Electronics Architecture of LHCb for Run3 and Future Upgrades

LHCb-PUB-2022-014
24/06/2022

LHCb Technical Note

Issue:
Revision: 1.0

Reference: LHCb-PUB-2022-014
Created: June 2022
Last modified: June 2022

Prepared By: Ken Wyllie, Federico Alessio

Abstract

The electronics architecture of the LHCb experiment from the Run3 era onwards is defined. This covers all data processing from the detectors up to the interface to the event-builder. Building blocks for implementation are also described.

Document Status Sheet

Table 1 Document Status Sheet

1. Document Title: Electronics architecture of the LHCb upgrade			
2. Document Reference Number: LHCb-PUB-2022-014			
3. Issue	**4. Revision**	**5. Date**	**6. Reason for change**
Draft			
	1.0	June 2022	

Table of Contents

1. Introduction

This note describes the electronics architecture of the LHCb experiment circa 2022 and operated during LHC Run 3 (U1). This architecture will also be implemented for the two future upgrades of the LHCb experiment, one a consolidation before LHC-Run 4 (U1b) and the other an upgrade of the full experiment before LHC-Run 5 (U2). These upgrades are described in a Framework Technical Design Report and references therein [1]. The readout electronics architecture of U1b/2 will be conceptually the same as that implemented for LHC-Run 3, where the detector systems transmit data collected for every bunch crossing to a farm of GPUs and CPUs and data filtering by triggering is entirely in software. LHCb does not use any hardware trigger.

Although there is no conceptual change to the architecture, additional features have been added during the implementation for LHC-Run 3 that will be carried forward to U1b/U2. Building blocks will also evolve to use new generations of common components. These are described in this document, which shall be used as a reference and guide for the implementation of systems to be compliant with the LHCb architecture.

The majority of the document describes the front-end electronics, which require custom design to match the requirements of each sub-detector. The back-end electronics and control infrastructure are common across the whole of LHCb. These are described in less detail here and more information can be found in the corresponding references.

2. General Architecture

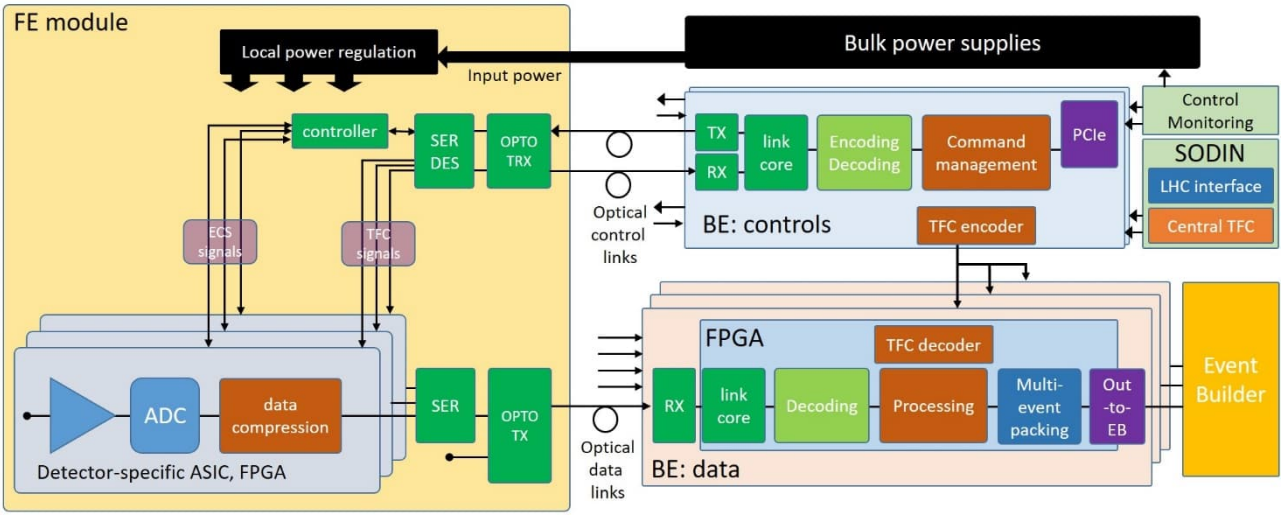

Figure 1: General architecture

The general architecture is shown in Figure 1. The front-end electronics (FE) amplifies, shapes and digitises the signals from the particle detectors. Digital data are then compressed before transmission on optical links. All components of the front-end electronics are located on or close to the detector, and are therefore exposed to some level of radiation depending on their position in the LHCb cavern.

The back-end electronics (BE) sit in a data-centre or counting-room and receive the data from the optical links. Data are decoded, pre-processed and formatted for transmission to the event-builder. The data flow can be controlled globally in the BE. The data-centre is a radiation-free environment where commercial components can be used for the implementation of the electronics.

Clocks and fast commands are distributed by a Timing and Fast Control (TFC) system synchronous to the LHC machine clock of 40.079 MHz. A set of fast control commands regulate the flow-control and synchronise the many components in the system. These commands are synchronous with the LHC clock and transmitted to the FE and BE by the TFC system. While the FE does not operate with a trigger, the BE receives a trigger from the TFC system based on predefined recipes allowing the selection of specific events (bunch crossings) expected to be of interest. Configuration and monitoring of the BE and FE electronics are via an interface to the Experiment Control System (ECS).

Low-voltage power is supplied to the FE electronics by a system of radiation- and magnetic-field-tolerant bulk supplies located 10s of metres away. Radiation-tolerant point-of-load regulators provide the local regulation of the power. High-voltages are provided by power supplies located in radiation-free zones around 60m away.

Generic optical links developed for the LHC experiment upgrades are conceived to provide data acquisition, TFC and ECS functionality in the same link. However, the LHCb system requires a much larger bandwidth for data transmission than for TFC/ECS. For reasons of cost and flexibility, data from the detectors are therefore transmitted on dedicated simplex links whilst the TFC and ECS communication to the FE electronics are merged onto a much smaller number of duplex links.

The LHC bunch structure consists of 3564 bunches spaced by 24.95ns. This defines the basis of all time tagging of data, which is coherent throughout the system and is critical to ensure the correct building of events. Loss of synchronisation with this structure must be recovered as quickly as possible. Regular resets or calibration commands can be issued to the system at convenient gaps in the bunch structure to guarantee synchronisation whilst also minimising data loss.

Data throughput and the required bandwidth are determined by the intensity of the proton-proton collisions at LHCb (luminosity) on top of a background of detector and/or electronics noise. The nominal instantaneous luminosity planned for LHCb during LHC-Run 3 and 4 (U1b) is 2×10^{33} $cm^{-2}s^{-1}$. This will increase to 1.5×10^{34} $cm^{-2}s^{-1}$ for LHC-Run 5 (U2). The electronics systems must be capable of sustaining the data throughput with margin above the target luminosity.

Each of the components in the global architecture is described in the following sections.

3. Front-end

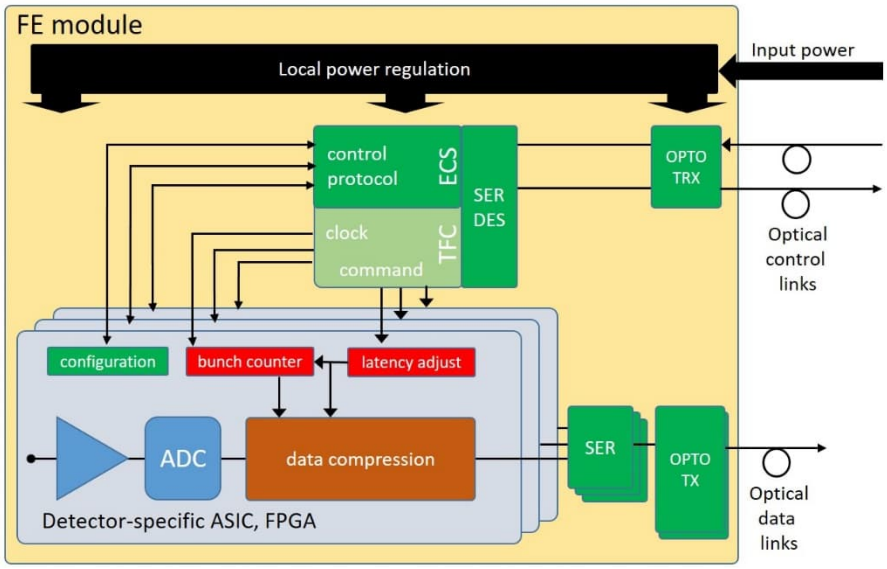

Figure 2: A typical front-end module

A typical FE module is shown in Figure 2. It consists of an FE chip(s) transmitting data to high-speed optical links, implemented using generic radiation-tolerant components. TFC signals arrive at the module on another independent 'master' optical link also used for ECS. The FE chip(s) must be tolerant to the radiation environment of the detector. High radiation environments generally require a custom-designed ASIC but commercial devices, such as FPGAs, 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 overall timing precision of this circuit must be good enough 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. Some detectors will require additional timing data per hit with a precision finer than the bunch crossing period. 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 digital data transmission.

Digital data are synchronised with the 40.079 MHz clock extracted by clock and data recovery from the master link. This clock should be phase adjusted to find the optimum efficiency with respect to the bunch collisions within the 25 ns clock period. Data are then compressed, buffered and formatted. The buffer will absorb statistical fluctuations in the compressed event size and allow an optimal use of the data bandwidth provided by the link. 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 required is the attachment of a header containing all or some bits of the bunch count number (BXID). This BXID must be generated within the FE module and increments synchronously with the clock. The data format is discussed in more detail below. The only

requirement on the FE related to latency is that the number of BXID bits transmitted with the data must be sufficient to unambiguously identify the events when they arrive at the BE.

After formatting, data are transferred on serial electrical links to the electrical interface of the optical link(s). These 'eLinks' operate at typically 160, 320, 640 or 1280 Mbit/s. Many eLinks congregate at the optical link interface(s) where data are serialised and transmitted optically from the module at multi-Gbit/s rates.

3.1. Data compression, zero-suppression and FE buffer

Data are compressed before transmission off the FE modules. In most cases, compression takes the form of zero-suppression. The size of the compressed data, and possibly the time used for compression, will statistically fluctuate 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 transmit information but with truncated (or no) data until the buffer is sufficiently empty. The activation of data truncation must be indicated in the information transmitted on the link.

The FE module must 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 requires that the input conditions must be well understood before the hardware is designed. 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 later optimisation of the implementation in real conditions. The use of such devices is addressed in Section 7. 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 BXID needed to prevent ambiguities in the time-stamping of packets.

If there is the possibility of many channels firing in the same bunch crossing (for example, an injection of common-mode noise) and no meaningful data can be extracted, then the FE module should truncate the event and send only header information. These are referred to later as 'Big-Events'. These are defined according to a threshold on the maximum number of hits per bunch crossing.

3.1.1. Transmitting non-zero-suppressed data

In certain regions of some sub-detectors, the hit occupancy reaches levels where there is no benefit gained from 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.

Where useful, the sub-detectors should include the possibility to transmit a non-zero-suppressed event as part of the normal data flow. This will allow diagnostics of, for example, the zero-suppression algorithm. Such an action is triggered through the TFC interface described later.

3.2. Synchronisation using bunch counter

Overall synchronisation of the system uses the BXID information included in the headers attached to the data. The FEs together with the TFC system together must allow for the calibration and correction of hardware delays to ensure coherent synchronisation across the full detector. Some examples are the following:

1. In some FE modules, signals from different parts of the detector may be collected at different times after the interaction because of time-of-flight or propagation delays;
2. Counter reset pulses (BXID reset) will have different arrival times from module to module due to fibre and cable delays.

A local bunch-counter of 12 bits should be used to generate the BXID. This counts from a minimum to a maximum. For the standard LHC bunch structure the count would be from 0 to 3563. When the counter reaches maximum, on the next clock it will wrap around to the minimum value. To help with synchronisation during commissioning, the minimum and maximum values of the counter are writeable through the ECS interface. When a BXID reset is received, the counter is reset to the minimum value.

3.3. Data framing

Data are transmitted in a packet format where each packet is attributable to a particular bunch crossing. The packet has a header field containing the BXID and any other information necessary for later data-processing, for example the length of the data field or a flag to indicated data truncation. The header is followed by the data collected for this BXID.

It is recommended that a header is transmitted for every BXID even if there is no data to transmit. Synchronisation is then easily checked in the BE by monitoring the incrementing BXID in the headers. This recommendation can be ignored if the required bandwidth cannot be justified or is impossible to implement due to a particular architecture. In these cases, data synchronisation must be checked regularly in the BE by means other than monitoring the increment per header of the BXID.

The choice of header/data structure in the packets must be optimised for efficient use of the links and the downstream resources required to unpack and merge the data from different link locations. This optimisation must be made taking into account the topology of the data which may vary across a sub-detector. For example, some sub-detectors have a non-uniform occupancy with a higher hit

density closer to the beam. In these intense regions, there may be no gain from zero-suppression so the data fields can be of a fixed length determined by the number of channels feeding each link. In areas of lower occupancy, zero-suppression can decrease the required bandwidth but with some statistical uncertainty on the number of hits per BXID. In this case, the length of the data field will vary from BXID to BXID.

Three different types of packet framing are defined in LHCb, and each sub-detector must choose one or a combination of these. The choice must be determined as the result of a co-design between FE and BE electronics, including optimised link use, data integrity, and resource usage in the BE logic. The types are listed below, and are discussed in more details in [2].

1. Dynamic packing with a variable header: This allows an efficient use of the links with a variable length of both the header and data fields. For example, a reduced header can be used for a BXID where there is no data. Unpacking in the BE is generally more complex than the other types and hence requires more logic resources.
2. Dynamic packing with a fixed header: The data field can vary in length, but not the header. This is less efficient for link usage, but requires fewer logic resources for unpacking in the BE.
3. Fixed packaging: Both header and data fields are fixed. Data may have to be truncated. The link usage is less optimal but this requires the least logic resources for unpacking in the BE.

Type 2 is illustrated in Figure 3.

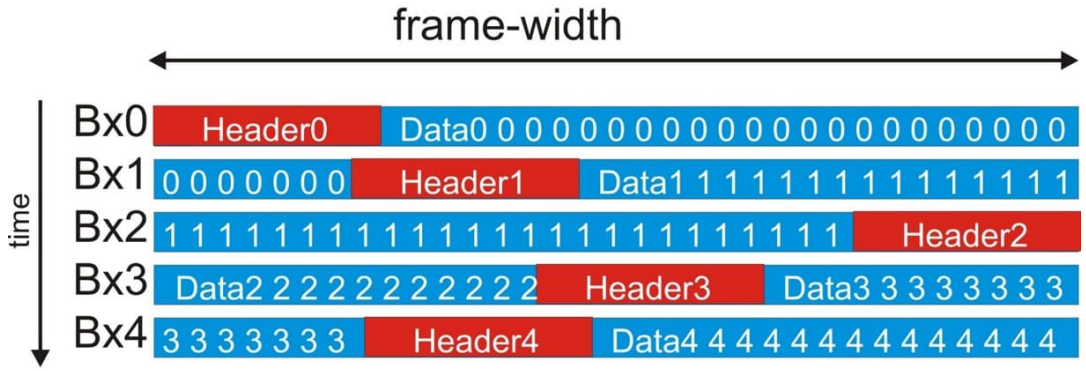

Figure 3: Example of data frame structure with dynamic packing and fixed header

The ‘frame-width’ indicates how many bits can be transmitted per bunch-crossing by the link(s). Note that this frame can also be split into mini-frames, each corresponding to an independent FE device. In this case, headers including BXID information must be transmitted per mini-frame.

In the example of Figure 3, Header0+Data0 exceed the frame-width and hence the remainder of the data field must be buffered in the FE and transmitted in the next frame. Buffer overflow will occur if many consecutive bunch crossings exceed this frame-width limit. At this point, headers continue to be sent but data are truncated until buffer space is available. The number of links (bandwidth) and buffer depth must be determined through simulation and must include margin.

3.3.1. Header definition & protection

The role of the header is to tag the data with the BXID and to send information describing the data field. A formal definition is in [2].

An example is shown in Table 2.

# Bits	Definition
12	BXID
1	Truncation Status: 0 = normal, 1 = data truncated
n	Length of data packet

Table 2: Header definition example

This example includes the full twelve bits of BXID. In certain cases, a reduced number of BXID bits can be transmitted. However, in this case it should be guaranteed by the sub-detector that this will not lead to ambiguities in the attribution of the data to the absolute bunch-crossing. The remaining bits of the BXID must then be added back into the data header by the BE module.

If truncation has occurred then the truncation bit is asserted.

The parameter ***n*** is defined by the sub-detector according to the maximum length of data it expects to transmit per bunch crossing. This field can also be used to indicate if the event data is non-zero-suppressed.

The BE electronics will use the contents of the data-length field to extract the data and will also monitor the increment of the BXID field as a cross-check of the synchronisation. Loss of synchronisation may occur if there are bit errors in the header field. Protecting or error-checking the header (for example, with parity) must be evaluated against the expense in bandwidth. This must be co-designed with the algorithm running in the BE, in particular to specify how many corrupt headers can be accepted and how to recover when this limit is exceeded.

3.4. Clocks in the FE modules

Beam data from the detectors are created by the particle collisions synchronous with the LHC bunch-crossing clock of 40.079 MHz clock. Hence all synchronous activities in the FE use this clock or a derivative of it. The master link regenerates and delivers this clock to the FE components. Within a certain precision, the frequency is the same as the LHC clock. The latency of the master link also guarantees a fixed phase with respect to the LHC clock, again within a certain precision. The precision of the frequency and phase of the delivered clock are discussed below.

This 40.079 MHz clock can be used for synchronous operations in the FE as well as for generating other clocks. This clock synthesis can be done in the FE components, but it must be ensured that all synthesised clocks have a fixed phase relationship with the LHC clock at all times. In general, a

careful analysis of the clocking of a FE module is required during the design process to ensure stable phase with respect to the LHC.

Phase adjustment is often necessary to find optimum working conditions. For example, the clock that samples the detector data before digitisation may require phase adjustment to optimise the sampling efficiency. Crossing between clocks domains without using FIFOs may also require phase adjustment. Current and future implementations of the master link can deliver other clocks with frequency multiples of 40.079 MHz and with fixed phases. Some (but not all) of these clocks can be phase-adjusted. If this is inadequate, then phase adjustment must be implemented in the FE components.

3.5. Resetting the FE module

The FE electronics must react to a set of reset signals in a well-defined fashion to ensure that the system can correctly synchronise at start-up and recover from errors after a malfunction. Synchronous resets are sent as part of the TFC frame via the master link interface as defined in Section 3.6. All of these reset signals are asserted for one clock cycle. None of the reset signals defined here should affect the configuration of the electronics.

3.5.1. BXID reset

The BXID reset is used to maintain a correct synchronization of the entire LHCb readout system with the bunch collisions of the LHC. When asserted, the FE modules should load the preset value into their bunch counters. No other electronics are affected by this action. The BXID reset is asserted at every turn of the LHC with a regular frequency of 11.245 kHz. Note that the BXID counter should not be sensitive to any other reset signal.

3.5.2. FE reset

The FE reset is used to reset the components on the FE modules through which data flows. This includes buffer address pointers and other digital logic related to the data flow, including operational counters and status bits. Moreover, the ECS configuration and control registers must not be changed, including time alignment delay settings of all types. The local bunch counter must also be insensitive to this signal.

The FE reset can be preceded by an extended period with the Header Only signal asserted to empty the FE output buffers and the BE input buffers. The FE reset can also be followed by an extended period with the Header Only signal asserted to allow the FE to put itself in a working state before passing data through. The duration of these is programmable in the TFC system.

3.6. Timing and Fast Controls (TFC) interface to the FE

Fast, synchronous signals are provided by the TFC system through the master optical link. This master link is driven by a common module sitting in the counting room. These signals include both the clocks to synchronously operate the system and commands to control the entire event flow from the FE. Clocks and commands are extracted from the data-stream on the master link by means of clock-data recovery with a fixed latency. Overall control is centralized in a TFC 'master' (or readout supervisor) which has a direct link to the LHC systems for radio-frequency and machine timing. This guarantees distribution of LHC beam-synchronous clocks and commands to the FE.

The master link protocol is designed to be as robust as possible with some user bandwidth sacrificed to allow forward-error correction. The remaining bandwidth is divided between ECS and TFC functionality.

There is one TFC frame transmitted down the master link every clock period. These are broadcasts so every FE device receives the same data. The broadcast commands instruct the FE components to execute synchronous tasks, the most important of which are the following:

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

These and other TFC commands are described below.

The bandwidth available in the master link for transmitting TFC (and ECS) information is large. Therefore, in many cases, it is efficient to use one master link to transmit TFC (and ECS) information to and from many FE components.

Figure 4 shows details of the TFC signals in an FE module. More details of the global TFC system can be found in [4].

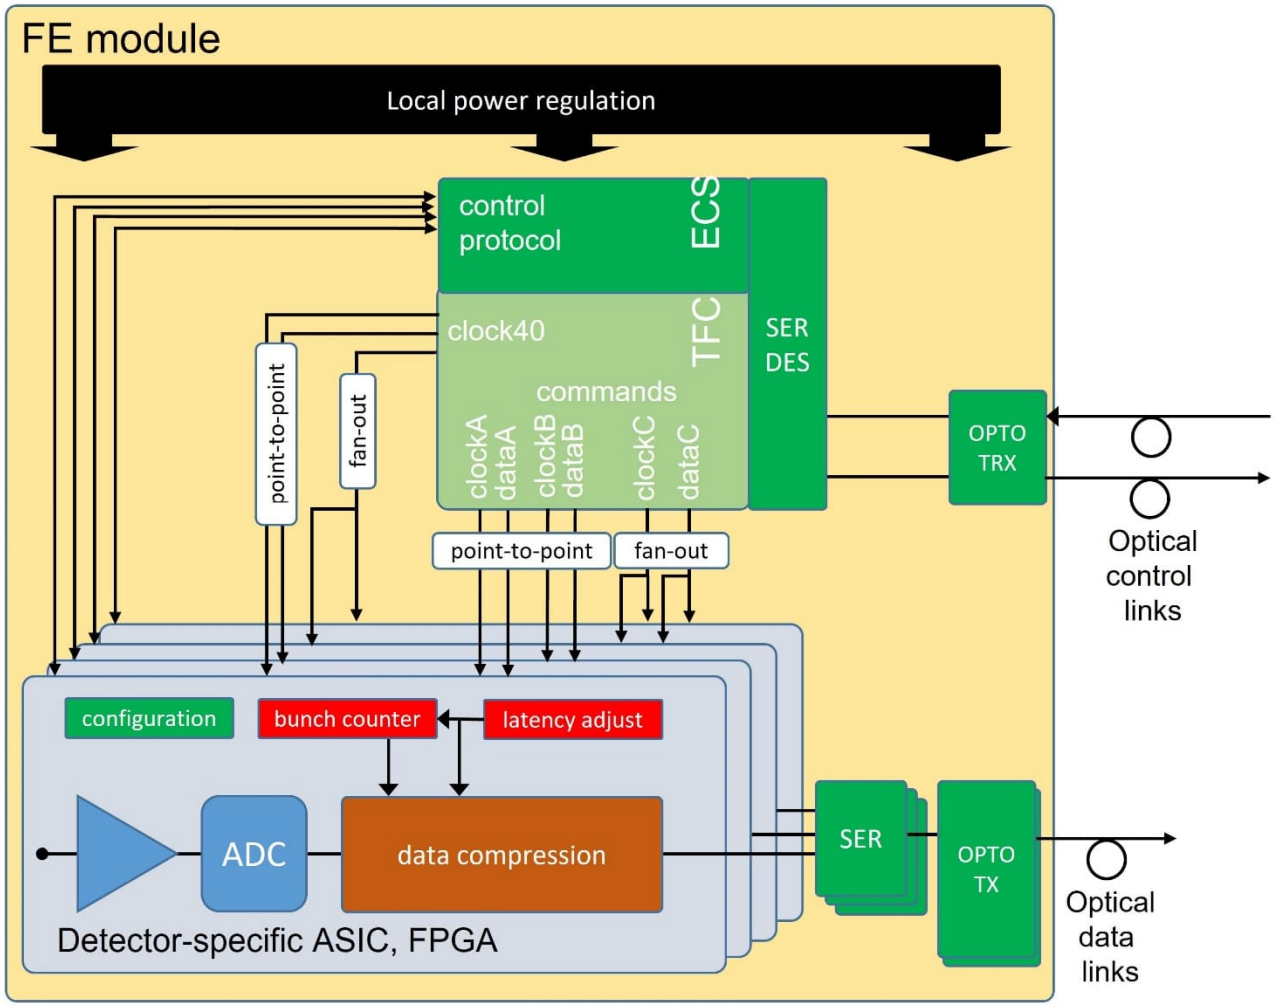

Figure 4: TFC (clocks and commands) distribution within an FE module.

3.6.1. Clocks

The master link provides the local 40.079 MHz clock used by the FE devices. The FE module must be capable of phase-shifting this clock if this is necessary to maximise detector efficiency. This can be done using phase-adjustable clock signals from the master link or by implementing clock phase adjustment in the FE chip. Note that the phase of the serial data containing TFC commands from the master link is not adjustable, so the bit-alignment of TFC commands may depend on the phase-adjustment of the clock.

Current and future implementations of the master link can deliver other clocks with frequency multiples of 40.079 MHz (80, 160, 320, 640, 1280 MHz). Some (but not all) clock outputs can be phase-adjusted in the master link. These clocks are low-swing differential signals that require termination at the receiver end. The easiest implementation is a point-to-point topology, although a fan-out topology (multi-drop) is possible with careful design of the line-termination. These two options are shown in Figure 4.

3.6.2. Commands

In current and future implementations, the master link transmits TFC commands to the FE components through serial electrical links running at a multiple of the 40.079 clock (for example 80, 160, 320 Mbit/s). The FE chip must deserialise these data streams to provide the correct TFC command per bunch crossing. This requires a mechanism in the FE chip to correctly bit-align the deserialised stream with the 40.079 MHz clock.

Figure 4 illustrates the idea. Pairs of signals (clockN and dataN) connect the master link interface to the FE components. Each pair transmits serial data and the corresponding clock. A typical example is a link operating at 320 Mbit/s and 320 MHz (or 160 MHz at double-data-rate). This rate supports a TFC broadcast command of 8 bits every LHC bunch-crossing. These are low-swing differential signals that require termination at the receiver end. The easiest implementation is a point-to-point topology, although a fan-out topology (multi-drop) is possible with careful design of the line-termination. These two options are shown in Figure 4.

The TFC data frame can be transmitted either in an encoded form or with each bit mapped to a command. Note that multiple commands may be sent per bunch-crossing (for example, a BXID reset together with a calibration command) so the choice of encoding must support this.

The FE must react to the TFC commands in a well-defined fashion. Table 3: Definition of TFC commands transmitted to the FE.Table 3 defines the commands.

Command	Function
BXID Reset	The FE modules load the preset minimum value into their bunch counters. No other electronics are affected by this action.
FE Reset	The components on the FE modules through which data flows are reset to a well-defined state. This includes buffer address pointers and other digital logic related to the data flow, including operational counters and status bits.
Calibration	The Calibration command is composed of one or more bits allowing different calibration types. The signals can be used to fire calibration systems either within the sub-detectors or within the FE electronics. The association between a calibration system of a sub-detector and a bit in the Calibration command must be defined together with the TFC system specification.
Header Only	The FE should only transmit header information to the BE for that particular event. This command is used for operational purposes only and it is not used to control the rate nor is it correlated with the bunch-filling scheme.
Bunch Crossing Veto (BX Veto)	The FE should only transmit header information to the BE for that particular event. This is used to reject the LHC crossings with no collisions already at the FE electronics and is based exclusively on the LHC Bunch Filling Schemes, which are loaded into the TFC system. Note that the "Header Only" and "Bunch Crossing Veto" commands are functionally equivalent from the point of view of the FE so can be considered as one command.
Non-Zero-Suppressed (NZS) mode	The FE must transmit non-zero suppressed data. In most cases this will require sending the event data across several frames of the data link so data from subsequent BXIDs will be lost. This mode can be ignored if NZS data cannot be accessed or if it is of no use to the sub-system.
Synch	This feature is used to synchronise the FE and BE systems. If the Synch command is sent, the FE electronics carries out two specific tasks. Firstly, the frame sent to the data link is replaced by a known fixed 'synch' pattern. Secondly, the pointers controlling the write and read operations of the FE buffer are zeroed. When the Synch bit is then de-asserted, the header of the next real data frame will be in a known position within the frame. These two actions will allow the correct frame alignment within the receiving logic of the BE modules. The receiver will be constantly looking for the synch pattern. When this is found, the receiver knows where to look for the header of the next real data frame. To avoid cases where detector data may accidently match the synch pattern, the Synch bit can be asserted for two or more cycles. It is not mandatory to fill the entire frame with a fixed synch pattern. It is recommended that the synch pattern is programmable.
Snapshot	The state of all status and monitoring registers in the FE device is instantaneously sampled in readable registers. These registers must be readable via the ECS interface.

Table 3: Definition of TFC commands transmitted to the FE.

3.6.3. Latency

The TFC commands corresponding to a certain BXID will be transmitted by the TFC master well before the actual bunch crossing takes place. This pre-scale time is chosen to ensure that TFC commands arrive at the FE module at the correct time with respect to the bunch-crossing.

All links in the TFC system have a deterministic latency. The total latency per link is determined during commissioning of LHCb and is used to calculate the pre-scale time. Additional delays can be added per master link to account for the latency spread across the experiment from, for example, different fibre lengths. To allow the local correction of latencies introduced by different link lengths beyond the master link, each FE chip must have the ability to delay the TFC commands by up to 16 clock cycles. This will allow the correction of any latency differences between FE components connected to the same TFC command interface. A larger delay is implemented in the TFC system at a link-by-link level. This delay can allow delaying the commands up to an entire orbit of 3564 clock cycles.

3.7. Precise Timing Distribution

At the higher luminosity of U2, the intensity of the collisions will be such that multiple events will pile-up within a bunch-crossing. In many sub-detectors, precise measurements of the arrival time of particles will be necessary to differentiate between these piled-up events. The precision required is of the order 20-50 ps per hit. The sub-detectors must design the FE components to meet this requirement where necessary.

One of the critical parameters for such a design is the precision of the clock delivered to the FE components from the master link. This is derived from the LHC master clock and re-generated through a cascaded system with many steps before arriving at the FE. Each step can introduce imprecisions in the clock, parametrised as jitter and phase noise. Current and future implementations of the master link are well characterised and measurements of these parameters can be used in the design of the FE. It is the responsibility of the LHCb sub-detectors to determine the timing precision of their design with these boundary conditions and implement additional features to improve the clock precision if required.

More details of the timing characterisation of typical master link implementations can be found in [5].

3.8. FE data flow

Figure 5 illustrates an example of the data flow through a FE module. This is for a system where the data are transported in a linear, time-ordered fashion. The application of the TFC commands are also shown. Note that the BXVeto and HeaderOnly commands are logically equivalent in the FE so HeaderOnly is ignored.

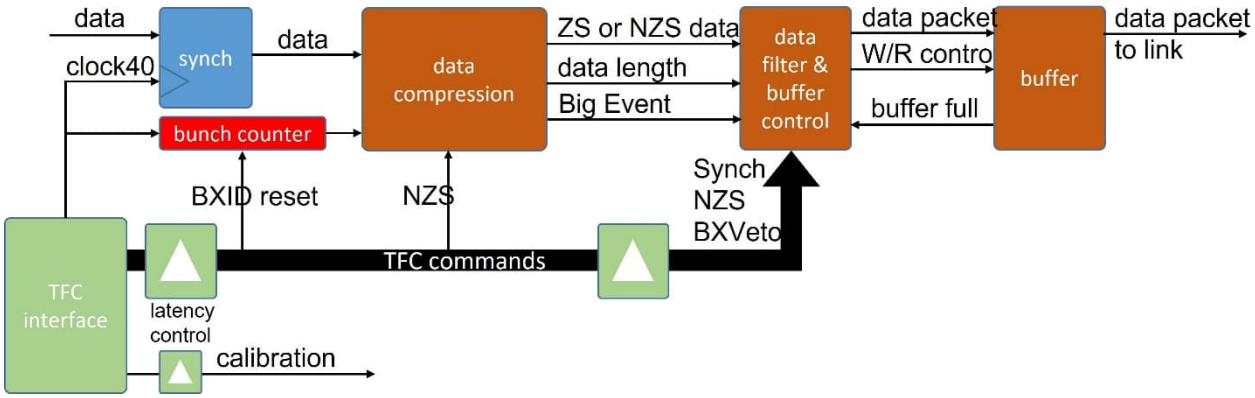

Figure 5: Data flow through the FE and the TFC signals

Data are first synchronised with the 40.079 MHz clock. They are then compressed by zero-suppression. This compression is bypassed if the NZS command is asserted for that bunch-crossing. The BXID is attached to the data packet at that point. The compression logic also calculates the data length and asserts the flag 'Big Event' if the number of hits exceeds a defined threshold. This information is transmitted to the filter/control block. Data are filtered according to the TFC commands arriving synchronously. Filtered data packets are then written into the buffer before transmission to the data link. The buffer status (full or almost full) is also used for filtering the data by activating truncation.

This example assumes that the latency of the data compression stage is fixed. In this case, the synchronicity of the TFC commands is ensured by tuning the delays with the built-in latency control and/or the master link. Then the data transport maintains a fixed latency all the way to the input of the buffer. If the latency of the data compression stage varies on an event-by-event basis, then synchronicity is already lost at that stage. In this case, the data packets must be tagged with the TFC data for each BXID at this point. Data are then filtered according to this tagged TFC information when the packets arrive at the filter.

An example of the data filtering output is shown in Table 4 as a function of the TFC commands and buffer status. The Synch command has highest priority, and in this case the full 12 bits of BXID must be transmitted.

The case of a Big Event must be defined by each sub-detector. For example, if no useful data can be extracted from such an event then the data can be truncated in the same way as for BXVeto or BufferFull.

Data read from the buffer are transferred via an output interface to the data link. If the buffer is not sufficiently full to fill a complete frame of the data link then 'idles packets' can be inserted. These idle packets must be clearly distinguishable from data packets.

Synch	**BXVeto or BufferFull**	**NZS**	**Output of data filter**			
			Header field			**Data field**
			BXID	**Truncation Flag**	**Length**	
1	X	X	BXID[11:0]	NA	NA	Synch pattern
0	1	X	BXID[n:0]	1	not transmitted	No data
0	0	1	BXID[n:0]	0	NZS code (unique)	Uncompressed data
0	0	0	BXID[n:0]	0	Data length	Compressed data

Table 4: Definition of Header and Data fields to write into buffer for different TFC commands. X signifies 'don't care'. 'n' is the number of bits (- 1) of BXID.

3.9. Special running mode

The FE must implement a special running mode for TFC commissioning and latency measurements. This mode is enabled by configuration through the ECS interface. The FE module continues to transmit data packets with the standard headers, but with the detector data replaced by the current value on the TFC command bus. This will allow the calculation of latencies in each sub-system and synchronisation checks.

3.10. Experimental Control System (ECS) Interface

The role of this interface is to configure the FE components, monitor their status and provide a means to measure parameters of the local detector environment. It is not conceived to be used as a Detector-Safety-System to protect the detector from damage.

The ECS interface to the FE shares the same physical master link as the TFC interface. All FE modules in LHCb implement this link, which is bi-directional link to allow writing and reading of configuration and monitoring data. The custom-protocol on this link is adapted to standard (or customized) slow-control protocols by electrical components on the FE module.

This TFC/ECS master link is driven by a common module sitting in the counting room. The link protocol is designed to be as robust as possible with some user bandwidth sacrificed to allow forward-error correction. The remaining bandwidth is divided between ECS and TFC functionality.

The generic scheme is shown in Figure 2. However, some detector topologies may allow the use separate modules (or 'Service Boards') dedicated to ECS (and TFC) functionality and serving a number of FE modules, as shown in Figure 6.

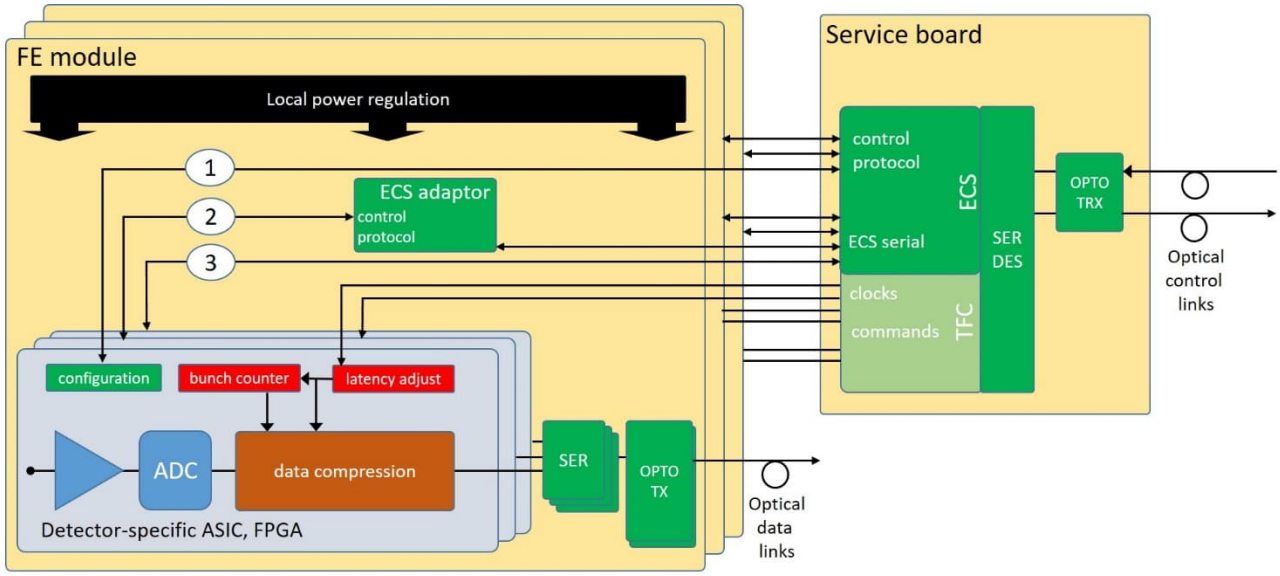

Figure 6: ECS (and TFC) distribution to FE modules by a Service Board.

Figure 6 shows three supported methods of implementing the ECS interface:

1. The ECS block of the master link interfaces directly to the FE components. Communication is via a standard protocol. The most widely supported is I2C.
2. An ECS adaptor is included on the FE module. This connects to the master link via differential eLinks. The adaptor then interfaces directly to the FE components. Communication is via a standard protocol. The most widely supported is I2C.
3. The FE components connect to the master link via serial differential electrical links. The protocol is custom-built based on bi-directional eLinks transmitting serial data at 80, 160 or 320 Mbit/s.

The use of differential signaling in implementations 2 and 3 allows longer distances between FE modules and services boards than I2C. Implementation 3 also allows higher rates of configuration-data transport.

The master links and ECS adaptors also include analog-digital and digital-analog convertors.

Many TFC/ECS master links are driven by a common BE module. This contains a powerful FPGA interfaced to a CPU. Several levels of firmware and software are centrally supported (bottom up):

- Firmware for the FPGA prepares and sends frames down the master links to execute the low-level protocols on the FE modules.
- Low-level libraries and command line tools in the software running on the CPU allow access to the FE components.
- Integration into the ECS via a WinCC-OA component provides the high level description and access to all electronics components.

Note that the time needed to configure a FE board does not depend directly on the bandwidth available and the volume of data. There are several points to consider:

1. Transport protocol overheads (headers, trailers, CRCs, etc.)
2. How the data is distributed, i.e. are there a few large registers or many small ones.
3. Are blocking operations needed? For example does a register needs to be set and read-back before the next one can be configured, or are read-modify-write operations required?
4. The choice of low-level protocol, how it is used and does it allow block transfers. For example, some I2C devices need an extra I2C operation to write a sub-address before the block of data can be written, and some devices do not implement auto-increment, so block transfer cannot be used.

3.10.1. Control and monitoring resource requirements

Control registers

All configurable ECS resources must support both read and write actions. The resources must be readable at all times without interfering with the operation of the detector. For efficiency reasons, several functional parameters may be located together in the same hardware register. The control registers must only be reset via a dedicated ECS command and not via any TFC command.

Status registers and counters

The FE modules must have ECS monitoring counters for all the received TFC commands, and all readout-related operational statistics. Moreover, all the TFC+ECS master-link internal registers must be accessible for monitoring.

With the exception of the BXID counters, the range of all counters should allow at least an hour operation at maximum counting rate (at least 34 bits at 40 MHz).

The FE system must implement a mechanism to obtain a set of consistent values for all monitoring counters in all devices. This requires that an update action or *snapshot* is implemented which is triggered either by an ECS write/read action to a particular register or by a TFC snapshot command. The action should trigger the registering of the values of all counters in the same bunch clock into

separate ECS registers that can then be read out sequentially through the ECS interface. Triggering the snapshot via the TFC gathers consistent information across the experiment for a particular BXID. The monitoring counters should continue undisturbed during the sampling operation.

Status registers reflecting the states of operational logic must be available through the ECS interface in three versions:

1. A live value readable at any time.
2. A value registered via the snapshot action described above in order to store values consistent with the counters.
3. A value asserted if the status register changes to an abnormal value during a single bunch clock cycle in order to be able to monitor instantaneous and rare erroneous state changes. This keeps its value until reset via the ECS.

Table 5 lists the counters to be implemented in the FE modules for monitoring purposes. These can be used to monitor the data flow and the distribution of TFC commands by the master control link.

Name	**Function**
BXIDReset_count	Increments if BXIDReset is asserted in a 40MHz clock cycle.
FEReset_count	Increments if FEReset is asserted in a 40MHz clock cycle.
HeaderOnly_count	Increments if HeaderOnly is asserted in a 40MHz clock cycle.
NZS_count	Increments if NZS is asserted in a 40MHz clock cycle.
BXVeto_count	Increments if BXVeto is asserted in a 40MHz clock cycle.
Calibration_count	Increments if the calibration command is asserted in a 40MHz clock cycle.
Snapshot_count	Increments if Snapshot is asserted in a 40MHz clock cycle.
Synch_count	Increments if Synch is asserted in a 40MHz clock cycle.
BufferFull_count	Increments if a buffer is (almost) full in a 40MHz clock cycle.
BigEvent_count	Increments if a ‘Big Event’ occurs in a 40MHz clock cycle, indicating that the data size for this bunch-crossing is above a threshold.
Truncation_count	Increments if data in a 40MHz clock cycle were truncated.

Table 5: Monitoring counters in the FE module

3.11. Local power regulation (point-of-load)

FE modules receive DC power from low-voltage bulk power-supplies. The input voltage at the modules is typically in the range 6 – 12 V. This must be locally regulated at the point-of-load and converted down to lower voltages required by the FE chips and link components. For example, opto-electronic components can typically require 2.5V whereas FE chips may require 1.2 or 1.5V. Commercial FPGAs often require many different power supply voltages. Hence the local regulation will consist of a number of point-of-load regulators with a common input voltage, each providing power at a particular voltage.

These local devices must be proven to withstand the radiation and magnetic-field environment of the detector. Power sequencing can implemented using the ECS interface except for components in the master link that must automatically activate when the input power is turned on.

The choice of local regulators (for example, DC-DC convertors or linear regulators) must be optimal for the performance of the FE components and the system infrastructure. For example, linear regulators provide precise voltages with little ripple but can dissipate significant heat. DC-DC convertors allow reductions in cable material but can create electromagnetic noise in the environment and often have to be shielded.

3.12. Radiation Tolerance

The FE electronics will be operating in a radiation environment. Each sub-system must understand the levels of radiation at the locations where their modules are situated and design the electronics accordingly. In general, hardening techniques combined with an appropriate technology choice can mitigate radiation effects in ASICs. FPGA logic can be protected by triple-modular-redundancy but configuration memory and hard-IPs are often sensitive. Many techniques are used as standard approaches for radiation hardening and these should be followed as much as possible.

Simulation is a useful tool for finding the sensitivities of a design, but radiation testing is mandatory.

The following radiation effects must be considered for the front-end electronics:

1. Cumulative damage from ionising radiation.
2. Cumulative displacement damage.
3. Single-event effects (SEEs): single-event upset (SEU), single-event latch-up (SEL), single-event functional interrupt (SEFI), etc.

The FE electronics must be tested and proven to tolerate these effects up to and beyond the expected levels in the environment of a sub-system. The following quantities are used to describe the environment and can be extracted from simulations of the LHCb cavern:

1. Total ionising dose.
2. 1 MeV-neutron equivalent fluence.
3. Hadrons above 20MeV in energy.

The statistical probability of SEEs must be estimated using measurements on a realistic implementation of the design. For example, if a design with an FPGA is tested then it must be programmed with the final configuration where all targeted hardware blocks are included. The results of such measurements can then be used to calculate the expected SEE rates in the detector environment.

Then the impact of observed SEEs must be assessed by the sub-detector. Complete protection against SEEs can come at a cost, for example in power consumption. A certain rate of SEUs in the data may be acceptable but must be accurately estimated based on the measurements and proven to have no impact on the efficiency of the detector. For example, frequent corruption of headers in the data packets may require regular re-synchronisation and hence down-time of the system.

SEFIs in the FE must be avoided. These can occur when, for example, a state-machine is upset and takes several clock cycles to recover. This has an impact on the efficiency of the detector.

All configuration memories must be fully protected against SEUs. The return of a system to the operational state after an SEU must not require a re-configuration.

SEL must be avoided as this can permanently damage the component and stresses the power infrastructure.

3.13. Monitoring and diagnostics

Many features can be added to the FE to allow monitoring of critical parameters and diagnosis of problems. Below is a non-exhaustive list of such features that have proven useful in previous system implementations.

1. Quality of master links: This can be tested by generating a pattern (eg fixed, pseudo-random) in the BE and transmitting this over the master link to the FE. The FE must be set into a specific mode to receive the pattern and count any discrepancies. This can be used to make a bit-error-rate test (BERT). This feature is often included in the master link components.
2. Received optical power on master link: Optical receivers often have an output which provides a measure of the received optical power. This output is often called the Received Signal Strength Indicator (RSSI). It can be measured with an ADC.
3. Quality of data links: This can be tested by generating a pattern (eg fixed, pseudo-random) in the FE and transmitting this over the data link to the BE. The BE must be set into a specific mode to receive the pattern and count any discrepancies. This can be used to make a bit-error-rate test (BERT). This feature is often included in the link components.
4. Voltage monitoring: It is often useful to monitor the voltage levels on the FE, particularly if there is little margin in their value for sensitive circuits. These can be measured with an ADC, often together with a voltage-divider.
5. Temperature monitoring: It is often useful to monitor the temperature on or close to the FE. Typical sensors (pt100, pt1000, NTC) can be read with an ADC with an accompanying current source.

4. Back-end

The back-end (BE) electronics are situated in the data-centre or counting room and act as an interface between the FE modules, event-builder-farm, 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 must be implemented in industrial standard formats to allow the use of standardised mechanics, cooling and power supplies.

All LHCb sub-detectors use the same hardware platform for the BE which implements the common interfaces for data transport to and from the FE, to the event-builder, and to the central TFC/ECS systems. The hardware is based on FPGAs and commercial optoelectronic transmitters and receivers. The firmware is also organised in a common structure but with customisable blocks for data processing to be implemented by each sub-detector.

The right hand side of Figure 1 shows the functional blocks of the BE, split between hardware modules for data-acquisition (BE-data) and for controls (BE-controls). These are described in more detail below.

4.1. BE for data-acquisition

As the transfer of data from the FE to BE is essentially asynchronous and with varying latencies, a major task for each BE-data unit is the merging of data fragments received on all its input links into coherent data packets corresponding to individual events. This is the first stage of event-building. Raw data can also be pre-processed if the FPGA resources allow, with the goal of minimising processing time in the later analysis. This could be clustering of hits, for example.

More details for the data-acquisition are shown in the firmware architecture of Figure 7, which is implemented for U1. This shows two identical parallel processing chains optimised for the input and output bandwidth limits of the hardware. A similar architecture will be used for U1b/2. Common blocks are in blue, and red blocks are customisable per sub-detector.

Each block of the BE for data-acquisition is discussed below and in references [2], [3].

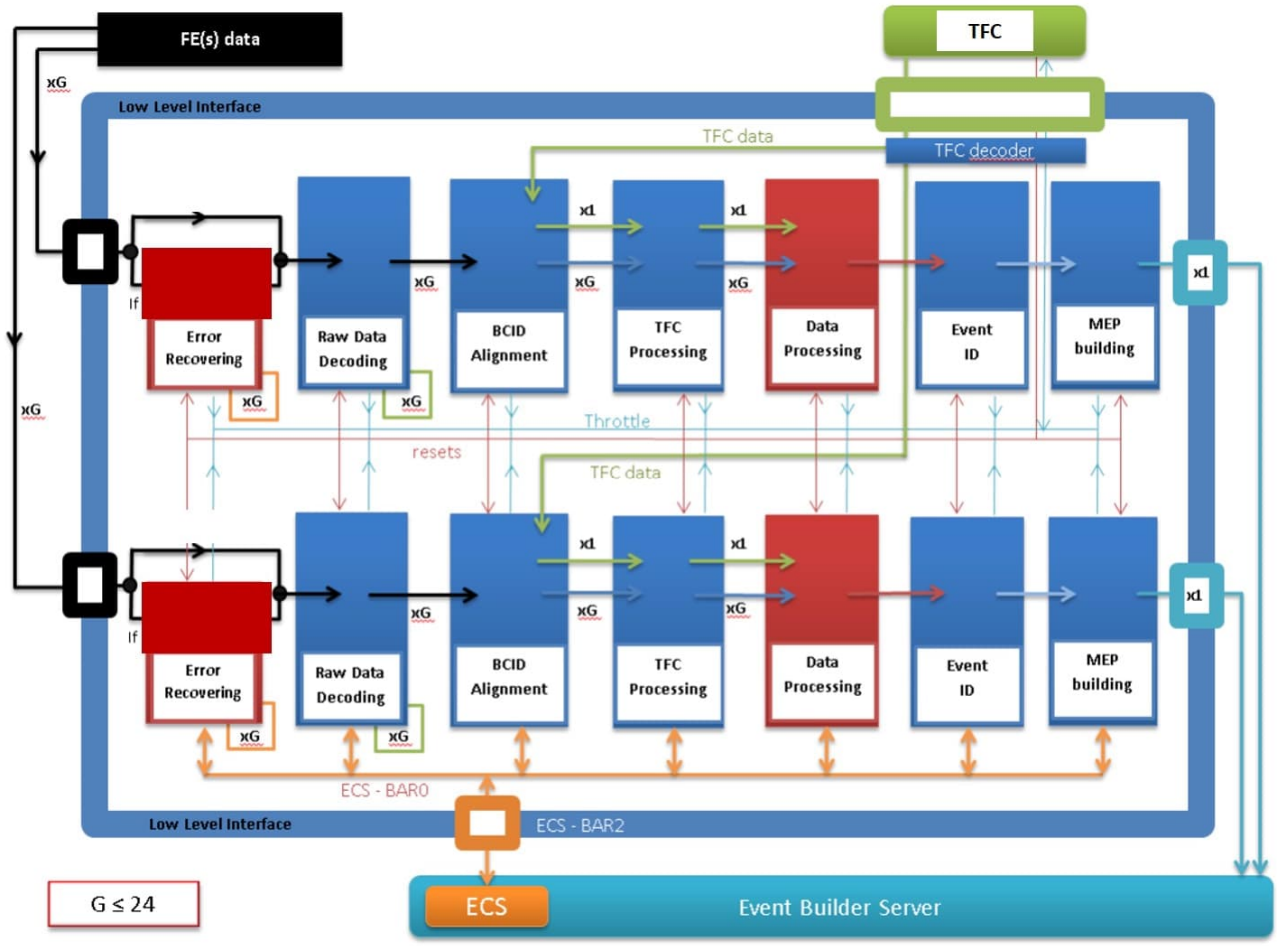

Figure 7: Firmware architecture for the BE data-acquisition in U1.

4.2. Low Level Interface

This includes all the common interfaces to the rest of the system:

1. FE data receiver (black): Receives and deserialises serial data from the optoelectronic components and decodes the low-level link protocol.
2. TFC interface (green): Receives synchronous commands from optoelectronic components connected to the TFC system, and transmits monitoring data.
3. Event-builder interface (blue): Formats and transmits data packets to the hardware interface to the event-builder farm.
4. ECS interface (orange): Receives and transmits configuration and monitoring data from/to the ECS system.

4.2.1. Error recovery, decoding, BXID alignment and TFC processing

The data integrity is checked and errors recovered, if possible. This includes, for example, checking the continuity of the BXIDs received in the headers. Validated data are then decoded and fragments from the different links with matching BXIDs are aggregated are aligned. At this stage, the full 12-bit BXID is recovered if a reduced version was transmitted from the FE.

Information from the TFC decoder is used for this alignment and then in the TFC processing. This will process and filter the data according to commands per BXID. For example, data from a particular BXID can be suppressed if required. This is used during commissioning for time alignment and other such purposes where a limited data throughput is useful.

4.2.2. Data processing

This is specific to each sub-detector. Typical processing tasks are zero-suppression, time re-ordering of data fragments or clustering of hits in physically adjacent channels. Any custom code must use the defined interfaces to the previous and next blocks, and must fit within the available logic resources in the FPGA.

4.2.3. Event ID and MEP building

Processed data are then tagged with an event ID, buffered, and incorporated into multi-event-packets (or multi-fragment-packets) to minimise overheads above the data-payload and hence make optimal use of the available bandwidth. The packets are then transmitted to the event-builder which builds events from fragments delivered from the full BE system.

4.2.4. TFC interface to BE

Fast, synchronous controls are provided by links from the central TFC system, one frame broadcast every LHC clock period. These commands are broadcasts so every BE device will receive the same data. The broadcast command is decoded by the BE module and then executed.

More details of the interface and commands are in [4].

4.2.5. Limits from I/O bandwidth & resources

To minimise the number of BE modules and hence cost, the resources must be used optimally to maximise the amount of FE data each module can process. The limits on data throughput in the BE must be considered when designing and dimensioning the system:

Input bandwidth for data from FE. This is often limited by the physical constraints on the module, for example the front-panel space.

FPGA resources. State-of-the-art FPGAs are well-equipped with transceivers, so the limits on data throughput are instead usually defined by the logic and clock resources. The optimal use of these must considered as part of the design of a sub-detector architecture. For example, the choice of FE data format can have a significant impact on the FPGA resources required for decoding. The

processing steps, such as clustering or time-reordering, must also be considered early and optimised by design choices in both FE and BE.

Output bandwidth for data to the event-builder. This is limited by the choice of interface to the event-builder. In past module implementations, this has defined the throughput of a module and the number of input links per module has been chosen accordingly.

4.3. BE for TFC/ECS distribution

TFC/ECS distribution in LHCb uses the same hardware as the BE for data-acquisition. Each BE-control module communicates with many FE units through point-to-point master links. Data for fast (TFC) and slow (ECS) controls are merged within the BE FPGA and transmitted to the FEs. BE control modules also distribute TFC information to the BE-data modules.

LHCb is sub-divided into partitions, which can operate independently or in common. The independence of these partitions must be maintained at all levels of the system. Partitions are typically a particular sub-detector or a part of a sub-detector.

Figure 8 shows the implementation for U1. A similar architecture will be used for U1b/2. Central control is performed by a module known as SODIN. This communicates directly to the LHC timing interface. In U1, SODIN connects to all BE-control modules through point-to-point links. In this way, SODIN broadcasts the LHC-synchronous clock and TFC commands. In turn, the BE-control modules re-transmit the TFC commands to the FEs after merging with the local ECS data.

Each BE-control module broadcasts the LHC-synchronous clock and TFC commands to a number of BE-data modules. In U1, these broadcasts are through Passive-Optical-Networks (PONs). To facilitate this, the BE-control modules are equipped with special SFP+ transceivers acting as optical line terminals (OLTs) and the BE-data module with SFP+ transceivers acting as optical network units (ONUs).

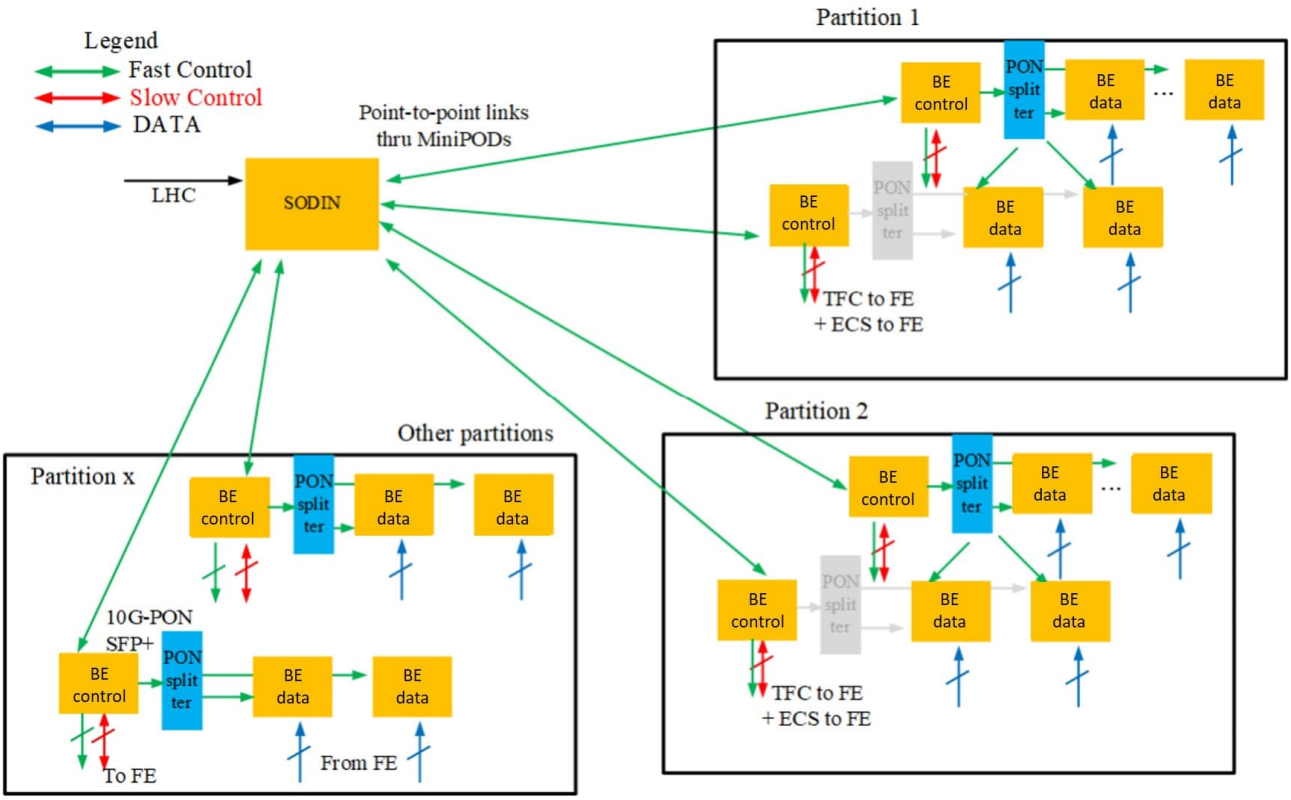

Figure 8: BE system for TFC/ECS distribution in U1.

5. Testing and Commissioning

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

1. Test pulse injection. A pulse is injected into the front-end amplifiers, triggered by a synchronous command issued by the FE TFC interface. The phase and amplitude of this pulse must be adjustable. This technique is principally for detector calibration and commissioning, but can also be used for the bunch-counter synchronisation if path differences in the distribution of the TFC command across the experiment are well known.
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 light source is triggered by a synchronous command from the TFC interface and its phase is adjustable.
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 patterns. The ability to transmit a fixed or deterministic pattern from the FE module is useful for the verification of the links and the receivers of the BE modules. The generation of useful patterns can be implemented in the FE digital logic and they can be injected into the data path in place of the signals from the detector channels.

6. Implementation using common components

Although the very-front-end components of the readout electronics must be custom-designed to meet the specific requirements of the particle detectors, LHCb has successfully followed the philosophy of using common components as standards across all sub-systems. These components are used to implement the data and control optical links, BE systems for data and controls, local power regulation and bulk power supplies.

These common components have been developed through multiple generations matching the timescales of the different LHCb upgrades. The following sections briefly describe these different generations and how they implement the functions in LHCb. This includes on-going developments that have not yet reached production maturity.

6.1. Optical Links

These are designed and qualified to be radiation-tolerant. They are all based on high-speed optical serial transmission (5, 10 Gbit/s or beyond) and can be operated as simplex or duplex links. Each link is a set of electronic and optoelectronic components. For transmission, the electronic components encode the data with a custom protocol including forward-error-correction (FEC), serialise the data, and drive signals to lasers coupled to optical fibres. For reception, the electronic components amplify signals received from photo-diodes, recover clock and data from the serial stream, deserialise the data, and decode the data from the custom protocol after FEC. The FEC enhances the robustness of the link against radiation effects but at the cost of user bandwidth.

The hardware for the FE is accompanied by firmware which implements the links in the FPGA environment of the BE using hard-IPs and core logic.

6.1.1. GBTX and Versatile Link for U1

The GigaBit Transceiver chip-set and Versatile Link are used to implement the links for U1. Details can be found in [6], [7], but the most relevant points for LHCb are summarised here. The components of the chip-set and the versatile link are shown in Figure 9. The trans-impedance amplifier (TIA), PIN diode (PD), laser driver (LD) and laser are mounted together in a bi-directional Small-Form-Pluggable (SFP) Package (VTRx). A dual-transmitter (VTTx) replaces the receiver components with a second transmitter channel.

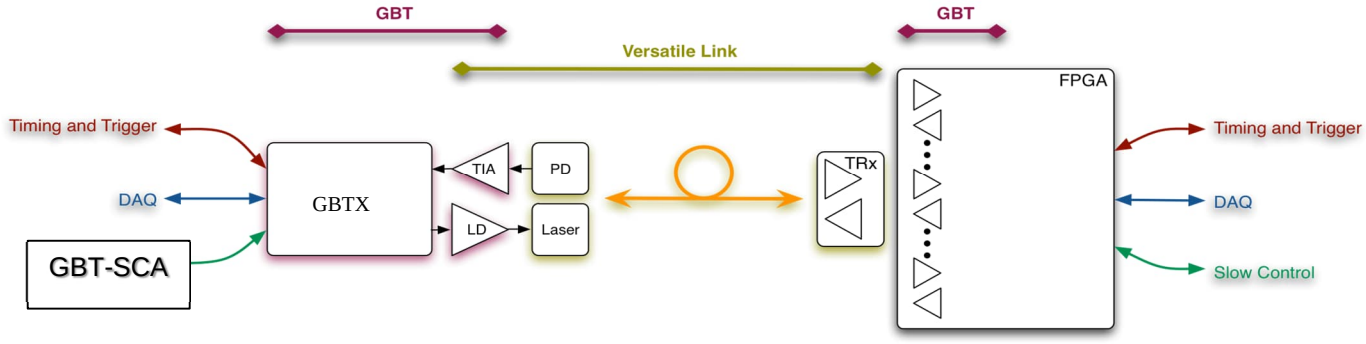

Figure 9: GBT chipset and Versatile Link with a VTRx. The VTTx can be driven by two GBTXs.

The GBTX is the serialiser/deserialiser chip operating at a serial rate of 4.8 Gbit/s in both directions. The data payload is 3.2 Gbit/s with FEC and 4.48 Gbit/s without. Some bits in the data frame are reserved for communication with a slow-control adaptor (the GBT-SCA) which translates into standard protocols such as I2C and JTAG.

The GBTX interfaces to the FE components via eLinks operating at 80, 160 or 320 Mbit/s.

The link can be operated in one of three different modes, two of which are used in LHCb:

1. Transceiver mode with the VTRx: The system is a duplex transceiver used to implement the master link for TFC/ECS. When powered-up, the GBTX auto-configures from built-in electrical fuses.
2. Transmitter mode with the VTTx: The system is a simplex transmitter used to implement the data links. The GBTXs are configured by and receive a reference clock from the master link.

6.1.2. lpGBT and Versatile Link +

These are the next generation of components designed for higher bandwidth and radiation tolerance. Details can be found in [8], [9], but the most relevant points for LHCb are summarised here. The components are shown in Figure 10Figure 9. The trans-impedance amplifier (TIA), PIN diode (PD), 4-channel laser driver (LD) and lasers are mounted together in a customized, compact, low-mass package (VTRx+). This is delivered with a pre-attached fibre pigtail. There is one receiver channel and four transmitter channels. Unused transmitter channels can be disabled by configuration.

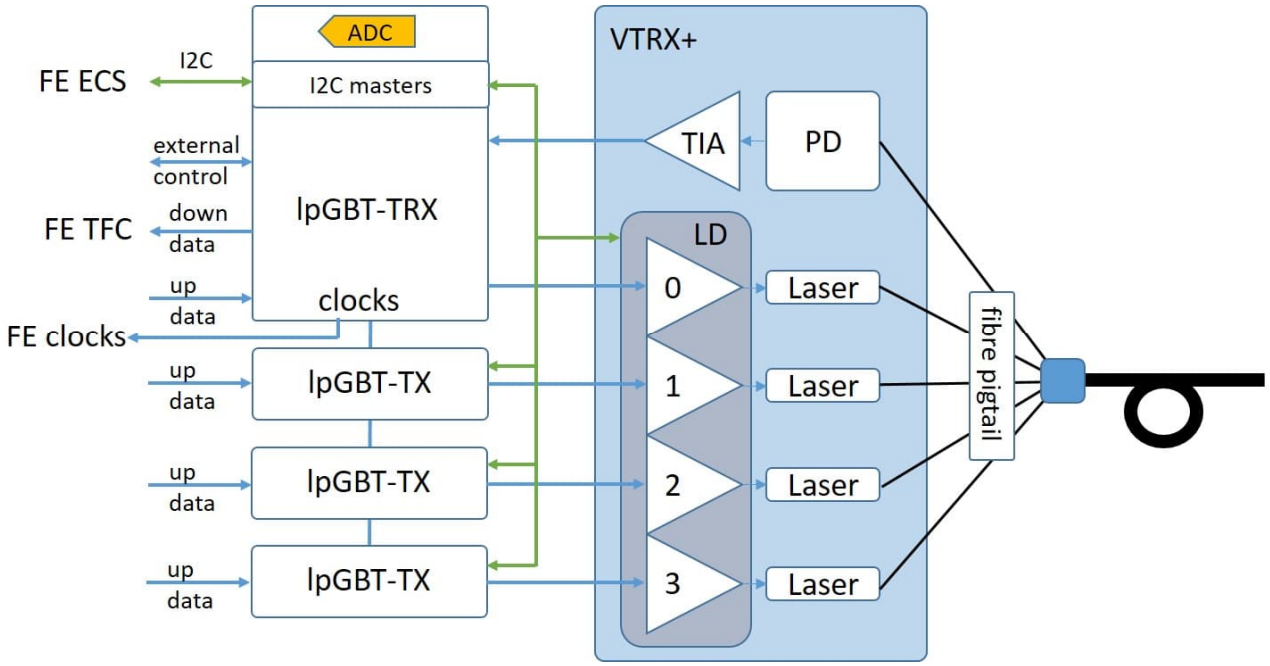

Figure 10: Links implemented using the lpGBT and VTRX+.

The lpGBT is the serialiser/deserialiser chip with the following features:

1. It can be configured as a duplex transceiver, simplex transmitter or simplex receiver. When powered-up, the lpGBT can auto-configure from built-in electrical fuses or a ROM.
2. The receiver (down-link) operates at 2.56 Gbit/s. FEC is implemented and the user bandwidth is 1.28 Gbit/s.
3. The transmitter (up-link) can operate at either 5.12 Gbit/s or 10.24 Gbit/s. Two types of FEC are implemented which can be selected according to the desired level of error-protection at the expense of user bandwidth. The user bandwidth is 3.84 or 4.48 Gbit/s (at 5.12 Gbit/s link-rate), and 7.68 or 8.96 Gbit/s (at 10.24 Gbit/s link-rate).
4. The lpGBT has three independent I2C masters for configuration of other local devices. The lpGBT can also communicate with other local devices through the external-control (EC) interface.
5. The lpGBT also has a general-purpose-IO interface that can be used for simple communication with local devices.
6. The data interface from/to the FE components is over eLinks operating at 80 (down-link only), 160, 320, 640 (up-link only) or 1280 (up-link only) Mbit/s. The lpGBT can generate clocks at 40, 80, 160, 320, 640 and 1280 MHz.
7. The lpGBT has a multi-channel ADC and a DAC.

Figure 10 shows a full implementation with a master (lpGBT-TRX) and three data links (lpGBT-TX). The master link provides a clock reference to the three data links and clocks to the FE. It configures the data lpGBTs and the VTRX+ over an I2C bus. The two remaining I2C masters implement the FE ECS and the eLinks from the master lpGBT are used for the FE TFC interface.

The GBT-SCA chip is no longer available. However, the lpGBT can itself be used a slow-control adaptor connected to a master lpGBT through the EC interface.

The lpGBT operates with 1.2V and the VTRX+ with 1.2V and 2.5V power supplies. The 1.2V defines the swing of the single-ended signals (I2C SDA and SCL, for example) and the common-mode of the differential signals (1.2/2 = 0.6 V).

6.1.3. Future link developments

An overview of on-going developments can be found here [10].

6.2. BE hardware platforms

As described in Section 4, all the functions of the BE are implemented with the same hardware platform. Current and future platforms are described below.

6.2.1. PCIe40 for U1

PCIe40 [11] is a generic module which implements the BE data-acquisition and control functionality. The module consists of a large FPGA, up to 48 bidirectional links with miniPOD optoelectronic modules for communication with the FEs, two bidirectional links with SPF+ form-factor for timing distribution from the TFC system, and a PCIe-Gen3 interface to a computer motherboard which is the interface to the event-builder.

The functionality is selected with different recipes of FPGA firmware. For data-acquisition, the FPGA is configured as a TELL40. For TFC/ECS, the FPGA is configured as a SOL40.

6.2.2. PCIe400 for U1b/U2

In preparation.

6.3. Local power regulation

In the past, linear regulators with low drop-out voltage (LDOs) were mounted on or close to the FE modules. These performed line and load regulation with input power from bulk supplies sited 20m away. Recently, DC-DC convertors have been adopted for this task. These convert high DC voltages (at low current) down to the lower voltages (at higher current) required by the FE. With a lower input current, higher resistance (low-cross-section) cabling can be tolerated which allows a reduction of material in the detector systems. These convertors are designed and qualified to be radiation and magnetic-field tolerant but in general require shielding to protect against the electromagnetic emissions they generate. All details of DC-DC convertors for current and future applications can be found in [12], and the main points are summarized below.

6.3.1. FEASTMP module for U1

These are DC-DC convertor modules integrating an ASIC (FEAST) and passive components on a small circuit board. The module can create electromagnetic interference so it is shielded with a metallized cover. The output voltage is set by the value of a resistor mounted on the module. The modules have an enable input and in LHCb U1 they will operate with an input voltage of 6 – 8V. The FEAST ASIC and hence the module are now obsolete and cannot be used beyond LHCb U1.

6.3.2. bPOL family

This is the current generation of DC-DC convertors with enhanced radiation tolerance beyond that of FEAST. These are delivered as packaged ASICs rather than modules. Hence they must be integrated directly onto the FE modules together with passive components, shielding and an efficient cooling mechanism. The bPOL12V ASIC is the most widely used and operates with an input voltage in the range 5.5 to 12V. It can source up to 4A with output voltage in the range 0.63 to 5V. The output voltage is selected by the value of a resistor mounted close to the convertor ASIC. The ASIC has an enable input. bPOL12V is designed to work with an air-core inductor so the circuit will tolerate a magnetic field of 4T. The ASIC has been tested to withstand radiation levels up to 150 Mrad and 4 x 10^{15}/cm^2 1 MeV-neutron equivalent fluence. Other ASICs exist in the bPOL family although bPOL12V is best suited to the LHCb applications in U1b and U2.

6.3.3. Future convertor developments

An overview of on-going developments can be found here [10].

6.4. Power supplies

Low and high-voltage power are generated by bulk power supplies situated outside the detector volume. There are a number of locations in the LHCb cavern where such supplies can be situated. These areas have relatively modest (but not negligible) levels of radiation and magnetic-field strength and hence can be used if the power supply is proven to tolerate this environment. This limits the cabling length to 10 – 20 metres, which is a big advantage for high current lines requiring large-cross-sections. Otherwise, power supplies must be situated in the D2 or D3 barracks behind the concrete shielding wall where there is neither radiation nor magnetic-field. Cable lengths to the detector are then 60 – 80 metres.

6.4.1. Power supplies for U1

The majority of low-voltage power supplies for U1 are Maratons manufactured by Wiener. This is a modular system, with a water-cooled front-end DC-DC convertor which is radiation and magnetic-field tolerant. These are situated in racks on the balcony beside the magnet or in the bunker below RICH2. Other modules in the system convert AC to DC power and provide control and monitoring. These are situated in D3 with cables connecting to the front-end convertor. Each channel of a Maraton supply typically provides 2 – 8V with maximum 50A.

High voltage power is mostly supplied by equipment from CAEN and ISEG. These are all situated in racks in D3.

6.4.2. Future power supply developments

In preparation.

7. Use of programmable commercial devices in FE modules

Data compression and formatting must be implemented in the FE electronics. In many cases, this requires complex algorithms that depend heavily on the expected occupancy of the detector. It is advantageous to have the freedom to change or tune these algorithms and using programmable devices allows such an optimization.

Such a device must 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. If the use of SEU-sensitive configuration memory cannot be avoided, the rate of upsets must be determined through a dedicated test campaign and its impact on the detector efficiency then evaluated. Appropriate steps must be put in place to mitigate down-time caused by corrupted configuration, such as continuous memory-scrubbing of the device.
4. The device must resist destructive single-event-effects, such as SEL.

Devices employing flash configuration memories have been shown to resist SEUs thanks to the inherent structure of the memory bits. SRAMs are more sensitive. However, many FPGAs now incorporate built-in circuits for error detection and correction.

Rigorous radiation testing is mandatory. As these are complex programmable devices with many features, radiation testing must be carried out with final firmware to guarantee a correct evaluation of their sensitivity.

8. References

[1] Framework TDR for the LHCb Upgrade II, CERN-LHCC-2021-012.

[2] G. Vouters et al., “Front-End and Back-End Data Format of the LHCb Upgrade”, LHCb Internal Note, LHCb-INT-2014-011, February 2014.

[3] F. Alessio et al., “Readout protocol for TELL40”, https://edms.cern.ch/document/1606939/1, April 2016.

[4] F. Alessio et al., “Readout Control Specifications for the Front-End and Back-End of the LHCb Upgrade”, CERN-LHCb-PUB-2012-017, July 2012.

[5] https://espace.cern.ch/HighPrecisionTiming

[6] https://espace.cern.ch/GBT-Project/default.aspx

[7] https://espace.cern.ch/project-versatile-link/public/default.aspx

[8] https://espace.cern.ch/GBT-Project/LpGBT

[9] https://espace.cern.ch/project-Versatile-Link-Plus

[10] LHCb upgrade electronics workshop, Nov 2020. https://indico.cern.ch/event/960837/

[11] J. P. Cachemiche et al., The PCIe-based readout system for the LHCb experiment, JINST 11(2016) P02013.

[12] https://espace.cern.ch/project-DCDC-new