Feasibility study to use Intel Xeon Phi and Xeon-FPGA computing accelerators with the Omni-Path network in the Event Filter Farm for the LHCb Upgrade

Christian Färber, Karel Ha, Niko Neufeld, Rainer Schwemmer, Paolo Durante
{christian.faerber, karel.ha, niko.neufeld, rainer.schwemmer, paolo.durante}@cern.ch
Physics Department, CERN, Geneva, Switzerland

LHCb Upgrade

Current situation

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

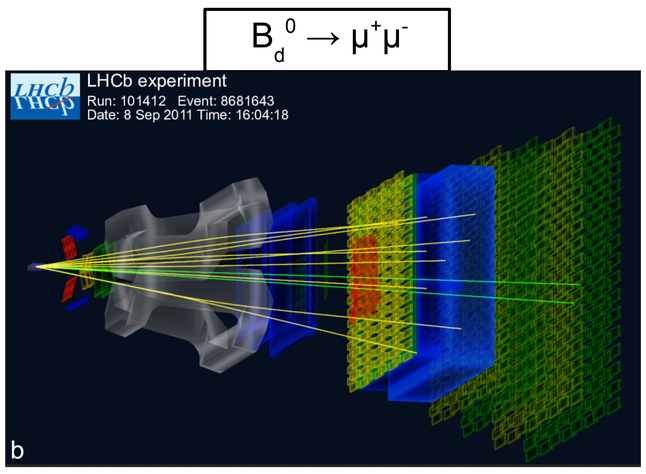

After the Upgrade

- 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.

High Throughput Computing Collaboration

Aims

- Feasibility of Intel technology for LHCb Upgrade in an Online context
- Data acquisition and event-building: move from firmware to software
- Accelerator-assisted processing for high-level trigger
- Primarily, integer computations for event building: lookup tables, offset calculations etc. (in contrast to float-intense operations for offline analysis and simulations)
- Cost vs. performance?

And...

- Code reviews
- Training and consultancy for developers
- Port LHCb code-base to *icc* compiler and provide maintenance

Omni-Path

- Intel's upcoming fabric technology suitable for DAQ systems
- 100 Gbit/s
- High-speed network across all Intel technology
  → "hand-tune" optimal ratio for HLT and Event Filter Farm (Xeon Phi vs. Xeon-FPGA vs. others)

Xeon Phi

The current version is called *Knights Corner* (KNC), main focus put on the upcoming *Knights Landing* (KNL).

Coprocessor especially suitable for high throughput computing

512-bit vector units for SIMD parallelization

60+ x86 cores on KNC, even more on KNL

16 GB of fast on-chip memory

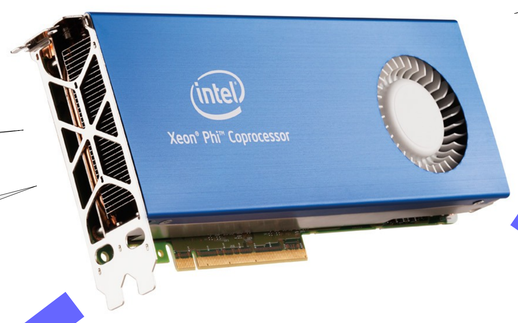

As stand-alone servers

Support for C/C++, Fortran via OpenCL, OpenMP, Intel Cilk Plus

Built-in interconnect (Intel Omni-Path) on KNL

Attractive for Trigger algorithms
- Track fitting
- Pattern recognition
- Particle identification

Interesting from the technical point of view
- EFF needs bandwidth of 40 Tbit/s.
- Many-to-many topology
  → Omni-Path support

Xeon-FPGA

First project: compare Muon Trigger performance on CPU and on FPGA

FPGAs are already used in the current system in the first trigger stage. This makes a combined system of CPU and FPGA very interesting for the future Event Filter Farm.

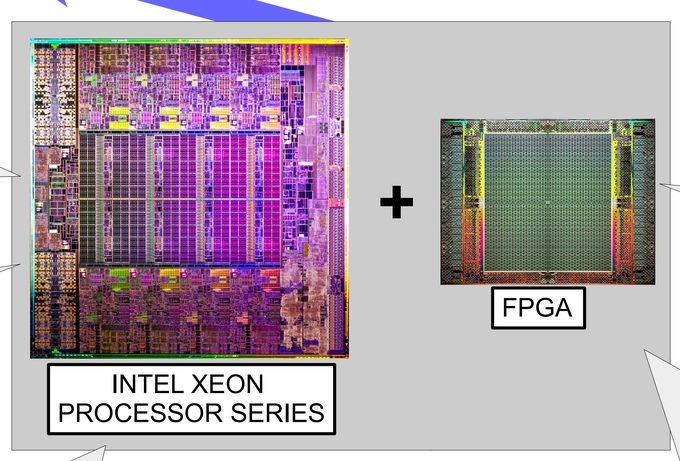

There is a computing accelerator implemented on the FPGA, which can be optimized for the algorithm.

Test system is a two-socket platform from Intel with a Xeon CPU and an FPGA.

CPU and FPGA are connected via the point-to-point interconnect QPI.

FPGA has cache-coherent memory access to the main memory of the server and it can collaborate with the CPU.