

Acceleration of an Particle Identification Algorithm used for the LHCb Upgrade with the new Intel® Xeon®-FPGA*

Christian Färber¹, Rainer Schwemmer¹, Niko Neufeld¹, and Jonathan Machen²

¹ Experimental Physics Department, CERN, CH-1211 Geneva 23, Switzerland,
{[Christian.Faerber](mailto:Christian.Faerber@cern.ch), [Rainer.Schwemmer](mailto:Rainer.Schwemmer@cern.ch), [Niko.Neufeld](mailto:Niko.Neufeld@cern.ch)}@cern.ch

² DCG IPAG, EU Exascale Labs, Intel® Corporation, Badenerstrasse 549, 8048
Zürich, Switzerland, Jonathan.Machen@intel.com

Abstract. The LHCb experiment at the LHC will upgrade its detector by 2018/2019 to a 'triggerless' readout scheme, where all the readout electronics and several sub-detector parts will be replaced. The new readout electronics will be able to read out the detector at 40 MHz. This increases the data bandwidth from the detector down to the event filter farm to 40Tb/s, which also has to be processed to select the interesting proton-to-proton collisions for later storage. The architecture of such a computing farm, which can process this amount of data as efficiently as possible, is a challenging task and several compute accelerator technologies are being considered for use inside the new event filter farm.

In the high performance computing sector more and more FPGA compute accelerators are used to improve the compute performance and reduce the power consumption (e.g. in the Microsoft Catapult project and Bing search engine). Also for the LHCb upgrade, the usage of an experimental FPGA accelerated computing platform in the event building or in the event filter farm (trigger) is being considered and therefore tested. This platform from Intel® hosts a general Xeon® CPU and a high performance Arria® 10 FPGA inside a multi-chip package linked via a high speed and low latency link. On the FPGA an accelerator is implemented. The FPGA has cache-coherent memory access to the main memory of the server and can collaborate with the CPU.

A computing intensive algorithm to reconstruct Cherenkov angles for the LHCb RICH particle identification was successfully ported to the Intel® Xeon®-FPGA platform and accelerated. The results show that the Intel® Xeon®-FPGA platforms, which are built in general for high performance computing, are also very interesting for the High Energy Physics community.

* This work is done in collaboration with Intel® Corporation in the High-Throughput Computing Collaboration (HTCC) and the authors would like to thank Intel® Corporation for the support which made this work possible.

1 The Intel® Xeon®-FPGA

The Intel® Xeon®-FPGA prototype is a server machine which hosts an Intel® Xeon® E5-2600 v4 CPU and an Intel® Arria® 10 GX 1150 FPGA (see Fig. 1). In the FPGA, 427.200 Adaptive Logic Modules (ALM) are available and, the FPGA hosts 1.708.800 registers and 1.518 DSP blocks. These digital signal processor (DSP) blocks are crucial for any algorithm using floating point calculations. These are hardened floating point add/mult blocks, which can add or multiply two floats in one clock cycle and the implementation needs no additional ALMs, like for the Intel® Stratix® V FPGA. Between the CPU and the FPGA a cache-coherent QPI bus is used to connect the two chips with a high bandwidth and low latency interconnect. This makes a tight collaboration of FPGA and CPU on same data in the main memory possible.

This is a real advantage of the Intel® Xeon®-FPGA versus the standard PCIe FPGA accelerator cards, because the PCIe FPGA card has to transfer the data from the main memory first to the local memory on-board, process the data with the FPGA and copy the results from the local memory back to the main memory. This reduces the filling of the FPGA pipelines dramatically[4], and makes only for the usage of iterative heavy algorithms or caching heavy algorithms sense.

GPUs and CPUs have an higher power consumption than FPGA accelerators[4], which makes FPGA acceleration in the metric of performance per Joule very interesting. Especially for data center with a large number of servers, where also the cost for energy and cooling makes a significant contribution to the total costs. For example the future LHCb computing farm has to be limited to 2 MW due to the cooling capacities, so the raw performance is important but also the performance per Joule.




Fig. 1. Intel® Xeon®-FPGA

The programming of the algorithms running on the FPGA can be written in Verilog and soon also in OpenCL.

2 Test Case: Sorting

The first test case for the new Intel® Xeon®-FPGA was the implementation of a sort algorithm for arrays of integer numbers on the FPGA. Sorting on an FPGA is interesting, because on a CPU the sorting of a n -element array scales with $n \times \log(n)$, but on a pipelined FPGA-design the time depends only on the clock frequency for a large number of arrays. The bottleneck is that the amount of resources on the FPGA scales with n^2 [5]. The block was tested with arrays filled with random numbers, the number of arrays was modified from 1 up to 32.000.

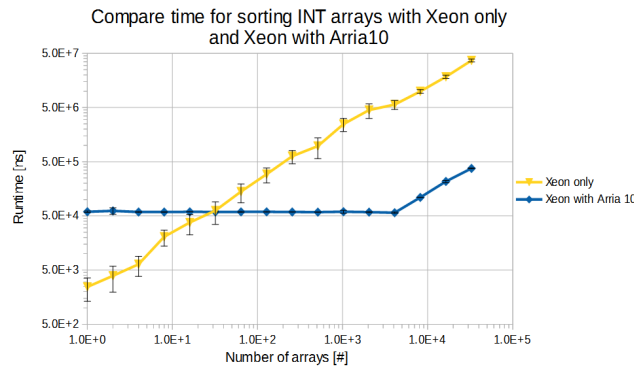


Fig. 2. Results of INT array sorting.

The achieved speed-up for sorting on the Intel® Xeon®-FPGA compared to the single threaded CPU, using "Quicksort", is for a large number of arrays (>4000) a factor 117 (see Fig. 2).

3 Cherenkov Angle Reconstruction

The RICH photon reconstruction is used for particle identification and crucial for the LHCb physics program. In the current version the RICH part of the trigger software takes too much time to be processed for every proton-proton collision (event). So it is a good candidate to be accelerated with the new Intel® Xeon®-FPGA. For each event hundreds of particle tracks are measured, each of them creates several Cherenkov photons. To find the corresponding Cherenkov photon cone to every particle track the same calculation has to be done for every photon found.

Charged particles travelling faster than the speed of light in a medium transmit Cherenkov irradiation. The photons are emitted in a cone of a certain angle around the particle momentum vector, this angle is called Cherenkov angle. The angle is correlated with the particle speed β and if the energy of the particle is

also known the particle can be identified. To calculate the Cherenkov angle a quartic equation has to be solved, a cube root, a rotation matrix, and several cross and scalar products have to be calculated [3].

The algorithm was implemented in a 259 clock cycle long pipeline, written in Verilog. The resources usage of the Arria® 10 FPGA after the Quartus synthesis optimization is shown in Table 1. The fast interface already uses 18% of the FPGA ALMs, and after the optimization the whole design takes 32% of all the ALMs. Only 15% of the DSP blocks are used to implement all floating point calculation blocks needed for one photon pipeline and in total 12% of all registers are used. The design runs at 200 MHz, which makes a calculation for a single photon within 5 ns possible if the pipeline is completely filled.

Table 1. FPGA resource usage for Cherenkov Angle reconstruction

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

The results are shown in Fig. 3. For a small number of photons the Xeon® CPU is faster due to the large latency of the photon pipeline on the FPGA. The break even is reached at roughly 200 photons. The time for the CPU version rises linearly and the time for the FPGA version stays constant until 8,000 photons, for more photons also the time for the FPGA version rises linearly. For larger number of photons the time ratio between CPU and FPGA is a factor 20 up to 35. On average, the photon pipeline processes a new photon only every second clock cycle. The bottleneck of the acceleration is the bandwidth between CPU and FPGA. In theory, using a higher bandwidth and using all FPGA resources for 5 photon pipelines, an acceleration factor of roughly 300 would be possible. This would be the limit for the used Arria® 10 FPGA. To reduce the difference between theory and measured acceleration, caching is tested right now, to increase the reuse of the data on the FPGA. This is for the Cherenkov angle algorithm possible, because many photon hits are combined with many particle tracks, and for all the combinations the same calculation is processed.

4 Summary

The LHCb experiment will be upgraded 2018-2019 to make a much more flexible 40 MHz detector readout possible. Afterwards no hardware trigger will be used any-more, only a complete software based trigger will be implemented to select the interesting proton-proton collision. The triggering will happen on a large Event Filter farm, which has to process and filter a input bandwidth of 40 TBits/s. This is a challenging task and several compute accelerator technologies are being considered to be used inside the new Event Filter farm.

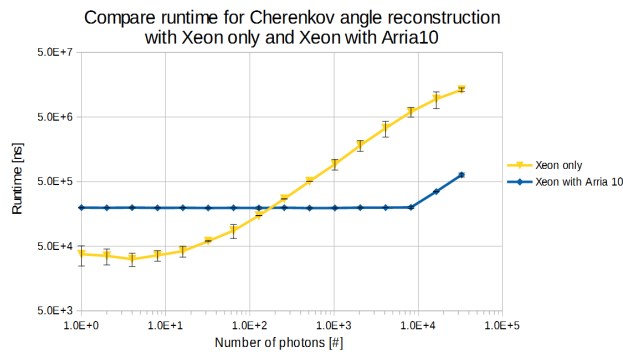


Fig. 3. Results of Cherenkov angle reconstruction.

Therefore, a study was especially here done to investigate the possible usage of the new Intel® Xeon®-FPGA to accelerate the Cherenkov angle reconstruction for the particle identification. A Verilog version was implemented on the Intel® Arria® 10 GX 1150 FPGA and an encouraging acceleration of 35x was measured for a large number of photons larger than 8000, compared to a single threaded Intel® Xeon® E5-2600 v4 CPU.

These results are very motivating and the High Energy Physics community could probably benefit similar from these new devices like the High Performance Computing community. It is expected that compared to other accelerators like GPUs the performance per Joule should be very interesting, which will be verified.

References

1. LHCb Collaboration, *Framework TDR for the LHCb Upgrade : Technical Design Report*. CERN, Geneva, 2012.
2. LHCb Collaboration, *LHCb Trigger and Online Upgrade Technical Design Report*. CERN, Geneva, 2014.
3. R. Forty and O. Schneider, *RICH pattern recognition*. CERN, Geneva, 1998.
4. S. Sridharan et al., *Accelerating particle identification for high-speed data-filtering using OpenCL on FPGAs and other architectures*. IEEE FPL16, Lausanne, Switzerland, 29 Aug - 2 Sept. 2016, DOI: 10.1109/FPL.2016.7577351
5. C. Färber et al., *Particle identification on a FPGA accelerated compute platform for the LHCb Upgrade*. Transactions on Nuclear Science, IEEE. 2017, DOI: 10.1109/TNS.2017.2715900