Acceleration of Cherenkov angle Reconstruction with the new Intel® Xeon/FPGA compute platform for the particle identification in the LHCb Upgrade

Christian Färber
CERN Openlab Fellow
LHCb Online group

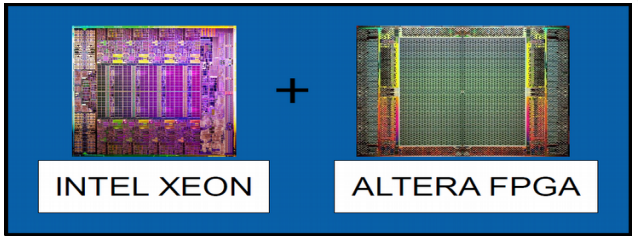

On behalf of the LHCb and HTC Collaboration

CHEP16
13.10.2016, San Francisco

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
  - KNL computing accelerator
  - Omni-Path 100 Gbit/s network
  - Xeon/FPGA computing accelerator

LHCb detector

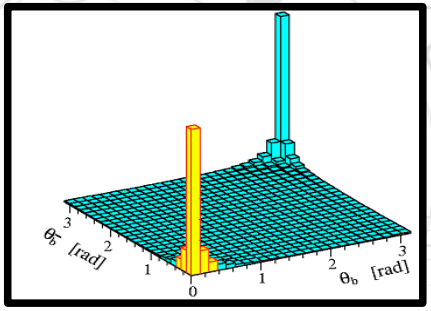

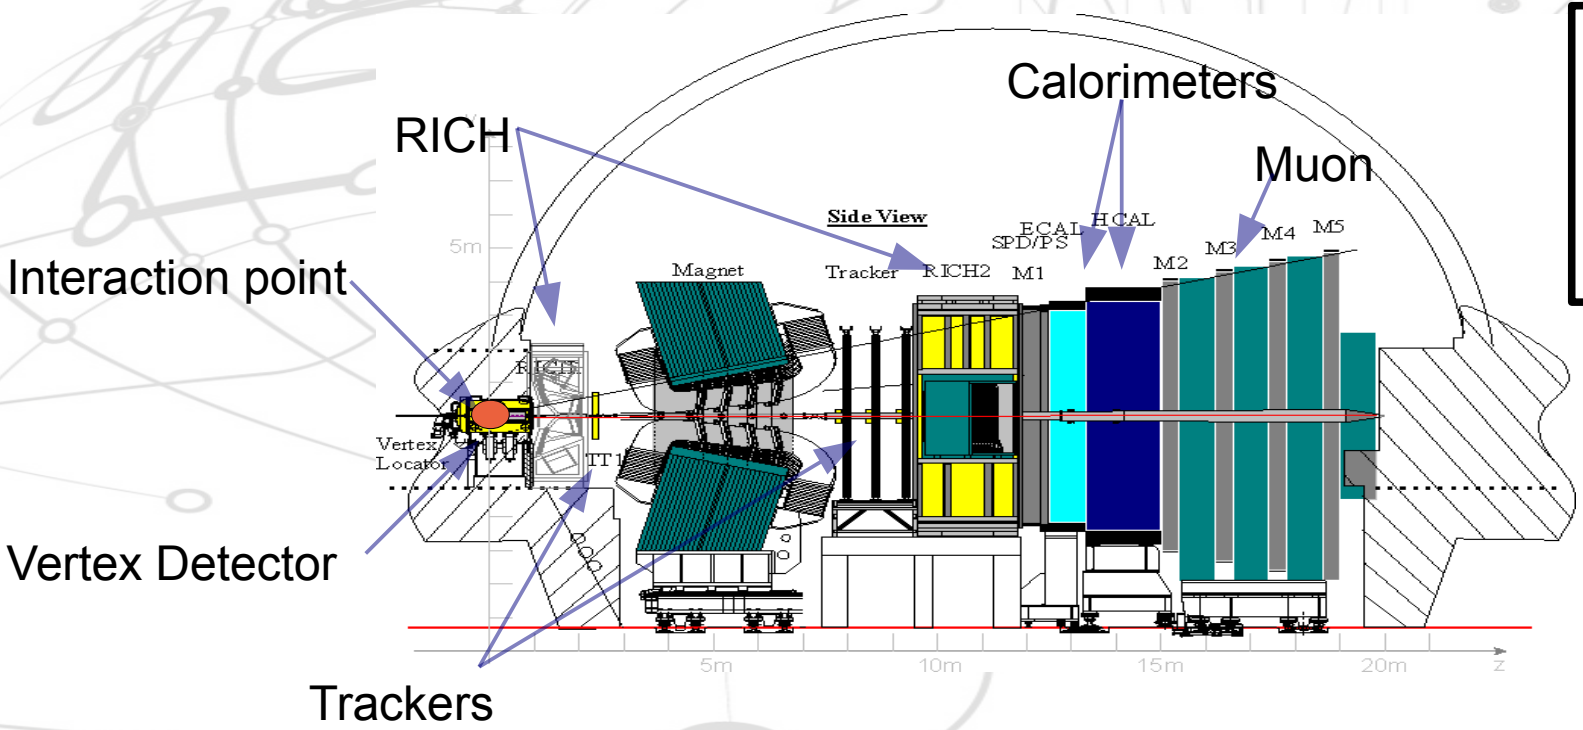

- Single-arm spectrometer designed to search new physics through measuring CP violation and rare decays of heavy flavour mesons.

- 40 MHz proton proton collisions
- Trigger with 1 MHz, upgrade to 40MHz
- Bandwidth after upgrade up to 40TBit/s

Nice overview of all the work:
Oral-250 C. Bozzi
Track 9, Tuesday

Run3 Upgrade LHCb

- New High Level Trigger farm for raw data input of ~ 40 Tbit/s!
- Different technologies are explored to realize fast and efficient processing of trigger algorithms.

- Test FPGA compute accelerators for the usage in:
  - Event building
    - Decompressing and re-formatting packed binary data from detector
  - Event filtering
    - Tracking
    - Particle identification

- Test system is the new Intel® Xeon/FPGA prototype!

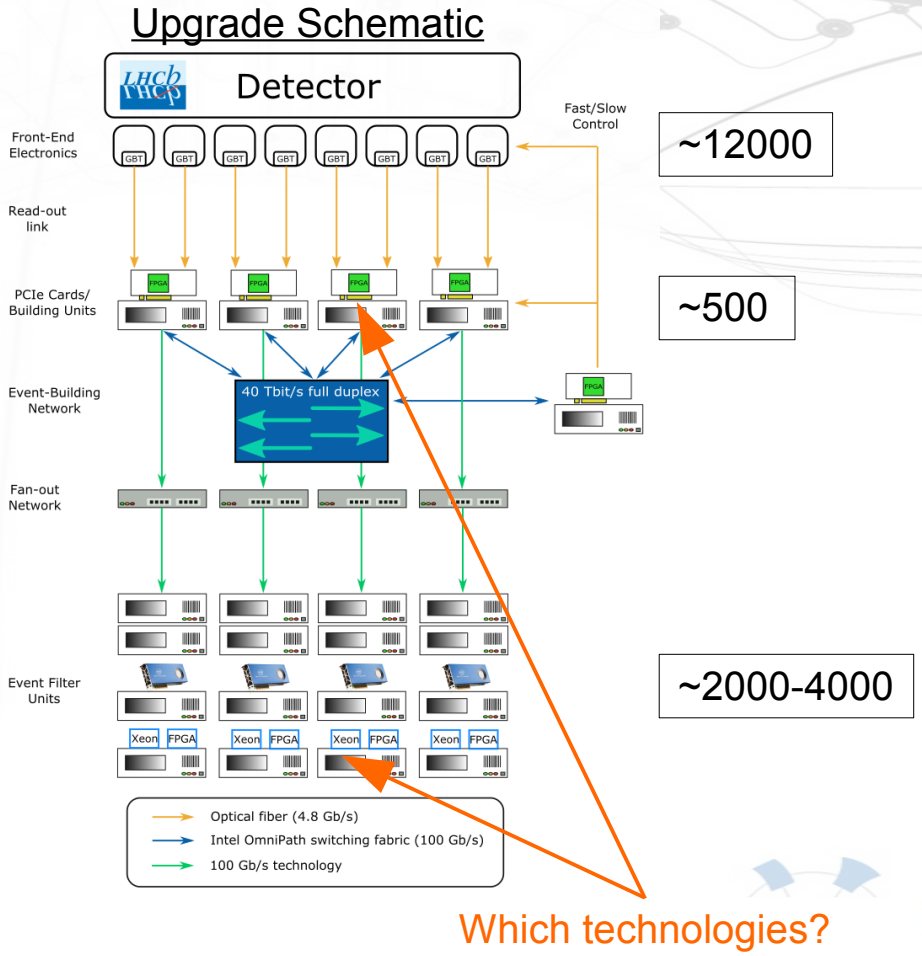

FPGAs as Compute accelerators

- Microsoft Catapult and Bing
  - Improve performance, 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 385
  - Intel® Xeon/FPGA

Intel® Xeon/FPGA

- Two socket system:

  First : Intel® Xeon® E5-2680 v2

  Second : Altera Stratix V GX A7 FPGA

  - 234'720 ALMs, 940'000 Registers, 256 DSPs

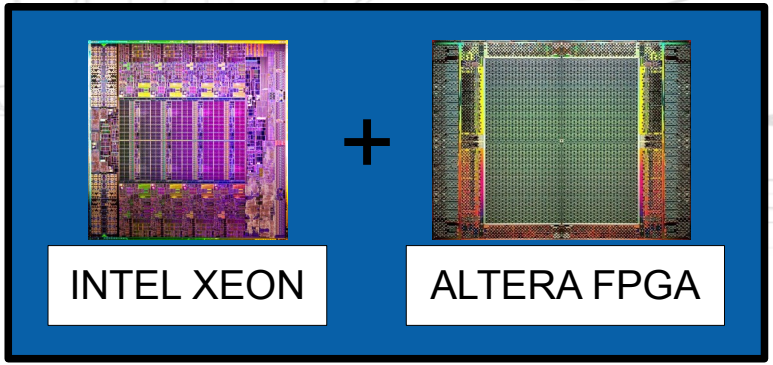

- Host Interface: high-bandwidth and low latency (QPI)
- Memory: Cache-coherent access to main memory
- Programming model : Verilog now also OpenCL
- Power usage: To be tested

Test case: RICH PID Algorithm

- Calculate Cherenkov angle Θ_c for each track **t** and detection point **D**
- RICH PID is not processed for every event, processing time too long!

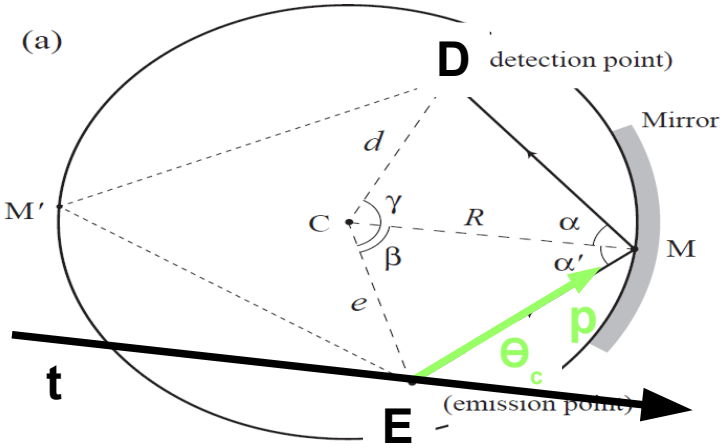

Reference: LHCb Note LHCb-98-040

Calculations:

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

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 for RICH kernel:

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

Intel® Xeon/FPGA Results

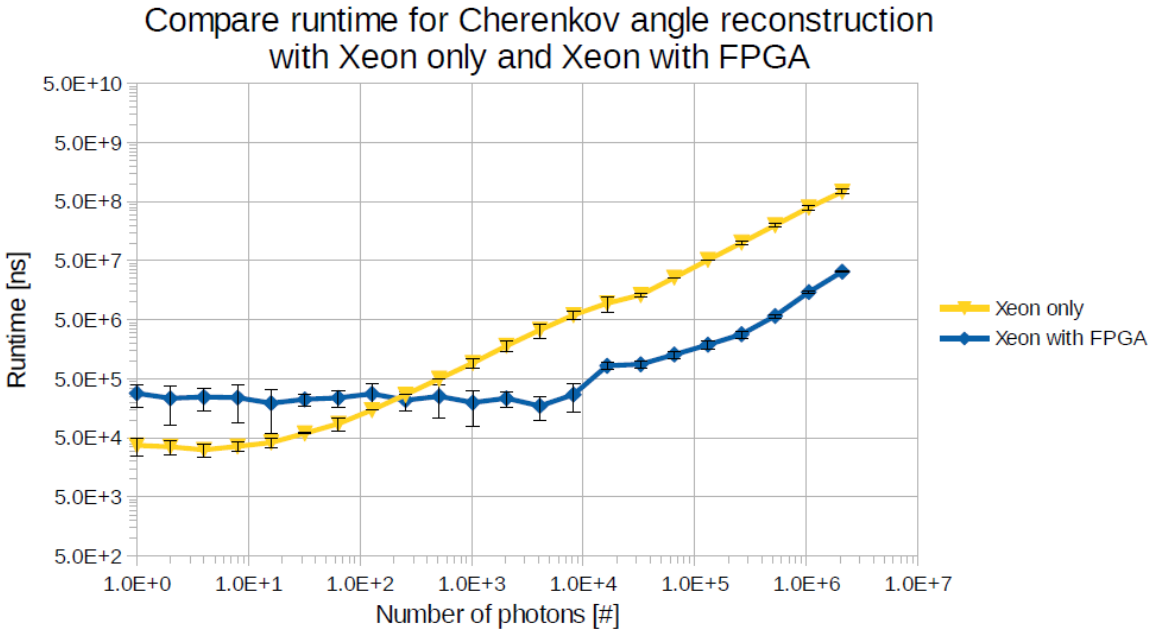

- Acceleration of factor up to 35 with Intel® Xeon/FPGA with respect to single Intel® Xeon thread
- Theoretical limit of photon pipeline: a factor 64 with respect to single Intel® Xeon thread
- Bottleneck: Data transfer bandwidth to FPGA

Compare Verilog - OpenCL

- Development time

  2.5 months – 2 weeks

  3400 lines Verilog – 250 lines C

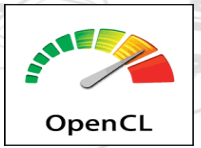

Faster

Easier

- Performance

  Cube root: x35 – x30

  RICH kernel: 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 Interconect

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

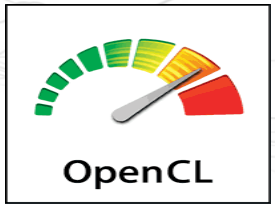

RICH Kernel

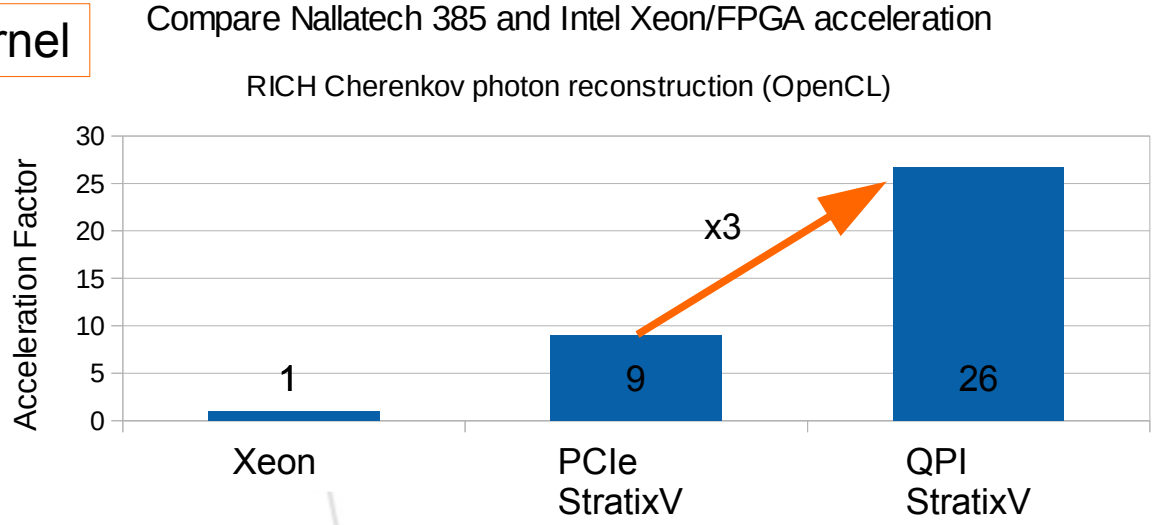

Future Tests

- Implement additional LHCb HLT algorithms
  - Tracking, decompressing and re-formatting packed binary data from detector, ...
- Compare performance with the new Intel® Xeon/FPGA compute platform with Arria 10 FPGA
  - Hardened floating point mult/accumulate blocks
- Test Nallatech CAPI (cache-coherent)
- Compare Verilog - OpenCL AFUs
- Power measurements
  - Compare with GPUs

Summary

- Results are very encouraging to use FPGA acceleration in the HEP field
- Intel® Xeon/FPGA compute platform performs better than server with Nallatech PCIe accelerator using the same FPGA
- FPGAs are strong in performance per Joule
  - Measurements and comparison to GPUs and others coming soon
- Programming model with OpenCL very attractive
  - Faster and easier algorithm implementation
- Also other experiments want to test the usage of the Intel® Xeon/FPGA with Arria10!

Backup

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: ≤ 25W (GPU up to 300W)

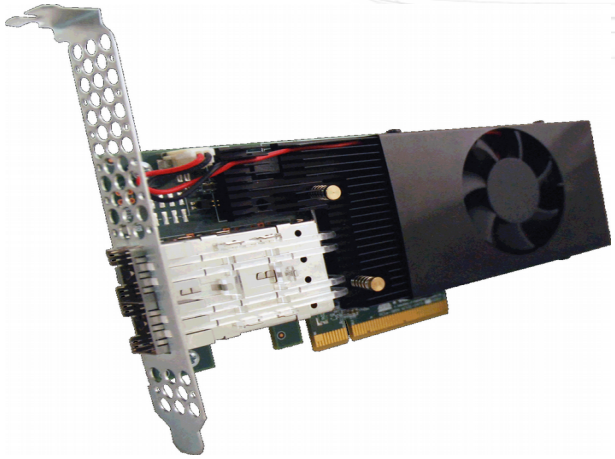

First results with Intel® Xeon/FPGA I

- Sorting of INT arrays with 32 elements
  - Implemented pipeline with 32 array stages
  - FPGA sort is x50 faster than single Xeon thread

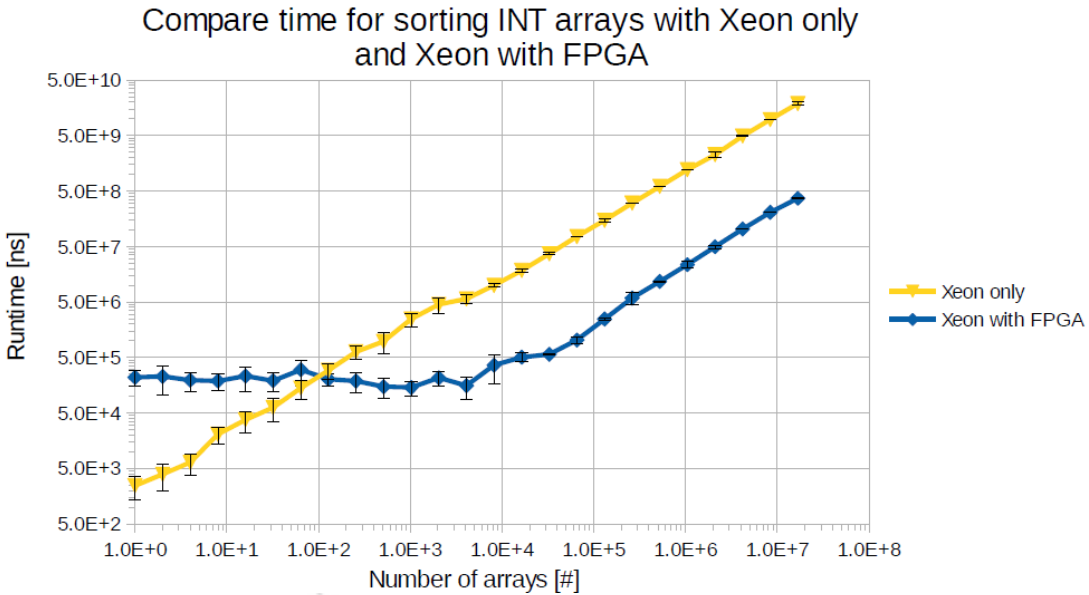

First results with Intel® Xeon/FPGA II

- Mandelbrot with floating point precision
  - Implemented 22 fpMandel pipelines running at 200MHz, each handles 16 pixels in parallel (total: 352 pixels).
  - FPGA is x12 faster as Xeon running 20 threads in parallel.
  - Used 72/256 DSPs
  - Reuse of data on FPGA high