Optimal Filtering on ARM for ATLAS Tile Calorimeter Front-End Processing

MITCHELL A. COX
UNIVERSITY OF THE WITWATERSRAND, JOHANNESBURG

HEPP 2015

Overview

- ATLAS TILE CALORIMETER (TILECAL)
- THE OFFLINE DATA PROBLEM
- TILECAL READ-OUT ARCHITECTURE
  - Phase II Upgrades
- OPTIMAL FILTERING WITH ARM

ATLAS Detector

- 40 MHz Bunch Crossings
  - 1 GHz Interaction Rate
  - Millions of Sensor Channels
  - Petabytes per Second!

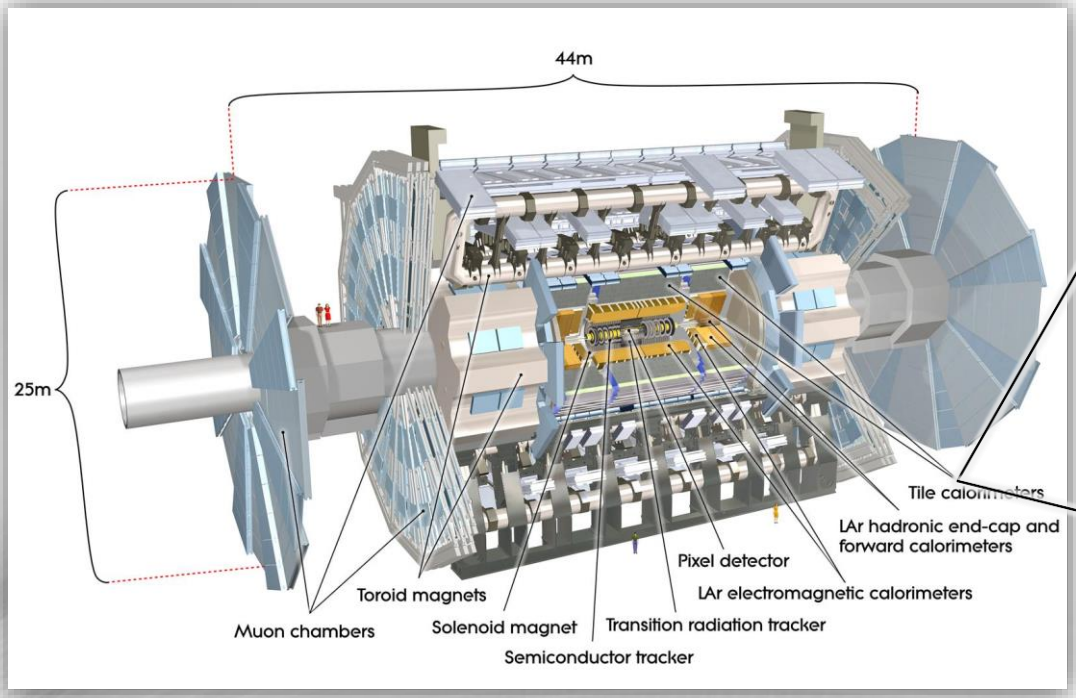

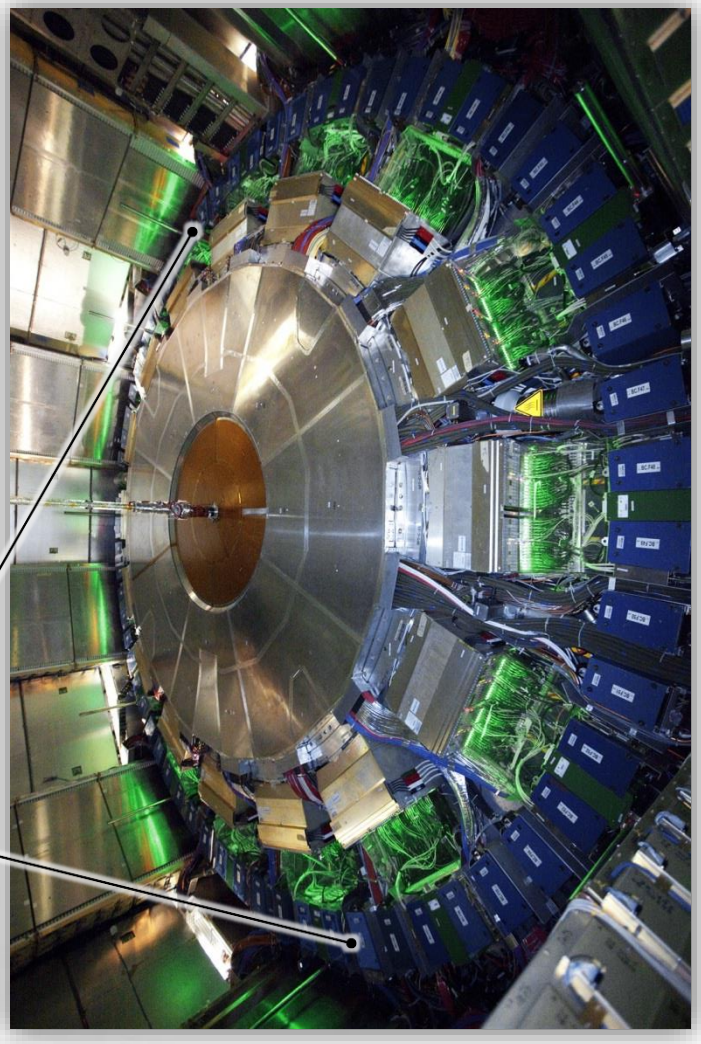

Tile Calorimeter

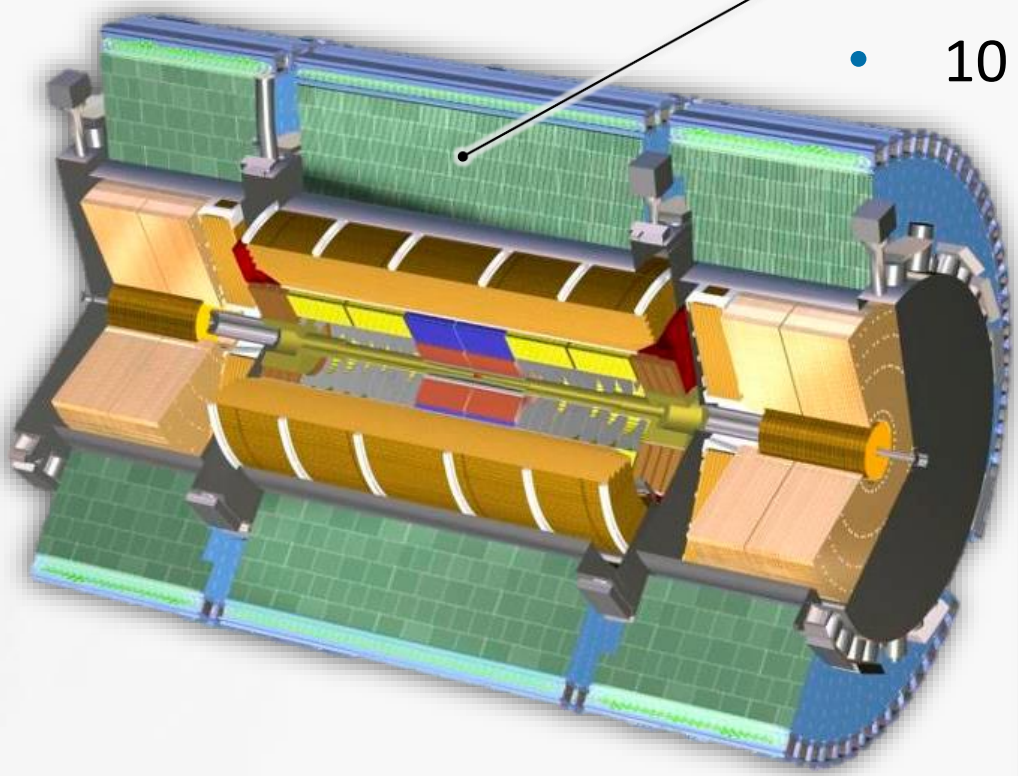

- 10 000 Photomultiplier Tubes

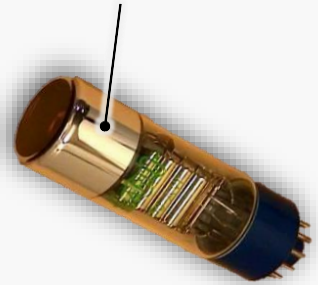

- Steel (absorber) and Scintillator
  - Wavelength shifting fibres

The Offline Problem

- PB/s storage is not feasible.

4 TB
200 MB/s
~5 hours

32 TB
200 MB/s
~2 days

1.3 PB
8.4 GB/s
~2 days

The Offline Problem

- PB/s storage is not feasible.

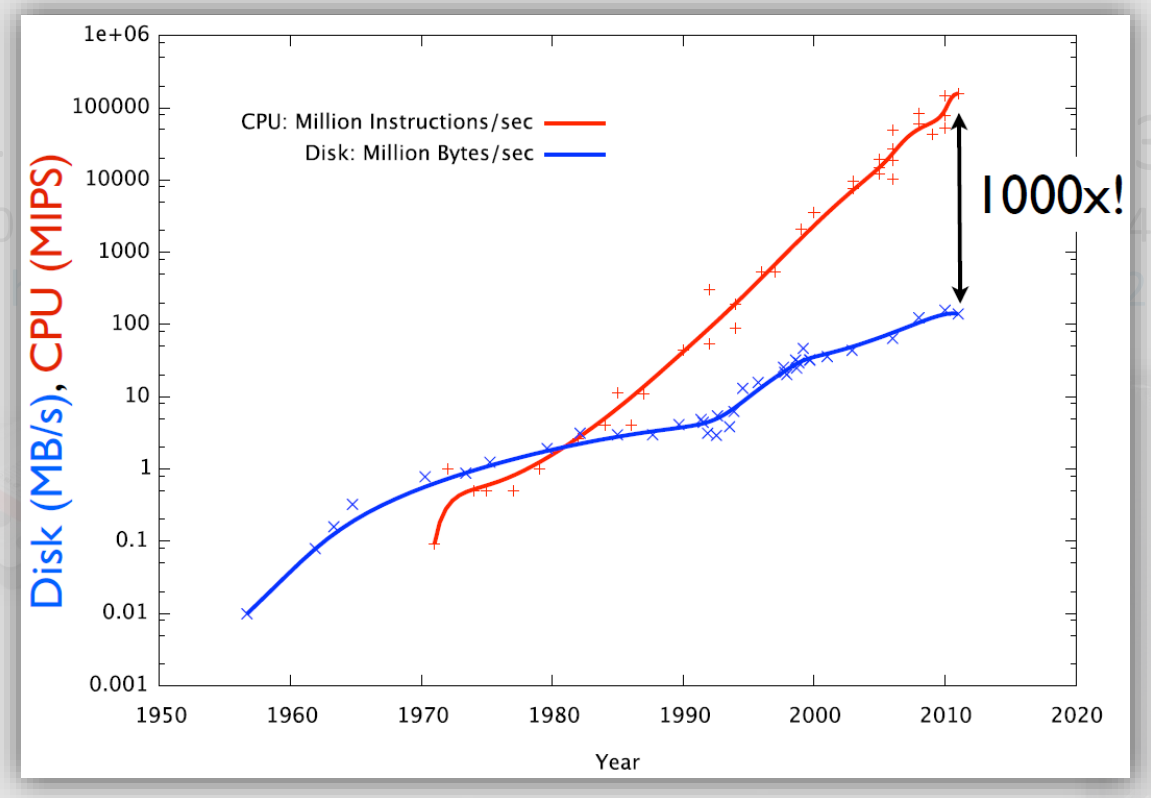

Reference: J Dursi. Parallel I/O doesn't have to be so hard: The ADIOS Library. 2012.

ATLAS Triggering and Data Acquisition System

- PB/s Raw reduced to MB/s Interesting Data

40 MHz
PB/s

100 kHz
GB/s

200 Hz
MB/s

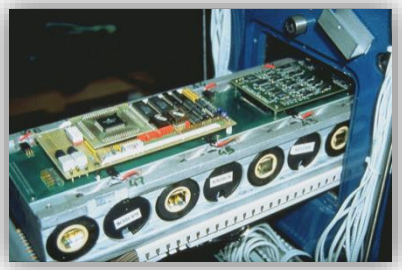

Algorithmic Intensity

ATLAS Triggering and Data Acquisition System

- PB/s Raw reduced to MB/s Interesting Data

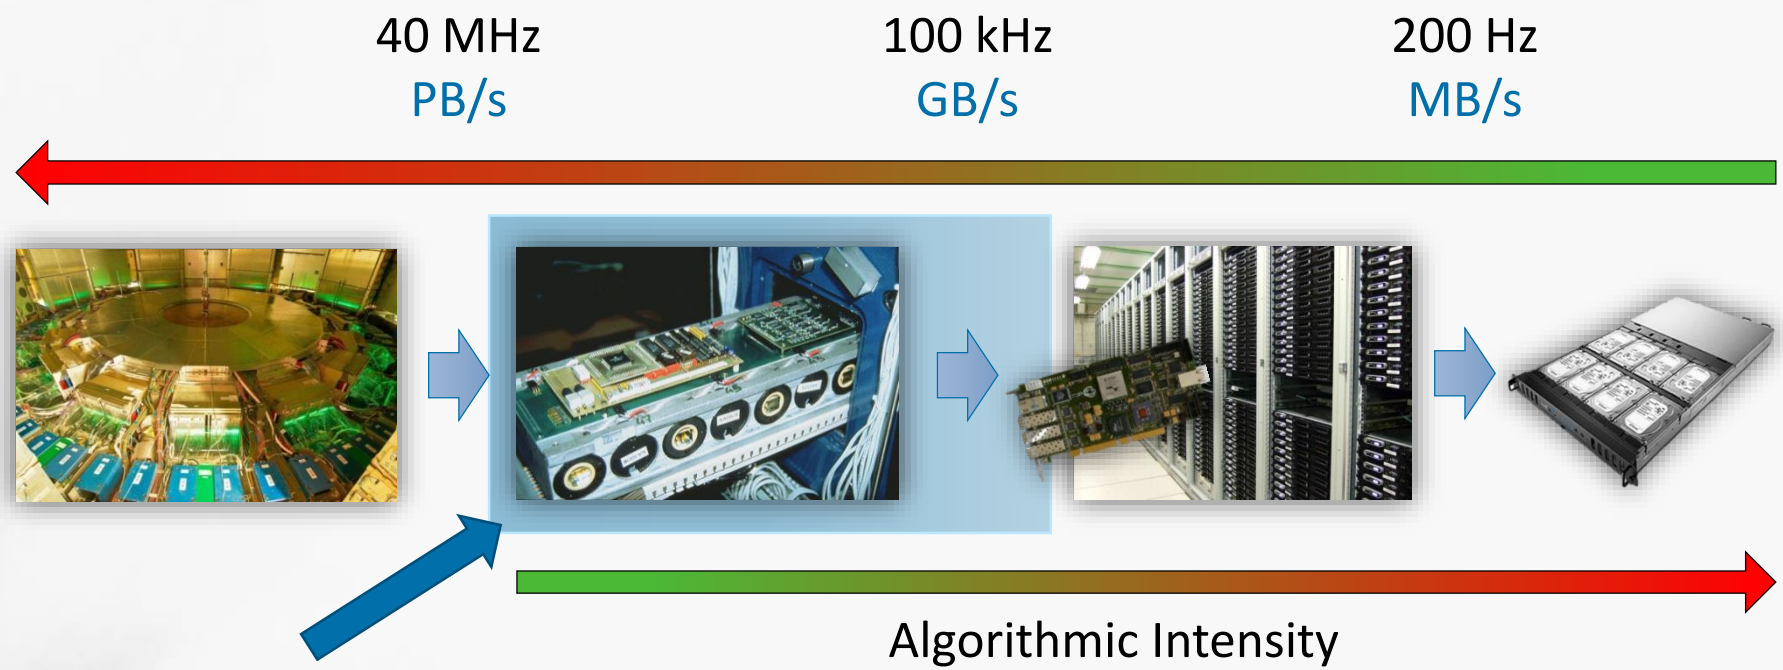

TileCal Read Out Architecture

Current:

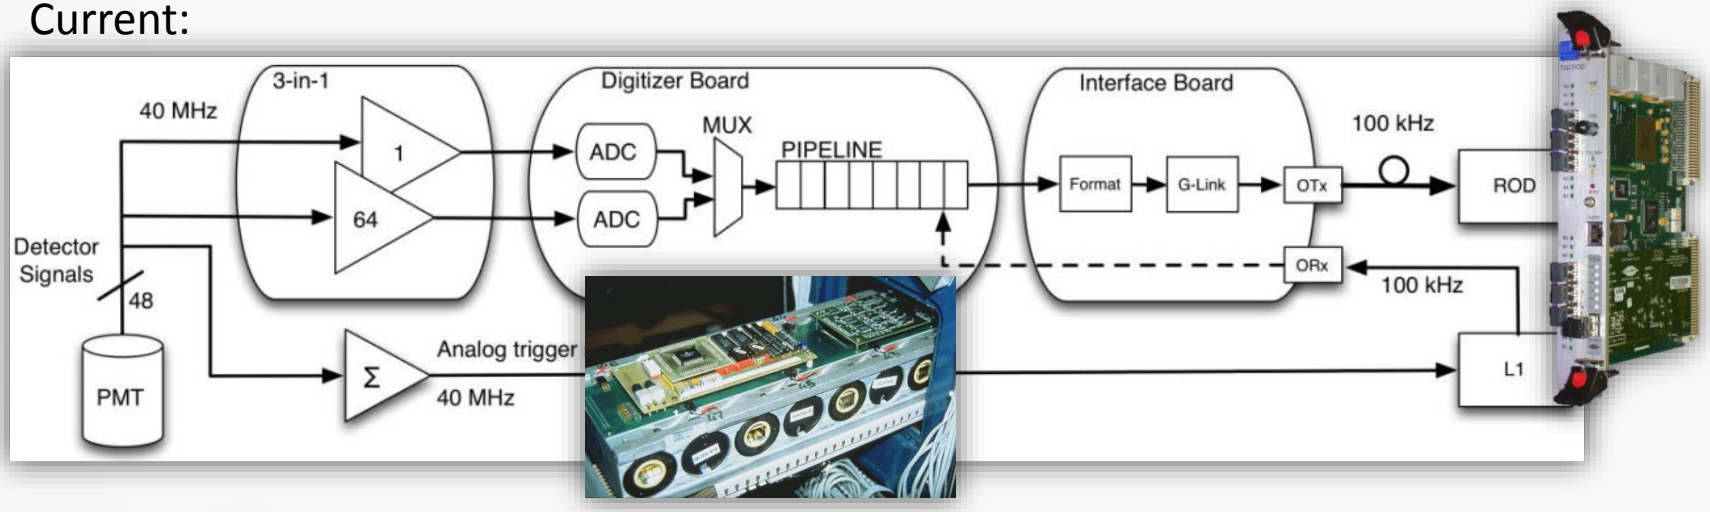

Future:

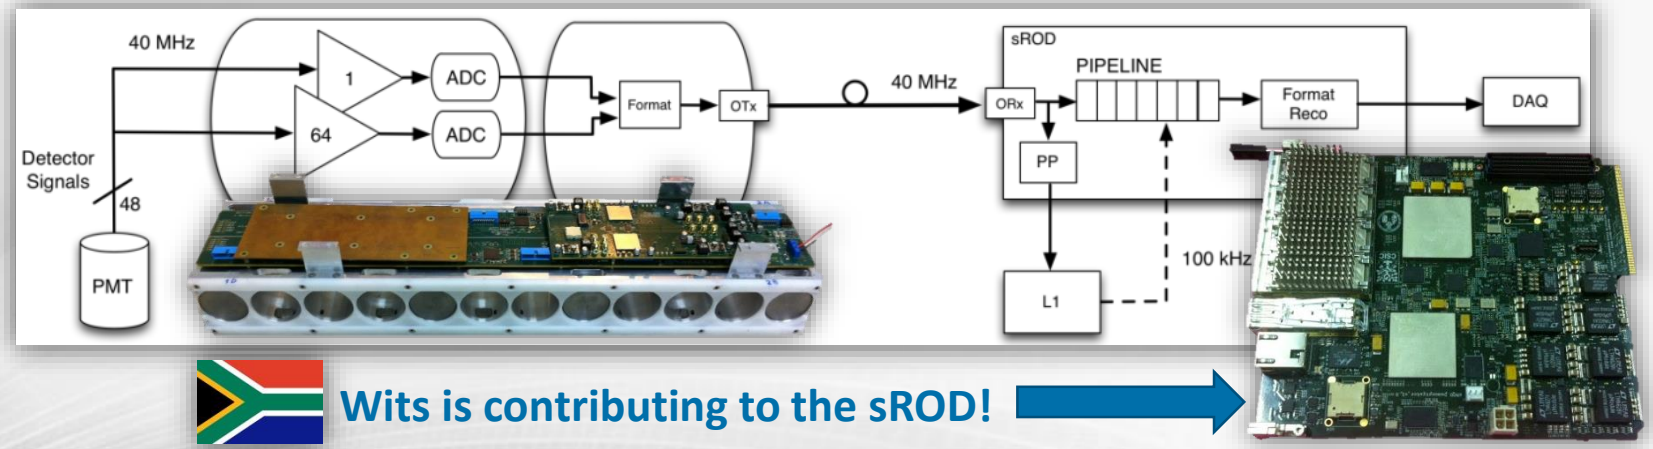

Wits is contributing to the sROD!

sROD PMT Energy Reconstruction

- PMT signal is conditioned, digitised and sent to sROD
- "Compresses" PMT samples to Amplitude, Phase and Pedestal with
  - Optimal Filtering
  - Matched Filter

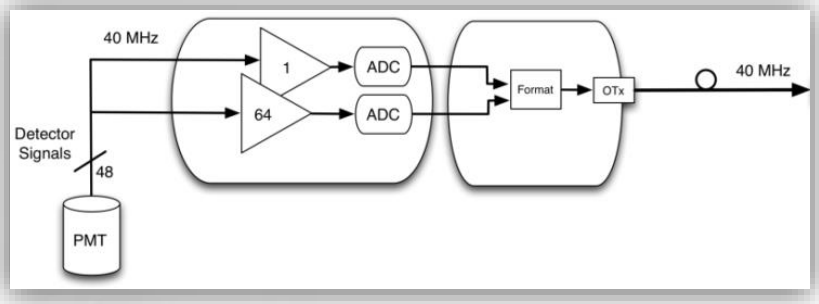

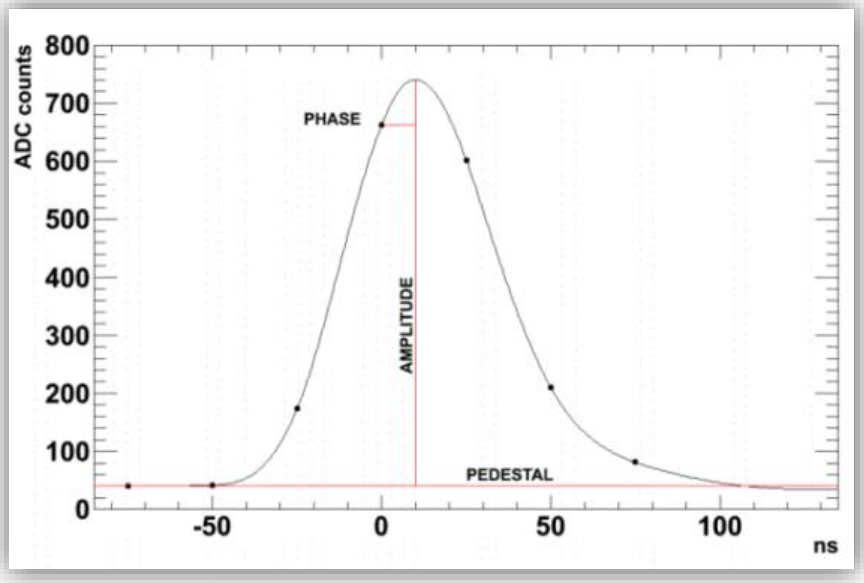

Figures Reference: B S Peralva. The TileCal Energy Reconstruction for Collision Data using the Matched Filter. 2013.

Processing Unit Integration

Current:

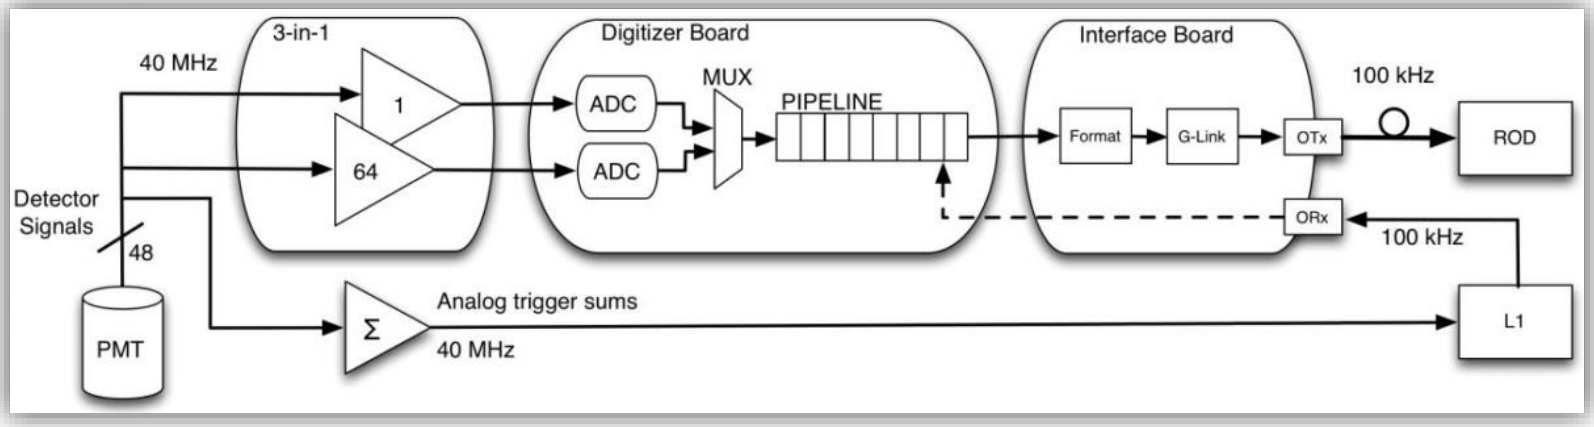

Future:

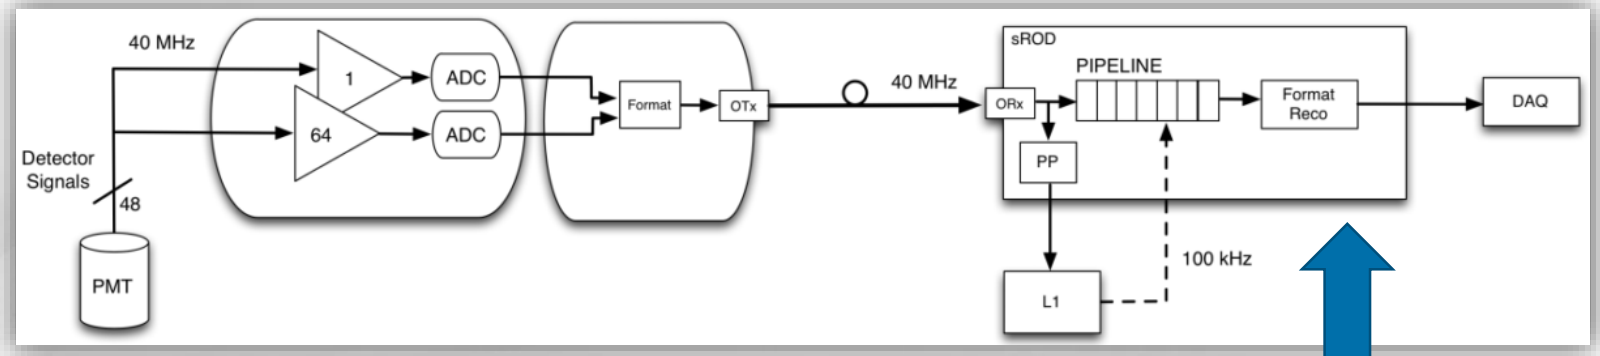

General purpose Processing Unit links to sROD

Processing Unit Integration

- Not in critical data path (for now)
- 80 Gb/s Ethernet Bandwidth
  - For general purpose usage
- 128 Gb/s Backplane Bandwidth
  - To sROD prototype

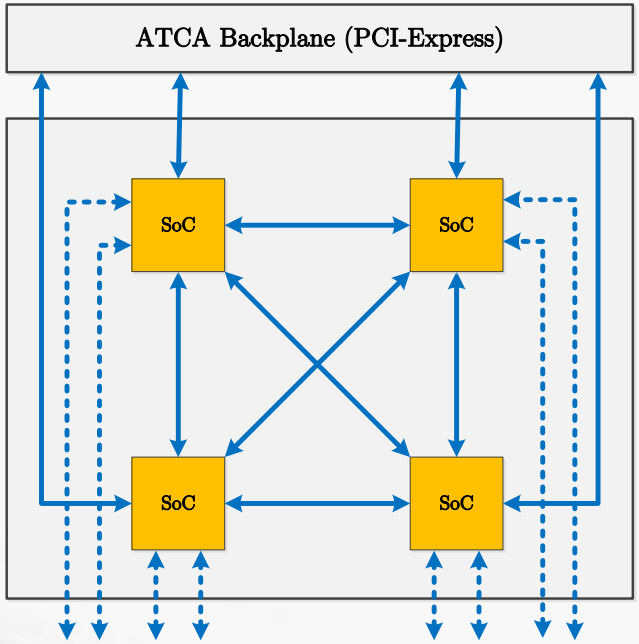

Processing Unit CPU

- ARM System on Chips
  - Low Power Consumption
  - Low Cost
  - High CPU Performance per Watt
- What about I/O performance?

Cortex-A9

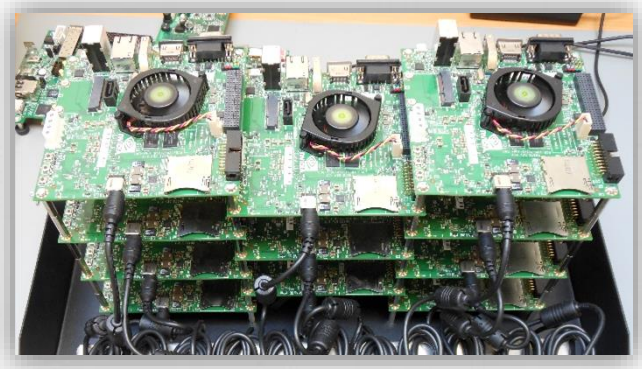

Cortex-A15 Cluster

Cortex-A57 (X-Gene)

System on Chip
External I/O Ports

Ethernet

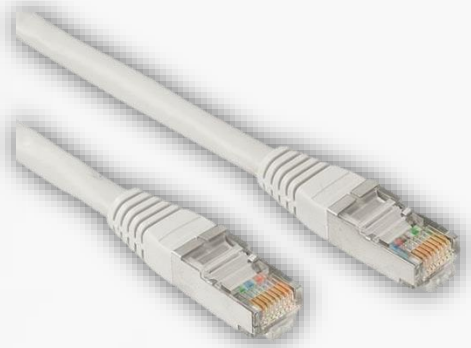

100 Mb/s - 1 Gb/s

12 - 125 MB/s

PCI-Express

N x 8 GT/s (Gen3)

≥ 985 MB/s

System on Chip
External I/O Ports

Ethernet

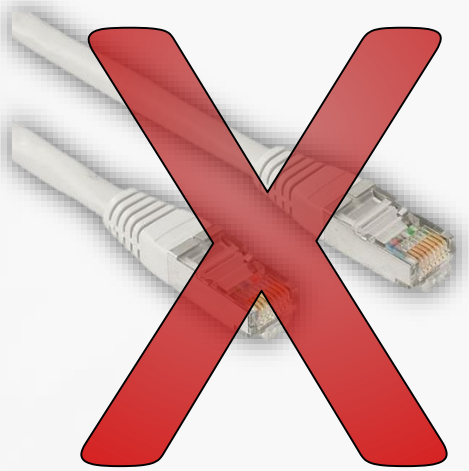

100 Mb/s - 1 Gb/s

12 - 125 MB/s

PCI-Express

N x 8 GT/s (Gen3)

≥ 985 MB/s

PCI-Express Benchmark Rig

- Test PCI-Express with a pair of SoCs:
  - Wandboard is a Quad-Core Cortex-A9 at 1 GHz
  - Freescale i.MX6 SoC

Manufactured in South Africa

PCI-Express Benchmark Rig

- Test PCI-Express with a pair of SoCs:
  - Wandboard is a Quad-Core Cortex-A9 at 1 GHz (i.MX6 SoC)

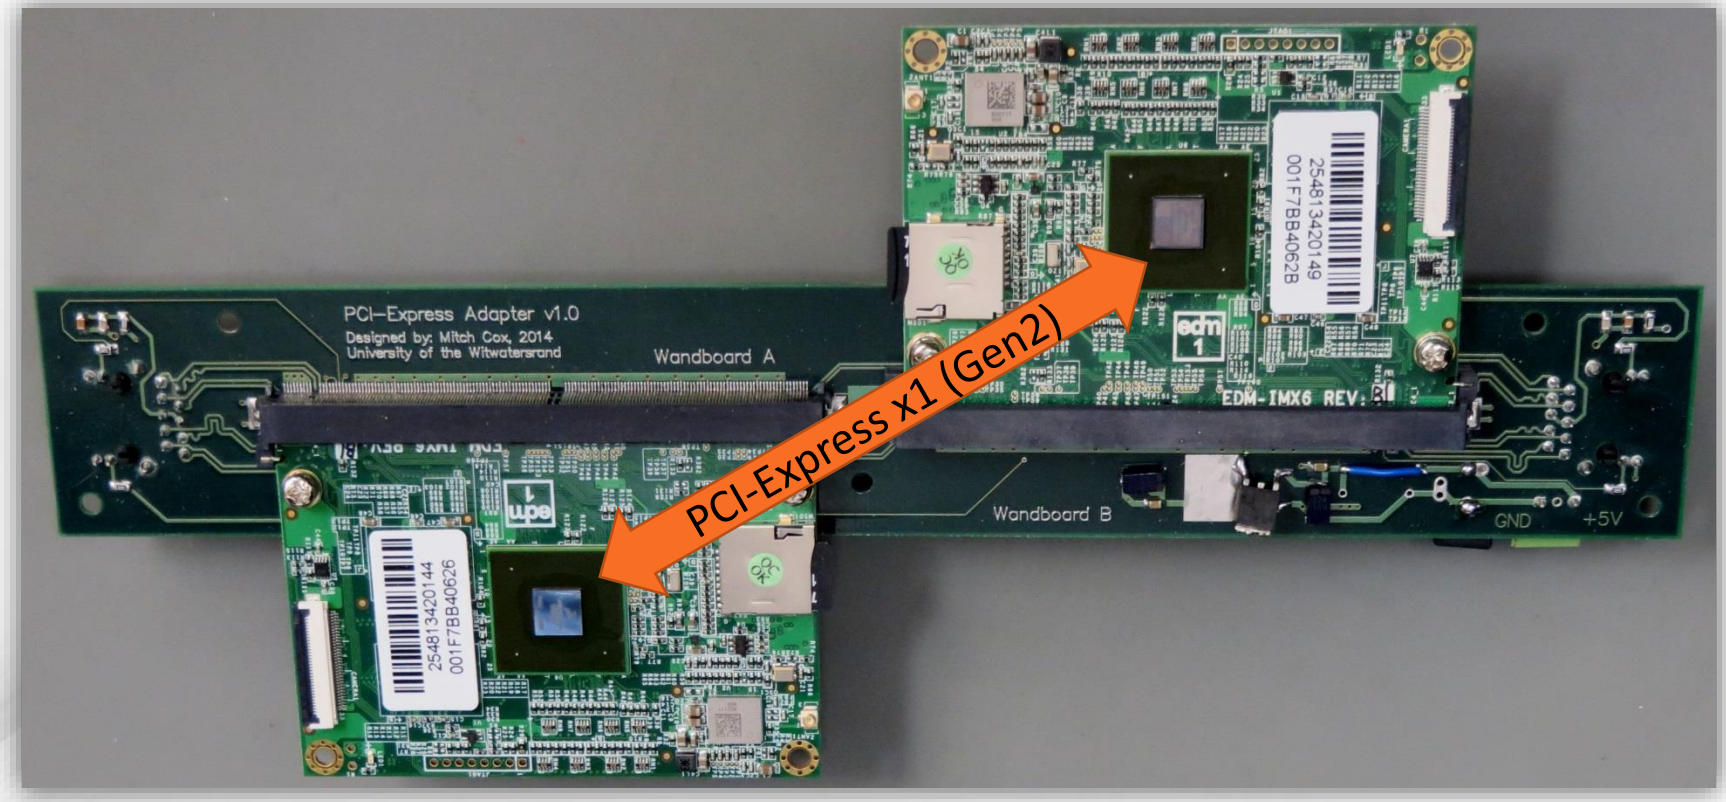

PCI-Express Test Results

- PCIe Gen2 x1 Link on i.MX6 SoC
  - 500 MB/s Theoretical
  - Uses CPU memcpy
    - No Direct Memory Access (DMA) on i.MX6

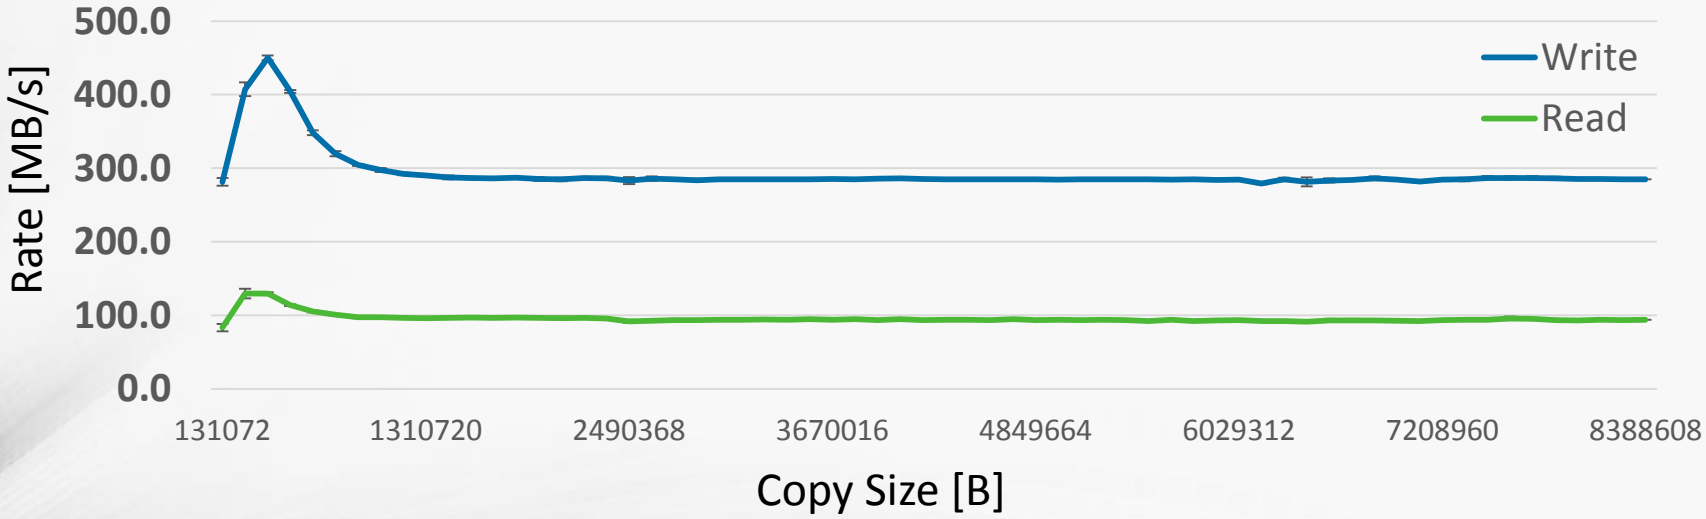

PCI-Express Test Results

- PCIe Gen2 x1 Link on i.MX6 SoC:
  - 500 MB/s Theoretical
  - Uses CPU memcpy
    - No Direct Memory Access (DMA) on i.MX6

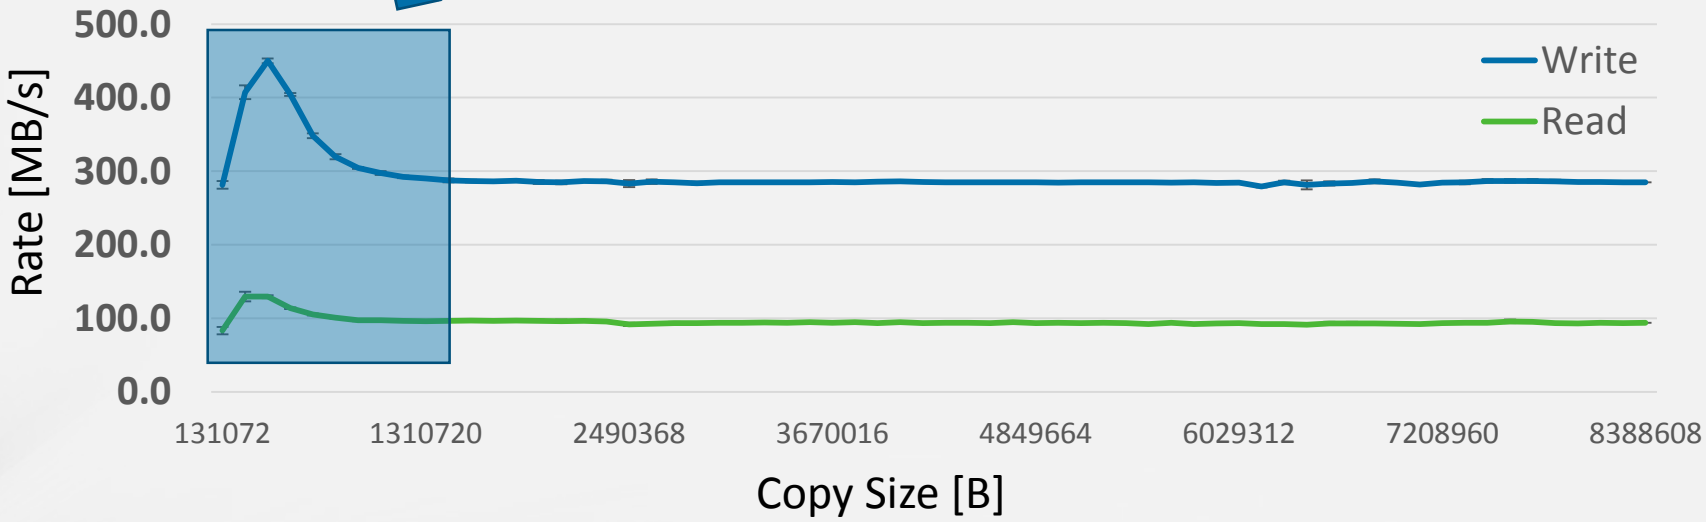

PCI-Express Test Results

- Can use i.MX6 IPU for “pseudo” DMA:

	CPU memcpy	DMA (Slave)	DMA (Master)
Read (MB/s)	94.8 ±1.1%	174.1 ±0.3%	236.4 ±0.2%
Write (MB/s)	283.3 ±0.3%	352.2 ±0.3%	357.9 ±0.4%

- Tested PCIe “ping-pong” latency:
  - 4 B packet
  - Average round trip time: 2 µs

PCI-Express Test Summary

- Throughput
  - 72 % of theoretical (500 MB/s) with DMA
  - 57 % of theoretical without DMA (200 MB/s)
    - 1 Gb Ethernet is max. 125 MB/s
- Latency
  - 2 µs (4 B), 42 µs (1500 B)
  - 1 Gb Ethernet is 50 – 125 µs
- Successful Proof of Concept
  - Superior to Ethernet

Optimal Filtering (Single Core)

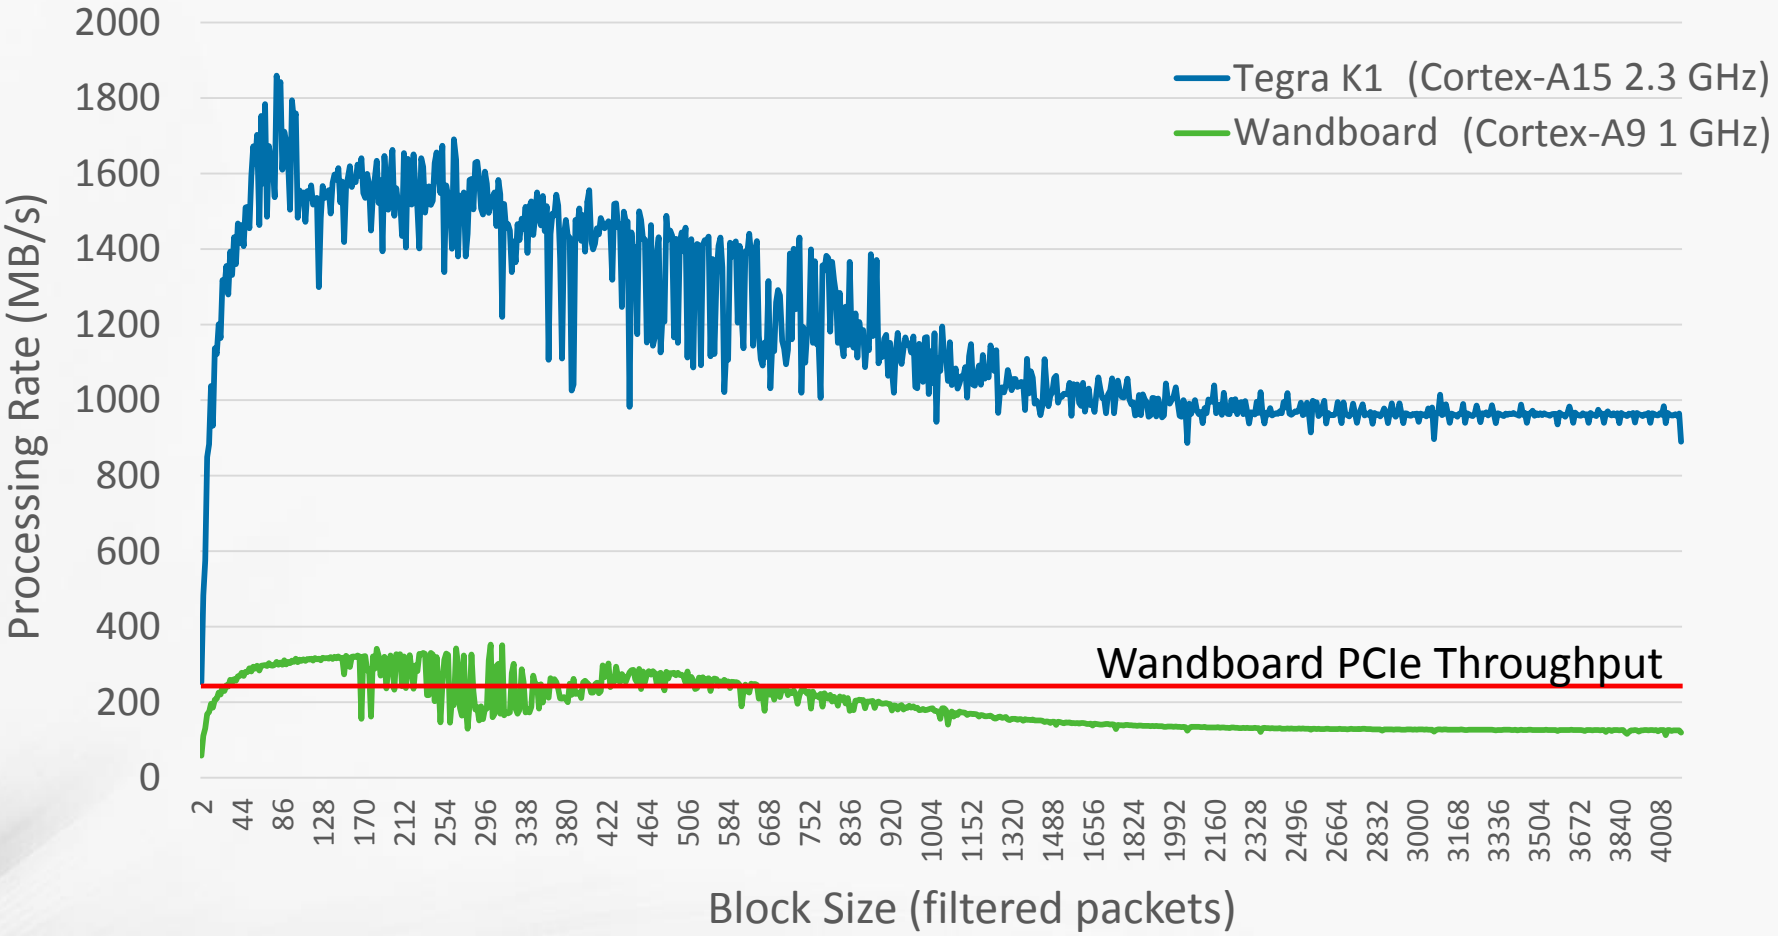

Optimal Filtering (Single Core)

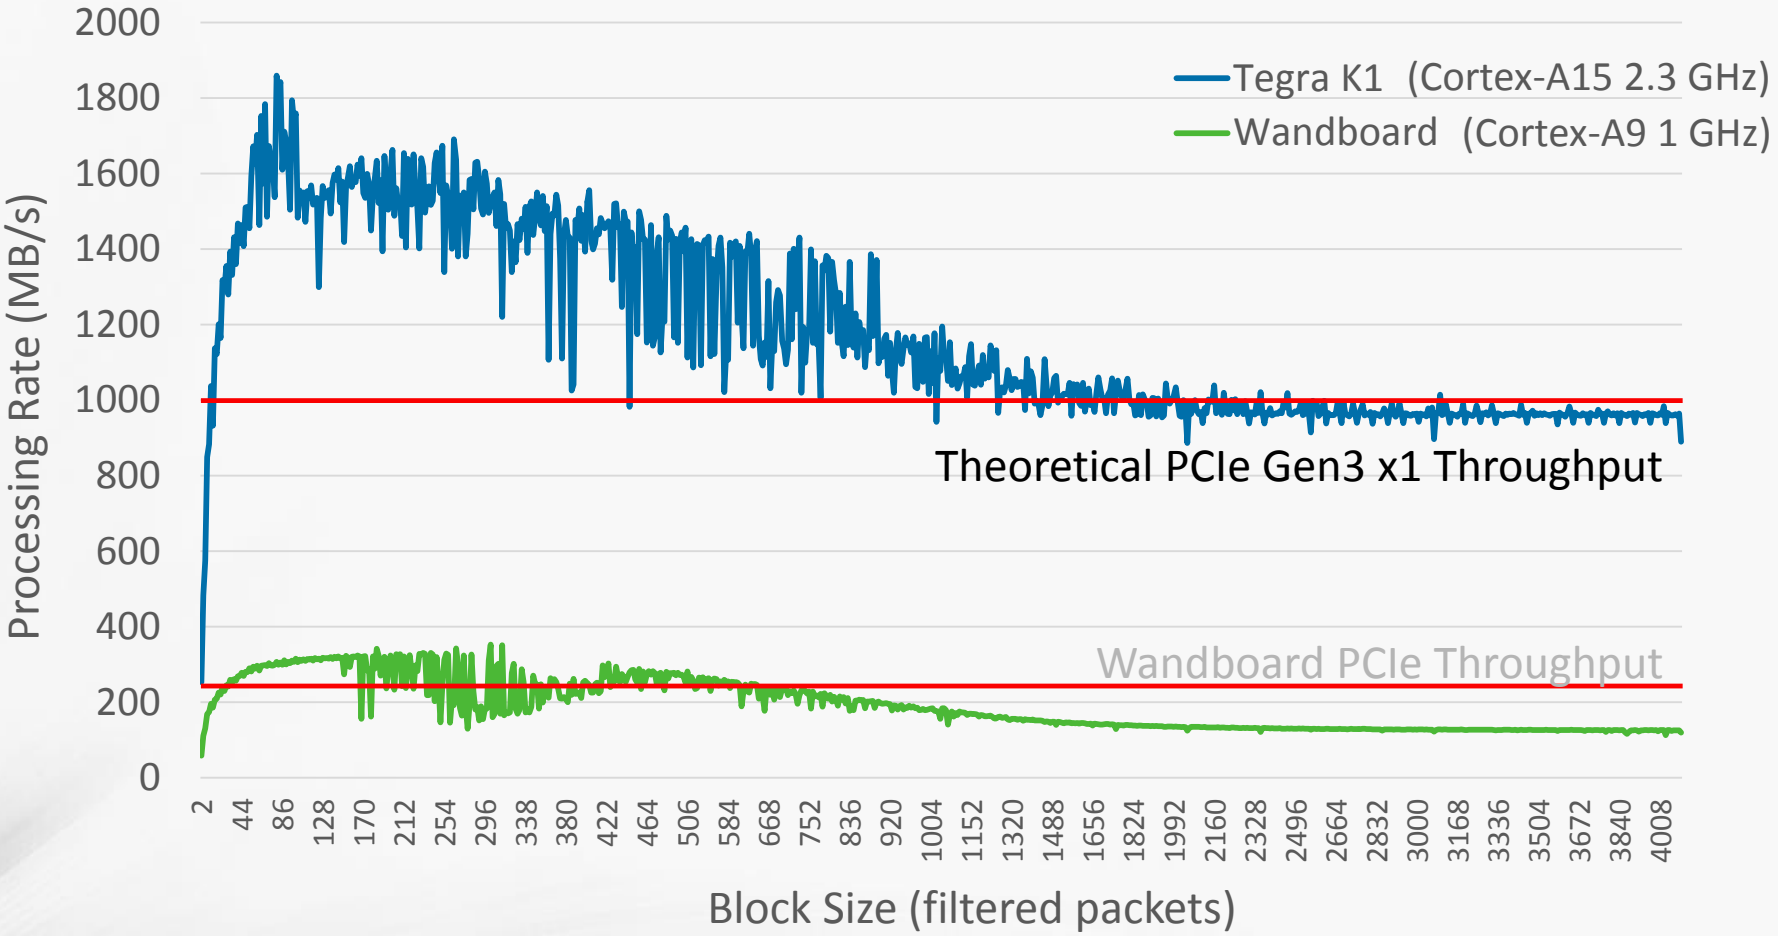

Summary

- High Data Throughput Computing
  - Required for Big Science
  - I/O is usually a bottleneck

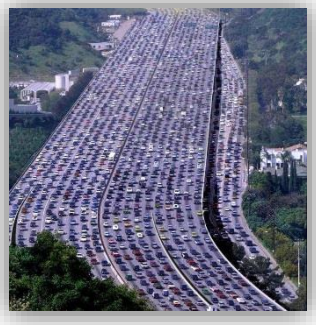

- ARM System on Chips
  - Affordable massive parallelism
- PCI-Express
  - High bandwidth External I/O
  - Alternative to Ethernet

Questions or Comments?

MITCHELL.COX@STUDENTS.WITS.AC.ZA

Acknowledgements

- The “Massive Affordable Computing Project” team:
  - Robert Reed, Matthew Spoor, Oscar Kureba (Physics)
  - Daniel O Kwofie, Ekow Etutu (Elec. Eng.)
  - Carlos Solans Sanchez, Alberto Valero Biot (Valencia – CERN)
- MSc Supervisors: Prof. Bruce Mellado, Prof. Ivan Hofsajer

- The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the authors and are not necessarily to be attributed to the NRF.
- I would also like to acknowledge the School of Physics, the Faculty of Science and the Research Office at the University of the Witwatersrand, Johannesburg.

Backup Slides

ARM Performance

	Cortex-A7	Cortex-A9	Cortex-A15
CPU Clock (MHz)	1008	996	1000
HPL (SP GFLOPS)	1.76	5.12	10.56
HPL (DP GFLOPS)	0.70	2.40	6.04
CoreMark	4858	11327	14994
Peak Power (W)	2.85	5.03	7.48
DP GFLOPS/Watt	0.25	0.48	0.81