

First-level trigger systems at LHC

N. Ellis

CERN, 1211 Geneva 23, Switzerland
Nick.Ellis@cern.ch

Abstract

Some of the challenges of first-level trigger systems in the LHC experiments are discussed. The first part of the paper gives an analysis of the requirements from the physics and technical points of view, including the need to make decisions with a latency of a few microseconds compatible with the length of pipeline memories in the front-end electronics. This is followed by a general discussion of the techniques that are used in designing first-level trigger systems and an overview of the first-level trigger systems of the LHC experiments. In the final part of the paper, some examples from the LHC experiments are used to illustrate how the triggers are realized in practice.

I. REQUIREMENTS

The role of the trigger is to make the online selection of particle collisions potentially containing interesting physics. Triggers must have high efficiency for selecting processes of interest for physics analysis, the efficiency should be precisely known, and the selection should not have biases that affect physics results. At the same time, a large reduction of rate from unwanted high-rate processes is required (capabilities of DAQ and also offline computers). Unwanted processes include instrumental backgrounds and high-rate physics processes that are not relevant for the analysis. Clearly, the trigger system must be also affordable which limits the complexity of the algorithms that can be used. It is not easy to achieve all the above simultaneously!

The LHC experiments [1–4] use multi-level triggers that provide rapid rejection of high-rate backgrounds without incurring (much) dead-time, and high overall rejection power to reduce the output to mass storage to an affordable rate. The fast first-level trigger (custom electronics) needs to have high efficiency, but its 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 the detector) pending the result. High overall rejection power is achieved by progressively reducing the rate after each stage of selection; this allows the use of more and more complex algorithms at affordable cost. The final stages of selection, running on computer farms, can use comparatively very complex (and hence slow) algorithms to achieve the required overall rejection power because they only see a greatly reduced input rate thanks to the earlier stages of selection. The multi-level trigger of ATLAS is shown in Fig. 1 as an example.

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. In this paper the term “event” is used to mean the record of the activity in a given bunch crossing (BC) — typically an event contains many proton–proton interactions. The 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 also 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!

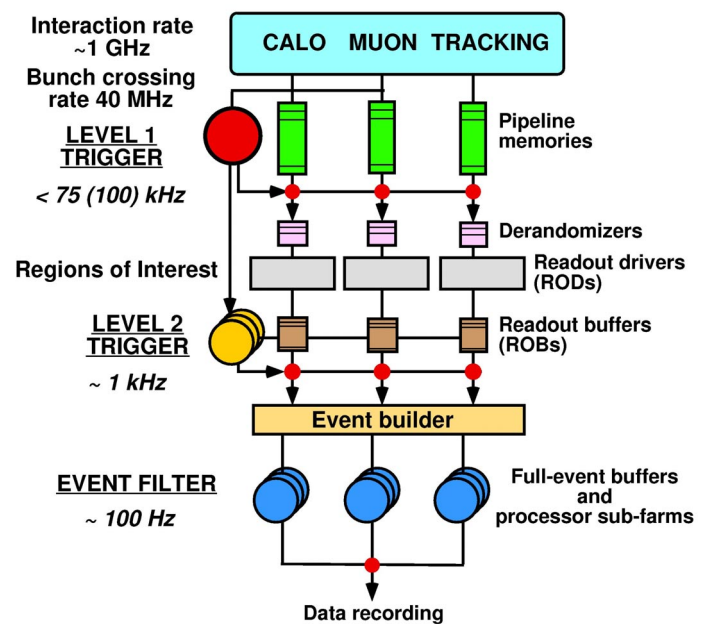


Figure 1: The ATLAS multi-level trigger

A. LHC physics

Discovery physics is the main emphasis for ATLAS and CMS. There is a vast range of new physics processes predicted in different theoretical models, with diverse signatures, and with very low signal rates expected in some cases. However, one should also try to be sensitive to new physics that has not been predicted — only experiment can tell us about the physics of the real world which may not coincide with the conjectures of our theoretical colleagues!

As can be seen in Fig. 2, there is a huge rate of Standard Model physics “backgrounds” to the new physics — the rate of proton–proton collisions is up to 10^9 Hz. (Substantially lower rates are anticipated for instrumental backgrounds such as beam–gas interactions.)




Figure 2: Rates of various physics processes at LHC

The triggers in ATLAS and CMS 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), $H \rightarrow \gamma\gamma$, and $H \rightarrow t\bar{t}$, $H \rightarrow b\bar{b}$.
- SUSY searches, e.g. producing jets and missing transverse energy (missing E_T).
- Searches for other new physics using inclusive triggers that it is hoped will be sensitive to unpredicted new physics.
- Studies of Standard Model processes that are of interest in their own right, and must be understood as backgrounds to new physics, e.g. W and Z boson, top and beauty quark production.

In contrast to the particles produced in typical proton–proton collisions (typical hadron $p_T \sim 1$ GeV), the products of new physics are expected to have large transverse momentum, p_T . For example, particles produced in the (two-body) decay of a new heavy object such as the Higgs boson have a typical p_T close to half of the mass; e.g. for $H \rightarrow \gamma\gamma$ with $m_H \sim 120$ GeV $\Rightarrow p_T \sim 60$ GeV for each photon.

Typical examples of first-level trigger thresholds for ATLAS (CMS similar) at LHC design luminosity are:

- Single muon $p_T > 20$ GeV (rate ~ 10 kHz) or a pair of muons each with $p_T > 6$ GeV (rate ~ 1 kHz);
- single e/γ $p_T > 30$ GeV (rate ~ 20 kHz) or a pair of e/γ each with $p_T > 20$ GeV (rate ~ 5 kHz);
- single jet $p_T > 300$ GeV (rate ~ 200 Hz), or compound triggers such as “jet $p_T > 100$ GeV and missing- $p_T > 100$ GeV” (rate ~ 500 Hz), or “four or more jets $p_T > 100$ GeV” (rate ~ 200 Hz).

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 ($\sim 2 \times 10^{32}$ $\text{cm}^{-2}\text{s}^{-1}$), giving an overall proton–proton 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 trigger will search for muons, electrons/photons and hadrons with $p_T > 1$ GeV, 2.5 GeV and 3.4 GeV respectively. The level-0 output rate is 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. The aim is to record an event rate of only ~ 200 Hz.

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 proton–proton experiments: $L \sim 10^{27}$ $\text{cm}^{-2}\text{s}^{-1} \Rightarrow$ Rate ~ 8000 Hz for Pb–Pb collisions (higher luminosity and rates are expected 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 (rate ~ 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, with up to $O(10,000)$ charged particles in the central region. An event size of up to ~ 40 MByte is anticipated when the full detector is read out.

As discussed above, the first-level trigger must identify events containing electrons (e), photons (γ), muons (μ), jets and large missing- p_T , e.g. due to neutrinos (ν). Fig. 3 illustrates how these signatures appear in the detector systems. Electrons and photons both give rise to localized clusters in the electromagnetic calorimeter (ECAL). The electrically charged electron gives a track in the inner detector (IDET), while the neutral photon does not. (As discussed below, the first-level triggers in ATLAS and CMS do not make any use of the IDET, so at this level electrons and photons cannot be distinguished.)

A jet, consisting of many particles, gives rise to an extended cluster with energy in both the ECAL and the hadronic calorimeter (HCAL). Muons give a track in the IDET, little energy deposition in the ECAL and HCAL, and a track in the external muon system (MuDET) – muons are the only charged particles to penetrate the calorimeters. The electrically neutral and weakly-interacting neutrino does not

interact in the detector system, but its presence may be inferred from transverse energy imbalance. The visible energy recoiling against a high- p_T neutrino appears to be unbalanced giving missing E_T .

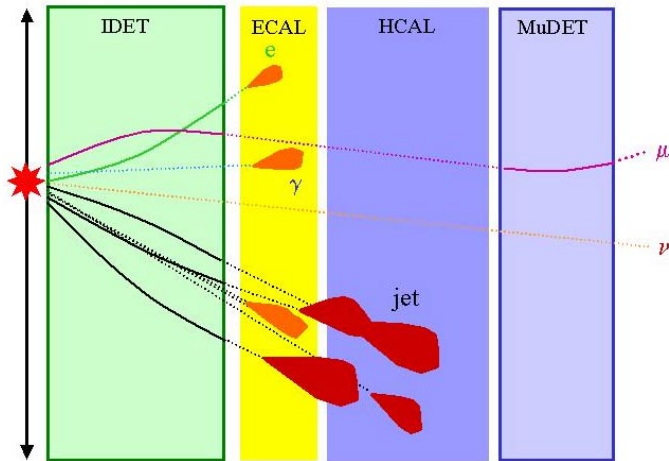


Figure 3: Signatures in a generic detector system

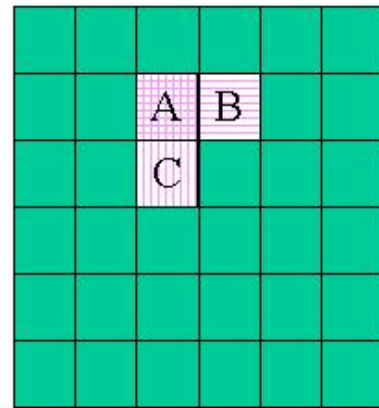
II. FIRST-LEVEL TRIGGER OVERVIEW

The first-level trigger typically must search for muons, electrons and photons, tau leptons and isolated hadrons, jets, and large missing or total transverse energy. It must form a trigger decision for each BC based on combinations of the above signatures. In ATLAS and CMS, the first-level trigger is based only on the muon detectors and the calorimeters — the inner tracking is used only in the higher-level triggers. Although a new trigger decision is issued every 25 ns (= BC period), the trigger latency is a few microseconds. This latency figure includes the time taken for signals to reach the trigger electronics and the time needed to distribute the decision to the detector front-end electronics using the TTC system [5]. Given the length of the signal path, typically ~ 100 m in each direction, a significant fraction of the latency is due to propagation delays on electrical or optical cables — taking a typical speed of ~ 0.2 m/ns, the round trip propagation delay is ~ 1 μ s. Until the trigger decision is received, the information from all detector channels must be retained in the front-end electronics.

The first-level trigger introduces a small amount of deadtime to avoid data loss or buffer overflow in the front-end systems, and to avoid the need to read out overlapping time-frames in the majority of detector systems.

Since the first-level trigger has to deliver a new decision every BC, while the trigger latency is much longer than the BC period, the trigger must concurrently process many events. This can be achieved by “pipelining” the processing in custom trigger processors built using modern digital electronics. This means breaking the 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.

The principle of pipelined processing is illustrated in Figs. 4 and 5. The simple algorithm shown in Fig. 4, which determines if the energy sum in a horizontal or vertical pair of trigger towers exceeds a programmable threshold, can be implemented using the logic shown in Fig. 5.



EM Calorimeter
(~ 3500 trigger towers)

Figure 4: Simplified trigger algorithm

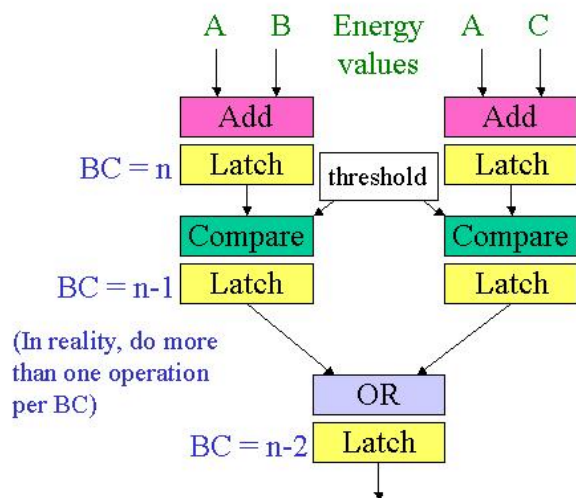


Figure 5: Illustration of pipelined processing

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.

A. Data-processing technologies

Field Programmable Gate Arrays (FPGAs) and other programmable devices now play a very important role in first-level trigger systems. Modern devices have a large gate count and many I/O pins; they operate at 40 MHz and above and their performance is sufficient for implementing many trigger algorithms. An important advantage of FPGAs is that they offer huge flexibility, including the possibility to modify algorithms as well as parameters of algorithms once the

experiments start running. Application-Specific Integrated Circuits (ASICs) are nevertheless used for some applications. They are more cost effective in some cases (e.g. where a large number of devices is needed), and they can offer higher speed performance than FPGAs. ASICs can also have better radiation tolerance and lower power consumption for on-detector applications.

B. Data-movement technologies

Data movement is an extremely important issue in the design of first-level trigger systems. High-speed serial links (electrical and optical) are used over a variety of distances, from a few meters to ~100 m. Comparatively inexpensive and low-power LVDS links are typically used for electrical transmission at ~400 Mbit/s over distances up to ~10 m. High-performance products such as HP G-link and Vitesse chipsets are used for Gbit/s transmission, with electrical transmission for short distances and optical transmission for longer distances. High density custom backplanes are needed for data exchange between modules within a crate. Some applications have very high pin counts (up to ~800 per 9U board), with data rates per (point-to-point) connection of ~160 Mbit/s. Data are often multiplexed beyond the LHC clock rate of 40 Mbit/s to reduce the connectivity problem to a level that can be managed. Data-movement considerations favour large (9U) boards — it is easier to handle interconnections on a board than between boards, and larger boards also have more space for connectors.

It is interesting to consider the flow of data in the first-level trigger that is illustrated in Fig. 6. The processing starts from a large amount of raw information, energies in calorimeter trigger towers or hits in the muon chambers. Because of the need for “environment information”, the raw information usually has to be fanned out between processing elements — the amount of data is therefore expanded. However, the subsequent processing then progressively reduces the information, finally to a single bit that flags if the event is to be retained or not. Of course, additional information is retained for the events that are selected by the first-level trigger (to guide the processing in the higher-level triggers and for monitoring).

LVL1 data flow

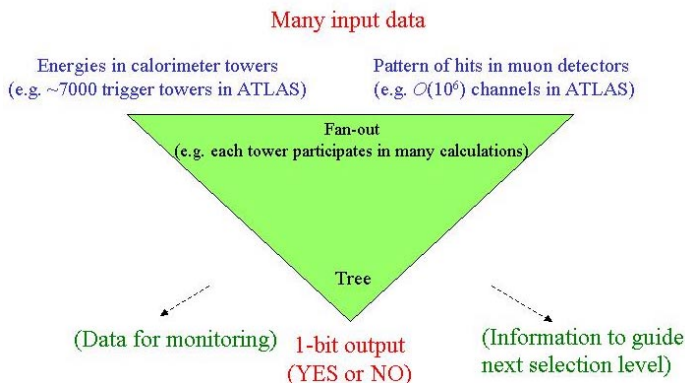


Figure 6: Dataflow in the first-level trigger

III. OVERVIEW OF LHC TRIGGER SYSTEMS

A. ATLAS

An overview of the ATLAS first-level trigger [1] is shown in Fig. 7. For the calorimeter trigger, discussed in more detail later in this paper, calorimeter cells in the ECAL and HCAL are combined into trigger towers by analogue summation of the corresponding signals; there are about 3500 such towers in each of the ECAL and the HCAL. For the muon trigger the raw data are the pattern of hits in $O(10^6)$ channels of the muon trigger chambers.

All of the calorimeter trigger system after the analogue-summation stage is off the detector — twisted-pair cables are used to transmit the analogue signals to an underground counting room that is shielded against radiation. Here people will be able to work even when there is beam in the LHC.

In the case of the muon trigger, the first stage of the trigger processing is mounted on the detector with implications for radiation tolerance and other factors (cooling, grounding, operation in magnetic field, very restricted access).

The results of the calorimeter and muon trigger processing, multiplicities of candidate objects for a variety of p_T thresholds, are passed to the Central Trigger Processor (CTP) that makes the final decision. The TTC system [5] is used to distribute the result to the detector front-end systems.

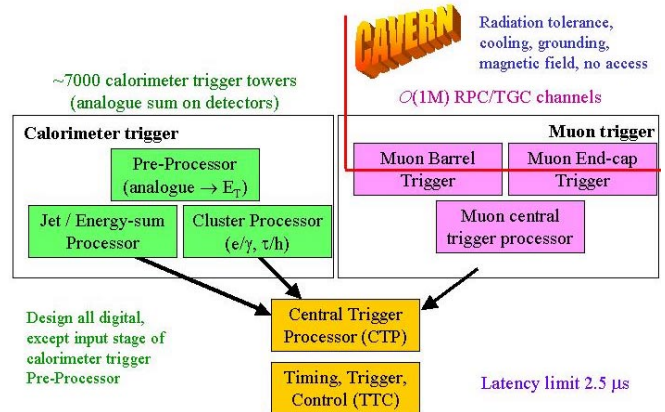


Figure 7: Overview of ATLAS first-level trigger

B. CMS

An overview of the CMS first-level trigger [2] is shown in Fig. 8. As discussed later in this paper, the trigger tower summation for the CMS calorimeter trigger is performed digitally using the same ADC system as for the precision readout. (In ATLAS, the full granularity Liquid-Argon (LAR) calorimeter information is retained in analogue memories pending the first-level trigger decision.) In the case of the CMS ECAL, the first part of the trigger processing, which includes BC identification of trigger-tower energies, is mounted on the detector [6].

The CMS muon trigger, discussed in more detail later in this paper, uses information from all of the muon detection

systems (RPCs, CSCs and DTs). As in ATLAS, the first part of the muon trigger processing is performed on the detector.

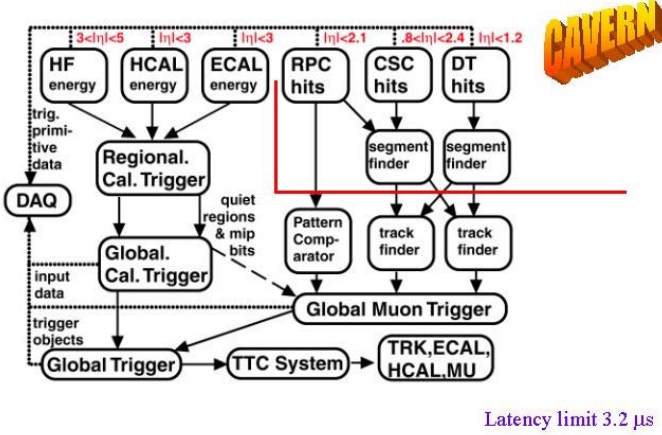


Figure 8: Overview of CMS first-level trigger

Whereas the ATLAS CTP works with multiplicity information from the calorimeter and muon trigger systems, the CMS global trigger, discussed in more detail later in this paper, retains detailed information on the individual candidate objects. This allows topological requirements to be made in the first-level trigger. (In ATLAS such requirements can be made in the second-level trigger based on the region-of-interest information provided by the first-level.)

C. LHCb

LHCb [3] uses two levels of buffering on the detector (c.f. one for ATLAS and CMS). The first-level trigger, called “level-0”, has an output rate of about 1 MHz and a latency of about 4 μ s, and is implemented using custom electronics. This is followed by a software-based “level-1 trigger” that reduces the rate further with a variable latency of \sim 2 ms.

The level-0 calorimeter trigger searches for high- p_T electrons/photons and hadrons; there is also a level-0 muon trigger. As discussed later in this paper, a pile-up veto rejects events with more than one proton–proton interaction vertex. As mentioned earlier, the p_T thresholds in LHCb are much lower than in ATLAS and CMS.

D. ALICE

In ALICE [4], trigger subsystems associated with many of the sub-detectors generate inputs to the first-level (“L0”) trigger. There are 24 L0 inputs (latency 900 ns; 2 μ s dead-time after each trigger). A very short latency is motivated by the fact that some of the detectors, which use track-and-hold rather than pipelined readout, need a prompt trigger signal. All of the trigger electronics is therefore mounted on the detector.

The first-level trigger is followed by L1 with 21 inputs (latency 6.2 μ s), and L2 with 6 inputs (latency 88 μ s \sim TPC drift time).

There is provision for the control of up to 24 independent sub-detectors, grouped into 6 detector clusters that are read out together.

In contrast to ATLAS, CMS and LHCb, ALICE does not always read all sub-detectors. Up to 50 trigger classes can be defined, specifying the following for each one: L0-L1-L2 patterns; pre-scale factor; and detector cluster for readout.

The use of slow detectors requires past–future protection logic. Different limits can be set for peripheral and semi-central interactions. Note that there are very different interaction rates in Pb–Pb, Ar–Ar and proton–proton running.

IV. CALORIMETER TRIGGERS

First-level calorimeter triggers are illustrated here by the example of the ATLAS electron/photon trigger. More information on the ATLAS trigger can be found in Ref. [1] with some recent developments described in a paper at this workshop [7]. Details of the CMS trigger can be found in Ref. [2], and in papers presented at this workshop [8], [6].

Level-1 Calorimeter Trigger Architecture

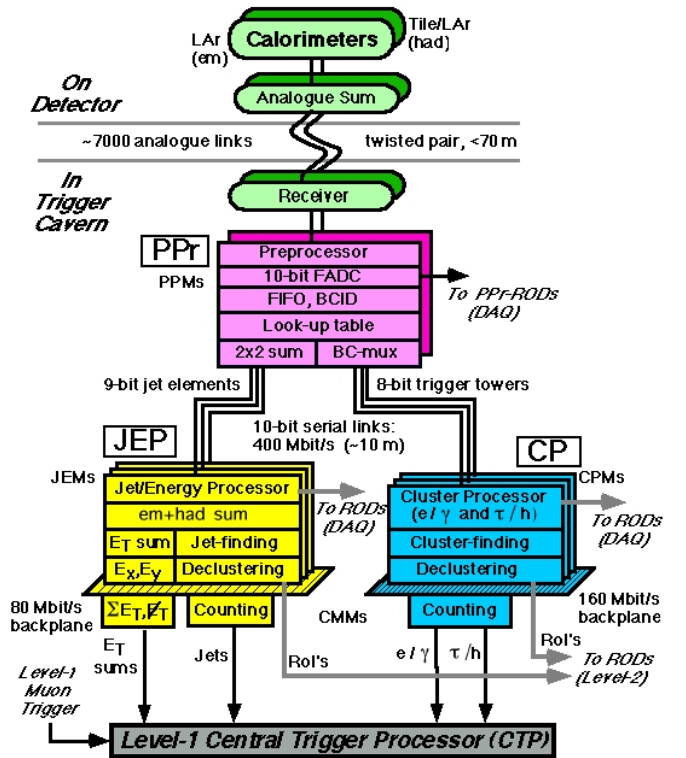


Figure 9: The ATLAS calorimeter trigger

The ATLAS calorimeter trigger system is illustrated in Fig. 9. Analogue electronics on the detector sums signals to form trigger towers. After transmission on twisted-pair cables of length \sim 70 m, the signals are received and digitised in the Pre-processor system. The digital data are then processed to measure E_T per tower for each BC, giving an E_T matrix for the ECAL and HCAL. Tower data are transmitted by the Pre-processor to the Cluster Processor (CP) that occupies four

crates, and to the Jet-Energy Processor (JEP) that occupies two crates. Values needed in more than one crate are fanned out by the Pre-processor system. The need to minimize the number of towers to be fanned out motivates a very compact design for the CP and JEP. Within the CP and 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.

A crucial aspect of the calorimeter trigger is the association of measured energies to the correct BCs. This is achieved using a combination of a Finite-Impulse Response filter and a peak finder, as shown in Fig. 10.

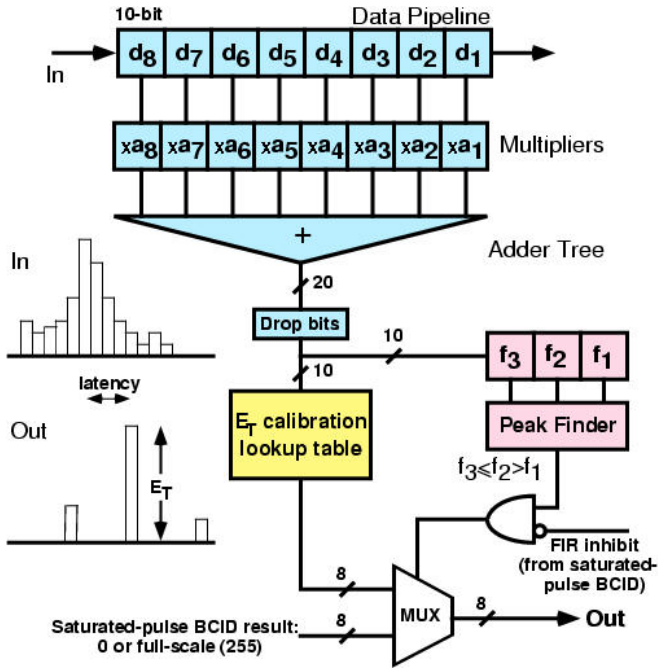


Figure 10: Pre-processing of calorimeter trigger-tower data in ATLAS

The calorimeter signals extend over many BCs, and one needs to combine information from a sequence of measurements to estimate the energy and identify the BC where the energy was deposited. This is achieved by using a Finite Impulse Response (FIR) filter, the result of which is fed to a look-up table to convert to E_T, and to a peak finder to determine the BC in which the energy was deposited. One needs to take care of signal distortion for very large pulses — clearly it is essential not to lose the most interesting physics!

A Pre-processor ASIC incorporates the above functions for four channels. It accepts 10-bit inputs from the ADCs at 40 MHz rate, applies calibration factors and converts to E_T units, and performs zero-suppression and BC identification. It also provides readout of the data for events that are selected by the first-level trigger.

The Pre-processor system is built using Multi-Chip Modules (MCMs) that each support four commercial 40 MHz ADCs, one Pre-processor ASIC and three LVDS drivers used to transmit the E_T data to the CP and JEP systems [9]. A photograph of a prototype Pre-processor MCM (currently

under evaluation) is shown in Fig. 11. From right to left one can identify the four ADC chips, the ASIC and three LVDS drivers.




Figure 11: Photograph of Pre-processor MCM prototype

Given the large number of channels to be equipped (~7000 in total), the use of an ASIC and MCM is a cost-effective solution that results in a very compact system.

Although the details of the pre-processing system described above are specific to ATLAS, similar functionality is implemented in CMS. As discussed at this workshop [6], the logic that prepares the trigger-tower data for the CMS ECAL will now be implemented on-detector. The CMS system digitises at 40 MHz with the full detector granularity. The trigger towers are formed by digital summation and using signal processing similar to that described above for ATLAS. The ATLAS and CMS digitisation schemes are illustrated in Fig. 12.

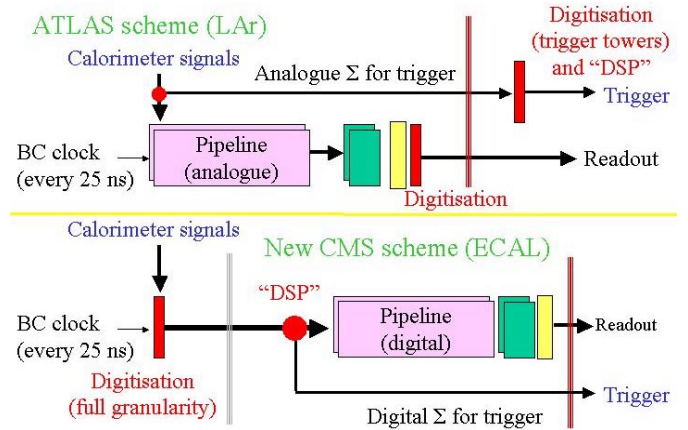


Figure 12: Digitisation in ATLAS and CMS

The ATLAS e/γ trigger (see Fig. 13) is based on 4×4 “overlapping, sliding windows” of trigger towers. Each trigger tower covers 0.1×0.1 in η×φ, where η is pseudo-rapidity and φ is azimuth. There are about 3500 such towers in each of the electromagnetic and hadronic calorimeters. Since the window can be centred on any tower, there are ~3500 such windows within which the algorithm must be performed for each BC. Also, each tower participates in the calculations for 16 windows, which is a driving factor in the trigger design. Note that the example of pipelined processing shown earlier in Fig. 5 actually implements part of this algorithm for a single window.

The array of E_T values computed in the Pre-processor has to be transmitted to the CP. This is done using digital electrical links to the CP modules, with about 5000 LVDS links operating at 400 Mbps. The data are converted on the CP modules to 160 Mbps single-ended signals before being fanned out to neighbouring modules over a very high-density

custom back-plane, using about 800 pins per slot in a 9U crate. Input/output to/from from FPGAs is performed at 160 Mbps.




Figure 13: ATLAS e/γ trigger algorithm

The e/γ (together with the τ/h) algorithm is implemented in FPGAs. This has only become feasible with recent advances in FPGA technology given the requirement for very large and very fast devices. Each FPGA handles 4×2 windows, and (allowing for the window size) needs data from $7 \times 5 \times 2$ towers ($\eta \times \phi \times \{E/H\}$). The algorithm is described in a language (VHDL) that can be converted into the FPGA configuration file. This provides 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.

V. MUON TRIGGERS

First-level muon triggers are illustrated here by the example of the CMS drift-tube based trigger. More details of the CMS trigger can be found in Ref. [2]. Information on the ATLAS muon trigger can be found in Ref. [1] with updates in papers presented at this workshop [10–13].

The CMS muon system (see Fig. 14) includes three detector technologies — Resistive Plate Chambers (RPCs) and Drift Tubes (DTs) in the barrel, and RPCs and Cathode Strip Chambers (CSCs) in the endcaps. As shown in Fig. 15, all three systems participate in the first-level trigger, with specific logic for each one that is followed by global logic that combines all the muon information.




Figure 14: Detector systems used in the CMS muon trigger

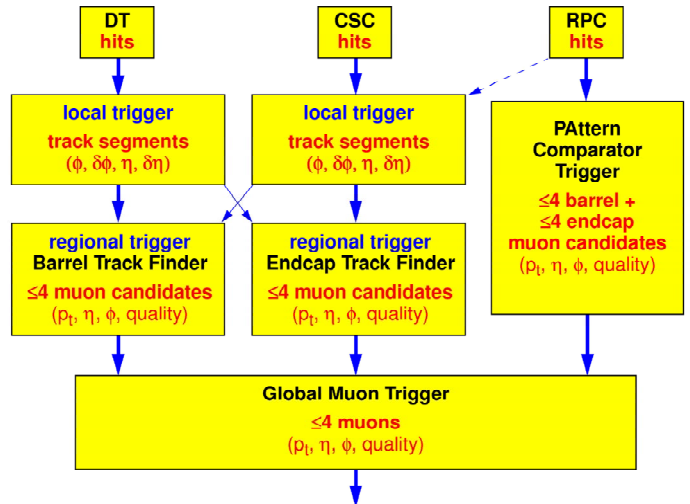


Figure 15: Block diagram of CMS muon trigger

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 (see Fig. 16). The deflection from 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. The derived information consists of local track segments or “superhits” (identified BC, position), and, in some cases (e.g. in CMS DT trigger) also direction information.




Figure 16: Principle of muon trigger

The DT trigger is based on a generalization of the mean-timer method as illustrated in Fig. 17. If a track traverses two drift chamber layers at normal incidence, one can reconstruct the time and position where the track passed through the chamber from the time of arrival of the signals at the two wires (two unknowns, two measurements) – v_d is the drift velocity. Extending the same principle to three layers, one can fit also the angle of incidence. In CMS, four layers are used allowing tracks to be found even if a hit is missing in one of the layers.




Figure 17: Bunch-time identification in the DT trigger

The local logic of the DT trigger is illustrated in Fig. 18. Each super-layer is used to find super-hits using the method discussed above for Bunch Time Identification (BTI). Within each station, pairs of super-layers are combined to find track segments using Track Correlator (TRACO) logic. All this electronics is mounted around the detector and ASICs are used for the implementation.

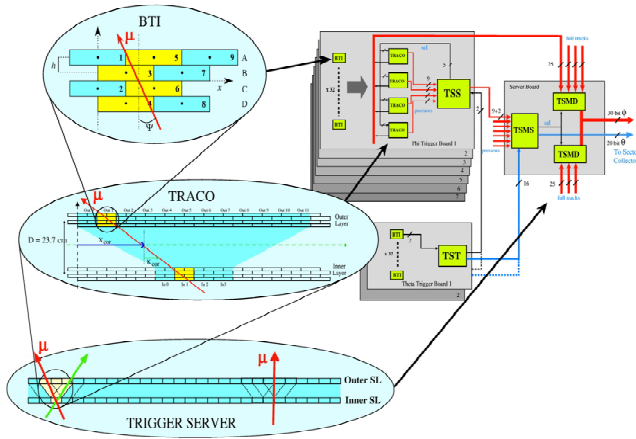


Figure 18: Local DT trigger logic

The on-detector logic sends track segments to the underground counting room where the remainder of the DT trigger is located. The algorithm, illustrated in Fig. 19, is implemented using FPGAs.

The track-finder electronics is mounted off detector. Look-up tables in the FPGAs contain limits of extrapolation windows. Track segments are combined to find the “best” two tracks within a sector. The track parameters (p_T , ϕ , etc) are then determined from the ϕ measurements in different stations.

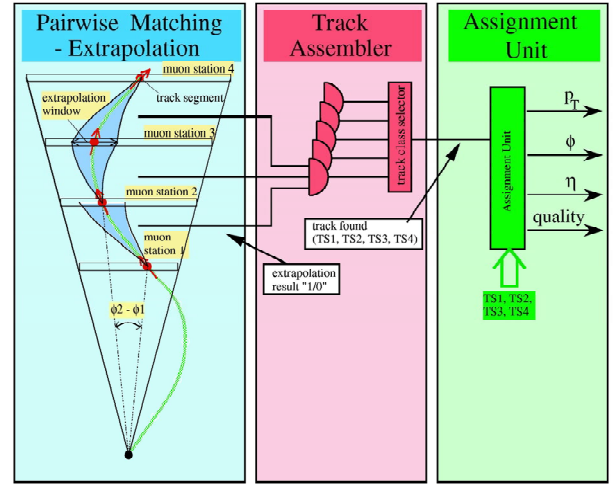



Figure 19: Off-detector DT trigger logic

VI. LHCb PILE-UP VETO

LHCb is designed to work with single-interaction events. It therefore operates at lower luminosity than ATLAS and CMS, $L \sim 2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, at which about 30% of BCs have a single proton–proton interaction. At this luminosity, about 10% of BCs have more than one interaction. LHCb therefore include a pile-up veto in their level-0 trigger. This avoids triggering on multi-interaction events that are not useful for the analysis and which may confuse the higher-level triggers.

As presented in detail at this workshop [14], the LHCb pile-up veto is based on two layers of the silicon tracker and uses a histogramming method (see Fig. 20). The principle of the algorithm is as follows: pairs of hits in each quadrant are combined to calculate the position, z , where a straight-line track giving rise to the hits in the two layers would intercept the beam-line; a histogram is made of the calculated z positions; the resulting histogram is analysed to find the position of highest peak; in a second pass, the process is repeated omitting hits that contributed to the first peak, and a search is made for a second peak indicating a second proton–proton interaction.



Tracks from same vertex give same k
Different vertices generate different k

Figure 20: Principle of operation of pile-up veto

The results of the pile-up veto algorithm applied to a simulated LHCb event containing two proton–proton interactions are shown in Fig. 21.

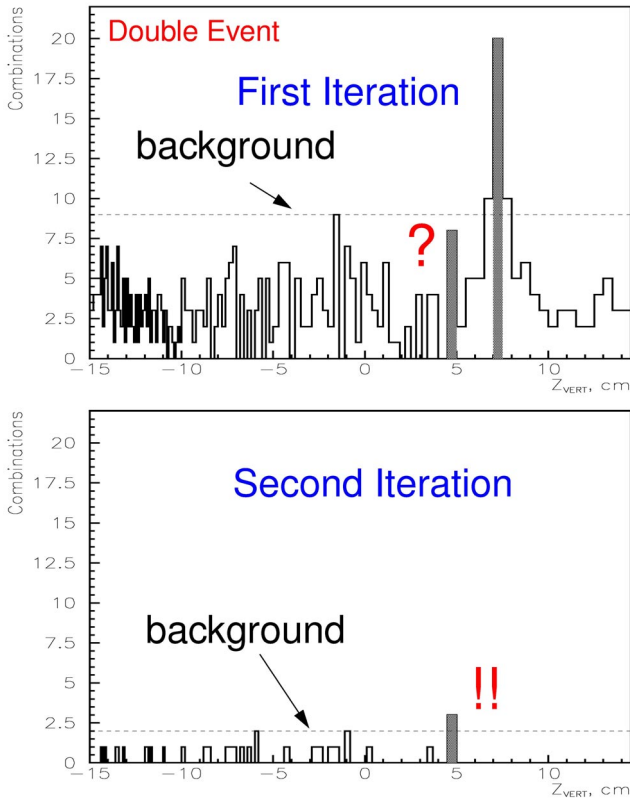


Figure 21: Histograms for a simulated LHCb event

As discussed in Ref. [14], the vertex finder will be implemented using FPGAs. As shown in Fig. 22, there will be a “farm” of four (plus one spare) FPGA-based vertex finders, each one handling one event in four. Data from different quadrants will be multiplexed into the vertex finders over a period of four BCs. This reduces the data rate into each finder by a factor of four. Each vertex finder uses parallel and pipelined processing to implement the algorithm.

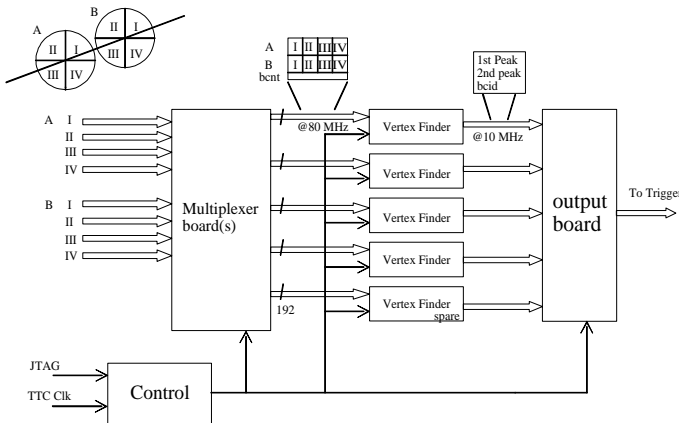


Figure 22: Farm of vertex finders

VII. CENTRAL/GLOBAL TRIGGERS

Central/global are illustrated here by the example of CMS. More details of the CMS trigger can be found in Ref. [2]. Note that a paper [15] presented at this workshop describes the LHCb level-0 decision unit.

The global trigger has to combine information from the different parts of the first-level trigger — local objects: μ , e/γ , τ/h , jet and global energy sums. It makes the overall decision based on combinations of conditions. These can be inclusive conditions, e.g. $p_T(\mu) > 20$ GeV, or more complex requirements, e.g. “ $p_T(\text{jet}) > 100$ GeV and $E_T^{\text{miss}} > 100$ GeV”. Topological conditions can also be imposed in CMS, e.g. “ $p_T(\mu_1) > 20$ GeV and $p_T(\mu_2) > 20$ GeV and $170^\circ < |\phi(1)-\phi(2)| < 190^\circ$ ”.

A block diagram of the CMS global trigger is shown in Fig. 23. It is implemented using FPGAs.

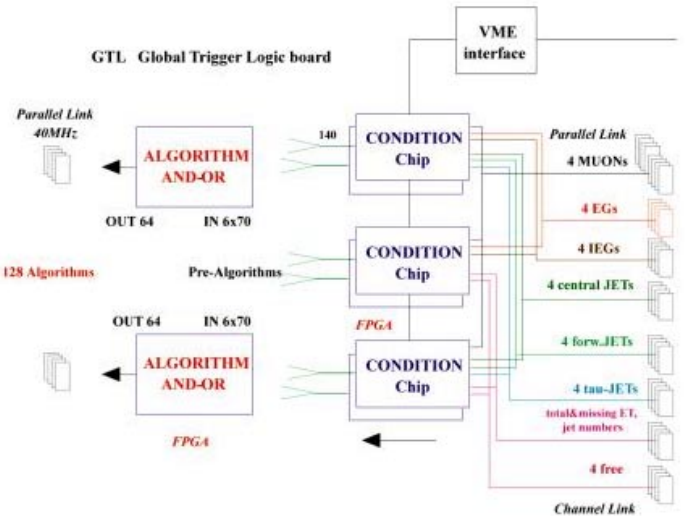


Figure 23: CMS global trigger processor

VIII. CONCLUDING REMARKS

First-level triggers for the LHC represent a huge challenge. They have a direct impact on the physics potential of the experiments since they make the first stage of physics selection. Note that 100 kHz is only $O(10^{-4})$ of the interaction rate in ATLAS and CMS, and the events rejected are lost forever.

The implementation of first-level triggers benefits from new technologies for processing and data movement. The latest generation FPGAs as well as ASICs and high-speed optical and electrical links are being used widely. A very nice aspect of being involved in first-level trigger developments is that there are many challenges that need to be addressed by engineers and physicists working together on algorithms, electronics and software.

A lot of design work and prototyping has been already done for the LHC first-level triggers, as evidenced by numerous reports at this workshop. But there is still plenty to do to complete the final designs and prototyping at module, subsystem and system level.

IX. ACKNOWLEDGEMENTS

I would like to thank numerous colleagues in ALICE, ATLAS, CMS and LHC*b* for their help in the preparation of this talk and lectures given in the CERN Technical Training programme upon which it was largely based. In particular, I would like to thank Hans Dijkstra, Eric Eisenhandler, Philippe Farthouat, Pedja Jovanovic, Peter Sharp, Wesley Smith and Orlando Villalobos Baillie.

X. REFERENCES

- [1] ATLAS Collaboration, First-level Trigger Technical Design Report, CERN/LHCC 98-14.
- [2] CMS Collaboration, Level-1 Trigger Technical Design Report, CERN/LHCC 2000-038.
- [3] LHC*b* Collaboration, LHC*b* Online System Technical Design Report, CERN/LHCC 2001-040 and references therein.
O. Schneider and T. Nakada, “LHC*b* Trigger”, in Proc. Int. Workshop on B-physics and CP violation, Ise-Shima, Japan 19–23 February 2001, eds. T. Ohshima and A. I. Sanda.
- [4] <http://www.ep.ph.bham.ac.uk/user/pedja/alice/>
- [5] B. G. Taylor, “Timing Distribution at the LHC”, plenary talk given at this workshop.
- [6] P. Busson, “Overview of the new CMS electromagnetic calorimeter electronics”, talk given at this workshop.
- [7] G. Mahout, “Prototype Cluster Processor Module for the ATLAS Level-1 Calorimeter Trigger”, talk given at this workshop.
- [8] W. H. Smith, “Tests of CMS Regional Calorimeter Trigger Prototypes”, talk given at this workshop.
- [9] W. Hinderer et al., “The Final Multi-Chip Module of the ATLAS Level-1 Calorimeter Trigger Pre-processor” in Proc. 7th workshop on electronics for LHC experiments, CERN-2001-005.
- [10] K. Nagano, “The ATLAS Level-1 Muon to Central Trigger Processor Interface (MUCTPI)”, talk given at this workshop.
- [11] R. Ichimiya, “An Implementation of the Sector Logic for the Endcap Level-1 Muon Trigger of the ATLAS Experiment”, talk given at this workshop.
- [12] H. Kano, “Results of a Slice System Test for the ATLAS Endcap Muon Level-1 Trigger”, talk given at this workshop.
- [13] R. Vari, “The Design of the Coincidence Matrix ASIC of the ATLAS Barrel Level-1 Muon Trigger”, talk given at this workshop.
- [14] L. Wiggers, “Pile-up Veto L0 Trigger System for LHC*b* Using Large FPGAs”, talk given at this workshop.
- [15] R. Cornat, “Level-0 Trigger Decision Unit for the LHC*b* Experiment”, talk given at this workshop.