

Specification of the LVL1 / LVL2 trigger interface

Version 2.0, April 11, 2002

**M. Abolins, R. Blair, J. Dawson, N. Ellis, Y. Ermoline, C.N.P. Gee,
Ph. Farthouat, M. Landon, J. Schlereth, C. Schwick, R. Spiwoks,
T. Wengler, F. Wickens**

??? Points requiring follow up are flagged with “???” in the text ???

??? Updates to author list ???



1 Introduction

This note specifies the interface between the LVL1 and LVL2 triggers in terms of the exchange of signals and data between the two subsystems and associated requirements. The specification includes the definition of the physical links, the content and format of the data to be transmitted, and signals used (e.g. for data-flow control). Details of the implementation of the various parts of the LVL1 trigger can be found in Ref. [1].

Figure 1 shows a block diagram giving an overview of the links between the two trigger subsystems. A number of links carry data from various parts of the LVL1 subsystem (e.g. the muon trigger) to LVL2. The “region-of-interest” (RoI) data actually include some global quantities (energy sums, trigger-type information) as well as the information for each localised RoI. The RoI data from LVL1 are sent to ROBs that can be read by the DAQ system in addition to being sent directly to the LVL2 trigger on dedicated links. (Note that the RoI data are only a small subset of the data that are read out to the ROBs from the LVL1 trigger.) For events that it selects, the LVL2 trigger will pass to the DAQ system the raw RoI information that it received from LVL1 together with quantities calculated in the LVL2 trigger processing, for use in monitoring and guiding of Event Filter processing.

On each link from LVL1, the LVL2 trigger receives the event number (L1ID), the bunch-crossing number (BCID) and an 8-bit summary of the trigger type, as well as the detailed RoI information or other information.

Also shown in Fig. 1 are the control signals exchanged between the LVL1 and LVL2 subsystems. Using the standard BUSY system, the LVL2 system can inhibit the generation of LVL1 triggers via a “throttle” signal. However, it is expected that flow control will normally be applied via the RODs using the standard ROD_BUSY mechanism.

Note that the data read out to ROBs from LVL1 contain much more information than is sent to the RoIB. Also, the data sent to the ROBs will generally be for a time frame of one or several bunch crossings, whereas the data sent to the RoIB are only for the bunch crossing that caused the trigger.

This note is structured as follows. After a section defining the terms used, the links between the LVL1 and LVL2 subsystems are specified. A number of general constraints and requirements on the two subsystems are then discussed. Next the data content and format on the each of the links from the LVL1 trigger to the LVL2 trigger are described. In a final section, miscellaneous other issues are discussed.

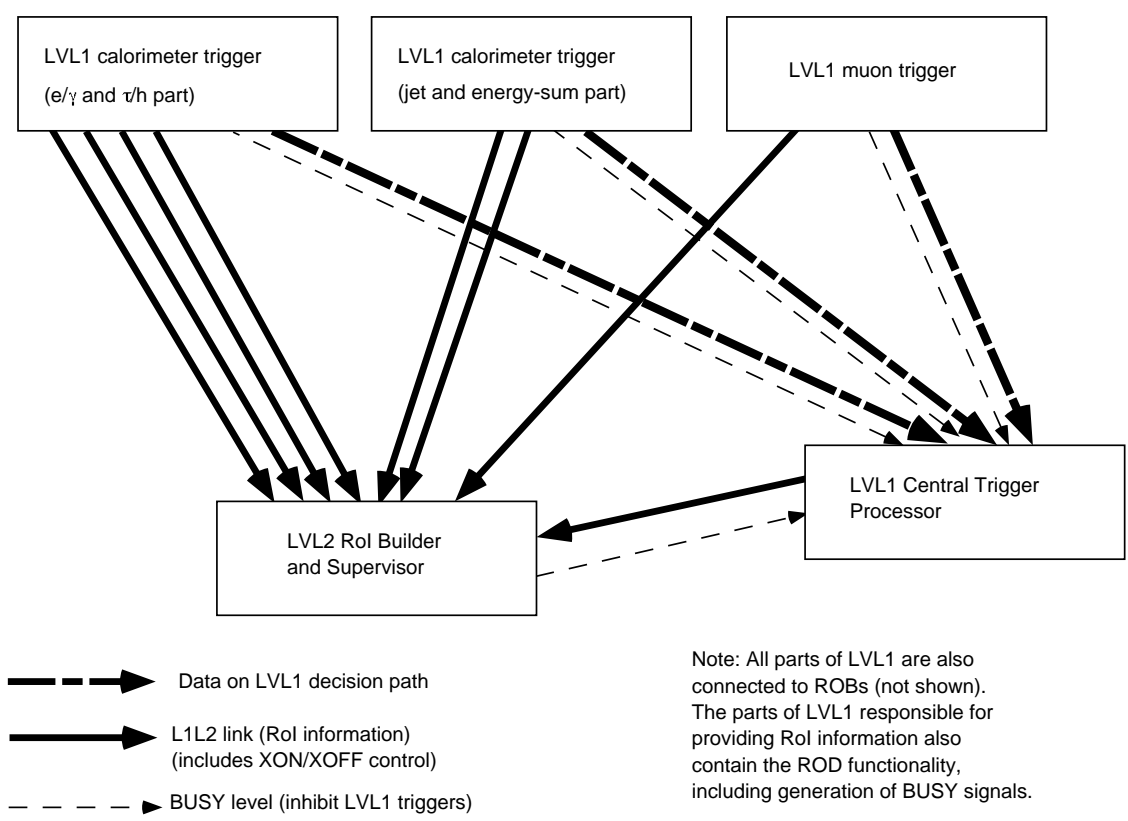


Figure 1: Block diagram of LVL1 – LVL2 interface.

2 Definition of terms

2.1 LVL1 trigger system

The LVL1 trigger system is composed of a number of elements as shown in Fig. 2. The main blocks are the calorimeter trigger, the muon trigger and the Central Trigger Processor (CTP), as well as the TTC system.

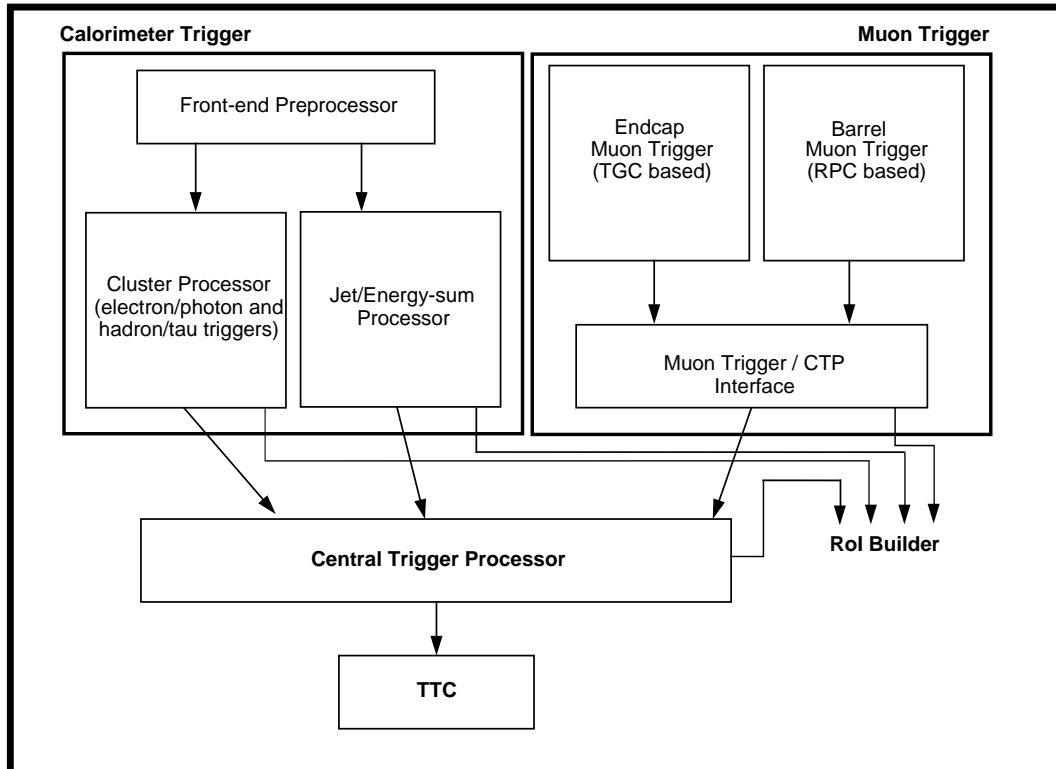


Figure 2: Block diagram of LVL1 trigger.

2.2 RoI builder (RoIB)

The RoIB [10] is the part of the LVL2 trigger that receives the RoI information from the LVL1 trigger. It also receives some global information — energy sums from the calorimeter trigger and trigger type from the central trigger processor (CTP). A block diagram of the RoIB/supervisor system is shown in Fig. 3.

On an L1A signal, the LVL1 trigger prepares and transfers to LVL2, via LVL1-to-LVL2 (L1L2) links (see Section 3), RoI fragments that pinpoint interesting features of the event within specific sub-detectors of ATLAS. These fragments are formatted as standard ROD data format records inside S-link control word headers and trailers (at least for the prototyping stage), and with the event number embedded in the record. The RoIB receives these RoI fragments, which may be considerably dispersed in time on the different links. It organises and formats a complete record for each event accepted by

LVL1 from these fragments, selects a processor in the LVL2 Supervisor (so-called RoI processor – see Section 2.3) to manage the event through the LVL2 trigger, and transfers via an S-link the assembled RoI record to the target processor. The RoIB performs these tasks at the maximum LVL1 trigger rate of 100 kHz. Note that the standard ATLAS readout link will be used for the final system if this meets the requirements of the RoIB/supervisor system.




Figure 3: Block diagram of RoIB/supervisor system.

2.3 LVL2 trigger supervisor

The Supervisor works in concert with the RoIB (see Fig. 3). It consists of a number of independent processors (the RoI processors). Each RoI processor receives a complete LVL1 RoI-data record from the RoIB, reformats it and sends it to a pre-selected LVL2 trigger system processor where the event is subjected to the LVL2 algorithms (each RoI processor manages a subset of the LVL2 system processors). When finished, the system processor sends a decision back to the RoI processor. The RoI processor then forwards the message to the Data Flow Manager (DFM) with instructions to either discard the event or to pass it on for processing by the Event Filter. The RoI processors monitor the LVL2 trigger system performance and report it to the run-control system.

3 LVL1-to-LVL2 (L1L2) links

Information will be sent to the RoIB on separate physical links from parts of the LVL1 trigger subsystem doing different aspects of the LVL1 processing. The L1L2 links are as follows:

- Muon trigger: a single link carries all of the information from the muon trigger to the RoIB. This is illustrated in Fig. 4.

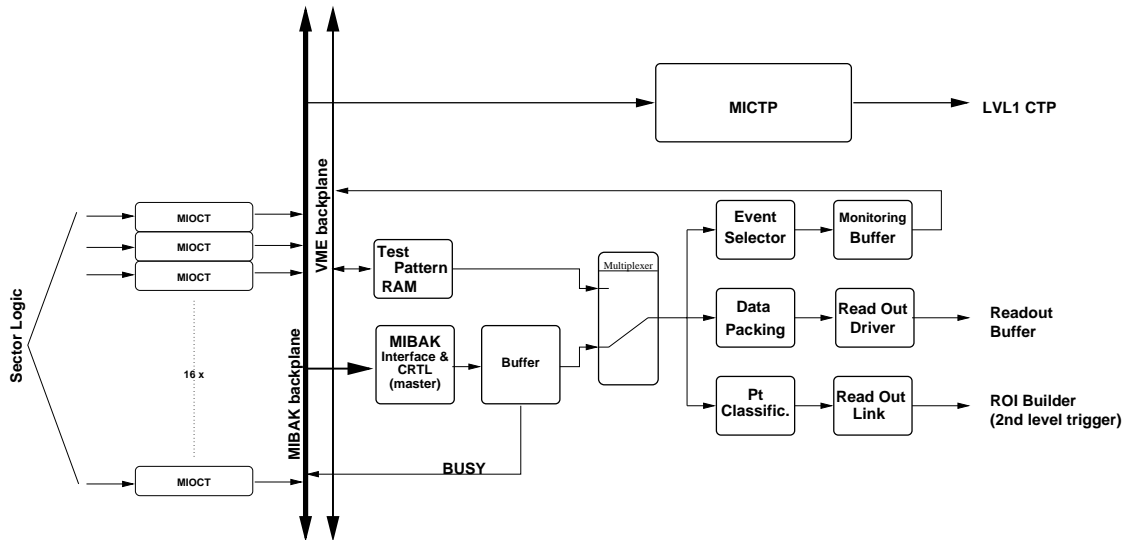


Figure 4: Logic driving the L1L2 link for the muon trigger.

- Calorimeter trigger: four separate physical links carry the electron/photon plus tau/hadron RoI information to RoIB (each link covers a different geographical region of the detector), and two separate physical links carry the combined jet and energy-sum RoI information. This is illustrated in Fig. 5.
- Central trigger processor: a single link carries all of the information from the CTP to the RoIB. See Fig. 6.

A standard physical link will be used for the L1L2 links. S-links [2] will be used for the prototyping phase; for the final system the standard readout links will be used if they satisfy the requirements. This will provide a bandwidth of about 160 MBytes/s per link, which is ample. Note that the RoIB does not receive any TTC input; the required information (L1ID, BCID, trigger-type summary) arrives on each of the L1L2 links allowing cross-checking between them.

??? We need to discuss if it makes sense to wait until ROL is defined (should we freeze on SLINK now? should we freeze on Gbit Ethernet version now??) Need follow-up discussion taking into account recommendations of Readout Link Task Force???

The data sent on the L1L2 links are required to conform to the standard ATLAS ROD data format [3], which is summarised in Table 1, since this existing format meets the

requirements for this application. The data content on the various links is specified in Section 6.

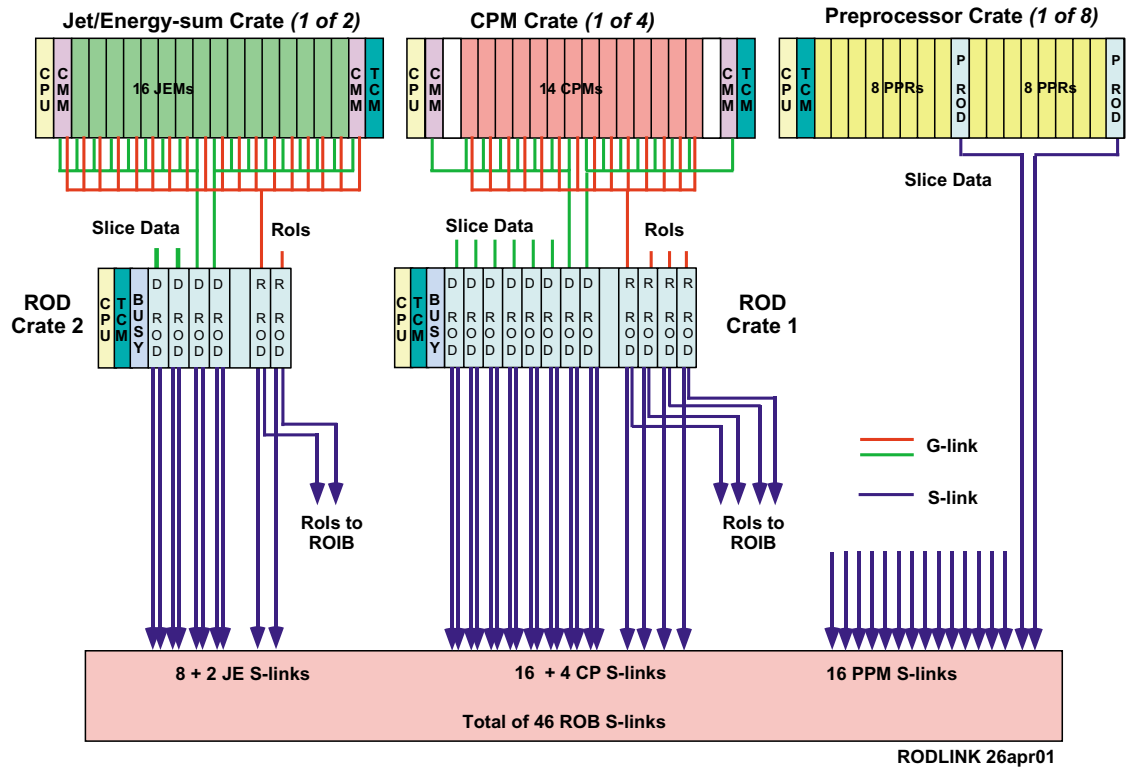


Figure 5: Organisation of RODs in the calorimeter trigger and links to LVL2. A modified architecture under discussion would reduce the number of DRODs but place extra input links on each. This would leave the number and format of L1L2 links unchanged.

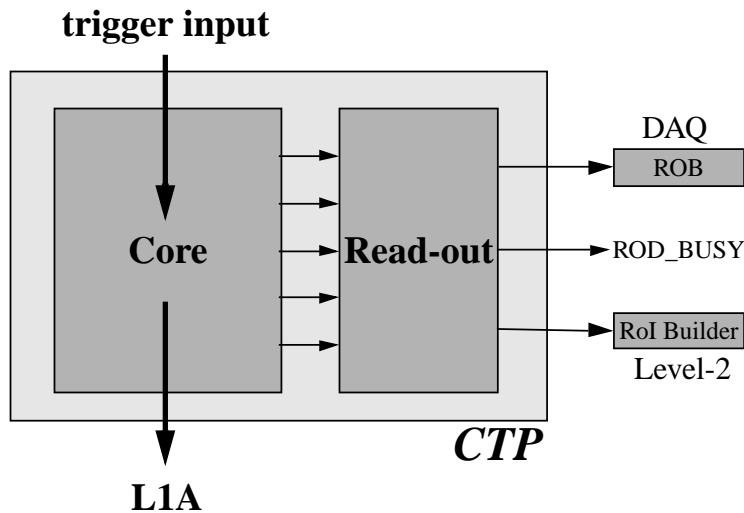


Figure 6: Link from the CTP to LVL2.

Table 1: Standard ROD data format inside S-link control-word markers.

Word	Content
	S-link fragment start marker (control word)
1	Start of header Marker
2	Header size in bytes [currently 32]
3	Format (header???) version number [currently 1]
4	Source identifier [Rod no, allocated centrally]
5	LVL1 ID [32-bit ROD_L1ID]
6	Bunch Crossing ID [12-bit ROD_BCID, zeroed each turn]
7	LVL1 trigger type [8 bits copied from TTC]
8	Detector event type
	<<< RoI data elements >>>
	<<< Status elements, e.g. error flags >>>
	Number of status elements
	Number of data elements
	Data or Status First flag [data first – value = 1]
	S-link fragment end marker (control word)

??? Need to check if Table 1 is consistent with the update of DAQ note 50. What about version number for the payload? DjF says that this will be discussed during April 2002 TDAQ week ???

3.1 Event type

The field “Detector event type” (this is the last field of the header before the data or status block starts) can be freely used to label the event type. Therefore we propose to use this field to label and distinguish:

- Physics events.
- Test events (generated artificially).
- Events which have been monitored throughout the ROD system with help of a special flag (this is currently being implemented in the MUCTPI system).
- Cosmic-ray events.
- Calibration events.

3.2 Status words

The only condition for the status word in the standard ATLAS ROD data format is that, if the first status word is non-zero, this indicates a corrupted/unusable/suspicious data block of the data fragment.

Two kinds of status information have to be transferred to LVL2:

- Error information that indicates if the data transmitted in the corresponding data fragment are corrupted, incorrect or in any sense unreliable.
- Status information useful or necessary to interpret the content of the data block.

For each kind of information a 32-bit word is proposed. The first one contains the error information. For non-corrupted events it is 0. Both words are split into two bit-fields of 16 bits each labelled “common” and “private”. The bits in the common field have the same meaning for every LVL1 subsystem involved. The private section contains bits that have meanings specific to the subsystem. They can be used to give more detailed information on error or status information. The proposed general format is discussed below; details for the specific cases of the different level-1 subsystems are discussed in Section 6.1.3 for the muon trigger, in Section 6.2.6 for the calorimeter trigger and in Section 6.3.2 for the central trigger processor.

3.2.1 Error-Bits (1st status word)

The following table lists the error bits of the common field of the first status word. For every error type a bit is reserved. If no error has been detected the data are most probably OK and the first status word is 0x00000000.

Bit position	Error Type	Abbreviation	Description
0	BCID mismatch	BCMM	An internal check of the BCID has failed.
1	EVID mismatch	EVMM	An internal check of the EVID has failed.
2	Time Out	TOUT	A time out in one of the modules has occurred. The data block may be incomplete.
3	Internal Data Transport Error	IDTE	Can be further explained in private section. Data may be incorrect.
4	Data Overflow	DOFL	An overflow in one of the internal buffers has occurred. (Should not occur if throttling works correctly.) The corresponding event fragments may contain incomplete data or status words.

3.2.2 Status-Bits (2nd status word)

The second status word contains bits that give general status information on the data of the data-block. They are to be seen as additional information concerning the whole event fragment. If set it does not mean that the data in the data-block are corrupted. Situations indicated by these bits are actually foreseen during normal data taking. Again the word is

divided into common and private fields.

Bit position	Information type	Abbreviation	Description
0	Internal Data Transport Inconsistency	IDTI	Occurs during internal processing. Is further explained in the private bit-field of the word.
1	Limited RoI Set	LRS	The number of RoIs sent to LVL2 has been limited in order to reduce the total data volume. The limit is subsystem dependent. It can be further explained in the private bit-field.

3.2.3 Definition of the protocol on the LVL1/LVL2 links

3.2.3.1 RoI Fragment Generation

There must be one and only one RoI fragment for every L1A, even if it contains no RoI data or if status or error flags are set. Note that this implies that RoI fragments may be generated with a maximum frequency of 100 kHz. Note that the RoIB must contain a mechanism to detect missing or corrupted RoI fragments.

3.2.3.2 RoI Fragment Order

The RoI fragments must arrive in the order of their generation with L1As, i.e. no RoI fragment is allowed to arrive at the RoIB before another RoI fragment (from the same processor) that was generated earlier.

3.2.3.3 RoI Fragment Numbering

Each RoI fragment must contain a 32-bit extended L1ID as required for RODs providing data to ROB. The least-significant 24 bits of this word are given by a counter that is reset by ECR and incremented by L1A; the most significant 8 bits are given by a counter of ECR that is reset under software control. The extended L1ID is used by the RoIB to build complete RoIs.

??? Note that we need to consider the implications of ECR resetting the extended L1ID both for event building and for checking of the integrity of the input data. What are the requirements on the absence of L1A being sent prior to and after ECR (veto period)? ROIB community ???

3.2.3.4 RoI Fragment Latency

The RoI fragments must be *available* for transmission on the LVL1/LVL2 links no later than 100 μ s after reception of the L1A. Note that the exact value of the maximum latency is to be discussed. Note that the latency with respect to arrival of the RoI fragments at the

RoIB can increase if the RoIB exerts flow control; this could be due to back-pressure from the LVL2 trigger or other parts of the HLT and DAQ system; eventually the ROD_BUSY mechanism will stop generation of L1As.

??? A short value of the maximum latency, such as 100 μ s, implies that the output buffer size in the LVL1 RODs must not be too large. Otherwise, e.g. in case of back-pressure or in case of a number of L1A's in quick succession, a long queue of RoI records could build up in the LVL1 RODs. When the XOFF gets removed, there could be a significant delay before the last record in the queue would be ready for transmission. A negative consequence of limiting the size of buffers in the LVL1 RODs is the possible introduction of deadtime through the ROD-BUSY mechanism. In view of this, a longer maximum latency value would be highly desirable. The ROIB community is asked to consider a value of 1 ms – is this feasible ???

3.2.3.5 RoI Fragment Skew

The RoI fragments must be available for transmission on the LVL1/LVL2 links with a time difference of less than 100 μ s between the first and the last RoI fragment becoming available. Note that the exact value of the maximum skew is to be discussed. Note that the skew with respect to the arrival of the RoI fragments at the RoIB depends on the flow control exerted by the RoIB on the different links. Note that the RoIB detects missing RoI fragments using a time-out mechanism; the value for the time-out must be chosen to be larger than the skew defined in this rule.

??? Is this a separate requirement compared to latency? Or should we assume that the maximum skew is equal to the maximum latency? ???

3.2.3.6 RoI Fragment Distance

The ROI fragments can arrive back-to-back, i.e. without any idle time of the LVL1/LVL2 link between them. The RoIB must provide sufficient input buffering (according to the S-Link standard) to handle this.

3.2.3.7 RoI Fragment Format

The ROI fragments must be formatted as defined in the standard ROD format (see Table 1). Status and error bits are extended according to the description in this document.

3.2.3.8 RoI Fragment Size

The ROI fragments must not exceed a total size of 63 32-bit words including the ROD format headers, trailers and S-Link control words, leaving a maximum of 48 RoI data words.

??? Is there any significant drawback in specifying a larger maximum size – this could be useful, e.g. in special modes of operation. ROIB community is asked to consider this ???

3.2.3.9 RoI Fragment Data

The data format of the RoI data words is defined in the following sections of this document.

4 General constraints and requirements on LVL2 from the LVL1 trigger

4.1 LVL2 data inputs

The inputs to LVL2 shall be able to sustain the maximum event, RoI and data rates expected from LVL1 with negligible introduction of dead-time (<1%) [4].

The inputs to the RoIB shall be able to accept the maximum LVL1 RoI data fragment rate (100 kHz per link), the maximum link transfer rate (e.g. 160 Mbytes/s) and the LVL1 RoI data-fragment size expected from LVL1 with negligible introduction of dead-time (<1%) [4].

The flow of data on each of the so-called L1L2 links from the LVL1 trigger system to the RoIB shall be controlled using the XON/XOFF mechanism included in the S-link / readout link protocol. Note that when asserting XOFF in order to disable the data, there will be a delay until the data flow is really interrupted. This latency depends on details of the implementation of the link as well as on the length of the transmission cable. Details of the behaviour for the S-link are described in “The S-link Interface Specification” [2]. It should also be noted that asserting XOFF will not prevent occurrence of further LVL1 triggers (triggers will continue, possibly for a considerable time, until a BUSY gets asserted as a consequence of buffers filling up in the LVL1 system or in the ROB and RODs).

4.2 Throttle

To deal with exceptional cases in which the LVL2 trigger cannot continue to accept further LVL1 triggers, the LVL2 trigger shall be able to assert a BUSY signal.

For this purpose, the LVL2 supervisor shall be able to assert a BUSY signal to LVL1. This signal will be produced in a programmable module using as inputs individual BUSY signals from each supervisor processor. A supervisor processor will raise a BUSY in a timely manner if it perceives problems with the string of processors under its control, e.g. their inability to keep up with ambient trigger rate. Note that asserting BUSY will promptly prevent the occurrence of further LVL1 triggers, but will not prevent the flow of incoming data that will continue until the buffers in the LVL1 system become empty.

4.3 Monitoring

For events selected by the LVL2 trigger, all of the data sent from LVL1 to LVL2 shall be passed on to the next trigger level.

This will allow monitoring of the L1L2 links by comparing data read out from LVL1 via the LVL2 trigger with those read out directly from LVL1 ROBs. (Such monitoring is in addition to checks that may be made on the link protocol error flags.)

The content of the RoI information needs to be available to monitoring programs running in LVL2 and beyond. To facilitate this, the RODs providing the RoI information also generate a duplicate copy that is sent to ROBs as part of the LVL1 trigger data, and so is collected and built into the complete events.

4.4 Error recovery

The LVL2 trigger shall be able to recover from situations in which, for an isolated event, incomplete or corrupted information is received from LVL1.

If an RoI record is to be built for an event and only a subset of the fragments is received from LVL1 something has to happen. One alternative is to start a timer when the first fragment is received and if and when the timer times out either abort the event or build the event with the fragments which have been received and send it to the target RoI processor. If and when the errant fragment is received it must then be ignored. An alternative is to let the S-link (readout link) channel and target processor wait and distribute events across the remaining channels until the missing fragment is received. However, this is really a variation of the timeout scenario because if the fragment is never received something has to happen.

Note that in order to recover from some classes of error it will be necessary for the run-control system to be involved, for example to perform reset actions.

LVL2 must provide a programmable option to force acceptance or rejection of events with incomplete RoI data.

4.5 Test and diagnosis of faults

It shall be possible to operate parts of the LVL1 and LVL2 triggers together in test mode.

This will be needed during the commissioning phase, and also for test and for diagnosis of faults.

The RoI Builder has provisions for taking its inputs from data buffers loaded via VME. Arbitrary patterns of bits can be loaded, processed through the RoI Builder/Supervisor chain to verify correct operation of the device. No test triggers are needed from LVL1 for this purpose. The only items remaining to be tested are the links (currently S-link) from LVL1 to the RoI Builder. These links may be tested individually by transmitting patterns of marching 0's and 1's. Provisions to do this should be built into the LVL1 system.

Test patterns including alternating 1's and 0's will be used to verify connections over the readout links. A more detailed test procedure will be needed to ensure that the LVL1 and LVL2 subsystems operate correctly together as a system, including correct generation of BUSY and proper recovery from error conditions. Details will be included in future revisions of this document, building on experience gained in joint tests of LVL1 and LVL2 prototypes.

5 General constraints and requirements on LVL1 from the LVL2 trigger

5.1 Constraints from LVL2 URD

A number of general requirements on the LVL1 trigger have been extracted from the LVL2 User Requirements Document [5].

- The LVL1 rate must not exceed 100 kHz averaged over short periods.
- The LVL1 system must be able to accept from LVL2 a BUSY level to inhibit further LVL1 triggers.
- The maximum time between the LVL1 trigger decision and the arrival of CTP and RoI information from LVL1 into the LVL2 system should not exceed 10 μ s for a time-isolated event. For a worst case condition, the delay should not exceed 100 μ s; a worst-case condition being defined as 8 L1A signals in 80 μ s (worst case allowed by dead-time logic in CTP). These figures assume that XOFF is not asserted on the L1L2 links.
- LVL1 shall provide event numbers that are unique within a given time-window of specified size (window of 100 μ s is sufficient).
- LVL1 must be able to generate special triggers to test the operation of LVL2.

??? Values quoted above are taken from LVL2 URD which is out of date – what do we do about updating it ???

5.2 Constraints from RoIB implementation

The data transferred from LVL1 to the RoIB are sent via S-link / readout link and accordingly must be organised following the S-link / readout-link specifications. Furthermore, the standard ATLAS ROD data format will be used. At least for the prototyping phase where S-link will be used, each RoI fragment must have a 32-bit control-word header and a 32-bit control-word trailer. The body of the fragment must be 32-bit data words with no embedded control words. The data source and the event number must be provided at fixed positions in the header of the RoI fragment.

There is a limit to the amount of skew the RoIB can tolerate between the first RoI fragment of an event and the last RoI fragment of that event received at the RoIB. If there are 8 RoI processors and LVL1 accepts events at a rate of 100 kHz then, on average, every 80 μ s the RoIB will want to send a record to the same RoI processor assuming a “round-robin” algorithm. Some reasonable amount of buffering can be provided, but a limit on the skew on the order of 100 μ s is indicated.

??? Value of skew limit to be discussed ???

There must be a limit on the number of data words in the transfers from LVL1 to LVL2 because these fragments must be buffered in the RoIB while waiting for all fragments

from an event to be received. Each RoI maps onto one data word. In the prototype hardware a fragment can consist of as many as 31 words including the S-link header and trailer control words. Subtracting the two control words, the 11 words of overhead in the standard event format and the two status words, leaves 16 words for RoI data. In the final system, the buffers in the RoIB will accommodate up to 63 words per fragment, leaving 48 words for RoI data after subtraction of overhead items. The LVL1 system will limit the number of RoIs of different types that can be sent on the L1L2 links as discussed in the next chapter – these limits are more restrictive than that given by the maximum fragment length.

5.3 Other requirements

It must be possible in a programmable way to disable sending of RoI information for certain trigger types (e.g. calibration or test triggers); in such cases empty event fragments must still be sent.

It should be noted that, in order to interpret the RoI information provided by LVL1, LVL2 will need access to a large number of configuration parameters from LVL1. Such information is identified in Section 6.1.4 for the muon trigger, in Section 6.2.7 for the calorimeter trigger and in Section 6.3.3 for the central trigger processor.

It must also be possible to operate the LVL1 trigger in a mode where the XON/XOFF protocol on the L1L2 links is ignored and where no RoI fragments or DAQ data are generated. This will be needed for commissioning the LVL1 system without higher-level triggers being operational, for example.

6 Data content, format and order on links from LVL1 to the RoI builder

6.1 Muon trigger


The muon-trigger to CTP interface (MUCTPI) does not define or modify the definition of the RoIs given in Ref. [6] (see Appendix A for a summary of the muon-trigger algorithms). The role of the interface to the LVL2 trigger is to check the consistency of the data coming from the preceding stages (sector logic for the barrel and end-cap muon trigger systems) and to select, according to several programmable options, the muon candidates that are presented to the LVL2 trigger. The data will be sent from the MUCTPI Read-Out Driver (MIROD) to the RoIB via a single L1L2 link (S-link driven by the 40 MHz system clock in the demonstrator prototype implementation).

Each sector can provide at most two candidates as described in Ref. [6]. Two independent p_T thresholds can be set for filtering the candidates that are passed to LVL2, one acting on the highest- p_T and the other on the second-highest- p_T candidate per sector. It is possible to set the threshold for the second-highest- p_T muon to the value 7, beyond the range of thresholds that exist (1-6), in which case only the highest- p_T candidate per sector will be sent to LVL2.

The stages in front of the MIROD module (one of the components for the MUCTPI – see Ref. [1]) have the possibility to send all muon candidates in a window of up to ± 2 BC width around the BC that caused the LVL1 trigger. The candidates sent to the LVL2 trigger are from a single BC that can be chosen (configuration parameter) within the five possible ones.

All overlap and overflow bits contained in the data coming from the sector logic (see MUCTPI section in the TDR [1] and Ref. [6]) are retained in the data sent to the LVL2 processor. Note that in overlap regions a single muon may give rise to more than one RoI. Such cases are resolved in the LVL1 muon counting logic, but all muon candidates are passed on to LVL2.

Figure 7 shows the proposed data format for the muon-trigger RoIs sent to the RoIB. The data words contain all information on the muon candidates which is provided by the Sector Logic to the MUCTPI. The Sector Logic can at most send two muon candidates per sector. A flag (labelled “1st”) indicates that a candidate has been the highest- p_T one in a sector. The encoding of the RoI-number and the overlap flags is also shown in Fig. 7 (the definition of RoI numbering is given in Ref. [6]). The overlap flags indicate that a candidate has been found in a region that overlaps with a neighbouring sector, and that the overlap has not been resolved by the Sector Logic. The MUCTPI takes into account these flags when forming the total muon candidate multiplicities for LVL1. LVL2 receives the full information sent by the Sector Logic. Figure 7a shows how the RoI records for the muon candidates are transmitted as an event fragment following the ROD data-format specification.



- CTRL : SLINK control word flag
- Pt : binary encoded Pt threshold (1 to 6)
- 1st : 1/0 most/second most energetic candidate in sector
- Ro : more than 1 candidate in ROI
- So : more than 2 candidates in sector
- cs : charge sign of muon candidate (Endcap only)

ROI and overlap :


Barrel :




E: overlap with neighbouring endcap sector

B: overlap with neighbouring barrel sector

Endcap :




Forward :



sector address:

BARREL :




Hemisphere selector (pos/neg rapidity region)

5 bit sector ID (0...31)

Barrel identifier

FORWARD :




Hemisphere selector (pos/neg rapidity region)

5 bit sector ID (0...23)

Barrel identifier

ENDCAP :



Hemisphere selector (pos/neg rapidity region)

6 bit sector ID (0...47)

Endcap identifier

Figure 7: Data format for individual RoIs from the muon trigger.

6.1.3 Status words


The status information is encoded in two different words. The first word indicates the presence of fatal errors. If this word is different from zero the data for the corresponding event may be corrupted and useless. In this case, the MUCTPI can be programmed to send an empty data frame (no candidates) to the RoI builder. The second word contains status information. The proposed encoding of both words is shown in Fig. 8.

As explained in Section 3.2, each word is divided into two bit-fields of which the least significant (bits 0 to 15) have a common meaning for all subsystems providing data to LVL2. Here only the MUCTPI specific part is explained.

The two bits in the first status word specify further an internal data transfer error (IDTE). If such an error occurred because the bus error line on the internal back-plane of the MUCPTI system (MIBAK) was raised, then MBERR is asserted. If on the other hand the MIROD detects a protocol error during the data collection over the MIBAK, then the MBPER flag is set.

The most significant bit in the second status word, labelled “M” indicates if the monitoring flag for the event has been set. The other four bits specify the case when a limited RoI set is sent to the RoIB (LRS flag asserted). If in the event a candidate has been suppressed because it did not exceed the programmable p_T threshold one of the flags SUP1 or SUP2 are activated. If the candidate was the highest- p_T candidate of the sector SUP1 is asserted, otherwise SUP2. COFL indicates that the maximum number of muon candidates which is sent to the RoIB has been exceeded (the number is programmable between 1 and 16). SOFL indicates that more than 10 candidates of one p_T have been found.

Muon Status Words



Errors:

MBPER : Mibak Protokol Error
 MBERR : Mibak Bus Error
 DOFL : Data Overflow
 IDTE : Internal Data Transfer Error
 TOUT : Time Out
 EVMM : Eventnumber Mismatch
 BCMM : BCID number Mismatch

Status:

SOFL : Muon Sorter Overflow
 COFL : Candidate Counter Overflow
 SUP2 : Candidate 2 has been suppressed
 SUP1 : Candidate 1 has been suppressed
 LRS : Limited ROI Set
 IDTI : Internal Data Transfer Inconsistency
 M : Monitoring Flag

Figure 8: Status words for muon trigger.

6.1.4 Configuration Information List

In order to interpret correctly the data sent to the LVL2 system (and to the DAQ) the actual configuration of the system must be saved in databases in order to be accessible for any unit that has to deal with the data. This configuration comprises, for example, the settings of all programmable options of the hardware or the mapping of logical channel numbers to the corresponding physical position in the detector.

More details of the configuration data for the muon trigger can be found in Ref. [11]. Examples of the information that is recorded are:

- Summary information on the six p_T -threshold settings.
- Parameters that control options in the trigger algorithms (e.g. use of Tile calorimeter to validate barrel muon candidates), and the readout (e.g. number of bunch crossings to be readout out with each L1A).
- Detailed record of data loaded into the electronics of the muon trigger for each data-taking run that allow the results of algorithm that is implemented in the hardware to be reproduced in software monitoring tasks. This includes the definition of the “roads” that are implemented in the barrel and endcap trigger processors.

6.1.5 Monitoring and test facilities

In the muon trigger system, different possibilities exist to monitor the correct functioning of the system and its interaction with LVL2. All muon candidates that arrive at the MUCTPI system are sent via a ROD to a ROB. Those candidates that were not sent to LVL2 because they did not pass one of the programmable thresholds are labelled with a flag so that they can be easily recognised by the monitoring software.

Among the programmable choices to select the events that are written into the monitoring buffers, there is the possibility to monitor events for which a dedicated flag is set. This flag is set according to different programmable criteria in the MICTP system and subsequently sent out via the internal back-plane with the L1A signal to the MIOCT modules. The MIROD module receives the flag along with the readout data of the MICTP. It is transferred in the status word to LVL2 so that the possibility exists to do synchronised monitoring of single events in the MUCTPI and the LVL2 system.

The MUCTPI contains playback memory that can be loaded via VME with arbitrary test data. These data will propagate to LVL2 which allows one to test the operation and interaction of the two systems under different conditions. Erroneous data can be fed into the memory in order to test the error handling in the MUCTPI and in the LVL2 system.

6.2 Calorimeter trigger

The LVL1 calorimeter trigger drives a total of six L1L2 links, four associated with the electron/photon plus tau/hadron trigger subsystem, and two associated with the jet plus energy-sum subsystems. The RoI information produced by each of these subsystems is described in the following sections. The algorithms used to define the RoIs are summarised in Appendix B; details can be found in Ref. [1].

Figure 9 shows the format of all RoI types produced by the calorimeter trigger. The different types are distinguished by a 4-bit data field, the RoI Type. The order in which RoI types may be presented is not defined here, and software should inspect the RoI type to determine the data type. This strategy allows some useful flexibility in cabling systems for testing. The higher-order bits of the RoI type field may also be used to distinguish RoIs from other data present on DAQ readout links.

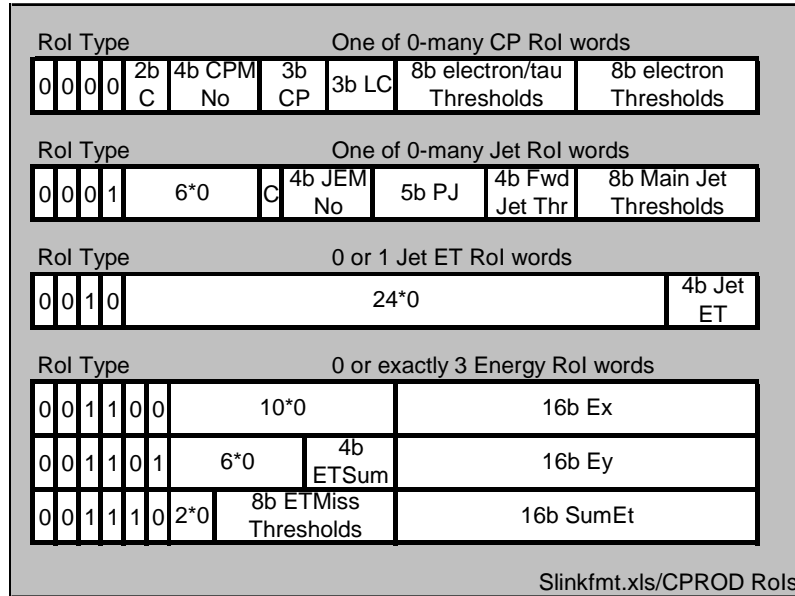


Figure 9: RoI formats from the Calorimeter Trigger

6.2.1 Electron/photon and tau/hadron trigger

The data content and format for the electron/photon (e.m.) and tau/hadron trigger RoIs is as follows. For each RoI, the following information is packed into a 32-bit word:

Threshold information:

- 8 bits When set, a bit indicates that the corresponding trigger E_T threshold set was passed for the e.m. trigger.
- 8 bits When set, a bit indicates that the corresponding trigger threshold set was passed for the e.m. or tau/hadron trigger. Each bit can be configured as an e.m. or tau/hadron trigger bit, and must be interpreted taking into account the LVL1 configuration.

Position information:

- 3 bits RoI Local coordinates within cluster-processing FPGA [0-7].
- 3 bits CP FPGA position within Cluster Processing Module [0-7].
- 4 bits Cluster Processing Module position in crate [1-14].
- 2 bits Cluster Processing Crate Number [0-3].

Each crate occupies a region of ± 2.6 in η by 1.6 in ϕ . Crate 0 starts at $\phi = 0$.

Each CPM occupies a region of 0.4 in η by 1.6 in ϕ . A possible convention places CPM 1 at the left-hand end of the crate, from $\eta = -2.6$ to -2.2 . Each CPM occupies the full crate range of ϕ .

Each CP FPGA occupies a region of 0.4 in η and 0.2 in ϕ . FPGA 0 starts at $\phi = 0$ relative to its CPM.

Within each CP FPGA, the eight possible RoI local coordinates are assigned thus:

6	7	2	3
4	5	0	1

This ordering derives from the encoding within the CP FPGA, and may be changed in the production version.

Other information:

- 4 bits RoI Type, set to 0000.

The RoIs are transmitted as the data part of an ATLAS ROD data fragment (see Table 1). RoIs from each of the four cluster crates are carried on a separate link. At the present design stage, RoIs will be geographically ordered, with the outer loop over CPM number, the next inner loop over CP FPGA ID, and the innermost readout loop over local RoI coordinates in the CP FPGA. This ordering is subject to change during prototyping, and cluster- E_t ordering may be considered.

6.2.2 Jet trigger and energy-sum triggers

Jet, energy-sum, and jet-energy-sum RoIs are carried on two links from the Jet/Energy subsystem. Each link will carry the jet RoI information from half of the Jet/Energy processor. The energy-sum and jet-energy-sum RoI information will also be carried on these links. It has not been decided which links will carry the energy-sum and jet-energy-sum data, nor if these data will be at the start or end of the event fragment. As indicated above, code interpreting the event fragments should use the “RoI type” fields to distinguish the various data blocks.

The format and content of all RoIs from the Jet/Energy-sum Processor is preliminary, and will be confirmed when the production JEM and ROD firmware is designed.

6.2.2.1 Jet RoIs

The data content and format for the jet trigger RoIs is as follows. For each RoI, the following information is packed into a 32-bit word:

Threshold information:

- 8 bits When set, indicates that the corresponding main jet threshold set was passed. Each bit can be configured as a square window of 0.4, 0.6, or 0.8 on each side, and must be interpreted taking into account the LVL1 configuration.
- 4 bits When set, indicates that the corresponding forward jet threshold set was passed. Each bit must be interpreted taking into account the LVL1 configuration.

Position information:

- 5 bits RoI Local coordinates within JEM [0-31].
- 4 bits JEM position in crate [0-15].
- 1 bit Jet Energy Processing Crate Number [0-1].

Each Crate processes two diagonally opposite quadrants in ϕ , with the first eight JEMs processing one quadrant, and the remaining JEMs the second quadrant. All JEMs contribute to the main jet trigger. Part of the processing capacity in the first and last JEMs within each group of eight is used to form the forward jet trigger. Crate 0 starts at $\phi = 0$.

Each JEM occupies a region of 0.8 in η by 1.6 in ϕ . A possible convention places JEM 0 at the left-hand end of the crate, from $\eta = -3.2$ to -2.6 .

Within each JEM, the least significant two bits determine the η value, (0 meaning lowest η), and the remaining three bits determine the ϕ position (0 meaning lowest ϕ).

Other information:

- 6 bits Reserved for future use. This may include an indication of overflow in the RoI calculation.
- 4 bits RoI type, set to 0001 for jet RoI.

At the present design stage, the jet RoIs will be geographically ordered, with the outer loop over crate number, the next inner loop over JEM position in crate, and the innermost readout loop over RoI position in each JEM. Again, this readout sequence may change during prototyping.

6.2.2.2 Summed- and missing- E_T RoI

For each L1A, a single energy-sum “RoI” is produced containing three 32-bit words. To distinguish these data from other RoIs and from each other, an extended RoI Type field is used, with the most significant four bits used to identify each word as an energy-RoI. The next two bits are used to identify the specific word, even though the words are always generated in the same sequence together.

The data are:

- 16 bits E_x Value
- 10 bits Reserved for future use
- 2 bits RoI sub-type, set to 00 for E_x datum
- 4 bits RoI type, set to 0011 for energy data block

- 16 bits E_y Value
- 4 bits Sum- E_T Thresholds passed
- 6 bits Reserved for future use
- 2 bits RoI sub-type, set to 01 for E_y and Sum- E_T datum
- 4 bits RoI type, set to 0011 for energy data block

- 16 bits Sum- E_T Value
- 8 bits Missing- E_T Thresholds passed
- 2 bits Reserved for future use.
- 2 bits RoI sub-type, set to 10 for Sum- E_T and Missing- E_T datum
- 4 bits RoI type, set to 0011 for energy data block

Unsigned arithmetic is used to make sub-sums of E_x and E_y in the separate quadrants of Jet-Energy crates. The sub-sums are then converted to twos-complement signed values conforming to the ATLAS coordinate system (see Technical Coordination TDR, Chapter 11) for further addition. E_x and E_y values sent in the RoI are the final signed twos-complement values obtained by addition of values from the separate crates.

6.2.2.3 Jet transverse-energy-sum RoI

For each L1A, a single jet transverse-energy-sum RoI is produced containing one 32-bit word. The data are:

- 4 bits Jet E_T -sum thresholds passed.
- 24 bits Reserved for future use.
- 4 bits RoI type, set to 0010

The two JEP crates both contain JEMs serving diagonally opposite quadrants in ϕ .

6.2.3 **Number of RoIs**

The number of RoIs that can potentially be generated by the calorimeter trigger is of the order of thousands, while the number expected from a typical event is less than ten. No useful function is served in sending large numbers of RoIs to LVL2. The maximum number to be sent will be subject to an overall limit determined by memory sizes, but a lower limit may be imposed by a configuration parameter loaded at the beginning of a run.

All RoI candidates generated in the calorimeter trigger are read out and zero-suppressed in RODs. There is no time penalty in reading small or large numbers of RoIs prior to the ROD. Inside the RODs, sufficient memory must be allocated to store RoIs prior to transmission on the L1L2 link. Each ROD would send up to this number of RoIs to LVL2, discarding any other non-zero RoIs and setting an indicator status bit to show that this had been done. The maximum number of RoIs to be sent to LVL2 should be programmable at run start, subject to the agreed upper limit (e.g. 32). The agreed RoI limit can be enforced only per ROD, not across the system as a whole, although the RoIB could apply a further cut.

6.2.4 **RoI Readout Latency**

There are four stages of data processing:

- (a) Generation of RoIs: RoI candidates are computed synchronously with trigger processing. On L1A, they are copied to FIFOs to await readout.
- (b) Collection of RoIs from CPMs, JEMs and CMMs to RODs: The links to the RODs will normally be empty, and the data are then read out serially by G-Link. The

readout in all cases is without zero suppression. For the Cluster and Energy RoIs, the readout time is expected to be around $1.6\mu\text{s}$ (for two 32-bit words). The JEM and CMM timings depend on data formats which are still in discussion, but could be similar.

- (c) Zero suppression in the ROD: This requires checking of data in each RoI, and selection of the occupied RoIs. Part of this can be done as the data are entering the ROD, so the added time is expected to be $< 1\mu\text{s}$. This will be measured during prototyping.
- (d) Transmission on L1L2 link: A separate L1L2 link is allocated for each e.m. or Jet/Energy-sum crate. The most pessimistic timing estimate is to assume the maximum fragment length (63 words) acceptable to the RoI builder, requiring $2.5\mu\text{s}$ to transmit.

The total RoI readout latency is thus expected to be less than $5\mu\text{s}$. If several L1As occur in quick succession, information is buffered in the processing modules and in the ROD, and sent as soon as link capacity is available. The ROD must await the arrival of the trigger type over the TTC system before the event header can be built. The trigger type normally arrives at the ROD within $1\mu\text{s}$ of the L1A (less than the data transfer time into the ROD). However the TTC system may be busy, so buffering is allowed in the ROD to allow for a short (programmable maximum) delay.

6.2.5 Status words

As explained in Section 3.2, each status word is divided into two fields, of which the least significant part (bits 0-15) have a standard meaning. The meaning of the other, calorimeter trigger specific bits, is as follows (see Fig. 12).

The two bits in the first status word indicate that an internal data transfer error was detected between the Cluster or Jet-Energy processor modules and the ROD (first bit), or between the Pre-processor and the Cluster or Jet-Energy modules (second bit).


The two bits in the second status word indicate that an incoming datum from the pre-processor was at the saturation value, or that saturation was reached during a subsequent arithmetic operation. The detailed handling of saturated data for RoI processing is still under discussion.

6.2.6 Configuration Information List

The following is a partial list of information describing the settings of the calorimeter trigger; more details can be found in Ref. [12]. In general, settings will be independently loadable into each processing element (FPGA or ASIC). Allowing thresholds to be adjusted with η and ϕ should help to reduce rate sensitivity in overlap areas of the calorimeters.

All thresholds are unsigned integers with a least significant bit precision of 1 GeV.

Calorimeter Status Words



Errors:

RLER : Error on the Link from the ROI Processor to the ROD
 PLER : Error on the Link from the Preprocessor to the tower processor
 DOFL : Data overflow
 IDTE : Internal Data Transfer Error
 TOUT : Time Out
 EVMM : Eventnumber Mismatch
 BCMM : BCID number Mismatch

Status:

IDS : Internal Data Saturation
 AOF : Arithmetic Overflow
 LRS : Limited ROI Set
 IDTI : Internal Data Transfer Inconsistency

Figure 12: Format of calorimeter trigger status words.

6.2.6.1 Electron/photon + Tau/Hadron Trigger

The configuration is separately loaded for each $0.4 (\eta) \times 0.2 (\phi)$ area. Each area requires 16 sets of:

- e.m./tau Cluster threshold;
- e.m. Isolation Threshold;
- Hadronic Ring isolation threshold;
- Hadronic Core isolation threshold;
- Electron/tau control bit. This bit when clear indicates that the threshold set will be treated as an electron, and when set that it will be treated as a tau. The bit should be provided in the database for all 16 threshold sets, even though only 8 can be re-configured at run time.

6.2.6.2 Jet Trigger

The configuration is separately loaded for each 0.8×1.6 module. Each area requires 8 sets of

??? Norman and Murrough to confirm “ 0.8×1.6 module” following discussion at meeting of 8 March 2002 ???

- Jet threshold;
- Jet size control bits. The possible values of 0, 1 or 2 indicate jet window sizes of 4×4 , 6×6 or 8×8 jet elements.

Each module also requires

- Threshold on elements entering Jet and Missing- E_T calculations;
- Threshold on elements entering Jet and Sum- E_T calculations;

Further information may be needed to configure the forward jet trigger.

6.2.6.3 Missing and sum- E_T Triggers

- Eight thresholds for Missing- E_T trigger;
- Four thresholds for Sum- E_T trigger.

6.2.6.4 Jet- E_T -sum triggers

- Four thresholds for jet- E_T -sum triggers

6.2.6.5 Preprocessor

The configuration is separately loaded for each 64-cell preprocessor module.

- BCID Coefficients.
- Look-up Table values.
- Channel enable/disable.
- A number of other parameters, including the FADC coarse and fine timing offset, have to be loaded but do not affect the digital trigger processing.

6.2.7 Error handling

Digital data from the LVL1 calorimeter trigger pre-processors to the cluster processor and jet/energy-sum algorithms may potentially be corrupted either due to link loss (LVDS or G-Link) or data bit errors on cables or back-planes. Using parity bits and link status bits, links within the calorimeter trigger system are monitored in hardware both on a tick-by-tick basis, and also over longer periods when there may be no triggers.

There is minimal bandwidth available for parity bits, so only about half of all possible data errors are detected. Absence of the error cannot be assumed to mean that the data are guaranteed to be good, although the links are highly specified and tested.

When a potential data-corruption error is detected at the cluster or jet-energy processor input:

- The received data are set to zero. In the case of the multiplexed data from the preprocessor to cluster processor subsystem, the data after demultiplexing are zeroed for two clock cycles.
- Bits are set in the data read (via RODs and ROB) from each module to indicate precisely which class of error was detected.
- If an RoI is generated whose algorithm covers the affected tower, an “Error in Incoming Data” bit could be set in the RoI.

Data travelling downstream of the calorimeter trigger algorithms may also be corrupted while travelling to the CTP. Such errors are at present detected only by monitoring.

Information read out via RODs to LVL2 could be corrupted before or after the ROD. It is likely that input to RODs will carry parity bits, with failures being reported in a status field in the event fragment. More details will be provided as the system is designed. Data errors on the L1L2 links and the readout links are detected only by monitoring. Note that in addition to data sent over the L1L2 links, the LVL2 trigger may, for a small subset of the events, access data from the ROB associated with the LVL1 calorimeter trigger.

6.2.8 Catastrophic errors

A class of errors exists which cannot be corrected locally and may need some system-wide corrective action. These are all associated with general system failures (such as missing ROD_BUSY) which should not happen if the system is operating correctly. Examples of the consequences of catastrophic errors include:

- Mismatch of BCID counters associated with data read at the ROD in response to a single L1A, i.e. synchronisation errors.
- Arrival of L1A when readout buffers (RoI or DAQ) are already full.
- Attempt to read event data from a trigger processor module when internal readout buffers are empty.
- Arrival of readout data at the ROD when ROD buffers are already full.

In each case, local corrective action allows processing to continue but may well result in all subsequent data or triggers being corrupt. A possible course of action is to flag such catastrophic errors in all RoIs sent to LVL2 and in all event fragments sent to the ROB from the affected region (probably a crate or module) until the condition is cleared.

6.2.9 Monitoring, testing and debugging facilities

Several options exist to monitor the correct functioning of the system. All calorimeter trigger modules are fitted with memories so that one or a few cycles of arriving and departing data from each module may be read into the DAQ data in response to each LVL1 accept. The RoI data is also copied into the DAQ stream. The number of data cycles read may be programmable and more data cycles may be read for selected trigger types or bunch crossing numbers or both. The same capability could allow no readout for certain trigger types.

Local monitoring uses a buffer in the ROD that is filled for all events triggered by LVL1, and then inspected by an onboard processor (DSP or CPU options are under consideration). Events are quickly scanned to select those required for monitoring (by pre-scaling, selection on trigger type or error type, etc). Chosen events may be analysed on the ROD or the data may be placed in a VME-accessible buffer for local analysis by level-1 monitoring computers. This mechanism also allows complete events to be stored for collection and monitoring.

All modules so far designed include data injection memories to allow test vectors to be loaded and injected in a controlled way. The injected data can simulate the various error conditions including link loss and parity errors.

A Data Source Sink (DSS) module has been produced [7]. By re-cabling the links concerned, this module allows data insertion or checking on any link within the calorimeter trigger system, including the readout links to the RoI Builder and ROBs. The module uses the L1A signal to achieve the necessary synchronisation.

6.3 Central trigger processor

The data sent from the CTP to the RoIB consist of the following in the present design:

- The 160 bits of input data after synchronisation and alignment.
- The prescaled-clock, random-trigger, turn number and BC-mask inputs to the trigger formation logic (so-called “internal inputs”).
- The 96 trigger items before any vetoes have been applied (“TBV”).
- The 96 trigger items after all vetoes have been applied (“TAV”).
- The 96 trigger items after prescaling (“TAP”).
- The EVID and BCID numbers (in header).
- The 8-bit trigger type summary word.

The format and content of the data from the CTP are summarised in Table 2, extracted from the CTP specification document [8]. The meaning of the readout from the different stages in the processing can be understood from Fig. 13.

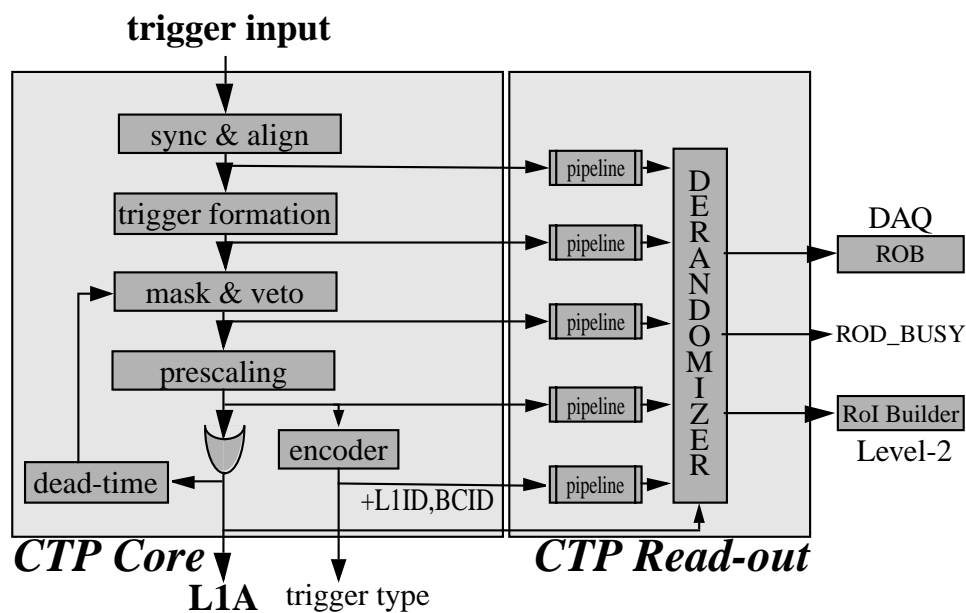


Figure 13: Block diagram showing readout from various points in the processing chain of the CTP.

??? Do we need revised figure with up-to-date view of the CTP – check with Ralf – e.g. show ???

6.3.1 Readout latency

Less than 500 ns is needed to transfer the data from one event to LVL2 assuming no queuing [1]. The CTP therefore comfortably meets the latency requirements.

Table 2: Format of data sent from the CTP to the RoIB. The terms TBV, TAV and TAP refer to the patterns of trigger items before the veto, after the veto and after pre-scaling.

32-bit Words
Input [1,32] of BC
Input [33,64] of BC
Input [65,96] of BC
Input [97,128] of BC
Input [129,160] of BC
Internal inputs (pre-scaled clock, random trigger, turn number, etc.)
TBV [1,32] of BC
TBV [33,64] of BC
TBV [65,96] of BC
TAV [1,32] of BC
TAV [33,64] of BC
TAV [65,96] of BC
TAP [1,32] of BC
TAP [33,64] of BC
TAP [65,96] of BC
Trigger_type

6.3.2 Status words

The CTP will provide the two status words described in Section 3.2. However, there is currently no expectation that the status bits will be set in the case of the CTP since it deals exclusively with fixed-length data and receives data over very simple parallel links from the calorimeter and muon trigger processors.

6.3.3 Configuration Information List

The following is a partial list of information describing the settings of the central trigger processor; more detail can be found in [9]. Trigger menu specifying the list of trigger items in terms of:

- External inputs provided by the calorimeter trigger, the muon trigger and other sources (e.g. calibration and test trigger inputs).
- Internally generated inputs to the combinatorial logic such as random triggers, prescaled clocks, bunch-crossing selection and LHC turn signals.

- Additional conditions such as prescale factors and masking of trigger items.

Note that interpretation of the external inputs requires information on the configuration of the calorimeter and muon trigger processors.

6.3.4 Error handling

The present CTP design does not include error detection on an event-by-event basis. However extensive monitoring facilities are provided (e.g. scalers and checking of results on a sampling basis).

6.3.5 Testing and debugging facilities

The CTP includes facilities to inject test data at the inputs that can be played through the system at the full clock speed. These test data can be taken from a memory that can be loaded via VME, or pseudo-random patterns can be generated. Intermediate results from the calculations performed in the CTP can be read out and checked against the expected values computed from the input data and configuration parameters. The data read out include all the information that is sent to LVL2.

7 Other issues

There are a number of issues for which it is not yet possible to go into details, but which will have to be addressed at a later stage. These include:

- Common control system for the LVL1 and LVL2 triggers. ??? for whole experiment ???
- Common monitoring system for the LVL1 and LVL2 triggers. ??? for whole experiment ???
- Common test and diagnostic system for the LVL1 and LVL2 triggers.

8 References

- [1] ATLAS first-level trigger Technical Design Report, CERN/LHCC/98-14.
- [2] S-link specification available from <http://www.cern.ch/HSI/s-link/>
- [3] “The event format in the ATLAS DAQ/EF prototype –1”, ATL-DAQ-98-129.
- [4] “Trigger and DAQ interfaces with front-end systems”, ATL-DAQ-98-103.
- [5] “ATLAS level-2 trigger User Requirements Document”, ATL-DAQ-2000-034.
- [6] “Interfaces and overlaps in the level-1 muon trigger system”, ATL-DA-EN-0001.
- [6a] ATL-DA-TR-0001.
- [7] DSS module – documentation is available from <http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html>
- [8] “CTP Technical Specification”, ATL-DA-ES-0006.
- [9] “Configuration data of the level-1 Central Trigger Processor”, ATL-DA-EP-0001.
- [10] The following documents were presented to the PDR for the ROIB that was held in February 2002:
 - Level 2 Supervisor Requirements, M. LeVine, J. Schlereth, Version 0.2, 27 Apr 2001.
<http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/DataFlow/DataCollection/docs/DC-009.pdf>
 - RoI Builder Requirements, J.Dawson, J.Schlereth, Version 0.2, 31 May 2001.
<http://atlasinfo.cern.ch/Atlas/GROUPS/DAQTRIG/DataFlow/DataCollection/docs/DC-014.pdf>
 - RoI Builder Control Software Design, J.Schlereth, Version 0.1, 20 Feb 2002.
<http://www.hep.anl.gov/reb/atlasLVL2/RoIB/RoIBControl-DD.pdf>
 - ATLAS Second Level Trigger Prototype RoI Builder, Project Specification, Version 0.5, 25 Feb 2002.
http://www.hep.anl.gov/reb/atlasLVL2/RoIB/RoIB_spec.ps
- [11] “Configuration data for the ATLAS level-1 muon trigger”, note in preparation.
- [12] “Configuration data for the ATLAS level-1 calorimeter trigger”, ATL-DA-EP-0002.

9 Update log

991007 Release of version 1.0 of document.

020307 Draft version 2.0 of document. Many changes to LVL1 calorimeter trigger part from Norman Gee and Murrough Landon. Update ROL link speed to 160 MB/s.

??? Need to fill in more complete list of changes based on change tracking ???

10 Appendix A: Muon Trigger RoI Algorithms

The Muon Trigger System consists of 3 layers of RPC-chambers in the Barrel region (RPC 1-3) and 3 layers of TGC chambers in the Endcap region (TGC 1-3). The layout is shown in Fig. A1. Similar algorithms are applied in both regions in order to find Muon candidates and to form RoIs for the LVL2 trigger.




Figure A1: Layout of muon trigger chambers.

10.1 Barrel Algorithm

RPC1, RPC2 and RPC3 all consist of chamber doublets. Muon candidates are found using the following algorithm:

- Look for hits in RPC2 (the so called Pivot-plane).
- Under the assumption that the hit in a) has been caused by a infinite momentum muon coming from the interaction point, determine the point where it would have hit RPC1.
- Look for hits in three different windows around the point determined in b). Because of the field of the toroid, each window width corresponds to a threshold in p_T .
- If hits are found in at least 3 of the 4 chamber layers in RPC1 and RPC2 in each of the two orthogonal projections, a muon candidate has been found.

- e) Similarly to the procedure above, for each candidate found in RPC1 and RPC2, define three search windows in RPC3 and look for hits therein. If a hit is in at least one of the doublet layers of RPC3 in each of the two orthogonal projections, the candidate muon is called a high p_T candidate. Otherwise candidate is called a low p_T candidate.
- f) Muons are classified in one of six p_T ranges depending on which of the 3 low p_T and 3 high p_T thresholds have been passed.

10.1.1 Granularity

The barrel muon system is divided into a positive and a negative rapidity region. Each of these consists of 32 sectors in ϕ which are further divided into either 7 or 6 pads depending on the architecture of the chamber. Each pad contains 4 possible RoI positions. The RoIs cover a region of 0.2×0.2 in the η - ϕ plane.

10.1.2 Restrictions

Due to limitations in the on-detector electronics, each sector can contain at most two muon candidates, and in each pad only one candidate can be processed. In case of more candidates being present, the highest p_T ones are processed and flags are set indicating that more candidates had been found in the sector or pad.

10.2 Endcap Algorithm

The TGCs of the end cap region are divided into three groups. TGC3 is the Pivot plane for the candidate search algorithm which is similar to the one used in the barrel. TGC2 and TGC3 are chamber doublets and TGC1 is a chamber triplet. In each chamber wires and strips are read out giving two orthogonal projections. For finding a muon candidate in TGC2 and TGC3, a 3 out of 4 majority is in each of the two projections. In order to define a high p_T candidate, TGC1 is used requiring a 2 out of 3 majority in the wire readout, and a 1 out of 2 majority in the strip readout (in TGC1 only two strip planes are read out).

10.2.1 Granularity

Each endcap is divided into two subsystems. The Forward region consists of "wheel" of TGC-chambers ranging from $\eta = 1.0$ to $\eta = 2.4$. They are divided into 24 sectors in ϕ each of which contains a matrix of 15×4 RoIs in the η - ϕ plane. The Endcap-region is divided into 48 sectors in ϕ each of which contains 38×4 RoIs in the η - ϕ plane.

10.2.2 Restrictions

Each Sector can process a maximum of 2 muon candidates. In case more candidates are present in a sector, the highest p_T ones are processed and a flag indicates the presence of non-processed candidates.

11 Appendix B: Calorimeter Trigger RoI Algorithms

This appendix provides a summary of and updates information from the level-1 TDR concerning the calorimeter trigger RoI algorithms. More detailed information may be found in Chapter 4 of the TDR, together with updates on the web at <http://hepwww.rl.ac.uk/Atlas-L1/TDR/Updates/TDR.html>.

11.1 *E/gamma and tau/hadron Cluster trigger*

The inputs to the algorithm are a set of ‘trigger towers’, of granularity 0.1×0.1 in $\eta \times \phi$. These are formed by analogue summation of calorimeter cells. There are separate sets of trigger towers from the (e.m.) and hadronic calorimeters. The cluster trigger uses data from the region $|\eta| < 2.5$, which is the limit of inner detector coverage and of high-granularity electromagnetic calorimetry.

The electron/photon and tau/hadron trigger algorithms are illustrated in Fig. B1. They are based on a window of 4×4 towers in the electromagnetic and hadronic calorimeters, and consist of six elements:




Figure B1: Trigger towers used in the electron/photon and tau/hadron triggers.

- Four 2-tower electromagnetic clusters, used to measure the E_T of electromagnetic showers.
- A square of 4 hadronic towers behind the electromagnetic clusters, which is used for isolation tests in the hadronic calorimeters.
- Four hadronic clusters summed in depth over the electromagnetic and hadronic towers, used to measure the E_T of hadronic showers. The hadronic clusters are formed

by adding the 2-tower electromagnetic clusters to the square of 4 hadronic towers centred in the algorithm window.

- A ring of 12 electromagnetic towers surrounding the clusters, which is used for isolation tests in the electromagnetic calorimeters.
- A ring of 12 hadronic towers behind the electromagnetic isolation ring, which is used for isolation tests in the hadronic calorimeters.
- A 2×2 -tower RoI cluster, summed in depth over the electromagnetic and hadronic towers, used to identify candidate RoIs.

The window slides in steps of one trigger tower in both the η and ϕ directions, and so there is one window for each tower within the acceptance of the electron/photon trigger.

The requirements for an e/gamma or tau/hadron object to be found within the window are:

- the most energetic of the four electromagnetic clusters (e/gamma trigger) or of the four hadronic clusters (tau/hadron) must pass the corresponding electromagnetic or hadronic cluster threshold;
- the total E_T in the electromagnetic isolation region must be less than the e.m. isolation threshold;
- the total E_T in the hadronic ring isolation region must be less than the hadronic ring isolation thresholds.
- the total E_T in the hadronic core isolation region must be less than the hadronic core isolation thresholds (e/gamma trigger only).
- the RoI cluster must be a local E_T maximum (see below).

If all of these conditions are met for an e/gamma or tau/hadron, then the window is considered to contain a trigger candidate. Sixteen sets of trigger E_T thresholds (combinations of cluster, e.m. isolation and hadronic isolation) are foreseen, with up to eight allocated for tau/hadron identification and the remainder for electron/photon identification.

11.2 Jet trigger

The main jet trigger algorithm is illustrated in Fig. B2. It is based on a window of 4×4 “jet elements”, which have a granularity of 0.2×0.2 in $\eta \times \phi$ and are summed in depth over the electromagnetic and hadronic calorimeters. The algorithm has two components:

- A 2×2 -element RoI cluster (0.4×0.4), used to identify candidate jet RoIs.
- A trigger cluster, used to measure the jet E_T . This cluster can be 2×2 , 3×3 or 4×4 jet elements (0.4×0.4 , 0.6×0.6 or 0.8×0.8 in $\eta \times \phi$).




Figure B2: Main Jet trigger algorithms based on a step of 0.2×0.2 . In the 0.6×0.6 algorithm, the RoI (shaded) may occupy any corner of the cluster.

The window slides in steps of one element in both the η and ϕ directions, and so there is one window for each jet element within the acceptance of the jet trigger, $|\eta| < 3.2$, which is the limit of the endcap calorimeter coverage.

The two requirements for a trigger object to be found within the window are:

- the RoI cluster is a local E_T maximum;
- the trigger cluster E_T passes the jet cluster threshold.

The optimum size of the jet cluster will depend on both the jet E_T and the luminosity. The choice of which cluster size can be used with which threshold will be programmable, and different jet cluster sizes may be used simultaneously at different E_T thresholds.

Eight main jet trigger thresholds are foreseen. Because it is possible for different trigger cluster sizes to be used for different trigger thresholds, a set of bits indicate which thresholds were passed by the RoI.

A trigger on jets in the FCAL is planned, and will be accommodated in JEMs 0, 7, 8 and 15 which serve the ends of the main jet η range in each JEP crate. Full details will be included in a future revision of this document.

11.3 Energy Algorithms

This section provides a summary of the algorithms used to compute the Jet transverse energy E_{Tj} , the total scalar transverse energy, and the missing transverse energy. Some of the details concerning energy summation and algorithms may change during engineering design of Jet/Energy modules and the common merger module.

11.3.1 Jet transverse energy

The total jet transverse energy is calculated using a simple multiplicity-based estimator. The inputs to the algorithm are the counts of the numbers of jets passing each of the eight main jet thresholds, and the values of those thresholds, together with the numbers of jets passing each of the four forward jet thresholds and the values of those thresholds. As the counts are rounded down to 7 (3 for forward jets), and the maximum incoming

calorimeter information is limited to 255 GeV per tower, the estimator may be low for events with high jet multiplicity or high jet energies.

To understand the estimator, consider an event containing n jets passing the threshold $E_T > x$ GeV and m jets passing $E_T > y$ GeV, with $n > m$ and $y > x$. There are $n - m$ jets in the E_T range $x \text{ GeV} \leq E_T \leq y \text{ GeV}$. These jets then have a total E_T of at least $(n - m)x$ GeV and at most $(n - m)y$ GeV. A number of possible estimators of the total jet E_T can be constructed in this way, such as lower limits on jet E_T (assuming each jet lies exactly at the lower end of the range, upper limits, or indeed both together).

The hardware implementation of the E_{Tj} calculation uses lookup tables to convert groups of jet counts to transverse energies. The energies are then summed and thresholds applied to yield the four E_{Tj} hit bits.

11.3.2 Scalar transverse energy

Each Jet/Energy module first sums the electromagnetic and hadronic parts of the input data to form jet elements, and then computes the scalar sum of the jet elements. The sum is transmitted to a Common Merger Module using an eight-bit compressed format with a six-bit mantissa and a two-bit exponent, which multiplies the mantissa by 1, 4, 16 or 64 to yield the full value transmitted.

On arrival at the Common Merger Module, the linear form of the summed energy is recovered, and the data from all Jet/Energy Modules are summed. The sum is compared to four thresholds with values up to 2047 GeV in 4 GeV steps.

11.3.3 Missing transverse energy

Each Jet/Energy module first sums the electromagnetic and hadronic parts of the input data to form jet elements. Components E_x and E_y are then obtained from lookup tables containing the appropriate geometrical constants. The E_x and E_y components are separately summed on each JEM, and the results transmitted to a Common Merger Module. An eight-bit compressed format data field is used for each datum, with a six-bit mantissa and a two-bit exponent multiplying the mantissa by 1, 4, 16 or 64 to yield the full value transmitted. The phi-quadrant architecture implies that the components from any one JEM always have a constant sign value, which is therefore not transmitted.

On arrival at the Common Merger Module, the linear form of the summed energy is recovered, and signed arithmetic used to sum the components from all Jet/Energy modules. A lookup table is used to perform the final quadrature addition and comparison with eight thresholds in a single step, yielding the eight hit bits.

11.4 RoI Generation

An RoI is generated for each trigger candidate. The RoI information consists of the location of the RoI, plus bits indicating which sets of thresholds it has passed. The RoI coordinate is therefore indicated at the granularity of the cluster or jet towers. Because each set of thresholds tests several variables (cluster and isolation E_T s, or jet energies and

sizes) there is no unique 'ranking' of the different sets, and so the classification word indicates all sets of thresholds passed by the RoI.

As the algorithm windows overlap, some additional logic is required to identify an unambiguous RoI coordinate. This logic identifies the window whose RoI cluster E_T is a maximum when compared with neighbouring RoI clusters. Since neighbouring RoI clusters may have exactly equal E_T , the requirement that the RoI cluster must be more energetic than its overlapping neighbours is relaxed slightly - requiring that the RoI cluster be more energetic than its neighbours along two connected edges, and at least as energetic as its neighbours along the opposite two edges (Fig. B3).

\cong	$>$	$>$
\cong	R	$>$
\cong	\cong	$>$

Figure B3: Local E_T maximum test for an RoI candidate.