

An Update of the 2x2 Implementation for the Level 0 Calorimeter Triggers

Christophe Beigbeder, Dominique Breton, Olivier Callot, Philippe Cros,
Bernard Jean-Marie, Jacques Lefrançois, Ioana Videau

Laboratoire de l'Accélérateur Linéaire (LAL), Orsay

1 Introduction

The 2x2 implementation for the Level 0 calorimetric triggers is described in Refs. [1, 2]. The purpose of this note is to present a detailed description of the data-flow and hardware implementation in answer to the questions of the Trigger Coordinator Hans Dijkstra. System issues are also addressed, such as the synchronisation of signals from the various detectors, the treatment of a "loss of clock" condition and aspects relating to the monitoring and debugging of the functioning of the trigger. We also provide details of cost estimates and the total latency of the trigger. Finally the principal milestones for reaching a final design are given in order to meet the deadline for the Technical Design Report (TDR) by the end of 2001.

2 Front-end Electronics and Detector Geometry

In this section we present the current status of the front-end electronics for the subdetectors involved in the calorimetric triggers, as well as their geometry, i.e. dimension and pad size. These parameters are only relevant to the trigger implementation because they play a role in the data-flow and hardware implementation.

2.1 Preshower, ECAL and HCAL

The Level 0 high P_T electron and hadron triggers are based on input from four subdetectors: the Pad Detector, the Preshower and ECAL for the electron trigger, HCAL and ECAL for the hadron trigger. The Pad Detector is discussed in Section 2.2. The Preshower, ECAL and HCAL have similar projective geometry and similar front-end electronics. The photomultipliers are mounted on the detector. The analog signals are sent over 10 m coaxial cables to the front-end cards. It was decided to use 9U VME boards in order to comply with the LHCb standard. Each board can handle 32 channels, instead of the 16 described in the Technical Proposal. For the Preshower it is foreseen to handle 64 channels on each board.

The ECAL and HCAL analog signals are shaped and digitized at 40 MHz, yielding a 12-bit output for the readout. For the trigger input, they are converted to an 8-bit word in a 4K RAM, which can perform the necessary coding, conversion or correction (e.g. transform E to

E_T). The ADC of the Preshower provides 10 bits for the readout. This value is compared to a threshold, providing 1 bit per channel for the trigger. All signals inside the Preshower, ECAL and HCAL are available at the same time (synchronous clock). However there are differences in timing between the three subdetectors.

The preferred solution for the location of the front-end crates is on top of the detector. Other less appealing solutions, would be to have the crates shared between the top and the bottom, or located on the side of the detector.

2.2 The Pad Detector

The role of the Pad Detector is to determine whether the electromagnetic cluster, which is giving a trigger signal, is charged or neutral. In the TP we assume that this is performed using the MU1 chamber, i.e. the first station of the Muon system. For the purpose of the electromagnetic trigger this chamber must have the appropriate geometry, i.e. pads projective to those of the Preshower, but displaced by half a cell in the horizontal plane. This is obtained in the TP by OR-ing two pads of smaller size. We also assume that MU1 will provide one bit of data per pad, with a synchronous clock, as in the case of the three other subdetectors. This solution has the disadvantage that MU1 fulfils two different functions, it participates to muon detection and identification on one hand, and is used to identify charged electromagnetic triggers. This may result in contradictory requirements.

An alternative solution would be to add a scintillator plane in front of the Preshower, to play the role of MU1 and having obviously the same readout structure as MU1. This solution will cost additional money, but has the advantage of disentangling the two functions of MU1, thus allowing more flexibility for both the muon and calorimeter systems. It would also make the connection between the Pad Detector and the other detectors involved in the electron trigger easier, since all the electronics can in this case be located on the same platform, avoiding many medium- or long-distance links. In the following we refer to this device as ‘Pad Scintillator’.

As this new Pad Scintillator has not yet been approved, the baseline option adopted in this note is to use the MU1 chamber. Using a Pad Scintillator instead, will not affect most of the trigger implementation. Any differences arising from this alternative approach will be mentioned whenever relevant.

2.3 Geometry

The exact geometry of the detector is not very relevant for the description of the proposed trigger. So far we have used the TP geometry, that has cells of 4, 8 and 16 cm in ECAL (labelled ‘4-8-16’). This geometry leads to 5952 channels for the Preshower and ECAL, and 1488 cells for HCAL (exactly 1 HCAL cell for 4 ECAL cells). With this geometry the ECAL readout requires 202 cards (some of which are incomplete) in 14 crates, and the HCAL readout has about 50 cards in 2 crates.

The calorimeter group is currently considering a ‘4-12’ geometry with only two cell sizes, nominally 4 and 12 cm, for a total of 5632 ECAL cells. The important point is that **the cell sizes of the various detectors are projective**, which means that photons originating at the interaction point will cross the same cell number in all detectors.

The proposed ECAL geometry is displayed in Figure 1. The proposed HCAL geometry is no longer a scaled version of ECAL. Due to the intrinsic size of the hadronic shower, and also due to the construction technique, it is better to have only two cell sizes in the ratio 1:2 for HCAL, called 12 and 24 cm (in fact, in order to achieve projectivity, the real size is 13.5 and 27 cm). The proposed boundaries, as displayed in Figure 2, allow an easy connection between ECAL

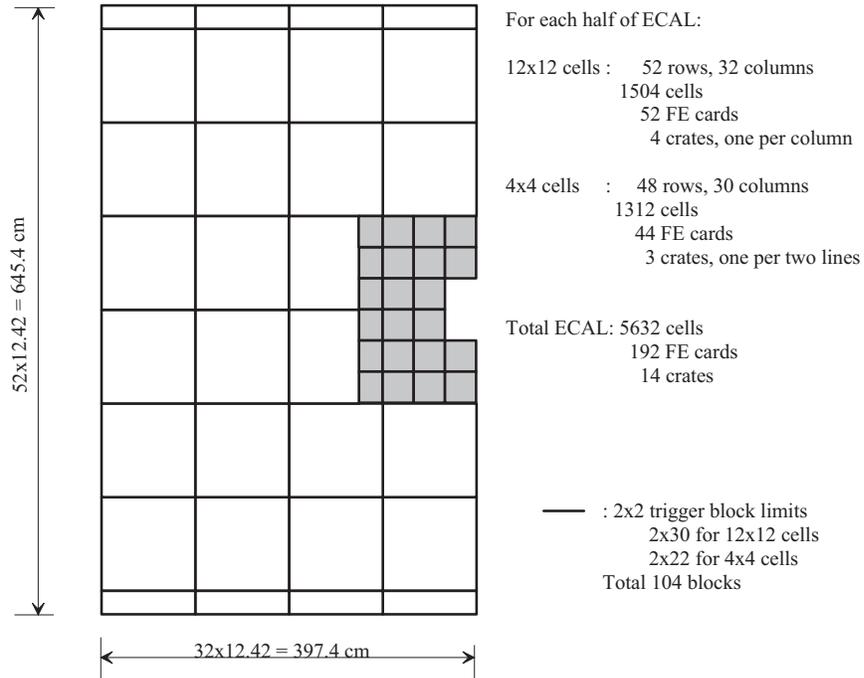


Figure 1: ECAL layout with the new geometry

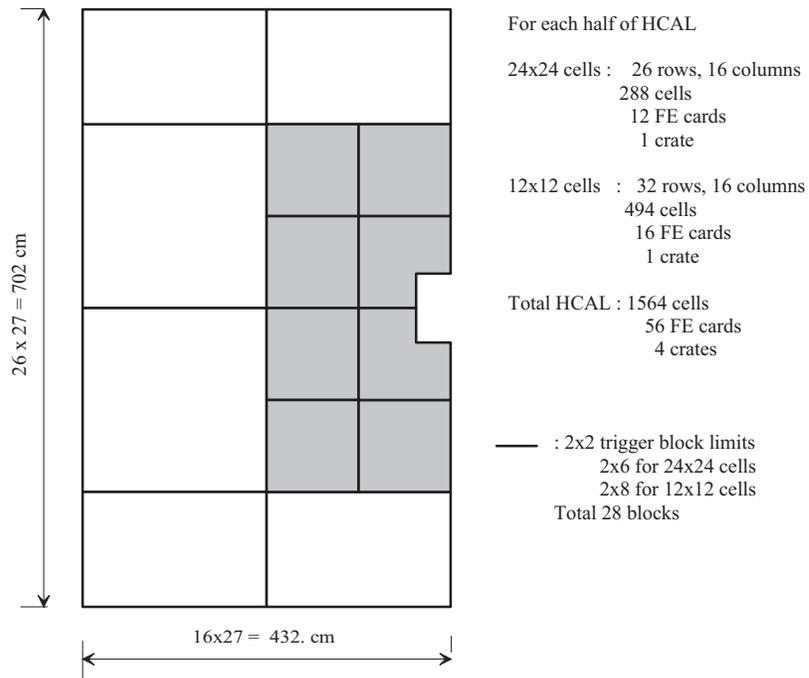


Figure 2: HCAL layout with the new geometry

and HCAL trigger blocks, but now the central HCAL trigger blocks will be connected to many ECAL blocks, up to 8. The number of cells in HCAL is in this case 1564. This geometry is still subject to changes, in particular the exact boundary between cells of different sizes is still being optimized.

3 Dataflow and Hardware Implementation

As described in Ref. [1], the underlying idea of the 2x2 trigger is to use the sum over groups of 2x2 cells, called **clusters**, to estimate the shower energy, and to search for local maxima of energy deposition in large **blocks** of 8x8 cells. This approach has the advantage of decreasing significantly the number of high-speed connections and also of selecting at an early stage a small number of candidates that need further processing. Appropriate validation using Preshower and Pad Chamber information is applied to this small number of candidates only, and after the final selection the result is sent to the Level 0 Decision Unit. The overall scheme is depicted on Figure 3 for the ECAL case, the HCAL one being very similar, except for inputs and outputs of the Summary Card and for the signal sent to the L0 Decision Box.

This scheme can be implemented by integrating the first step of the trigger on the front-end boards, thus minimizing again the number of connections. A first level of processing is performed on the board, then subsequent work is performed at the crate level in a ‘Summary Card’. Finally the trigger candidates are obtained in a single ‘Selection Crate’, which sends the results to the Level 0 Decision Unit.

Each front-end card hosts the cells of 4 columns of 8 rows, therefore a block is built of 2 cards. A crate contains up to 16 FE-cards, i.e. 8 blocks. In order to build the 2x2 sums, one needs access to the neighbouring cells on top and on the right side. This is trivial when the cells are located on the same board. The cells located on the neighbouring card in the same crate are accessed via the backplane, using point-to-point dedicated short lines which introduce no delay, i.e. the data are available at the same time as when they are on the same board. For the up neighbours, a short distance (1 to 2 m) connection between boards is needed, and is performed by twisted pair cables without multiplexing, allowing the data to be available in one clock cycle. A board receives 4x8 bits from the top and send 4x8 bits to the bottom, which requires 64 twisted pairs. These connections will be located on the back of the card, since the front of the card handles the analog signals (Shaper and ADC), and we don’t want to have digital lines too close to the analog inputs. The connections between front-ends cards are shown schematically in Figure 4. Note that the meaning of ‘top’ and ‘right’ assumes that a card holds a vertical slice of 4 columns and 8 rows. For the outer part of ECAL, it seems more convenient to put an horizontal slice in each card, and then one has to rotate Figure 1 to find what is ‘top’ and ‘right’. This rotation allows to have more links on the backplane, and less cables between crates.

3.1 On the front-end card

The first step is for the available (on board and right neighbours) signals, to wait for the up neighbours to arrive. Then the 2x2 sums are performed in two steps, first the vertical sum then the horizontal sum of the vertical sums, keeping the result on 8 bits with saturation at 255 if needed. The highest of the 32 clusters on the card is selected, in two steps: highest cluster in each column, then the highest cluster of the 4 columns. The output of the card is one value of E_T and its address in the card, on 5 bits, 13 bits in total. The whole operation is performed synchronously in pipeline mode. This is done using 5 FPGA’s, one per group of 8 channels, plus one for the inter-group selection. It takes 5 clock cycles to obtain this value. The front-end

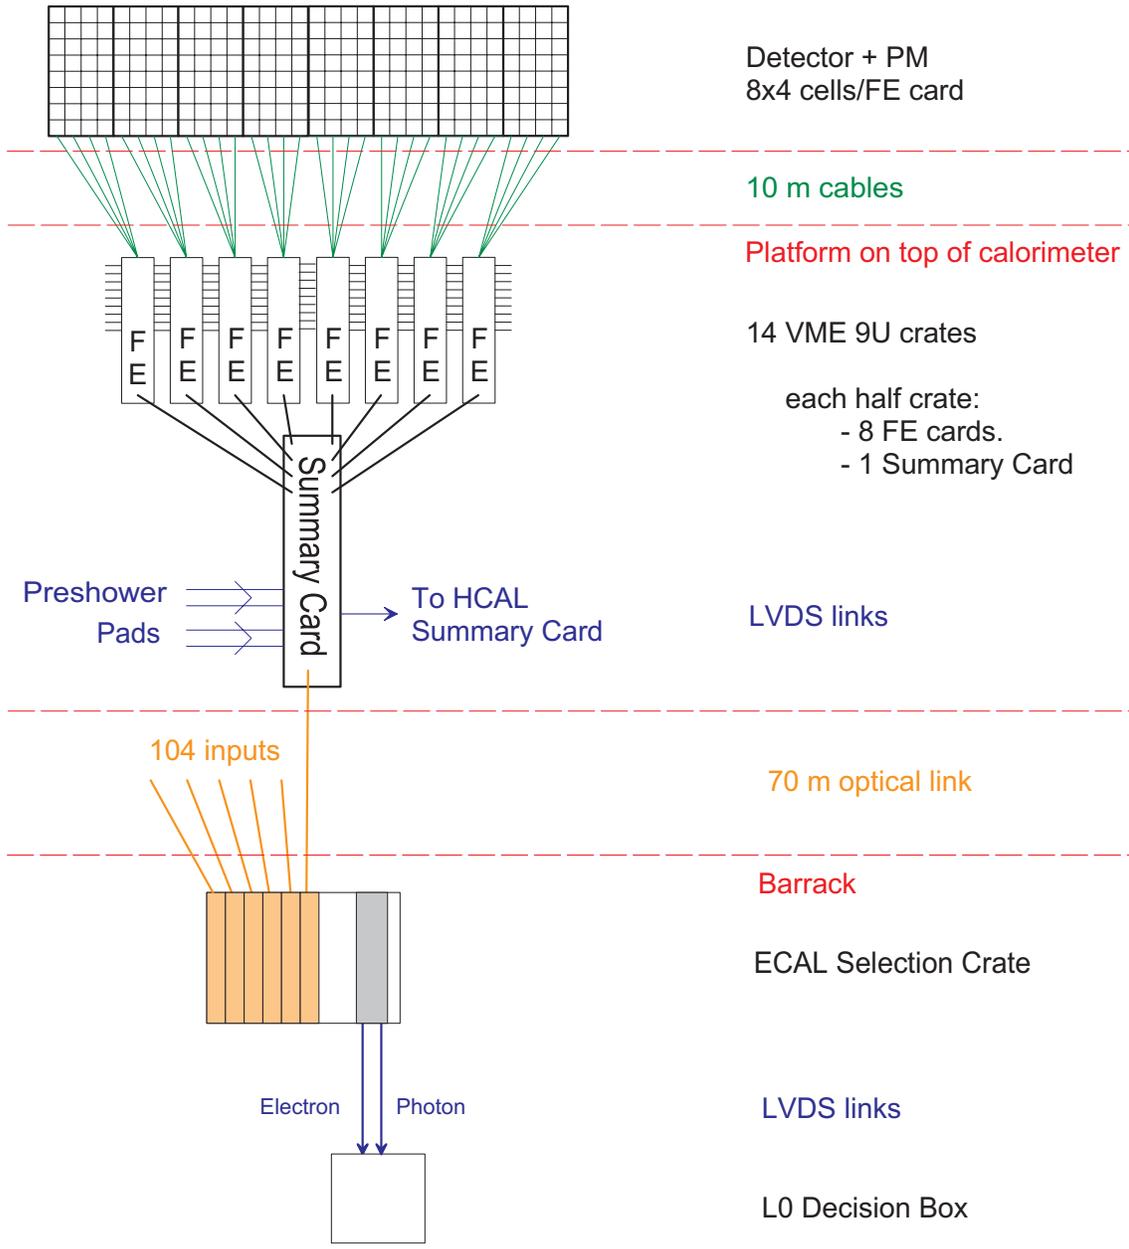


Figure 3: Overall view of the ECAL 2x2 trigger implementation.

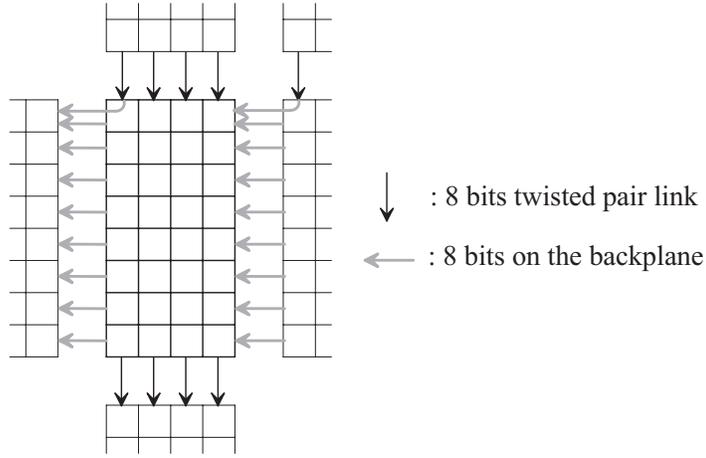


Figure 4: Schematic view of the links between front-end cards. Note the longer path for the neighbour in the ‘up-right’ corner, which is sent via an up link and then arrives on the backplane.

card, including the trigger part, is identical for ECAL and HCAL.

3.1.1 Backplane layout

Each card uses 32 pairs for sending and 32 pairs for receiving the vertical neighbours. It has also 9x8 backplane lines for sending and 9x8 lines for receiving the horizontal neighbours. It also needs 13 lines for sending the 13-bit results to one of the two Summary Cards of the crate, using point-to-point lines. The Summary Cards are located in the center of the group of 8 front-end cards to which they are connected, therefore the space required on the backplane corresponds to 4 times 13 lines. There is then enough backplane space left for the readout needs. The backplane of ECAL and HCAL crates is in principle identical. The only difference may occur at the location of the Summary Cards, since they have different external inputs. However, we will try to have a unique design for both types of crates.

A schematic view of the connections needed on a front-end card is shown on figure 5.

3.1.2 Position estimate (optional)

If required by the Level 1 trigger, it is possible to provide a more accurate estimate of the cluster position, derived from the signal ratio in the horizontal and vertical directions. This means that the two sums per column and the two sums per row (4 times 8 bits) have to be computed for each cluster, kept during the selection process, and then presented to two (horizontal and vertical) Look Up Tables (LUT) after the last operation on the front-end card. From the two values (8 bits) of E_T in the two rows (respectively two columns), each LUT produces a relative vertical (respectively horizontal) position inside the cell for the highest cluster in the card. Each LUT has 64K 3 bit values. The result of each card takes in this case $13+2 \times 3=19$ bits, and requires more backplane space. The extra hardware is most probably only the two LUT’s, the rest should fit inside the FPGA used for the normal processing. The 6 bits of position information are transported unchanged in the rest of the processing, up to the Level 0 Decision Unit, which passes them to the Level 1 processors if needed. This operation adds one clock cycle (for LUT access) to the estimated latency in the front-end card. The size of each intermediate result has

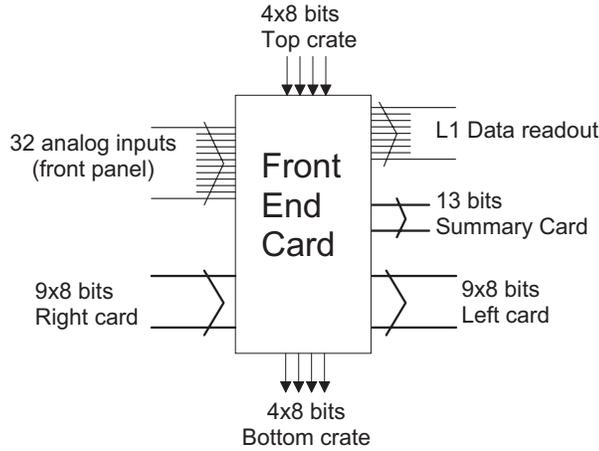


Figure 5: Schematic view of the connections on a front-end card

to be increased by 6 bits, which is indicated in parenthesis in the following.

3.2 The Summary Card

In principle the processing on this card is simple: It handles 4 blocks independently, and for each block it selects the cluster with the highest E_T from the two front-end cards, and sends it to the Selection Crate with one more address bit. However, the Summary Card also collects the information from the other detectors. This functionality is more complicated and it is different for ECAL and HCAL. Therefore, while the front-end cards are identical for ECAL and HCAL, there are two types of Summary Cards.

3.2.1 ECAL Summary Card

This card must verify the electromagnetic nature of the selected clusters, and provide the electron and photon signatures. This is performed using the Pad Detector and Preshower information. An electron has one of the 4 corresponding Preshower cells above threshold, and the Pad Detector area corresponding to this cell has a hit. The correspondence is 2 pads per Preshower cell, covering the same height and twice the width, since pads are shifted horizontally by half a cell. The Summary Card receives then 9x9 Preshower cells and 10x9 pads for each block. This means 81+90 bits, and will typically be provided over LVDS short distance links, as the Preshower (and Pad Scintillator) electronics sit on the same platform. This distribution of the pad signals is delicate for the MU1 chamber, as there is a big problem in dispatching the proper signal to the proper destination (12000 channels patch panel ?), and as the distance between the MU1 electronics and the ECAL Summary Cards may be larger than 10 m. A Pad Scintillator detector is by far simpler to handle.

The first operation is to build a 9x9 pad validation array, by OR-ing two physical pads and requiring a coincidence with the Preshower cell. Then a second step is to build 8x8 validation tables, by computing the OR of each 2x2 area, for the Preshower cells and for the validated pads. Lastly, one has to extract the bit corresponding to the address of the cluster, in the two tables, to obtain the ‘electromagnetic’ and ‘electron’ bits. If the electromagnetic bit is not set, the cluster is ignored by the Selection Crate, i.e. its E_T can be zeroed. The ‘electron’ bit is

added to the information passed to the Selection Crate.

The ECAL Summary Card also sends the address and the E_T of each cluster to the HCAL Summary Cards. This means 4 times 14 bits per Summary Card, over medium distances since the HCAL summary cards should be less than 10 m away from the ECAL ones, therefore one can use LVDS links.

The output to the Selection Crate is 4 times 15(21) bits, as it includes the electron bit. Note that there is still no need to add a Bunch ID, as we are completely synchronous, nor a block ID, as each block is transmitted over its own cable.

To conclude, the ECAL Summary Card requires 1 cycle to select the highest of the two front-end card candidates, and 3 cycles to obtain the electron bit. This card must also synchronize the information from the three detectors, by delaying (pipe-line) at least two of them. The amount of information to handle is not too big, and this delay could be performed inside the FPGA's. We expect the processing of one block to fit in a single FPGA, therefore the card contains only 4 FPGA's, plus the LVDS receivers and the drivers for inputs and outputs. The total input is 8 times 13(19) bits on the backplane, and 4 times 171 bits on LVDS links from the Pad Detector and Preshower, which means about 4 times 20 pairs. The outputs are 4 LVDS links to HCAL, and 4 optical links to the Selection Crate. A schematic view of the connections needed on the ECAL Summary Cards is shown on Figure 6.

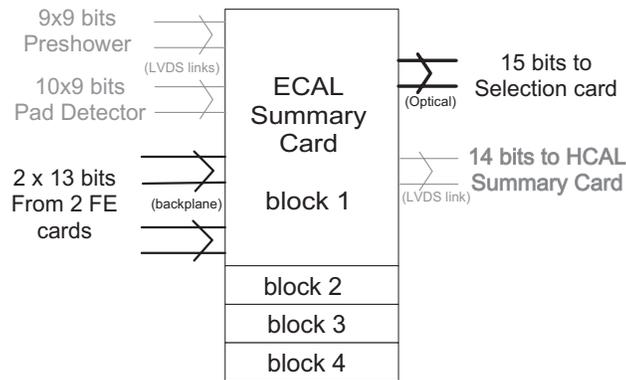


Figure 6: Schematic view of the connections on the ECAL Summary Cards. The connections are shown only for one of the 4 blocks.

3.2.2 HCAL Summary Card

This card performs the selection of the highest energy cluster in each of the 4 blocks it handles independently, and adds the corresponding ECAL energy. For this purpose it receives 14 bits of information (address plus E_T) from the related ECAL blocks. As can be seen from Figures 1 and 2, one HCAL block is related to one or more ECAL blocks the maximum being 8. An ECAL cluster is added to an HCAL cluster if their address matches, i.e. if the ECAL cluster is in front of the HCAL cluster. This is performed by look-up tables. Since the matching addresses can be found in different ECAL blocks, we select only the highest of the matching ECAL clusters before adding the ECAL and HCAL E_T . Typically, the matching of addresses takes one cycle, the selection of the highest in the up to eight ECAL clusters takes another cycle, and the addition a third cycle. Here also, the relative timing of ECAL and HCAL should be adjusted by appropriate pipe-lined delays on one of the two signals, but as this is essentially

at the same step in the signal processing chain, the relative delay should be no more than one or two cycles. The card has 8 times 13(19) bit inputs on the backplane, up to 8 times 4 LVDS inputs from ECAL, and 4 optical output links to the HCAL Selection Crate. A schematic view of the connections needed on HCAL the summary cards is shown on Figure 7.

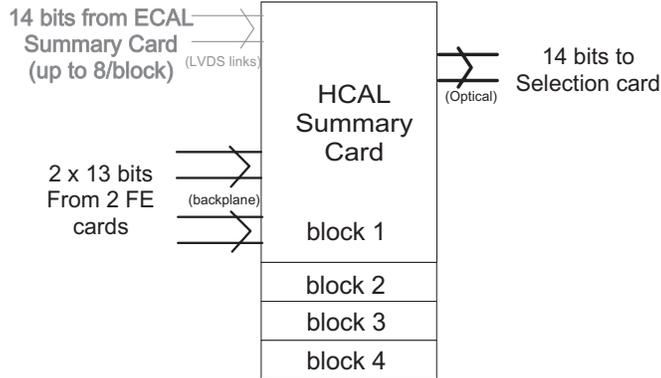


Figure 7: Schematic view of the connections on the HCAL Summary Cards. The connections are shown only for one of the 4 blocks.

3.3 Selection Crate

The Selection Crate handles the selection of the interesting clusters. The proposed method is to find the value of E_T for which there are no more than ‘n’ clusters having an equal or higher E_T , where ‘n’ is a small number, which depends on whether we select independently electrons and photons or not. A possible value is $n=4$. The selection method is described in [2], and requires 9 clock cycles to find and memorize the candidates. Note that a minimum value of E_T should also be required, typically the lowest trigger threshold for the given type of particle. Therefore the average number of selected clusters is very small compared to one, typically of the order of the Level 0 rejection rate, let’s say $1/20$. It is important to note that up to this step, the whole processing is synchronous, and does not depend on the history. At this stage, one sends the candidates to the Level 0 Decision Unit, **ordered** by beam crossing, but not sorted inside one crossing. The average rate is one candidate every 20 cycles, and the maximum rate would be 4 candidates for one cycle (for ‘n’=4), even in the case of catastrophic events. Note that we do not send data to the Level 0 Decision Unit if there is no candidate for a given beam crossing. ECAL and HCAL Selection Crates are independent. The ECAL crate will produce Electron and Photon candidates at the same time. If convenient, one may use separate links to send the trigger candidates for electrons and photons to the Level 0 Decision Unit. Note that the block address should now be combined with the 6 bit internal address to obtain the global address of the selected clusters. The amount of data per candidate is 13 address bits, 8 E_T bits, 1 ‘electron’ bit, optionally 6 position bits, and 8 bits for the beam crossing ID, giving a total of 30(36) bits.

The method described in [2], which uses an analog line to count the number of candidates, could be replaced by a simpler, more brutal method: one could apply a threshold and output the first ‘n’ candidates over this threshold. This would bias events with high multiplicity, but if it is applied on already validated electrons or photons, it should not create trigger inefficiencies. The only loss is due to rejection at a later stage, for example if all the selected electrons do not fulfil the Level 1 track trigger while another block in the event would have passed this cut.

The Selection Crates have to be in a location accessible from both halves of the calorimeters. A solution could be the balcony, but a better solution is to put these crates in the barracks, as this will make the electronics accessible.

3.3.1 Ghost cleaning

In the TP, the calorimeter triggers use only the highest E_T . If we want to use also the second highest, then removal of ‘ghosts’ is required. As the same calorimeter cell can participate to several blocks, a very high cell can produce 2 or 4 candidates, in 14/64 and 1/64 of the cases respectively. The ‘ghost’ cleaning consists of removing the lowest of two candidates if they share one cell, i.e. if their address in both x and y differs by 0 or 1. Part of this cleaning can be performed in the Summary Card, as horizontal neighbours are available on the same board. Vertical comparison is difficult to perform before the Selection Crate. One could imagine to put the vertical neighbours on the same card in the Selection Crate and perform the second half of the ghost cleaning there, before the selection of the highest blocks.

3.4 Hardware summary

The hardware required for this implementation is summarized in Table 1.

Item	Quantity	Remarks
(These items are needed for the readout and are only slightly modified)		
9U VME crates (front-end)	26	14 ECAL + 4 HCAL + 8 Preshower
Calorimeter FE cards	248	192 ECAL + 56 HCAL (identical)
Preshower FE cards	96	64 channels per board
ECAL summary cards	28	
HCAL summary cards	8	
Selection crates 9U VME	2	
Selection cards	18	14 ECAL + 4 HCAL (8 channels each)
Selection controller cards	2	One ECAL and one HCAL
Preshower to ECAL links	9x104	9 bit LVDS links, 9/block, 104 blocks
Pad Detector to ECAL links	10x104	9 bit LVDS links, 10/block, 104 blocks
ECAL to HCAL links	104	LVDS 14 bits
Summary Cards to Selection Crates	104+28	1Gbit/s long distance links
Selection Crate to Level 0 Decision Unit	3	1.5 Gbit/s links, per type of particle

Table 1: Hardware for this implementation

4 Cost estimate

The cost estimate presented in this section is based on 1999 prices for existing or announced commercial components. No spares are included in the cost.

Since part of the trigger hardware is integrated on the front-end electronics boards, we have counted the cost of the specific components, and separately the fraction of the cost of the PC board used for these components. The short-distance links between the front-end boards of different crates (Fig. 4) are included in the cost as well as the price of the specific backplane

connections. The Summary Cards are included in their totality. The cost of the specific part of the backplane is estimated but no additional cost for the crates is counted. The results of this exercise are given in Table 2.

The high-speed links have been identified, their length and band-width are summarised in Table 3. No cost estimate is attempted, according to the general agreement.

Item	unit cost CHF	quantity	Total cost kCHF
(Trigger part of front-end boards)			
Hardware components	640	248	160
10% of the PC board	50	248	12
ECAL Summary card	3000	28	84
HCAL Summary card	3000	8	24
25% of front-end backplane	250	18	5
9U VME Selection crates and backplane	13000	2	26
Selection cards	3000	18	54
Selection controller card	6000	2	12
Total			377

Table 2: Cost estimate for the hardware components of the Level 0 Trigger

Item	Comment	distance	quantity
Preshower links to Summary card	9-bit LVDS	5-10 m	9 x 104
Pad chamber links to Summary card	9-bit LVDS	10 m	10 x 104
ECAL to HCAL links	14-bit LVDS	10 m	104
Summary card to Selection Crates	1 Gbit/s long distance	70 m	132
Selection crates to Level 0 Decision Unit	1.5 Gbit/s	10 m	3

Table 3: High-speed links involved in the Level 0 trigger

5 Asynchronous processing

As indicated in Section 3, the proposed trigger is FULLY SYNCHRONOUS up to the sending of the candidates to the Level 0 Decision Unit. Provided the average number of candidates per beam crossing stays well below 1, there is no problem. With a proper minimum threshold, every candidate will trigger the Level 0, except if rejected by the pile-up VETO. If the pile-up VETO rejects not more than half the candidates, the number of beam crossings producing at least one candidate should be less than 1/20. With an average of 1.5 candidates for beam crossings which have candidates, the occupancy of the output bus stays below 10 %. Fluctuations in the latency due to this bus occupancy is very small, and the probability to wait more than 10 cycles, for example, to output some information is extremely low. Therefore we think that there is no need for simulation of the asynchronous part for the 2x2 implementation.

6 Latency

The latency was estimated and the number given to Ioana Videau before the December 1998 LHCb week. Some slight changes have occurred, all numbers are explained in Section 3. The time taken from the availability of the data in the FE card to the reception by the Level 0 Decision Unit is described in Table 4 and amounts to 1300 ns. Adding the time before (233 ns) and after (750+t), according to Ioana's presentation, the 2x2 total latency is $\sim 2300+t$ ns, 't' being the time allowed for a smart Level 0 decision.

Step	Nb of cycles	source of delay	delay (ns)
Processing in FE card	5-6		150
Summary Card	4 + synchro		150
Sending to Selection Crate		30m + drivers	250
Selection Crate processing	10		250
Latency for extraction	less than 10		250
Sending to Level 0		30m + drivers	250
Total for 2x2 proposal			1300

Table 4: Latency in the 2x2 implementation

7 Synchronization

The basic assumption is that all signals from ECAL (resp HCAL/Preshower/Pad Detector) are in phase after digitization. The front-end electronics includes some clock adjustment so that the ADC strobe can be adjusted channel by channel to compensate the time differences due to either the particle time of flight, which differs by up to 3 ns depending on the location in the calorimeter, or the difference in PM transit time, which depends on the PM high voltage. The same global clock is used after the ADC, so that all signals are exactly in phase when entering the trigger part of the front-end card. As described in Section 3, the processing on the front-end card is performed using the same clock, and therefore all signals are synchronous by construction.

The various detector signals (ECAL/HCAL/Preshower/Pad Detector) are not in phase with each other, due to location, cable length and PM properties. The synchronization between them is performed at the input of the Summary Cards, with two functions: adjust the phase of the signals, and delay some of them by enough clock cycles so that data from the same event are processed together. As this is a constant delay, we don't need to tag the data with a bunch ID in each front-end electronics system and to check the bunch ID at the entrance of the Summary Card. The phase and the delay can be easily computed once the properties of the various PMs and of the final electronics are known. As all signals of a given detector are in phase, the adjustment is only a global parameter, i.e. there are only 3 values to find: Pad Detector vs. ECAL, Preshower vs. ECAL, ECAL vs. HCAL.

7.1 Loss of clock

The whole system is working with a clock, generated from the TTC signals and distributed to all boards in the crate and to virtually all the chips on the boards. Inside the chips (e.g.

FPGA) the clock is distributed to almost all functional units. The effect of ‘missing a clock cycle’ depends clearly on where the clock cycle is missed. It can range from the TTC having missed a cycle, to one gate misbehaving inside an FPGA. Of course, we have not listed nor simulated all possible failures, but they can be classified as follows: either a missed clock cycle is automatically recovered after some delay, or it is not. In this case the question to ask is whether this can be detected and/or automatically recovered.

- In all cases where a chip acts as a hardware pipe line, i.e. physically moves data from one location to another one at every cycle, a missed clock will affect the validity of its output for only a few cycles, the number of cycles being the number of physical moves performed by the circuit. Typically, this is less than 10 cycles as described in Section 3. The system is self-resynchronizing. This is the case for ALL the trigger processing up to the extraction of the results in the Selection Crate.
- In the case a circuit acts as a counter, a missed clock changes the validity of the counter forever, until it is reset. This situation occurs in the Selection Crate, where counters are used to generate the bunch ID. As shown in Figures 5 and 6 of [2], there are independent counters on each card. Clearly the counter value could be distributed on the backplane, as these counters should always have the same value. This will make the system more robust. This counter should always be in phase with the global LHCb bunch ID, and can be re-synchronized using the TTC distributed information. With this re-synchronization, the loss of a clock cycle on the bunch counter in the Selection Crate will affect events only up to the next accepted trigger, as for each trigger the TTC system broadcasts a bunch ID that can be used for re-synchronization.

7.2 Monitoring and debugging

The basic method of monitoring the trigger is to compare the decision of the trigger hardware with what can be computed from its inputs. As all inputs are digital, and readout by a different path, it is possible to exactly predict the answer, i.e. the list of candidates sent to the Level 0 decision box and their properties. This box will store the information for readout by the DAQ system, so that checks can be made online by the monitoring software. A task in the Level 2 (or Level 3) farm can monitor the trigger system, and detect faults. This is the most efficient way to detect problems, as it runs permanently with a 40 (5) kHz input.

For debugging, we foresee several tools. Firstly, the input of the trigger system is a RAM, with plenty of unused location. These can be used to store and send on request a sequence of known patterns to the trigger in place of the E_T computed from the ADC value. We also intend to have FIFO’s dedicated to debugging, logging the output of various elements on each board. This was implemented in the Babar boards built at Orsay, and has proven to be very useful. The filling of these FIFO’s is triggered by an external signal, and they are read out afterwards by the ‘slow’ control port of the card. It is also planned to be able to use these FIFO’s as input, i.e. to load them with a sequence of known values and use these values as input to the next stage of the trigger electronics. The combination of injecting known inputs and of logging intermediate results allows examination of the proper behaviour of each sub-system, and faults to be pin-pointed very efficiently.

A comprehensive check of the behaviour of the various elements of the trigger can be performed during commissioning, and between fills. The duration of such a complete test has not been estimated, but it is likely that 5-10 minutes should be sufficient to detect most hardware problems.

8 Simulation

The 2x2 trigger algorithm is not identical to the one implemented in SICB so far. A detailed simulation has been implemented in the framework of SICB since about one year, but was available only in a private area. Version 117 of SICB will contain the official version of the 2x2 trigger, implementing several hardware constraints: no connection between cells of different sizes, no connection between right and left part of the calorimeter, work with 8 bit E_T , implementation of the trigger blocks as described in Figures 1 and 2.

As the geometry of the calorimeters has changed, and is projective only starting from version 117, we have to wait for enough data to be available before comparing the performance of the 2x2 trigger to the official SICB implementation with the ‘final’ calorimeter geometry. However, a first comparison has already been performed with the TP geometry, with a non projective Preshower and a semi-projective connection between ECAL and HCAL for the hadron trigger. The performance was reported at the trigger meeting during the September 1998 LHCb week. The results for the channel $B_d^0 \rightarrow J/\psi(e^+e^-)K_s^0$ are summarized in Table 5.

Rejection	SICB			2x2		
	E_T	Efficiency		E_T	Efficiency	
		All	Selected		All	Selected
50	1.75	50.2	59.7	1.98	51.6	61.3
100	2.14	39.2	48.8	2.42	39.9	47.0

Table 5: Efficiency vs. rejection for $B_d^0 \rightarrow J/\psi(e^+e^-)K_s^0$, electron trigger

As can be seen from the table, the performance is very similar for both ‘all’ and ‘selected’ events, i.e. events which can be reconstructed and tagged. The error on these efficiencies is a few %, due to the error on the threshold for 1% acceptance defined with 30,000 minimum bias events, and the limited statistics on good events, around 2500 events.

A similar study has been performed for non-selected $B_d^0 \rightarrow \pi^+\pi^-$ and is given in Table 6, showing again similar if not slightly better performance.

Rejection	SICB		2x2	
	E_T	Efficiency	E_T	Efficiency
10	2.06	49.9	2.12	55.9
15	2.34	43.2	2.39	47.4
20	2.51	39.0	2.57	42.0
25	2.65	36.5	2.71	38.3

Table 6: Efficiency vs. rejection for $B_d^0 \rightarrow \pi^+\pi^-$, hadron trigger

9 Milestones

In order to reach a final design for the proposed implementation and to have the TDR ready for the end of 2001, we foresee the following scenario:

- Summer 1999: A prototype front-end card is built for the calorimeter test beam. Although this card will not have the trigger part implemented for the test beam, the possibility is provided to add this part later, for tests in the lab which will take place during the fall.
- End 1999: A complete design of the calorimeter front-end card will be available. The design will include the trigger part, as described in this note. The details on the test and monitoring sub-systems may be finalised later.
- Mid 2000: A complete design of the two summary cards, ECAL and HCAL, will be available. The exact date depends on the choice of the Pad Detector, and on the progress in the design of the Preshower front-end card.
- End 2000: The Selection Crate and its two types of cards are fully specified and a complete design is available.
- Mid 2001: A description of the methods and tools for testing the various cards is released.
- End 2001: Release of the TDR

The above does not represent a detailed work plan, rather it is meant to provide a reasonable estimate of the order of priorities. Work will also continue in parallel on the simulation of the trigger in the LHCb Monte-Carlo, to try to improve it's performance before the design is frozen.

References

- [1] D.Breton et al, 'An Alternative High p_T Electron Trigger for LHC-B'
LHCb 97-021 (3 December 1997)
- [2] The LAL LHCb group, 'An Alternative High P_T Electron and Hadron Trigger for LHCb'
LHCb 98-058 (4 September 1998)