Chapter 3 Data Handling

3.1 Trigger/DAQ architecture

Events with B mesons can be distinguished from other inelastic pp interactions by the presence of secondary vertices, and particles with high transverse momentum (p_T). This is illustrated in Fig. 3.1 and Fig. 3.2. However, events with fully-reconstructed interesting $b\overline{b}$ final states represent only a small fraction of the total $b\overline{b}$ sample, due to the small branching ratios and the limited detector acceptance. The LHCb trigger system must therefore be selective and efficient in extracting the small fraction of interesting events from the large number of $b\overline{b}$ and other pp inelastic events.

Figure 3.3 shows the general architecture of the trigger and data acquisition (DAQ) system, including the main data and control flows involved in transporting the subdetector data from the front-end electronics to the event storage. The main parameters are summarised in Table 3.1.

The trigger strategy is based on four levels:

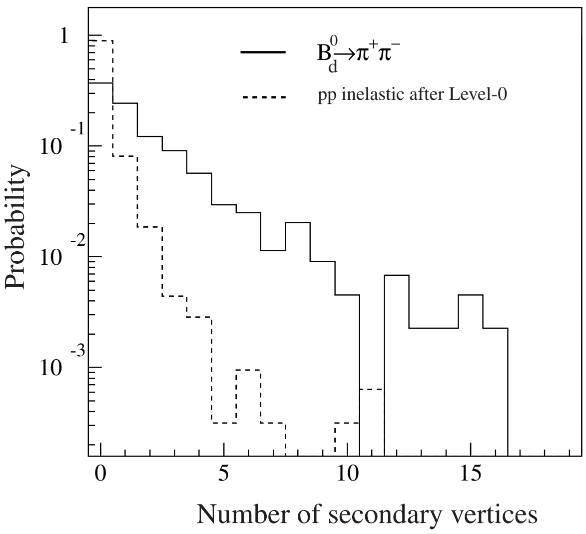

Figure 3.1: Number of secondary-vertex candidates reconstructed in the Level-1 trigger, for inelastic pp events (dashed) and $B_d^0 \to \pi^+\pi^-$ events.

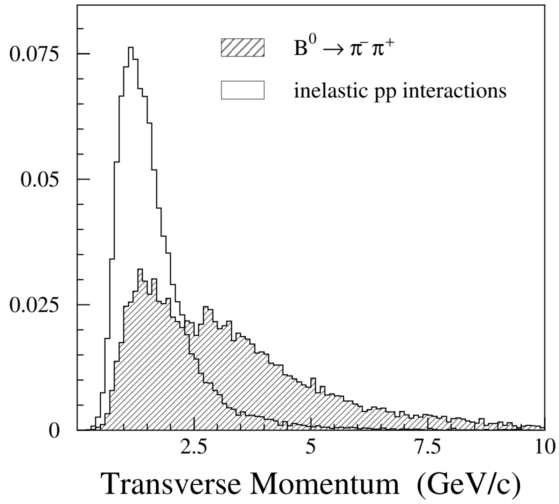

Figure 3.2: The p_T distributions of the charged hadrons with the highest p_T in the event, for pp inelastic events and $B_d^0 \to \pi^+\pi^-$ events.

Level-0 comprises three high p_T triggers, operating on muons, electrons and hadrons, together with a pile-up veto that suppresses bunch-crossings with more than one pp interaction. Level-0 operates at the bunch-crossing frequency of 40 MHz, and is designed to achieve a total suppression factor of 40. This includes a factor of about 10 suppression of single inelastic pp interactions by the high p_T trigger.

The Level-0 trigger has a fixed latency of 3.2 μs. This is sufficient to cover the execution time of the Level-0 algorithm ($\sim$ 2 μs), and the time to collect trigger input data and deliver the Level-0 decision to the front-end electronics (~ 1 μs). The choice of latency is compatible with the standard LHC design of pipelines (128 steps) in the front-end electronics.

Level-1 has two components. The first, the vertex trigger, selects events containing one or more secondary vertices. The second, the track trigger, seeks to confirm the high p_T trigger by searching for high p_T tracks in the tracking system. Level-1 operates at the Level-0 accept rate, nominally 1 MHz, and has a suppression factor of 25.

The transfer of data from the front-end electronics to the DAQ system is initiated by a positive Level-1 decision. The average event length is ~100 kB and thus the DAQ must be able to cope with data rates of order 4 GB/s.

The Level-1 latency is variable and is largely determined by the execution time of the Level-1 algorithm. The processing-time distribution has an average of 120 μs, assuming a CPU power of 1000 MIPS, and has long tails. The time for transferring data and broadcasting the Level-1 decision is $\sim$ 30μs. The maximum latency of Level-2 is set to 256 μs. The buffer space required in the front-end electronics to bridge the Level-1 latency is not a serious issue, since it can be implemented as a digital FIFO and can be of arbitrary length.

Level-2 eliminates events with fake secondary vertices, by using momentum information. Fake vertices are typically caused by multiply-scattered low-momentum tracks. Level-2 operates at the Level-1 accept rate, nominally 40 kHz, and achieves a suppression factor of 8. It correlates information from the vertex detector and tracking stations.

The Level-2 latency is determined by the execution time of the algorithm. There is clear interest to make this as efficient as possible, in order to minimise the required installed CPU power, as the trigger operates at 40 kHz. An average latency of 10 ms is required, assuming a CPU power of 1000 MIPS. The buffering requirements in the readout units are estimated to be $\sim$ 50 events for every ms of latency, and are therefore not a critical issue.

Level-3 uses full and partial reconstruction of final states to select events associated with specific b-hadron decay modes. It makes use of information from all detector components, including the RICH. The suppression factor of Level-3 is ~25 and the data recording rate is ~200 Hz.

The Level-3 latency is estimated from existing reconstruction algorithms to be $\sim$ 200 ms. At an input rate of $\sim$ 5 kHz, this gives a deployed CPU power requirement of 10^6 MIPS. Again, buffer capacity is not a critical issue, as each CPU receives only $\sim$ 30 Hz of events.

Table 3.1: Trigger/DAQ parameters.

Parameter	Value
Readout channels	950,000
Average event size	100 kB
Average Level-0 accept rate	1 MHz
Level-0 latency (fixed)	3.2μs
Average Level-1 accept rate	40 kHz
Level-1 latency (variable)	< 256 μs
Front-end Links	165
Readout Units (RUs)	~120
Event Building bandwidth	2-4 GB/s
Sub-farm Controllers (SFCs)	~ 120
Average Level-2 processing time (on a 1000 MIPS CPU)	10 ms
Average Level-2 accept rate	5 kHz
Average Level-3 processing time (on a 1000 MIPS CPU)	200 ms
Average Level-3 accept rate	200 Hz
Data storage rate	20 MB/s

3.2 Trigger system

3.2.1 Level-0 calorimeter triggers

The electron trigger requires isolated showers in the ECAL, together with correlated hits in the pad chamber (M1) in front of the calorimeter, and energy deposition in the Preshower. Simulation shows that a suppression factor of 100 for inelastic pp interactions can be achieved with an E_T threshold of 2.34 GeV, whilst maintaining an efficiency of more than 19% for B $\rightarrow$ e + X events, for which the electron hits the calorimeter. A photon trigger is implemented by requiring a higher E_T threshold and no correlated pad chamber hits. For example, a suppression factor of ~150 can be achieved for inelastic pp events with a cut of 4 GeV, whilst retaining 22 % of $B_d^0 \rightarrow K^{*0}\gamma$ where the decay products are within the detector acceptance. Events with high E_T in the HCAL are accepted in order to select B decays to hadronic final states. A cut of E_T > 2.4 GeV retains 60% of the $B_d^0 \rightarrow \pi^+\pi^-$ events that can be reconstructed offline, whilst suppressing the background by a factor of 17.

Three different approaches to the implementation of these triggers are under investigation.

The first approach (3D-Flow) uses programmable ASIC's assembled into planar layers, each processor being directly associated with one or more detector elements. Several layers are needed to achieve a pipeline mode of operation, where the number of layers required depends on the latency

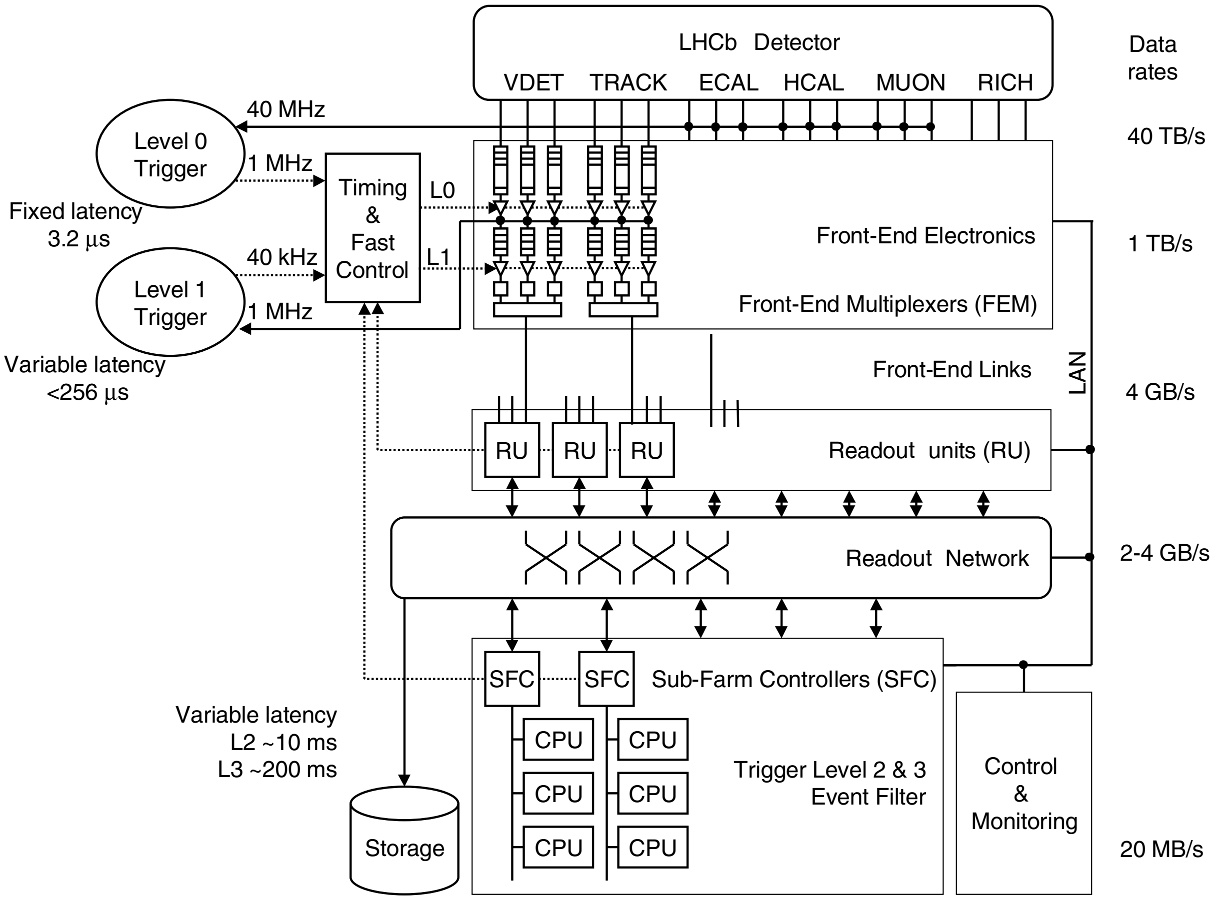

Figure 3.3: General architecture of the trigger and DAQ system.

of the algorithm. Each processor exchanges information with the 4 nearest neighbours in its layer in order to implement the cluster algorithm. Studies based on a 3×3 clustering show that the first stage of the electron and photon algorithms can be implemented with a 3D-Flow system containing 4 layers, each layer comprising 6000 processors. The hadron algorithm requires 4 layers of 1500 processors. The number of additional processors required to collect the results and apply cuts is small by comparison. Simulation studies show that the total execution time of the Level-0 algorithms is less than 1.5 μs, which falls within the allocated time.

The second implementation study, also based on 3×3 clustering, seeks to exploit the similarity of the LHCb and HERA-B electron triggers. Energy information that has been digitised and calibrated is sent from the front-end electronics into processing units, which make extensive use of look-up tables. Specialised units look for correlated hits in the pad chamber and Preshower, and apply the E_{T} cut.

The third approach uses an alternative algorithm, based on 2×2 clustering and searches for energy clusters in the front-end electronics. The aim is to reduce the large number of connections. The performance obtained is comparable to that of the algorithm employed by the two previous implementations.

3.2.2 Level-0 muon trigger

The muon trigger must maintain good efficiency for detecting $\mathrm{B} \rightarrow \mu + \mathrm{X}$, whilst suppressing muons from π and K decays. The algorithm searches for tracks in all five stations of the muon detector and imposes a minimum p_{T} requirement. Simulations show that a p_{T} resolution of 230 MeV/c can be achieved, which gives signal efficiencies of 30% for muons within the acceptance of the spectrometer, with a retention of < 2% for inelastic pp interactions.

Two implementation studies have been conducted for the muon trigger. The first uses the same 3D-Flow architecture as that used in the calorimeter trigger. Data from all 45,000 channels are fed to a stack of processors comprising three layers for making the track search, each layer containing 1300 processors. A second, much smaller, stack of processors is needed to apply the final steps of the algorithm.

The second approach aims to exploit the low track multiplicity in the muon chambers by already identifying muon-track candidates with a coarse granularity in the front-end electronics. Detailed hit information is sent to the trigger units only from those regions in the vicinity of the candidate muon tracks. This results in a data suppression factor of about 20. Full muon track processing is then performed on these pad hits.

3.2.3 Level-0 pile-up veto

Since primary vertices are spread along the beam axis with $\sigma_z = 5.3$ cm, multiple primary vertices, in a bunch crossing with more than one pp interaction, can be identified by reconstructing them with a modest resolution. An algorithm has been developed to exploit the fact that the z coordinate of the primary vertex can be determined from only the radial position of the emerging tracks. Simulation studies show that a vertex resolution of 1 mm along z can be obtained using the Pile-up Veto Counter. The performance of this veto is sufficient to give an 80% rejection of bunch-crossings with two or more pp interactions, whilst retaining 95% of single pp interactions. The vertex-based veto can be complemented by a calorimeter total energy measurement.

3.2.4 Level-0 decision unit

The results from all Level-0 trigger components are fed to a single unit where they are combined and a final decision is made. The acceptable combinations can be programmed to give flexibility and allow the trigger to be adapted to running conditions.

3.2.5 Level-1 vertex trigger

The r-ϕ geometry of the Vertex Detector is chosen to facilitate the implementation of the vertex trigger algorithm. After tracks are found in the r-z projection within the 61 degree sectors ("2D tracks"), they are combined with other 2D tracks in the opposite sector to build two-track vertices. The position of the primary vertex is determined by combining all two-track vertices, giving a resolution of 80 μm in z and 20 μm in x and y. The impact parameter with respect to the primary vertex is then calculated for each 2D track, such that tracks coming from the primary vertex can be eliminated. The ϕ-cluster information is subsequently added for all remaining 2D tracks and the track is reconstructed in three dimensions. Finally, the algorithm finds secondary vertices that are significantly separated from the primary vertex. Simulation studies show that a factor of 25 suppression of inelastic pp interactions can be achieved with a 45% efficiency for $\mathrm{B}^0_\mathrm{d} \to \pi^+\pi^-$ events that can be reconstructed offline.

The following implementation is proposed for the vertex trigger. All Vertex Detector data for a specific event, at a reduced resolution, are collected and dispatched to a single processor in a processor farm, where the trigger algorithm is executed entirely in software. Each event contains about 1000 clusters on average, which results in a data collection rate of 2 GB/s. Processing requirements are estimated from benchmarks made by running the algorithm on a 180 MHz Pentium CPU. Assuming the Level-1 trigger operates at 1 MHz, $\sim$120 processors of 1000 MIPS CPU capacity would be needed to implement the vertex-trigger algorithm. The maximum latency of the algorithm would be of the order of 300 μs on a 1000 MIPS processor, with the current version of the software. Optimisation of the code has still to be done and the goal of 256 μs should be easily achievable. A solution for the readout network has been devised and elements of the computer farm can be assembled from commercial components.

3.2.6 Level-1 track trigger

The tracking detectors are used to reject fake high-p_T tracks that passed Level-0, which originate from background due to secondary interactions in detector material, particle decays and overlapping showers. The tracking algorithm uses high-E_T clusters and tracks found by Level-0 as seeds, thereby eliminating the need for full pattern recognition. Tracks are identified by a series of backward extrapolations, starting from the seed. After each extrapolation to the next tracking station, the track parameters are recalculated using a simplified Kalman filter technique. It is sufficient to stop the extrapolation at station T6 where the track momentum is already measured with a resolution of approximately 3%. Simulation shows that the track trigger reduces the number of inelastic collisions by as much as a factor of 10 whilst retaining between 50% and 80% of signal events, depending on the B decay mode.

Each tracking station will be equipped with a processing unit for collecting data and executing the algorithm. Application-specific hardware would be used for the pre-processing steps and, since the maximum latency is 256 μsec, commercial Digital Signal Processors can be used for executing the algorithm.

3.2.7 Level-1 decision unit

The results of the vertex and track triggers are sent to the Level-1 decision unit. This unit also receives information from the Level-0 decision unit. Optimal combinations can then be made based on all information available before the final Level-1 decision is formed. This unit will be implemented using look-up tables, so that the trigger can be adapted to running conditions and physics requirements.

3.2.8 Level-2 trigger

Level-2 is designed to match the vertex information provided by the Vertex Detector with the momentum information provided by the tracking detectors. Firstly, all tracks are reconstructed in 3D in the Vertex Detector, and an initial estimate of the momentum is taken from the polar angle of the track. The primary vertex position is found with a resolution of 9 μm in x and y and 38 μm in z. Tracks in the forward hemisphere having an impact parameter larger than twice the resolution ($\sim$6 tracks per event) are followed to station T5. The impact parameter and its error are then recalculated using the measured momenta. Events are kept if there are at least three tracks with an impact parameter greater than 3 times the resolution. Background from $\mathrm{K_S}$ and Λ decays is reduced by requiring that the impact parameter is <2 mm. Simulations show that Level-2 loses very few B-meson events, while it gives good rejection of light quark events. Another positive feature is that it also suppresses $\mathrm{b\overline{b}}$ events where the b-hadron decay products are only partially contained inside the detector acceptance.

The Level-2 algorithms are executed on a farm of general-purpose processors. The goal of a Level-2 latency of $\sim$10 ms/event should be easily achievable, and the total CPU requirement for Level-2 is 4×10^5 MIPS.

3.2.9 Level-3 trigger

A set of independent algorithms is used to select exclusive b-hadron decays according to their topology. The filter algorithms are based on reconstruction and analysis codes, but since they are applied in real-time, final calibration and alignment constants are not necessarily available and cuts are deliberately loosened. Fitted tracks are subjected to a precise vertex reconstruction, and, for those tracks having a large impact parameter, a fast RICH pattern-recognition algorithm is executed to search for kaons. Finally all information is used to make a flavour tag and the invariant masses are calculated for the b-hadron decay of interest. It has been found that a rejection factor of 25 can be achieved. Benchmarks of the Level-3 algorithms indicate a total CPU requirement of 10^6 MIPS.

3.3 Data acquisition system

The role of the DAQ system is to read zero-suppressed data from the front-end electronics, to assemble complete events and to provide sufficient CPU power for execution of the Level-2 and Level-3 algorithms. The flow of data through the DAQ system is being studied using simulation data. The input rates are determined by the average event size, which is $\sim$100 kB, and the Level-1 accept rate, nominally 40 kHz. This results in a total data-acquisition rate of 4 GB/s. The output rate is governed by the nominal Level-3 accept rate ($\sim$200 Hz), leading to a data storage requirement of $\sim$20 MB/s.

The current design contains the following functional components. Readout Units (RU's) receive data from one or more front-end links and assemble the fragments belonging to each event. Full event building is achieved by having all RU's dispatch their data into a readout network such that fragments belonging to the same event arrive at the same destination. Complete events are assembled at the destination by a unit called the Sub-Farm Controller. This unit also has the role of allocating each event to one of the free CPU's it manages, and the Level-2 and Level-3 algorithms are executed on this CPU. Accepted events are written to storage devices.

Two protocols are being studied for managing the flow of data and for ensuring events are assembled correctly. In the first of these (full readout), data are immediately routed through the readout network to the destination as soon as they appear at the RU. Thus complete events are immediately made available, and there is complete flexibility in defining and applying the high-level trigger algorithms. In the second approach (phased readout), data from subsets of the RU's are sent in several phases corresponding to the requirements of the high-level triggers. In this approach use is made of the fact that high-level triggering is performed in a sequence of steps, each step requiring a subset of the data. Thus the bandwidth required of the readout network can be reduced (by $\sim$50%), since data need not be transmitted from all detectors if the event can be rejected at an early stage. The full readout protocol is superior from the point of view of simplicity of protocol, and the flexibility in application of the high-level trigger algorithms. However, the future availability of high-bandwidth network technologies will determine whether this solution can be realised at an affordable cost.

The clock, Level-0 and Level-1 decisions must all be transmitted to the front-end electronics. The Trigger, Timing and Control distribution system (TTC), developed in the context of the RD12 project and used by the other LHC experiments, can also be used for LHCb, despite the added requirement of distributing the Level-1 decision.

The Detector Control System (DCS) will be used to monitor and control the operational state of the LHCb detector, and to acquire slowly-changing data from the detector to keep a permanent record of environmental parameters. The DCS will be fully integrated with the DAQ system.

3.4 Computing

A data-flow analysis has been made of the computing system, in order to identify the key tasks and data flows. Estimates of computing requirements are made by analysing existing LHCb simulation, reconstruction and analysis programs, and, in cases where they do not yet exist, by extrapolations from similar codes developed by other experiments (such as HERA-B). Results show that the total CPU requirements needed for simulation, for Level-2 and Level-3 triggers, and for reconstruction and analysis of simulated and real data, amounts to $\sim 5\times10^6$ MIPS. The data-storage requirements for both real and simulated data amount to $\sim$500 TB per year.

At present, LHCb depends almost entirely on central computing facilities for satisfying its computing needs. A modest increase in these needs ($\sim$50% per year) is expected up to the end of 2000, and the use of central services, inside and outside of CERN, will continue for this period. There will then be a significant increase in activity in test beams, and investment in private CPU capacity is therefore be envisaged. Investment in full scale facilities will be left as late as possible, to benefit from improvements in the price/performance ratio of computing equipment. It is considered that the computing requirements are within the limits of current technology trends, and that the equipment will be affordable.

Attention will be given to the quality of software, particularly for the high-level trigger algorithms. The extra effort that this implies can be recovered by extracting as much use as possible out of each software component by ensuring its reuse in all application domains. This implies investment in sound engineering and management practices. The LHCb simulation program is written in FORTRAN, but will be re-engineered using Object Technologies.