1 Introduction

1.1 Purpose and scope of the TDR

The purpose of this document is to describe the design of the ATLAS level-1 (LVL1) trigger, as well as the plans for its construction, installation and commissioning, and operation. The design is based on an extensive programme of R&D and demonstrator prototyping work performed within the ATLAS, RD12 [1-1], and RD27 [1-2] projects since 1991. Most of the critical elements of the design have already been successfully demonstrated. Where this is not the case, the TDR describes the remaining demonstrator prototyping work required to validate the design. The design builds and improves upon the one presented in the ATLAS Technical Proposal [1-3].

This TDR documents the LVL1 trigger design at the system level, describing functional components at the level of electronic modules (printed circuit boards, multichip modules, application-specific integrated circuits) and their interconnections. The TDR includes a summary of trigger performance results that justify the choices that were made in the design. More details on the trigger performance of the LVL1 trigger, and of the Trigger, DAQ and Event Filter system as a whole, can be found in an accompanying document [1-4]. Another document [1-5] presents a technical progress report and workplan for the rest of the Trigger / DAQ / Event Filter / DCS system.

In order to keep the TDR to a reasonable size, not all of the details are given in this document. Where appropriate, reference is made to ATLAS internal notes and other documents.

1.2 Organization of the TDR

The level-1 trigger TDR is organized as follows. Chapter 2 gives a general description of the LVL1 trigger system, giving an overview of its architecture in the wider context of the ATLAS Trigger/DAQ system as a whole. It explains some of the general requirements on the LVL1 trigger, and discusses some of the more important implementation issues. Chapter 2 can be considered as an ‘executive summary’.

Chapters 3–16 describe the main hardware elements of the LVL1 trigger system: the calorimeter trigger (Chapters 3–8), the muon trigger (Chapters 9–14), the central trigger processor (Chapter 15) and the timing, trigger and control distribution system (Chapter 16). These are followed by a number of chapters that address general issues related to the LVL1 trigger: software (Chapter 17), a summary of the overall trigger latency (Chapter 18), the strategies for setting up the timing of the experiment (Chapter 19) and for handling deadtime (Chapter 20).

The remaining chapters address the issues of installation, access and maintenance (Chapter 21), safety (Chapter 22) and project organization and management (Chapter 23). Chapter 23 also summarizes the schedule and cost.

1.3 References

1-1 *RD12 Status Report*, CERN / LHCC / 97–29, April 1997.

1-2 *RD27 Status Report*, CERN / LHCC / 97–57, October 1997.

1-3 *ATLAS Technical Proposal*, CERN / LHCC / 94–43, December 1994.

1-4 *ATLAS Trigger Performance Status Report*, CERN / LHCC / 98–15, June 1998.

1-5 *ATLAS DAQ, EF, LVL2 and DCS Technical Progress Report*, CERN / LHCC / 98–16, June 1998.

2 General description of the level-1 trigger system

2.1 ATLAS trigger and data-acquisition system overview

The ATLAS trigger and data-acquisition system is based on three levels of online event selection [2-1]. Each trigger level refines the decisions made at the previous level and, where necessary, applies additional selection criteria. Starting from an initial bunch-crossing rate of 40 MHz (interaction rate ~10^9 Hz at a luminosity of 10^{34} $\mathrm{cm}^{-2}\mathrm{s}^{-1}$), the rate of selected events must be reduced to ~100 Hz for permanent storage. While this requires an overall rejection factor of 10^7 against ‘minimum-bias’ processes, excellent efficiency must be retained for the rare new physics, such as Higgs boson decays, that is sought in ATLAS.

Figure 2-1 shows a simplified functional view of the Trigger/DAQ system. In the following, a brief description is given of some of the key aspects of the event-selection process.

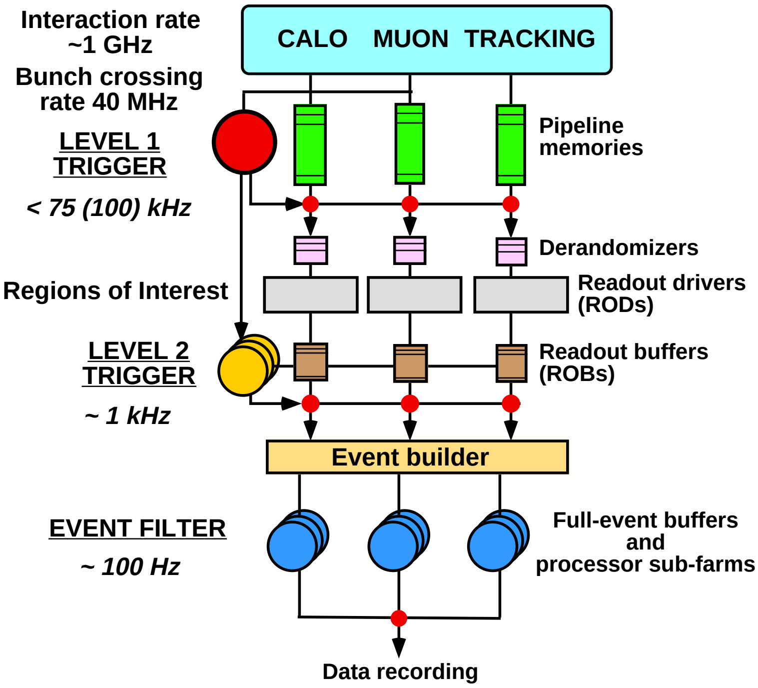

Figure 2-1 Block diagram of the Trigger/DAQ system.

The level-1 (LVL1) trigger described in this TDR makes an initial selection based on reduced-granularity information from a subset of detectors. High transverse-momentum (high-p_T) muons are identified using only the so-called Trigger chambers, resistive-plate chambers (RPCs) in the barrel, and thin-gap chambers (TGCs) in the endcaps [2-2]. The calorimeter selections are based on reduced-granularity information from all the ATLAS calorimeters (electromagnetic and hadronic; barrel, endcap and forward) [2-3], [2-4]. Objects searched for by the calorimeter trigger are high-p_T electrons and photons, jets, and taus decaying into hadrons, as well as large missing and total transverse energy. In the case of the electron/photon and hadron/tau triggers, isolation can be required. Information is available for a number of sets of p_T thresholds (generally 6–8 sets of thresholds per object type).

The missing and total scalar transverse energies used in the trigger are calculated by summing over trigger towers. However, a trigger on the scalar sum of jet transverse energies is also available.

The LVL1 trigger decision is based on combinations of objects required in coincidence or veto. As discussed in an accompanying document [2-5], most of the physics analyses that have been considered by ATLAS can be made using, at the trigger level, fairly simple selection criteria of a rather inclusive nature. However, the trigger implementation is flexible and it can be programmed to select events using more complicated signatures. Examples of LVL1 trigger 'menus' for 'high' (10^{34} cm^{-2}s^{-1}) and 'low' (10^{33} cm^{-2}s^{-1}) luminosity are shown in Tables 2-1 and 2-2. While the simple menus shown cover almost all of the mainstream discovery physics that ATLAS plans to study, in practice it is expected that real trigger menus will be significantly more complicated.

Table 2-1 Example of LVL1 trigger menu ($L = 10^{34}$ cm^{-2}s^{-1}).

Trigger	**Rate (kHz)**
Single muon, $p_T > 20$ GeV	4
Pair of muons, $p_T > 6$ GeV	1
Single isolated EM cluster, $E_T > 30$ GeV	22
Pair of isolated EM clusters, $E_T > 20$ GeV	5
Single jet, $E_T > 290$ GeV	0.2
Three jets, $E_T > 130$ GeV	0.2
Four jets, $E_T > 90$ GeV	0.2
Jet, $E_T > 100$ GeV AND missing $E_T > 100$ GeV	0.5
Tau, $E_T > 60$ GeV AND missing $E_T > 60$ GeV	1
Muon, $p_T > 10$ GeV AND isolated EM cluster, $E_T > 15$ GeV	0.4
Other triggers	5
Total	**~40**

The maximum rate at which the ATLAS front-end systems can accept LVL1 triggers is limited to 75 kHz (upgradable to 100 kHz). As indicated in Tables 2-1 and 2-2 (and discussed in Ref. [2-5]), the rates estimated in trigger performance studies, using trigger menus that meet the needs of the ATLAS physics programme, are about a factor of two below this limit. Given that there are large intrinsic uncertainties in the calculations, this safety factor is not over-generous. However, if necessary, rates could be significantly reduced without major consequences for the physics programme, for example by increasing the thresholds on some of the inclusive (single-object) triggers when operating at the highest luminosities, and by relying more heavily on multi-object triggers.

An essential requirement on the LVL1 trigger is that it should uniquely identify the bunch-crossing of interest. Given the short (25 ns) bunch-crossing interval, this is a non-trivial consideration. In the case of the muon trigger, the physical size of the muon spectrometer implies times-of-flight comparable to the bunch-crossing period. For the calorimeter trigger, a

Table 2-2 Example of LVL1 trigger menu ($L = 10^{33}$ cm^{-2}s^{-1}).

Trigger	**Rate (kHz)**
Single muon, $p_T > 6$ GeV	23
Single isolated EM cluster, $E_T > 20$ GeV	11
Pair of isolated EM clusters, $E_T > 15$ GeV	2
Single jet, $E_T > 180$ GeV	0.2
Three jets, $E_T > 75$ GeV	0.2
Four jets, $E_T > 55$ GeV	0.2
Jet, $E_T > 50$ GeV AND missing $E_T > 50$ GeV	0.4
Tau, $E_T > 20$ GeV AND missing $E_T > 30$ GeV	2
Other triggers	5
Total	~40

serious challenge is that the pulse shape of the calorimeter signals extends over many bunch crossings.

It is important to keep the latency (time taken to form and distribute the trigger decision) to a minimum. During this time information for all detector channels has to be retained in ‘pipeline’ memories. These memories are generally contained in custom integrated circuits, placed on or close to the detector, usually in inaccessible regions and in a high-radiation environment. The total number of detector channels, excluding the pixel detectors, exceeds 10^7. For reasons of cost and reliability, it is desirable to keep the pipeline lengths as short as possible. The LVL1 latency, measured from the time of the proton–proton collision until the trigger decision is available to the front-end electronics, is required to be less than 2.5 μs. In order to achieve this, the LVL1 trigger is implemented as a system of purpose-built hardware processors. The target latency for the LVL1 trigger is 2.0 μs (leaving 500 ns contingency).

Events selected by LVL1 are read out from the front-end electronics systems of the detectors into readout buffers (ROBs); present estimates foresee about 1700 ROBs in total. A large number of front-end electronics channels are multiplexed into each ROB. Intermediate buffers, labelled ‘derandomizers’ in Figure 2-1, average out the high instantaneous data rate at the output of the pipeline memories to match the available input bandwidth of the readout drivers (RODs).

All of the data for the selected bunch crossing from all of the detectors are held in the ROBs either until the event is rejected by the level-2 (LVL2) trigger (in which case the data are discarded) or, in case the event is accepted by LVL2, until the data have been successfully transferred by the DAQ system to storage associated with the Event Filter (which makes the third level of event selection). The process of moving data from the ROBs to the Event Filter (EF) is called event building. Whereas before event building each event is composed of many fragments, with one fragment in each ROB, after event building the full event is stored in a single memory accessible by an EF processor.

The LVL2 trigger makes use of ‘region-of-interest’ (RoI) information provided by the LVL1 trigger. This includes information on the position (η and ϕ, where η is pseudorapidity and ϕ is azimuthal angle) and p_T range of candidate objects (high-p_T muons, electrons/photons, hadrons/taus, jets), and energy sums (missing-E_T vector and scalar E_T value, where E_T is transverse energy). The RoI data are sent by LVL1 to LVL2, for all events selected by the LVL1 trigger, using a dedicated data path. Using the RoI information, the LVL2 trigger selectively accesses data from the ROBs, moving only the data that are required in order to make the LVL2 decision. The LVL2 trigger has access to all of the event data, if necessary with the full precision and granularity. However, typically only data from a small fraction of the detector, corresponding to limited regions centred on the objects indicated by the LVL1 trigger, are needed by the LVL2 trigger. Hence, usually only a few per cent of the full event data are required thanks to the RoI mechanism.

Some of the signatures used at LVL2 to reduce the rate are discussed in the following; the reader is referred to Ref. [2-5] for more details. It is expected that LVL2 will reduce the rate to ~1 kHz. However, in contrast to the 75 kHz (upgradable to 100 kHz) limit for LVL1 that comes from the design of the detector front-end electronics, this is not a hard number. Optimization of the sharing of the selection task between LVL2 and the EF is being studied in the next phase of the project. The latency of the LVL2 trigger is variable from event to event; it is expected to be ~1–10 ms.

In the case of muon triggers, rejection power at LVL2 comes from sharpening (and where necessary raising) the p_T threshold compared to LVL1, and from applying isolation requirements. Sharper p_T thresholds are possible by using the precision muon chambers (MDTs and CSCs) [2-2] and also the inner detector. The isolation requirements use calorimeter information, demanding little E_T in a region around the muon candidate.

For isolated electrons, rejection power at LVL2 comes from using the full-granularity calorimeter information and requiring a matching high-p_T charged track in the inner detector; the transition-radiation signature gives additional rejection power. For photons, less rejection power is possible than in the case of electrons since the inner detector cannot be used. (Given the relatively high probability for photon conversion in ATLAS, it is not planned to use a track veto for the photon trigger.) However, for the important physics channel $H \rightarrow \gamma\gamma$, the trigger can require a pair of photons, with a rejection factor for each γ compared to LVL1 due to the use of high-precision, high-granularity calorimeter information.

For the hadron / tau trigger, rejection at LVL2 is achieved using the full-granularity calorimeter information and the inner detector. A localized, isolated (hadronic) calorimeter cluster with a matching high-p_T track is required.

In the case of jets, much less rejection power is possible. Jets are the dominant high-p_T process at the LHC, and the threshold behaviour of the LVL1 trigger is reasonably sharp. Hence for jet triggers LVL2 must either increase the threshold or make additional requirements in order to significantly reduce the rate. The possibility of identifying b-quark jets at LVL2 using inner-detector information for tagging is under study.

Concerning the energy-sum triggers (missing E_T, total scalar E_T), only limited improvement is possible using the RoI mechanism. The energy-sum values from LVL1 are provided to LVL2 and refinements can be made to correct, e.g., for high-p_T muons (the LVL1 missing-E_T trigger uses only calorimeter information, so muons contribute to the observed missing E_T) or for saturated trigger-tower signals. The possibility of performing a full missing-E_T recalculation at LVL2 for a small subset of events remaining after other LVL2 selection criteria have already been applied is being investigated.

The LVL1 trigger makes available RoI information for all of the objects that contributed to the event being selected; these are called primary RoIs. Furthermore, in order to allow additional requirements to be made at LVL2, the LVL1 trigger can provide RoI information for objects that did not contribute to the selection of the event. Such RoIs, typically for objects of relatively low p_T, are called secondary RoIs.

After LVL2, a last stage of selection is performed in the EF. Here the algorithms will be based on offline code. The EF must reduce the rate to a level suitable for permanent storage, currently assumed to be ~100 Hz for full events of size ~1 Mbyte.

The status of work on the LVL2 trigger, DAQ, Event Filter and Detector Control systems, together with plans for work to be completed before submission of the Technical Proposal for these subsystems, is documented in a Technical Progress Report [2-6]. Details of the trigger performance studies for LVL2, as well as for LVL1, are also documented [2-5].

An important consideration is to define interfaces with other systems sufficiently early to avoid wasting effort on incompatible designs. The interface between the ATLAS front-end systems and the Trigger/DAQ system is specified in a requirements document [2-7] that has been endorsed by all systems. This defines (at a functional level) the interaction between the LVL1 trigger and the detector systems, including issues of trigger latency, rate capability, and permitted levels of data loss and deadtime. It does not cover the special cases of the provision of trigger signals from the muon Trigger chambers and from the calorimeters to the LVL1 trigger; these are addressed in this document and also in the corresponding detector TDRs that have already been approved, and in User Requirements Documents (URDs) [2-8], [2-9]. The requirements document [2-7] also defines the interaction and sharing of responsibilities between the detector systems on the one hand, and the LVL2 trigger and DAQ on the other hand, at the level of the ROBs. The interface between the LVL1 and LVL2 triggers is discussed in this TDR and also in Ref. [2-6].

2.2 Overview of level-1 trigger system

Figure 2-2 is a context diagram for the LVL1 trigger system, showing the interfaces to other systems. Input data/signals are received from the muon trigger chambers and from the calorimeters. The trigger decision has to be made available to the front end and readout systems for all the ATLAS subsystems — as shown, this is done using the Timing, Trigger and Control (TTC) distribution system (see Chapter 16). Additional information that can be sent over the TTC system are the event identifier number, the bunch-crossing number, and a trigger word that summarizes why the event was selected by the LVL1 system.

The LVL1 trigger is a system of synchronous, pipelined processors running at 40 MHz or multiples thereof. The clock signals used to drive the processing are derived from the LHC machine clock. Another signal received from the machine is a bunch-zero signal, sent once per turn of the machine (period 88 μs). This allows one to make an absolute identification of bunches within the machine.

The LVL1 trigger system can accept externally-generated trigger signals. These include calibration trigger signals associated, for example, with test-pulse generation in the calorimeters. Various other facilities included are related to calibration procedures, including the possibility to fire test pulses at well-defined points in the LHC cycle, and the arbitration of the use by detector systems of gaps in the LHC bunch structure for calibration triggers.

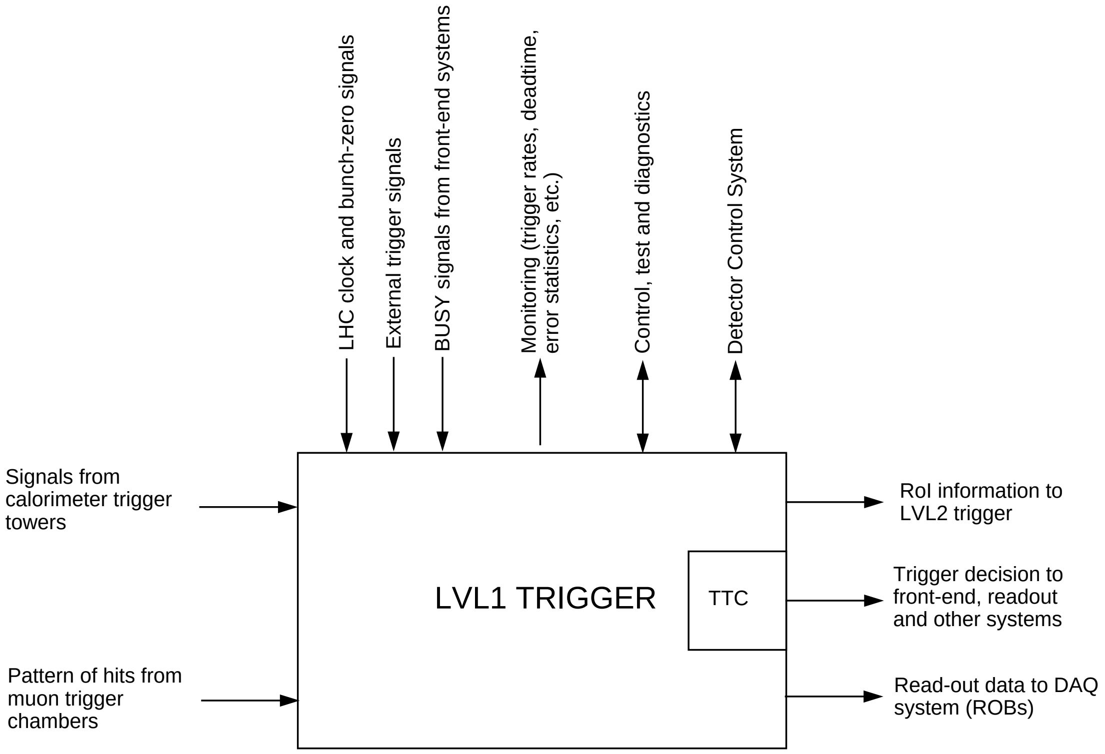

Figure 2-2 Context diagram for the LVL1 trigger system.

The LVL1 system receives inputs ('BUSY' signals) that can impose deadtime, for example if buffers become full in the readout systems. These are also used during initialization to prevent triggers being generated until the readout systems are ready.

The LVL2 trigger receives the accept/reject signal from LVL1. For accepted events, it receives in addition information on the candidate objects found by the LVL1 trigger which it uses to drive the RoI processing. More detailed information from LVL1, including input data, and intermediate and final results of the trigger processing, are stored in ROBs. These data are read out by the DAQ system for events that are selected by LVL2; the data may also be selectively accessed by the LVL2 trigger, as for detector data.

In addition, the LVL1 trigger is connected to the run control (an element of the DAQ system) and to the Detector Control System (DCS). The run control infrastructure is used for testing, monitoring and fault diagnosis, for running calibration programs, and for configuring the trigger system (e.g. loading all the parameters that control the operation of the trigger and specify the selection criteria). The run control system will be used to acquire data that are not associated with individual events, such as scaler information needed to monitor rates and deadtime. The functions of the DCS include control and monitoring of the electronics racks and crates (voltages, temperature, etc.).

As shown in Figure 2-3, the LVL1 trigger system is composed of a number of building blocks — the calorimeter trigger, the muon trigger, the Central Trigger Processor (CTP) and the Timing, Trigger and Control (TTC) system. These subsystems are functionally well-defined, and the interaction between them is limited. The interfaces are presently defined at the functional level, and their detailed physical implementation will be specified in the next phase of the project (this is considered to be straightforward). The compatibility of the different subsystems will be checked, prior to the start of construction, within the context of the ATLAS review process (i.e.

via Production Readiness Reviews). The functions of each of the subsystems are described in a bit more detail in the following, considering only the LVL1 trigger function *per se*.

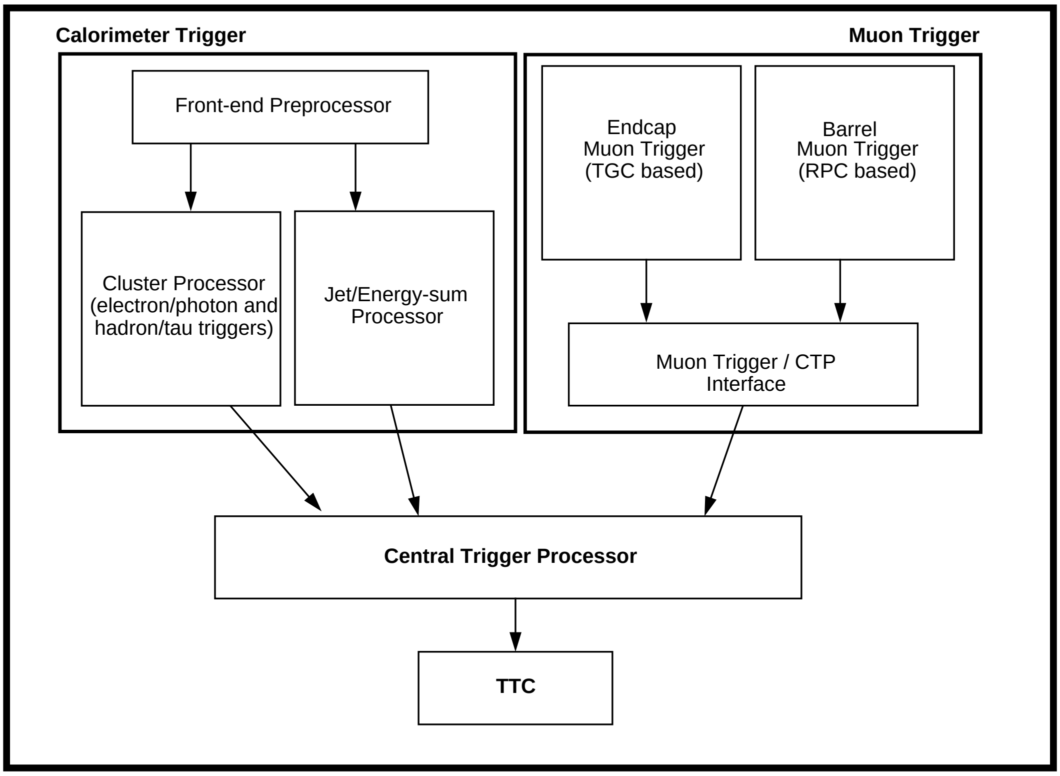

Figure 2-3 Block diagram of the LVL1 trigger system.

2.2.1 Calorimeter trigger

The inputs to the calorimeter trigger system are analogue signals from trigger towers that have a typical granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ (the granularity becomes coarser beyond $|\eta| \approx 2.5$). Separate tower signals are received from the electromagnetic and hadronic calorimeters, corresponding to a total of ≈ 7200 analogue input signals to the trigger system. The trigger towers are formed in the calorimeter front-end electronics by analogue summation over the corresponding calorimeter cells. Transmission to the calorimeter trigger system, located in USA15, is on individually-shielded twisted-pair cables.

In the calorimeter trigger system, the trigger-tower signals are digitized using a dedicated ADC system. Digital signal processing is applied to extract the E_T for calorimeter pulses and to assign it to the correct bunch crossing, since the shaped pulses from the calorimeters extend over several bunch-crossing periods. This part of the calorimeter trigger system is labelled ‘front-end Preprocessor’ in the figure.

The subsequent calorimeter-trigger processing is fully digital, and is divided into two parts. One part (labelled ‘Cluster Processor’ in the figure) performs a search for high-p_T electrons / photons and hadrons / taus using the full-granularity trigger-tower information from the Preprocessor. The other part (labelled ‘Jet / Energy-sum Processor’ in the figure) searches for high-E_T jets and calculates the missing-E_T and total scalar-E_T values.

For the electron / photon trigger, there are eight sets of thresholds that can be programmed independently; each set consists of a threshold on the E_T of the cluster, an isolation threshold on the surrounding E_T in the electromagnetic calorimeter, and a ‘hadron-veto’ threshold on the E_T in the associated hadron-calorimeter towers. Similarly, for the hadron / tau trigger, there are eight sets of thresholds that can be programmed independently; each set consists of a threshold on the E_T of the cluster, and isolation thresholds for the surrounding E_T in the electromagnetic and hadronic calorimeters. For the jet trigger there are eight thresholds that can be programmed independently and to which the E_T in ‘jet windows’ is compared. The jet-window size is also programmable.

Summation is performed over the trigger towers to calculate the missing-E_T vector and the total scalar E_T for the event. Thresholding is performed, with eight threshold values for the missing-E_T trigger (thresholds applied to modulus of missing-E_T vector), and four for the scalar-E_T trigger.

For each object type, information is sent to the CTP for each bunch crossing. In the case of the electron / photon, hadron / tau and jet triggers, this is the multiplicity of objects for each set of thresholds. For the energy-sum triggers, information is sent indicating which thresholds have been passed.

2.2.2 Muon trigger

The muon trigger receives as input the pattern of hit strips (and wire groups in the case of the TGC detectors) in the muon Trigger chambers. These data are produced by amplifier-shaper-discriminator circuits in the RPC and TGC front-end electronics. There are a total of more than 800k input signals to the muon trigger system. The timing resolution is sufficiently good that the trigger can, with very high probability, identify the bunch crossing that contained the muon. The trigger searches for patterns of hits consistent with high-p_T muons originating from the interaction region. The logic provides six independently-programmable p_T thresholds. The output sent to the CTP for each bunch crossing is the multiplicity of muons for each of the six p_T thresholds (i.e. six multiplicity values).

As indicated in Figure 2-3, the muon trigger system is subdivided into a part specific to the RPC detectors, a part specific to the TGC detectors, and a part that combines information from the full system and prepares the input to the CTP. Some of the RPC- and TGC-specific logic is mounted on or near the detectors in the experimental cavern; the rest is located in the underground shielded counting room, USA15.

2.2.3 Central Trigger Processor

The role of the CTP is to combine the information for the different object types and to make the overall LVL1 accept / reject decision. Trigger menus can be programmed with up to 96 items, each item being a combination of requirements on the input data. The requirements can be very simple, demanding for example at least one muon for a given p_T threshold, or more complicated, for example requiring at least one electron above a specified p_T threshold, missing E_T above another specified threshold, and no high-p_T muons. The CTP can also combine the information on jet multiplicity versus threshold to estimate the jet-E_T sum, to which it can apply a threshold. The overall LVL1 decision will be to accept an event if any of the 96 menu items is satisfied, subject to prescale and deadtime considerations discussed below.

The CTP allows a prescale factor to be programmed individually for each of the menu items. This facility will be used to down-scale high-rate triggers, for example allowing data to be collected concurrently for jet production over a large E_T range; this is useful for QCD studies and also for monitoring the detector performance. The CTP is also responsible for deadtime control. Deadtime may be generated externally via a busy signal, for example if the LVL2 trigger is close to saturation. It will also be generated internally using algorithms that predict when overflow conditions are likely to occur in front-end systems. The simplest element of the deadtime logic is to prevent the occurrence of two triggers separated by fewer than five bunch crossings (125 ns). The CTP also contains a system of scalers used to monitor rates and deadtime.

2.2.4 Timing, Trigger and Control distribution system

The TTC system is responsible for distributing to the front-end systems a number of signals, including the LHC clock and the LVL1 trigger decision. The TTC backbone, which is described in this TDR, is based on the optical-broadcast system developed in the RD12 Collaboration [2-10]. The TTC system receives the LVL1 accept signal from the CTP and the LHC clock and bunch-zero signals from the machine. These and other signals are encoded and transmitted optically to the front-end systems, or to intermediate points at which a change is made to detector-specific protocols for TTC distribution.

2.3 Summary of requirements for the level-1 trigger

Detailed evaluations have been made of the requirements for the different parts of the LVL1 trigger system — the calorimeter trigger [2-9], the muon trigger [2-8], the CTP [2-11] and the TTC system [2-12]. A brief and mainly qualitative summary of some of the more important general requirements is given here. More details and quantitative requirements can be found in subsequent chapters of this TDR, and in the requirements documents for each of the subsystems.

2.3.1 General requirements

The primary function of the LVL1 trigger is to provide, for each bunch crossing, a signal specifying if the bunch crossing should be retained for further analysis, typically because it contains a potentially interesting physics signature. As discussed above, the decision is based on input signals received from the calorimeters and muon Trigger chambers. The trigger decision has to be distributed, along with other signals, to the front-end electronics of the detector systems. A secondary function of the LVL1 trigger is to provide RoI information to guide the LVL2 trigger.

2.3.2 Physics requirements

The physics performance of the trigger is discussed in Chapters 4 and 14 of this TDR for the LVL1 calorimeter and muon triggers, respectively, and in more detail in Ref. [2-5] which also addresses the LVL2 trigger. As stated above, the objects upon which the LVL1 trigger is based are high-p_T muons, electrons/photons, hadrons/taus and jets, and large missing and total scalar E_T. Considerations include the p_T range over which the trigger must be able to operate, and the required acceptance and efficiency to find the objects, taking into account constraints on the acceptable trigger rate. Other issues are the ability to resolve nearby objects while not double-counting single objects; this is important for multi-object triggers such as the di-electron and di-muon triggers. Also important is the number of thresholds that can be concurrently used; typically different thresholds are used for single and multi-object triggers, and further thresholds are needed for prescaled triggers.

An indication of the required range of threshold settings can be obtained by considering the numerous studies that have been made to assess the physics capabilities of ATLAS. It turns out that the mainstream discovery physics is covered by fairly simple LVL1 trigger menus based on rather inclusive signatures, as shown in Tables 2-1 and 2-2. In practice, much more extensive menus will be used, including triggers for specialized physics studies (including Standard Model physics), and to collect samples for calibration and monitoring of the detectors and physics control samples (e.g. for background evaluation studies). However, for these additional menu items, often only limited statistics will be required, and prescale factors can be applied to limit the rate to acceptable levels.

The geometrical acceptance of the LVL1 trigger is driven by the design of the detector, where precision measurements in the calorimeters and the coverage of the inner detector are limited to pseudorapidity, $|\eta| < 2.5$. The muon, electron / photon and hadron / tau triggers are required to cover this range ($|\eta| < 2.4$ in the case of the muon trigger). For the jet trigger, the calorimeter trigger towers that are used extend up to $|\eta| < 3.2$, the edge of the endcap calorimeters. The missing and total scalar E_T triggers use all the calorimeters, giving a coverage of $|\eta| < 4.9$.

The LVL1 trigger has to take into consideration the need to provide RoI information to LVL2, with thresholds for flagging objects lower in p_T than those used actively in making the LVL1 trigger decision. The exact threshold requirements here will depend on the scope of the LVL2 trigger (compared to the Event Filter). Also important for LVL2 is the accuracy with which the position of the object within the detector can be specified by LVL1. More details on the requirements on the LVL1 trigger from LVL2 can be found in Ref. [2-13].

The trigger conditions (e.g. thresholds and multiplicity requirements) must be programmable to adapt to different luminosity conditions and changing physics requirements. There must be sufficient flexibility to cope with unforeseen background conditions or new physics.

Beyond the normal physics triggers, consideration has to be given to handling various kinds of special triggers needed for calibration and monitoring, and to understand the detector and background conditions. These include calibration triggers, such as test-pulse triggers associated with the different detector systems. They also include special triggers on cosmic rays and beam-halo particles (implemented using the LVL1 muon trigger system). In addition, provision has to be made for random triggers and for prescaling high-rate triggers as discussed above.

2.3.3 Measurement of trigger acceptance and efficiency

It will be important for physics studies to correct for acceptance and efficiency losses of the trigger. In contrast to the situation in e^+e^- machines such as LEP, where the interaction rates are relatively small and where all high-p_T events can therefore be accepted, the triggers at the LHC have to be extremely selective. In general, one cannot rely on a high level of redundancy in the trigger, with signal events being selected on the basis of multiple independent signatures.

(Some channels, such as H → $l^+l^-l^+l^-$, would be selected by multiple independent signatures, but this is generally not the case.)

Despite the limitations imposed by the difficult environment of a hadronic machine, it is necessary to be able to measure the trigger efficiency. This can be achieved by determining the efficiency to trigger on objects as functions of p_T, η and ϕ. For multi-object triggers, these efficiencies can then be combined. Object efficiencies can be determined by using large samples of two-object events, where either object could have triggered. Such samples can, if necessary, be collected using prescaled single-object triggers with low p_T thresholds.

2.3.4 Rate and background rejection

The design of the front-end electronics for the detector systems imposes constraints on the rate of LVL1 triggers that can be accepted. The average rate must not exceed 75 kHz (upgradable to 100 kHz). Furthermore, the interval between successive LVL1 triggers must be at least five bunch crossings (125 ns), corresponding to 1% deadtime. As discussed later in this document in Chapter 20, more complicated deadtime logic is provided that allows for high- and low-priority triggers.

The design of the trigger must take into account the need to reduce the rate of triggers due to background processes to acceptable levels, while maintaining the required high efficiency for signal processes that are to be analysed. Adequate safety margins are required to allow for the large uncertainties in the rate calculations (typically of a factor of two or more).

The dominant background in the LVL1 trigger often comes from high-p_T jet production (except for the jet trigger where this is the signal). For example, jets may fake isolated electrons / photons if they fragment into leading π^0 particles. Similarly, background to the muon trigger comes from decays in flight of high-p_T charged pions and kaons. In the case of the muon trigger, the background radiation in the cavern also requires careful attention. In addition, possible false triggers due to cosmic rays and beam-halo particles have to be considered.

Also relevant to the discussion on rates is the threshold sharpness. Given that most of the processes that dominate the LVL1 rate, such as high-p_T jets and muons from b-quark decays, have steeply falling p_T distributions, it is desirable to make a sharp cut on p_T. Otherwise the rate may be dominated by objects that, with more detailed analysis using the full detector data, are found to have p_T below the nominal threshold.

2.3.5 Latency

A requirement that was already mentioned is that the latency should not exceed 2.5 μs. All the ATLAS front-end electronics systems are being designed and implemented with pipeline lengths sufficient to accommodate this latency. In order to keep some contingency, the target latency in designing the LVL1 trigger system is 2.0 μs. This contingency is needed to allow for unforeseen steps in the processing pipeline that may emerge in the detailed design, and for remaining uncertainties in the routing of cables carrying signals to or from the trigger. Note that care has been taken in specifying the latency requirements (see Ref. [2-7]) to allow for detector-specific treatment of TTC signals that may, in some cases, add extra latency but bring overall benefits.

2.3.6 Bunch-crossing identification

Another requirement that was already mentioned is the need for the trigger to uniquely identify the bunch crossing of interest. This is achieved for all the triggers that are considered; the probability to trigger on the wrong bunch crossing is expected to be negligibly small.

2.3.7 Timing calibration

An important consideration for the LVL1 trigger design (and also for all ATLAS front-end electronics systems) is the need to set up the timing under various operating conditions. The short bunch-crossing period, the large physical size of the apparatus, the huge number of electronic channels and units, and the inaccessibility of much of the electronics make this very challenging. Well thought-out strategies for setting up the timing need to be developed at an early stage so that the necessary tools can be included in the system designs.

It must be possible to set up the timing of the experiment for physics running and also for special runs, for example with cosmic rays or beam-halo particles. The timing also has to be set up for test and calibration runs with test pulses. These issues are discussed in more detail in Chapter 19.

2.3.8 Processing and transmission errors

False triggers may be caused by processing or transmission errors within the LVL1 system. For example, a single high-order bit being wrongly set in the E_T value for an electromagnetic-calorimeter trigger tower, due to a transmission error, would very likely result in a false electron / photon trigger. The rate of such false triggers is required to be much smaller than the rate of true triggers. Similarly, the loss of triggers due to processing or transmission errors is required to be very small.

An analysis of the effects of processing and transmission errors is required. Where necessary, error detection and protection are included in the design of the system. Also included in the design is the ability to disable parts of the system that have become faulty. This facility can be used to allow continued operation of the trigger, typically with only a small degradation in performance, until repairs can be made.

2.3.9 Operation of equipment in the experiment cavern

The special requirements for equipment that is installed in the ATLAS cavern are addressed. These include the need for radiation tolerance and for operation in a significant magnetic field, and constraints on the amount of heat that can be dissipated. They also include issues of reliability and maintainability for electronics that cannot be accessed easily (and not at all during running). Note that the only equipment in the cavern under the responsibility of the LVL1 trigger project is the on-detector part of the muon-trigger electronics. The calorimeter groups are responsible for sending to the USA15 shielded counting-room the analogue signals for the trigger towers.

2.3.10 Other requirements on the LVL1 trigger

There are numerous other requirements on the LVL1 trigger relating to commissioning and operating it *in situ*, including its connections to other parts of the Trigger / DAQ / DCS system. These are documented in the various requirements documents referred to above, and include the following:

- The system must read out the input data upon which the trigger decision was based, intermediate results, and the final trigger decision. This is required for use in monitoring the operation of the trigger and the quality of the input data to it, and is also used in trigger efficiency determination. The readout system has to respect the same requirements as for detector front-end systems (pipelined readout, etc.).
- A subset of the data that are read out has to be passed, via a separate path, to the LVL2 trigger. These data include information on all the candidate objects used in the RoI mechanism, as well as global quantities such as the missing-E_T vector and the total scalar-E_T value, and trigger words that specify why the event was selected by the LVL1 trigger.
- An extensive system (hardware and software) is required for control. This must be done in an integrated way across the different parts of the LVL1 trigger and more generally with the rest of the Trigger / DAQ system. A record of the trigger configuration, including details of selection criteria, calibration constants used, and other parameters, must be maintained. The time required to set up the trigger system or to change the trigger conditions must not be too long, so as to minimize loss of data.
- Provision has to be made for calibration of the system. Generally this requires interacting with the detector groups, for example for the use of test-pulse systems.
- Extensive capability must be built into the system for monitoring so that faults can be detected quickly. Online monitoring software can make extensive checks using the data that are read out from the trigger system. Scalers are required to monitor trigger rates and deadtime; provision has to be made to monitor some rates for individual bunches in the machine. The scaler data have to be read out separately from the event data. Finally, monitoring (using the DCS) is required of voltage levels, currents, temperatures, etc.
- The system must be designed with built-in test facilities to allow efficient and rapid detection and diagnosis of faults. This will be important both in the commissioning phase and during maintenance and repair operations.
- The failure rate and time required to perform repairs must be such that the resulting loss of data to the experiment is acceptably small. The system must be designed to be robust against localized failures. For example, it must be possible to prevent defective detector channels (noisy calorimeter cells or muon chamber strips) from seriously degrading the performance of the trigger.
- Maintenance requirements have been considered. Sufficient spares must be available taking into consideration the long lifetime of the experiment, and provision must be made to have adequate experts at CERN when the experiment is running.
- Standard safety requirements must be followed, for example concerning fire safety of cables.
- The ATLAS standard grounding rules should be followed where applicable.

2.3.11 Requirements on external systems

The requirements analyses that have been performed also identified numerous requirements on external systems, such as the calorimeter and muon detectors that provide the input signals, and also the DAQ and DCS systems that provide the infrastructure used to control, monitor, and read out the LVL1 system.

The requirements on the calorimeter and muon systems are discussed in detail in Chapters 3–8 and Chapters 9–14 of this TDR, respectively. These include a specification of the boundary of responsibility between the trigger and the detector systems, and a definition of the interface (including issues of detector granularity). Requirements are also specified on the properties of the input signals to the trigger system as discussed in the following paragraphs.

In the case of the calorimeters, requirements on the input signals cover issues of dynamic range, linearity, pulse shape, amplitude / transverse-energy calibration and timing, from the points of view of normalization, uniformity between cells, and stability. The requirements also address the issues of noise in and crosstalk between trigger towers, as well as the need to be able to remove individual ‘noisy’ cells from the trigger towers. It is required that the test-pulse systems of the calorimeter be available to generate test signals at the input to the calorimeter trigger for diagnostic and calibration purposes.

Requirements on the muon detector system include issues of optimizing the layout and granularity of the Trigger chambers taking into account the needs of the trigger, and of making sure that the absolute and relative positioning and alignment of the chambers is sufficiently precise that the performance of the trigger is not degraded. Other requirements relate to the timing uniformity of the input signals. It is also required that the test-pulse system of the Trigger chambers be available to generate test patterns at the input to the muon trigger.

2.4 Implementation

In the following chapters of this TDR, a detailed description is given of the LVL1 trigger architecture and the plans for implementing it, as well as results from an extensive R&D programme that has demonstrated the feasibility of the key elements of the design presented.

As discussed in Section 2.2 above, the LVL1 system is partitioned into the calorimeter trigger, the muon trigger, and the CTP and TTC systems (Figure 2-3). The following sections discuss briefly some of the implementation issues associated with these subsystems.

The LVL1 trigger system presents enormous challenges in terms of the required processing power and data movement capacities. Only a system of custom-electronics processors, with ‘hard-wired’ algorithms, can perform the necessary calculations within the allowed 2 µs latency, with new data arriving every 25 ns. The very high data rates within the trigger system, for example a total of about 2000 Gbit / s between the calorimeter trigger front-end Preprocessor and Cluster Processor subsystems, require the use of advanced, high-speed links.

The LVL1 system is implemented as a synchronous, pipelined, parallel processor driven by the 40 MHz LHC machine clock; the logic is clocked at 40 MHz or multiples thereof. At the start of the calculation many input channels are processed in parallel (e.g. data from all ~7200 calorimeter trigger towers); at the end the result is a single bit signalling if the event is to be accepted or not. Every clock tick one step in the calculation is performed and the results are passed on to the next step. While 2 µs corresponds to a total of 80 25-ns clock ticks, an appreciable fraction of the available time is used for signal transmission on cables and optical fibres.

The processor designs make extensive use of application-specific integrated circuits (ASICs). These are employed where large numbers of units, high density and / or ultimate performance are required, justifying the development costs. Elsewhere, use is made of configurable logic (e.g. field-programmable gate arrays) and fast memories. For the on-detector part of the muon-trigger electronics, ASICs have the advantage that it is easier than with commercial off-the-shelf components to ensure that the required radiation tolerance will be achieved.

High-speed data transmission is required in many parts of the system. Within both the calorimeter and muon trigger systems large numbers of high-speed serial links are needed. For long links, such as those between the on-detector and off-detector parts of the muon trigger (length up to 80 m), optical links are used. For short links, such as those between different parts of the calorimeter trigger processor (length ~5 m), cheaper electrical links are used. For inter-module communication within electronics crates, very high-density, high-speed custom backplanes are required. For example, the calorimeter trigger Cluster Processor uses backplanes with >300 signal connections per board, each working at 160 Mbits / s.

A detailed system-level design exists for all parts of the LVL1 trigger; in many areas more detailed designs already exist. Extensive demonstrator prototyping work has been performed over a period of several years, including evaluations of prototypes for the muon, calorimeter and central trigger processors using signals from prototype detectors in test beams. In addition to proving the critical elements of the design, the demonstrator programme has allowed the LVL1 trigger community to acquire the necessary expertise and experience to build the final system.

2.4.1 Calorimeter trigger

The calorimeter trigger system is partitioned into three major parts as shown in Figure 2-3: the front-end Preprocessor, the Cluster Processor, and the Jet / Energy-sum Processor. Each of these subsystems consists of several crates of custom electronics — eight for the front-end Preprocessor, six for the Cluster Processor and five for the Jet / Energy-sum Processor. The full system is rack-based and is located in the underground counting room USA15 that is shielded against radiation.

2.4.1.1 Front-end preprocessor

The Preprocessor receives the analogue signals from the electromagnetic and hadronic calorimeters for ~7200 trigger towers, with a typical granularity of 0.1×0.1 in pseudorapidity–azimuth space. (Analogue summing is performed within the calorimeter front-end electronics to form the trigger tower signals from the finer granularity cell signals.) The input signals are digitized at 40 MHz rate using fast 10-bit analogue-to-digital converters (ADCs), with a least count corresponding to transverse energy $E_T \sim 0.25$ GeV. Digital signal processing is applied to the ADC data to perform bunch-crossing identification (BCID), apply a tower threshold and extract E_T for hit towers; the E_T value is set to zero unless the BCID assigns the pulse to the bunch-crossing under consideration.

The tower threshold, used to reduce the effects of electronic and pile-up noise, together with pedestal subtraction and the final calibration in transverse-energy units, are applied in a lookup table. The 10-bit data from the ADC and BCID logic are reduced to eight bits at the output of the lookup table, giving E_T in units of ~1 GeV. These data are transmitted to the Cluster Processor that searches for isolated high-p_T electrons and photons (isolated electromagnetic clusters), and hadrons and taus (isolated hadronic clusters). As discussed above, the total rate for transmitting these data is ~2000 Gbits / s. The system uses ~2000 serial data links operating in parallel; since the links are short, the signals are sent electrically.

The Preprocessor performs some additional processing for the Jet / Energy-sum Processor. Groups of 2×2 tower 8-bit E_T values are summed, separately for the electromagnetic and hadronic calorimeters. The resulting data are rounded down to nine bits (full-scale $E_T \sim 512$ GeV) before transmission using ~1000 serial data links similar to the ones used for transmission to the Cluster Processor.

2.4.1.2 Cluster Processor

The Cluster Processor acts on input data that can be viewed as an array of 50×64 E_T values in each of the electromagnetic and hadronic calorimeters. The algorithms for both the electron / photon trigger and the hadron / tau trigger search for isolated clusters using information from 4×4 trigger-tower windows (16 towers from each of the electromagnetic and hadronic calorimeters). A search is made for clusters for all possible window positions (so-called overlapping, sliding windows). As a consequence, each tower is used in the calculation for 16 different windows. This has very important implications for the design since data have to be fanned out between processing elements that sometimes are in different ASICs, boards or even crates.

The data fan-out task is addressed at several levels in different ways. Firstly, the mapping of the calorimeter onto the electronic modules of the Cluster Processor has be optimized to simplify as far as possible the sharing of data between modules. Some data have to be shared between crates, in which case the fan-out is done by duplicating the signals at the output of the Preprocessor. Within crates, data have to be shared between at most two electronic modules — this is done on a custom 160-Mbit / s backplane.

A fundamental ingredient to limiting data sharing between modules is to perform a lot of processing on each module. This is achieved using large 9U printed circuit boards with high density input / output and processing logic. High-density input using high-speed serial links is achieved by mounting commercial link receivers in die form on multi-chip modules (MCMs). High-density processing results from implementing in a single ASIC the logic required to process eight windows for both the electron / photon and tau / hadron triggers.

The results from the window processing are combined within the Cluster Processor subsystem to calculate multiplicity values for electrons / photons and tau / hadrons, with eight threshold sets per cluster type. The resulting 16 multiplicity values, rounded down to three bits per value, are passed to the CTP.

2.4.1.3 Jet/Energy-sum Processor

The Jet / Energy-sum Processor receives an array of ~30×32 9-bit E_T values in each of the electromagnetic and hadronic calorimeters from the Preprocessor. In contrast to electrons, photons, and taus and hadrons, a jet is not a well-defined object. Different jet algorithms can be considered, both at the analysis level and at the trigger level. An important consideration is the window size used to search for jets. Here, the optimum choice will depend on many factors — the jet E_T of interest, the luminosity (level of pile-up within the window), and the need to resolve nearby jets in multi-jet events. Since there is no clear best choice, and to provide maximum flexibility to react to conditions at the LHC, the jet-processor design has a programmable window size. A choice can be made (for each set of thresholds) between a window size of 2×2, 3×3 and 4×4 jet elements, where each jet element is 2×2 trigger towers (i.e. each jet element has a size of $\sim 0.2 \times 0.2$ in pseudorapidity–azimuth space. In each case, the windows slide in steps of one jet element in each direction and hence are overlapped.

Many of the comments that were made in the context of the Cluster Processor also apply here. In particular, the fan-out of data is a major issue that is addressed in a similar way in both systems (high-density input/output electronics, optimized mapping of detector onto electronic modules). However, instead of using ASICs for the window processing, here FPGAs are used since the number of chips needed does not justify the cost of developing an ASIC.

The results from the window processing are combined within the Jet/Energy-sum Processor to calculate multiplicity values for jets with eight thresholds. The resulting eight multiplicity values, rounded down to three bits per value, are passed to the CTP.

The other function of the Jet/Energy-sum Processor is to calculate the missing E_T vector and total scalar E_T value for the full event. This is done by summing the E_T values over all of the jet elements and the forward calorimeters. In the case of the missing-E_T calculation, the vector energy components are calculated from the E_T values, using lookup tables to multiply by $\sin(\phi)$ and $\cos(\phi)$. After summation of E_x and E_y separately, a look-up table is used to compute the scalar missing-E_T value.

The missing-E_T value is compared with eight thresholds and the resulting eight-bit comparator pattern is transmitted to the CTP. The same is done for the scalar-E_T value, except that there only four thresholds are provided.

2.4.1.4 Demonstrator programme

The critical aspects of the design of the calorimeter trigger processor have been addressed in a demonstrator programme. This has shown the viability of implementing trigger algorithms in ASICs, and has demonstrated that a high density of high-speed serial links can be achieved by mounting commercial link components on multi-chip modules. The use of custom backplanes using single-ended signals at 160 Mbit/s per link for inter-module communication has also been validated.

Many specific evaluations have been done in the laboratory, for example detailed measurements of error rates for data transmission between modules. In addition, a ‘full-slice’ demonstrator prototype system has also been constructed, containing prototype ADC and front-end modules, cluster-processor modules, and a jet-processor module, as well as the CTP demonstrator module (see below). This has been successfully operated together with ATLAS prototype calorimeters at the test beam.

2.4.2 Muon trigger

As shown in Figure 2-3, the muon trigger is divided into three parts — one part associated with the RPC (barrel) detector, a second part associated with the TGC (endcap) detector and the third part that combines information from the two systems and forms the interface to the CTP.

For both the RPC and TGC based subsystems, part of the electronics is mounted on or close to the detectors. This is necessary for the logic that performs operations that are local to groups of chambers. However, as much as possible of the subsequent logic is placed in the USA15 counting room which is shielded against radiation. The on-detector logic will be radiation tolerant.

Although the algorithms used in the barrel and endcap are conceptually similar, there are significant differences, and different electronics is used in the two cases. Reasons for this include the following:

- The properties and sizes of the RPC and TGC detectors are different. In the case of the RPCs, the spread in time of signals at the output of the discriminator is dominated by the range of propagation delays along the strips, depending on the position of the muon in the chamber. As a consequence, there is a correlation in the timing of hits in different detector layers. In contrast, the timing spread for the TGCs is limited by the response of the detectors.
- The magnetic field map in the endcap is complex due to the interference between the fields generated by the barrel and endcap toroids. This means that there is not a clear separation between the ‘bending’ and ‘non-bending’ projections in the endcap system. In contrast to the barrel, information on bending in the two projections has to be combined before making a classification in p_T.

In both the barrel and the endcap the logic is based on coincidence-matrix ASICs that will be implemented using deep-submicron CMOS technology. These matrices search for patterns of hits in successive detector planes consistent with high-p_T muons originating from the interaction region. In each projection, one matrix implements the low-p_T trigger based on two planes of doublet chambers. A second matrix implements the high-p_T trigger using the output of the low-p_T matrix and an additional plane of chambers.

The coincidence matrices discussed above search for track candidates in projection. Further logic combines the information from the two projections. At this stage, track candidates in space are identified in small regions of the detector with a granularity in pseudorapidity–azimuth space smaller than or about 0.1×0.1, classified into six p_T ranges. These track candidates are collected together over sectors of the detector, retaining the two highest-p_T candidates per sector. In the barrel there are 32 sectors in azimuth per half barrel. Each endcap is divided into two pseudorapidity ranges, $1.05 < |\eta| < 1.92$ and $1.92 < |\eta| < 2.40$, with 48 and 24 sectors respectively.

The information from all the sectors is combined in the muon-trigger interface to the CTP (MUCTPI). This counts the number of muon candidates for each of the six p_T thresholds, and passes the multiplicity information on to the CTP. The MUCTPI is responsible for detecting cases where muons traverse more than one sector due to chamber overlaps, making sure that they are counted only once in the multiplicity calculation. Overlaps within sectors are handled by the logic specific to the barrel and endcap subsystems.

2.4.2.1 RPC-based subsystem

In the RPC-based subsystem, the coincidence matrices and the ‘pad’ logic that combines information from the two projections within regions, of size $\Delta\eta \times \Delta\phi \sim 0.2 \times 0.2$, are located on the detector. About 800 high-speed optical links are used to transmit the data from the pads to the sector logic that is located in the USA15 counting room.

The coincidence matrices operate with a clock speed of 320 MHz to match the intrinsic time resolution of the RPC detectors, so that the ‘gate’ width can adjusted to minimize the rate of fake triggers due to accidental coincidences between random hits, e.g. due to radiation in the cavern.

2.4.2.2 TGC-based subsystem

In the TGC-based subsystem, the coincidence matrices are preceded by ‘patch panels’ that synchronize the hits to the bunch-crossing clock using an adjustable gate width. The subsequent logic is clocked at 40 MHz. The on-detector system includes the low and high-p_T coincidence logic for each projection, but not the so-called r–ϕ coincidence logic that combines the two projections, which is located in the USA15 counting room.

The data from the on-detector logic are encoded before transmission on ~600 high-speed optical links to the r–ϕ coincidence circuits in the sector logic. They include the position of the track candidate within the region covered by the associated on-detector logic, and a measurement of the deflection compared to a straight-line extrapolation from the nominal interaction point. The information from the two projections is combined in order to determine the p_T range of the muon candidate, and the two highest-p_T candidates per sector are retained for transmission to the MUCTPI.

2.4.2.3 Muon–CTP interface

The MUCTPI consists of a single crate of electronics modules located close to the CTP in the USA15 underground counting room. It receives the input data electrically from the sector logic of the RPC- and TGC-based trigger subsystems. Given the small size of the system, it will be implemented using field programmable logic and fast memories. A prototype version of the MUCTPI is currently under design.

2.4.2.4 Demonstrator programme

Demonstrator prototypes have already been tested for both the RPC- and the TGC-based systems. These have proven the design concepts and the performance of the muon trigger. In the case of the RPC logic, demonstrator prototype ASICs have been incorporated in a ‘full-slice’ test system and operated at the test-beam connected to full-size prototype detectors.

2.4.3 Central trigger processor and TTC system

The CTP is responsible for combining information from the calorimeter and muon trigger systems. As discussed above, it receives the multiplicities of candidate electrons/photons, hadrons/taus, jets and muons, as well as information on the missing and total scalar E_T. It forms the overall LVL1 trigger decision, accepting events on the basis of up to 96 menu items. Each menu item is a list of criteria that, if satisfied, will result in the event being selected, subject to deadtime and prescale conditions. In the CTP the input data are combined to form decisions for each menu item using lookup tables followed by combinatorial logic. This gives a very high degree of flexibility in specifying the criteria for each menu item.

The CTP takes care of handling the deadtime of the experiment, preventing the overflow of data buffers elsewhere in the experiment by inhibiting triggers. Scalers monitor this deadtime as well as giving information on the rates for each type of trigger. Each menu item can be prescaled by a programmable factor, allowing data to be collected with high-rate triggers concurrently with low-rate ones. In addition to triggers based on information from the calorimeter and muon processors, the CTP accepts external inputs (e.g. for calibration and test triggers) and can generate random triggers and prescaled bunch-crossing triggers.

The CTP, located in the USA15 counting room, will be implemented on one or possibly two 9U printed circuit boards. The use of ASICs is not justified in such a small system. The logic is therefore implemented using a combination of configurable logic (FPGAs and CPLDs) and fast memories. A reduced-size version of the CTP has been constructed as a demonstrator prototype. This has been successfully tested in the laboratory and, together with the calorimeter trigger demonstrator system, at the test beam.

The TTC system is responsible for distributing the LVL1 trigger decision (L1A), together with the 40 MHz clock and other signals, to the detector front-end and readout systems. It uses the system developed in the RD12 collaboration, in which members of the ATLAS LVL1 trigger community are actively participating. This system uses an optical-broadcast network connecting up to about 1000 destinations to each source. The clock, L1A and other signals are encoded and broadcast as a single optical signal to all destinations. The TTC system is separated into zones, with a few zones for each detector system. All of the TTC sources (one source per zone) are located in the USA15 counting room.

Each TTC destination will include an optical/electrical converter followed by a receiver ASIC (TTCrx) that decodes the signal, providing the clock, L1A and other signals electrically. Prototype TTCrx chips have already been produced with performance close to that required for ATLAS, and the next version, implemented in a radiation-hard technology, will be available soon.

2.5 References

2-1 *ATLAS Technical Proposal*, CERN/LHCC/94–43, December 1994.

2-2 *ATLAS Muon Spectrometer Technical Design Report*, CERN/LHCC/97–22, May 1997.

2-3 *ATLAS Liquid Argon Calorimeter Technical Design Report*, CERN/LHCC/96–41, December 1996.

2-4 *ATLAS Tile Calorimeter Technical Design Report*, CERN/LHCC/96–42, December 1996.

2-5 *ATLAS Trigger Performance Status Report*, CERN/LHCC/98–15, June 1998.

2-6 *ATLAS DAQ, EF, LVL2 and DCS Technical Progress Report*, CERN/LHCC/98–16, June 1998.

2-7 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps

2-8 *LVL1 Muon Trigger User Requirements Document (Draft version 1.4)*, ATLAS working document, ATL–DA–ES–0002, March 1998.

2-9 *ATLAS Level-1 Calorimeter Trigger User Requirements Document*, version 1.1.0, ATLAS working document, ATL–DA–ES–0001, April 1998.

2-10 *RD12 Status Report*, CERN / LHCC / 97–29, April 1997.

2-11 *LVL1 Central Trigger Processor (User) Requirements Document (Draft version 1.1)*, ATLAS working document, ATL–DA–ES–0003, February 1998.

2-12 *Timing, Trigger and Control System (TTC) (User) Requirements Document (Draft version 0.2)*, ATLAS working document, ATL–DA–ES–0004, February 1998.

2-13 *ATLAS Level-2 Trigger User Requirements Document*, ATLAS note DAQ–NO–79, November 1997.