

Implementing the L1 trigger path

R. Jacobsson

ABSTRACT

This note discusses two issues which have emerged recently related to the transmission of data to the DAQ system by the L1 Front-End electronics and the implementation of the L1 trigger path. Firstly the note describes a scheme to centrally broadcast the IP destinations for the L1 data channel and the High-Level Trigger channel. The scheme is already supported by the TFC system, in particular by the current version of the Readout Supervisor “ODIN”.

Secondly, the note discusses the consequences of eliminating the L1 Trigger Sorter module located between the L1 trigger processing in the CPU farm and the TFC system. Two possible implementations are described: sorting the L1 triggers in the Readout Supervisor “ODIN” or handling unsorted L1 triggers fully. The two solutions require a redesign of the Readout Supervisor which implies a delay of about one year with respect to the current planning. However, the current version can still be used for tests and in the experiment provided a provisory L1 Trigger Sorted module exist.

LHCb Technical Note

Issue: LHCb 2003 – 80
Reference: LHCb Online
Created: 18 July 2003
Last modified: 22 August 2003

Prepared By: R. Jacobsson, CERN, Geneva, Switzerland
LHCb Online Group

INDEX

1	INTRODUCTION	3
2	CURRENT USE OF THE TTC BROADCASTS.....	5
3	IP DESTINATION BROADCASTING.....	5
4	ELIMINATING THE L1 TRIGGER SORTER MODULE	6
4.1	L1 TRIGGER DECISION PACKET	6
4.2	SORTING IN THE READOUT SUPERVISOR.....	8
4.3	HANDLING UNSORTED TRIGGERS FULLY	10
5	CONCLUSION	12
6	REFERENCES	13

1 Introduction

Figure 1 shows a simplified picture of the LHCb readout system [1]. The Level-1 Front-End electronics (L1 FE) [2] has two channels to the event building network, one of which is used to transmit event data to the L1 trigger processing which takes place in the CPU farm, and the other which is used for the complete readout after the L1 trigger decision [3]. Thus, upon receiving event data from the Level-0 Front-End electronics (L0 FE) after a positive L0 trigger decision, the L1 FE electronics sends a part of the data over the event building network for the L1 trigger processing and buffers the complete event data during the L1 trigger latency (currently about 58 ms). At this level the events are identified by a L0 Event ID, which is simply the number of L0 trigger accepts. Finally upon receiving a positive L1 trigger decision the L1 FE electronics de-randomize the events, zero-suppress the data and send the complete event data to the CPU farm for the High Level Trigger (HLT) processing. In order to improve the performance of the readout network and the event building in the Sub-Farm Controllers (SFC), the L1 FE electronics pack the data fragments for several consecutive events into Multi-Event Packets (MEP) for both the L1 and the HLT channel [4]. The data transport format is based on the Internet Protocol (IP). A maximum packing factor of 32 and 16 events have been defined for the L1 and the HLT channel, respectively.

The Timing and Fast Control (TFC) system [5], shown in detail in Figure 2, drives the readout by distributing synchronously timing, the two levels of trigger decisions and control information to the Front-End electronics via the TFC distribution network. The entire TFC mastership is located in the Readout Supervisor “ODIN” [6].

Since the L1 trigger processing time varies, the L1 trigger decisions will come out of order with respect to the L0 Event ID. A L1 trigger sorter module [7] ensures that the L1 triggers are handled by the TFC system and distributed to the L1 FE electronics in order.

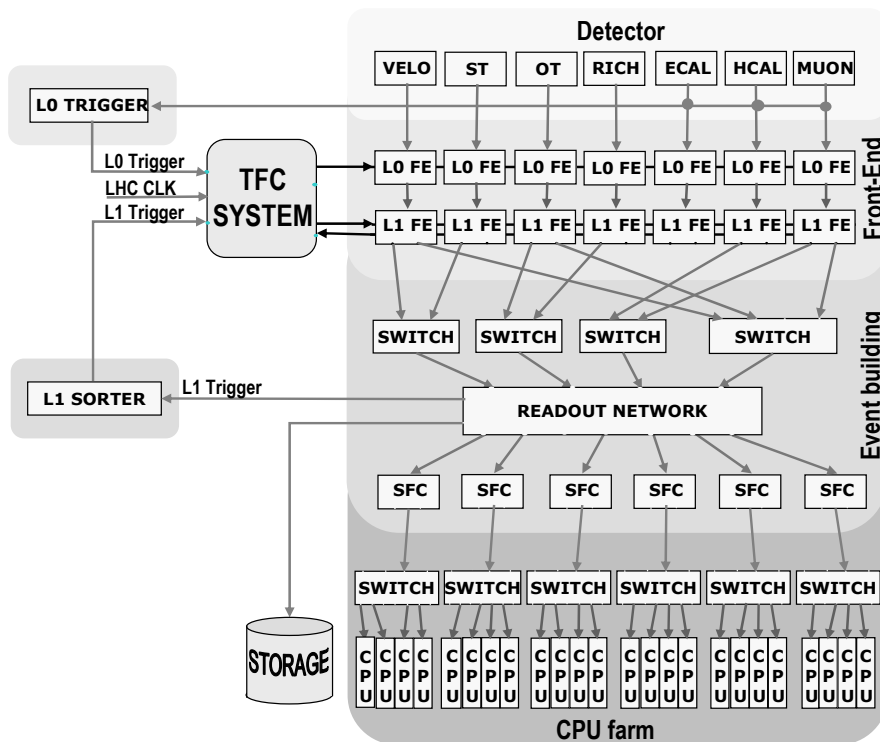


Figure 1: Simplified picture of the LHCb readout system architecture.

The architecture requires that all the L1 FE boards assign the same IP destination to the Multi-Event Packets containing the same events. In order to control the assignment centrally this report describes a scheme to broadcast synchronously the destination for the two channels from the Readout Supervisor “ODIN” using the standard TFC distribution network. The scheme has the advantage that it imposes a certain level of concurrency in the Front-End electronics and that it allows rapid update of the IP destination table in case of local breakdowns in the CPU farm. The IP destination assignment scheme requires no hardware modifications to the TFC system and the function has been implemented in the Readout Supervisor. The only side effect is an increased load on the TFC distribution network since the L1 trigger decisions and the synchronous control commands have to share the bandwidth with the IP destination broadcasts.

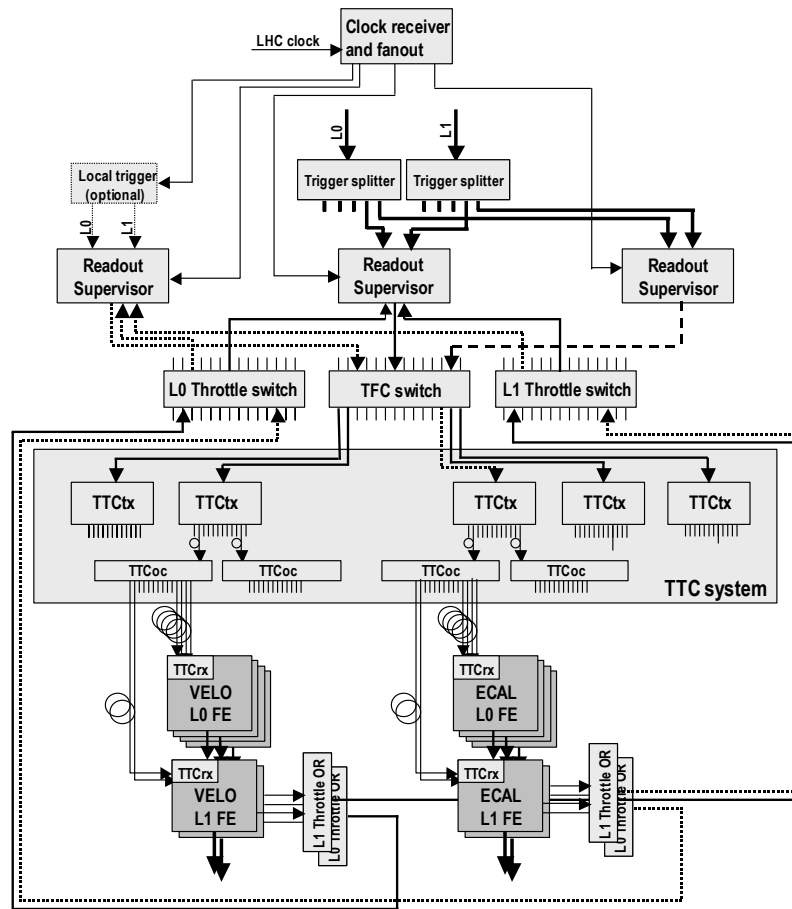


Figure 2 : Overview of the TFC architecture.

It has also been suggested that the L1 trigger sorter module could be eliminated, implying that the Readout Supervisor receives unsorted L1 trigger decisions. This has several consequences for the TFC system. This report also describes a scheme to accommodate this architecture. Two solutions are envisaged depending on the capability of the L1 FE electronics:

- Implement decision sorter at the L1 trigger input stage of the Readout Supervisor “ODIN” and thus maintain ordered L1 triggers on the TFC distribution network.
- Run fully unsorted.

Whereas the first solution could be handled by adding a function hardware block to the “final prototype” of the Readout Supervisor currently in production [8], the second solution demands for a major redesign of the RS hardware.

2 Current use of the TTC broadcasts

The TFC distribution network consists of the TFC Switch and the standard CERN Timing, Trigger and Control (TTC) system developed in the CERN RD12 project [9].

The TTC distribution network is based on fiber optics carrying two communication channels: a low-latency accept/reject signal (channel A) and framed and formatted broadcasts including Hamming code (channel B). Two types of broadcasts are available: 16 bit frames which have 8 bits of user information, so called short broadcasts, and 42 bit frames which have 16 bits of user information (8 bit data/8 bit address), so called long broadcasts. The short broadcasts and the long broadcasts take 400 ns and 1050 ns, respectively, to transmit.

In LHCb channel A is used to transmit the L0 trigger decision at 40 MHz. Channel B is used to transmit:

- The L1 trigger accepts and rejects.
- Commands resetting event related counters in the Front-End electronics used to identify the accepted events and to check synchronisation.
- Commands resetting the Front-End electronics in order to prepare it for data taking or to recover from an error condition.
- Calibration commands activating specific calibration systems in the Front-End electronics or in the sub-detectors.

The commands are encoded using only the short broadcast format as summarized in Table 1 [10].

Table 1 : Summary of the short broadcasts.

	7	6	5	4	3	2	1	0
L1 trigger	1	Trigger type			L0 EvID		0	0
Reset	0	1	R	L1 EvID	L1 FE	L0 FE	ECR	BCR
Calibration	0	0	0	1	Pulse type		0	0
Command	0	0	1	Command type			0	0

3 IP destination broadcasting

The IP transport format used for the L1 and the HLT channel requires the 48-bit Ethernet destination address and the 32-bit IP destination address. In order to only broadcast ten bits of destination address allowing up to 1024 different destinations, the SFCs' will all have the same Ethernet base address and IP base address consisting of the 38 and the 22 most significant bits, respectively. The base addresses could also be different for the L1 and the HLT channels. Each L1 FE electronics board stores the base addresses in programmable registers [4].

The Readout Supervisor broadcasts the IP destinations using the long broadcast format of the TTC system (Table 2). Bit [15..14] = "10" identify the destination broadcasts and bit 13 distinguishes between a HLT and a L1 destination broadcast. The RS incorporates a lookup table of 10-bit

destinations and a broadcast state machine for each data channel. The two state machines count the number of L0 accepts and L1 accepts, respectively, and transmit a destination broadcast every n trigger accepts, where n is a programmable interval corresponding to the Multi-Event Packet packing factor for each data channel. The destination is retrieved by stepping through the lookup table in order to assign the MEPs to the SFCs in a round-robin manner. In order to implement a simple load balancing scheme to take into account possible differences in the number of CPU associated with each SFC, the destination table has a depth equal to the total number of CPUs and contains repeatedly the same destination address for each SFC in proportion to its number of CPUs.

Stopping the data recording is carried out by permanently inhibiting the L0 trigger. In order to read out all events pending in the system, bit 12 in the destination broadcast defines a flush command. Even if the L1 FE board has not received enough events to fill an MEP with the predefined number of events, the packet is transmitted upon receiving a flush command. The L1 flush command will be sent with a long delay following the last L0 trigger accept to ensure that the events have been read out of the L0 de-randomizer. The HLT flush command is transmitted after yet another delay.

Table 2 : Summary of the long broadcasts for the IP destination assignments.

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
L1 IP destination	1	0	0	Flush	R	R	Ethernet/IP address									
HLT IP destination	1	0	1	Flush	R	R	Ethernet/IP address									

Table 3 shows the load relative to the TTC bandwidth of each broadcast type. The contribution from the destination broadcasts for the HLT channel assumes an HLT MEP packing factor of one. Thus, given the TTC bandwidth, the minimum event packing factor for the L1 channel is three, leading to a total load of 86%.

Table 3 : Relative load on channel B of the TTC network adding the IP destination broadcasts.

Broadcast	Rate	Time/broadcasts	TTC load
Short commands (BCR/ECR, calibrations, other)	~4/LHC turn = 44.8 kHz	16 cycles = 400 ns	1.8 %
L1 trigger accepts, short broadcasts	40 kHz	16 cycles = 400 ns	1.6 %
L1 trigger rejects, short broadcasts	1 MHz	16 cycles = 400 ns	40 %
HLT destinations, long broadcast	Packing = 1 → ~40 kHz	42 cycles = 1.05 μs	4.2 %
Total			47.6 %
L1 destinations, long broadcasts	Packing ≥ 3 → ~330 kHz	42 cycles = 1.05 μs	34.6 %
Total			82.2 %

4 Eliminating the L1 Trigger Sorter Module

As mentioned above, eliminating the L1 trigger sorting performed in a separate module before the TFC system leads to two possibilities: performing the sorting in the Readout Supervisor or handling unsorted triggers fully in the entire system. The following sections outline a possible implementation for each scheme.

4.1 L1 Trigger Decision Packet

The Readout Supervisor will receive the L1 trigger decisions via Gigabit Ethernet as IP packets using the Gigabit Ethernet mezzanine card developed in LHCb [11]. There are two possibilities for the transport format:

- Sending one L1 trigger decision per IP packet.
- Packing the L1 trigger decision into Multi-Decision Packets (MDP) in a similar way to the Multi-Event Packets on the L1 and HLT data channel.

There is no particular preference with respect to the implementation of the data handling in the Readout Supervisor. The second has the advantage of reducing the network load and the MDPs would naturally contain all the decisions corresponding to the L1 Multi-Event Packets. However, it requires some computing power from the SFC. It would have to do a simple pre-sorting by reserving a fixed memory space, slot in the decisions as they come from the processing in the CPUs and, upon receiving the last decision of the packet, prepare a header and send it. In addition, the SFC would have to incorporate a timeout mechanism to handle cases in which a decision is not received from a CPU. Since the Readout Supervisor will always incorporate a timeout mechanism, it will handle missing decision in the first case where only one decision is sent per IP packet.

Table 4 : Single-Decision Packet format. The eight bytes of preamble are not shown.

32-bit word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Ethernet + IP header (see [4])			
...				
8				
9	L0 Event ID			
10	Error and status block			L1 decision
11	L1 decision information			
...				
16				
17	CRC			

Table 5 : Multi-Decision Packet format. The eight bytes of preamble are not shown.

Decision	32-bit word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
	0	Ethernet + IP header (see [4])			
	...				
	8				
	9	MDP Error block			# of events
0	10	L0 Event ID			
	11	Error and status block			L1 decision
	12	L1 decision information			
	...				
	17				
1	18	L0 Event ID			
	19	Error and status block			L1 decision
	20	L1 decision information			
	...				
	25				
N	N * 8 + 10	L0 Event ID			
	N * 8 + 11	Error and status block			L1 decision
	N * 8 + 12	L1 decision information			
	...				
	N * 8 + 17				
	N * 8 + 18	CRC			

Table 4 and 5 shows two possible formats for the Single-Decision Packet (SDP) and the Multi Decision Packet. The formats are analogous to the format of the Multi Event Packets. The L1

trigger data including the L0 event ID is proposed to be 32 bytes, which includes 24 bytes of L1 decision information. With the current L1 latency (~58 ms), it allows using a 2 Mb memory for the sorting of the L1 triggers, or a 4Mb memory for a doubled latency.

Although it would be enough to include the L0 Event ID once for the first event in the Multi-Decision Packet, since the packet contains consecutive decisions, a L0 Event ID before each decision is more generalized and simplifies slightly the processing in the Readout Supervisor. In order to handle missing or late decisions in one location, i.e. in the Readout Supervisor, the SFC should on a timeout, flag the error in the MDP error block and in the error block of the empty decision and transmit the MDP with a zero L0 Event ID and all zeros in the decision information of the missing decision.

4.2 Sorting in the Readout Supervisor

Figure 3 shows a block diagram of a redesigned Readout Supervisor which includes a function block performing the L1 trigger sorting (Q_L1). Only the functions related to the handling of the L1 trigger and the Readout Supervisor Front-End are discussed here. The implementation of the other functions remains largely untouched. See References [6] and [8] for a full description of the Readout Supervisor functionality.




Figure 3 : Simple block diagram showing the function blocks and the data flow of a Readout Supervisor that supports unsorted L1 triggers.

The L0 Accept Buffer (L0 ABUF) is based on a RAM and stores information at each L0 trigger accept which is needed to process the L1 triggers. It should contain the L0 Event ID of the event, a 3- or 4-bit trigger type and a force bit. The L0 Event ID is used to cross-check the incoming L1 trigger decision and to know which the next trigger decision is to broadcast to the L1 Front-End electronics. The trigger type specifies the source of the trigger and it is transmitted to the L1 Front-End electronics in the L1 decision broadcast to allow the Front-End electronics to react differently to different types of triggers. The force bit indicates that the trigger has been forced at L0 and that consequently the L1 trigger decision should be made positive, irrespective of the L1 physics trigger decision. This ensures that calibration triggers etc are accepted.

A L1 input state machine in the Q_L1 block reads the L1 trigger decision packets from the Gigabit Ethernet mezzanine card (GbE RX) via a 32-bit 80 MHz bus. Assuming the L1 triggers are received

as Single Decision Packets, which is the worst case, implies a data load of 576 Mb/s. This is well below the bandwidth of the bus. The state machine drops the ten 32-bit words of Ethernet/IP protocol and writes the eight 32-bit words of L1 trigger data into the L1 Decision Buffer (L1 DBUF) via a 40 MHz 32-bit bus using the L0 Event ID as the address.

A second state machine compares the L0 Event ID of the incoming L1 triggers with the L0 Event ID of the next trigger to broadcast stored in the L0 Accept Buffer. Upon receiving the next L1 trigger to broadcast, the state machine prepares the broadcast frame, transmits it together with a broadcast request to the TTC broadcaster in the Q_MP module. The state machine subsequently looks for the next L1 trigger in the L1 Decision Buffer by comparing the L0 Event IDs. If it has been received the broadcast frame is prepared, otherwise the state machine waits for the L1 trigger to arrive on the network. The state machine also ensures that the L1 trigger accept broadcasts are spaced by a minimum of 20 μ s and that the first L1 trigger reject broadcast after a L1 accept is transmitted after 900ns and otherwise at intervals of 400 ns.

The maximum latency of the total L1 trigger path including processing has been set to 58 ms. The Readout Supervisor must guarantee that this is not exceeded due to a missing L1 trigger. Since the maximum input rate to the L1 buffer is fixed by the readout time of the L0 Derandomizer in the L0 FE electronics, the L1 buffer is constantly closed to full at maximum rate. The Readout Supervisor can thus detect a missing L1 trigger decision by monitoring the occupancy of the L1 buffer in the L1 FE electronics. This is carried out by an emulator running in the Q_L1 module. The same could also be achieved by monitoring the occupancy of the L1 Accept Buffer. An excess above a certain level which is defined with a safe margin indicates that a triggers decision is missing. The Readout Supervisor then uses the information in the L0 Accept Buffer to prepare a fake L1 reject broadcast. In case several consecutive triggers are missing, the output state machine will continue sending fake rejects. In this case, an error bit is set which indicates that a global reset by the Experiment Control System is needed. Clearly at a lower rate the missing trigger would be detected later, which in practice means that the scheme automatically allows for a longer latency at lower rate. It only guarantees that the L1 buffer never overflows.

It may also happen that the missing trigger arrives late. This is detected by checking the trigger contiguity using a 32-bit L0 event ID. Since a fake trigger has already been generated in place of the late trigger, the Readout Supervisor discards the trigger and counts the occurrence.

The L1 throttle is also applied in the Q_L1 module converting L1 trigger accepts into rejects to prevent buffer overflows in the system.

A first version of the Q_L1 code performing the functionality described above has been written and simulated [12]. Receiving the L1 trigger decisions as Multi-Decision Packets only involves a small modification of the L1 input state machine.

The Readout Supervisor should prepare a RS data block which contains event identifiers, type of trigger, GPS time, the L1 decision information etc. Data related to the L0 trigger is sampled at the L0 trigger accepts and data related to the L1 trigger is added at the L1 trigger accepts. The block is appended to the event data like the data block of any detector front-end. On the current version of the Readout Supervisor the RS Front-End is located on the board. Although the design of the RS FE on the current version of the Readout Supervisor would work in this scheme, it would be beneficial to take advantage of the common L1 FE board [13] as an RS FE in a redesign. It could be achieved by implementing two serial high-speed data links similar to the links used by the L0 FE to transfer data to the L1 FE, one for the L0 data (L0 LINK) and a second for the L1 data (L1 LINK). Two

developments are available: using the same design as the L0 FE based on the CERN Gigabit Optical Link (GOL)[14] or using the HOLA mezzanine card[15] developed by Atlas. The HOLA card conforms to the SLINK standard and is based on the same TLK2501 transceiver used as receiver on the optical receiver cards of the common L1 FE boards. The L0 data would be prepared immediately upon a L0 trigger accept by the Q_L0FE block, transmitted via the serial link and be stored in the L1 Buffer (L1B) in the common L1 FE board (Figure 4). Upon broadcasting a L1 trigger accept, the Q_L1 block would transmit the L1 data, including the L1 decision information, over the second link. The Pre-Processor FPGA (PP-FPGA) on the common L1 FE board would store the L1 data temporarily while waiting for the corresponding L1 trigger accept. Upon receiving the trigger accept, the PP-FPGA would read out the L0 data from the L1 Buffer, concatenate it with the L1 data, and send it to the SyncLink-FPGA via the HLT bus. The SyncLink-FPGA would pack the RS event data in the same manner as all other front-ends and transmit it via the HLT channel to the event building.

In this scheme the short broadcast format in Table 1 can still be used to transmit both L1 trigger accepts and L1 trigger rejects.




Figure 4 : Simple block diagram showing the dataflow in the common L1 FE board. Making use of the common L1 FE board as a Readout Supervisor Front-End would only require moderate changes in the standard firmware.

4.3 Handling unsorted triggers fully

The most obvious consequence of handling unsorted triggers is that the implementation of the L1 buffer in the Front-End electronics must be based on a RAM from which the events can be read out using the L0 Event ID as address. This implies that the L1 trigger broadcast must contain the L0 Event ID, at least wide enough to uniquely identify all events in the system at any given time. A L0 Event ID of 20 bits would give a safe margin. Together with a 4-bit trigger type, the L1 broadcast could consist of two consecutive long format broadcasts as defined in Table 6. However, the bandwidth of the TTC channel B implies that only L1 trigger accept broadcasts can be transmitted. L1 trigger rejects are handled by implementing a circular writing to the L1 buffer in the L1 Front-

End. The events that have not been read out are thus automatically overwritten. The scheme has the drawback that it reduces the chances of detecting L1 errors in the L1 Front-End electronics.

Although the TFC system does not allow transmitting both the L1 trigger accepts and the rejects to the Front-End, the Readout Supervisor should receive both from the L1 trigger processing in order to be able to check the synchronism of the L1 trigger processing and the L1 path up to the Readout Supervisor.

Table 7 shows the load relative to the TTC bandwidth of each broadcast type in this scheme. The contribution from the destination broadcast for the HLT channel assumes a Multi-Event Packet packing factor of one. Therefore, given the TTC bandwidth, the minimum event packing factor for the L1 channel is two, leading to a total load of 67%.

Table 6 : Handling unsorted L1 triggers fully require sending two consecutive long broadcasts for the L1 trigger accepts.

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
L1 trigger (1)	0	1	0	R	L0 EvID(11 downto 0)												
L1 trigger (2)	0	1	1	R	Trigger type		R	L0 EvID(19 downto 12)									

Table 7 : Relative load on channel B of the TTC network transmitting IP destination broadcasts and only L1 trigger accepts as long broadcasts.

Broadcast	Rate	Time/broadcasts	TTC load
Short commands (BCR/ECR, calibrations, other)	~4/LHC turn = 44.8 kHz	16 cycles = 400 ns	1.8 %
L1 trigger accepts, two long broadcasts	40 kHz	2 x 42 cycles = 2.1 μs	8.4 %
HLT destinations, long broadcast	Packing = 1 → ~40 kHz	42 cycles = 1.05 μs	4.2 %
Total			14.4 %
L1 destinations, long broadcasts	packing >= 2 → < ~500 kHz	42 cycles = 1.05 μs	52.5 %
Total			67 %

The hardware implementation of the Readout Supervisor is the same as shown in Figure 3. A L1 input state machine in the Q_L1 block reads the L1 trigger decision packets from the Gigabit Ethernet mezzanine card (GbE RX). The incoming L0 Event ID is used to retrieve the L0 trigger information stored in the L0 Accept Buffer (L0 ABUF). In order to de-randomize the L1 trigger accepts and transmit them with a minimum spacing of 20 μs, the L1 triggers accepts as given by either the incoming decision bit or the force bit in the L0 Accept Buffer are stored in the L0 Decision Buffer (L1 DBUF) in a FIFO-like manner.

A second state machine reads out the first pending event in the L1 Decision Buffer (L1 DBUF), prepares the broadcast frame and transmits it together with a broadcast request to the TTC broadcaster in the Q_MP module. As described above, the state machine also transmits the L1 data including the L1 decision information to the common L1 FE board.

A drawback of this scheme is that it is complicated to implement a scheme to detect missing or late triggers (a single SDP or an entire MDP). Although a missing trigger would simply be treated as a reject and not cause any problems to continuous running, it is imperative that they are detected in the Readout Supervisor. Any simple detection mechanism involves storing the L0 Event IDs in a sorted fashion. Implementing a detection mechanism without large storage and sorting can be done by increasing the L0 latency with a few cycles in the Readout Supervisor. The Q_L0 module which is writing the L0 trigger information into the L0 Accept Buffer on every L0 trigger accept could detect that an entry has not been readout by the L1 trigger handling. Before overwriting the contents of the next L0 trigger it could force the L1 handling to read out the trigger and handle the error condition.

5 Conclusion

Assuming that the baseline online architecture includes a L1 Trigger Sorter module, the current version of the Readout Supervisor is the complete final prototype. It supports the scheme of transmitting at regular interval the L1 and the HLT destination broadcasts and it includes the RS Front-End on-board.

Supporting an architecture without a L1 Trigger Sorter module requires redesigning the Readout Supervisor. The same hardware design could support both on-board sorting and handling unsorted triggers fully. Whereas sorting on-board provides the same level of performance as with the current architecture, handling unsorted L1 trigger fully reduces the error checking capabilities in the L1 FE and requires increasing the internal L0 latency of the Readout Supervisor. Although the latter must be verified, it should not pose a problem with the current L0 latency budget. If the Readout Supervisor has to be redesigned it would also be advantageous to make use of the common L1 FE board for the RS Front-End.

With respect to the schedule, redesigning the Readout Supervisor would imply a delay of about one year. However, if the delay is acceptable, the current version of the Readout Supervisor is still sufficient for testing and even for starting the experiment, provided a provisory L1 Trigger Sorter module is implemented.

Table 8 : Summary of the long broadcasts. It is not decided yet whether the L1 triggers are transmitted as long or short broadcasts.

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
L1 IP destination	1	0	0	Flush	R	R	Ethernet/IP address										
HLT IP destination	1	0	1	Flush	R	R	Ethernet/IP address										
L1 trigger (1)	0	1	0	R	L0 EvID(11 downto 0)												
L1 trigger (2)	0	1	1	R	Trigger type			R	L0 EvID(19 downto 12)								
Other	0	0	R	R	R	R	R	R	R	R	R	R	R	R	R	R	

Table 8 shows a summary of the new broadcasts. The first two must already be handled by the L1 Front-End electronics. For the L1 triggers the baseline is still to use the short broadcast format and transmit both L1 trigger accepts and rejects until a decision has been taken on the issue of the L1 trigger sorting. It should be noted that the common L1 FE board already supports the scheme of only transmitting unsorted L1 trigger accepts as long broadcasts.

6 References

- [1] The LHCb Collaboration, *LHCb Online System TDR*, CERN/LHCC 2001-040, 19 December 2001.
- [2] J. Christiansen, *Requirements to the L1 front-end electronics*, LHCb note 2003-078 Revision 2.0, August 15, 2003.
- [3] Artur Barczyk, Jean-Pierre Dufey, Beat Jost, Niko Neufeld, *The New LHCb Trigger and DAQ Strategy: A System Architecture based on GbEthernet*, presented at the Real Time conference, Montreal, May 18-23, 2003.
- [4] B. Jost and N. Neufeld, *Raw-data transport format*, LHCb note 2003-063 DAQ, 28 July 2003.
- [5] Z. Guzik, R. Jacobsson, B. Jost, *Driving the LHCb Front-End Readout*, proceedings of the Real Time 2003 conference, May 2003, Montreal, Canada. Submitted to IEEE TNS.
- [6] Richard Jacobsson, Beat Jost, Zbigniew Guzik, *Readout Supervisor Design Specifications*, LHCb 2001-012 DAQ.
- [7] A. Barczyk, J.-P. Dufey, B. Jost, N. Neufeld, *A common implementation of the Level-1 and High Level Trigger Data Acquisition*, LHCb note 2003-079.
- [8] R. Jacobsson, P. Koenig, Z. Guzik, A. Chlopik, *The Final LHCb Readout Supervisor "ODIN"*, proceedings of the 8th Workshop on Electronics for LHC Experiments, September 2002, Colmar, France.
- [9] RD-12 Documentation on WWW (<http://www.cern.ch/TTC/intro.html>) and references therein.
- [10] B. Jost, TFC Broadcast Format, LHCb note 2001-017, 7 March 2001.
- [11] H. Muller, F. Bal, A. Guirao, *Gigabit Ethernet mezzanines for DAQ and Trigger links in LHCb*, LHCb note 2003-21, 28 April 2003.
- [12] Grzegorz Kasprowicz, report by summer student to come.
- [13] A. Bay et al., *Common read out board for LHCb – specification v. 2.0*, LHCb note 2003-007, 3 July 2003.
- [14] P. Moreira et al., *G-Link and Gigabit Ethernet Compliant Serializer for LHC Data Transmission*, proceeding of the NSS conference, 2000.
- [15] Atlas collaboration, *HOLA - High-speed Optical Link for Atlas*, <http://hsi.web.cern.ch/HSI/s-link/devices/hola/>