



Olivier Callot

on behalf of the LHCb collaboration



Implementation and Performance of the LHCb trigger

- The challenge
- The solution: a three level system
 - Level-0 is synchronous, custom hardware
 - Level-1 is software, limited latency, on part of the data
 - HLT is software, same infrastructure, on the complete event
- Performance
- Future developments

The challenge

◆ LHC environment

- 40 MHz crossing rate
- Only about 3×10^7 crossing per second have interacting beams.

◆ LHCb environment

- Luminosity around $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$
 - Single interaction preferred, to match B decay to its production vertex
- Visible cross-section is about 70 mb
 - Producing at least 2 tracks in the LHCb acceptance
- About **12 millions visible events per second.**

- $b\bar{b}$ production cross-section around 500 μb .
 - ~100,000 $b\bar{b}$ pairs per second...

- But interesting decays have low branching fraction
 - $5 \cdot 10^{-6}$ for $B \rightarrow \pi\pi$
 - $20 \cdot 10^{-6}$ for $B \rightarrow J/\Psi(\rightarrow \mu\mu)K_s^0$
- **Interesting events** have a total rate of a **few Hz**.
 - Acceptance and selection keep only order of 1 %.
- We foresee to write on storage **about 200 Hz**.
 - Including side bands and reasonably wide cuts

◆ Requirements


- Reduction factor **around 20,000**
- Best possible efficiency on physics signal
 - E.g. lifetime cuts are applied offline to remove background
 - No need to trigger on events that will be rejected
 - No dead time
- Affordable, debug-able, robust, ...

The solution

◆ Three level system

- Level-0 : **hardware system** with fixed latency
 - Custom electronics
 - Pipe-line operation, synchronous, fixed latency : 4 μ s
 - Reduce the rate to 10^6 events per second
- Level-1 : **Software** analysis on **part of the data**
 - Reduced and packed data from only few detectors
 - Run on a PC farm, about 1800 CPU
 - Average latency 1 ms, maximum 58 ms.
 - Reduce the rate to 40,000 events per second
- High Level Trigger (HLT) : **Software** on the **complete event**
 - Tracking with almost final accuracy
 - Select interesting physics decays

◆ Coordination by the Timing and Fast Control system



Level-0

◆ Select high p_T particles

- B meson have a high mass. At least one decay product should have a few GeV p_T .

◆ Work for all types of particles

- Hadrons, using large local deposits in HCAL
- Electrons, photons, π^0 using large local deposits in ECAL
- Muons, using tracks in the Muon system

◆ Reject busy events

- Multiple interactions, re-interaction, ...
 - Easier to trigger on, more complex to process, not used for physics
 - Better to lower a bit the thresholds, and have more useful events.

Level-0 Calorimeter

◆ Local deposits




- Use 2 cells x 2 cells area
 - ~6000 in ECAL, 8x8 to 24x24 cm²
 - ~1500 in HCAL, 26x26 or 52x52 cm²

◆ Work with E_T

- Converted from ADC value, on 8 bits
- Need to access neighbours
 - Easy when on the same board
 - Dedicated backplane for most of the connections
- Only highest E_T interesting
 - Select locally

◆ Sign particle type with PRS-SPD

- Same geometry
 - Same address → point to point link
 - But card-to-card communication → LVDS links



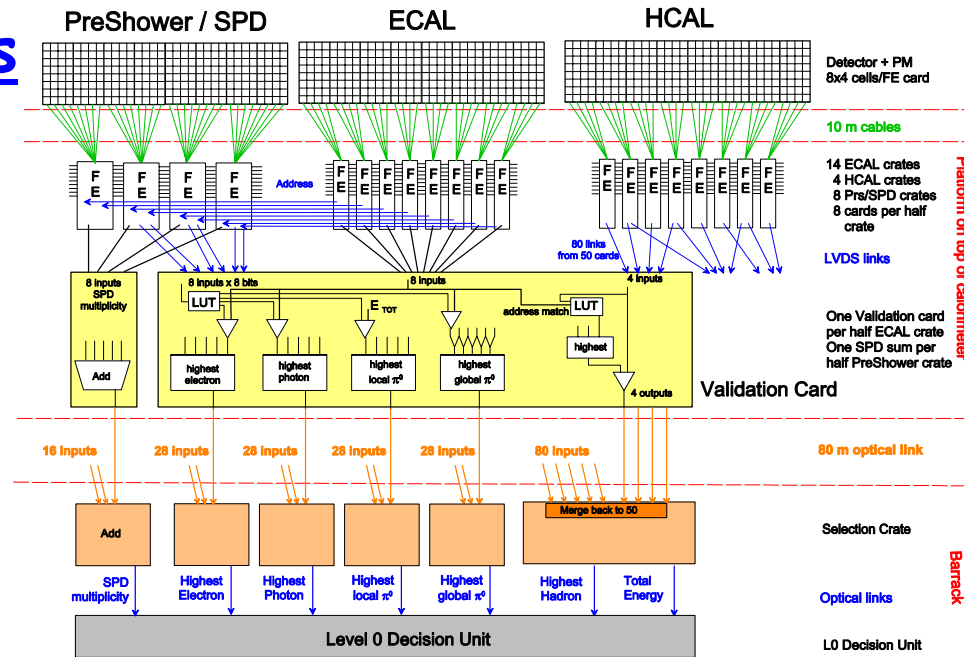
↓ : 8 bits LVDS multiplexed link
← : 8 bits on the backplane

◆ HCAL: add ECAL deposits

- Only when high enough
 - The geometry is not the same

◆ Performed in the cavern

- 10 Gy/year, SEU
 - anti-fuse PGA
 - triple voting



◆ Send candidates to Selection Crate

- Via ~200 optical links to the barrack
- Single crate to select the highest E_T of each type
- Send to L0 Decision Unit

◆ Produce also SPD multiplicity and Total HCAL E_T

- Reject busy or empty events

Level-0 Muon trigger

◆ Straight line search in M2-M5



- Chambers with projective geometry
- Infinite momentum \rightarrow same cell
 - FOI to allow non-infinite momentum !

◆ Extrapolation to M1

- Before shielding \rightarrow high occupancy
- But OK with prediction from M2-M3


◆ Momentum from M1-M2

- Assuming from Primary Vertex
- Using LUT
 - No computation



◆ Processed chamber data sent to the barracks

- Each chamber has 4 layers
 - OR of the chambers → 100% efficiency
- Logical pads created:
 - from physical pads, or
 - from crossing strips
- Optical link
 - 148 ribbons of 12 fibres
 - Binary data → BER not too important
 - But measured $< 10^{-15}$ from link, estimated 10^{-11} with SEU.




◆ 4 crates, each handling $\frac{1}{4}$ of the chamber

- 15 cards, handling each a sub-region
- Data exchanged via a dedicated backplane
 - High speed point-to-point links
- 4x2 muon candidates to L0 Decision Unit

Level-0 Pile-Up system


◆ 2 silicon planes backwards

- Measure Z of the Primary Vertex
 - From the ratio of the radii
- Histogram and search for second peak



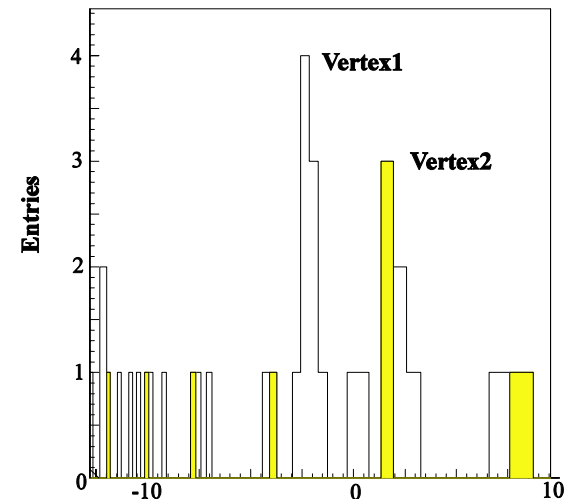
◆ Digital data sent to the barracks

- Same optical link technology



◆ Dedicated boards

- 4 identical, one event / 100 ns to each.
- Peak parameters sent to L0 Decision Unit.



Level-0 Decision Unit

◆ Synchronisation of the data

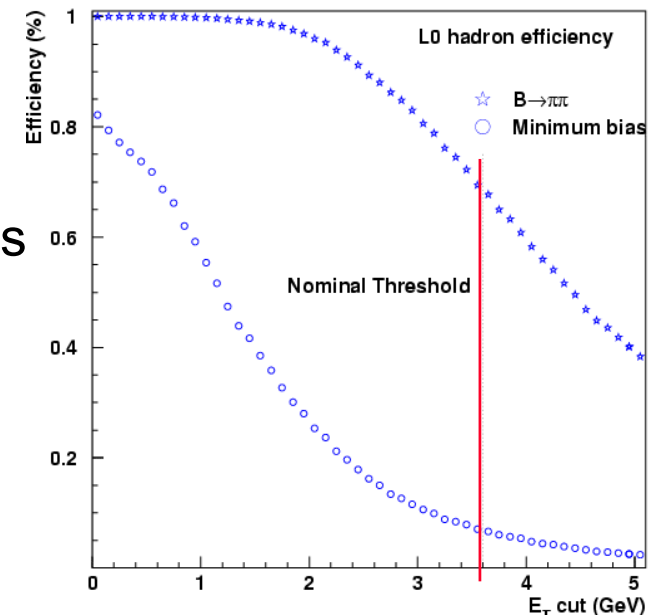

- All inputs at 40 MHz

◆ Apply thresholds

- Five types of particles for Calorimeter
- Eight muons select highest and pair

◆ Global variables


- Reject busy events or multiple interactions
- Not used for very clean channels
 - Di-muons are always accepted



Level-1

◆ Select long lived particles

- Tracking in the Velo
 - Select high Impact Parameter tracks
- But multiple scattering can fake IP
 - Estimate p_T from deflection up to TT
- Confirm that there are **high p_T particles** with high IP



◆ Software on a PC farm

- Send at 1 MHz events fragments to a CPU !
 - 126 sources, around 30 bytes/event each
- Trick: **Group 25 events in the same packet**
 - Fit in a single Ethernet frame, this reduce the overhead
 - Less interrupts at destination
- Big (but available today) IP router to connect sources and CPUs.


◆ Same infrastructure for HLT

- Same Ethernet network, same sub-farms and computing nodes
- Nodes are running HLT / reconstruction in the background
- When a Level-1 event arrives, it interrupts, and run at high priority
 - Minimal latency for Level-1, maximal use of the CPU.

◆ Separate HLT inputs

- Available from “TELL1”
 - Common readout board
- Allows to group links
- Events also grouped by packet of ~10.


◆ TFC system



Level-1 software

◆ High speed Velo tracking


- R sensors, 45° sectors, 512 strips/sector, 40 μm to 100 μm pitch



- **Straight line search** in R-Z view, forward and backward tracks
 - Around 58 (+ ~30 backward) tracks
- Vertexing, using the sector number as Phi measurement
 - $\sigma_z \sim 60 \mu\text{m}$, $\sigma_{x,y} \sim 25 \mu\text{m}$
- Select tracks with **high impact parameter, 0.15 to 3 mm**
 - about 8.5 per event
- **Full space tracking** for those tracks


◆ Momentum measurement using TT

- Only 0.15 T.m between Velo and TT
- Measures p_T with 20-40% accuracy
- Reject most low momentum tracks
 - They can fake high impact parameter.



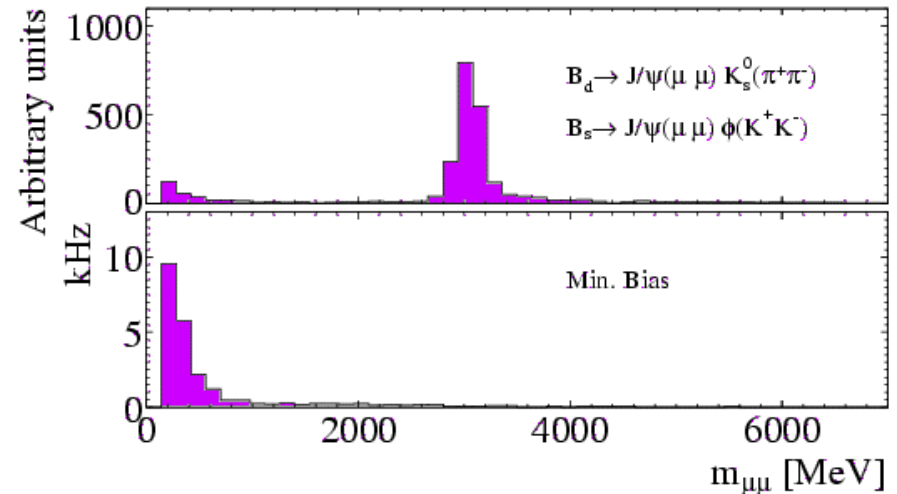
◆ Decision using two highest p_T tracks

- For those with high IP.




◆ In fact, add some “bonus” for identified particles

- Di-muons are easy to tag
 - Mass \rightarrow J/ Ψ peak
- Photons events also
 - 3 GeV E_T photons are rare.



◆ The whole processing takes around 8 ms NOW

- We know how to improve it by $\sim 20\%$
- Moore's law should give a factor 6
- 1 ms in 2007 seems feasible
 - We need 1000 CPU for Level-1.



HLT algorithms

◆ Tracking with (almost) full accuracy

- Extend Velo tracks across the magnet $\rightarrow \sigma(p)/p \sim 0.6 \%$

◆ Confirm Level-1 decision


- Select from the 3D Velo the tracks to extend, measure accurately the momentum, confirm the decision
 - Gain a factor 2 without sensible loss
 - In about 25 ms on today's CPU

◆ Full reconstruction

- All tracks, around 50 ms now

◆ Physics selection

- Not yet coded or timed



Timing and Fast Control

◆ Distribute decisions using RD45's TTC system

- Level-0 decision at 40 MHz
- Level-1 decision at 1 MHz
- Events' destination at 40 kHz (L1) / 4 kHz (HLT)
 - Simple round-robin assignment, not a central event manager

◆ Readout Supervisor as 'chef d'orchestre'

- **Emulate the Level-0 occupancy** front-end to avoid buffer overflow
 - 0.5% dead time at 1 MHz, guarantee to never exceed 1.11 MHz
- Receive **throttle signals** from the readout boards in case they can't cope with the rate
 - In this case, disable temporarily the trigger, until the system is in a less busy state
- Critical component
 - Prototype exists and works well.

Performance

◆ Measured on selected events

- This means used for physics analysis
- Minimum bias retention adjusted
 - 1 MHz at Level-0
 - Sharing also optimised
 - 40 kHz at Level-1


◆ ~30% for hadronic channels

- 60-70 % with di-muons
- Also around 30% for $e/\gamma/\pi^0$ channels

L0 efficiency

L1 efficiency

Total L0×L1 efficiency



Future developments

◆ TDR has been submitted last month


- Approval expected in November

◆ Level-0 is mature

- Hardware implementation has started
 - Fully synchronous → independent of history
 - Many testing features included

◆ Level-1 and HLT hardware being prototyped

- Commercial components (switches, router, PCs)
 - Routing available today, even if too expensive
 - Relies on Moore's law to get powerful enough PC in 2007
- Connection and **control of the farm is the issue.**



◆ Level-1 and HLT software being developed

- Prototypes show that **we can stay within the time budget**
 - Level-1 code is in the correct speed range, need ~1000 CPU
 - HLT code is expected to take around 500 CPU
 - L1 confirmation in about 25 ms today → 4 ms in 2007 = 160 CPU
 - Full tracking in 60 ms today → 10 ms in 2007 = 200 CPU
 - Event selection should be < 40 ms today → 140 CPU
- **Event selection** is still at the prototype stage
 - We have code for a few channels, few ms today.
 - But we have many channels to select, and want to be quite open for the selection (if possible)

LHCb will be able to efficiently select the events useful for CP violation studies in B decays from the beginning of LHC operation