LHCb and its electronics.

J. Christiansen, CERN
On behalf of the LHCb collaboration
jorgen.christiansen@cern.ch

Abstract

The general architecture of the electronics systems in the LHCb experiment is described with special emphasis on differences to similar systems found in the ATLAS and CMS experiments. A brief physics background and description of the experiment are given to understand the basic differences in the architecture of the electronics systems. The current status of the electronics and its evolution since the presentation of LHCb at LEB97 is given and critical points are identified which will be important for the final implementation.

I. PHYSICS BACKGROUND

LHCb is a CP (Charge & Parity) violation experiment that will study subtle differences in the decays of B hadrons. This can help explaining the dominance of matter over antimatter in the universe. B hadrons are characterised by very short lifetimes of the order of a pico-second, resulting in decay lengths of the order of a few mm. A typical B Hadron event is shown in Fig. 1 and Fig. 2 to illustrate the types of events that must be handled by the front-end and DAQ systems in LHCb. A typical B Hadron event contains 40 tracks inside the detector coverage close to the interaction point and up to 400 tracks further down stream.
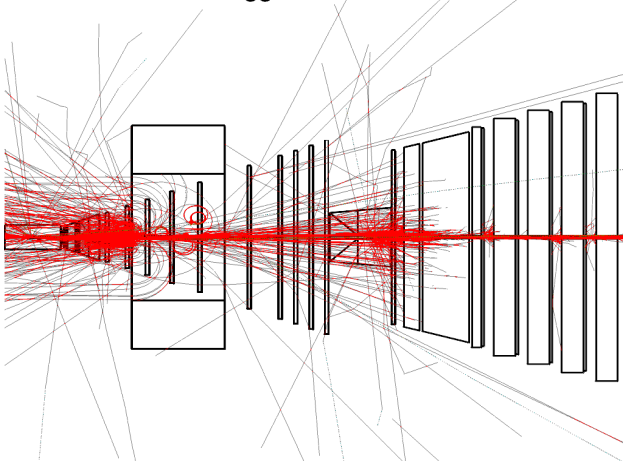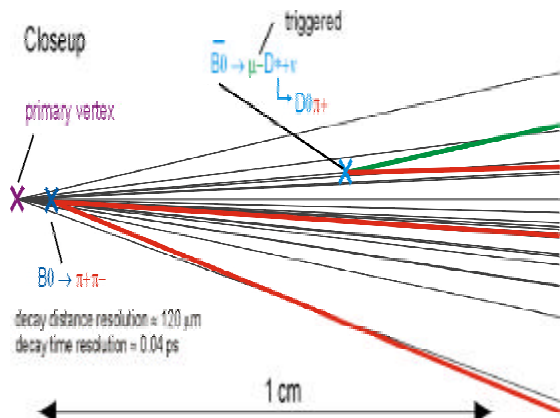

Figure 1: Typical B event in LHCb

Figure 2: Close-up of typical B event

II. MAJOR DIFFERENCE FROM CMS AND ATLAS

The LHCb experiment is comparable in size to the existing LEP experiments and it is limited in size by the use of the existing DELPHI cavern. The size, the budget and the number of collaborators in LHCb are of the order of 1/4 of what is seen in ATLAS and CMS. It consists of ~1.2 million detector channels distributed among 9 different types of sub-detectors. Precise measurements of B decays, close to the interaction point, requires the use of a special Vertex detector, located in a secondary LHC machine vacuum a few cm from the interaction point and the LHC proton beams. The need for very good particle identification requires the inclusion of two Ring Imaging CHerenkov (RICH) detectors.

The layout of sub-detectors resembles a fixed target experiment with layers of detectors, one after the other as shown in Fig. 3 and Fig. 4. This layout of detectors, which can be opened as drawers to the sides, ensures relatively easy access to the sub-detectors, compared to the enclosed geometry of ATLAS and CMS.

B hadrons are abundantly produced at LHC with a rate of the order of 100kHz. Efficient triggering on selected B hadron decays is though especially difficult. This has enforced a four level trigger architecture, where the buffering of data during the two first trigger levels is taken care of in the front-end. The first level trigger, named L0, has been defined with a 4.0 us latency (ATLAS/CMS: 2.5 us and 3.2 us) and an accept rate of 1 MHz (ATLAS/CMS: 50 – 100 kHz). To obtain this trigger rate, in the hardware driven first level trigger system, it has been required to limit the interaction rate to one in three bunch crossings to insure "clean events" with single interactions (ATLAS/CMS: ~30 interactions per

bunch crossing). It has even been required to have a special veto mechanism in the trigger system to prevent multiple interactions to saturate the available trigger bandwidth. The difficulty of triggering on B events can be illustrated by the fact that the first level trigger in LHCb is 3x30x1MHz/100KHz = 1000 times more likely to accept an interaction than what is seen in ATLAS/CMS. The high trigger rate has forced a tight definition of the amount of data that can be extracted for each trigger, and made it important to be capable of accepting triggers in consecutive bunch crossings (ATLAS/CMS: gap of 3 or more). It is also necessary to buffer event data during the second level trigger in the front-end electronics, to prevent moving large amounts of data over long distances (ATLAS/CMS: Only one trigger level in front-end).
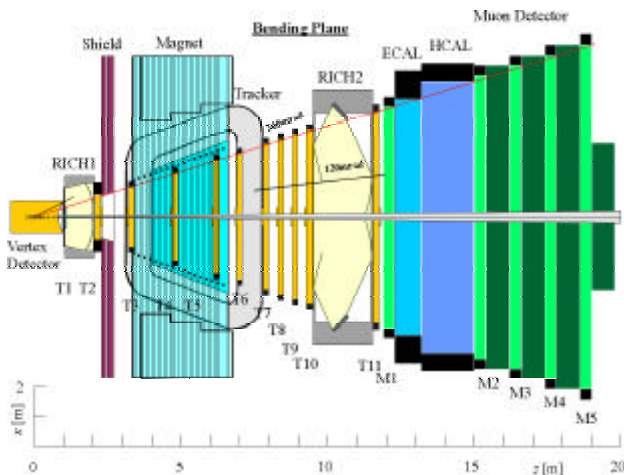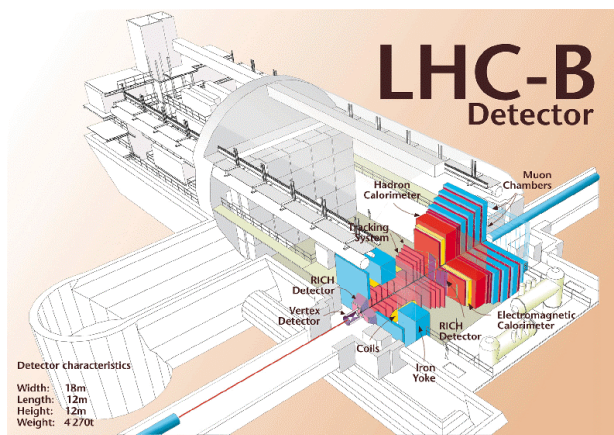

Figure 3: Configuration of sub-detectors in LHCb

Figure 4: LHCb detector in DELPHI cavern

III.   LHCb EVOLUTION SINCE LEB97

Since the presentation of the LHCb experiment at the LEB97 workshop in London significant progress has taken place. LHCb was only officially accepted as a LHC experiment in September 1998. The main architecture of the experiment and its electronics has been maintained and most detector technologies have now been defined.

Detailed studies of the triggering have shown that trigger latencies were required to be prolonged. A L0

trigger latency expansion from 3.0 μs to 4.0 μs was found appropriate, as compatibility with existing ATLAS and CMS front-end implementations was not an issue. The L1 trigger processing was found to be more delicate and sensitive to background rates than initially expected. The latency has been prolonged from 50 μs to 1000 μs, as the cost of additional memory was found to be insignificant. The Architecture of the trigger implementations has now been chosen after studying several alternative approaches.

The two levels of triggering in the front-end, with high accept rates, has called for a tight definition of buffer control and buffer overflow prevention schemes that work across the whole experiment. For the derandomizer buffer, related to the first level trigger, a scheme based on a central emulation of the occupancy has been adopted. For the second level triggering and the DAQ system an approach based on trigger throttling has been chosen, as data here in most cases will be zero-suppressed and therefore can not be predicted centrally.

In a complicated experiment the front-end, trigger and DAQ systems rely on a partitioning system to perform commissioning, testing, debugging and calibration. A flexible partitioning system has been defined such that each sub-system can run independently or be clustered together in groups.

For the DAQ system a push based event building network, distributing event data to the DAQ processing farm, has been maintained after simulating several different event-building schemes on alternative network architectures.

IV.   FRONT-END AND DAQ ARCHITECTURE

LHCb has a traditional front-end and DAQ architecture with multiple levels of triggering and data buffering as illustrated in the figure below.
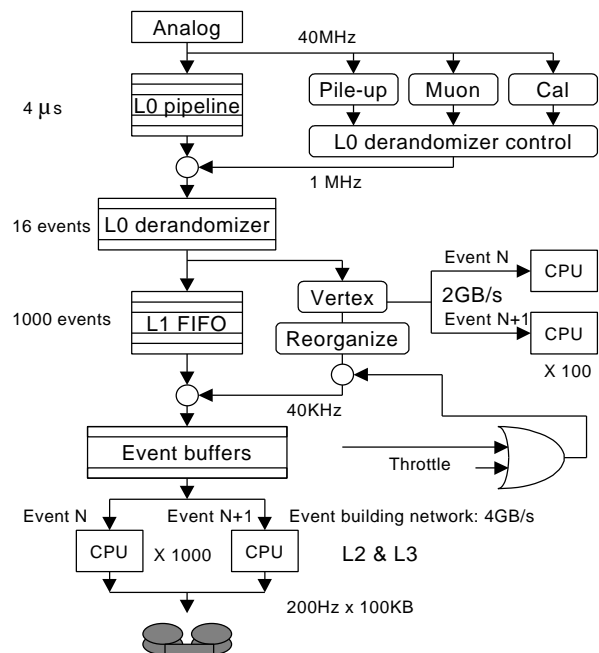

Figure 5: General front-end, trigger and DAQ architecture.

Traditional pipelining is used in the first level trigger system together with a data pipeline buffer in the front-end electronics. A special L0 derandomizer control function is used in the final trigger decision path to prevent any derandomizer overflows.

A second level trigger (L1) event buffer in the front-end is a peculiarity of the LHCb architecture. The inclusion of this buffer in the front-end was forced by the 1MHz accept rate of the first level trigger. The associated L1 trigger system determines primary and secondary vertex information in a system based on high performance CPU's interconnected by a high performance SCI network. For each event, accepted by the L0 trigger, Vertex data is sent to one out of hundred processors that makes the decision of accepting or rejecting the event. The processing of individual events is performed on individual processors and the processing time required varies significantly with the complexity and topology of the event. This results in trigger decisions to be taken out of order. To simplify the implementation of the L1 buffers in the front-end, the trigger decisions are reorganised into their original order. This allows the L1 buffers in the front-ends to be simple FIFO's, at the cost of increased memory usage. An interesting effect of the reorganisation of the L1 trigger decisions is that the L1 latency, seen by the front-end, is nearly constant even though the processing time of events have large variations as illustrated in Fig. 10. After the L1 trigger all buffer control is based on a throttle of the L1 trigger (enforcing L1 trigger rejects).

Finally event data is zero suppressed, properly formatted and sent to the DAQ system over a few hundred optical links to a standard module called the readout unit as shown in Fig. 6. This unit handles the interface to a large processor farm of a few thousand processors via an event building network. The processor farm takes care of the two remaining software driven trigger levels. An alternative configuration of the readout unit enables it to be used to concentrate data from up to 16 data sources and generate a data stream which can be passed to a read out unit or an additional level of data multiplexing. The readout unit is also used as an interface between the Vertex detector and the L1 trigger system.

The general architecture has been simulated with different simulation tools. The front-end and first level trigger systems have been simulated at the clock level with hardware simulation tools based on VHDL. The CPU based systems (L1 trigger and DAQ) have been simulated with high level simulation tools like Ptolemy.
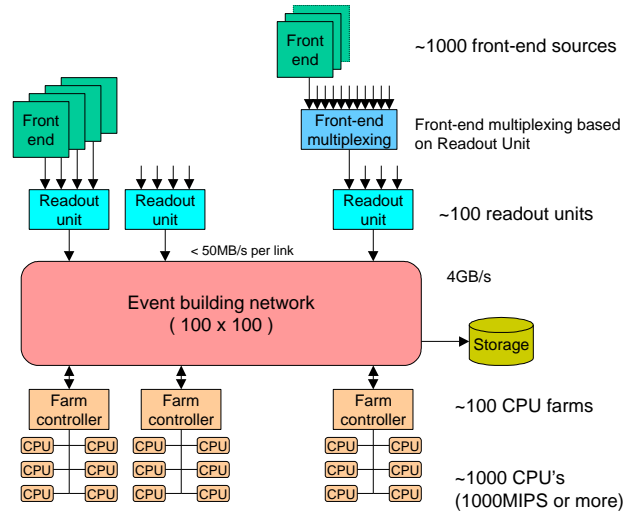

Figure 6: DAQ architecture.

A. L0 derandomizer control

At a 1 MHz trigger rate it is critical that the L0 derandomizer buffer is used efficiently and overflows in any part of the front-end system are prevented. The high accept rate also dictates the high bandwidth required from the L0 derandomizers to the L1 buffer. This bandwidth must though for cost reasons be kept as low as possible, enforcing additional constraints on the L0 derandomizer buffer. To ensure that all front-end implementations are predictable, it was decided to enforce a synchronous readout of the L0 derandomizer at 40MHz. A convenient multiplexing ratio of data from 32 channels is appropriate at this level. To be capable of identifying event data and verifying their correctness additional 4 data tags (Bunch ID, Event Id and Error flags are obligatory) are appended to the data stream, resulting in a L0 derandomizer readout time of maximum 900 ns per event. To obtain a dead time below 1% a L0 derandomizer depth of 16 events is required as illustrated in Fig. 7.
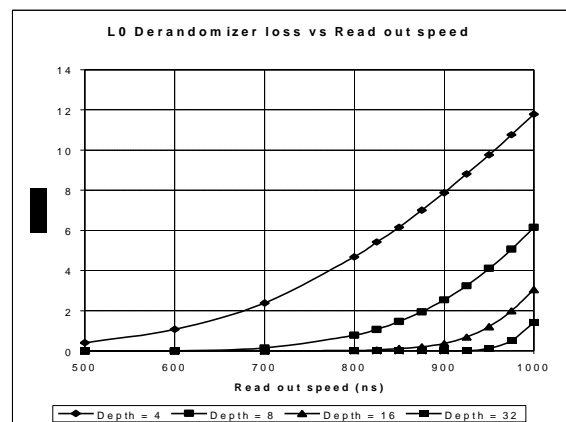

Figure 7: L0 derandomizer dead time as function of readout time and buffer depth

To simplify the control and prevent overflows of the derandomizers, it is defined that all front-end implementations must have a minimum buffer depth of 16

events and a maximum readout time of 900ns. With such a strict definition of the derandomizer requirements, it is possible to emulate centrally the L0 derandomizer occupancy. This is used to centrally limit trigger accepts as illustrated in Fig. 8. Remaining uncertainties in the specific front-end implementations are handled by assuming a derandomizer depth of only 15 in the central emulation.
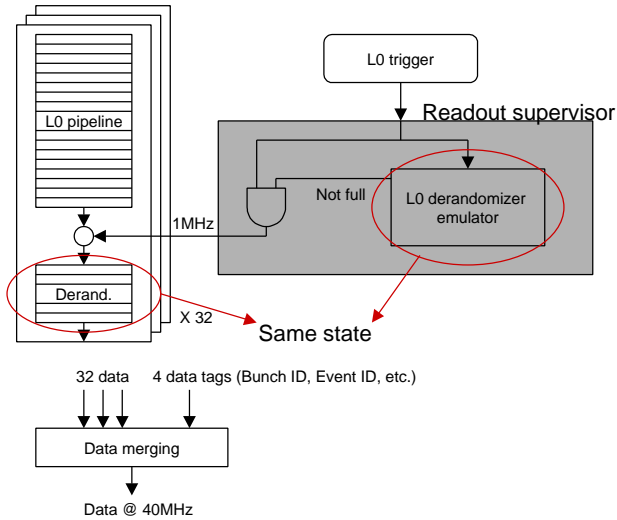

Figure 8: Central L0 derandomizer overflow prevention.

In addition a simple throttle mechanism is available to enable the following L1 trigger system to gate the L0 accept rate if it encounters internal saturation effects.

B. Consecutive triggers

Triggers in consecutive bunch crossings are in all other LHC experiments not supported, to simplify the front-end implementations for sub-detectors that need to extract multiple samples per trigger. In LHCb, with a much higher trigger rate, a 3% physics loss is associated with each enforced gap between trigger accepts. It has been determined that all sub-detectors can be made to handle consecutive triggers without major complications in their implementation.

Consecutive triggers can also have significant advantages during calibration, timing alignment or verification of the different sub-detector systems. With the defined size of 16 events in the L0 derandomizer, all detectors can handle a sequence of up to 16 consecutive triggers. Firing such a sequence of triggers can, as illustrated in Fig. 9, be used to monitor signal pulses from the detectors, study spill-over effects (in some cases called pile-up) and ease the time alignment of detector channels.

To insure the best possible use of this feature in the LHCb front-end system, a simple and completely independent trigger based on a few channels of scintillators is under consideration. Such a simple trigger can be programmed to generate triggers with any combination of interactions within a given time window

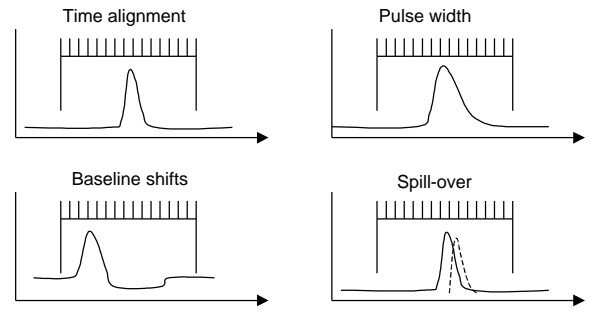(No interaction, Single interaction, two interactions in consecutive bunch crossings, etc.).

Figure 9: Use of consecutive triggers for calibration, time alignment and spill-over monitoring.

C. L1 and DAQ buffer control

The control of L1 buffers in the front-end and buffers in the DAQ system can not be performed centrally based on a defined set of parameters. Event data are assumed to be zero-suppressed before being sent to the DAQ system and large fluctuations in buffer occupancies will therefore occur. A throttle scheme is the only possible means of preventing buffer overflows at this level. A throttle network with up to 1000 sources and a latency below 10 us can throttle the L1 trigger when buffer space become sparse. A highly flexible fan-in and switch module is used to build the throttle network. The throttle switch allows the throttle network to be configured according to the partitioning of the whole Front-end and DAQ system. In addition it has a history buffer, keeping track of who was throttling when, to be capable of tracing sources of system dead time.

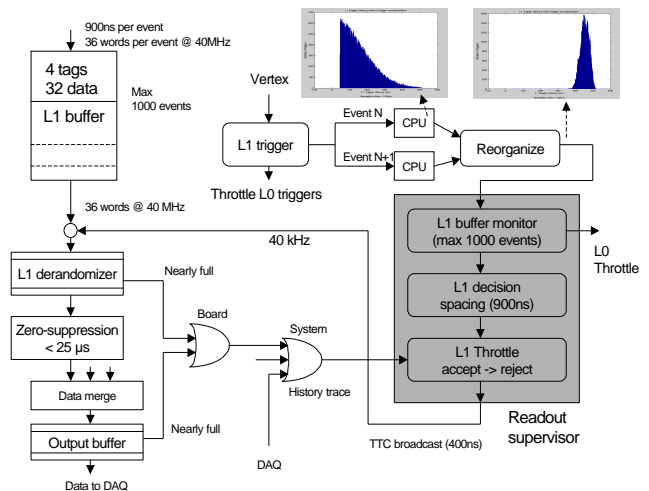Figure 10:   L1 and DAQ buffer control with plots of estimated L1 trigger latency distributions before and after reorganisation.

Because of the CPU based L1 trigger system, where the trigger decision latency varies significantly from event to event, an additional monitoring of the number of events in the L1 buffers is implemented centrally. If it is seen that the L1 buffer occupancy gets close to its

maximum (1000 events) a central throttle of the L0 trigger will be generated to prevent overflows.

The input data to the L1 buffers, from the L0 derandomizers, are as previously state specified to occur with a spacing of 900ns between each event. To prevent constraining the L1 buffers on their readout side, the L1 trigger decisions are specified to arrive via a TTC broadcast with a minimum spacing of the same 900 ns. This simplifies the implementation of the L1 buffers and their control in the different front-end systems.

D. Readout supervisor

The readout supervisor is the central controller of the complete front-end system via the TTC distribution network. It receives trigger decisions from the trigger systems but only generates triggers to the front-end that are guaranteed not to overflow any buffer, according to the control mechanisms previously described. In addition it must generate special triggers needed for calibration, monitoring and testing of the different front-end systems. Generation of front-end resets on demand or at regular intervals is also specified. During normal running only one readout supervisor is used to control the complete experiment. During debugging and testing a bank of readout supervisors are available to control different sub-systems independently, via a programmable switch matrix to the different branches of the TTC system. This allows a very flexible partitioning of sub-systems down to TTC branches. Each sub-detector normally consists of one or a few branches.
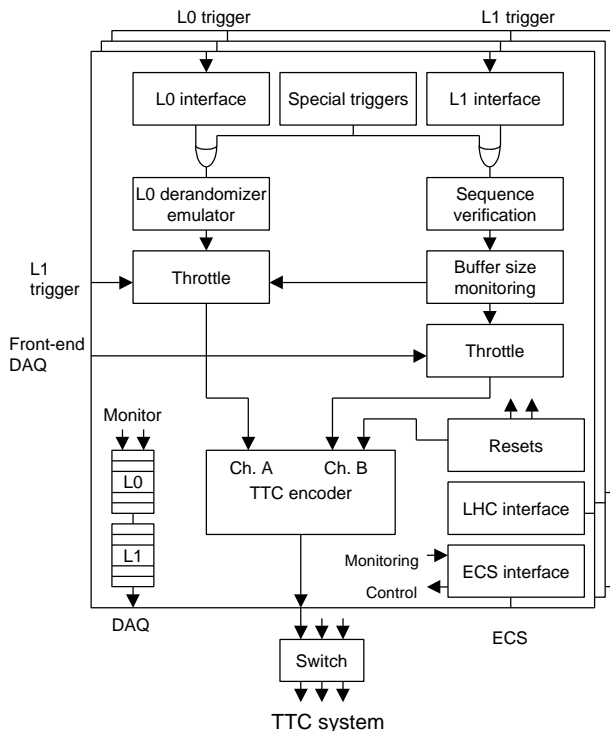

Figure 11: Architecture of readout supervisor.

The readout supervisor also contains a large set of monitoring functions used to trace the function of the system and the effective dead times encountered. This is read out on an event by event basis to the DAQ system together with the normal event data and is also accessible from the ECS system (Experiment control system).

V. RADIATION ENVIROMENT

The LHCb radiation environment is in first approximation less severe than what is seen in ATLAS and CMS because of the much lower interaction rate (factor ~ 100 lower). On the other hand, the LHCb detector is a forward angle only detector, where the radiation levels are normally the highest. The less massive and less enclosed detector configuration also allows more radiation to leak into the surrounding cavern. The total dose seen inside the detector volume ranges from ~1Mrad/year in the Vertex detector to a few hundred rad/year at the edge of the muon detector as shown in Fig. 12. The electronics located inside detectors is limited to the analogue front-ends and in some cases the L0 pipeline and the accompanying L0 derandomizer.
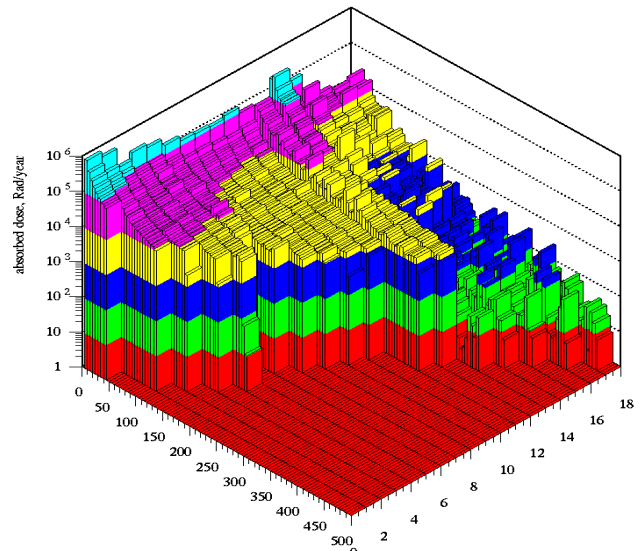

Figure 12: Radiation levels inside the LHCb detector.

At the edge of most detectors and in the cavern the total ionising dose is of the order of a few hundred rad per year and a Neutron flux of the order of 10^{10} 1Mev neutrons/cm^2 per year. This can be considered to be sufficiently low radiation levels that most electronics can support without significant degradation. The electronics located in the cavern are in general the front-end electronics with the L0 pipelines and the L1 buffers, and the first level trigger systems. This electronics consists of boards located in crates, where individual boards can be exchanged with short notice. It is assumed that short time access to the LHCb cavern, of the order of one hour, can be granted with a 24-hour notice. The installation of power supplies in the cavern is a special critical issue as their reliability in several cases have been seen to be poor, even in low dose rate environments.

Because of the additional trigger level in LHCb, the electronics located in the cavern is of higher complexity than seen in the other LHC experiments. Significant amounts of memories will be needed for this electronics and the use of re-programmable FPGA's is an attractive solution. The radiation level is though sufficiently high that single event upsets can still pose a significant problem to the reliability of the experiment.

A hypothetical front-end module handling ~1000 channels, with L1 buffering and data zero-suppression, could use 32 FPGA's with 300Kbit each for their configuration. The estimated Hadron flux with an energy above 10Mev is of the order of $3 \cdot 10^{10}$ cm^{-2} per year. With a measured SEU cross-section of $4 \cdot 10^{-15}$ cm^2/bit for a standard Xilinx FPGA, one can estimate each module to have a SEU upset every few hours. At the single module level this could possibly be acceptable. In a system with the order of 1000 modules, the system will be affected by SEUs a few times per minute!. Unfortunately it will be relatively slow to recover from such failures as the approximate cause must be identified and then the FPGA has to be reconfigured via the ECS system. This will most likely require several seconds to accomplish. This simplified example clearly shows that SRAM based FPGA's only can be used with great care, even in the LHCb cavern where the radiation level is quite low.

VI.   ERROR MONITORING AND TESTING

The use of complex electronics in an environment with radiation requires special attention to be paid to the detection of errors and error recovery.

Frequent errors can be feared to occur, making it vital that these can be detected as early as possible to prevent writing corrupted data to "tape". The format definition of data accepted by the first level trigger has been made to include data tags that allows data consistency checks to be performed. The use of Bunch ID and Event ID tags is enforced. Additional two data tags are available for error flags and data checksums when found appropriate. Up through the hierarchy of the front-end and the DAQ system the consistency of these data tags must be verified when merging data from different data sources. This should ensure that most failures in front-end systems are detected as early as possible. All front-end buffer overflows must also be detected and signalled to the ECS system, even though the system has been made to prevent such problems. The use of continuous parity checks on all setup registers in the front-end is strongly encouraged and the use of self-checking state machines based on one-hot encoding is also proposed.

To be capable of recovering quickly from detected error conditions, a set of well defined reset sequences have been specified for the front-end system. These resets can be generated on a request basis from the ECS system or can be programmed to occur with predefined intervals by the readout supervisor. To recover from corrupted setup parameters in the front-ends a relative fast

download of parameters from the ECS system has been specified. Local error recovery, e.g. from the loss of a single event fragment in the data stream, is considered dangerous as it is hard to determine on-line if the event fragment is really missing or the event identification has been corrupted. Any event fragment with a potential error must be flagged as being error prone.

To be capable of performing efficient testing and debugging of the electronics systems, the normal triggering path, the readout data path and the control/monitoring data path have been specified to be independent. All setup registers in front-end implementations must have read-back capability to be capable of confirming that all parameters has been correctly downloaded. The use of JTAG boundary scan testing is also strongly encouraged. To be capable of performing efficient repairs of electronics in the experiment, within short access periods, it is important that failing modules can be efficiently identified. It is also important that it can be confirmed quickly if a repair has actually solved the encountered problem.

VII.   EXPERIMENT CONTROL SYSTEM

The traditional slow control system, or now often called the detector control system, has in LHCb been brought one level higher to actually control the complete experiment. This has given birth to a new system name: Experiment Control System (ECS). In addition to the traditional control of gas systems, magnet systems, power supplies, and crates the complete front-end, trigger and DAQ system is under ECS control. This means that the downloading of all parameters to front-end and trigger systems is the responsibility of ECS. These parameters include large look up tables in trigger systems and FPGA configuration data on front-end modules and will consist of Gbytes of data. The active monitoring of all front-end and trigger modules for status and error information is also the role of ECS. In case of errors, ECS is responsible for determining the possible cause of the error and perform the most efficient error recovery. The DAQ system, consisting of thousands of CPU's, is also under the control of the ECS system which must monitor/insure its correct function during running. With such a wide scope, the ECS system will be highly hierarchical with up to one hundred PCs, each controlling clearly identified parts of the system in a nearly autonomous fashion.

The ECS being the overall control of the whole experiment requires it to have extensive support for partitioning. The whole front-end, trigger and DAQ system must have hardware and software support for such partitioning and the ECS will need a special partitioning manager function. The software framework for the ECS system must be a commercially supported set of tools with well defined interfaces to standard communication networks and links.

The physical interface from the ECS infrastructure to the hardware modules in the front-end and trigger systems

is a non-trivial part. The bandwidth requirements vary significantly between different types of modules and it must reach electronics located in environments with different levels of radiation. It has been emphasised that the smallest possible number of different interfaces can be supported in LHCb. Currently an approach based on the standardisation of maximum three interfaces is pursued. For environments with no significant radiation (underground counting room) a solution based on a so-called credit card PC has been found very attractive (small size, standard architecture, Ethernet interface, commercial software support, and acceptable cost). For electronics located inside detectors special radiation hard and SEU immune solutions are required. The most appropriate solution for this hostile environment has been found to be the control and monitoring system made for the CMS tracker, being developed at CERN. An interface to electronics boards located in the low-level radiation environment in the cavern is not necessarily well taken care of by the two mentioned interfaces. Here a custom 10Mbits/s serial protocol using LVDS over twisted pairs is being considered with a SEU immune slave interface implemented in an anti-fuse based FPGA. All considered solutions supports a common set of local board control busses: I^2C, JTAG and a simple parallel bus.
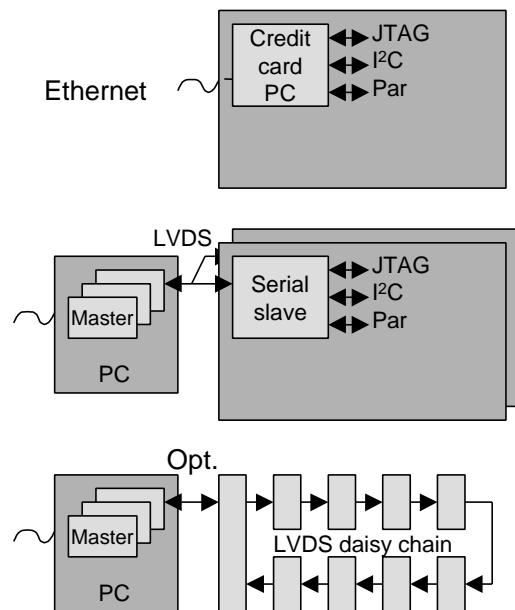

Figure 13:   Supported front-end ECS interfaces.

VIII.   STATUS OF ELECTRONICS

As previously mentioned the architecture of the front-end, trigger and DAQ systems are now well defined. Key parameters are fixed to allow the different implementations to determine their detailed specifications. After the LHCb approval in 1998, many beam tests of detectors have been performed and the final choice of detector technology is in most cases done. The electronics systems are currently being designed. LHCb is currently in a state where the different sub-systems are

progressing against the Technical Design Reports (TDR) within the coming year. The choice of network technology for the event building in the DAQ system is on purpose delayed as much as possible to profit from the fast developments in this area from industry. The required bandwidth of the network is though already available with today's technology, but prices in this domain are expected to decrease significantly within the coming years.

In most sub-detectors the architecture of the front-end electronics is defined and is progressing towards their final implementation.

For the Vertex detector the detailed layout of the detector in its vacuum tank is shown in Fig 14 and parts of its electronics is shown in Fig 15.
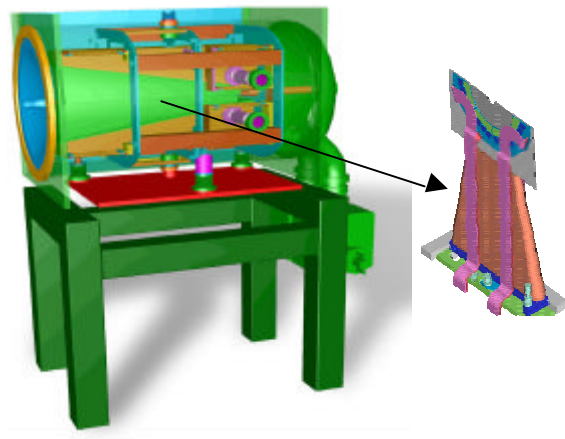

Figure 14:   Vertex detector vacuum tank and detector hybrid.

Figure 15:   Vertex detector hybrid prototype

A critical integration of electronics is required in the RICH detector. A pixel chip (developed together with ALICE) has to be integrated into the vacuum envelope of a hybrid photon detector tube as shown in Fig. 16 and 17. Here a parallel development of a backup solution based on commercial Multi Anode Photo Multiplier Tubes is found necessary, in case serious problems in the

complicated pixel electronics and its integration into the vacuum envelope is found.
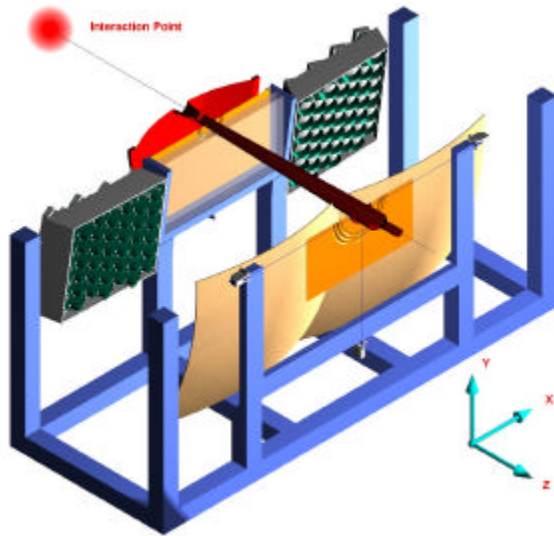

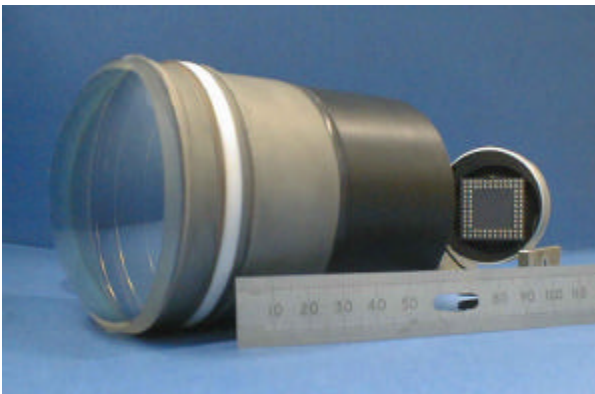Figure 16:   RICH detector with HPD detectors

Figure 17:   Pixel HPD tube

For the Calorimeter system, consisting of a Scintillating pad detector, a Preshower detector and an Electromagnetic and Hadron calorimeter, most of the critical parts of the electronics have been designed and tested in beam tests. For the E-cal and the H-cal the same front-end electronics is used to minimise the design effort.
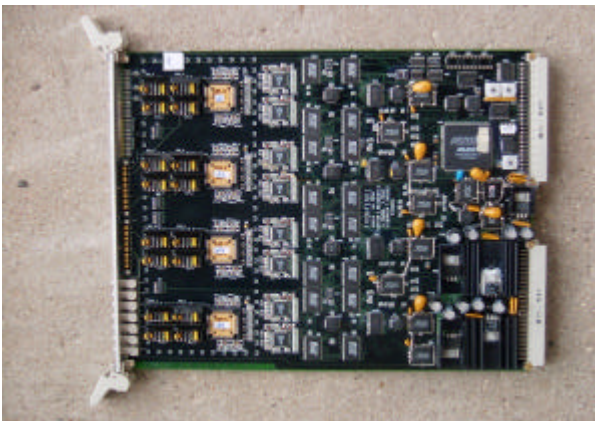

Figure 18:   Common Ecal and Hcal 12 bits digitising front-end.

IX.   APPLICATION SPECIFIC INTEGRATED CIRCUITS

The use of application specific integrated circuits is vital for the performance and acceptable cost of all sub-detectors in LHCb. ASIC's are therefore critical components onto which the feasibility of the whole experiment is based. ASIC design is complicated and expensive and any delays in their finalisation will often result in delays for the whole system. Delays of the order of 1 year can easily occur in the schedule of mixed-signal integrated circuits as no quick repairs can be made. In our environment, which is not accustomed to the design/testing and verification of large complicated integrated circuits, the time schedules are often optimistic and the time needed for proper qualification of prototypes is often underestimated.

The rapidly evolving IC technologies offer enormous performance potentials. The use of standard sub-micron CMOS technologies being radiation hardened using enclosed gate structures has been a real strike of "luck" to the HEP community. Rapid changes in IC technologies can on the other hand pose a significant risk that designs made in "old" (~5 years) technologies can not be fabricated when finally ready. This fast pace in the IC industry must be considered seriously, when starting an IC development in the HEP community, as the total development time here is often significantly longer than what can be allowed in industry. The production of IC's also pose an uncomfortable problem. Any IC ready for production today can most likely not be produced again in a few years. It is therefore of outmost importance that designs are properly qualified to be working correctly in the final application and that sufficient spares are available.

As a rule of thumb, one can assume that each sub-detector in LHCb relies on one or two critical ASIC's. For harsh radiation environments the use of sub-micron CMOS with hardened layout structures is popular. The total number of ASIC designs in LHCb is of the order of 10, spanning from analogue amplifiers to large and complicated mixed signal designs. The production volume for each ASIC is of the order of a few thousand. This low volume also poses a potential problem as small volume consumers will get very little attention from the IC manufactures in the coming years. The world-wide IC production capacity is expected to be insufficient over the coming two years with a general under-supply as a consequence. In such conditions it is clear that small consumers like HEP will be the first to suffer (not only for ASIC's).

X. HANDLING ELECTRONICS IN LHCB

The electronics community in LHCb covering front-end, trigger, ECS and DAQ is sufficiently small that general problems can be discussed openly and decisions can be reached. There is a general understanding that common solutions between sub-systems and between

experiments must be used to manage to build the required electronics systems with the limited recourses available (manpower and financial). One common ASIC development is made between the Vertex and the inner tracker detector. This is also assumed used in the RICH backup solution. A common L1 trigger interface, DAQ interface and data concentration module is designed to be used globally in the whole experiment.

Regular electronics workshops of one week have insured that the general architecture of the front-end, trigger, DAQ and ECS systems are well understood in the community and a set of key parameters has been agreed upon. In addition, a specific electronics meeting, half a day during each LHCb week, is held with no other concurrent meetings. The electronics co-ordination is an integral part of the technical board with co-ordinators from each sub-system. In the technical board it is understood that the electronics of the experiment is a critical (and complicated and expensive and - - -) part of the experiment that requires special attention in the LHC era because of its complexity, performance and special problems related to radiation effects.

XI.   CHALLENGES IN LHCB ELECTRONICS

Special attention must be paid to certain critical points in the development and production of the electronics systems for the LHCb experiment. Many of these will be common to problems encountered in the other LHC experiments, but some will be LHCb specific. The fact that the electronics systems in LHCb in many cases still are in a R&D phase will also bias the current emphasis put on specific problems.

For problems common with the other LHC experiments the most efficient approach is obviously to make common projects as has been promoted by the LEB committee. Funding for such projects though seem to be quite difficult to find. Specific areas where common support are crucial are the TTC system with its components and support for design of radiation hardened ASIC's in sub-micron (0,25 um) CMOS technology. In addition it would be of great value if the question of using power-supplies in low – medium level radiation environments (cavern) could be evaluated within such a common project.

Time schedules of ASIC's are critical as further progress can be completely blocked before working chips are available. This is especially the case where the front-end chips are an integral part of the detector (RICH HPD). The out-phasing of commercial technologies may also become critical in certain cases.

LHCb has a special need of using complicated electronics in the experiment cavern. The total dose is here sufficiently low that the use of COTS can be justified. The problem of SEU effects on the reliability of the total system must though be carefully analysed. The use of power supplies in the cavern is also a question that must be considered.

It is clear that there is a lack of electronics designers in the HEP community to build (and make working) the increasingly complicated electronics systems needed. The electronics support that can be given by CERN, to all the different experiments currently under design, is also limited. Initiatives in LHCb have been taken to involve other electronics institutes/groups in the challenges involved in our systems. Engineering groups though often prefer to work on industrial problems or specific challenges within their own domains. There is also a continuous political push for these groups to collaborate with industry. With the currently profitable and expanding electronics industry it is also increasingly difficult to attract electronics engineers and computer scientists to jobs in research institutions.

A new potential problem is surfacing in the electronics industry. The consumption of electronics components is currently increasing because of the success of computers/internet and mobile phones. Many small electronics companies have serious problems obtaining components, as large customers always have precedence. This problem of under-supply in the electronics industry is expected to get even worse and potentially last the coming few years. The need of small quantities of specialised circuits for the electronics systems for LHC experiments may therefore bring unexpected delays in the final production.

The verification and qualification of electronics components, modules and sub-systems, before they can be considered ready for production, is often underestimated in our environment. The complexity of the systems has increased rapidly with the last generations of experiments and the time needed for proper qualification often grows exponentially with complexity. This problem is, as previously stated, especially critical for ASIC's. For programmable systems based of FPGA's or processors this is to a large extent less critical. One must though not forget that a board based on a processor or FPGA's is not worth much without the proper programming, which may take a significant amount of time to get fully functional.

We also have to worry about the usual problem of documentation and maintenance of the electronics systems. The LHC experiments most likely have to be kept running, with continuous sets of upgrades, during ten years or more. A set of schematics without any other additional documentation is for complex systems far from sufficient. In many cases the schematics are not even available in a usable form, as the design of many electronics components will be based on synthesis. In some cases the tools used for the design will even not be available after a few years.