13 Physics selection and HLT performance

13.1 Introduction

A preliminary view of the on-line event-selection scheme and the corresponding physics coverage was presented in the HLT, DAQ and DCS Technical Proposal (TP) [13-1]. Since then the studies have evolved to cope with a new scenario for the start-up of the LHC machine and in response to funding constraints. The LHC start-up scenario has a target to deliver 10 fb^{-1} in one year, now assuming a peak luminosity per fill of 2×10^{33} cm^{-2} s^{-1}, which is a factor of two higher than assumed in the TP. At the same time as having to address this higher initial luminosity, financial resources from the HLT/DAQ project had to be re-assigned to meet urgent needs elsewhere in the experiment. As a consequence, the construction of the HLT/DAQ system will have to be staged to match the available funding, so that only a reduced system will be available for the initial running. These changes required a major revision of the Physics and Event Selection Architecture of the experiment, including new ideas for reducing event rates and sizes while retaining as many as possible of the ATLAS physics goals. Needless to say, only the availability of real data will allow this proposal to find a final implementation and a fine tuning of the relative weights of the selection signatures. However, it is important to be able to face this phase with the most complete set of tools and a versatile selection architecture, in order to cope with the obvious unknowns that will likely show up at the time of LHC start-up.

As it has been described in Chapter 9, the High Level Trigger (HLT) system of the experiment is composed of two separate event-selection steps, LVL2 and the Event Filter (EF), each with distinctive and complementary features. A common denominator is that they will operate using software algorithms running on commercial computers to test hypotheses of particle identification and apply event-selection criteria. LVL2 will do this with special-purpose algorithms that need to operate in about 10 ms and use only part of the detector information at full granularity, while the EF will have the fully-built events at its disposal and work with an execution time of the order of a second. It is important to maintain a flexible scheme able to adapt easily to changes in machine conditions (e.g. luminosity or background rates). The modularity of the HLT will allow the implementation of different reduction steps at different stages.

A essential input to the HLT process is the seeding of the selection with the results from LVL1. When making performance studies for the HLT, a detailed simulation of the LVL1 result is therefore needed. This identifies the regions of the detector (Regions-of-Interest) where potential candidates for interesting physics objects are found. This simulation, described in Section 13.2, allows for a realistic use of the information coming from LVL1, using the same algorithms and procedures that will be implemented in custom hardware in the experiment.

Given the commonalities between LVL2 and the EF, it was recognized that a coherent and organized approach to the software components of the trigger was needed. The work presented in Section 13.3, which represents an important step forward with respect to the TP and is the result of a large effort, has concentrated on this issue. Common tools have been developed for the event selection, and common data-model components and methods have been identified that can be shared by different algorithms, in particular at LVL2. This approach will ease the implementation of different selection schemes and also simplify the migration of algorithms between trigger levels.

Another important point for new developments has been compliance with the updated detector geometry and with the format of the raw data as it will come from the ReadOut System. This implies that the algorithms must operate starting from streams of bytes organized according to the readout structure of each detector, in exactly the same way as in the real experiment. This has allowed one to study and understand the implications of converting these byte-streams to the objects needed by algorithms in order to perform trigger selections, including making preliminary measurements of the computing time needed for these conversions.

In Section 13.4 the outcome of the present studies is presented. It should be noted that the detailed results on trigger rates and efficiencies presented in the TP and in related papers and presentations are still considered a valid baseline approach. Present work has concentrated on implementing a realistic approach to data handling and preparation for the algorithms and on assembling a proper set of software selection tools. Emphasis has been put on the selection of electrons and photons, both of which stem from an electromagnetic cluster identified at LVL1, and muons. For each of the electron / photon and muon event-selection "vertical slices", a thorough implementation of the approach described above has been attempted. For events that passed the LVL1 selection, bytestream raw data organized according to the actual detector readout format are used by LVL2 algorithms operating within the PESA Steering and Control framework (see Section 9.5). Simulated RoI information from LVL1 is used to initiate the LVL2 processing. Trigger elements are built using detector information and used to test particle identification hypotheses. For events retained by LVL2, the EF reconstruction and analysis are performed (the EF may or may not be seeded using the LVL2 result) and the final-selection result is made available for off-line use. Results on rejection against the dominant backgrounds and on the efficiencies for typical signals are reported, as well as the rates deriving from each of the selections.

In order to span fully the ATLAS physics coverage, signatures involving jets, hadronic decays of tau leptons, large missing transverse energy and also jets with b-quark content are needed, in addition to the electron, photon and muon ones discussed above; these are also discussed in Section 13.4. As described in Chapter 4, spare on-line resources will be used for B-physics studies, e.g. for luminosities below the peak one. Results are based on preliminary analyses of samples of the order of 10^6 events and they will be extended in the near future to the larger samples of several 10^7 events produced thanks to the effort of the ATLAS Data Challenge Group [13-2].

The global assessment of the event rates to be recorded for off-line analysis, based on the present knowledge for each signature, is made in Section 13.5. A sketch of issues related to the initial phase of the experiment, seen from the selection architecture point of view, is given in Section 13.6.

13.2 The LVL1 trigger simulation

An important ingredient to many HLT tests and studies is the simulation of the LVL1 trigger, the results of which serve as input to the HLT trigger processing. The ATLAS LVL1 trigger [13-3] is itself a complex system consisting of the calorimeter trigger, the muon trigger and the Central Trigger Processor (CTP) that makes the final LVL1 event decision. Figure 13-1 gives an overview of the LVL1 trigger; the various components mentioned in the figure will be explained later in this section except for the TTC system (trigger, timing and control) which has no equivalent in the simulation.

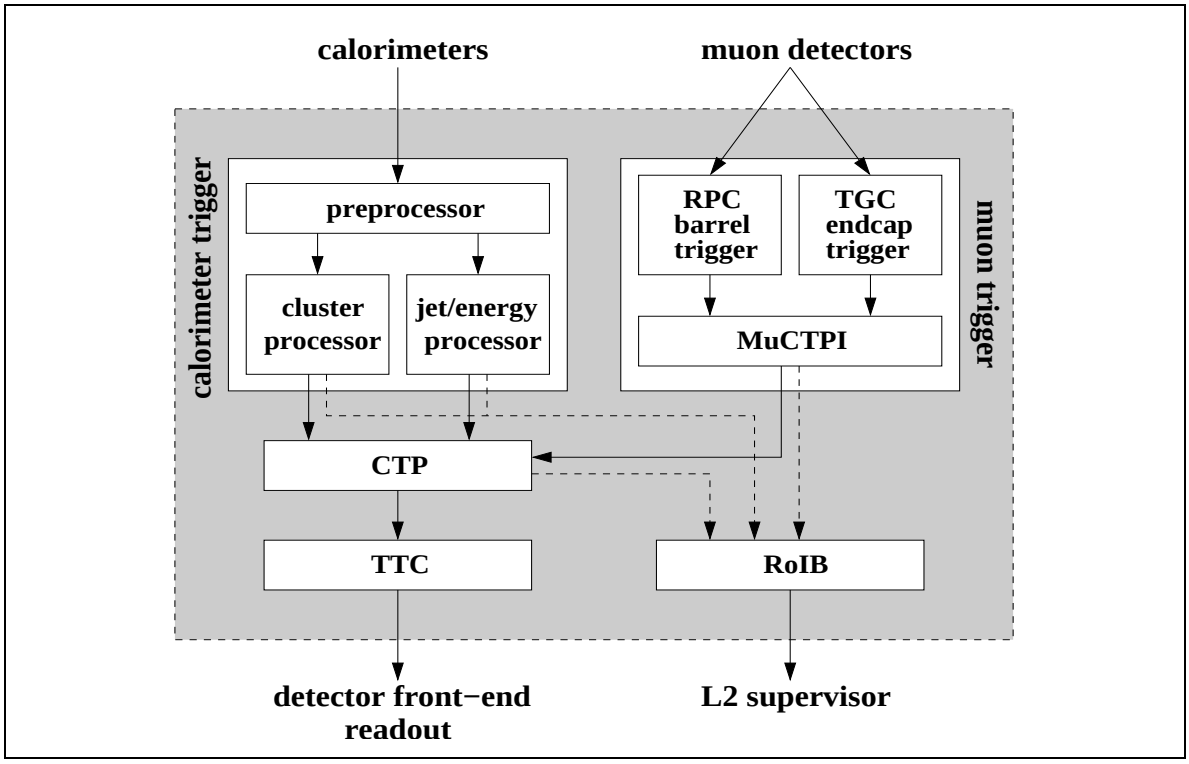

Figure 13-1 An overview of the ATLAS LVL1 trigger system. The Region-of-Interest Builder (RoIB) formally is not a part of the LVL1 trigger. However, it is simulated together with the LVL1 trigger.

The LVL1 trigger simulation is implemented in C++ in the ATLAS offline computing framework Athena. It relies heavily on the ATLAS offline data storage implementation, the so-called Transient Event Store (TES). The structure of the simulation follows closely the structure of the LVL1 trigger hardware. Figure 13-2 shows a package view of the LVL1 simulation. The simulation consists of packages simulating various hardware blocks: the resistive plate chamber (RPC) muon trigger (indicated by the package TrigT1RPC in Figure 13-2), the Muon-to-CTP Interface (MuCTPI, package TrigT1Muctpi), the calorimeter trigger (package TrigT1Calo) and the Central Trigger Processor (package TrigT1CTP). The LVL1 configuration (package TrigT1Config) and the simulation of the Region-of-Interest Builder (package TrigT1RoIB) are provided as additional packages. There are also packages for the definition of the LVL1 result raw data object (package TrigT1Result), for classes used by more than one package (package TrigT1Interfaces), and for the conversion of the LVL1 result into the hardware format, the so-called *bytestream* conversion (package TrigT1Result-Bytestream). The various parts of the simulation shown in Figure 13-2 will be explained in the next sections. The simulation of the muon trigger in the endcaps, the signals for which are provided by the thin-gap chambers (package TrigT1TGC), so far exists only as a stand-alone program and will not be discussed in detail in this document.

The interfaces and data formats to be used in the simulation [13-4] were designed to follow as closely as was practical the data formats used in the LVL1 trigger hardware which are documented in [13-5]. Additional information on the LVL1 simulation can be found in [13-6].

13.2.1 Configuration of the LVL1 trigger

The first task of the LVL1 trigger configuration package is to translate the trigger menu, i.e. the collection of event signatures LVL1 is supposed to trigger on, into something that the simula-

tion of the CTP can understand and use in making the event decision based on logical combinations of the inputs delivered by the calorimeter and muon triggers. The LVL1 signatures, or *trigger items*, are combinations of requirements (or *trigger conditions*) on the multiplicities of various kinds of candidate objects found by the calorimeter and muon triggers in the event. (See later subsections for details about the calorimeter and muon triggers and their simulation.)

A simple example of a trigger item is 'one (or more) electron / photon candidate(s) with transverse energy above 10 GeV and one (or more) muon candidate(s) with transverse momentum above 15 GeV'. In a frequently-used and obvious notation this reduces to '1EM10+1MU15', where the string 'EM' ('MU') represents the electron / photon (muon) candidate, and the numbers in front of and behind the string symbolize the required multiplicity and the required transverse momentum, respectively. The combination of a string and a threshold value (like 'EM10') is called a *trigger threshold*.

The second task of the LVL1 configuration package is to configure the simulation of the calorimeter and muon triggers to deliver the information required to make the event decision using the trigger menu, i.e. to deliver the multiplicities for the required trigger thresholds. For the example mentioned above, the calorimeter trigger has to be configured to deliver the multiplicity for the threshold 'EM10', i.e. the number of electron / photon candidate objects with transverse momentum above 10 GeV, to the CTP. It is obvious that the trigger menu and the trigger thresholds for the calorimeter and muon triggers have to be defined consistently. In particular, all thresholds used in the definition of any trigger condition in any trigger item must be delivered to the CTP by the calorimeter and muon triggers.

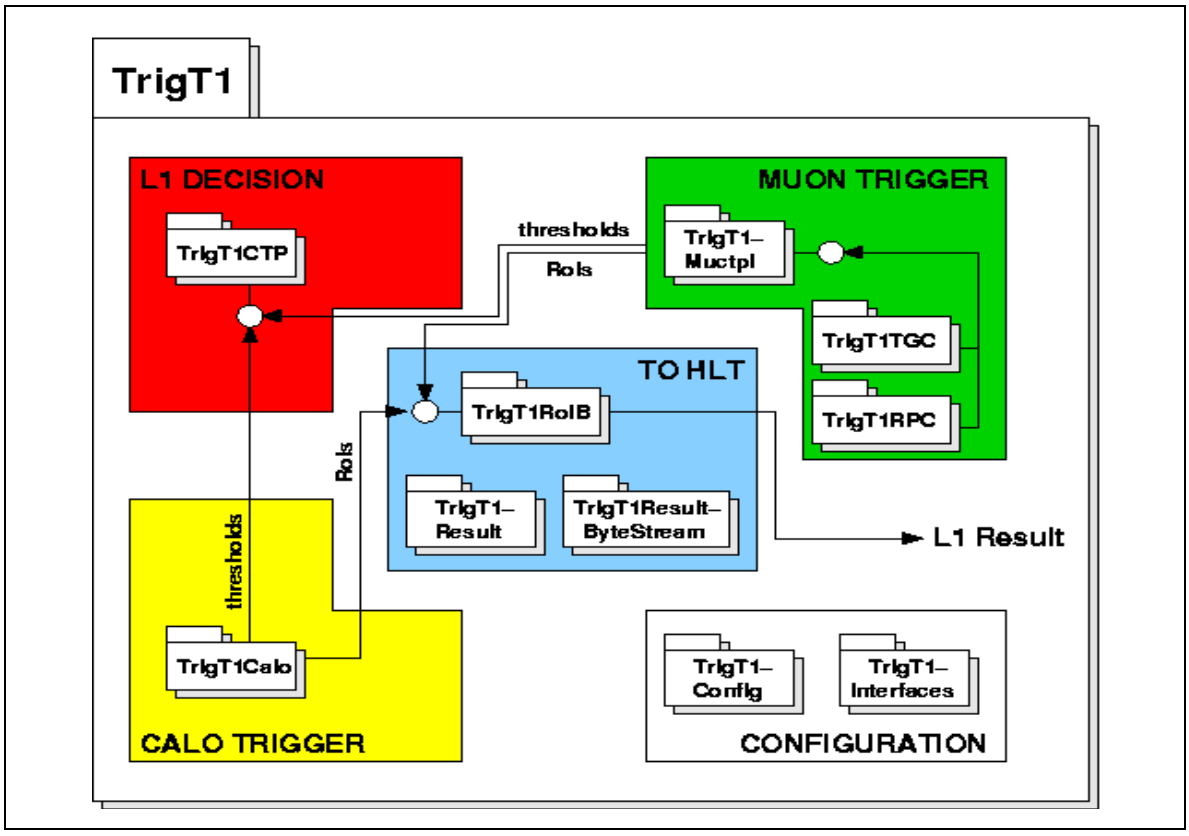

Figure 13-2 A package view of the LVL1 trigger simulation.

Both the trigger menu and the list of required trigger thresholds are defined using XML and are parsed into instances of C++ classes using the Xerces DOM API [13-7]. The parsing of the trigger menu creates an object which contains the information on how the CTP simulation has to discriminate the calorimeter and muon trigger inputs (trigger conditions) and what items have to be built from these conditions. In addition, configuration objects for the calorimeter and muon triggers are created in the configuration process and are stored in the TES for later retrieval by the calorimeter and muon trigger simulations. These objects contain the list of thresholds for which the subsystems have to provide multiplicity information to the CTP simulation.

The LVL1 trigger configuration software is currently being adapted to also be able to configure the LVL1 trigger hardware by deriving the necessary look-up table files and FPGA configuration files from the XML trigger menu and trigger threshold list. Such a common configuration scheme will allow for cross-checks between hardware and software.

13.2.2 The calorimeter trigger and its simulation

The LVL1 calorimeter trigger [13-8] searches for localized energy depositions in the calorimeters due to electrons and photons (electromagnetic clusters), single hadrons and hadronic decays of tau leptons (narrow hadronic clusters) or jets (broader hadronic clusters). For each type of cluster, a number of programmable transverse-energy (E_T) thresholds are provided. With the exception of the jet clusters, isolation requirements can be imposed — these are implemented by applying thresholds on isolation variables associated with the cluster. The multiplicities of the candidate objects of each type are counted for each threshold set, and the multiplicity values are passed on to the CTP to be used in making the LVL1 event decision.

In addition to the local-energy triggers discussed above, the calorimeter trigger calculates global energy sums (total transverse energy and missing transverse energy) which are discriminated against a number of programmable thresholds; the results of the threshold comparisons are also passed to the CTP to be used in making the LVL1 event decision.

The calorimeter trigger uses signals from ~7200 trigger towers which are analogue sums over calorimeter cells in the liquid-argon or scintillator-tile calorimeters. The trigger tower signals are digitized in the preprocessor electronics, which also performs digital signal processing to calculate transverse energy and make bunch-crossing identification. The resulting data are passed on to two processor subsystems. The Cluster Processor searches for electron / photon and tau / hadron candidates within 3200 trigger towers of granularity $\Delta\eta\times\Delta\phi = 0.1\times0.1$ in each of the electromagnetic and hadronic calorimeters in the pseudorapidity range $|\eta| < 2.5$. The Jet / Energy Processor searches for jets and makes the global energy sums using coarser (jet) elements of granularity $\Delta\eta\times\Delta\phi = 0.2\times0.2$ over a larger rapidity range ($|\eta| < 3.2$ in case of the jet trigger; $|\eta| < 4.9$ for forward jets and the energy-sum triggers). Note that it is the preprocessor that sums the trigger towers, independently for the electromagnetic and hadronic calorimeters, to form the larger elements used in the Jet / Energy Processor.

In the electron / photon trigger a candidate object is defined by a local E_T maximum in a region of 2×2 trigger towers (electromagnetic plus hadronic), corresponding to a 0.2×0.2 region in η–ϕ space. Vetos may be applied on the hadronic transverse energy in this region and / or on the transverse energies in the rings of towers surrounding the 2×2 region in the electromagnetic and hadronic calorimeters. The cluster thresholds are applied on the maximum transverse energy in any edgewise-adjacent pair of electromagnetic-calorimeter trigger towers within the central 2×2 region. (See Refs. [13-3] and [13-9] for a more detailed description of the various calorimeter trigger algorithms.)

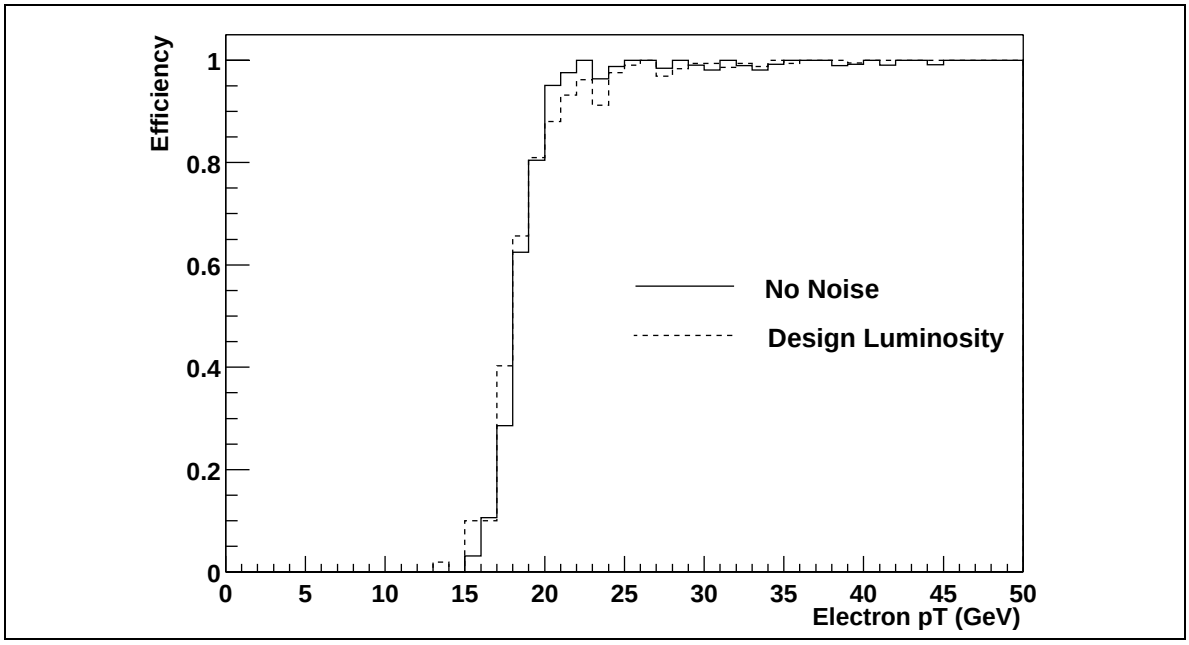

Figure 13-3 Single electron trigger efficiency as a function of the electron transverse energy E_T, using a 17 GeV threshold for different scenarios. See text for details.

In addition to counting multiplicities of candidate objects, Regions-of-Interest (RoIs) are identified and passed, via the Region-of-Interest Builder, to the LVL2 trigger for events that are retained by the LVL1 trigger. The RoIs, which contain information on the geographical location of the object in the calorimeters and on the E_T thresholds passed, are used to seed the HLT event-selection process. More detailed information from the LVL1 calorimeter trigger is sent to the DAQ using standard readout links.

The LVL1 calorimeter trigger simulation is designed to reproduce in detail the functionality of the hardware, but it does not entirely duplicate the dataflow. The primary reason for this is efficiency — the hardware trigger will process large amounts of data in parallel, which does not translate well to offline simulation software.

Currently the simulation of the LVL1 calorimeter trigger starts from calorimeter cell energies which are summed to form tower energies. (A more complete simulation which is in preparation will include details of the signal-processing chain from the calorimeters to the output of the preprocessor, but this is not yet available.) The cell energies, which can be taken from the ATLAS fast simulation or from the detailed GEANT simulation of the calorimeters, are used to build, in a simplified geometrical approach, trigger tower signals to which calibration and a gaussian noise can be applied. The tower data are passed to the Cluster Processor simulation (electron / photon and tau / hadron finding), and are summed into the coarser jet elements used in the simulation of the Jet / Energy Processor (jet finding and calculation of energy sums).The results from the simulation of the calorimeter trigger are stored in the TES for later retrieval. These include the multiplicity outputs that are used by the CTP in making the event decision, and the details of candidate objects that are used to guide the LVL2 trigger processing (the simulation of the output from the calorimeter trigger simulation to the DAQ has not yet been implemented.). Note that the RoI data from the simulation of the calorimeter trigger are stored in the same bytestream format as will be used in the hardware.

The HLT steering software requires the RoIs to be given with coordinates in η–ϕ space and a value (in GeV) of the E_T threshold that was passed, not in terms of the LVL1 internal data for-

mat (which reflects the details of the electronic implementation of the trigger). Software converters are provided to translate the raw 32-bit RoI data words [13-5] into objects, complete with methods to return the required data.

An effort has been made to validate the performance of the calorimeter trigger simulation [13-10]. Here we present some examples of efficiency and rate results for the electron / photon trigger.

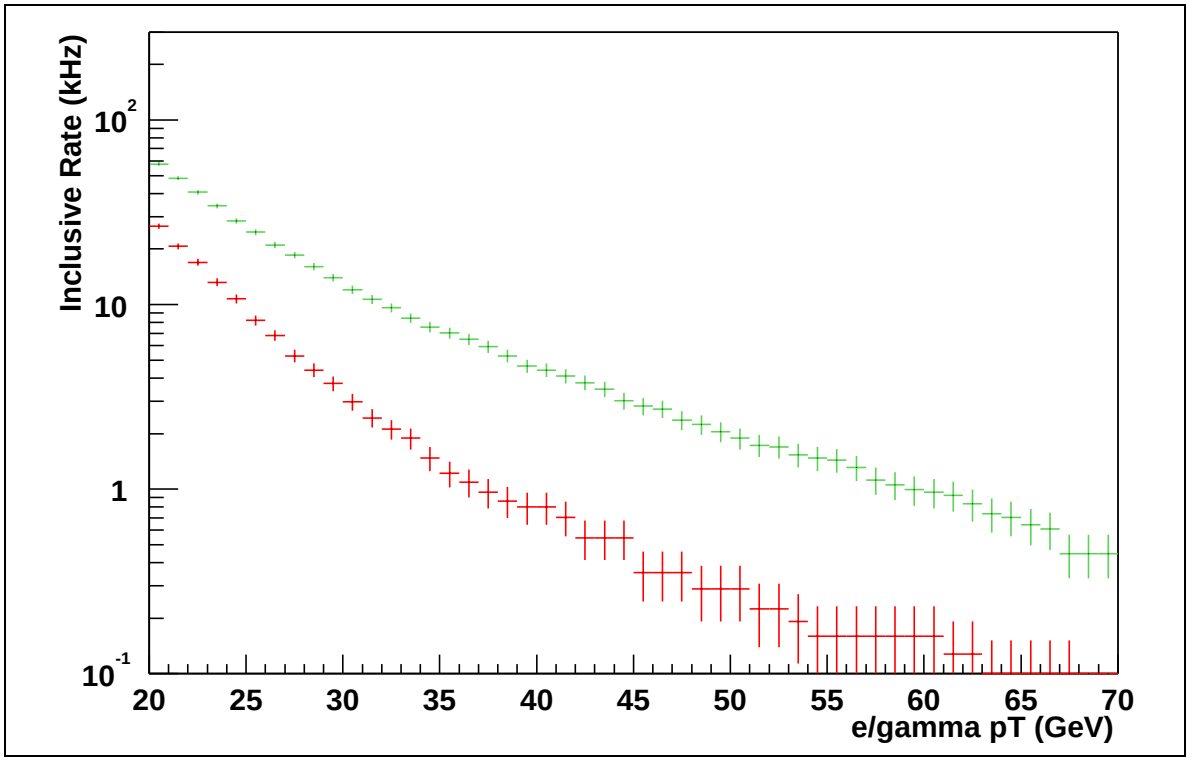

Figure 13-4 Electron/photon trigger rate versus E_T threshold without (top) and with (bottom) isolation requirements at low luminosity 2×10^{33} cm^{-2} s^{-1}. See the text for details.

Figure 13-3 shows the electron trigger efficiency as a function of p_T when requiring that the trigger cluster E_T be greater than 17 GeV. The electrons used in this study came from simulated $Z \rightarrow e^+e^-$ decays, and were required to be well isolated from other high-p_T particles (to ensure that the electron shower was responsible for the trigger) and restricted to a fiducial range $|\eta| < 2.45$. Results are shown without electronic noise or pileup, and also at design luminosity (1×10^{34} cm^{-2} s^{-1}) with electronic noise added. A sharp rise in the efficiency is observed around the threshold value in both cases, with little degradation due to noise and pileup

The rate of the LVL1 electron / photon triggers is dominated by misidentified jets. Figure 13-4 shows the estimated trigger rate for the single electron / photon trigger as a function of transverse-momentum threshold for a luminosity of 2×10^{33} cm^{-2} s^{-1}. The upper and lower bands give the rate before and after applying the isolation requirements. The E_T-threshold scale in the plot is defined so that the efficiency for selecting genuine electrons with transverse-energy equal to the quoted value is 95%. Figure 13-5 presents the rate for the electron / photon-pair trigger for the high-luminosity scenario of 1×10^{34} cm^{-2} s^{-1}, with and without isolation.

We have compared the results obtained using the new object-orientated software against those presented in the LVL1 TDR [13-2] which were obtained using the previous FORTRAN-based

software. The results are not expected to be identical since the studies used different version of Pythia for the event generation, there have been significant changes to the detector model, and some changes have been made to the RoI-selection and isolation algorithms. Note also that new simulation does not yet include the details of the signal-processing chain for the trigger-tower summation, digitization and bunch-crossing identification. In general there is fair agreement between the old and new results. For example the estimated rates for the single electron / photon trigger for a luminosity of 2×10^{33} cm^{-2} s^{-1} agree with the earlier simulations to within 20%. In view of the potential sensitivity of the isolation cuts to details of the trigger-tower simulation and to be conservative, we have based our estimates of event and data-rate requirements for the HLT / DAQ system on an extrapolation of the LVL1 TDR results.

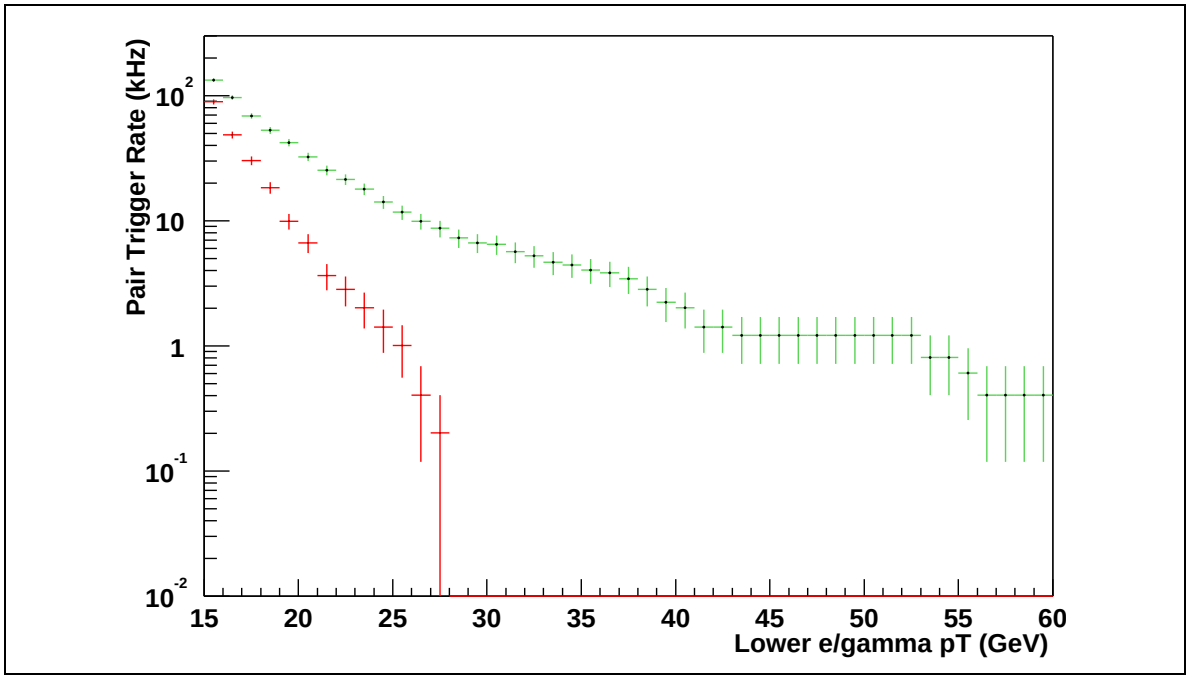

Figure 13-5 Predicted trigger rates for electron pairs without (top) and with (bottom) isolation at high luminosity 1×10^{34} cm^{-2} s^{-1}. See the text for details.

13.2.3 The RPC muon trigger and its simulation

The barrel muon trigger ($|\eta| < 1.05$) uses information from six layers of Resisitive Plate Chamber (RPC) detectors that are organised in three stations [13-11]. The middle RPC station (RPC2) is called the *pivot plane*. The principle of the algorithm that finds muon candidates is as follows [13-12]: each hit found in the pivot plane is extrapolated to the innermost RPC station (RPC1) along a straight line through the interaction point, and a *coincidence window* is defined around the point of intersection. Since muons are deflected in the magnetic field, the size of the coincidence window defines the p_T threshold of the trigger. A low-p_T muon candidate is found if there is at least one hit in the coincidence window and if in at least one of the stations RPC1 and RPC2 there are hits in both planes of the doublet. This condition must be satisfied for the measurements in the bending plane and also for those in the non-bending plane. If, in addition, there is a hit within a coincidence window in at least one of the two planes of the outermost station RPC3, a high-p_T candidate is found. Again the coincidence condition must be satisfied for the measurements in the bending plane and also for those in the non-bending plane.

In the hardware implementation of the trigger, the algorithm is implemented using logic that is programmed to recognize valid patterns of hits in the RPCs. All pivot-plane hits are processed in parallel, allowing a trigger with a very short fixed latency. There are three independently programmable thresholds for each of the low-p_T and high-p_T triggers. For each of the 64 sectors of the RPC trigger, up to two muon candidates can be selected and sent to the MuCTPI. If there are more than two candidates in a sector, the two highest-p_T candidates are retained and a flag indicates that additional candidates were suppressed.

The input to the simulation of the muon trigger logic was provided by a package that simulates the RPC detector system; this was done with the ATLSIM program. The muon detector layout used for this simulation was version “P03” [13-12]. The geometry of the RPC detectors and the stations built from them, as well as the positioning within the muon spectrometer, were reproduced in great detail according to the latest engineering drawings. The material composition and geometry of the single-RPC units were also correctly simulated. The simulation of the RPC detector response was based on results obtained in test-beam experiments. The hits produced in the simulation of charged particles crossing the RPC detectors were collected and stored in output files for use in downstream packages for the simulation of the trigger logic and also for the event reconstruction.

The detector simulation stage is followed by a number of packages which, using Athena algorithms and services, simulate in detail the hardware core and the overall architecture of the LVL1 muon barrel trigger. The hardware simulation is built from a set of classes which reproduce, up to the level of the internal data registers, the behaviour of the basic hardware components: the Coincidence Matrix ASIC (CMA), the Pad Board, the Sector Logic and the Read-Out Driver. The detector data are first accessed by the CMA, in which the RPC digits are translated into bits indicating which input channels fired. Channel masking, time alignment and the introduction of an artificial dead time for the fired channels are possible, although not used yet in the present implementation. The outputs of the eight CMAs belonging to each Pad Board are collected and are searched for valid trigger coincidences in the r–z and r–ϕ views. The Sector Logic then identifies the two highest-p_T muon candidates among all the candidates from all the Pad Boards in the corresponding sector. The output of the Sector Logic, including the addresses and p_T-threshold information for the RoIs, is finally stored in the transient event store of the Athena framework from where it can be retrieved by the MuCTPI simulation package. The CMAs supply information also to the read-out system; this data path is also simulated. The resulting data are organized in a structure that follows exactly the one implemented in the hardware (i.e. bytestream format). This, together with software converters of the interpretation of the bytestream data, allows the use of the RPC data in LVL2 selection algorithms such as muFast ([13-12] and [13-13]).

The architecture model takes care of connecting together the different simulated hardware components and of simulating the data flow in the RPC muon trigger. It is built using a data-driven approach: the event data are arranged in a hierarchical structure going from the RPC digits up to the Sector Logic output. Each level in the hierarchy corresponds to a complete processing step of one hardware component, and the result of each step is available on demand. The “on-demand” approach is very flexible and allows one to save processing time since simulation is only requested for the sectors that contain data.

The architecture simulation makes use of two services that describe the setup of the trigger system: the geometry of the trigger chambers and the cabling of the hardware components. The geometry service accesses the RPC engineering parameters via the AMDB database and builds a model in which the chambers are organized in a suitable way for the LVL1 trigger (i.e. in trigger stations, building a continuous detector surface at fixed radius, and not as appendices of the

MDT chambers as it is in the database). The cabling service provides the mapping between the trigger components (RPC readout strips to CMA channels, CMAs to Pad Boards, and Pad Boards to Sector Logic) and holds the configuration of the CMAs. This configuration depends on the muon p_T thresholds which are required. The cabling data as well as the CMA configuration data are read from ASCII files.

The packages that implement the architecture model also provide a fast simulation of the trigger logic that is completely disentangled from the hardware simulation and does not take the timing of signals into account. The main application of this fast simulation is inside the LVL2 selection algorithm muFast, where it is used to emulate the functionality of the LVL1 trigger. Some information on the performance of the RPC muon trigger can be found in [13-13] and [13-14].

13.2.4 The Muon-to-CTP interface and its simulation

The Muon-to-CTP Interface (MuCTPI, see Ref. [13-15]) receives up to two muon candidates from each of the 208 sectors of the barrel (RPC) and endcap (TGC) muon triggers. From these candidates, the multiplicities of muons are calculated for up to six different muon p_T thresholds and sent to the CTP for use in making the trigger decision. If the event is retained by the LVL1 trigger, RoI information on the muon candidates is sent via the Region-of-Interest Builder to the LVL2 trigger (if there are more than 16 candidates, the 16 highest-p_T candidates are retained and a flag is set). More detailed information is sent on a separate link to the DAQ. The data sent to LVL2 and the DAQ conform to the standard ROD bytestream data format. The MuCTPI avoids double-counting of muons which pass through more than one pivot plane, e.g. in the barrel–endcap transition region.

The MuCTPI simulation follows the hardware scheme as closely as possible, down to the data formats used in the hardware. The data flow is emulated using the same stages of processing as in the hardware, including the propagation of error and status bits. Access functions are provided for every type of information available in the hardware. The simulation was originally developed as a stand-alone program for testing the prototype MuCTPI hardware. It has recently been ported to the Athena framework and integrated with the simulation of the RPC trigger on the input side, and with the simulations of the CTP and the RoIB on the output side. The output to the readout is also simulated although this is not yet used within the LVL1 simulation efforts.

13.2.5 The LVL1 CTP and its simulation

The LVL1 trigger event decision is made in the Central Trigger Processor (CTP, see Ref. [13-16] and [13-17]) in the two-step procedure that was discussed above:

The multiplicities of candidate objects from the calorimeter and muon triggers for various p_T thresholds are compared with the *trigger conditions* introduced in Section 13.2.1, checking against simple multiplicity requirements. Depending on the inputs from the calorimeter and muon trigger, each trigger condition takes a logical value TRUE or FALSE.

The trigger conditions (or rather their logical values) are combined using AND, OR and NOT operations to give complex *trigger items*. Each trigger item corresponds to a signature to be triggered on by LVL1 as defined in the trigger menu; gates and prescales can be applied to each individual item. The LVL1 trigger result is the logical OR of all trigger items.

The logical relations between the conditions and the items on one side, and the conditions and the input multiplicities on the other side, are provided by the LVL1 trigger configuration (Section 13.2.1). The CTP provides outputs to the RoIB and to the readout; the information that is sent includes bit patterns for the input signals and for the trigger items before and after prescales and vetos, as well as the final event-decision result.

In the currently-existing prototype hardware implementation of the CTP, the CTP-D ('D' for demonstrator, see Ref. [13-18] and [13-19]), the selection procedure is implemented using two look-up tables (LUT) to compute the trigger conditions and two programmable devices (CPLDs) for the combination of conditions to trigger items. The design of the final CTP is in progress based on the use of large FPGAs.

The current CTP simulation follows closely the CTP-D design — conversion to the final CTP design will be done in due course. First, the input multiplicities from the calorimeter trigger and MuCTPI simulations are collected. The multiplicities that are used in the trigger menu are checked against the respective conditions (the conditions are taken from the C++ object representing the trigger menu that is provided by the configuration step). The conditions are then combined to trigger items in a recursive algorithm. All trigger items are passed through a prescale algorithm, and the logical OR of all items is formed, resulting in the LVL1 event decision (the LVL1 accept or L1A signal which can be expressed as 0 or 1, FALSE or TRUE). The deadtime algorithms that exist in the hardware have not yet been implemented in the simulation. Finally, the CTP result object, which in content and format corresponds precisely to the one provided by the hardware, is formed and stored in the TES for later use by the RoIB.

13.2.6 Interface to the HLT

The interface between the LVL1 trigger and the HLT is the Region-of-Interest Builder (RoIB, see Ref. [13-20]). This device, which is formally part of the DataFlow system, collects the information relevant for LVL2 from the calorimeter and muon triggers and from the CTP. It combines all the data into a single block which it then passes to the LVL2 supervisor assigned to the event in question. The LVL1 data are received in S-LINK format (four links from the calorimeter trigger Cluster Processor, two from the calorimeter Jet/Energy Processor, and one each from the MuCTPI and the CTP). The RoIB has to operate at the highest foreseen LVL1 output rates without introducing additional deadtime.

The RoIB simulation picks up the information stored in the TES by the calorimeter trigger, MuCTPI and CTP simulations, and constructs a LVL1 result raw-data object (RDO). Converters are provided for translating the bytestream format into objects which serve to seed the LVL2 trigger. These contain the value (in GeV) of the threshold that has been passed and the location of the RoI in $\eta-\phi$ space.

13.3 Common tools for on-line HLT selection

The HLT algorithms are the basic software components that provide results from which the trigger decision is derived. These algorithms operate within the context and environment of the Event Selection Software (ESS) which was already discussed in Chapter 9. Section 13.3.1 provides an overview and description of the ESS from the viewpoint of the HLT algorithms. The objects of a common Event Data Model, i.e. the objects that the algorithms exchange and manipulate, are described in Section 13.3.2. A description of the HLT algorithms intended to operate

in the LVL2 environment is given in Section 13.3.3, while the algorithms for the EF are described in Section 13.3.4.

13.3.1 Algorithmic view of the Event Selection Software

Unlike their counterparts from the offline reconstruction, HLT algorithms must allow themselves to be guided by the Steering of the Event Selection Software. The Steering is configured with sets of Sequence and Menu Tables and guides the HLT algorithm processing using so-called Trigger Elements. Trigger Elements characterizing abstract physical objects with a label (e.g., 'electron', 'jet', or 'EMRoI') and effectively decouple the Steering of the trigger selection from details of the reconstruction Event Data Model. The ESS uses the seeding mechanism (see Section 9.5) to guide the HLT algorithm processing to the those subsets of event data that correspond to the LVL1 RoIs.

The HLT trigger processing is data driven, because the set of LVL1 RoIs determines which of the predefined Sequences are to be executed for this event. Each of the LVL1 RoI objects are associated to a Trigger Element on which the Steering acts upon. For each Trigger Element, the Steering executes the required HLT algorithms as defined in the Sequence Table. Hence, it is possible that a given algorithm may be executed several times per event. This is fundamentally different to the 'Event Loop' approach of the offline reconstruction paradigm where a given offline algorithm acts only once upon each event.

During the LVL2 trigger latency event data remain within ROBs unless and until they are actively requested. The LVL2 algorithms request and process only a small fraction of the event data, leading to a very substantial reduction in the network and computation resources that would otherwise be required.

The first step in the process of obtaining data for a RoI is the conversion of a geometrical region (e.g. a cone with an extent η and ϕ) into Identifiers; this is accomplished with the HLT Region Selector [13-21]. Presently these Identifiers are the `IdentifierHash`es used in the offline software. Each Identifier corresponds to a Detector Element in a sub-detector, e.g., a Pixel Module or a sampling of a LAr Trigger Tower. The Region Selector uses information from the detector description during its initialization phase to build an so-called `EtaPhiMap` for each layer (or disk) of a sub-detector. This map is essentially a two-dimensional matrix in η and ϕ. Each element of the matrix consists of a list of Identifiers. An HLT Algorithm uses the Region Selector interface, which is the name of the sub-detector to request data from (*i.e.*, Pixel, SCT, TRT, LAr, Tile, MDT, RPC, CSC, or TGC) and the physical extent of the geometrical region, to request the list of identifiers. Given the vastly different designs of the sub-detectors, different procedures are used dependent on sub-detector.

Interactions with the Data Collection system are hidden from the HLT algorithm behind a call to Storegate. Within Storegate, sub-detector data are aggregated into collections within a container. Each collection is labelled with the Identifier corresponding to the Detector Element. Algorithms request event data from Storegate using the set of Identifiers obtained by the HLT Region Selector. If the collections are already within Storegate, it returns them to the HLT algorithm. If not, Storegate uses the Raw Data Access scheme (Chapter 9) to request the corresponding ROB data from the Data Collection system. A converter translates the Raw Data into either Raw Data Objects (RDOs) or, by invoking an Data Preparation algorithm, into Reconstruction Input Objects (RIOs). The obtained RDOs or RIOs are stored within the collections within the Container within Storegate.

13.3.2 Event Data Model Components

During 2002 and 2003, there has been a substantial effort within the HLT, offline, and sub-detector communities to establish a common Event Data Model (EDM). In Section 13.3.2.1 the concept of Detector Elements is discussed as an organizing and identifying principle for most EDM objects. The raw-data model, that forms the present agreed organization of trigger objects and their relationship, is described in Section 13.3.2.2 and reconstruction data-model classes specific to LVL2 and EF algorithms are described in Section 13.3.2.3.

It is worth noting that a large amount of effort has been invested by ATLAS as a whole (TDAQ, offline and detector communities) in making as realistic as possible a simulation of the data that the HLT/DAQ will have to deal with.The detector response is simulated using a detailed GEANT3-based simulation of the ATLAS detectors. The simulation of the raw data includes the effects of pile-up in the bunch crossing that gave rise to the LVL1 trigger, and also in other bunch crossings nearby in time, with an average number of interactions per bunch crossing chosen according to the relevant luminosity scenario. Electronic noise is also included in the simulation. The “time-frame” of digitized data (i.e. measurements in successive bunch crossings) are then processed to reproduce in as much detail as possible the raw data as they will appear at the output of the Readout System. For example, in the LAr calorimeter digitizations from five successive bunch crossings are combined to extract measurements of energy, time and quality. In the experiment, this signal processing will be performed at the level of the Readout Driver.

13.3.2.1 Event Data Organization

Event Data (e.g., Raw Data Objects (RDOs) and Reconstruction Input Objects (RIOs)) are aggregated into collections corresponding to adjacent readout channels within the physical detector. These collections reside in an Containers with collection Identifier labels corresponding to the unit of aggregation. For most sub-detectors, the organizing principle is that of the Detector Element.

In the Pixel detector a Detector Element is a module, equivalent to a single Silicon wafer; there are 1744 such elements. For the SCT, a Detector Element is one side of a module, equivalent to a bonded pair of wafers whose strips are oriented in a single direction (*i.e.*, axial or stereo); there are 8176 SCT Detector Elements. For the TRT no such direct correspondence to detector module exists. There are 19008 TRT Detector Elements [13-22].

For the calorimeters, the concept of Detector Element is difficult to define. Instead, the organizing principle for event data is that of the LVL1 Trigger Tower and sampling.

Within the barrel muon spectrometer an MDT Detector Element corresponds to one of the two multi layer chambers in a station. For the LVL2, the RPC data is organised in Detector Elements according to the LVL1 coincidence logic. A Detector Element is identified with the RPC pad associated with the second station. For the EF the RPC data is reorganised into a Detector Element for each station. A TGC Detector Element is one TGC η division, or chamber, in a TGC station; there are 24 forward stations in a ring and 48 endcap stations in a ring, and there are four rings at each end of the ATLAS detector. Finally, the CSC Detector Element is a single CSC chamber in a station. A CSC station typically has two multilayers.

13.3.2.2 Raw Data Model Components

Byte Stream Raw Data are ROB-formatted data produced by the ATLAS detector or its simulation [13-22]. The format supports a set of hierarchical fragments, where only the bottom level, the ROD fragment, is defined by the sub-detector group. Although the format of the Byte Stream has not been completely finalized, the detector groups have implemented the Raw Data simulation according to the best available information.

A Raw Data Object (RDO) is uncalibrated Raw Data converted into an object representing a set of readout channels. Historically these data have been referred to as *Digits*. It is the representation of Raw Data which is put into the Storegate and is potentially persistifiable.

The purpose of the RDO converters is dual: first a Byte Stream file can be created by taking the information from already created RDOs (in the transient store, from ZEBRA); second, a Byte Stream file can be read by the converters to fill the RDOs (or the RIOs for LVL2). The same converters are used for decoding the Raw Data from the online Data Collection. Since the RDOs are a representation of the specific detector output, their content can change during the lifetime of the sub-detectors.

A detailed description of the Raw Data Model components is available elsewhere [13-23].

13.3.2.3 Reconstruction Data Model Components

Algorithms interact with Reconstruction Input Objects (RIOs) as opposed to RDOs. For each sub-detector system, classes of RIOs have been defined and are described in the following subsections.

13.3.2.3.1 Inner Detector

For the Inner Detector the RIOs are organised into Detector Element collections.

The Pixel and SCT RIOs are Clusters. A Cluster in the Pixel detector is a two-dimensional group of neighbouring readout channels in a Detector Element. A Cluster in the SCT is a one-dimensional group of neighbouring readout channels in a Detector Element. For Pixel and SCT, there are currently two implementations of the Cluster class: one used for the EF and offline, and one used for LVL2. The one used in the EF has Pixel and SCT sharing the same class. For LVL2 there is a common structure for Pixel, SCT and TRT, but each subdetector has its own concrete class. Both the LVL2 and EF sets of cluster classes contain a list of RDO identifiers from which the cluster is built. The number of member functions is limited in both set of classes and the member functions follow the Inner Detector Requirements [13-22].

At LVL2, Pixel and SCT RIOs are converted to 3-dimensional coordinates in the ATLAS global coordinate system using an HLT algorithm to create Space Points. These algorithms accept as input a Container of Cluster Collections of the appropriate type and return an Container of Space Points.

For the Pixels, the creation of Space Points consists of combining the local coordinates of Clusters with information on the position and orientation of the Detector Element to give the global coordinates. The process for the SCT is more complicated since a single SCT detector provides only a one-dimensional measurement. However, an SCT module, consisting of two detectors in a stereo-pair, provide two-dimensional information. One species of SCT Detector Element, phi-

layer, has strips orientated parallel to the beam axis, the other, *u* or *v* layer, is rotated by ± 40 mrad with respect to those of the phi-layer Detector Elements. The formation of Space Points consists of the following steps:

- Associate each phi-layer Cluster Collection[1] with the corresponding stereo-layer Cluster Collection.
- For each pair of Collections (phi + stereo), take each phi-layer Cluster and search for associated stereo-layer Clusters. If there is more than one associated stereo layer Cluster, a Space Point is formed for each one (in this case, at most one will be a correct measurement, the others will form 'ghost' points). If no associated stereo-layer hit is found, a point is created from the phi-layer information alone.
- Calculate the second coordinate (*z* for the barrel, or *R* for the end-caps).
- Using information on the position and orientation of the Detector Element transform to global coordinates.

Note that for the LVL2 Space Points a simplifications is made in the interest of speed. No attempt is made to form Space Points from Tracks passing close to the edge of a module, where the corresponding stereo-layer Cluster is in a different module.

The TRT RIO gives the drift circle of a straw. Here, the same classes are used for LVL2, EF and offline. The granularity of the TRT RIO is the same as for the RDO: that of a straw, thus the RIO contains an identifier which is the offline identifier for a straw. In the case of the RDO the straw information is uncalibrated, while in the case of the RIO the straw information is calibrated, i.e. the drift radius is calculated from the drift time. For the current prototype, the time–radius function is the same for all straws, but in the future the parameters of the function will come from the Interval of Validity Service [13-24].

13.3.2.3.2 Calorimeters

For the Calorimeters, the RIOs are calibrated calorimeter cells (`LAr` and Tiles), the same as in the offline reconstruction.

Both LAr and Tile have a common `CaloCell` base class which represents the measurement in the calorimeters of energy, position, time and quality. A `CaloCell` has been calibrated so that it returns the physical energy deposit *in the cell* with units of GeV, but without any kind of leakage corrections. Time is given in nanoseconds and refers to when the deposit occurred, relative to the trigger; it should be consistent with zero for good hits. Quality reflects how well the input to the system matched the signal model on which the algorithm is based.

13.3.2.3.3 Muon Spectrometer

For the barrel Muon Spectrometer it was found expedient to use in part RDOs instead of RIOs as input to the HLT muon-selection algorithms. The RDOs are organized inside Storegate in identifiable collections and can be accessed in the same way as RIOs. For the MDTs, the RDOs and RIOs are ordered in collections that correspond each to an MDT Detector Element, i.e. an MDT chamber. Each RDO [13-25],[13-26],[13-27] contains the information of one MDT read-out channel, *i.e.* one tube. The information is the time of the leading edge of the MDT pulse and the

1. There is a Cluster Collection per Detector Element.

charge integrated over the first 20 ns. From this, the uncalibrated drift time can be calculated. The RIOs contain the calibrated drift time.

The definition of RPC RDOs is complicated by the fact that the RPCs are trigger chambers. Their read-out is optimised for the LVl1 trigger task and does not reflect an easily identifiable geometrical structure such as an RPC chamber with its strips. Consequently, RPC RDOs are ordered following the read-out structure, by PADs and CMAs. Each RDO corresponds to a fired CMA channel [13-28], [13-29]. The RDOs are organized in CMA objects, i.e. in collections corresponding to one CMA each. The CMA objects in turn are organized per PAD, i.e. in collections corresponding to one PAD each.

One RPC RIO corresponds to an RPC Detector Element. An RPC RIO contains the information of a collection of RPC strips that fired. There is no simple correspondence between RPC strips and RPC read-out channels. In order to translate a fired RPC read-out channel, which is a CMA read-out channel, into a RPC strip, a cable map and processing of the information of the CMAs for the opposite view is required.

Corresponding class definitions are being developed for the TGC and CSC detectors.

13.3.3 HLT Algorithms for LVL2

LVL2 is the most demanding step, in terms of required system performances (latency, data transfer, etc.), of the HLT trigger selection. In the following the present view of the algorithms needed to implement LVL2 selection is given. It is worthwhile noticing that several options for using in the best possible way detector information are taken into account, hence more than one algorithm is available to accomplish a defined task. This will allow to implement a robust, flexible and redundant selection scheme, which will be studied with present and future simulations.

13.3.3.1 IDSCAN

IDSCAN (see Refs. [13-30] and [13-31]) is a track-reconstruction package for LVL2. It takes as input Space Points found in the Pixel and SCT Detectors. A series of sub-algorithms (Z-Finder, Hit Filter, Group Cleaner, Track Fitter) then processes these inputs and output Tracks and the Space Points associated with them.

The Z-Finder determines the *z*-position of the primary interaction vertex. The algorithm assigns the input hits to narrow ϕ-bins. For each bin it then extrapolates lines joining pairs of hits back to the beam-line, entering in a histogram the *z* values obtained for all hit-pairs. After integrating over all ϕ-bins, it takes as the *z*-position of the primary vertex the histogram region with the most entries.

The Hit Filter finds groups of hits compatible with Tracks coming from the *z* position found by Z-Finder. It puts all hits into a histogram binned in ϕ and η. It then finds clusters of hits within this histogram. It creates a *group* of hits if such a cluster has hits in more than a given number of layers.

The groups of hits found by the Hit Filter are used by the Group Cleaner which splits them into Tracks and removes noise hits. Each triplet of hits forms a potential track for which p_T, ϕ_0, and d_0 are calculated. It forms groups from these triplets with similar parameters, applying certain quality cuts. It accepts a track candidate if a group contains enough hits.

Finally, the Track Fitter verifies track candidates and calculates the track parameters by using a fitting algorithm adapted from SCTKalman [13-32][13-33][13-34]. It returns a list of Space Points on the Track, the Track parameters, and an error matrix.

13.3.3.2 SiTrack

SiTrack is a track-reconstruction package for LVL2 which extends and upgrades a previous algorithm called PixTrig. SiTrack takes Pixel and SCT Space Points as input, and outputs fitted reconstructed Tracks each storing pointers to the Space Points used to build it. SiTrack is implemented as a single main algorithm SiTrack which executes a user-defined list of sub-algorithms (chosen among ST-Space Point Sorting, ST-Muon Vertex, ST-Track Seeding, and ST-Three Point Fit).

The ST-Space Point Sorting collects the Space Points coming from the Pixel and SCT detectors and sorts them by module address, storing the result in a Standard Template Library (STL) map. This processing step is performed in order to speed-up data access for the other reconstruction sub-algorithms.

The ST-Muon Vertex is a primary-vertex identification algorithm mostly suitable for low-luminosity events with a high-p_T muon signature. It is based on track reconstruction inside a LVL1 muon RoI — the highest-p_T track is assumed to be the muon candidate or, failing that, to come from the same primary interaction as the muon. The primary-vertex position along z is taken from the point of closest approach of the track to the beam line.

The ST-Track Seeding uses the sorted Space Point map and a Monte Carlo look-up table (MC-LUT) linking each B-layer Pixel module to modules belonging to other logical layers. It builds track seeds formed by two Space Points and fits them with a straight line; one or more logical layers can be linked to the B-layer, the latter option being particularly useful if robustness to detector inefficiencies must be ensured. If the primary vertex has already been reconstructed by ST-Muon Vertex, a fraction of fake track seeds can be rejected during their formation, applying a cut on their z distance from the primary vertex.

If no vertex information is available, a histogram, whose resolution depends on the number of seeds found, is filled with the z impact parameter of each seed. The z position for the primary vertex is then determined from the position of the maximum in the histogram. This vertexing algorithm, which can be operated in both RoI and full-scan modes, is best suited for high-luminosity events containing many high-p_T tracks (e.g., b-tagging). Independent cuts on r-ϕ and z impact parameters are eventually applied to the reconstructed seeds to further reduce the fake fraction.

ST-Three Point Fit extends track seeds with a third Space Point; it uses a map (which is built using Monte Carlo data) associating to each seed a set of module roads[1] the track could have hit passing through the Pixel or SCT detectors. A subset of modules is extracted from each road according to a user-defined parameter relating to their ‘depth’ inside it (e.g., the user can decide to use modules at the beginning or in the middle of each road, etc.). Space Points from the selected modules are then used to extend the seed, and candidate tracks are fitted with a circle. Ambiguities (e.g., tracks sharing at least one Space Point) can be resolved on the basis of the track quali-

1. A road is a list of modules ordered according to the radius at which they are placed starting from the innermost one.

ty, to obtain an independent set of tracks that can be used for trigger selection or as a seed for further extrapolation.

13.3.3.3 TRTLUT

TRT-LUT is a LVL2 algorithm for track reconstruction in the TRT. It is described in detail elsewhere [13-35][13-36]. The algorithm takes as input Hits in the TRT. The algorithmic processing consists of Initial Track Finding, Local-Maximum Finding, Track Splitting, Track Fitting and Final Selection. It outputs the Hits used and Tracks with their parameters.

During the Initial Track Finding a histogramming approach is employed. A two-dimensional histogram is constructed with bins in ϕ and Q/p_T for the track parameters at the primary vertex (Q indicates the electric-charge sign). For each hit in the event, the contents for all ϕ and Q/p_T bins consistent with the hit position are incremented. The bin corresponding to the true parameters of a track will receive entries for all hits on that track. A very fast implementation of the algorithm is achieved by using a Look-Up Table (LUT) to obtain the list of bins associated with given hit. Having built the histogram, tracks can be found by Local Maximum Finding.

The Track Splitting stage of the algorithm analyzes the pattern of hits associated to a track candidate. By rejecting fake candidates composed of hits from several low-p_T tracks, the track splitting step results in an overall reduction by a factor of roughly two in the number of track candidates. For roads containing a good track candidate, it identifies and rejects any additional hits from one or more other tracks. The result of the overall Track Splitting step is a candidate that consists of a sub-set of the straws within a road.

The final step of TRT-LUT, Track Fitting and Final Selection, performs a fit in the r-ϕ (z-ϕ) plane for the barrel (end-caps) using a third-order polynomial to improve the measurement of ϕ and p_T. Only the straw position is used (*i.e.*, the drift-time information is not used). The track is assumed to come from the nominal origin. After the fit, a reconstructed p_T threshold of 0.5GeV is applied.

13.3.3.4 TRTKalman

TRT-Kalman [13-37][13-38] is a package based on xKalman++ (see Section 13.3.4.1). The name is in fact a misnomer since the Kalman-filter component of xKalman++ is not used for the TRT; a histogram search and least-squares fit is used instead.

13.3.3.5 T2Calo

T2Calo (see Refs. [13-39], [13-40], [13-41] and [13-42]) is a clustering algorithm for electromagnetic (EM) showers, seeded by the LVL1 EM trigger RoI positions [13-43]. This algorithm can select isolated EM objects from jets using the cluster E_T and shower-shape quantities.

The RIOs from which it starts are calibrated calorimeter cells (LAr or Tiles), as in the offline reconstruction. The output is a LVL2-specific class containing the cluster energy and position, and the shower-shape variables useful for the selection of EM showers.

The first step in T2Calo is to refine the LVL1 position by finding the cell with highest energy in the second sampling of the EM calorimeter. This position (η_1, ϕ_1) is later refined by calculating the energy-weighted-average position (η_c, ϕ_c) in a window of 3×7 cells (in $\eta \times \phi$) centred

around (η_1, ϕ_1), using as weights the energy in the second sampling. As described in Ref. [13-40], a number of parameters are calculated for use selecting EM clusters and rejecting background due to jets:

- In sampling two, $R^{shape}_{\eta} = E_{37} / E_{77}$ is calculated. The expression E_{nm} stands for the energy deposited in a window of $n \times m$ around (η_1, ϕ_1).
- In sampling one, $R^{strip}_{\eta} = (E_{1st} - E_{2nd}) / (E_{1st} + E_{2nd})$ is obtained in a window of $\Delta\eta \times \Delta\phi = 0.125 \times 0.2$ around (η_c, ϕ_c). Here E_{1st} and E_{2nd} are the energies of the highest and second-highest local maxima found after summing over the two strips in ϕ for each position in η within the window. A local maximum is defined as a single strip with energy greater than each of its two adjacent strips.
- The total energy, E, deposited in the EM calorimeter is calculated in a window of 3×7 cells around (η_1, ϕ_1), summing over all three samplings.
- Finally, the energy that leaks into the hadron calorimeter E^{had} is calculated in a window of size $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ around (η_c, ϕ_c).

13.3.3.6 muFast

The muFast algorithm is a stand-alone LVL2 tracking algorithm for the Muon Spectrometer. Previous versions are described in detail elsewhere [13-13].

The algorithm is steered by the RoI given by the LVL1 Muon Trigger and uses both RPC and MDT measurements. At present this algorithm is limited to the barrel region and is based on four sequential steps:

1. LVL1 emulation: pattern recognition in the MDT system is initiated by the RPC hits that formed the LVL1 track candidate. The RoI information obtained directly from LVL1 (via the RoI Builder) gives only the position of the track at the pivot plane. Details of the RPC hits that made up the LVL1 track therefore have to be obtained by running a fast algorithm that repeats the basic logic of the LVL1selection.
2. Pattern recognition is performed using the RPC hits that made up the LVL1 track candidate to define a road in the MDT chambers around the muon trajectory. MDT tubes lying within the road are selected and a contiguity algorithm is applied to remove background hits not associated with the muon track.
3. A straight-line track fit is made to the selected tubes (one per tube monolayer) within each MDT station. For this procedure the drift-time measurements are used to exploit fully the high measurement accuracy of the muon-tracking system. The track sagitta is then evaluated by combining the measurements from the different stations.
4. A fast p_T estimate is made using LUTs. The LUT encodes, as a function of η and ϕ, the linear relationship between the measured sagitta and Q/p_T.

The output of this algorithm is a measurement of the muon p_T at the production point, together with η and ϕ. The detailed description of the implementation of muFast in the trigger framework is given in Ref. [13-44].

13.3.3.7 muComb

The combination of the features of tracks measured at LVL2 in the muon spectrometer and the Inner Detector (ID) provides rejection power against decays in flight of charged pions and kaons (i.e. $\pi^+ \rightarrow \mu\nu$ and $K^+ \rightarrow \mu\nu$), and of fake muon-spectrometer tracks composed of hits induced by the cavern background. In addition, the combination of the two measurements improves the momentum resolution of reconstructed muons over a large momentum range.

The matching of tracks between the muon spectrometer and the ID can be performed extrapolating the ID track to the muon detectors. The procedure needs to take into account the detector geometry, the material distribution and the inhomogeneity of the magnetic field. An accurate extrapolation would require the use of detailed geometry and magnetic-field databases, together with a fine tracking. All this would be expensive in terms of CPU time and therefore not acceptable for the LVL2 trigger.

To provide a fast tracking procedure, the effects of the geometry, material and magnetic field have been parameterized using simple analytic functions of η and ϕ. The extrapolation of ID tracks to the entrance of the muon spectrometer is performed using straight-line extrapolation in two independent projections, the transverse and the longitudinal views, corrected for bending in the magnetic field and energy-loss effects. In the transverse projection the ID track extrapolation in ϕ is corrected as follows:

$$\Delta\varphi = \frac{Q\alpha}{p_T - p_T^0} \qquad \text{13-1}$$

where α is related to the field integral, and p_T^0 accounts for the transverse-energy loss in the material of the calorimeter that is approximately independent of the track transverse momentum, p_T. Both α and p_T^0 have been determined by fitting $\Delta\varphi$ for simulated muons as a function of p_T. It is found that $p_T^0 \sim 1.5$ GeV, *i.e.* about half of the transverse-energy loss of low-momentum muons, as naively expected. A similar approach has been followed for the extrapolation of the *z*-coordinate in the longitudinal view.

The matching is done using cuts on the residuals on the positions in the orthogonal coordinates *z* and $R\phi$. For matching tracks the transverse momentum of the muon is estimated using a weighted average of the independent p_T measurements in the muon spectrometer and the ID. For each combined track, a χ^2 function is used to evaluate the quality of the p_T matching. Reconstructed muons from π and K decays in flight often give high χ^2 values (e.g. because the ID measures the hadron momentum while the muon-spectrometer measures that of the muon which is lower) and thus can be rejected thanks to the excellent resolution of both the ID and muon-spectrometer measurements.

13.3.4 HLT Algorithms for EF

The baseline option for the selection at the Event Filter stage is to adopt algorithms from the offline software suite. This will allow the online selection to easily benefit from improvements stemming from offline analyses and to avoid duplication of efforts. In the following, a preliminary list of the presently available algorithms and their characteristics can be found, which have been used in the EF selections described in Section 13.4.

13.3.4.1 xKalman++

xKalman++ is a package for global pattern recognition and Track fitting in the Inner Detector for charged tracks with transverse momentum above 0.5GeV. A more detailed description of this algorithm is available elsewhere [13-45].

The algorithm starts the track reconstruction in the TRT using a histogramming method or in the Pixel and SCT detector layers using segment search.

The first reconstruction method outputs a set of possible track-candidate trajectories defined as an initial helix with a set of parameters and a covariance matrix. As a second step the helix is then used to define a track road through the precision layers, where all the measured clusters are collected. xKalman++ attempts to find all possible helix trajectories within the initial road and with a sufficient number of clusters.

The other method, where track finding starts in the Pixels or SCT, outputs a set of Space Points as an initial trajectory estimation. In the next step the Space Points serve as an input for the Kalman filter-smoother formalism that will add the information from the remaining precision layers. Each reconstructed Track is then extrapolated into the TRT, where a narrow road can be defined around the extrapolation result. All TRT Clusters together with the drift-time hits found within this road are then included for the final track-finding and track-fitting steps.

There are three seeding mechanisms available in the offline environment: the reconstruction of the full event; the reconstruction of a region-of-interest and (available soon) EM calorimeter seeding. In the HLT environment, used as an EF algorithm, xKalman++ will be seeded by the LVL2 result.

After the pattern-recognition and Track-fitting steps, xKalman++ stores the final Track candidates in a collection in Storegate. The Track candidate contains the following information:

- fit procedure used;
- helix parameters and their covariance matrix at the end-points of the filter procedure in the precision layers (point on the trajectory closest to the vertex) and in the TRT (point on the trajectory closest to calorimeter);
- total χ^2 resulting from final fit procedure;
- list of all hits on track from all subdetectors;
- total number of precision hits N_p.
- total number of straw hits N_s, empty straws crossed N_e, and of drift-time hits N_t.

Furthermore, a track candidate is stored in the final output bank if it passes the following cuts:

- the number of precision hits is larger than 5 to 7;
- the ratio $N_s/(N_s+N_e)$ is larger than 0.7 to 0.8;
- the ratio N_t/N_s is larger than 0.5 to 0.7;
- no track accepted previously has the same set of hits as the current one — this last cut removes full *ghost tracks.*

13.3.4.2 iPatRec

iPatRec [13-46] is a pattern-recognition algorithm used in the Event Filter that searches for tracks starting from Space Point combinations. Pixel and SCT Space Points are used to form Track candidates. Candidates are extrapolated to the TRT and drift-time hits are added. At the initialization phase, iPatRec creates a geometry data-base describing the properties of each detector module in the precision tracker plus the module's relationship to a simplified material model. This model consists material “layers” assumed to be either concentric cylinders in the barrel region or planes normal to the beam-axis in the end-cap regions. Additional “layers” represent the TRT detector, beam-pipe and inert support / service material. Track finding, following and fitting procedures make extensive use of this data-base. Another initialization task is to parameterize the magnetic field to enable a fast propagation of track parameters between layers.

In the first step of event reconstruction, adjacent Raw Data channels are clustered, and Space Points are produced from these Clusters. Each Space Point is assigned to one of seven partitions according to its distance from the beam intersection region. Within each partition the points are ordered according to their azimuthal coordinate. The general procedure is to form Track candidates using Space Point combinations from three different partitions, subject to criteria on maximum curvature and crude vertex region projectivity. Candidates then undergo a Track Fit procedure to give Track parameters with covariance at the point of closest approach to the beam-line (perigee parameters). The Track Follower algorithm propagates these parameters to form an intersect with error ellipse at each “layer” in turn. Clusters from the traversed detectors are associated to the Track. Final allocation of Clusters are taken after a further track-fit. During this fit, energy loss and Coulomb scattering are taken into account by allocating a scattering centre (with associated extra fit parameters) to each “layer” traversed. An active detector region traversed without an associated cluster is classified as a “hole” and retained for material and quality information. Tracks with a fit probability greater than 0.001 and a maximum of three holes are extrapolated to the TRT, where a histogramming technique is used to select the TRT hits to be added to the Track. Tight cuts are made on the straw residual and on the ratio of found to expected straws, in order to limit high-luminosity occupancy effects.

Tracks with clusters in the two innermost Pixel layers plus TRT association are called primary Tracks. Otherwise at most one hole is allowed: truncated Tracks start in the innermost layers but cannot be followed to the outermost layers or TRT; secondary Tracks start further out and are required to have TRT association. Various partition combinations are taken to maintain track-finding efficiency for the three types of Track even in the event of a higher than expected detector inefficiency. To avoid track duplication, only candidates with two unallocated Space Points are initiated, and tracks sharing more than 50% of their Clusters are deemed ambiguous whence only the one with higher quality is retained.

To speed up the execution, a preliminary track-finding pass looks only for high-quality primary Tracks and finishes as soon as one is found. The vertex from this Track is then used to subdivide the Space Point partitions into projective slices, greatly reducing the combinatorial load. Impact-parameter criteria are adjusted according to the distance to the first cluster to ensure there is no bias against b-, c- or s-particle decays. The code iterates to allow for several interaction vertices, very necessary at high luminosity, and reverts to a slower algorithm when no high quality tracks are found. A special fit procedure is available to better handle electron bremsstrahlung. This is invoked from the subsequent combined reconstruction for tracks associated to an EM-calorimeter cluster.

13.3.4.3 LArClusterRec

LArClusterRec is the reconstruction package for electromagnetic clusters in the calorimeter; it is organized in two steps as described below.

In the first step, towers are created by summing the cells of the electromagnetic calorimeter and the pre-sampler in depth using a granularity of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ corresponding to the granularity in the second sampling of the EM calorimeter. The inputs to the tower building are the calibrated calorimeter cells which are produced by the package LArCellRec.

In the next step, a sliding-window algorithm is used. If a local maximum is found with total E_T in the window above a given threshold, a cluster is created which is subsequently stored in the cluster container. To reconstruct the cluster energy and position is calculated in a given window.[1] The cluster energy is corrected for η and ϕ modulations of the calorimeter response and for leakage outside the cluster in a given window. In the region between the barrel and end-cap calorimeters, the cluster energy is in addition corrected for energy losses using the energy deposit in the crack scintillators. The η position in the first and second sampling is corrected for s-shapes, the ϕ position is corrected for an offset: both of them are geometry effects.

13.3.4.4 egammaRec

EgammaRec is an algorithm designed to calculate quantities that can be used to separate electrons and photons from jets. To do so, electromagnetic-cluster and tracking information are used.

In the electromagnetic calorimeter electrons are narrow objects, while jets tend to have a broader profile. Hence, shower shapes can be used to reject jets. This is handled by the EM Shower Builder which calls a number of algorithms which calculate diverse quantities using information from the first and second samplings of the electromagnetic calorimeter, as well the first sampling of the hadronic calorimeter.

Cluster and track information are combined in the Track Match Builder. For a given cluster all tracks are examined in a window around the cluster position. In case more than one track is found, the one with the highest p_T is retained. If the E/p ratio is $0.5 < E/p < 1.5$, the track match is successful. In the subsequent particle-identification step the information provided by egammaRec can be used. In the case of an electron hypothesis, jets can be rejected by analysis of the EM shower shape, tight track-quality cuts, E/p matching, and the position match in η and ϕ directions between the Cluster and the Track. Photons can be selected by analysing the EM shower shapes, using reconstruction of photon conversions in the Inner Detector, and, possibly, using a Track veto for non-converted photons.

13.3.4.5 Moore

Moore is an offline track reconstruction package for the Muon Spectrometer. A detailed description of Moore is available elsewhere [13-47]. Moore reconstructs tracks in the full η range (barrel+endcaps) of the muon spectrometer. However, we will restrict the description to the barrel since at present only this region is included in the trigger chain (Muon vertical slice).

1. This window can be different from the one used for the sliding window algorithm.

In the standard offline configuration Moore takes as input collections of digits or clusters inside the muon spectrometer and outputs fitted reconstructed tracks whose parameters are expressed at the first measured point inside the muon spectrometer. The reconstruction flow is desegregated into sequential steps, and each step is driven by an algorithm module that builds partial or final reconstruction objects. Each algorithm retrieves objects created by the previous modules from Storegate and builds transient objects that are subsequently recorded in Storegate from where they are available for other algorithms. Data and algorithms are strictly separated: the algorithms should know the structure of the data objects that they are accessing or producing, but the objects do not depend on the algorithms. The existence of only a dependence from algorithms to data and the flowing sequence of the reconstruction steps allows to establish which algorithm will produce an object at run-time. At present the overall reconstruction starts from the searches of regions of activity performed in two consecutive steps: one searches for activity regions in the ϕ trigger hits from RPC and another performs the search on the *r-z* view considering the precision hits of MDTs.

Inside the MDTs the drift distance is calculated from the drift time by applying various corrections to the simplest linear formula. These corrections concern the time of flight, the second coordinate and the propagation along the wire, and the Lorenz effect. The best Track segment is chosen form the possible four lines tangential to the drift circles. All the MDT segments of the outer station are combined with those of the Middle layer. The MDT hits in the segments are combined with the RPC ϕ information, forming outer Track candidates. All the successfully fitted candidates are kept for further processing.

Each outer Track that is successfully reconstructed is subsequently used to associate inner-station MDT hits. A final track is defined as a collection of RPC hits and MDT hits successfully fitted from at least two layers. The parameters of the fitted Tracks are specified at the first measured point inside the Muon Spectrometer. In order to be used for physics studies an extrapolation of the Track parameters to the production point is needed. To accomplish this task a different offline package, Muon Identification (MuId), is used. In MuID the multiple scattering is parameterized as scattering planes in the calorimeters and the energy loss is evaluated from the calorimeter measurements or from a parameterization as a function of η and the muon momentum. After a refit at the vertex the Track parameters are used for further analysis.

When dealing with data already selected by the trigger the first two steps can be substituted by *ad hoc* makers that seed the track search in the regions selected by the trigger, by using the HLT Region Selector. We refer to this way of executing as “seeded mode”. More details regarding the implementation of the Moore package in HLT are given in Ref. [13-48].

13.4 Signatures, rates and efficiencies

In the following subsections, we present a preliminary assessment of the physics performance of the trigger using algorithms for LVL2 and EF applied to representative classes of final-states: electrons and photons; muons; jets, taus and missing transverse energy; b-jets; and B-physics. This broad classification stems from the physics goals of the ATLAS experiment, as explained in Chapter 4 and in Ref. [13-49].

Whenever possible, realistic raw data formats have been used as a starting point for the HLT processing, using a very detailed simulation of the raw data as they will appear when accessed from the Read Out System. This is deemed to be extremely important since, as described in the following, the steps that prepare the data for the algorithmic processing represent a sizeable

fraction of the total computing time, even when using optimized data-preparation software. We would like to underline that this has been made possible thanks to the close collaboration between the TDAQ and the Detector groups.

Steering control (as described in Section 9.5) has also been employed, highlighting the flexible boundary between LVL2 and EF. Selection schemes are then derived, which contain the signatures used to decide whether or not to reject events.

In order to maximize the discovery potential, the selection schemes generally only use inclusive signatures. With the exception of B-physics, exclusive reconstruction of particle decays is not required and topological variables[1] are not used in the selection, although this is technically feasible both at LVL2 and in the EF (e.g. to select $Z \rightarrow l^+l^-$ decays exclusively).

It must be emphasised that system performance (e.g. minimizing execution time and volume of data needed) is a major requirement in implementing the HLT selection, to comply with the constraints imposed by the on-line environment and available resources. In this chapter also some indication of the compliance with the requirements will be given for representative selections. Since, obviously, system and physics performance are correlated, all results have been obtained after trying to optimise both simultaneously.

13.4.1 e/gamma

In the present view of the ATLAS trigger menus, the inclusive electron and photon triggers are expected to contribute an important fraction of the total high-p_T trigger rate. After the selection in LVL2 and the EF, the remaining rate will contain a significant contribution due to events from Standard Model physics processes containing real isolated electrons or photons (b, c $\rightarrow$eX, W $\rightarrow$eν, Z $\rightarrow$ee, direct-photon production, etc.). According to the general guidelines indicated in Chapter 4, the identification of these trigger objects is made using only inclusive selection criteria.

The electron and photon triggers can be viewed as a series of selection steps of increasing complexity. At LVL1, electrons and photons are selected using calorimeter information on a reduced granularity. After receiving the LVL1 electromagnetic (e.m.) trigger RoI positions, the LVL2 trigger performs a selection of isolated e.m. clusters using the full calorimeter granularity within the RoI of size $\Delta\eta \times \Delta\phi$=0.2 $\times$0.2. To calculate the cluster E_T, the most up-to-date available calibration data are used (see Section 13.3.3.5). Electrons and photons are selected based on the cluster E_T and shower-shape quantities that distinguish isolated e.m. objects from jets. A further, more refined calorimeter-based selection may classify the e.m. cluster as a candidate photon trigger object. In the next step, electrons are identified by associating the e.m. cluster with a track in the Inner Detector. In general, track candidates are found by independent searches in the TRT and SCT/Pixel ('Precision') detectors in the region identified by the LVL1 RoI. Details of the different LVL2 tracking algorithms that are currently being studied are described in Section 13.3.3.1, Section 13.3.3.2 and Section 13.3.3.4. For the results presented in the next section, however, only IDscan (Section 13.3.3.1), has been used; work is in progress to evaluate the other algorithms. For electron candidates, matching in position and momentum between the track and the cluster is required. Electron candidates passing the LVL2 selection criteria are retained to be examined by the EF.

1. E.g., variables, such as invariant masses, constructed by combining of several high-p_T objects.

In the EF, electrons are selected with a very similar strategy to LVL2 using information from the calorimeters and the Inner Detector. For track reconstruction, results obtained with xKalman++ are presented below; work is in progress to evaluate the alternative track-reconstruction program iPatRec. First results show very similar behaviour between these two algorithms. The main differences with respect to LVL2 arise from the availability of calibration data and the possibility to use more sophisticated reconstruction algorithms with access to the detector data for the full event. This results in sharper thresholds and better background rejection. In order to avoid biases introduced by using different reconstruction algorithms for online and offline selection, the EF will select events using as far as possible the offline reconstruction algorithms. However, using the offline reconstruction in the EF implies that the offline algorithms must comply with the stricter EF requirements in terms of robustness and system performance. This has not yet been achieved, but work is in progress to change the algorithms accordingly.

The present study uses the currently-available ATLAS offline reconstruction software (see Section 13.3.4) as a prototype of the future EF code. The criteria used to identify electrons and photons need to be softer in the EF than in the offline in order not to lose events prematurely. In previous EF studies [13-50], [13-51], the electron and photon selections used the same cuts as in the offline selection, but did not use some “critical” criteria. For example a track veto for non-converted photons is not applied at the EF level because it requires good control of the fake tracks in the Inner Detector and thus a very good understanding of the tracking performance, especially in the presence of pile-up. A more realistic EF electron selection has been used for the studies presented here.

In the following, the physics performance of the selection of electrons by the HLT is reviewed. Only the main performance issues are discussed — a more complete report on this work can be found in Ref. [13-52]. The results were obtained using the new framework as described in Section 9.5. Work is in progress to assess more fully the physics performance, including the photon selection which will be reported on at a later stage. The system-performance aspects of the electron selection are discussed briefly in Section 13.4.1.2. Note that the electron and photon selections presented here build on earlier studies [13-50],[13-51],[13-53].

13.4.1.1 HLT Electron Selection Performance

The performance of the electron and photon triggers has been estimated for single electrons and photons, and for some standard physics channels (e.g. $Z \rightarrow ee$, $W \rightarrow e\nu$, $H \rightarrow 4e$). The physics performance is characterized in terms of efficiency for the signal channel, and the rate expected for the selection. Here only the single-electron trigger is considered; for the two-object triggers sufficiently high-statistics datasets could not be reconstructed and analysed in time for the TDR. The efficiencies were obtained using fully-simulated events with single electrons, and the rates were estimated with $\sim 2 \times 10^6$ fully-simulated di-jet events; pile-up was included corresponding to low and design luminosity (see Ref. [13-2]). A higher-statistics sample of 10^7 di-jets is available, but this still needs to be analysed.

The events were generated using PYTHIA 6.203 [13-54][13-55]. It should be noted that compared to older studies the overall cross section for QCD di-jets with E_T>17GeV has increased by 50% [13-56]. However, a comparison made using fast simulation [13-56] has shown that, at the same time, the rate of isolated electrons has gone down by 50% due to a decrease in the rate of electrons from b and c-decays.

Compared to previous studies a more up-to-date detector geometry has been used. An important change is the amount of material in the Inner Detector that has increased significantly, in

particular due to the support of the first b-layer. Therefore, bremsstrahlung effects for electrons have become more important, especially in the end-caps. In addition, a realistic magnetic field map is now used in the simulation instead of a constant solenoidal field of 2T. This change affects the tracking performance in the end-caps. More details on the datasets that have been produced can be found in Ref. [13-2].

In general, events with electrons and photons are selected on the basis of single high-p_T objects or of pairs of somewhat lower-p_T objects. The physics performance of the single-electron triggers is summarized below and documented in more detail in Ref. [13-10] and Ref. [13-52].

The currently-achieved performance for the single isolated-electron HLT algorithms is summarized in Table 13-1 as a function of the main steps in the LVL2–EF trigger chain for the 2×10^{33} cm^{-2} s^{-1} luminosity scenario. The trigger steps have been factorized by detector in order to show the rejection that each stage contributes to the trigger. It should be noted that there are strong correlations between the different selection criteria.

Table 13-1 Performance of the single-electron HLT at 2×10^{33} cm^{-2} s^{-1}, obtained in a preliminary study with the new HLT software chain. The results are presented in a single sequence. 'Matching' refers to position and energy–momentum matching between calorimeter clusters and reconstructed tracks (at LVL2 only precision-detector tracks are used in the present study). The efficiencies are given for single electrons of p_T = 25 GeV already selected by LVL1, averaged over the full pseudorapidity range $|\eta| < 2.5$. For reference, the efficiency of the LVL1 selection is 95%. The rates are normalized to a LVL1 rate for e.m. clusters of ~12 kHz. Note that the quoted errors are statistical — as discussed in the text, there are very large systematic uncertainties on the rate estimates, e.g. from cross-section uncertainties.

Trigger Step	Rate [Hz]	Efficiency [%]
LVL2 Calo	2114 ± 48	95.9 ± 0.3
LVL2 Tracking	529 ± 24	88.0 ± 0.5
LVL2 Matching	137 ± 12	86.6 ± 0.6
EF Global	30 ± 5	76.2 ± 0.7

The overall reduction in rate achieved by LVL2 is ~90 for a loss of efficiency of 14% with respect to LVL1. The additional rate reduction provided by the EF is a factor of about five for a further efficiency loss of ~10%. This shows the power of the two-tier HLT selection.

Comparing the new results with previous studies [13-1], the rate reduction at LVL2 was found to be almost identical for the same electron p_T. This is a useful cross-check and a significant achievement given the completely different selection-software architectures. It also gives us confidence that the new software that became available only shortly before the submission of this document, is working correctly. Moreover, comparing this to previous studies [13-3], we find the rate reduction at LVL2 is almost identical: at the EF, the rate reduction is currently 4.6, compared to 5.2 as reported in [13-3], These numbers still need to be understood better, and further optimization and cross-checks are needed to ensure that the selection scheme does not reject too many 'real' electrons prematurely.

As a cross-check of the above results, the electron efficiencies have also been extracted using fully-simulated $W \rightarrow e\nu$ events. With respect to LVL1, efficiencies of $(90.6 \pm 0.8)\%$ after LVL2 and $(83.2 \pm 1.0)\%$ after EF have been obtained. The simulation included pile-up for a luminosity of 2×10^{33} cm^{-2} s^{-1} and used W events in which the electron had p_T>25 GeV at the generator level. These efficiencies are slightly higher than the ones given in Table 13-1 for single electrons of p_T=25 GeV. This is expected since the efficiency is better for electrons with p_T well above the

threshold value than for those at the threshold value. Preliminary studies of the full offline electron / jet separation analysis, that uses even more powerful selections than the ones we can exploit at the EF, also confirm the rates and efficiencies found here.

The analysis of the performance of the single-electron trigger e30i at high luminosity is in progress and the results are still very preliminary. First results using only the EF selection give a rate of 165 Hz for an efficiency of ~72% with respect to LVL1. Further work is needed to optimize the signal efficiency and rate reduction, including also the LVL2 trigger.

Based on our previous experience, a reasonable target at the EF is to accept electrons with an overall efficiency of about 80% w.r.t. LVL1 in order not to cut too hard on physics. For the rate obtained in the previous studies (40 Hz), the electron efficiency with the new software would be ~78% w.r.t LVL1, as can be seen in the table extrapolating the results given after the EF calorimeter selection. From this we conclude that, with an improved tuning of the different selection steps, a final rate of ~40 Hz at 2×10^{33} cm^{-2} s^{-1} luminosity for an 80% electron efficiency, as obtained in previous studies [13-57], is still valid using the new selection algorithms and updated detector geometry, and allowing for the changes at generator level. The results are preliminary and work is in progress to optimise and cross-check the current selection cuts. For example, the cuts will be tuned in a more optimised way as a function of pseudorapidity, and additional criteria such as isolation in the Inner Detector around the electron track are being studied.

13.4.1.2 HLT Strategy, Algorithm Optimisation and the LVL2–EF Boundary

A first assessment of the system performance has been made for LVL2 in view of its highly demanding environment in terms of execution time and latencies for data transfers. The algorithm execution times per RoI measured on a 2 GHz PC are ~2 ms for the calorimeter reconstruction using T2Calo and ~ 3 ms for the track reconstruction using IDscan — note that the latter is only called for RoIs retained after the calorimeter-cluster selection. The data-preparation code for the calorimeters is currently being optimized to minimize the computing time per RoI. Timings of less than 10 ms per RoI have been achieved on a 2 GHz PC, and efforts will continue to reduce this further.

It is worth noting that the evaluation mentioned above of the data-preparation time was only possible thanks to the realistic simulation of the raw data that the HLT will receive from the Readout System. It appears that the data-preparation time could be a very significant part of the total LVL2 processing time, even though it is restricted to the RoIs.

For the EF, we do not yet have results at the same level of detail as for LVL2, although it is planned to make measurements for both data preparation and algorithmic processing. As an example of the preliminary results obtained so far, data preparation in the EF for the whole LAr calorimeter takes about 0.5 s on a 2 GHz PC. Results from the studies that are under way will be documented in future ATLAS notes.

The use of system resources in the electron HLT can be minimized by exploiting the modularity of the trigger. By ordering the trigger steps so that events are rejected as early as possible, both the overall processing times and data transfer rates may be reduced. Factorizing the trigger-algorithm components also provides flexibility to move the rejection power from LVL2 to the EF or vice versa, to optimize efficiency of the physics selection, rejection of high-rate backgrounds and use of system resources. These issues have been extensively studied in the past and are reported in Ref. [13-53].

If the rate for a trigger item from any level in the selection is too high, one can reduce it either by raising the E_T threshold of the item or by tightening the selection criteria. However, this results in a loss in efficiency for physics signals. The loss in physics may partly be recovered by adding more exclusive trigger selections for the channels of interest, but the contribution to the overall rate of these extra items must be taken into account. There are long-term ongoing studies, performed together with the ATLAS physics community, to assess the impact of such changes in order to prepare alternative scenarios.

As an example, Figure 13-6 illustrates the impact on the selection of $W \rightarrow e\nu$ events of raising the threshold in the single-electron HLT selection at 2×10^{33} cm^{-2} s^{-1} (nominal threshold of 25 GeV). The impact on other physics signal such as $Z \rightarrow ee$, and $H \rightarrow 4e$ is discussed in Ref. [13-52].

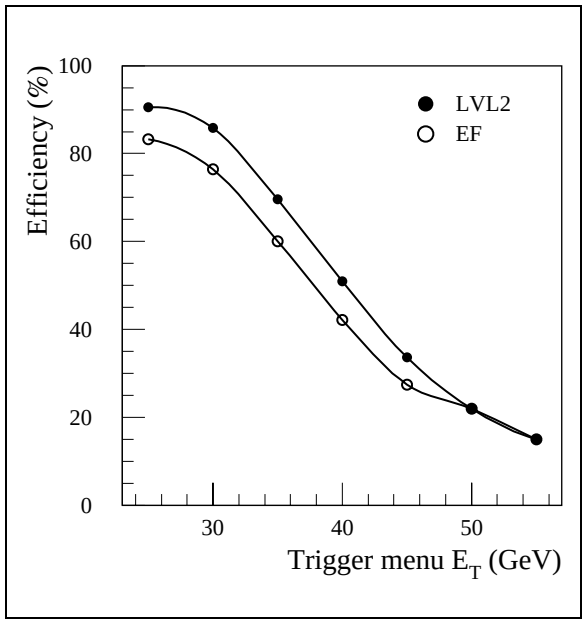

Figure 13-6 Efficiency to select $W \rightarrow e\nu$ events at LVL2 and in the EF as a function of the E_T threshold in the trigger menu at 2×10^{33} cm^{-2} s^{-1} luminosity. Only events that passed the LVL1 selection and for which the electron has p_T>25 GeV at the generator level have been considered.

As illustrated above, the HLT strategy contains considerable flexibility. Various possibilities exist to adjust the balance between minimizing the required computing resources and maximizing the physics performance. For many channels of interest, the selection scheme also provides considerable redundancy. More details on the trigger selection strategy can be found in [13-1] and in Chapter 4.

13.4.2 Muon selection

The main purpose of the high-level muon trigger is the accurate reconstruction of muon tracks within RoIs indicated by the LVL1 muon trigger. LVL2 and the EF must reject low-p_T muons (i.e. muons with p_T below the threshold that are initially selected due to the limited resolution in the first trigger level), secondary muons produced in decays in flight of charged pions and kaons, and fake muon tracks composed of hits from the cavern background. The EF must be able to reconstruct additional muons present in the event that were not reconstructed or selected by the LVL1 and LVL2 triggers.

Whilst the LVL1 trigger system uses only hits from the dedicated trigger detectors (RPCs in the barrel and TGCs in the endcap), LVL2 and the EF have access to the full measurements of the Muon Spectrometer, in particular the data from the Monitored Drift Tubes (MDTs). This allows very good track reconstruction in the muon spectrometer. The high-background environment in the muon spectrometer demands algorithms with robust and fast pattern recognition capable of rejecting hits induced by the cavern background.

The tracks found in the LVL2 muon trigger are extrapolated for combination with measurements in the Inner Detector and the calorimeters. Matching between muon tracks measured independently in the muon system and the Inner Detector selects prompt muons and rejects many fake and secondary muons. This is especially important for the B-physics trigger in low-

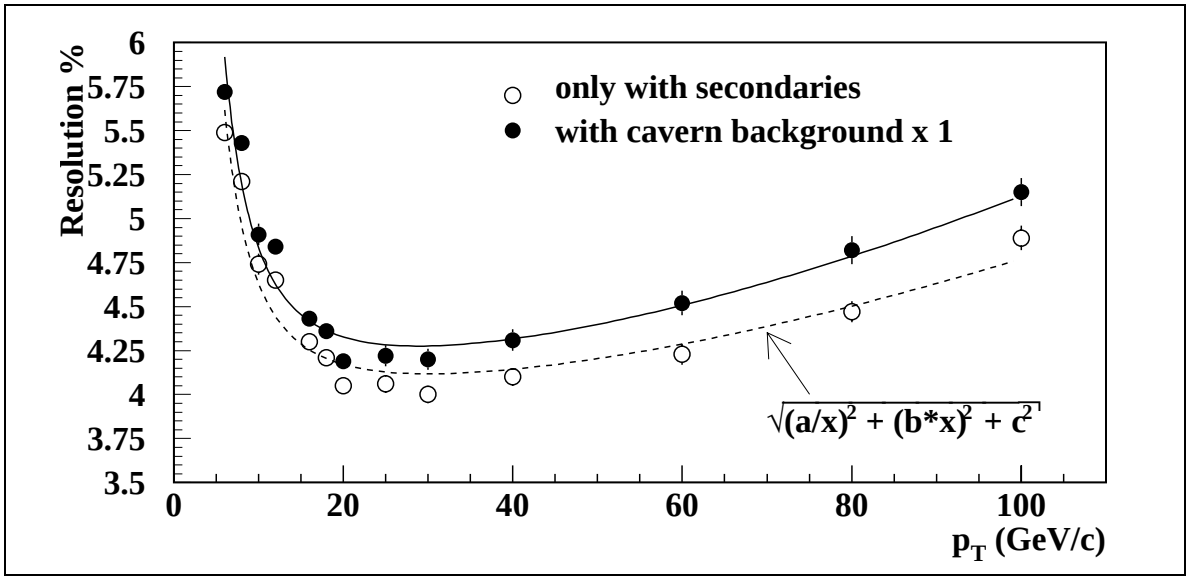

Figure 13-7 The p_T resolution of the muFast algorithm as a function of muon p_T.

luminosity running, for which the selection of relatively low-p_T prompt muons events is the first stage in B-physics trigger selections. The studies presented in this section are limited to the barrel region ($|\eta|<1$) — future work will extend them for the full pseudorapidity range.

13.4.2.1 The Physics Performance of LVL2 Muon algorithms

The physics performance of the LVL2 muon trigger algorithms was presented and discussed in detail in [13-1]. The algorithm muFast has been implemented in the new framework with no important changes in the code used to reconstruct the muon tracks. The main difference of the present version of the code, compared to the one studied previously, is the use of "LVL1 Emulation" to identify among the RPC hits the ones used by the LVL1 trigger to select muon candidates. We expect this to have very little impact on the overall muon-reconstruction efficiencies at LVL2, and no effect at all on the muon p_T resolution. The following summarizes the most relevant results obtained in the previous studies [13-13].

Table 13-2 Total output rates [kHz] of the stand-alone LVL2 muon trigger after application of the muFast algorithm for the 6 GeV low-p_T threshold at 1×10^{33} cm^{-2} s^{-1} and 20 GeV threshold at the design luminosity.

Physics Process	low-p_T	high-p_T
π/K decays	3.00	0.07
b decays	0.90	0.09
c decays	0.50	0.04
$W \rightarrow \mu\nu$	0.003	0.02
cavern background	negligible	negligible
Total	**4.40**	**0.22**

Given the steeply-falling muon-p_T spectrum, the rate of muons measured with transverse momenta above a given threshold depends strongly on the p_T resolution of the reconstructed muons that is therefore a critical parameter in the LVL2 trigger. The resolution of the muFast algorithm is shown as a function of p_T in Figure 13-7. The resolution ranges between 4.0% and 5.5% for muons in the p_T interval 6–20 GeV. These results are comparable with the p_T resolution obtained by the offline muon reconstruction program MUONBOX [13-58].

The total rates after this algorithm for 6 GeV and 20 GeV thresholds have been evaluated by convolving the efficiency as a function of p_T with the muon differential cross-sections of the dominant muon-production processes. Account was taken of the rejection provided by the LVL1 selection. The rates after LVL2 are shown in Table 13-2[1].

The rates from π/K decays were calculated using the predicted cross-sections from the DMPJET program; the rates would be lower by about 50% if the PYTHIA prediction were used instead. Preliminary studies of the trigger rate arising from the cavern background have been made using predictions of the cavern background from the FLUKA package. The probability that a fake LVL1 muon trigger is accepted by the LVL2 is below 10^{-2}. With this upper limit it is safe to neglect the contribution from fake muons.

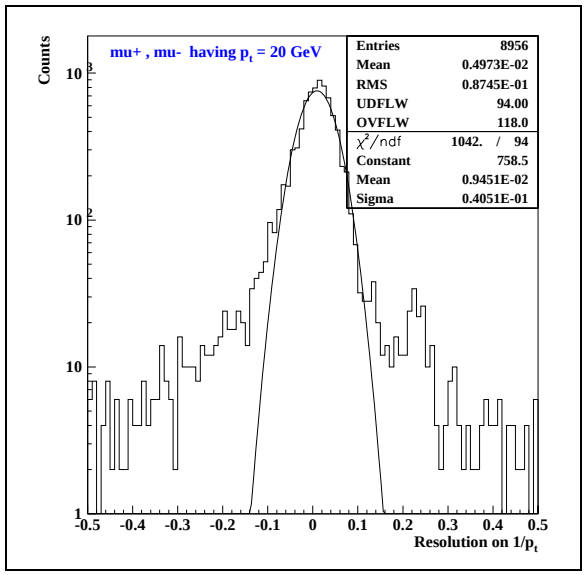

Figure 13-8 Transverse-momentum resolution $(1/p_T\text{-}1/p_T^{true})/(1/p_T^{true})$ for 20 GeV muons reconstructed with the muFast in the new selection framework (no cavern background). Tails arise from events with large-angle Coulomb scattering or from poor momentum reconstruction where delta-rays were emitted by the muon.

The implementation of muFast in the new software framework has been used to reconstruct events with single muons of p_T = 20 GeV. Figure 13-8 shows the distribution of $(1/p_T\text{-}1/p_T^{true})/(1/p_T^{true})$. The resolution is 4.0% which is identical to the result obtained with muFast in the old implementation. This result supports the expectation that the physics performance does not change with respect to the previous results [13-13].The implementation of the muComb algorithm in the new framework is still ongoing. Also in this case the algorithm under implementation is the same one discussed in [13-1]. We can therefore rely on the results already obtained. Figure 13-9 shows as a function of the muon p_T the efficiency for the combined (muon spectrometer plus inner detector) reconstruction of prompt muons and of secondary muons from π/K decays in flight.

The requirement of good matching (i.e. z/ϕ and p_T matching) between the muon-spectrometer and inner-detector tracks reduces contribution from π/K decays in flight to the low-p_T trigger rate to 1.0 kHz: a factor three reduction compared to the rate after the muFast algorithm. Taking into account the reduction in rate from the improved p_T resolution for prompt muons, the total rate after the muComb algorithm is 2.1 kHz for muons with p_T>6 GeV and $|\eta|$<1. Note that these numbers are calculated for a luminosity of 1.0×10^{33} cm^{-2}s^{-1}.

13.4.2.2 The Physics Performance of the Muon Event Filter Algorithm

The physics performance of the MOORE EF package has been evaluated with simulated single-muon samples, with no cavern background, in the p_T range 3–1000 GeV. Here we have considered a *fully-wrapped* version of the offline code that is completely equivalent, from the recon-

1. $W\rightarrow\mu\nu$ cross section found with PYTHIA for p_T^{μ}>3 GeV and $|\eta_{\mu}|$<2.7 : 9.56 pb.This corresponds to 100 Hz at 1×10^{34} cm^{-2} s^{-1}

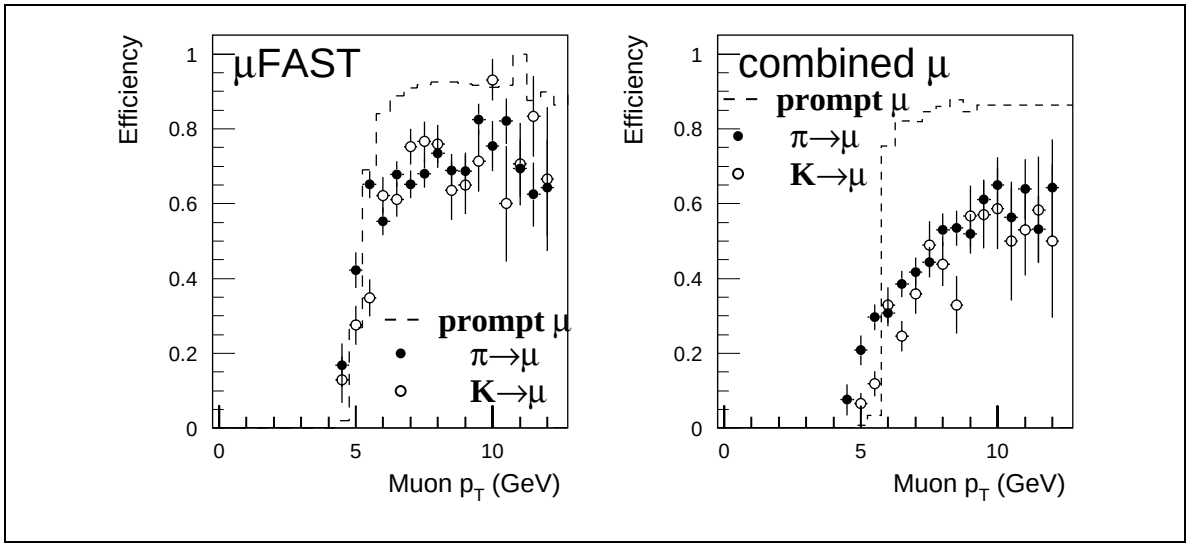

Figure 13-9 The efficiency with respect to LVL1 of the combined (muon spectrometer plus inner detector) reconstruction at LVL2 for prompt muons and for muons from π/K decays in flight. The left-hand plot shows the efficiency of the stand-alone algorithm muFast and the right-hand plot shows the efficiency of the combined muon algorithm muComb.

struction point of view, to the offline version. This version does not support seeding of the track search, e.g. by the results of the LVL2 trigger (see below).

Figure 13-10 shows the track-reconstruction efficiency as a function of muon p_T in the barrel region. The triangle-up symbols show the efficiency of the reconstruction in the muon spectrometer ("Moore"), while the triangle-down symbols show the efficiency of the reconstruction after the track has been extrapolated to the interaction region ("Muid"). The energy loss in the calorimeter has been parameterized as a function of η and muon momentum. The loss of efficiency of "Muid" at low p_T is due to the failure of the extrapolation of muons that exhibit in the spectrometer a transverse momentum of a few GeV. Muons with p_T larger than 8 GeV are reconstructed with efficiencies above 95%, which is equivalent to the results from MUONBOX shown in [13-49], [13-58].

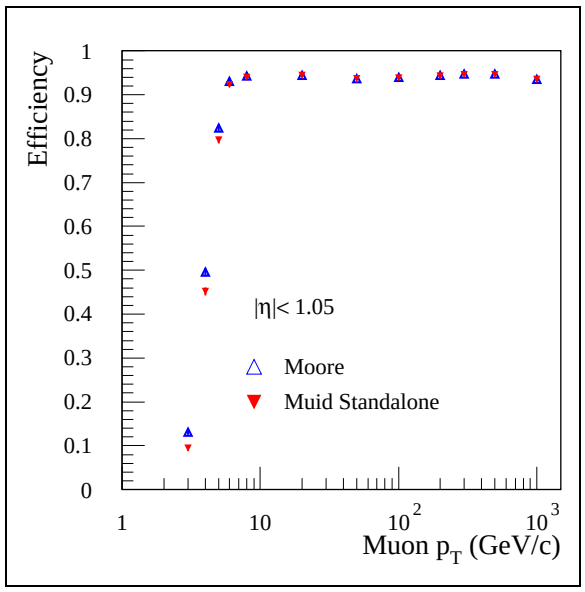

Figure 13-10 Single-muon reconstruction efficiency as a function of muon p_T in the barrel region; triangle-up: stand-alone reconstruction (Moore); triangle-down: reconstruction at the nominal interaction vertex point.

Figure 13-11 shows the muon p_T resolution of the Moore package in the stand-alone muon spectrometer and after extrapolation to the interaction vertex. Again, the results are consistent with earlier studies. The physics performance has been checked for a *seeded* [13-48] version of the algorithm. We don't expect significant differences with respect the full version, since the reconstruction and fitting methods are the same as the ones used by the offline package. This is supported by Figure 13-12 where the transverse-momentum resolution of the seeded version is

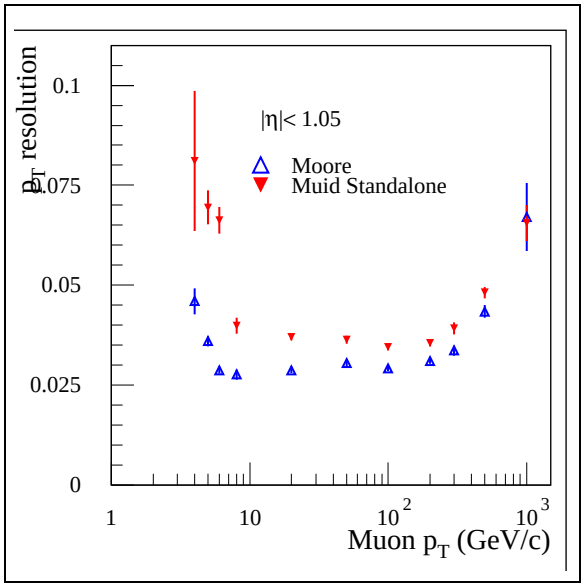

Figure 13-11 *Fully-wrapped* version: transverse-momentum resolution as a function of the generated muon p_T.

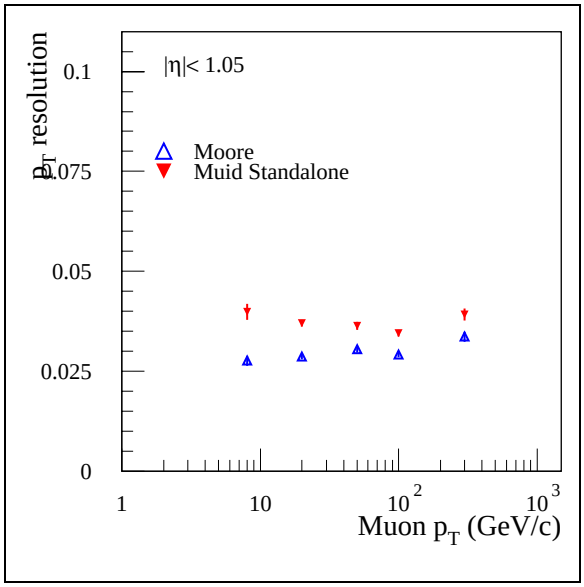

Figure 13-12 *Seeded* version: transverse momentum resolution as a function of the generated p_T.

shown for p_T above 8 GeV. A more complete presentation of the Moore physics performance is given in Refs. [13-47] and [13-48].

13.4.2.3 The Timing Performance of the Muon Algorithms

The muFast trigger algorithm has been benchmarked on a 2 GHz machine. The event sample consisted of about 430 events with single muons in the barrel region of p_T=100 GeV, with and without simulation of the cavern background. As shown in Figure 13-13 the average processing time of muFast is about 1.0 ms per RoI with an r.m.s. of 0.2 ms, and with no event taking more than 4.1 ms.

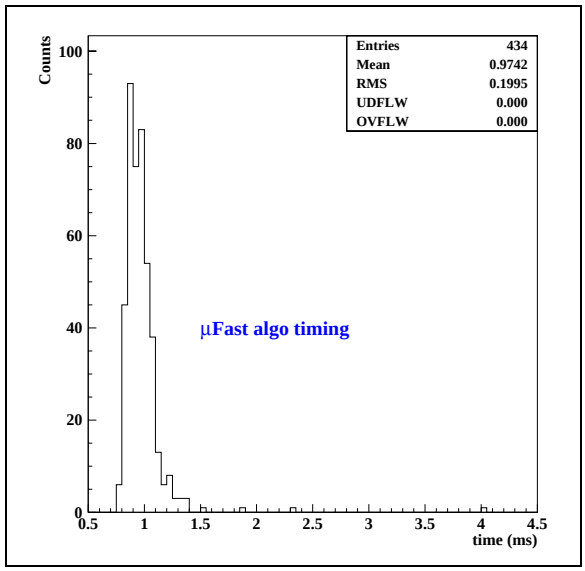

Figure 13-13 Processing time of muFast.

Figure 13-14 shows the strong correlation of the processing time with the amount of RPC activity. The composition of this overall timing is shown in Figure 13-15 — as can be seen, the processing time is dominated by the “Level-1 Emulation” that contributes about 0.8 ms.

A very preliminary simulation of the cavern background has been made to check the robustness and the latency of the selection algorithms. This background has been simulated assuming the nominal luminosity 1×10^{34} cm^{-2} s^{-1} and boosting by a factor two the predictions provided by the GCALOR package [13-59]. In presence of the cavern background the muFast timing increases by only about 100 μs. The data-preparation time has been evaluated on the same sample of events. Without cavern background the total time is 3.4 ms; this increases to 8.9 ms if the cavern background is added. The time has been found to depend linearly on the

muon-hit multiplicity. This allows one to extrapolate the data-preparation time to more severe background conditions. Considering a safety factor of five, the data-preparation time is expected to increase to 22.5 ms, while the muFast algorithm processing time remains below 1.5 ms. Therefore, on a 2 GHz processor, the overall CPU time is expected to be below 25 ms.

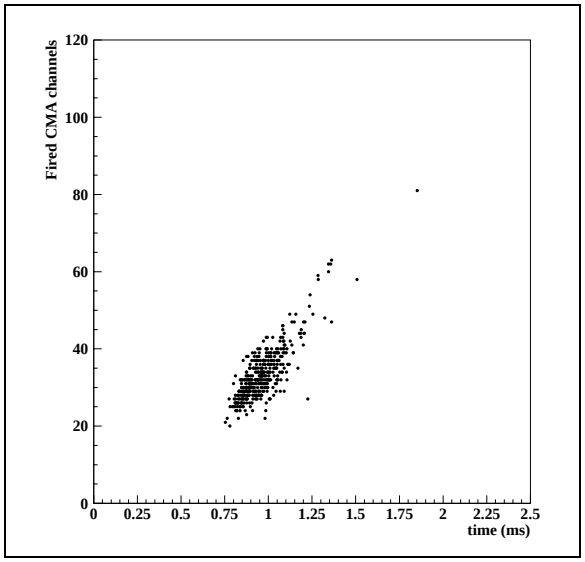

Figure 13-14 Correlation between the muFast processing time and the number of the fired channels in the RPC readout.

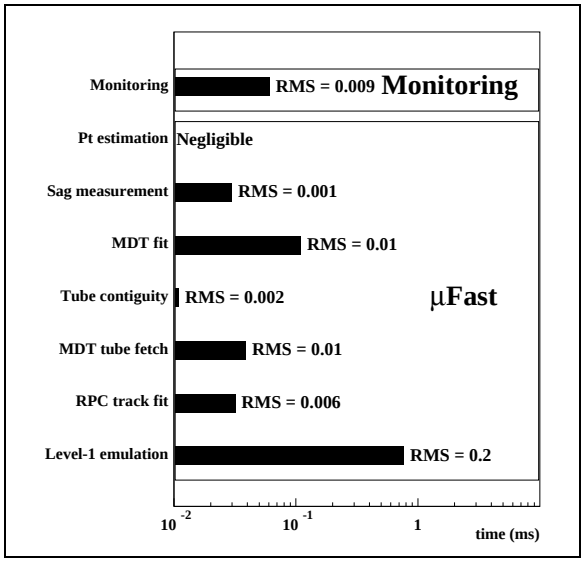

Figure 13-15 Contributions to the muFast processing time. The bar length shows the average processing time; the RMS is given in ms.

The timing performance of the Moore algorithm, for both seeded and fully-wrapped modes, has also been evaluated on a 2 GHz machine. The code was built in optimized mode. The event samples consisted of around 500 events each, fully simulated. Table 13-3 shows the average execution times per event both for seeded and wrapped modes. The timing includes the procedure of extrapolation to the vertex of the reconstructed track. In these results a conservative approach has been adopted where the timings include also the accesses to the data, the data preparation, and the region-selector accesses.

Table 13-3 Summary of Moore timing tests for single-muon events, without and with cavern background at 1×10^{34} cm^{-2} s^{-1}: (a) nominal background intensity, (b) twice the nominal background. intensity

Sample	Time (seeded mode) ms		Time (wrapped mode barrel) ms	
	< >	RMS	< >	RMS
8 GeV	86	69	80	63
20 GeV	65	36	66	43
50 GeV	68	68	68	51
100 GeV	71	60	76	69
300	75	60	81	88
100 GeV (a)	775	70	2800	600
100 GeV (b)	1500	600	3600	2800

The general behaviour of the average execution time for the different samples is represented in Figure 13-16 for both seeded and wrapped mode. In order to show the impact of the extrapolation to the vertex we have plotted the execution times for the track reconstruction inside the MuonSpectrometer (Moore), the execution times for extrapolating the track to the vertex (MuID), and the sum of the two (total). The execution times are rather flat over the analysed p_T range.The time for the whole reconstruction of single muon, including data access and data preparation, is on average below 100 ms. For those events we do not expect a large difference running in wrapped and seeded mode, since the data set that is accessed is quite similar.

The points in Figure 13-16 were obtained by averaging only on events for each sample that register execution times below one second. In order to show the impact of events that register longer execution times, we define a time efficiency as the ratio between the number of reconstructed tracks in one second and the number of Regions of Interest. The plot of time efficiency is shown in Figure 13-17 for both seeded and wrapped mode.

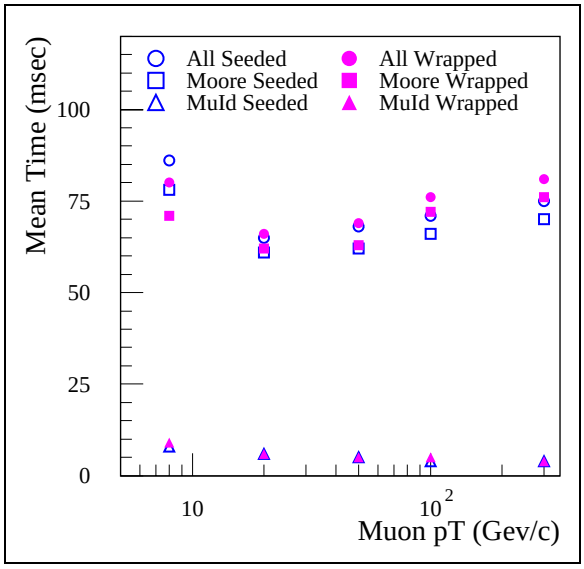

Figure 13-16 Average execution time (in ms) for different p_T values (in GeV), obtained with Moore, with MuId stand-alone and with both of them, in seeded and wrapped mode.

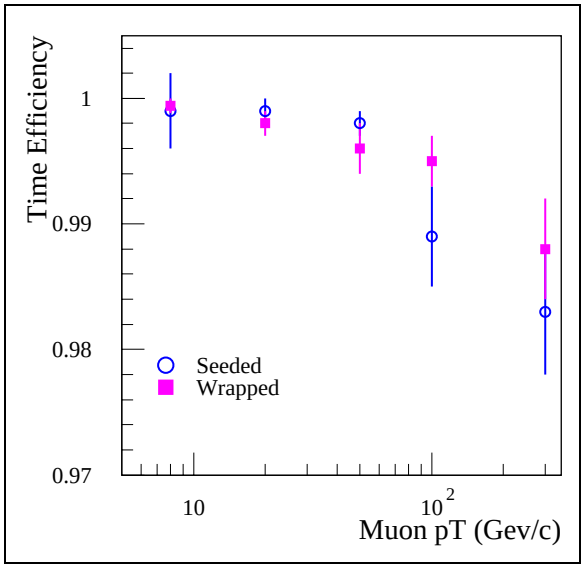

Figure 13-17 Time efficiency for seeded and wrapped mode for different p_T values.

13.4.3 Tau/Jets/$E_T$miss selection

13.4.3.1 The Tau Trigger

A major Standard Model source of tau leptons in ATLAS will be $W \rightarrow \tau\nu$ and $Z \rightarrow \tau\tau$. The tau lepton will also play a key role in the search for new physics. In the MSSM the heavy-scalar (H) and pseudo-scalar (A) Higgs-boson decays to tau-pairs are strongly enhanced with respect to Standard Model Higgs boson case. Also, a key decay channel for the charged Higgs boson is $H^{\pm} \rightarrow \tau\nu$.

The identification of the hadronic decays of tau leptons is based on the selection of narrow isolated jets with low multiplicity in the tracking system. The shower isolation and shape are calculated for both the e.m. and hadronic calorimeters separately. The fraction of energy deposited

by the tau-jet in the e.m. calorimeter has a mean value around 60%. The hadronic shower is broader in the hadronic calorimeter than in the e.m. calorimeter. Thus the jet information obtained from the e.m. calorimeter is more selective than that from the hadronic calorimeter. A detailed description of the tau trigger studies presented below is given elsewhere [13-60].

13.4.3.1.1 The LVL1 Tau Trigger

The motivation for a LVL1 tau calorimeter trigger is manifold, both in conjunction with electron, muon or missing-E_T signatures to enhance Z, A or W coverage, and for calibration purposes. Narrow tau jets containing one (three) charged particles give rise to narrow isolated energy depositions in the calorimeters. It is envisaged that an isolation requirement will be a valuable part of the tau trigger at all levels.

The e/γ and τ/h LVL1 algorithms are described in detail elsewhere [13-1],[13-3]. The LVL1 τ/h calorimeter trigger is based on a 4×4 array of "trigger towers" in the electromagnetic and hadronic calorimeters (within the region $|\eta| < 2.5$) where the tower granularity is $(0.1 \times 0.1)\ \Delta\eta \times \Delta\phi$. A core E_T is defined in the trigger algorithm as the sum of the electromagnetic and hadronic E_T in a 2×2 trigger-tower region. The trigger algorithm is based on four elements — the trigger cluster(s), an e.m. isolation region, a hadronic isolation region and an "RoI cluster".

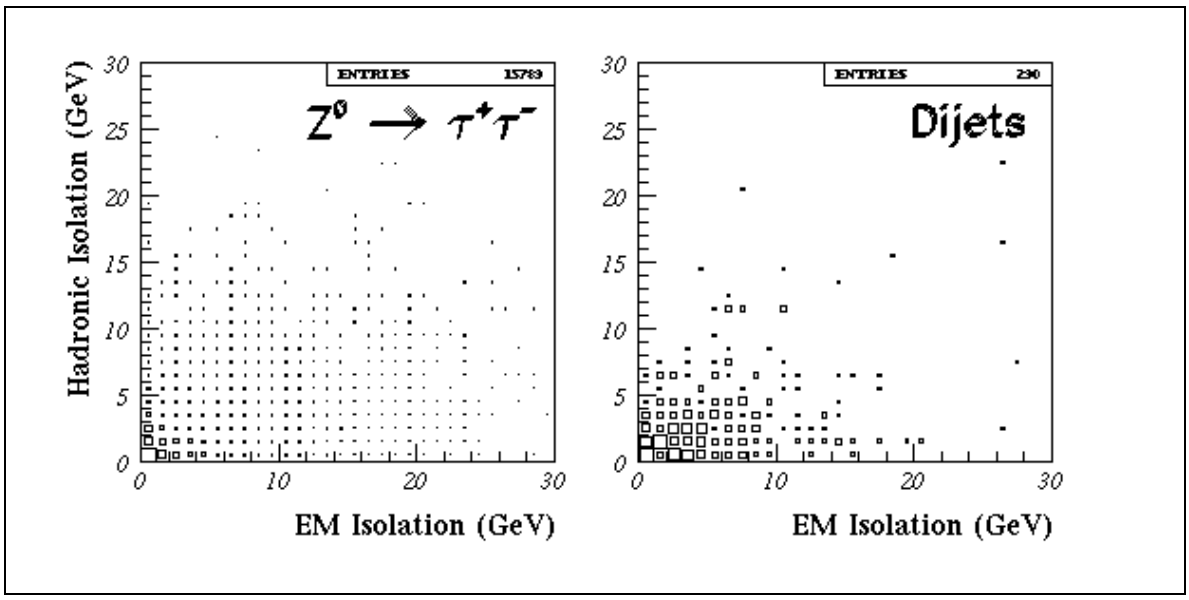

Figure 13-18 The hadronic isolation E_T vs. the e.m. isolation E_T (rings of 12 towers) for tau and QCD di-jets using the LVL1 trigger algorithm.

The distributions of the two isolation variables of the LVL1 τ/h trigger algorithm (in which the cluster E_T is measured in a region of 2×1 e.m. + 2×2 hadronic towers) are shown in the scatter plots of Figure 13-18 for simulated $Z \rightarrow \tau^+\tau^-$ events and QCD di-jet events. Looking at the corresponding projections, it is found that the EM isolation is more powerful that the hadronic one. In both plots a LVL1 algorithm has been employed with a core-E_T threshold of 20 GeV. The correlation between the isolation E_T values indicates that the hadronic isolation is of limited use in rejecting QCD jets. For example, the stand-alone tau trigger rate from QCD di-jets, at a luminosity of $1 \times 10^{33}\ cm^{-2}\ s^{-1}$ and neglecting the effect of pile-up, for a core-E_T of 20 GeV, a jet threshold of 20 GeV and an isolation cut at 10 GeV would be about 19 kHz.

13.4.3.1.2 The High-Level Tau Trigger

The tau trigger utilizes the e/gamma slice tools described in Section 13.4.1. The signal selection is tuned using events of the type $Z \rightarrow \tau^+\tau^-$. Background evaluation is performed using fully-simulated di-jet events. The LVL2 studies involve the verification of the LVL1 decision and, subsequently, tau identification using parameters that describe the shower shape in layers one and two of the e.m. calorimeters. The LVL2 variables used in a previous analysis [13-61], [13-62] could not be used in the present due to changes in the e/gamma software implementation. Additional rejection of background jets can be achieved by using the information from tracks associated to the tau RoI.

The LVL2 algorithm is applied to LVL1 tau RoIs. Loose LVL1 cuts are chosen for the study presented here: a cluster E_T in excess of 20 GeV is required, and the e.m. and hadronic isolation thresholds are set to 10 GeV. LVL2 jet calibration is applied to the cells within the LVL1 RoI window. The energy-weighted position ($\Delta\eta_\tau \times \Delta\phi_\tau$) of the tau-jet candidate is computed from all calorimeter cells within the LVL1 window. The first part of the LVL2 algorithm is the confirmation of the LVL1 decision. In order to do this the LVL1 algorithm described above is repeated using fine-grained cell information instead of the coarse-grain information available at LVL1.

As a first step, a check was made to ensure that LVL1 RoI coordinates are good approximations of the LVL2 tau coordinates, by measuring the distance between LVL1 RoI and the associated LVL2 cluster. After this successful check, for the LVL2 clusters associated with LVL1 RoIs, the next step is to look at the LVL2 e/gamma calorimetric variables that have some power to select taus over QCD jets.

Three variables were identified to discriminate between τ-jets and the di-jet background. The performance of the algorithm as a function of the shower-shape variable R_{37} was examined first. This variable is defined as the ratio of the E_T contained in a 3×7-cell cluster to the E_T in a 7×7-cell cluster centred on the same seed cell, calculated for the second layer of the e.m. calorimeter (Section 13.3.3.5). The second variable studied, $F_{12,}$ is defined in the first layer of the e.m. calorimeter and functions as an isolation variable: $F_{12} = (E_{max1} - E_{max2})/(E_{max1} + E_{max2})$ where E_{max1} and E_{max2} are the highest and second highest strip energies, respectively, in the first layer of the e.m. calorimeter: this quantity is defined for each LVL2 cluster. Is should be noted that the variables R_{37} and F_{12} are highly correlated. The third variable was, $F_{e.m.}$, the e.m fraction of the total energy.

Table 13-4 Trigger efficiencies for taus and for QCD jets, for different R_{37} cuts and the LVL1 trigger condition defined in the text. The right-hand column shows the corresponding stand-alone tau trigger rate based on QCD di-jets with a Pythia generation threshold E_T of 35 GeV.

Cut	Tau Efficiency % (wrt LVL1 accepts)	QCD jet efficiency% (wrt LVL1 accepts)	LVL1/2 tau trigger rate using calorimetry (kHz)
$R_{37} > 0.75$	88 ± 1	71 ± 1	11.1 ± 0.2
$R_{37} > 0.80$	83 ± 1	59 ± 1	9.4 ± 0.2
$R_{37} > 0.85$	75 ± 1	44 ± 1	7.0 ± 0.2
$R_{37} > 0.90$	62 ± 1	26 ± 1	2.3 ± 0.2

The variable with the most power to select taus and reject QCD jets was found to be R_{37}. Since the F_{12} variable is highly correlated with R_{37}, it was not found possible to obtain useful further

improvements in tau efficiency and jet rejection using F_{12}. Likewise, it was not possible to employ $F_{e.m.}$ to gain further QCD jet rejection whilst maintaining a good tau efficiency.

Three possible cuts in the R_{37} variable were studied: $R_{37} > 0.75, 0.80, 0.85$. These cuts were chosen to maintain a tau efficiency greater than 75%. Table 13-4 shows the efficiency for triggering on $Z \rightarrow \tau^+\tau^-$ and QCD di-jets, along with the corresponding stand-alone tau trigger rate at LVL2 estimated using QCD di-jets, for the three R_{37} cuts mentioned. The values are also given for reference for $R_{37} > 0.90$, although this gives an unacceptably low efficiency for selecting taus. The chosen set of LVL1 conditions, denoted 20-10-10, are the loose LVL1 cuts defined for the LVL1 study: in this case the core E_T is 20 GeV and the e.m. and hadronic isolation thresholds are both 10 GeV.

Additional rejection of background QCD jets can be achieved by using the information from tracks associated with the tau LVL1 RoI. The track information available at LVL2 was used to associate Inner Detector tracks, found using the "IDscan" algorithm, with the tau RoI by requiring that ΔR between the track and the LVL2 tau cluster direction, associated with the tau LVL1 RoI, obeyed the relation, $\Delta R < 0.3$. The inner-detector track-multiplicity distributions obtained for $Z \rightarrow \tau^+\tau^-$ and QCD di-jets are shown in Figure 13-19. The resulting LVL2 tau and QCD di-jet efficiencies are shown in Table 13-5 for a few useful combinations of calorimeter and track-based cuts.

Table 13-5 The tau and QCD-jet efficiencies for various useful combinations of R_{37} and LVL2 track-multiplicity cuts, for the specified LVL1 trigger condition. The right-hand column shows the corresponding stand-alone tau trigger rate based on QCD di-jets with a Pythia generation threshold E_T of 35 GeV.

Cuts	Tau efficiency % (wrt LVL1 accepts)	QCD di-jet efficiency % (wrt LVL1 accepts)	Stand-alone tau/hadron trigger rate (kHz)
LVL1: 20-10-10			15.6 ± 0.2
R_{37}>0.8, #LVL2 tracks<5	74 ± 1	34 ± 1	5.3 ± 0.2
R_{37}>0.8, #LVL2 tracks<4	70 ± 1	30 ± 1	4.7 ± 0.2
R_{37}>0.9, #LVL2 tracks<5	57 ± 1	18 ± 1	2.8 ± 0.2
R_{37}>0.9, #LVL2 tracks<4	54 ± 1	15 ± 1	2.3 ± 0.2

13.4.3.1.3 Tau Selection in the Event Filter

At the Event Filter stage, access to the complete, calibrated, event is possible for the first time. In addition, the tracking information has been further refined to reduce the number spurious track segments and minimize tracking inefficiency. Thus, it is possible to refine the LVL2 decision. Existing off-line studies of tau/hadron identification and jet rejection [13-63] provide the basis for the EF trigger decision. Typical trigger criteria for tau/hadron jets with $E_T > 30$GeV and $|\eta| < 2.5$ are as follows:

- The jet radius computed using only the e.m. cells contained in the jet, R_{em}, must obey the inequality: $R_{em} < 0.07$.
- The difference between the E_T contained in cones of sizes $\Delta R = 0.2$ and $\Delta R = 0.1$, normalized to the total jet E_T, ΔE_T must obey the inequality: $\Delta E_T < 0.1$ (isolation fraction).
- The number of reconstructed charge tracks pointing to the cluster (within a ΔR of 0.3), N_{tr}, is equal to one or three.

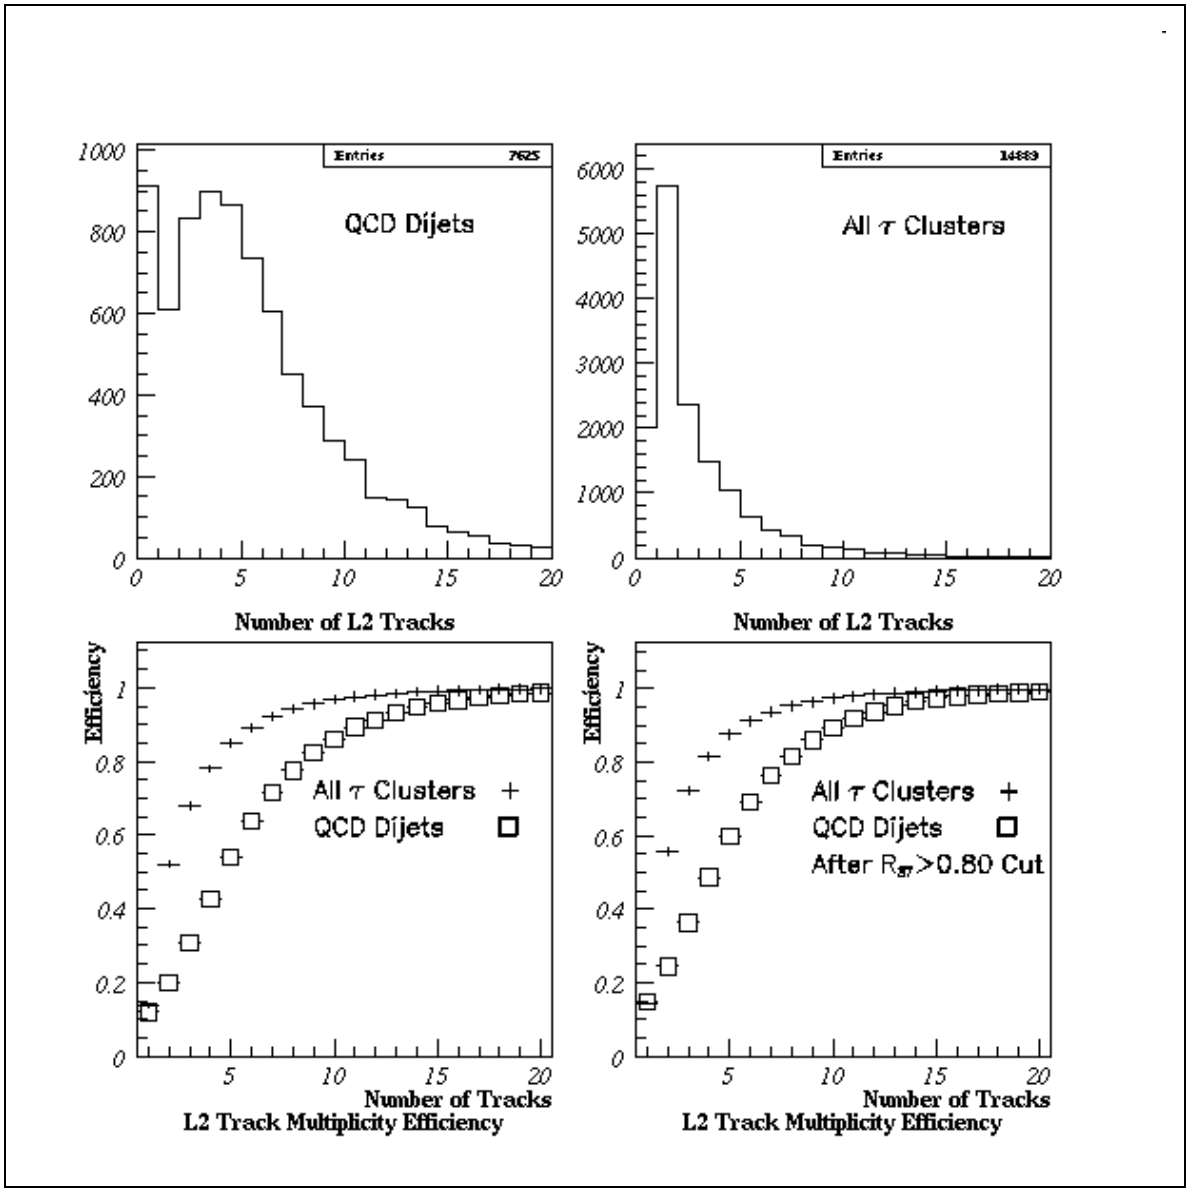

Figure 13-19 The top plots show the inner detector track multiplicities determined at LVL2 for $Z \rightarrow \tau^+\tau^-$ &QCD di-jets The bottom plots show the variation of tau & jet finding efficiency with upper cut on track multiplicity.

Other tau-identification variables that have been considered are the tau hadronic E_T and a tau-likelihood variable included in the "Taurec" reconstruction package in the ATHENA framework. The likelihood function uses simple analytical fits to the distributions of the variables mentioned above, as well as the p_T of the highest p_T track. The variables mentioned above still have power at the EF level to reject QCD jets whilst retaining an adequate efficiency for selecting taus. An analysis of the rejection of QCD jets at the EF level is currently under way [13-60].

13.4.3.1.4 Jet Rejection at LVL2 Following a Tau +E_T-miss trigger

A method of improving the signal acceptance for final states involving taus, as well as retaining an acceptable trigger rate, is to combine the stand-alone tau trigger with an E_T^{miss} trigger. The effect of adding a LVL1 E_T^{miss} requirement on stand-alone tau HLT rates is shown in Table 13-6 for one set of LVL2 trigger criteria and four trigger-threshold configurations: τ20 + xE25; τ20 +

xE35; τ30 + xE25; and, τ30 + xE35. As can be seen from the table the LVL1 E_T^{miss} requirement substantially enhances the jet rejection and thus reduces considerably the tau trigger rate coming after LVL2.

Table 13-6 Tau and QCD di-jet efficiencies for various LVL1 & LVL2 trigger criteria. In the right-hand column the corresponding tau stand-alone LVL2 trigger rates are given including the LVL1 E_T^{miss} requirement.

LVL1 Tau + E_T-miss trigger (GeV)	LVL2 Cuts	LVL2 tau eff. wrt LVL1 (%)	LVL2 jet eff. wrt LVL1 (%)	LVL2 stand-alone tau trigger rate (kHz)
τ(20-10-10) + xE25	-	29 ± 1	10 ± 1	1.56 ± 0.16
τ(20-10-10) + xE35	-	11 ± 1	1.9 ± 0.2	0.30 ± 0.03
τ(20-10-10) + xE25	R_{37}>0.9 #LVL2tracks<5	15 ± 1	1.6 ± 0.2	0.25 ± 0.03
τ(20-10-10) + xE35	R_{37}>0.9 #LVL2tracks<5	5.1 ± 0.2	0.5 ± 0.1	0.07 ± 0.01
τ(30-10-10) + xE25	-	43 ± 1	18 ± 1	0.85 ± 0.05
τ(30-10-10) + xE35	-	20 ± 1	4.3 ± 0.4	0.20 ± 0.02
τ(30-10-10) + xE25	R_{37}>0.9 #LVL2tracks<5	23 ± 1	3.6 ± 0.4	0.17 ± 0.02
τ(30-10-10) + xE35	R_{37}>0.9 #LVL2tracks<5	9.6 ± 0.5	1.2 ± 0.2	0.06 ± 0.02

13.4.3.2 E_T-miss Trigger

Missing transverse energy will provide a distinct and important signature for new physics at the LHC. A precise and reliable measurement of E_T^{miss} requires good calorimeter performance and energy resolution, good linearity of response and hermetic coverage over a wide range of pseudorapidity.

The E_T^{miss} + jet trigger is an example of a trigger based on the combination of a global variable (E_T^{miss}) and localized RoIs in the detector. The bulk of the trigger rate will result from fluctuations in the energy measurements of QCD jets, partly as the result of the presence of material in front of the calorimeters and in the regions between the various calorimetry sections. The main instrumental effects arise from the difference in response between the various calorimeter technologies and from the fact that the calorimetry is highly non-compensating.

The contribution of the EF to reducing the LVL1 E_T^{miss} trigger rate should be important for three main reasons. Firstly, accurate calorimeter calibration and inter-calibration are available [13-49]. Secondly, a separate calibration for cells between clusters can be utilized. Thirdly, the cell E_T cutoff applied to suppress noise can be tuned accurately [13-66]. Initial studies of the E_T^{miss} + jet trigger rate used samples of QCD di-jets (from Pythia) made with both fast and full simulation. Excellent agreement was obtained between these simulation methods. Details of the simulation can be found elsewhere [13-67].

The E_T^{miss} + jet trigger-rate calculation has been revisited using the latest ATHENA object-oriented reconstruction software framework. In addition, the following improvements were included in the analysis: an updated version of the Pythia event generator; an updated model for minimum-bias interactions; a new algorithm for E_T^{miss} calculation utilizing the H1 calibration approach as well as tuned cell cuts. Initially the low-luminosity scenario has been considered without pile-up but with electronic noise added to the simulation. The preliminary results at 2×10^{33} cm^{-2} s^{-1} give an E_T^{miss} + jet trigger rate of 50±12 Hz for (E_T^{jet}, E_T^{miss})>(60, 60) GeV, and

40±12 Hz for (E_T^{jet}, E_T^{miss})>(70, 70) GeV. These rates are consistent with earlier results [13-1] and can be reduced further by applying topological cuts as in the past. A plot of the trigger rate versus the E_T^{miss} threshold is shown in Figure 13-20 for both jet threshold settings.

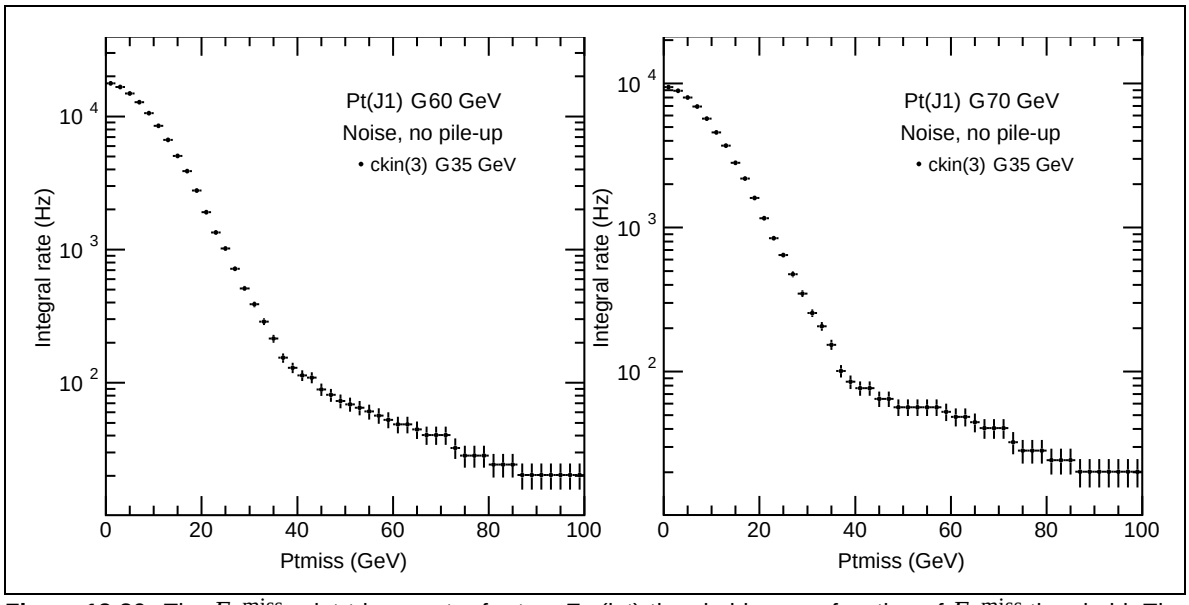

Figure 13-20 The E_T^{miss} + jet trigger rate, for two E_T (jet) thresholds, as a function of E_T^{miss} threshold. The plots are for a luminosity of 2×10^{33} cm^{-2} s^{-1} but without pile-up included. Electronic noise has been added.

Recently estimates of the tau + E_T^{miss} HLT rate have been made using the ATHENA reconstruction software. In this case the HLT rate was estimated by assuming that the existing offline tau and E_T^{miss} code to simulate the performance of the tau + E_T^{miss} HLT. The effects of electronic noise in the calorimetry were included in the simulation. This preliminary analysis of the low-luminosity trigger rates did not include the effects of pileup. The same offline tau-identification criteria bulleted in Section 13.4.3.1.3 were used. Also, a cut of ~2σ in the electronic noise was applied to all calorimeter cells immediately after reconstruction and before any further analysis. In this case an improvement was made to the standard offline analysis by making the tau-identification criteria dependent on the E_T range [13-68]. In this way the cuts could be optimised for the highest rejection for a given efficiency. It was seen that the jet rejection came mainly from the E_T^{miss} cuts and after that from the tau-identification criteria. This preliminary analysis had only limited statistics and thus only a rough estimate of the trigger rate is possible. However, fixing our cuts to give a tau efficiency of 55%, and taking the tau E_T >100 GeV, the estimated tau + E_T^{miss} rate is ~5Hz.

13.4.3.3 Triggering on Jets

The purpose of the high-level jet trigger is to reduce the rate of events containing jets, compared to the LVL1 one, by exploiting the improved E_T measurement. This improvement is achieved by applying a more refined energy calibration and jet definition, and using the fine-granularity and high-precision readout data. As jets are the dominant high-p_T process rate reduction cannot be expected by removing fake jet objects. In contrast to e.m. clusters, jets cannot be flagged as background (unless they are due to instrumental effects). Jets are searched for in the region of $|\eta| < 3.2$ in the e.m. and hadronic calorimeters. Hereafter, preliminary results will be presented using the following threshold values and luminosity scenarios: j180, j75x3, and j55x4 at

$1 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and j290, j130x3, and j90x4 at $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Work is in progress to better adapt the studies to the present start-up scenario.

The LVL2 jet finding is guided by the LVL1 jet regions of interest (RoI) — jets are sought in a region around these RoIs. In the current view the EF performs jet finding across the whole pseudorapidity region with the improved jet calibration and inter-jet calibration allowed at this level. The performance of the high-level jet trigger is described in detail elsewhere [13-69]. The jet rates in this report were obtained using fully-simulated QCD di-jet events.

Table 13-7 The top portion of the table shows HLT rates, obtained using fully-simulated events, for different trigger-menu items [13-69]. The middle portion of the table shows the estimated rates using a fast simulation [13-70]. The bottom portion of the table shows the jet-trigger rates for the two menu values indicated.

item	$1 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ LVL2+EF (Hz)	Item	$1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ LVL2+EF (Hz)
j180	253 ± 26	j290	275 ± 87
3j75	286 ± 28	3j130	440 ± 110
4j55	127 ± 15	4j90	165 ± 67
j180	228 ± 8	j290	270 ± 33
3j75	314 ± 10	3j130	270 ± 33
4j55	153 ± 7	4j90	149 ± 25
2j180	101 ± 6	2j290	109 ± 21
2j130	377 ± 11	2j230	358 ± 38

The ATLAS fast-simulation program ATLFAST (version 00-02-22) was used to generate 13 million di-jet events that were utilized to provide an estimate of the HLT (LVL2+EF) jet-trigger rates for one, two, three and four-jet final states [13-70]. In this case the jet-finding algorithm required a jet-initiator E_{T} of 5 GeV, a jet-cone radius R=0.4 and a jet E_{T} threshold of 10 GeV. The trigger rates were normalized, using a single scaling factor, to the rates predicted, for low and high-luminosity running, in a previous analysis based on fully-simulated QCD di-jets [13-69]. The one, two, three, and four-jet trigger rates are reported in Table 13-7 for the thresholds given in reference [13-69] and obtained using the ATLFAST simulation.

13.4.4 b-tagging

The selection of b-jets at trigger level can improve the flexibility of the HLT scheme and possibly extend its physics performance. In particular, for topologies containing several b-jets, the ability to separate b-jets from light-quark and gluon jets could increase the acceptance for signal events (if the use of jet-E_{T} thresholds lower than those discussed in Section 13.5 is feasible) or reduce the background (and hence the rate) for events containing b-jets that have already been selected by other triggers.

The study presented in this section defines and characterizes, in the low-luminosity case, a b-jet selection for LVL2 [13-71] based on Inner Detector information. The extension of the selection to the EF and the definition of an overall strategy for the b-tagging will be addressed in future studies.

13.4.4.1 LVL2 track reconstruction for b-tagging selection

The track reconstruction and the precise determination of the track parameters (in particular the impact parameter (d_0) in the transverse plane) are crucial ingredients of the b-jet trigger.

Several tracking algorithms based on the silicon detectors have been presented in Section 13.3. This study is based on SiTrack [13-72], a LVL2 algorithm that selects and builds track candidates using triplets of space points, because of its good impact-parameter resolution. An early version of SiTrack, PixTrig (based solely on the pixel detector), was used to perform a similar study for the Technical Proposal [13-1].

The recent change in the pixel detector geometry (increased B-layer radius) caused a significant degradation in the impact-parameter resolution for tracks built using the pixel detector alone. The impact-parameter resolution has been recovered, and slightly improved at high p_T (see Figure 13-21), using the capability of SiTrack to use any layer of the inner detector — the comparatively large lever arm granted by an SCT layer ensures a better resolution of the track parameters in transverse plane.

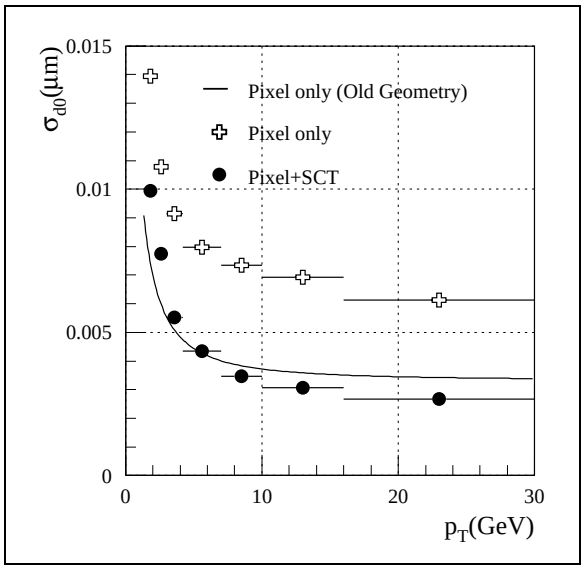

Figure 13-21 Comparison of d_0 resolution as a function of p_T for different configuration of the LVL2 tracking algorithm SiTrack. The label “Old Geometry” refers to the results obtained in Ref. [13-1] with an earlier geometry of the pixel detector (the so-called Physics-TDR layout).

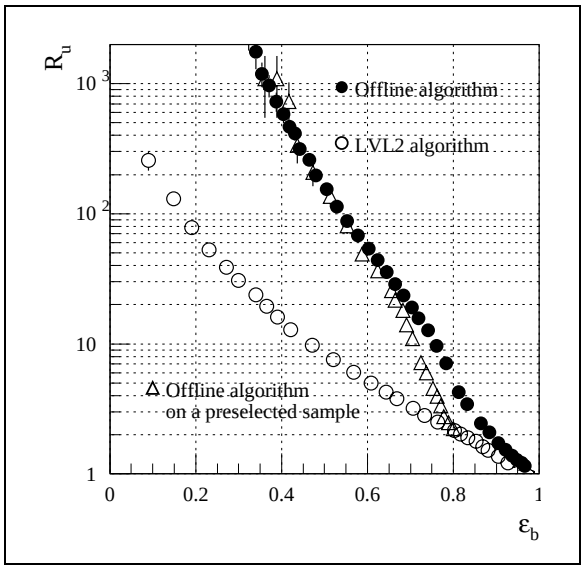

Figure 13-22 he u-jet rejection as a function of the b-jet efficiency for jets from the decay of Higgs bosons with m_H=120 GeV for trigger and offline algorithms. The bias of the LVL2 selection on the offline selection is shown.

13.4.4.2 b-tagging algorithm

The b-tagging selection starts with track reconstruction performed by SiTrack within the LVL1 jet RoI. For each reconstructed track the significance of the transverse impact parameter $S = d_0 / \sigma(d_0)$ is computed; the error on the impact parameter $\sigma(d_0)$ is parameterized, using simulated events, as a function of p_T.

The b-jet discriminator variable is then built using the likelihood-ratio method; for each track (i) in the jet, the ratio of the probability densities for the track to come from a b-jet or a u-jet, $f_b(S_i) / f_u(S_i)$, is calculated; the product $W = \Pi\, f_b(S_i) / f_u(S_i)$ of these ratios over all tracks in the jet is com-

puted; finally the tagging variable $X = W/(1+W)$ is defined. Jets are tagged as b-jets if $X \sim 1$ and u-jets if $X \sim 0$. The selection efficiency for b-jets and the rejection of light-flavour jets can be tuned varying the cut on the X variable.

The processing time of the b-tagging selection algorithm is dominated by the track-reconstruction phase and today is less than 8 ms per jet on a 2 GHz processor for dijet events. Note that the figure of 8 ms does not include the data-preparation time that remains to be assessed in the new HLT software framework.

13.4.4.3 Results on single b-jet tagging

The b-tagging algorithm has been characterized on single b-jets coming from $H \rightarrow b\bar{b}$ decays (m_H = 120 GeV) produced in association with a W boson at low luminosity, and corresponding u-jets (taken as representative of the light-flavour jets) obtained by artificially replacing the b-quarks from the Higgs decay with u-quarks; the LVL1 jet RoI was simulated by selecting a region $\Delta\phi \times \Delta\eta = 0.4 \times 0.4$ centred around the direction of the quarks coming from the decay of the Higgs boson.

The efficiencies for b-jets (ε_b) and rejection factors (R_u) against u-jets (defined as the inverse of the efficiency for u-jets) are given in Table 13-8. The modest rejections obtained do not spoil the

Table 13-8 Rejection of the LVL2 b-tagging algorithm against u-jets for three different values of the b-jet efficiency: 60% (top), 70 % (middle) and 80 % (bottom) at low luminosity. The results are shown for different intervals of E_T and η of the jet.

	40 GeV < E_T < 70 GeV	**70 GeV < E_T < 100 GeV**	**E_T > 100 GeV**
$\|\eta\| < 1.5$	5.9 ± 0.4	5.1 ± 0.6	4.8 ± 0.7
	3.6 ± 0.2	3.9 ± 0.4	3.1 ± 0.4
	2.3 ± 0.1	2.8 ± 0.3	2.4 ± 0.3
$\|\eta\| > 1.5$	3.7 ± 0.4	4.9 ± 1.0	3.0 ± 0.8
	2.6 ± 0.2	2.9 ± 0.5	2.6 ± 0.6
	1.5 ± 0.1	1.6 ± 0.2	1.7 ± 0.4

interest of applying a b-tagging selection at LVL2: in events with multiple b-jets, loose selections (necessary to minimize the bias on the selected sample) applied to several jets can still produce significant rejections

13.4.4.4 Comparison with Offline b-tagging

The performance of the LVL2 b-jet tagging algorithm has been compared to that of an offline algorithm based on impact parameter measurements in the transverse plane. This choice provides a coherent comparison of online and offline algorithms; more exhaustive comparison studies will be performed on specific physics selections in due course.

Figure 13-22 demonstrates that the LVL2 and offline selection are well correlated and that, with an appropriate tuning of the LVL2 selection, it is possible to provide subsequent analyses with an unbiased sample.

Different combinations of working points for the LVL2 trigger selection and offline analysis could be chosen depending on the topology of the events and on the required offline b-tagging efficiency. Additional flexibility can be added to the b-tagging selection by integrating the LVL2 selection with a selection at EF, where tracking with offline quality will be available.

13.4.5 B-physics

About one collision in every hundred will produce a $b\bar{b}$ quark pair. Therefore, in addition to rejecting non-$b\bar{b}$ events, the B-trigger must have the ability to identify and select those events containing B-decay channels of specific interest. Important areas of study include: CP-violation measurements with the channels $B_d \to J/\psi K_S$ (with both $J/\psi \to e^+e^-$ and $J/\psi \to \mu^+\mu^-$) and $B_d \to \pi^+\pi^-$ (and more generally $B_{d,s} \to h^+h^-$, where h denotes a pion or kaon); measurements of B_s oscillations using $B_s \to D_s\pi$ and $B_s \to D_s a_1$ decays with $D_s \to \phi\pi$; analysis of $B_s \to J/\psi\phi$ and $B \to J/\psi\eta$ final states, including the search for CP-violation in $B_s \to J/\psi\phi$ decays where new physics could enhance the effect significantly beyond Standard Model expectations; rare decays of the type $B_{d,s} \to \mu^+\mu^-(X)$; and *B*-hadron production measurements. High statistics are required for the precision measurements. As documented in Ref. [13-49], ATLAS is well placed to make significant and competitive contributions in many of these areas.

Since the HLT/DAQ/DCS Technical Proposal [13-49] the B-trigger has been re-assessed in the light of a number of developments, including the likelihood of a reduced inner-detector layout at the start of running, a factor of two increase in the target start-up luminosity and various trigger scenarios for deferring spending on the HLT/DAQ system. The aim is to retain the maximum possible coverage of key B-physics channels within the available resources.

It is important to study a range of scenarios since the actual LHC machine start-up conditions are uncertain, the luminosity is expected to vary from fill-to-fill, and there are large uncertainties in the physics cross-sections and in the calculation of required resources. A flexible trigger strategy has, therefore, been developed based on a di-muon trigger (with p_T thresholds, which may be η dependent, of about 3–6 GeV) at luminosities of $\sim 2 \times 10^{33}\ cm^{-2}\ s^{-1}$ or above, and introducing other triggers at lower luminosities. Lower luminosities will occur since some LHC fills will have lower luminosities than others, and the luminosity will fall during each beam-coast (over the period of a beam-coast the luminosity is expected to fall by about a factor of two).

Two strategies have been investigated for these additional triggers, as follows:

- The preferred strategy is to require at least one LVL1 jet or EM RoI in addition to a single-muon trigger (e.g. $p_T > 8$ GeV). After validating the LVL1 muon at LVL2, and in the EF, track reconstruction is performed within the RoIs using pixel, SCT and, optionally, TRT information. The reconstructed tracks form the basis of selections for, e.g., $J/\psi(e^+e^-)$, $B(h^+h^-)$ and $D_s(\phi\pi)$. Since track reconstruction is performed only inside the RoIs, the resources required are modest. More details are given in Ref. [13-73].
- If the LVL1 RoI multiplicity proves too high (or the efficiency too low) for the above approach, a fall-back solution is to perform track reconstruction within the full acceptance of the SCT and pixel detectors (so-called full-scan) for events with a LVL1 single-muon trigger (e.g. p_T>8 GeV), after confirmation of the muon at LVL2. In order to minimize execution time, the TRT is not presently included and so the $J/\psi(e^+e^-)$ trigger is not possible. The reconstructed tracks form the basis of selections for, e.g., $B(h^+h^-)$ and $D_s(\phi\pi)$. This requires somewhat greater resources than the method described above, in order to perform

the full-scan, but promises better efficiency. This strategy is described in detail in Ref. [13-74].

In all cases, at least one LVL1 trigger muon is required to initiate the B-trigger. Since the inclusive cross-section for muon production from decays in flight of pions and kaons falls more rapidly with p_T than that for prompt muon production from B-hadron decays, see Figure 13-23, an appropriate choice of p_T threshold gives a powerful reduction of the trigger rate due to background processes. For example, a p_T threshold of 8 GeV gives a single-muon trigger rate of 10 kHz at LVL1 for a luminosity of 1×10^{33} cm^{-2} s^{-1}. Most of this rate is due to muons with true p_T below threshold, originating from pion and kaon decays in flight, a large proportion of which can be rejected at LVL2 on the basis of more precise measurements in the muon spectrometer and inner detector. After the LVL2 selection the single-muon trigger rate is about 2 kHz at a luminosity of 1×10^{33} cm^{-2} s^{-1}; about one third of this rate is due to b→μ decays. It is important not to set the muon p_T threshold too high as this would significantly reduce the statistics in the signal channels of interest and render the measurements un-competitive. The rate is further reduced at LVL2 and in the EF by requiring other triggers in addition to the muon, as described in the following sections.

13.4.5.1 Di-muon triggers

A di-muon trigger provides a very effective selection for several important B-physics channels, e.g. $B_d \rightarrow J/\psi(\mu^+\mu^-)K_s$ and rare decays of the type $B \rightarrow \mu^+\mu^-(X)$. The LVL1 muon trigger is efficient down to a p_T of about 5 GeV in the barrel region and about 3 GeV in the end-caps. However the actual thresholds used for the di-muon trigger will be determined by rate limitations. For example, a p_T threshold of 6 GeV would give a di-muon trigger rate of about 200 Hz after LVL1 at a luminosity of 2×10^{33} cm^{-2} s^{-1}. These triggers are mostly due to muons from heavy-flavour decays, see Figure 13-23, plus some single muons which are doubly counted when they traverse more than one pivot plane in the muon trigger chambers (in most cases such double counting is already resolved at LVL1). Doubly-counted single muons are removed at LVL2 which also sharpens the p_T threshold, reducing the rate to about 200 Hz. In the EF, the tracks are reconstructed using offline-quality algorithms and specific selections are made on the basis of mass and decay-length cuts. These consist of semi-inclusive selections, for example to select $J/\psi(\mu^+\mu^-)$ decays with a displaced vertex, and in some cases exclusive selections such as for $B_{d,s} \rightarrow \mu^+\mu^-$. Estimated trigger rates are shown in Table 13-9..

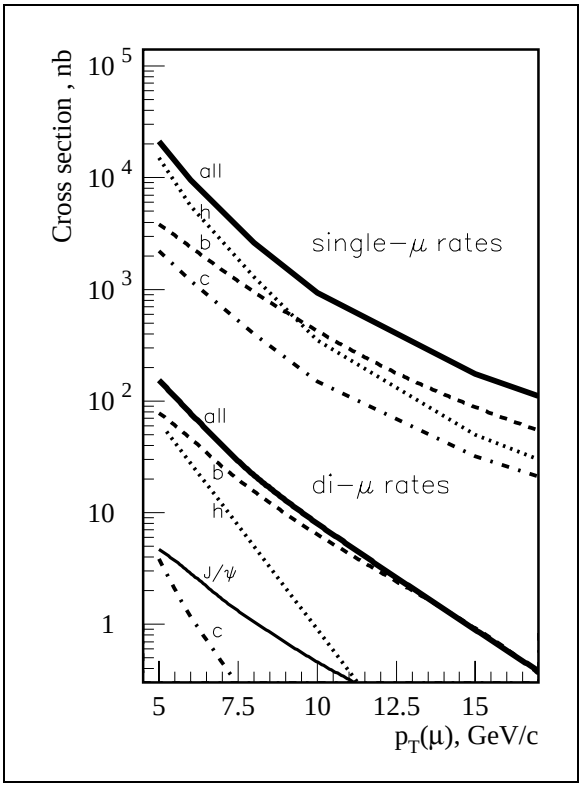

Figure 13-23 Single-muon and di-muon cross-sections. Curves are shown for muons from K and π decays in flight (labelled "h"), b and c decays, and for the sum of these sources ("all"). Muons are considered for $|\eta| < 2.5$. For di-muons, the horizontal axis shows the p_T of the ;ower p_T muon. At least one muon must have $p_T > 6$ GeV and $|\eta| < 2.4$ corresponding to LVL1 trigger conditions..

13.4.5.2 Hadronic final states

For events with a muon trigger, two strategies have been studied for selecting hadronic final states based on either an ID full-scan or an RoI-based method. The ID full-scan consists of track reconstruction at LVL2 within the entire volume of the SCT and pixel detectors and, optionally, the TRT. The alternative, preferred, strategy uses low-E_T LVL1 jet clusters to define RoIs for track reconstruction in the ID. By limiting track reconstruction to the part of the ID lying within the RoI, only about 10% of the detector on average, there is potential for a significant saving in data movement and processing time compared to the full-scan (up to a factor of ten, depending on the RoI multiplicity per event).

Preliminary studies of efficiency and jet-cluster multiplicity at LVL1 have been made using a fast simulation program. This uses a rather detailed model of the calorimeter which includes a parametrization of longitudinal and transverse shower profiles, and takes into account the effects of pulse history, digitization and Bunch Crossing IDentification (BCID). These studies indicate that a cut $E_T > 5$ GeV on the transverse energy reconstructed in the LVL1 jet window gives a reasonable mean jet-cluster multiplicity of about two RoIs per event for events containing a muon with $p_T > 6$ GeV[1], see Figure 13-24. .

Table 13-9 Estimated trigger rates for examples of B-physics trigger selections at luminosities of 2×10^{33} cm^{-2} s^{-1} and 1×10^{33} cm^{-2} s^{-1}. In these examples the di-muon selection at LVL2 consists only of confirming the muons. The division of selections between LVL2 and the EF remains to be optimised.

	2×10^{33} cm^{-2} s^{-1}		1×10^{33} cm^{-2} s^{-1}	
Trigger	**LVL2**	**EF**	**LVL2**	**EF**
$B_{d,s} \rightarrow \mu^+\mu^-(X)$	200 Hz	small	100 Hz	small
$J/\psi(\mu^+\mu^-)$		10 Hz		5 Hz
$D_s(\phi\pi)$	–	–	60 Hz	9 Hz
$B(\pi\pi)$	–	–	20 Hz	3 Hz
$J/\psi(ee)$	–	–	10 Hz	2 Hz
Total	200 Hz	10 Hz	190 Hz	20 Hz

These studies have been repeated using a full GEANT-based detector simulation followed by a simulation of the LVL1 calorimeter trigger. All sources of electronic noise are simulated, both for the LAr and Tile calorimeters, as well as for the trigger towers. The effect of the BCID system was not simulated, however, and as a result there is a significant contribution to the RoI multiplicity from small signals from other bunch crossings. The mean number of jet RoIs found per event is shown in Figure 13-24 as a function of the LVL1 E_T threshold. The mean multiplicities are significantly higher than those obtained with the fast simulation — this is most probably explained by the lack of BCID in the simulation. Once BCID is included, the measured multiplicities are expected to be much closer to the results of the fast simulation, although this remains to be verified. The full-simulation results, therefore, represent an upper limit, with the anticipated multiplicity lying closer to the fast-simulation results. The full simulation shows that a LVL1 jet

1. These studies used a muon-p_T threshold of 6 GeV. The LVL1 jet and EM RoI multiplicities can be assumed to be similar for events selected with an 8 GeV muon threshold, within the accuracy of these estimates.

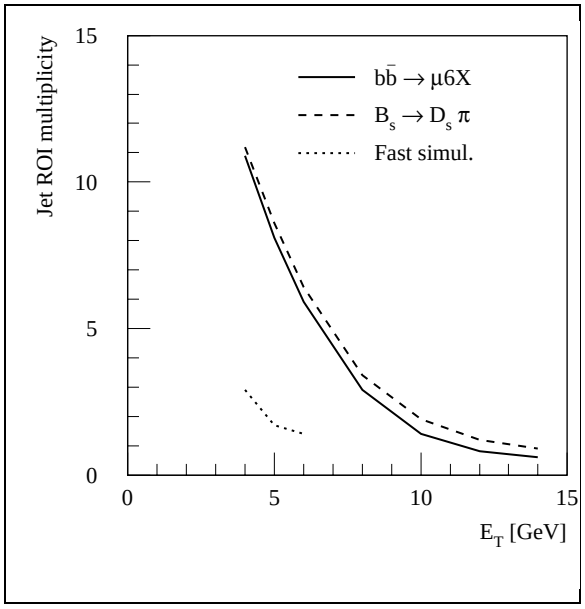

Figure 13-24 The mean number of jet RoIs per event shown as a function of the E_T cut. Results are shown for bb→μX events with a muon of p_T > 6 GeV. Results are shown for both the fast simulation (dotted curve) and full simulation (solid line). The mean multiplicity is also shown for $B_s \rightarrow D_s \pi^+$ signal events with full simulation.

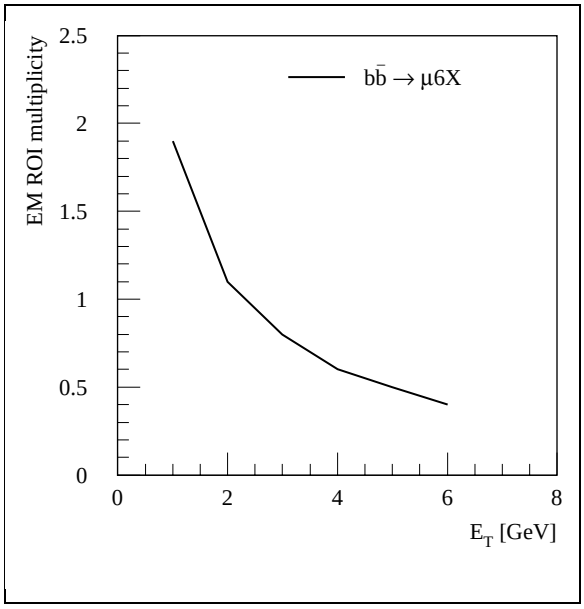

Figure 13-25 The mean number of EM RoIs per event shown as a function of LVL1 E_T cut. Results are shown from the fast simulation for bb→μX events with muon p_T > 6 GeV.

trigger requiring E_T > 6 GeV would give 80% efficiency for finding the RoI for $B_s \rightarrow D_s \pi^+$ events containing a B_s meson with p_T > 16 GeV

Based on the ID tracks reconstructed by one of the above methods (either full-scan or RoI-based) further selections are made for specific channels of interest. These are kept as inclusive as possible at LVL2, with some more exclusive selections at the EF. For example, samples of $B_s \rightarrow D_s \pi$ and $B_s \rightarrow D_s a_1$ events can both be triggered by selecting events containing a $D_s(\phi\pi)$ candidate.

Tracks are reconstructed in the EF inside RoIs defined from the results of LVL2. Using LVL2 to guide the EF reconstruction reduces the amount of data to be processed. For example, a region encompassing all LVL2 tracks forming $D_s(\phi\pi)$ or $B(\pi^+\pi^-)$ candidates corresponds to about 10% of the ID acceptance, on average. Tighter mass cuts can be applied in the EF than at LVL2 since there is time available for a more refined track reconstruction which yields better track-parameter resolutions. In addition, EF selections may include precise decay-vertex reconstruction, allowing further cuts on vertex-fit quality and decay length.

Studies using a full-detector simulation have shown that an overall HLT efficiency of about 60%[1] can be obtained for $B_s \rightarrow D_s \pi$ signal events where all final-state particles have p_T > 1.5 GeV. The corresponding trigger rates are shown in Table 13-9; further details are given in Ref. [13-75]. It has been shown that, with appropriate tuning of selection cuts, there is very little degradation of trigger performance if the pixel layout is changed from three barrel layers to the two layers expected at the start of LHC running. For example, studies based on the LVL2 IDSCAN algo-

1. These studies were based on a full-scan of the SCT and pixels at LVL2 with a p_T threshold for reconstructed tracks of 1.5 GeV. The efficiency for an RoI-based trigger has yet to be measured.

rithm show that reducing the cut on the number of SCT and pixel space-points to a minimum of four for the initial layout (compared to a minimum of five for the full layout) yields the same efficiency as for the full layout with only a 10% increase in trigger rate. Other studies have shown that the trigger is insensitive to the anticipated levels of misalignment in the ID [13-76].

13.4.5.3 Muon–electron final states

A muon–electron trigger is used to select channels such as $B_d \rightarrow J/\psi(e^+e^-)K_s$ with an opposite-side muon tag, or $B_d \rightarrow J/\psi(\mu^+\mu^-)K_s$ with an opposite-side electron tag. The LVL1 trigger is used to find low-E_T electron / photon clusters which define RoIs to be investigated at LVL2. Preliminary studies, using a fast simulation, show that a reasonable compromise between RoI multiplicity and electron efficiency might be obtained requiring E_T>2 GeV for the LVL1 cluster. This gives a mean RoI multiplicity of about one for events containing a muon with $p_T > 6$ GeV, see Figure 13-25. These studies give an efficiency of about 80% to find an RoI for both the e^+ and the e^- from $J/\psi(e^+e^-)$ in events where both final-state particles have $p_T > 3$ GeV. At LVL2, the electron / photon RoIs are confirmed in the calorimeter, using full-granularity information and including the pre-sampler. Inside the RoI, a search is then made for tracks in the SCT, pixels and, optionally, the TRT. The RoI about each electron candidate can be quite small, about $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$. This gives a large saving in reconstruction time, compared to a full-scan, but requires a higher p_T threshold than is possible with the ID full-scan. The tracks are reconstructed again in the EF, including a vertex fit, and cuts are applied to the decay length and fit quality. The estimated trigger rates are shown in Table 13-9.

13.4.5.4 Resource estimates

In order to estimate the computing resources required for the B-trigger, measurements of algorithm execution times have been combined with estimates of trigger rates at each step of the selection. Various reconstruction algorithms have been timed on several different platforms in order to determine the mean execution time at a given luminosity, the scaling of execution time with the number of hits in an event, and hence the scaling with luminosity. These timing measurements have been combined with the estimates of trigger rates and RoI multiplicities to give an estimate of the computing resources required for the B-physics trigger [13-77].

It is estimated that ~15 additional 2 GHz processors are required for the B-physics LVL2 trigger, using the preferred RoI-guided strategy. The corresponding number is about 50 processors for the fall-back full-scan strategy. These estimates are based on algorithm execution times measured in a previous software framework, CTrig. Preliminary measurements in the new Athena framework give algorithm execution times (excluding data access overheads) that are a factor 2–3 higher than the CTrig measurements. It is anticipated that with further optimization, algorithm speeds will approach those obtained in the CTrig framework.

In summary, the use of low-E_T LVL1 RoIs to guide reconstruction at LVL2 promises a full programme of B-physics for very modest resources. However, RoI multiplicities and efficiencies need to be verified in studies using a full-detector simulation and including a detailed simulation of the LVL1 trigger, with BCID. This will be followed up in further studies.

13.5 HLT output rates to off-line

After the submission of the HLT/DAQ Technical Proposal [13-1], the baseline scenario for the start-up of the LHC was changed to a target initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. The information already available from the detailed HLT performance studies, as documented in the TP, was used to derive the ATLAS HLT trigger menu for this new scenario; the work concentrated on the parts of the menu contributing most of the rate, consisting mainly of inclusive high-p_T object selections. In addition, the impact of constraints coming from delayed availability of financial resources was taken into account by restricting the B-physics trigger selection at peak luminosity to di-muon signatures only (see Section 13.4.5). The resulting trigger menus for LVL1 and for the HLT are presented in Table 13-10 and Table 13-11 respectively. The notation used in these tables is explained in Chapter 4.

The signatures shown in the trigger menus are needed to guarantee complete coverage of the ATLAS physics programme for observations of new physics and precision measurements. As discussed in Chapter 4, the aim is: to be as open as possible to (perhaps unpredicted) signatures of new physics; to avoid as much as possible any bias introduced in the trigger selection; and to allow refined selection criteria to be used in the offline analyses. It should be noted that these menus assume stable operation of the LHC machine and the ATLAS detector, and thus do not apply to the initial start-up phase of ATLAS and the LHC, which is addressed in more details in the following section.,

Table 13-10 LVL1 trigger menu with rates for a luminosity of 2×10^{33} cm^{-2} s^{-1}

LVL1 signature	rate (kHz)
EM25I	12.0
2EM15I	4.0
MU20	0.8
2MU6	0.2
J200	0.2
3J90	0.2
4J65	0.2
J60+XE60	0.4
TAU25I+XE30	2.0
MU10+EM15I	0.1
others (pre-scaled, exclusive, monitor, calibration)	5.0
Total	~25.0

Table 13-10 lists the main LVL1 selection signatures together with the rates expected for an initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. For most of the signatures, the nominal threshold given is defined to give an efficiency of about 95% for a physics object (μ, e/γ, etc.) with E_T equal to the threshold value. This is achieved by an appropriate tuning of the selection cuts at LVL1. The selection signatures at LVL1 include single- and di-object isolated electromagnetic-calorimeter clusters (which provide candidates for electron and photon object reconstruction at the HLT), single- and di-muon candidates, and the combination of an isolated electromagnetic cluster and

a low-p_T muon. In addition, hadronic jets are selected for various multiplicity requirements and E_T thresholds. It is expected to also include a di-jet trigger (with a threshold lower by several tens of GeV than the one of the inclusive single-jet trigger), which is not shown in the table. Finally, signatures requiring large missing transverse energy, in association with either a jet or a tau[1] candidate are present. About 20% of the total rate of 25 kHz is allocated for pre-scaled, exclusive, calibration, monitor and other triggers, examples of which are described in Chapter 4**.**

It should be noted that the LVL1 trigger rate estimates are obtained using a physics event Monte Carlo generator and a detailed simulation of the ATLAS detector response. They do not contain any safety factor against uncertainties on the rates, for which possible sources are discussed detail later. For the design value of the maximum LVL1 accept rate of 100 kHz, this trigger menu thus gives an effective safety factor of four against possible uncertainties affecting the LVL1 trigger rate.

Table 13-11 HLT trigger menu with rates for a luminosity of 2×10^{33} cm^{-2} s^{-1}

HLT signature	rate (Hz)
e25i	40
2e15i	<1
γ60i	25
2γ20i	2
μ20i	40
2μ10	10
j400	10
3j165	10
4j110	10
j70+xE70	20
τ35i+xE45	5
2μ6 with vertex, decay-length and mass cuts (J/ψ, ψ', B)	10
others (pre-scaled, exclusive, monitor, calibration)	20
Total	~200

Table 13-11 shows the HLT output rate corresponding to the above LVL1 trigger menu; as for before, it does not include any safety factor against uncertainties on the rates. The rates shown were obtained using the selection algorithm steps at LVL2 and in the EF for the various objects, as described in detail in the previous sections of this chapter. The LVL1 electromagnetic selections separate at the HLT into single-/di-electron and single-/di-photon signatures, which together account for about 1/3 of the HLT output rate. About 25% of the HLT output rate is from the single- and di-muon triggers, whereas single- and multi-jet triggers constitute about 15% of

1. In past documents [13-49], [13-3], [13-1] the LVL1 tau trigger always implicitly required isolation criteria to identify narrow hadronic jets from tau decays, which was not signalled by the letter 'i' in the physics object shown in trigger menus. In this document, the notation has been updated to indicate more correctly the selection conditions.

the total rate. Selections involving missing transverse energy contribute about 15% of the rate. Only 5% of the HLT rate at peak luminosity is allocated to B-physics related triggers (the low-p_T di-muon signature with additional mass cuts to select J/ψ, ψ' and rare B-meson decays). About 10% of the total rate is allocated for pre-scaled and other triggers.

As already mentioned for the LVL1 trigger menu, it is important to note that the rate estimates are the result of simulations, which start with a physics event generator (mostly PYTHIA) and then involve a detailed GEANT-based simulation of the ATLAS detector response. The rate estimates are thus subject to several sources of uncertainties:

- **Knowledge of cross-sections:** some of the cross-sections (e.g. the ones for multi-jet production) have big uncertainties (factors of two or more), which directly affect the corresponding contributions to the trigger rates.
- **Realistic detector behaviour and performance:** the results presented in this document are obtained using the simulated detector behaviour, thus they apply only for stable running conditions with a commissioned and well-understood detector.
- **Background conditions:** unforeseen beam-related background conditions could have an impact on the trigger rates; such backgrounds could lead to increased rates (possibly only from certain regions of the detector).
- **Resources constraints and software performance:** the timing performance of the selection software when faced with real data might limit the rate capabilities of the HLT (or intermediate stages of it) given constraints from the availability of computing resources — the choice of LVL1 thresholds for the physics objects used in the on-line selection might have to be adjusted as a result, with implications also for the HLT thresholds.

As mentioned already, the menus discussed above are shown for the target initial peak luminosity of 2×10^{33} cm^{-2} s^{-1}. It is expected that during a beam coast, as the luminosity drops by a factor of about two or more, or for machine fills which do not reach the target peak luminosity, the selection will be enlarged by assigning a larger fraction of the available rate to pre-scaled triggers (and / or to include additional exclusive or topological selections, which are not activated at the target luminosity). This would result in the writing of events at a constant rate to mass storage, making full use of the online data-movement and computing resources.

13.6 Start-up scenario

The TDAQ system will have to cope with an evolving data-taking environment, in particular during the start-up phase of the detector operation. It will have to accommodate:

- quickly-evolving luminosity;
- variable beam-related background conditions;
- variable electronic-noise conditions;
- changing configuration of sub-detector electronics, subject to final adjustment iterations;
- limited data-handling capacity of the initial HLT / DAQ system imposed by financial limitations and the need to add computing and data-movement capacity as new resources become available;
- learning stages of the detector operation.

The TDAQ system commissioning thus has to be considered as a self-adjusting process, optimised to use efficiently the data available at each stage of the detector operation — at first data collected in sub-detector stand-alone runs, then data collected in the cosmic-ray muon runs, subsequently data collected in single-beam operation of LHC, and eventually beam–beam collision data.

13.6.1 LVL2 trigger and EF commissioning

The commissioning phase of the HLT/DAQ system will follow that of the subdetectors and of the LVL1 trigger. Many issues will need to be addressed at that stage, to achieve proper synchronization of LVL1 signals from trigger detectors (calorimeters, RPC, TGC) and of the CTP. This relates as well to the operations of Data Collection, Data Flow and Event Builder components, that feeds the algorithms at the LVL2 and EF levels with the proper data fragments or fully-assembled events. Also data consistency from all ATLAS sub-detectors will need to be established, and pre-processing and formatting will need to be studied since they directly impact the behaviour of the data-preparation step of the HLT algorithms.

Once the sub-detector data have been understood sufficiently (e.g. in terms of calibration and alignment), it will be necessary to debug and tune the selection of events at LVL2 and in the EF. The commissioning of LVL2 and the EF can be facilitated by splitting it into several phases. Initially the HLT systems will probably be used only to flag events, allowing offline debugging and testing of the selection before it is used actively in the online environment.

Needless to say, as many preparations as possible will made before the first beam–beam collisions occur, starting with test data, followed by an extended period of cosmic-ray running. It may be possible to carry out further commissioning during running with a single proton beam, e.g. using beam–gas interactions in the ATLAS detector.

Once the HLT is operational, there will be a period in which the selection criteria at all trigger levels are adapted, perhaps radically, in the light of experience from the first running. Account will have to be taken of any limitations in the performance of the initial HLT/DAQ system (e.g. any bottlenecks in the Dataflow or HLT systems), and also any problems with the detectors or the LVL1 system (e.g. ‘hot channels’). The trigger menus will have to be tuned, taking into account the measured rates for different LVL1 signatures, to optimize the physics coverage within the available HLT/DAQ capacity. Account will also have to be taken of triggers due to beam-related backgrounds that might be particularly significant in the early running.

13.7 References

13-1 *ATLAS High-Level Triggers, DAQ and DCS Technical Proposal*, CERN/LHCC/2000-17 (2000)

13-2 ATLAS Data Challenge 1, http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/DC/DC1/DC1-Report-V1.0-211002.pdf

13-3 *ATLAS Level-1 Trigger Technical Design Report*, CERN/LHCC/98-14 (1998)

13-4 T. Schörner-Sadenius and T. Wengler, *Formats and Interfaces in the Simulation of the ATLAS First Level Trigger*, ATL-DA-ES-0029 () https://edms.cern.ch/document/345463/

13-5 M. Abolins et al., *Specification of the LVL1 / LVL2 trigger interface*, ATL-D-ES-0003 () https://edms.cern.ch/document/107485/

13-6 T. Schoerner-Sadenius, *The ATLAS Level-1 Trigger offline Simulation: Overview, Core Parts, and Use*, ATLAS Internal Note, ATL-COM-DAQ-2003-008 (2003)

13-7 *Xerces: XML parser in C++*, http://xml.apache.org/xerces-c/

13-8 E. Eisenhandler, *Level-1 Calorimeter Trigger URD*, ATL-DA-ES-0001 (1998) https://edms.cern.ch/document/ATL-DA-ES-0001/

13-9 A.T. Watson, *Updates to the Level-1 e/gamma and tau/hadron Algorithms*, ATLAS Internal Note, ATL-DAQ-2000-046 (2000)

13-10 E.J.W. Moyse and A.T. Watson, *Performance and Validation of TrigT1Calo, the offline Level-1 Calorimeter Trigger Simulation*, ATLAS Internal Note, ATL-COM-DAQ-2003-010 (2003)

13-11 ATLAS Collaboration, *Muon Spectrometer TDR*, CERN/LHCC/97-22 (1997)

13-12 A. Di Mattia and L. Luminari, *Performances of the Level-1 Trigger System in the ATLAS Muon Spectrometer Barrel*, ATLAS Internal Note, ATL-DAQ-2002-008 (2002) http://atlas.web.cern.ch/Atlas/GROUPS/MUON/layout/muon_layout_P.html

13-13 A. Di Mattia et al., *A muon trigger algorithm for Level-2 feature extraction*, ATLAS Internal Note, ATL-DAQ-2000-036 (2000)

13-14 A. Di Mattia et al., *RPC Trigger Robustness: Status Report*, ATLAS Internal Note, ATL-DAQ-2002-015 (2002)

13-15 S. Veneziano, *Preliminary Design Report of the MUCTPI Electronics*, ATC-RD-ER-0001 () https://edms.cern.ch/document/310152/

13-16 N. Ellis, *Central Trigger Processor URD*, ATL-DA-ES-0003 (1999) https://edms.cern.ch/document/ATL-DA-ES-0003/

13-17 P. Farthouat, CTP Technical Specification, ATL-DA-ES-0006 (1999) https://edms.cern.ch/document/107447

13-18 G. Schuler et al., Central Trigger Processor Demonstrator (CTPD), ATL-DA-ER-0005 (2000) https://edms.cern.ch/document/115660/

13-19 I. Brawn et al., A Demonstrator for the ATLAS Level-1 Central Trigger Processor, ATL-DA-ER-0008 (2000) https://edms.cern.ch/document/249113/

13-20 R. Blair et al., *ATLAS Second Level Trigger Prototype RoI Builder*, ATL-D-ES-0011 (2003) https://edms.cern.ch/document/367638/

13-21 A.G. Mello, S. Armstrong, and S. Brandt, *Region-of-Interest Selection for ATLAS High Level Trigger and offline Software Environments*, ATLAS Internal Note, ATL-COM-DAQ-2003-005, ATL-SOFT-2003-005 (2003)

13-22 S. Armstrong et al., *Requirements for an Inner Detector Event Data Model*, ATLAS Internal Note, ATL-DAQ-2002-011 (2002)

13-23 PESA Core Algorithms Group, S. Armstrong editor, *Algorithms for the ATLAS High Level Trigger*, ATLAS Internal Note, ATL-COM-DAQ-2003-013 (2003)

13-24 C. Leggett and A. Schaffer, ATLAS EDM-DD Workshop, CERN, 2003

13-25 K. Assamagan et al., *Definition of Raw Data Objects for the MDT Chambers of the Muon Spectrometer,* ATLAS Internal Note, ATL-COM-MUON-2003-020 (2003)

13-26 T.A.M. Wijnen, *The MROD Data Format and the towering partioning of the MDT chambers,* ATLAS Internal Note, ATL-COM-MUON-2002-011 (2002)

13-27 P. Bagnaia et al., *Online versus offline identifiers for the MDT chambers of the Muon Spectrometer,* ATLAS Internal Note, ATL-COM-MUON-2003-017 (2003)

13-28 K. Assamagan et al., *Raw Data Object Definition for the RPC chambers of the ATLAS Muon Spectrometer,* ATLAS Internal Note, ATL-COM-MUON-2003-19 (2003)

13-29 G. Aielli et al., *Data Format of the RPC Detector of the Muon System,* ATLAS Internal Note, ATLAS-COM-MUON-2003-018 (2003)

13-30 H. Drevermann and N. Konstantinidis, *Determination of the z position of primary interactions in ATLAS,* ATLAS Internal Note, ATL-DAQ-2002-014 (2002)

13-31 H. Drevermann and N. Konstantinidis, *Hit Filtering and Grouping for Track Reconstruction,* ATLAS Internal Note, ATL-DAQ-2003-XXX (2003), document in preparation

13-32 I. Gaines, S. Gonzalez and S. Qian, *Implementation of an Object Oriented Track Reconstruction Model into Multiple LHC Experiments*, Computing in High Energy and Nuclear Physics Conference, Padova, 2000

13-33 D. Candlin, R. Candlin and S. Qian, *Update of an Object Oriented Track Reconstruction Model for LHC Experiments*, Computing in High Energy and Nuclear Physics Conference, Beijing, 2001

13-34 J. Baines et al., *B-Physics Event Selection for the ATLAS High Level Trigger,* ATLAS Internal Note, ATL-DAQ-2000-031 (2000)

13-35 J. Baines et al., *Global Pattern Recognition in the TRT for B-Physics in the ATLAS Trigger,* ATLAS Internal Note, ATL-DAQ-99-012 (1999)

13-36 M. Sessler and M. Smizanska, *Global Pattern Recognition in the TRT for the ATLAS LVL2 Trigger,* ATLAS Internal Note, ATL-DAQ-98-120 (1998)

13-37 S. Sivoklokov, presentations made in PESA Core Algorithms Group meetings, December 2002 and January 2003

13-38 S. Sivoklokov, *High pT Level-2 Trigger Algortihm for the TRT Detector in ATRIG*, ATLAS Internal Note, ATL-DAQ-2000-043 (2000)

13-39 M.P. Casado, S. González, and T. Shears, *TrigT2Calo package*, http://atlas-sw.cern.ch/cgi-bin/cvsweb.cgi/offline/Trigger/TrigAlgorithms/TrigT2Calo/

13-40 S. González, T. Hansl-Kozanecka and M. Wielers, *Selection of high-pT electromagnetic clusters by the level-2 trigger of ATLAS*, ATLAS Internal Note, ATL-DAQ-2000-002 (2000)

13-41 S. González, B. González Pineiro and T. Shears, *First implementation of calorimeter FEX algorithms in the LVL2 reference software,* ATLAS Internal Note, ATL-DAQ-2000-020 (2000)

13-42 S. González and T. Shears, *Further studies and optimization of the level-2 trigger electron/photon FEX algorithm*, ATLAS Internal Note, ATL-DAQ-2000-042 (2000)

13-43 *ATLAS Level-1 Trigger Technical Design Report*, CERN/LHCC/98-14 (1998)

13-44 A. Di Mattia, S. Falciano and A. Nisati, *The implementation of the muFast algorithm in the new PESA framework*, ATLAS Internal Note, ATL-COM-DAQ-2003-024 (2003)

13-45 I. Gavrilenko, *Description of Global Pattern Recognition Program (xKalman)*, ATLAS Internal Note, ATLAS-INDET-97-165 (1997)
http://maupiti.lbl.gov/atlas/xkal/xkalmanpp/index.en.html.

13-46 R. Clifft and A. Poppleton, *IPATREC: Inner Detector Pattern Recognition and Track Fitting*, ATLAS Internal Note, ATLAS-SOFT-94-009 (1994)

13-47 D. Adams et al., *Track Reconstruction in the ATLAS Muon Spectrometer with MOORE*, ATLAS Internal Note, ATL-COM-MUON-2003-012, ATL-COM-SOFT-2003-008 (2003)

13-48 D. Adams et al., *MOORE as Event Filter in the ATLAS High Level Trigger*, ATLAS Internal Note, ATL-COM-MUON-2003-013, ATL-COM-SOFT-2003-009, ATL-COM-DAQ-2003-012 (2003)

13-49 *ATLAS Detector and Physics Performance Technical Design Report*, CERN/LHCC/99-14 and CERN/LHCC/99-15 (1999)

13-50 R. Mommsen et al., *Performance studies for electron and photon selection at the event filter*, ATLAS Internal Note, ATL-DAQ-2000-007 (2000)

13-51 R. Mommsen, *Electron trigger performance studies for the event filter*, ATLAS Internal Note, ATL-DAQ-2001-004 (2001)

13-52 J. Baines et al., *Performance Studies of the High Level Electron Trigger*, ATLAS Internal Note, ATL-COM-DAQ-2003-0020 (2003)

13-53 J. Baines et al., *First study of the LVL2-EF boundary in the high-p_T e/gamma high-level-trigger*, ATLAS Internal Note, ATL-DAQ-2000-045 (2000)

13-54 T. Sjostrand, Comput. Phys. commun. 82 (1994) 74

13-55 T. Sjostrand et al., Comput. Phys. Commun. 135 (2001) 238

13-56 E. Richter-Was, *Samples from the DC1 production: dataset 2000, dataset 2001, dataset 2003; few items from physics content evaluation*, ATLAS Internal Note, ATL-COM-PHYS-2003-026 (2003)

13-57 S. Tapprogge, ATLAS Rates, Presentation at LHCC Meeting, May 2002,
http://agenda.cern.ch/askArchive.php?base=agenda&categ=a02443&id=a02443s1t2/transparencies

13-58 A. Virchaux et al., *MUONBOX: a full 3D Tracking Programme for Muon Reconstruction in the ATLAS Spectrometer*, ATLAS Internal Note, ATL-MUON-1997-198 (1997)

13-59 GCALOR Package,
http://wswww.physik.uni-mainz.de/zeitnitz/gcalor/gcalor.html

13-60 B. Caron, J. L. Pinfold and R. Soluk, *A Study of the ATLAS Tau/hadron Trigger*, ATLAS Internal Note, ATL-COM-DAQ-2003-030 (2003)

13-61 *ATLAS Trigger Performance Status Report*, CERN/LHCC/98-15 (1998)

13-62 B. G. Piniero, *Tau Identification in the Second Level Trigger*, ATLAS Internal Note, ATL-DAQ-1998-127 (1998)

13-63 D. Cavalli and S. Resconi, *Combined Analysis of $A/H \rightarrow \tau\tau$ Events from Direct and Associated bbA Production*, ATLAS Internal Note, ATL-PHYS-2000-005 (2000)

13-64 D. Cavalli et al., *Search for $A/H \rightarrow \tau\tau$ Decays*, ATLAS Internal Note, ATL-PHYS-94-051 (1994)

13-65 D. Cavalli and S. Resconi, *Tau Jet Separation in the ATLAS Detector*, ATLAS Internal Note, ATL-PHYS-98-118 ()

13-66 D. Cavalli, *Missing Transverse Momentum Reconstruction in ATLAS*, ATLAS Internal Note, ATL-PHYS-96-080 (1996)

13-67 B. Caron et al., *Event Filter Rates for the E_T-miss + jet Trigger at Low Luminosity*, ATLAS Internal Note, ATL-DAQ-2000-016 (2000)

13-68 D. Cavalli et al., *Status Report on the τ-Identification Performance*, ATLAS Physics Workshop, Athens, 2003

13-69 M. Wielers, *Performance Studies of Jets in the High Level Trigger*, ATLAS Internal Note, ATL-DAQ-2000-015 (2000)

13-70 J. de Jong and J. L. Pinfold, *Estimating the ATLAS Multi-jet Trigger Rate Using ATLFAST*, ATLAS Internal Note, ATL-COM-DAQ-2003-029 (2003)

13-71 M. Cervetto et al., *b-tagging Event Selection for the ATLAS High Level Trigger: an update*, ATLAS Internal Note, ATL-COM-DAQ-2003-034 (2003)

13-72 M. Cervetto et al., *SiTrack: a LVL2 track reconstruction algorithm based on Silicon detectors*, ATLAS Internal Note, ATL-COM-DAQ-2003-25 (2003)

13-73 A. Watson and V. Ghete, *A study of the use of Low E_T calorimeter RoI in the ATLAS B-Trigger*, document in preparation

13-74 J. Baines et al., *B-Physics Event Selection for the ATLAS High Level Trigger*, ATLAS Internal Note, ATL-DAQ-2000-031 (2000)

13-75 B. Epp, V.M. Ghete and A. Nairz., *Event Filter Rate for the Ds Trigger*, ATLAS Internal Note, ATL-DAQ-2001-003 (2001)

13-76 J. Baines et al., *Effects of Inner Detector Misalignment and Inefficiency on the ATLAS B-physics Trigger*, ATLAS Internal Note, ATL-DAQ-2001-006 (2001)

13-77 J. Baines et al., *Resource Estimates for the ATLAS B-physics Trigger*, ATLAS Internal Note, ATL-COM-DAQ-2002-001 (2001)