Lecture 2

Questions from last lecture?

Plan for today's lecture

- Requirements and constraints for triggering at the LHC
  - Driven by the physics objectives of the experiments
    - ATLAS and CMS (general-purpose, proton-proton, discovery physics)
    - LHC*b* (B physics, proton-proton)
    - ALICE (specialized for heavy-ion collisions)
- Case study from LHC illustrating how the T/DAQ implementation follows the ideas presented yesterday
  - Example of electron/photon trigger in ATLAS
- Commissioning of the T/DAQ systems for LHC in 2008
  - With single bunch, single beam at the injection energy of 450 GeV
    - No collisions or acceleration yet, unfortunately!
  - And using cosmic rays (won't have time to cover this)

Requirements – ATLAS and CMS

- Triggers in the general-purpose proton–proton experiments, ATLAS and CMS, will have to:
  - Retain as many as possible of the events of interest for the diverse physics programmes of these experiments
    - Higgs searches (Standard Model and beyond)
      - e.g. $H \rightarrow ZZ \rightarrow$ leptons, $H \rightarrow gg$; also $H \rightarrow tt$, $H \rightarrow bb$
    - SUSY searches
      - With and without R-parity conservation
    - Searches for other new physics
      - Using inclusive triggers that one hopes will be sensitive to any unpredicted new physics
    - Precision physics studies
      - e.g. measurement of W mass
    - B-physics studies (especially in the early phases of these experiments)
  - N.b. selections often need to be made at analysis level to suppress backgrounds, so focus especially on events that will be retained

Constraints – ATLAS and CMS

- However, they also need to reduce the event rate to a manageable level for data recording and offline analysis
  - L = 10^{34} $cm^{-2}s^{-1}$, and $\sigma \sim 100$ mb $\Rightarrow$ 10^9 Hz interaction rate
    - Even rate of events containing leptonic W and Z decays is *O*(100 Hz)
  - The size of the events is very large, *O*(1) MByte
    - Huge number of detector channels, high particle multiplicity per event
  - Recording and subsequently processing offline, *O*(100) Hz event rate per exp^t with *O*(1) MByte event size implies major computing resources!
  - Hence, only a tiny fraction of proton–proton collisions can be selected
    - Maximum fraction of interactions triggering at full luminosity *O*(10^{-7})
- Have to balance needs of maximising physics coverage and reaching acceptable (i.e. affordable) recording rates

LHC*b*

- The LHC*b* experiment, which is dedicated to studying B-physics, faces similar challenges to ATLAS and CMS
  - It will operate at a comparatively low luminosity (~2×10^{32} $cm^{-2}s^{-1}$), giving an overall proton–proton interaction rate of ~20 MHz
    - Chosen to maximise the rate of single-interaction bunch-crossings
  - The event size is comparatively small (~100 kByte)
    - Fewer detector channels
    - Less occupancy due to lower luminosity
  - However, there is a very high rate of beauty production
    - Given $\sigma \sim 500$ μb, bb production rate ~100 kHz
  - The trigger must therefore search for 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 also very demanding, particularly from the DAQ point of view
  - The total interaction rate will be much smaller than in the pp experiments
    - $L \sim 10^{27}\ cm^{-2}s^{-1} \Rightarrow R \sim 8000$ Hz for Pb–Pb collisions
  - 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 particle 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 the full detector is read out
  - Even more than in the other experiments, the volume of data to be stored and subsequently processed offline will be massive
    - Data rate to storage ~1 GByte/s (limited by what is possible/affordable)

Signatures used for triggers

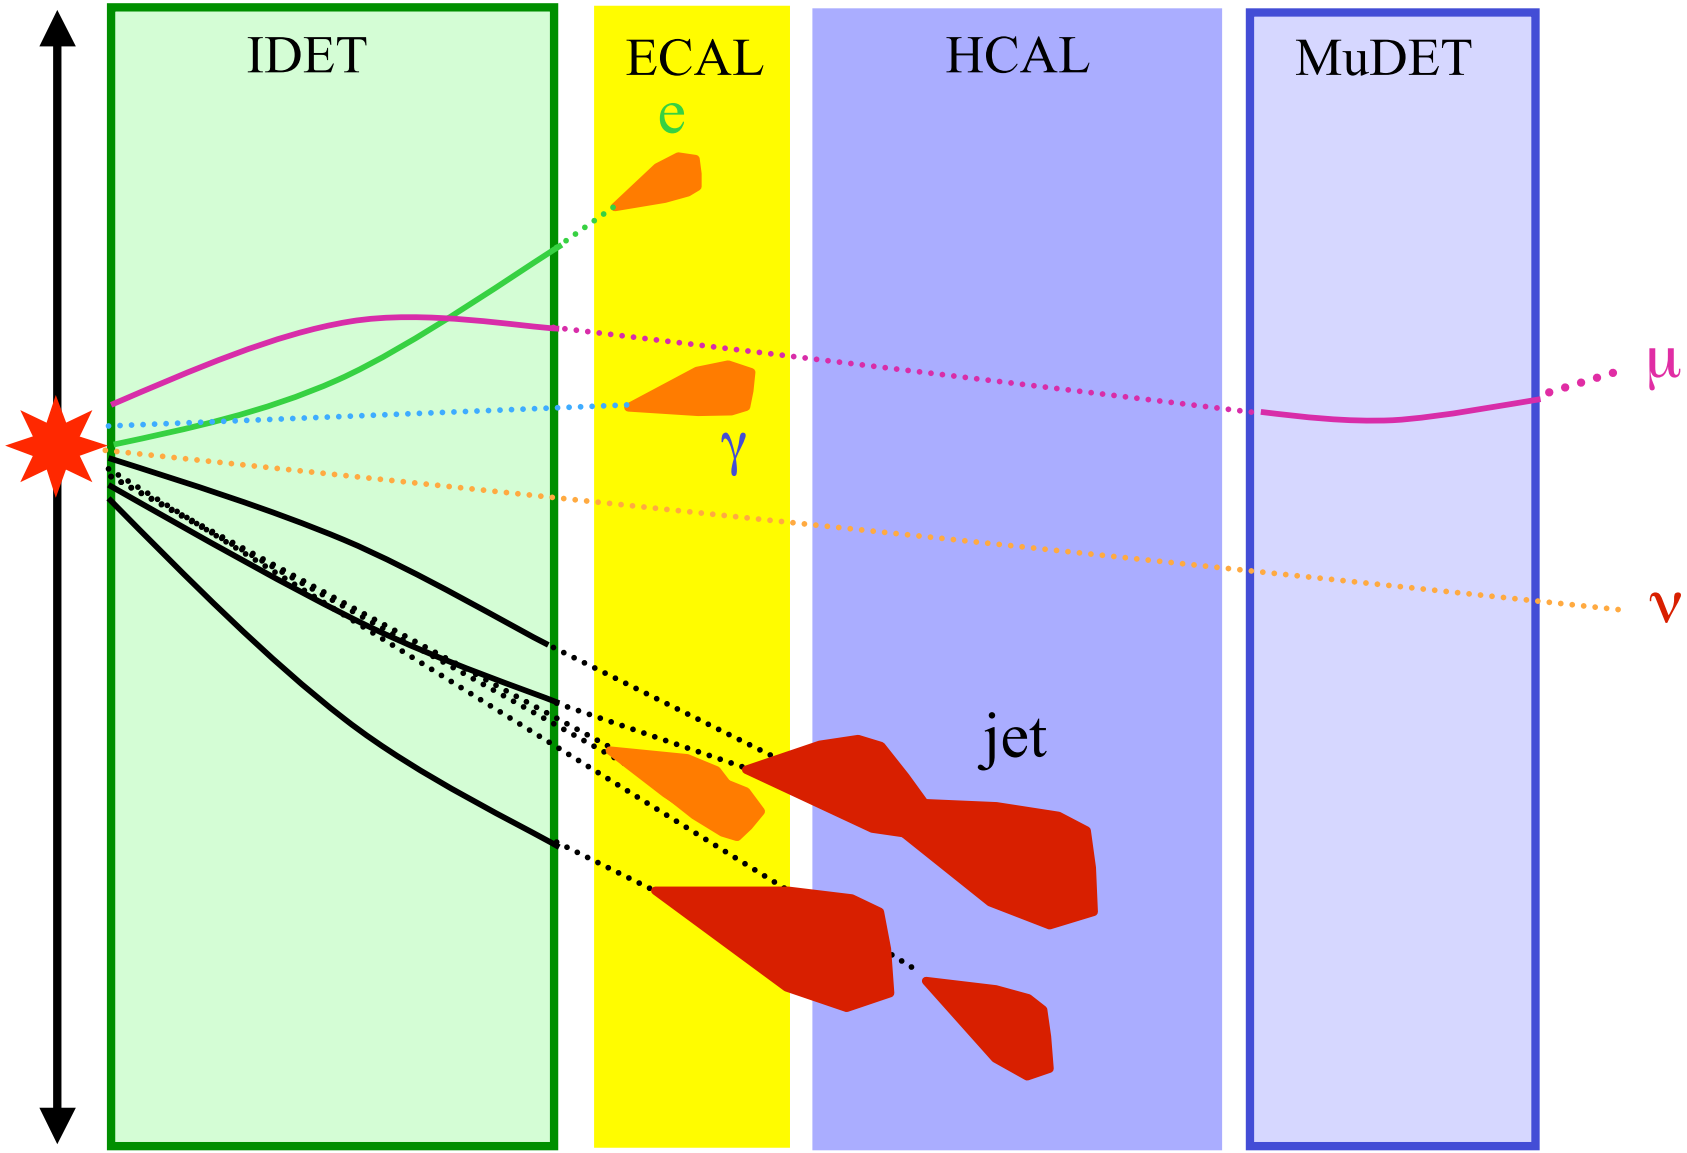

LVL1 selection criteria

- Features that distinguish new physics from the bulk of the cross-section for Standard Model processes at hadron colliders are
  - In general, the presence of high-p_T particles (or jets)
    - e.g. these may be the products of the decays of new heavy particles
      - In contrast, most of the particles produced in minimum-bias interactions are soft ($p_T \sim 1$ GeV or less)
  - More specifically, the presence of high-p_T leptons (e, μ, τ), photons and/or neutrinos
    - e.g. the products (directly or indirectly) of new heavy particles
      - These give a clean signature c.f. low-p_T hadrons in minimum-bias case, especially if they are "isolated" (i.e. not inside jets)
  - The presence of known heavy particles
    - e.g. W and Z bosons may be produced in Higgs particle decays
      - Leptonic W and Z decays give a very clean signature
        - » Also interesting for physics analysis and detector studies

LVL1 signatures at hadron colliders

- LVL1 triggers therefore search for
  - High-p_T muons
    - Identified beyond calorimeters; need p_T cut to control rate from $\pi^+ \rightarrow \mu\nu$, $K^+ \rightarrow \mu\nu$, as well as semi-leptonic beauty and charm decays
  - High-p_T photons
    - Identified as narrow EM calorimeter clusters; need cut on E_T; cuts on isolation and hadronic-energy veto reduce strongly rates from high-p_T jets
  - High-p_T electrons
    - Same as photon (some experiments require matching track already at LVL1)
  - High-p_T taus (decaying to hadrons)
    - Identified as narrow cluster in EM+hadronic calorimeters
  - High-p_T jets
    - Identified as cluster in EM+hadronic calorimeter — need to cut at very high p_T to control rate (jets are dominant high-p_T process)
  - Large missing E_T or total scalar E_T

LVL1 trigger menu

- An illustrative trigger menu for LHC at high luminosity includes:
  - One or more muons with $p_T > 20$ GeV (rate ~ 11 kHz)
  - Two or more muons each with $p_T > 6$ GeV (rate ~ 1 kHz)
  - One or more e/γ with $E_T > 30$ GeV (rate ~ 22 kHz)
  - Two or more e/γ each with $E_T > 20$ GeV (rate ~ 5 kHz)
  - One or more jets with $E_T > 290$ GeV (rate ~ 200 Hz)
  - One or more jets with $E_T > 100$ GeV & $E_T^{miss} > 100$ GeV (rate ~ 500 Hz)
  - Three or more jets with $E_T > 130$ GeV (rate ~ 200 Hz)
  - Four or more jets with $E_T > 90$ GeV (rate ~ 200 Hz)
- Full menu will include many items in addition (>100 items total)
  - Items with τ (or isolated single-hadron) candidates
  - Items with combinations of objects (e.g. muon & electron)
  - Pre-scaled triggers with lower thresholds
  - Triggers for technical studies and to aid understanding of data
    - e.g. trigger on bunch-crossings at random to collect unbiased sample

HLT menu

- Illustrative menu for LHC at 2×10^{33} $cm^{-2}s^{-1}$ luminosity (CMS):
  - $p_T^e > 29$ GeV or 2 electrons $p_T^e > 17$ GeV
    - Rate ~ 34 Hz
  - $p_T^\gamma > 80$ GeV or 2 photons $p_T^\gamma > 40, 25$ GeV
    - Rate ~ 9 Hz
  - $p_T^\mu > 19$ GeV or 2 muons $p_T^\mu > 7$ GeV
    - Rate ~ 29 Hz
  - $p_T^\tau > 86$ GeV or 2 taus $p_T^\tau > 59$ GeV
    - Rate ~ 4 Hz
  - $p_T^{jet} > 180$ GeV and missing $E_T > 123$ GeV
    - Rate ~ 5 Hz
  - $p_T^{jet} > 657$ GeV or 3 jets $p_T^{jet} > 247$ GeV or 4 jets $p_T^{jet} > 113$ GeV
    - Rate ~ 9 Hz
  - Others (electron•jet; b-jets, etc.)
    - Rate ~ 7 Hz
  - Total ~ 100 Hz of which a large fraction is "physics" – large uncertainty on rates!
    - Need to balance physics coverage against offline computing cost

Case study (e/γ trigger in ATLAS)

- Some general LVL1 trigger issues
  - The ATLAS LVL1 trigger system
- LVL1 e/γ algorithm and implementation
  - Example of pipe-lined trigger processing
- High-level electron trigger
  - Example where sophisticated event-selection algorithms need to be used online to get the required separation, with good signal efficiency and with high background rejection
    - Background primarily comes from high-p_T jets that are misidentified as electrons
      - There are lots of jets!

General LVL1-trigger design goals

- Need large reduction in physics rate already at the first level (otherwise readout system becomes unaffordable)
  - $O(10^9)$ interaction rate → less than 100 kHz in ATLAS and CMS
    - Require complex algorithms to reject background while keeping signal
- An important constraint is to achieve a short latency
  - Information from all detector channels ($O(10^8)$ channels!) has to be held in local memory on detector pending the LVL1 decision
    - The pipeline memories are typically implemented in ASICs (Application Specific Integrated Circuits), and memory size contributes to the cost
  - Typical values are a few μs (e.g. less than 2.5 μs ATLAS, 3.2 μs CMS)
- Require flexibility to react to changing conditions (e.g. wide luminosity range) and — hopefully — new physics
  - Algorithms must be programmable (adjustable parameters at least)

Overview of ATLAS LVL1 trigger

Radiation tolerance, cooling, grounding, magnetic field, no access

~7000 calorimeter trigger towers

O(1M) RPC/TGC channels

Calorimeter trigger

Pre-Processor (analogue → E_T)

Jet / Energy-sum Processor

Cluster Processor (e/γ, τ/h)

Muon trigger

Muon Barrel Trigger

Muon End-cap Trigger

Muon central trigger processor

Design all digital, except input stage of calorimeter trigger Pre-Processor

Central Trigger Processor (CTP)

Local Trigger Processors (LTP)

Timing, Trigger, Control (TTC)

Latency limit 2.5 μs

ATLAS LVL1 e/γ trigger

- 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

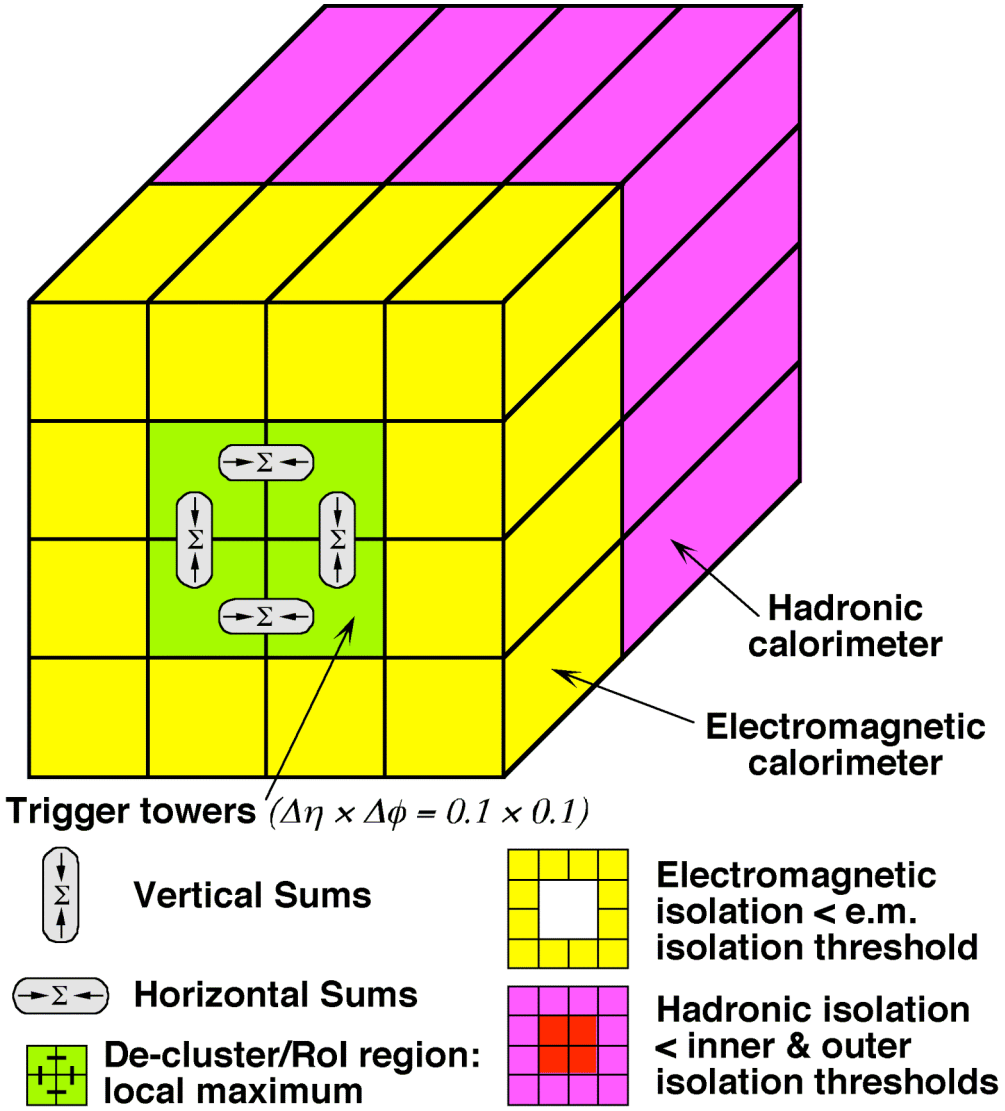

ATLAS LVL1 calorimeter trigger

- 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 Cluster Processor (only 4 crates in total)
  - Fan out values needed in more than one crate
    - Motivation for very compact design of processor
- Within CP crate, values need to be fanned out between electronic modules, and between processing elements on the modules
- Connectivity and data-movement issues drive the design

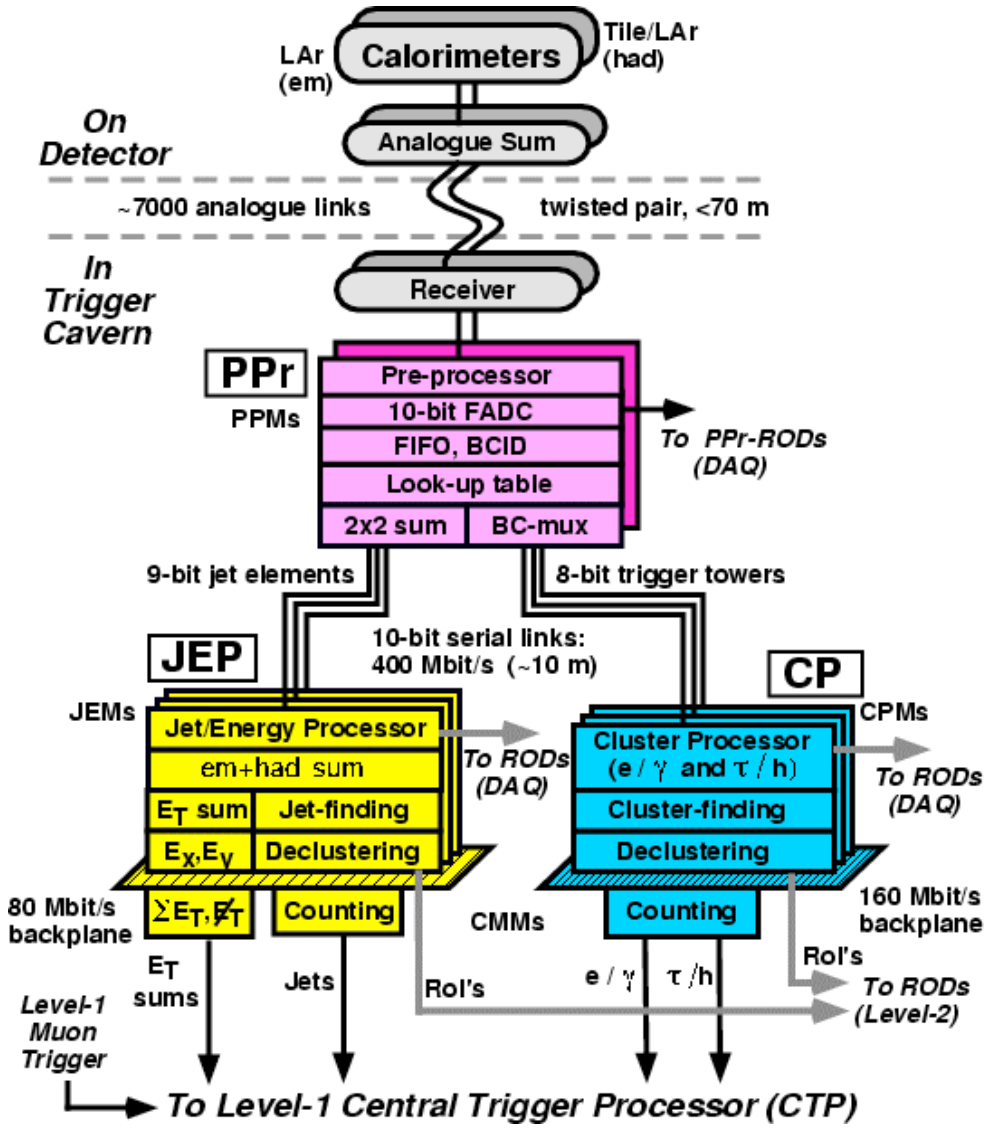

Bunch-crossing identification

- 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 energy was deposited
- Apply Finite Impulse Response filter
  - Result → LUT to convert value to E_T
  - Result → 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

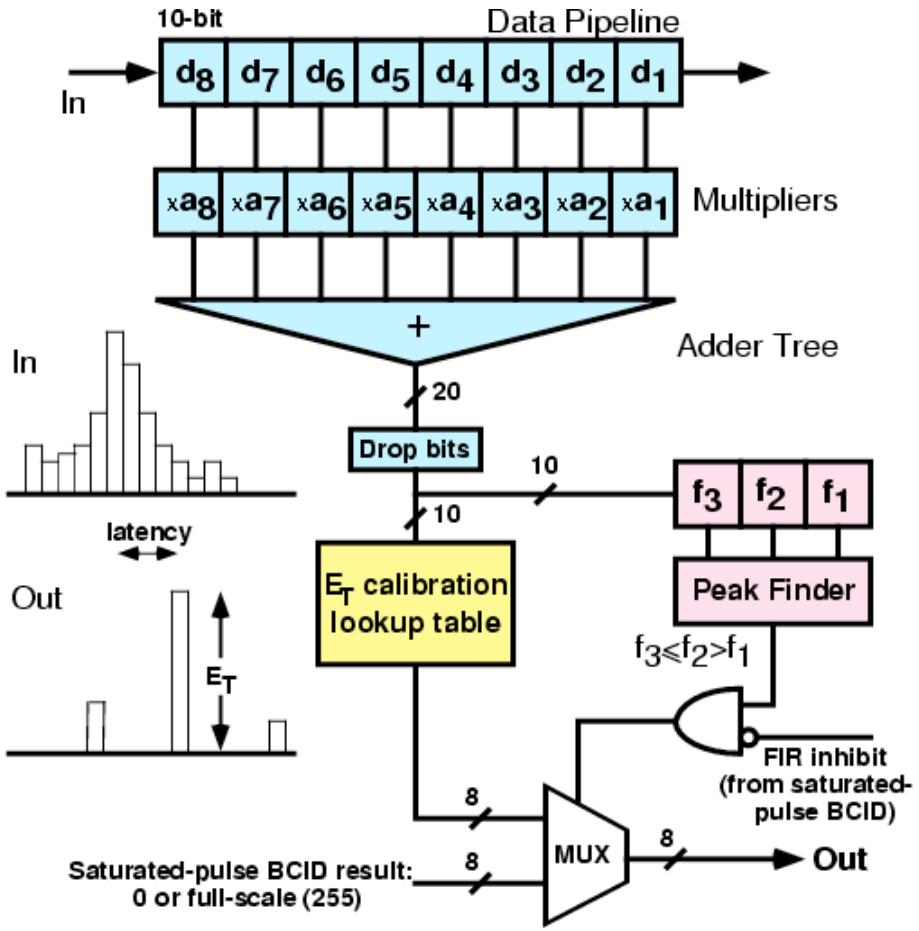

Data transmission and Cluster Processor
(numbers for ATLAS)

- The array of E_T values computed in the previous stage has to be transmitted to the CP
  - Use digital electrical links to Cluster Processor modules
    - ~5000 links @ 400 Mbps
  - 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 † per module
    - On-board fan out is comparatively straightforward

- 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 (η×ϕ×{E/H})
  - Algorithm is described in a programming language that can be converted into 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

† FPGA = Field Programmable Gate Array
i.e. reprogrammable logic

HLT electron trigger

- LVL1 e/γ trigger is already very selective
  - Need to use complex algorithms and full-granularity detector data in HLT
- Calorimeter selection
  - Sharpen E_T cut
  - Use shower-shape variables
    - Laterally and in depth
- Associated track in inner detector
  - Matching calorimeter cluster
    - Energy – momentum consistent
- Much work to develop algorithms and tune their many parameters to optimize signal efficiency and background rejection
  - Efficiency depends on signal definition
- HLT is implemented in software, running on farms of PCs
  - Almost full flexibility within the constraints of the available computing resources
  - Available time per event is 10s of milliseconds to a few seconds (second and third

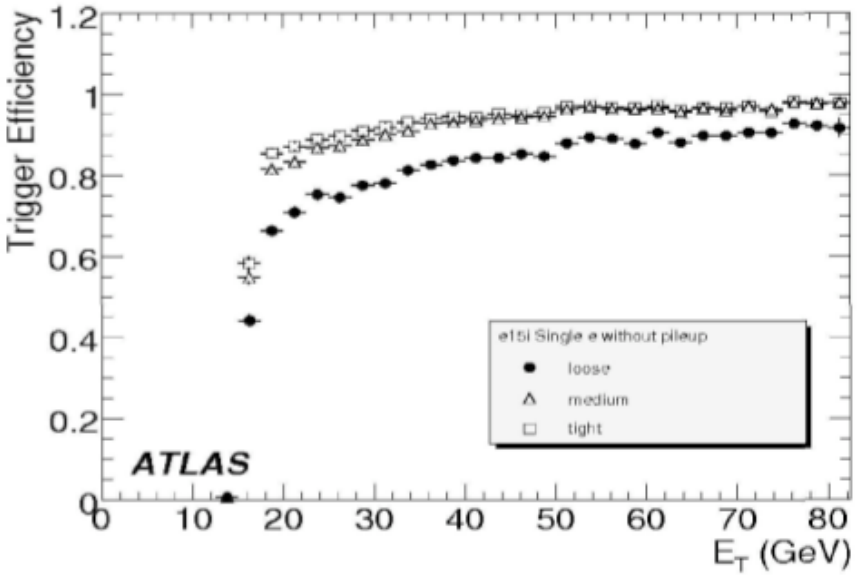

Commissioning of the T/DAQ systems during start-up of LHC in 2008

- Some history
  - 10 September 2008, first beam in the LHC (1 bunch at a time, 450 GeV)
    - Beam on collimators – "beam splash" events
    - Beam circulating for a few turns
    - Beam circulating for tens of minutes
    - No collisions (just single beam), no acceleration (injection energy)
    - Lots of media attention!

  - 19 September 2008, serious incident required shut-down of LHC
    - Will restart commissioning the T/DAQ systems with beam this autumn
      - After repairs and improvements to the LHC machine
    - In meantime, much experience gained running experiments with cosmic rays

Commissioning of the T/DAQ systems

- First objectives (plans)
  - Establish a stable time reference
    - Trigger on incoming beam (beam pick-ups, see later)
  - Time-in the experiment
    - Adjust programmable delays to read out the correct BC over the full detector
      - Note that Time-of-Flight corrections are different for outgoing collision products, downward-going cosmic-ray muons, and beam-halo (see later)
    - Adjust programmable delays to align all other LVL1 triggers to the reference
      - Minimum-bias trigger
      - Calorimeter triggers (including e/γ)
      - Muon triggers
  - Provide minimum-bias trigger (for single beam / collisions)
    - Trigger on activity in detector as well as (in time coincidence with) beam
  - Provide more selective LVL1 triggers, then progressively add HLT

Some examples from ATLAS and CMS

- Splash event and associated distribution
  - These events had a huge amount of activity in the detectors and fired many triggers
- Beam pick-up signals
  - See beam signals correlated with timing signals provided by machine
- Beam-halo event and associated distribution
  - Small number of beam-halo muons in each event, see correlations between different detector subsystems
- Illustration of progress in tuning trigger timing
  - Example of work done in all experiments
- Muon time of flight (consistent with speed of light)

Beam-splash event in ATLAS

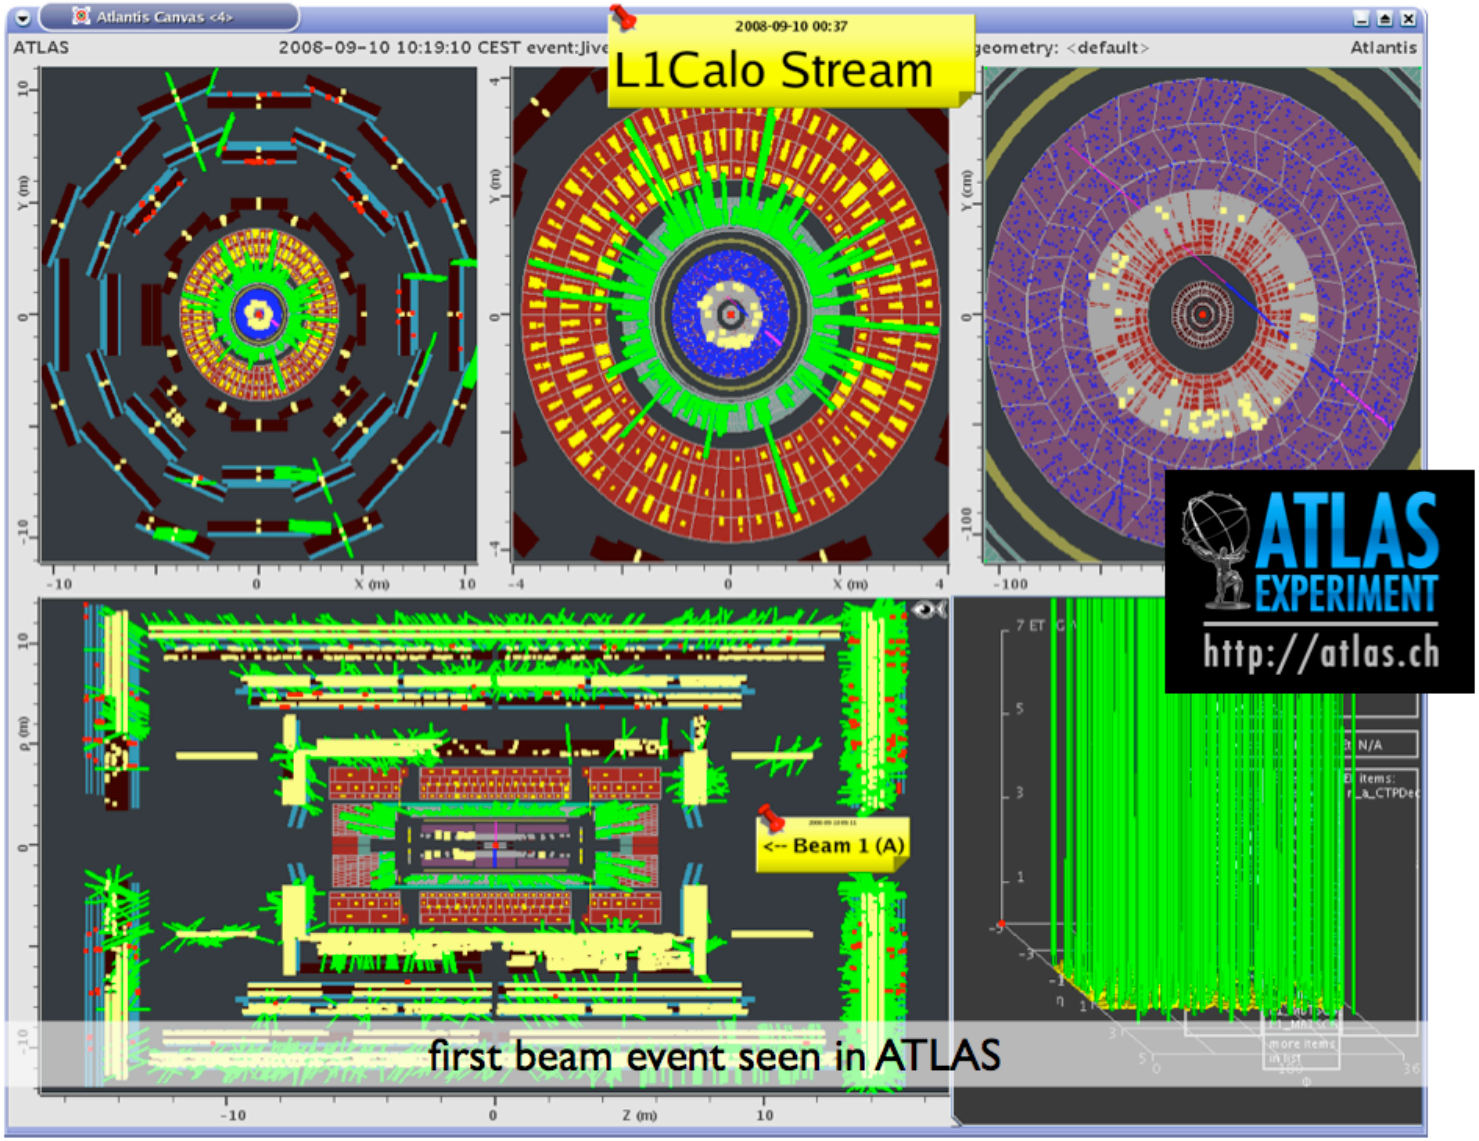

Study E_T and time versus η, ϕ

- Beam-splash events very useful
  - Identify (very few) problem channels, e.g. bad time calibration
- Note that ToF and 8-fold symmetry of ATLAS detector in ϕ can be seen in this plot of LVL1 calorimeter trigger E_T versus η, ϕ

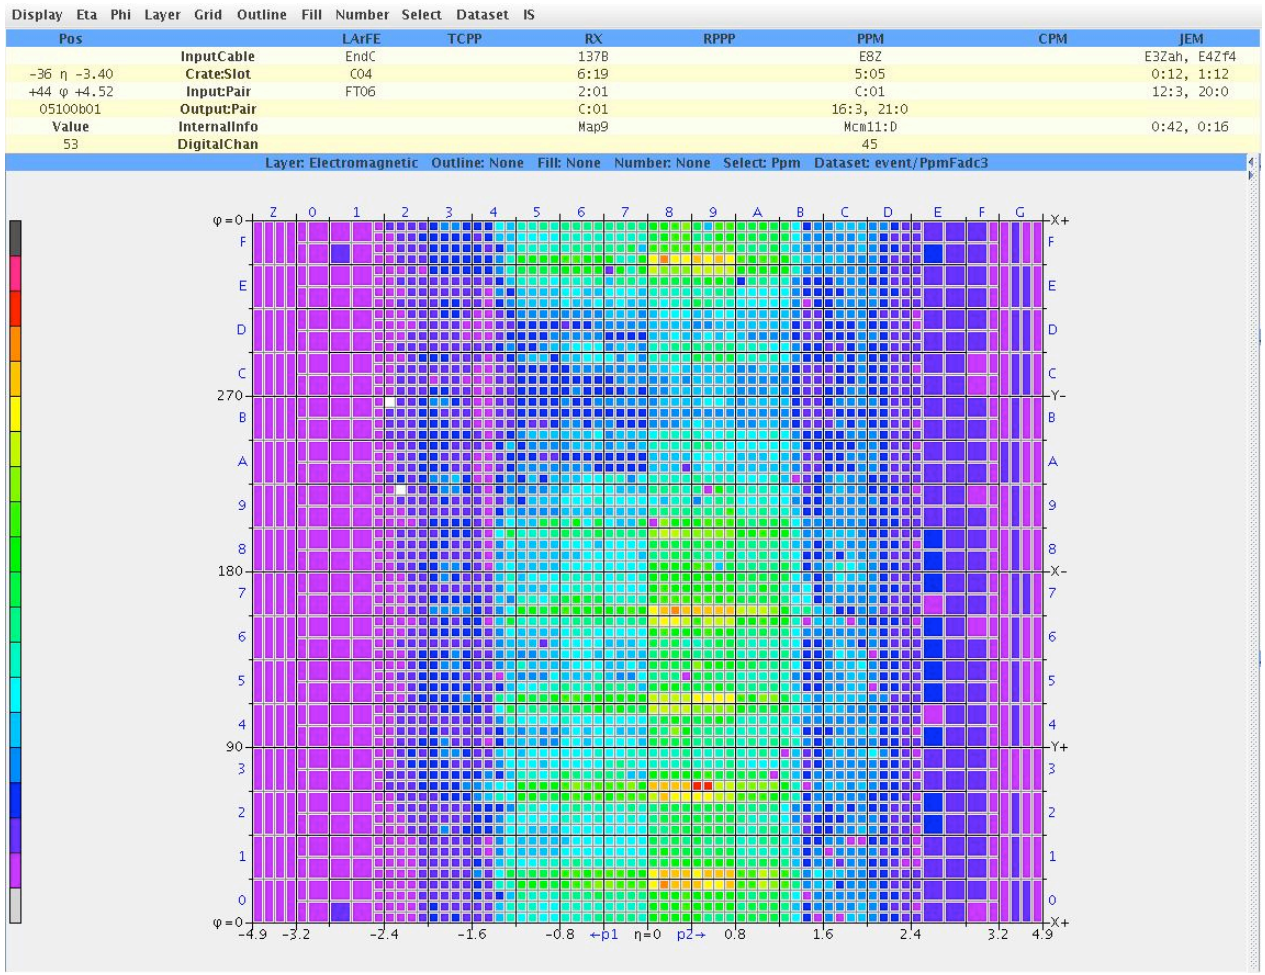

Use beam pick-ups as time reference

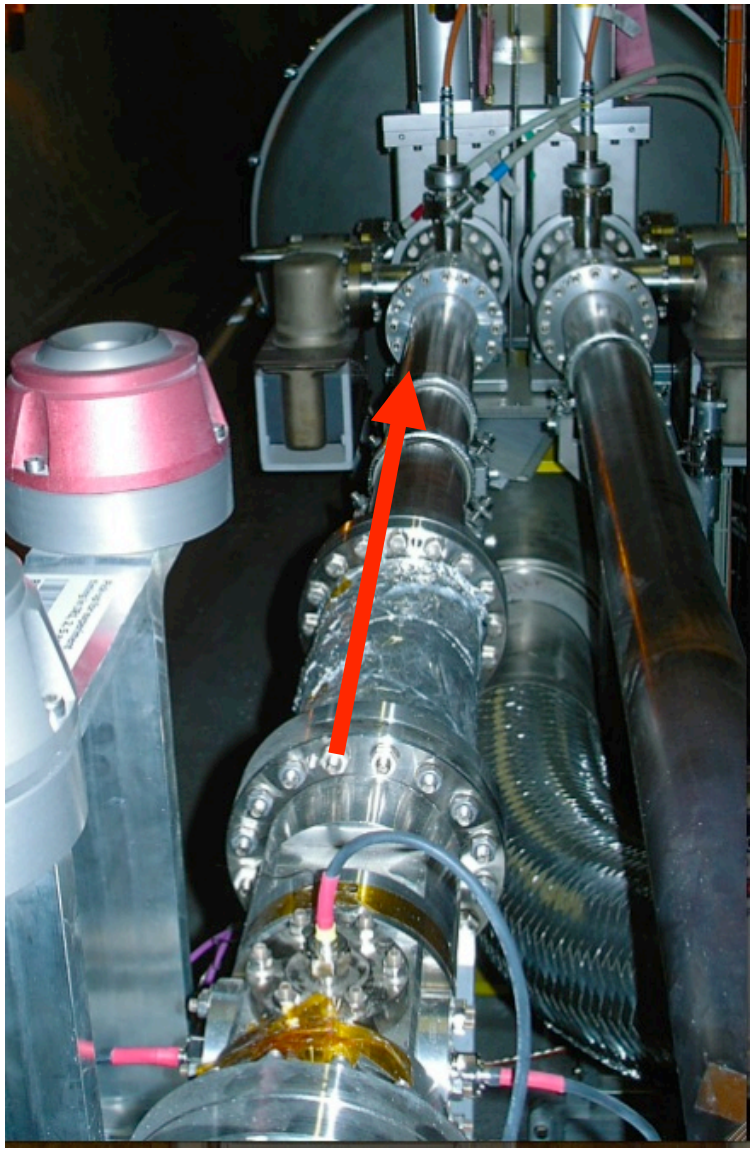

Plot from CMS

A beam-halo event in CMS (muon seen in CSCs and in HCAL)

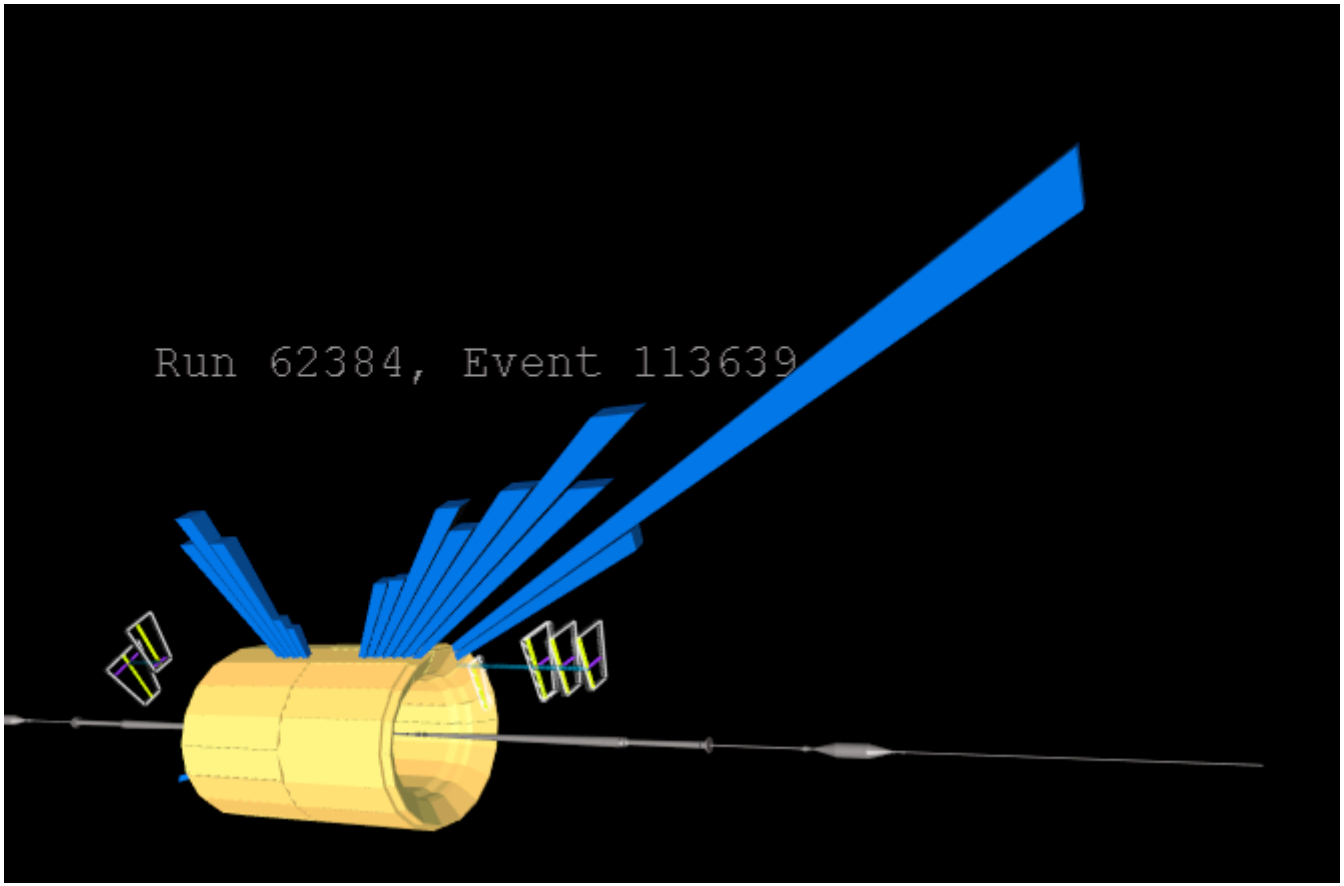

CMS – beam halo and cosmic muons

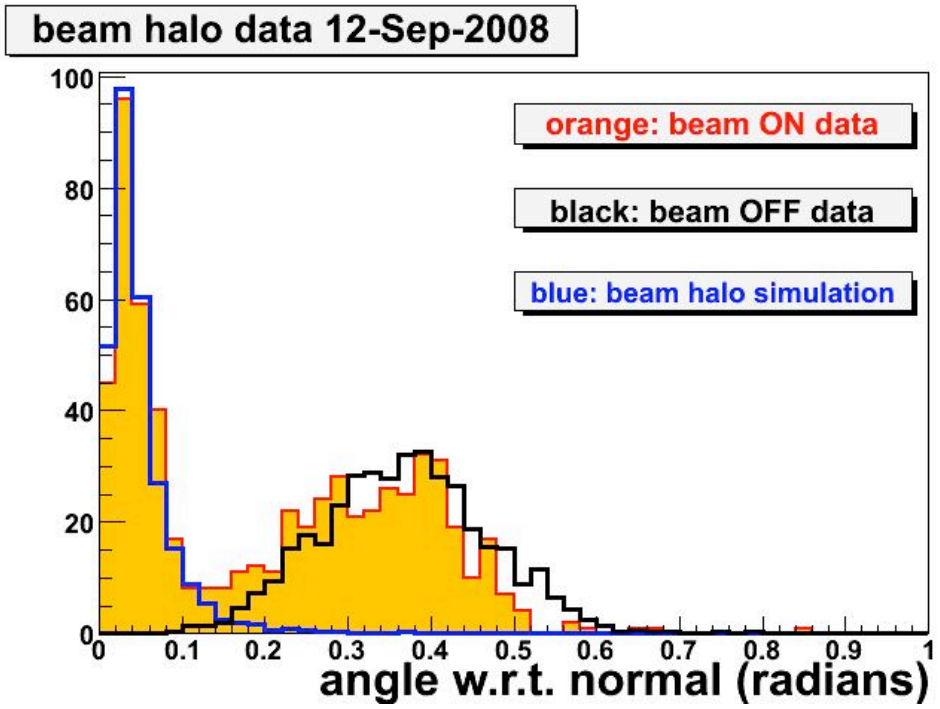

Tuning the timing of ATLAS trigger (progress between 10 and 12 Sep 2008)

- Note change in horizontal scale
- Note logarithmic vertical scale

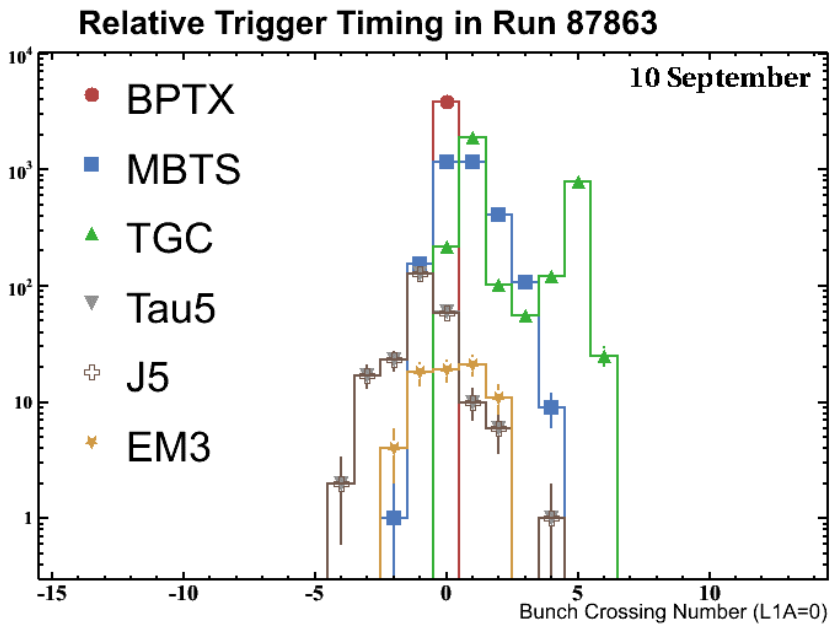

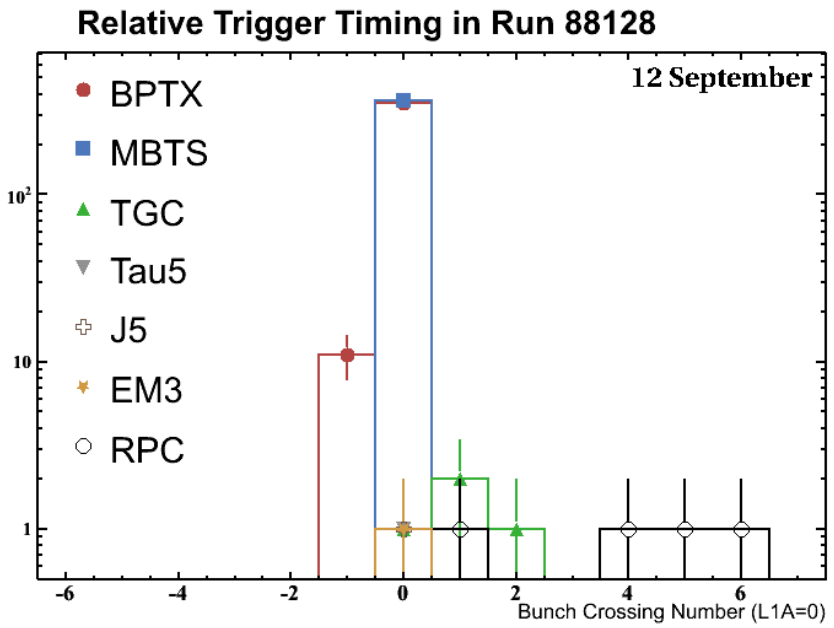

See Time-of-Flight of beam-halo muons (~100 ns to traverse ATLAS)

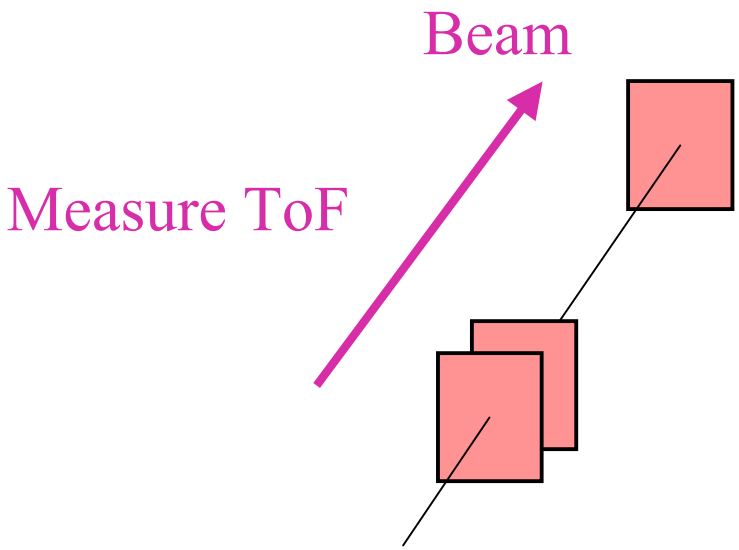

Distance between muon TGC detectors at opposite ends of ATLAS is 28 m (92 ns) at speed of light

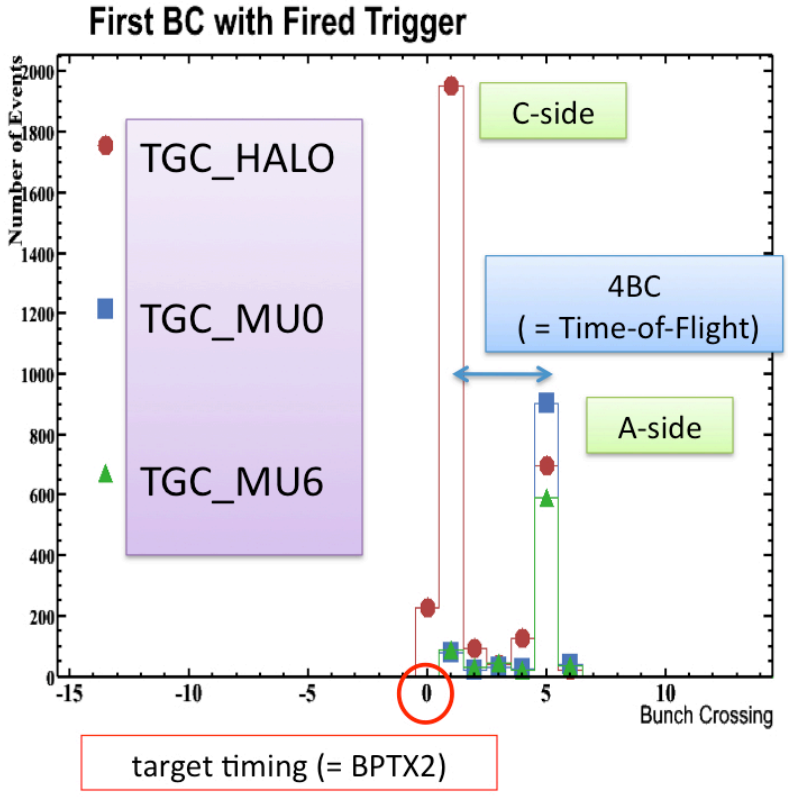

Final Remarks

- I hope I have succeeded to give you some insight into the challenges of building T/DAQ systems for HEP experiments
  - Challenges in physics (inventing algorithms that are fast, efficient for the physics that we want to do and give good rate reduction), and challenges in electronics and computing
- Also how the subject has evolved to meet the increasing demands, e.g. LHC compared to LEP
  - New ideas exploiting new technologies
- Finally, I hope that more of you will participate actively in this exciting field in the years to come!

Spares

Triggers at LEP

- The triggers at LEP aimed to select any e^+e^- annihilation event giving a visible final state
  - Including events with little visible energy
  - Plus some fraction of two-photon events
  - Plus Bhabha scattering events
- Furthermore, they aimed to select most events by multiple, independent signatures
  - Maximize efficiency
    - Probability to pass trigger A OR trigger B $\sim 1 - \delta_A\delta_B$
      - δ_X = inefficiency of trigger X, assuming losses uncorrelated
  - Allow measurement of the trigger efficiency from the data
    - Use events selected by trigger A to measure the efficiency for trigger B, and vice versa

PHYSICS REQUIREMENTS — TWO COMPLEMENTARY EXAMPLES

- LEP
  - Precision physics was main emphasis
  - Absolute cross-section determination was critical
    - E.g. determination of the number of neutrino species
- LHC
  - Discovery physics will be main emphasis
    - Vast range of predicted processes with diverse signatures
      - Very low signal rates expected in some cases
    - Should also be sensitive to new physics that has not been predicted!
  - Huge rate of Standard Model physics backgrounds
    - Rate of proton–proton collisions ~ 10^9 Hz (σ = 100 mb, L = 10^{34}cm^{-2}s^{-1})

LEP requirements (1)

- Trigger had to:
  - Identify all events coming from e^+e^- annihilations with visible final states
    - Including at LEP-I: $Z \rightarrow$ hadrons, $Z \rightarrow e^+e^-$, $Z \rightarrow \mu^+\mu^-$, $Z \rightarrow \tau^+\tau^-$
    - Including at LEP-II: WW, ZZ, single-boson
    - Including cases where there is little visible energy
      - e.g. in Standard Model: $e^+e^- \rightarrow Z\gamma \rightarrow \nu\nu\gamma$
      - e.g. in new particle searches such as $e^+e^- \rightarrow \chi^+\chi^-$ (with small $\chi^+ - \chi^0$ mass difference), giving only low energy visible particles (χ^0 LSP)
  - Retain some fraction of two-photon collision events
    - Used for QCD studies
  - Identify Bhabha scatters
    - Needed for precise luminosity determination

LEP requirements (2)

- Could retain events with any significant activity in the detector
  - Even when running at Z peak, rate of Z decays was only *O*(1 Hz)
    - Physics rate was not an issue
- Challenge was in maximising efficiency/acceptance of trigger
  - And also, making sure that the efficiency and acceptance could be determined with very high precision
    - Absolute cross-section determination depends on knowing:
      - Integrated luminosity (efficiency to trigger on Bhabha events)
      - Experimental efficiency and acceptance for process in question (efficiency to trigger on physics process)
        - » Events selected by many redundant triggers (high efficiency; cross-checks)
    - A major achievement at LEP was to reach per-mil precision
- The trigger rates and also the DAQ data rates were modest

Selection criteria at LEP

- Details depend on the exp^{t}, e.g. ALEPH "menu" was as follows:
  - Triggers implemented within segments (60 regions in θ, ϕ)
    - Single muon trigger
      - Track seen penetrating the hadron calorimeter & seen in the Inner Tracker
    - Single charged EM energy trigger
      - EM calorimeter cluster and track in Inner Tracker
    - Single neutral EM energy trigger
      - EM calorimeter cluster
        - » Higher threshold than above to keep the rate down
  - Total-energy triggers
    - Threshold on energies summed over large regions (barrel or a full endcap)
  - Back-to-back tracks trigger
  - Bhabha luminosity-monitor triggers

Trigger implementation at LEP

- In general, LVL1 triggers were implemented using a combination of analogue and digital electronics
- Details depend on the experiment, e.g. ALEPH was as follows:
  - Calorimeter triggers were implemented using analogue electronics to sum signals, applying thresholds on the sums
  - LVL1 tracking trigger looked for patterns of hits in the Inner Tracking Chamber (ITC) consistent with a track with $p_T > 1$ GeV/c
    - At LVL2, information from the TPC was used instead
  - Final decision was made by combining digital information from the calorimeter and tracking triggers
    - Local combinations within segments of the detector
    - Global combination (logical OR of conditions)

HLT/DAQ Software

- A major challenge lies in the HLT/DAQ software
  - Algorithms
    - HLT can be subdivided into LVL2 and LVL3
      - Separate processor systems (e.g. ATLAS) or two distinct processing steps in the same processor system (e.g. CMS)
  - Framework that manages the flow of data and supports the algorithms
    - Supervising an event from when it arrives at the HLT/DAQ system until it is rejected or accepted and recorded to permanent storage
    - Transporting data to the algorithms as required
  - Large amount of associated online software
    - Run control
    - Databases (description of hardware & software, calibration constants, etc.)
    - Book-keeping (run conditions, log of errors, etc)
    - Etc

DAQ at LEP (e.g. ALEPH)

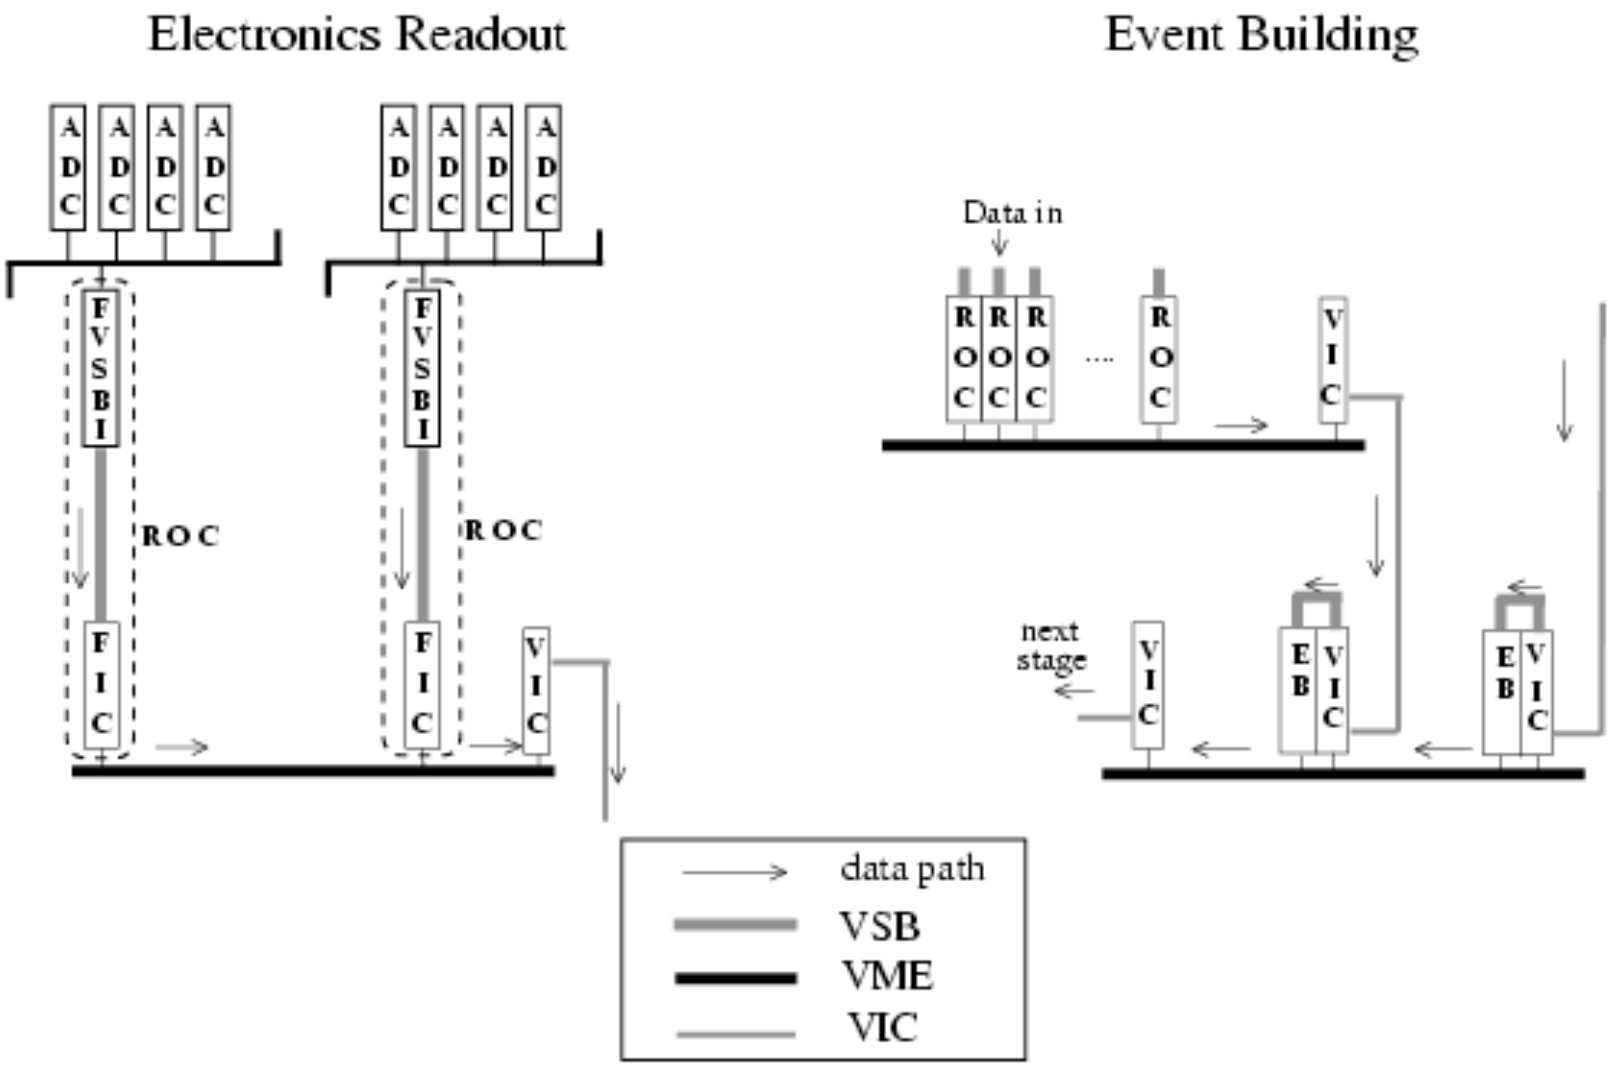

Minimum-Bias Trigger Scintillators

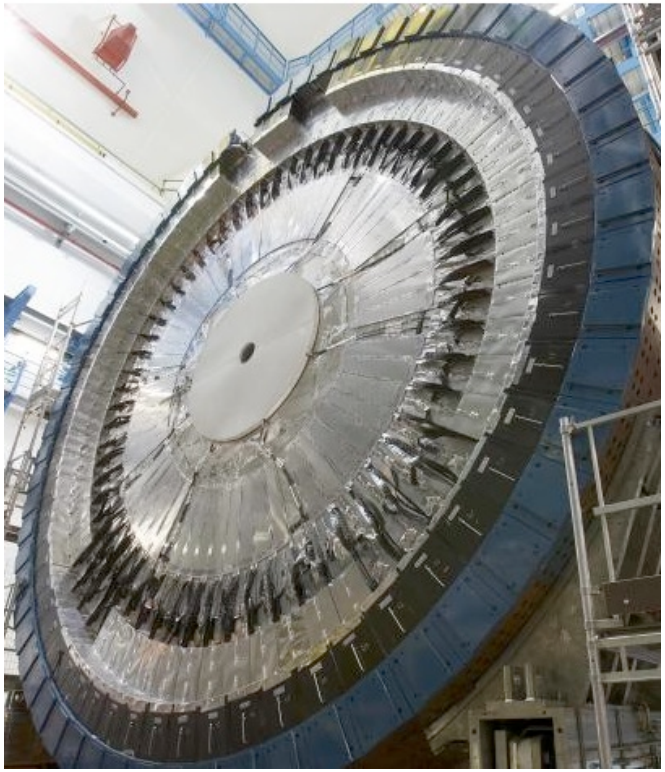

DAQ at LEP (e.g. ALEPH)

- Following a LVL2 trigger, events were read out as follows:
  - Data were transferred within each crate to the readout controller (ROC)
    - After this, further LVL1 and LVL2 triggers could be accepted
  - Events were "built" in two stages
    - Event
      - Sub-event × n
        - » ROC data block × m
- Once events were in the main readout computer, the LVL3 trigger made a final selection before data were recorded

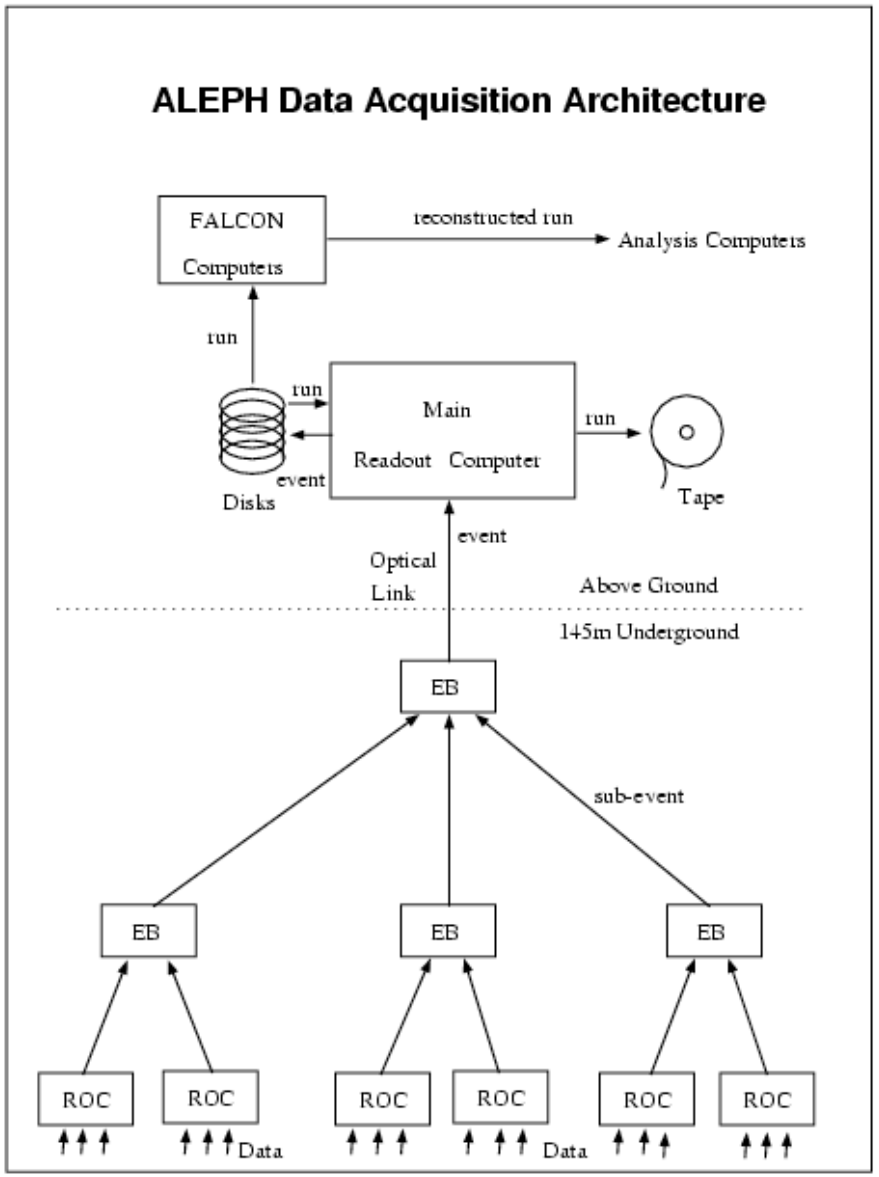

DAQ at LEP (e.g. ALEPH)

- The ALEPH DAQ used a hierarchy of computers
  - Local readout controllers (ROCs) in each crate
    - In addition to reading out the data from ADCs, TDCs, etc., these performed processing (e.g. applying calibration to convert raw ADC values to energies)
      - Zero suppression was already performed at the level of the digitizers where appropriate
  - Event builders (EBs) for sub-events
    - These combined data read out from the ROCs of a given sub-detector into a sub-event
  - Main event builder
    - Combined data from the EBs for the different detectors
  - Main readout computer
    - Received full events from main EB; performed LVL3 trigger selection; recorded selected events for subsequent offline analysis

Size of detectors and the speed of light

$p_T \equiv$ transverse momentum $\perp$ beams

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

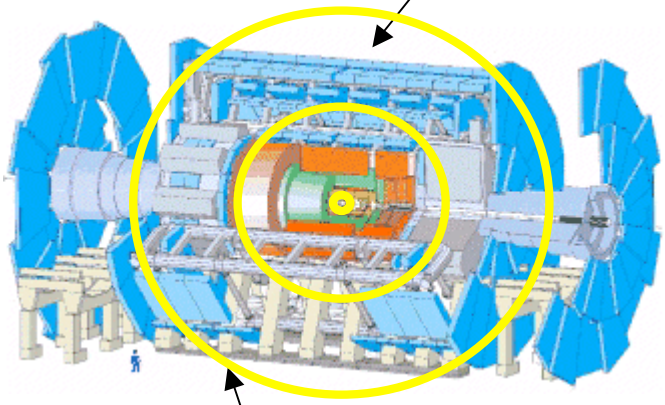

Need to read out also here

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 m × 3.3 ns/m = 73 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 given that the readout pipelines are mounted on the detector