ATLAS Internal Note
DAQ-NO-103
09 June 1998

TRIGGER & DAQ INTERFACES WITH FRONT-END SYSTEMS: REQUIREMENT DOCUMENT

VERSION 2.0

Atlas Trigger-DAQ Steering Group

Abstract

The Trigger/DAQ system is the central place where all signals and data coming from each sub-detector front-end electronic system will be integrated in a uniform way. This document describes the functional partitioning of the overall system and specifies the interfaces between the front-end electronics, and the trigger, DAQ and detector control systems.

Table of Contents

History

Version 1.0

Date: 22 January 1996. First version approved during Atlas week of November 1995.

Draft 2.0

Date: February 1998.

Update the dead-time value for operating at 75 kHz and 100 kHz rates.

Update TTC crate figure.

ROD_BUSY handling.

Glossary of terms

BC

Bunch Crossing.

BCID

Bunch Crossing Identifier. Number that defines the bunch crossing at which an event occurred. ÔPotentialÕ bunch crossings are numbered 0 to 3563 per LHC orbit, starting with the first following the LHC extractor gap.

BCR

Bunch Counter Reset. Signal broadcast by the TTC system once per LHC orbit (89 µs) to control the phase of the local bunch counters.

CTP

Central Trigger Processor. The place where the LVL1 trigger is generated.

DAQ

Data Acquisition System.

DCS

Detector Control System.

Derandomizer

Place where the data corresponding to a LVL1 trigger accept are stored before being read out.

ECR

Event Counter Reset. Signal broadcast by the TTC system to reset the local event counters.

Fast Controls

Control signals broadcast by the TTC system in synchronism with the LHC clock and orbit signals, such as BCR, ECR, Clear, Run, Pause, Restart, Abort, Stop, selective Resets and synchronous Test commands.

FE

Front-End.

FE Data Stream

Virtual or real link between the derandomizer and the ROD, through which the FE data are transmitted.

LCS

Local Control Station.

Level -1 Buffer

Buffer (analogue or digital) included in the FE electronics which retains the data until the L1A is received.

L1ID

LVL1 Trigger Accept Identifier. A L1ID is built at different levels of the read-out system. The TTCrx provides a 24 bit L1ID with each L1A signal. In conjunction with the BCID, it defines uniquely an event.

LVL1, LVL2, LVL3

Abbreviations for level-1, level-2 and level-3 and for the associated trigger systems.

L1A

LVL1 trigger Accept signal produced by the Central Trigger Processor (CTP) when an event has met the LVL1 trigger criteria.

Raw Data

Data provided by the front-end electronics to the read-out buffer.

ROB (Read-Out Buffer)

Standard module which receives data from the ROD, stores them and makes them available to the LVL2 trigger processors and, for LVL2-selected events, to the LVL3 trigger processors.

ROD (Read-Out Driver)

Functional element which gathers data from the derandomizers over one or more data streams and builds pieces of events to be sent to the ROB.

ROI

Region Of Interest.. A region limited in eta and phi, indicated by the level-1 trigger to contain candidates for objects requiring further computation. In the case of B-physics triggers at low luminosity, some ROI's may be defined internally within the level-2 trigger system.

ROL (Read-Out link)

Physical link between ROD and ROB through which the data are sent at the event rate of the LVL1 trigger.

Sub-detector

Physical partition of the system that corresponds to a whole element of detector using the same FE electronics (ex: calorimeter, TRT ...)

TDD

Technical Design Document.

Trigger Type

An 8-bit word transmitted with the L1A and giving information about the type of event.

TTC (Timing, Trigger and Control)

Standard system which allows the distribution of timing, trigger and control signals, using the technology developed in RD12 and described in TTCdoc. The system delivers the standard TTC signals such as the LHC clock, L1A and Fast Controls, and provides for the distribution of other detector-specific commands and data.

TTCdoc

Documentation on the TTC system. Available at http://www.cern.ch/TTC/intro.html.

TTCrx

TTC receiver ASIC which delivers the decoded and deskewed TTC signals, bunch and event counters required by FE electronics controllers.

T/DAQ

Trigger-DAQ.

1. Overview

1.1. Introduction

The main goal of this document is to define minimum requirements to:

- Allow the design of a Trigger/DAQ (T/DAQ) system that is able to run smoothly and is easy to maintain.
- Ensure easy integration of sub-detectors in the DAQ system.
- Allow the different detectors to run together with the minimum possible overall deadtime (a small local deadtime for each sub-detector level may lead to a large overall deadtime if deadtime is not handled in a uniform way).
- Allow concurrent and independent calibrations.
- Allow each sub-detector to work autonomously and concurrently for setting up.
- Allow efficient error-detection and error-recovery systems.
- Standardise within the experiment as many elements as possible.
- Optimise the LVL2 trigger tasks that retrieve data from Regions Of Interest (ROI).
- Define a global strategy for timing, trigger and control distribution.
- Enable safe and coherent operation of the detector.

This document sets out:

- Statements concerning the boundaries between the front-end (FE) and T/DAQ systems.
- Specified and justified requirements to the FE groups.
- Functional specifications of the physical interfaces between the FE systems and the T/DAQ system.
- Initial specifications of the protocols used between the FE systems and the T/DAQ system.
- Initial specifications of a minimum raw-data format.
- In a later version, specification of the interfaces between the sub-detector hardware and the Detector-Control System (DCS).

1.2. Document organisation

The ATLAS T/DAQ read-out system, the Timing, Trigger and Control system and the Detector Control system are partitioned into subsystems, typically associated with sub-detectors, which have the same logical components and building blocks. Overall standardisation of all basic common elements is the only way to minimise cost and maintenance, and evolution problems over the next 15Ð20 years corresponding to the lifetime of the experiment. Effort should be made from an early stage to break the project down into functional units and their interfaces, and to define standard common features and protocols for each logical element.

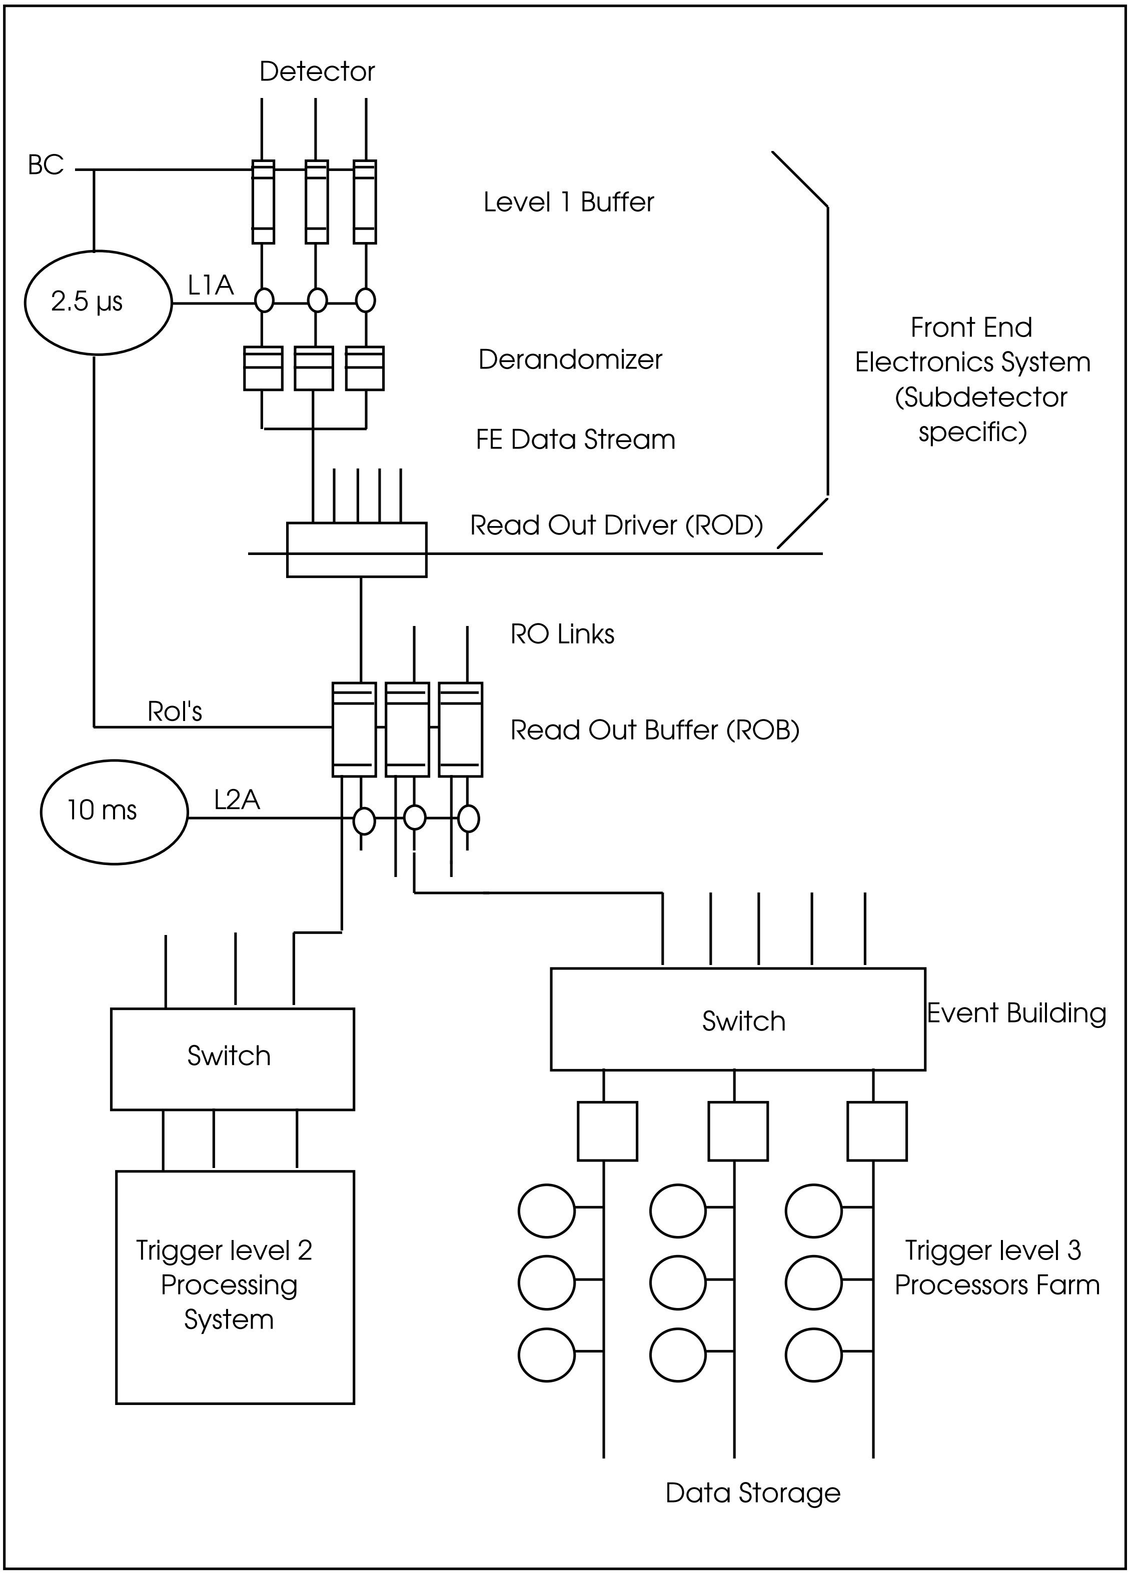

Fig. 1 Functional view of the read-out system

The Technical Proposal [1] presented the concept of the read-out system (Fig. 1). This ÒrequirementsÓ document attempts to define very clearly the interface boundaries between the T/DAQ system and the front-end electronics subsystems in order to allow completion of front-end electronics design soon, as required for some detectors. A future step will be a Technical Design Document (TDD) that will propose an implementation of the T/DAQ system to meet these requirements.

In Section 2 of this document, the different functional components of the system, from the FE up to the read-out buffer (LVL2 & LVL3 trigger buffer) will be identified, including the Timing, Trigger and Control system and the Detector Control System:

- FE electronics including LVL1 Buffer and derandomizer.
- FE data stream.
- Read-out Driver (ROD).
- RO links to Read-out Buffers (ROB).
- Read-out Buffers (ROB).
- Timing, Trigger and Control system (TTC).
- Detector Control System (DCS).

In Section 3, the main issues will be addressed:

- Partitioning and integration:

  T/DAQ, TTC and DCS systems should be partitioned in a consistent way which allows concurrent and independent calibration, stand-alone data-taking and integrated running modes.

- Physical read-out partitioning:

  When involved in the LVL1 trigger, the detector groups have to provide trigger signals directly to the LVL1 trigger system.

  When involved in the LVL2 Trigger, the detector groups have to group channels in order to match as much as possible the LVL1 trigger segmentation and associated LVL2 trigger ROIs.

- System-engineering aspects like redundancy and reliability must be considered as well as "self testability" of each element of the read-out chain.
- Error recovery:

  The overall system will contain tens of thousands of elements running at high speed. The risk of error is therefore not negligible, and a consistent error-recovery mechanism should be foreseen for all the elements to avoid having to stop data-taking each time an element fails.

- Strategy for timing in the experiment:

We need a single coherent backbone for distributing timing, trigger and synchronous control signals. The system must provide for the programmable adjustment of delay compensation elements, to allow operation with cosmic ray and calibration triggers as well as normal running.

We have to consider the practicalities of setting up the experiment.

In Section 4, the boundaries between the FE and T/DAQ systems will be stated for each element. The main considerations are:

- Detector groups have to provide input data to the ROB in a standard format, through a standard interface, and they are financialy responsible for the system up to the input of the ROB, including the link but excluding the receiver integrated in the ROB.
- Detector groups will be provided by T/DAQ with one or more standard TTC transmitter crates and interface modules based on the RD12 system [2] which will deliver the Timing, Trigger and Control signals and commands by optical fibres. The signals can be received and decoded by optical and electronic components developed in RD12 for this purpose, but detector groups will be responsible for the cost of these components and other components required for any subsequent distribution of the signals.
- Detector groups will be provided with one or more standard DCS crates. They are responsible for the interfaces to the detector control electronics.
- Detector groups will be provided with standard DAQ crates housing the ROBs and associated control processors and modules.

In Section 5, necessary requirements and interface and protocol specifications will be given. Detector responsibility for FE electronics is subject to constraints described in this document (such as pipeline length, derandomizer size, ...)

1.3. Main parameters

Some parameters that are crucial for designing each front-end electronics system are given in Table 1.

Clock frequency	40.08 MHz
Maximum LVL1 trigger rate	75 kHz upgradable to 100 kHz
LVL1 pipeline length	> 2.5 µs
Minimum interval between two L1A	125 ns (see text)
Raw data read-out time & derandomizer size	To introduce deadtime (or data loss) less than 1% @75 kHz 6% @ 100 kHz

- Table 1 Some critical FE parameters.

The following gives some explanations concerning the items in the table:

- 100 kHz is the maximum average LVL1 trigger rate, integrated over many full machine cycles, including gaps. The design of the read-out electronics has to be made for this value. However, following the descoping discussions that took place in November 1995 [5], it is acceptable if the initial system can handle 75 kHz LVL1 rate, provided it is upgradable to the full rate.
- The design of the LVL1 system aims at a latency less than 2 μs, to which a contingency of 500 ns has been added (cable-length uncertainties, alternative schemes to provide data to the LVL1 trigger, details of the electronic design, etc.). This timing includes everything up to the output of the standard TTC receivers, including cable delays back and forth between the detector and the LVL1 trigger system, under assumptions explained in Section 5.1.1. Detector groups are responsible for the computation of their required LVL1 buffer or pipeline length and it is stressed that they must foresee contingency for the part of the system under their responsibility. They should consult with the T/DAQ group before freezing the pipeline length.
- Two consecutive LVL1 triggers will be separated by at least four (4) untriggered crossings, as required by some detector groups.
- The global LVL1 trigger system will guarantee that no more than M L1A will occur in any given period, T. M and T will be adapted in order to maintain the dead-time as low as possible.
- A n% dead-time has to be understood as n% of the time lost or n% of the data lost. It includes at least the following sources:
  - 4 BC without L1A after a L1A occured.
  - Dead-time introduced by the central trigger processor (M, T algorithm).
  - Dead-time introduced at the derandomiser level.
  - Dead-time introduced at the ROD level.
- Detector groups have some freedom to define their own strategy concerning the derandomizer size and the raw data read-out speed, provided they remain within an acceptable dead time (see Fig. 6 in Chapter 5).[1]

[1]. However, it is pointed out that it is intended to contribute to the management of the derandomizers in the LVL1 Central Trigger Processor (CTP) by evaluating, on a statistical basis, the occupancy of the derandomizers. It would therefore be better to have a more-or-less unified scheme within the experiment to keep the LVL1 CTP reasonably simple and to minimise the overall deadtime. The overall deadtime will at best be the worst subsystem deadtime, but it could well be the sum of all subsystems deadtime in worst case. It is also of no use for some subsystems to spend a lot of effort and money to optimise their local deadtime if the overall one is bad. Cross-detector discussions will therefore be essential to tackle this point.

2. Identification of components

A functional view of the system from the FE to the ROB (LVL2 trigger buffer) is given in Fig. 2. This chapter defines the different components. The constraints placed on these components are set out in Chapter 5.

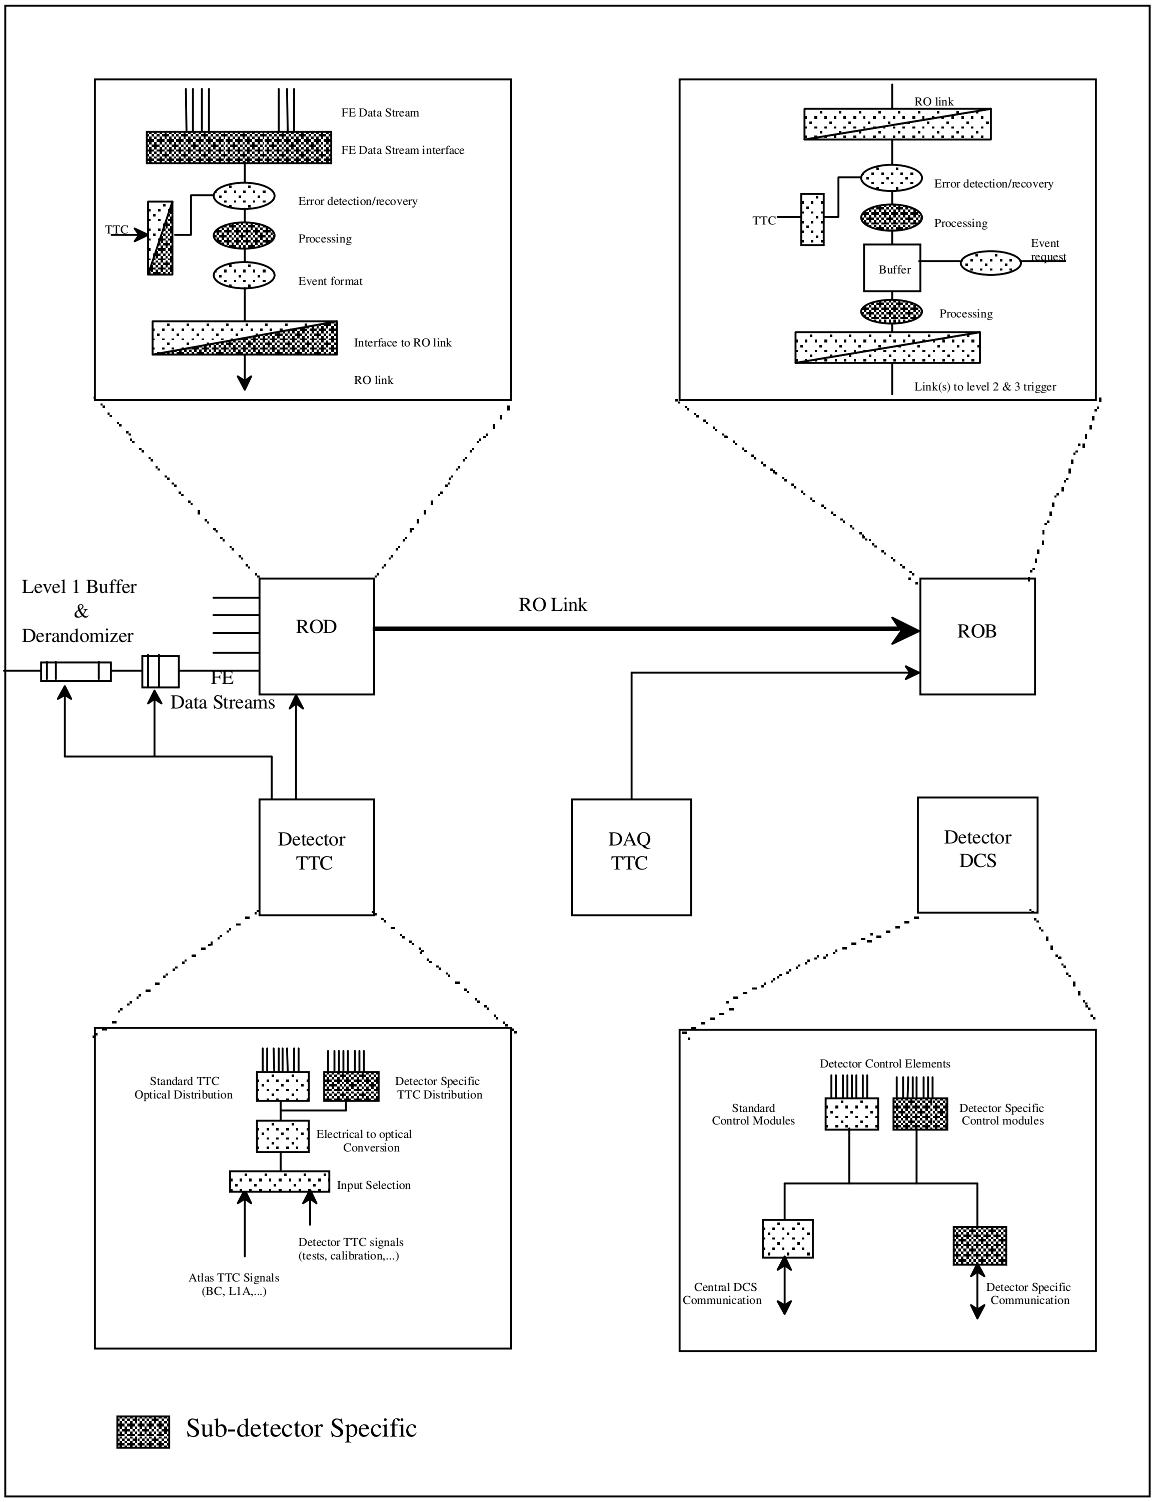

Fig. 2 Functional view of the components

2.1. Front-end electronics subsystem

The front-end electronics subsystem includes different functional components:

- The front-end analogue or analogue-digital processing.
- The LVL1 buffer in which information (analogue or digital) is stored and is retained for a time long enough to accommodate the LVL1 trigger latency. This element is subject to constraints (size, extra information, use of TTC signals).
- The derandomizer in which the data corresponding to a LVL1 trigger accept (L1A) are stored before being sent to the following level. This element is necessary to accommodate the maximum instantaneous LVL1 rate without introducing deadtime, and is subject to constraints.
- Dedicated links or buses are used to transmit the front-end data stream to the next stage. Their speed, in conjunction with the derandomizer size, is subject to constraints.

2.2. Read-Out Driver (ROD)

The ROD is the functional element of the FE system where one can reach a higher level of data concentration and multiplexing by gathering information from several FE data streams. Elementary digitised signals are formatted as raw data prior to be transferred to the ROB. The data format of the event, as well as the error detection/recovery mechanisms are subject to constraints.

2.3. Read-Out Link

Data are transmitted on this link from the ROD to the ROB at the L1A event rate. This link is subject to constraints (physical interface to ROB and protocol).

2.4. Read-Out Buffer (ROB)

The ROB includes the LVL2 buffer and receives and stores the raw data from the ROD. It performs the following tasks:

- Data storage during LVL2 and, for LVL2-selected events, at least until read-out to LVL3 is completed.
- Buffer management.
- Error detection and recovery, diagnostics and on-line monitoring.
- Data server for LVL2 and LVL3.

It may also performs local pre-processing for LVL2 data (ROIs) and LVL3 event fragments (formatting).

We aim to have a single standard ROB design for the whole experiment.

2.5. Level-1 trigger system

The LVL1 trigger function is provided by a 40 MHz synchronous system that receives directly from the front-end electronics dedicated coarse data, organised and segmented in an appropriate way. These data are extracted from the relevant FE electronics systems and sent to the sub-trigger processors through dedicated LVL1 links. The functions associated with the LVL1 trigger are:

- LVL1 trigger processors for the calorimeter and muon systems.
- LVL1 central trigger processor including deadtime logic.
- Connection to TTC system.

These elements are subject to constraints.

2.6. Timing, Trigger and Control (TTC) system

The TTC system interfaces with the LVL1 CTP, DCS (in some cases), LVL2 and/or DAQ systems and the LHC machine and provides timing, trigger and control signals to the FE Electronics Systems, RODs, LVL1 trigger processors and ROBs. It delivers all the fast signals which are synchronised to the LHC machine clock and orbit phase [3], and which have to be compensated for different detector, electronics and propagation delays. These include the LHC clock, L1A, BCR, ECR, selective resets and synchronous test commands. The system can also transmit asynchronous controls and data such as trigger-type information and individually-addressed channel enable, deskew-adjust and calibration parameters.

2.7. Detector Control System (DCS)

The DCS will be the tool to operate all sub-detectors in a homogenous way and to monitor continuously all their operational parameters. It will also supervise the common experimental infrastructure. Both functions are needed, not only for the safety of the detector, but also to ensure the quality of the physics data. The DCS will span the full range of functions from global operations performed by the shift operator to detailed interactions done by the equipment expert.

A detailed description of the DCS will be provided. A "Detector Control System User Requirements" [4] is under preparation.

3. Main issues

3.1. Partitioning and Integration

Before having the whole T/DAQ system working, each detector group will have to make autonomous tests, and to integrate their read-out system with the global one. Later, during physics running periods, detector groups will still need to work independently and concurrently, for instance in order to calibrate. This requires that:

- The DAQ is partitioned in a coherent way and that the insertion or removal of a subsystem from the global DAQ should be performed in a non-destructive and easy way.
- The TTC system is partitioned in the same coherent way within the experiment. The use of the same TTC backbone by each sub-detector will help this operation.
- The DCS is partitioned in the same coherent way.

To achieve the above, we must:

- Define clearly the interface between the sub-detector read-out systems and the global DAQ, by making available a standard ROB.
- Define and make available a standard TTC system.
- Define and make available a standard DCS system.
-

3.2. Physical read-out partitioning

Both the LVL1 and LVL2 triggers will have to deal with data corresponding to well-defined physical segments of the detector (towers, regions of interest). Although one could imagine having a system able to access pieces of data distributed in an arbitrary way over a large number of ROBs, it seems extremely important to minimise the number of locations in which the required data are located. This will facilate the data collection and make the system cheaper and more reliable.

As a consequence, it is necessary for both LVL1 and LVL2 trigger that the data are grouped in order to match the trigger segmentation of ROIs or logical subsets of them, as far as is possible.

The overall number of ROBs also has to be minimised. This requires that as many channels as possible are gathered at the ROD level before being sent to the ROBs. This implies the use of fast ROLs.

Matching the trigger segmentation and minimising the number of ROBs might be antagonistic for some sub-detectors (for instance RPCs). In this case, priority should be given to the trigger segmentation matching if the sub-detector participates to LVL2 trigger.

3.3. Error detection and error recovery

We are facing a system composed of tens of thousands of elements running at high speed and requiring thousands of links to read out the data. Similarly, timing signals (BC, ...) and trigger signals (L1A) are broadcast at high rate to a large number of destinations. The probability per unit time to get a failure, either in the sub-detector electronics or in distributing the timing and trigger signals, may be rather high compared to today's experiments. As we cannot stop the data taking each time an element fails, one has to foresee an error detection and recovery mechanism at a very early stage in the read-out chain.

One can distinguish several kinds of errors: FE errors (e.g. a PLL which unlocks), data transmission errors and synchronisation errors (e.g. missing or extra L1A signal at the FE level).

For the FE errors, one has to:

- Transmit error flags with the data, so that each level of the read-out chain can detect and flag the presence of an error and take the appropriate action. The range of error protection may be different from one sub-detector to an other. The level of error detection is a trade-off between safety on one hand and extra complexity in the FE or extra bandwidth on the other.
- Foresee the possibility to reset the electronics locally in order to restart in proper conditions, without making a general reset of the whole experiment.
- For FE synchronisation errors, one has to:
- Transmit with the raw data a unique identifier (BC number, LVL1 number) which allows a check to be made of the alignment of the data whenever different pieces of an event are gathered (for the same trigger number, one should have the same BC number).
- Automatic local recovery is required for simple errors.

3.4. Initialisation and reset

In general, initialisation of the front-end electronics needs a non-negligible time. It has to be done at the start of a run in a fully-handshake way (one has to wait for the completion of the initialisations before starting the generation of triggers).

It may also happen that an error in part of a sub-detector FE system requires a reset to recover. In this case, the reset has to be local and should not disturb the data taking of the other parts of the sub-detector.

For sub-detectors requiring frequent FE-reset signals, a global mechanism for generating regular resets is foreseen: at periodic intervals (frequency of the order of 1 Hz to be defined) a time window (length to be defined) will be provided by the CTP, during which there will be no L1A signals. This will allow one to empty the LVL1 buffers and derandomizers, and to reset the FE electronics without disturbing the data taking. A broadcast command will be sent by the TTC system at the beginning of this time window.

3.5 Timing of the experiment

Facilities have to be provided to adjust the timing signals at several levels[2]:

- Master phase adjustment relative to the LHC orbit.
- Fine deskew of the BC clock phase at all destinations.
- Global relative timing adjustment between BCR and L1A to associate the trigger decision with the appropriate BCID.
- Timing adjustment of BCR and L1A (together) to compensate the delays due to time-of-flight, detector, electronics and signal-propagation delays. This will also be needed to allow cosmic ray runs in order to take data, align the detector, etc., before beams are available.
- Separate timing adjustment for fast controls to allow the generation of test signals with different delays.

Each detector group has to develop a strategy which will allow the adjustment of the timing (within their detector and relative to the other detectors) for cosmic ray, calibration/alignment and physics runs. The practicality of this strategy is crucial.

3.6. Calibration and monitoring aspects

These tasks can be done at the ROD and/or ROB level. Spying on the data at the ROD or ROB level during a run may be possible, but only under the condition that it does not significantly increase read-out dead time or latency. The trigger type might be necessary at the ROD level for signal conditioning in some case, and will be made available by the TTC system.

It is desirable to have monitoring capability at the ROD level to allow running independent of the main DAQ system (especially during the integration phase or for debugging purpose). This does not imply to develop another data acquisition system but to foresee access to the data at the ROD level. Monitoring at the ROB level is also desirable as one has access to more information (e.g. LVL2 trigger decision).

2. The standard TTC system provides the facilities that allow these timing parameters to be adjusted at the level of the outputs from the timing receiver (TTCrx) chips. Detector groups must provide any additional delay adjustments necessary (e.g. to compensate differential delays in a subsequent detector-specific distribution scheme). The subaddress/data transmission facility of the TTC system may be used to control these additional adjustments.

3.7. Engineering issues

We will have to produce a very large number of electronics elements and maintain them for a long time. Furthermore, it is very likely that we will face upgrades during the lifetime of the experiment. The system has to be flexible, modular and sufficiently decoupled enough to allow easy upgrades of different parts.

It is very desirable to standardise elements as much as possible in the experiment; this will reduce the design, production and maintenance cost.

Accessibility is an important point which should be kept in mind.

Testability of each elements of the read-out system has to be foreseen.

4. T/DAQ Front-end Boundaries

This part of this document defines the boundaries between the FE and T/DAQ systems in terms of responsibility (system design, cost accounting, maintenance and operation) for the different components of the system.

4.1. Front-end electronics system

4.1.1. Level-1 Buffer

Detector groups are responsible for providing the buffer. The requirements given in Chapter 5, concerning its length or depth and the information to be added to the raw data, must be satisfied.

The T/DAQ group is responsible for making available to the detector groups the standard TTC signals (BC, L1A and fast control signals) from which the signals necessary for the operation of the buffer can be derived. The L1A signal from the CTP will be valid a fixed number of bunch crossings after the interaction occurred and there will be no significant fluctuation in this timing after the experiment has been set up.

The TTC system can be used by the detector groups to adjust the timing to match their LVL1 buffer length.

4.1.2. Derandomizer

Detector groups are responsible for providing the derandomizers. The requirements given in Chapter 5 concerning the maximum deadtime introduced at the derandomizer level must be satisfied.

The T/DAQ group is responsible for introducing two untriggered bunch crossings after a L1A, as requested by some detector groups.

The T/DAQ group is responsible for introducing deadtime when derandomizers are estimated to be nearly full. This means that the Central Trigger Processor runs an algorithm to estimate the occupancy of the derandomizers. This is done on a statistical basis only and detector groups have to protect against overflow, especially if they perform zero suppression or data compression before the derandomizer stage. In addition, the CTP will guarantee that no more than M L1A signals will be issued within any given period, T.

As mentioned earlier, cross-detector discussions will be necessary to define a uniform strategy at the derandomizer level in order to obtain the lowest possible overall deadtime.

4.1.3. Front-end data stream

Detector groups are responsible for the transmission of the data from the FE electronics to the rest of the system. The speed of the data stream, in conjunction with the derandomizer size, have to be sufficient to fulfil the requirements concerning deadtime given in Chapter 5.

4.2. Read-Out Driver

Detector groups are responsible for providing this part of the system. A local error-detection and error-recovery mechanism has to be implemented in order to avoid data-taking blockages.

Detector groups are responsible for providing a BUSY signal to force the LVL1 Central Trigger Processor to introduce deadtime when necessary. The time necessary for this BUSY signal to reach the LVL1 CTP may be long (e.g. if the RODs are far away from the LVL1 CTP) and several L1A signals can still be issued after the ROD asserts it.

Detector groups are responsible for formatting the data sent to the ROB, as specified in Chapter 5.

4.3. Read-Out Link

Detector groups are responsible for providing the RO link. It has to comply with the specifications given by the T/DAQ group. A standard RO link will be proposed.

4.4. Read-Out Buffer

The ROB is the first place in the read-out where full uniformity and standardisation can be achieved across the whole detector. This is why this function should be separated from the Front-End Electronics System. Another advantage of this separation is increased flexibility in the design of higher-level event-selection schemes:

- By allowing later decisions than for the FE electronics and as a consequence, allowing the use of better technologies.
- By allowing changes in LVL2 latency and later upgrades without interfering with the FE electronics.

The T/DAQ group is responsible for providing the standard ROB of ATLAS.

Detector groups are expected to contribute, within the T/DAQ project, to the specifications and the design of the ROB to make sure that their requirements are satisfied.

4.5. Timing, Trigger and Control system

For each detector the T/DAQ group will provide and be responsible for maintaining a standard TTC system based on RD12 technology. The system will make available the timing, trigger and control signals of the experiment through an optical distribution system. It will comprise the required number of TTC laser transmitters, and TTCvi modules.

The system allows detector groups to run in stand-alone mode for development, test, calibration runs, etc. It also has internal clock and orbit generators to allow operation in the development phase when the LHC signals are not available.

The T/DAQ group will be responsible for the cost up to the destination ends of the standard TTC system optical fibre distribution network. The detector groups will be responsible for the cost of the opto-electronic receivers and TTCrx ASICs (or equivalent) for subsequent decoding of the signals, and their integration in the FE electronics.

The point at which conversion is made from the standard TTC distribution system to any additional detector-specific one is to be decided by the detector groups. Any detector-specific system must satisfy the requirements defined in Chapter 5.

Detector groups are encouraged to discuss their requirements with the T/DAQ group, so that the standard system can be optimised.

4.6. Detector Control System

The T/DAQ group will provide each sub-detector with one or more Local Control Stations (LCS), each of which consists of a crate with processor and network interface and an operator console. The skeleton software and general services (configuration data base, operator interface, alarm system, diagnostics system, data base for data logging, logging of interactions and incidents, etc.) will also be provided. Sub-detector groups can be supplied (at their expense) with standardised modules for the LCS crate (e.g. fieldbus interface, multiplexed ADC, I/O units) together with the corresponding software.

Sub-detector groups have to provide the modules for the LCS crate, interfacing to their specialised hardware (front-end electronics, gas system, calibration system). They are responsible for the corresponding software.

5. Requirements and rules

This chapter gives rules to be followed at each level of the read-out system, from the FE electronics system up to the ROB.

5.1. FE Electronics System

5.1.1. Level-1 Buffer

Rule 1.1.

The LVL1 buffer has to be big enough to accommodate the maximum possible LVL1 trigger latency. If it is implemented as a pipeline, the information from the detectors has to be stored at the 40 MHz clock rate (or a multiple of it) until the LVL1 trigger decision is available.

Remark.

Where LVL1 pipelines are longer than the LVL1 trigger latency, the L1A signal can be delayed to match the pipeline length. Provision is made for such delays in the TTCrx ASIC.

Rule 1.2.

It is the responsibility of the sub-detector groups to calculate the required size of the LVL1 buffers. They must consult with the T/DAQ group before freezing the buffer size, to allow a final check that the LVL1 latency, including the part under the responsibility of the sub-detector group, has not been underestimated. The following notes are given to help the sub-detector groups in calculating the buffer size; the T/DAQ group should be consulted for clarification and advice, where required.

Notes.

The LVL1 trigger latency has been computed on the following basis:

- The target latency for the LVL1 trigger is 2 μs[3], to which 500 ns of contingency has been added to allow for changes in cable routing, alternative schemes for providing data to the LVL1 trigger, details of electronic design, etc. For sub-detectors using the standard TTC distribution system and the shortest cables possible for returning the L1A signal to the FE electronics, this gives a maximum latency 2.5 μs.
- The quoted latency is measured from the (average) time when protons collide. It includes all delays up to the output of the standard TTCrx ASIC, including propagation delays on cables and fibres between the CTP and the on-detector

[3]. In the Technical Proposal, a longest fibre length of 75 m from the CTP via the TTC to the standard TTC receivers was assumed. Correcting for an updated cable-length estimate of 81 m and based on an average propagation delay of 5 ns per meter, the present best estimate of the latency is 2.06 μs assuming an independent trigger ADC system for the LVL1 calorimeter trigger.

electronics. Fig. 3 shows the different parts of the LVL1 trigger processor and TTC distribution system.

- In the following, a number of items are listed which lie outside the control of the T/DAQ group, and which are therefore excluded from the above latency calculation.
- The quoted latency DOES NOT INCLUDE ANY EXTRA DELAY INTRODUCED TO DISTRIBUTE THE L1A SIGNAL BY MEANS OTHER THAN THE STANDARD TTC SYSTEM.
- The quoted latency DOES NOT INCLUDE DELAYS WITHIN THE ON-DETECTOR FE SYSTEMS BEYOND THE OUTPUT OF THE TTCRX ASIC.
- The quoted latency assumes that the cables returning the L1A signal to the on-detector electronics follow the shortest possible path. It DOES NOT INCLUDE ANY EXTRA DELAYS INTRODUCED BY LONGER CABLING ROUTINGS DUE TO SPECIAL IMPLEMENTATIONS (E.G. RETURN PATH VIA REMOTELY-LOCATED RODS).
- The quoted latency DOES NOT INCLUDE THE NEED FOR SOME SUB-DETECTORS TO RECORD SAMPLES PRIOR TO THE BC WHICH GAVE RISE TO THE LVL1 TRIGGER.
- In some implementations of LVL1 buffers, provision is made to remove defective cells by programming the buffer manager. In such cases, the number of cells required should be calculated so that the buffer size meets the latency requirement after the expected (worst case) number of defective cells has been subtracted.

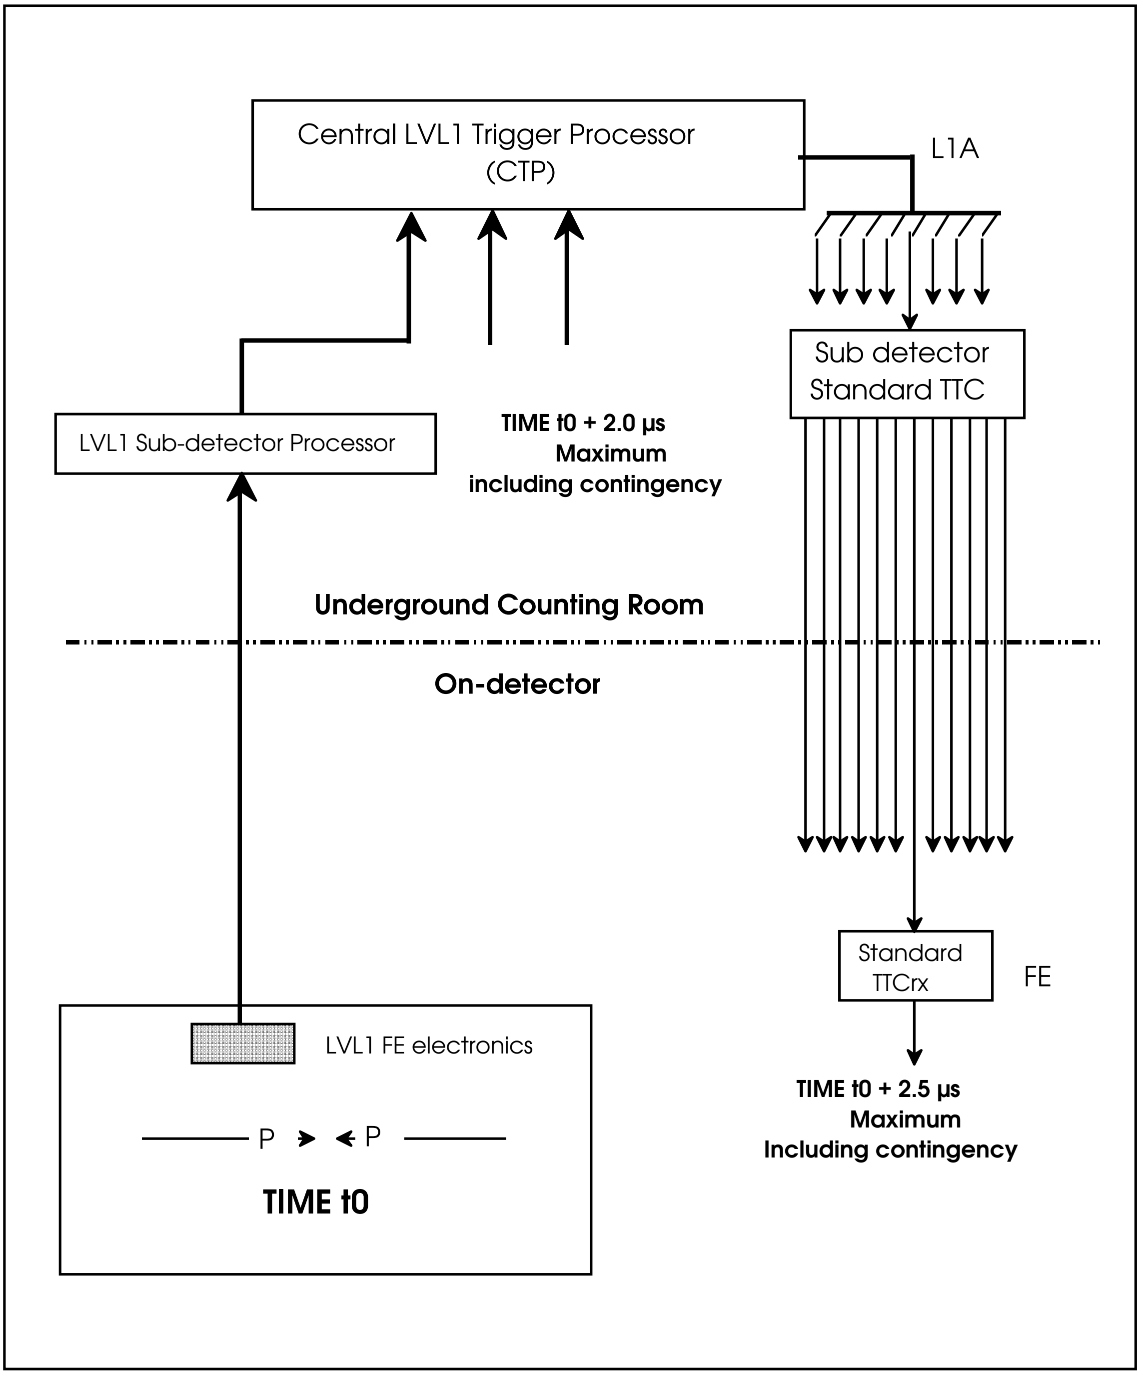

Fig. 3 LVL1 trigger latency shown for a sub-detector using the standard TTC components up to the FE electronics level

5.1.2. Derandomizer size

Rule 2.1.

When a L1A signal occurs, the data stored in the LVL1 pipeline, corresponding to the bunch crossing which produced the trigger, have to be transferred to the next level of buffering (ROD). Due to the random time of arrival of the L1A signal and the limited transmission speed, a derandomizer stage has to be foreseen. The size of the derandomizer and the speed of the FE links have to be chosen in order to limit the deadtime introduced at the level of the derandomizer to a value less than 1% in the worst case for the sub-detector as a whole with a L1A rate of 75 kHz. When running at 100 kHz, it is acceptable to introduce up to 6% of dead-time (or data loss). Fig. 4 gives the percentage of events lost as a function of these two parameters, calculated for 100 kHz average L1A rate. As mentioned earlier, it is important to have some uniformity in the derandomizer stage in order to minimise the overall deadtime. Cross-detector discussion will be necessary before the derandomizer designs are frozen.

Rule 2.2.

If zero suppression or data compression is performed before the derandomizer stage, possibly huge fluctuations in event sizes have to be foreseen. Detection and flagging of possible overflows should be considered.

Rule 2.3.

The occupancy of the derandomizers has to be estimated in order to limit losses of events when they become full. The traditional way of asserting a BUSY signal to force the CTP (LVL1 Central Trigger Processor) to introduce deadtime is not practicable in ATLAS because of the distance between the FE electronics and the CTP and the high LVL1 trigger rate: several hundred ns typically would be necessary for this BUSY signal to reach the CTP; during this time several L1A could still arrive. Such a scheme would therefore require large derandomizers or introduce extra overall deadtime.

In view of the above, the occupancy of the derandomizers is monitored on a statistical basis in the CTP. Deadtime necessary to limit the loss of data is introduced in the CTP.

Remarks.

In order to limit the complexity of the CTP and the overall deadtime, a uniform approach (derandomizer size, read-out speed) between the sub-detectors is desirable[4].

---

[4]. The evaluation of the derandomizer occupancy could also be done at the ROD level, but here again the time necessary to transmit a BUSY signal to the CTP could be excessive.

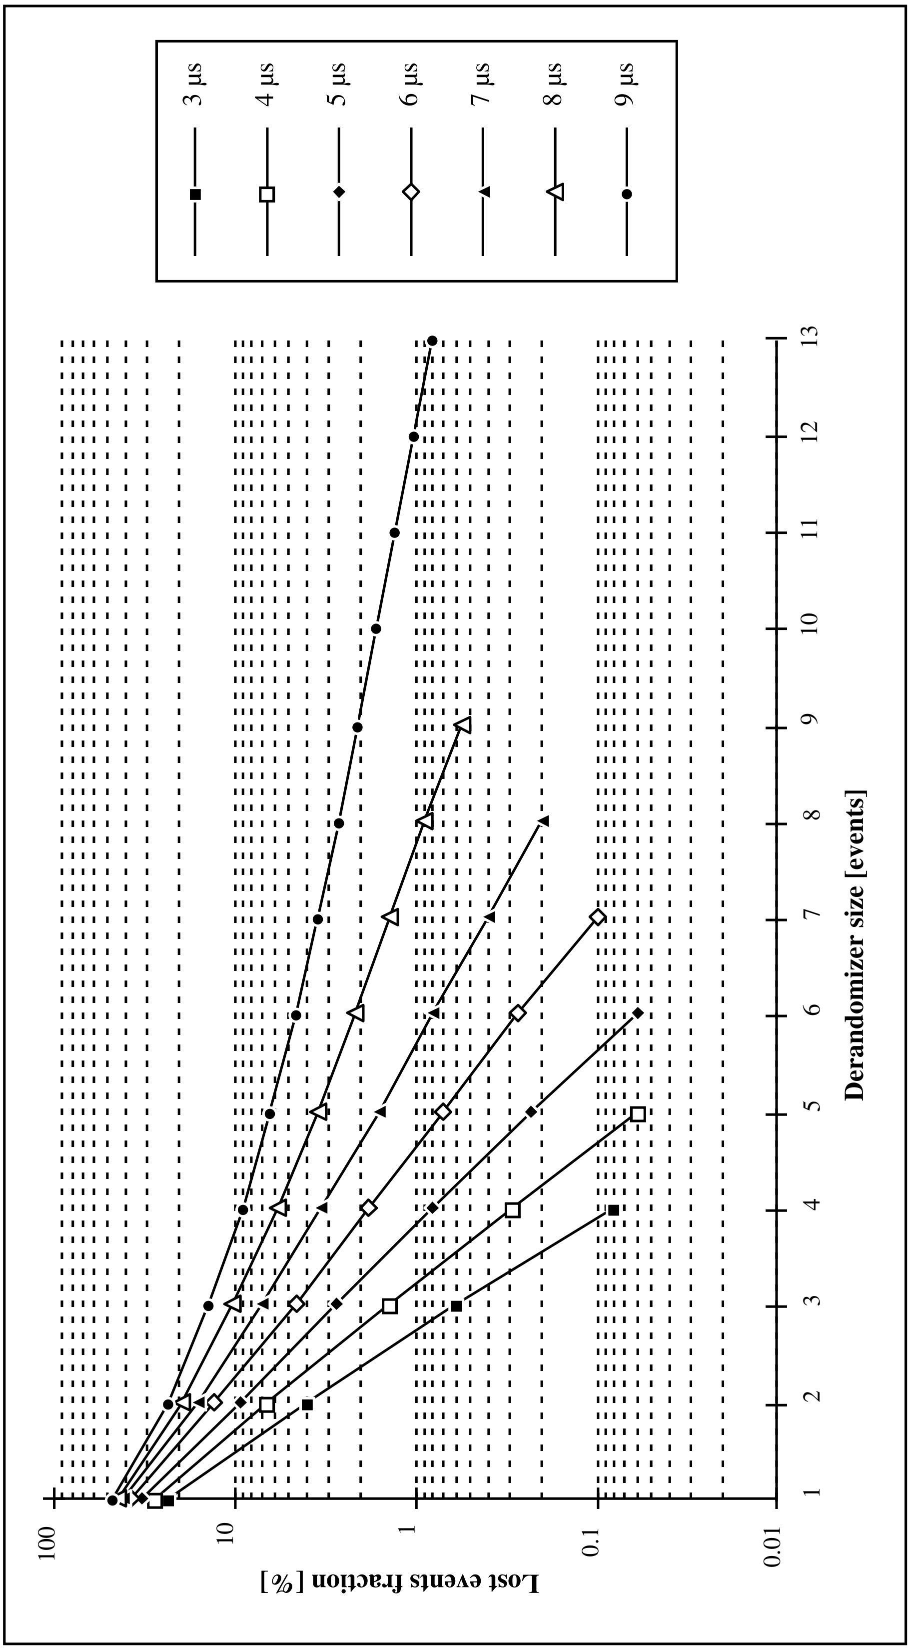

Fig. 4 Percentage of events lost versus derandomizer size and read-out speed, for an average L1A rate of 100 kHz and an exponential distribution of the time between two consecutive L1A signals.

5.1.3. Bunch-crossing and trigger identifiers in FE electronics

It is planned to identify uniquely the events with two identifiers: the Bunch Crossing Identifier and the L1A identifier. The BCID is a 12 bit number which gives the BC number within an LHC turn at which the event occured. The L1ID is a 24 bit number which gives the L1A number. From the ROD level, the full length BCID and L1ID must be provided. It is possible to have shorter BCID and L1ID at the FE level.

Rule 3.1.

In addition to the raw data, sub-detector systems will have to provide a unique bunch crossing number with absolute meaning to allow subsequent error detection and recovery.

This FE_BCID number should be a 12-bit number (which covers a full machine turn), output of a BC_Counter incremented by the BC signal and reset by the BCR signal. The BCR signal has to be adjusted in time to get the correct absolute meaning of the FE_BCID.

Remark.

Although the introduction of the full 12-bit FE_BCID at the input of the FE buffer is the safest way to detect alignment errors, it could be that the additional cost to the FE electronics is deemed too high. In this case, a FE_BCID with a smaller number of bits could be used initially. In this case the full FE_BCID must be formed at the derandomizer level or, at the latest, at the ROD level. The probability to have wrongly synchronised events that are not detected has to be kept at an acceptable level (to be defined).

Rule 3.2.

In addition to the raw data and the FE_BCID, sub-detector systems must add a FE_L1ID number (event number formed by counting L1A). This FE_L1ID must have a minimum number of bits and must be resettable. A minimum of 4 bits seems reasonable.

Remark.

Although the introduction of a FE_L1ID number at the input of the derandomizer is the safest way to detect alignment errors, it could be that the additional cost to the FE electronics is deemed too high. In this case, the full FE_L1ID must be formed at the ROD level.

5.2. Read-Out Driver

Rule 4.1.

The ROD receives data from derandomizers at the L1A event rate. Whenever possible, the ROD output should match the ROI segmentation or a subset of it; the sub-detector groups must discuss with the T/DAQ group on this point.

Rule 4.2.

The ROD receives all the standard TTC signals including ROD_L1ID (event number) and ROD_BCID (BC number corresponding to L1A), and the trigger type. ROD_L1ID must have 24 bits to accommodate about 160 seconds of running with a unique ROD_L1ID per event (sufficient to identify uniquely any event in the T/DAQ system). ROD_L1ID must be resettable. ROD_BCID must have 12 bits to accommodate a complete machine turn with a unique bunch crossing number. ROD_BCID should be reset each time BCR occurs so that it has an absolute meaning. If the ROD uses the RD12 TTCrx ASIC to decode the TTC signals, these two pieces of information are available.

Rule 4.3.

When gathering data coming from different derandomizers, the ROD must check that the FE_L1ID numbers are consistent with each other and with the low-order bits of ROD_L1ID. It must also check that the different FE_BCID numbers are identical with each other and are consistent with the ROD_BCID number.

If a mismatch occurs, the ROD must flag the event with an error code. It must then either try to recover or continuously send error-flagged events until a FE_Reset occurs. It must not stop the data taking, as it might be that only a small part of the data are bad.

Rule 4.4.

From the different pieces of event coming from the derandomizers, and after processing (if any), the ROD must build an event fragment including at least the following information:

ROD_L1ID	24 bits
ROD_BCID	12 bits
ERROR/STATUS	
ROD_MODULE_ID	
DATA	

ROD_MODULE_ID is a module identifier which is unique and which could be used to check connections. The content will be defined later (it could be for instance Type Identifier and Serial Number).

ERROR/STATUS is used to flag error conditions and gives status information such as:

- Synchronisation error detected.
- FE error detected.
- Physics data.
- Fake data.
- Empty event.
- ...

Standard codes will be defined later.

This event format does not include the event separator, which will be defined later. Even if there is no data for a given event, the other information must be sent.

A complete event format cannot yet be fully defined. It will be done in due course.

Rule 4.5.

The event data have to be sent to the ROB through the RO links at the L1A event rate (i.e. all the data of all the events have to be transmitted to the ROB). The event ordering has to be maintained to facilitate event handling and checking in the ROB. An out-of-order arrival of events would require, for instance, a time-out mechanism to detect missing events; although feasible, it would complicate the input part of the ROB.

Rule 4.6.

Due to the limited speed of the RO link, a buffer could be necessary before the link; this buffer must be protected against overflow. If this buffer is nearly full, the ROD has to provide a ROD_BUSY signal in order to stop L1A generation. A global ROD_BUSY (logical OR of the ROD_BUSYs) per sub-detector has to be provided to the CTP. The CTP will monitor the global ROD_BUSY signals it receives. Individual ROD_BUSY signals have to be monitored by the sub-detector systems and must be resettable and maskable in order to avoid data-taking blockages.

The transmission of the ROD_BUSY to the CTP will be specified in due course (electrical specifications, connectors, cables,...). A standard module handling several ROD_BUSYÕs, allowing their control (status, enable/disable) and monitoring (duration) and their transmission in a hierarchical way will be made available to the sub-detector groups (at their expense).

It is the responsibility of the sub-detector groups to define the time at which the ROD must assert ROD_BUSY. They must take into account the time for this signal to reach the CTP, the number of L1A which can occur during this time and the number of events already stored in the derandomizers.

Rule 4.7.

A test mode where fake events with known data (either fixed pattern or known pattern loaded by some means) are produced, has to be foreseen in order to be able to test the read-out chain from the ROD level up to ROB.

Rule 4.8.

If access to the data by means other than the normal read-out is possible (for monitoring for instance), it should not introduce deadtime or additional latency. It should not interfere with the main read-out (by freezing part of the buffers or slowing down the read-out, for instance).

Rule 4.9.

During the initialisation phase (before the start of data taking), the ROD has to assert ROD_BUSY until the initialisation is completed, in order to start the data taking only when the initialisation is finished.

A full handshake mechanism will be defined.

Rule 4.10.

If, because of an error in the FE electronics connected to the ROD, a local reset has to be generated, the ROD should not assert the ROD_BUSY during this operation. It should produce empty events with the proper error flag until it is able to re-establish synchronisation and send normal data.

5.3. Read-Out Link

Rule 5.1.

A standard RO link specification, with which the sub-detector groups have to comply, will be provided by the T/DAQ group.

Rule 5.2.

The RO links are used to transmit the data from the full sub-detector for all the events from the RODs to the ROBs at the L1A event rate. The bandwidth must be large enough to accommodate the average data rate expected for the maximum LHC luminosity and a L1A rate of 100 kHz. Sufficient safety margin must be included for uncertainties in the occupancy (taking into account the protocol overhead).

Remarks

This link must include a test mode in order to check it easily in situ.

We expect to implement an XON/XOFF mechanism on this link, so that the ROD knows when the ROB occupancy is too high and can assert its ROD_BUSY. If this mechanism were not included, the ROB would have to assert its own BUSY in order to introduce deadtime, at a time which is very dependent on the ROD structure (buffer size, etc.); this might harm the concept of having a standard ROB module.

Rule 5.3.

The output port of this link (input of the ROB) has to comply with the ROB specification.

Rule 5.4.

If it is foreseen that the ROB will accept non-event data (e.g. calibration data), the bandwidth of the link must accommodate the possibly large data-transfer rate.

5.4. Read-Out Buffer

Rule 6.1.

A standard ATLAS ROB will be provided. The functional specification will be described in a separate document.

Rule 6.2.

The raw data of a given event fragment should be available in each ROB within a reasonable time following the L1A signal, compatible with the overall L2 latency (typically the delay should not exceed 100 μs). The format of raw data must follow the standard rule 4.4.

Remarks

The ROB will be located on the ground floor and be physically decoupled from the FE electronics.

All the data of all the events are received at the L1A event rate. The buffer size has to accommodate the maximum LVL2 trigger latency and LVL3 trigger data read-out latency for at least a few % of the events for the full LHC luminosity and L1A rate of 100 kHz. Sufficient safety factors must be included for uncertainties on the LVL2 latency and rejection factor[5], and for the time needed for the data to be sent to the LVL3

5. The Technical Proposal assumes a LVL2 rejection factor of 100. This means that only 1% of the events have to be retained in the ROB. Nevertheless it is necessary to foresee a reasonable safety margin.

trigger processors. In particular, it has to be foreseen that during B-physics data taking the number of events which have to be retained until the end of LVL3 trigger read-out may be higher (possibly 10% of the LVL1-accepted events).

The ROB receives all the TTC signals and maintains a table of expected events with a ROB_L1ID number (24 bits) and a ROB_BCID number (12 bits) in order to check the events coming from the ROD.

In case of an error it should follow the same rules as the ROD.

Communication with the LVL2 and LVL3 trigger processors (event requests and clear, data links,...) will be subject to specifications (to be defined) which will have to be followed.

A test mode, in which fake events with known patterns (either fixed or loaded by some means) are generated, must be foreseen in order to test the read-out chain from the ROB level upwards.

If access to the data by means other than the normal read-out is possible (for monitoring for instance), it should not introduce any significant deadtime or additional latency and should not interfere with the main read-out (by freezing part of the buffers or slowing down the data transmission to the LVL2 and LVL3 trigger processors, for instance).

If the RO link does not implement an XON/XOFF mechanism, the ROB has to produce a ROB_BUSY signal when the buffer is nearly full, so that the CTP can introduce deadtime. A global ROB_BUSY (logical OR of the ROB_BUSYs) per sub-detector has to be provided to the CTP. The CTP will monitor the global ROB_BUSY signals that it receives. Individual ROB_BUSY signals have to be monitored by the T/DAQ system and must be resettable and maskable in order to avoid data-taking blockages.

5.5. Timing, Trigger and Control system

Rule 7.1.

One or more standard TTC systems (transmitter crates, modules, optical distribution network) will be provided by the T/DAQ group for each sub-detector group. The system (see Fig. 5) delivers through optical fibres the standard TTC signals such as the LHC clock, L1A and fast control signals. Some of the timing signals are shown in Fig. 6 and LHC timing parameters are summarised in Table 2. Further details are given in [3].

The system will also provide for the distribution of sub-detector-specific commands and data provided by the sub-detector groups.

The system has internal timing generators to allow test running when the LHC machine signals are not available, and can accept a local trigger or trigger emulation source instead of L1A.

Rule 7.2.

The following signals will be available from a standard TTC timing receiver ASIC (TTCrx) or equivalent:

- 40.08 MHz clock of programmable phase, derived from the LHC machine clock. A non-deskewed clock is also available for processor use.
- BCR Ð Bunch Counter Reset. Signal broadcast once per LHC orbit (88.924 μs) to reset the local bunch counters. By means of a programmable delay, the BCR timing is adjusted to match the bunch counter phase to the arrival of the corresponding sub-detector data at the output of the pipeline, compensating for the different particle flight times and detector, electronics and propagation delays.
- L1A Ð Level-1 Trigger Accept. The LVL1 CTP generates the L1A signal a fixed time interval after the interaction occurred, for events passing the LVL1 trigger criteria. The TTC system broadcasts the L1A and BCR signals with a relative timing such that, when they are deskewed by the same delay, L1A is delivered by the TTCrx concurrently with the BCID to which it corresponds.

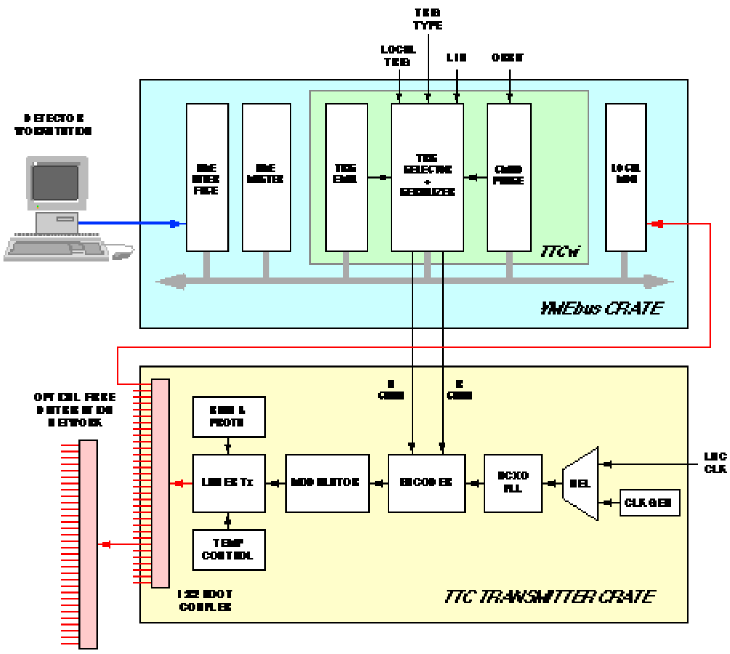

Fig. 5 TTC transmitter system components

- ECR Ð Event Counter Reset. Signal broadcast to reset all event counters.
- Synchronous Command Signals. These fast control signals are broadcast by the TTC system to control the running of the experiment (Clear, Run, Pause, Restart, Abort, Stop, selective Resets, window for FE resets, etc.). Once all the required commands have been established a uniform set will be defined for use by all sub-detectors. This facility also provides for the broadcasting of synchronous commands to sub-detector-

specific test generators. The commands are deskewed by a separate delay from that of L1A and BCR.

- Asynchronous commands and data. For the transmission of individually-addressed slow controls and data such as channel-enables, calibration parameters or deskew values for additional programmable delays external to the TTCrx. Also used for internal TTCrx tests, control and deskew programming.
- Error. TTC signals are checked by the receiver for various parity, phase and overrun errors to detect defective system components.

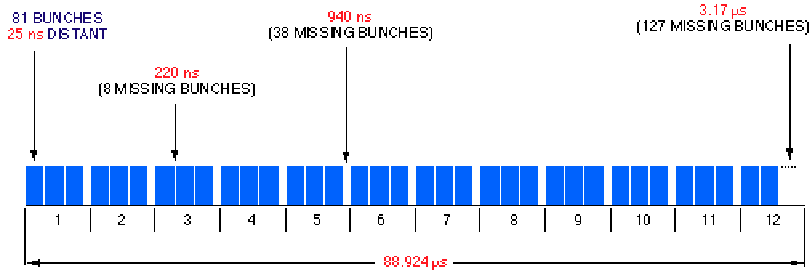

Fig. 6 Timing scheme

Revolution time	88.924 µs
Bunch crossing rate	40.08 MHz
Bunch spacing	24.9501 ns
Bunch train LHC extraction gap	81 3170 ns = 127 missing BC
LHC injection gap	950 ns = 38 missing BC
SPS injection gap	220 ns = 8 missing BC
Number of real Bunch Xings Number of CLK periods	2835 3564

Table 2 LHC timing

Rule 7.3.

The fibres, optical tree couplers and connectors which distribute the standard TTC signals to the modules using them are provided by the T/DAQ group.

Rule 7.4.

The sub-detector groups are responsible for the cost of the opto-electronic receivers and ASICs for the reception and decoding of the TTC signals, and for their integration in the FE electronics. (Standard components for these functions are being developed by RD12). Any additional sub-detector-specific systems for the further distribution of the TTC signals are also the responsibility of the sub-detector groups employing them.

Rule 7.5.

The standard TTC system provides for timing adjustment up to the TTCrx outputs. In the case that a TTCrx services a group of channels having significantly different time-of-flight or sub-detector-specific distribution delays, sub-detector groups have to foresee the additional delay compensation elements required. These elements may be controlled by the individually-addressed asynchronous command and data-transmission facility of the standard TTC system.

Rule 7.6.

Detector-specific extensions of the standard TTC system should have the same basic capabilities as the standard system, namely:

- Phase or timing adjustment of BC, L1A, BCR, ECR and synchronous commands at the module or channel level as appropriate.
- Capability to send commands without introducing significant deadtime.

5.6. Detector Control System

Rules to be defined.

References

[1] ATLAS technical proposal December 94

[2] RD12 B.G. Taylor. Timing, Trigger and Control distribution for LHC detector. http://www.cern.ch/TTC/intro.html

[3] Brouzet etal. The LHC Proton injection chain
CERN PS 93-39 (DI)
CERN SL 93-37 (AP)
LHC Note 256, October 93

[4] H. Burckhart - DCS user requirement document (in preparation)

[5] ATLAS-GEN-14. Report of the Global Descoping Task Force, Nov. 95.