Architecture of a level 1 track trigger for the CMS experiment

2010 JINST 5 C08004

(http://iopscience.iop.org/1748-0221/5/08/C08004)

PUBLISHED BY IOP PUBLISHING FOR SISSA

RECEIVED: *March 30, 2010*
ACCEPTED: *July 1, 2010*
PUBLISHED: *August 12, 2010*

WORKSHOP ON INTELLIGENT TRACKERS 2010,
FEBRUARY 3rd–5th 2010,
LAWRENCE BERKELEY NATIONAL LABORATORY, BERKELEY, U.S.A.

2010 JINST 5 C08004

Architecture of a level 1 track trigger for the CMS experiment

E. Hazen,[a] U. Heintz,[b,1] M. Johnson,[c] R. Lipton,[c] M. Narain[b] and S.X. Wu[a]

[a]*Physics Department, Boston University,*
590 Commonwealth Ave, Boston, MA 02215, U.S.A.
[b]*Physics Department, Brown University,*
182 Hope Street, Providence, RI 02912, U.S.A.
[c]*Fermi National Accelerator Laboratory,*
P.O. Box 500, Batavia, IL 60510, U.S.A.

E-mail: Ulrich_Heintz@brown.edu

ABSTRACT: The luminosity goal for the Super-LHC is 10^{35}/cm^2/s. At this luminosity the number of proton-proton interactions in each beam crossing will be in the hundreds. This will stress many components of the CMS detector. One system that has to be upgraded is the trigger system. To keep the rate at which the level 1 trigger fires manageable, information from the tracker has to be integrated into the level 1 trigger. Current design proposals foresee tracking detectors that perform on-detector filtering to reject hits from low-momentum particles. In order to build a trigger system, the filtered hit data from different layers and sectors of the tracker will have to be transmitted off the detector and brought together in a logic processor that generates trigger tracks within the time window allowed by the level 1 trigger latency. This paper describes a possible architecture for the off-detector logic that accomplishes this goal.

KEYWORDS: Trigger concepts and systems (hardware and software); Particle tracking detectors (Solid-state detectors); Si microstrip and pad detectors

[1]Corresponding author.

doi:10.1088/1748-0221/5/08/C08004

Contents

1 Motivation

The design luminosity for the SuperLHC (SLHC) is $10^{35}/\text{cm}^2/\text{s}$, an order of magnitude higher than the design luminosity for the LHC. At a bunch spacing of 25 ns there will be about 300 proton-proton interactions per beam crossing on average. At this interaction rate trigger rates will increase drastically. The level 1 trigger rate for the CMS detector must remain below 100 kHz to avoid having to rebuild the entire DAQ system. In order to achieve substantial reductions of the level 1 trigger rate without raising thresholds rejection power must be moved from the high level trigger (HLT) into level 1. In the current CMS trigger system tracking information is used only in the HLT and not in the level 1 trigger. Thus adding tracking information into the level 1 trigger is an obvious way to achieve this goal. The level 1 track trigger will not only have to be efficient for tracks with high transverse momenta (p_T) but it will also have to identify isolated tracks. This is important, for example, for triggering on τ-leptons. For this the trigger must have reasonable efficiency for softer tracks with p_Ts as low as 2 GeV and it must be able to assign the tracks to a specific primary collision vertex.

2 Hit rejection and detector geometry

The CMS tracker will have to be upgraded for SLHC as well. In order to arrive at quantitative statements about hit rates and resources needed for the off-detector processor, we use simulations based on a specific detector geometry. However, many of the ideas and conclusions presented here will also apply to alternative detector geometries.

We assume that the upgraded tracker will again consist of silicon detectors. At the peak luminosity, the occupancy in the tracker will be too high for all hits to be read out to an off-detector trigger processor. Thus there will have to be some on-detector rejection of hits from very soft tracks. This rejection can be based on the angle at which tracks are incident on the sensor planes in the tracking detector. Soft tracks hit the detectors at shallower angles than hard tracks.

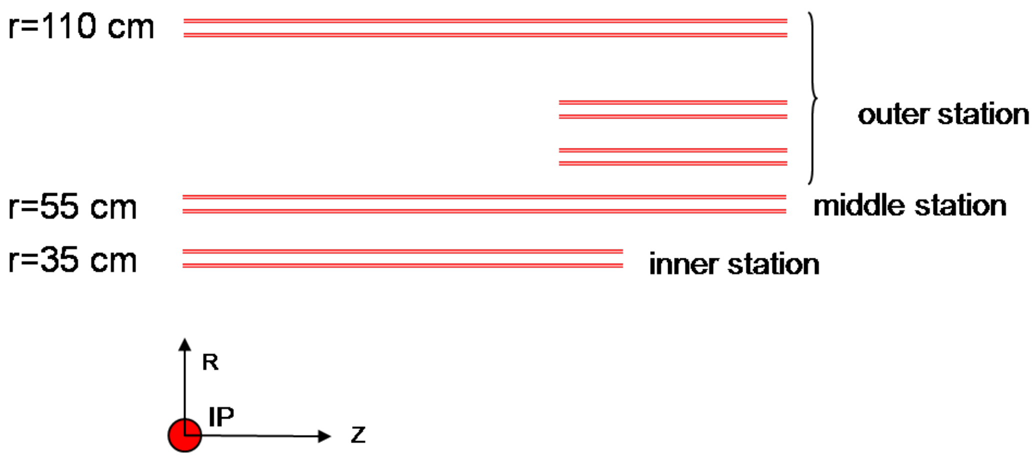

Figure 1. Cut away view of one quarter of the tracker. The horizontal lines indicate the sensor planes.

We employ stacked layers of silicon detectors to detect the incident angle of the tracks. Each stacked layer consists of two sensor planes, closely spaced in radial position. Readout chips (ROC) cluster hit strips in each of the planes independently and discard clusters that are more than two strips wide in azimuth. Then they look for coincidences between clusters in both planes that are separated in azimuth by less than some maximum value $\Delta\phi$. The cut on azimuthal separation effectively removes clusters created by tracks below some transverse momentum threshold. We call such coincidences doublets. These are the trigger primitives that the off-detector track finding is based on.

We assume silicon sensors with an active area of about 100 mm $\times$ 100 mm in $r\varphi\times z$, a strip pitch in azimuthal direction of about 100 μm, and strips that are 1 to 5 mm long in z direction. Two such sensors are stacked on top of each other, separated by a spacer about 1 mm thick, and read out by a single ROC [1]. These sensor stacks are then arranged to form concentric cylindrical layers centered on the beam axis [2]. We consider a tracker geometry that consists of six such cylindrical layers. Pairs of these layers are spaced at a radial distance of about 4 cm. We refer to these pairs as stations. The inner station is positioned at a radius r = 35 cm, the middle station is at r = 55 cm, and the outer station is at r = 110 cm. At this radius the outer station would have to be longer than the available space to have the same coverage in polar angle as the middle and inner stations. We therefore break it into multiple concentric pieces at high z as shown in figure 1. In the inner station one sensor covers a sector of 15^o in φ. Every such φ sector has 42 sensors along z to make a detector 4.2 m in length and there are 24 such φ sectors.

3 Data transmission off detector

We base our rate estimates on a Monte Carlo simulation, using minimum bias interactions generated with PYTHIA [3]. The number of interactions per event follows a Poisson distribution. For this study the mean number of interactions generated per event is 200. table 1 gives the average doublet rate for these events at z = 0. The rates are strongly peaked at z = 0 and the rate averaged over the en-

Table 1. Hit rates for events with an average of 200 minimum bias interactions at z=0. Doublet rates are given for a transverse momentum threshold of 2 GeV.

radius r	35	55	110	cm
cluster rate	4.0	1.6	0.2	MHz/cm^2
cluster rate	10	4.0	0.5	/crossing/module
doublet rate	0.3	0.13	0.025	MHz/cm^2
doublet rate	0.7	0.3	0.06	/crossing/module

tire length of the detectors is less than half the rate at z = 0. The table shows the order of magnitude reduction in data rate from clusters to doublets with a transverse momentum threshold of 2 GeV.

Each doublet requires 20 bits of information, 10 bits to encode the φ centroid, 7 bits to encode the z centroid and 3 bits to encode the difference between the φ centroids of the two clusters which is a measure of the p_T of the track.

For the two stacked layers of the inner station the average doublet rate for one $15^o \varphi$ sector, with a sensor area of two times 4200 cm^2, is about 25 Gb/s. Because our estimates have not been confirmed with actual data from LHC collisions, we add in a factor 10 for safety and to account for event to event fluctuations in the number of doublets. We further assume optical links with a bandwidth of up to 10 Gb/s which appears achievable even with today's technology. Then 30 such links per station and sector would be sufficient to transfer the doublet data from the detector into an off-detector processor.

We assume (conservatively) that the doublet multiplicities in the middle and outer stations are the same as in the inner stations and that we therefore need the same number of optical links for all three stations, which leads to 2160 optical links for the entire tracker. If each driver dissipates a power of about 1 W then the total power dissipated by the optical links would be around 2 kW.

4 Off-detector processing

The principle for the off-detector pattern recognition is to test the pattern of doublets for sets consistent with trajectories of charged particles through the detector. We intend to program the logic in a field programmable gate array (FPGA) such that every valid combination of doublets corresponds to a logic block. If all doublets required for a given trajectory are present the corresponding logic block returns true. This requires that all doublets are loaded into the FPGA and that all possible combinations are programmed into its logic. Then all logic blocks are evaluated in parallel and all tracks found simultaneously. Of course the number of possible doublet combinations for the entire detector is much too large to be programmed into a single FPGA. We therefore have to break the problem down into smaller blocks. This is achieved in several ways.

Doublets from a given track only populate a limited range in azimuth. Thus we can divide processing into azimuthal sectors. Then the processor for each sector would have to receive the data from its home sector and the two neighbouring sectors in azimuth. The lowest p_T for which the trigger has full acceptance is defined by the lowest p_T for which a track can cross at most one sector boundary. If we divide the detector into 15^o sectors this value is 2.4 GeV.

One whole sector still requires too many combinations and too much input data to be processed in a single FPGA. We therefore divide each sector into 12 subsectors. We assign each track to a subsector based on the azimuth at which it hits an anchor station. Each of the three stations can serve as anchor station.

Each station consists of two layers of stacked detectors, separated by 4 cm in radial position. These are sufficiently close together that a given track can only give rise to doublets from a very limited region in z and within one azimuthal sector. Detectors adjacent in azimuth must be made to overlap sufficiently so that any track that crosses sector boundaries hits both layers in at least one of these sectors. The first step of the track finding algorithm consists of connecting doublets in the two stacked layers of a station. We call the resulting track stubs tracklets. In order to find tracklets an FPGA needs to only receive doublets from one station and one azimuthal sector which are transmitted over 30 optical links. This can be accommodated by currently available FPGAs. We expect to find on average about 40 tracklets per sector per station. For each tracklet we will need 30 bits to encode the information and thus each of the tracklet finders output data at an average rate of 48 Gb/s.

The second step of the track finding algorithm is to connect the tracklets in a sector to make tracks. The tracklets are extrapolated to the anchor layer and distributed to FPGAs, corresponding to the subsector they extrapolate to. Tracklets may extrapolate to subsectors in the same sector or in one of the two adjacent sectors. Thus the 48 Gb/s from each tracklet finder will be distributed over 36 target subsectors and one optical link should be sufficient for each of these connections.

Each subsector processor receives tracklet data from two stations in three sectors plus the data from the anchor station in the home sector. Since there is no need to extrapolate the tracklets in the anchor layer we can relax the requirement to find tracklets and also allow single unmatched doublets. This makes the algorithm robust against detector failures by allowing for one missing doublet.

We replicate the subsector processing units three times, with each of the three stations serving as anchor layer in one of the units, thus allowing one missing doublet in any layer without causing inefficiency. Figure 2 shows a schematic of one such unit that uses the outer station as the anchor station.

Each track can then be checked for z-consistency and tracks for which the pattern of the z-coordinates of the doublets is not consistent with the trajectory of a single particle are discarded.

Since we are finding tracks with different anchor stations in parallel there is a potential for finding duplicate tracks which need to be removed. Therefore all tracks are routed to a duplicate removal processor. This will receive p_T-ordered lists of tracks from 12 subsectors in each of three anchor stations. Duplicate tracks will be discarded and a list of tracks will be send over optical links to the level 1 trigger.

The off-detector processing can be accomplished in approximately eleven pipeline stages. Even if this number should increase somewhat in a more detailed design processing can be completed in under 1 μs. The pipeline stages are as follows:

1. load doublet data from sensors into tracklet finder;
2. find tracklets from doublets;
3. sort tracklets by destination subsector;

2010 JINST 5 C08004

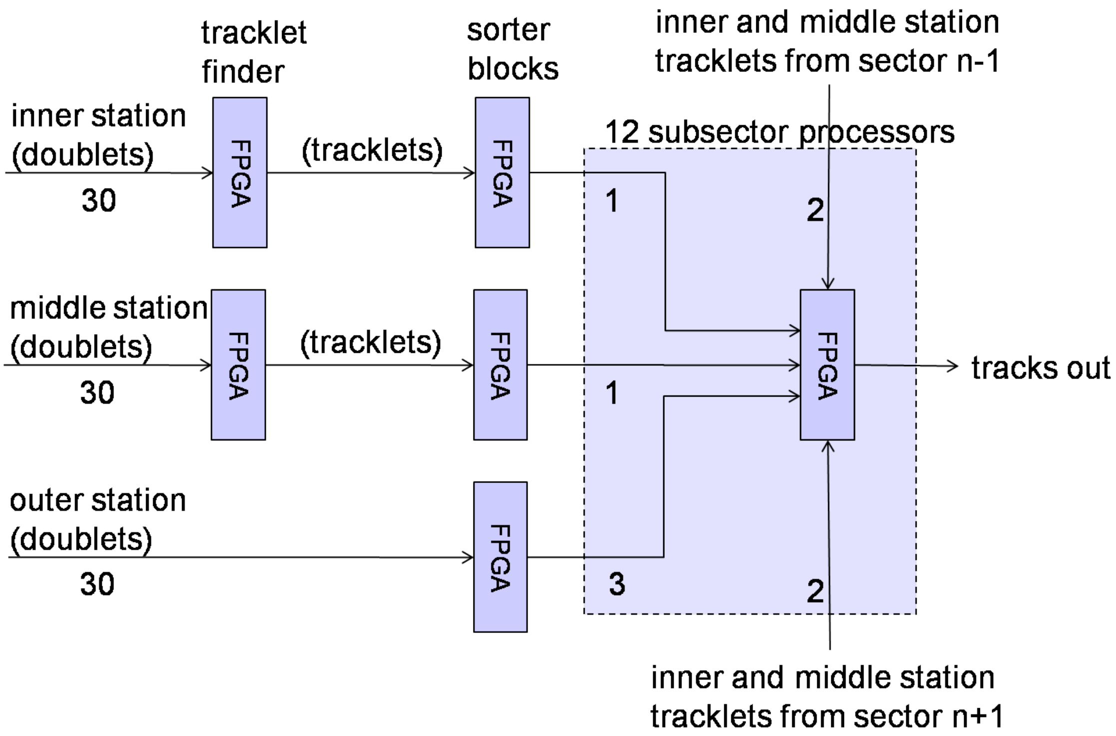

Figure 2. Schematic of 1/3 of the off-detector processor for one sector. Shown is the part that uses the outer station as the anchor station. The numbers give the number of optical links for each connection.

4. transfer to subsector processors;
5. receive tracklet data from tracklet finder in subsector processor;
6. find tracks from tracklets and doublets;
7. check z consistency;
8. transfer tracks to duplicate eliminator;
9. receive track data from all 3 stations in duplicate eliminator;
10. compare inputs, eliminate duplicates;
11. send track data to L1 trigger.

5 Summary and conclusion

We have designed a strawman architecture for the off-detector processor of a level 1 track trigger for the CMS detector at the SLHC. The design is robust against detector failures, e.g. even if a hit is lost in any layer tracks can still be found. This is a crucial property of a trigger algorithm so that individual sensor or chip failures do not impact the trigger efficiency. The design could be implemented in today's technology and is not constrained by the size of available FPGAs. Getting all required data into one place is an important constraint and requires a very large number of high speed links. Reliability, mass, and power dissipation of fiber links are therefore critical.

Acknowledgments

We thank Laura Fields from Cornell University for providing the Monte Carlo simulations that our rate estimates are based on.

References

[1] See R. Lipton et al., *3D Technology for Intelligent Trackers*, in *Workshop on Intelligent Trackers 2010*, February 3rd–5th 2010, Lawrence Berkeley National Laboratory, Berkeley, U.S.A.

[2] See W. Cooper et al., *The design of stable, low-mass support and cooling structures*, in *Workshop on Intelligent Trackers 2010*, February 3rd–5th 2010, Lawrence Berkeley National Laboratory, Berkeley, U.S.A.

[3] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026.