CPPM, IN2P3−CNRS et Université d'Aix−Marseille II

A synchronous architecture for the L0 muon trigger

LHCb Technical Note

Issue: 1
Revision: 1

Reference: LHCb 2001−010
Created: January 23, 2001
Last modified: February 8, 2001

Prepared by: E Aslanides, J.P. Cachemiche, P.Y. Duval, R. Le Gac, O. Leroy, P.L. Liotard, M. Menouni, R. Potheau and A. Tsaregorodtsev

Abstract

In this note we describe a new implementation for the L0 muon trigger. It is based on a fully synchronous and pipeline architecture. The 25920 logical pads and strips data, produced by the muon detector, are sent at once to the muon trigger. The data transfer is performed at a frequency of 40 MHz and relies on ~1250 high speed optical links running at 1.6 Gb/s. The muon trigger processor is distributed over 4 crates. Each of them is connected to a quarter of the muon system. The implementation is based on a unique processing board. Data transfers between boards are performed at 80 MHz via a custom backplane running only point to point connections. Synchronization between data and the bunch crossing identifier is preserved during the whole processing cycle. In this note, we present the data flow, hardware implementation, latency and cost estimate.

Document Status Sheet

1. Document title: A synchronous architecture for the muon trigger			
2. Document Reference Number: LHCb 2001−010			
3. Issue	**4. Revision**	**5. Date**	**6. Reason for change**
1	1	February 8, 2001	First release

Table of Contents

1 Introduction

The muon trigger architecture proposed in reference [1] was designed to minimize the number of connections between the muon detector front-end electronics and the trigger. To achieve this goal the processing is performed in two steps. In the first one, we identify a muon and localize its position using a coarse information. In the second step, the fine tracking and transverse momentum computation are performed using the pad information selected in the first step. This baseline architecture uses ~200 optical links running at 1.6 Gb/s, macro FPGAs and custom backplanes with multidrop and multipoint busses. It does not preserve the synchronization between the data and the *BC identifier*. The extraction of the pad information from the coarse data is done in the FE electronics. It is performed for each candidate one after the other.

In March 2000, the muon group finalized the logical layout of the detector [2]. A detailed description can be found elsewhere in reference [3]. It is projective and contains pads and strips. At that time, the size of the *Field Of Interests* were reduced along the *x* axis and completely closed along *y*.

The reduction in the number of channels and the reduction in the size of the *Field Of Interest* open the road to an architecture fully synchronous and pipelined. Full pad and strip hit maps are transferred to the muon trigger every 25 ns. The muon finding is performed immediately at the pads and strips level. The new architecture is based on ~1250 optical links running at 1.6 Gb/s, macro FPGAs and custom backplane with point to point connections. Synchronization between data and the *BC identifier* is preserved during the whole processing.

In the next section of this document, the overall scheme of the data flow is described together with the corresponding hardware implementation. In section 3 and 4, the latency and the cost are estimated.

2 The synchronous architecture

Our main guidelines to design the muon trigger architecture have been the following:

1. Preserve synchronization between data and *BC identifier* during the whole processing;
2. Custom backplane contains only point to point connections;

3. The hardware is located in the counting room to be insensitive to *Single Event Upset* (SEU);
4. The hardware is simple, compact, flexible and easy to debug, monitor and repair;
5. The implementation can evolve as far as the physics requirements and the technological progress are concerned.

Figure 1 shows an overview of the muon system. The data flow between the muon stations, the *off detector electronics* (ODE) and the muon trigger is described. The muon trigger is connected to the FE electronics through 1248 high speed optical links. The trigger processor is distributed over 4 crates. Each of them is connected to a quarter of the muon system.

The optical links allow to place the trigger in the counting room, 50 meters away from the detector. This place is radiation free. Thus, the muon trigger is immune to *Single Event Upset* and off the shelf components can be used. However, the trigger interface is sensitive to *Single Event Upset* through the optical link drivers. It is implemented in the *ODE Electronics* located in the cavern near the detector.

In the next sections, we describe the organization of the ODE electronics, the trigger interface located in the ODE, the data transfer between the ODE and the trigger, the trigger processing board and the final selection of the candidates.

2.1 The ODE Electronics and its trigger interface

The muon FE electronics is described in reference [4]. It contains three components: the front−end boards, the intermediate boards and the ODE electronics. The *front−end boards* are mounted on the chambers. They contain amplifiers, shapers and discriminators and drive physical channels. The *Intermediate board*s are located on the left and right sides of each station. They transform physical channels into logical ones. The *Off Detector Electronics* are in the middle of the left and right side of each station. They synchronize logical channels with a *BC identifier*, store the information into the L0 and the L1 buffers. They also contain trigger and DAQ interface.

The organization of the ODE boards is shown in Figure 2 for stations M_1, M_2 (M_3) and M_4 (M_5). The design criteria have been the following:

1. The number of inputs per board is limited to 192: the maximum number of LVDS connections affordable on a 9U board with a low cost;
2. Perfect match between the ODE board and the *Logical Unit* structure;
3. Minimal number of ODE boards.

The obtained result is summarized in Table 1.

At the end of the *ODE* processing, the logical channels are stored in the L0 buffer. They are also given to the trigger interface. It pushes the data on high speed optical links connected to the muon trigger. The number of links required per *ODE* board is determined by the organization of the trigger processor.

A trigger processing unit, shown in Figure 5, receives data from station M_1 to M_5 for 1/192 of the muon system. Its input data are carried on separated links to avoid a complicated data distribution in the trigger board. The number of bits per link is equal to 32. It contains the pads/strips information and the *BC identifier* encoded on 4 bits. A single optical link transports only the information of 1 or 2 *Logical Units*. This rule was introduced to preserve the *Logical Unit* structure and to avoid complicated data distribution in the *ODE* board. The number of links per ODE board and their contents are summarized in Table 2.

In this scheme, the trigger interface implemented in the *ODE* is extremely simple. It only drives 6, 8 or 12 optical links. There is no interconnection between the ODE boards. However the matching between the ODE data and the optical links depends on the station and on the region.

A schematic view of this interface is shown on Figure 3 when the information of 2 *Logical Units* are pushed on a single link.

Table 1: Basic numbers describing the ODE boards

Region	Station	Number of boards per station	Inputs per board	Logical Units per board	Optical links per board	Optical Links pattern[1]
1	1	12	192	−	8	4+24
	2−3	8	168	12	6	4+28
	4−5	8	144	−	6	4+24
2	1	12	192	−	8	4+24
	2−3	8	192	12	12	4+16
	4−5	4	168	24	12	4+14
3	1	12	192	−	8	4+24
	2−3	8	168	6	6	4+28
	4−5	4	120	12	12	4+10
4	1	12	192	−	8	4+24
	2−3	8	168	6	6	4+28
	4−5	4	120	12	12	4+10

2.2 Data transfers between the ODE and the muon trigger

The *ODE Electronics* and the muon trigger are connected via ~120 ribbon cables containing 12 optical fibers each. At one end, it is terminated by a ribbon connector while at the other end it is split as shown in Figure 4.

The first and the last muon station are separated by ~7 meters. The collection of data coming from the 5 stations is obtained by splitting the ribbon cable at one end. Thus complicated patch panels are avoided. On the other side, the ribbon connector provides a high density of connections: 24 optical fibers are plugged on a 9U board using only 10% of the height.

2.3 The muon trigger

The muon trigger is distributed over 4 crates. Each of them is connected to a quarter of the muon system. A trigger crate contains 15 processing boards. A

1 The first number is the number of bit used to encoded the BC identifier. The second one is the number of pads / strips information carried on the optical links.

trigger board maps either 1/12 of region R_i (i=1,3,4) or 1/24 or region R_2. Each board is subdivided in 4 *Processing Units*. The organization of the trigger board and of the processing unit are shown in Figure 5 for a quarter of the muon system. It is also summarized in Table 2.

Table 2: Basic numbers describing the muon trigger

Region	Trigger Processing boards	Processing Units per board	Optical links per board	Optical links from					Optical links/ Processing Unit
				M1	M2	M3	M4	M5	
1	12	4	24	8	4	4	4	4	6
2	24	2	16	4	4	4	2	2	8
3	12	4	24	8	4	4	4	4	6
4	12	4	24	8	4	4	4	4	6

A *Processing Unit* performs the muon finding in 1/192 of the muon system. It receives pads and strips data from all muon stations at a frequency of 40 MHz. It looks for a straight track pointing towards the interaction point and crossing all the five muon stations. The search starts with a "virtual pad" hit in station M_3. It is performed in parallel for the 96 "virtual pads" of station M_3 contained in the *Processing Unit*. A "virtual pad" is defined by crossing horizontal and vertical strip.

The track finding is based on the projectivity of the logical layout. The starting point is given by a "virtual pad" hit in station M_3. A road is open in station M_4, M_5, M_2 and M_1. It is shown in Figure 6. Its size depends on the dimension of the *X Fields Of Interest*. Their current values are given in Table 3. For each seed the *Processing Unit* collects all the pads/strips involved in the road.

From the seed position in station M_3, the position of the pad is extrapolated to station M_2, M_4 and M_5 assuming a straight line coming from the interaction point. In station M_4 and M_5, we look for at least one hit within the *Field Of Interest* centered on the extrapolated pad. In station M_2, we look for the closest hit to the extrapolated pad within the *Field Of Interest*.

The position of the track is extrapolated to station M_1, assuming a straight line based on the pad positions in stations M_2 and M_3. Then, the nearest hit is searched within the *Field Of Interest*.

There is a one to one correspondence between the pad hit in station M_3 and the extrapolated one in station M_2, M_4 and M_5 since the logical layout is projective. There is also a one to one correspondence between the pair of pads hit in station M_2, M_3 and the extrapolated pad in station M_1. The corresponding relation are illustrated in Figure 6. They are used to implemented the track finding algorithm using only logical operations (AND, OR,...).

At the end of the processing, the number of muon candidates is limited to 2 per *Processing Unit*. If the number of candidates exceeds this limit, the event is rejected. The transverse momentum is computed for the 2 remaining candidates.

The selection of the 2 candidates with the highest transverse momentum is performed in 2 steps. The first one runs in each trigger processing board. The 2 candidates with the highest P_T are selected from the values produced by the 4 *Processing Units*. In the second step, the final candidates are selected from the 30 muons produced by the 15 processing boards. They are then sent to the L0 decision box by the controller.

Table 3: Size of the Fields Of Interest. The centre of the field is given by the position of the extrapolated pad. The units are pads.

Station	Size of the Field Of Interest	
	X coordinate	*Y* coordinate
1	±3	±0.5
2	±4.5	±0.5
3	–	–
4	±1.5	±0.5
5	±1.5	±0.5

2.3.1 The muon trigger processing board

A scheme of the muon trigger processing board is shown in Figure 7. The board is unique and contains 4 *Processing Units*. Inputs of the board are 24 optical links, grouped in 2 ribbons of 12 fibers each. A *Processing Unit* is connected to 6 optical fibers: 2 for station M_1 and 1 for each station M_j where j=2,...,5 (except for region R_2 of station M_2 (M_3) where 2 optical fibers are needed, see below)

The processing units are arranged following a 2x2 matrix. The left and right columns are interconnected to allow the data exchange required by the muon tracking along the x−axis. The top and bottom rows are also interconnected to solve particular cases appearing in station M_2 (M_3) for region R_2 and R_1:

1. Data from region R_2 of station M_2 (M_3) are carried on 2 optical links while only one is used in the other regions. The additional link is connected to the processing unit on the bottom row. It is the only region where the bottom *Processing Units* are not involved in the track finding. The top−bottom interconnection allows the additional data transfer to the active top *Processing Units.*
2. A *Logical Unit* is shared by two *Processing Unit* belonging to the same column in region R_1 of station M_2 (M_3). The top−bottom interconnection allows the corresponding data distribution.

Processing Units are also interconnected to the backplane to insure data exchange between boards. This exchange is performed via point to point connections running at 80 MHz.

The *Processing Unit* includes 2 look−up tables (LUT) to compute the transverse momentum of the 2 muon candidates. The transverse momentum is encoded on 8 bits and the address of the candidate on 14 bits: the M_3 pad address within the *Processing Unit* (7b), the M_2 pad address within the *Field Of Interest (*4b), The M_1 pad address within the *Field Of Interest* (3b). The size of one LUT is equal to 16 kbytes.

The 4 *Processing Units* are connected to the module performing the final selection of the candidates. It contains 2 functions:

1. *Select the 2 candidates with the highest P_T in the processing board.*
   A maximum of 4×2 candidates are produced every 25 ns. We select two of them with the highest P_T. Thus, we have a maximum of 30 candidates distributed over 15 processing boards at the end of this stage.
2. *Select the 2 candidates with the highest P_T through the processor.*
   The final selection relies on a daisy chain. The 2 candidates of board #1 are joined with those of board #2. The latter select the 2 out of 4 with the highest P_T. They are then transmitted to the board #3 and so on so forth. At the end of this chain, the controller gets the 2 candidates with the highest P_T within a quarter of the muon system. They are sent to the L0 decision box.

The last module are the L1 buffer and the interface to ECS. The latter is based on a credit card PC since we are in the counting room. This interface allows to download FPGA using JTAG protocol . It is also used for debugging and monitoring.

2.3.2 The Processing Unit

A scheme of the muon processing board is shown in Figure 8. The number of inputs and outputs for a *Processing Unit* is mainly governed by the pads/strips information exchanges between processing units and between boards. To minimize their number, the data coming from the optical links and the custom backplane are transferred at a speed of 80 MHz. The I/O for a *Processing Unit* are summarized in Table 9 in Appendix 1.

A *Processing Unit* contains 6 main functions required by the track finding:

1. *Synchronization*.
   Muon data corresponding to a given bunch crossing arrive at different times at the output of the optical link.
2. *Muon identification*.
   It is performed using logical operations applied to pads and strips coming from station M_2, M_3, M_4 and M_5. The 96 seeds of station M_3 are treated in parallel. For each seed, 5 bits are output: Valid candidate (1b), M_2 address within the *Field Of Interest* (4b). Data for the 96 seeds are given to the M_1 pad finding unit.
3. *M_1 pad finding.*
   It receives the information coming from station M_1 and from the muon identification box. It performs the extrapolation in station M_1 and finds the nearest hit. The corresponding logical operations are performed in parallel for the 96 seeds.
4. *Selection of the 2 candidates over 96.*
   The selection function finds a maximum of 2 candidates out of 96. If the number of candidates is above this limit a flag is set. It will be used later on to reject the event. The address of the each candidate appears on the output encoded in a 14−bit word.
5. *P_T computation*.
   It is performed using 2 look−up tables of 16 kbytes each.

During the whole processing the *BC Identifier* is carried with the data.

A *Processing Unit* also contains a L0 buffer, error detection of synchronization and an interface to the debugging/monitoring circuits.

The simulation of the main functions and the estimation of the number of I/O show that a Processing *Unit* can be implemented in an APEX20K600 FPGA chip containing ~600 000 gates, 38 kbytes of RAM and 488 I/O pins.

2.3.3 Data exchange between boards

Pads and strips hit maps are exchanged between boards. The number of data being exchanged is driven by the size of the *Field Of Interest*. A map is given in Table 10 of Appendix 2.

Data are merged when they move from region R_i to R_{i+1}. Thus all data collected in region R_{i+1} have the same granularity. This rule minimizes the number of data exchanges on the backplane.

To evaluate the feasibility of the backplane, we count the number of pins required to connect a trigger processing board and a controller. Links embedded in the backplane are of point to point type. They are implemented with GTL+ technology. It requires 2 physical pins for each signal in order to minimize crosstalk and noise. A summary of the number of pins is given in Table 4 for a processing board and in Table 5 for a controller. They are equal to 492 and 630† respectively. These numbers are below the limit, given by the total number of physical pins available on the VME64 9U backplane: ~100 for power supplies and ground, ~650 for user applications.

Table 4: Number of pins required to connect a processing board

	Signal	Speed	Techno.	Pin/Signal	Pins
Synchronization	6	40 MHz	GTL+	2	12
Data exchange	334	80 MHz	GTL+	2	334
P_T Daisy Chain	120	80 MHz	GTL+	2	120
Services	11	40 MHz	GTL+	2	22
L1 to DAQ [90 kHz]	2	80 MHz	GTL+	2	4
Total					**492**

† Our estimation is based on the synchronization protocol described in [1]. A new one based on circular buffers is under study. It will reduce the number of pins to 450, for the controller.

Table 5: Number of pins required to connect the controller

	Signal	Speed	Techno.	Pin/Signal	Pins
Synchronization	90	40 MHz	GTL+	2	180
P_T Daisy Chain	60	80 MHz	GTL+	2	60
Services	165	40 MHz	GTL+	2	330
L1 to DAQ [90 kHz]	30	80 MHz	GTL+	2	60
Total					**630**

3 Latency

The estimation of the processing time is based on the simulation of the basic components implemented in a *Processing Unit*. We don't take into account the time to transport the signal from the cavern to the counting room since it is already included in the fixed time of the L0 latency [5]. A summary of the muon processing time is given in Table 6. It is equal to ~1 µs well below the 1.2 µs authorized for the muon trigger.

Table 6: The muon trigger processing time

	Clock	Time [ns]
Optical link synchronization	3	
Data exchange + synchronization	7	
Muon identification	1	
M1 pad finding	1	
Extraction of the 2 candidates (PU)	4	
P_T computation	2	
Candidate selection (board)	4	
Candidates selection (Processor)	16	
Transmission to L0 decision	2	
Total	**40**	**1000**

4 Cost estimate

Table 7: Cost of the trigger interface implemented in the ODE Electronics

	Quantity	Spare	Price/unit [KCHF]
Optical transceivers	1248	120	0.447
Total			**306−611**

Table 8: Cost of the muon trigger processor.

	Quantity	Spare	Price/unit [kCHF]
Trigger Processing Board	60	6	11.1
Controller	4	2	10.6
backplane	4	1	7
Optical ribbon cables	120	12	0.8
Crates	4	2	10
Monitoring			25
Total			**900−1052**

The cost of the synchronous architecture is detailed in Table 7 and Table 8. The former contains the additional cost which has to be included in the muon FE electronics while the latter is the price of the muon trigger processor.

Our estimation is based on the price of the optical link components paid in the end of 1999 with a discount of 20%. The latter is expected when more than a thousand of pieces are bought. It is also based on a real quotation for the PCB manufacturing proposed by Techci in the end of 2000. The price of the FPGA for a *Processing Uni*t is extrapolated to 2003. We take the value of 300 $ provided by Altera for the APEX20K600 chip.

The cost estimate was done for two cases. A pessimistic approach based on the current price of the optical links. An optimistic approach where the price of the optical links is divided by a factor 2. A reduction of the cost is expected in the coming years with the increased number of manufacturers and the size of the market: selected components will be used in high speed networking.

The additional cost of the muon FE electronics varies between 306−611 kCHF. It strongly depends on the price of the optical transceivers.

The cost of the muon trigger processor amounts to ~1 MCHF taking into account spares. It is rather insensitive to the price of the optical components.

5 Conclusions and prospects

The proposed synchronous architecture for the L0 muon trigger is extremely attractive. It is synchronous, insensitive to *Single Event Upset*. The trigger interface is simple. The muon trigger is based on a single processing board design where all functions are well separated. The custom backplane contains only point to point connections. Debugging, final integration and maintenance will be rather straight forward.

The implementation requires a high level of integration. It has to cope with a large number of data exchanged between boards. It can be built using available technologies like high speed optical links running at 1.6 Gb/s, FPGA's with 600 000 gates, point to point connections at 80 MHz based on GTL+.

References

[1] E. Aslanides et al., The L0(μ) processor, LHCb Note, LHCb 99−008

[2] P. Colrain et al., Optimization of the muon system logical layout, LHCb Note, LHCb 2000−016

[3] B. Schmidt et al., LHCb muon system by numbers, LHCb Note, LHCb 2000−089

[4] A. Lai et al., Muon detector front−end architecture, LHCb Note, LHCb 2000−017

[5] J. Christiansen et al., The latency of the level 0 trigger, LHCb Note, LHCb Note 99−015

Appendix 1

Table 9: I/Os of a generic Processing Unit.

		Signal	Clock	Valid
OL from M1 (80 MHz)		32	2	2
OL from M2 (80 MHz)		16	1	1
OL from M3 (80 MHz)		16	1	1
OL from M4 (80 MHz)		16	1	1
OL from M5 (80 MHz)		16	1	1
Top−bottom interconnection (80 MHz)		32	2	2
To backplane (80 MHz)	M1	20	1	
	M2	4	1	
	M4	3	1	
	M5	3	1	
From backplane (80 MHz)	M1	20	1	
	M2	6	1	
	M4	4	1	
	M5	4	1	
PU left to right exchange (40 MHz)	M1	40		
	M2	8		
	M4	5		
	M5	5		
PU right to left exchange (40 MHz)	M1	40		
	M2	8		
	M4	5		
	M5	5		
PT candidates (40 MHz)		60		
Master clock (40 MHz)			1	
Synchronization (40 MHz)		6		
BC identifier (40 MHz)		8		
L0 accept (40 MHz)		1		
Monitoring bus		16	1	
JTAG		3		
Total			**428**	

Appendix 2

To	Station	11	12	13	21u	21d	22u	22d	23u	23d	31	32	33	41	42	43	Emission
11	1 2 4 5		80 16 8 8														112
12	1 2 4 5	80 16 8 8							40 12 8 8								180
13	1 2 4 5									40 12 8 8							68
21u	1 2 4 5						40 8 5 5										58
21d	1 2 4 5							40 8 5 5									58
22u	1 2 4 5				40 8 5 5								20 6 3 3				90
22d	1 2 4 5					40 8 5 5							20 6 3 3				90
23u	1 2 4 5		20 8 6 6										20 6 3 3				72
23d	1 2 4 5			20 8 6 6									20 6 3 3				72
31	1 2 4 5											80 16 10 10					116
32	1 2 4 5										80 16 10 10					40 12 6 6	180
33	1 2 4 5						10 3 3 3	10 3 3 3	10 3 3 3	10 3 3 3						40 12 6 6	140
41	1 2 4 5														80 16 10 10		116
42	1 2 4 5													80 16 10 10			116
43	1 2 4 5											20 6 6 6	20 6 6 6				76
Reception		**112**	**152**	**40**	**58**	**58**	**77**	**77**	**87**	**87**	**116**	**154**	**166**	**116**	**116**	**128**	
Em+Rec		**224**	**332**	**108**	**116**	**116**	**167**	**167**	**159**	**159**	**232**	**334**	**306**	**232**	**232**	**204**	

Table 10: Map of the data exchange between trigger processing boards

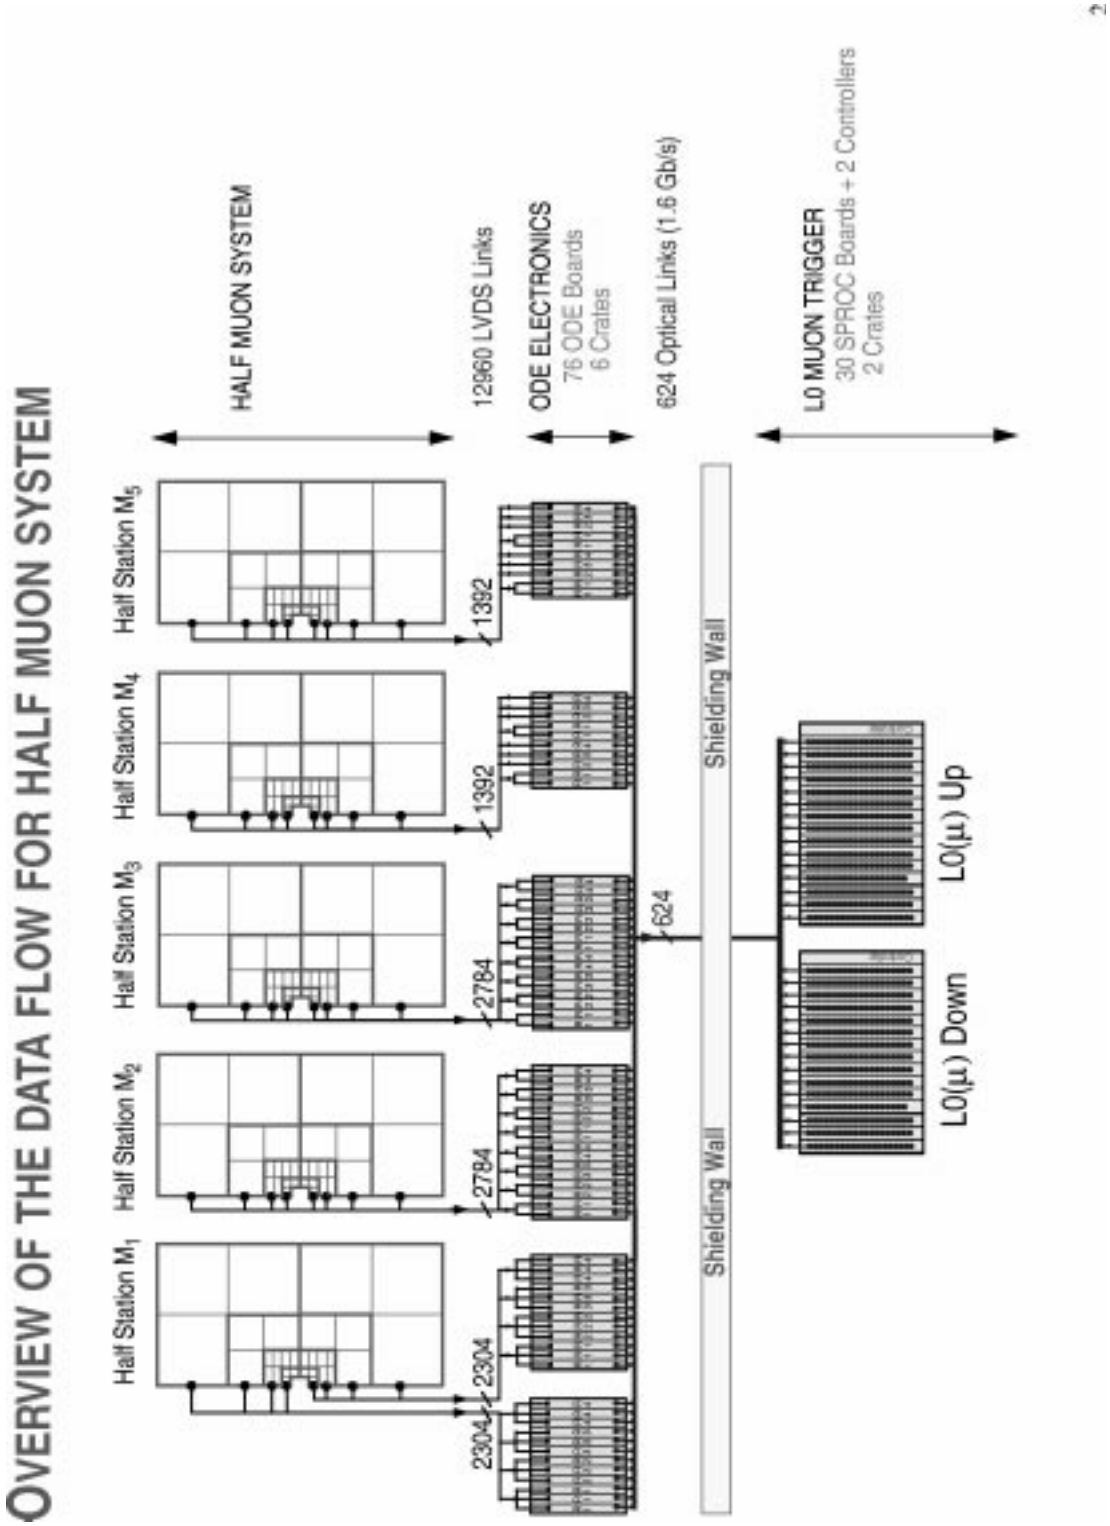

Figure 1: Overview of half a muon system

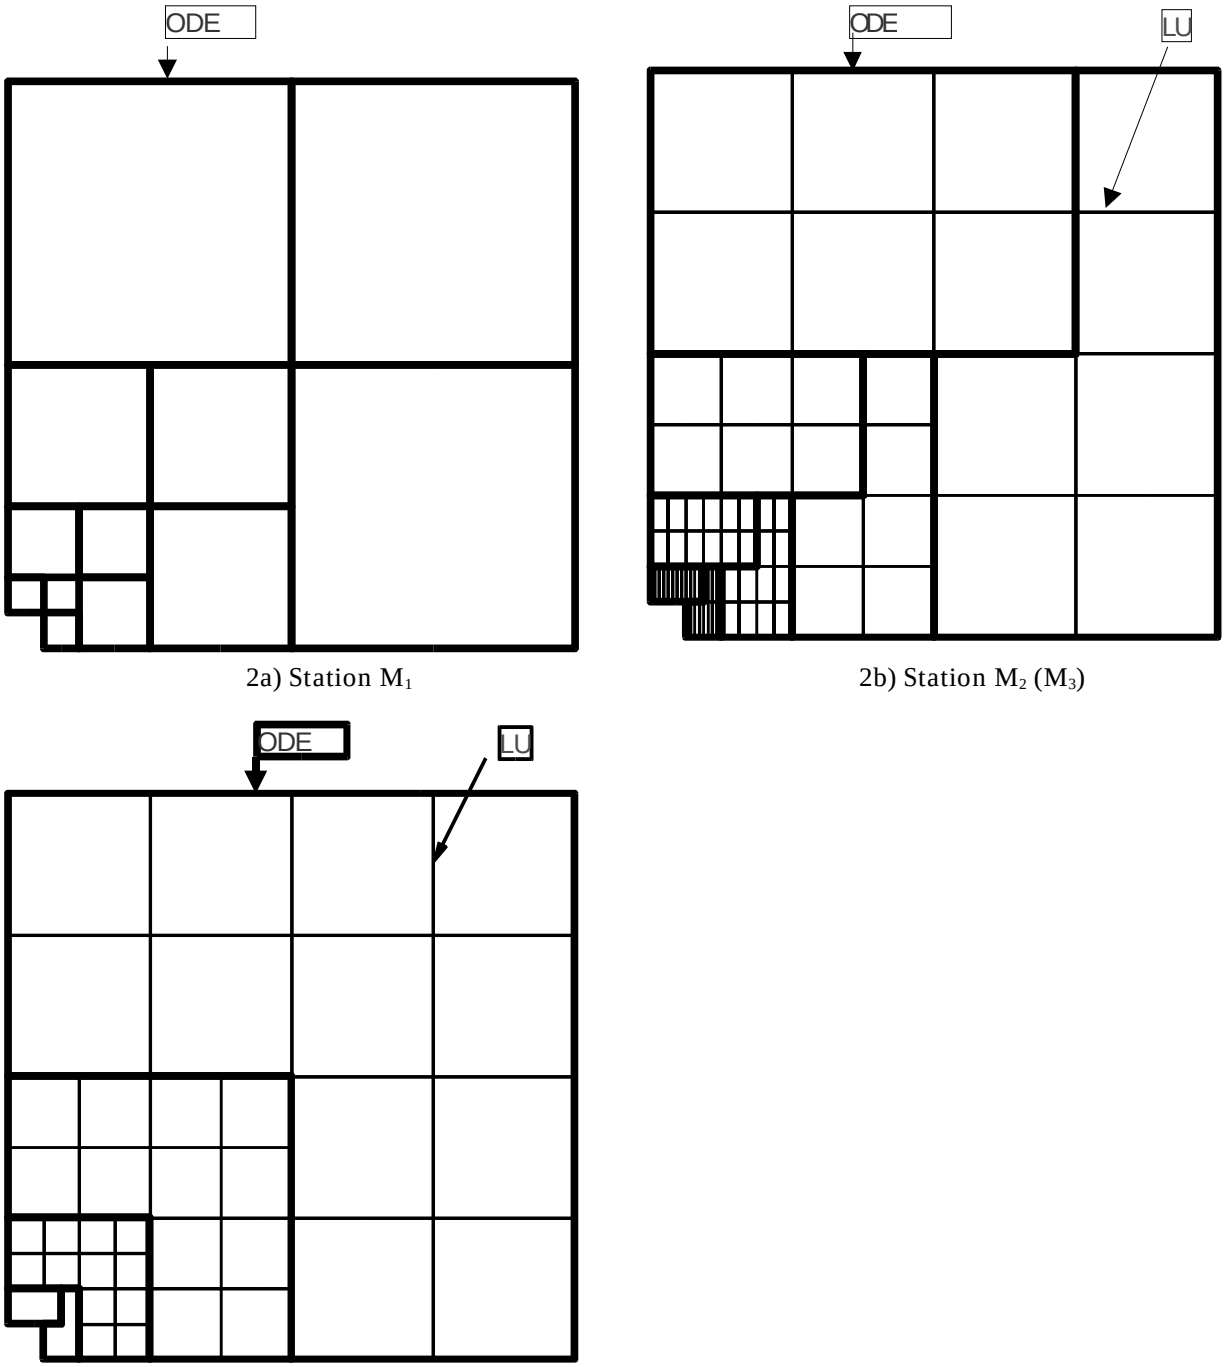

2a) Station M_1

2b) Station M_2 (M_3)

2c) Station M_4 (M_5)

Figure 2: Organization of the ODE Electronics for a quarter of station: 2a) M_1 ; 2b) M_2 or M_3 ; 2c) M_4 or M_5. The thick lines delimit the boundaries of the ODE boards while the thin lines show the Logical Unit structure

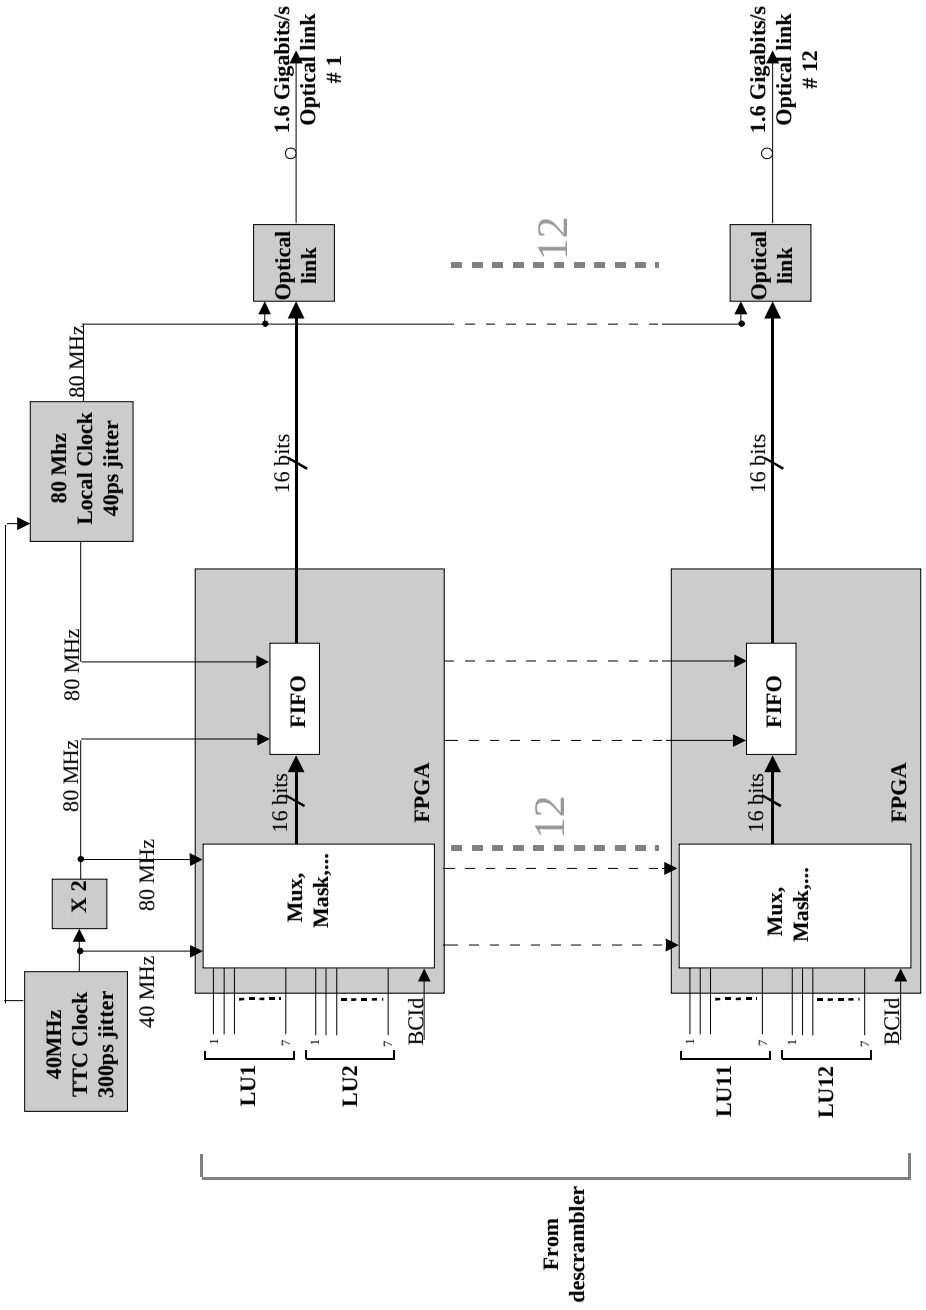

Figure 3: Scheme of the trigger interface implemented in the ODE Electronics. In this figure the strip hit maps of 2 Logical Units and the BC Identifier are merged on a single optical link. This interface can drive up to 12 optical links

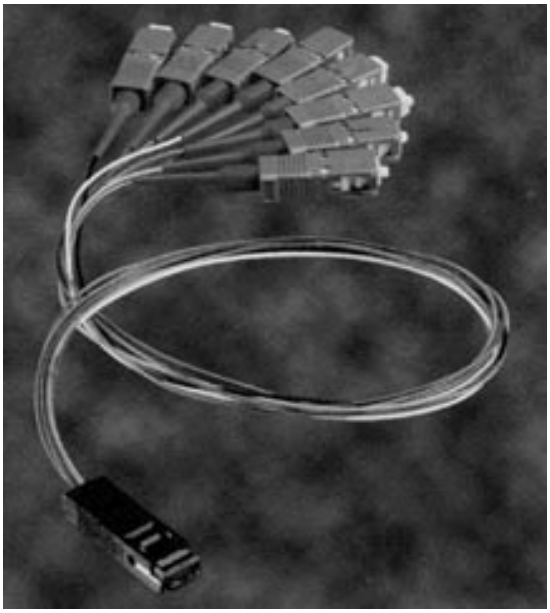

4a)

4b)

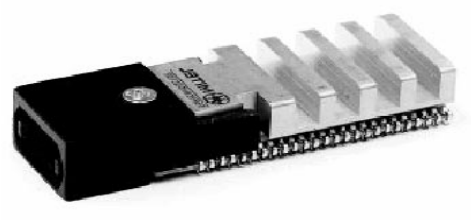

4c)

Figure 4: View of Optical components: 4a) ribbon of 12 optical fibers split at one end (AMP 492512−3); 4b) and 4c) ribbon connector including optical transceivers running up to 2.5 Gb/s. Their typical size is 1.7 cm × 1.6 cm × 4.7 cm (e.g. MITEL MFR62340−J)

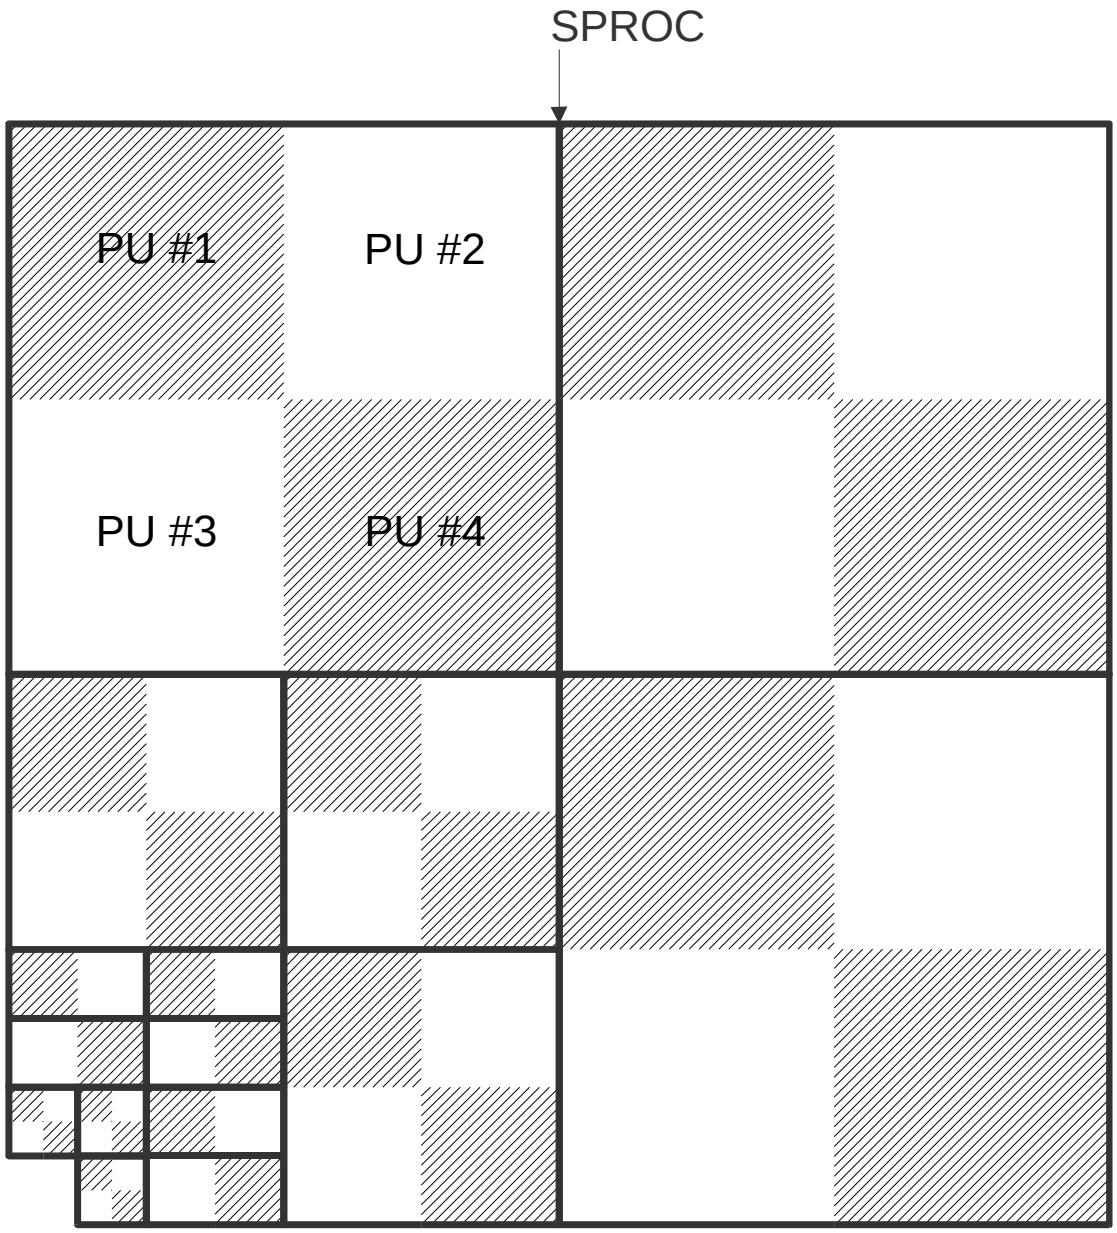

Figure 5: Organization of the trigger processing boards and of the Processing Units. The thick lines delimit the trigger processing boards. The hatch and white areas show the structure of the Processing Units

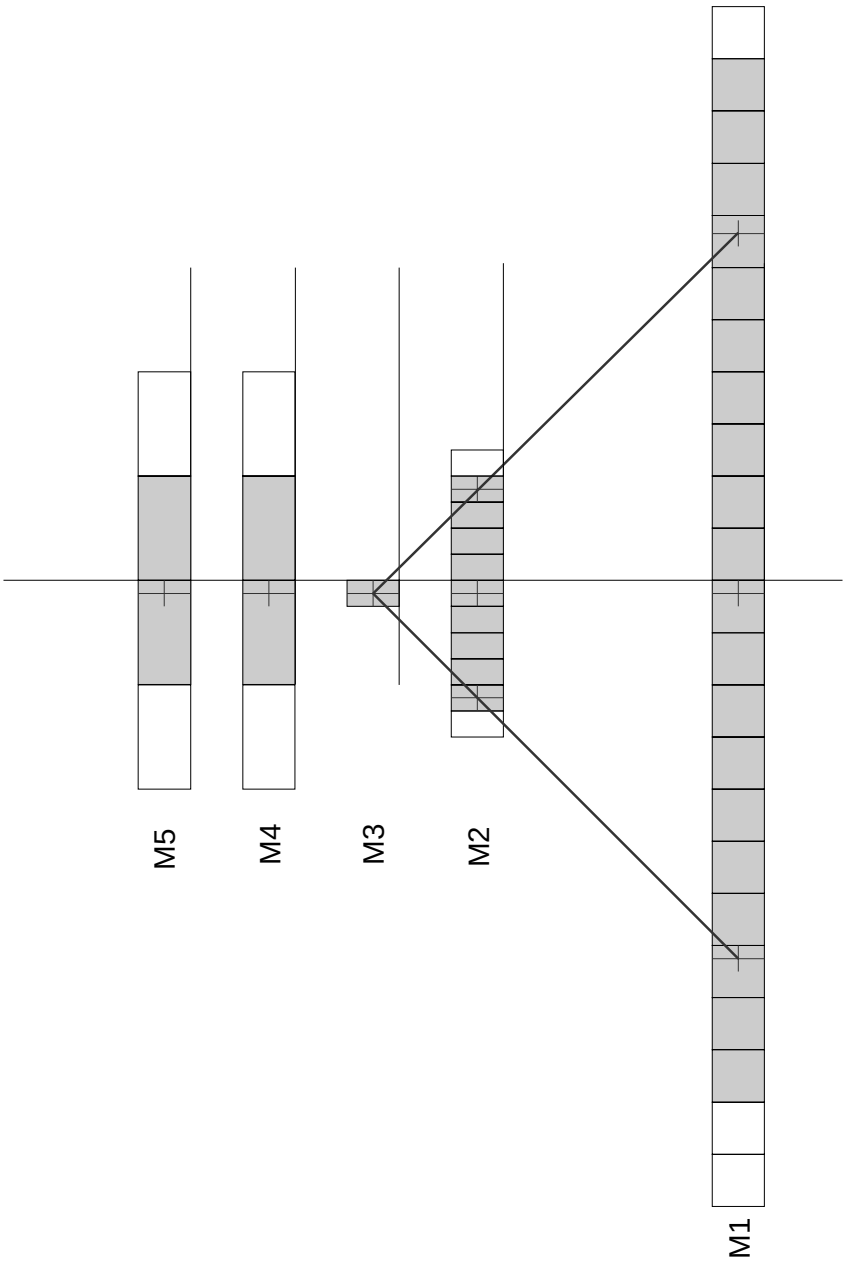

Figure 6: The track finding road. The starting point is given by a "virtual pad" hit in station M_3. The pad in the road are in gray. The crosses show the extrapolated position in station M_2, M_4 and M_5. The line joining the centres of pads in station M_2 and M_3 points to the extrapolated pad in station M_1

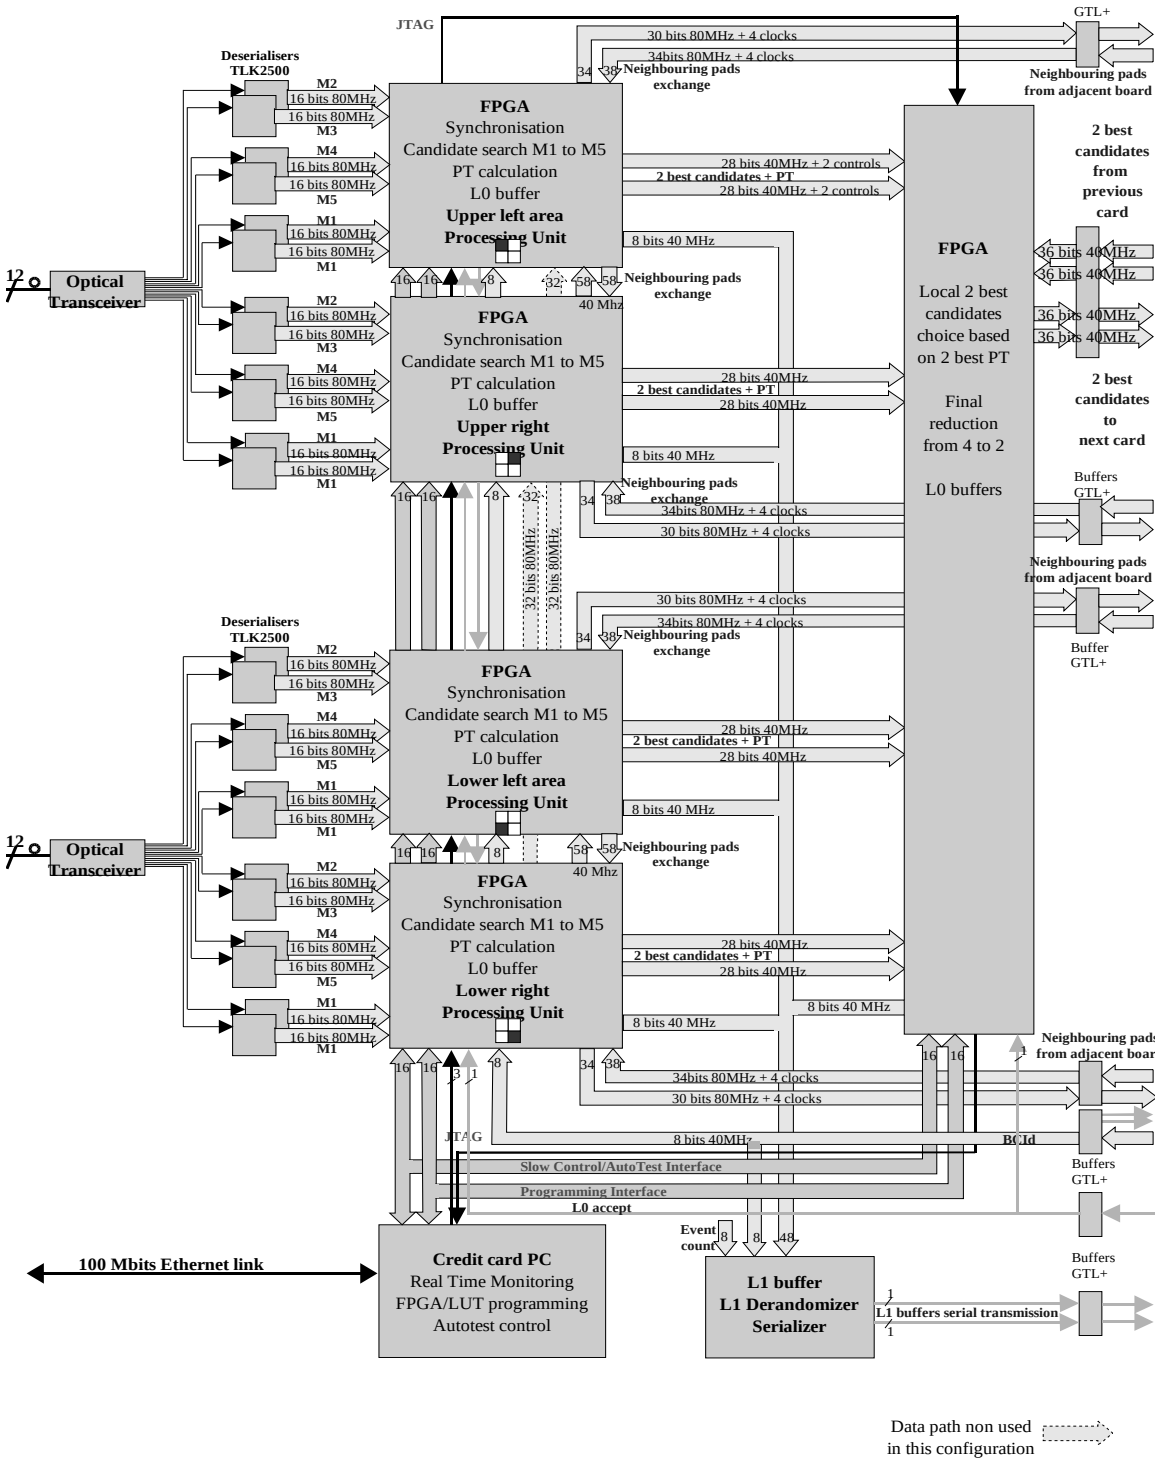

Figure 7: Scheme of the trigger processing board

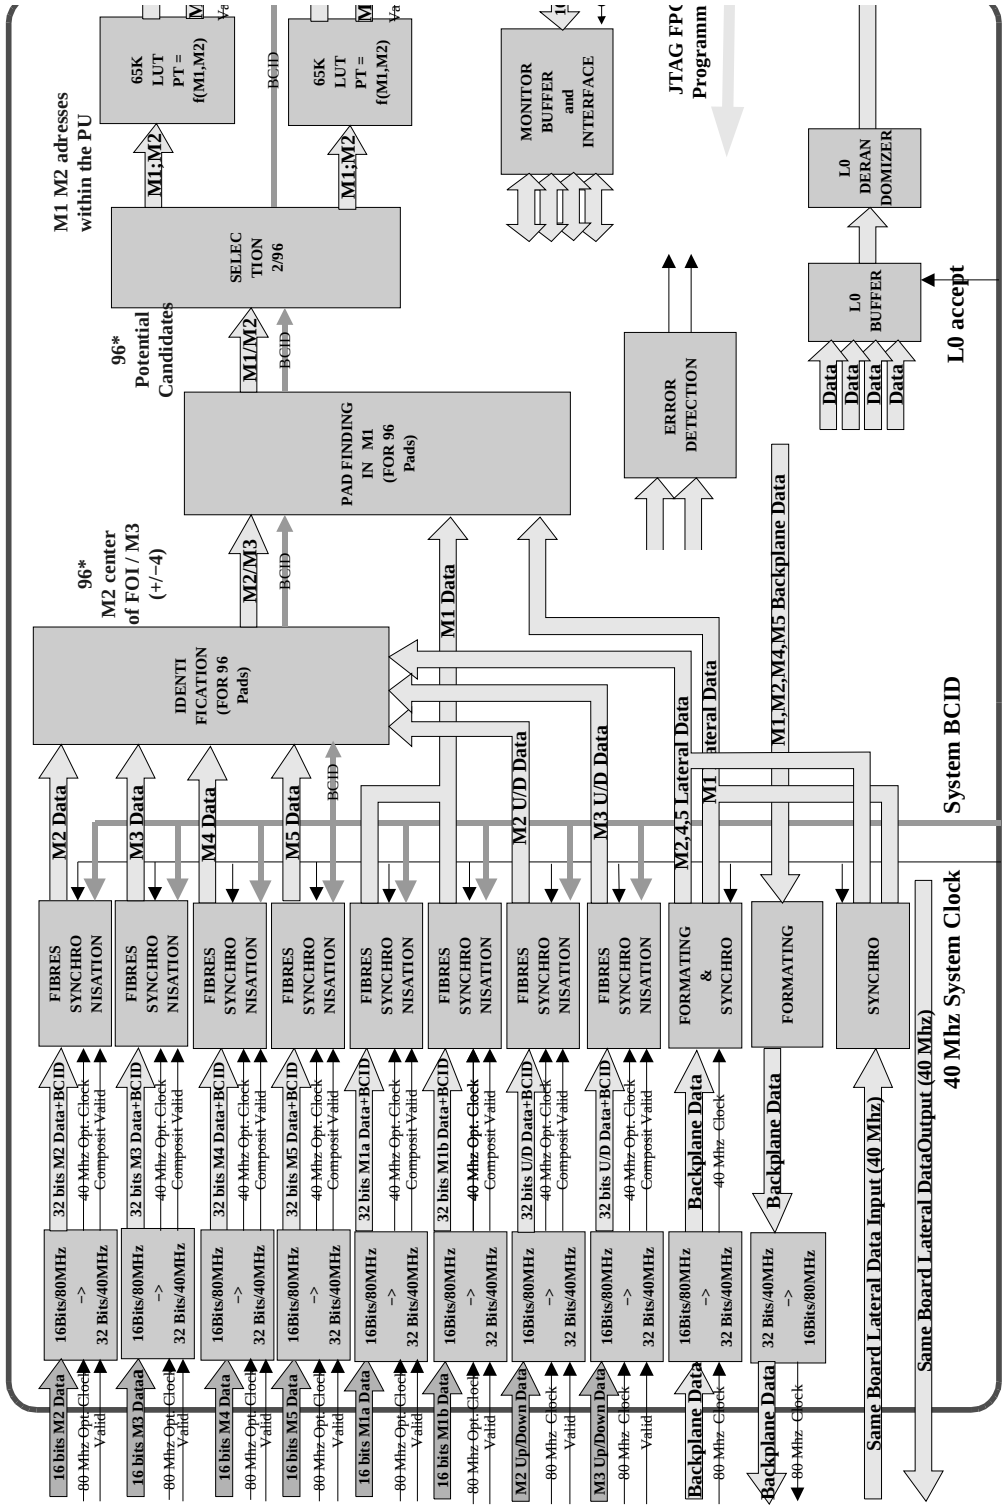

Figure 8: Scheme of the Processing Unit