SELECTIVE READOUT IN THE CMS ELECTROMAGNETIC CALORIMETER

T.Monteiro, CERN, Geneva, Switzerland, on behalf of the CMS ECAL Upper Level Readout and Trigger Group (email:teresa.monteiro@cern.ch)

Abstract

Algorithms suitable to reduce the volume of ECAL data passed to the CMS data acquisition (DAQ) system are investigated in terms of their performance. Various implementation scenarios are analysed which demonstrate the feasibility of the selective readout (SR) techniques proposed.

1. INTRODUCTION

The CMS Electromagnetic calorimeter (ECAL) [1] is a finely segmented detector, consisting of approximately 80000 $PbWO_4$ crystals, organised, in the barrel region, in supermodules of 1700 crystals each. Each CMS ECAL upper level readout and trigger card, envisioned as a 9U VME module, is designed to handle 100 crystals. A single VME crate, located in the CMS counting room and consisting of 17 such cards, effectivelly handles an entire ECAL supermodule. The data from these 17 cards are concatenated in each crate, on an event-by-event basis, by the Data Concentrator Card (DCC), and sent to a memory which is read by the DAQ system.

If, for all channels, 10 time samples per channel (i.e., the preamplifier response for 10 consecutive 25ns bunch crossings in a window around the trigger) are to be read for each Level-1 triggered event (L1A), the ECAL will generate 2.4MBytes of data per L1A. To reduce the data volume down to 120kBytes per L1A, as required by the central CMS DAQ system, a data reduction factor of 20 must be achieved. A channel-by-channel zero suppression scheme might not be the best way to fully achieve this reduction. Only the reading of all neighbouring crystals in a shower, at least down to the noise-level of the electronics, guarantees the benefits of the high energy resolution of this detector. Also, the online extraction of an energy value from the crystal time samples would require a very good understanding of the pulse shape and pile-up effects. As the actual number of ECAL crystals containing information of use in reconstructing any given triggered event is small, a substantial reduction in data volume is possible, without any loss of information meaningful to the physics, by means of a selection of the channels to be readout per event.

An algorithmic data reduction in the ECAL has to be based on principles of calorimetric energy measurement. The energy of a showering particle is contained primarily within a compact array of crystals. Therefore, the energy reading in a compact array of crystals is a means to base a selection of crystals to be readout in an event. The region of crystals to be readout needs to be large enough to achieve the limiting angular and energy resolution of the ECAL.

2. TOWER BASED SELECTIVE READOUT ALGORITHM

To motivate the choice of a selective readout algorithm, the requirements imposed by the CMS physics goals to this detector can be summarised as follows:

-Clean identification of energetic (Et>5Gev) electrons and photons, i.e. showers contained within a compact array of crystals. For this purpose the fine granularity of the ECAL, allowing a detailed shower shape analyses, and its high intrinsic energy resolution are equally important.

-Identification of low energy clusters, ≈1GeV, for measurements of particles within jets or to be used in isolation criteria. This task requires the use of the fine granularity available but is less demanding in terms of energy resolution.

-Precise measurement of jets and missing transverse energy which require the readout of the full ECAL, eventually with a coarser granularity.

Considering the above, and keeping in mind that, at the Level-1 calorimeter trigger, the ECAL is organised into trigger towers, consisting, in the barrel region, of $5\eta \times 5\phi$ crystals, the ECAL data can be split into two potential data blocks. On one hand, a coarse-grain data block, consisting of the energy sums of all existing trigger towers, always passed to the data acquisition system; on the other hand, a fine-grain data block, consisting of the time samples of the individual crystals selected by an algorithm.

In the tower based selective readout scheme, proposed below, the energy sums in the trigger towers are used to generate the readout information. A tower transverse energy reading above a programmable threshold Th_{High} will force the readout of 10 samples for all the channels in a 3x3 array of trigger towers centred on the trigger tower of interest. An intermediate tower transverse energy, $Th_{Low} < Et < Th_{High}$, will force the readout of 10 samples for all the channels in that particular trigger tower. For towers with low transverse energy, only the coarse-grain data will be available.

Complementary, zero suppression is applied on a channel-by-channel basis with a threshold close to the noise level, independently of the selective readout algorithm. The zero suppression calculation is based on the peak amplitude of a prompt pulse arriving after the L1A.

3. POSSIBLE SELECTIVE READOUT IMPLEMENTATIONS

3.1 ECAL Upper level readout and trigger system overview

The ECAL upper level readout and trigger boards, called ROSE100 and described in [2], receive data digitised at 40MHz in the calorimeter very front end electronics via serial optical links operating at 800 Mbit/s. These signals are converted from optical to electrical, deserialised, and split into two parallel paths. On one hand, they are pipelined into memories, to wait for a trigger decision to proceed to the data acquisition system. Dedicated point-to-point links allow the transmission of the ROSE100 data from each module to the Data Concentrator Card in the same crate. On the other hand, they are used as input to the trigger primitive generation block, where the total trigger tower transverse energy, among other quantities, is computed.

Due to different flight paths to different parts of the detector plus different link lengths, temperature variations or mechanical stretching of links, the various trigger channels are not necessarily synchronized with each other. The synchronization circuits, placed at the output of the trigger primitive generation block, guarantee that for every subsequent processing stage the data are synchronised and belong to the same bunch crossing. The synchronization method [3] uses the LHC bunch structure to achieve synchronization. The timing adjustments are computed based on trigger data accumulated (histogramed) in each synchronization circuit, that should match the expected bunch structure. A few minutes of running should be enough to determine the settings needed for synchronization, which can be continuously monitored via the same histograms.

3.2 The Data Concentrator card

The Data Concentrator Card performs the tasks illustrated in the functional diagram shown in figure 1: it is responsible for collecting the ECAL event data per readout crate, formatting it and transmitting it to the data acquisition system, via high speed links. The DCC allows data quality monitoring and, optionally, provides stand-alone triggers. Furthermore, the DCC plays an important role in the selective readout implementation, as described in the section below.

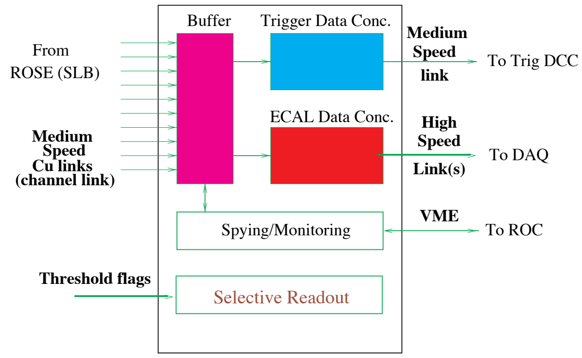

Figure 1: the DCC functional block diagram.

The DCC receives the data from the readout modules and stores it in input FIFOs (iFIFO) controlled by the Input Handlers, one for each iFIFO. The event fragments stored in the iFIFO are assembled and stored in output FIFOs (oFIFO) – the Event-oFIFO stores the data to be sent to the ECAL DAQ path, the Trigger-oFIFO stores trigger primitive data to be sent to the Trigger DAQ path.
The data transfer from input to output FIFOs is controlled by the Data Concentrator (DC) Handler, which receives the selective readout information for each event as well as the L1 trigger. A list of L1 triggers waiting data transfer is kept in the DC Handler. Associated with each FIFO is also a list of events in the FIFO, updated by the relevant readout handlers.

3.3 Selective Readout implementations

In the present design, the energy of each trigger tower is compared with two (programmable) thresholds, and a 2-bit status word ('threshold flag') is set. This calculation is executed as part of the trigger logic and runs synchronously at 40MHz. These bit pairs are transferred on a dedicated bus to the corresponding DCC. Due to the selective readout algorithm presented, in which a particular trigger tower might force the readout of all crystals from a neighbouring one, there is a need to communicate selective readout information across crates, in order to cover all the supermodule borders.

A possible implementation scenario is that the DCC acts as a local 'reflector', receiving the 2-bit words from boards within its own crate and from neighbouring DCCs, and transferring back to the boards the results of the tower selective readout decision, after performing the algorithm. By using the CMS Level-1 processing latency to exchange the selective readout information, readout may commence immediately upon the receipt of a L1 accept signal. This implementation is represented in figure 2.

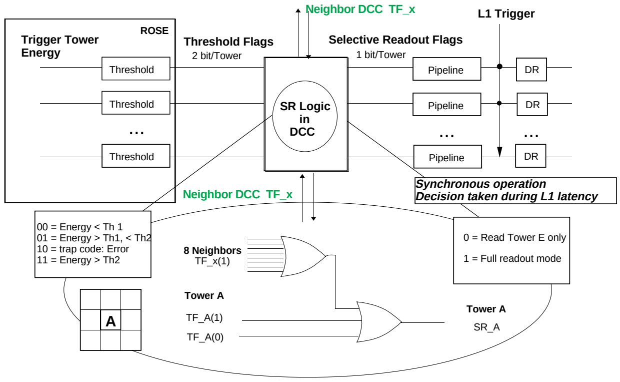

Figure 2: Schematic representation of a possible tower based selective readout implementation.

Another possibility, currently under investigation and depicted in figure 3, is to implement a dedicated Centralised Selective Readout Processor (CSRP) in order to facilitate the information exchange across the hardware boundaries. The DCCs will identify the seed regions and communicate them to the CSRP where a map is created of the channels to be read. This information is returned to the DCCs where the selective readout is performed.

This scheme would allow for more flexibility of new selective readout algorithms, including the eventual use of long range correlations.

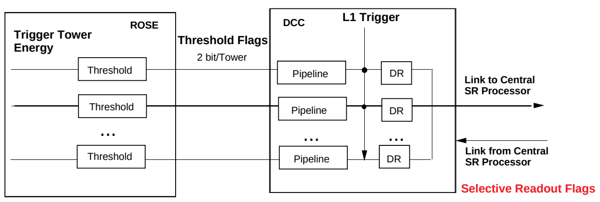

Figure 3: Schematic representation of an alternative, centralised, tower based selective readout implementation

However, care has to be taken to guarantee that the selective readout decision is taken in a fraction of time small compared to the average time between L1 triggers, 100kHz, and also so that dead-time situations are avoided or minimised.

4. TOWER BASED SELECTIVE READOUT PERFORMANCE

The effectiveness of selectively reading out regions of the ECAL was studied with a full-detector Monte Carlo simulation of hard-QCD scattering events with a minimum generated parton transverse momentum of 100 GeV/c and in the presence of both in-time and out-of-time minimum bias pile up (luminosity 10^{34} $cm^{-2}s^{-1}$). A realistic model for the (preamplified) pulse height for a series of time samples and subsequent extraction of an energy value were included, and so was the effect of noise.

This simulation allows the study of the distributions of the number of channels to be transferred from a ROSE board to the DCC, and of the total number of crystals readout per DCC. It also allows the optimisation, in terms of data volume, of the assignment of channels per board and crate in the ECAL endcaps, where the geometry of the trigger towers is irregular.

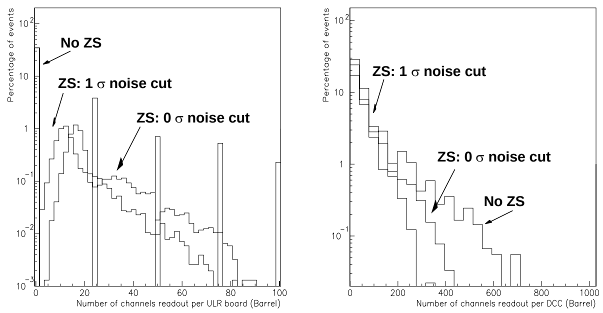

Figure 4: ECAL readout occupancy after a tower based selective readout, in terms of the percentage of events in which a certain number of channels is to be readout per board (left) and per crate (right). See text for details.

In this studies the thresholds were set to Th_{High}=2.5Gev and Th_{Low}= 1GeV. The use of thresholds in Et rather than in energy insures that similar values can be used both in the barrel and endcap regions.

Results of the simulation are shown in figure 4 for the cases where no zero suppression was applied and where complementary channel-by-channel zero suppression was applied with a threshold equivalent to 0 or 1σ of noise (30MeV in the ECAL barrel).

The overall percentage of channels readout out for the ECAL barrel and endcaps are given in table 1 – in the case where no zero suppression is applied, the average event size per DCC will be of the order of 2.5 Kbytes (10 samples per channel, 3 byte format), to be read at a L1 event rate of 100kHz.

Table 1: Data volumes after tower based SR

Zero Suppression	ECAL Barrel	ECAL Endcaps
None	4.9%	3.9%
0σ	2.9%	2.5%
1σ	2.2%	1.9%

5. CONCLUSIONS

The tower based selective readout scheme presented, inspired by the characteristics of an electromagnetic deposit in the ECAL and by the physics requirements of this detector, allows the desired ECAL data volume reduction, a factor ≈20 with respect to the total data produced by the detector. The impact of such algorithm in particular physics channels in still under investigation, but it is believed to be quite small.

The tower based selective readout algorithm uses the level 1 ECAL trigger information, which is available in the ECAL upper level readout and trigger system.

Two possible hardware implementations for this (or similar) selective readout algorithm were described, with particular emphasis on the role of the Data Concentrator Card, demonstrating the feasibility of this technique.

Prototypes validating the concepts discussed here are expected by the summer 2000.

6. REFERENCES

1. CMS collaboration, Technical Design Report, CERN/LHCC 97-33.
2. B. Lofstedt, these proceedings.
3. J.Varela, CMS note-96/011.