First-level trigger systems at LHC

Nick Ellis

EP Division, CERN, Geneva

Outline

- Requirements from physics and other perspectives
- General discussion of first-level trigger implementations
  - Techniques and technologies
- Overview of first-level triggers for the LHC experiments
  - ATLAS, CMS, LHCb, Alice
- Calorimeter triggers
  - Illustrated with example of ATLAS e/γ trigger
- Muon triggers
  - Illustrated with example of CMS drift-tube based trigger
- Pile-up veto in LHCb
- Central/global triggers
  - Illustrated with example of CMS global trigger
- Conclusions

General trigger requirements

- The role of the **trigger** is to make the online selection of particle collisions potentially containing interesting physics
- Need **high efficiency** for selecting processes of interest for physics analysis
  - Efficiency should be precisely known
  - Selection should not have biases that affect physics results
- Need **large reduction of rate** from unwanted high-rate processes (capabilities of DAQ and also offline computers)
  - Instrumental background
  - High-rate physics processes that are not relevant for analysis
- System must be **affordable**
  - Limits complexity of algorithms that can be used
- Not easy to achieve all the above simultaneously!

Why do we need *multi-level* triggers?

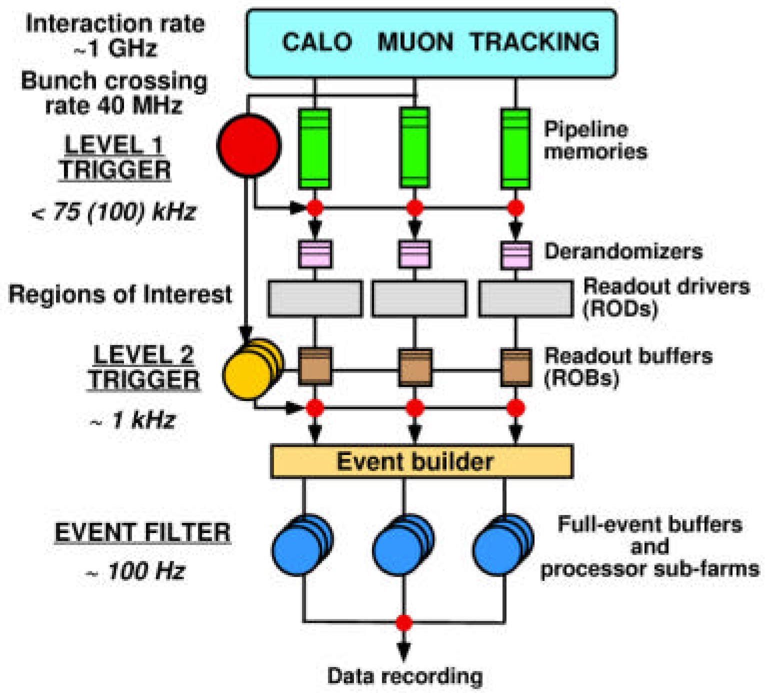

Example: ATLAS

- Multi-level triggers provide:
  - Rapid rejection of high-rate backgrounds without incurring (much) dead-time
    - Fast first-level trigger (custom electronics)
      - Needs high efficiency, but rejection power can be *comparatively* modest
      - Short latency is essential since information from all (up to $O(10^8)$) detector channels needs to be buffered (often on detector) pending result
  - High overall rejection power to reduce output to mass storage to affordable rate
    - Progressive reduction in rate after each stage of selection allows use of more and more complex algorithms at affordable cost
    - Final stages of selection, running on computer farms, can use comparatively very complex (and hence slow) algorithms to achieve the required overall rejection power

Requirements from physics perspective

- Typically, trigger systems select events according to a "trigger menu", i.e. a list of selection criteria
  - An ***event*** is selected by the trigger if one or more of the criteria are met
    - I use the term "event" to mean the record of the activity in a given bunch crossing — typically an event contains many proton–proton interactions
      - First-level trigger has to identify uniquely the BC of interest
  - Different criteria may correspond to different signatures for the same physics process
    - Redundant selections lead to high selection efficiency and allow the efficiency of the trigger to be measured from the data
  - Different criteria may reflect the wish to concurrently select events for a wide range of physics studies
    - HEP "experiments" — especially those with large general-purpose "detectors" (detector systems) — are really experimental facilities
- Remember that events rejected by the trigger are lost forever!
  - In contrast to offline processing and physics analysis, there is no possibility of a second chance!

LHC physics

(see talk of P. Sphicas)

- Discovery physics is the main emphasis for ATLAS and CMS
  - Huge range of predicted new physics processes with diverse signatures
    - Very low signal rates expected in some cases
  - But should also try to be sensitive to new physics that has not been predicted!
- Huge rate of Standard Model physics backgrounds
  - Rate of proton–proton collisions up to 10^9 Hz
  - Much lower rates predicted for instrumental backgrounds such as beam–gas interactions

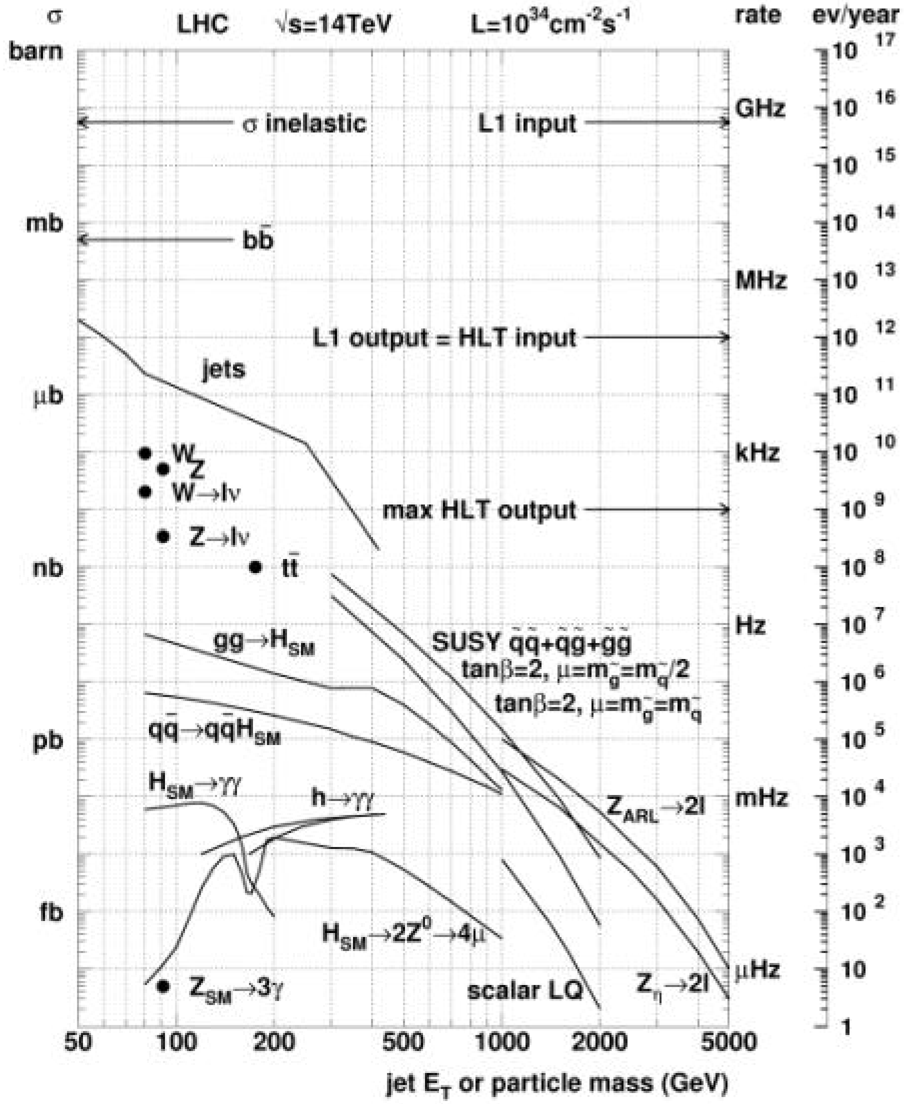

ATLAS and CMS

- The trigger will have to retain as many as possible of the events of interest for the diverse physics programmes of these experiments, including:
  - – Higgs searches (Standard Model and beyond)
    - E.g. $H \rightarrow ZZ \rightarrow$ leptons (e or μ), $H \rightarrow \gamma\gamma$; also $H \rightarrow \tau\tau$, $H \rightarrow bb$
  - – SUSY searches
    - E.g. producing jets and missing E_T
  - – Searches for other new physics
    - Using inclusive triggers that one hopes will be sensitive to unpredicted new physics
  - – Studies of Standard Model processes which are of interest in their own right, and must be understood as backgrounds to new physics
    - W and Z bosons, top and beauty quark production

ATLAS and CMS (continued)

- In contrast to the particles produced in typical pp collisions (typical hadron $p_T \sim 1$ GeV), the products of new physics are expected to have large transverse momentum, p_T
  - E.g. if they were produced in the decay of new heavy particles such as the Higgs boson; e.g. $m \sim 100$ GeV $\Rightarrow p_T \sim 50$ GeV
- Typical examples of first-level trigger thresholds for LHC design luminosity are:
  - Single muon $p_T > 20$ GeV (rate ~ 10 kHz)
    - Pair of muons each with $p_T > 6$ GeV (rate ~ 1 kHz)
  - Single e/γ $p_T > 30$ GeV (rate ~ 20 kHz)
    - Pair of e/γ each with $p_T > 20$ GeV (rate ~ 5 kHz)
  - Single jet $p_T > 300$ GeV (rate ~ 200 Hz)
    - Jet $p_T > 100$ GeV and missing-$p_T > 100$ GeV (rate ~ 500 Hz)
    - Four or more jets $p_T > 100$ GeV (rate ~ 200 Hz)

Effect of p_T cut in minimum-bias events

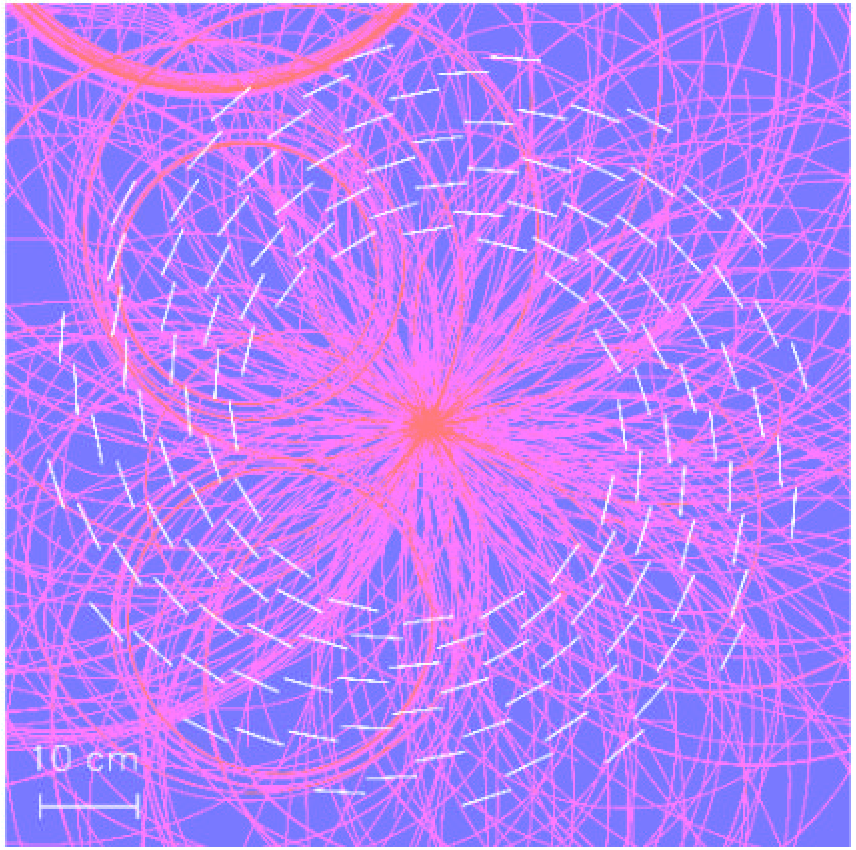

All tracks

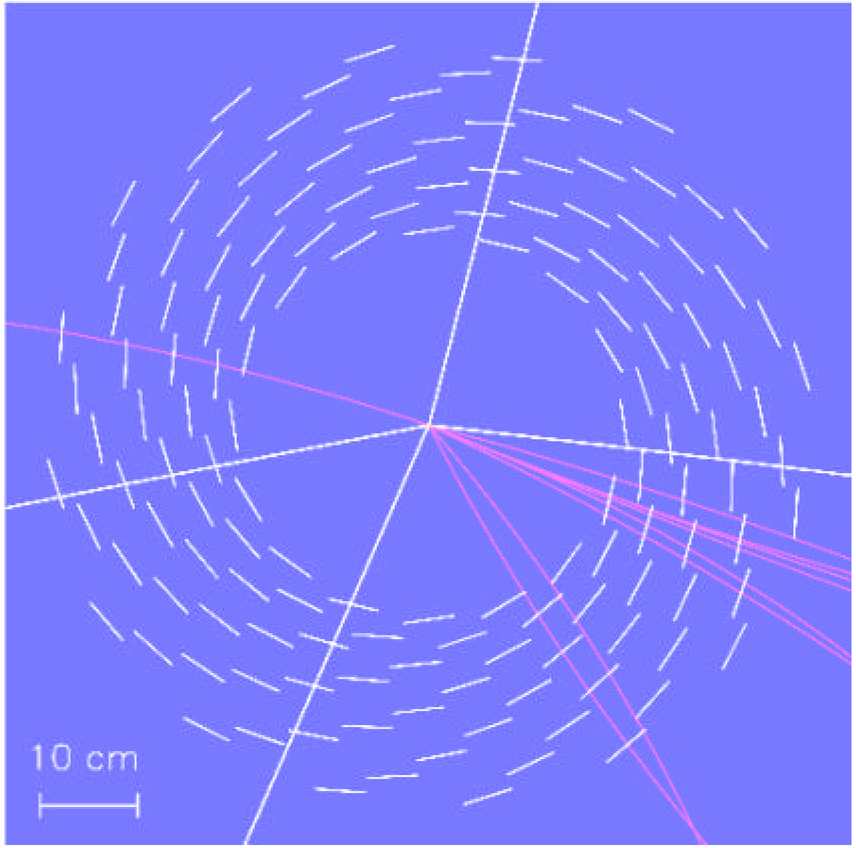

$p_T > 2$ GeV

Simulated $H \rightarrow 4\mu$ event + 17 minimum-bias events

LHCb

- The LHCb experiment, which is dedicated to studying B-physics, faces similar challenges to ATLAS and CMS
  - It will operate at a relatively low luminosity (~2×10^{32} $cm^{-2}s^{-1}$), giving an overall pp interaction rate of ~20 MHz
    - Chosen to maximise the rate of single-interaction bunch crossings
  - However, to be sensitive to the B-hadron decays of interest, the trigger must work with comparatively very low p_T thresholds
    - The first-level ("level-0") trigger will search for muons, electrons/photons and hadrons with p_T > 1 GeV, 2.5 GeV and 3.4 GeV respectively
      - Level-0 output rate up to ~1 MHz
  - Higher-level triggers must search for displaced vertices and specific B decay modes that are of interest for the physics analysis
    - Aim to record event rate of only ~200 Hz

ALICE

- The heavy-ion experiment ALICE is very demanding from the DAQ point of view, but the trigger is simpler than for the other experiments
  - The total interaction rate will be much smaller than in the pp experiments
    - $L \sim 10^{27}\ \mathrm{cm}^{-2}\mathrm{s}^{-1} \Rightarrow R \sim 8000$ Hz for Pb–Pb collisions (higher rates for lighter ions and protons)
  - The trigger will select “minimum-bias” and “central” events (rates scaled down to total ~40 Hz), and events with dileptons (~1 kHz with only part of the detector read out)
  - However, the event size will be huge due to the high multiplicity in Pb–Pb collisions at LHC energy
    - Up to *O*(10,000) charged particles in the central region
    - Event size up to ~ 40 MByte when full detector is read out

What do μ, e, γ, jets, etc "look like"?

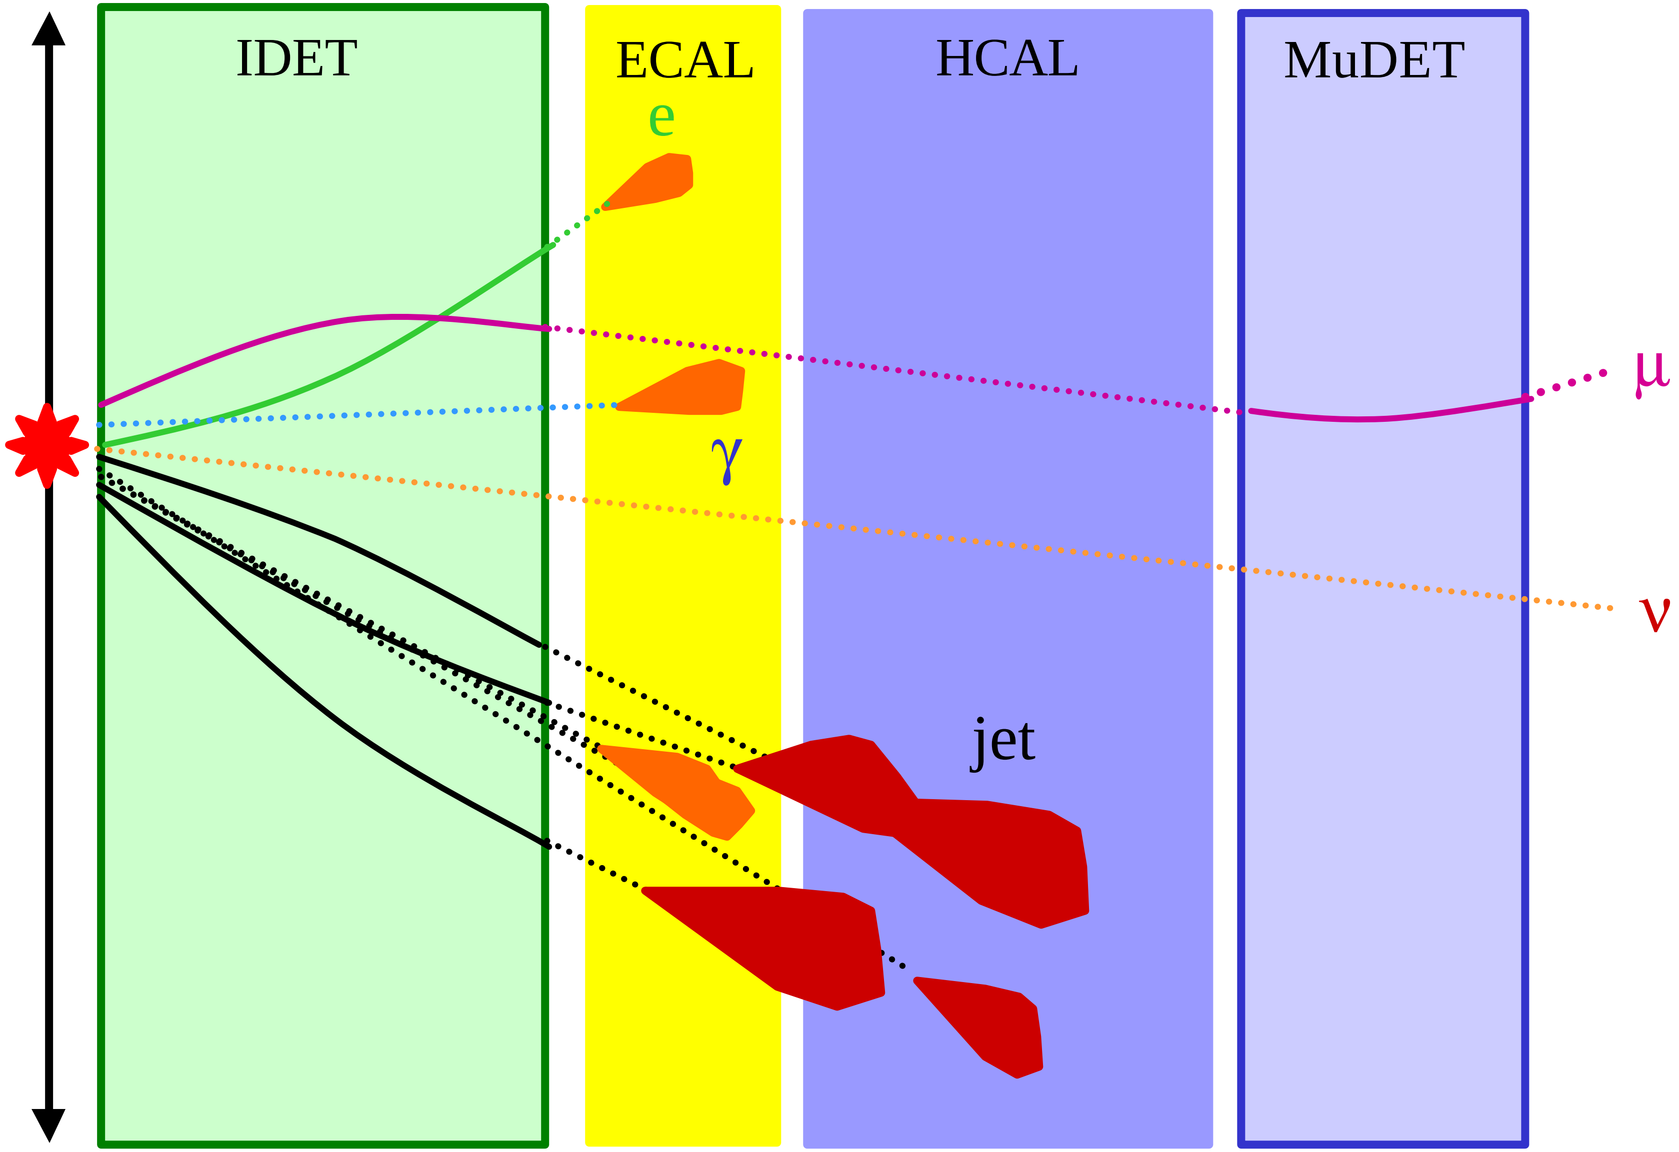

FIRST-LEVEL TRIGGER OVERVIEW

Muon detector signals

Calorimeter signals

New data every 25 ns

BUSY signals

Search for high-p_T:

- muons
- electrons/photons
- taus/hadrons
- jets

Calculate:

- ΣE_T
- missing E_T

Form trigger decision for each BC based on combinations of above

First-level Trigger

Introduce deadtime to avoid data loss or buffer overflow in front-end electronics

Decision every 25 ns
Latency ~few μs

Yes/No

Distribute first-level trigger decision to front-end electronics

- See talk of B.G. Taylor

Size of detectors and the speed of light

Trigger finds high-p_T muon here $\Rightarrow$ select event

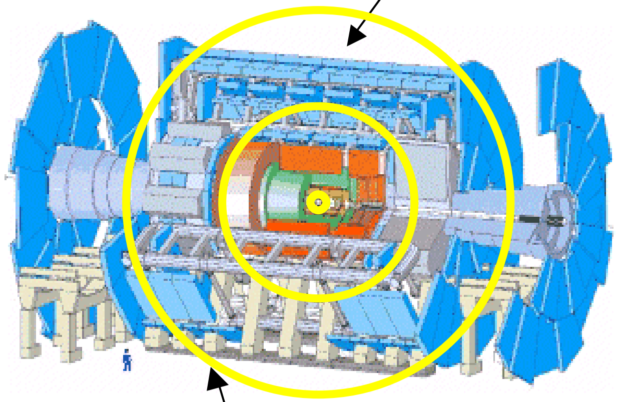

Need to read out also here

ATLAS, the biggest of the LHC detectors, is 22 m in diameter and 46 m in length

speed of light in air 0.3 m/ns

$22 \text{ m} \times 3.3 \text{ ns/m} = 73 \text{ ns}$
c.f. 25 ns BC period

The other LHC detectors are smaller, but similar considerations apply

It is **impossible** to form and distribute a trigger decision within 25 ns
(in practice, latency is at least ~ 2 μs)

Pipelined first-level triggers

- First-level trigger has to deliver a new decision every BC, but the trigger latency is much longer than the BC period
  - First-level trigger must concurrently process many events
  - This can be achieved by “pipelining” the processing in custom trigger processors built using modern digital electronics
    - Break processing down into a series of steps, each of which can be performed within a single BC period
    - Many operations can be performed in parallel by having separate processing logic for each one
  - Note that the latency of the trigger is fixed
    - Determined by the number of steps in the calculation plus the time taken to move signals and data to and from the components of the trigger system
      - Signals have to pass from the detector to the trigger electronics and back, with a round trip distance of about 200 m (1 μs delay)

Pipelined first-level trigger (illustration)

Note that logic must be duplicated for all ~3500 positions in calorimeter!

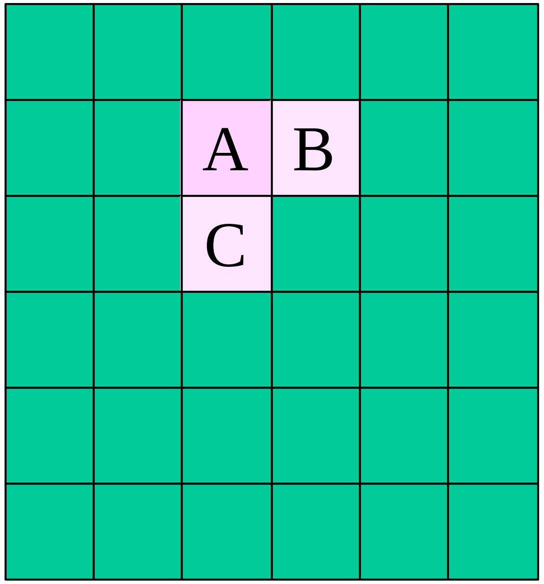

EM Calorimeter
(~3500 trigger towers)

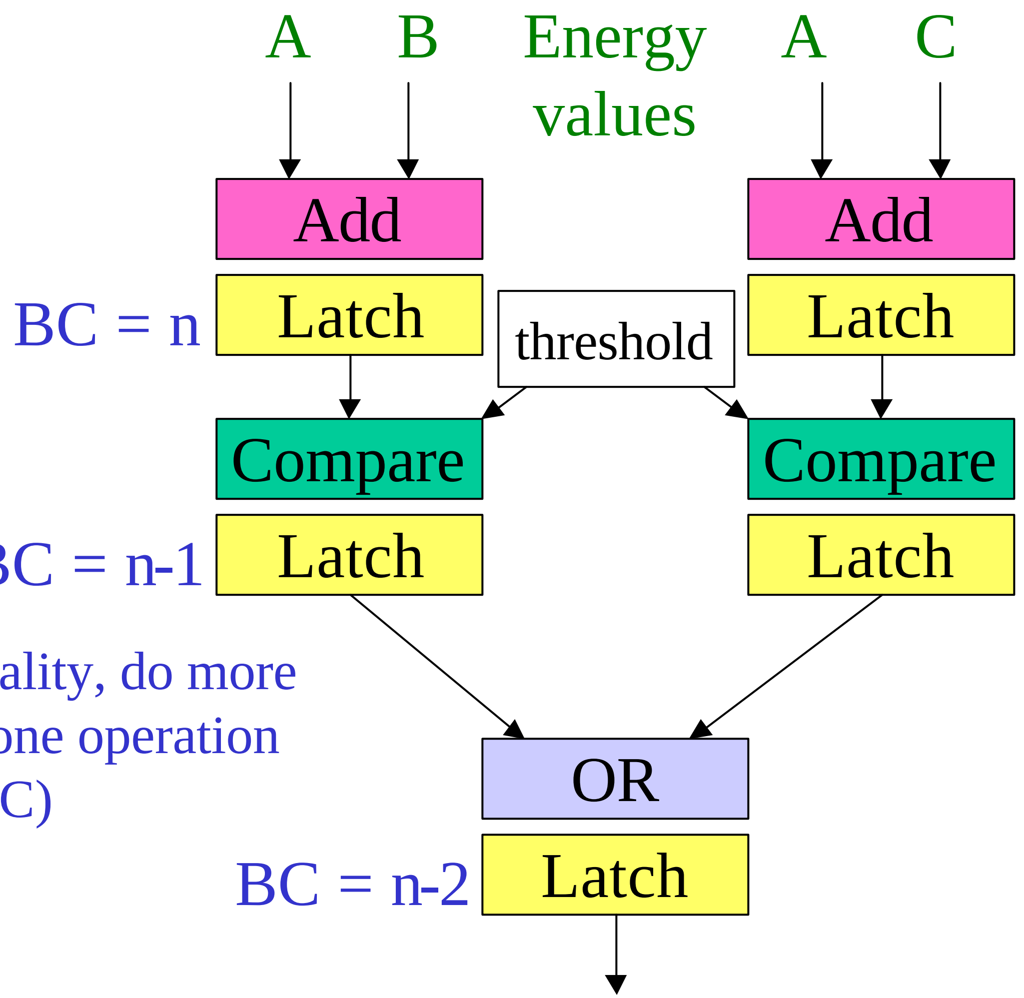

(In reality, do more than one operation per BC)

Data-processing technologies

- FPGAs (and other programmable devices) now play a very important role
  - Large gate count and many I/O pins available; operate at 40 MHz and above; performance sufficient for implementing many trigger algorithms
    - Offer huge flexibility
    - Possibility to modify algorithms as well as parameters of algorithms once experiments start running
- ASICs used for some applications
  - More cost effective in some cases (e.g. large number of devices)
  - Offer higher speed performance than FPGAs
  - Can have better radiation tolerance and lower power consumption for on-detector applications

Data-movement technologies

- High-speed serial links (electrical and optical)
  - Comparatively inexpensive and low-power LVDS links for electrical transmission at ~400 Mbit/s over distances up to ~10 m
  - Products such as HP G-link and Vitesse chipsets for Gbit/s transmission; using optical transmission for longer distances
- Very high density custom backplanes
  - High pin counts (up to ~800 per 9U board)
  - Data rates per (point-to-point) connection ~160 Mbit/s
    - Multiplex data beyond 40 Mbit/s to reduce connectivity problem to a level that can be managed
- Use large (9U) boards
  - Easier to handle interconnections on board than between boards

LVL1 data flow

Many input data

Energies in calorimeter towers
(e.g. ~7000 trigger towers in ATLAS)

Pattern of hits in muon detectors
(e.g. $O(10^6)$ channels in ATLAS)

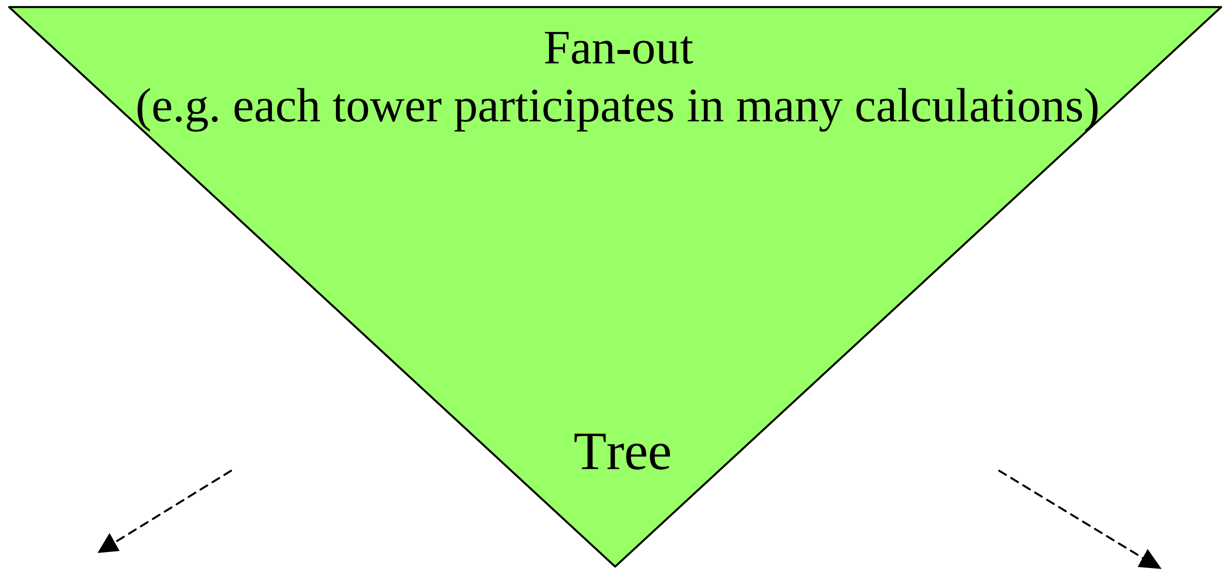

(Data for monitoring)

1-bit output
(YES or NO)

(Information to guide
next selection level)

Overview of ATLAS first-level trigger

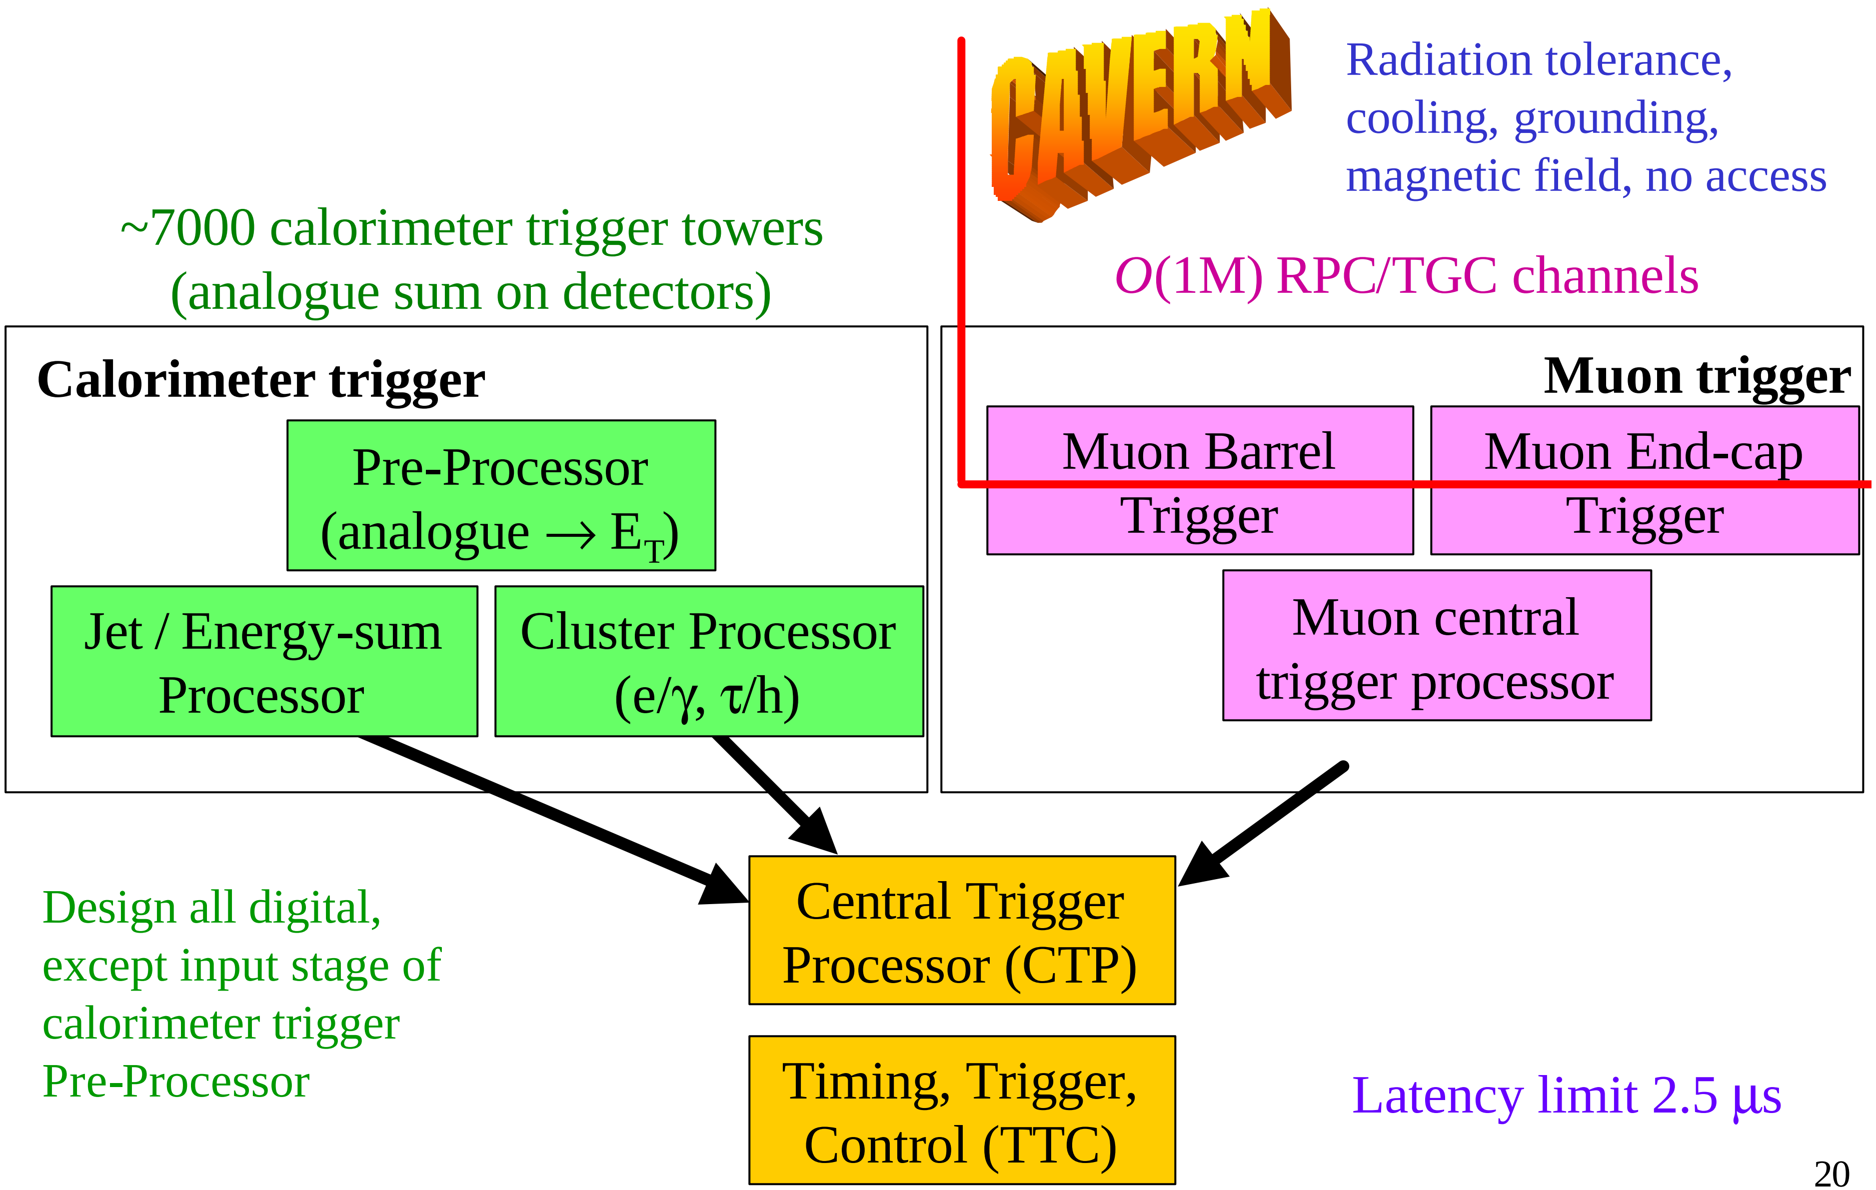

Overview of CMS first-level trigger

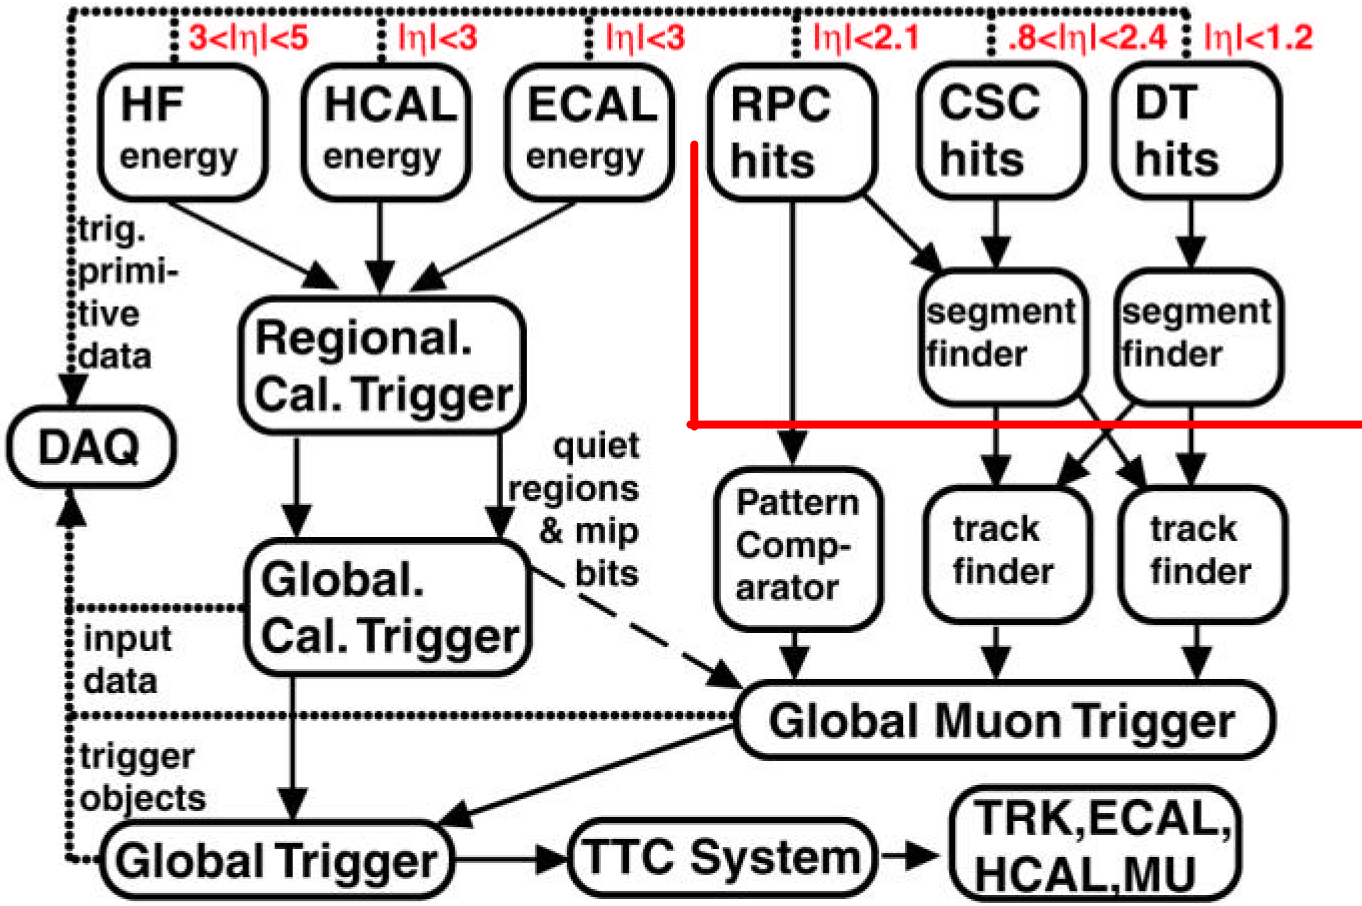

Latency limit 3.2 μs

Overview of LHCb first-level trigger

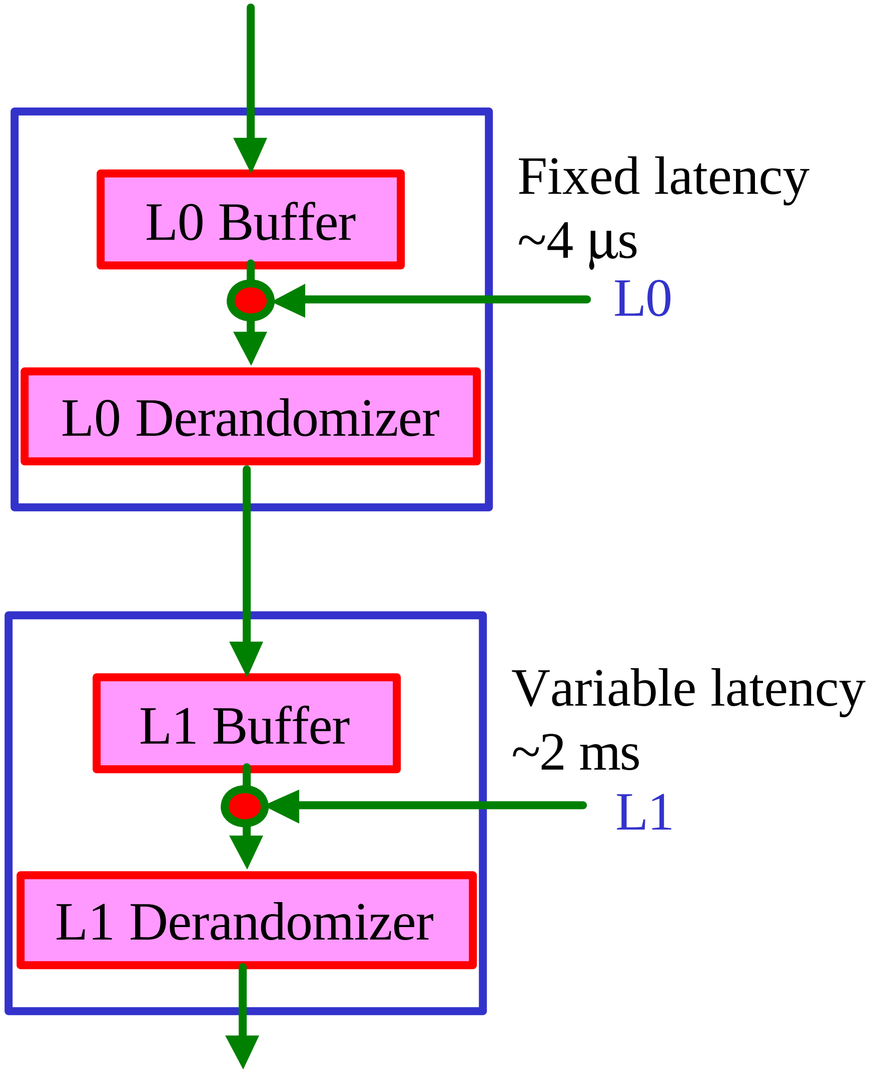

- Two levels of buffering on the detector (c.f. one for ATLAS and CMS)
- L0 (≡first-level) trigger (electronics)
  - Calorimeter
    - Electrons/photons, hadrons
  - Muon detectors
    - Muons
  - Pile-up veto
    - Reject events with more than one pp interaction vertex
- L1 trigger (software)
  - Vertex detector
    - Secondary vertices

Overview of ALICE first-level trigger

- Logic associated with subdetectors generates trigger inputs
  - 24 L0 inputs (latency 900 ns; 2 μs deadtime after each trigger)
    - Some detectors need prompt trigger signal
      - Track-and-hold rather than pipelined readout
      - All trigger electronics on detector
  - 20 L1 inputs (latency 6.2 μs)
  - 6 L2 inputs (latency 88 μs ~ TPC drift time)
- Provision for control of up to 24 independent subdetectors
  - Grouped into 6 detector clusters that are read out together
    - In contrast to ATLAS/CMS/LHCb, don't always read all subdetectors
- Define up to 50 trigger classes, specifying for each one
  - L0-L1-L2 patterns, prescale factor and detector cluster for readout
- Use of slow detectors requires past–future protection logic
  - Different limits for peripheral and semi-central interactions
    - Note very different interaction rates in Pb–Pb, Ar–Ar and p–p cases

CALORIMETER TRIGGERS

- Illustrate with example of ATLAS e/γ trigger
  - Will also discuss briefly the different trigger digitisation schemes in ATLAS and CMS
- See related talks in parallel sessions:
  - ATLAS
    - G. Mahout: Prototype cluster-processor module for the ATLAS level-1 calorimeter trigger
  - CMS
    - W.H. Smith: Tests of CMS regional calorimeter trigger prototypes
    - P. Busson: Overview of the new CMS electromagnetic calorimeter electronics

ATLAS first-level calorimeter trigger

- Analogue electronics on detector sums signals to form trigger towers
- Signals received and digitised
  - Digital data processed to measure E_T per tower for each BC
    - E_T matrix for ECAL and HCAL
- Tower data transmitted to CP (4 crates) and JEP (2 crates)
  - Fan out values needed in more than one crate
    - Motivation for very compact design of processor
- Within CP & JEP crates, values need to be fanned out between electronic modules, and between processing elements on the modules
- Connectivity and data-movement issues drive the design

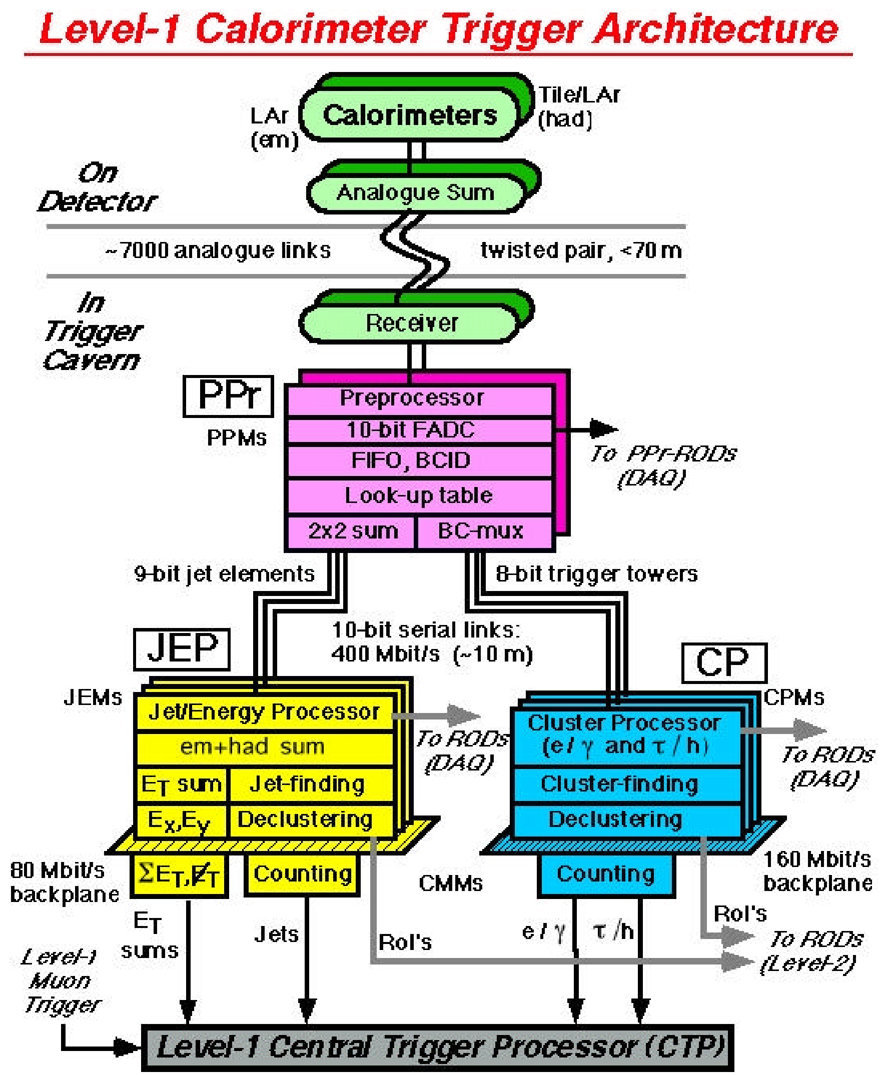

Digitisation options (ATLAS c.f. CMS)

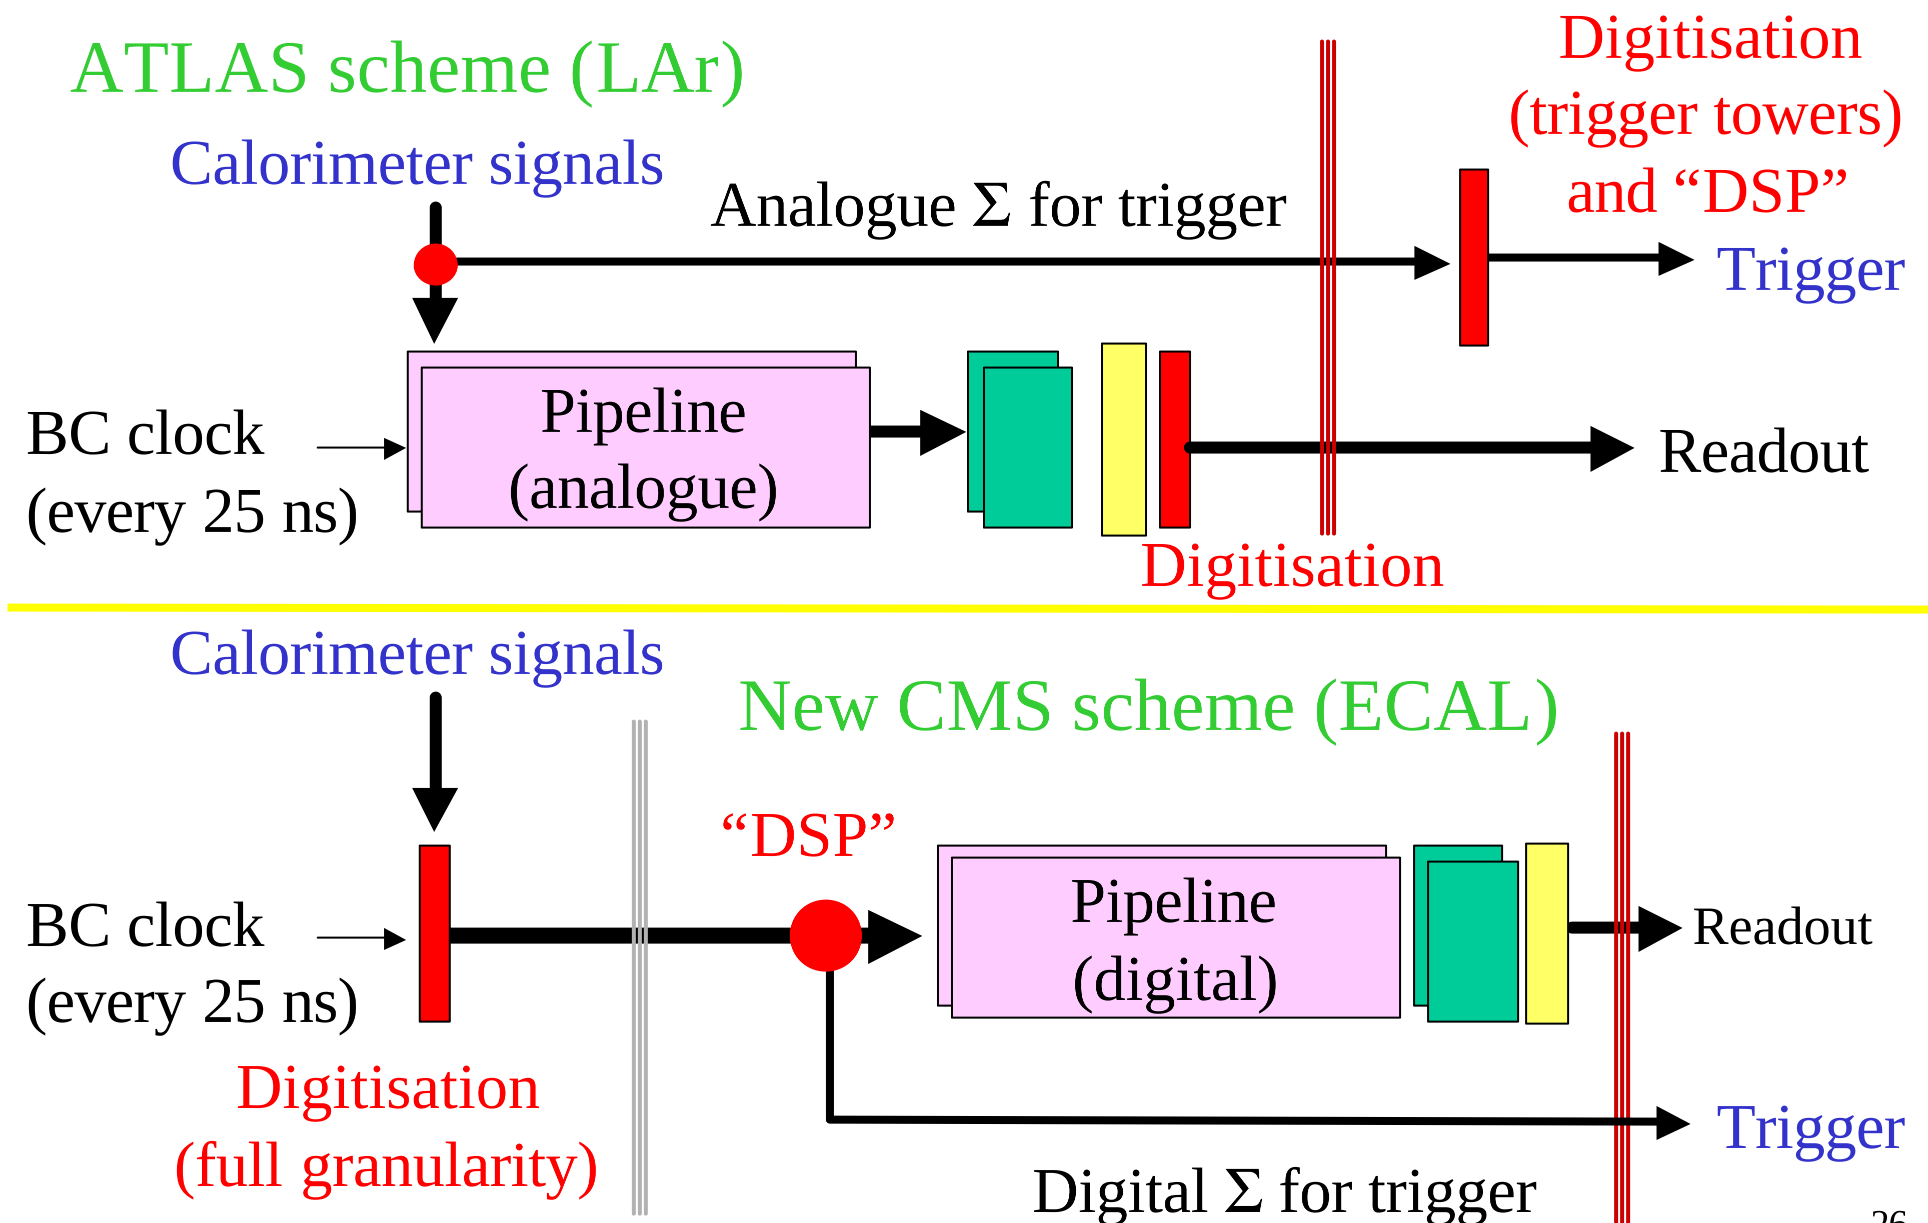

Bunch crossing identification

- Calorimeter signals extend over many bunch crossings
  - Need to combine information from a sequence of measurements to estimate the energy and identify the bunch crossing where the energy was deposited
- Apply Finite Impulse Response filter
  - Result $\rightarrow$ LUT to convert to E_T
  - Result $\rightarrow$ peak finder to determine BC where energy was deposited
- Need to take care of signal distortion for very large pulses
  - Don't lose most interesting physics!
- An ASIC incorporates the above

e.g. ATLAS

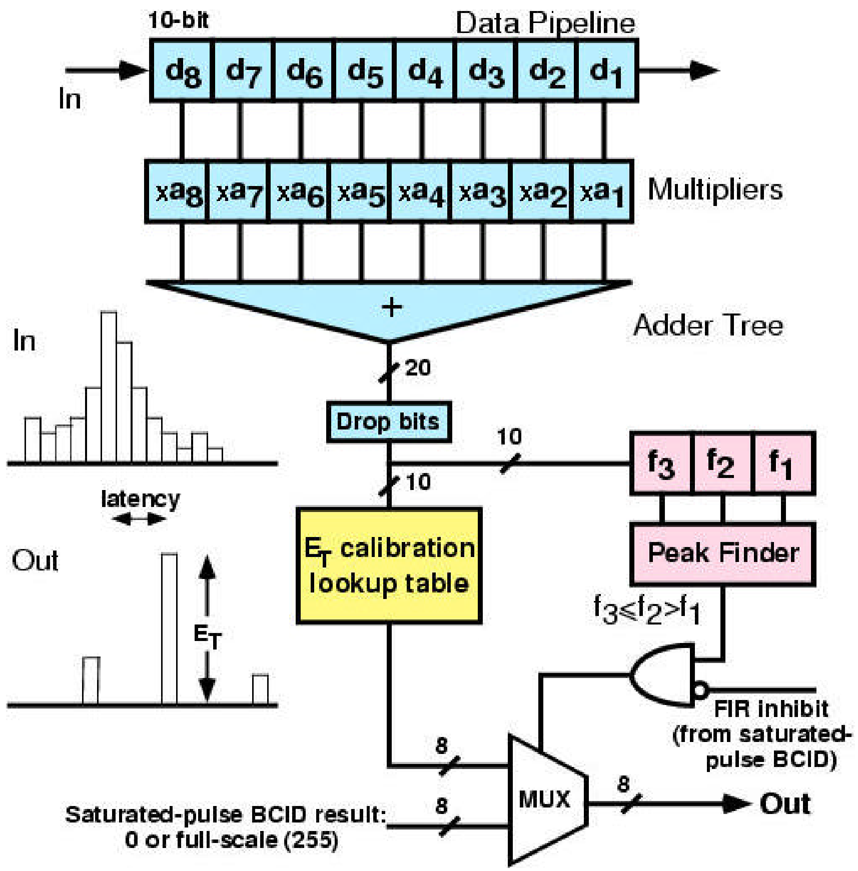

ATLAS Pre-Processor MCM and ASIC

- ADC
  - Use commercial 40 MHz ADCs
- ASIC (the only one in the calorimeter trigger)
  - ASIC handles 10-bit inputs from four commercial 40 MHz ADCs
    - Calibration, zero-suppression, BC identification, readout, etc
    - Cost effective solution given quantity needed
- MCM
  - Contains 4 ADCs, PPr ASIC and LVDS drivers
    - Allows high-density, cost-effective implementation

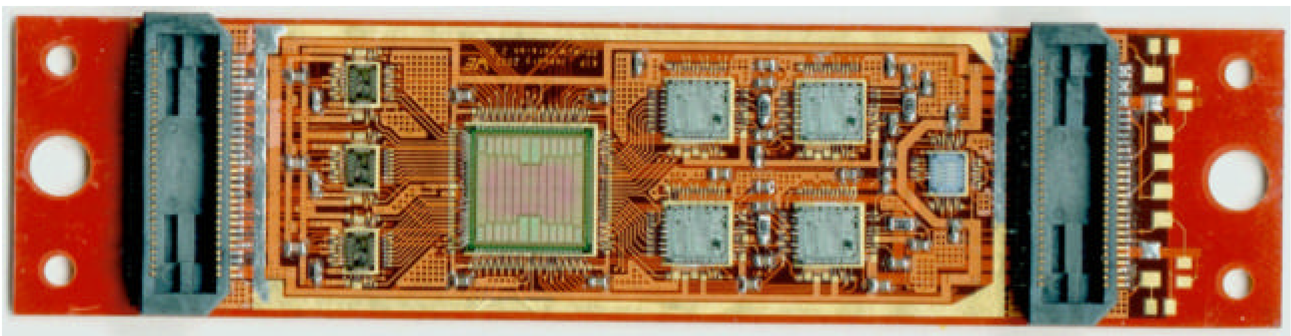

ATLAS e/γ trigger (implemented in CP)

- ATLAS e/γ trigger is based on 4×4 "overlapping, sliding windows" of trigger towers
  - Each trigger tower 0.1×0.1 in η×ϕ
    - η pseudo-rapidity, ϕ azimuth
  - ~3500 such towers in each of the EM and hadronic calorimeters
- There are ~3500 such windows
  - Each tower participates in calculations for 16 windows
    - This is a driving factor in the trigger design

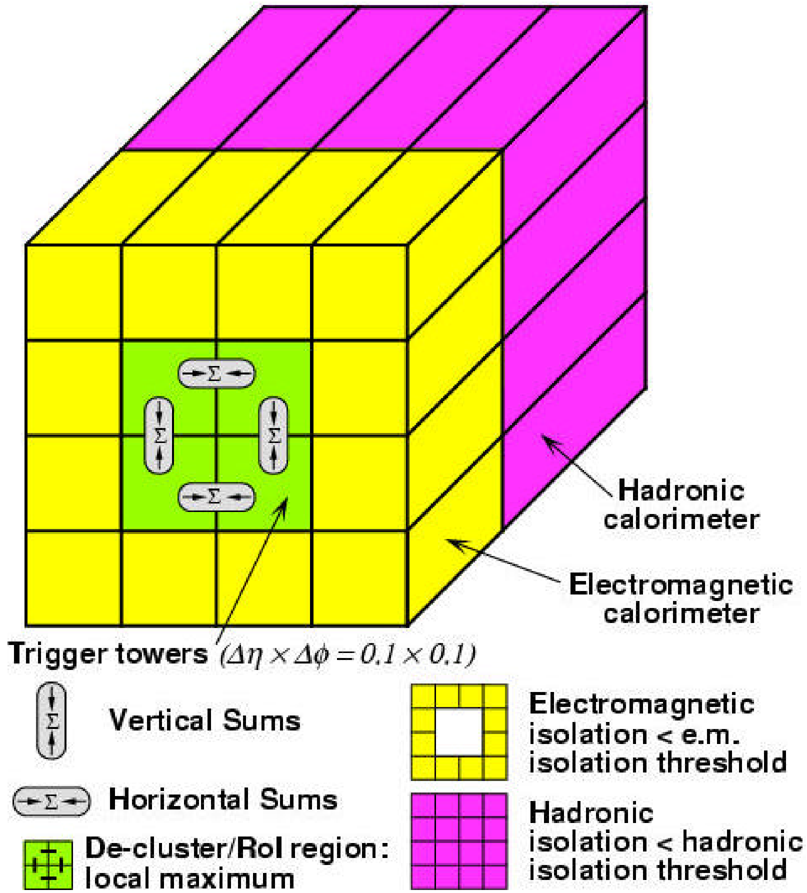

Slide shown earlier illustrates part of the processing for each window position

Note that logic must be duplicated for all ~3500 positions in calorimeter!

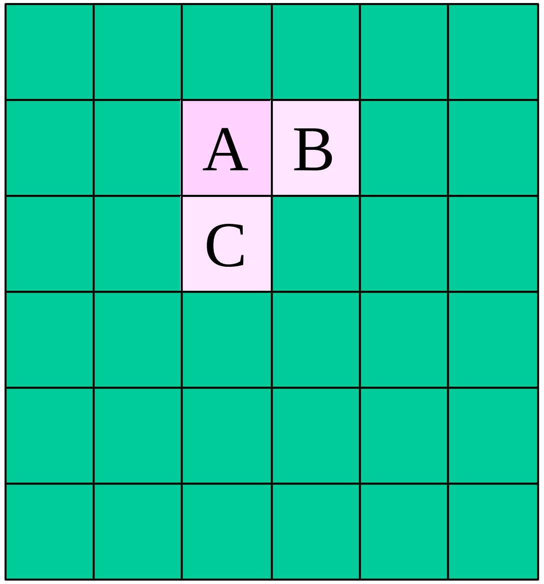

EM Calorimeter
(~3500 trigger towers)

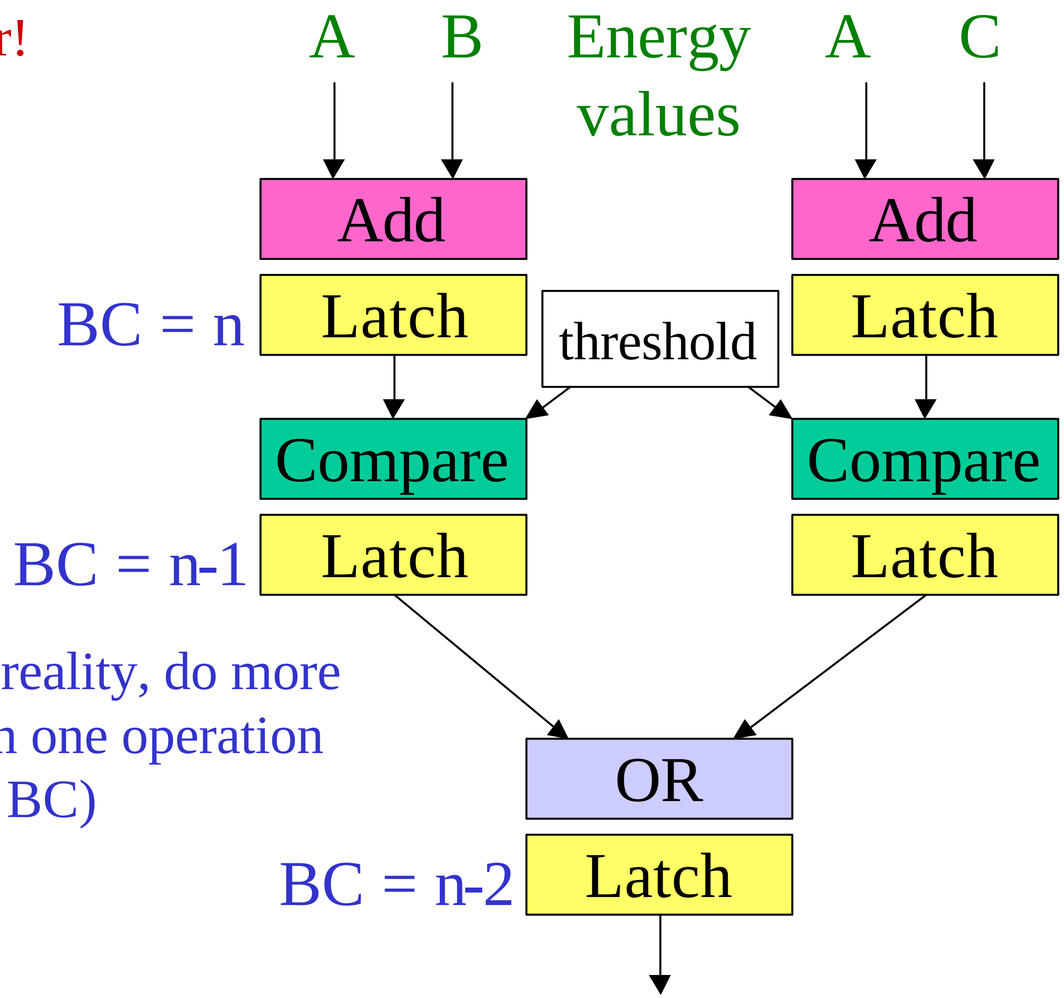

(In reality, do more than one operation per BC)

Data transmission and Cluster Processor

- The array of E_T values computed in the Preprocessor has to be transmitted to the CP
  - Use digital electrical links to CP modules (LVDS)
    - ~5000 links @ 400 Mbps
    - Convert to 160 Mbps single-ended signals on CP modules (LVDS rx; serializer FPGA)
  - Fan out data to neighbouring modules over very high density custom back-plane
    - ~800 pins per slot in 9U crate
    - 160 Mbps point-to-point
  - Fan out data to 8 large FPGAs in each CP module
    - Receive data at 160 Mbps in FPGAs that implement the algorithms
- The e/γ (together with the τ/h) algorithm is implemented in FPGAs
  - This has only become feasible with recent advances in FPGA technology
    - Require very large and very fast devices
  - Each FPGA handles 4×2 windows
    - Needs data from 7×5×2 towers ($\eta \times \phi \times$ {E/H})
  - Algorithm is described in a language (VHDL) that can be converted into the FPGA configuration file
    - Flexibility to adapt algorithms in the light of experience
  - Parameters of the algorithms can be changed easily
    - E.g. cluster-E_T thresholds are held in registers that can be programmed without reconfiguring the FPGAs

MUON TRIGGERS

- Will illustrate with example of CMS drift-tube trigger
- See related talks in parallel sessions:
  - ATLAS
    - K. Nagano: The ATLAS level-1 muon to central-trigger processor interface (MUCTPI)
    - R. Ichimiya: An implementation of the sector logic for the endcap muon trigger of the ATLAS experiment
    - H. Kano: Results of a slice system test for the ATLAS endcap muon level-1 trigger
    - R. Vari: The design of the coincidence matrix ASIC of the ATLAS barrel level-1 muon trigger

CMS muon system

- CMS muon system includes three detector technologies
  - RPC and DT in barrel
  - RPC and CSC in endcaps
- All three detector systems participate in the first-level trigger
  - Specific logic for each system
  - Global logic that combines all the muon information
- After some general introductory remarks on muon triggers, I will discuss as an example the Drift Tube (DT) trigger
  - Combines information from four DT muon stations (see figure)

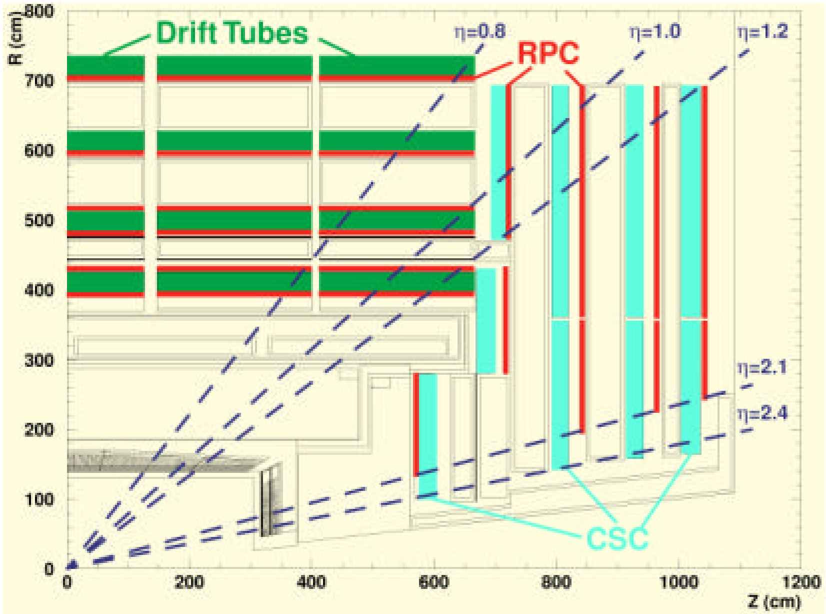

CMS muon trigger overview

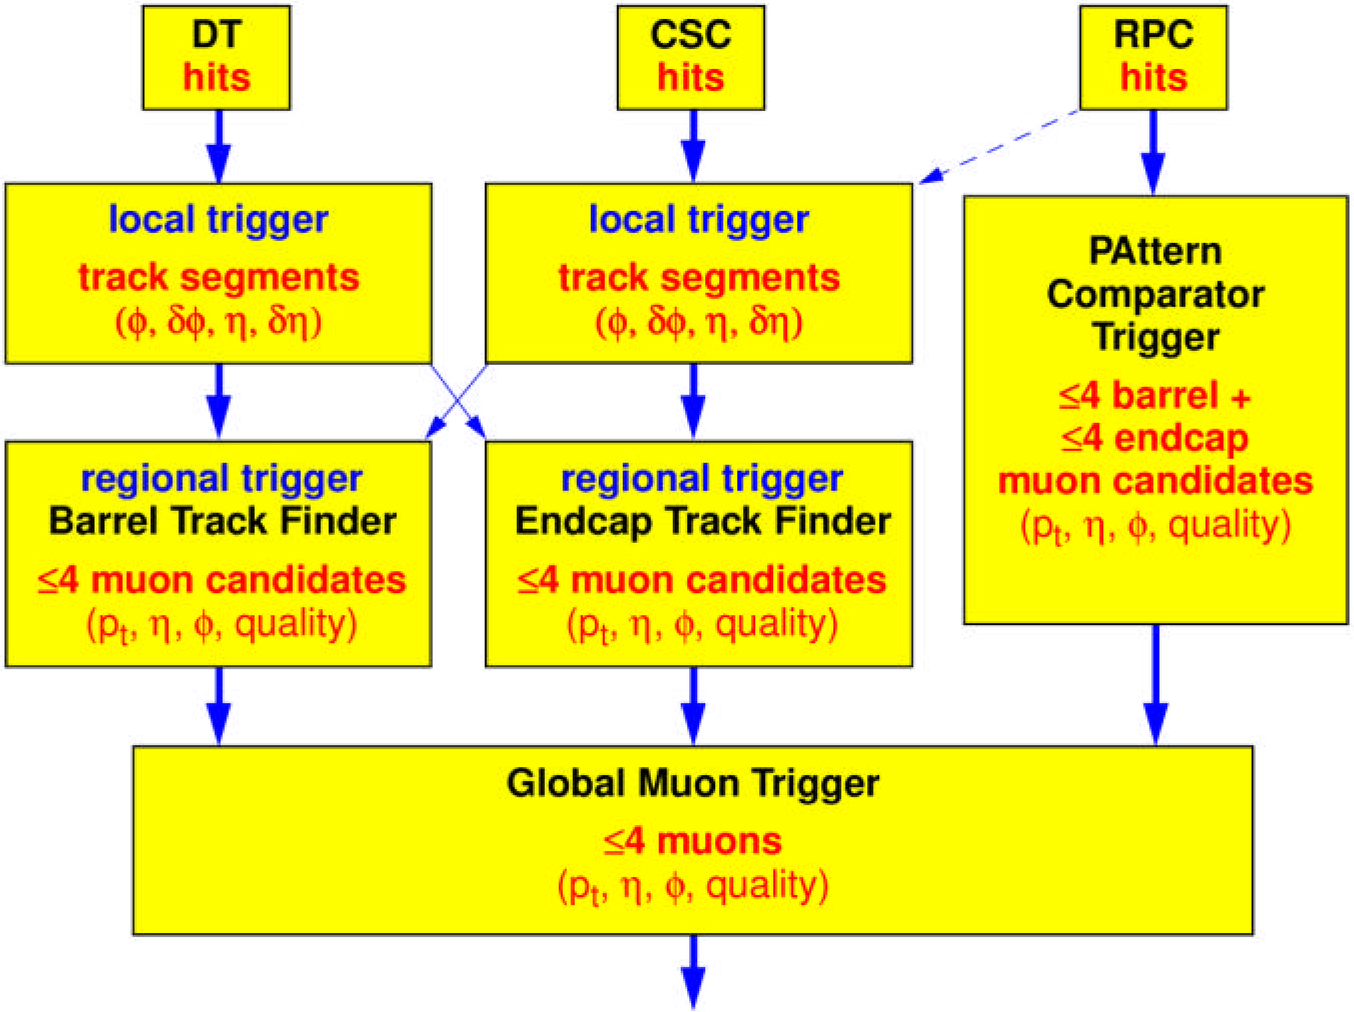

CMS global trigger receives p_T, η, ϕ information for candidate e/γ, μ, etc.
(ATLAS central trigger works with multiplicity information only)

Muon triggers

- In general, muon triggers look for a pattern of hits in the muon chambers consistent with a high-p_T muon originating from the collision point
  - The deflection in the magnetic field is inversely proportional to p_T
    - An infinite-momentum muon follows a straight-line trajectory
- Some of the detectors used in the triggers have a response time below 25 ns (e.g. RPCs)
- For slower detectors, information from several chamber layers has to be combined to identify locally which bunch crossing gave rise to the hits, as well as giving the position of the muon in the chambers
  - Local track segments or "superhits" (identified BC, position)
    - In some cases, e.g. DT, also direction information

Illustration — principle of DT trigger

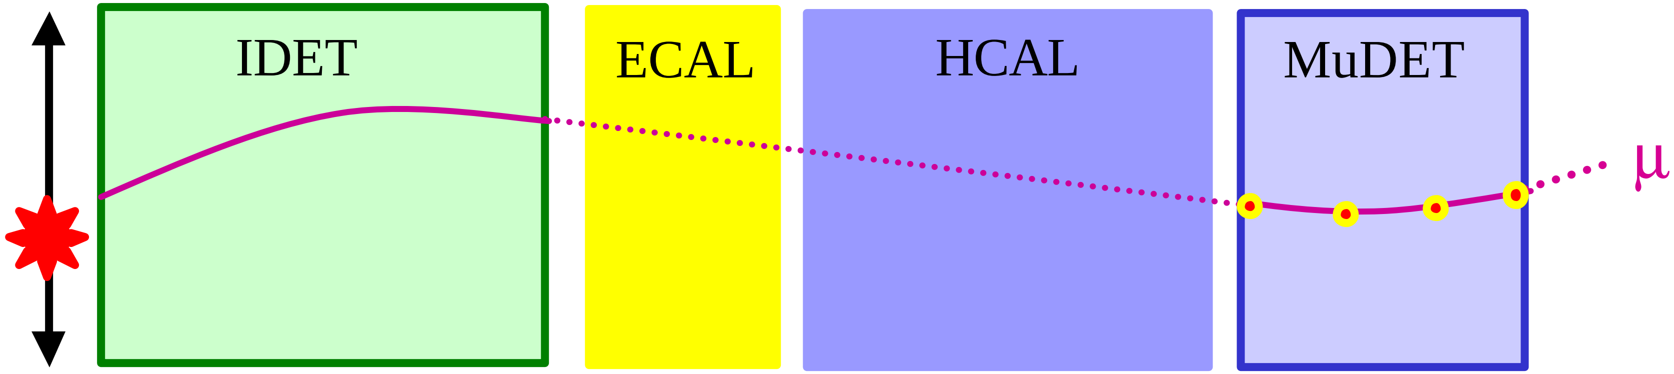

2 chamber layers

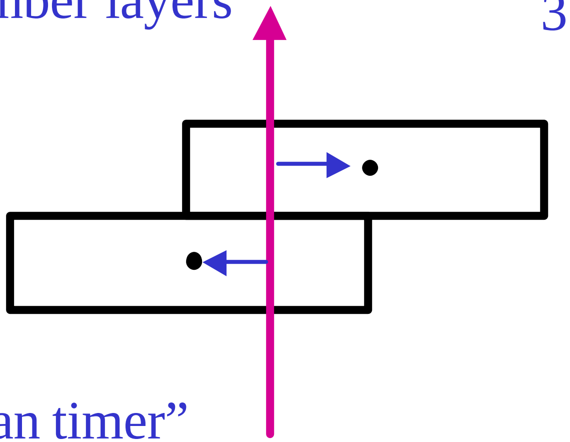

“mean timer”

$T_1+T_2 = T_{max}$

$(T_1-T_2)/2v_d = x$

3 chamber layers - inclined tracks

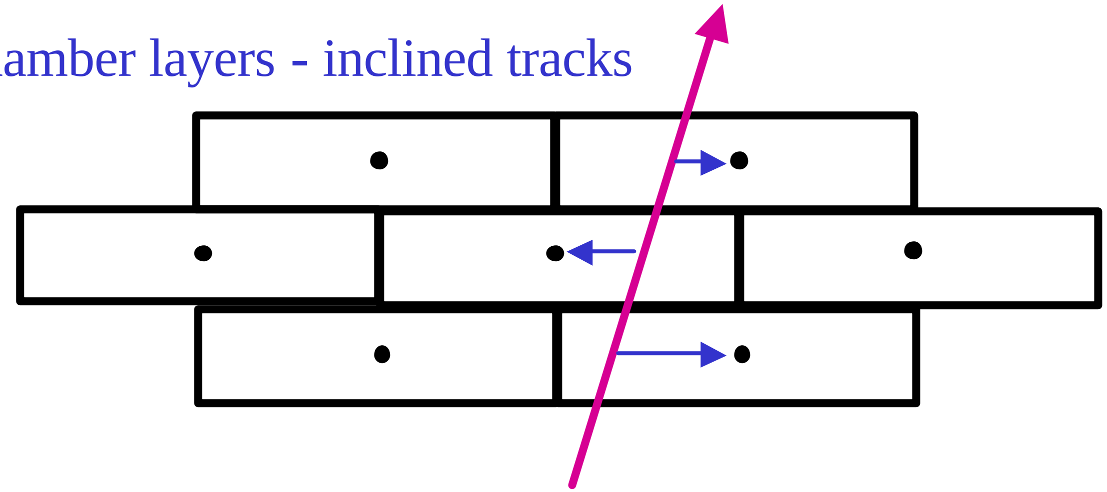

Extending the scheme to 4 DT layers, can handle inclined tracks even if 1 hit lost due to inefficiency or dead region

- provides identified BC, position, angle with high efficiency

Maximum DT 380 ns >> 25 ns

CMS local Drift Tube muon trigger

Bunch & **T**ime **I**dentification

TRAck **CO**rrelator

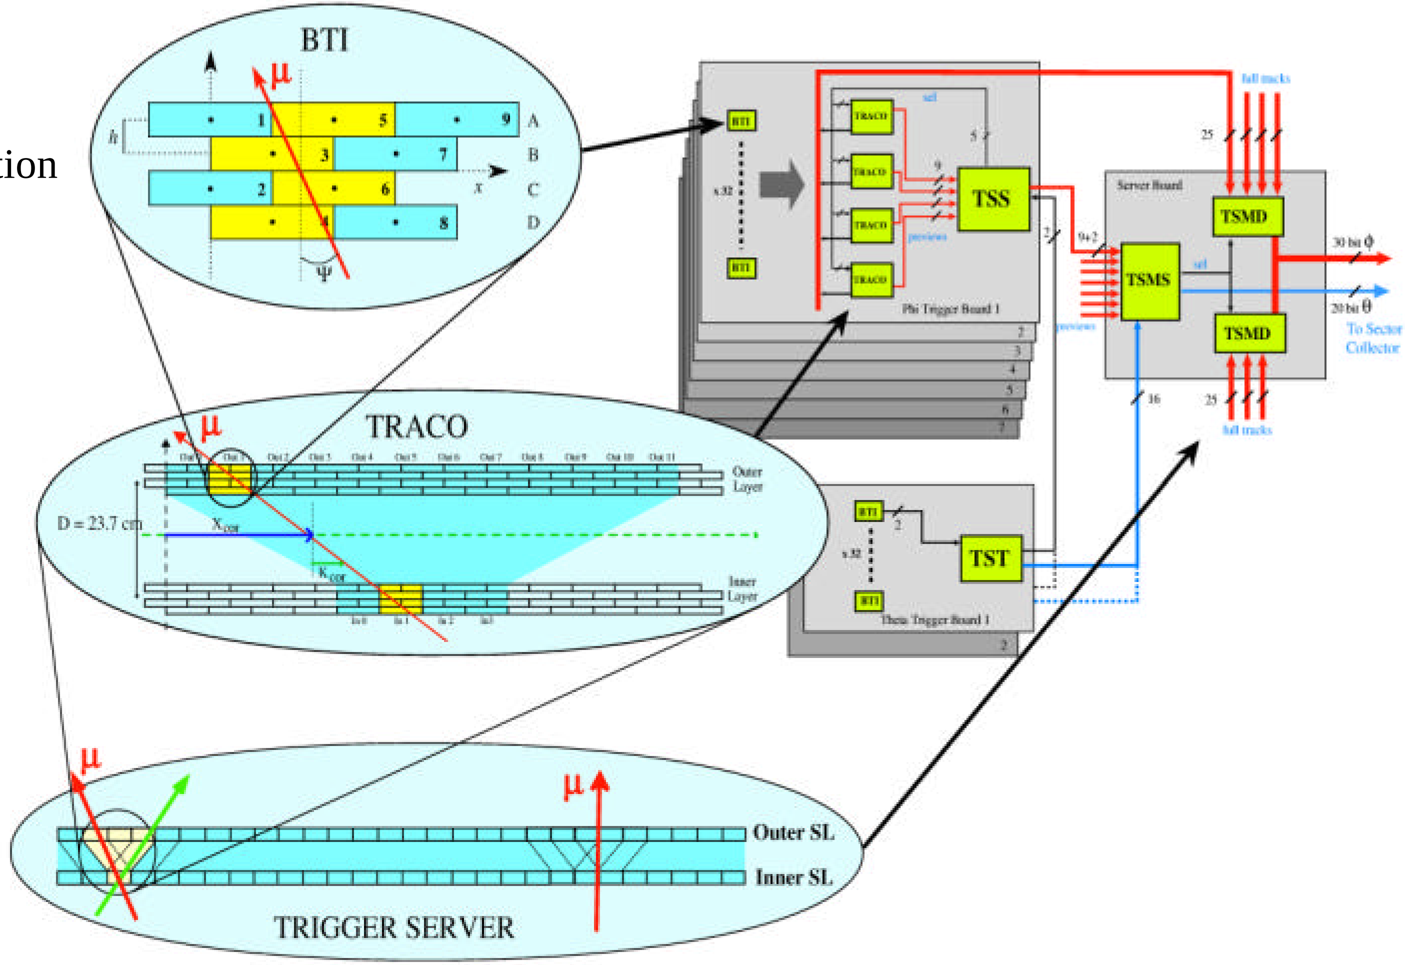

Local trigger electronics associated with each Super Layer
is mounted on the detector and implemented using ASICs

DT trigger - prototype

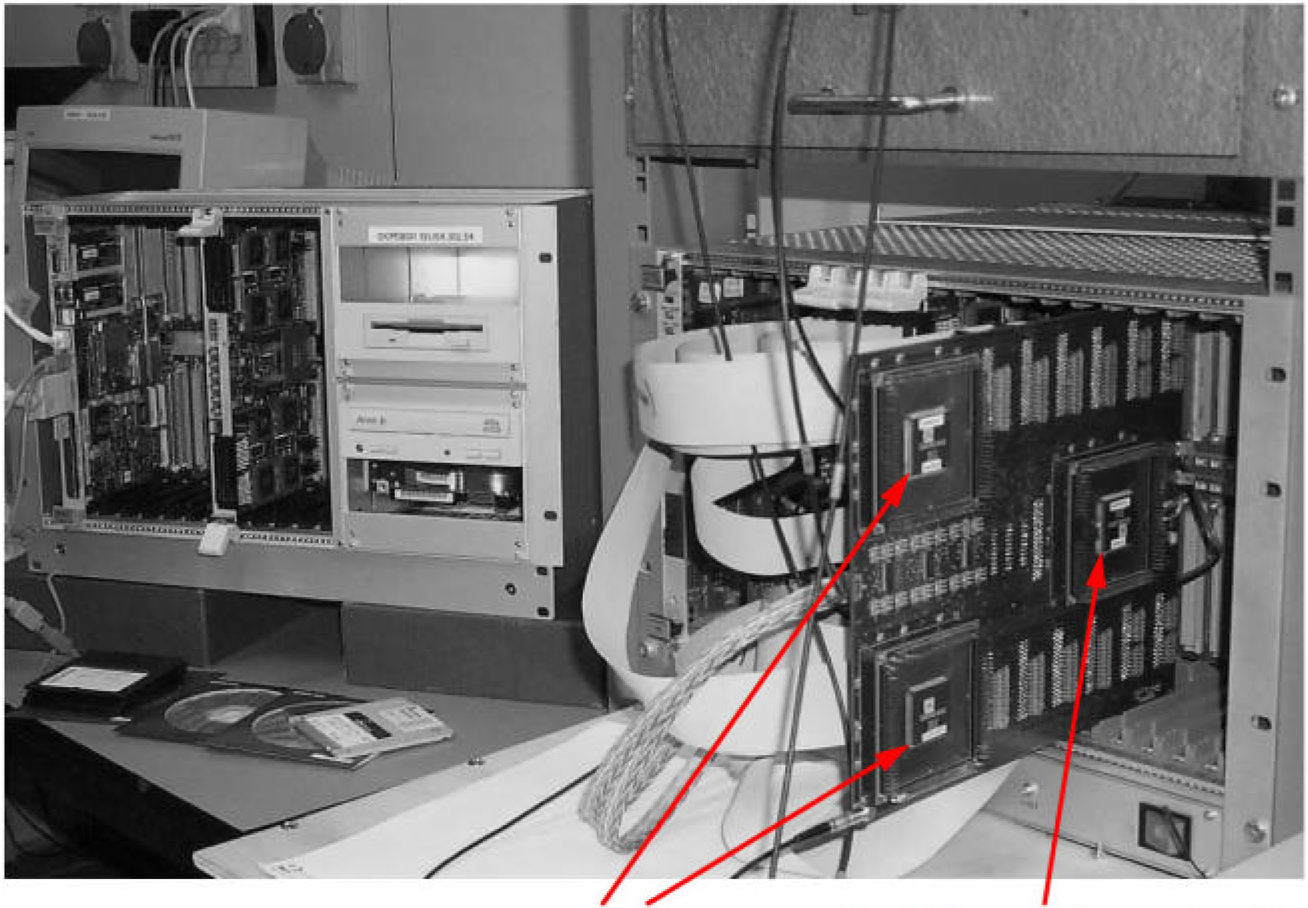

Track Sorter Master (data) Track Sorter Master (sorting)

CMS DT track finder

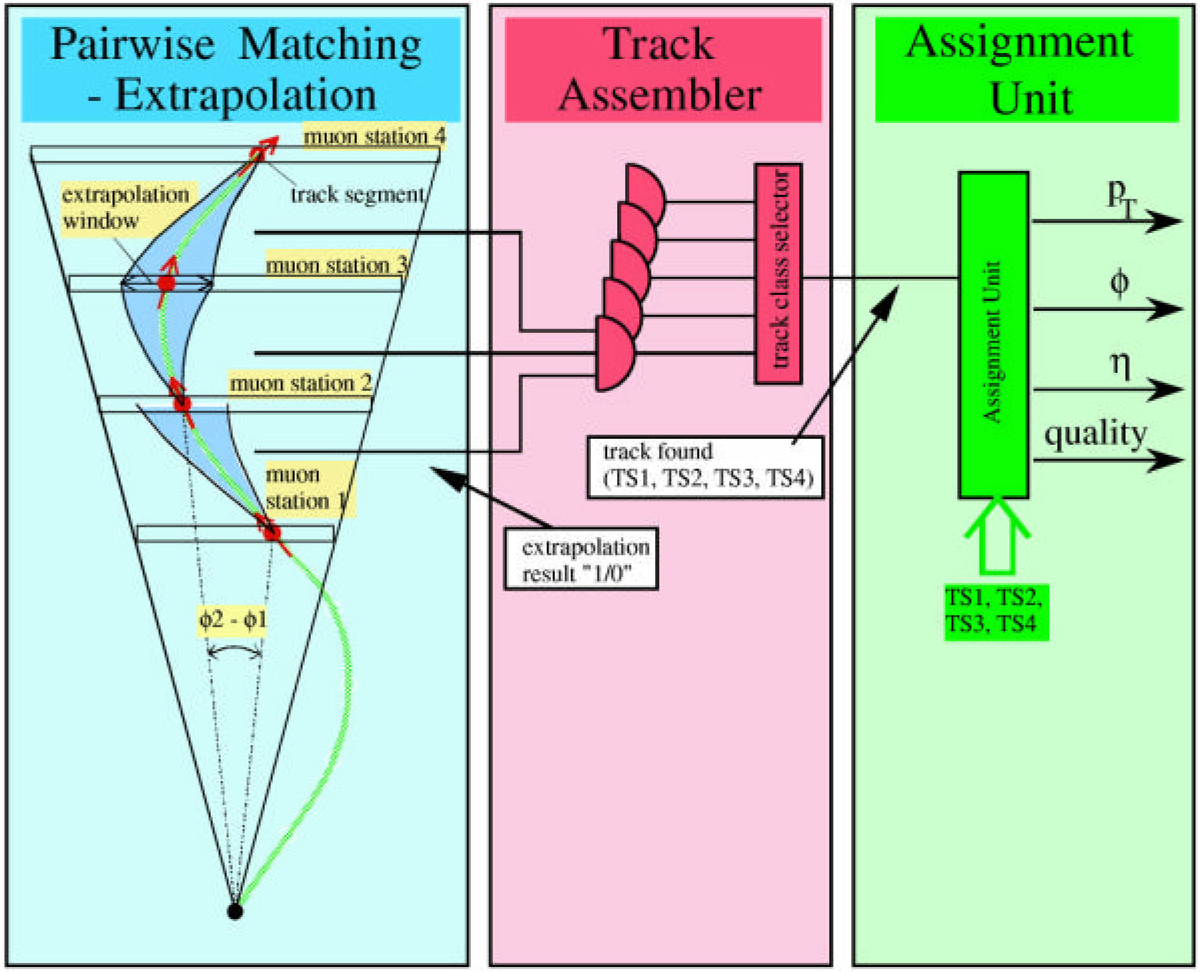

Track-finder electronics is mounted off detector and is implemented using FPGAs

- LUTs in FPGAs contain limits of extrapolation windows
- Track segments are combined to find the "best" two tracks within a sector
- The track parameters are then determined from the ϕ measurements in different stations

LHCb PILE-UP VETO

(see L. Wiggers' talk in parallel session)

- LHCb is designed to work with single-interaction events
  - Operate at lower luminosity ($L = 2\times10^{32}$ $cm^{-2}s^{-1}$)
    - 30% of BCs have single interaction
    - 10% of BCs have >1 interaction
  - Include pile-up veto in "level-0" trigger
    - Avoid triggering on multi-interaction events that are not useful for the analysis
- Trigger on muons, electrons/photons and hadrons
  - Much lower p_T thresholds than in ATLAS and CMS
  - Possible thanks to absence of pile-up and high input-rate capability of second level of triggering
    - Second triggering level ("level-1") designed for 1 MHz input rate
      - Reduce rate to ~40 kHz with latency up to 2 ms (software)
      - Includes secondary-vertex trigger

LHCb pile-up veto algorithm

Si strip detectors

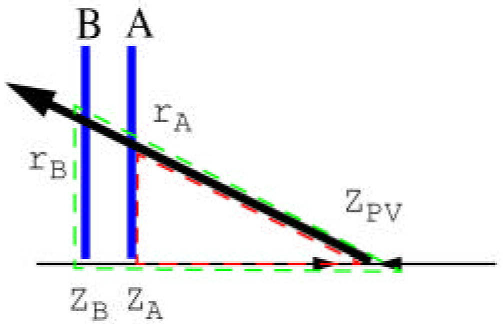

$$\frac{r_A}{r_B} = \frac{z_{PV} - z_A}{z_{PV} - z_B} = k$$

X=Y=0

Tracks from same vertex give same k

Different vertices generate different k

Histogramming method

- Histogram z for combinations of hits
- Find position of highest peak
- In second pass, omit hits that contributed to the first peak

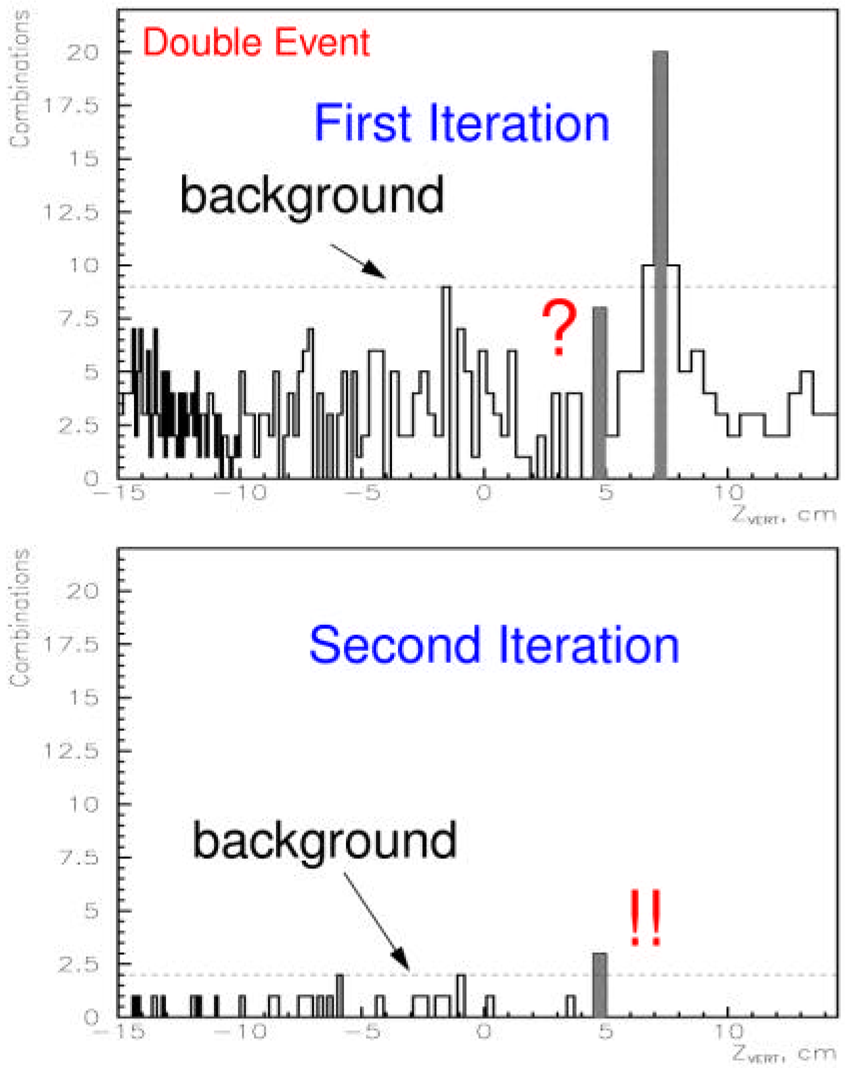

Farming in the first-level trigger

- Vertex finder will be implemented using FPGAs
  - Use “farm” of 4 (+ 1 spare) FPGA-based vertex finders, each one handling one event in four
  - Multiplex data from different quadrants into the vertex finders over a period of 4 BCs
    - Reduces the data rate into each finder by a factor of 4
    - Each vertex finder uses parallel and pipelined processing

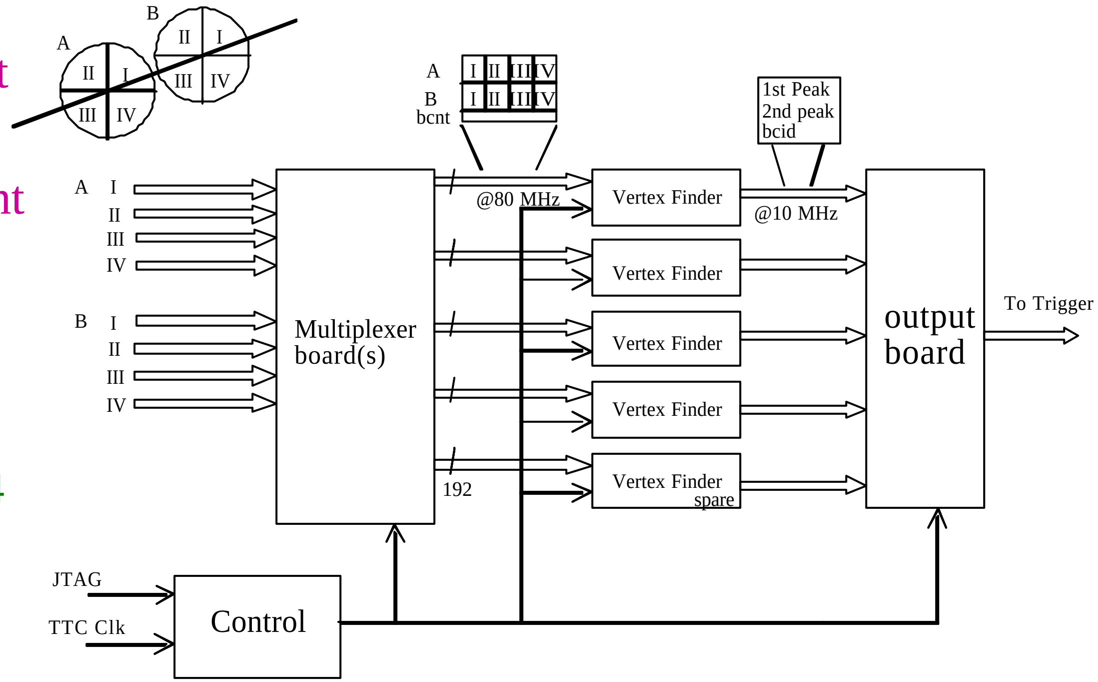

CENTRAL/GLOBAL TRIGGERS

- Will illustrate with example of CMS Global trigger
- See related talk in parallel sessions:
  - LHCb
    - R. Cornat: Level-0 trigger decision unit for the LHCb experiment

Global trigger decision

- Global trigger has to combine information from the different parts of the first-level trigger
  - Local objects: μ, e/γ, τ/h, jet
  - Energy sums
- Makes overall decision based on combinations of conditions
  - Inclusive triggers
    - E.g. $p_T(\mu) > 20$ GeV
  - More complex requirements
    - E.g. $p_T(\text{jet}) > 100$ GeV **and** $E_T^{miss} > 100$ GeV
  - Topological conditions (CMS)
    - E.g. $p_T(\mu_1) > 20$ GeV **and** $p_T(\mu_2) > 20$ GeV **and** $170^o<|\phi(1)-\phi(2)|<190^o$

Example: CMS global trigger

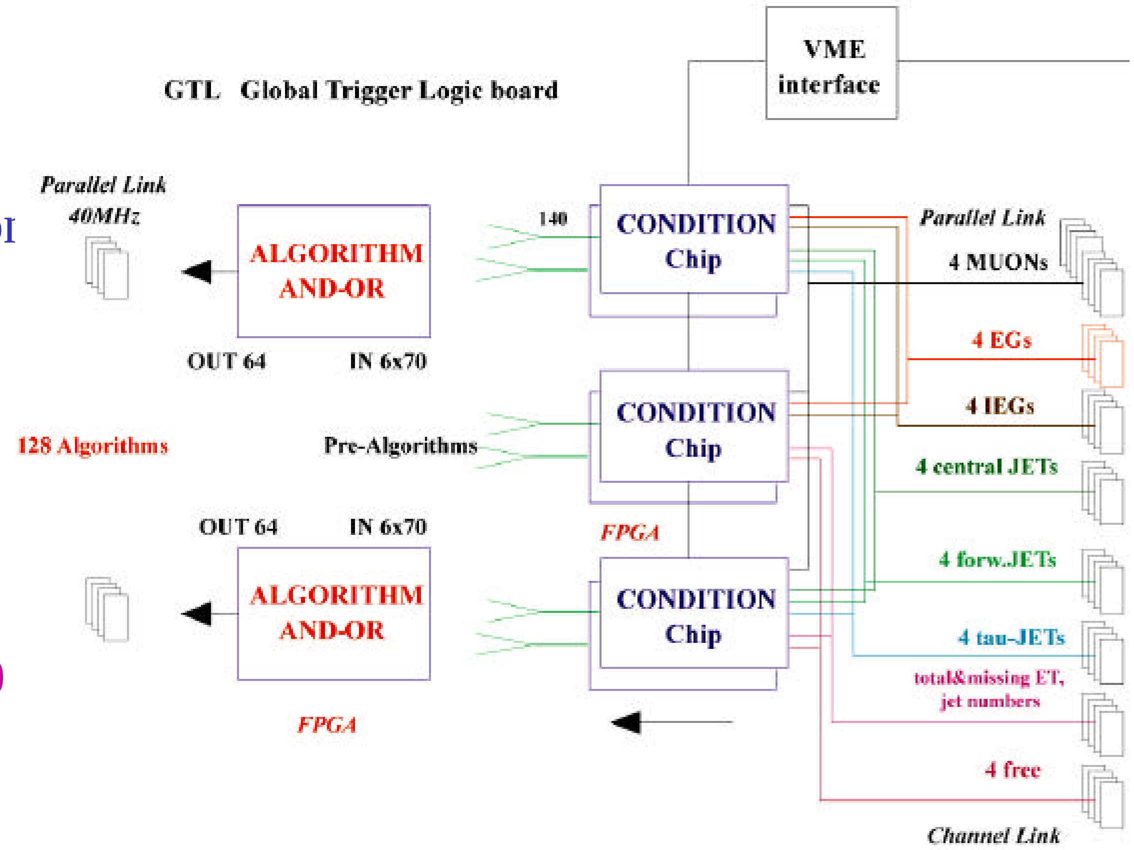

Implemented in FPGAs

Concluding remarks

- First-level triggers for LHC represent a huge challenge
  - Direct impact on the physics potential of the experiments
    - First stage of physics selection
      - 100 kHz is $O(10^{-4})$ of interaction rate in ATLAS and CMS
    - Events rejected are lost forever
  - Benefit from new technologies for processing and data movement
    - Latest generation FPGAs and ASICs
    - High-speed optical and electrical links
  - Lots of challenges for engineers and physicists working together
    - Algorithms, electronics and software
- A lot of design work and prototyping has been done
  - But there is still plenty to do!
    - Final design and prototyping at module, subsystem and system level