

Muon Detector Front-end Architecture

LHCB Technical Note

Issue: version 2
Revision: 1

Reference: LHCb 2000-17
Created: 30th June 2000
Last modified: 30th July 2000

Prepared By: A. Lai
I.N.F.N. Cagliari
V. Bocci, G. Martellotti, S. Martinez
I.N.F.N. Roma I

Abstract

This document describes the present envisaging of the front-end electronics of the LHCb muon detector. When talking about muon front-end electronics, we consider the whole system from the Amplifier-Shaper-Discriminator (ASD) outputs, up to the trigger and DAQ interfaces. The very front-end is not considered in the document apart from its system's implications. The document will explore different architecture options and illustrates in more detail one particular implementation option, which, as of today, seems the most suitable one.

Document Status Sheet

Table 1 Document Status Sheet

1. Document Title: Muon Detector Front-end Architecture			
2. Document Reference Number: 2000-17			
3. Issue	4. Revision	5. Date	6. Reason for change
Draft	1	25 May 2000	draft
1 st version	1	20 June 2000	1 st complete version
2 nd version	1	30 th July 2000	Suggestions/remarks from muon group members

Table of Contents

1. Introduction	1
2. System overview: requirements, constraints and design guidelines	3
3. System tasks	5
3.1 Logical Channels generation.....	5
3.2 Synchronization.....	6
3.3 Trigger and DAQ interfaces.....	9
4. Choice of a baseline architectural scheme	11
5. Front-end Architecture	14
5.1 ASD boards.....	14
5.2 Intermediate Boards.....	15
5.3 Off -Detector Electronics.....	18
5.4 Detector Control System.....	22
6. Costs	24
7. Conclusions	25
8. Acknowledgements	36
9. References	26

1. Introduction

The main task of the LHCb Muon Detector is to provide the whole information for the $L0(\mu)$. For every Bunch Crossing (BX), the $L0(\mu)$ trigger identifies the muon tracks and calculates their transverse momentum (p_T). The information used by the $L0(\mu)$ trigger consists in a certain number of “Logical Channels”. These are supplied to the trigger circuitry at the BX frequency of 40 MHz and must be time-stamped with the time information about the exact BX they belong to (BX identifier = BX Id).

As of today, the hardware structure of the LHCb Muon Detector is not completely defined up to its more intimate details. However, its basic structure in terms of segmentation, number of channels and signals characteristics was frozen in the so-called “March 2000 layout”. This is enough to start designing a scheme for the downstream electronics system. On the other hand, the $L0$ trigger structure is already defined in good detail [1].

The starting point of the system consists in the outputs of the ASD (Amplifier-Shaper-Discriminator) chips. These are digital signals and are called “Physical Channels”. An up to date distribution of the number of Physical Channels per muon station is given in Table 1.1. The total number of physical channel to be processed is about 150k.

Physical Channels are generally different from the Logical Channels used by the $L0$ trigger. This happens because the segmentation needed by the $L0(\mu)$ trigger is generally coarser than what is possible to obtain on the detector itself, owing to affordable single channel detector capacitance (and noise) and single channel occupancy.

Physical Channels (Number)	Station 1	Station 2	Station 3	Station 4	Station 5	Sum
Region 1 (wire pad)		144	144			288
Region 1 (cath pad)	1152	192	192	288	288	2112
Region 2 (wire pad)		288	288			576
Region 2 (cath pad)	2304	384	192	288	288	3456
Region 3 (cath pad)	2304	1152	1152	576	576	5760
Region 4 (wire pad)	2304	1152	1152	1152	1152	6912
Sum/Quad/Layer	8064	3312	3120	2304	2304	19104
Sum/Quad	16128	6624	6240	4608	4608	38208
Sum	64512	26496	24960	18432	18432	152832

Table 1.1. Distribution of Physical Channels in the 5 stations [2].

As a consequence, Physical Channels are merged together to generate logical ones. Naturally, the degree of merging varies according to granularity and to the geographical position of the channels inside the detector. Table 1.2 gives the Logical Channels distribution throughout the 5 stations. We have about 26k Logical Channels as an input to the L0(μ) and DAQ electronics.

Logical Channels (Number)	Station 1	Station 2	Station 3	Station 4	Station 5	Sum
Region 1(h-st/pad)	576	192	192	288	288	1536
Region 1(v-strip)		144	144	0	0	288
Region 2 (h-st/pad)	576	96	96	96	96	960
Region 2 (v-strip)		288	288	72	72	720
Region 3 (h-st/pad)	576	48	48	48	48	768
Region 3 (v-strip)		288	288	72	72	720
Region 4 (h-st/pad)	576	48	48	48	48	768
Region 4 (v-strip)		288	288	72	72	720
Sum/Quadrant	2304	1392	1392	696	696	6480
Sum	9216	5568	5568	2784	2784	25920

Table 1.2. Distribution of Logical Channels in the 5 stations [2].

Thus, the problem we are faced to is the following: what is the best way to:

- Gather the 150k Physical Channels,
- merge them into Logical Channels,
- align them in time, minimizing inefficiencies due to time misalignment,
- finally, distribute 26k Logical Channels equipped with their BX identifier.

Next chapter gives a more detailed list of the overall system requirements. Chapter 3 analyses the issues associated to the each task to be performed, with a particular attention to synchronization, which turns out to be the most tricky and delicate operation of our system. Chapter 4 explores different implementation options that can be considered. The options are compared and their pros and cons examined. This gives us a baseline solution for system architecture. Chapter 5 is the central part of the document and gives a description of the baseline architectural scheme. The different system stages are highlighted and described. Chapter 6 gives some consideration on system cost. Finally (chapter 7), we give our conclusions and we consider the next steps to be done towards the coming-up date of the Technical Design Report editing.

2. System overview: requirements, constraints and design guidelines.

The requirements for the muon electronics system are here summarized.

- The Muon detector readout is based on a binary answer, that is no ADC and TDC information is required. We will see later that this assumption is not completely true.
- Starting from the 150k Physical Channels, about 26k Logical Channels have to be generated for the L0 trigger and the DAQ. The logical channel generation is actuated as logical ORs (and majority ORs) of selected groups of Physical Channels.
- Before being shipped to the L0 trigger, Logical Channels must be time-aligned. We identify two levels of time alignment:
 - i) The BX alignment, at the level of 25 ns. This corresponds to a BX identifier (BXId) of 8 bits, which must be attached to data being shipped.
 - ii) The Fine Time alignment, inside one BX. This corresponds to find the t_0 for each channel at a level of 3 ns time resolution about, in order to keep inefficiencies due to time misalignment between channels at an acceptable level. The choice of a 3 ns time resolution has two reasons. First, it is fairly comparable to the detector signals jitter. Secondly, it corresponds to dividing the clock period into about 8 parts, that is using a 3 bits TDC. Two bits (6 ns) would be not sufficient. Four bits (1.5 ns) would exceed the intrinsic detector resolution.

Other important requirements are the following:

- The system should also contain the L0 and L1 buffers and the interfaces to the trigger and DAQ systems.
- The system should also take care of writing, reading back and monitoring internal configuration data and in particular those of the ASD boards. These last parameters consist of: discriminator thresholds, which are mandatory, detector environment conditions (as, for example, temperatures), which would be advisable to monitor. Moreover, a reasonable sample of amplifiers's analogue outputs would be very useful for checking purposes. The measured t_0 channel should be written on local dedicated registers for real-time correction at the physical channel level.

In defining our architectural scheme we must consider a number of important constraints. First of all, the technical choice of electronics components is faced to the amount of accumulated dose all around the detector, which puts a severe constraint on the placement of the electronics itself and, as a consequence, on the data-path. Recent estimation of the LHCb detector dose has been given in reference [3]. A safe place for commercial electronics, corresponding to an accumulated dose of 1 krad in 10 years, can be found only in the immediate proximity of the detector itself. If moving inside the detector, radiation hard or at least radiation tolerant electronics have to be used.

All these considerations suggest placing even the early stages of the system off detector, by putting on detector the minimum amount of components (basically only the ASD circuits). This contrasts with the opportunity to reduce the number of channels links exiting the detector (150k LVDS cables) as soon as possible.

The other important boundary condition about the front-end architecture is the already well-established L0 trigger scheme [1].

Finally, we want to point out a number of design guidelines, which, although to some extent obvious, are important to be kept in mind.

In defining our architecture we will try to:

- Minimize the number of different kinds of boards.
- Avoid the use of custom electronics as a first-attempt solution. This is done both for cost convenience and also for practical reasons (development and realization time).
- Avoid the use of radiation-hard technologies whenever possible, unless very strong practical reasons are present (significant increase in system integration and reduction in system complexity).
- Add monitor, control and diagnostics facilities about system operations and functionalities. This is a general concern, which is of particular importance, considering the strong access restrictions we will have with respect to system maintenance.

3. System tasks.

In the present chapter we describe in more detail the operations implied in each of the main tasks of the front-end system: Logical Channels generation, Synchronization and Trigger Interface.

3.1. Logical Channels generation.

The first stage of our system combines the outputs of the ASD chips (~150k Physical Channels, LVDS) into logical channel information (~26k LVDS). This stage has to introduce a minimal additional electronic jitter to the signal. This implies careful board layout and circuit implementation. Moreover, it must be configurable/programmable, in order to avoid designing different circuits for each different topological region of the detector, since the logics to generate logical channels varies according to the channels' position inside the detector.

The generation of Logical Channels has to be realized both by means of Majority ORs (combining different gaps belonging to the same "x-y coordinate") and by means of simple logical ORs, when grouping different Physical Channels onto a Logical one of coarser spatial resolution. The exact grouping is not completely defined at the moment and has still a certain unresolved dependence on the detector technology choices in detector design.

After merging, the original "physical" information would be lost and only the information related to logical channels retained. This is to be avoided for the following reasons:

- i) We loose remote access to the single ASD outputs. This has serious consequences in monitoring and maintaining the system.
- ii) The Logical Channels can be built up from Physical Channels belonging to different physical chambers. The possibility to measure the time jitter and regulate the delay for physical channels is of fundamental importance in time calibration of the system.

Consequently, a suitable way to keep the single Physical Channels accessible even after logical combination must be foreseen. This can be done according to the conceptual sketch given in Fig. 3.1. A group of Physical Channels enter a programmable masking box on one side and a multiplexer on the other side. Whenever needed, the Masking Box allows to select each physical channel one at a time before entering the logic. On the other hand, the Multiplexer allows to cycle among the input channels in order to monitor the single channel activity on a built-in scaler circuit, even during data taking. For this purpose, the registers inside this monitoring circuitry should be accessible via DCS at run time (see paragraph 5.4).

For a correct fine-time alignment, delay adjusting should be performed at the level of physical channels, to correct delays between channels and between chambers, before Logical Channels formation. So a suitable technique of delay adjustment has to be implemented at this stage. This drives us directly to consider the problem of time alignment, which we will address in the next paragraph.

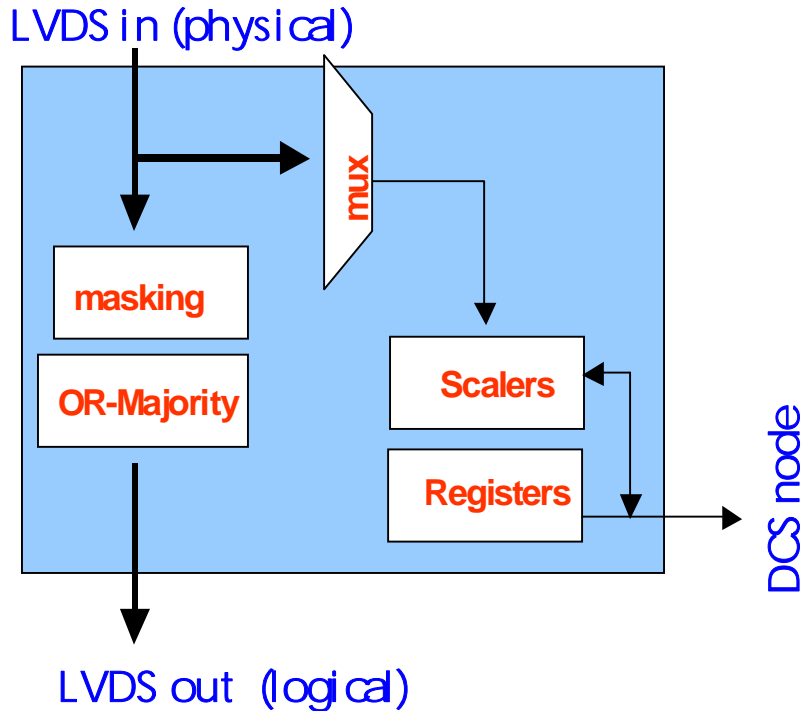


Fig. 3.1. Sketch of logical functions in Logical Channels generation.

3.2. Synchronization.

Before trying to elaborate a suitable algorithm for system synchronization, it is necessary to understand clearly the time behaviour of signals from detector. In particular, the time behaviour of background is of concern in our case. Indeed, it is expected a certain degree of time un-correlation for signals originated by background hits, especially in the last stations. Indeed, while the muon hits are well timed within the 25 ns BX cycle, most of the event hits belong to low energy tracks from showering processes and can introduce signals with long delays.

In order to understand the effects, the relevance of background hits in synchronization operations has been studied by simulation. The time spread of the Logical Channels signals has been simulated. The detector layout used in the simulation is the one described in the March 2000 layout [4] for what concerns granularity and strip implementation. It has been assumed that each station contains only one layer of double-gap chambers.

The time associated to a fired logical channel has been evaluated by taking into account the following effects:

- Time of flight associated to each hit (T_{flight}), as given by the SICBMC program [5]
- Time jitter of the chamber response: it corresponds to the measured time distribution of a single gap wire chamber. The response of the two gaps of the same chamber to a crossing particle have been assumed uncorrelated.

- When forming Logical Channels from Physical Channels, we neglect the jitter which might be introduced by the combination of Physical Channels belonging to the same gap of the same chamber. We have assumed one front-end for the two corresponding pads in the two gaps, and the time associated to a given logical pad corresponds to the first arriving signal.
- When forming logical strips out of logical pads in stations M2 to M5, the corresponding pads have been OR-ed.
- Logical Channels have been time-equalized: in order to time-align muon signals from all chambers and stations, the quantity:

$$T_{chamber}^0 = \sqrt{X_{chamber}^2 + Y_{chamber}^2 + Z_{station}^2} / c,$$

is subtracted from the time-of-flight associated to the hit.

The time equalization (and synchronization) of Logical Channels is assumed to be done with a discrete process with steps of 3 ns. To take into account this effect, we have added a random number generated according to a step function of 3 ns duration.

As mentioned above, most of the event hits belong to low energy tracks from showering processes and can arrive with long delays [6], contributing to the background of subsequent events.

In the time distribution histograms reported in Figs. 3.2 and 3.3, two cases are considered:




Fig. 3.2 Time distribution of hits in the four regions of station M2 for the default background level. Shaded histogram: hits produced by muons from B-decay. Solid line: time measured with a 3-bit TDC for all hits in MB events under standard running conditions. Dashed line: as solid line, but assuming bunch crossings well separated in time.

1. Standard running conditions with continuous 25 ns bunch crossing repetition and a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.
2. Isolated bunch crossings (or interactions at very low luminosity).

The hit time distribution is shown in Fig. 3.2 for the four regions of station M2, and in figure 3.3 for the region 4 in stations M1 to M5. The shaded histograms represent the distribution for hits generated by muons from B-decay. The solid and dotted lines represent the time distribution of all fired Logical Channels in Minimum Bias (MB) events as measured by a 3-bit TDC (solid line for standard running conditions and dotted line for isolated bunch crossings). It can be seen that the peak of the distribution for Minimum Bias is wide and displaced in respect to that for muons. The shape of the distribution changes depending on regions and stations. Under standard running conditions, the distribution for station M5 is almost flat. The situation gets worse with a higher background level and introducing the intrinsic chamber noise (which has been neglected in the simulation, but might be relevant in regions covered by RPC's).

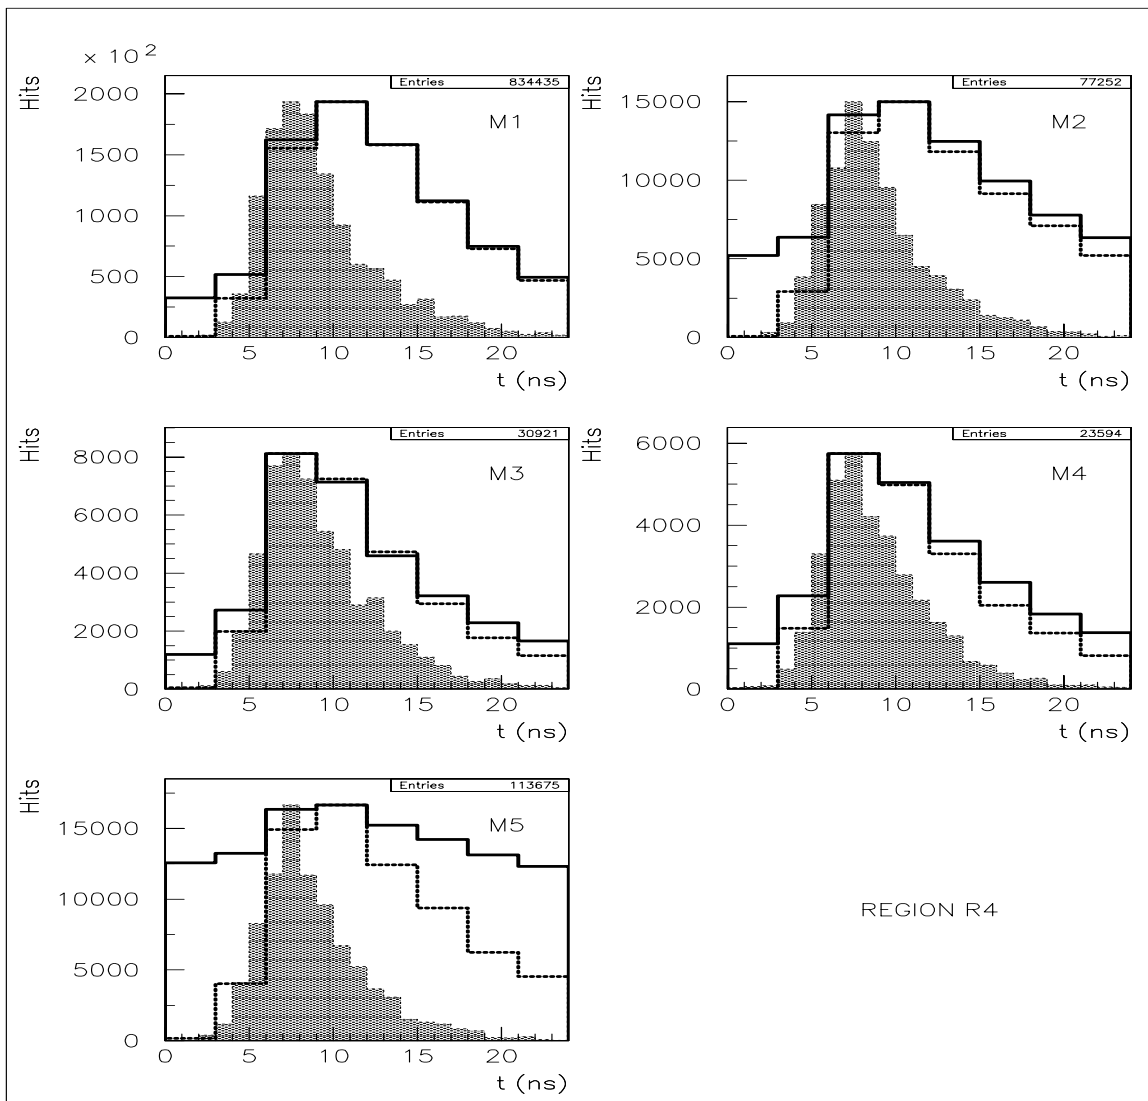


Fig. 3.3 Time distribution of hits in region R4 for the five stations M1--M5 for the default background level. Shaded histogram: hits produced by muons from B-decay. Solid line: time measured with a 3-bit TDC for all hits in MB events under standard running conditions. Dashed line: as solid line, but assuming bunch crossings well separated in time.

In Fig. 3.4 the solid line shows the time distribution (region 4 of station M5) of hits from high- p_T reconstructed tracks (standard field of interest, $p_T > 1\text{GeV}/c$) in MB events under standard running conditions for (a) default background level and (b) high background level (maximum safety factor). It can be seen that the solid line distributions reproduce quite well those corresponding to muons from B decays in both cases.

Since the L0 muon trigger requires 5 stations fired out of 5, small time misalignments could result in non-negligible inefficiencies. Therefore, it is advisable to acquire the time information with the use of a TDC. This would render the system much safer and would strongly speed up any procedure for time calibration and time monitoring of the muon detector. The results presented above indicate that a time resolution of 3 ns (3-bit TDC) is adequate for timing.




Fig. 3.4 Time distribution of hits in region R4 of station M5. a) default background level. b) The maximum safety factor to background is applied. Shaded histogram: hits produced by muons from B-decay. Solid line: time measured with a 3-bit TDC for hits from high- p_T reconstructed tracks in MB events.

A fine time tuning of the Physical/Logical Channels is possible taking data under standard running conditions and then reconstructing tracks from aligned hits, or alternatively selecting “isolated” events (non preceded by other interactions for a sufficiently long period of time). In this case, the sharp rise of the resulting distribution for all hits could be used for synchronization purposes.

3.3. Trigger and DAQ Interfaces.

The structure of the Trigger Interface is basically driven by the already established L0(μ) trigger scheme [1]. The reader is invited to read the cited reference [1] for a detailed description. Here it is just necessary to remind what follows. The L0 trigger algorithm is presently organized in a two-step procedure.

- i) The first step consists in fast (and coarse) track identification (FIP – Fast Identification Processor – stage). This is based on the Sectors information. A Sector is an appropriate logical combination of Logical Channels. The number of Logical Channels per Sector depends on the Detector Station/Region considered. 5 to 12 Logical Channels form a Sector. The front-end

electronics sends to the FIP the sector information as soon as available, with the corresponding BXId.

- ii) The second step consists in detailed track finding and ends up with the transverse momentum calculation. This step is named DMP (Detailed Muon Processor). It implies an interrogation of a DPRAM by the DMP itself. The detailed logical channel information is extracted for the Sector selected by the FIP. The operation must be performed within the L0 latency of 4 μ s.

This scheme implies a strict dialog between the trigger and the front-end electronics. In particular the trigger processor should be able to address and access the L0 DPRAM during execution of its p_T -calculating algorithm.

Recently, a different trigger organization option has started to circulate [7]. In this scheme the trigger architecture is completely synchronous and data-driven. Here the front-end interrogation is eliminated and the trigger is supposed to receive the complete BX information (Logical Channels) of the whole detector every 25 ns. This scheme has attractive features and will be considered later. However, unless explicitly stated, in the following we will always refer to the baseline scheme. When needed, the “new” version will be referred to as the “data-driven” version.

Here we list the operations performed by the Trigger Interface more schematically. Every 25 ns, the trigger interface has to:

1. dispatch the Sector hit map and the corresponding BXId to the FIP, for stations M2-M5 and for each BX.
2. store on a DPRAM the logical information for each sector, the storing address being given by the BXId.
3. receive sector addresses and the corresponding BXId for each FIP candidate at a maximum rate of 40 MHz. The FE electronics will interrogate the DPRAM for the selected sector and the given BXId. It will return the sector address, the corresponding logical information and the BXId to the DMP.

This dialogue between the FE electronics and the trigger require a direct communication and data exchange between FE boards.

At a lower frequency, but concurrently with the above operations, the DAQ Interface can perform the following operations:

- A. It can receive the L0-yes from the L0 Decision Box (1 MHz average frequency).
- B. It can be asked to transfer the L0-triggered information from the L0 pipelines onto the L0 derandomizers.
- C. It can multiplex and write the L0 derandomizers outputs onto the L1 buffers.
- D. It can receive the L1 decision and write the corresponding data onto the L1 derandomizers (40 kHz average frequency).
- E. It can multiplex and zero-suppress the L1 derandomizers’ output data towards the DAQ links.

While the operations A-E imply basically a “local” (on board) data transfer, up to the DAQ link, the operations 1-3 imply a direct communication with the L0(μ) trigger.

4. Choice of a baseline architectural scheme.

In the present chapter, we:

- i. Identify the system components performing the different tasks.
- ii. Find a physical place for the components with respect to the detector. This defines the amount of communication links needed and puts important constraints on the possible communication protocols.
- iii. Segment the whole system both “horizontally” (functional stages per kind of boards), and “vertically” (number of channels per board).

This will give us a number of possible solutions, among which we will choose a baseline one.

We identify and highlight the following functional stages:

1. Amplifiers-Shapers-Discriminators (ASD) stage, outputting the Physical Channels.
2. Front-end controls (discriminator thresholds, temperature, synchronization delays).
3. Field Bus node, associated to writing and reading back control/configuration parameters. (This should be present on every system unit containing writable/readable data and/or parameters).
4. Logical channel generation, containing masking and monitoring facilities.
5. Synchronization of Logical Channels.
6. L0 and L1 pipeline memories.
7. Trigger/DAQ interface.

Considering now the different possible location for electronics with respect to the detector, we can identify basically 3 zones:

- A. On detector. 10 years dose ranges from 1 krad (periphery) up to about 1 Mrad in the inner part of station M1. So, on-detector electronics need to be radiation hard and/or at least radiation tolerant.
- B. Outside detector, but attached to it (right and left sides). We call this location “Intermediate”. The accumulated dose here should be safely below 1 krad. Consequently, commercial components could be used. The maximum distance for signal links between zone A and zone B is expected to be below 6-7 m.
- C. Off detector. Always inside the cavern, but some meters away from detector. This is obviously considered a radiation-safe zone. The maximum distance for signal links between zone B and zone C is expected to be 7-8 m.

Assuming that the stages 1 to 7 have to be executed in sequence, we can map them in locations A to C. We will consider 4 different mappings. They constitute our 4 basic options for the muon electronics architecture. They are illustrated in fig. 4.1.

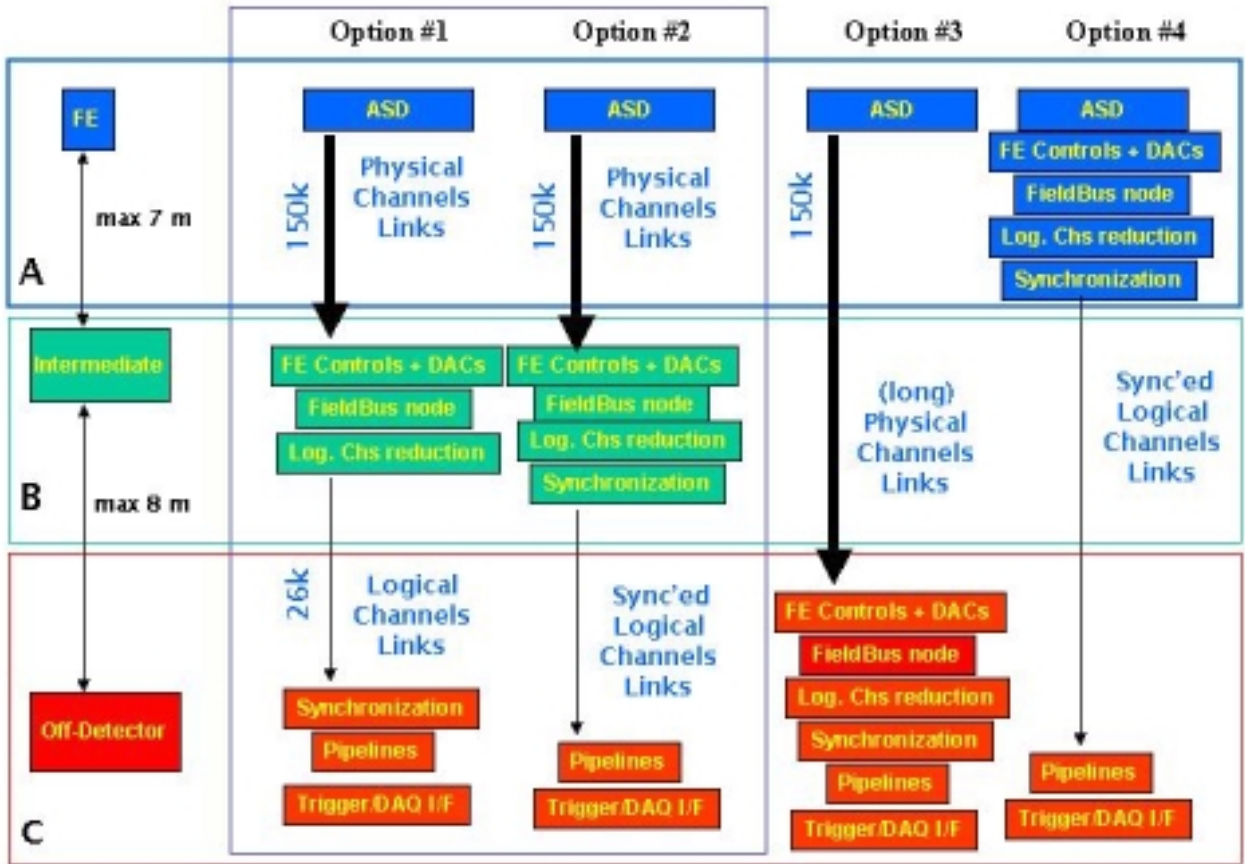


Fig. 4.1. Four architectural options.

In the baseline trigger scheme, we foresee a DPRAM access with the use of a custom bus. This suggests us to put stages 6 and 7 either in zone B or C and not A. On the other hand, it seems difficult to crowd everything on one single kind of boards. Moreover, owing to the high number of Physical Channels (150k), it is not practical to address memories scattered over a high number of boards in zone B. So the most natural choice is to place both stages 6 and 7 in zone C. At this point, we are left with stages 2 to 5, as stage 1 is naturally attached to the detector. In order to minimize the number of different kinds of boards, the obvious solution would be to push stages 2-5 either on zone B (option #4), or on zone C (option #3).

Option #4 strongly reduces the number of LVDS links needed to transport the Physical Channels out from the ASD boards, as the information can be serialized after synchronization. However, i) it implies the need of complex radiation-hard electronics, in order to place electronics in the hot inner regions; ii) it introduces complex digital electronics for synchronization and Logical Channels generation tasks on the same boards as the ASD chips. This would result in complex boards whose layout is difficult and that could result to be critical.

At the opposite extreme, Option #3 implies the use of a large amount of cabling travelling from the ASD boards up to the C zone. In order not to make the number of the resulting complex boards explode, a very large number of connections should be put on each single board. The analogue connections needed to control the front-end analogue boards would be found to be longer than strictly necessary. Finally, it is difficult to match a suitable segmentation of Logical Channels and sectors per board as requested by the trigger scheme.

Option #2 and #1 are very similar, with the only exception that the synchronization task is performed (#2) or not (#1) in the intermediate zone. Option #2 could be convenient because data can be serialized after synchronization and LVDS links for Logical channel transportation drastically reduced. On the other hand, considering boards' occupation, the electronics located in zone C would be badly exploited and the overall complexity would be concentrated in zone B. In comparison, option #1 can be seen as a more balanced structure. Moreover, option #2 will have a considerable amount of boards in zone B, due to the fact that these boards must house about 150k Physical Channels. Integrating also the synchronization stage on these boards will have a strong impact in increasing system complexity and costs. So, option #1 is expected to be less expensive and more maintainable than option #2.

As a result of the above description, options #3 and #4 would be ruled out owing to relevant drawbacks, while option #2 results to be more critical, unbalanced and more expensive than option #1.

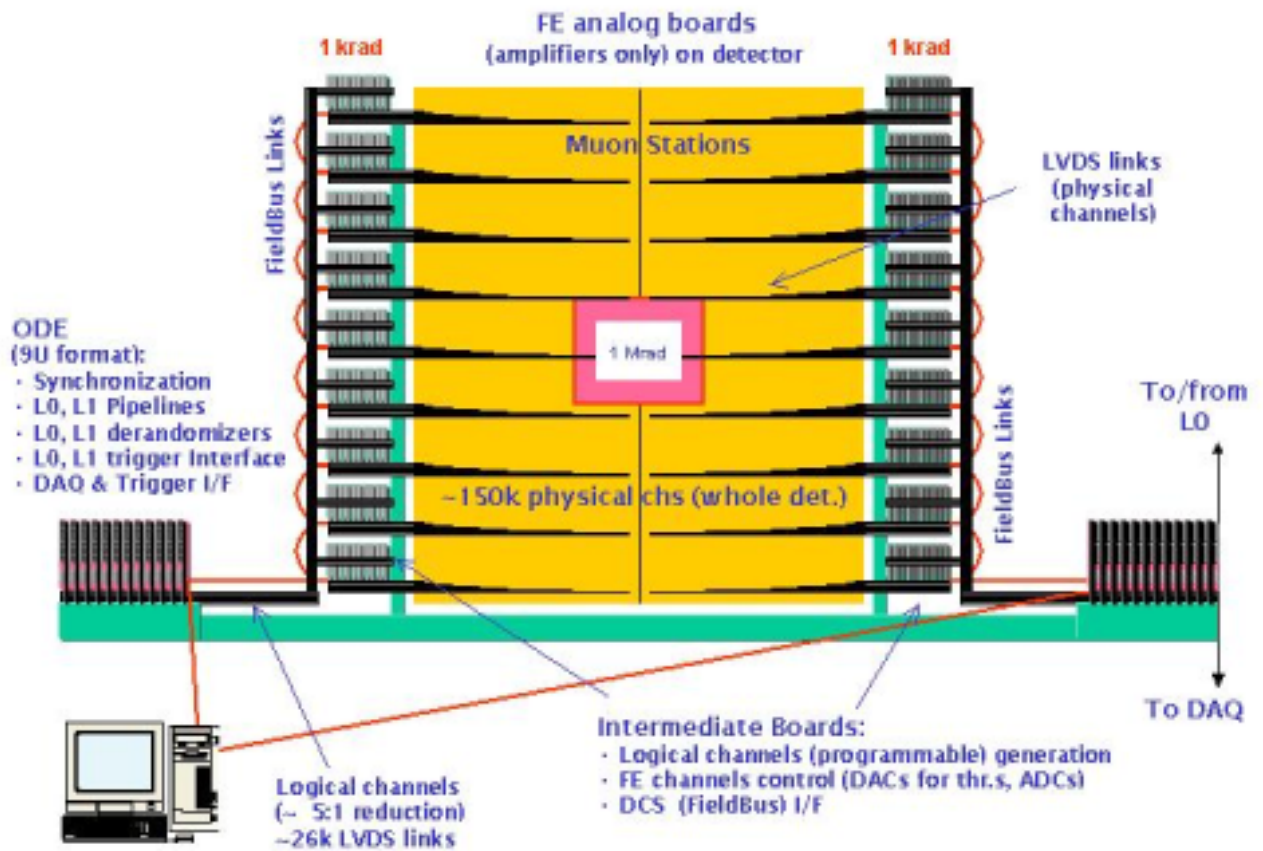


Fig. 4.2 Muon electronics architecture according to option#1. The ODE boards location in the figure does not refer strictly to a definite physical place.

Therefore, we will now use option #1 as our baseline solution for the Muon detector electronics. In the next chapter we will define in further details the structure of this architecture. Fig 4.2 gives a pictorial sketch of the location of electronics with respect to the detector, according to our baseline solution. In the following, electronics situated in the B zone is referred to as **Intermediate Boards (IBs)**, while electronics placed in the C zone is called **Off Detector Electronics (ODE)**.

It is important to say that the choice of option #1 is not yet a definitive one. Accurate feasibility studies are required to validate this choice. They are related to a detailed definition of the detector structure, component costs, high density connections and other basic implementation issues.

5. Front-end Architecture.

In the present chapter we describe in further detail our baseline Front-end architecture. The description is organized in three parts. We describe the very front-end first (ASD boards), containing only stage 1 (see chap. 4), the Intermediate Boards (second part), containing stages 2 to 4. Finally (third part) we describe the Off Detector Electronics, containing stages 5 to 7.

5.1. ASD Boards.

ASD Boards (ASDB) result to be very simple. They are completely analogue boards containing only the ASD chips and passive components (resistors, protection diodes etc.). This simplifies a lot the board layout. A temperature sensor could be worth adding.

The ASDB is a little boards containing 1 or 2 chips. Using small size boards allows a better match between the electronics boards and the detector geometry. The number of channels per board strictly depends on the components used. The component choice depends on the type of detector and, as of today, it is not completely defined. A number of candidates are presently under development and test.

We will have around 8-16 digital outputs per board (1-2 ASD chips). The output signals are expected to be LVDS. When this is not the case, a suitable translator stage has to be added. It is preferable to add the translator stage on the ASDB itself, otherwise different input stages have to be foreseen in the Intermediate Boards.




Fig 5.1 Sketch of the ASDB

In addition to LVDS outputs, a subset of the analogue signals (at least one per board) is useful to better monitor the behaviour of the detector and of the electronics itself.

The ASDB boards inputs the discriminator's thresholds from the IB. Fig. 5.1 gives a sketch of the ASDB and in particular shows the foreseen IO.

5.2. Intermediate Boards

The main problem when exiting the ASDB is the huge number of LVDS links to be handled (around 150k). Consequently, a considerable number of signals has to be input on a single board. Very high-density connectors are available on the market [9], allowing to connect about 350 LVDS pairs on one 9U (VME style) board.

On the other hand, it is important to respect the segmentation of the detector in Logical units and Logical Channels, in order to have a definite mapping of the detector on the electronics boards. Therefore, a maximum number of 192 pairs per board seems preferable. This allows us to obtain a better match between boards and Logical Units [10]. Table 5.1 demonstrates how the channels distribution would be per region (R1-R4) and per station (M1-M5).

	M1	M2	M3	M4	M5
R1	96/24	56/28	56/28	48/24	48/24
R2	192/24	112/32	80/32	96/28	96/28
R3	192/24	192/28	192/28	192/20	192/20
R4	192/24	192/28	192/28	192/10	192/10

Table 5.1 Inputs/outputs per IB [10].

The main task of the IB is to generate the Logical Channels. One logical channel can be the result of the combination of Physical Channels belonging to different chambers. As a consequence, the fine time adjustment must be performed on physical channel rather than on logical ones. One programmable delay should be foreseen for each of the 192 input channels of the IBs. The delay to be added to signals is written on dedicated registers through the Fieldbus node, after completion of the time calibration procedure (see par. 3.2 and 5.3). A single adjustable delay for each physical channel can be avoided by fine-time-adjusting at the level of groups of OR-ed Physical Channels whenever they come from the same chamber. These channels are expected to have an homogeneous time behaviour. This complicates the IB scheme, but allows saving a relevant amount of components and a relevant amount of data to be uploaded onto the boards at set-up time. Figs. 5.2a and 5.2b give an example of a logical channel generation scheme performing time adjustment on groups of Physical Channels.

In order to limit the effect of single channels dead time on the OR combination, it is fundamental to work on re-shaped signals: physical channels should be shaped to last 3-5 ns from their rising edge.

In this time-adjustment scheme, the use of the TTC system [11] is not strictly necessary at the level of the IB. The time adjustment is done by adding a suitable number of time steps to the input signals, without the use of any time reference.

In order to make it possible to access the single physical channel also after logical channel generation, the possibility to mask physical channels is foreseen, as well as the facility of cycling on the inputs themselves to monitor the single channels' rates (see par. 3.1). The programmable mask registers and the channel scalers must be accessible for read/write operations through the bus node.




Fig. 5.2a. Example of physical and logical entities. Logical strips are formed by logical pads, which are formed by the OR of 2 layers (L1, L2). Each layer consists of 2 gaps (G1, G2). 4 physical channels per layer are shown. Horizontal logical strips are formed by pads of the same chamber, vertical ones are formed by pads of different chambers.

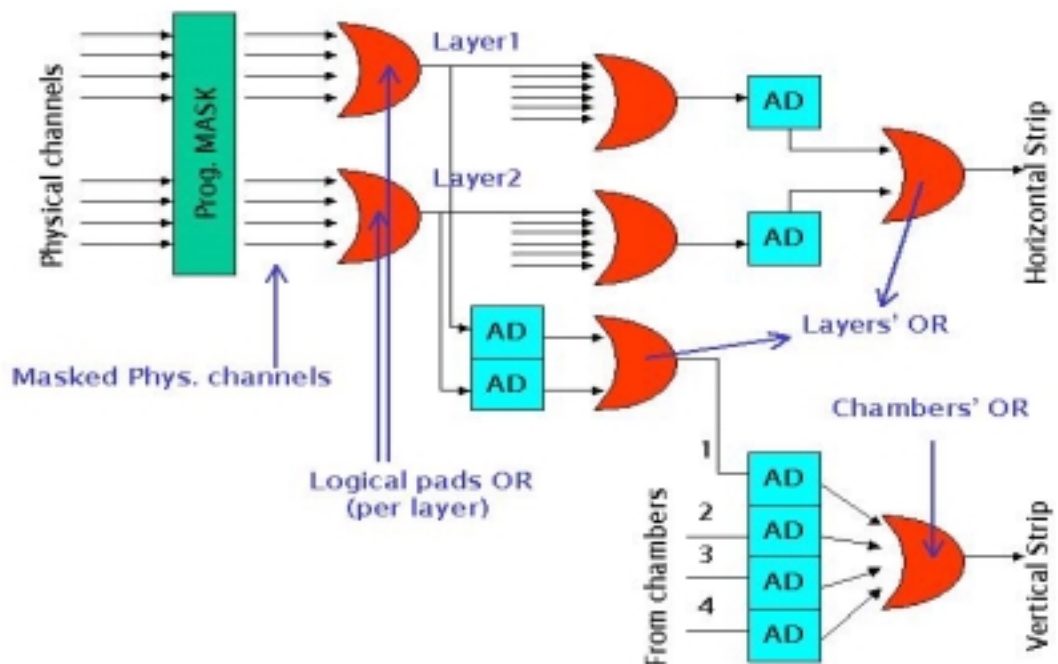


Fig. 5.2b. Example of logical channels generation. Adjustable Delays (AD) are placed only to combine signals from different layers or different chambers.

The logics necessary for Logical Channels generation and physical channel masking are fitted inside FPGAs. A range of 24-48 input Physical Channels per component seems at the moment a reasonable quantity. This would imply 8 to 4 FPGAs inside each IB. By exploiting the masking feature, scalars can be

placed to spy the logical channel rates, either on the ODE or on the IB itself. Also this function can be integrated inside FPGA.

The IB would be also used to configure and monitor the ASDBs. The configuration and monitoring information are exchanged through the Bus node.

Discriminator thresholds are passed to the ASDB through dc signals, provided by DACs placed on the IB. Considering 8 channels per ASD chip, putting one threshold level per ASD chip results in 24 DACs inside one IB, in order to configure 192 Physical Channels. This number can be lowered to 12 (6) if the same threshold is used for 2 (4) ASD chips. A reasonable compromise must be reached here, according to the chamber characteristics.

The monitoring information from the ASDB consists in temperature control signals and analogue outputs from the ASD chips. In order to minimize the number of ADCs, analogue multiplexing can be used. Considering one analogue output per chip, we have 24 signals to A/D convert per IB. This number could be a bit too high, so most probably a lower amount of analogue outputs can be monitored. 12 analogue signals within 192 (1 analogue output every two 8-inputs-chips) is a more affordable number. They could be multiplexed (2:1) onto, for example 6 ADC placed inside the IB. This would be useful to perform significant checks on the ASDs' and chambers' operation remotely. An alternative approach is to keep the analogue outputs on the ASDB but eliminate ADCs from the IB. Dedicated ADC boards can be used separately during system tests and setup. The same boards can be used at run time to monitor a selected set of analogue outputs.

Fig. 5.3 summarizes the resulting structure of the IB. The main components and IO counts are shown.

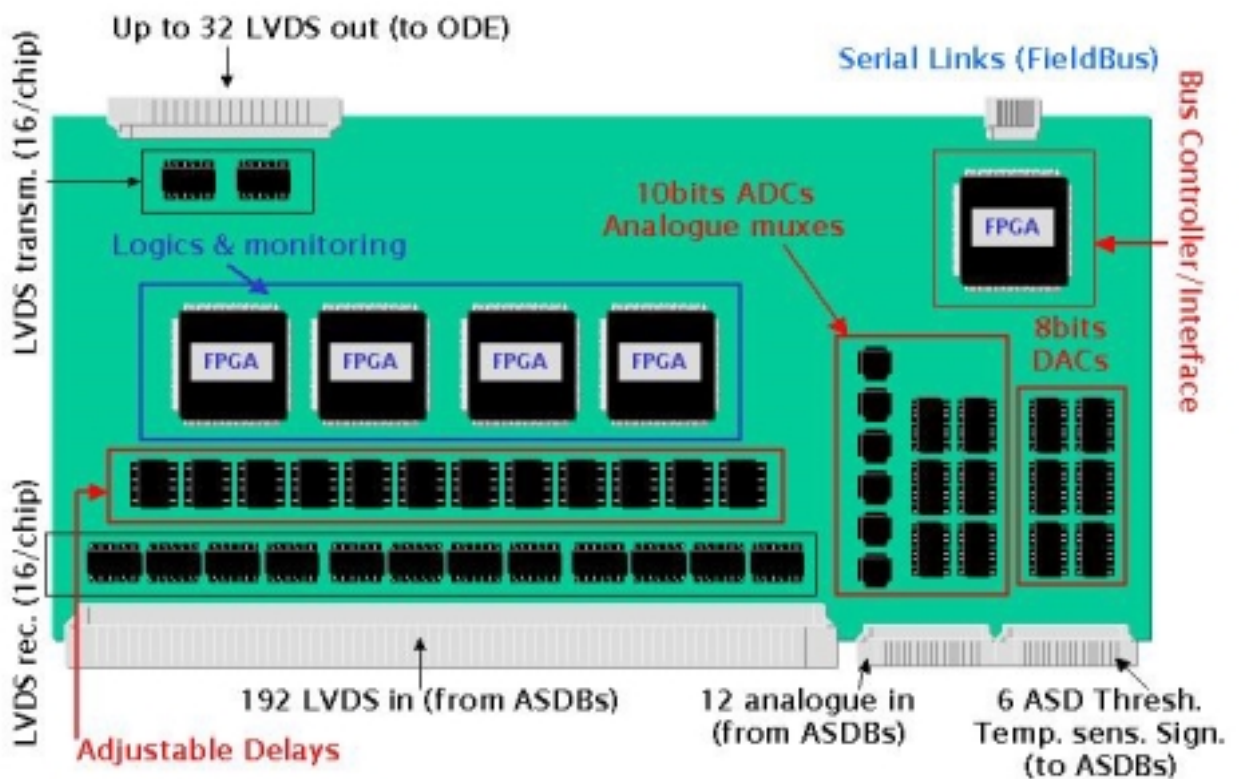


Fig. 5.3. Sketch of the IB.

The logics of the IB would be integrated inside FPGAs. Consequently, a special care should be spent about the problem of Single Event Upsets (SEUs), which is expected to be relevant at the radiation level present in the intermediate zone.

SEUs depend on particle rates rather than accumulated dose. They normally appear as transient pulses in logic or support circuitry, or as bit-flips in memory cells or registers. They corrupt the configuration data, the data itself and the normal function of states machine. These are "soft" bit errors in the sense that a reset or reprogramming of the device restores normal behavior.

The FPGAs allow the use of flexible high integration electronics, avoiding the development of expensive and not upgradable ASIC chips. As a drawback, the FPGAs architecture uses 50 times more flip-flops for device programming. Thus it shows a dependence 50 times worst on SEUs with respect to an ASIC.

The program substrate of FPGAs is in SRAM technology, with a typical cross-section/bit of 10^{15}cm^{-2} for SEUs events [12] in a typical case. Considering a configuration of 1 Mbit, we have a cross-section per device of 10^{-9}cm^{-2} .

In case of Single Event Upset, the device must be reprogrammed *in loco*, to avoid bandwidth consuming at level of DCS. A local copy of FPGAs configuration can be stored in Flash Memory. This kind of memory shows a cross-section/bit for SEUs events of $10^{-18(19)}\text{cm}^{-2}$ [12]. If we consider, as an example, a flux of 10^{11} proton $\text{yr}^{-1}\text{cm}^{-2}$ and a typical LHCB year of 1000 hours, we have a SEU inside one FPGA per hour and a SEU inside the Flashrom per year of running. This numbers must be better verified by tests. However, they already show how that SEU effects have to be carefully evaluated and kept under serious control in system design. A special circuit for recovering these errors must be included in the logic. The circuit reads-back continuously the FPGA configuration and, when a wrong configuration is detected, it resets the device and reloads the FPGA. When also the Flash Memory has a wrong bit, the circuit asks the DCS to download a fresh bitstream.

5.3. Off-Detector Electronics.

ODE is dedicated to synchronization of signals and trigger and DAQ interfacing. The task of synchronization (see par. 3.2) can be seen as subdivided in two parts:

- Time measurement (fine and coarse);
- Time adjustment (fine and coarse).

Fine-time adjustment is performed inside the IB. It is left to the ODE to measure time during dedicated time-calibration runs and to assign the BX Id to the incoming signals ("coarse" time adjustment). The BX Id corresponds to the 8 bits address of the pipeline memory.

From a logical point of view the time-calibration procedure would imply the following sequence:

1. Measure the arrival time of signals (Logical Channels). A time resolution of about 3 ns should be used in time measurements. This number matches with the degree of signals time jitter and also corresponds to subdivide the 25 ns clock period into 8 parts (3 bits TDC).
2. Accumulate enough statistics to build time distributions (one per channel) and calculate the number of 3-ns-steps to add/subtract to the arrival time in order to include the time distribution inside one 25 ns clock period (BX). We call this number as Fine Time Offset (FTO). The FTOs must be communicated

to the IBs, where they are used for correction by means of programmable delays. This phase of the procedure is illustrated in fig. 5.4.

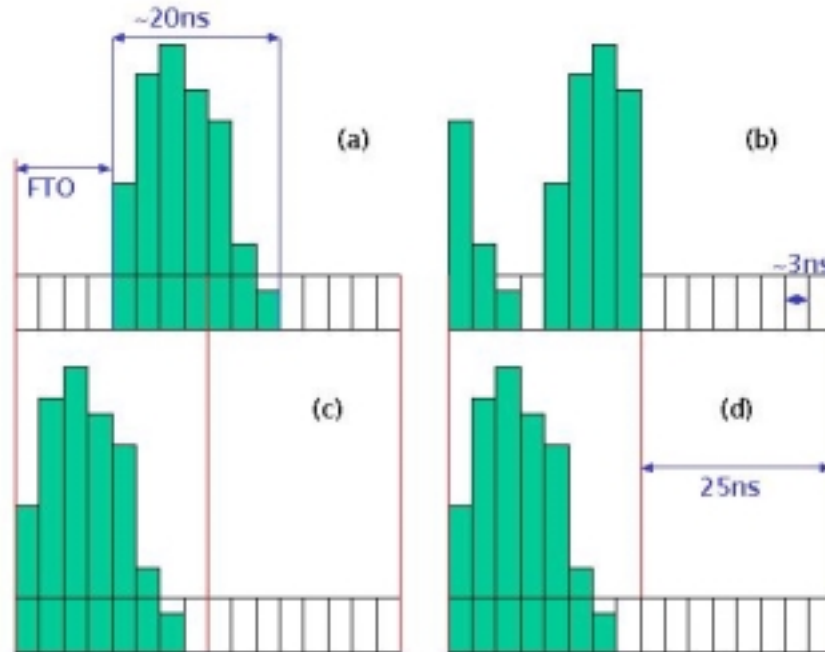


Fig. 5.4 FTO histograms and correction. (a) Displaced time distribution with respect to the 40 MHz clock. (b) Accumulated histogram on a 25 ns gate. (c) Time distribution after FTO correction. (d) Accumulated histogram on a 25 ns gate after FTO correction.

- After fine time correction, keep acquiring a suitable number of subsequent time slots, fill a histogram (number of events per BX Id) and compare the difference in number of time slots between the acquired histogram and the expected structure of BX, due to the LHC machine spill sequence. This difference is called BX Offset (BXO) and must be used at the level of the ODE to suitably correct the BX Id of the channels. The LHC bunch crossing structure is given in Fig. 5.5. The regular sequence of full and empty machine cycles provides an absolute time reference. Fig 5.6 illustrates of comparison between the acquired BX histograms and the absolute reference (for example, the BX reset at the beginning of the BX sequence itself).

According to the discussion given in par. 3.2, the FTO could be measured from fine-time histograms either taking into account reconstructed tracks (in order to separate background), or acquiring time hits at low luminosity. In both cases the rise of the distribution can be identified and the FTO determined. Reconstructed tracks are possible only using a track finding algorithm, which could be applied starting from a coarse BX alignment and looking at a number of subsequent time slots, in order to take into account possible BX misalignments. On the other hand, the use of low luminosity hits would allow to better automatize the procedure, because histograms could be realized directly on dedicated memories inside the ODE.

The main functional units inside the ODE boards are shown in fig. 5.7. Incoming LVDS signals are received and converted into TTL levels (Input unit). Their phase with respect to the system clock is measured (TDC unit). They are synchronized at 40 MHz and time-stamped (Sync unit). The time-stamp is the BX Id itself. Data are written onto pipeline memories (L0 Buffer Unit) at the address corresponding to the assigned BX Id. Data will consist in the hit value (1 or 0) and the fine time information as an output of the TDC unit. The L0 Buffer unit will contain also the L0 derandomizers. In order to reduce the data

bandwidth for the DCS, it could be useful to add a specific memory and a specific unit, for accumulating both the Fine Time and the BX structure histograms (Histogram unit). This would allow uploading the complete histograms directly, rather than each event one at a time. A built-in automatic algorithm for FTO and BXO calculation could also be easily implemented. However, the same algorithm can be executed by software programs on acquired data during dedicated runs. Thus we consider this unit as an option. The L1 pipelines are also part of the ODE (L1 Buffer unit), as well as the Trigger and DAQ interface units, which end up with a proper number of optical links for communication to L0, L1 trigger and DAQ systems.

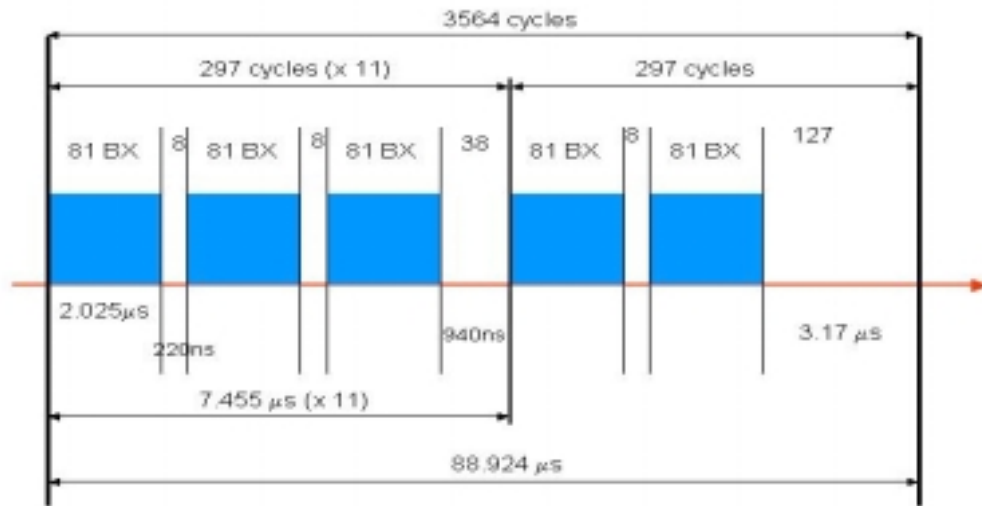


Fig. 5.5 LHC BX structure: regular sequence of full (coloured zones) and empty cycles.

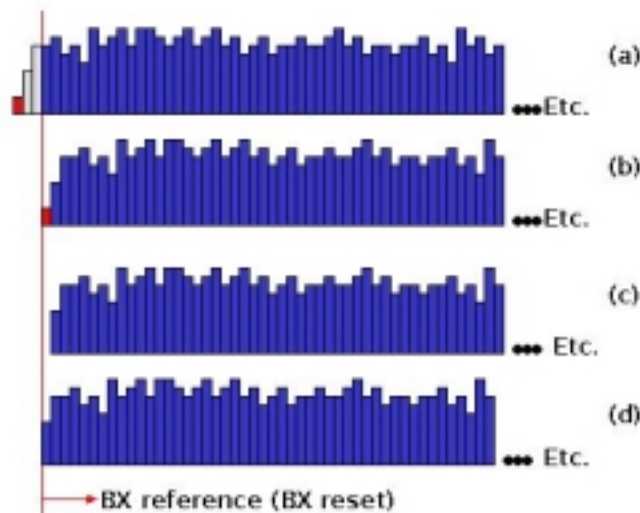


Fig. 5.6 Example of BX alignment. (a) BX histogram where the considered channel is found to be misaligned by 3 cycles: the leftmost is due to fine time misalignment, the other 2 to clock misalignment. (b) Calculated BXO is wrong by 1 clock cycle (FT misalignment). (c) Fine time correction cleans the leftmost slot and (d) the BXO is re-corrected properly.

The last studies about detector layout mappings on the trigger scheme, along with the opportunity to fit as many Logical Channels as possible on one single board, suggest a number of Logical Channels in the range of 200 per board, which results in about 150 ODE boards. In particular, a good segmentation (and mapping) of the detector can be achieved by placing 192 Logical Channels per board [10]. The 192 channels are then subdivided into 8 (6) integrated components, called Sync-IC. They can integrate the functionalities required for 24 (32) channels. This corresponds to the Sync unit, the L0 buffer unit and the Histogram unit. These units could be implemented also with a high performance FPGA. On the other hand, it is worth considering also the integration of the TDC unit inside the same components. This corresponds to the design of 24 (32) 3-bits integrated TDC, with 3 ns time resolution. In such a case the above units, along with the TDC one, would be integrated on a custom IC. The choice between the custom and the FPGA solution must be made evaluating carefully the convenience in terms of costs and system reliability and compactness. Feasibility studies and simulations on the custom solution are presently on progress.

The L1 buffer unit and the IO interfaces (trigger and DAQ) are realizable on FPGAs. Fig. 5.8 gives a pictorial sketch of the ODE board, highlighting the main hardware components and the number of IO connections.

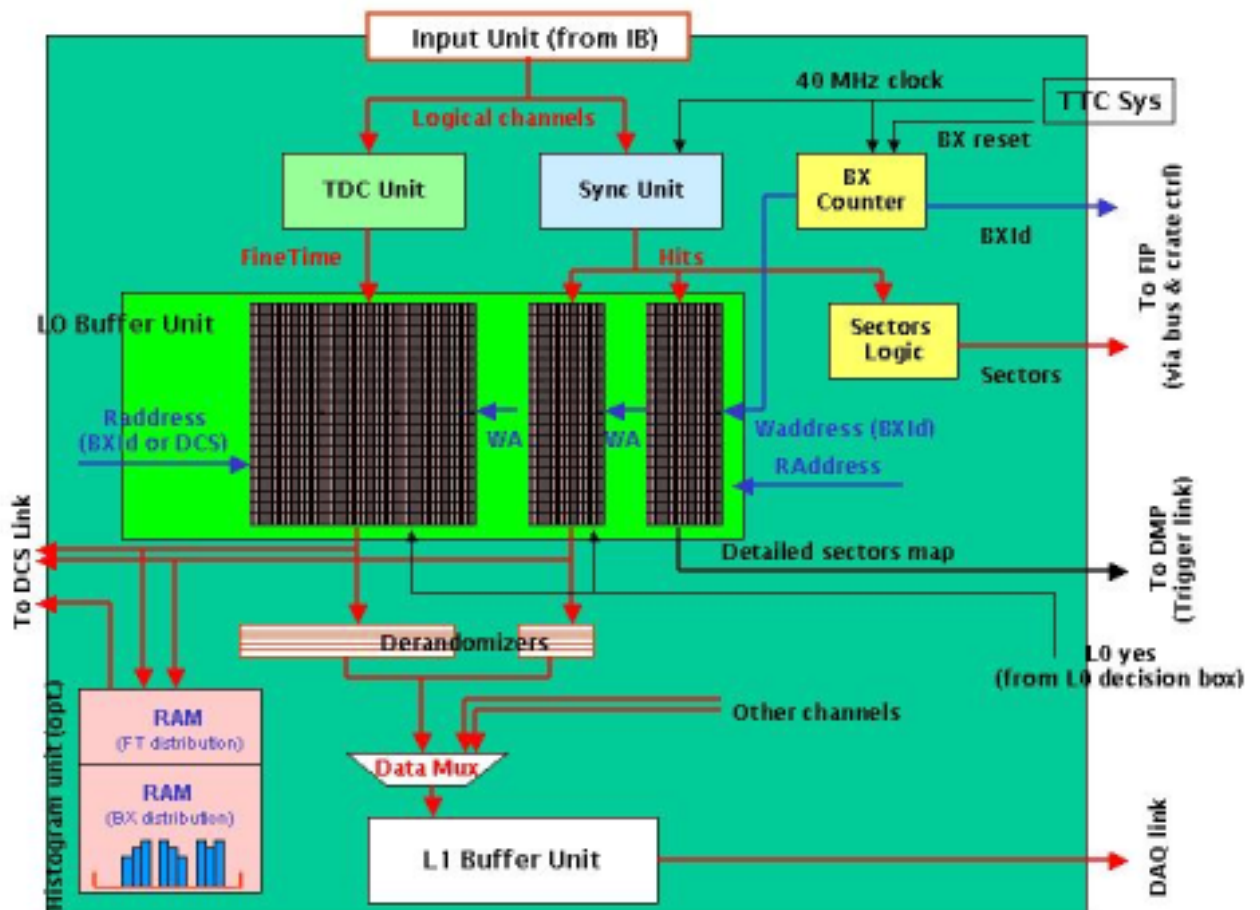


Fig. 5.7 Functional units of ODE

In the baseline trigger scheme, an additional module is necessary in the ODE system. This is the Crate Controller, communicating to the ODE boards on one side and on the FIP and DMP on the other. It receives the sector addresses from the FIP for each FIP candidate. They are broadcast to ODE. ODE will return the sector address and content to the Controller. They are then transported to the DMP. One ODE controller is needed per crate, so around ten of them are needed.

On the other hand, in the data-driven scheme, no crate controllers would be needed (and no sector-OR either). In such a case each ODE boards communicate directly to the trigger system, without any dialogue among the other ODE boards. Data are transmitted using optical links, along with their channel addresses and time specifications. This simplifies the front-end architecture and communication protocol, but increases the number of links between the trigger and the front-end electronics. The choice is still opened and under discussion. A carefully evaluation from the trigger implementation point of view is presently in progress.

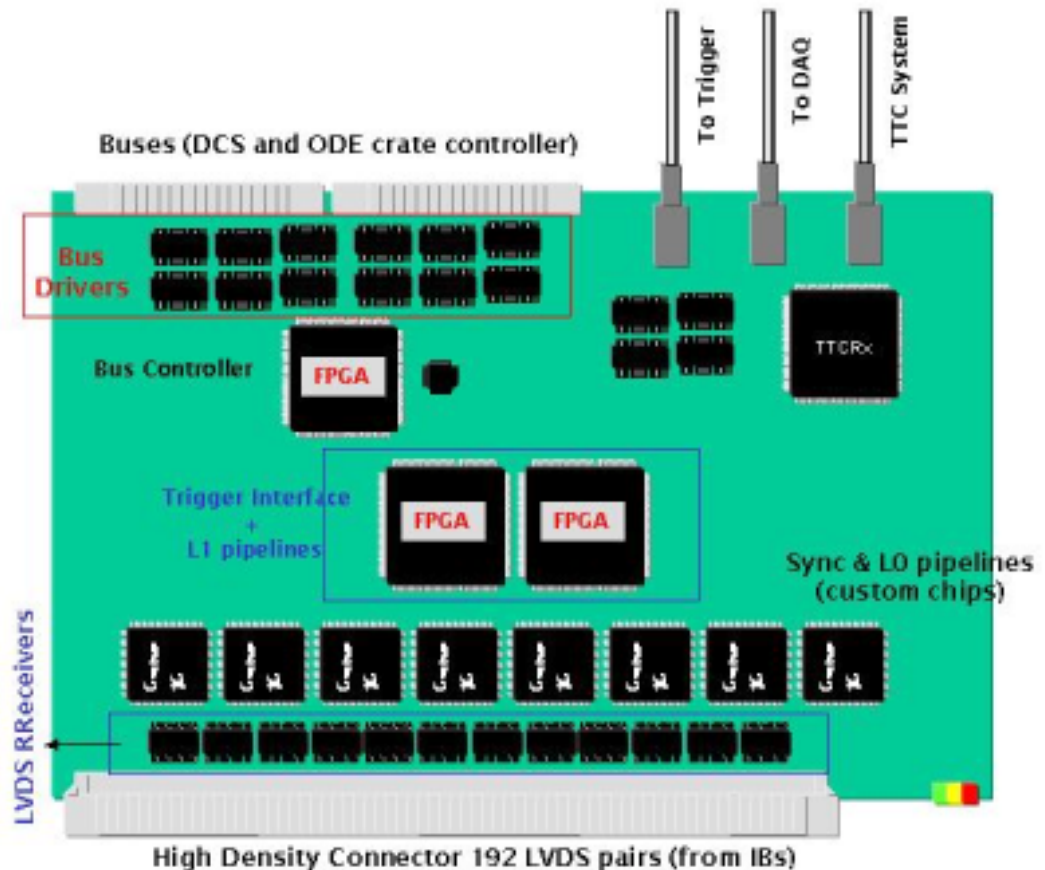


Fig. 5.8 Pictorial sketch of the ODE board.

5.4. Detector Control System.

A distributed control system based on a specific fieldbus will provide the basic control and monitoring functions of the LHCb muon detector. A fieldbus is a simple cable bus, connecting nodes using a specific protocol. The nodes usually contain a microprocessor. The “intelligence” of the node can be used to handle the fieldbus protocol and also to execute simple local tasks.

In the LHCb muon detector we need a system with one node per IB and per ODE board. It performs the following operations:

- 1) Control of ASD thresholds.
- 2) Configuration of physical channels masks during calibration runs or when a noisy channel is detected.
- 3) Upload and download delay settings for physical channels.
- 4) Acquire the histograms (ODE) and rates (IB) for each physical channel.
- 5) Reprogram or update the FPGAs configuration on Flash ROMs (see also the SEU effects – par. 5.2).
- 6) Run the Scan test of the board using JTAG.

To minimize connections inside the single board a local serial bus can be used. The I²C local bus can be the proper choice. The I²C bus was developed in the early 1980's by Philips semiconductors. Its purpose was to provide an easy way to connect a CPU to peripheral chips in a TV-set with enough bandwidth (typically 100 Kbit/s). The DAC for the threshold the mask for each single channel and the delay settings can be connected to one or more I²C branches inside each single board.

The SEU effect causes the corruption of FPGA configuration during normal run. If we consider a system with the new configuration downloaded centrally from the DCS, every time a SEU corrupts the FPGA configuration there is need of a considerable bandwidth dedicated to this purpose. We want to use a feature of Flash ROM devices, consisting in a tolerance to single event bit corruption three order of magnitude better than the FPGA with SRAM architecture. In this way the downloading of bit configuration for the FPGA is done using the local Flash ROM copy inside the board.

If we think to split the FE-boards (IB and ODE) in 45 DCS chains each one with maximum of 32 boards we can consider a maximum of 160 FPGA for each DCS branch. With 160 FPGA we have a mean number of wrong configuration for chain of about 3 FPGA for minute. The problem can be solved locally but every 6 hour we need to reprogram the correspondent Flashrom. Using as Fieldbus like CAN with 1Mbit/s the time needed to reprogram the Flashrom is 2-3 seconds. Controller Area Network (CAN) is designed to provide an efficient, reliable and very economical link between sensors and actuators. CAN communicates at speeds up to 1Mbit/s with up to 40 devices. Originally developed to simplify wiring in automobiles, its use spreads in machine and factory automation products. The CANbus is used inside the LHC machine and, due to the requirements of the automotive market, it will be available for a long period of time.

We use the DCS in the apparatus to control the functions of each single board, for the ASD threshold we have one DAC for each eight physical channel than we need 19 Kbytes for the DAC 8 bits registers.

For masking each single physical channel in the muon apparatus we need the same amount of data of the DAC registers (19 Kbytes), and for the delays we need about 50 Kbytes. The parameters to upload from the FE boards are in the order of 300 kbytes and come from the 89 bits for each logical channel. With 40 Kbytes/s of available true bandwidth (only payload without header and handshaking) using 40 CANbus chains we need 22 ms to download all the threshold and the mask registers, 31 ms to download all the delays and 187 ms to acquire all the histograms. The time needed to upload all the FPGA Flash ROMs at the same time is less than 5 minutes.

The reduced time of single access in the cavern and the inaccessibility of each single board during running time impose an efficient method to diagnostic the behavior of each failing board to plan an apparatus intervention. Embedded test, emulation and maintenance circuitry are well defined and understood in the test community. The IEEE standard 1149.1, known as JTAG, gives the possibility to perform Boundary-scan test of a single PCB board. Boundary scan is a special type of scan path with a register added at every I/O pin on device. The most obvious benefit offered by the boundary-scan technique is allowing fault isolation at the component level. Every DCS node must have a JTAG interface for remote diagnostic purpose, the JTAG can be used also as backup solution for FPGA programming.

6. Costs

We make some considerations about system costs. We do not give a detailed estimate, because a detailed components choice has not been done yet. However, the following points can be stated:

1. Concerning the ASDB, we keep giving the cost of 6 CHF per channel. This cost comprises both the preamplifier and the PCB.
2. We have a remarkable number of IB (around 1060). As already stated, this is due to the high number of physical channels and to the need of:
 - a. Avoiding the use of rad-hard electronics;
 - b. Endowing the system of a sufficient set of monitoring facilities;
 - c. Having sufficient flexibility in configuring the system, though using a limited set of boards' types.

As a result, the IB are complex boards. They are based on the use of FPGAs. A reasonable cost estimate for the IB cannot go much below 3 kCHF. This represents a heavy amount of money on the muon detector budget. It can be necessary to drop one of the 3 conditions above, in favour of lowering the system cost. In particular, one can consider to partially reduce the number of physical channels already at the level of the ASDB, OR-ing physical channels of the same chamber. Channel masking should be kept at the same level.

3. The number of ODE boards is minimized (152). The ODE boards are complex and large boards. They contain a relevant number (around 10) of high density FPGAs and/or custom ICs. They also contain 3 optical links transmitters. They are estimated in the range of 5-6 kCHF. The number of ODE Boards could be increased if it helps to drastically reduce the number of IBs. This can be another remedy to the high number of IB. In other words, this means to re-consider the present detector mapping in logical and physical channels.

7. Conclusions.

This document has given a description of our present envisaging of Muon front-end electronics. The document should be considered as a status report of the recent studies about architecture and its related issues. The work is currently in progress. Major changes in system design might be expected within the completion of the TDR. Nevertheless, the basic characteristics and crucial points of the system have been highlighted. In particular, it has been seen that the task of system synchronization turns out to be crucial, both for its importance within trigger operations and for its difficulties in implementation. The opportunity to measure the time arrival of signals to be fed to the L0 trigger has been demonstrated. A resolution around 3 ns fills our needs. The need to time-adjust signals at the level of discriminator's outputs (before logical combination) has also been demonstrated.

A baseline architecture has been identified. In this baseline scheme, three different kinds of boards are necessary from the architectural point of view. They are:

- The ASD Boards, containing only Amplifiers, Shapers and Discriminators, placed directly on detector.
- The Intermediate Boards (IB), placed very near to detector, right outside from its sensible area, in a radiation-safe place. These are used for front-end configuration and monitoring, fine time adjustment and generation of the Logical Channels from the physical ones. About 1060 IBs are foreseen, housing a maximum number of 192 inputs each.
- The Off-Detector Electronics (ODE) Boards, also placed near to detector. They contain time-measurement functions, L0 and L1 buffers and associated circuitry, trigger and DAQ interfaces. About 150 ODE boards are foreseen.

The critical points of the system, which require further thinking and work to be better finalized, are the following:

- Implementation of the synchronization algorithm (time measurement and adjustment). A large number of TDCs and programmable delay is required. Also a custom solution could be worth considering.
- Hardware tests are also important to detail our system. In particular, high speed signals' transmission and the use of not-too-expensive high density connectors must be proved.
- The problem of SEU effects has to be more precisely evaluated and understood in its implications on system design, especially concerning electronics location and the massive use of FPGAs.
- A special effort has to be done to minimize costs, without spoiling system performance and flexibility.

We can consider the above points as the main gaps to be filled in view of the TDR completion date.

8. Acknowledgements.

We wish to thank J. Christiansen and R. Le Gac for their carefully reading of the document and their useful suggestions.

9. References.

- [1] E. Aslanides et al., *The L0(μ) processor*, LHCb 99-008
- [2] B. Schmidt (editor), *LHCb Muon System by Numbers*, LHCb 2000-60.
- [3] V. Talanov, *Estimation of absorbed dose levels at possible locations for LHCb detector electronics*, LHCb 2000-15.
- [4] P. Colrain and B. Schmidt, *Muon System Optimization*, LHCb 2000-16 (in preparation).
- [5] A. Tsaregorodtsev, *SICB User Guide: a GEANT3-based simulation package for the LHCb experiment*. June 1999.
- [6] A. Tsaregorodtsev. *Muon System parameterised background (algorithm and implementation)*. LHCb 00-011 (internal note).
- [7] R. Le Gac. Presentation given at the Muon meeting. March 2000.
- [8] Texas Instrument Application Note: <http://www.ti.com/sc/psheets/slla053/slla053.pdf>
- [9] AMP Catalog. AMP Z-PACK™ 2mm Hard Metric Connector System.
- [10] R. Le Gac. Presentation given at the Muon meeting. May 2000.
- [11] A large set of TTC-related documentation can be found at: <http://ttc.web.cern.ch/TTC/intro.html>.
- [12] 1998 J.R.Coss, R.F.Miyahira, L.E.Selva, G.M.Swift., *Device SEU susceptibility update: 1996-1999*, IEEE Radiation Effects Data Workshop Record, held in conjunction with the IEEE NSREC conference in Norfolk, Virginia, 12-16 July 1999.