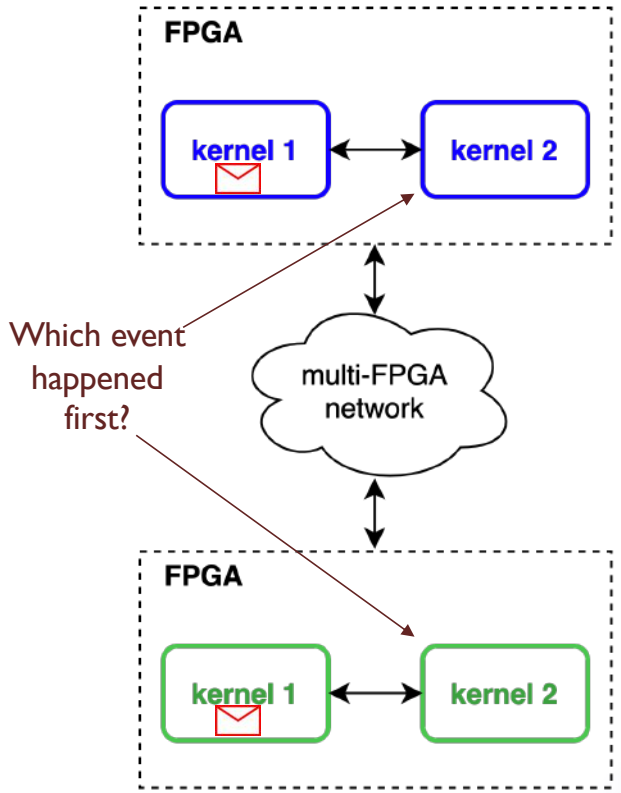

Motivation

- Monitoring events in a single FPGA
  - Use the same timer to generate timestamps
- Monitoring events in a multi-FPGA system
  - How to order events in the right sequence?
  - How to find point-to-point latency?

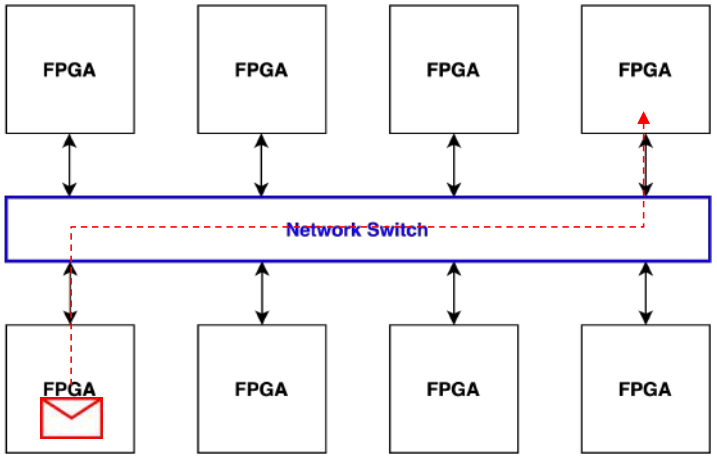

The Pharos Performance Monitor

- Uses the idea of "global time"
- Measures unidirectional point-to-point latency
- Collects traffic data
- Logs events in a multi-FPGA system
- Independent from lower-level communications

The Pharos Performance Monitor

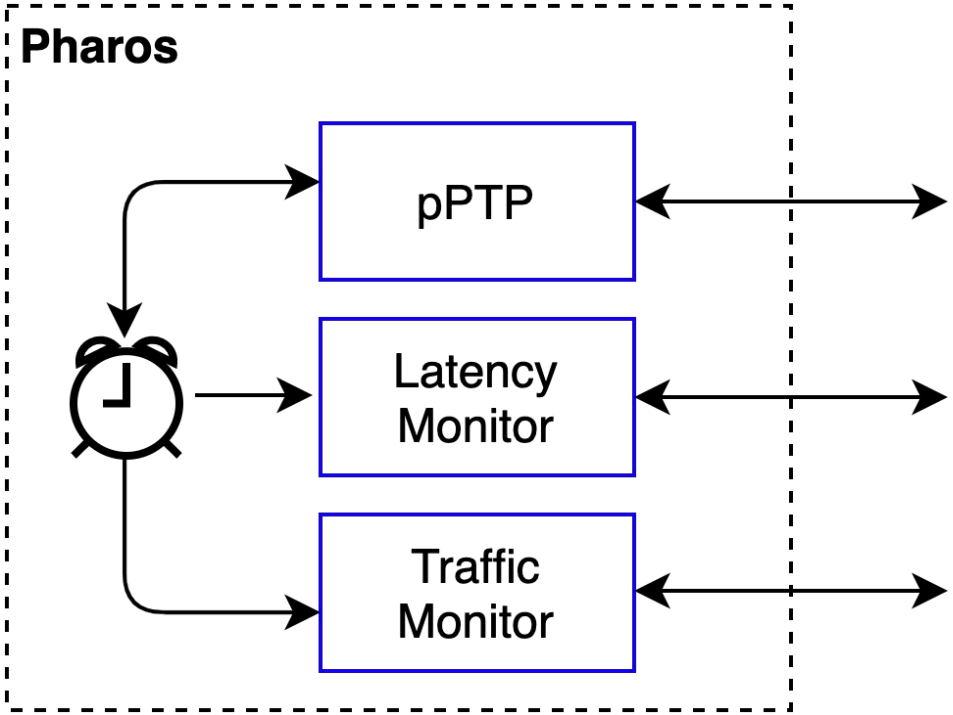

BUT WHAT WORKS BEST SO FAR?

Simulation!

If you don't simulate, it won't work

If you do simulate, it might work

Co-Simulation using Galapagos

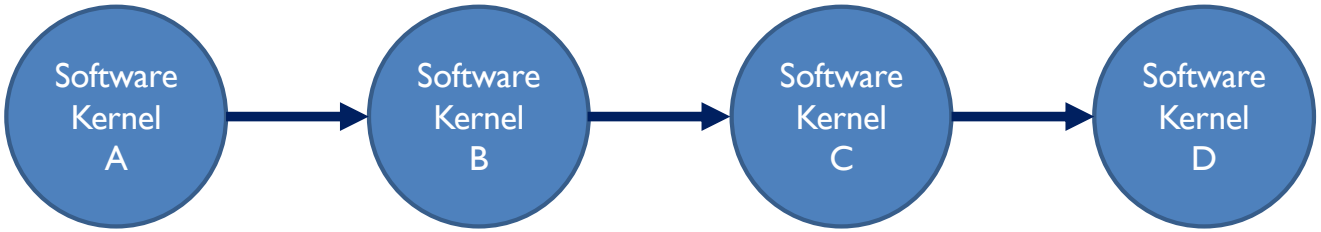

Start with everything running in software

Incremental Hardware Implementation

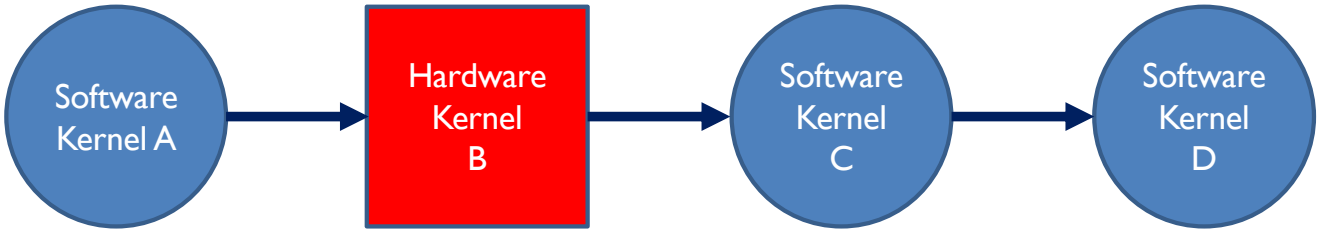

Test each kernel in hardware individually

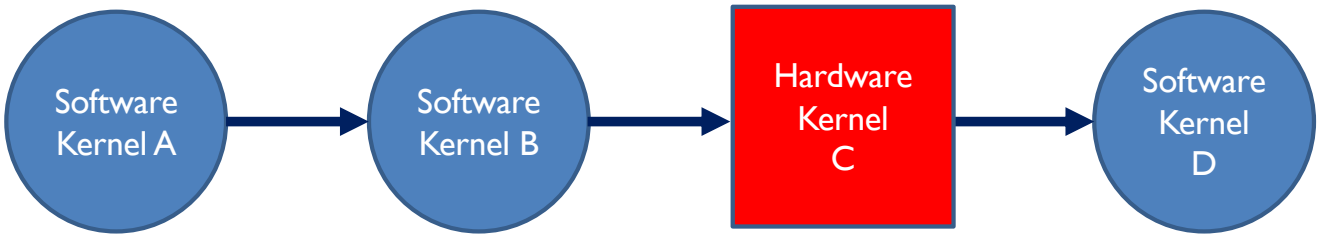

Test each kernel in hardware individually

Full System Integration

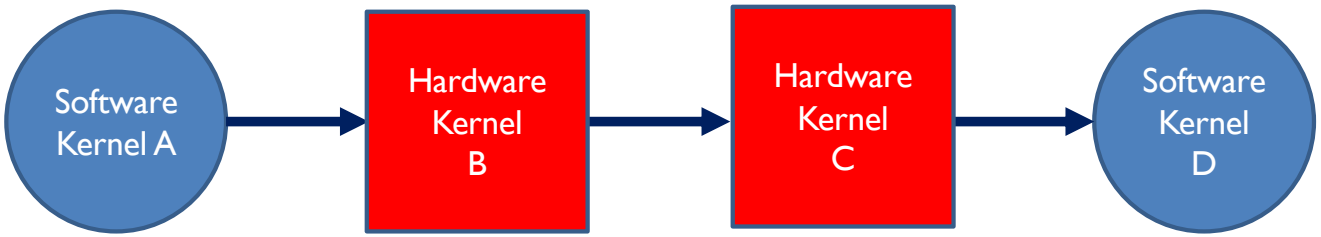

Then put it all together

Co-simulation with RTL Model

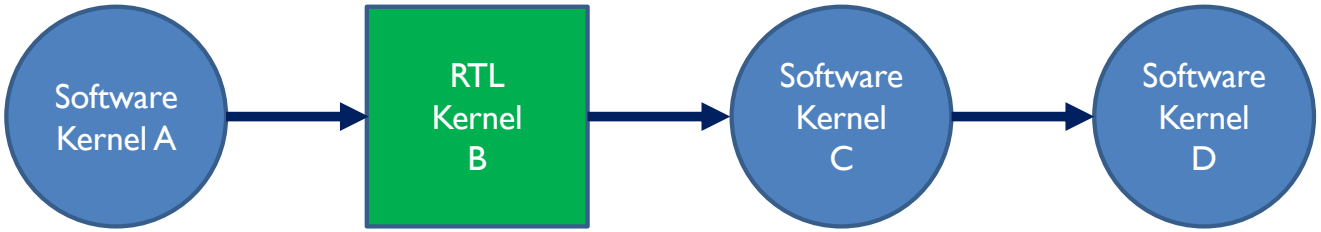

If you really need to…

FINAL THOUGHTS

- We need a collection of tools to help us debug
  - Have described some possibilities we've considered
  - Catapult Flight Data Recorder
- Should be usable by everyone, i.e., not one-off
  - Build on the work of others, don't re-invent
- How and what to standardize?
- We need to think more about debugging!!

THANKS FOR LISTENING