CMS Internal Note 99-xxxxx
July 29, 1999
Version 1.0

DRAFT DRAFT DRAFT DRAFT DRAFT DRAFT

Conceptual Design
CMS HCAL Upper Level Readout System

D. Baden, R. Bard, S. Eno, J. Giganti,
N. Hadley, G.W. Sullivan
University of Maryland

J. Elias
Fermi National Accelerator Laboratory

E. Halzen, E. Machado, L. Sulak
Boston University

Abstract

This document serves as a Design Report for the HCAL upper level readout, defined as that set of electronics (and functions) which receive (but does not inlcude) QIE fiber output in preparation for Level 1, Level 2, and DAQ processing. This includes all readout cards, trigger primitive processing, formatting, and delivery to the (and including) the HCAL DDU.

Contents

List of Tables

List of Figures

1 Introduction

This document presents the current conceptual design of the CMS Hadron Calorimeter (HCAL) upper level readout system. The upper level readout system is defined as everything between the optical fibers carrying the HCAL data from the detector, and the input to the L2/L3 processor Farm.

Originally, the design of the HCAL readout system closely followed the design of the ECAL readout system, which was based on the FERMI project ASICs. However, considering the current (and projected) availability of large affordable FPGA chips with 10^6 gates and 10^3 I/O pins, we have redesigned the readout and trigger cards. We believe that the current design is much less risky (and more likely to succeed) due to the programmability of the FPGA. In addition, this enables us to merge all readout and trigger preprocessing into a single HCAL Trigger and Readout (HTR) card, which decreases the engineering and production cost (fewer different cards) and allows for fewer channels per card (more cards with readout capability).

2 Overall System Design

The overall system design and its context within CMS HCAL and Trigger is shown in figure 1. Parameters are as in table 1. The full readout system consists of 24 VME crates. Each crate will be responsible for receiving HCAL data, maintaining the 40MHz pipeline, maintaining synchronization via the TTC, and preparing and transmitting the data for both Level 1 and Level 2/DAQ. To do this, the HCAL readout crate will contain 18 HTR cards, 1 in-crate processor card for monitoring and administration (CPU), and one Data Concentrator Card (DCC). Crates will reside in a room near the detector, but below the surface.

Each crate will service a single Device Dependent Unit (DDU) for Level 2/DAQ feeding. The optical fibers supplying the readout crates are not included in the system. Once data has arrived, it is processed for Level 1 and Level 2/DAQ transmission in parallel while maintaining the 40MHz pipeline. Buffer sizes and transmission requirements will all be based around an expected maximum Level 1 accept rate (L1A) of 100kHz.

The optical fibers supplying the readout crates are not included in the upper level system. The trigger links, however, are included, and possibly the HCAL DDU as well although this has not been entirely settled as of this time.

3 Upper Level Readout Crate

The readout crate is mechanically a 9U x 400 mm VME crate. The P3 area is completely user definable, and will be used by HTR cards for transmission of trigger primitive data to Level 1 (described below) via auxiliary backplane cards. We anticipate that we will need the middle row of P2 (for VME 32) for calibration and monitoring

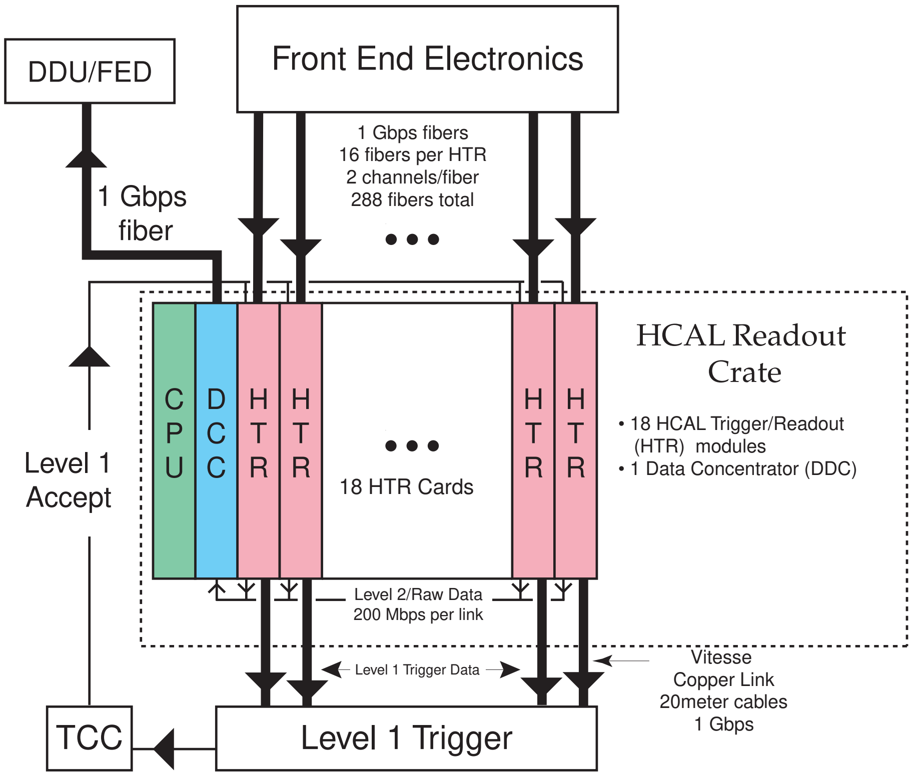

Figure 1: HCAL Readout and Trigger Crate.

I/O. We are investigating bussing the 40MHz pipeline clock (as well as several other signals from the TCC**, but which ones?**), also over P2.

HCAL data arrives at each crate via a maximum of 288 high speed fiber optical links implementing the HP G-links protocol. Each fiber link delivers data from 2 HCAL channels to the HTR card. The crate will contain 18 HTR cards, enough for 576 HCAL channels, 16 fibers and 32 HCAL channels per HTR card. Once arrived, the data is transformed, buffered, and pipelined. While in the pipeline, data is sent through separate Level 1 and Level 2/DAQ paths in parallel. The Level 1 Path prepares the trigger primitives to be sent to the level 1 trigger system over a maximum of 144 links per crate, where a L1 decision is made. The average L1 accept rate is required to be less than approximately 100kHz. If the L1 trigger passes the event, the data stored within the buffers in the Level 2/DAQ Path is collected, processed and transmitted to the DAQ system via DCC input cards.

4 HCAL Readout and Trigger cards

In this section we discuss the HTR card, which consist of an input section, a Level 1 Path, and a Level 2/DAQ Path.

HCAL readout crates	24
HTR cards per crate	18
Input fibers per HTR	16
HCAL channels per fiber	2
HCAL channels per HTR	32
HCAL channels per crate	576
TOTAL HCAL channels	13,824
TOTAL HTR cards	432
TOTAL DCC cards	24

Table 1: HCAL upper level readout system parameters.

4.1 Input Section

The HCAL calorimeter signals are digitized at each crossing (40 MHz). The digital signals for two channels (the two depth segments of a single eta-phi tower) are combined into 16 bits on one optical fiber. The data from each channel consists of an 8 bit number (5 bits of resolution, 3 bits of range). Since there are 4 capacitors per QIE output, a single bit ("cap bit") will be added to encode which of the 4 capacitors the data originates from. Note that since the two combined channels are from different QIE chips, the cap bit will be the result of front-end logic which considers whether both cap bits are synchronized. This "cap bit" will cycle through the sequence 1-0-0-0 every 4 crossings, and will be reset to ensure synchronization every orbit ($\sim 88\mu$sec). If synchronization is lost on either of the 2 channels, the logic will flip the corresponding state. The readout state machine will therefore be able to detect such errors. The fibers will make use of the HP G-links protocol runing in encoded mode with 20 bit frames, 17 of which are data, at 40 Mframes/sec for an total data transmission bandwidth of approximately 700Mbit/sec.

The fibers which deliver data to the HTR cards will plug into the front panel. This scheme alleviates engineering R&D necessary for successful input over the P3 backplane, and takes advantage of the G-links built-in receiving and deserializing capability. As shown in figure 2, the receiver section of the HTR card will add a "frame bit" which will tag transmission errors, and using the 8 bits of data and the cap bit will transform the data from 8-bit non-linear to 16 bits linear. Logic will determine wether the not the cap bit was correct (based upon the beam crossing signal from the TCC).

A PLD will receive the 17 bits of G-links (2 8-bit channels plus a cap bit) data along with the frame bit. A "cap synch error" bit is derived from the cap bit itself, and carried along in the pipeline with the data. The cap sequence ($1^{\rm st}$,$2^{\rm nd}$,$3^{\rm rd}$,$4^{\rm th}$) is turned into a 2-bit number, and each 2-bit number will be combined with the 8 bits of channel data to form a 10-bit "address". This address will be used as the input to a lookup table (LUT) or ASIC, resulting in a 16-bit linear tower energy **(E or ET?)**. Whichever of the two schemes are employed (LUT or ASIC), both will allow the downloading of 1024 patters into the memory and clock them through the system at 40 MHz to perform diagnostic testing of the system.

The 16 bit linear data value plus the two error bits (frame and cap) are sent to the Level 1 and Level 2/DAQ paths for processing.

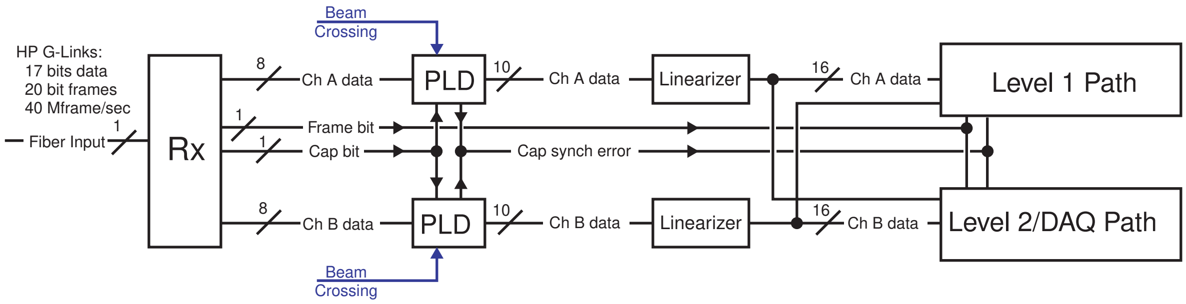

Figure 2: HTR input section.

4.2 Level 1 Path

The L1 tap needs to be processed and summed with its other depth segments for that eta-phi tower before being sent to the L1 trigger. The HCAL signals are collected over a period of 5 crossings. The first two are for baseline subtraction, the last three are the time it takes to collect the energy. The tapped L1 signals must be processed to giver the energy in the single beam crossing which created the energy. The MCM contains a simple processor for this purpose. However, we should develop an algorithm for implementation in a gate array as an alternative should we not use the MCM. The summing is done by a combination of the readout card and trigger card. For the barrel, there are only two depth segments for each tower. This sum is performed on the readout card. For the endcap calorimeters there can 2,3 or 4 segments which need to be added. No more then two are done on the readout card, further summing is done by the trigger cards. The forward calorimeter is even more complex. After summing, the 10 bits with least count of .5 GeV are picked out of the 16 and sent to the adjacent trigger card by parallel connection across the front panel.

The Level 1 Path is fully pipelined with no buffering. **(is it here that we would use the error bits to zero the energy before summing? after summing?)** As shown in figure 3, channels which correspond to the two HCAL depths in a single HCAL tower will be added together into a 17-bit data word and passed through a Level 1 "filter" which will determine what energy to be associated with what beam crossing, producing a 16-bit word. The filter will implement a scheme involving the energy in five consecutive buckets. The most likely scheme will involve adding the energy in the 5 towers A, B, C, D, and E with the weights such that the output will be given by

$$\text{Output} = 1.0 \cdot (A + B + C) - 1.5 \cdot (D + E),$$

thus integrating over the three buckets which had real energy, subtracting any common baseline. We plan to implement this filter using an FPGA, since these weights involve only factors of 2. **(but what about the negative sign?)** The resultant 16 bits of data is then assigned to a particular crossing (maintaining the pipeline). A "muon bit", the result of a window calculation for the presence of a muon, is then produced and added to the data as a "feature bit". The 16 bit HCAL tower energy is then subject to a truncation and overflow calculation, eliminating 6 bits, and then sent through a non-linear transformation further reducing the tower energy to 8 bits. **(hey guys, is this energy or transverse energy???)** Along with the feature bit, these 9 bits are combined with 9 bits of data from another tower and merged into a 24 bit Level 1 format for transmission.

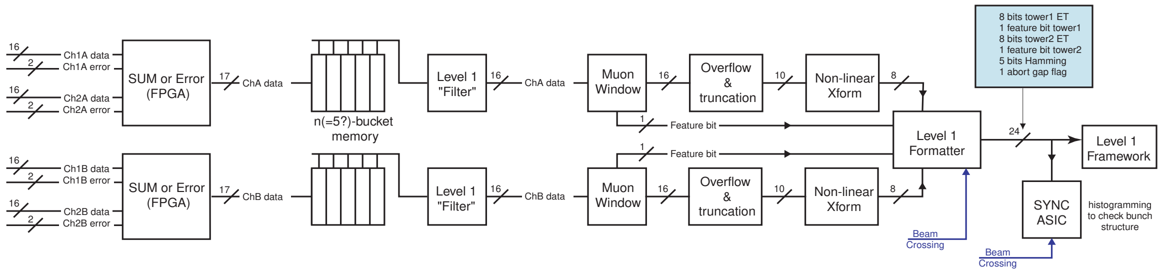

Figure 3: HTR Level 1 Path.

As detailed in table 2, the 24 bits of Level 1 data will include 9 bits each from 2 towers, plus 5 bits of Hamming code for error recovery, plus an addition bit signifying the abort gap. **(i forget why we need to have an abort gap bit? is it so that we can use this for calibration data?**

Purpose	Bits
Tower 1 energy	8
Tower 1 muon flag	1
Tower 2 energy	8
Tower 2 muon flag	1
Abort gap flag	1
Hamming bits	5
Total	**24**

Table 2: Format for the 24 bits of Level 1 data prepared by the HTR cards.

4.3 Level 2/DAQ Path

As mentioned in section 4.1, all data enters both the Level 1 and Level 2 /DAQ paths in parallel, and data consists of 16 bits HCAL energy **(or ET?)** and 2 bits for cap and frame errors. Figure 4 shows a schematic of the Level 2/DAQ path. Data is first put into RAM awaiting a Level 1 decision from the Level 1 framework. **(is it here that we would use the error bits to zero the energy before summing? if so, then do we need to continue to carry the extra 2 bits through to the derandomizing buffers? after summing? how does the level 1 decision get into the card? what kind of protocol? copper or fiber? and what is the nature of the l1 decoder?)** A decoder will determine from the level 1 answer, beam counter signal, and pipeline delay which data is associated with which level 1 accept, producing a address pointer to the RAM for the HCAL channels. An appropriate set of 10 data values will then be moved into a "derandomizing buffer" which serves as a vector buffer holding all 10 data values and associating these values with a particular level 1 accept. This buffer will have to be able to hold 10 channels of 18 bits each and some number of level 1 accepts. The architecture of the HTR card will have all 32 channels' derandomizing buffers read out over a common internal data bus, and passed into a Level 2 filter which has the same function as that for Level 1, with more precision (hence the need for the additional 5 HCAL channels) Level 2 filter constructs the appropriate sum, the data is ready for formatting and sending to Level 2/DAQ. Each of these components are discussed below.

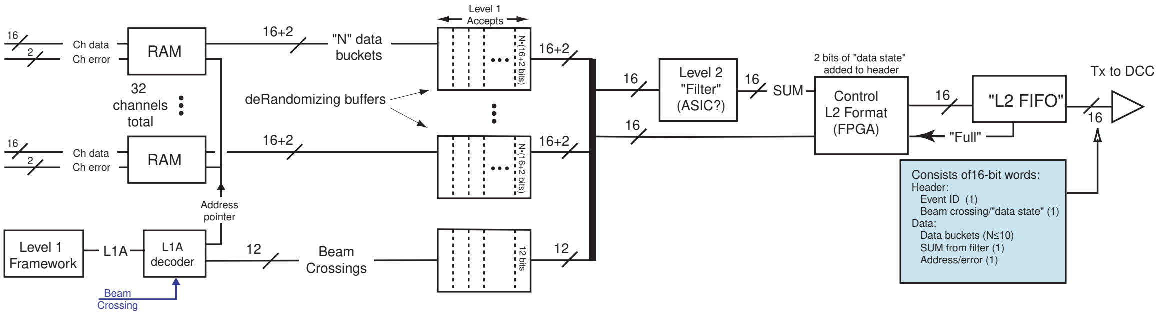

Figure 4: HTR Level 2/DAQ Path.

4.3.1 Derandomizing Buffer

4.3.2 Level 2 Filtering

As discussed above, the Level 2 filtering is expected to be more precise than that use for Level 1. The filter will use as its input the values of a single channel for some number of buckets, and combine them into an energy **(ET?)** and associate that energy with a particular beam crossing. The filter will have the form:

$$\text{Output} = \sum_{i=1}^{N} w_i D_i$$

where N is the number of channels involved in the sum (*e.g.* 10), w_i corresponds to the weights and D_i corresdponds to the 16 bit sum in channel i. For instance, for the Level 1 filter, $N = 5$, $w_1 = w_2 = w_3 = 1.0$ and $w_4 = w_5 = -1.5$. A Monte Carlo study will determine the optimum weights and number of channels. Note that we plan to implement this filter using the ECAL ASIC **(i need references for this!!!)**

4.3.3 Level 2/DAQ Control and Formating

5 DCC Module