Readout of the CMS experiment during the 2010 heavy ion run

Ivan Amos Cali
(for the CMS collaboration)

Quark Matter 2011, Annecy, France

The CMS Readout

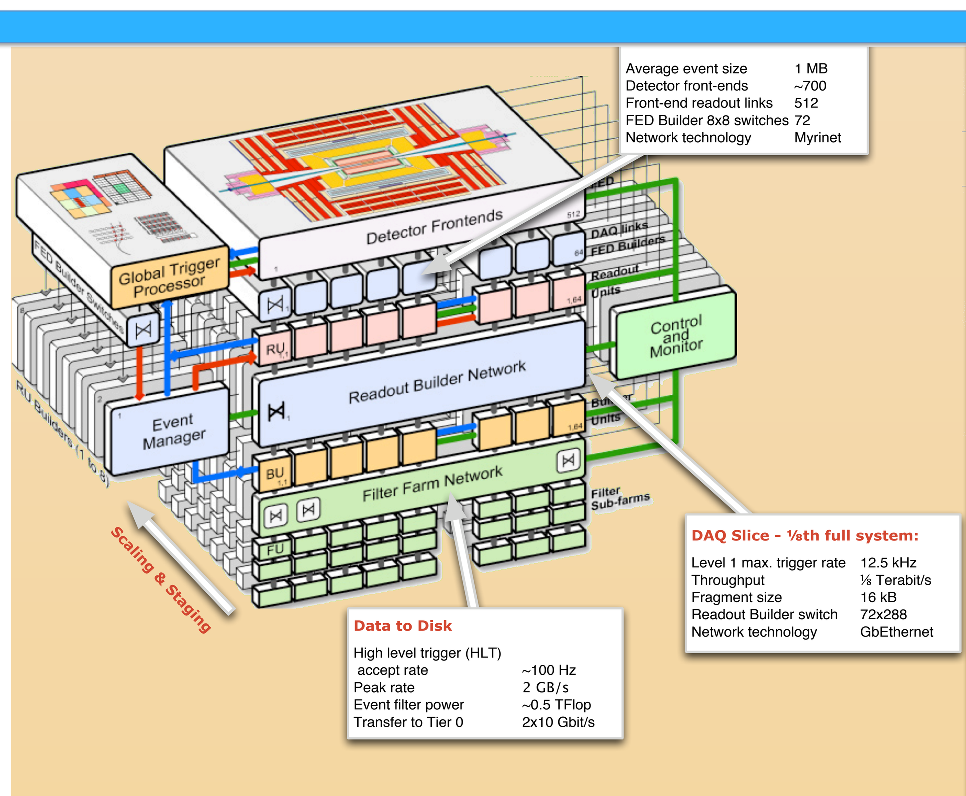

The CMS readout system is designed and optimized for pp collisions with an L1 rate of 100 kHz and an event size at the DAQ input of roughly 2 MB per event.

The large multiplicities expected in PbPb collisions require a different optimization of the zero suppression algorithms. The optimization could only be done after the data was taken. To make sure that the collected data is of highest quality, the CMS collaboration decided to disable the zero suppression algorithms for the silicon strip tracker and the electromagnetic and hadron calorimeters for the duration of the first PbPb run. This resulted in an event size of about 24 MB at the DAQ input and 12 MB to tape. About 11 million channels were recorded for each event.

DAQ performance:
The DAQ system performance has been studied in detail and a new software configuration has been designed for PbPb collisions.

It has been measured that the system was able to handle up to an L1 trigger rate of 3.5 kHz that corresponds to roughly 350 MB/s/link of throughput. At the HLT the data were compressed to 12 MB per event and sent to the Storage Manager (SM). The latter was also responsible for the data transfer to the Tier-0 computing center. The CMS detector was recording data at up to 220 Hz using all the available SM bandwidth in terms of disk I/O and ethernet. The bandwidth to tape was over 2.2 GB/s, which is well beyond the design specifications (more than 6 times the data volume per second recorded during the pp running).

The CMS HI Trigger

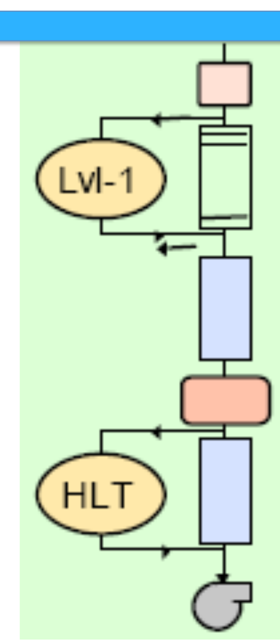

CMS has a two-level DAQ/Trigger architecture: L1 and HLT. The L1 trigger consists of low level hardware using muon track segments, calorimeter cells and scintillators. The HLT is comprised of a powerful online computing farm performing event building and triggering. The HLT accesses the full event information and it uses the standard "offline" reconstruction software to trigger.

The 2010 HI run had a peak instantaneous luminosity of 3x 10^{-5} cm^{-2} s^{-2} and CMS has been able to collect data with an L1 rate up to 220 Hz. The L1 trigger was generated using information from the Beam Scintillator Counters (BSC), Electromagnetic and Hadron Calorimeters (ECAL, HCAL, HF), and the Zero Degree Calorimeters (ZDC).

The HLT was used to produce two main primary datasets (PD). The HIAllPhysics PD contains jets, muons, photons, minimum bias, and zero bias triggered events. The HICorePhysics PD is a subset of the HIAllPhysics PD and it contains jets, photons and muons triggered events for prompt and low-statistics analysis.

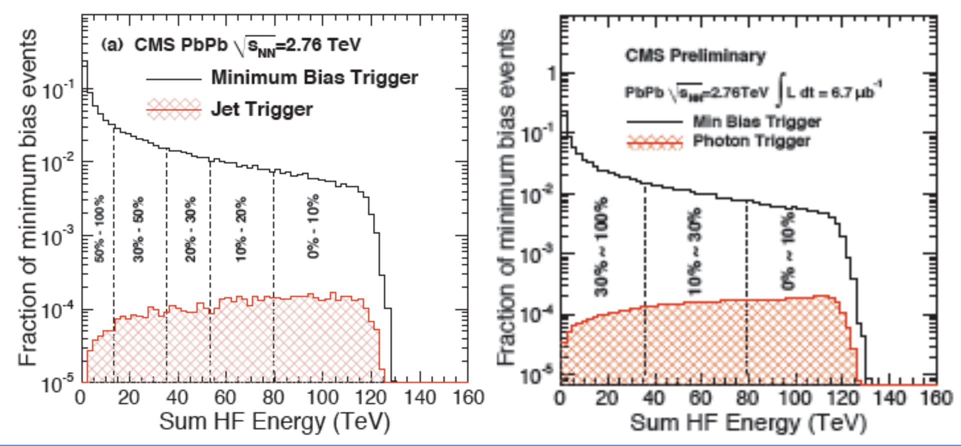

The plots on the left show the fraction of jet and photon triggers with respect to the minimum bias trigger as function of the total HF deposited energy (centrality).

The Silicon Tracker Zero Suppression Algorithm

The CMS silicon tracker readout chain was primarily designed, optimized and tested for pp collisions. In PbPb collisions the expected tracker occupancy was up to 80%. Each set of 128 silicon strip detectors is connected to a single APV readout chip which may introduce electronics effects common to the entire set of detectors. These effects include baseline shift, baseline deformation, and APV saturation generated by highly ionizing particles, which became significant during HI data taking. The standard zero suppression algorithms were not designed to operate in such a high multiplicity environment or to cope with the electronics effects described above. During the 2010 HI data taking the tracker was read out without zero suppression and the actual zero suppression has been performed offline using a baseline follower algorithm developed specially for the PbPb data.

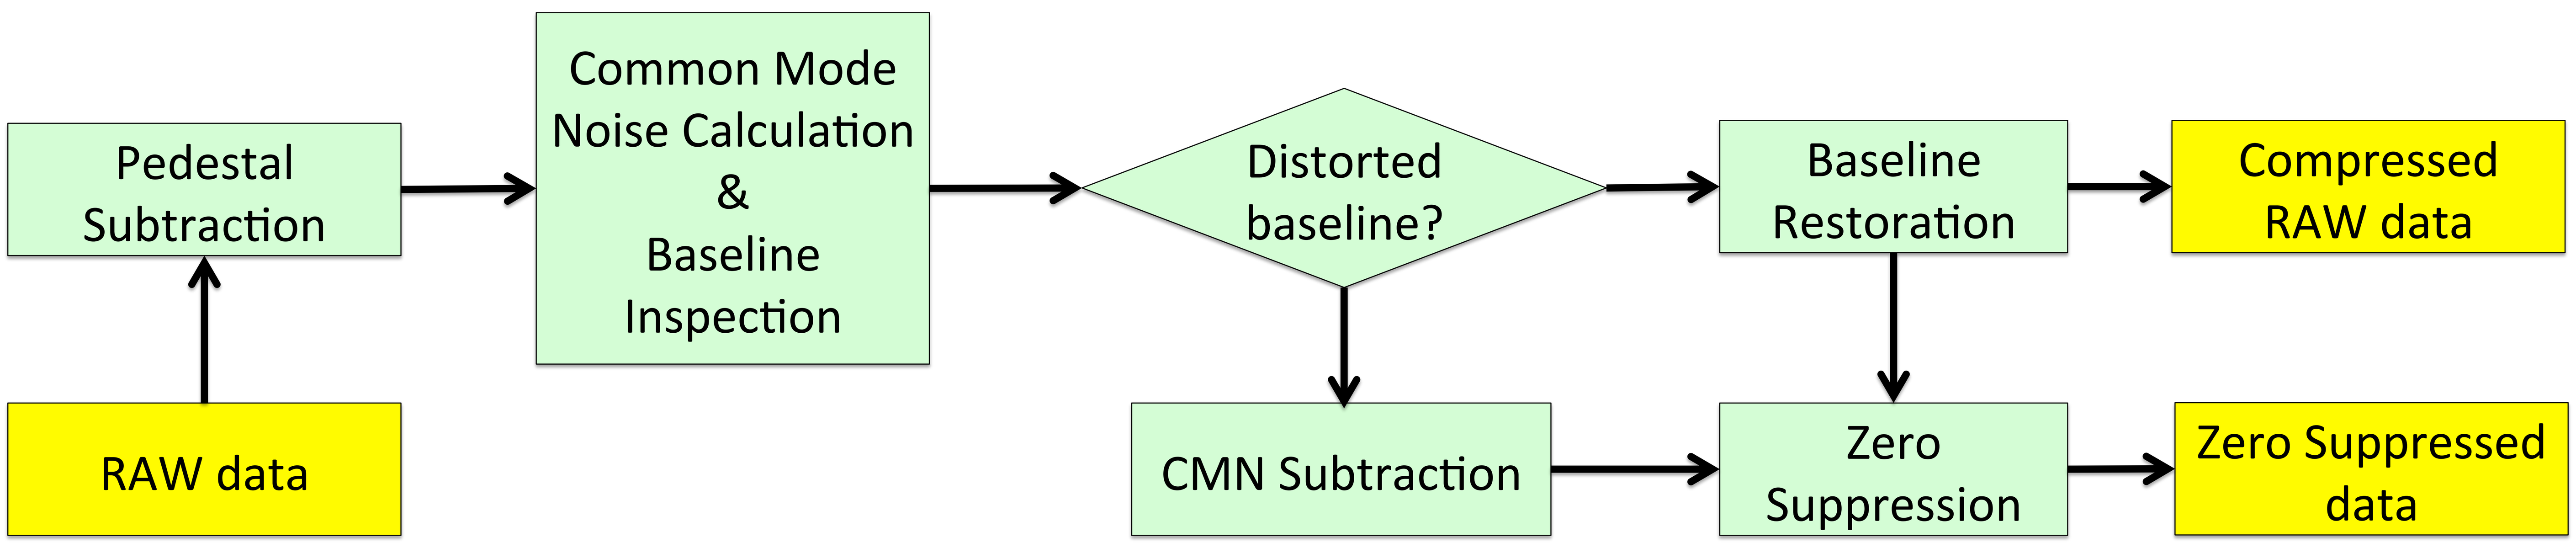

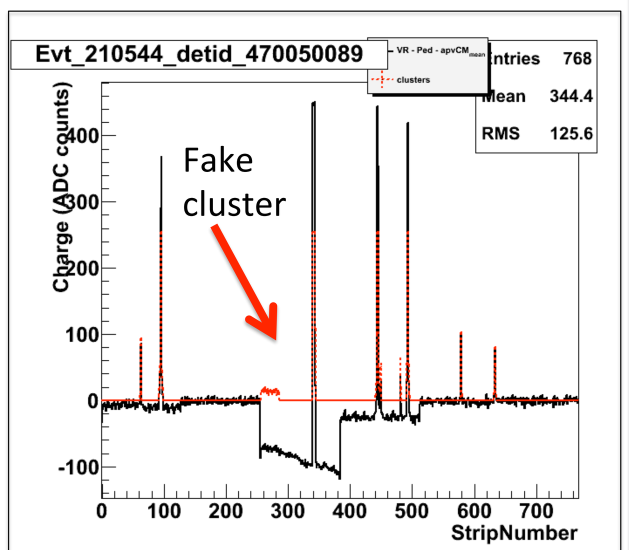

Fig. 1 shows the behavior of the standard zero suppression algorithm in case of baseline deformation. Shown in black is the ADC distribution over the 768 strips of a module. The corresponding reconstructed clusters are shown in red.

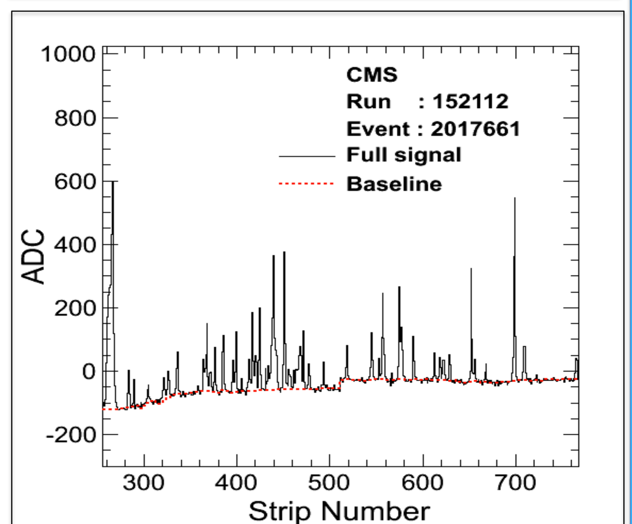

Fig. 2 shows an example of baseline reconstruction at high multiplicity in the presence of a deformation caused by highly ionizing particle. Shown in black are the ADC distribution over the 768 strips of a module. Shown in red is the corresponding reconstructed baseline that is subtracted before performing zero suppression.

The Baseline inspection and restoration algorithm:

1. The baseline inspection looks for flat regions in the distribution of ADC counts arranged by position for a given APV. The ADC value of the strips in flat regions are replaced with the average ADC value in that region.
2. The baseline is considered to be distorted if the vertical distance between the various flat regions is bigger than a certain threshold . The restoration of the baseline is applied only in case of a distorted baseline. Otherwise, the calculated common mode noise is considered to be the baseline.
3. The baseline restoration algorithm interpolates between the previously-calculated flat regions and reconstructs the distorted baseline associating a new ADC value to each strip. The baseline is then cleaned and subtracted from the raw data.

The Common Mode subtraction algorithms:

- Median Algorithm: the baseline ADC value is calculated as the median ADC count from all strips in the APV.
- Iterated Median Algorithm: the median ADC count of all strips in the APV is calculated, then recalculated using only strips with an ADC count below the calculated median plus 2-sigma times the strip noise, and then recalculated a third time in the same manner.

In just a few weeks of PbPb collisions, CMS collected about 890 TB of data (roughly 79 M events). During the HI data taking, the new zero suppression algorithm optimized for heavy ions collisions was applied to the prompt reconstructed events. After the run was over, an updated version of the algorithm was applied to the full dataset. A new set of compressed and zero suppressed raw data was then produced. The new collection occupies only 150 TB and was reconstructed for analysis. The new RAW data collection hosts the zero suppressed tracker readouts, as well as the unsuppressed readouts of the modules determined to have distorted baselines on at least one APV. The full dataset reconstruction took 12 days using an average of 3000 CPU cores.

In the plot on the right, the efficiency, fake rate, and momentum resolution of charged tracks using the new tracker zero suppression algorithm are shown.