



The Intel way to future Online processing: the High Throughput Computing collaboration

Niko Neufeld, CERN/PH-Department
niko.neufeld@cern.ch

Oct. 5th 2015

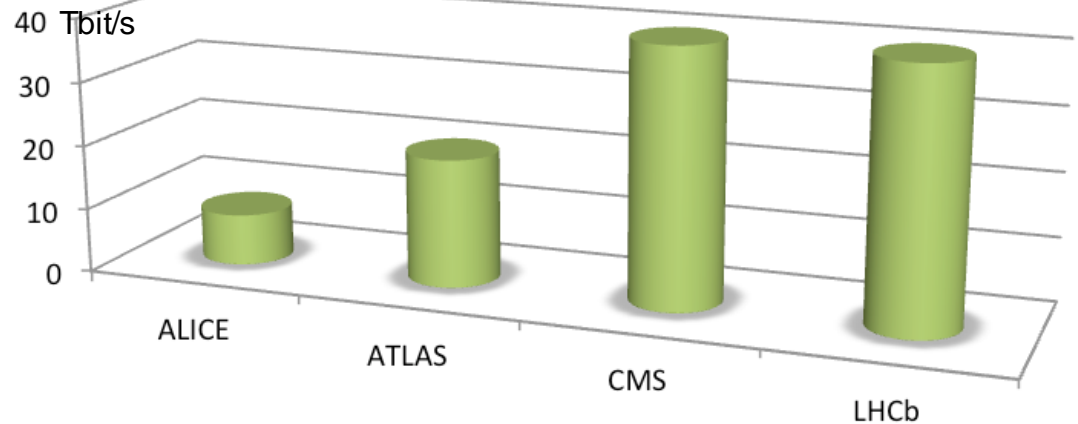


HTCC in a nutshell

- Apply upcoming Intel technologies in an Online context
- L1-trigger, data acquisition and event-building, accelerator-assisted processing for high-level trigger

Online computing challenges

- More and more data
- Limited money
- Limited manpower
- Limited power





First level selection

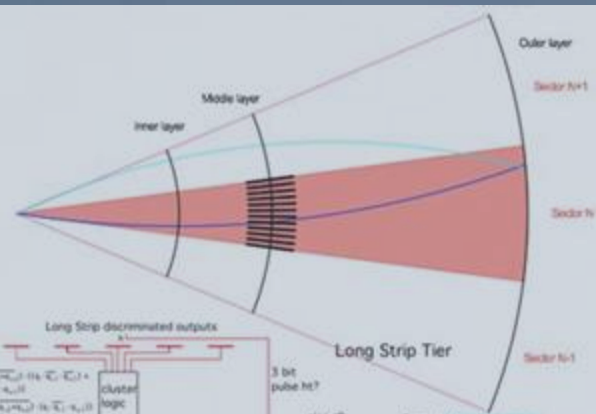
Challenges for Level-1 trigger

- Keep high efficiency in the face of many overlapping collisions → add tracking information
- Remain flexible, robust and easy to reproduce

Level 1 – a domain of custom hardware

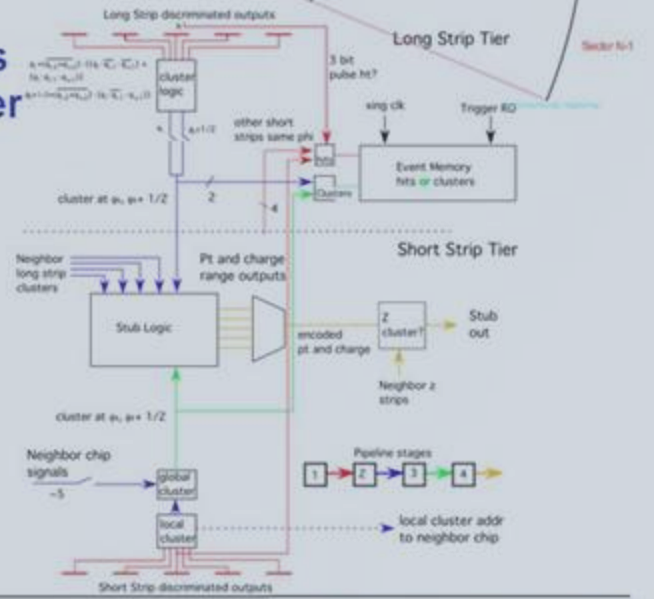
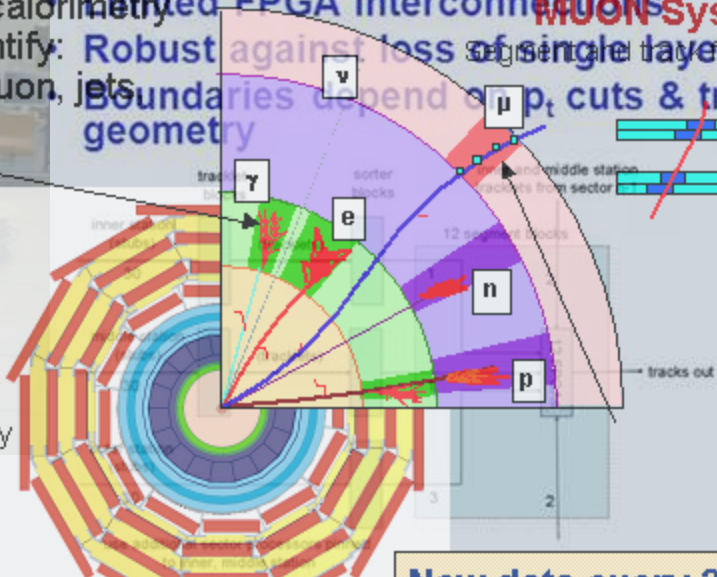


...designed to ... hits with ...
 ... through ...
 ... architecture ...
 Read out ... data
 Tracker organized into pni segments



Use prompt data (calorimetry and muons) to identify:
 High p_t electron, muon, jets
 missing E_T

limited FPGA interconnections
 Robust against loss of single layer hits
 Boundaries depend on p_t cuts & tracker geometry



CALORIMETERS

Cluster finding and energy deposition evaluation

New data every 25 ns
 Decision latency ~ μ s

Wesley Smith, University of Wisconsin, April 30, 2011

Texas A&M Seminar: Triggering CMS - 56



New (and old) L1 challenges

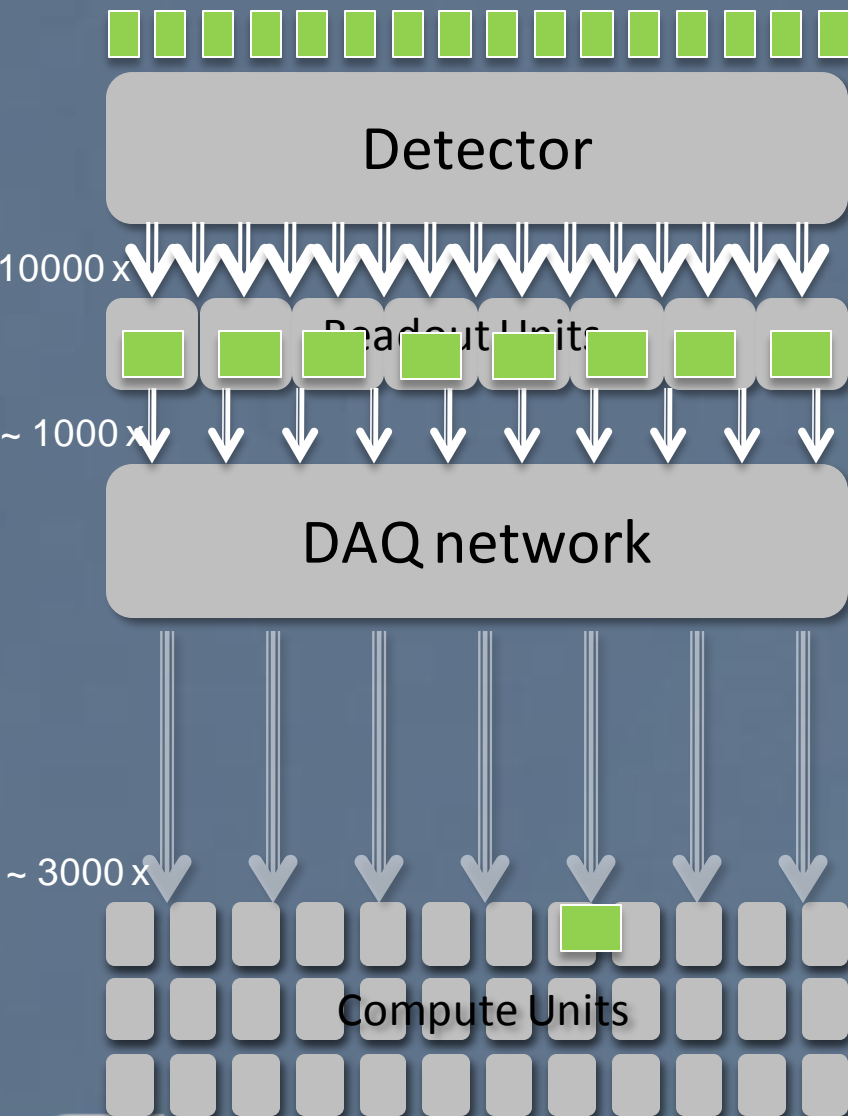
- A combination of (radiation hard) ASICs and FPGAs
- Sophisticated algorithms need more time, bigger FPGAs, more data
- Long-term maintenance issues with custom hardware and low-level firmware
 - Upgrades usually mean replacing all the hardware
- Exact reproducibility of results without the custom hardware challenging and/or computationally intensive

Intel / CERN HTC Collaboration

- Intel has announced plans for the first Xeon with coherent FPGA providing new capabilities
- We want to explore this to:
 - Move from firmware to software
 - Custom hardware → commodity
- Need real-time characteristics for L1:
 - algorithms must decide in $O(10)$ microseconds or force default decisions
- FPGA can provide hard-realtime, cache-coherent access to memory and can collaborate with CPU(s) (probably need to take CPU cores out of scheduler for this)
- We use existing FPGA versions of common algorithms (e.g. Hough transform, muon trigger)
- Later will compare with OpenCl or other higher-level synthesis (TBD)

Core DAQ

Working with full collision data



- Pieces of collision data spread out over 10000 links
- All pieces must be brought together into one of thousands compute units
- Compute units running complex filter algorithms (today dual-socket Xeon servers)

↓ custom radiation- hard link from the detector
3.2 Gbit/s

↓ DAQ (“event-building”) links – some LAN
(10/40/100 Gbit/s)


DAQ challenge

- Transport multiple Terabit/s reliably and cost-effectively
- Integrate the network closely and efficiently with compute resources (be they classical CPU or “many-core”)
- Multiple network technologies should seamlessly co-exist in the same integrated fabric (“the right link for the right task”)

Intel / CERN HTC Collaboration

- Explore Intel's new fabric OmniPath to build a DAQ network
 - Have ported LHCb event-builder exerciser to libfabric and run tests on Intel Truescale (the omnipath predecessor) on an HPC site
- Use OmniPath to integrate Xeon, XeonPhi and Xeon/FPGA concept in optimal proportions as compute units
 - E.g. move the accelerator “out of the box” → see KNL

High Level Trigger



“And this, in simple terms, is how we find the Higgs Boson”

High Level Trigger

- The challenges here share a lot with “offline” workloads, in particular reconstruction

HTCC and KNL

- How can KNL speed up typical big consumers
 - Pattern reco & Tracking
 - Particle Identification (done on LHCb money 😊)
- KNL as event-builder / event-sorter
 - Currently running on KNC, porting to KNL

KNL and Xeon/FPGA as accelerators

- Use existing “kernels” for important work-loads
 - Compare performance on KNL and Xeon/FPGA
- Xeon/FPGA should have an advantage over PCIe based accelerators (both FPGA and GPGPU) because of the cache-coherent, low-latency access to main-memory and CPU (no PCIe bottle-neck)
- We have demonstrated first offload in simulation, will test on real hardware soon

Who is HTCC

- Omar Awile
- Christian Färber
- Karel Ha (student) 100%
- Sebastien Valat
- Rainer Schwemmer
- Paolo Durante
- Olof Barring 10 - 25%
- Pawel Szostek
- Niko Neufeld
- Jon Machen (Intel Corp) 50%

Summary

- The LHC experiments need to reduce 100 TB/s to ~ 25 PB/ year
- Today this is achieved with massive use of custom ASICs and in-house built FPGA-boards and x86 computing power
- Finding new physics requires massive increase of processing power, much more flexible algorithms in software and much faster interconnects
- The CERN/Intel HTC Collaboration will explore Intel's Xeon/FPGA concept, Xeonphi and OmniPath technologies to build the LHC trigger of the future