

An overview of the LHC-B experiment and its electronics

Jørgen Christiansen on behalf of the LHC-B collaboration
 CERN/ECP 1211 Geneva 23, Switzerland
 Email: Jorgen.Christiansen@cern.ch

Abstract

A general overview of the LHC-B detector and its associated electronics is given. The basic Beauty physics being searched for is introduced and a four level trigger architecture capable of extracting the rare events of interest is shown. The general architecture of the front-end electronics is described and implementation specific issues like radiation hardness and on/off -detector location of electronics are discussed. Finally the implementation of the LHC-B detector and electronics is compared to other similar experiments.

1. Introduction

LHC-B is a dedicated experiment to study CP violation and other rare phenomena in the decays of Beauty particles. The b hadron decays are characterised by short lifetimes of a few pico seconds giving average decay lengths of the order of millimetres. This makes it vital to identify multiple vertices with high precision. The experiment is built as a single arm detector covering a forward angle between ~ 10 mrad and ~ 400 mrad (0.5 - 23 deg.) giving a geometrical acceptance of b final states of $\sim 15\%$.

The interaction rate is on purpose kept at a level where multiple interactions in the same bunch crossing are rare, to enable the triggering system to perform reliable primary and secondary vertex detection. The LHC-B experiment can run already in the early phase of LHC at its optimal efficiency and at higher LHC intensity the beam at the LHC-B interaction point will be “de-focused” to keep the interaction rates at ~ 15 MHz. Multiple interactions are rejected in the trigger system by a pile-up veto mechanism.

Charged particle densities in the LHC-B detector follows a general $1/r^2$ dependency on the distance to the beam line requiring the detector granularity to vary accordingly. A typical b event contains 40 (first tracker station) to 350 (last tracker station) charged tracks per detector plane.

The total size of the LHC-B detector is comparable to a typical LEP experiment. It will be housed in the existing DELPHI cavern (IP8) with a minimum of civil engineering required. The interaction point is located off-

centre at one side of the cavern.




Fig. 1: Layout of LHC-B experiment.

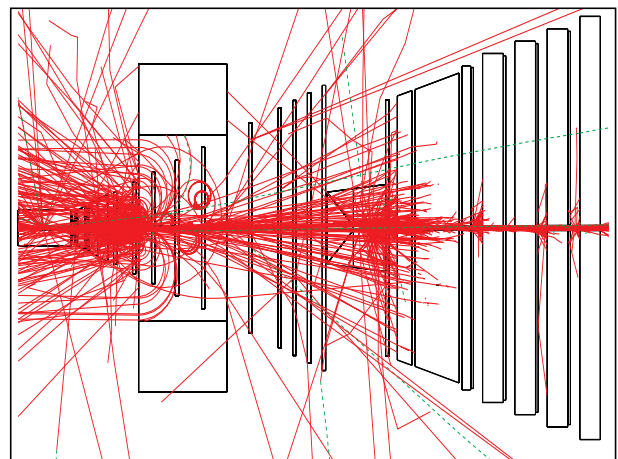


Fig. 2: Charged tracks from typical b event.

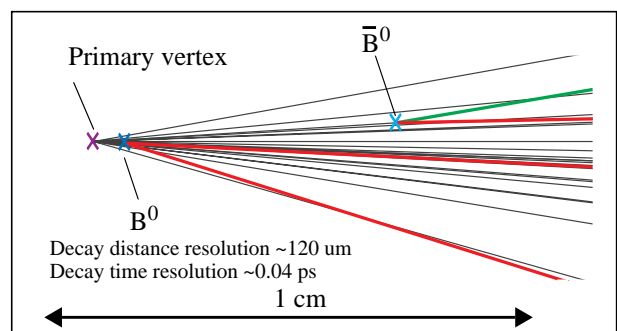


Fig. 3: Typical b event in Vertex detector

2. Detector

The LHC-B detector layout shown in figure 1 resembles a fixed target experiment with layers of vertical detector planes along the beam line increasing in size with the distance to the interaction point.

Detector	Channels K	Occupancy %	Event size Kbyte
Vertex	233	0.5	5
Inner tracker	70	8	22
Outer tracker	95	3.5	14
RICH	500	1	20
ECal	10	5	2
Hcal	4	5	1
Muon	33	1	2
Total	~ 1000		66 (~100)

Table 1: LHC-B sub-detectors with number of channels and estimated occupancies.

A silicon vertex detector performs primary and secondary vertex identification. It consists of 18 detector stations covering a radial distance of 1 cm - 6 cm from the beam. Each station consists of 6 sectors with two planes sensitive to r and ϕ respectively. The $r - \phi$ geometry is optimized for fast identification of primary and secondary vertices in the trigger system. The vertex detector is placed in a secondary vacuum separated from the main LHC machine vacuum.

12 Tracker stations, located at distances of $z = 1\text{m} - 12\text{m}$ from the interaction point, perform general tracking and measure particle momentum in the field of a dipole magnet. The outer part of the tracker stations, with modest particle densities, consists of honey comb drift chambers. For the inner part three options are being considered: Micro-Cathode Strip Chambers, Micro-Strip Gas Counters or Silicon Strips.

Two Ring Imaging Cerenkov detectors (RICH 1,2), covering different solid angles, perform particle identification. The characteristic rings of Cerenkov light is focused by mirrors onto arrays of single photon counting detectors.

Electromagnetic and hadron calorimeters perform energy measurements and identification of electrons and

hadrons. The calorimeters are in the current thinking implemented as absorber/scintillator towers in a Shashlik and tile geometry, all being read out by photo multipliers.

Finally 4 layers of muon detection are used for muon identification. The muon detector will be implemented with Resistive-Plate-Chambers or Cathode-Strip-Chambers in a projective pad geometry with pad sizes increasing with the distance to the beam line.

3. Triggering

The LHC-B read-out system is a 4 level trigger architecture as shown in figure 4. Four levels are required because of the difficulty of accurately identifying b events of interest. The first two levels of triggering can be considered hardware driven and the last two levels software driven.

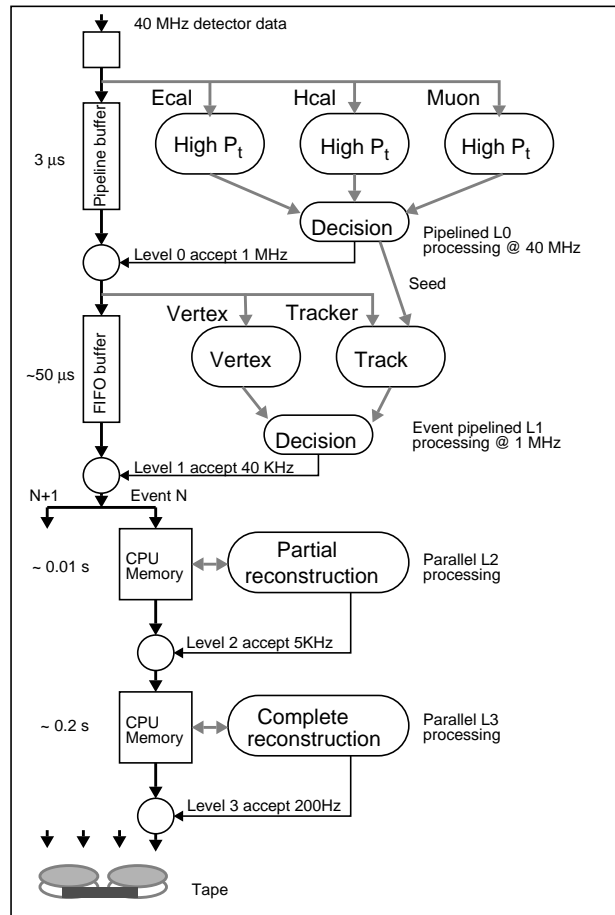


Fig. 4: Triggering architecture of LHC-B.

The level 0 trigger (L0) makes a first event selection based on the measurement of transversal momentum (p_t) from the calorimeters and the muon detector. It will be implemented in specialised hardware with a pipeline depth of 128 clock cycles ($3 \mu\text{s}$) in a similar fashion as done in CMS and ATLAS. The use of a large array of

specialised pipelined processors (3D flow) is being evaluated.

Trigger level 1 (L1) includes tracking information from the vertex and the tracker detectors. The vertex trigger selects events with large impact parameter tracks in relation to the primary vertex. The track trigger performs track reconstruction of a selected set of tracks through the tracker stations and makes a selection based on p_t and origin of tracks. The implementation of the trigger 1 algorithms will use a mixture of special purpose hardware (likely FPGA's) to perform the low level functions and farms of high performance processors (e.g. DSP's) to make the final selection. The L1 algorithms are estimated to require a processing time of tens of micro seconds for an average event. Complicated events may require 2 - 4 times the average processing time. Using farms of processors to process events with varying complexity means that trigger decisions will be taken out of order. A special trigger supervisor is envisaged to collect trigger decisions and deliver them to the L1 buffers in correct order. The ordering of trigger decisions significantly simplifies the control of the L1 buffers in the front-end at the cost of increased memory size.

The level 2 (L2) and level 3 (L3) trigger selection will be based on partial and complete reconstruction of events at full resolution. This will be performed on a large set of high performance micro processors.

An option which is currently being investigated is to allow the trigger 1 processing to take significantly longer (several hundreds of micro seconds). This would enable the tracks from the vertex and the tracker detectors to be linked to each other. Preliminary results indicate that this may reduce the L1 accept rate from 40 KHz to 5KHz giving a significant cost reduction in the data acquisition system.

The buffering of event data during trigger level 0 and trigger level 1 until its transmission over read-out links is considered to be a part of the front-end electronics of the experiment. From the read-out links through event building, trigger level 2 and 3 selection until writing data on tape is considered a part of the data acquisition system (DAQ).

4. Front-end electronics

A general view of the front-end electronics is shown in Fig 5. Detector signals are, after proper amplification and shaping, stored in pipelined L0 buffers. Data in the L0 buffers can be analog for detectors with large number of channels and limited dynamic range (Vertex, RICH, inner tracker), or digital for detectors with limited number of channels, large dynamic range, binary data or

TDC data (calorimeter, outer tracker, muon). For detectors participating in the L0 decision, binary (muon) or reduced dynamic range (~8 bit for calorimeters) data are sent to the L0 trigger hardware.

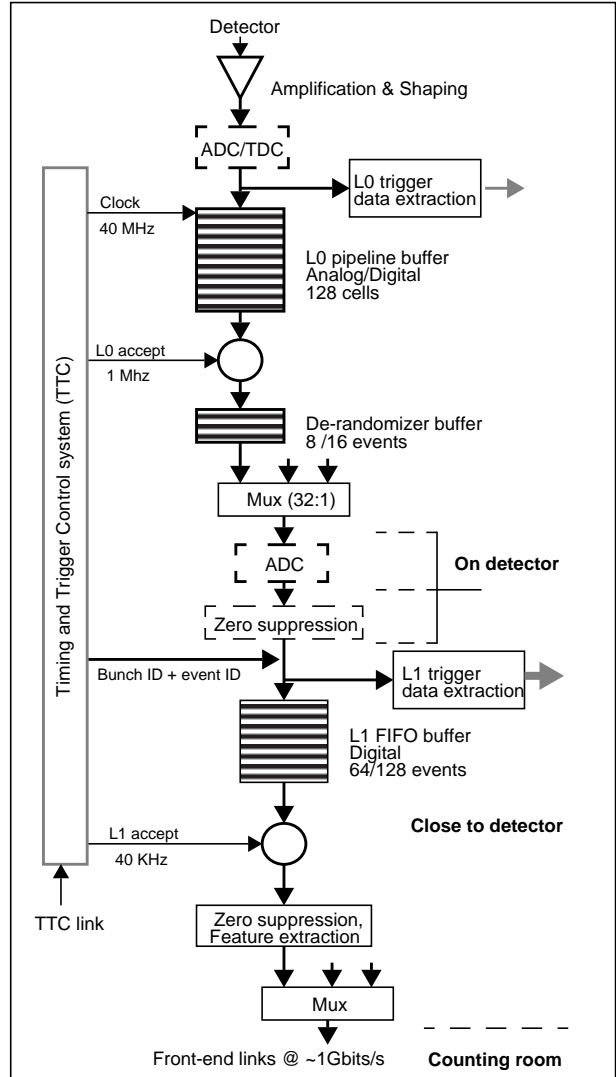


Fig. 5: Front-end electronics architecture

Events are accepted at the output of the L0 pipeline at an average rate of 1 MHz. After the L0 accept all data must be digitized before being stored in the L1 buffer. The relatively high L0 rate makes it necessary to smoothen the statistical fluctuations of the L0 rate by a combination of de-randomizer buffers and a set of trigger restrictions. Some detectors may require several samples to be extracted from the L0 pipeline for each positive trigger.

It is important to find a good compromise between trigger restrictions and the required size of the de-randomizer buffers. At a trigger rate of 1 MHz a minimum bunch spacing of 3 between triggers already results in a

potential loss (dead time) of 5% of the triggers (only 0.5% for ATLAS and CMS). It is currently not determined if all LHC-B sub-detectors can handle consecutive L0 triggers. A set of specific trigger restrictions will be chosen such that each sub-detector can optimize their implementation within a well defined frame work. A set of different trigger restrictions are shown in table 2 with their potential loss of triggers. Currently a combination of C and E is considered a good compromise. An inefficiency of a few percent is considered acceptable if it leads to significant simplifications and cost reductions in the front-end.

Restriction	Loss
A: Min. spacing 3 bunches	5%
B: De-randomizer of 4, read-out in 500ns	0.33%
C: De-randomizer of 8, read-out in 500ns	0.003%
D: De-randomizer of 8, read-out in 800ns	0.6%
E: De-randomizer of 16, read-out in 800ns	0.02%
F: De-randomizer of 32, read-out in 1000ns	1.5%
G: De-randomizer of 64, read-out in 1000ns	0.8%

Table 2: Trigger inefficiency

Trigger restriction D and E were specifically introduced as it suits a 40 MHz ADC being shared by 32 channels ($32 * 25 \text{ ns} = 800 \text{ ns}$) or a 80 MHz ADC shared by 64 channels.

The positive L0 triggers distributed to the front-end buffers are centrally restrained according to the chosen restrictions. This prevents any de-randomizer buffer from overflowing and makes it unnecessary to provide a method for each de-randomizer buffer to reduce the trigger rate (requiring large network).

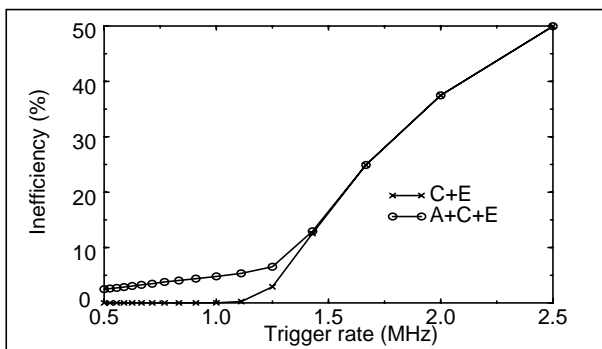


Fig. 6: L0 inefficiency as function of trigger rate

For a given set of trigger restrictions the inefficiency

depends strongly on the actual trigger rate. For the chosen set of C+E and A+C+E the inefficiency is shown as function of trigger rate in Fig. 6.

The Vertex and the tracker detectors must, after digitization, supply binary data to the trigger level 1 processing system.

One of the special characteristics of the LHC-B front-end electronics is the hard-wired L1 buffer. The data storage time of $\sim 50 \mu\text{s}$ requires that the L1 buffer is implemented as a digital buffer. Guaranteeing that L1 decisions always arrive in order enables the L1 buffers to be implemented as simple FIFO's with separators between blocks of event data. It is currently considered to implement the L1 buffer using discrete memory devices controlled by a controller ASIC common to the whole experiment. Using discrete memory devices for the L1 buffer makes it relatively easy to extend the L1 latency if required.

Some kind of de-randomization will also be required at the output of the L1 buffer. In this case the de-randomization can be performed on the triggers, instead of the accepted data, because the L1 buffer is not a simple pipeline buffer. In fact the L1 trigger decisions need to be de-randomized before the TTC system, which has only a limited bandwidth available for the L1 distribution.

For each positive L1 trigger, data must be extracted from the L1 buffers. Proper zero suppression and data formatting is done before local event building merges data from many sources into structured blocks of data. Data is finally transmitted from the front-end electronics to the data acquisition system over standardized serial links.

The physical location of the front-end electronics is still being investigated for each sub detector. The LHC-B detector being an "open" structure enables a significant part of the electronics to be located off detector. The general philosophy is that the analog front-end and the L0 buffer can be located inside the detector. The L1 buffer must be off detector, in a location where standard commercial components can be used. The 1MHz L0 accept rate though requires a significant bandwidth at this boundary.

Radiation levels inside the LHC-B detector during normal operation is estimated to be of the order of 20 Krad/year at a distance of 30 cm from the beam pipe. At distances of 10 cm and 1m the radiation levels are estimated to be respectively 200 Krad and 2 Krad per year using the general $1/r^2$ dependency. This clearly indicates that the front-end electronics used inside the detector must be radiation hard, or at least radiation tolerant, and that electronics should be located as far away from the

beam line as practically possible.

The first level of electronics off-detector is placed in crates a few meters away from the detector. The radiation level at this location must be kept below 1 Krad/year possibly by some local shielding. The effect of the magnetic field from the dipole must also be minimized to enable the used of standard power supplies and fans. These crates will not be accessible while LHC is running but should be accessible when LHC is stopped. The next level of electronics will be ~25 meters away from the detector behind a thick concrete wall to allow access during LHC operation.

In the current thinking the distribution of the 40 MHz bunch crossing clock and the level 0 and level 1 triggers to the front-end electronics is based on the TTC system also used by the other LHC experiments. The problems related to the distribution of the level 1 trigger at a rate of 1 MHz, for which the TTC was not specifically designed, are currently being studied. Using a general broadcast for this seems feasible but a dedicated broadcast type would offer significant advantages.

5. Data Acquisition System

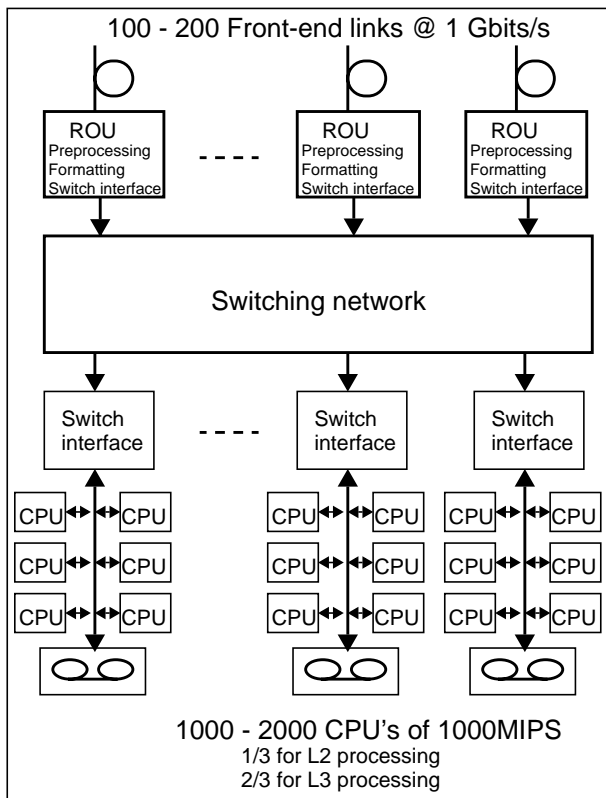


Fig. 7: Data acquisition system

Data from all sub detectors are received by Read-Out Units (ROU) in the DAQ system on several hundred op-

tical links at an event rate of 40 KHz. The total event size after proper zero suppression and formatting is estimated to be of the order of 100 Kbytes giving a total input rate to the DAQ system of 4 Gbytes/s. The read-out units will perform the necessary preprocessing and buffering of data and finally make event data accessible to the high level triggers via a large switching network.

The characteristics of the switching network will have a major impact on the performance and cost of the data acquisition system. Two implementations using ATM (Asynchronous Transfer Mode) or SCI (Scalable Coherent Interface) are being studied. ATM uses serial data streams (155Mbits/s or 620Mbits/s) routed by a network of switching nodes with internal buffers to store data cells in congested parts of the network. SCI is in principle a shared memory mapped bus with a very high raw bandwidth (0.5 - 2 Gbytes/s) using a parallel data path.

The final level 2 and 3 processing will be performed on 1000 - 2000 commercial high performance (1000 MIPS) processors connected to the switching network in local farms of processors. The processing will take full advantage of event parallelism, letting one processor perform all processing required for one event. The level 2 algorithms will perform a partial reconstruction of events and reduce the event rate to ~ 5 KHz. Level 3 will need to perform full event reconstruction and finally reduce the storage rate to 200 events/s.

Currently a push architecture is favoured where all data for one event is pushed to one processor. This can be done under the control of a central supervisor or by using pre-defined event mapping tables in each data source. This scheme requires an intelligent switch interface which can assemble event fragments without disturbing the processing on the CPU's.

A pull architecture, where only partial data is accessed by the level 2 processing, is also being considered. This approach can likely reduce the required bandwidth of the switch at the cost of a significantly more complex protocol handling.

6. Comparison to other experiments

LHC-B and the other LHC experiments (mainly ATLAS and CMS) have many similarities but also quite a few differences. LHC-B has one interaction per ~3 bunch crossings compared to several interactions per bunch crossing. The LHC-B detector is on the other hand a forward angle only detector, where the particle densities are the highest. The layout of the LHC-B detector is an open structure compared to the massive and fully enclosed ATLAS and CMS detectors.

The front-end electronics has many similarities com-

ing from the 40 MHz bunch crossing rate and similar trigger level 0 latency (level 1 for ATLAS and CMS). Electronics inside the LHC-B detector must also be designed for radiation levels requiring the use of expensive radiation hard/tolerant technologies. However the significantly higher L0 accept rate in LHC-B will in practice mean that most existing front-end chips must be modified or completely redesigned. The open detector structure gives much better access to electronics inside the detector and enables more of the front-end electronics to be moved off detector.

The Hera-B experiment also has many similarities to the LHC-B detector. The major difference is the lower bunch crossing rate of 10 MHz. This makes it necessary also to introduce major changes to Hera-B electronics to adapt it for LHC-B.

7. LHC-B collaboration

The LHC-B collaboration consists today of approximately 300 physicists from 45 different institutes. The required engineering resources needed in the future to build the detector with its associated electronics and software are in the process of being identified. It is considered very important that solutions adopted benefit as much as possible from existing developments within the high energy physics community, to minimize the time and resources required.

LHC-B is still a young and comparatively small collaboration. Additional optimizations of the general architecture will possibly be found in the near future. Currently the collaboration is preparing the technical proposal for the spring of 1998.

8. References

- [1] LHC-B Web page: <http://wwwlhcb.cern.ch>
- [2] LHC-B Letter of intent: CERN/LHCC 95-5, LHCC/I8. August 25, 1995
- [3] LHC-B Trigger and Data Acquisition Progress Report. LHCB97-005 TRIG. Feb. 26 97.
- [4] LHC-B Trigger architecture and data flow. DAQ96, RCNP, Osaka