

# TIPP'17

## Acceleration of an particle identification algorithm used for the LHCb Upgrade with the new Intel<sup>®</sup> Xeon<sup>®</sup>-FPGA



Christian Färber  
CERN Openlab Fellow  
LHCb Online group



On behalf of the LHCb Online group and the HTC Collaboration


TIPP2017, Beijing  
25.05.2017

# 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<sup>®</sup>-FPGA computing accelerator



# Detector Example: LHCb



- 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



# Future challenges

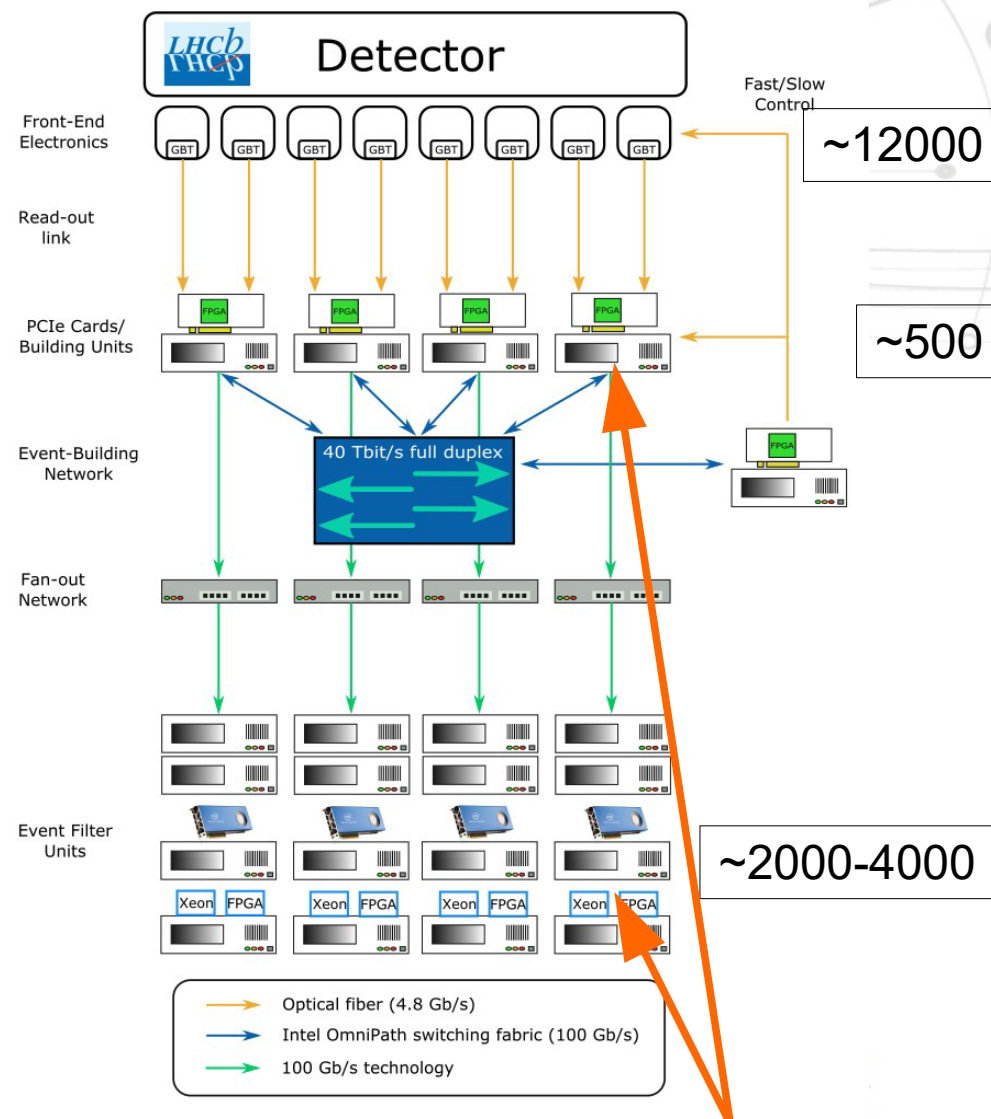
- Higher luminosity from LHC
- Upgraded sub-detector Front-Ends
- Removal of hardware trigger
- Software trigger has to handle:
  - Larger event size (50KB to 100KB)
  - Larger event rate (1MHz to 40MHz)

Data Network - Throughput



# Upgrade Readout Schematic

- Raw data input  $\sim 40$  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<sup>®</sup> Xeon<sup>®</sup>-Phi and other compute accelerators



Which technologies?

# FPGAs as Compute accelerators

- Microsoft Catapult and Bing
  - Improve performance,  
reduce power consumption
- Reduce the number of Von-Neumann abstraction layers
  - Bit level operations
- Power only logic cells and registers needed
- Current Test Devices in LHCb
  - Nallatech PCIe with OpenCL
  - Intel<sup>®</sup> Xeon<sup>®</sup>-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
- Memory: Cache-coherent access to main memory
- Programming model: Verilog and OpenCL

# New Intel® Xeon®-FPGA with Arria10 FPGA

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






# Sort with Xeon® and FPGA


- Sorting of INT arrays with 32 elements
  - Implemented pipeline with 32 array stages
  - FPGA sort is up to x117 faster than single Xeon® thread
  - Bandwidth to the FPGA is the bottleneck

Time ratio for sorting with Xeon only to Xeon with FPGA



# Test case: RICH PID Algorithm

- Calculate Cherenkov angle  $\Theta_c$  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 Stratix V

- 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



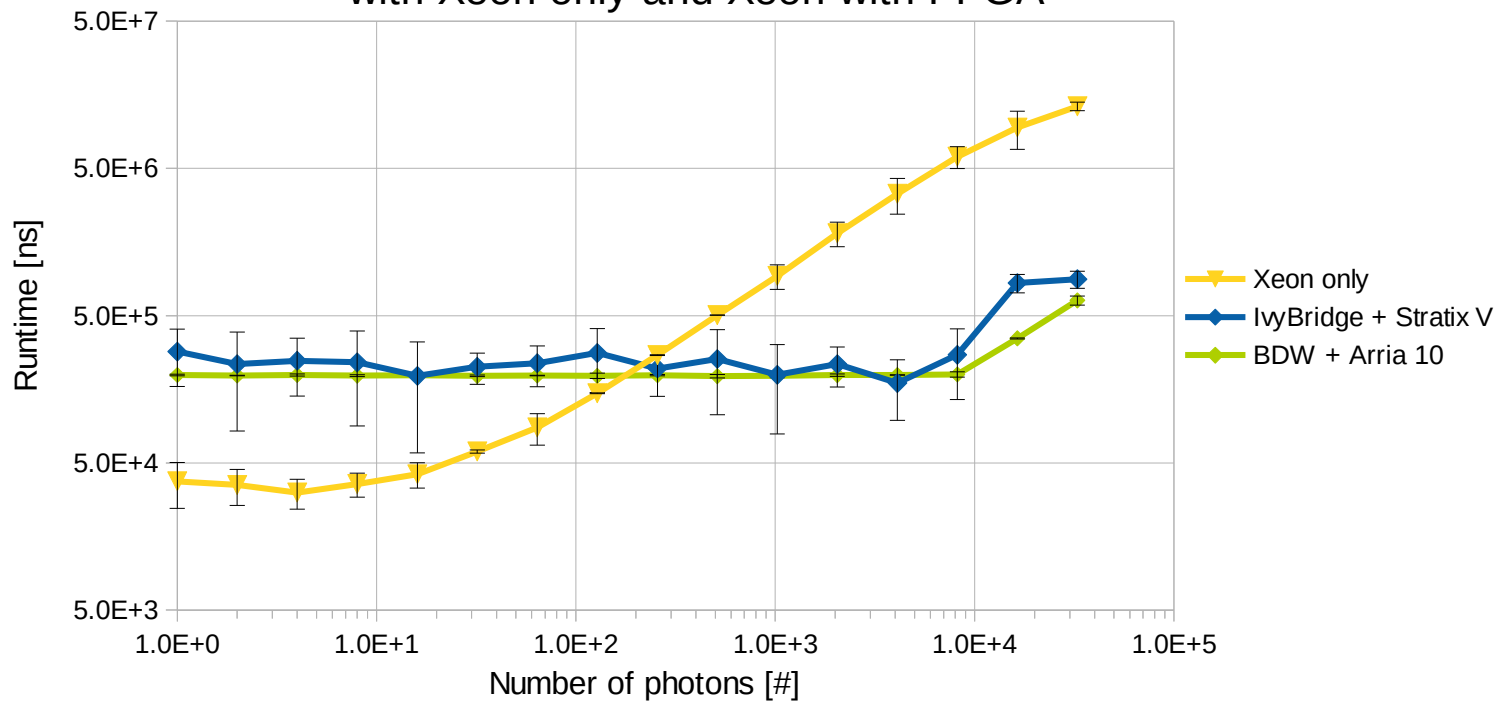
# Implementation of Cherenkov Angle reconstruction Arria 10

- 259 clock cycle long pipeline written in Verilog
  - Stratix V blocks ported using HFB: complex square root, rot. matrix, cross/scalar product,...
- Pipeline running with 200MHz → 5ns per photon
  - With Arria 10 GT FPGA 400 MHz possible
- FPGA resources:

FPGA Resource Type	FPGA Resources used [%]	For Interface used [%]
ALMs	32	18
DSPs (HFBs)	15	0
Registers	12	5

# Intel® Xeon®-FPGA Results

Compare runtime for Cherenkov angle reconstruction  
with Xeon only and Xeon with FPGA




- 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 for Stratix V FPGA, for Arria 10 FPGA a factor ~ 300
- Bottleneck: Data transfer bandwidth to FPGA, caching can improve this, tests ongoing

# Compare Verilog - OpenCL

- Development time

2.5 months – 2 weeks

3400 lines Verilog – 250 lines C



**Faster**

**Easier**

- Performance

Cube root : x35 – x30

RICH : x35 – x26

**Comparable performance**

- FPGA resource usage Stratix V


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<sup>®</sup>-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)



Christian Färber,  
TIPP2017, Beijing – 25.05.2017


# Future Tests

- Implement additional LHCb HLT algorithms
  - Kalman filter, decompressing and re-formatting packed binary data from detector, ...
- Compare performance with Intel<sup>®</sup> Xeon<sup>®</sup>-FPGA system with Skylake + Arria 10 FPGA
  - Skylake + Arria 10 arrived in our lab now 😊
- Measurements of Arria10 PCIe accelerators
- Compare Verilog vs. OpenCL
- Power measurements
  - Compare with GPUs, ...



# Summary

- Results are very encouraging to use FPGA acceleration in the HEP field
- Intel<sup>®</sup> Xeon<sup>®</sup>-FPGA accelerator performs better than the Nallatech PCIe board using the same FPGA
- Programming model with OpenCL very attractive and convenient for HEP field
- Also other experiments want to test the usage of the Intel<sup>®</sup> Xeon<sup>®</sup>-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!







# 谢谢



Christian Färber,  
TIPP2017, Beijing – 25.05.2017



# 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:  $\leq 25W$  (GPU up to 300W)

