

HARDWARE TRIGGERS AT THE LHC

Eric Eisenhandler

Physics Department, Queen Mary & Westfield College, University of London, London E1 4NS, UK

email: e.eisenhandler@qmw.ac.uk

Abstract

This paper gives an overview of hardware triggers, variously called level-0 and level-1, at the two LHC general-purpose experiments, CMS and ATLAS, and at the two specialized experiments, LHCb and ALICE. The emphasis will be on techniques, technologies and special features chosen to be able to handle the huge numbers of detector channels, unprecedented event rates, and very short bunch-crossing time that characterize experiments at the LHC.

1. INTRODUCTION

Triggering of LHC experiments presents enormous and unprecedented technical challenges. The two general-purpose experiments, CMS and ATLAS, must be capable of running at the LHC's extremely high design luminosity of $1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, which produces an inelastic collision rate of $\sim 1 \text{ Ghz}$. The bunch-crossing time of 25 ns is extremely short, requiring that most of the electronics be pipelined, and which implies that on average there are ~ 20 inelastic collisions per bunch-crossing.

LHCb must confront the long-standing problem of triggering on B-meson production at hadron colliders in the difficult conditions of the LHC, in such a way as to allow it to do high-precision physics. ALICE, on the other hand, does not need a very selective trigger. However, it has to handle a huge volume of data, and also find a way to identify and record events in which its Time Projection Chamber is unusable due to pile-up but useful physics could still be extracted from other parts of the detector.

All four experiments are huge undertakings having enormous numbers of detector channels, both in order to achieve high precision and to cope with the high rates. All use multi-level trigger architectures in order to reduce the raw event and readout-data rates to a level that can be stored and analysed. The first level or two of these trigger systems must work far too fast to rely on general-purpose microprocessors, but instead must consist of custom hardware to carry out specific tasks as quickly as possible. Yet at the same time they must be programmable at the level of thresholds, operating parameters and modes so as to be as versatile as possible. This is necessary in order to be able to adapt to both unexpected operating conditions and to the challenge of new and unpredicted physics that may well turn up. In this brief review the custom 'hardware' triggers of all four experiments will be described briefly and, where relevant and interesting, compared.

All of the experiments have higher-level triggers based on software running in processor farms, in order to refine further the event selection and to reduce the rate to a feasible level for permanent storage. Unfortunately, space does not permit discussion of these; nor does it allow any discussion of the physics performance of the hardware triggers described.

2. ATLAS LEVEL-1 TRIGGER

The ATLAS level-1 trigger [1] is based entirely on muon detector and calorimeter information. Two separate trigger systems produce results that are combined for decision-making in a Central Trigger Processor (CTP), as shown in fig. 1.

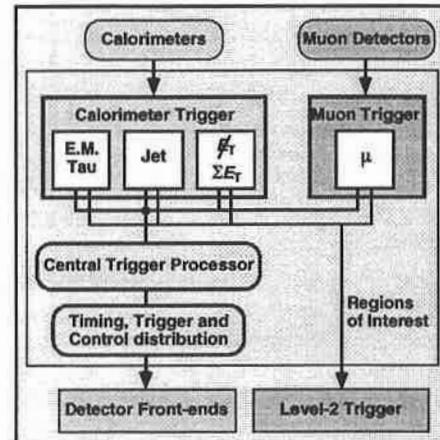


Fig. 1. Block diagram of the ATLAS level-1 trigger.

The ATLAS level-1 trigger must reduce the rate from the bunch-crossing value of 40 MHz to 75 kHz (with the possibility of a future upgrade to 100 kHz). The latency allowed between the interaction time and the trigger decision reaching the detector front-ends is 2.5 μs . For safety about 0.5 μs is preserved as contingency, and almost half of the remaining 2 μs is consumed in cables or fibres from and back to the detectors. Since the trigger obviously needs more than 25 ns to do its work, deadtime-free operation demands pipelined operation. The current estimate of level-1 trigger latency is 2.05 μs .

Other requirements on the level-1 trigger include unique bunch-crossing identification (BCID), which is a particular problem with the calorimeters (see sect. 2.2), and the provision of 'regions-of-interest' (RoIs) to the level-2 trigger so that it only has to read in data around all the trigger objects found at level-1.

2.1 Muon trigger

The muon trigger uses dedicated, fast muon detectors in order to achieve the required speed of operation. In the barrel these are resistive-plate chambers (RPC), and in the

endcap thin-gap chambers (TGC). The layout of the three muon ‘stations’ is shown in fig. 2. Each station has a chamber doublet, except for the inner endcap station which has a triplet. The RPCs cover $|\eta| < 1.05$, and since they have no wires are relatively easy to build and can cover large areas inexpensively. The TGCs cover $1.05 < |\eta| < 2.4$, and need finer granularity since the trigger stations are closer together than in the barrel, momenta are higher, and because of higher backgrounds in the forward region, especially in areas outside the toroidal magnetic field. Both types of chamber are fast enough to give unique BCID, and they also provide the second coordinate to $\sim 5\text{--}10$ mm precision. There are ~ 800 k trigger channels to handle.

As illustrated in fig. 2, the inner two muon stations are used in coincidence for the low- p_T trigger, with a p_T threshold range of 6–10 GeV, while all three stations are used in coincidence for the high- p_T trigger, which provides a threshold range of 8–35 GeV.

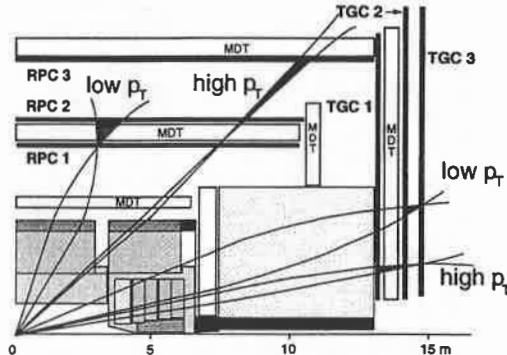


Fig. 2. ATLAS muon-trigger detectors.

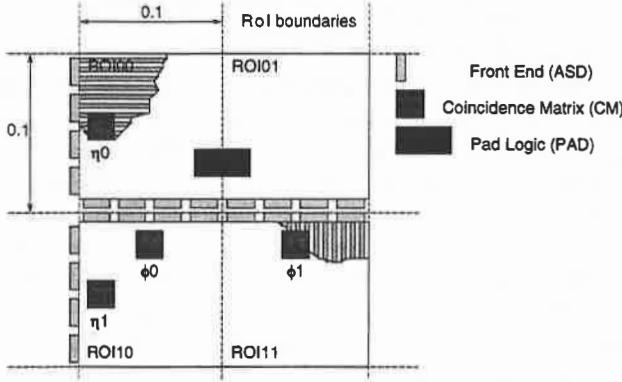


Fig. 3. Barrel muon trigger on-detector electronics.

The barrel muon-trigger logic is mounted on the muon-trigger detectors, as shown in fig. 3. There are a total of 55,000 front-end boards and 3,328 coincidence-matrix boards. These are based on coincidence-matrix ASICs (fig. 4) that synchronize the signals, then look for tracks inside ‘roads’ in one coordinate view. The matrix is $32 \times 48 \times 3$ to give three programmable p_T thresholds each for low- and high- p_T triggers. Deep submicron CMOS is used for radiation tolerance. The ASIC has a working frequency of 320 MHz, ~ 120 k gates, 210 I/O pins, and consumes 1 W. A prototype has performed well.

The two coordinate views and the low- p_T and high- p_T triggers are combined in pad logic boards, which assign candidates to RoIs and resolve overlapping objects.

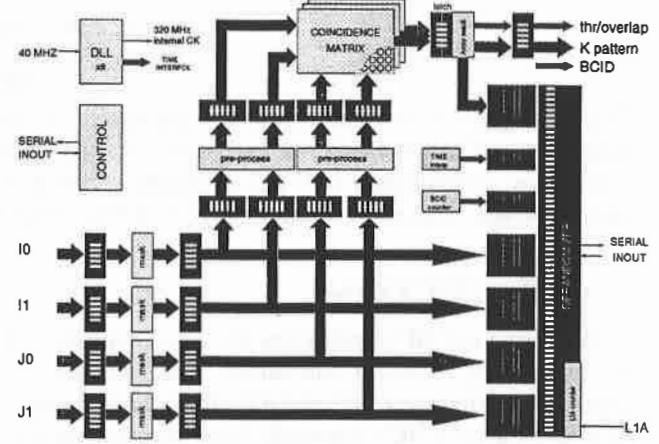


Fig. 4. Coincidence-matrix ASIC for barrel muon trigger.

The endcap muon-trigger logic is shown in fig. 5. The low- p_T trigger requires coincidence matrices of 72×88 with 3 out of 4 coincidence logic, while the high- p_T trigger needs a 256×288 matrix with 2-fold logic. Prototypes have used FPGAs, but ASICs are planned.

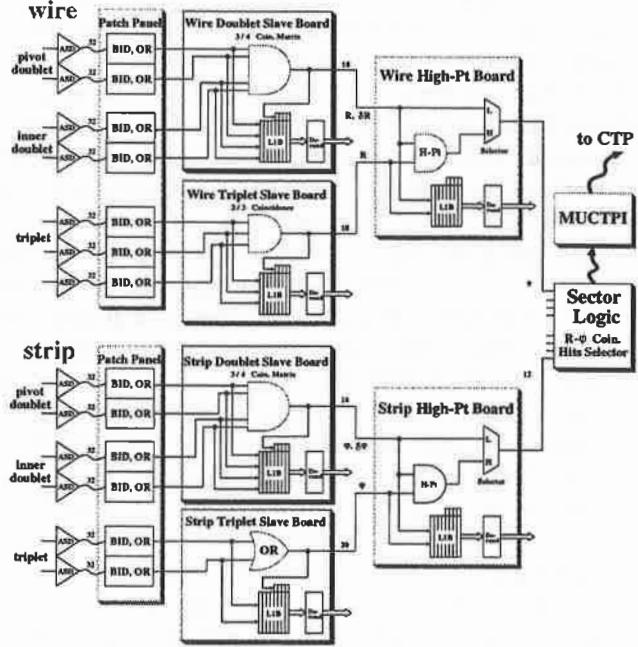


Fig. 5. Endcap muon trigger logic.

Both the barrel and endcap muon triggers send results off-detector optically to sector logic, which examines 64 sectors in the barrel and 72 per endcap and passes the two highest- p_T candidates per sector to the muon-CTP interface. This combines the sector results to produce the total multiplicity passing each of the three low- p_T and the three high- p_T muon thresholds to the CTP. The muon-CTP interface also removes double-counting in muon-chamber overlap regions.

2.2 Calorimeter trigger

The calorimeter trigger uses trigger-tower signals summed on the detector and transmitted in analogue on twisted pairs to the trigger, whose architecture is shown in fig. 6. There are three subsystems: the Preprocessor, the Cluster Processor that finds electron/photon and hadron/tau candidates exceeding any of eight E_T thresholds each, and a Jet/Energy-sum Processor that finds jets and missing- E_T exceeding any of eight thresholds and total scalar E_T exceeding four thresholds. The results are sent to the CTP in the form of multiplicities of each type of trigger object, and as RoIs giving the coordinates of each object found to level-2.

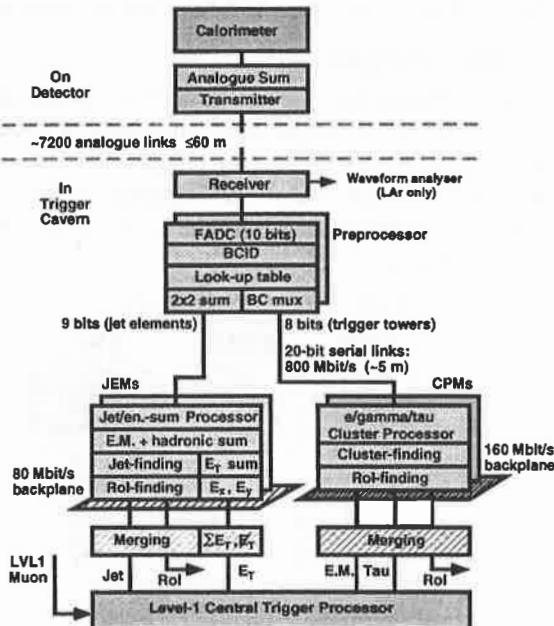


Fig. 6. Block diagram of the calorimeter trigger.

Trigger towers are summed over the full depth of each calorimeter, and laterally in η - ϕ to 0.1×0.1 for $|\eta| < 2.5$ and typically 0.2×0.2 beyond. The preprocessor digitizes the signals to 10 bits, with ~ 0.25 GeV/count. After preprocessing, the trigger algorithms use ~ 1 GeV/count. A summary of the trigger-element granularity and coverage is given in table 1.

Table 1. ATLAS calorimeter trigger parameters.

Trigger type	Granularity	Coverage	No. of elements
electron/photon	$\sim 0.1 \times 0.1$	$ \eta < 2.5$	~ 6400
hadron/tau			
jet	$\sim 0.2 \times 0.2$	$ \eta < 3.2$	~ 1920
missing- E_T	$\sim 0.2 \times 0.2$	$ \eta < 4.9$	~ 1986
sum- E_T			

The Preprocessor consists of eight crates of 16 preprocessor modules, each module handling 64 trigger towers. In order to achieve this, most of the electronics is on multi-chip modules (MCM), and much is done on an ASIC, as shown in fig. 7. Memories are provided for

reading out trigger data to DAQ. Since the calorimeter pulses are several bunch-crossings wide, a crucial issue is bunch-crossing identification, which also requires that an accurate E_T value is extracted. A programmable digital algorithm using a finite-impulse response filter and a peak-finder is implemented on the ASIC, as well as separate logic for BCID on saturated pulses so that they always produce a trigger. A lookup table calibrates E_T , subtracts pedestals, and applies a noise threshold. Results are transmitted serially to the cluster processor on HP G-links, and by using the fact that the BCID forbids pulses on two successive bunch-crossings it is possible to transmit four trigger towers per serial link at less than 1 Gbaud. Jet/energy-sum information is pre-summed to 0.2×0.2 before serial transmission.

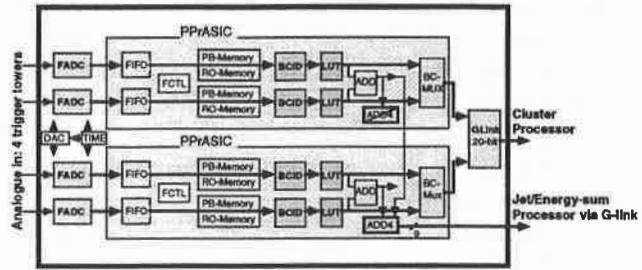


Fig. 7. Functional diagram of Preprocessor MCM.

The electron/photon algorithm is illustrated in fig. 8. Two-tower sums are compared to E_T thresholds, and independently-programmable e.m. and hadronic isolation thresholds are available. The overlapping windows slide by 0.1 in both η and ϕ , so a localized shower produces hits in more than one window. The ambiguity is resolved and RoIs identified by also demanding that the inner 4×4 towers contain a local E_T maximum compared to the eight overlapping neighbours. This algorithm is executed in ASICs, each of which handles eight such windows. The hadron/tau algorithm is very similar, except that the hadron isolation region is the outer 12 cells and the threshold is done on a sum of e.m. and hadronic towers; this is performed in parallel in the same ASICs.

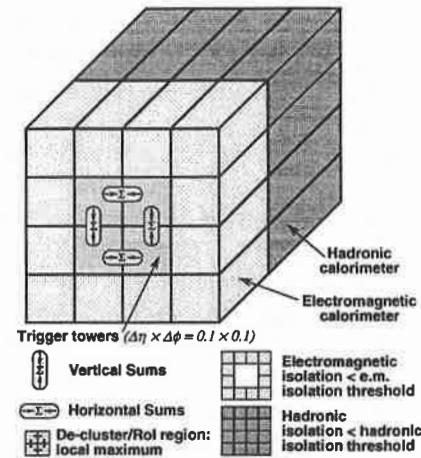


Fig. 8. ATLAS electron/photon algorithm.

The jet algorithm is shown in fig. 9. For each of eight thresholds, the size of jet window can be independently selected to be 4×4 , 3×3 , or 2×2 jet elements of 0.2×0.2 each, in order to be able to optimize on inclusive triggers or to resolve multiple jets. The RoI and de-clustering mechanism again uses local maxima. The windows slide and overlap by 0.2 in η and ϕ . The jet, missing- E_T and total- E_T triggers use FPGAs extensively.

In all of these overlapping-window algorithms, each trigger element participates in 16 windows. This implies massive data fanout. In order to keep the number of connections manageable, inputs to modules in both types of trigger processor use serial inputs carrying multiple elements per link. Backplane fanout between modules uses semi-serialized single-ended data at 160 Mbit/s in the cluster processor and 80 Mbit/s in the jet/energy-sum processor. For both processors, the architecture is as shown in fig. 10, with crates fully covering quadrants in ϕ and modules covering slices in η . This requires fanout only to the nearest-neighbour modules, which greatly simplifies the backplanes. Fanout between crates is done by duplicating signals at the Preprocessor outputs.

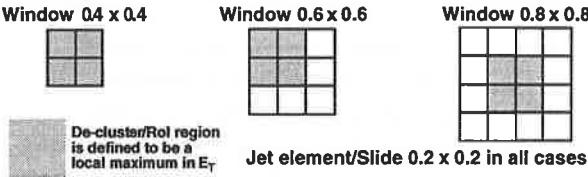


Fig. 9. ATLAS jet algorithm. The window size is programmable for each choice of threshold.

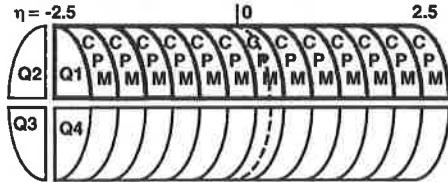


Fig. 10. Calorimeter trigger ϕ -quadrant architecture.

2.3 Central Trigger Processor

The CTP receives results from the calorimeter and muon triggers in the form of 3-bit multiplicities above thresholds for electron/photons, hadron/taus, and jets, as well as bits flagging missing- E_T and total- E_T above thresholds. The 128 input bits also allow calibration and test triggers. Combinatorial logic forms up to 96 different types of trigger, permitting combinations such as: at least two jets of $E_T > 50$ GeV AND missing $E_T > 30$ GeV.

Outputs go to the Timing, Trigger and Control system for distribution to detector front-ends, DAQ, etc. as well as telling level-2 what caused the trigger. Other functions of the CTP include deadtime control, prescaling of high-rate triggers, and monitoring of rates and deadtime. The logic is based on FPGAs and CPLDs.

3. CMS LEVEL-1 TRIGGER

The CMS level-1 trigger [2–4] has very similar requirements to ATLAS, so much is familiar. However, there are also some interesting differences of approach. Once more, there are separate muon and calorimeter triggers, with a combined requirement of reducing the rate to 75 kHz. The latency permitted is somewhat longer, at 3.2 μ s, and the current estimate for the design is 3.0 μ s.

One difference of philosophy is that ATLAS compares objects to E_T or p_T thresholds locally and sends *hit multiplicities* to the CTP, while CMS *sorts* objects both locally and globally and sends E_T or p_T together with coordinate and quality information to the Global Trigger where thresholds and other requirements are imposed.

3.1 Muon trigger

As in ATLAS there are low- p_T and high- p_T triggers, but in CMS the low- p_T trigger uses dedicated RPCs while the high- p_T trigger uses the main muon detectors — drift tubes (DT) in the barrel and cathode-strip chambers (CSC) in the endcaps, as shown in fig. 11 — to refine the measurement of p_T .

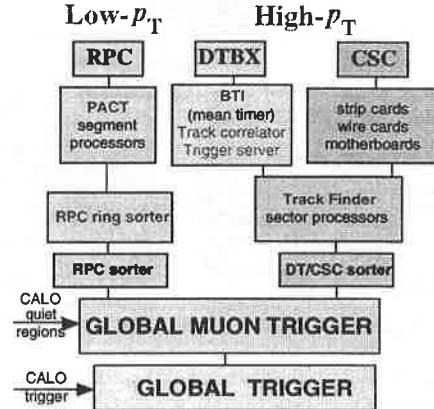


Fig. 11. Block diagram of the CMS muon trigger.

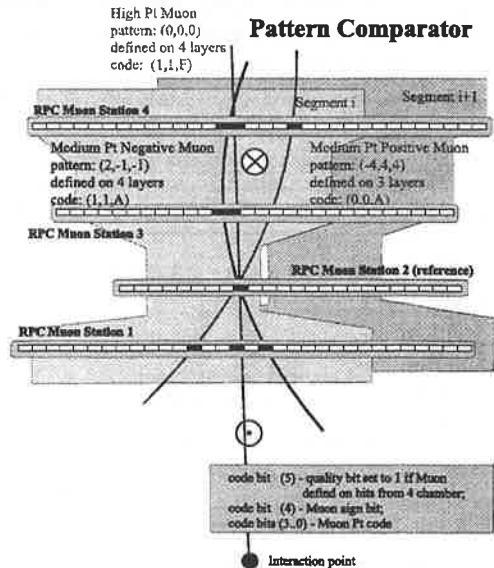


Fig. 12. RPC trigger concept.

The RPCs cover $|\eta| < 2.1$, in $\eta-\phi$ strips of $0.1 \times 5/16^\circ$. Hits in the four RPC stations are compared to predefined templates covering different p_T ranges in Pattern Comparator ASICs, as shown in fig. 12. The modularity of the trigger, which can measure p_T up to 50–100 GeV, is 38 rings each divided into 144 ϕ segments.

The layout of the DTs and CSCs is shown in fig. 13. Each barrel DT station has one z and two ϕ super-layers. Six rings of 0.35 in η are each divided into 12 ϕ segments. The Bunch and Track Identifier forms $r-\phi$ vectors for each super-layer by solving linear equations for the hits, then the Track Correlator combines the two ϕ super-layers to form a vector for each station — see fig. 14 (left). Finally, the Trigger Server sorts these vectors by quality and p_T , and outputs the two highest to the Track Finder.

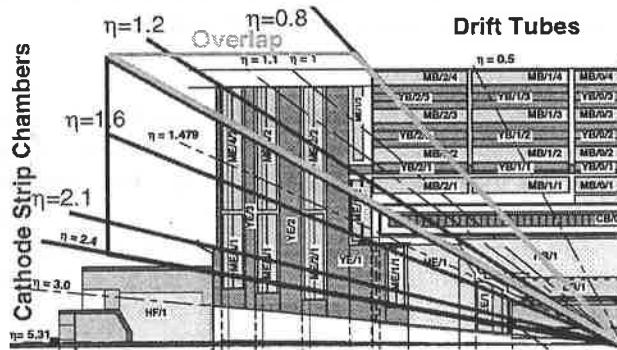


Fig. 13. CMS muon detector layout.

The CSCs have six layers per station, with readout on radial strips and orthogonal wires. The Local Charged Track processor finds coincidences inside predefined roads in ≥ 4 out of six layers, and sends the vector to the Track Finder — see fig. 14 (right).

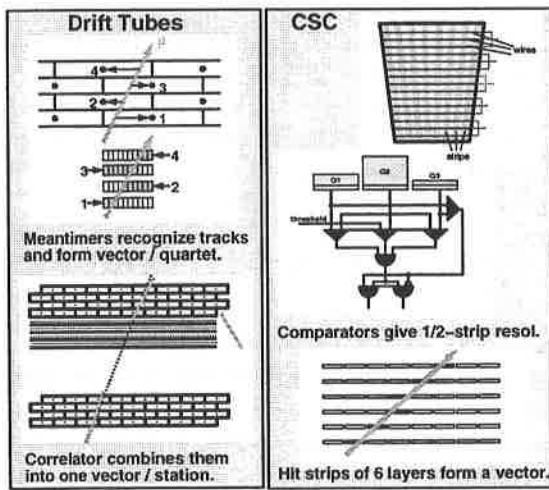


Fig. 14. Drift tube and CSC muon-trigger concepts.

The Track Finder (fig. 15) combines DT and CSC track segments into full tracks, deals with the DT/CSC overlap region, assigns p_T and quality to each one, and sorts them.

The Global Muon Trigger takes in the four highest- p_T muon candidates from both the RPC Pattern Comparator and the DT/CSC Track Finder, removes ghosts, looks at

'quiet' bits from the hadron calorimeter for isolation, and outputs the four highest muons along with their p_T and location to the Global Trigger.

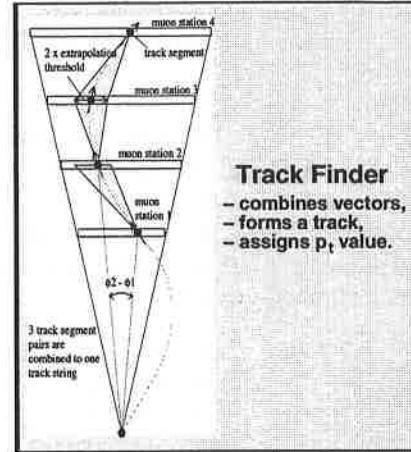


Fig. 15. Drift tube and CSC muon-trigger Track Finder.

The muon trigger logic has been prototyped using FPGAs, but depending on market developments in FPGAs a number of ASICs might be used; several prototype ASICs either exist or are being designed.

3.2 Calorimeter trigger

The overall architecture of the calorimeter trigger is shown in fig. 16. Trigger towers are 0.087×0.087 in $\eta-\phi$ for $|\eta| < 2.1$, and in general twice as big in η for $2.1 < |\eta| < 2.6$. The total number is 54×72 towers for each of the e.m. and hadronic calorimeters. Towers are formed on Trigger Primitives Boards, which transmit the tower data to the Calorimeter Regional Trigger on an 8-bit quad-linear scale plus error bits. These links use serial 1.2 Gbaud copper links with Vitesse Gigabit Ethernet chips.

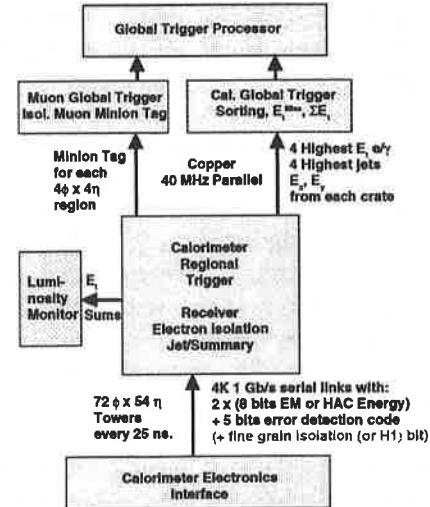


Fig. 16. Block diagram of the calorimeter trigger.

In addition to forming trigger towers, the Primitives Boards also compare pairs of crystal strips with their neighbours to produce very fine-grained isolation bits, as

shown in fig. 17. The Primitives Boards will do most of their work on an ASIC.

The Calorimeter Regional Trigger carries out the algorithms for electron/photons and jets, and begins global energy sums before passing the information to the calorimeter global trigger.

The electron/photon algorithm is explained in fig. 17. As in ATLAS, pairs of towers are examined. Hadronic veto logic is done separately for the tower behind the peak and for its neighbours. Unlike ATLAS, the e.m. isolation covers corners rather than a full ring in order to minimize the fanout required, but this is compensated by the fine-grained shower-profile cut.

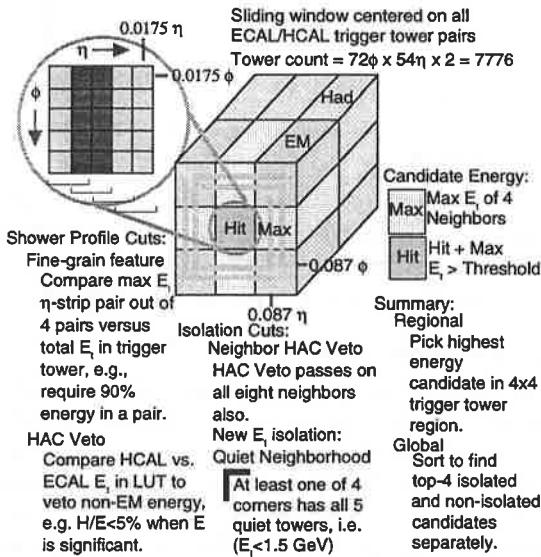


Fig. 17. CMS electron/photon algorithm.

The jet algorithm uses 4×4 non-overlapping windows of 0.35×0.35 in η - ϕ , a size optimized for resolving multi-jet triggers. It is claimed that the non-overlapping windows do not compromise physics performance.

The Calorimeter Regional Trigger uses 19 9U double-depth crates (one is for forward calorimetry needed in energy sums), modularized as two in η and nine in ϕ . The crates (see fig. 18) contain eight Receiver Cards which linearize the data to 7-bit precision and do the first stage of jet and energy sums. Eight Electron Isolation Cards carry out the e.m. algorithm using ASICs. Both electron and jet data are sent to a Jet/Summary card, which begins the process of sorting out the best candidates and forms energy sums for the crate, to send on to the Global Calorimeter Trigger. Data transfers within the crate are 160 Mbit/s differential point-to-point; the backplane exists and works.

In order to achieve low latency for this part of the trigger, two ASICs will be used. A GaAs adder ASIC that sums eight 13-bit numbers in 25 ns using a 160 MHz clock has already been produced, and a sort ASIC that will produce the four highest of 32 8-bit input values is being worked on.

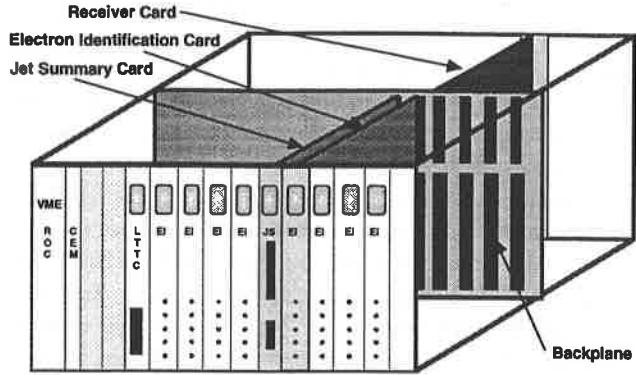


Fig. 18. Calorimeter Regional Trigger crate.

The Global Calorimeter Trigger sorts out the four highest- E_T isolated and the four highest- E_T unisolated electron/photons, the four highest jets, and the missing and total E_T for passing to the global trigger.

3.3 Global Trigger

The Global Trigger (see fig. 19) takes in the trigger objects having the highest E_T or p_T and quality: four muons, four isolated electron/photons, four unisolated electron/photons, four jets, as well as total- E_T and missing- E_T . There are 32 inputs, with possible expansion to 40. Combinatorial logic allows up to 128 trigger combinations.

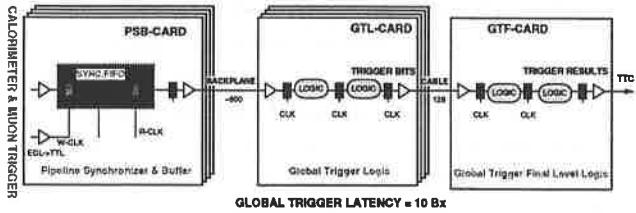


Fig. 19. Global Trigger.

Unlike the ATLAS CTP, it is here that thresholds are applied. The additional information accompanying each object also allows cuts in quality and in location, e.g. in η . It is clear that there is potential for future expansion of capabilities, such as topological triggers — the main limitation is trigger latency.

4. LHCb LEVEL-0 AND 1 TRIGGERS

LHCb is a smaller experiment dedicated to b-quark physics [5], and like its antecedents at hadron colliders, triggering is both very difficult and absolutely crucial. As shown in fig. 20, it has a ‘level-0’ trigger based on calorimetry and muons, and a level-1 trigger on secondary vertices that characterize b-decays, and tracking. As will be seen, the level-1 vertex trigger looks more like a typical level-2 software trigger than others discussed here, but it must be done quickly and is utterly essential to LHCb so it is included for those reasons.

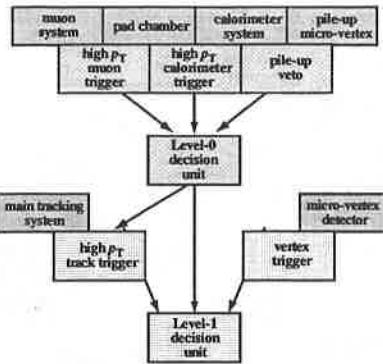


Fig. 20. Block diagram of LHCb level-0 and 1 triggers.

The trigger requirements are that level-0 should have a fixed latency of $< 3.2 \mu\text{s}$ and reduce the rate from $\sim 9 \text{ MHz}$ (see below) to $< 1 \text{ MHz}$. The level-1 trigger has a variable latency of $< 256 \mu\text{s}$ with an average of $\sim 120 \mu\text{s}$ while reducing the rate to $< 40 \text{ kHz}$.

Unlike ATLAS and CMS, LHCb cannot analyse bunch-crossings producing more than one p-p interaction, so its running luminosity will be $\sim 2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, yielding a single-interaction rate of $\sim 9 \text{ MHz}$ and a multiple interaction rate of $\sim 3 \text{ MHz}$. A special pile-up veto at level-0 will be used to eliminate multiple interactions.

4.1 Level-0 trigger

This looks for high- p_{T} electrons, photons, hadrons and muons, although it must be borne in mind that what LHCb regards as ‘high- p_{T} ’ tends to be an order of magnitude lower than ATLAS or CMS.

4.1.1 Calorimeter triggers

Electromagnetic calorimeter information is used to select isolated e.m. showers, with the preshower helping to reject hadrons, and tracker pads in front of the calorimeter used to discriminate between electrons and photons. Hadrons are selected by first examining the hadronic calorimeter, then adding e.m.-calorimeter energy in matching regions.

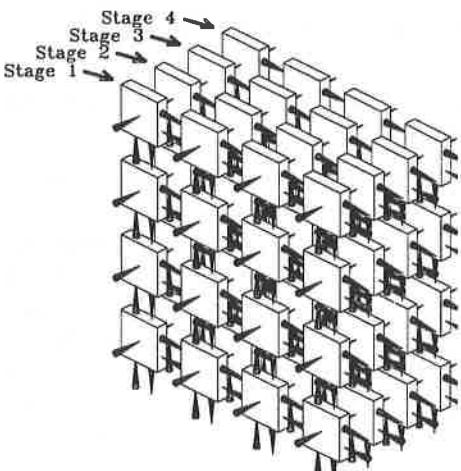


Fig. 21. Concept of 3D-Flow processor.

There are several competing options for these triggers. One option, whose principle is illustrated in fig. 21, is 3D-Flow with 3×3 clustering of calorimeter cells. Programmable processor ASICs running at 80 MHz are arranged in planar layers. To allow 40 MHz pipelined operation, several layers are needed. Cluster logic is done by nearest-neighbour ASICs exchanging data. It is estimated that the electron/photon trigger would need four layers, with ~ 6000 processors per layer, and that the algorithm would take $< 1.5 \mu\text{s}$ to execute.

Another 3×3 clustering option is based on what is used in HERA-B, and uses regions-of-interest and a lookup-table technique.

Finally, there is also a proposal to use 2×2 clustering instead of 3×3 to simplify the logic and to reduce the necessary connectivity.

4.1.2 Muon trigger

The muon trigger will use all five muon stations. Two-dimensional pad readout is used to give the necessary trigger speed. Again, there are still options to be decided.

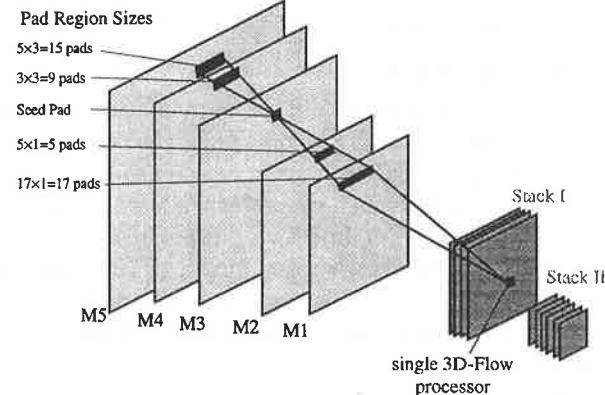


Fig. 22. Level-0 muon trigger 3D-Flow option.

One proposal would once more use 3D-Flow, as shown in fig. 22. In this case 45,000 readout channels would have to be processed, and this would need three processor layers with ~ 1300 processors per layer.

A less ‘heavy’ solution would first use coarse track-finding to limit the number of track candidates needing to be examined in detail. The proposal is to base such logic on FPGAs and DSPs.

4.1.3 Pile-up veto

As already mentioned, bunch-crossings producing multiple interactions cannot be analysed since a unique primary vertex is needed. Multiple interactions are vetoed at level-0 using two dedicated silicon microstrip planes with very fast readout in the backward direction. 3600 circular strips of pitch 120–240 μm and covering 60° in azimuth are processed in parallel to find projected vertex coordinates. The principle is shown in fig. 23, and the layout of the entire vertex detector including the two veto-counter planes is drawn in fig. 24.

A fast processor based on FPGAs finds the z -coordinate of potential vertices to $\sim 1 \text{ mm}$ and histograms

them. It then finds and counts peaks in the histogram to obtain an estimate of the number of interaction vertices. Since primary vertices have $\sigma_z \sim 5$ cm, this veto can retain $\sim 95\%$ of single interactions while rejecting $\sim 80\%$ of multiple interactions.

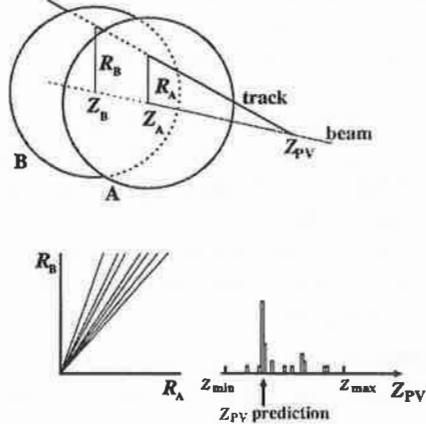


Fig. 23. Concept of level-0 pile-up veto.

4.2 Level-1 trigger

4.2.1 Vertex trigger

This vital trigger should produce a sufficient rate-reduction on its own. It has been facilitated by a redesign of the silicon-microstrip vertex detector (see fig. 24) to use r - ϕ geometry, which simplifies the logic greatly. The procedure is first to find two-dimensional r - z tracks starting from three consecutive hits in r . Then two-track vertices and histograms are used to find z of the primary vertex to ~ 80 μ m, and finally x and y of the primary vertex to ~ 20 μ m.

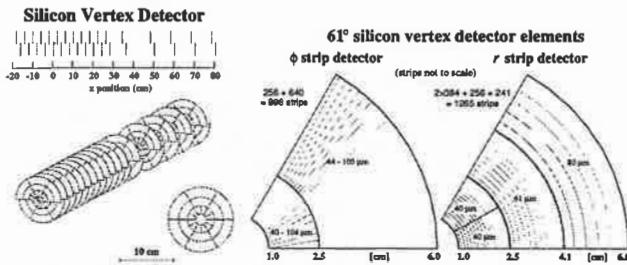


Fig. 24. Layout of silicon vertex detector.

Once the primary vertex has been found, the impact parameter of all tracks with respect to the primary vertex can be evaluated, and then the ϕ data is used to reconstruct the tracks having large impact parameters fully in three dimensions. A search is then made for two-track secondary vertices.

The implementation will be more like a higher-level software trigger than the others discussed here. Vertex-detector events must be built at ~ 1 MHz, and a sustained data throughput of ~ 2 Gbyte/s is required. A number of event-building options are being examined, including the use of dual-port RAMs, as shown in fig. 25. Sub-farms of processors, most likely based on PC-like boards, will be used.

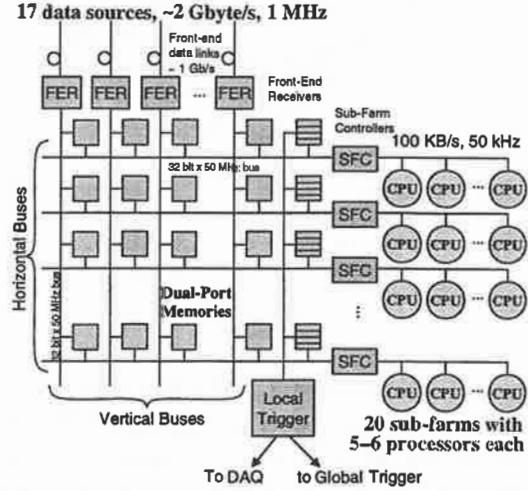


Fig. 25. Dual-port RAM option for level-1 vertex trigger.

4.2.2 Track trigger

A further level-1 trigger, to be staged, uses information from the main LHCb tracking chambers to try to reject false high- p_T level-0 triggers due to decays, secondary interactions, etc. This is based on ideas used in HERA-B. Seeds from the level-0 muon and calorimeter triggers are used to search for tracks (see fig. 26), then a cut is made on the reconstructed p_T . The implementation would be based mainly on DSPs, with some custom electronics. Similar logic is being used for a vertex trigger in the H1 upgrade, and LHCb will benefit from this experience.

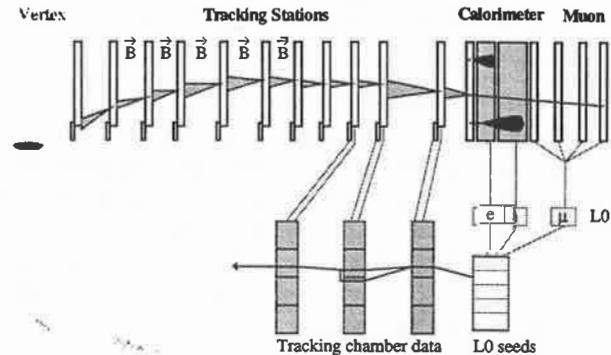


Fig. 26. Concept of level-1 track trigger.

5. ALICE LEVEL-0 AND 1 TRIGGERS

The heavy-ion experiment ALICE [6, 7, 3] is very different from the other LHC experiments. A selection of relevant parameters is shown in table 2. Some of the most notable ones are the huge charged-particle density, the relatively low trigger selectivity required, and the enormous data volume, due mainly to the large Time Projection Chamber (TPC). In fact, due to the long drift time in this device ALICE also foresees doing physics using other parts of the detector and other triggers — mainly dimuons — while the TPC is unavailable, and this adds to the job of the trigger logic. Note that the discussion here will mainly concern ALICE's lead-lead running.

Table 2. Comparison of ALICE and CMS/ATLAS parameters.

	ALICE			CMS/ATLAS
	$Pb-Pb$	$Ca-Ca$	$p-p$	$p-p$
Bunch-crossing period (ns)	125	125	25	25
Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	10^{27}	3×10^{27} 10^{29} ($\mu\mu$)	10^{30}	10^{34}
σ minimum bias (barn)	8	3	0.1	0.1
$dN(\text{charged})/d\eta$	8000	1200	8	8 ($\times 18$)
Minimum-bias rate (Hz)	8000	8000 3×10^5 ($\mu\mu$)	10^5	10^9
Level-1 trigger rejection		10^{-1}		10^{-4}
Event storage rate (Hz)	40 1000 ($\mu\mu$)	150 1000 ($\mu\mu$)	1000	100
Event size (bytes)	33–39 M 0.25 M ($\mu\mu$)	5–6 M 0.1 M ($\mu\mu$)	0.5 M	1 M
Data storage rate (bytes/s)		10^9		10^8
Data storage (bytes/year)		10^{15}		10^{15}

The hardware triggers are divided into level-0 and level-1. An overall block diagram of ALICE triggering is shown in fig. 27. Level-0 has a relatively short, fixed latency of $< 1.2 \mu\text{s}$ and reduces the rate by about a factor of 10, while level-1 has a latency of $< 2.7 \mu\text{s}$ with a rate reduction of only about a factor of two. The main effect of the two levels of trigger is to select central events.

The reason for this small difference in latencies is that ALICE has some detectors with track-and-hold electronics that need to be strobed very quickly, hence level-0, whose short latency requires the level-0 trigger logic to be in the experimental cavern to minimize cable length. However, the detector used in level-1 is too far downstream to fit inside this latency due to the length of its signal cables, as shown in fig. 28.

Another important ingredient is the ability to associate some of the detector with triggers whose physics analysis does not require the TPC, and logic to select this mode of operation within $\pm 100 \mu\text{s}$ of any activity in the TPC in order to prevent pile-up. This is called past-future protection.

Many members of the ALICE collaboration also work on NA57, and the trigger for NA57 is being used as a test bed for a number of concepts needed for ALICE.

5.1 Level-0 trigger

5.1.1 Minimum-bias trigger

This aims to select real interactions from backgrounds. It uses the Forward Multiplicity Detector, a device based on microchannel plates. Their signals have a pulse width of $\sim 1 \text{ ns}$ and a time resolution of $\sim 50 \text{ ps}$, so timing differences between the forward and backward directions can select vertex z -coordinates and thereby reject beam-gas interactions. The timing logic is based on fast passive summation of pads.

The Forward Multiplicity Detector is also used to trigger on charged-particle multiplicity in specific ranges of rapidity, using the pulse-heights of the signals.

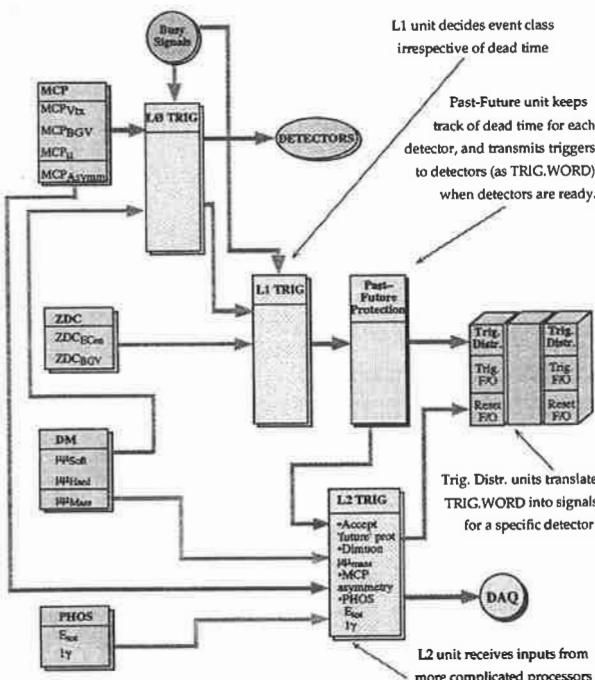


Fig. 27. Block diagram of ALICE triggers.

5.1.2 Dimuon trigger

This trigger is used with the TPC, but is also the cornerstone of triggering in events when the TPC is not used. It was originally in level-1, but rearrangement of the cabling now permits it to be in level-0 (see fig. 28).

The dimuon trigger is based on two RPC stations of two planes each, and finds muon tracks using coincidence matrices that will use either FPGAs or ASICs. Simply demanding $p_T > 1$ GeV already reduces the rate by a factor of ~ 10 .

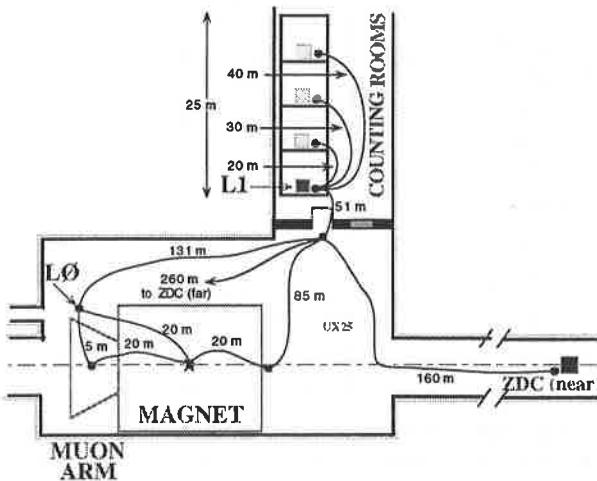


Fig. 28. Layout of ALICE trigger cables.

5.2 Level-1 trigger

The level-1 trigger is now entirely based on the Zero-Degree Calorimeters, a system of small calorimeters in the LHC tunnel at ± 92 m. In each arm there is one calorimeter for protons and another for neutrons. Readout uses scintillating fibres and photomultiplier tubes. This trigger helps assure the centrality of events, though it only reduces the rate by a factor of ~ 2 .

5.3 Past-future protection

The past-future protection, already mentioned above, is logic that can keep track of all significant interactions, not just triggers, and avoid TPC events over a period of ± 100 μ s. In Pb-Pb running, it rejects 63% of all potential triggers needing the TPC.

In this dead period, other events that do not need the TPC, such as dimuon triggers, are taken. Thus, there are two types of events recorded, huge events at low rate (with TPC readout), and small events at high rate (without TPC readout). Some parameters of both types of events are shown in table 2.

6. SUMMARY

In this brief overview, we have seen that although the CMS and ATLAS level-1 triggers have the same requirements and thus many similarities in their approach, there are also significant differences in both the choice of algorithms and in the design philosophy of the hardware.

We have also seen how LHCb is tackling its triggering problems head-on with a secondary-vertex trigger, and also how ALICE plans a variety of triggers and readout options in order to deal its long TPC drift time.

ACKNOWLEDGEMENTS

I would like to acknowledge my friends and colleagues in all four experiments for their help in compiling and correcting this review. Special thanks to Emilio Petrolo (ATLAS), Wesley Smith (CMS), Ueli Straumann (LHCb), and Orlando Villalobos-Baillie (ALICE).

REFERENCES

1. *ATLAS First-Level Trigger Technical Design Report*. CERN/LHCC 98-14.
2. *CMS Technical Proposal*. CERN/LHCC 94-38.
3. Third Workshop on Electronics for LHC Experiments, London, UK, September 1997. CERN/LHCC 97-60.
4. Second Workshop on Electronics for LHC Experiments, Balatonfured, Hungary, September 1996. CERN/LHCC 96-39.
5. *LHCb Technical Proposal*. CERN/LHCC 98-4.
6. *ALICE Technical Proposal*. CERN/LHCC 95-71.
7. *ALICE Technical Proposal Addendum 1, The Forward Muon Spectrometer*. CERN/LHCC 96-32.