

Testing, time alignment, calibration and monitoring features in the LHCb front-end electronics and DAQ interface

J. Christiansen

CERN, 1211 Geneva 23, Switzerland

Jorgen.Christiansen@cern.ch

Abstract

An overview of time alignment, testing, calibration and monitoring features in the front-end electronics of LHCb is given. General features for this are defined and examples are given of how this has been implemented in the LHCb front-end electronics and DAQ interface.

I. INTRODUCTION

A sufficient level of timing alignment, monitoring and built-in testing features in the front-end electronics system of a large scale experiment will be vital during the different phases of a HEP experiment to obtain a reliably working system. The electronics systems of a large experiment consist of thousands of complicated modules interconnected by thousands of communication interfaces (LHCb: ~7.000 optical links). It will be a significant challenge to get all these different types of modules (LHCb: ~50 different board types for a total of ~25.000 modules) to work correctly together during a long time period (10 years) in a hostile environment (noise, magnetic fields, radiation, etc.). Modules used to build a local sub-system will be tested together in a well controlled lab. environment beforehand but may still show problems when installed in the final “hostile” environment. The global integration of individual sub-systems will in many cases not occur before the whole experiment has been fully installed and final commissioning is started. Calibration and monitoring features are needed to continuously verify that the systems work correctly with the required precision. Extensive test and debugging features are needed during the initial commissioning phase and also to perform long term maintenance and repairs [3].

II. FRONT-END AND DAQ ARCHITECTURE

The LHCb front-end architecture [1,2] shown in Figure 1 consists of analogue front-ends for analogue data treatment followed by a 160 clock cycles deep first level pipeline buffer (called L0 in LHCb). All LHCb front-ends have programmable length latency buffers so local adjustments of this can be made if needed. L0 trigger accepted events are stored in a 16 deep derandomizer buffer before readout to the DAQ system at an average rate of up to 1 MHz. Event data is at this level transferred on optical links (with the exception of the Vertex detector using multiplexed analogue copper links) to a DAQ interface located in the counting house. Each event block consists of 32 words of detector data and 2 – 3 words of event header information and finally an event separator adding up to a maximum event readout time of 900ns. The FPGA based DAQ interface module verifies the received event fragments and perform sub-detector specific zero-suppression and/or data compression before sending it to the

DAQ system. The DAQ system is based on a large CPU farm (~ 2000 CPUs) and a large Gigabit Ethernet based readout network as shown in Figure 2. Multiple event fragments (8 – 16) are merged into Multi Event Packets (MEP) in each DAQ interface module to assure good GBE link utilization at the high LHCb trigger rate. Each DAQ interface module has four GBE outputs to cope with the high readout rate (total LHCb data bandwidth: ~0.5Tbit/s). The front-end, trigger and DAQ systems are controlled from the Experiment Control System (ECS).

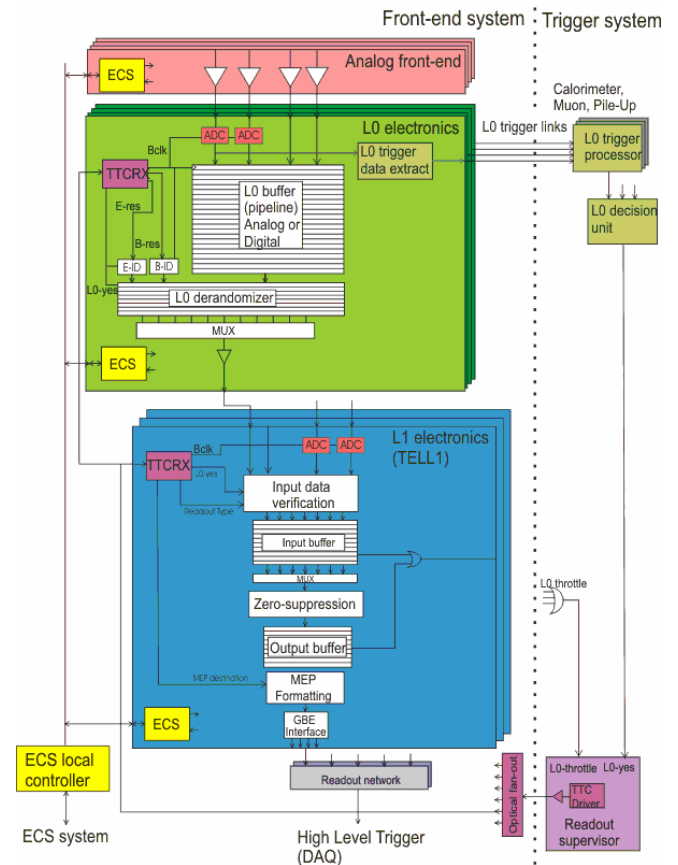


Figure 1. LHCb front-end architecture.

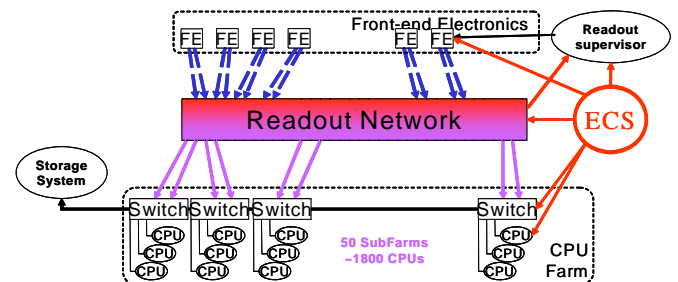


Figure 2. LHCb DAQ architecture

A. Timing and Fast control

Timing and fast control signals are distributed to all front-end electronics and DAQ interface modules via the TTC system based on optical fibre fan-outs and the TTCrx receiver chip. The TTC fibre distribution is driven from a bank of up to 16 readout supervisor modules that generates the necessary timing and trigger signals to the front-ends in a way that assures that all front-ends maintain full synchronization (e.g. applying restrictions to trigger preventing buffer overflows). Each of the available readout supervisors can drive local partitions or the global experiment via a programmable TTC switch/fan-out. LHCb has 16 TTC partitions that can be used to test and commission individual sub-detectors independently. During normal physics running only one readout supervisor drives the whole LHCb front-end and readout system. This implies that all the different sub-systems must interpret all TTC signals and broadcast messages in a unified fashion across the whole system. This is different from some of the other LHC experiments where each partition is controlled from individual TTC controllers, even when working as one global system. Such an approach can in principle allow sending different signals and broadcasts to different partitions but this must be done with great care to assure that different partitions are maintained fully synchronized to each other

III. FRONT-ENDS

Specific features in the LHCb front-end architecture related to timing alignment, calibration, testing and monitoring are indicated in Figure 3 below.

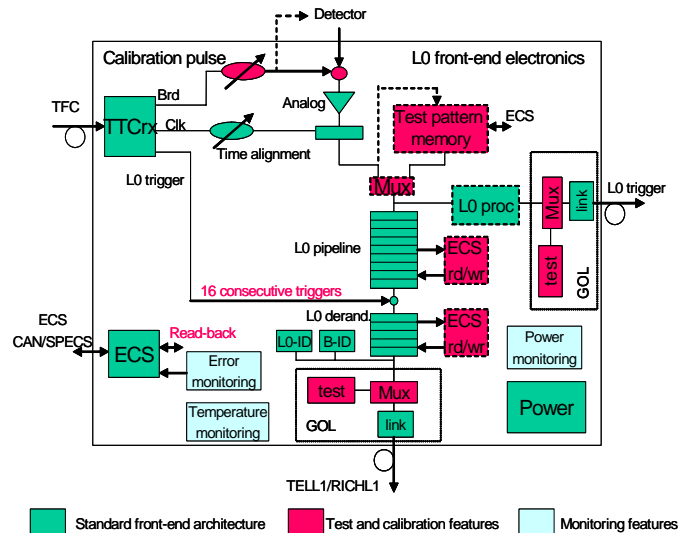


Figure 3. Test and monitoring features in the front-end.

Major features related to time alignment, calibration, testing and monitoring are described below. Some sub-detectors have included additional testing and debugging features which allows the Experiment Control System (ECS) to have direct read/write access to the L0 latency buffer and/or the derandomizer buffer. ECS read access to the derandomizer buffer has allowed some sub-systems to

perform extensive front-end module tests via the ECS interface alone (no DAQ system needed).

A. Timing alignment

All sub-detectors have basic time alignment features based on the programmable delays of the TTCrx chip to capture detector signals with the correct phase and in the correct bunch cycle. The LHCb experiment located in an underground cavern and with a horizontal orientation can unfortunately not rely on cosmic muons to perform initial time alignment and detector alignment as the rate of muons traversing multiple sub-detectors in LHCb is extremely low. LHCb therefore depends critically on real beam collisions (or beam gas events) to perform the global time alignment. For small sub-detector systems (Vertex, Inner tracker, etc.), or local regions of large detector systems, only very small local time differences between channels will exist and can therefore to a first approximation be assumed to be time aligned between local channels.

To perform time alignments on real beam interactions a basic interaction trigger will be needed that is internally well time aligned from the beginning. In LHCb such an interaction trigger will be made from the Hadron calorimeter that has an internal time alignment system based on LED light pulse injection into its PMT's via known lengths of fibre. Global fine time alignment will be based on this trigger using detailed timing histograms to obtain a fully time aligned experiment. This will require a significant number of real interactions and software tools to obtain time alignment will be critical to minimize the use of sparse beam time available in the initial running of LHC during 2007 – 2008. The initial verification without LHC beams of such a time alignment system can be made using the basic pulse injection scheme with programmable delays as described below.

B. Calibration pulse injection

A common calibration pulse injection scheme has been defined across all sub-detectors based on a short TTC broadcast message [1]. The message encoding used allows up to four different types of calibration pulses. One of these calibration pulse types has been defined as a common type with a pre-defined timing in relation to the corresponding trigger accept as shown in Figure 4 below.




Figure 4. Common calibration pulse injection timing.

The three remaining calibration pulse injection types are available for sub-detector specific use if needed during local testing and commissioning. The corresponding trigger to the common calibration pulse has been defined to have a delay of 16 clock cycles plus the L0 trigger latency (160 clocks). This allows all sub-detectors to have enough time in the front-ends for the generated calibration pulse to enter their pipeline buffers for correct extraction with the corresponding trigger accept. All LHCb sub-detectors have implemented calibration pulse injection schemes for all detector channels based on the common type. This allows the common calibration pulse injection to be used globally in the whole LHCb system. The detailed characteristics of the calibration pulse injection is very sub-detector dependent caused by the specific nature of each sub-detector (Light injection in detector, pulse injection in analogue front-end, fixed or variable amplitude, etc.). All sub-detectors can adjust the time of the pulse injection locally over a dynamic range of 16 clock cycles allowing to verify that timing alignment features (hardware and software) is correctly working before having final beam collisions.

Certain sub-detectors rely on the use of the common calibration pulse type injection during normal physics running to closely monitor the gain and stability of the detector and its analogue front-end electronics (e.g. calorimeters). Only a limited number of detector channels can be exercised per calibration pulse injection to correctly monitor the detector performance without problematic crosstalk effects. Such events can if needed skip local zero-suppression in certain sub-detectors. Local round robin schemes have been implemented for such detectors to scan all detector channels within a reasonable time window (few minutes).

C. Consecutive triggers

LHCb has a first level trigger accept rate of up to 1.1MHz which is a factor ~ 10 higher than ATLAS/CMS. The high trigger rate is caused by the difficulty of making efficient hardwired trigger systems for B physics at LHC. This has implied the use of a relative large derandomizer buffer in the front-ends (16 compared to ~ 4 in other LHC experiments) and the use of a significant number of radiation hard optical links to transport acquired data from the experiment to the counting house. At this high trigger rate enforced spacings between trigger accepts results in a $\sim 2.5\%$ loss in physics per enforced gap. It was also in the early phase of LHCb considered very useful to trigger a whole sequence of consecutive triggers (max 16) to have an efficient way to perform a first coarse time alignment between channels and also directly measure pulse width, spill-over and baseline shift effects as indicated in Figure 5.

Handling of consecutive triggers is particular difficult for detectors with relative long pulse shapes and detectors with drift times covering several bunch crossing periods. Specific features in the front-end electronics of such sub-detectors have though allowed consecutive triggers to be supported (calorimeter and outer tracker). During recent system tests it has though unfortunately been discovered that one of the LHCb sub-detector front-end ASIC's does not handle correctly consecutive triggers. Consecutive triggers are still planned to be used extensively during first steps of local and global commissioning tests with dedicated software tools to

acquire data and analysis it. The specific sub-detector requiring a single gap between triggers will have to handle such commissioning tests in a dedicated manner. Final physics running will be made with a minimum trigger gap of one to assure correct readout across all sub-detectors.

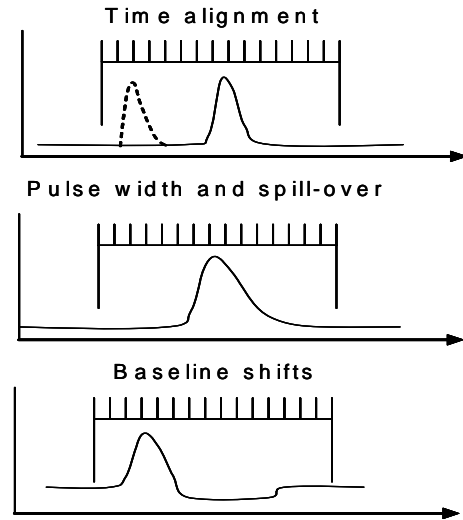


Figure 5. Use of consecutive triggers

D. Synchronization and Data monitoring

To assure correct data taking during extended physics running it is extremely important that the whole front-end and readout system is synchronized. A local de-synchronization could pass unnoticed for extended periods if the front-end and trigger systems are not capable of continuously verify their correct synchronization. In the first level trigger system, and in the extraction of accepted events, the systems must be perfectly synchronized at the clock level. For the interface to the DAQ system it must be assured that only event fragments from the same event are merged in the global event building. Careful monitoring of this is particular needed when the front-end electronics is located in locations with high radiation levels and may get de-synchronized by single event upsets (SEU).

For communication interfaces (e.g. optical links) in the L0 trigger systems specific idle/synchronization patterns are introduced each machine cycle to allow local systems to verify their synchronization and re-synchronize if needed. In addition all trigger data carries a few bits of bunch ID information as indicated in the figure below.




Figure 6. Idle/synchronization pattern and bunch ID information on trigger links.

For the readout path it is enforced that all event readout after the first level trigger accept must have an event header with bunch ID and Event ID information. This allows the DAQ interface to check the correct event synchronization of each data source, comparing the received information with reference information received from its TTC interface (this part not in radiation area so no SEU effects). To ensure event synchronization even in case of bit or word errors on readout links, it is enforced to transmit an idle/synchronization pattern between each event fragment.




Figure 7. Event tagging and event separation.

It has been verified with radiation tests that trigger and readout links can re-synchronize on a single word idle/synchronization pattern when the transmitter or receiver PLL's have not lost frequency lock (word consisting of four 8B/10B idle characters). Tests have been made to confirm that the transmitter and receiver PLL's do not lose lock from single bit or word errors.

E. Pattern and spy memories

In sub-detector systems involved in the first level trigger (calorimeter, muon and pileup veto) testing features based on built in pattern generation and pattern acquisition (spy) memories are used. Specific hit patterns can be written from ECS to pattern generation memories in the front-ends as indicated in Figure 4. These hit patterns can then be applied in a pattern by pattern scheme trigger by TTC broadcasts or can be used in a continuous circular fashion in special testing modes. Spy memories at the output of the local trigger systems and at key internal locations can capture the detailed response of the systems to verify their correct function or to determine the cause of mal functions.

F. Optical links

~7000 optical links are used in LHCb for first level trigger systems and for data readout after the first level trigger accept. All link transmitters on the detector are based on the radiation hard serializer chip GOL from the CERN Microelectronics group [4]. The global reliability of the LHCb experiment will depend critically on the reliability and error rates of these links. Assuming the commonly accepted Bit Error Rate (BER) of 10^{-12} - 10^{-13} for optical links, the total system will have transmission errors at the rate of 1-10 per second. It is therefore clear that the effective error rate of the links must be much lower than this and that basic link transmission errors must not generate problems at the system level. The enforced use of regular idle/synchronization patterns ensures that the sub-systems can remain functional and recover by them self

within very short time intervals (e.g. single event fragment corrupted or part of L0 trigger system not fully functional until next machine cycle).

The GOL chip itself is assumed SEU immune as it uses full internal triple redundant logic. It has been seen that 12way optical transmitters used in some of the LHCb sub-detectors has a small SEU upset rate caused by some simple internal circuitry. The use of regular idle/synchronization patterns have been demonstrated to resolve this problem in a fully satisfactory fashion [5].

To assure that all optical link receivers, transmitters and installed fibres are working in a fully satisfactory fashion it is required that all links are verified to work with a BER below 10^{-12} (test takes ~10 min) when an additional 6db optical attenuation is inserted (in addition to fibre and patch panel losses). This has been seen to assure that the optical links work with a BER rate lower than what can practically be measured. Measurements of BER with 9 and 12db optical attenuation are also required in the design qualification of transmitter and receiver modules [6]

It is required that all optical links must be capable of measuring BER in situ. A simple pattern generator function in the GOL allows to send a continuous counting test pattern and all receiver modules are required to have a simple pattern verification and error counting function on all optical link inputs.

During normal operation single bit errors (that translates to the loss of a whole word) are normally detected by the 8B/10B encoding used on the link and word resynchronization is a previously mentioned in section D obtained by the use of regular idle/synchronization patterns.

IV. DAQ INTERFACE

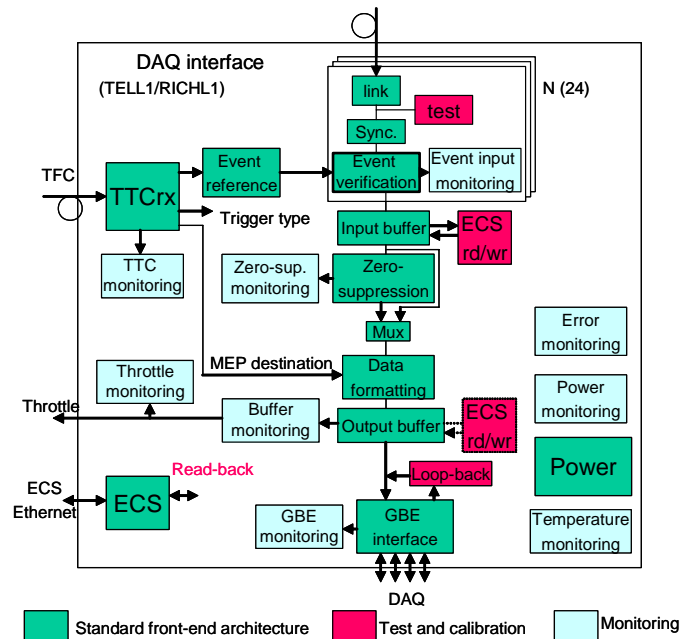


Figure 8. Data flow in DAQ interface.

A generic data flow model of the DAQ interface is shown above with the marking of specific testing and calibration features and general monitoring. The DAQ interface is implemented with the common module called the TELL1 [7]

for all sub-detectors, except the RICH detector that have chosen to make a dedicated module. The TELL1 is based on power full FPGA's with a generic firmware VHDL framework handling the global data flow, buffering and system interfaces. Only the zero-suppression is sub-detector specific as it depends strongly on the type of detector data.

Received event data are on each input verified against an event reference from the Local TTCrx receiver, as described in paragraph III.D, before being passed to the zero-suppression processing. Zero-suppression can be disabled in a static fashion from the ECS or can be disabled in a dynamic fashion for specific event types (e.g. calibration events). It is also possible for specific trigger types to read out both non zero-suppressed data and zero-suppressed data to allow extensive verifications of the zero-suppression function. Event data is then buffered and a given number of events are merged into a Multi Event Package (MEP) to be sent to the DAQ system. The destination address of the MEP is received via a long TTC broadcast message that also contains a few bits of event identification verification information. MEP's are finally sent to the DAQ processing farm over a large GBE based readout network via a standardized GBE plug-in card with 4 GBE ports [8].

A. Testing and debugging features

Extensive testing and debugging features are built into the generic VHDL framework and are therefore available across all sub-detectors in a unified fashion. The simple BER test function is available in all optical inputs as previously described. Raw received input data can in dedicated testing modes be accessed via the ECS interface and emulated raw data can also be inserted to allow the full processing of the board to be verified. The zero-suppression can be enabled or disabled in a static or dynamic form as described above. Fully formatted MEP packets in the output buffer can be accessed from the ECS interface for verification or for injecting specific MEP packets for system testing. The quad GBE interface plug-in card, based on commercial MAC and PHY chips, can make data loops at multiple levels between incoming and outgoing GBE traffic on each port. In addition a LHCb specific packet mirroring function is implemented in the FPGA's that allow a received encapsulated MEP packet to be retransmitted to any of its four ports to a destination defined in the packet itself. This packet mirroring allows extensive verification of the DAQ readout network at high rates driven directly by the CPU farm itself.

B. Monitoring and local ECS interface

Data, event and error monitoring counters are used extensively throughout the data flow to allow detailed monitoring and tracing of system failures. The number of events and data words that have passed key locations in the data flow is counted (input links, TTC messages, output of zero-suppression, data buffering and output to the readout network). Sub-detector specific monitoring of the zero-suppression is included according to the needs of each sub-detector.

General monitoring, testing and debugging of the TELL1 module are performed via the ECS interface based on an

small plug-in credit card PC running Linux. Direct access to all control and monitoring features of the board is made via a DIM server on the CC-PC. Local running monitoring programs can also perform local intelligent monitoring and built locally histograms without loading the global ECS system with such trivial tasks. Local monitoring tasks are though in general not yet well defined but several sub-detectors plan to use such features.

V. SUB-DETECTOR SPECIFICS

Main features for time alignment, calibration, and monitoring have been defined globally but large differences exist between sub-detectors in the detailed function of these. Many of these differences have been dictated by specific features of the individual detector technologies (e.g. binary or analogue readout) and some have been determined by design choices for the front-end electronics (e.g. ECS access to pipeline and derandomizer buffers). Such implementation choices normally have very good justifications seen from the local sub-system point of view, but may in some cases be impractical at the global system level. Some sub-systems have implemented specific features to allow the pulse injection to be usable during normal physics running were other systems have been made such that this will be difficult or impractical. An example of this is the question of how the pulse injection is made. In the Beetle silicon strip front-end chip [9] the calibration pulse injection has opposite polarity between neighbour channels and changes polarity every second calibration pulse. This calibration pulse injection scheme has been found useful for the detailed characterization of the front-end chip, but must be considered impractical at the system level to generate specific hit patterns on channels to test and verify the Vertex pileup veto trigger system.

The muon detector has a particular problem with low hit rates and a relative large fraction of the hits being out of time background hits. To ensure efficient time alignment of all muon detector channels dedicated time histogram memories have been built into the TDC front-end ASIC.

VI. CONCLUSIONS AND LEASSONS LEARNT

Features for testing, time alignment, calibration and monitoring have been defined to have a global set of tools to be used during commissioning, running and maintenance of the LHCb experiment. It is important to define such features as early as possible to allow the necessary support for it to be built into the front-end electronics of each sub-system. Such global features must though also have sufficient flexibility to cover the specific characteristics of the different sub-detectors and to cover potential problems and imperfections in all the different sub-systems. Key features must be strongly enforced across all the sub-systems, to allow efficient system running, even in cases where local sub-systems insist that they do not see the need for such features.

It was in LHCb realized relatively late that optical links could become a major system reliability problem. Fortunately an un-documented pattern generation feature of the GOL link serializer was identified which is now part of a standardized link test and qualification procedure.

All testing, calibration and monitoring features need extensive software support to allow its efficient use at the global system level. Much of this software still needs to be finalized at the local and global level. Final verification of all the required features across sub-detectors will to a large extent not be made before starting the final commissioning of the experiment.

VII. REFERENCES

- [1] LHCb-2001-014, Requirements to the L0 front-end electronics:
<http://documents.cern.ch/archive/electronic/cern/others/LHB/public/lhcb-2001-014.pdf>
- [2] EDMS 715154, Requirements to the L1 front-end electronics: <https://edms.cern.ch/document/715154>
- [3] EDMS 692583, Test, time alignment, calibration and monitoring in the LHCb front-end electronics:
<https://edms.cern.ch/document/692583>
- [4] GOL chip: <http://proj-gol.web.cern.ch/proj-gol>
- [5] EDMS 486833, Detailed Specification of the ODE-Muon Trigger interface: <https://edms.cern.ch/comment/486833>
- [6] EDMS 680438, Qualification of the optical links for the data readout in LHCb: <https://edms.cern.ch/document/680438>
- [7] TELL1 web page: <http://lphe.epfl.ch/~ghaefeli/>
- [8] EDMS 520885, Quad gigabit Ethernet plug-in card:
<https://edms.cern.ch/document/520885>
- [9] Beetle reference manual:
http://wwwasic.kip.uni-heidelberg.de/lhcb/Publications/BeetleRefMan_v1_3.pdf