Chapter 6 Front-End Electronics

6.1 Introduction

Features of the front-end electronics that are common to all LHCb detectors, and which are constrained by the requirements of the LHCb Trigger and DAQ system, are described here. Features specific to the requirements of a particular sub-detector are discussed in the appropriate chapter.

The front-end electronics receive analogue signals from the detector that are synchronised with the LHC bunch crossing at 40 MHz, and must deliver zero-suppressed and formatted data to the DAQ system over the read-out links. The detector data must be buffered in the front-end electronics for a time corresponding to the latency of the Level-0 and Level-1 triggers. Each trigger-level decision results in a reduction of the amount of data that must be transmitted downstream, and therefore after each stage, data from many buffers can be multiplexed onto the same output stream.

The impact of the LHCb trigger scheme on the buffering requirements in the front-end electronics is shown in Fig. 6.1. The Level-0 trigger makes the first event selection using signals from the calorimeters, the muon detector and the pile-up veto. It is implemented in specialised hardware with a latency of 128 clock cycles (3.2 μs). The Level-1 trigger uses tracking information from the vertex and the main tracking detectors. The implementation of the L1 algorithms uses a mixture of special purpose hardware and farms of high-performance processors. The latency of the Level-1 algorithms is variable, but limited to a maximum of 256 μs. A special trigger decision unit is foreseen to deliver both Level-0 and Level-1 trigger decisions to the front-end electronics in chronological order, thereby simplifying the implementation of the L1 buffers. To ensure that the Level-1 front-end buffers do not overflow, they must be capable of storing up to 256 events accepted by Level-0.

The goal is to develop a model for the front-end electronics that fulfils the general LHCb requirements and that can be adapted to the specific needs of each sub-detector. In this way, common components can be more easily identified and developed for all appropriate sub-detectors.

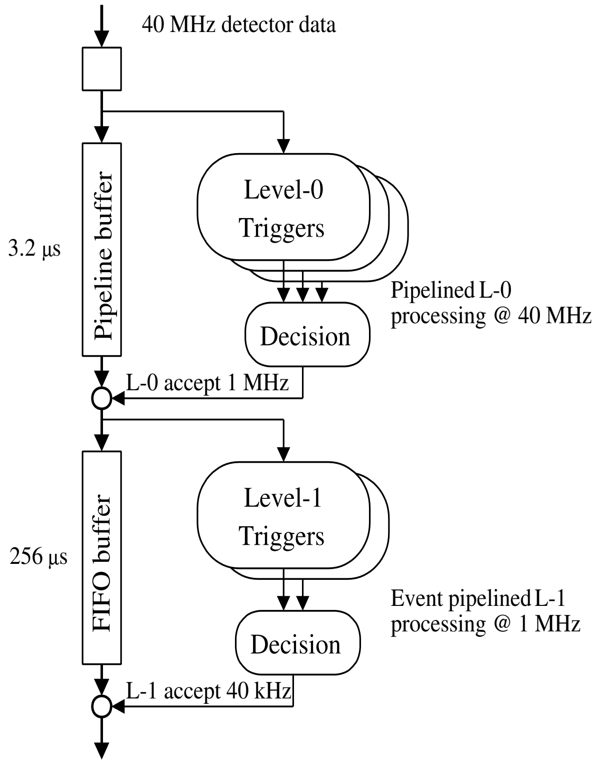

Figure 6.1: Trigger logic and front-end electronics.

6.2 Front-end model

A general view of the front-end electronics is shown in Fig. 6.2. Detector signals are amplified and shaped, when appropriate, and are stored in pipelined Level-0 buffers. Data in the Level-0 buffers can be analogue for detectors with limited dynamic range (such as the Vertex Detector, RICH, and inner tracker), digital for detectors with large dynamic range, binary or TDC data (such as the calorimeter, muon detector, and outer tracker). For detectors participating in the Level-0 decision, binary data (e.g. muon), or digital data with a reduced dynamic range data (e.g. 8 bit for calorimeters), are sent to the Level-0 trigger hardware.

Events are accepted at the output of the Level-0 pipeline at an average rate of 1 MHz. Following a Level-0 "accept" all data must be digitized, if not

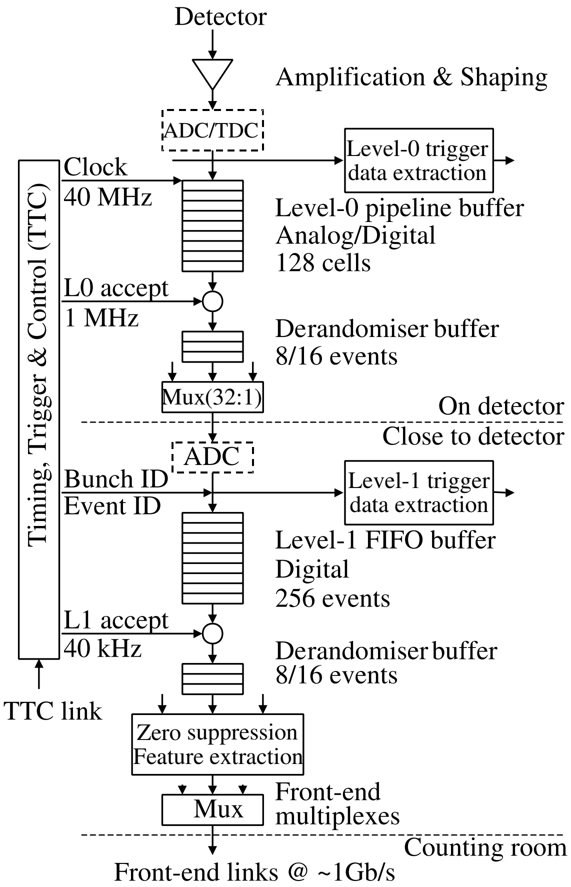

Figure 6.2: Front-end electronics architecture.

Table 6.1: Inefficiencies for different read-out scenarios.

Scenarios	Loss
A: De-randomiser depth 4, read out in 500ns	0.33%
B: De-randomiser depth 8, read out in 500ns	0.003%
C: De-randomiser depth 8, read out in 800ns	0.6%
D: De-randomiser depth 16, read out in 800ns	0.02%
E: De-randomiser depth 32, read out in 1000ns	1.5%
F: De-randomiser depth 64, read out in 1000ns	0.8%

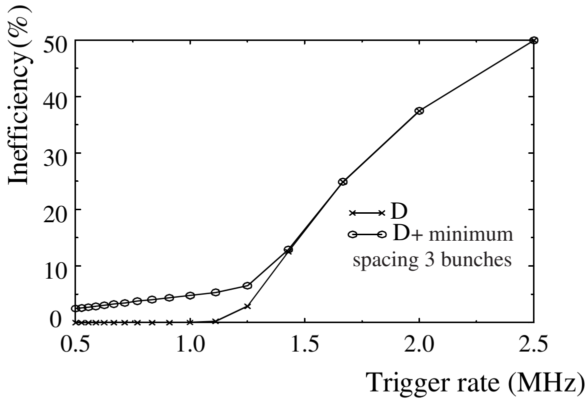

Figure 6.3: Level-0 inefficiency as function of trigger rate.

already done, before being stored in the Level-1 buffer. Level-0 "accepts" can be very close in time, in principle as close as 25 ns. A de-randomiser buffer (FIFO) is needed at the output of the pipeline and before the multiplexing stage, in order to match the data flow traffic with the speed capabilities of the output path. The optimal depth for the de-randomiser buffer is chosen to minimise the inefficiency that can arise should the buffer overflow, within constraints imposed by the cost of implementation. Table 6.1 shows the inefficiencies for different read-out time scenarios and buffer depths. The read-out time of 800 ns is of special significance, since it corresponds to the time needed to transmit multiplexed data from 32 channels using the 40 MHz clock. It can be seen that scenario C or D would both be acceptable as they result in a very small loss rate. Solution C uses the optimal readout time of 800 ns and an 8-deep de-randomiser, the depth implemented on many existing chips. The deadtime can be made negligible by increasing the depth of the de-randomiser to 16 (Solution D).

Some detectors may require samples from more than one beam crossing to be extracted from the Level-0 pipeline for each accepted trigger. With front-end electronics designs developed by ATLAS and CMS [1], this can be implemented only by inhibiting the formation of triggers from consecutive bunch crossings. As LHCb has a high Level-0 trigger rate this would result in a relatively high deadtime of $\sim$ 5% per MHz. This is clearly undesirable and therefore solutions will be sought to avoid this.

Fig. 6.3 shows the inefficiency as a function of trigger rate for scenario D. The contribution to inefficiency from the limited buffer depth becomes significant only for trigger rates above 1.25 MHz. The figure also shows that adding the minimum bunch spacing requirement adds a loss of 5% per MHz, and that this dominates for trigger rates below 1.25 MHz.

Buffer overflows must be prevented, in order to simplify event assembly and preserve the integrity of the data. Clearly the buffer depths and read-out speeds must be well matched across all components of the front-end electronics, such that a global scheme can be implemented to prevent triggers appearing at too fast a rate. The central readout supervisor can be programmed not to pass more than a given number of triggers in a given time.

The vertex detector and the tracking detectors must, after digitisation, supply binary data to the Level-l trigger processing system. During the latency of Level-1, the event data will be stored in the Level-1 buffer, which comprises a fixed number of event slots. The data storage time of 256 μs requires that the Level-1 buffer is implemented as a digital buffer. Guaranteeing that Level-1 decisions arrive in correct sequence enables the Level-1 buffers to be implemented as simple FIFO's, with separators between blocks of event data. It is currently envisaged to implement the Level-1 buffer using discrete memory devices controlled by an ASIC, and this design could be common for all subdetectors. An advantage of using discrete memory devices is that the Level-1 latency can be relatively easily extended, if required.

There is also a need for de-randomisation of the dataflow at the output of Level-1, following the same arguments as used for the Level-0 pipeline. The read-out and the subsequent data processing after the Level-1 buffer must be able to cope with an average event rate of 40 kHz. Since the positive Level-1 decisions can happen as close as 1 μs, another FIFO is needed to store temporarily the Level-1 accepted events.

Full zero-suppression and data formatting must be performed before data are dispatched over the front-end links to the data acquisition system. The front-end multiplexers merge data from many modules or cards into structured blocks of data, and these are transmitted over standardised serial links.

6.3 Location of the electronics

The optimal physical location of the front-end electronics is still being investigated for each subdetector. The LHCb detector, being an "open" structure, enables a significant part of the electronics to be located "off-detector". The general philosophy is that the analogue front-end and the Level-0 buffer can be located inside the detector. The Level-1 buffer should be off-detector, in a nearby location where radiation is sufficiently low that standard commercial components can be used. The 1 MHz Level-0 accept rate requires a significant bandwidth at this boundary.

Radiation levels inside the LHCb detector during normal operation are estimated to be of the order of 20 krad/year at a distance of 30 cm from the beam pipe. Because of the approximate $1/r^2$ dependence, standard components may be used for radial distances from the beam greater than 4 metres. At this distance the integrated dose over 10 years is below 1 krad. Front-end electronics used inside the detector must be radiation-hard, or at least radiation tolerant.

The first level of off-detector electronics is placed in crates a few metres away from the detector. The possibility of damage to the electronics caused by thermal neutrons still has to be evaluated. It is planned to use standard power supplies and fans, so the effect of the magnetic field from the dipole must also be considered. These crates will only be accessible while LHC is stopped. The next level of electronics will be 60 metres away from the detector behind a thick concrete wall, allowing access during LHC operation.

6.4 Distribution of clock and triggers

In the current scheme, 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 same TTC (Trigger, Timing and Control) system as that used by the other LHC experiments [1]. The distribution of the Level-1 trigger, not provided in the present design of the TTC, is currently being studied and preliminary results indicate that the system is able to handle Level-1 decision rates up to 1 MHz.

References

[1] Proceedings of "Workshop on Electronics for LHC Experiments", LHCC/97-60.