LHCb BASE-LINE LEVEL-0 TRIGGER 3D-FLOW IMPLEMENTATION

Dario B. Crosetto

900 Hideaway Pl.
DeSoto, TX 75115
Crosetto@vxcern.cern.ch; Crosetto@physics.rice.edu

Original version 17 February 1999; added Appendix upon request by L0 trigger group 2 June 1999

Abstract

The LHCb (Large Hadron Collider Beauty Experiment at CERN, Geneva, Switzerland) Level-0 trigger implementation with the 3D-Flow system is described in detail using components and technology available today. It offers full programmability, allowing it to adapt to unexpected operating conditions and enabling new, unforeseen physics.

The 3D-Flow Processor system is a new, technology-independent concept in very fast, real-time system architectures. Based on the replication of a single type of circuit of approximately 100K gates, which communicates in six directions: bi-directional with North, East, West, and South neighbors, unidirectional from Top-to-Bottom, the system offers full programmability, modularity, ease of expansion and adaptation to the latest technology.

A complete study of its applicability to the LHCb Calorimeter triggers is presented. Full description of the input data handling, either in digital or mixed digital-analog form, of the data processing, and the transmission of results to the global level-0 trigger decision unit are provided.

Any level-0 trigger algorithm (2x2, 3x3, 4x4, etc.) with up to 20 steps, can be implemented with zero dead time, while sustaining input data rate (up to 32-bit per input channel, per bunchcrossing) at 40 MHz. For each step, each 3D-Flow processor can exchange up to 26 operation, inclusive of compare, ranging, finding local maxima, and efficient data exchange with neighboring channels. (One to one correspondence between input channel and trigger tower).

It is shown how the whole Level-0 calorimeter trigger can be accommodated into 6 crates (9U), each containing 16 identical boards carrying only two main types of components, front-end FPGAs (Field Programmable Gate Array) and 3D-Flow processors.

All 3D-Flow inter-chip Bottom-to-Top port connections are all contained on the board (data are multiplexed 2:1, printed circuit board (PCB) traces are shorter than 6 cm); all 3D-Flow inter-chip North, East, West, and South port connections between boards and crates are multiplexed (8+2):1 and are shorter than 1.5 meters.

Full implementation of a 3D-Flow system, for the most complex trigger algorithm, requires 320 cables to North and South crates and 40 cables to East and West crates (cable cost = \$2 each).

For applications requiring a simpler real-time algorithm (e.g., requiring less than 20 steps, which is equivalent to 10 layers of 3D-Flow processors), the number of connections for the inter-boards (North and South), and inter-crates (East and West) will also be reduced to the number of layers used by the simpler algorithm, thus not requiring all cables to be installed (e.g., applications requiring only 9 layers of 3D-Flow processors will save 32 cables to the North, 32 to the South, 4 to the East, and 4 to the West crates).

Details are also given on timing and synchronization issues, ASIC (Application Specific Integrated Circuit) design verification, real time performance monitoring and design (software and hardware) development tools.

TABLE OF CONTENTS

1 INTRODUCTION	4
2 THE 3D-FLOW: A SINGLE TYPE OF CIRCUIT FOR SEVERAL ALGORITHMS	4
2.1 SYSTEM LEVEL	4
2.2 SYSTEM ARCHITECTURE	5
2.3 PROCESSOR ARCHITECTURE	5
2.4 INTRODUCING THE THIRD DIMENSION IN THE SYSTEM.....	6
2.5 THE 3D-FLOW ARCHITECTURE OPTIMIZED FEATURES FOR THE FIRST LEVELS OF TRIGGERS.	7
3 A SINGLE TYPE OF COMPONENT FOR SEVERAL ALGORITHMS	7
3.1 THE EVOLUTION OF IC DESIGN.....	8
3.2 TECHNOLOGY-INDEPENDENT 3D-FLOW ASIC.....	8
4 FIRST-LEVEL TRIGGER ALGORITHMS	9
5 LHCB LEVEL-0 TRIGGER OVERVIEW	10
5.1 PHYSICAL LAYOUT.....	10
5.2 LOGICAL LAYOUT.....	11
5.3 ELECTRONIC RACKS (FUNCTIONS/LOCATIONS)	12
6 A SINGLE TYPE OF BOARD FOR SEVERAL ALGORITHMS	13
6.1 3D-FLOW MIXED-SIGNAL PROCESSING BOARD (OPTION 1)	13
6.2 3D-FLOW DIGITAL PROCESSING BOARD (OPTION 2)	15
6.3 FRONT-END SIGNAL SYNCHRONIZATION/PIPELINING/DERANDOMIZING/TRIGGER WORD FORMATTER.....	18
6.3.1 Level-0 front-end, VHDL coding of the I/O and signals definition	19
6.3.2 Level-0 front-end, VHDL coding of registering the input data on rising edge	21
6.3.3 Level-0 front-end, VHDL coding of the updating contents of variable delays.....	21
6.3.4 Level-0 front-end, VHDL coding of selecting a variable delay	22
6.3.5 Level-0 front-end, VHDL coding of the formatting of the trigger word and multiplexing the output.	22
6.3.6 Level-0 front-end, VHDL coding for the 128 pipeline buffer	23
6.3.7 Level-0 front-end, VHDL coding of moving accepted L-0 event from 128 pipe to FIFO.....	23
6.3.8 Mapping the Level-0 front-end circuits into ORCA OR3T30 FPGA.....	24
6.3.9 Front-end circuits 'Real-Estate' and cost considerations.....	24
6.4 LOGICAL-TO-PHYSICAL LAYOUT OF 64 CHANNELS/10 LAYERS ON THE 3D-FLOW BOARD	25
6.5 ON-BOARD DATA-REDUCTION, CHANNEL-REDUCTION AND BOTTOM-TO-TOP LINKS	25
6.6 DETAILS OF THE ON-BOARD BOTTOM-TO-TOP LINKS (6 CM)	27
7 CRATE(S) FOR 3D-FLOW SYSTEMS OF DIFFERENT SIZES.....	28
7.1 CRATE BACKPLANE LVDS LINKS NEIGHBORING CONNECTION SCHEME	28
7.2 NUMBER OF NEWS LINKS FOR THE CHIP-TO-CHIP, BOARD-TO-BOARD, CRATE-TO-CRATE	29
7.3 IMPLEMENTATION OF THE BACKPLANE CRATE-TO-CRATE LVDS LINKS (OPTION 1).....	29
7.4 IMPLEMENTATION OF THE BACKPLANE CRATE-TO-CRATE LVDS LINKS (OPTION 2).....	30
7.5 THE 3D-FLOW CRATE	31
8 GLOBAL LEVEL-0 TRIGGER.....	32
9 TIMING AND SYNCHRONIZATION ISSUES OF CONTROL SIGNALS.....	33
10 ASIC DESIGN VERIFICATION.....	34
11 HOST COMMUNICATION AND MALFUNCTIONING MONITOR.....	35
12 SOFTWARE TOOLS.....	36
APPENDIX A: COST/PERFORMANCE COMPARISON OF THE CALORIMETER FRONT-END AND LEVEL-0 TRIGGER	38

APPENDIX B: FRONT-END ELECTRONICS FOR THE PRESHOWER DETECTOR	41
APPENDIX C: HIGH SPEED BACKPLANES	41
Figure 1. The 3D-Flow Processing Element (PE) or "logical unit."	5
Figure 2. One layer (or stage) of 3D-Flow parallel processing	6
Figure 3. General scheme of the 3D-Flow pipeline parallel-processing architecture.	7
Figure 4. The evolution of IC design.	8
Figure 5. Technology-independent 3D-Flow ASIC	8
Figure 6. Overview of the use of the 3D-Flow System in particle identification in HEP.	9
Figure 7. LHCb Level-0 Trigger - Physical Layout.	10
Figure 8. LHCb Level-0 Trigger - Logical Layout.	11
Figure 9. On-Detector Electronics for Level-0 trigger.	12
Figure 10. Off-Detector Electronics for level-0 trigger.	12
Figure 11. Electronics in the Control Room for the Calorimeter Level-0 Trigger Monitoring.	12
Figure 12. Mixed-signal processing board (front view).	14
Figure 13. Mixed-signal processing board (rear view).	14
Figure 14. Digital processing board (front view).	17
Figure 15. Digital processing board (rear view).	17
Figure 16. Front-end signal synchronization, pipelining, derandomizing, and trigger word formatting.	19
Figure 17. VHDL code and graphical representation of registering input data.	21
Figure 18. VHDL code and graphical representation for the updating of the variable delays.	21
Figure 19. VHDL code and graphical representation for the selection of the variable delays.	22
Figure 20. VHDL code and graphical representation for formatting and multiplexing the trigger word.	22
Figure 21. VHDL code and graphical representation of the 128 pipeline buffer.	23
Figure 22. VHDL code and graphical representation for moving accepted data from the pipeline to the FIFO.	23
Figure 23. 3D-Flow layer interconnections on the PCB board.	25
Figure 24. Bottom-to-Top Links on the PCB board.	26
Figure 25. Position of the bypass switches for the data flow (Input/Output) from Top-to-Bottom ports.	27
Figure 26. Bottom-to-Top Links on the PCB (details).	27
Figure 27. 3D-Flow System LVDS Links Neighboring Connection Scheme.	28
Figure 28. 3D-Flow North, East, West, and South LVDS Links.	29
Figure 29. Crate-To-Crate Backplane LVDS Links (Option 1).	29
Figure 30. Crate-To-Crate Backplane LVDS Links (Option 2).	30
Figure 31. The 3D-Flow Crate.	31
Figure 32. LHCb Programmable Global Level-0 Trigger Decision Units.	32
Figure 33. Scheme of the control signal distribution with minimum skew.	33
Figure 34. ASIC design verification. From user's system algorithm down to the gate-level circuit.	34
Figure 35. Demonstrator of a System Monitor for 128 3D-Flow Channels.	35
Figure 36. Interrelation between entities in the Real-Time Design Process.	37
Figure 37. Design Real-Time software tools	37
Table 1. The 3D-Flow architecture optimized features for first-level trigger algorithms	7
Table 2. VHDL code for the definition of the architecture of the front-end circuit.	20
Table 3. Mapping the Level-0 front-end circuit into ORCA OR3T55 FPGA.	24
Table 4. System Monitor Demonstrator test results for 128 channels.	36
Table 5. System Monitor estimated timing for 1024 channels.	36
Table 6. Calorimeter Front-end and Level-0 trigger cost implementation comparison (the cost of the front-end electronics of the PreShower and of the Pad Chamber are not included in this table).	38
Table 7. Cost comparison of the Preshower Front-end electronics.	39
Table 8. Calorimeter Level-0 Trigger Features/Performances.	40

1 INTRODUCTION

The importance of flexibility and programmability for the trigger systems of today's sophisticated High Energy Physics (HEP) experiments has been recognized repeatedly. As a recent example, in an article presented at the 1998 workshop on electronics for LHC experiments¹, Eric Eisenhandler states that "Triggering of LHC experiments presents enormous and unprecedented technical challenges [and that].... first level or two of these trigger systems must work far too fast to rely on general-purpose microprocessors... Yet at the same time must be programmable. ... This is necessary in order to be able to adapt to both unexpected operating conditions and to the challenge of new and unpredicted physics that may well turn up."

The 3D-Flow system was conceived to satisfy exactly such stringent requirements. The result was a system suitable for application to a large class of problems, extending over several fields in addition to HEP, for which it was originally devised.

In the following, after a description of the general architecture and properties of the 3D-Flow concept, all the aspects of its application to LHCb Level-0 trigger are discussed in detail. In particular, all the details of the circuits, components and assembly, as they can be achieved with today's technology, are provided. When compared with competing proposals, the 3D-Flow solution offers system sizes and costs at least 50% lower than the alternatives, while maintaining the important advantages of full programmability, modularity, scalability and ease of monitoring.

The style of the description is in a bottom-up fashion: circuit, architecture vs. trigger needs (see Table 1), chip, board, crate, system, global trigger decision unit, timing and synchronization of control signals, real-time malfunctioning monitor, development and design verification tools.

All sections and figures are related to Figure 8 in the following manner:

- Figures 9, 10, and 11: physical aspects of the logical elements shown in Figure 8.
- Section 2: functional information on the conceptual architecture of the 3D-Flow indicated by 8 in Figure 8.
- Section 3: physical implementation of the basic element 3D-Flow that was conceptually described in Section 2.
- Section 4: general information on fully programmable first level triggers as the part indicated by 8 in Figure 8.
- Section 6.1: physical implementation in 9U boards of what is indicated by the numbers 6, 7, 8, and 9 in Figure 8.
- Section 6.2: alternative detailed implementation of Section 6.1 with the information of the physical implementation in 9U boards of what is indicated by the numbers 7, 8, and 9 in Figure 8.
- Section 6.3: logical functions indicated by 7 in Figure 8.
- Sections 6.4, 6.5 and 6.6: relationship between physical and logical layout of what is indicated by 8 in Figure 8.
- Section 7: physical implementation in 9U crates of what is indicated by the numbers 6, 7, 8, 9, and 10 in Figure 8.
- Section 8: logical function indicated by 10 in Figure 8.
- Section 9: control signals distribution among what is indicated by 6, 7, 8, 9, and 10 in Figure 8.
- Section 11: description of the box to the right of Figure 8, L0 CAL (Level-0 Calorimeter real-time monitoring system). The other box, L0 MUON has the same functionality and implementation.
- Section 12: tools which allow simulation of the programmable system indicated in Figure 8 by 8, 9, and the L0 Calorimeter and Muon monitor.

2 THE 3D-FLOW: A SINGLE TYPE OF CIRCUIT FOR SEVERAL ALGORITHMS

The system is based on a single type of replicated circuit called 3D-Flow processing element² (PE) consisting of about 100K gates. Several PEs can be put into a single component. The 3D-Flow PE circuit is technology-independent. Implementation with the current technology of 0.25 μm which has a gate count of $\sim 30\text{K}$ gates per mm^2 requires about 3 mm^2 of silicon per PE. A chip accommodating 16 PEs requires a silicon area of about 50 mm^2 in 0.25 μm technology (leading to a chip @ 2.5 Volt, 600-pin EBGA, 4 cm x 4 cm) and about 25 mm^2 in 0.18 μm technology available next year (leading to a chip @ 1.8 Volt, 676-pin EBGA, 2.7 cm x 2.7 cm). However, the later technology dissipates ten times less power with respect to the 0.25 μm .

The main characteristics of the 3D-Flow³ system architectures based on a single 3D-Flow component are the following:

2.1 System level

Objective:

Oriented toward data acquisition, data movement, pattern recognition, data coding and reduction.

Design considerations:

- Quick and flexible acquisition and exchange of data, bi-directional with North, East, West, and South neighbors, unidirectional from Top-to-Bottom.
- Small on-chip area for program memory in favor of multiple processors per chip and multiple execution units per processor, data-driven components (FIFOs, buffers), and internal data memory. (Most algorithms that this system aims to solve are short and highly repetitive, thus requiring little program memory.)
- Balance of data processing and data movement with very few external components.

- Programmability and flexibility provided by download of different algorithms into a program RAM memory.
- Strong emphasis on modularity and scalability, permitting solutions for many different types and sizes of applications using regular connections and repeated components.

2.2 System architecture

The goal of this parallel-processing architecture is to acquire multiple data in parallel (up to the maximum clock speed allowed by the latest technology) and to process them rapidly, accomplishing digital filtering on the input data, pattern recognition, data moving, and data formatting.

The system is suitable for "particle identification" applications in HEP (calorimeter data filtering, processing and data reduction, track finding and rejection).

The compactness of the 3D-Flow parallel-processing system in concert with the processor architecture (its I/O structure in particular) allows processor interconnections to be mapped into the geometry of sensors (such as detectors in HEP) without large interconnection signal delay, enabling real-time pattern recognition. This work originated by understanding the requirements of the first levels of triggers for different experiments, past, present, and future. A detailed study of each led to the definition of system, processor, and assembly architecture suitable to address their recognized common features. To maintain scalability and simplify the connectivity, a three-dimensional model was chosen, with one dimension essentially reserved for the unidirectional time axis and the other two as bi-directional spatial axes (Figure 1).

The system architecture consists of several processors arranged in two-orthogonal axes (called layers; see Figure 2), assembled one adjacent to another to make a system (called a stack; see Figure 3). The first layer is connected to the input sensors, while the last layer provides the results processed by all layers in the stack.

Data and results flow through the stack from the sensors to the last layer. This model implies that applications are mapped onto conceptual two-dimensional grids normal to the time axis. The extensions of these grids depend upon the amount of flow and processing at each point in the acquisition and reduction procedure as well as on the dimensionality of the set of sensors mapped into the processor layers.

Four counters at each processor arbitrate the position of the bypass/in-out switches (Top-to-Bottom ports. See Figure 25) responsible for the proper routing of data. Higher-dimensional models were considered too costly and complex for practical scalable systems, mainly due to interconnection difficulties.

2.3 Processor architecture

The 3D-Flow processor is a programmable, data-stream pipelined device that allows fast data movements in six directions with digital signal-processing capability. Its cell input/output is shown in Figure 1.

Architecture (processor and system)

- ◆ Modularity (the same "logical unit," the 3D-Flow PE, replicated several times -- in a chip, on a board, or on a system)
- ◆ Powerful I/O

"logical unit" = 3D-Flow Processing Element (PE)

Figure 1. The 3D-Flow Processing Element (PE) or "logical unit."

The 3D-Flow can operate on a data-driven, or synchronous mode. In data-driven mode, program execution is controlled by the presence of the data at five ports (North, East, West, South, and Top) according to the instructions being executed. A clock synchronizes the operation of the cells. With the same hardware one can build low-cost, programmable first levels of triggers for a small and low-event-rate detector, or high-performance, programmable higher levels of triggers for a large detector. The multi-layer architecture and automatic by-pass feature from Top-to-Bottom ports allow for event input to be sustained at the processor clock rate, even if the actual algorithm execution requires many clock cycles, as described below.

The 3D-Flow processor is essentially a Very Long Instruction Word (VLIW) processor. Its 128-bits-wide instruction word allows for concurrent operation of the processor's internal units: Arithmetic Logic Units (ALUs), Look Up Table memories, I/O busses, Multiply Accumulate and Divide unit (MAC/DIV), comparator units, a register file, an interface to the Universal

Asynchronous Receiver and Transmitter (UART) used to preload programs and to debug and monitor during execution, and a program storage memory.

The high-performance I/O capability is built around four bi-directional ports (North, East, South, and West) and two mono-directional ports (Top and Bottom). All of the ports can be accessed simultaneously within the same clock cycle. N, E, W, and S ports are used to exchange data between processors associated with neighboring detector elements within the same layer. The Top port receives input data and the Bottom port transmits results of calculations to successive layers.

A built-in pipelining capability (which extends the pipeline capability of the system) is implemented using a "bypass mode." In bypass mode, a processor will ignore data at its Top port and automatically transmit it to the Top port of the processor in the next layer. This feature thus provides an automatic procedure to route the incoming events to the correct layer. Several 3D-Flow processing elements, shown in Figure 1, can be assembled to build a parallel processing system, as shown in Figure 2.

Figure 2. One layer (or stage) of 3D-Flow parallel processing.

2.4 Introducing the third dimension in the system

In applications where the processor algorithm execution time is greater than the time interval between two data inputs, one layer of 3D-Flow processor is not sufficient.

The problem can be solved by introducing the third dimension in the 3D-Flow parallel-processing system, as shown in Figure 3.

In the pipelined 3D-Flow parallel-processing architecture, each processor executes an algorithm on a set of data from beginning to end (e.g., the event in HEP experiments, or the picture in graphics applications).

Data distribution of the information sent by the external data sources as well as the flow of results to the output are controlled by a sequence of instructions residing in the program memory of each processor.

Each 3D-Flow processor in the parallel-processing system can analyze its own set of data (a portion of an event or a portion of a picture), or it can forward its input to the next layer of processors without disturbing the internal execution of the algorithm on its set of data (and on its neighboring processors at North, East, West, and South that are analyzing a different portion of the same event or picture. The portion of event or picture is called "Frame A1, Frame A2, etc.," in Figure 3.).

The programming of each 3D-Flow processor determines how processor resources (data moving and computing) are divided between the two tasks or how they are executed concurrently.

A schematic view of the system is presented in Figure 3, where the input data from the external sensing device are connected to the first layer of the 3D-Flow processor array.

The main functions that can be accomplished by the 3D-Flow parallel-processing system are:

- Operation of digital filtering on the incoming data related to a single channel;
- Operation of pattern recognition to identify events of interest; and
- Operations of data tagging, counting, adding, and moving data between processor cells to gather information from an area of processors into a single cell, thereby reducing the number of output lines to the next electronic stage.

In calorimeter trigger applications, the 3D-Flow parallel-processing system can identify patterns of energy deposition characteristic of different particles type, as defined by more or less complex algorithms, so reducing the input data rate to only a subset of candidates.

In real-time tracking applications, the system can perform pattern recognition, calculate track slopes, and intercepts total and transverse momenta as well.

Figure 3. General scheme of the 3D-Flow pipeline parallel-processing architecture.

2.5 The 3D-Flow architecture optimized features for the first levels of triggers.

The following list of Table 1 shows the most important features of the 3D-Flow that make it very efficient to solve algorithms of the first level of triggers in High Energy Physics.

Table 1. The 3D-Flow architecture optimized features for first-level trigger algorithms

A Typical Level-0 Algorithm Requires:	The 3D-Flow Architecture Offers:
100% of the time during algorithm execution it is required to input data and output results.	Top and Bottom ports are: multiplexed only 2:1, propagating, by means of the by-pass switches, either input data or output results at each cycle. Outputs are required to drive only up to 6 cm.
Only 10% of the time of the algorithm execution it is required to exchange data with neighbors.	North, East, West, and South ports are: multiplexed 10:1, do not require many cables, have very low power consumption with LVDS (Low Voltage Differential Signaling) I/O requiring to drive only up to 1.5 meters.
Operation of comparing with different thresholds, finding local maximum.	A special unit with 32 registers/comparators can compare 4 values, find their range, or find the local maximum, or the greater between pairs, all in one cycle.
Short programs.	128 words of program memory.
Lookup table to convert ADC values.	Four data memories, each for lookup tables of 256 locations of 16-bit, or for buffering.
Arithmetic and Logic operations (multiplying by calibration constants, adding to calculate cluster energies).	All Arithmetic, Logical and data move operations are provided by parallel units executing up to 26 operations per cycle. (Including Multiply-Accumulate and Divide at variable precision).

3 A SINGLE TYPE OF COMPONENT FOR SEVERAL ALGORITHMS

The fundamental concept of the 3D-Flow system, based upon the replication of a single, interconnecting component, can be maintained in spite of the rapid advances in Integrated Circuits technology. For this purpose, the 3D-Flow "System on a Chip" (SOC) complete circuit description is provided in generic VHDL in IP⁴ form, so that it can be implemented at any time using the technology of the day. This in turn will allow to achieve, at any moment in time, the best performance in terms of power dissipation, size, speed, and, consequently, cost.

SOCs (System On a Chip) utilizing IPs (Intellectual Property) Virtual-Components (VC) are redefining the world of electronics. A catalog⁵, of 200 VC IP's was introduced at DAC '98.

3.1 The evolution of IC Design

The evolution of ICs is illustrated in Figure 4. All current indications and projections confirm that the evolution will continue to increase rapidly in the years to come. Furthermore, the traditional way of designing systems will change: the current productivity of about 100 gates per day (*EE Times*, Oct. '98) will need to improve substantially, in order to resist competition. The evolution of ICs reflects what is illustrated in Figure 4. Many statements in this regard have been reported by specialized

magazines. Using today's methodology, a 12-million-gate ASIC would require 500 person-years to develop, at a cost in excess of \$75M. Companies will not be able to afford this cost, unless one develops IP blocks in order to build System-On-a-Chip. Analog design retains its investment for several years, while digital design becomes outdated in about one year.

The 3D-Flow System digital design based on a single replicated circuit

- allows for implementation of the user's conceptual algorithm, at the gate circuit level, into the fastest High-Speed, Real-Time programmable system.
- retains its value because of its powerful 'Design Real-Time' tools that allow the user to quickly design, verify, and implement a System-On-a-Chip (SOC) based on a single replicated circuit (the 3D-Flow processing element [PE] in IP form [C++, VHDL, and netlist]), that can be targeted to the latest technology at any time.

	Current Industry		Forecast	
	1997	1998	1999	2001
Logic Gates	300K	1M	5M	12M
Technology	0.5 μm	0.35 μm	0.25 μm	0.18 μm
Speed	1.4 x '96	1.4 x '97	1.4 x '98	1.4 x '2000
Cost/100K gates	\$25	\$5	~ \$2	???
3D-Flow SOC (100K gates/PE)	450K (4-PEs)	1.7M (16-PEs)	6.5M (64-PEs)	13M (128-PEs)

Figure 4. The evolution of IC design.

3.2 Technology-independent 3D-Flow ASIC

Figure 5 shows the main characteristics of the 3D-Flow chip that is technology-independent. As the technological performance increases⁶, the multiplexing of the I/O can also increase. For example the (8+2):1 of the LVDS (Low Voltage Differential Signaling) serial links can increase to 16:1 or (16+2):1 when the LVDS serial link speed reaches 1.2 Gbps or higher. (16 PEs, for example, requires a silicon area of about 50 mm² in 0.25 μm technology, leading to a chip @ 2.5 Volt, 600-pin EPGA, 4 cm x 4 cm, while it requires about 25 mm² in 0.18 μm technology that will be available next year, and it will lead to a chip @ 1.8 Volt, 676-pin EPGA, 2.7 cm x 2.7 cm. Please see the Web site of LSI-Logic⁶ as an example of technology currently available).

Figure 5. Technology-independent 3D-Flow ASIC

4 FIRST-LEVEL TRIGGER ALGORITHMS

Typical first-level trigger algorithms at LHC⁷ experiments need to sustain the input data rate at 40 MHz with zero dead time, providing a yes/no global level-0 (or level-1) trigger output at the same rate; need to exchange --for about 10% of the duration of the algorithm-- data with neighboring elements; need to find clusters with operations of multiply/accumulate; and need to have a special unit that should be a combination of registers/comparators capable of executing in one cycle operations such as ranging, local maximum, and comparing different values to different thresholds. While short, the first-level trigger algorithms need a good balance between input/output operation and several other operations of moving data, data correlation, arithmetic, and logical operation performed by several units in parallel. Typical operations also include converting ADC values into energies or a more expanded 16-bit nonlinear function that is quickly accomplished by lookup tables. The internal units of the 3D-Flow have all these capabilities, including powerful I/O.

The desired performance, programmability, modularity, and flexibility of the 3D-Flow⁸ are represented schematically in Figure 6. With a 3D-Flow processor running an 80 MHz clock speed, it has been shown that the calorimeter trigger requirements can be met by a 3D-Flow system of 10 layers, each layer comprising about 6000 Processing Elements (PEs), one element per ECAL block (sometimes referred to as "trigger tower," that is corresponding to all signals from ECAL, HCAL, PreShower and Muon detectors contained in a specific view angle from the interaction point). Each PE executes the user's defined trigger algorithm on the information received from the detector, at the bunch crossing 40 MHz rate (requiring a time interval ranging from 100 ns to 300 ns, depending on the complexity of the algorithm.). The ten-layer stack is then followed by a data collection "pyramid," where the information from any trigger tower (3D-Flow input channel) where an event of interest was found is routed to a single exit point. The data routing that provides channel reduction is accomplished via the NEWS ports within a time of the order of a microsecond, depending on the size and number of channels in the system.

The details of algorithms devised to perform the selection process, schematically shown in Figure 6, and their realization within the 3D-Flow systems have been described elsewhere^{2, 9, 10, 11}. The present document provides a detailed description of all the components, and their layout, required to build the 3D-Flow system appropriate for the implementation of the calorimeter trigger (the muon trigger implementation details cannot yet be fully defined, since the actual detector configuration is still under discussion, and it will be the subject of a future paper).

While utilizing existing technology in each individual step, the resulting system is very compact in the total number of crates (e.g., 6 crates for the calorimeter trigger) and is less costly than other proposed solutions^{12, 13}. And this, while conserving the intrinsic properties of full programmability and ease of expansion.

The full simulation of the algorithm can be verified from the system level to each component gate level by comparing the bit-vectors generated by system simulation tools and the tools provided by the Electronic Design Automation (EDA).

Figure 6. Overview of the use of the 3D-Flow System in particle identification in HEP.

5 LHCb LEVEL-0 TRIGGER OVERVIEW

5.1 Physical Layout

The preferred layout for the LHCb¹¹ level-0 trigger is to have all decisions made in electronics racks located on the “balcony” at some 40 meters from the detector. In this configuration, the only link from the control room, located about 70 meters from the detector, to the level-0 trigger electronics is given by the trigger monitor, operating through slow control on RS-422 links. Figure 7 shows the path of the signals from the different sub-detectors to the electronics, and the corresponding time delays (the numbers identifying each step in Figure 7 correspond to the same numbers in Figure 8).

An alternative scheme would call for locating all the level-0 trigger electronics in the control room. This scheme would have the advantage of easier access for maintenance; its disadvantage is that it would be necessary to run longer cables which will require longer latency. What follows is the first option having the level-0 trigger electronics on the balcony.

Another choice has to be made on whether to convert the signals from analog-to-digital on-detector or off-detector. The selection of one scheme instead of another will consequently require some changes in the electronics. The current preferred solution among the LHCb collaboration it seems to be the one which foresees a mixture of analog and digital signals to be received from the detector; however, for maximum flexibility, a 3D-Flow level-0 trigger system that foresees receiving signals from the detector solely in digital form is also reported (see Section 6.2 --3D-Flow digital processing board-- Option 2). LHCb current approach is more similar to that used in the Atlas¹⁴ experiment, in which the analog signals are transported for about 60 meters and are converted to digital in a low-radiation area. On the contrary, the first-level trigger of the CMS experiment receives all digital information. The conversion is being made on-detector by means of the radiation-resistant QIE analog-to-digital converter (Q for charge, I for integrating, and E for range encoding), which was developed at Fermi National Laboratory.

After the particles have traveled from the interaction point to the calorimeter, and the signal is formed by the photomultipliers (steps 1 and 2), a minimum of analog electronic circuit with line driver will be installed close to the photomultiplier. The signal is then transported through a coaxial 17 position ribbon cable (part number AMP 1-226733-4) to the 3D-Flow mixed-signals processing board (shown in figure 12).

These analog signals are foreseen to be converted to 12-bit digital form with standard components such as Analog Device AD 9042¹⁵. For the analog signals available at the PreShower sub-detector it will be desirable, because of lower cost, to use a shorter cable set from the different sensors to a location where the signals can be grouped together in sets of 20 bits or more. The above analog signals, as well as the ones from the muon stations, are foreseen to be converted to only one-bit digital value. Once the digital signals have been grouped, they can be sent in digital form on standard copper cables (e.g. equalized cables AMP 636000-1), through one of the available serializers at 1.2 Gbps. (Serializers at 2.4 Gbps are also available; however, they are limited to 10 meters in copper or at longer distances in optical fiber and are more expensive). In case the radiation is too high where the transmitter (or serializer) has to be installed, radhard components¹⁶ should be considered.

In the event the preferred solution by LHCb would be that of keeping the analog cables shorter than 10 meters (see Figure 7), since in LHCb there is a very low radiation (fewx100 rad/year @ 10m), than the entire 6 racks of front-end and level-0 calorimeter trigger can be moved to the top of the detector.

Figure 7. LHCb Level-0 Trigger - Physical Layout.

5.2 Logical Layout

The scheme of the entire Level-0 trigger system for the event selection ("trigger") for the LHCb¹¹ High Energy Physics experiment is summarized in Figure 8.

Figure 8 shows the logical function performed by the different signals and electronics previously shown in Figure 7 (see also the timing information indicated by the number inside the circle in Figure 7). It is divided into three sections. The section at the left shows the electronics and signals on the detector. The center section shows the electronics and signals in the racks located off-detector (where all decision electronics for the level-0 trigger are located). The section on the right shows the cables/signals carrying the information to the DAQ and higher-level triggering system that are received at the control room. In this scheme, only the monitoring electronics of the level-0 trigger is located in the control room.

In regard to the analog-to-digital front-end electronics for the preshower detector, the design assumes to receive 1-bit digital information from each preshower detector channel. Since the baseline solution described in the TP (see TP Chapter 10.2.3) foresees this electronics accommodating close to the multi-anode PMTs, a printed circuit board with the size of 9 cm x 9 cm is estimated to be used for each 16-channels PMT, on each of the 375 PMT installed on the preshower detector. The design of this section of electronics is described briefly in Appendix B.

Figure 8. LHCb Level-0 Trigger - Logical Layout.

The LHCb detector, consisting of several sub-components (ECAL, HCAL, PreShower, Muon, VDET, TRACK and RICH) monitors the collisions among proton bunches occurring at a rate of 40 MHz (corresponding to the 25 nsec bunch crossing rate). At every crossing, the whole information from the detector (data-path) is collected (indicated in the figure by the number 4), digitized (indicated by the number 6), synchronized and temporarily stored (indicated by 7) into digital pipelines (conceptually similar to 128 deep, 40 MHz shift registers), while the Trigger Electronics (indicated by 8 and 9), by examining a subset of the whole event data (trigger path), decides (indicated by 10) whether the event should be kept for further examination or discarded. In the LHCb design, the input rate of 40 Tbytes per sec¹¹ (see top of the figure), needs to be reduced, in the first level of triggering, to 1 Tbytes/sec, i.e. a 1 MHz rate of accepted events. The selection is performed by two trigger systems (indicated by 8) running in parallel, the Calorimeter Trigger, utilizing mainly the information from the ElectroMagnetic and Hadronic Calorimeters (ECAL and HCAL) to recognize high transverse momentum electrons, hadrons and photons; and the Muon Trigger, utilizing the information from five planes of muon detectors to recognize high transverse momentum muons.

The resulting global level-0 trigger accept signal (indicated by 10 in the figure) enables the data in the data-path to be stored first into a derandomizing FIFO and later to be sent through optical fiber links to the higher-level triggers and to the data acquisition (see in Figure 8 the signal Global L0 distributed to all front-end 128 bunch crossing “bx” pipeline buffers). Real-time monitoring systems (L0 CAL monitor and L0 MUON monitor) supervise and diagnose the programmable level-0 trigger from the distant control room.

5.3 Electronic Racks (Functions/Locations)

Figure 9 shows the estimate of the type of electronics that will be needed on-detector for Level-0 trigger. Figure 10 shows the number and functionality of the crates and racks located off-detector that will be required to accommodate the level-0 electronics. A fully programmable calorimeter Level-0 trigger implemented with the 3D-Flow requires 6 crates (9U). This is to be compared with the, less flexible 2x2 trigger implementation option¹², requiring 20 crates (9U), or with a third, HERA-B like solution¹³ requiring 40 crates (9U). Figure 11 shows the monitoring system for the 3D-Flow calorimeter trigger. This, together with any other monitoring of the level-0 muon trigger and of the global level-0 decision unit, should be accommodated in the control room

Figure 9. On-Detector Electronics for Level-0 trigger. Figure 10. Off-Detector Electronics for level-0 trigger.

Figure 11. Electronics in the Control Room for the Calorimeter Level-0 Trigger Monitoring.

6 A SINGLE TYPE OF BOARD FOR SEVERAL ALGORITHMS

The modularity, flexibility, programmability, and scalability of the 3D-Flow system is kept all the way from the component to the crate(s). This is valid also for the type of board used in the system. Only a single type of board is needed in a 3D-Flow system of any size. This board can change for each application from mixed signals analog and digital to a purely digital board, depending on the nature of the input signals received from the sensors.

Following are described a mixed-signal 3D-Flow processing board (option 1), and a purely digital processing board (option 2). The only difference among the two boards is the front-end electronics. In one case, there are preamplifiers and analog-to-digital converters, in the second case there are high-speed optical fiber links.

The board design presented here, based upon an 80 MHz processor, accommodates 64 trigger towers and 10 processing layers. With a 16-bit wide word processor, such a board can sustain an input bandwidth of 10.24 Gbyte/s (80 MHz x 2 bytes x 64) and process the received information on each of the 64 channels with zero dead time and a real-time algorithm of the complexity up to 20 steps. (It should be considered that up to 26 different operations can be executed at each step, including efficient operations of data exchange with neighboring channels).

With today's technology, it is not a problem to feed a 9U x 5HP board from the front panel with digital information at 10.24 Gbyte/s, e.g., the information could be received by the board using currently available deserializer/receiver links from several vendors at 1.2 GHz. Possible choices for such deserializer devices include Hewlett Packard HDMP-1024¹⁷, HDMP-1034 @ 1.2 Gbps, AMCC quad serial backplane serializer/deserializer device with single and dual I/O S2064¹⁸/S2065 @ 1.25 GHz, and from VITESSE¹⁹). Alternatively, using the deserializer from AMCC-S3044²⁰. @ 2.4 GHz (this device requires a minimum network interface processor that can be implemented in FPGA), Lucent Technologies²¹ TC16-Type 2.5 Gb/s optical transmitter/receiver with 16 channels 155 Mb/s serializer/deserializer, or using links soon to become available for the short range at 10 GHz that are already available for the long range in telecommunications (see Lucent Technologies²² and . or Nortel²³) may also solve this problem.

Should the transmission distance exceed 30 meters @ 1.2 GHz (only 10 meters can be achieved with acceptable Bit Error Rate –BER– for transmission over copper @ 2.4 Gbps), then the more expensive optical fiber receivers should be coupled to the components mentioned above. As one can notice from the type of components listed above, non all vendors provide devices with functions of deserializing/receiving/demultiplexing, separated from the functions of serializing/transmitting/multiplexing. The same situation occur when one of the above components has to be coupled with a fiber optic receiver. Also in this case we may find vendors that offer both functions (optical fibers receiver/transmitter) in a single component at a lower cost in some cases then the price of a component with a single functions. Some example of matching the previous deserializer/receiver with optical fiber receivers (or receiver/transmitter) are: Hewlett Packard HDMP-1024 with the optical transreceiver HFCT-53D5²⁴, AMCC-S3044²⁰ with the fiber optic receiver SDT8408-R²⁵, and Lucent Technologies deserializer TRCV012G5 with the optical fiber transreceiver Netlight1417JA. Connectors carrying several fibers are provided by many vendors (e.g. from Methode²⁶).

The above deserializing/receiving components have matching components that can be supplied by the same vendors, which have the function of serializing/transmitting/multiplexing and optical fiber transmitter that are needed for transmission of the input data from the front-end electronics, or for the transmission of the output results from the 3D-Flow digital (or mixed-signal) processing board to the data acquisition system and higher-level triggers. A few examples are: deserializer HDMP-1034, matched with serializer HDMP-1032, deserializer HDMP-1024, matched with serializer HDMP-1022, deserializer AMCC-S3044²⁰, coupled with the fiber optic receiver SDT8408-R²⁵, matched with the serializer AMCC-S3043²⁷, coupled with the fiber optic transmitter SDT8028-T²⁵ (this device requires a minimum network interface processor that can be implemented in FPGA).

In the mixed signal application (option 1), only 80 analog signals (64 ECAL + 16 HCAL, since each HCAL is equivalent to an area of 4 ECALs), converted to digital with 12-bit resolution in addition to 192 bits (1 PreShower + 2 Pads from muon station 1 x 64) are received by each board every 25 ns. This is not saturating the bandwidth of the 32-bit x 64 channels = 2048 bits every 25 ns bunch crossing that the 3D-Flow system could sustain.

However, the front-end electronic FPGA chips on the same board described in detail in Section 6.3 (see Figure 16) increase the input bandwidth to the 3D-Flow system by formatting and generating the input trigger word to be sent to each of the 64 channels. More precisely, the FPGA trigger word formatter (see Section 6.3.5 and Figure 20) reduces the ECAL information from 12-bit to 8-bit, and increases by duplicating information to different channels (e.g. sending the same 8-bit HCAL information to each of the 4 subtended ECAL blocks, and sending the same 2-bit Pads to 4 neighboring blocks), in order to save some bit-manipulation instructions to the 3D-Flow processors.

6.1 3D-Flow mixed-signal processing board (Option 1)

Features of the 3D-Flow mixed-signal processing board built-in standard 9U x 5HP x 340 mm dimensions (see Figures 12, and 13):

- Converts 80 analog inputs (ADC 12-bit resolution), and produce 4 copies of each HCAL digitized value;
- synchronizes 1152 inputs (16 x 4 x 12 bits ECAL, 16 x 12 bits HCAL, 16 x 4 x 1 PreSh, 16 x 4 x 2 Pads) every 25 ns;
- saves 1152 raw-data every 25 ns in a 128x1152 pipeline-stage digital buffer;

Figure 12. Mixed-signal processing board (front view).

Figure 13. Mixed-signal processing board (rear view).

- processes data received from 64 trigger towers (or data received at a continuous input data stream of 10 Gbyte/s) and sends to the global level-0 trigger the information (tower ID, bunch crossing ID, and energy) of the clusters that passed the level-0 trigger algorithm;
- receives the global level-0 trigger accepts and sends out the raw data of the corresponding accepted events;
- derandomize accepted raw data into a FIFO;
- all 3D-Flow inter-chip Bottom-to-Top port connection are within the board (data are multiplexed 2:1, PCB traces are shorter than 6 cm); all 3D-Flow inter-chip North, East, West, and South port connections between boards and crates are multiplexed (8+2):1 and are shorter than 1.5 meters;
- communicate with the host monitoring/control system via 16 RS-422 links to download user's algorithms into the processors and upload performance data (the status of all processors during 8 consecutive cycles) for monitoring purposes;
- communicates with the host monitoring/control system to download the FPGA programming, to adjust signal synchronization, pipeline stages, FIFO buffer, and trigger word formatter;
- communicates through 160 LVDS²⁸ links to North, East, West, and South neighboring boards.

What follows is a description of the board with its component list and assembly information.

The 3D-Flow mixed-signal processing board has on the front panel:

- three connectors for receiving digital raw data from the PreShower and muon M1 detectors through six copper twisted pair links at 1.2 Gbps. The receiver from Hewlett Packard HDMP-1034 (or HDMP-1024, dimension: 23 mm x 17 mm) could be used;
- five 17-conductor coaxial ribbon cables²⁹ for analog input (see Figure 12) from the electromagnetic, hadronic calorimeter, and from the control signal (reset, control A1, clear, clock, and global level-0 accept);
- 17 bidirectional RS-422 links for monitoring the on-board 3D-Flow system and loading different circuits into the FPGA;
- one RJ45 connector carrying four high-speed LVDS output signals to the global level-0 trigger decision unit;
- one optical fiber carrying out raw data relative to the event accepted by the level-0 trigger decision unit (e.g., Hewlett Packard transmitter at 1.2 Gbps HDMP-1022 (dimension: 23 mm x 17 mm) coupled with the fiber optic transceiver HFBR-53D5 (dimension: 39.6 mm x 25.4 mm).

On the rear of the board are assembled alternately four 200-pin AMP-9-352153-2³⁰ connectors with three 176-pin AMP-9-352155-2³⁰ connectors. The latter connectors have a key for mechanical alignment to facilitate board insertion. Of these, 1280 pins carry LVDS signals to neighboring 3D-Flow chips residing off-board in the North, East, West, and South direction; 48 pins are used for power and ground.

Starting from the left of the board, we have 80 analog preamplifiers P (half of the components are on the rear of the board as shown in Figure 13), 80 analog-to-digital converters A (e.g., Analog Device¹⁵ AD9042 converting each analog input channel to 12-bit at 40 MHz). The converted data are then combined with the other digital information received from the other detectors (PreShower and muon stations) into 16 FPGAs (4 channels fit into an ORCA³¹ 256-pin BGA OR3T30) for the purpose of synchronization, pipelining, derandomizing, and trigger word formatting.

Formatted data are then sent to the processor stack (see Figures 23, and 24), to be picked by the first available layer, according to the setting of the bypass switches (see Figure 25), where the trigger algorithm is then executed.

At the bottom of the stack (see Figure 24), the first layer of the pyramid checks whether a valid particle (electron, hadron, or photon) was found.

The entire board (64 channels) is designed to send to the global trigger decision unit an average of 40-bits of information of clusters validated by the trigger algorithm (tower ID, time stamp, and energies) at each bunch crossing, through four LVDS⁶ links at 400 Mbps on the J1 connector.

If the detector has higher occupancy so that any region of 64 channels could be expected to transmit to the global level-0 decision unit more than 40-bit per bunch crossing, then it would be sufficient to select a higher-speed link (e.g., 1.2 Gbps). If the occupancy is still higher, the number of output links to the global trigger decision unit can be increased to the required level.

If, on the other hand, 40-bit per bunch crossing per group of 64 were sufficient, then it would be simpler not to use the National Semiconductor serializer DS92LV1021, but rather have the North, East, West, or South ports of the 3D-Flow chip driver send the information directly to the global level-0 decision unit. In the present board, these serializer chips from National Semiconductor have been considered in order to make a conservative choice in terms of driving capabilities to three meters, while the 3D-Flow chip is required to drive only 1.5 meters on the LVDS I/O.

The board consists of surface-mounted devices assembled on both sides, with some free space not covered by components.

6.2 3D-Flow digital processing board (Option 2)

The digital processing board carries on the motherboard 16 high-speed receiver links at 2.4 Gbps (e.g., the set from AMCC-S3044²⁰ and the SDT8408-R²³ optical fiber receiver, which contains 16 sockets for mezzanine boards with the same set of

components, or with the transmitter set AMCC-S3043²⁴ and the SDT8028-T²³ (these devices require a minimum network interface processor that can be implemented in FPGA).

The user can install as many mezzanines as required (up to 16) for the application in order to optimize the cost. For example, one could use 16 x receivers set on the mother board to sustain a rate of data input to the board of 5 Gbyte/s, and install 16 x transmitter mezzanine boards that provide 5 Gbyte/s output. Another application may need instead to install 15 x receiver mezzanine boards that together with the 16 on-board receivers provide 9.92 Gbyte/s input bandwidth, and only one transmitter mezzanine board for 320 Mbyte/s output data. This configuration will satisfy many high-energy physics experiments where the real-time trigger algorithm achieves a substantial reduction.

As another example, the CMS calorimeter level-1 trigger³² (currently implemented in 19 crates (9U) using a different approach, while it will require only 5 crates (9U) if the 3D-Flow approach would be used), requires to receive only 18-bits from each trigger tower (electromagnetic, hadronic, fine grain, and characterization bit). Thus only 5 additional mezzanine fibers and receiver modules must be installed. One board can process 64 trigger towers and send to the global level-1 trigger decision unit the particles ID, time stamp, and energy information of the particles validated locally by the trigger algorithm. Subsequently, it can provide the raw-data of the particles validated by the global level-1 trigger. This scheme has the advantage of flexibility: if the experiment later requires not only changing the level-0 (or level-1) trigger algorithm, but also increasing the number of bits (information) used in the level-0 (or level-1) trigger algorithm, this can be done without redesigning the hardware. In the case of the CMS calorimeter trigger algorithm, by using the digital processing board of the 3D-Flow approach, the user can in the future increase the number of bits from each trigger tower from 18 to 31 before being required to redesign the hardware.

Features of the 3D-Flow digital processing board, built-in standard 9U x 5HP x 340 mm size (see Figures 14 and 15):

- Inputs 1024 digital inputs and outputs 1024 digital outputs every 25 ns, or any combination of I/O having a total of 2048 I/Os and a minimum of 1024 inputs every 25 ns;
- synchronizes up to 2048 inputs every 25 ns from different detectors (electromagnetic, hadronic, PreShower, and M1)
- saves up to 2048 raw-data every 25 ns in a 128 x 2048 pipeline-stage digital buffer;
- processes data received from 64 trigger towers (or data received at a continuous input data stream of 9.92 Gbyte/s) and sends to the global level-0 (or level-1) trigger the information (trigger tower ID, time-stamp, and energies) of particles that passed the level-0 trigger algorithm;
- receives the global level-0 trigger accepts and sends out the raw data of the corresponding accepted events;
- derandomizes accepted raw data into FIFO;
- all 3D-Flow inter-chip Bottom-to-Top port connection are within the board (data are multiplexed 2:1, PCB traces are shorter than 6 cm); all 3D-Flow inter-chip North, East, West, and South ports connections between boards and crates are multiplexed (8+2):1 and are shorter than 1.5 meters;
- communicate with the host monitoring/control system via 16 RS-422 links to download user's algorithms into the processors and upload performance data (the status of all processors during 8 consecutive cycles) for monitoring purposes
- communicates with the host monitoring/control system to download the FPGAs programming, to adjust signals synchronization, pipeline stages, FIFO buffer and trigger word formatter;
- communicates through 160 LVDS links to North, East, West, and South neighboring boards.

What follows is a description of the board with its component list and assembly information.

The 3D-Flow digital processing board has on the front panel:

- 16 optical fibers of receivers, each at 2.4 Gbps installed on the mother board and 16 optional optical fibers (transmitter or receiver) installed on the mezzanine boards. (receiver SDT8408-R²⁵, dimension: 15.24 mm x 36.4 mm, with the deserializer AMCC-S3044²⁰, dimension: 17 mm x 17 mm, both at 2.5 Gbps and transmitter SDT8028-T²⁵, dimension: 15.24 mm x 36.4 mm, with the serializer AMCC-S3043²⁷, 17 mm x 17 mm. (These devices require a minimum network interface processor that can be implemented in FPGA).
- 17 bidirectional RS-422 links for monitoring the on-board 3D-Flow system and loading different circuits into the FPGAs;
- one RJ45 connector carrying four high-speed LVDS output signals to the global level-0 trigger decision unit.

On the rear of the board are assembled alternately four 200-pin AMP-9-352153-2³⁰ connectors with three 176-pin AMP-9-352155-2³⁰ connectors. The latter connectors have a key for mechanical alignment to facilitate board insertion. Of these, 1280 pins carry LVDS signals to neighboring 3D-Flow chips residing off-board in the North, East, West, and South direction; 48 pins are used for power and ground.

The mezzanine board is built with four PAL16P8 (high-speed, 5n pin-to-pin, or fast PLD) for the purpose of demultiplexing the 16-bit at 155 MHz provided by the AMCC-S3044²⁰ into 32-bit at 77.5 MHz. This additional PALs are needed at least until next year when the FPGAs at 160 MHz will become available and the signals from the AMCC chip could be sent directly to the

Figure 14. Digital processing board (front view).

Figure 15. Digital processing board (rear view).

FPGA chip. The reason for installing the 4 PAL on the mezzanine board is to lower the high frequency through connectors (77.5 MHz instead of 155 MHz). This will allow to use lower cost connectors to be used.

The digital data (from the electromagnetic and hadronic calorimeter, preshower, and muon) are sent into 16 FPGAs (4 channels fit into an ORCA³¹ 256-pin BGA OR3T30) for the purpose of synchronization, pipelining, derandomizing, and trigger word formatting.

Formatted data are then sent to the processor stack (see Figures 23, and 24), to be picked by the first available layer, according to the setting of the bypass switches (see Figure 25), where the trigger algorithm is then executed. At the bottom of the stack, the first layer of the pyramid checks whether a valid particle (electron, hadron, or photon) was found.

The output of the particle found locally by the trigger algorithm (tower ID, time stamp, and energies) are sent out to the global level-0 decision unit through an RJ45 connector carrying four LVDS⁶ links at 400 Mbps. The same consideration applies to this board that was made for the mixed-signal processing board described in Section 6.1 on the number of bits sent to the global level-0 decision unit that is related to the detector occupancy.

The raw data of the events validated by the global level-0 trigger are sent out to the higher-level trigger system and DAQ, through the installed transmitter mezzanine boards. The necessary number of transmitter mezzanine boards should be installed in order to sustain the volume of raw data information to be sent out.

Boards contain surface-mounted devices assembled on both sides, with some free space not covered by components.

6.3 Front-end Signal Synchronization/pipelining/derandomizing/trigger word formatter

The complete detailed study for the overall level-0 front-end electronics has been performed and described elsewhere³³. Detailed circuits that can be downloaded in the ORCA OR3T30 FPGA are provided in the above referenced document, together with testbenches for easy verification of the correlation between signals and their timing performance.

For the mixed-signal processing board, after the task of amplification and conversion of analog signals to digital by means of an ADC such as Analog Devices AD9042¹⁵ converting to 12-bit at 40 MHz, all digital information is sent to 16 FPGAs. Each FPGA can implement all functions described below for four channels out of 64 channels in a board. The study has been made referring to the component from Lucent Technologies ORCA OR3T55³¹ with 256-pin BGA, dimension 27 mm x 27 mm.

The digital information relative to four trigger towers is sent to the input of one FPGA. If a PAD from the muon station is used by more than one trigger tower, it will be sent to all the appropriate FPGA units.

All data are strobed into a register inside the FPGA at the same time; however, the present design allow for the possibility that data from different detectors (e.g., muon Pad vs. ECAL) be out of phase by one or two bunch crossings.

Next, a delay from 0 to 2 clock counts at each bit received at the input of the FPGA needs to be inserted. This function, called “variable delay,” is shown in Figures 16, 18, and 19.

For each channel we have, then, 12-bit information from the electromagnetic calorimeter, 12-bit information from the hadronic calorimeter, 1-bit information from the preshower, and 2-bit information from the muon pad chamber, for a total of 27-bits per input channel.

The above 27-bits input channels need to be stored into a level-0 pipeline buffer of 128 clocks (or bunch crossings) while the trigger electronics verifies whether the event should be retained or rejected. This function is called “128 pipeline.” (See Figure 21)

When an event is accepted, the global level-0 trigger decision unit sends a signal to all the “128 pipeline” bits buffer to move the accepted bit (corresponding to an accepted event) to a derandomizing FIFO buffer (see Figure 22). This function is called “FIFO.” For each channel we will have a 27-bit FIFO containing the full information relative to the accepted event. Even though the whole process is synchronous, it is safer to extend the width of the FIFO in each FPGA. At present, 8-bit has been reserved for the time-stamp bunch-crossing counter, however, the length of this field is defined in the configuration file, and can be changed at any time

Each FPGA handles the information of four trigger tower channels, memorizes the information for 128 clock cycles, stores the information relative to the accepted events (at an average of 1 MHz) into 32-bit deep (this parameter can be changed at any time), 80-bit wide FIFO. The width of the output FIFO in each FPGA is calculated as follows: 4 x 12-bit electromagnetic, 12-bit hadronic, 4 x 1-bit preshower, 4 x 2-bit PADs of muon station, and 8-bit time-stamp from a bunch crossing counter that will allow one to verify partial event information at different stages of the data transmission (optical fibers, deserializer, etc.). Thus for each accepted event, each FPGA will send 80-bit through the serializer and the optical fiber to the upper level trigger and DAQ.

A strobe signal received from the upper level decision units and DAQ (called EnOutData in Figures 16 and 22) will read all output FIFOs from the FPGAs at an estimated rate of 1 MHz.

Besides the synchronization, 128-pipeline storage, and derandomization of the full data path, it is also necessary to generate the trigger word to be sent to the 3D-Flow trigger processor. In order to save some 3D-Flow bit-manipulation instruction, the function of formatting the input trigger word can also be implemented into the FPGA (see Figure 20).

At present, the input trigger word is defined as:

- 8-bit electromagnetic calorimeter
- 8-bit hadronic calorimeter

- 1-bit preshower
- 8-bit Pad from muon station M1.

Obviously, the format could be redefined at any time by reprogramming the FPGA through the RS-422 link.

Following are the details of the functions listed above written in “generic VHDL” code suitable to several FPGAs or ASICs.

Likewise a breakdown of the above functions mapped to the ORCA³¹ Programmable Function Units (PFUs) is also provided.

Figure 16. Front-end signal synchronization, pipelining, derandomizing, and trigger word formatting.

6.3.1 Level-0 front-end, VHDL coding of the I/O and signals definition

For each of the units shown in Figure 16, its corresponding excerpt of the VHDL code is provided as follows (the complete VHDL code with the equivalent synthesized version for FPGA OR3T30 and the relative test bench are provided in reference 33).

The design is synchronous. All inputs and outputs are registered with the clock (to avoid code repetition the registering of input and output are not shown for all I/Os in this document; however, there are two examples, one registering the input data on the rising edge, another registering the output data).

This code is implementing one out of the 64 trigger tower channels on the board. The other channels will be identical except for the change of the signal name. In addition, there is the extension of the time-stamp bunch crossing counter to the raw-data output FIFO accepted by the global trigger decision unit. The trigger word is going out of the FPGA as 8-bit every 6.25 ns, since the Top port of the 3D-Flow is multiplexed 2:1.

Table 2. VHDL code for the definition of the architecture of the front-end circuit.

```

---
-----
-- file name : lhcb_FE.vhd
-- author    : Dario Crosetto
--
-- project   : LHCB
-- date      : 2/12/99
-- purpose    : This file implements the front-end signals synchronization,
--              pipelining, derandomizing, trigger word formatter.
--              The initial code is for one trigger channel (4 trigger
--              channels can be accommodated into an OR3T30 FPGA)
--
-----

LIBRARY IEEE;
USE IEEE.std_logic_1164.ALL;
USE IEEE.std_logic_arith.ALL;

-----
--Entity Definition
-----

ENTITY pipeline IS
  PORT (
    clock, reset      : IN STD_LOGIC;
    EM_1              : IN STD_LOGIC_VECTOR(11 DOWNTO 0);
    HD_1              : IN STD_LOGIC_VECTOR(11 DOWNTO 0);
    PS_1              : IN std_logic;
    M1_1              : IN STD_LOGIC_VECTOR(1 DOWNTO 0);
    G_L0              : IN std_logic; -- Global Level-0
    EnInData          : IN std_logic; -- Enable In Data
    EnOutData         : IN std_logic; -- Enable Out Data

    TO_3DF_1          : OUT STD_LOGIC_VECTOR(15 DOWNTO 0);
    DataOut           : OUT STD_LOGIC_VECTOR(26 DOWNTO 0);
  );
END pipeline;

-----
--ARCHITECTURE Definition
-----

ARCHITECTURE rtl OF pipeline IS

  CONSTANT PS_del    : std_logic_vector(1 DOWNTO 0) := "10"; --select delay 2
  CONSTANT HD_del    : std_logic_vector(1 DOWNTO 0) := "00"; --select delay 0
  CONSTANT EM_del    : std_logic_vector(1 DOWNTO 0) := "00"; --select delay 0
  CONSTANT M1_del    : std_logic_vector(1 DOWNTO 0) := "01"; --select delay 1

  SIGNAL Select_del_EM : std_logic_vector(1 DOWNTO 0);
  SIGNAL Select_del_HD : std_logic_vector(1 DOWNTO 0);
  SIGNAL Select_del_PS : std_logic_vector(1 DOWNTO 0);
  SIGNAL Select_del_M1 : std_logic_vector(1 DOWNTO 0);
  SIGNAL EM_1_clkd     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL HD_1_clkd     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL PS_1_clkd     : STD_LOGIC;
  SIGNAL M1_1_clkd     : STD_LOGIC_VECTOR(1 DOWNTO 0);
  SIGNAL Delay1_EM     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL Delay2_EM     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL Delay1_HD     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL Delay2_HD     : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL Delay1_PS     : STD_LOGIC;
  SIGNAL Delay2_PS     : STD_LOGIC;
  SIGNAL Delay1_M1     : STD_LOGIC_VECTOR(1 DOWNTO 0);
  SIGNAL Delay2_M1     : STD_LOGIC_VECTOR(1 DOWNTO 0);
  SIGNAL TEMP_3DF_1    : STD_LOGIC_VECTOR(31 DOWNTO 0);
  SIGNAL SHORT_3DF_1   : STD_LOGIC_VECTOR(15 DOWNTO 0);
  SIGNAL sync_data_EM  : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL sync_data_HD  : STD_LOGIC_VECTOR(11 DOWNTO 0);
  SIGNAL sync_data_PS  : STD_LOGIC;
  SIGNAL sync_data_M1  : STD_LOGIC_VECTOR(1 DOWNTO 0);
  SIGNAL PIPE_EM0      : STD_LOGIC_VECTOR(127 DOWNTO 0);
  SIGNAL PIPE_EM1      : STD_LOGIC_VECTOR(127 DOWNTO 0);
  .....
  SIGNAL PIPE_EM11     : STD_LOGIC_VECTOR(127 DOWNTO 0);
  SIGNAL PIPE_HD0      : STD_LOGIC_VECTOR(127 DOWNTO 0);
  SIGNAL PIPE_HD1      : STD_LOGIC_VECTOR(127 DOWNTO 0);

```

```

.....
SIGNAL PIPE_HD11      : STD_LOGIC_VECTOR(127 DOWNT0 0);
SIGNAL PIPE_PS        : STD_LOGIC_VECTOR(127 DOWNT0 0);
SIGNAL PIPE_M1_0       : STD_LOGIC_VECTOR(127 DOWNT0 0);
SIGNAL PIPE_M1_1       : STD_LOGIC_VECTOR(127 DOWNT0 0);
SIGNAL TEMP_OPIPE      : STD_LOGIC_VECTOR(26 DOWNT0 0);
SIGNAL TO_IN_FIFO      : STD_LOGIC_VECTOR(26 DOWNT0 0);

-----
--

```

6.3.2 Level-0 front-end, VHDL coding of registering the input data on rising edge

```

--Rising Edge triggered
EDGETRI: PROCESS(reset, clock)
BEGIN
    IF (reset = '0') THEN
        EM_1_clkd <= (others => '0');
        HD_1_clkd <= (others => '0');
        PS_1_clkd <= '0';
        M1_1_clkd <= (others => '0');

    ELSIF (clock'EVENT AND clock = '1') THEN
        IF (EnInData = '1') THEN
            EM_1_clkd <= EM_1;
            HD_1_clkd <= HD_1;
            PS_1_clkd <= PS_1;
            M1_1_clkd <= M1_1;

        ELSE
            END IF;
        ELSE
            END IF;
    END PROCESS EDGETRI;

```

Figure 17. VHDL code and graphical representation of registering input data.

6.3.3 Level-0 front-end, VHDL coding of the updating contents of variable delays

```

--adding delays
ADD_DLY: PROCESS (clock, reset)
BEGIN
    IF (reset = '0') THEN
        delay1_EM <= (others => '0');
        delay2_EM <= (others => '0');
        delay1_HD <= (others => '0');
        delay2_HD <= (others => '0');
        delay1_PS <= '0';
        delay2_PS <= '0';
        delay1_M1 <= "00";
        delay2_M1 <= "00";

    ELSIF (clock'EVENT AND clock = '1')
        delay1_EM <= EM_1_clkd;
        delay2_EM <= delay1_EM;
        delay1_HD <= HD_1_clkd;
        delay2_HD <= delay1_HD;
        delay1_PS <= PS_1_clkd;
        delay2_PS <= delay1_PS;
        delay1_M1 <= M1_1_clkd;
        delay2_M1 <= delay1_M1;

    ELSE
        END IF;
    END PROCESS ADD_DLY;

```

Variable Delay

Updating

Figure 18. VHDL code and graphical representation for the updating of the variable delays.

6.3.4 Level-0 front-end, VHDL coding of selecting a variable delay

```
-- change constant based on detector, and/or
electronics, and/or cables length
    Select_Del_EM <= EM_del;
    Select_Del_HD <= HD_del;
    Select_Del_PS <= PS_del;
    Select_Del_M1 <= M1_del;

-- This synchronises EM
    sync_data_EM <= delay2_EM WHEN (Select_Del_EM = "10")
                     ELSE delay1_EM WHEN (Select_Del_EM = "01")
                     ELSE EM_1_clkd;

-- This synchronises HD
    sync_data_HD <= delay2_HD WHEN (Select_Del_HD = "10")
                     ELSE delay1_HD WHEN (Select_Del_HD = "01")
                     ELSE HD_1_clkd;

-- This synchronises PS
    sync_data_PS <= delay2_PS WHEN (Select_Del_PS = "10")
                     ELSE delay1_PS WHEN (Select_Del_PS = "01")
                     ELSE PS_1_clkd;

-- This synchronises M1
    sync_data_M1 <= delay2_M1 WHEN (Select_Del_M1 = "10")
                     ELSE delay1_M1 WHEN (Select_Del_M1 = "01")
                     ELSE M1_1_clkd;
```

Variable Delay

Figure 19. VHDL code and graphical representation for the selection of the variable delays.

6.3.5 Level-0 front-end, VHDL coding of the formatting of the trigger word and multiplexing the output.

This code is defining the trigger word to be sent to the 3D-Flow processor. Any combination of bits available in the FPGA can be used, the same signal can be sent to several 3D-Flow processors, and the format can be changed at a later time

```
-- Multiplexer to output to the trigger only
-- on 16-bit data line the 32 signals
--
    SHORT_3DF_1 <= TEMP_3DF_1(15 DOWNT0 0)
                     WHEN (clock = '0')
                     ELSE TEMP_3DF_1(31 DOWNT0 16);
```

```
-- Format the 32-bit output word to
-- the trigger decision processor
    TEMP_3DF_1 <= sync_data_EM(7 DOWNT0 0)
                  & sync_data_HD(7 DOWNT0 0)
                  & sync_data_PS & "0000000"
                  & sync_data_M1(1 DOWNT0 0)
                  & "000000";
```

At the place of these last & "000000"; the user can insert names of neighboring PADs used in the trigger signature of this trigger channel

Figure 20. VHDL code and graphical representation for formatting and multiplexing the trigger word.

6.3.6 Level-0 front-end, VHDL coding for the 128 pipeline buffer

```
-- Pipelining 128 Stages
PIPE: PROCESS (clock, reset)
BEGIN
  IF (reset = '0') THEN
    PIPE_EM1(127 DOWNTO 0) <= (others => '0');

    ELSIF (clock'EVENT AND clock = '1') THEN

      PIPE_EM0(127 DOWNTO 0) <= PIPE_EM0(126 DOWNTO 0) & sync_data_EM(0);
      PIPE_EM1(127 DOWNTO 0) <= PIPE_EM1(126 DOWNTO 0) & sync_data_EM(1);
      PIPE_EM2(127 DOWNTO 0) <= PIPE_EM2(126 DOWNTO 0) & sync_data_EM(2);
      PIPE_EM3(127 DOWNTO 0) <= PIPE_EM3(126 DOWNTO 0) & sync_data_EM(3);
      PIPE_EM4(127 DOWNTO 0) <= PIPE_EM4(126 DOWNTO 0) & sync_data_EM(4);
      PIPE_EM5(127 DOWNTO 0) <= PIPE_EM5(126 DOWNTO 0) & sync_data_EM(5);
      PIPE_EM6(127 DOWNTO 0) <= PIPE_EM6(126 DOWNTO 0) & sync_data_EM(6);
      PIPE_EM7(127 DOWNTO 0) <= PIPE_EM7(126 DOWNTO 0) & sync_data_EM(7);
      PIPE_EM8(127 DOWNTO 0) <= PIPE_EM8(126 DOWNTO 0) & sync_data_EM(8);
      PIPE_EM9(127 DOWNTO 0) <= PIPE_EM9(126 DOWNTO 0) & sync_data_EM(9);
      PIPE_EM10(127 DOWNTO 0) <= PIPE_EM10(126 DOWNTO 0) & sync_data_EM(10);
      PIPE_EM11(127 DOWNTO 0) <= PIPE_EM11(126 DOWNTO 0) & sync_data_EM(11);
      PIPE_HD0(127 DOWNTO 0) <= PIPE_HD0(126 DOWNTO 0) & sync_data_HD(0);
      PIPE_HD1(127 DOWNTO 0) <= PIPE_HD1(126 DOWNTO 0) & sync_data_HD(1);
      PIPE_HD2(127 DOWNTO 0) <= PIPE_HD2(126 DOWNTO 0) & sync_data_HD(2);
      PIPE_HD3(127 DOWNTO 0) <= PIPE_HD3(126 DOWNTO 0) & sync_data_HD(3);
      PIPE_HD4(127 DOWNTO 0) <= PIPE_HD4(126 DOWNTO 0) & sync_data_HD(4);
      PIPE_HD5(127 DOWNTO 0) <= PIPE_HD5(126 DOWNTO 0) & sync_data_HD(5);
      PIPE_HD6(127 DOWNTO 0) <= PIPE_HD6(126 DOWNTO 0) & sync_data_HD(6);
      PIPE_HD7(127 DOWNTO 0) <= PIPE_HD7(126 DOWNTO 0) & sync_data_HD(7);
      PIPE_HD8(127 DOWNTO 0) <= PIPE_HD8(126 DOWNTO 0) & sync_data_HD(8);
      PIPE_HD9(127 DOWNTO 0) <= PIPE_HD9(126 DOWNTO 0) & sync_data_HD(9);
      PIPE_HD10(127 DOWNTO 0) <= PIPE_HD10(126 DOWNTO 0) & sync_data_HD(10);
      PIPE_HD11(127 DOWNTO 0) <= PIPE_HD11(126 DOWNTO 0) & sync_data_HD(11);
      PIPE_PS(127 DOWNTO 0) <= PIPE_PS(126 DOWNTO 0) & sync_data_PS;
      PIPE_M1_0(127 DOWNTO 0) <= PIPE_M1_0(126 DOWNTO 0) & sync_data_M1(0);
      PIPE_M1_1(127 DOWNTO 0) <= PIPE_M1_1(126 DOWNTO 0) & sync_data_M1(1);

    ELSE
      END IF;
    END PROCESS PIPE;
```

128 PIPELINE

Figure 21. VHDL code and graphical representation of the 128 pipeline buffer.

6.3.7 Level-0 front-end, VHDL coding of moving accepted L-0 event from 128 pipe to FIFO

```
-- Save last word of the 128 pipeline buffer in TEMP_OPIPE in case the event
-- was accepted by the global level-0 trigger and needs to be sent to the
-- output derandomizing FIFO
TEMP_OPIPE(26 DOWNTO 0) <= PIPE_EM0(127) &
  PIPE_EM1(127) & PIPE_EM2(127) & PIPE_EM3(127) & PIPE_EM4(127) &
  PIPE_EM5(127) & PIPE_EM6(127) & PIPE_EM7(127) & PIPE_EM8(127) &
  PIPE_EM9(127) & PIPE_EM10(127) & PIPE_EM11(127) & PIPE_HD0(127) &
  PIPE_HD1(127) & PIPE_HD2(127) & PIPE_HD3(127) & PIPE_HD4(127) &
  PIPE_HD5(127) & PIPE_HD6(127) & PIPE_HD7(127) & PIPE_HD8(127) &
  PIPE_HD9(127) & PIPE_HD10(127) & PIPE_HD11(127) & PIPE_PS(127) &
  PIPE_M1_0(127) & PIPE_M1_1(127);
```

```
-- clocking the raw data globally accepted by G_L0 to the FIFOs
CLK_2FIFO: PROCESS (clock, reset)
BEGIN
  IF (reset = '0') THEN
    TO_IN_FIFO <= "00000000000000000000000000000000";
  ELSIF (clock'EVENT AND clock = '1') THEN
    IF (G_L0 = '1') THEN
      TO_IN_FIFO <= TEMP_OPIPE;
    ELSE
      END IF;
    ELSE
      END IF;
    END PROCESS CLK_2FIFO;
```

```
begin
PROCESS (reset, WRclock, WRen)
BEGIN
  IF (reset = '1') THEN
    WRaddr <= "00000";
  ELSIF (WRclock'event AND WRclock = '1') THEN
    IF (WRen = '1') THEN
      WRaddr <= WRaddr + 1;
    END IF;
  END IF;
  END IF;
  FIFO write address (or counter)
END;
```

128 PIPELINE

VHDL NOT SHOWN

- FIFO read counter (similar to FIFO write counter) .
- code for FIFO is generated by ORCA tools “Scuba”
- clocking all signals and output register (see example on registering input data), and attaching ID counter

Figure 22. VHDL code and graphical representation for moving accepted data from the pipeline to the FIFO.

6.3.8 Mapping the Level-0 front-end circuits into ORCA OR3T30 FPGA

The above “generic VHDL” style suitable for any FPGA or ASIC, if kept as is, will be technology-independent. The synthesis tools of different vendors will translate into gates for their technology. However, the user may further improve the layout for a particular technology in order to best optimize the silicon. (This effort is not convenient for large designs such as the 3D-Flow chip because of the portability and the fact that having a technology-independent design is more important. In the long run, given the rapid advances in technology, it will also be cost effective, eliminating the need to spend many hours to save a few gates in an environment where the gates cost less every year.)

Since this is a small design, and the architecture of the ORCA Programmable Function Unit was known, the exercise of mapping the function into logic was not very complex.

The basic elements of the ORCA architecture used to implement the above functions are a programmable Logic Cell (PLCs), and Programmable Input/Output Cells (PICs). An array of PLCs is surrounded by PICs. Each PLC contains a programmable Function Unit (PFU) containing 8 registers, a supplemental Logic and interconnect Cell (SLIC), local routing resources, and configuration RAM (used in our case to implement the 128 pipeline buffer).

Following is the resulting optimization, calculated for four trigger channels that can be implemented in a OR3T30 FPGA device (which the synthesis tools from ORCA may not recognize from the above code).

Table 3. Mapping the Level-0 front-end circuit into ORCA OR3T55 FPGA.

Function	# of PFU	Comment
Input register	0	Use PIC registers
Variable delay	20	1 PFU per 4 input bits
3DF interface	32	
128-clock pipeline	80	1 per input bit
Counters (for 128-clock pipeline)	9	
32x80 FIFO	20	4 bit per PFU (use dual-port memory)
80-bits Parallel In, Serial Out regs	10	
5-bit read pointer	4	For FIFO read pointer
5-bit write pointer	4	For FIFO write pointer
Miscellaneous	3	

The total number of PFUs required is 182. The OR3T30 contains 196 PFU.

6.3.9 Front-end circuits 'Real-Estate' and cost considerations

The result of this study of mapping the front-end circuits to the ORCA OR3T30³¹ provides useful information to make a decision of cable length, location of electronics, etc.

This study originated from the need of generating the trigger word for the 3D-Flow level-0 trigger system. It would have been neither technically efficient, nor cost effective to design a separate circuit to extract the trigger information from the full granularity of the DAQ (or higher-level trigger) signals.

The extraction of the level-0 trigger word, is well integrated (as shown in Section 6.3.5) into the circuit of the front-end that is performing the functions of input data synchronization, pipelining, derandomizing (FIFO). In fact, a design is more likely to turn out to be harmonious if all problems, from the front-end to the global level-0 decision unit are considered as a whole problem to be solved rather than to split it into different stages with the risk of losing performance in the interface between two circuits designed by two different persons.

The result of this study provides the following information:

- 16 FPGA per board would be needed for the front-end electronic and trigger word extraction of 64-trigger tower. The total calorimeter and muon station 1, front-end electronics will requires 1536 FPGAs
- Only about 375 additional OR3T30 FPGAs are required to complete the FE for all subdetectors participating in the level-0 trigger. The calculation is as follows: the remaining subdetectors are the muon station 2, for 12,000 bits, and stations 3, 4, 5 for 6000-bits for a total of 30,000-bits. Assuming that the above function be implemented for 80-bit per FPGA OR3T30, about 375 additional components are needed.
- The constraints of mapping the circuit into the FPGA are: a) the ORCA PFU architecture is well optimized if the range of the variable delay that performs synchronization is limited from 0 to 2, b) the pipeline depth should not be greater than 128.
- Purchasing about 2000 FPGA chips will provide maximum flexibility in downloading different circuits at a later stage. A Masked Array Conversion for ORCA (MACO) version of the circuit that was designed in FPGA is more cost-effective for a volume of a few thousands of pieces; however, it will not allow for future changes to be made to the circuit.

6.4 Logical-To-Physical Layout of 64 channels/10 layers on the 3D-Flow board

The optimized layout of the 3D-Flow PC board needs to take into account the need to communicate both with neighboring processors in the same layer (NEWS ports), as well as along the successive layers (Top and Bottom ports). In the current implementation, each layer is represented by 4 IC's (64 channels per board, 16 processors per IC). Each stack consist of 12 layers, i.e. 10 layers of actual pipelined algorithm execution (as discussed in Section 2 and in Section 6.5) followed by two more layers to provide the first stages of data funneling (the "pyramid").

One key element to keep in mind is that, while data transfer among layers occurs at every clock cycle, only about 10% of the time data are exchanged within the same layer. These considerations (see Table 1) have led to the layout shown in Figure 23. Sequential numbers of chips on the board physical layout (Figure 23, bottom left) indicate chips in the same x/y position in the logical scheme (right of Figure 23) corresponding to the position in subsequent layers, while chips numbered 1, 13, 25, and 37 correspond to the 64 processors of the first layer of the 3D-Flow system that are connected to the FPGAs which send the formatted trigger word of the data of the detector.

The chips corresponding to the first layer (labeled 1, 13, 25, and 37) are positioned in the central column of the board, while the remaining elements of each stack (2 to 12, 14 to 24, etc.) follow the arrowhead pattern shown in Figure 23 (note that chips 9-12, 21 to 24, etc., are positioned on the board's opposite side, as shown in Figure 24).

This layout allows for each group of 16 processors to keep the minimum PCB trace distance for the Bottom-to-Top connection between chips belonging to different layers.

All 3D-Flow inter-chip Bottom-to-Top ports connection are within the board (data are multiplexed 2:1, PCB traces are shorter than 6 cm), while all 3D-Flow inter-chip North, East, West, and South port connections between boards and crates are multiplexed (8+2):1 and are shorter than 1.5 meters.

Figure 23. 3D-Flow layer interconnections on the PCB board.

6.5 On-board Data-Reduction, Channel-Reduction and Bottom-To-Top Links

Figure 24 shows the relation between the logical layout of a stack of 3D-Flow chips, its implementation in hardware, and the functionality performed by processors in different layers in a stack.

The left bottom part of Figure 24 shows the top part of the mixed-signal processing board (front and rear), whereas the chip arranged in a logical position are shown in the right part of the figure.

The layout immediately shows that Bottom-to-Top connections can be kept within 6 cm, allowing minimum latency in data propagation in a synchronous system at 80 MHz.

Processor number 1 receives the trigger word data from the FPGA (or detector data). Up to two 16-bit words of information can be received by processor 1 at each bunch crossing. During the subsequent clock cycles, processor 1 executes the user trigger algorithm (including data exchange with its neighbors on the same layer, on-board, off-board, or off-crate).

The interconnection between neighboring elements, typical of the 3D-Flow architecture, allows to implement searches for energy deposition in 2x2, 3x3, 4x4, 5x5, 7,7, etc., clusters of neighboring calorimeter elements. It will not be required to redesign the board, but it could be achieved simply by reprogramming the processors..

After a layer of processors has received the data relative to one bunch crossing (or, more in general, one “frame”), further incoming data are bypassed (according to the setting of the bypass switches) to the next layer of processors (as shown in Figure 25). After 10 bunch crossings, the next set of data is fetched again by the processor of layer 1, which in the mean time has finished the execution of the algorithm and placed the result in a local output FIFO buffer. The same clock cycle used to fetch the input data is also used to transmit the results of the previous calculation to the bottom port.

This same board design could easily be adapted to situations where, because of simpler algorithms, fewer than ten layers are required to keep up with the incoming data. In this case, one would have a not fully populated board, with jumpers to bypass the unused locations (See Figure 26 and Section 6.6 for more details). The number of connections for the inter-boards and inter-crates North, East, West, and South will also be reduced to the number of layers used by the simpler algorithm, thus not requiring all cables with RJ45 connectors to be installed.

.Figure 24. Bottom-to-Top Links on the PCB board.

As the outcome of the process described above, the results applying the trigger algorithm to the data of each bunch crossing arrive every 25 ns to the processors in the first layer of the pyramid (layer 11). Their task is to check whether an event of interest (high P_T electron, photon or hadron) has been reported. In the affirmative case, time stamp and block ID are attached to the results, and the full information is forwarded to the next layer (layer 12).

Layer 12, the “base” of the channel-reduction pyramid, receives at most a few validated candidates at every bunch crossing. Only two of the four layer-12 chips are connected, via the Bottom-to-Top ports, to the next layer 13, containing only two chips.

The accepted candidates are first routed internally, within layer 12, to the “exit points,” from where they are transmitted to the next layer 13 (see center of Figure 13, and Figure 15, 3D-Flow chips). The channel-reduction process is going, to layers of fewer and fewer channels, until the results are sent to the global level-0 trigger unit.

Figure 25. Position of the bypass switches for the data flow (Input/Output) from Top-to-Bottom ports.

6.6 Details of the On-Board Bottom-To-Top Links (6 cm)

In order to keep the distance from the bottom port to the top port to a minimum, the pin assignment of the 3D-Flow needs some consideration.

There are 16 processors on a chip; all 16 processors have top and bottom port signals multiplexed 2:1 connected to the pins of the chip (600-pin EBGA @ 2.5 Volt, with dimensions of 40 mm x 40 mm and a pitch spacing between balls of 1.27 mm could be reduced next year to 1 mm pitch providing a 676-pin EBGA @ 1.8 Volt, with dimensions of 27 mm x 27 mm). Moreover, 12 processors also have some of the North, East, West and South ports connected to the pins. (The other connections between NEWS ports are internal to the chip).

For each of the 16 processors (see Figure 26), the top-bottom ports are kept within a group of 25 pins (8-data lines and 2 control lines for the top port, and 8 data lines and 2 control lines for the bottom port; the remaining 5 pins are reserved for VCC and GND). Furthermore, the pin of Bit-0 of the top port is adjacent to the pin of bit-0 of the bottom port, and so on for all bits.

Figure 26. Bottom-to-Top Links on the PCB (details).

This could be of some advantage to the user who might not need to populate the entire board of 3D-Flow chips because of a simpler and faster trigger algorithm. In such a case, a simple jumper between the top and bottom ports would eliminate the need to redesign the entire board.

For the 12 processors that have some NEWS ports connected to the pins of the chip, only a group of five pins is necessary, two transmit and two receive; the remaining are used either for VCC or for GND, depending on whether there are more neighboring pins of one type or another in a given area. The presence of two twisted-pair links enables simultaneous communication of data in both directions. In the case of very complex algorithms requiring little neighboring communication but longer programs, one could limit the communication to one direction at a time, saving 50% of the links and thus having twice as many layers in the 3D-Flow system for the same number of connections on the backplane.

7 CRATE(S) FOR 3D-FLOW SYSTEMS OF DIFFERENT SIZES

A 3D-Flow system of any size can be built even if it exceeds the number of channels that can be accommodated into a single crate.

7.1 Crate Backplane LVDS Links Neighboring Connection Scheme

Figure 27, bottom right, shows how the actual layout of the calorimeter block (or trigger tower) is mapped onto the boards in the needed set of crates, while on the left is shown the corresponding physical layout of the boards within the crate.

In order to minimize the connection lengths, the first board in a crate is followed immediately by the board containing the “below” processors (which was called “south” in the 3D-Flow nomenclature), and then by the “rights” ones (e.g., the board 18, to the right of 17, in the physical layout occupies the position below board 17 in the logical layout. The next board (19) will be to the left of 18 in the physical layout and to the left of 17 in the logical layout, and so on). The corresponding backplane connectors link the bottom part of each odd-numbered board (3D-Flow south) to the top (3D-Flow north) of the even-numbered to its right, while the East-West links run between either even-to-even or odd-to-odd board-locations.

Since there are 10 layers of processors in a stack and each layer has four links in each direction (for a total of 16 links per layer), the 160 LVDS links are required from one board to its neighbour in any NEWS direction. Each LVDS link has two wires, requiring than a total of 320 pins in each direction.

Figure 27. 3D-Flow System LVDS Links Neighboring Connection Scheme.

7.2 Number of NEWS Links for the chip-to-chip, board-to-board, crate-to-crate

Figure 28 summarizes the number of LVDS links between chip-to-chip, board-to-board, and crate-to-crate.

Figure 28. 3D-Flow North, East, West, and South LVDS Links.

7.3 Implementation of the Backplane Crate-To-Crate LVDS Links (Option 1)

One option in the implementation of the interconnection scheme shown in Figure 29 is to use AMP³⁴-646372-1 and AMP³⁴-646373-1 long feedthrough pins (through the backplane printed circuit) connectors.

At the rear part of the backplane one can insert female connectors into the long feedthrough pins, as shown at the left-bottom of Figure 29 (Courtesy of AMP. Catalog 65911). The male shroud is fitted with snap latches for firm retention of cable connector housings.. Even though this solution is compact and elegant, it is not very practical; it is difficult to find parts because it is not of a standard construction, and it is also very expensive.

Figure 29. Crate-To-Crate Backplane LVDS Links (Option 1).

7.4 Implementation of the Backplane Crate-To-Crate LVDS Links (Option 2)

This solution of option 2 is very low in cost and it is practical because it makes use of parts that are widely used in consumer computer electronics. The final aspect, however, will not look much different from the racks of the local area network (with many panels with female RJ-45 connectors and many RJ-45 cable/connectors) of a large company or of an internet service provider.

At the rear connector of each board (front-board), a second board (rear-board) is inserted into the long feedthrough pins of connectors AMP³⁴-646372-1 and AMP³⁴-646373-1. There will be no electronics on this rear-board -- just female connectors RJ-45. Since the RJ-45 are widely used, they come in blocks of 8, or 4 assembled for printed circuit mounting. For each rear board four rows (positioned as shown in Figure 30 to allow insertion of the male connector in between the four rows) of RJ-45 connectors (each with 20 female RJ-45 connectors) are needed. Each row is made of two parts AMP³⁵ 557573-1 and one part 557571-1.

The rear-board will have only four blocks (out of seven male connectors installed on the backplane) of female connectors AMP³⁴ 646372-1, or AMP³⁴ 646373-1 on the backplane side, since only 320 pins are needed to carry 160 LVDS links to the board on the neighboring crates.

Should the overall 3D-Flow system need to be expanded to the east and west, the two boards at the far right and at the far left of the crate will make exceptions in having RJ-45 female connectors assembled on both sides, and they will have two more female connectors AMP³⁰-9-352153-2, or AMP³⁰-9-352155-2 on the backplane side, since they have to carry 160 links to the West and/or to the East crates.

The total number of cables to the North and South crates will then be 640, while the cables to the East and West crates will number only 40. In the case of applications requiring simpler real-time algorithm (e.g., requiring less than 20 steps, that is equivalent to 10 layers of 3D-Flow processors), then the number of connections for the inter-boards (North, and South), and inter-crates (East, and West) will also be reduced to the number of layers used by the simpler algorithm, thus not requiring all cables with RJ45 connectors to be installed (e.g., applications requiring only 9 layers of 3D-Flow processors will save 64 cables to the North, 64 to the South, 4 to the East, and 4 to the West crates).

The cable used for this solution can be found at any computer store. Such cables come assembled at different lengths (in our case, a standard 3 feet is needed), with two male connectors at both ends and tested at different categories for different speeds. The cost would be about \$2 each.

Figure 30. Crate-To-Crate Backplane LVDS Links (Option 2).

7.5 The 3D-Flow Crate

The 3D-Flow crate is built in such a way that allows connection of several crates to the four sides (North, East, West, and South--NEWS) in order to allow the user to build 3D-Flow systems of any size while keeping the maximum distance between components to less than 1.5 meters. It is very important to keep the maximum distance as short as possible in synchronous systems and where the overall performance depends on the data exchange with neighboring elements.

Figure 31 shows the 3D-Flow crate as a modular part of a larger 3D-Flow system made of several crates. The overall features of a crate are based on the number of channels and the 3D-Flow processor speed. A conservative choice of components and technology sets the number of channels at 1024 (64 per board) and the processor speed at 80 MHz.

Figure 31. The 3D-Flow Crate.

In summary, a 3D-Flow crate, built in standard 9U x 84HP x 340 mm dimensions, accommodating 16 mixed-signals processing 3D-Flow boards has the following features:

Backplane communications within the crate:

- The backplane of the crates establish the communication of four groups of 320 pins from the connectors of each of the 16 boards with the neighboring (and off-crate) boards. The above connections implement the North, East, West, and South 3D-Flow connection scheme. The backplane connectors link the bottom part of each odd-numbered board (3D-Flow South) to the top (3D-Flow North) of the even-numbered (board or connector) to its right, while the East-West links run between either even-to-even or odd-to-odd board-locations (See Figure 27).

Off-crate communications:

- Each crate communicates through 2560 LVDS links to North and South crates when several of them are required because the entire application cannot fit into a single crate. In the case of applications requiring a simpler real-time algorithm (e.g., requiring less than 20 steps, which is equivalent to 10 layers of 3D-Flow processors), then the number of connections for the inter-boards (North and South), will also be reduced to the number of layers used by the simpler algorithm, thus not requiring all cables with RJ45 connectors to be installed (e.g., applications requiring only 9 layers of 3D-Flow processors will save 32 cables to the North and 32 to the South crates).
- Each crate communicates through 160 LVDS links to East and West crates. For the same reason explained above, a simpler algorithm which does not require all 10 layer of 3D-Flow PEs, will reduce the number of cables required to the East and West crates (e.g., applications requiring only 9 layers of 3D-Flow, will save 4 cables to East and 4 cables to West)

8 GLOBAL LEVEL-0 TRIGGER

Figure 32 shows the Global Level-0 trigger decision units. It consist of two rear-boards with no electronics, but only connectors. The board receiving the candidates particles from the calorimeter level-0 trigger crates has 96 cables (one per mixed-signal processing board). The information goes through the back-panel connector through connectors AMP³⁴ 646372-1 and AMP³⁴ 646373-1 to the board at the front of the crate called CALO L-0. This board is shown at the bottom right of Figure 32. The programmable global level-0 trigger decision board for the calorimeter (or the candidates that need to be validated by the other muon global level-0 decision unit) sends out the calorimeter information through the front-panel connector RJ-45 to the Global level-0 calorimeter board the calorimeter information. The CALO L-0 board contains 3D-Flow chips and FPGA chips that allow to implement a global level-0 algorithm in a programmable form. The Muon L-0 board has the same functionality as the CALO L-0 board. Finally, the Global L-0 decision unit shown at the bottom-left of the Figure 32 receives the data through two RJ45 connectors on the front panel from Calo L-0 and Muon L-0; it performs further sorting and global level-0 trigger algorithms in order to generate a single yes/no signal that will be sent to all the units in the calorimeter crates and to the muon crates. These signals are sent through AMP³⁶ 200346-2 connectors on the same coaxial ribbon cable used at the front panel of each mixed-signal processing board. (Only one coax cable out of the 17 in each coax ribbon cable is attached to this connector from each mixed-signal processing board. See how coax cables are split at one end in Figure 12).

Figure 32. LHCb Programmable Global Level-0 Trigger Decision Units.

9 TIMING AND SYNCHRONIZATION ISSUES OF CONTROL SIGNALS

The 3D-Flow system is synchronous. This makes it easier to debug and to build.

The most important task is to carry the clock, reset and trigger signals to each 3D-Flow component pin within the minimum clock skew. (The overall task is easier if each component accommodates 16 processors.)

This task can be accomplished without using special expensive connectors, delay lines, or sophisticated expensive technology since the processor speed required to satisfy the design is running at only 80 MHz. The expected worst clock skew for the distribution of one signal to up to 729 chips using components PECL 100E111L or DS92LV010A Bus LVDS Transceiver, is less than 1 ns (This parameter is calculated by using the worst skew between different components that is reported in the component data sheet). The maximum skew for a system of 11,664 processors will be of 450 ps. Fanout to 104,976 3D-Flow processors could be accomplished by adding one stage in the clock distribution, increasing the maximum signal skew to 650 ps.

Figure 33. Scheme of the control signal distribution with minimum skew.

Designing equal-length printed circuit board traces is not difficult to achieve with the aid of today's powerful printed circuit board layout tools such as Cadence Allegro.

The other consideration in building the 3D-Flow system is that all input data should be valid at the input of the first layer of the 3D-Flow system at the same time. This goal is achieved as described in Section 6.3, and 6.3.5.

All other signals in the 3D-Flow system are much easier to control than for any other system (given the modularity of the 3D-Flow approach) because they are of short distance, reaching only the neighboring components.

10 ASIC DESIGN VERIFICATION

Figure 34 shows the ASIC verification design process. The user's real-time algorithm is simulated on the SYSTEM TEST-BENCH. Expected results (top right) are checked versus different input data set (top left). Bit-vectors for one or more PEs (for any PE in the system) are saved to a file (center bottom). Test-bench parameters for any PE(s) are generated by the system test-bench for software (center left) and hardware (center-right) simulator. All bit-vectors are compared for design validation.

SW/HW Verification (4/16 PEs)

Figure 34. ASIC design verification. From user's system algorithm down to the gate-level circuit.

11 HOST COMMUNICATION AND MALFUNCTIONING MONITOR

An essential part of the 3D-Flow design is that every single processor is individually accessible by a supervising host, via an RS-232 line (or through an RS-422 that is subsequently converted to RS232 if long distances not reachable by RS232 is required). One RS-232 serial port controls a group of four 3D-Flow PEs, including all PEs in subsequent layers behind the first layer (also called 3D-Flow stack. See Figure 24). In addition to providing the ability to download and initialise the system, this feature also provides the capability to periodically test the processor's performance by downloading test patterns and/or test programs. A continuous monitoring can be performed by reading through RS232 the status of eight consecutive cycles of all processors and comparing them with the expected ones. These status bits are simultaneously saved into a silicon scratch pad register of all processors at a pre-recorded trigger time corresponding to a selected line of the program executing the filtering algorithm in a selected layer.

In the case of a suspected or detected malfunction, the processor performance could be tested remotely and its performance diagnosed. In the event of catastrophic malfunction (e.g. a given processor completely failing to respond, or a broken cable), normal operation, excluding the offending processor (or connection), can still be maintained by downloading into all the neighbours a modified version of the standard algorithm, instructing them to ignore the offending processor.

Obviously, physics considerations would dictate whether such a temporary fix is acceptable, but it is a fact that the system itself does contain the intrinsic capability of fault recovery, via purely remote intervention. Figure 35 shows the cost of one IBM PC workstation and peripherals/cables required to monitor one 3D-Flow crate.

Demonstrator SM for 128 Trigger Channels

(Using off-the-shelf components)

SM for 1024 Trigger Channels

$$1 \times \left[\begin{array}{l} 1 \text{ x IBM-PC @ } \$1500 \\ 8 \text{ x Rockets @ } \$561 \\ 16 \text{ x Panels @ } \$200 \\ 256 \text{ x RS232 cables @ } \$2.6 \end{array} \right] = \$9,853$$

SM for 6144 Trigger Channels

= \$59,118

Figure 35. Demonstrator of a System Monitor for 128 3D-Flow Channels.

Table 4 shows the performance of the System Monitor (SM) tested on 128 channels connected via 32 x RS232 @ 230.4Kbaud. The connection was made between two IBM-PC computers using one PCI RocketPort³⁷ board with 32 x RS-232 installed on the System Monitor and one ISA RocketPort³⁷ board with 32 x RS-232 installed on the Virtual Processing System (VPS) computer. The cost of each board was \$561. Four 16-port switch-selectable (RS-232/RS-422) interface box of the cost of \$200 each and 32 cables with 32 null-modem were necessary to make the connections between the two computers.

Even if the board setting of the communication speed at each port allowed 460.8Kbaud, the test were carried at 230.4Kbaud because it was detected a bottleneck resulting from the multiplexing of the signals on the cable connecting the 16-port switch to the ISA (or PCI) board. When all 32 ports were used at the same time, there was a minimal increase in throughput performance if 230.4 Kbaud, or 460.8 Kbaud were selected.

On one computer was installed the System Monitor (SM) program, while on the second computer was installed the Virtual Processing System (VPS) program. The SM was initializing and monitoring the VPS only through the 32 RS-232 serial ports.

Control signals (3D-Flow system reset, input data strobe, etc.) to the VPS were generated by the SM and sent through the standard COM1: of the two computers. The time one PC computer could execute all functions (loading, monitoring, etc.) on 1024 PEs was estimated by extrapolation.

Table 4. System Monitor Demonstrator test results for 128 channels.

FUNCTION	# of PEs	Current [sec]	Ideal [sec]	Reachable [sec]
Loading & Initializing	1280	112	2	6
Monitoring	4	1.6	0.001	0.5
Monitoring one Layer	128	8.65	0.1	4.8 (0.8)*
Monitoring the entire System	1280^	86	1	30 (8)*

^ The system under test was made of 10 layers, each RS-232 is addressing a stack of 4 PEs (4 PEs x 32 RS-232 x 10 layers = 1280 PEs)

* In parentheses is the timing using the 3D-Flow hardware at the place of the VPS.

Table 5. System Monitor estimated timing for 1024 channels.

Function	# of PEs	Estimated time [sec]
Loading & Initializing	10,500	~ 60
Monitoring	4	~ 0.5
Monitoring one Layer	1024	~ 2
Monitoring the entire System	10,500^	~ 20

^ The estimated 3D-Flow system includes: 4 PEs x 256 RS-232 x 10 = 10,240 + 3D-Flow pyramid = 10,500 PEs.

12 SOFTWARE TOOLS

The 3D-Flow Design Real-Time is a set of tools that allows the user to:

1. create a new 3D-Flow application (called project) by varying size, throughput, filtering algorithm, and routing algorithm, and by selecting the processor speed, lookup tables, number of input bits, and output results for each set of data received for each algorithm execution;
2. simulate a specified parallel-processing system for a given algorithm on different sets of data. The flow of the data can be easily monitored and traced in any single processor of the system and in any stage of the process and system;
3. monitor a 3D-Flow system in real-time via the RS232 interface, whether the system at the other end of the RS232 cable is real or virtual;
4. create a 3D-Flow chip accommodating several 3D-Flow processors by means of interfacing to the Electronic Design Automation (EDA) tools.

A flow guide helps the user through the above four phases.

A system summary displays for a 3D-Flow system created by the Design Real-Time tools, the following information:

1. characteristics, such as size, maximum input data rate, processor speed, maximum number of bits fetched at each algorithm execution, number of input channels, number of output channels, number of layers filtering the input data, number of layers routing the results from multiple channels to fewer output channels;
2. time required to execute the filtering algorithm and to route the results from multiple channels to fewer output channels.

A log file retains the information of the activity of the system when:

1. loading all modules in all processors;
2. initializing the system;
3. recording all faulty transactions detected in the system (e.g., data lost because the input data rate exceeded the limit of the system or because the occupancy was too high and the funnelling of the results through fewer output channels exceeded the bandwidth of the system);
4. recording any malfunction of the system for a broken cable or for a faulty component.

A result window can be open at any time to visualize the results of the filtering or pattern recognition algorithm applied to the input data as they come out at any layer of the system.

The generation of test vectors for any processor of the system can be selected by the user at any time to create the binary files of all I/Os corresponding to the pins of a specific FPGA or ASIC chip. These vectors can then be compared with those generated by the chip itself or by the VHDL simulation.

Figure 36 shows the interrelation between the entities in the Real-Time Design Process, while Figure 37 shows some of the windows available to the user to create, debug, and monitor a 3D-Flow system with different algorithms of different sizes and simulate it before construction

Figure 36. Interrelation between entities in the Real-Time Design Process.

Figure 37. Design Real-Time software tools

Acknowledgments

The warmest acknowledgment goes to the SBIR office of the DOE of Dr. Rober Berger, for all its support, including a lot of moral support. This was demonstrated most recently in December 1998 and earlier along with Dawnbreaker in the commercialization assistance program. Great support during all this project was also received from Albert Werbrouck, Sergio Conetti (who also reviewed this article), and Billy Bonner, to whom I am very grateful. I would also like to acknowledge Jeannine Robertazzi for providing hints on VHDL, J. Vorgert for his help in optimizing the mapping of the front-end design to ORCA FPGAs, P. Fathi from National Semiconductors for information regarding LVDS, and J. Naples for editing assistance.

APPENDIX A: COST/PERFORMANCE COMPARISON OF THE CALORIMETER FRONT-END AND LEVEL-0 TRIGGER. (Answers to questions by L0 trigger group).

Table 6. Calorimeter Front-end and Level-0 trigger cost implementation comparison (the cost of the front-end electronics of the PreShower and of the Pad Chamber are not included in this table).

ITEM	LAL Orsay (9U)		3D-Flow (9U)		Bologna (9U)	
	Number of units	Cost [KCHF]	Number of units	Cost [KCHF]	Number of units	Cost [KCHF]
Board Design	4	340	1	85	4	340
Backplane Design	3	165	1	55	3	165
Crates inter-cabling			40	10		
Boards	304	1216	96	480	477	2052
Backplanes	20	80	6	24	40	160
Crates	20	200	6	60	40	400
Total		2001		714		2973

Legenda:

- LAL board design (4): front-end card (248 units) - (LHCb 98-058, Sec. 4.1; LHCb 99-007³⁸, Sec. 2); ECAL summary card (28 units) - (LHCb 98-058, Sec. 4.2.2, LHCb 99-007, Sec. 3.2.1); HCAL summary card (8 units) - (LHCb 98-058, Sec. 4.2.3; LHCb 99-007, Sec. 3.2.2); selection card (18 units), selection controller card (2 units) - (LHCb 98-058, Sec. 4.3; LHCb 99-007, Sec. 3.3);
- 3D-Flow board design (1): 3D-Flow mixed-signal processing board (96 units) - (see this note, Sec. 6.1, 6.4, 6.5, and 6.6; LHCb 99-006, Sec. 5, and 6);
- Bologna board design (4): front-end card (212 units. Since L0 trigger is implemented for 3x3 algorithm in separate cards rather than 2x2 as for LAL, two cards per each LAL front-end crate are not used) - (LHCb 98-034³⁹, Sec. 3.2; LHCb 6 May 99⁴⁰, Sec. 3, and Table 7); ECAL L0 card (208 units) - (LHCb 98-034, Sec. 3.2; LHCb 6 May 99, Sec. 3, and Table 7); HCAL L0 card (56 units) - (LHCb 98-034, Sec. 3.2; LHCb 6 May 99, Sec. 4, and Table 7); Message dispatcher card (1 unit) - (LHCb 6 May 99, Sec. 3.6, and Table 7);
- LAL backplanes (3) - front-end cards backplane (LHCb 98-058, Sec. 4, Figure 2; LHCb 99-007, Sec. 2, Figure 3, Table 1 and 2); summary cards backplane (LHCb 98-058, Sec. 4.2; LHCb 99-007, Sec. 3.2); selection cards backplane (LHCb 98-058, Sec. 4.3; LHCb 99-007, Sec. 3.3);
- 3D-Flow backplane (1) - (see this note, Sec. 7.1, and 7.4);

3D-Flow crates inter-cabling (40) - (see this note, Sec. 7.1, 7.2, and 7.4 Figure 27, 28, and 30); each unit consist of a) two boards only with connectors and no components as shown on the bottom right of Figure 30, and b) 40 cables with RJ-45 connectors 3 feet long, which are shown at the bottom of Figure 30.

One unit is used to make a connection between two boards in two crates, e.g. board "2" is connected to board "17" in Figure 27. The cost of the cables are CHF 2.5/each, for a total of CHF 100, the cost of the 4 Z-pack connectors is CHF 40, the cost of the RJ-45 PCB connectors is CHF 35, and the cost of two dummy boards accommodating the connectors is CHF 75, for a total of CHF 250 per unit.

- Bologna backplanes (3) - front-end cards backplane (LHCb 98-058, Sec. 4, Figure 2; LHCb 99-007, Sec. 2, Figure 3, Table 2); ECAL L0 cards backplane (LHCb 98-034, Sec. 3.2; LHCb 6 May 99, Sec. 3, and Table 7); HCAL L0 cards backplane (LHCb 98-034, Sec. 3.2; LHCb 6 May 99, Sec. 4, and Table 7);
- LAL crates (20) - (Transparencies by F. Harris LHCb collaboration week 16 Feb 1999: 14 crates FE-logic for ECAL, 4 crates for HCAL; LHCb 99-007, Table 2, 2 crates for L0)
- 3D-Flow crates (6) - (see this note, Sec. 7.5)
- Bologna crates (40) - (Transparencies by F. Harris LHCb collaboration week 16 Feb 1999: 14 crates FE-logic for ECAL, 4 crates for HCAL; LHCb 6 May 99, Table 7, 22 crates for L0)

Prices have been uniformed with the exception of the "3D-Flow mixed-signal processing board," which has been considered more expensive.

LAL and Bologna boards (36.6 cm x 40 cm) have been estimated in average CHF 4000/board

The "3D-Flow mixed-signal processing board," (36.6 cm x 34 cm) even if smaller, it has been considered more expensive because, as it is shown in this note Figures 12, and 13, it has some components assembled on the rear of the board. The estimated cost for this board in table 1 is CHF 5000/board.

If it will be required to conform to a standard size, the board could be built also 36.6 cm x 40 cm as the others.

For the 3D-Flow solution the interconnecting cables (see this note, Sec. 7.2), and the PCB rear boards of the 6 crates (See this note, Sec. 7.4) carrying only connectors have also been included in the cost. The equivalent cost has not been included in the other two alternative solution because it was difficult to estimate them from the current LHCb notes (several part number were not indicated).

- The cost to design a 9U board (400 mm in depth) has been estimated CHF 85000.
- The cost to design a backplane has been estimated CHF 55000.
- The cost of a backplane has been estimated CHF 4000 (at any given speed, the backplane of the 3D-Flow system will cost less than the backplanes of LAL, or BO systems, because of the 3D-Flow architecture which simplifies things in many area. See Appendix C for details).
- The cost of a 9U crate has been estimated CHF 10000.

The cost of the board of the alternative systems (LAL-Orsay and Bologna) has been estimated in Table 6 to be CHF 2.7/cm², while it has been estimated at CHF 4/cm² for the 3D-Flow mixed-signal processing board and CHF 5.5/cm² for the 3D-Flow front-end preshower board of Table 7.

Table 7. Cost comparison of the Preshower Front-end electronics.

ITEM	LAL Orsay = Bologna (9U) 14 m ² of PCB in total		3D-Flow (9 cm x 9 cm) 3 m ² of PCB in total	
	Number of units	Cost [KCHF]	Number of units	Cost [KCHF]
Board Design	1	85	1	20
Backplane Design	1	55		
Boards	96	384	375	169
Backplanes	8	32		
Crates	8	80		
Total		636		189

Legenda:

- LAL (or Bologna) board design (1) - front-end preshower cards with 64-channels each (96 units) (LHCb 98-058, Sec. 4.1; however, modified for 64-channels 9U boards; LHCb 99-007, Sec. 2.1, Figure 3, Table 1 and 2);
- LAL (or Bologna) backplane design (1) - front-end backplane for the 64-channels preshower cards (8 units) - (LHCb 98-058, Sec. 4, Figure 2, however, modified for 64-channels 9U boards; LHCb 99-007, Sec. 2, Figure 3, Table 1 and 2);
- LAL (or Bologna) crates (8) - (LHCb 98-058, Sec. 4, Figure 2; however, modified for 64-channels 9U boards; LHCb 99-007, Sec. 2, Figure 3, Table 1 and 2);
- 3D-Flow front-end preshower board is described in Appendix B.

Considering that the PCB size of the LAL front-end board for preshower with 64-channels is 18 times the size of the 3D-Flow front-end board for preshower with 16-channels, the cost difference of CHF 450 for 3DF and CHF 4000 for LAL, that is only 8.8 times, provides a realistic term of cost comparison among the two systems.

Besides the importance of the cost in one experiment, the more important is the performance of the level-0 trigger, and its flexibility to accommodate future changes. Here is a list with references of the features/performances

Table 8. Calorimeter Level-0 Trigger Features/Performances.

ITEM	LAL	3DF	BO	REFERENCES
2x2 level-0 Algorithm	X	X		LHCb 98-058, LHCb 99-007, LHCb 96-2, TRIG 96-1, See this note
3x3 level-0 Algorithm		X	X	LHCb 98-034, LHCb 6 May 99, LHCb 96-2, TRIG 96-1, See this note
Fully programmable L0		X		LHCb 96-2 TRIG 96-1, See this note, lhcB 99-006
Trig-word changes Add L0 subdetectors		X		See this note, Sec. 6.1, 6.2, 6.3.5 LHCb 99-006, Sec. 5.2, 5.2.2
No boundary limitation		X		See this note Sec. 2, 6, 7.1, 7.2, 7.4, LHCb 99-006, Sec. 5, 6
Modular		X		See this note Sec. 6, LHCb 99-006 Sec. 5, 6
Scaleable		X		See this note Sec. 7.5
Technology-independent		X		See this note Sec. 3. LHCb 99-006 Sec. 4, Table 3

Legenda:

- 2x2 level-0 trigger algorithm is the preferred algorithm by LAL group. This algorithm can also be executed on the 3D-Flow system as a simplified version of the 3x3 algorithm. Instead of having each processor sending/receiving data from its 8 neighbors, each one sends only one data to south and east and fetches from north and west. The algorithm is much simpler requiring less steps.
- 3x3 level-0 trigger algorithm is the preferred algorithm by the Bologna group, however it is fixed and even slight changes in speed requires the redesign of the complete system, not only using faster components to meet the increased speed, but the functionality has to change. The 3x3 level-0 algorithm has been simulated in detail on the 3D-Flow system.
- Fully programmable level-0 trigger: Table 1 of this note Sec. 2.5 describes the typical level-0 trigger algorithm requirements and the 3D-Flow features aimed to provide full programmability for 2x2, 3x3, 4x4, 5x5, etc.

LAL proposal implements only 2x2 trigger algorithm, Bologna group proposes only 3x3 trigger algorithm, any small change of the requirements (as it is the case from HERA_B to LHCb requires to redesign completely the system, even if the 3x3 is the same basic function used in both). LAL proposal, besides not having the programmability, it does not have the connections to two ports. See also reference notes LHCb 96-1, TRIG 96-2, LHCb 99-006, IEEE Trans. Nucl. Sc. 43. 170(1996), the LHCb LOI, the LHCb TP, and the papers referenced in the above documents.

- Allowing changes of the trigger word, and allowing to add subdetectors at level-0. Due to budget limitation, the LHCb experiment in some cases has to consider staging the construction of the detector. If in the future it will be desired to build an additional subdetector for increasing the performance at level-0 trigger, or if it would be desired to use some information from the vertex detector (or to increase the information from the existing detectors) for the level-0 trigger, the

3D-Flow allows not only to redefine the trigger-word, but also to add bits from the current 23-bit up to 32-bit (see this note, Sec. 6.1, 6.2, 6.3.5, and LHCb 99-006 Sec. 5.2.2) The insertion of new bits and the formatting of the new trigger word will be done in the 3D-Flow front-end FPGA, the algorithm change will be done in the 3D-Flow processor).

- No boundary limitation. Each 3D-Flow processors sends the data to 8 neighbors (or 24 neighbors, etc.) and receives the information from 8 neighbors (or 24 neighbors, etc.), regardless if the processor is on a chip, on a board, or off-board (see this note Sec. 6.1, 6.2, 6.4, 7.1, 7.2, 7.4).

APPENDIX B: FRONT-END ELECTRONICS FOR THE PRESHOWER DETECTOR. (Answers to questions by L0 trigger group).

LHCb TP describes at Chapter 10.2.2. "... This method allows the photodetectors and readout electronics to be sited away from the beam, thus reducing the radiation dose, if necessary." In Chapter 10.2.3. of the TP is reported "The baseline solution for reading out the fibers uses multi-anode photomultiplier tubes (PMT's). The 16-pixel Hamamatsu H6568 is suited for coupling the 16-fibre bundle of a detector channel to a single pixel the readout electronics will be the same as for the calorimeters, but an ADC with only 8 bits can be used resulting in a slight cost reduction of the chain). The electronics can be placed closed to the PMT. If necessary, the PMT's and the electronics could be placed as far from the beam axis as necessary by making the fibers longer..."

This last sentence of the TP, suggests that it is more economical and desirable to have all analog electronics on the PMTs and to send out from each channel only 1-bit digital information @ 40 MHz to the trigger electronics, and 8-bit digital information @ 1 MHz to the DAQ and higher level trigger.

There is a discrepancy in regard to the resolution of the ADC conversion from the preshower analog signals from the LAL - Orsay group and the LHCb TP (8-bit in the TP, and 10-bit in the note LHCb 99-007). The present, short description of the 3D-Flow front-end electronics for preshower detector design is flexible to accommodate both resolutions with a small cost difference (mainly the difference in cost of the ADC, while one could use the same FPGA and serializer/transmitter).

This solution foresees accommodating all analog front-end electronics in a 9 cm x 9 cm printed circuit board at the base of each 16-pixels multi-anode photomultiplier. By using 375 multi-anode photomultipliers, the 6000 preshower electronic channels should be accommodated in a total of 3 m² of printed circuit board.

The FE-preshower PCB will have at the input the following signals: 16 analog inputs from the PMT, a 40 MHz clock, one EnInData, a reset, a global level-0 accept/reject, and an EnOutData.

The output signals could be the following: one LVDS digital serial line @ 152 Mbps to the DAQ and higher level triggers (calculated as 16 x 8-bit + 24-bit of event-ID and bunch crossing-ID), two LVDS digital serial lines each @ 320 Mbps, each carrying 8 x 8-bit @ 40 MHz to the level-0 trigger electronics, one start_burst_out signal, three FIFO out_status_flag signals.

All output signals are generated by an OR3T55 FPGA which contains the same circuit as the one described in the note LHCb 99-006 in Sections 5.2.1, 5.2.3, and 5.2.4.

The estimated size of 9 cm x 9 cm of the PCB is accurately estimated with safely margin in the following manner:

16 x ST-44 package ADC (or equivalent) = 19.2 cm², an equivalent area for the preamplifiers, LVDS serializers (e.g., one DS90CR215, or three DS92LV1212) = 2.10 cm², one FPGA OR3T55-352-pin, 35 mm x 35 mm = 12.25 cm², one FPGA configuration EPROM = 1 cm², one connector for 17-conductor coaxial ribbon cable = 5 cm², two RJ-45 connectors = 2.5 cm², one connector for control signals = 0.6 cm² (the latter area for connectors must be doubled because they occupy both sides of the PCB) for a total of 69.95 cm².

Considering that two surfaces of a 9 cm x 9 cm is equivalent to 162 cm², the estimate to use only 69.95 cm² for components, leaves a safe margin for the connections and via.

APPENDIX C: HIGH SPEED BACKPLANES. (Answers to questions by L0 trigger group).

Even if at first it seems an impressive number of connections on the 3D-Flow crate backplane (1328-pin per connector x 16 connectors = 21248 connections @ 400 Mbps), the concerns should however be eliminated after the explanation that among these 21248 pins, 768 are power and ground which do not require traces, 6400 are connections to the neighboring crates which do not require traces. One set of 2240 short segment traces (about 5 cm long) in one layer, non intersecting each other will connect two pins of two even connectors. Another set of 2240 short segment traces on another layer is connecting pins of odd connectors, and a third set of 2560 non intersecting traces on another layer are connecting pins of adjacent connectors.

The apparently at first complicated problem, turns out to be simpler and solvable with 6 layers PCB. The reason for the simplicity is the original approach of the 3D-Flow that simplifies things in many areas, another architecture that would have to route traces connecting 20480 pins @ 400 Mbps in a "spaghetti" fashion, would have required maybe more than 30 layers.

The other two implementations (LAL and BO) which do not have a regular pattern of short distance traces, but rather have (as described in LHCb 99-007, Sec. 3.1.1.) short and long traces (since the summary cards are located in the center of the group of 8 front-end cards to which they are connected), will require to build a more complicated backplane. At a given speed, the 3D-Flow backplane will be of simpler construction because of the regular pattern evenly distributed over the entire surface of the backplane, while the other implementation will have area, close to the summary cards, with a higher density of connections.

The following describes some consideration that will show feasibility of the specific backplane repoted in figure 27. The border connectors have less pins connected to traces because signals go to neighboring crates through cables.

Let us consider the backplane connector 19 of figure 27. The total number of pins with signals on this connector is 1280. Of these,

- 320 do not require traces because they go to the north crate through cables receiving the signals through the feedthrough AMP 646372-1 connector;
- 320 traces go to connector 20 to the right. The above traces can be accommodated in layer 1 (layer 2 being ground);
- 320 traces go to connector number 21 to the right (assume two traces going in between two rows of pins of connector 20. The distance of the rows of all connectors is 2 mm, each backplane connector has 166 rows);
- 320 traces go to connector 17 to the left (assume two traces going in between two rows of pins of connector 18). The above traces can be accommodated in layer 3 (layer 4 being ground).

The connector 20 (see Figure 27) has:

- 320 traces connected to connector 19 on layer 1 as mentioned above;
- 320 connections go to the south crate which does not require traces;
- 320 traces go to connector number 22 to the right (assume two traces going in between two rows of pins of connector 21);
- 320 traces go to connector 18 to the left (assume two traces going in between two rows of pins of connector 19). The above traces can be accommodated in layer 5 (layer 6 being ground).

This pattern in six layers per pair of connectors will be repeated for the 8 blocks of the entire backplane.

An appropriate signal assignment to the pins of the connector (such as all signals going to the connector to the right are in the right columns of the connectors, same for the left) will avoid via and crossing signals.

A back-of-the-envelop calculation (simulation tools will provide us precise results) will tell us that a 0.3 mm wide trace, at a distance of 0.3 mm from a ground plane with the appropriate dielectric will provide an impedance of about 75 ohms.

There will be a possibility to draw traces with a controlled impedance at our choice of a value in the range of 50 to 100 ohms. Increasing trace width or decreasing trace distance to the ground layer will lower the impedance.

The entire backplane could be built in only six layers for speed up to 400 Mbps using two traces of about 0.1 mm width in between the 2 mm rows of the connector (VCC can be accommodated on signals layers). This implementation will have for each layer 320 traces of about 5 cm in length, laid side-by-side on a total width of about 32 cm (the width of the backplane is about 36 cm).

In the event higher speed are required (currently are under development backplanes up to 2.5 Gbps), only one trace should be used in between two rows of connector pins. (The board will have no more than 12 layer in total). In this case, the implementation will have for each layer only 160 traces of about 5 cm in length, laid side-by-side on a total width of 320 mm.

References

¹ Eisenhandler, E., Hardware Triggers at the LHC. Fourth Workshop on Electronics for LHC Experiments, Rome – Italy September 21-25, 1998, pp. 47-56.

² Crosetto, D. "High-Speed, Parallel, Pipelined, Processor Architecture for Front-End Electronics, and Method of Use Thereof." LHCb 96-2, TRIG 96-1

³ Crosetto, D. Massively Parallel-Processing System with 3D-Flow Processors." Published by IEEE Computer Society. 0-81816-6322-7194, pp. 355-369.

⁴ <http://www.rapid.org> (see also ../aboutrapid.html for IP information)

- ⁵ <http://rapid.rapid@catalog.rapid.org/scripts/isynch.dll?panel=TclScript&file=sipcTop.tcl> (for Virtual Components catalog)
- ⁶ <http://www.lsilogic.com/products/PRchart.html> and [../unit5_2.html](http://www.lsilogic.com/unit5_2.html).
- ⁷ <http://www.lhc01.cern.ch> (Large Hadron Collider Project at CERN, Geneva – Switzerland)
- ⁸ Crosetto, D. "Programmable Level-1 Trigger with 3D-Flow array," Computing in HEP, San Francisco, CA, 21-27 April 1994. Editor: S.C. Loren, pp 57-61
- ⁹ S. Conetti, D. Crosetto, "The LHCb Calorimeter Trigger and its implementation with the 3D-Flow system," LHCb 98-13
- ¹⁰ G. Corti, B. Cox, and D. Crosetto, "An Implementation of the L0 Muon Trigger Using the 3D-Flow," note LHCb 98-001
- ¹¹ LHCb Technical Proposal, CERN/LHCC 98-4 20 February 1998.
- ¹² <http://lhcb.cern.ch/notes/98-058.ps> The LAL group, An alternative high PT electron and hadron trigger for LHCb, LHCb 99-007, TRIG. And transparencies from the February 1999 LHCb collaboration meeting. LHCb level-0 trigger hardware implementation in 20 crates (9U): 14 crates FE-logic for ECAL, 4 crates for HCAL, 1 crate for ECAL Selection L-0 trigger, and 1 crate for HCAL Selection L-0.
- ¹³ Marco Bruschi, INFN - Bologna. A modified version of the Hera-B Level-0 calorimeter trigger. Transparencies from the LHCb collaboration meeting on 16 February 1999. Note LHCb 6 May 99 level-0 trigger implementation in 14 crates FE-logic for ECAL from LAL, 4 crates FE-logic for HCAL from LAL, 22 crates from Table 7.
- ¹⁴ <http://atlasinfo.cern.ch/atlas/groups/daqtrig/tdr/tdr.html> (The Atlas Level-1 Trigger Technical Design Report)
- ¹⁵ <http://www.analog.com/pdf/ad9042.pdf>
- ¹⁶ G. Mirabelli, et al., "A GaAs Transceiver Chip for Optical Data Transmission," IEEE Transaction on Nuclear Science, Vol. 40, N. 4, August 1993, pp. 713-716.
- ¹⁷ <http://www.hp.com/hp-comp>
- ¹⁸ <http://www.amcc.com/pdfs/S2064.pdf>
- ¹⁹ <http://www.vitesse.com/>
- ²⁰ <http://www.amcc.com/pdfs/S3044.pdf>
- ²¹ <http://www.lucent.com/micro/opto/document.html>
- ²² <http://www.lucent.com/opticalnet>
- ²³ <http://www1.nortelnetworks.com/technology/opto/home.html>
- ²⁴ <http://www.hp.com/hp-comp/promo/sff.html>
- ²⁵ <http://www.sel-rtp.com/products> then select: Electro-Optic Products, then select: Telecommunication
- ²⁶ <http://www.methode.com/fopusr/datasheets/808mp.pdf> (see series MP with 2-12 fibers per connector)
- ²⁷ <http://www.amcc.com/pdfs/S3043.pdf>
- ²⁸ <http://www.national.com/an/AN/AN-1040.pdf> (application on LVDS Bit Error Rate measurements)
- ²⁹ <http://www.amp.com> or AMP catalog number 82158, page 5 and page 12
- ³⁰ <http://www.amp.com> or AMP catalog number 65911 page 14
- ³¹ <http://www.lucent.com/micro/fpga/orcapdfs/DS98-163-1.pdf>
- ³² http://cmsdoc.cern.ch/doc/notes/docs/NOTE1998_074 W. Smith, et al. "CMS Calorimeter Level-1 Regional Trigger Conceptual Design." CMS NOTE-1998/074
- ³³ Crosetto, D., LHCb 99-006, TRIG, 30 March 1999. http://lhcb.cern.ch/electronics/simulation/3dflow_FE/lhcb99_006.pdf
- ³⁴ <http://www.amp.com> or AMP catalog number 65911 page 12.
- ³⁵ <http://www.amp.com> or AMP catalog number 82066 page 13
- ³⁶ <http://www.amp.com> or AMP catalog number 82003 page 45
- ³⁷ <http://www.comtrol.com/sales/specs/rocket.htm#Specs> PCI and ISA 32 x RS-232 ports for IBM PC.
- ³⁸ C. Beigbeder, et al., An Update of the 2x2 Implementation for the Level 0 Calorimeter Triggers. LHCb 99-007 29 April 1999
- ³⁹ V. Alberico, et al. The HERA-B electromagnetic pre-trigger and its possible adaptations to the LHC-B Level-0 calorimeter trigger. LHCb 98-034, 10 February 1998.
- ⁴⁰ The LHCb Bologna Group. Proposal for a Level 0 calorimeter trigger system for LHCb. LHCb note 6 May 1999.