



# RICH reconstruction on FPGAs

Christian Färber  
CERN Openlab Fellow  
LHCb Online group

On behalf of the LHCb Online group and the HTC Collaboration

8th LHCb Computing Workshop  
17.11.2016



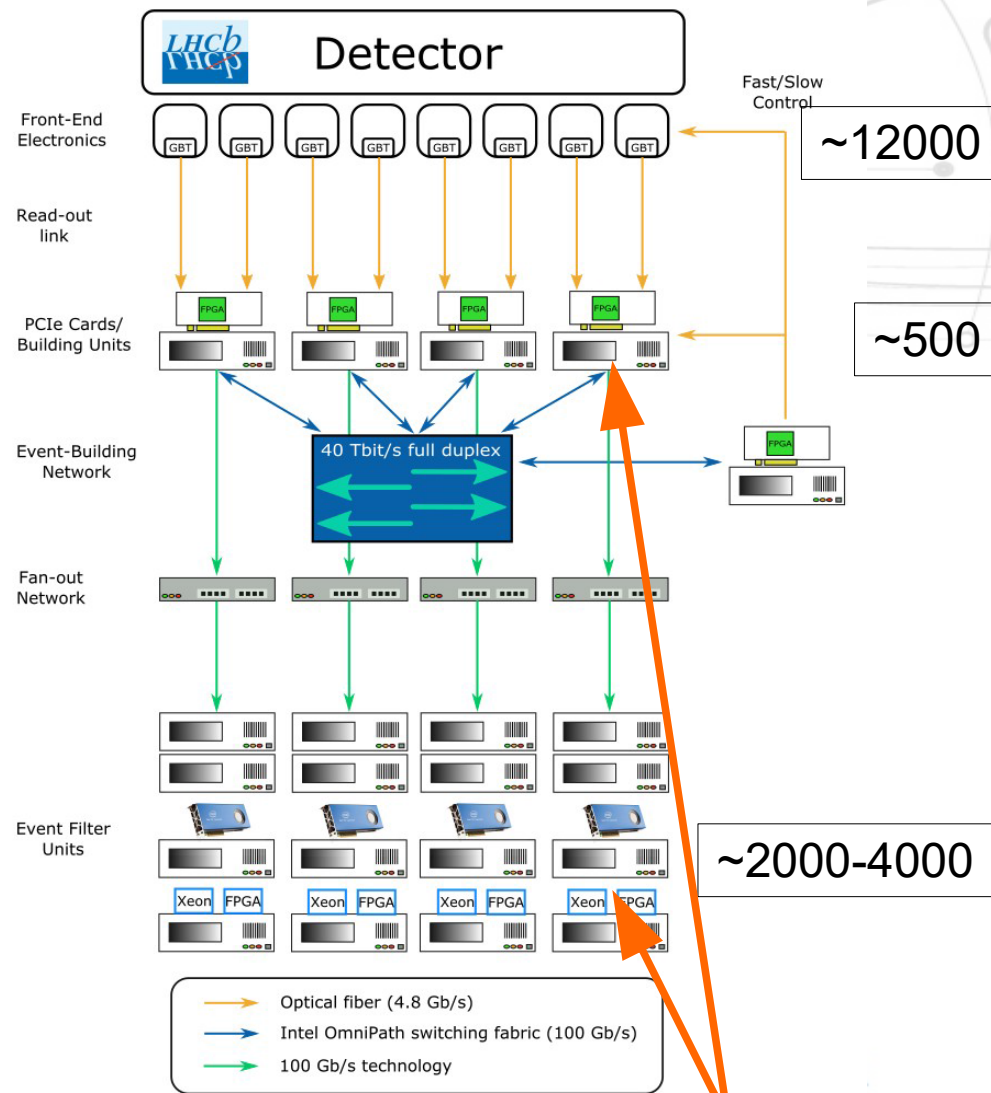
# HTCC

- High Throughput Computing Collaboration
- Members from Intel and CERN LHCb/IT
- Test Intel technology for the usage in trigger and data acquisition (TDAQ) systems
- Projects
  - Intel<sup>®</sup> KNL computing accelerator
  - Intel<sup>®</sup> Omni-Path 100 Gbit/s network
  - Intel<sup>®</sup> Xeon/FPGA computing accelerator



# Upgrade Readout Schematic


- Raw data input ~ 40 Tbit/s
- EFF needs fast processing of trigger algorithms, different technologies are explored.
- Test FPGA compute accelerators for the usage in:
  - Event building
    - Decompressing and re-formatting packed binary data from detector
  - Event filtering
    - Tracking
    - Particle identification
- Compare with: GPUs, Intel® Xeon/Phi and other compute accelerators



Which technologies?


# FPGAs as Compute accelerators

- Microsoft Catapult and Bing
  - Improve performance, reduce power consumption
- LHCb: Test for future usage in upgraded HLT farm:
  - Event building
  - Track fitting, pattern recognition, PID algorithms
- Current Test Devices in LHCb
  - Nallatech PCIe with OpenCL
  - Intel<sup>®</sup> Xeon/FPGA



# Test case: RICH PID Algorithm

- Calculate Cherenkov angle  $\Theta_c$  for each track  $\mathbf{t}$  and detection point  $\mathbf{D}$
- RICH PID is not processed for every event, processing time too long!




Reference: LHCb Note LHCb-98-040

## Calculations:

- solve quartic equation
- cube root
- complex square root
- rotation matrix
- scalar/cross products


# Nallatech 385 Board

- FPGA: Altera Stratix V GX A7
  - 234'720 ALMs, 940'000 Registers
  - 256 DSPs
- Programming model : OpenCL
- Host Interface: 8-lane PCIe Gen3
  - Up to 7.5GB/s
- Memory: 8GB DDR3 SDRAM
- Network Enabled with (2) SFP+ 10GbE ports
- Power usage:  $\leq 25W$  (GPU up to 300W)



# Nallatech 385 Board Results I


- Performance reference:
  - Intel Core i7-4770 CPU single thread vectorized



- Acceleration of factor up to 6 with Nallatech 385
- FPGA kernel faster, bottleneck data transfer

# Nallatech 385 Board Results II

- Energy efficiency comparison of three devices



Estimate: S.Sridharan


- It is estimated that the FPGA accelerator is a factor 4.3 more energy efficient than the GPU
  - Power measurements are planned to check!



# Intel<sup>®</sup> Xeon/FPGA

- Two socket system:

First: Intel<sup>®</sup> Xeon<sup>®</sup>  
E5-2680 v2



Second: Altera Stratix V GX A7 FPGA

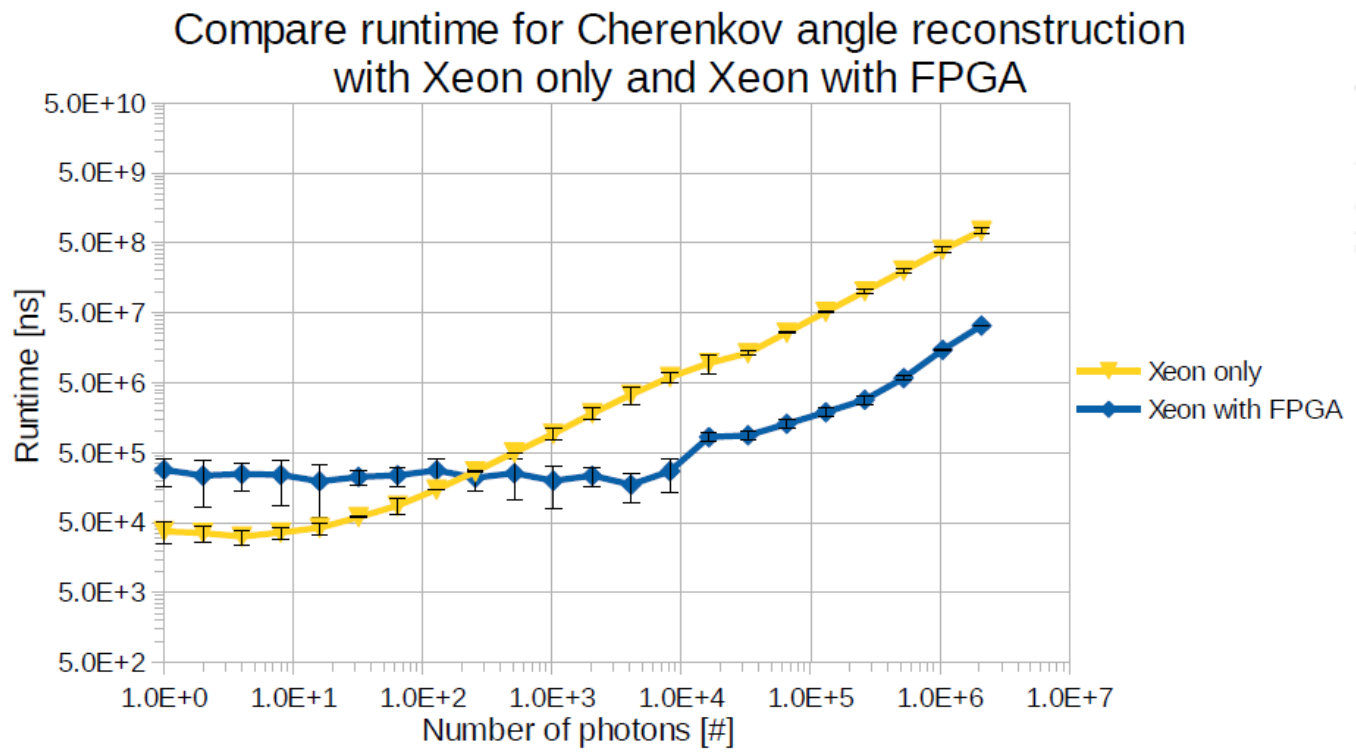
- 234'720 ALMs, 940'000 Registers, 256 DSPs
- Host Interface: high-bandwidth and low latency
- Memory: Cache-coherent access to main memory
- Programming model: Verilog now also OpenCL
- Power usage: Will be tested with next version

# Implementation of Cherenkov Angle reconstruction

- 748 clock cycle long pipeline written in Verilog
  - Additional blocks developed: cube root, complex square root, rot. matrix, cross/scalar product,...
  - Lengthy task in Verilog with all test benches (implementation took 2.5 months)
- Pipeline running with 200MHz → 5ns per photon
- FPGA resources:

FPGA Resource Type	FPGA Resources used [%]	For Interface used [%]
ALMs	88	30
DSPs	67	0
Registers	48	5

# Intel<sup>®</sup> Xeon/FPGA Results




- Acceleration of factor up to 35 with Intel<sup>®</sup> Xeon/FPGA
- Theoretical limit of photon pipeline: a factor 64 with respect to single Intel<sup>®</sup> Xeon<sup>®</sup> thread
- Bottleneck: Data transfer bandwidth to FPGA

# Compare Verilog - OpenCL

- Development time

2.5 months – 2 weeks

3400 lines Verilog – 250 lines C



OpenCL

**Faster**

**Easier**

- Performance

CBRT : x35 – x30

RICH : x35 – x26

**Comparable performance**


- FPGA resource usage

RICH Kernel	Verilog RTL	OpenCL
FPGA Resource Type	FPGA Resources used [%]	FPGA Resources used [%]
ALMs	88	63
DSPs	67	82
Registers	48	24

**Similar resource usage**

# Compare PCIe – QPI Interconnect


- Nallatech 385 PCIe vs. Intel<sup>®</sup> Xeon/FPGA QPI
- Both Stratix V A7 with 256 DSPs
- Programming model: OpenCL
- Reconstruct 1'000'000 photons



RICH Kernel

Compare Nallatech 385 and Intel Xeon/FPGA acceleration

RICH Cherenkov photon reconstruction (OpenCL)




# Future Tests

- Implement additional LHCb HLT algorithms
  - Tracking, decompressing and re-formatting packed binary data from detector, ...
- Compare performance with first multichip Intel<sup>®</sup> Xeon/FPGA system with Arria 10 FPGA
  - Broadwell+Arria10 arrived in our lab now 😊
- Measurements of Arria10 PCIe accelerators
- Implement the FPGA program in our framework, Gaudi, ...
- Power measurements
  - Compare with GPUs, ...



# New Intel<sup>®</sup> Xeon/FPGA with Arria10 FPGA

- Multichip package including:
  - Intel<sup>®</sup> Xeon<sup>®</sup> E5-2600 v4
  - Intel<sup>®</sup> Arria10 GX 1150 FPGA
    - 427'200 ALMs, 1'708'800 Registers, 1'518 DSPs
- Hardened floating point add/mult blocks!
- Host Interface: Bandwidth 5x higher than Stratix V version
- Memory: Cache-coherent access to main memory
- Programming model : Verilog soon also OpenCL



# Power Measurements

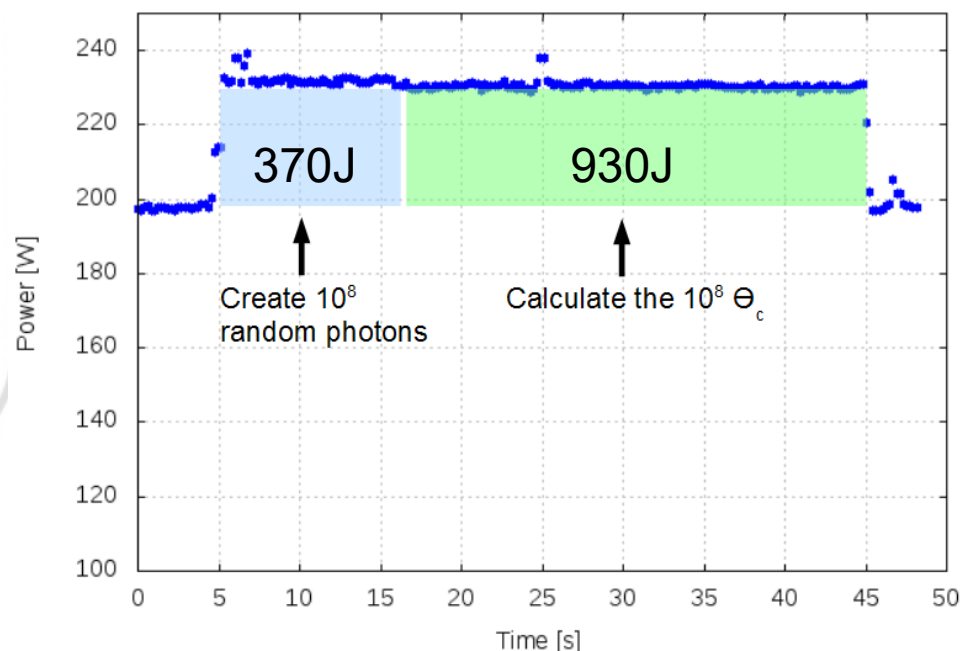
- Just started with power meter
- Run and analyse scripts are finished
- We have measurement permission for Intel<sup>®</sup> Xeon/FPGA Skylake + Arria10



– We will get this  
2017 – Q1

- Also Xeon CPUs, GPUs and other PCIe FPGA accelerators will follow
- Interesting metric is: photons/(s\*J)


Power measurement: Xeon CPU (E5-2680 v2) single thread






# Summary

- Results are very encouraging to use FPGA acceleration in the HEP field
- Intel<sup>®</sup> Xeon/FPGA accelerator performs better than the Nallatech PCIe board using the same FPGA
- Programming model with OpenCL very attractive and mandatory for HEP field
- Also other experiments want to test the usage of the Intel<sup>®</sup> Xeon/FPGA with Arria10!
- High bandwidth interconnect and modern Arria10 FPGA lets expect high performance and performance per Joule for HEP algorithms! Don't forget Stratix10!





# Thank you



Christian Färber,  
8th LHCb Computing Workshop – 17.11.2016





# Backup



hutterstock

Christian Färber,  
8th LHCb Computing Workshop – 17.11.2016



CERN openlab