The Massively Affordable Computing Project: ARM System on Chips for High Data Throughput Scientific Computing

MITCHELL A. COX
UNIVERSITY OF THE WITWATERSRAND, JOHANNESBURG, SOUTH AFRICA

ACAT 2014

Overview

- DATA IN MODERN TIMES
- CONVENTIONAL COMPUTING PARADIGMS
- DATA STREAM COMPUTING PARADIGM
- SYSTEM ON CHIP BASED PROCESSING UNIT
- PCI-EXPRESS I/O BENCHMARKS

Data is getting BIGGER!

Big Data

LHC

NICA

SKA

Data is getting BIGGER!

TeraByte/s

Massive Data Processing

- Data volume must be reduced before storage

Massive Data Processing

- Data volume must be reduced before storage
  - Generic processing complements existing FPGA's

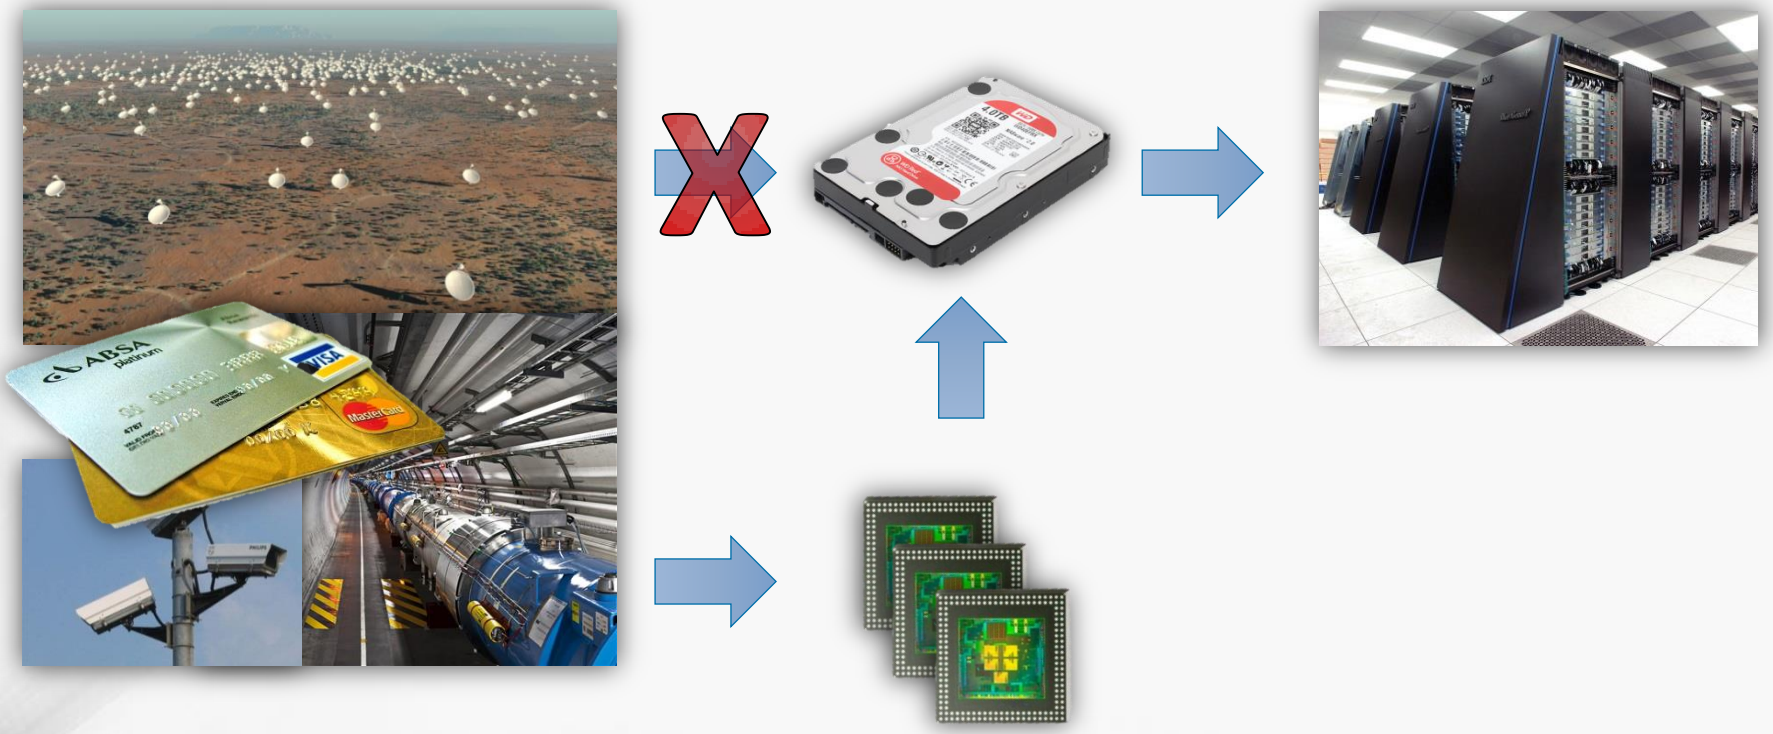

Massive Data Processing

- Data volume must be reduced before storage
  - Generic processing complements existing FPGA's

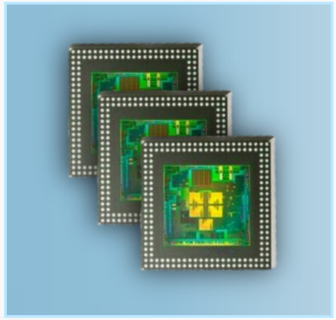

My focus is here…

Conventional Computing Paradigms

- **H**igh **P**erformance **C**omputing
  - Tightly Coupled
  - FLOPS
- **H**igh **T**hroughput **C**omputing
  - Loosely Coupled
  - Jobs/Day (FLOPS)
- **M**any **T**ask **C**omputing
  - Tightly or Loosely Coupled
  - FLOPS or I/O Throughput

Data Stream Computing

- Three important constraints:

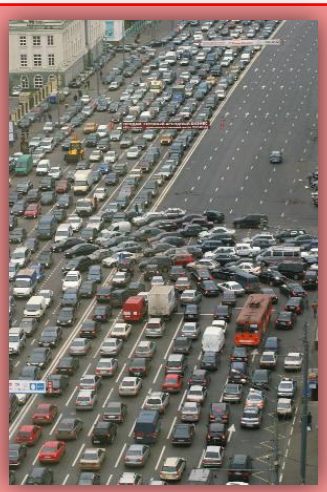

High Data Throughput

High Data Throughput

- CPU and External I/O must be balanced.

Unbalanced

(Conventional Systems)

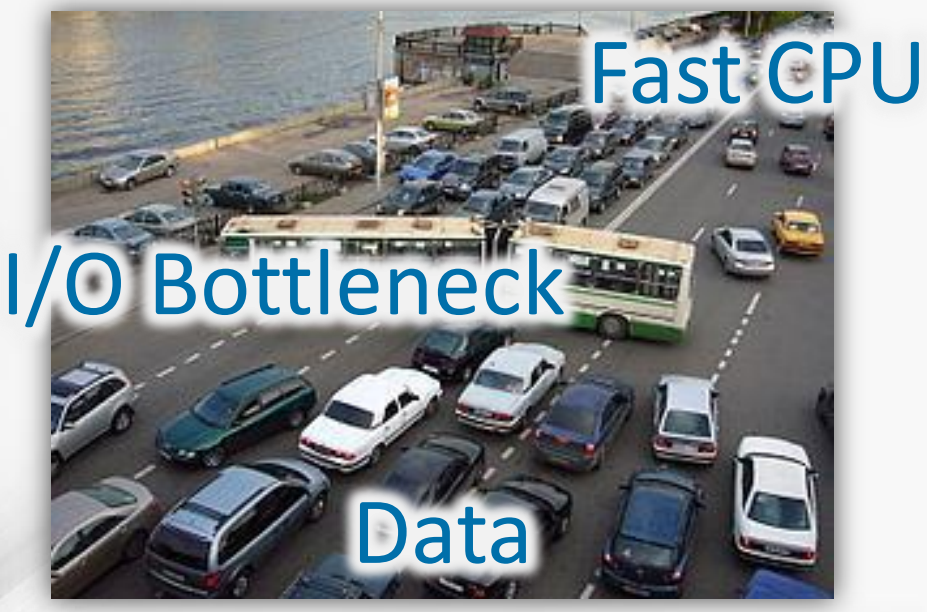

Balanced

(Data Stream Computing)

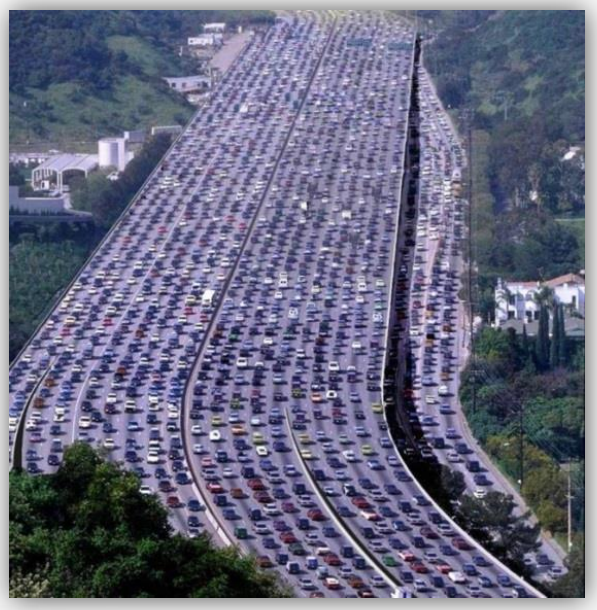

Data Stream Computing

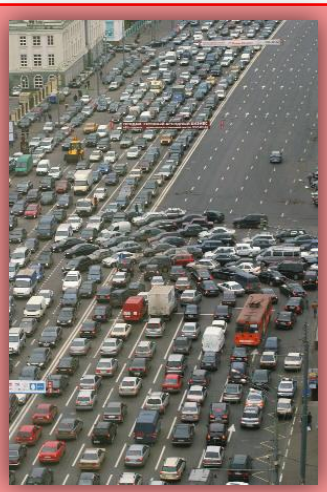

High Data Throughput

Data Stream Computing

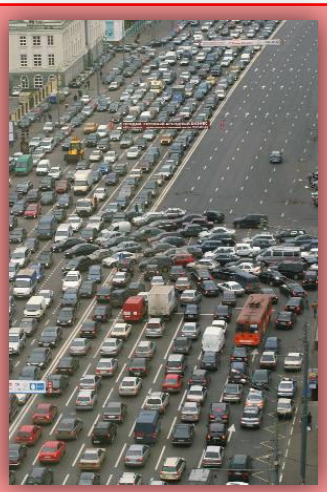

High Data Throughput

No Offline Storage Allowed

The Offline Problem

- TB/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.

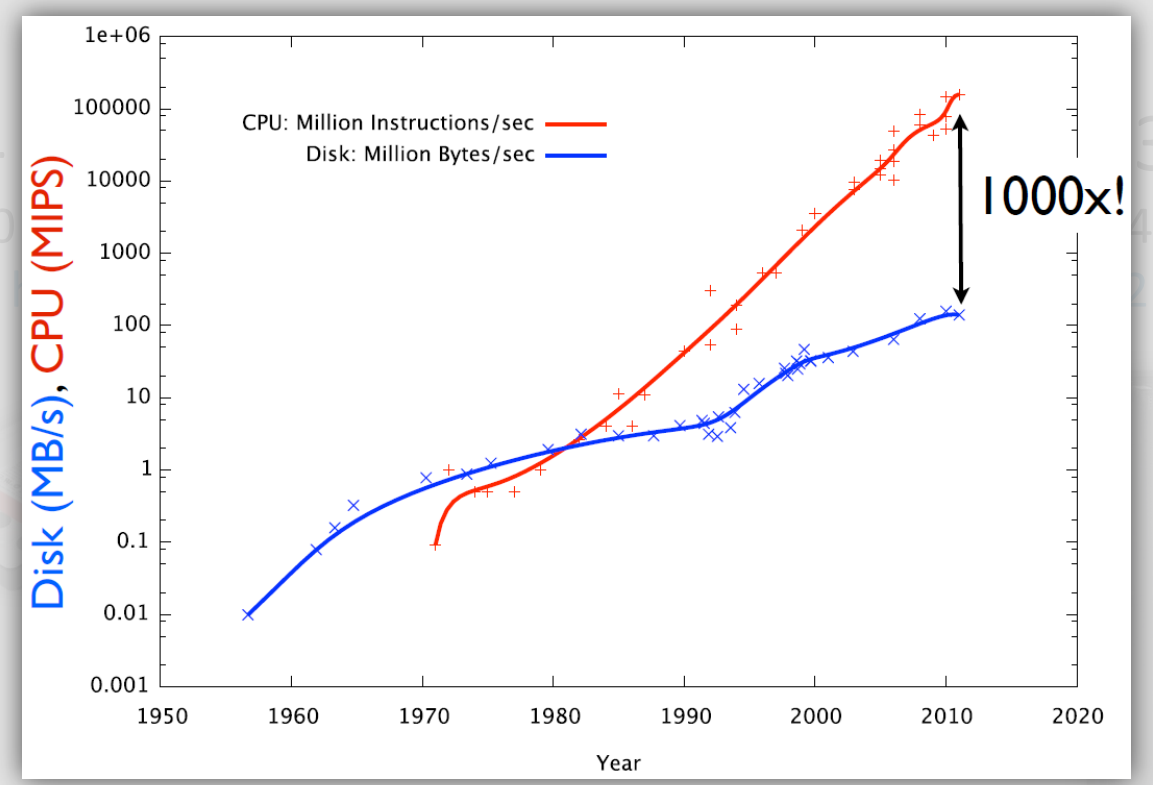

Reference: J Dursi. Parallel I/O doesn't have to be so hard: The ADIOS Library. 2012.

Data Stream Computing

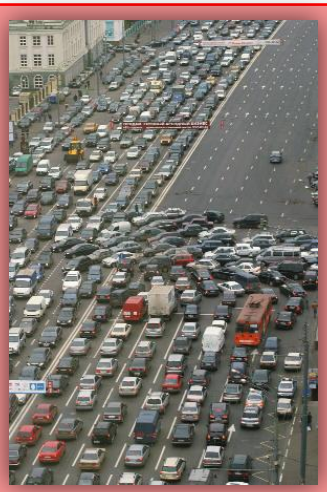

High Data Throughput

No Offline Storage Allowed

Data Stream Computing

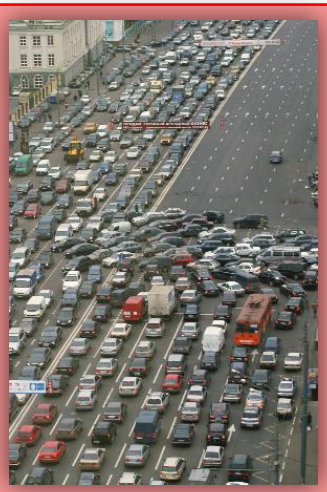

High Data Throughput

No Offline Storage Allowed

Programmer Friendly

System on Chips

- ARM or Intel Atom SoC
  - Low Power Consumption
  - Low Cost
  - High CPU Performance per Watt
- What about I/O performance?

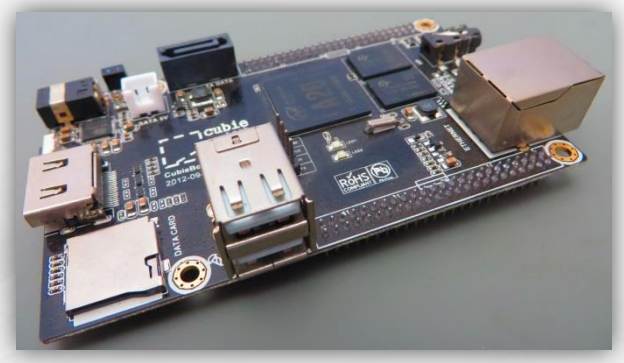

Cortex-A7

Cortex-A9

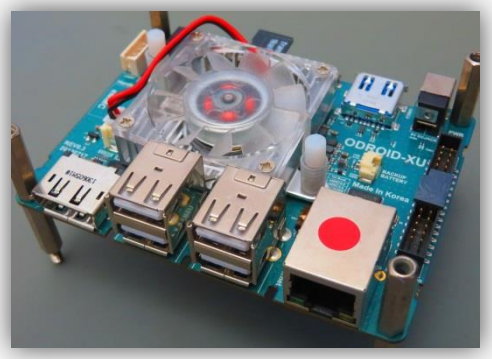

Cortex-A15

ARM SoC Performance Overview

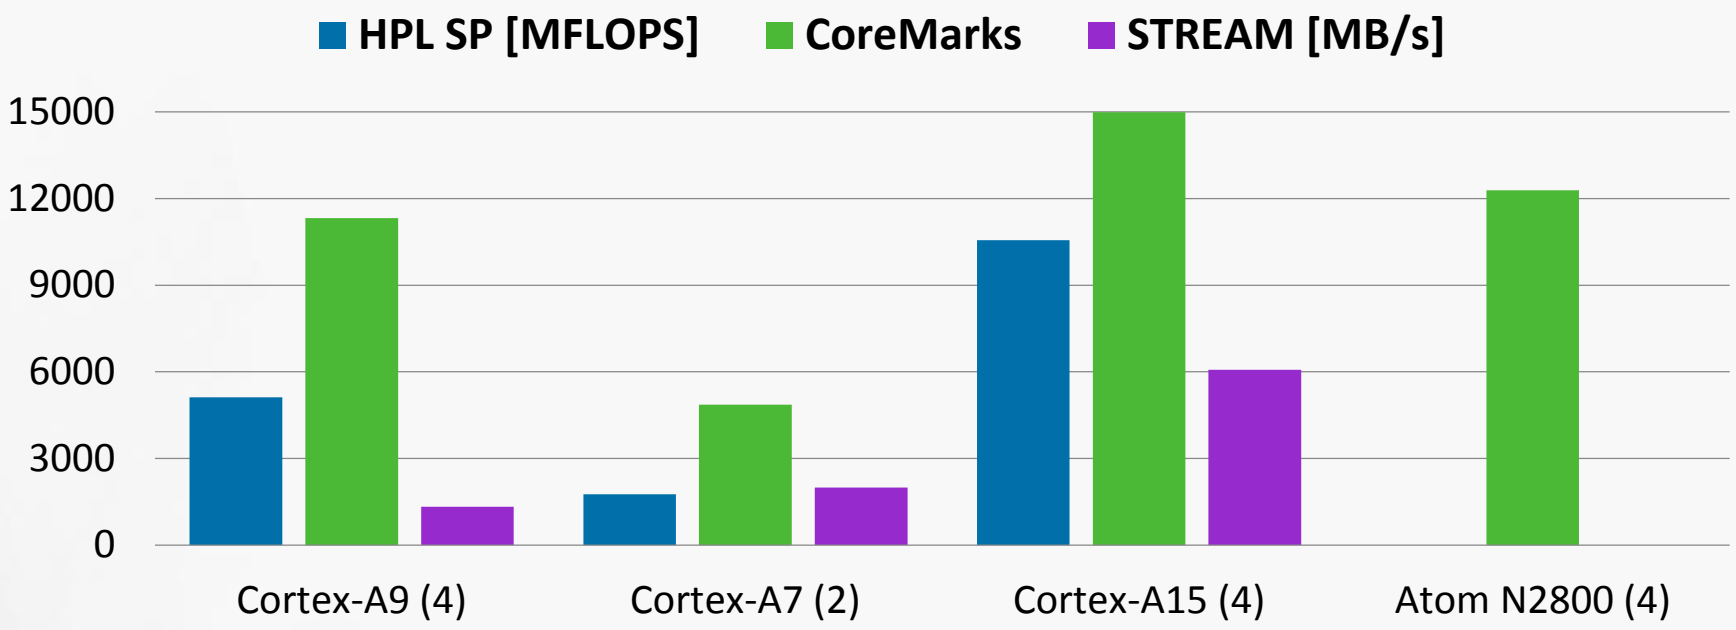

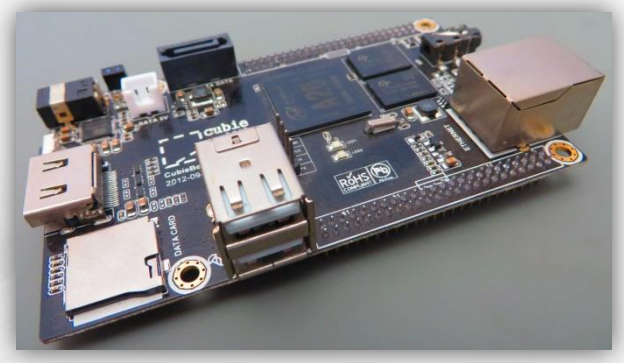

Cortex-A7

Cortex-A9

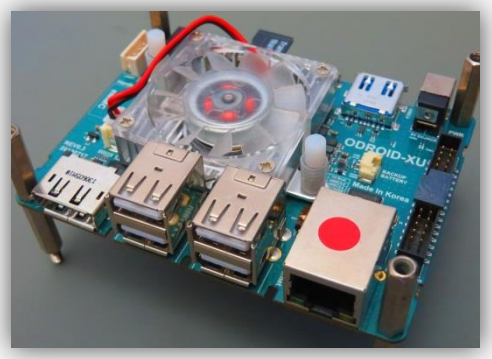

Cortex-A15

ARM SoC Performance Overview

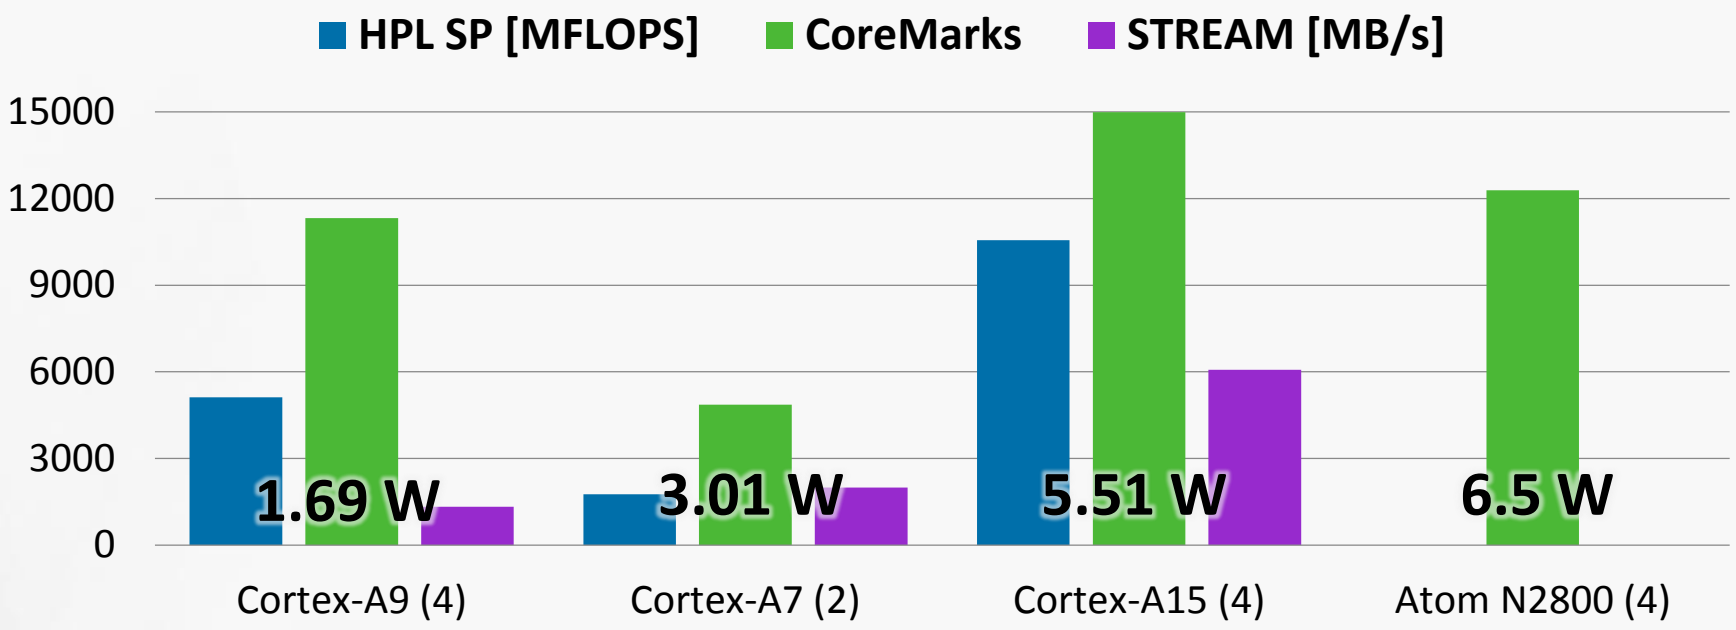

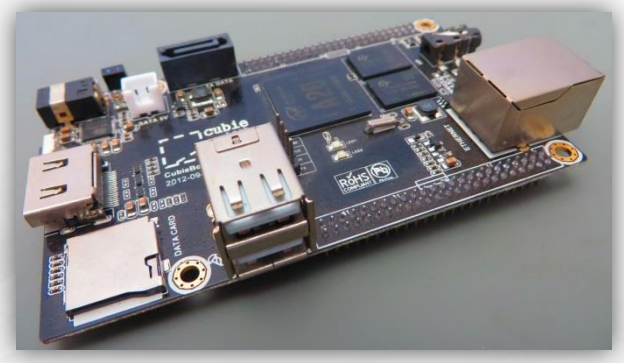

Cortex-A7

Cortex-A9

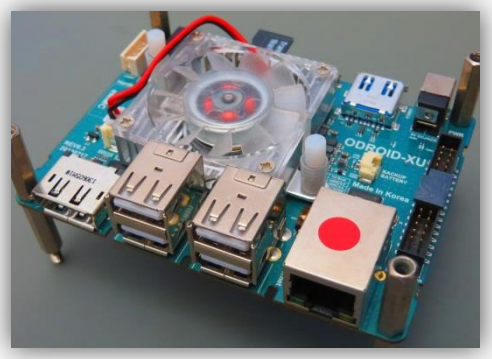

Cortex-A15

System on Chip
External I/O Ports

Ethernet

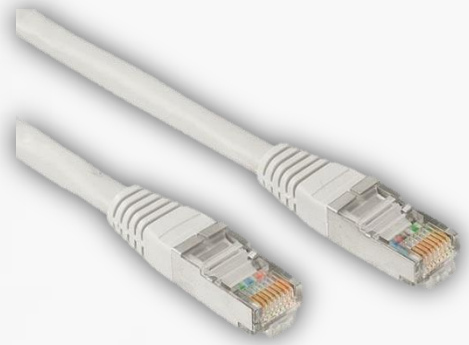

100 Mb/s - 1 Gb/s

12 - 125 MB/s

PCI-Express

N x 5 GT/s

≥ 500 MB/s

System on Chip
External I/O Ports

Ethernet

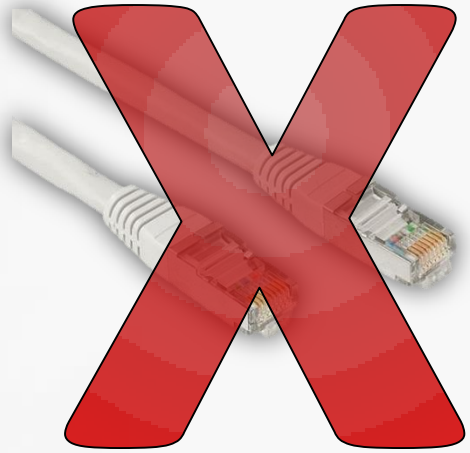

100 Mb/s - 1 Gb/s

12 - 125 MB/s

PCI-Express

N x 5 GT/s

≥ 500 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)

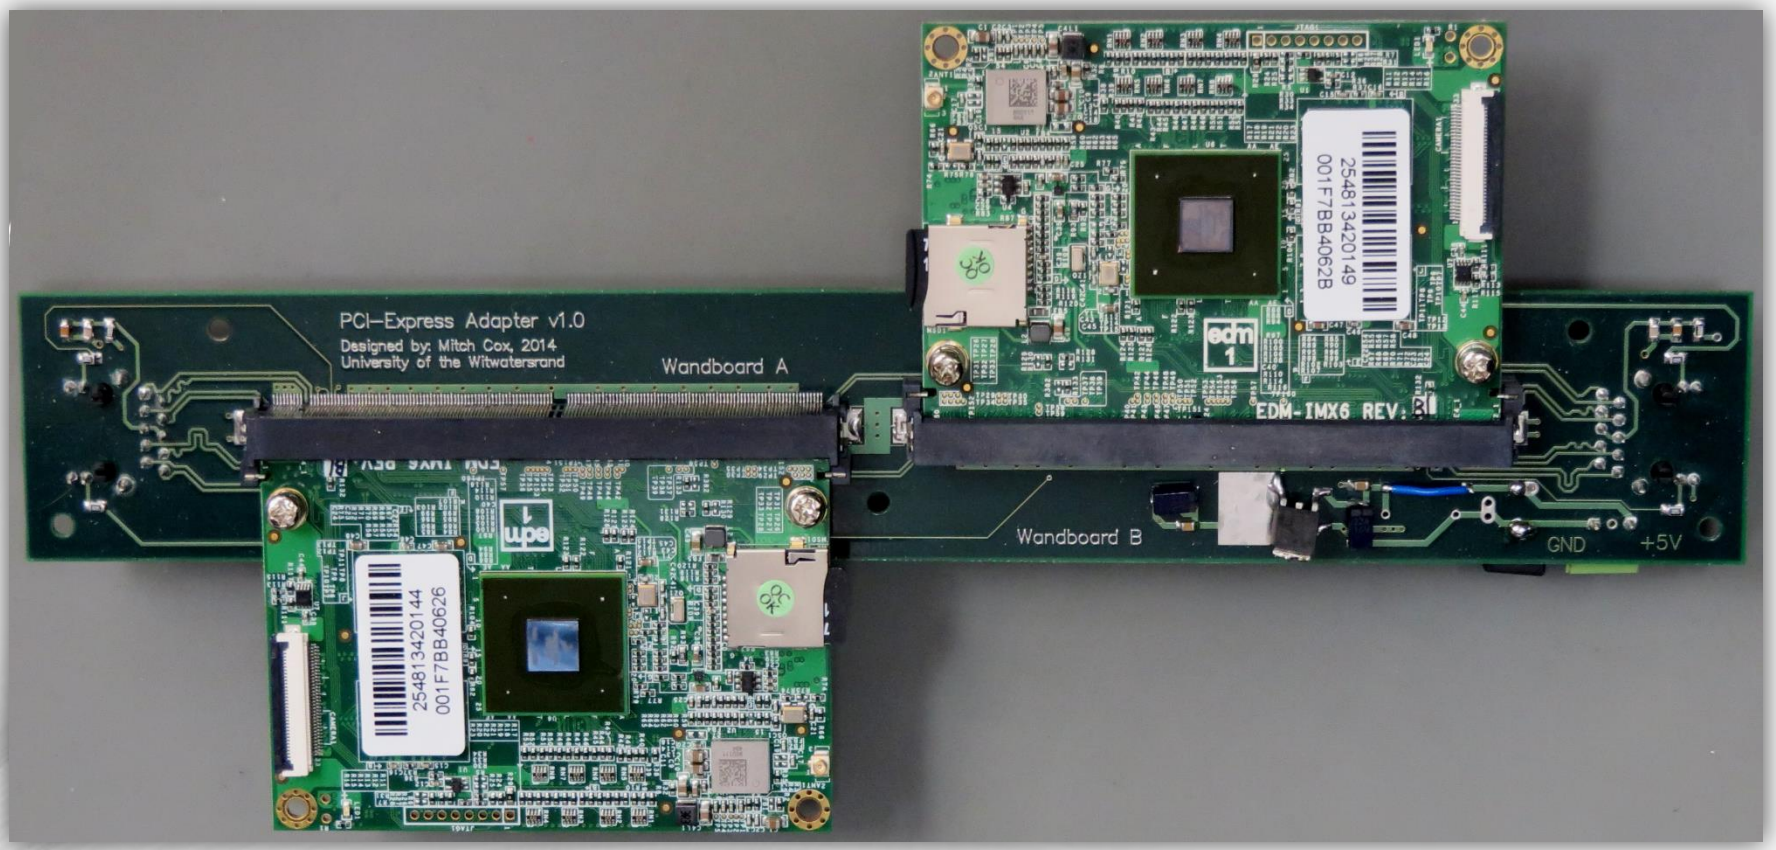

PCI-Express Test Results

- PCIe x1 Link on i.MX6 SoC:
  - 500 MB/s Theoretical

	CPU memcpy	DMA (EP)	DMA (RC)
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%

- 72 % of theoretical with Direct Memory Access (DMA)
  - Superior to Ethernet
  - Successful Proof of Concept
- 40 Gb/s PU needs 12 Freescale i.MX6 SoCs

Data Stream Computing Processing Unit

High Data Throughput
Ethernet Interface
40 Gb/s

Multiple
System on Chips
> 60 GFLOPS

Appears as a
Single System

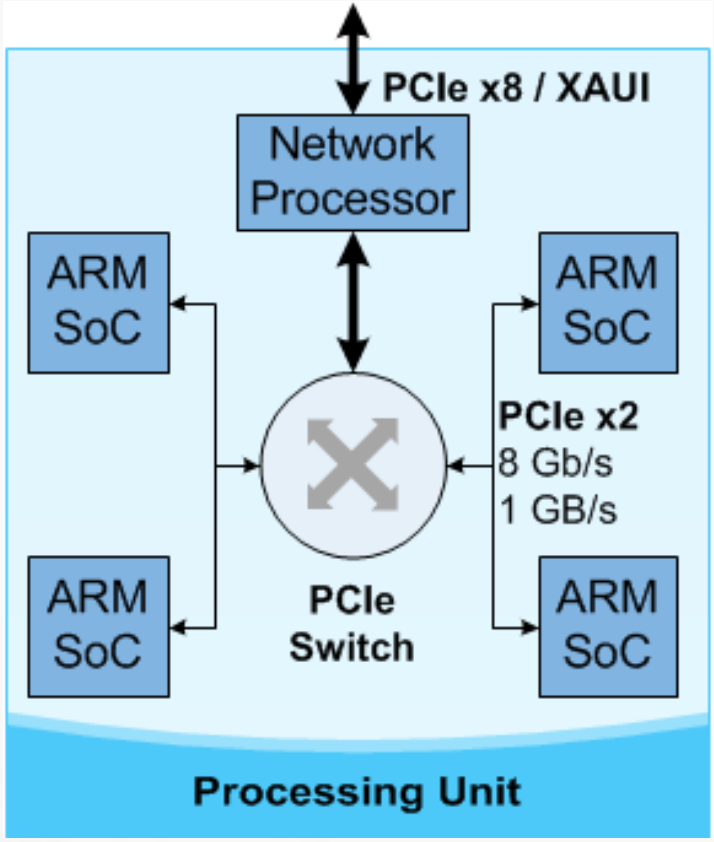

PCI-Express Adapter

- Designed and Built Adapter for PCIe Cluster

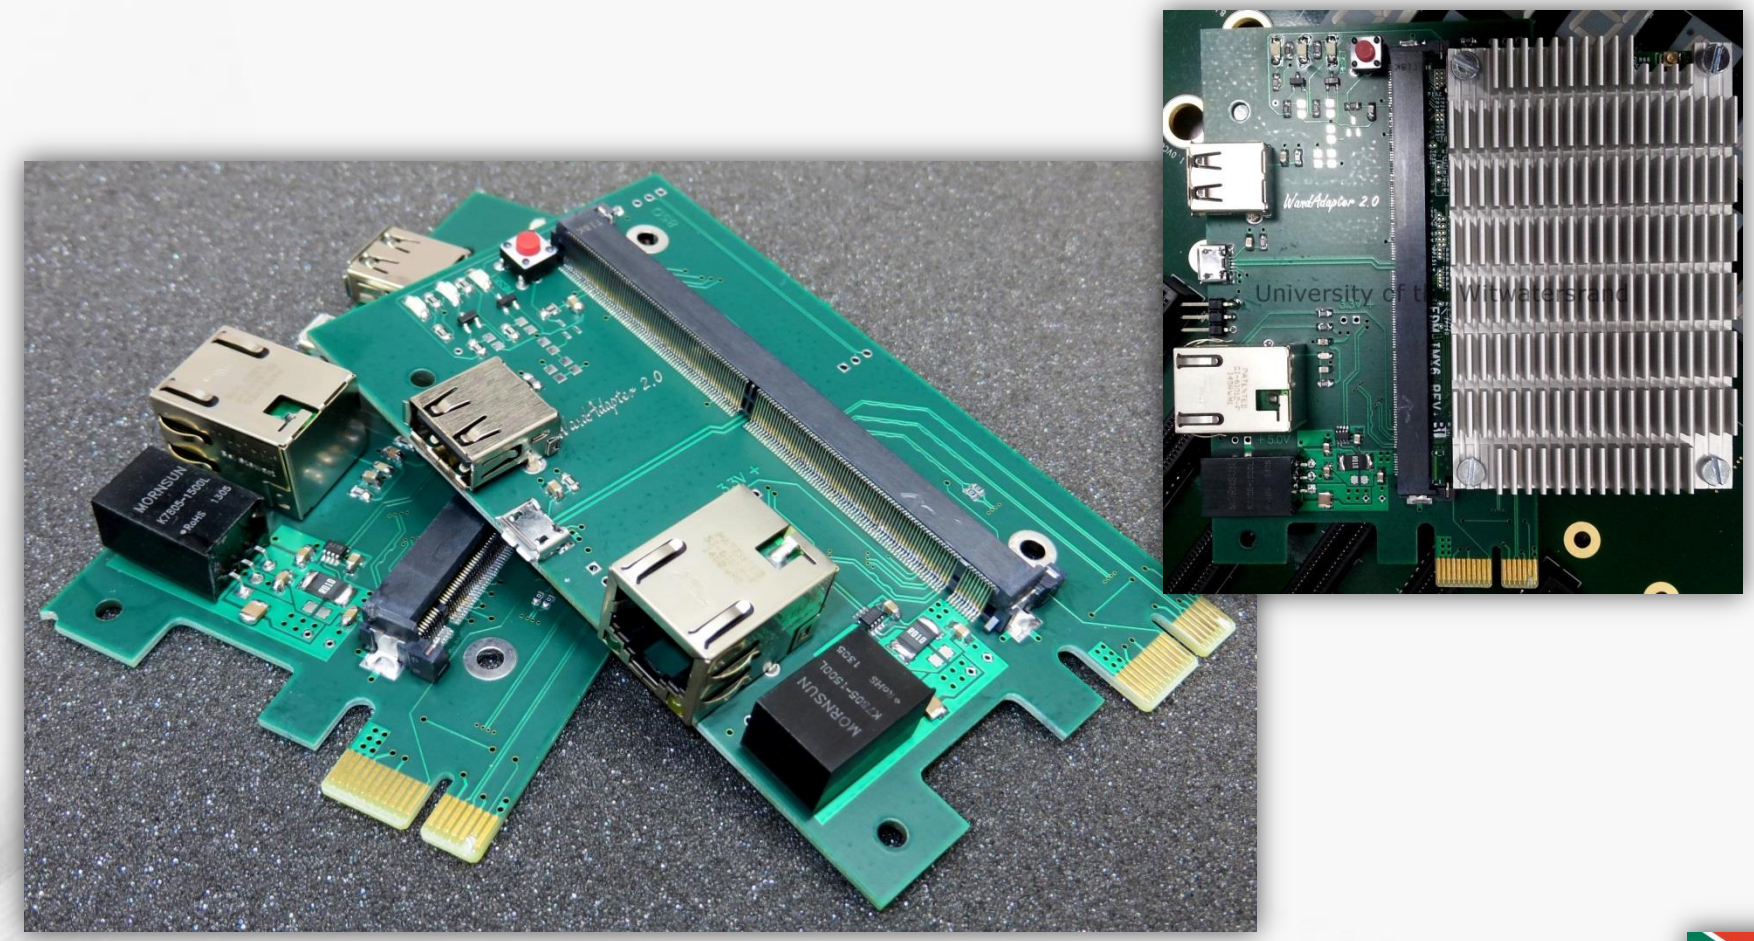

Manufactured in South Africa

PCI-Express Cluster

- 16 Gb/s Data Throughput (Theoretical)
  - 4 Wandboards + Freescale QorIQ Network Processor

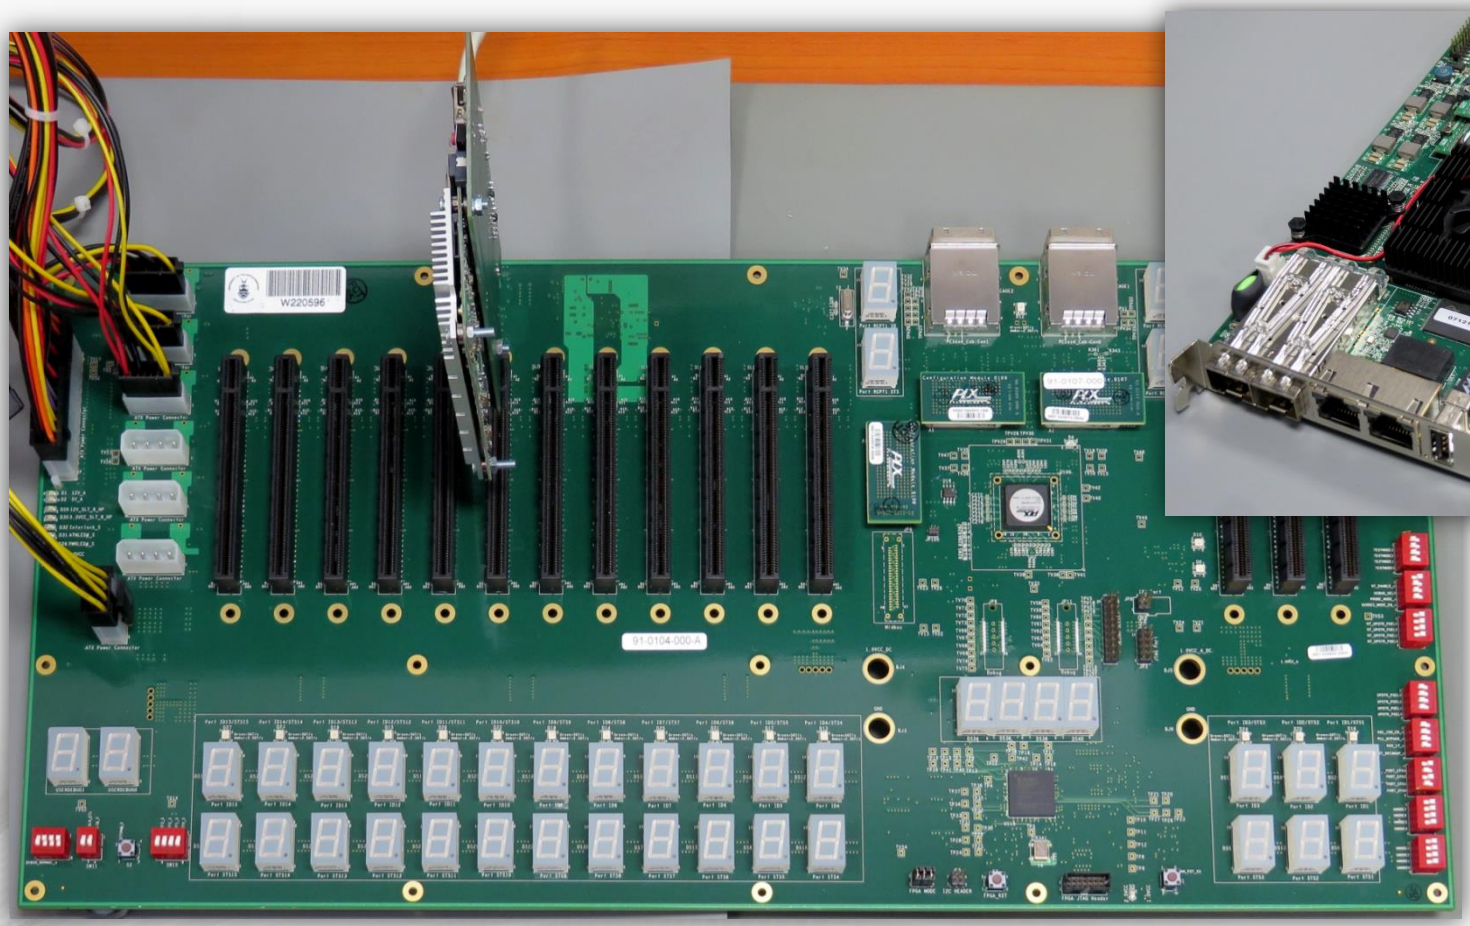

PCIe Ethernet Emulation Driver

- Transmit and Encapsulate Ethernet packets over PCIe
- Standard Ethernet Device (eth2)
- Communication:
  - Via the Network Processor (PCIe Root Complex)
  - Peer to Peer in Future
- Currently in Development

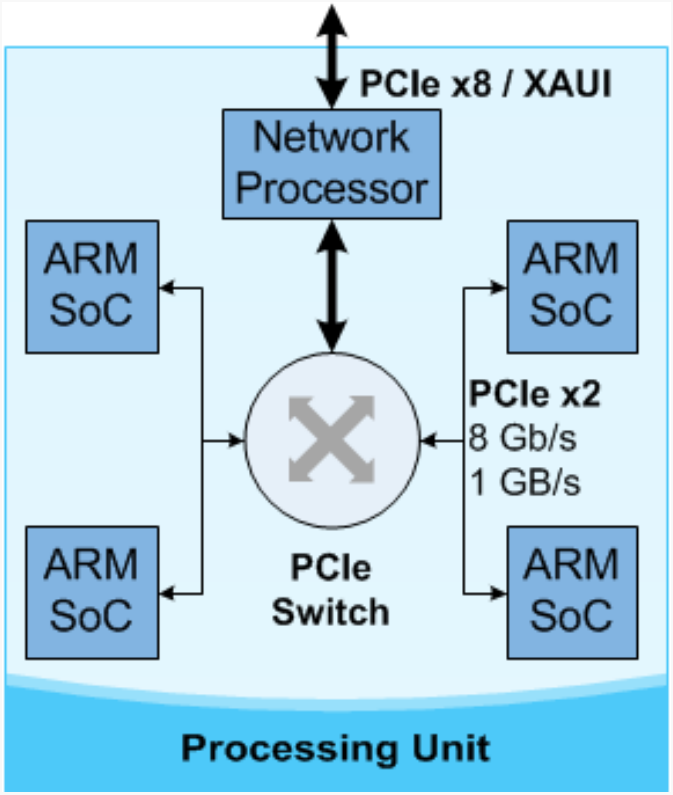

Summary

- Data Stream Computing

- ARM SoC-based Processing Unit

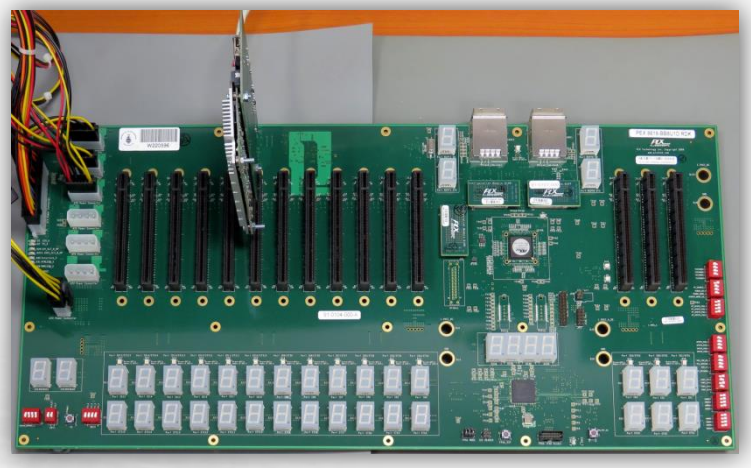

Questions or Comments?

MITCHELL.COX@CERN.CH

Acknowledgements

- The “Massive Affordable Computing Project” team:
  - Robert Reed, Thomas Wrigley, Matthew Spoor
  - MSc Supervisor: Prof. Bruce Mellado

- 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