Particle identification on an FPGA accelerated compute platform for the LHCb Upgrade

Christian Färber[1], Niko Neufeld[1], Rainer Schwemmer[1], Jonathan Machen[2]

{christian.faerber, niko.neufeld, rainer.schwemmer }@cern.ch, jonathan.machen@intel.com

[1]Experimental Physics Department, CERN, Geneva, Switzerland

[2]DCG IPAG, EU Exascale Labs, Intel, Switzerland

On behalf of the HTC Collaboration - June 2016

LHCb Upgrade

Current situation

- Raw data output ~ 10 Tbit/s (not zero-suppressed)
- Hardware and Software Trigger for selecting events
- Application example: Process ~ 10^{11} B_d^0 decays for detecting one $B_d^0 \rightarrow \mu^+\mu^-$
- Upgrade program foresees 10x higher pp collision rate.

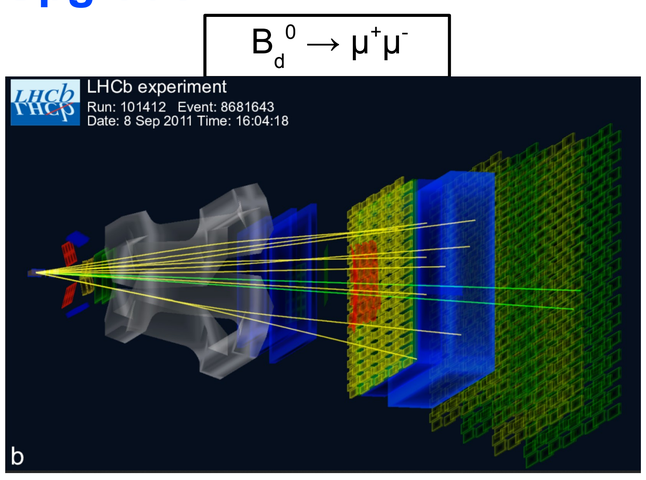

After the Upgrade

- 2018 LHCb will change its detector to a trigger-free readout, reading every collision (one every 25 ns) and a much more flexible software-based trigger system, the Event Filter Farm (EFF).
- Events will be processed and triggered on an event-by-event basis by the Event Filter Farm.

The Event Filter Farm

- Raw data input ~ 40 Tbit/s (already zero-suppressed by the front-end electronics)
- Only Software Trigger for selecting the events
- EFF needs fast processing of trigger algorithms (decision within O(10) µs). → Different technologies have to be explored.
- High-speed interconnect technology has to be investigated and used.
- Test FPGA compute accelerators for the usage in Event building, Tracking and particle identification of the upgraded High-Level-Trigger farm and compare with: GPUs, Intel® Xeon/Phi and other computing accelerators

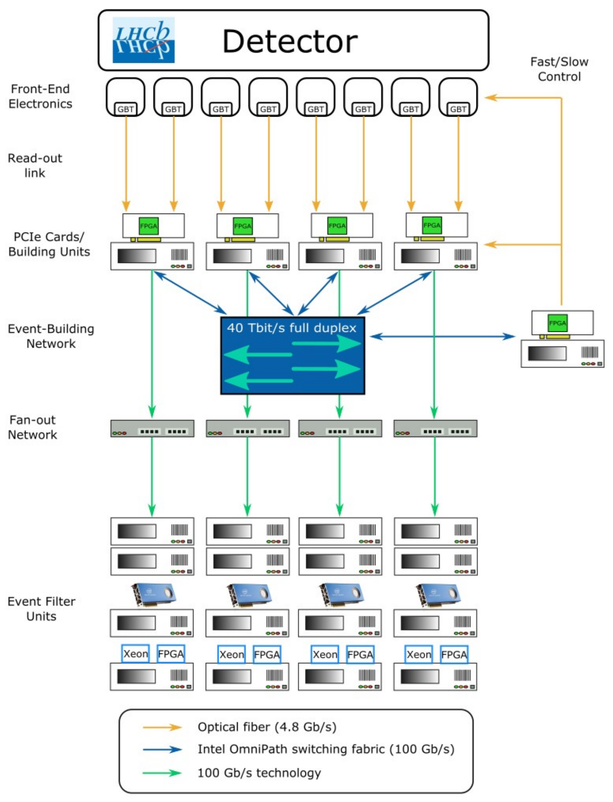

First Test Cases

Sorting

- Runtime scales on CPU with n x log(n) n = number of elements
- On FPGA with pipeline and parallel compare it depends only on pipeline clock frequency
- FPGA sort is a factor x50 faster than single Intel® Xeon thread

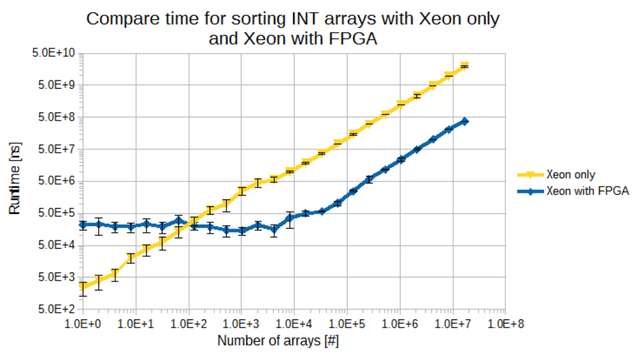

Cubic root

- Implementation for floats with shifting n-th root algorithm
- Implemented 7 root pipelines for parallel processing (200MHz)
- FPGA cubic root is a factor x35 faster than single Intel® Xeon thread

Mandelbrot

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

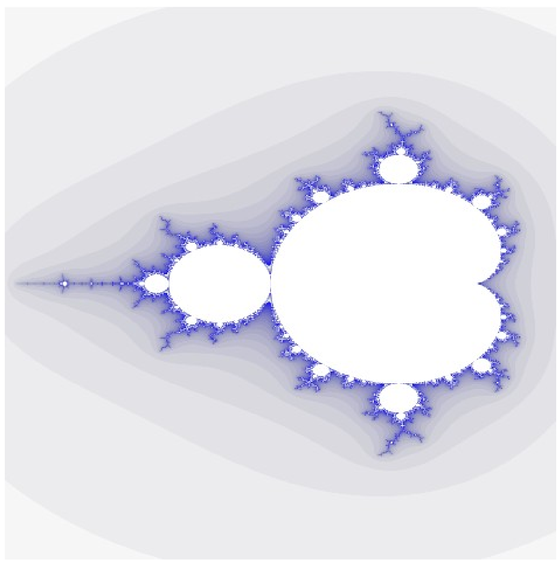

High Throughput Computing Collaboration

HTCC

- Members from Intel and CERN LHCb/IT
- Test Intel technology for the usage in trigger and data acquisition (TDAQ) systems
- Projects
  - Intel® Omni-Path 100 Gbit/s network
  - Intel® Xeon/Phi computing accelerator
  - Intel® Xeon/FPGA computing accelerator

Intel® Xeon/FPGA

Prototype

- 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)

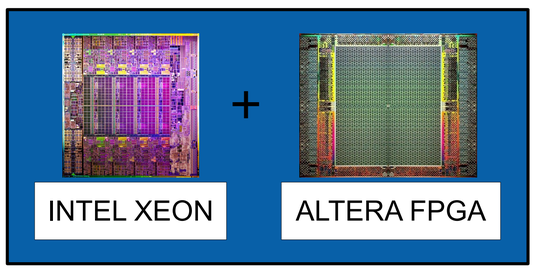

- Memory: Cache-coherent access to main memory
- Programming model: Verilog now also OpenCL
- Power usage: FPGAs are very power efficient up to a factor x10 lower than GPUs → Measurements will follow soon

Future system

- Intel® Xeon CPU and FPGA in one package
- Including newest high performance Altera FPGA: Arria 10
- Faster interface for interconnect of CPU and FPGA

Cherenkov Angle Reconstruction

Algorithm

- Particle travelling faster as speed-of-light in medium emitting cherenkov radiation in an angle depending on the particle speed
- Calculate Θ_c knowing points **D,C,E** and particle track **t**

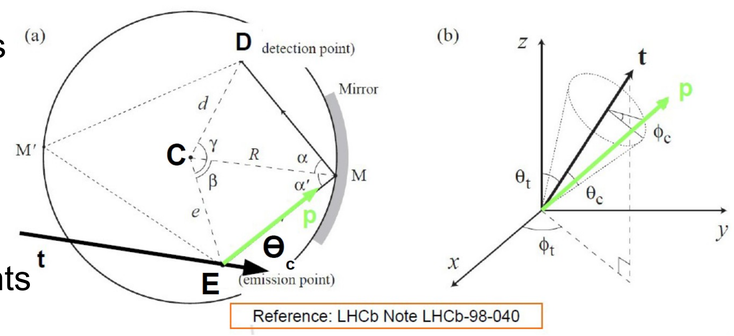

Reference: LHCb Note LHCb-98-040

Implementation of Cherenkov Angle reconstruction

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

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

Results so far very encouraging

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

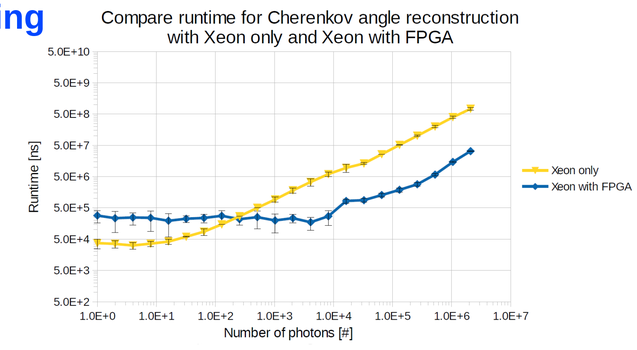

Future Tests

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