

Electronics Architecture of the LHCb Upgrade

LHCb Technical Note

Issue:

Revision: 0.7

Reference: EDMS

Created:

Last modified:

Prepared By: Ken Wyllie

Abstract

The electronics architecture of the upgraded LHCb experiment is defined.

Document Status Sheet

Table 1 Document Status Sheet

1. Document Title: Electronics architecture of the LHCb upgrade			
2. Document Reference Number: EDMS			
3. Issue	4. Revision	5. Date	6. Reason for change
Draft	0.7	June 2010	

Table of Contents

1.	INTRODUCTION.....	3
2.	GENERAL ARCHITECTURE	4
3.	FRONT-END.....	5
3.1.1.	<i>Zero-suppression.....</i>	<i>6</i>
3.1.2.	<i>Transmitting non-zero-suppressed data.....</i>	<i>7</i>
3.2.	SIGNAL SYNCHRONISATION USING BUNCH COUNTER.....	7
3.3.	DATA FRAMING.....	8
3.3.1.	<i>Header definition</i>	<i>8</i>
3.3.2.	<i>Orbit reset.....</i>	<i>9</i>
4.	BACK-END	10
4.1.	DATA RECEIVER AND SYNCHRONISATION.....	10
4.2.	INPUT BUFFER AND INTERACTION TRIGGER.....	10
4.3.	DAQ FORMATTING AND TRANSMISSION	11
4.4.	TRIGGER THROTTLING	11
4.5.	IMPLEMENTATION OF BACK-END	11
5.	INTERACTION TRIGGER.....	13
6.	TIMING AND FAST CONTROL (TFC)	15
7.	EXPERIMENT CONTROL SYSTEM (ECS) INTERFACE.....	17
7.1.	FE ECS INTERFACE	17
7.2.	BE ECS INTERFACE.....	17
8.	DATA ACQUISITION (DAQ) INTERFACE.....	18
9.	TESTING AND COMMISSIONING.....	19
10.	RESETS	20
10.1.	BCNT RESET	20
10.2.	EVENT-ID RESET	20
10.3.	FE RESET.....	20
10.4.	BE RESET	20
10.5.	CONFIGURATION RESET	21
11.	USE OF THE GBT AND VERSATILE LINK	22
11.1.	IMPLEMENTATION IN LHCb	23
12.	USE OF PROGRAMMABLE COMMERCIAL DEVICES IN FE MODULES.....	25

13. REFERENCES.....26

1. Introduction

This note describes the electronics architecture adopted for the upgrade of the LHCb experiment. The upgrade is described in an Expression-of-Interest [1] and is based on a highly efficient trigger algorithm run on a farm of CPUs and using data from all bunch-crossings. No hardware trigger is applied to the on-detector electronics, so sub-detectors must transfer the data from each bunch crossing at 40MHz. It is expected that only seventy-five percent of bunch-crossings will contain interactions leading to interesting physics, so an ‘interaction trigger’ based on a simple algorithm will be applied to the data at the level of the off-detector electronics in the counting room. This will reduce the event rate to about 30MHz.

The following chapters will describe the architecture chosen to meet these general requirements.

2. General Architecture

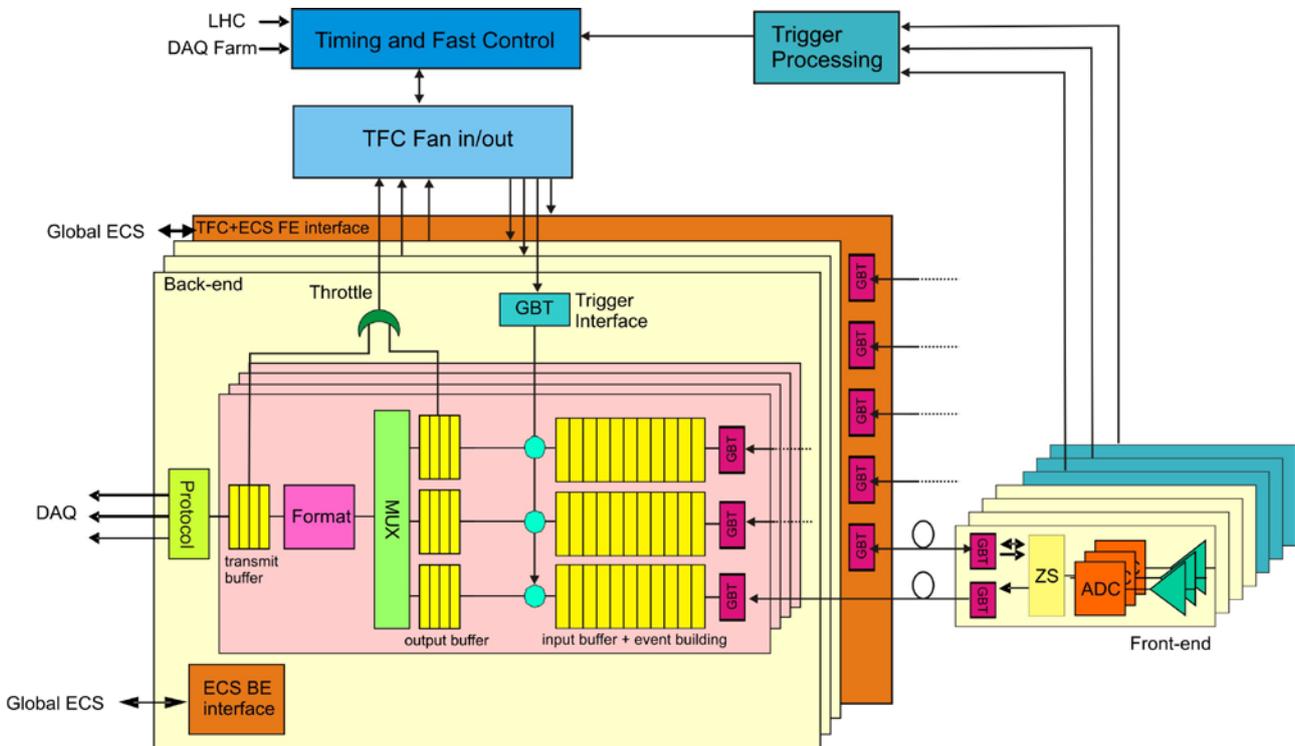


Figure 1: General architecture

The general architecture is shown in Figure 1. The front-end (FE) amplifies and shapes the small signals generated within the particle detectors. These signals are digitised, zero-suppressed and then transmitted down a high-speed optical link running at a serial rate of 4.8Gbit/s. All components of the front-end electronics are located on or close to the detector, and are therefore exposed to some level of radiation.

The back-end electronics (BE) sit in the counting room and receive the data from the optical links. After buffering and filtering by the interaction trigger, data are formatted for transmission to the data acquisition system. The counting-room is a radiation-free environment and commercial components can be used for the implementation of the electronics.

Data from certain sub-detectors are extracted via an independent transmission system to the trigger processors to generate the interaction trigger. This is then distributed to the back-end electronics by a Timing and Fast Control (TFC) system. Configuration and monitoring of the BE and FE electronics are through an interface to the Experiment Control System (ECS).

Each of these components is described in the following chapters.

3. Front-end

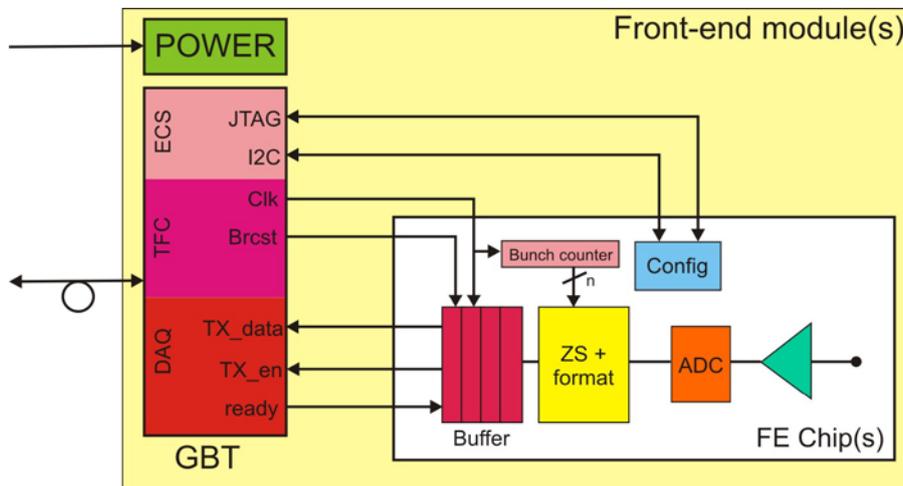


Figure 2: A typical front-end module

A typical FE module is shown in Figure 2. It consists of an FE chip(s) connected to a high-speed bi-directional optical link, implemented using the Gigabit Bidirectional Trigger and Data link (GBT) [2]. The FE chips should be tolerant to the radiation environment of the detector. High radiation areas often require a custom design but commercial devices can be used as long as they are proven tolerant to the expected radiation levels.

The FE chip(s) should amplify, shape and digitise the detector signals according to the requirements of the detector. The rise time of the shaped pulse must be sufficiently short to allow the correct identification of a hit with its bunch crossing while the recovery time must meet the detector specifications on, for example, double-pulse resolution or pile-up. Each sub-detector must choose its digitisation scheme as an optimisation of performance, complexity and cost. For example, the simplest case is a one-bit digitisation implemented using a discriminator. Digitisation at this stage is mandatory to allow high speed digital transmission.

Digital data are synchronised with the 40MHz clock provided by the GBT. The GBT allows fine phase adjustment of this clock and this feature should be used to find the optimum efficiency within the 25ns clock period. Data are then zero-suppressed, buffered and formatted. The buffer will absorb statistical fluctuations in the zero-suppressed event size and allow an optimal use of the data bandwidth provided by the GBT. However, this implies that data from different FE modules will arrive asynchronously at the BE modules. Additional information is therefore added to the data frames to allow reconstruction of the complete events. The minimum formatting is the attachment of a header containing some bits of the bunch-counter (Bcnt). This Bcnt must be generated within the front-end chip synchronously with the clock. These bits are used to track the data later in the system. The data format is discussed in more detail below.

After formatting, data are transferred to the DAQ interface of the GBT chip. Data are then transmitted off the detector only when the GBT is in the READY state.

Fast, synchronous controls are provided by the GBT link transmitting from the counting room, one frame broadcast every 25ns clock period. The broadcast frame is decoded by the FE chip to carry out the following tasks:

1. generation of a FE reset
2. generation of a calibration pulse
3. generation of the Bcnt reset

The encoding of the frame is described in Section 5.

Configuration of the FE chip(s) is done through the ECS interface of the GBT chipset. The FE chip(s) should support one of the standard protocols provided by the GBT.

The bandwidth available in the GBT for transmitting ECS/TFC information is large. So in many cases, it is efficient to use one duplex GBT link to provide ECS/TFC information to many front-end components.

3.1.1. Zero-suppression

Data are zero-suppressed before transmission off the FE modules. The amount of data and the time to zero-suppress will have some statistical fluctuation based on the detector occupancy in each bunch crossing. A buffer should therefore be implemented to allow for these fluctuations and the FE module should carefully monitor the behaviour of this buffer. If the buffer occupancy reaches its limit, the FE module must truncate the data and allow the buffer to recover. The module continues to send header information but with no data until the buffer is sufficiently empty. The activation of data truncation must be indicated in the header.

The FE module should be capable of masking noisy or faulty channels that could limit the efficiency of a zero suppression algorithm.

The implementation of zero-suppression in the FE electronics implies that the parameters must be well understood before the hardware is designed and constructed. For example, the efficiency of the algorithms and the amount of buffering must be tested carefully using data from realistic physics simulations. Depending on the detector radiation environment, it may be possible to use programmable devices for the zero-suppression, which will allow some flexibility in the implementation. The use of such devices is addressed in Section 11. The important parameters that should be quantified are:

1. Inefficiency = number of truncated bunch-crossing packets/number of bunch-crossings.

2. Maximum size of the data packets. This has a large impact on the buffering implemented in the BE modules.
3. Maximum readout latency. Depending on the algorithms, there could be a wide spread in the latency between the bunch-crossing and the corresponding data packet arriving at the BE module. This also has a large impact on the buffering in the BE modules and on the number of bits of Bcnt needed to prevent ambiguities in the time-stamping of packets.

[For discussion: Definition of buffer limits

Buffer implementation – can all sub-detectors implement a large RAM buffer as proposed by the Calo, buffer a sequence of ~ 300 BXs and always truncate those at the end of this sequence]

3.1.2. Transmitting non-zero-suppressed data

In some regions of some sub-detectors, the hit occupancy can be high enough that there is no benefit in zero-suppressing the data. Transmitting non-zero-suppressed data is allowed in these cases as long as the header protocol is consistent with zero-suppressed data.

Instantaneous switching between zero-suppression and non-zero-suppression is discouraged because of the complexity it introduces in the front-end.

All sub-detectors must include a special running mode where the front-end module transmits non-zero-suppressed data at a limited rate. This can be used during calibration and commissioning. Such a running mode implies a reconfiguration of the readout system, including the BE and TFC, and hence adequate partitioning to allow a particular sub-detector to run alone in this mode.

3.2. Signal Synchronisation using bunch counter

Overall synchronisation of the system should be done using the Bcnt information included in the data. Some possible sources of de-synchronisation are the following:

1. In some FE modules, signals from different parts of the detector may arrive at different times after the interaction
2. Counter (BCnt) reset pulses may have different arrival times from module to module due to cable delays.

To allow synchronisation, a local offset of the Bcnt counter should be implemented. This offset is loaded into the counter on the arrival of a Bcnt reset and its value should be programmable via the ECS interface.

3.3. Data framing

Data are transmitted in a format containing a header of fixed length followed by the bunch crossing data of varying length. This is illustrated by the example in Figure 3, which shows the transmission of a number of data packets of varying sizes and the use of the front-end buffer.

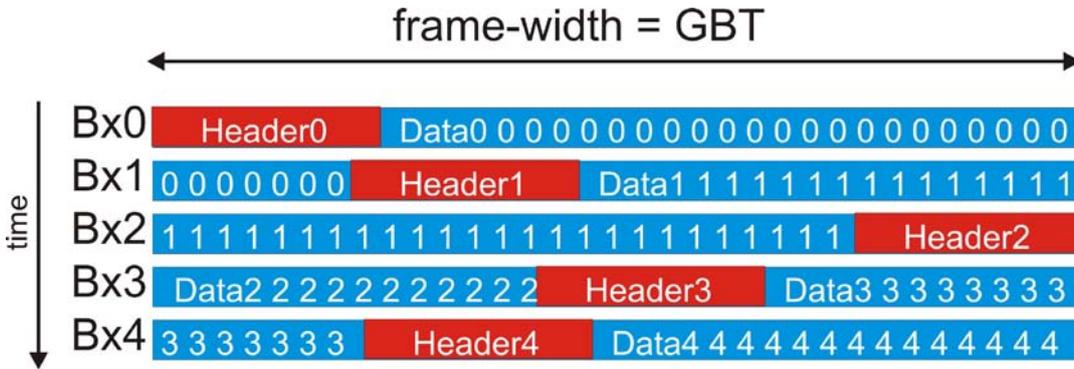


Figure 3: Example of data frame structure

In this example, the width of the frame is limited by the bandwidth of the GBT. The front-end buffer allows the transmission of data packets longer than the GBT frame. However, buffer overflow will occur if a number of consecutive bunch crossings exceed this limit. At this point, headers continue to be sent but transmission of data frames should stop until the buffer occupancy falls below the limit.

3.3.1. Header definition

The header contains m bits and these are defined in Table 2.

Bits	Definition
$[(m-1):5]$	Length of the data for this BX
$[4:1]$	Four LSBs of the BCnt
$[0]$	Status: 0 = normal, 1 = data truncated

Table 2: Header definition

The parameter m is defined by the sub-detector according to the maximum length of data it expects to transmit per bunch crossing. If, for example, a sub-detector will send a maximum of 64 bits of data per bunch crossing, the field $[(m-1):5]$ is 6 bits, and hence m is 11. Note that if the data size for

a particular bunch crossing exceeds the maximum allowed for this sub-detector, then the data must be truncated and header bit [0] asserted.

The back-end electronics will use the contents of the data-length field to extract the data and can also monitor the increment of the BCnt field as a cross-check of the synchronisation.

Note that this header definition is the minimum that should be implemented. For example, some systems may need to transmit more bits of BCnt to avoid ambiguities in the data packets or a bit to indicate if the data is zero-suppressed or not.

3.3.2. Orbit reset

The LHC bunch structure will consist of 3564 bunches spaced by 24.95ns. The final 119 bunch-crossings are empty thus leaving a gap of $\sim 3\mu\text{s}$. This gap will be used to reset the front-end electronics and guarantee correct synchronisation for the next orbit.

After the final filled bunch crossing in the orbit at LHCb, the TFC system will wait for a fixed number of crossings to allow the sub-detectors to fully transmit any remaining data from their front-end buffers. The TFC will then broadcast BCnt and front-end resets to reset the bunch counters and buffers of the front-end modules respectively. Resets are discussed in more detail in Section 10.

[Discussion

1. Note sure if 119 BXs are enough to empty fully the front-end buffers. Maybe just reset BCnt in the gap at the end of the orbit
2. Have to make sure that the latency of this BCnt reset is not longer than 119 bunch crossings!
3. We assume that the orbit structure will not change.....]

4. Back-end

The back-end (BE) electronics are situated in the counting room and act as an interface between the FE modules, DAQ, TFC and ECS systems. The counting room is a radiation-free environment and commercial components can be used without the need for radiation qualification. Electronics modules should be implemented in industrial formats to allow the use of standard mechanics, cooling and power supplies.

The left hand side of Figure 1 shows the blocks of the BE, and these are described below.

4.1. Data Receiver and synchronisation

This is based on the GBT system and uses the GBT functionality implemented in an FPGA. Data are received on the optical links arriving from the FE.

[Discussion:

The latency for data to arrive at the BE will vary (see 3.1.1). A maximum will be defined. We have to decide where this is enforced: FE or BE. Doing it on the BE is more flexible, but means that the FE will waste bandwidth in transmitting packets that will be immediately suppressed by the BE....]

4.2. Input buffer and interaction trigger

The interaction trigger is applied to remove data from empty bunch-crossings. When running at full rate, its frequency is about 30MHz. It will also be used to reduce the data rate to the DAQ if, for example, the system cannot run at full capacity or buffers are close to overflowing. It is transmitted from the TFC system by means of a GBT interface.

Data arriving from the FE are stored in the input buffer to allow some latency for making the interaction trigger decision. The occupancy of this buffer should be continuously monitored. If it reaches its limit, the buffer controller must truncate the data and allow the buffer to recover. The buffer continues to send header information but with no data until the buffer is sufficiently empty. The activation of data truncation must be indicated in the header.

With zero-suppression being performed in the FE modules, there is no longer a fixed time between the bunch-crossing and the arrival of its data at the BE input buffer. This requires that the interaction trigger decisions include the Bcnt number of that bunch crossing. Accepts and rejects are done by comparing the Bcnt in the data header with that of the interaction triggers stored in the trigger buffer. At this point, data accepted by the trigger are tagged with an event number.

Another consequence of zero-suppression on the FE is that there may be a wide spread in the latencies of data from the same bunch crossing but from different channels arriving at the input buffer. This spread should be incorporated into the size of the input buffer and the trigger buffer.

Data accepted by the interaction trigger are given an Event-ID extracted from a local Event-ID counter.

[Discussion points:

trigger latency – we should define a maximum

Event-ID – how many bits?

Should we buffer the interaction triggers?

Do we want to build complete events on TELL40 and then transmit to DAQ? ie What is the allowed time spread on transmitting events from the different TELL40 units to the DAQ?]

4.3. DAQ formatting and transmission

Data are merged from different sources according to the bandwidth of the link to the DAQ. To optimise the use of the bandwidth, it may be advantageous to construct multi-event packets (MEPs) before transmission. This formatting strongly depends on the event size and the occupancy of the output buffers at this stage will be monitored. If they are close to being full, then an alarm is back-propagated to initiate throttling.

Packets of data are then transmitted to the DAQ system on one or more links whose total bandwidth can sustain the data throughput.

4.4. Trigger throttling

All buffers shall monitor their status and signal an alarm when they are close to being full. The transmission of this alarm will be handled by a GBT interface sending data from the BE module back to the TFC system. Any BE module in the system can throttle the interaction trigger.

4.5. Implementation of back-end

[To be expanded by TELL40 people]

The GBT link offers bi-directionality and is conceived to provide DAQ, TFC and ECS in one link. However, the large imbalance between the required bandwidths from and to the detector together with the relatively coarse modularity of the LHCb sub-detectors (in contrast to the trackers of ATLAS and CMS) means that it is more efficient to use dedicated simplex GBT links for DAQ and dedicated duplex links for ECS/TFC. This implies the implementation of a BE module dedicated and optimised for receiving DAQ data and another module dedicated to the ECS/TFC interface of the front-ends.

[Discussion:

1. The two modules could be the same motherboard, but with different mezzanine implementations eg DAQ mezzanines to received data from FE and ECS/TFC mezzanines to both receive and transmit data from FE]

5. Interaction Trigger

The interaction trigger is included in the system to provide the following functions:

1. Allow staging of the DAQ system by controlling the rate.
2. Allow a throttling mechanism if buffers overflow or the CPU farm has insufficient resources to process the data

The interaction trigger will make decisions based on detector data and is based around the existing LHCb L0 trigger [3]. It will be formed with data from the Calorimeter and the Muon detectors. The trigger will be broadcast from the TFC system with a fixed latency.

Figure 4 illustrates the principle. The ECAL and HCAL detectors are readout by FE boards equipped to handle the new architecture but with a separate data path for the trigger. This transmits data to a Validation Card where initial selection is carried out. This is the same module used in the existing experiment. Data from the Validation Card are then transmitted off-detector to two BE cards (TRIG40) equipped with the appropriate input links running at 1.6 Gbit/s with 8bit/10bit encoded data. Decision-making algorithms are carried out on these BE boards who then transmit the results directly to the TFC system. The Muon trigger is based almost entirely on the existing L0-muon system, where data from off-detector electronics (ODE) are transmitted to the counting room at 40MHz. The existing L0-muon system then processes this data and transmits decisions to a third TRIG40. Final muon decisions are made in this TRIG40 which interfaces directly to the TFC. Note that the TRIG40s play the role of the L0 decision unit in the existing LHCb. Note that the TRIG40 modules are also equipped with links to transmit data to the DAQ in a manner identical to the standard BE modules.

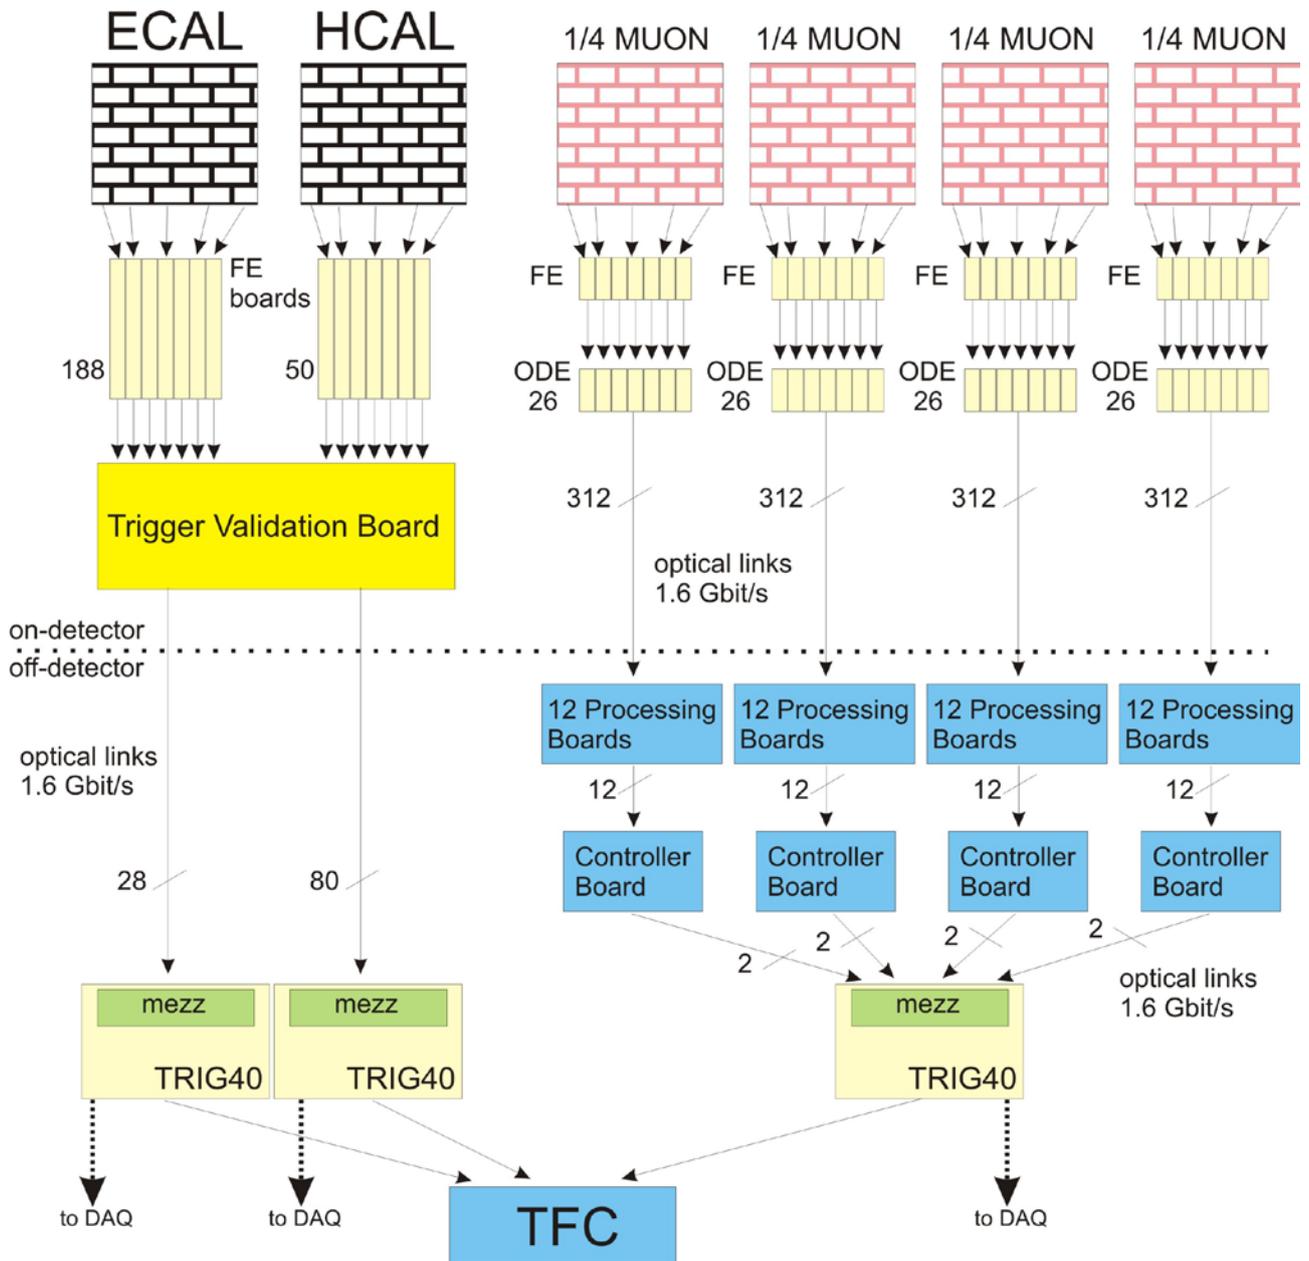


Figure 4: Interaction trigger system

[Discussion points: trigger latency – we should define a maximum

Should we buffer the interaction triggers?]

6. Timing and Fast Control (TFC)

[Section to be expanded by Richard and Federico]

Triggers, clock and fast control signals are distributed by the TFC system.

The interaction trigger is transmitted as a bunch-crossing number of 12 bits to ensure no ambiguities within an orbit cycle of 3564 crossings.

The FE and BE modules receive clock and fast controls (broadcasts) via their GBT interfaces. The data part of the GBT frame contains the TFC broadcast which is then decoded by the modules. The frame-rate and hence broadcast bandwidth is 40 MHz. The commands and their corresponding codes are listed in Table 3.

Command	GBT frame bit[X]	GBT frame bit[X]	GBT frame bits [X:X]	GBT frame bits [19:8]	GBT frame bits [7:0]								
					7	6	5	4	3	2	1	0	
No operation	0	0	0	0	0	0	0	0	0	0	0	0	0
Bcnt reset	x	x	x	x	x	x	x	x	x	x	x	x	1
Event-ID reset	x	x	x	x	x	x	x	x	x	x	x	1	x
FE reset	x	x	x	x	x	x	x	x	x	1	x	x	x
BE reset	x	x	x	x	x	x	x	x	1	x	x	x	x
Calibration pulse	x	x	x	x	x	x	x	1	x	x	x	x	x
Readout Type	x	x	x	x	0,1 ... ,7			x	x	x	x	x	x
MEP transmit	x	1	MEP event ID	MEP destination address	x	x	x	x	x	x	x	x	x
Interaction trigger	1	x	x	x	x	x	x	x	x	x	x	x	x

Table 3: Definition of TFC broadcast bits in GBT frame

[To be included:

Definition of interaction trigger transmission ie BCnt

Discussion:

An interesting proposal is for the TFC to broadcast the ‘collision state’ of each BX, ie flagging BXs where there are empty bunches or no collisions. This could be useful for the detectors where noise hits can form a significant part of the data bandwidth or where they rely on empty bunches (no data) to avoid saturation of their buffers.]

7. Experiment Control System (ECS) Interface

[Section to be expanded by Beat & Niko]

7.1. FE ECS interface

The ECS interface to the FE components will be implemented using the GBT, as described in Section 10. This bi-directional link allows the writing and reading of configuration and monitoring data. The GBT-SCA chip provides an interface between the GBT and standard protocols such as I2C and JTAG and can be mounted on an FE module.

[To be defined:

The 'master' module and its protocol.]

7.2. BE ECS interface

The ECS interface to the BE components will be implemented using a credit-card PC (CCPC) mounted on a mezzanine daughter-card. This will provide standard control interfaces for hardware configuration and monitoring. The communication protocol between the CCPC and the central ECS system will be TCP/IP.

[Discussion:

1. Proposal to use NIOS]

8. Data Acquisition (DAQ) Interface

[Section to be expanded by Beat & Niko]

9. Testing and Commissioning

All FE and BE modules should contain functionality for in-situ testing, calibration and commissioning. In this architecture, a crucial task will be the correct synchronisation of the bunch counters across the experiment and mechanisms should be included in the electronics to allow this. Examples of these are:

1. Test pulse injection. An analog pulse is injected into the front-end amplifiers, triggered by a synchronous command issued by the GBT-TFC interface. The phase of this pulse can be adjusted by the GBT but the amplitude has to be set by the FE module. This technique can be used for the bunch-counter synchronisation as long as path differences in the distribution of the TFC command across the experiment are well known. The timing of this TFC command must be done at a known time after transmission of the BCnt reset, described in Section 10.
2. Triggered light source. If the detector is sensitive to light, then a single pulsed light source illuminating many channels can be used to test and time align these channels relative to each other. The trigger of the light source is generated by a synchronous command from the GBT-TFC interface and its phase is adjustable. The timing of this TFC command must be done at a known time after transmission of the BCnt reset, described in Section 10.
3. Cosmic rays. These can be used for an absolute time alignment of the bunch counters, but care has to be taken because of the asynchronous arrival time of the particles.
4. Digital pattern. The ability to transmit a fixed or predictable pattern from the FE module can be useful for the verification of the links and the receivers of the BE modules. The generation of this pattern can be implemented in the FE digital logic and should run independently of the actual signals from the detector elements.

10. Resets

The FE and BE electronics must react to a set of reset signals in a well-defined fashion to ensure that correct error recovery and synchronization can be obtained at start-up or after a detected malfunction. Synchronous resets are sent as part of the TFC frame via the GBT interface. These are defined in Table 3. None of the reset signals defined in 9.1 – 9.4 should affect the configuration of the electronics.

10.1. Bcnt reset

The Bcnt reset is used to maintain a correct synchronization to the bunch collisions. When asserted, the FE and BE electronics modules should reset (or preset) their bunch counters. No other electronics are affected by this action. The Bcnt reset will be asserted once every machine orbit to maintain overall synchronization of the system.

10.2. Event-ID reset

The event ID appended to data accepted by the interaction trigger is based on a counter incremented for each trigger. This counter is reset by the Event-ID reset.

10.3. FE reset

The FE reset is defined to reset the components on the FE modules. This includes buffer address pointers and other digital logic related to the data flow.

10.4. BE reset

The BE reset is defined to reset the components on the BE modules. This includes buffer address pointers and other digital logic related to the data flow. It should only be asserted after any valid data stored in the output buffer has been transmitted.

10.5. Configuration Reset

[To be discussed if we need it]

This will reset the configuration parameters of the modules and is applied through the ECS interface.

adapter) which translates into standard protocols such as I2C and JTAG. SC falls under the heading of ECS data in LHCb jargon.

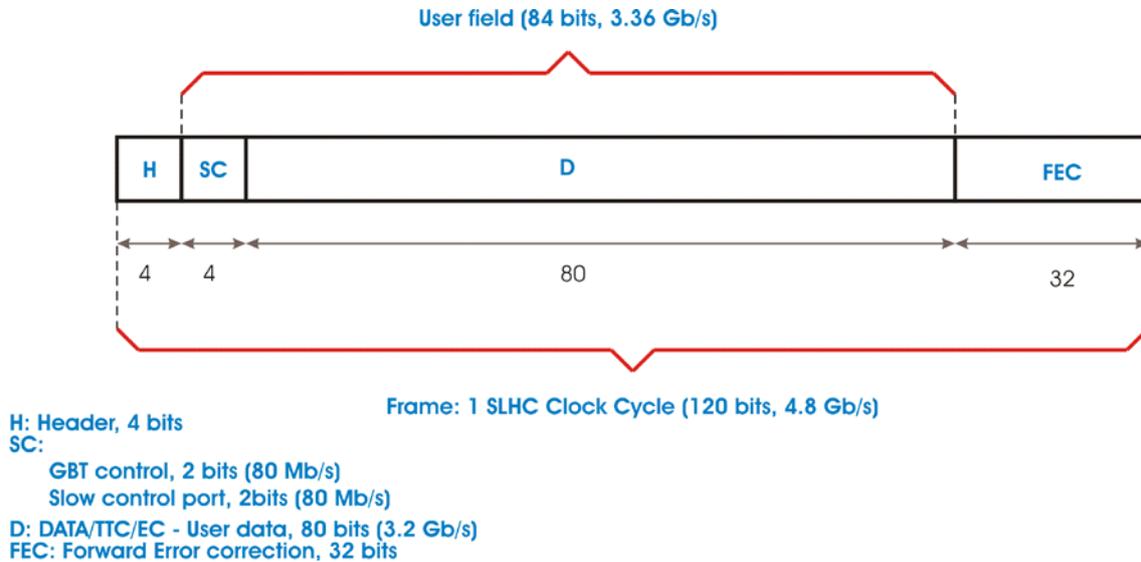


Figure 6: GBT transmission field

The GBTX can be used in one of three different modes. In all cases, and external clock source (for example, a crystal oscillator) is needed to help in the initial locking procedure. Note that unused blocks are powered off in cases 2 and 3.

1. Transceiver mode: The system is a duplex transceiver. Configuration (write/read) of the GBTX is done by the receiver circuitry. The clock output of the GBTX is recovered from the receiver data.
2. Receiver mode: The system is a simplex receiver. Configuration (write/read) of the GBTX is done by interfacing the chip to an auxiliary system that has bi-directional capability. The clock output of the GBTX is recovered from the receiver data.
3. Transmitter mode: The system is a simplex transmitter. Configuration (write/read) of the GBTX is done by interfacing the chip to an auxiliary system that has bi-directional capability. An external clock source must be used, for example another GBTX in one of the other two modes above.

11.1. Implementation in LHCb

The data bandwidth from the LHCb detector to the counting room will be many times larger than the TFC and ECS traffic in the opposite direction. So it is likely that the GBTX will be widely used in simplex-transmitter mode. However, duplex transceivers will be required for writing and reading ECS information and the receiver mode for TFC operations (clock, calibration, resets).

An example of a possible implementation is shown in Figure 6 for a system with a FE module containing many channels and hence a high data bandwidth. The FE chips drive multiple GBTXs configured in simplex-transmitter mode. These then drive a number of dual-transmitter optical packages connected to fibres that fan into a 12-way fibre ribbon. This arrives at a 12-way receiver on the BE board.

A single GBTX on this FE module is configured in duplex-transceiver mode to accept TFC and ECS data from a TFC/ECS module. This distributes TFC information to the GBTXs and other FE components. It also provides the local ECS interface through the GBT-SCA. For example, the GBT-SCA has a number of I2C busses that can be used to configure and monitor the front-end components.

The implementation of the links has to be optimized according to the bandwidth and integration requirements of each sub-detector.

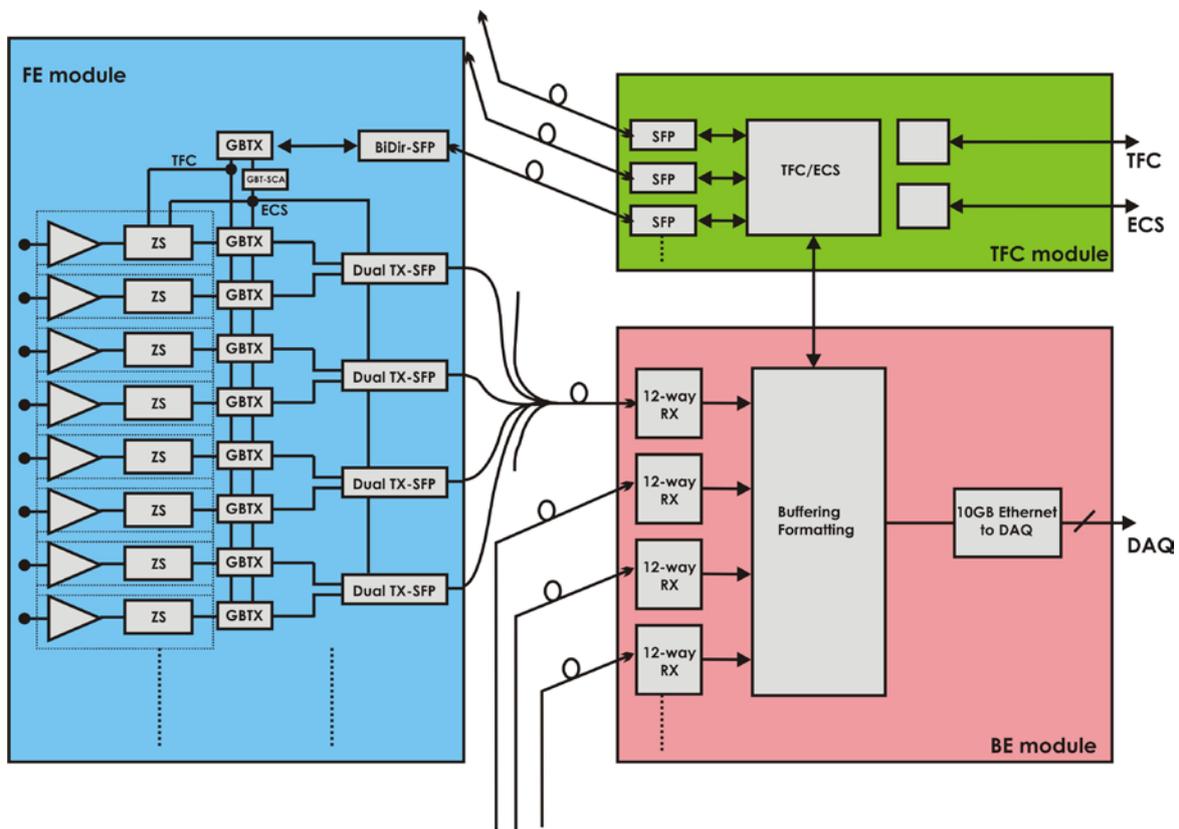


Figure 7: A possible link implementation in the LHCb upgrade

12. Use of programmable commercial devices in FE modules

Zero-suppression and data packing must be carried out in the FE electronics. In many cases, this requires sophisticated algorithms that depend heavily on the expected occupancy of the detector. It is advantageous to have the possibility of changing or tuning these algorithms and using programmable devices allows such an optimization.

Such a device has to be proven to resist the radiation environment of the detector according to the following requirements:

1. Any increase in power consumption due to total-ionising-dose must be estimated and the power infrastructure should be capable of supporting this.
2. The device must be tolerant to single-event-upsets (SEUs) to a level acceptable for the efficiency of the detector. For example, it may be acceptable to have occasional bit errors in the data, but critical circuits, such as buffer pointers, must be much more tolerant to SEUs. Mitigation techniques, such as triple-redundancy, should be used.
3. The configuration of the device must resist the radiation environment.
4. The device must resist destructive single-event-effects.

Devices employing flash programming from ACTEL have been shown to resist upsets in their configuration. The latest versions of these devices are currently under study for their radiation resistance.

13. References

- [1] Expression of Interest for an LHCb Upgrade, CERN/LHCC/2008-007, LHCb 2008-019, 22nd April 2008.
- [2] LHCb Trigger System, Technical Design Report, CERN-LHCC-2003-031, September 2003.
- [3] <https://espace.cern.ch/GBT-Project/default.aspx>
- [4] xxx
- [5] xxx
- [6]