



High Throughput Computing Collaboration

A CERN openlab / Intel collaboration

Niko Neufeld, CERN/PH-Department

niko.neufeld@cern.ch

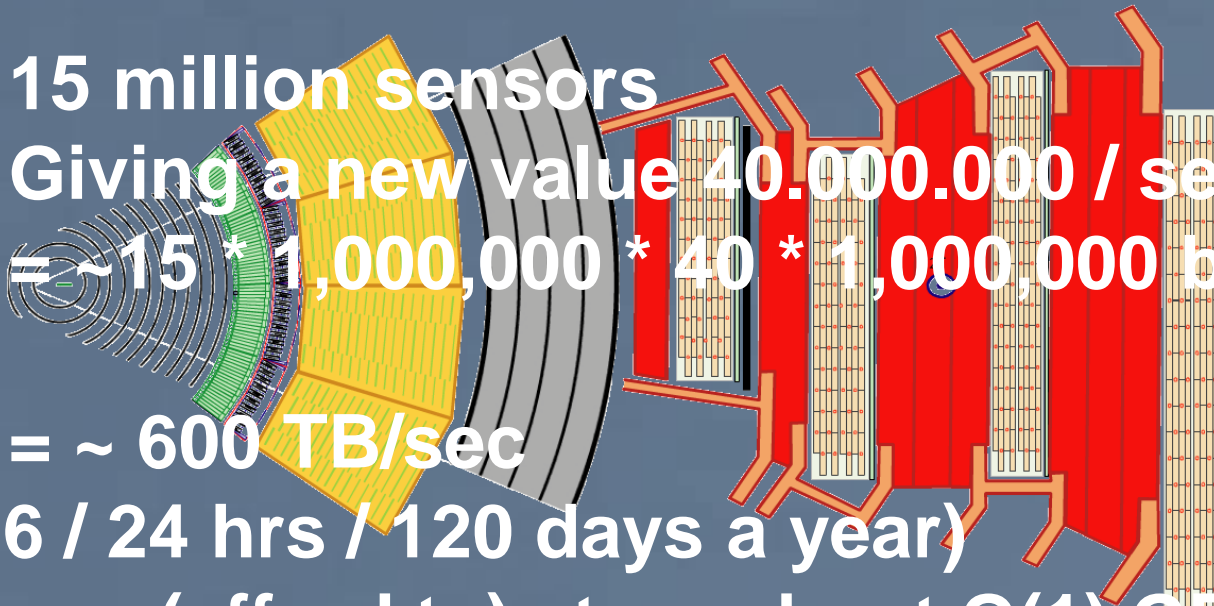


HTCC in a nutshell

- Apply upcoming Intel technologies in an Online / Trigger & DAQ context
- Application domains: L1-trigger, data acquisition and event-building, accelerator-assisted processing for high-level trigger

40 million collisions / second: the raw data challenge at the LHC

- 15 million sensors
- Giving a new value 40.000.000 / second
- $\approx 15 * 1,000,000 * 40 * 1,000,000$ bytes
- ≈ 600 TB/sec
(16 / 24 hrs / 120 days a year)
- can (afford to) store about O(1) GB/s



Defeating the odds


1. Thresholding and tight encoding
2. Real-time selection based on partial information
3. Final selection using full information of the collisions

Selection systems are called “Triggers” in high energy physics

Challenge #1

First Level Triggering

Selection based on partial information



- A combination of (radiation hard) ASICs and FPGAs process data of “simple” sub-systems with “few” $O(10000)$ channels in real-time

- Other channels need to buffer data on the detector

- this works only well for “simple” selection criteria

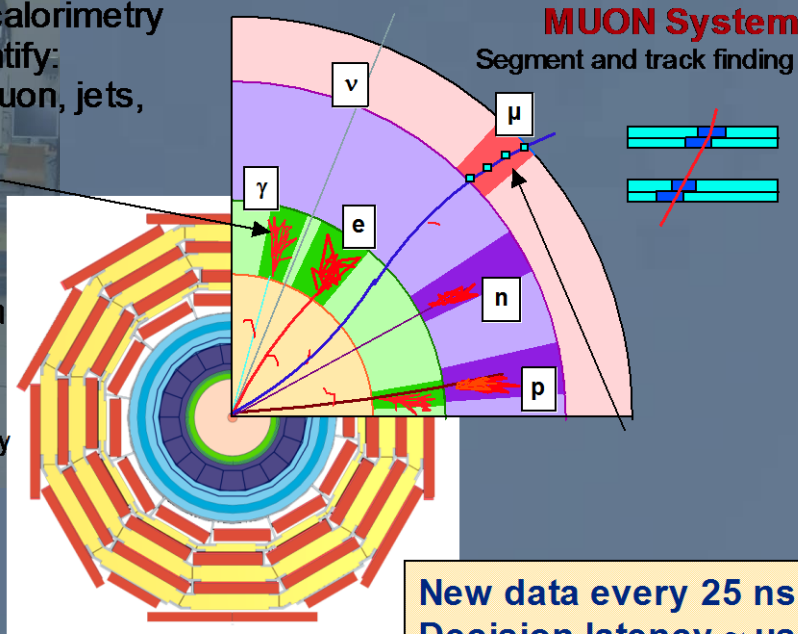
- long-term maintenance issues with custom hardware and low-level firmware

- crude algorithms miss a lot of interesting collisions

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

CALORIMETERS

Cluster finding and energy deposition evaluation



FPGA/Xeon Concept

- Intel has announced plans for the first Xeon with coherent FPGA concept providing new capabilities
- We want to explore this to:
 - Move from firmware to software
 - Custom hardware → commodity
- Rationale: HEP has a long tradition of using FPGAs for fast, online, processing
- Need real-time characteristics:
 - algorithms must decide in $O(10)$ microseconds or force default decisions
 - (even detectors without real-time constraints will profit)


HTCC and the Xeon/FPGA concept

- Port existing (Altera ☺) FPGA based LHCb Muon trigger to Xeon/FPGA
 - Currently uses 4 crates with > 400 Stratix II FPGAs → move to a small number of FPGA enhanced Xeon-servers
- Study ultra-fast track reconstruction techniques for 40 MHz tracking (“track-trigger”)
- Collaboration with Intel DCG IPAG -EU
 - Data Center Group, Innovation Pathfinding Architecture Group-EU

Challenge #2

Data Acquisition

Working with full collision data event-building



- Pieces of collision data spread out over 10000 links received by $O(100)$ readout-units
- All pieces must be brought together into *one* of thousands compute units → requires very fast, large switching network
- Compute units running complex filter algorithms

Future LHC DAQs in numbers

	Data-size / collision [kB]	Rate of collisions requiring full processing [kHz]	Required # of 100 Gbit/s links	Aggregated bandwidth	From
ALICE	20000	50	120	10 Tbit/s	2019
ATLAS	4000	500	300	20 Tbit/s	2022
CMS	4000	1000	500	40 Tbit/s	2022
LHCb	100	40000	500	40 Tbit/s	2019


HTCC and data acquisition

- Explore Intel's new OmniPath interconnect to build the next generation data acquisition systems
 - Build small demonstrator DAQ
- Use CPU-fabric integration to minimise transport overheads
- Use OmniPath to integrate Xeon, Xeon/Phi and Xeon/FPGA concept in optimal proportions as compute units
 - Work out flexible concept
- Study smooth integration with Ethernet (“the right link for the right task”)

Challenge #3

High Level Trigger


High Level Trigger



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

- Pack the knowledge of tens of thousands of physicists and decades of research into a huge sophisticated algorithm
- Several 100.000 lines of code
- Takes (only!) a few 10 - 100 milliseconds *per collision*

Pattern finding - tracks







Same in 2 dimensions

VELO RZ

VELO RZ

7.6.2009 4:18:49
Run 50416 Event 473

7.6.2009 4:18:49
Run 50416 Event 473



- Can be much more complicated: lots of tracks / rings, curved / spiral trajectories, spurious measurements and various other imperfections



HTCC and the High Level Trigger

- Complex algorithms
 - Hot spots difficult to identify → cannot be accelerated by optimising 2 -3 kernels alone
 - Classical algorithms very “sequential”, parallel versions need to be developed and their correctness (same physics!) needs to be demonstrated
- Lot of throughput necessary → high memory bandwidth, strong I/O
- There *is* a lot of potential for parallelism, but the SIMT-kind (GPGPU-like) is challenging for many of our problems
- HTCC will use next generation Xeon/Phi (KNL) and port critical online applications as demonstrators:
 - LHCb track reconstruction (“Hough Transformation & Kalman Filtering”)
 - Particle identification using RICH detectors

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, Xeon/Phi and OmniPath technologies for building future LHC TDAQ systems