20 Strategy for handling deadtime in the experiment

20.1 Introduction

The ATLAS front-end electronics and readout systems contain many levels of buffering. Information may be lost at any of a number stages of the readout chain if buffers become saturated. Different strategies can be adopted to handle this situation, the two extreme ones being:

- introduce deadtime to avoid uncontrolled information loss;
- accept information loss and build a readout system able to accept incomplete events and possible loss of synchronization.

The first of these strategies has been chosen and it has been decided to introduce deadtime in the Central Trigger Processor (see Chapter 15) in order to:

- easily control and monitor the deadtime of the experiment;
- have a relatively simple and safe readout system relying on the presence of data for every event;
- simplify the front-end electronics systems by imposing an upper limit on the event rate and a minimum time between consecutive events.

Nevertheless, the front-end and readout systems are being designed to be robust against buffer overflow and loss of synchronization. Note that a few front-end systems, such as the semiconductor tracker, have variable-length data and, for these, it cannot be excluded that buffers will occasionally overflow with some local loss of data.

As described in Ref. [20-1] and shown in Figure 20-1, there are buffers at different places in the readout chain:

- At the front-end level, a derandomizer buffer, located just after the level-1 pipeline, can store a few events. This is needed in order to match the limited bandwidth of the front-end links, designed to cope with the average LVL1 trigger rate, to the random arrival time of the Level-1 Accept (L1A) signal.
- In the Readout Drivers (ROD) there are buffers at the input stage before the data processing is done and / or at the output stage before the readout link.
- In the Readout Buffers (ROB) the data are stored until the LVL2 trigger has made a decision.

Any of these buffers can become full so some deadtime must be introduced if data loss is to be avoided. The deadtime is introduced in three ways:

- Systematic short deadtime of four bunch crossings after each L1A signal is introduced in the CTP to accommodate front-end electronics limitations.
- The CTP limits the number of L1A signals that can be generated within a given period of time to prevent derandomizer overflows.
- The CTP can be vetoed with an external signal to handle the occupancy of the ROD and ROB buffers.

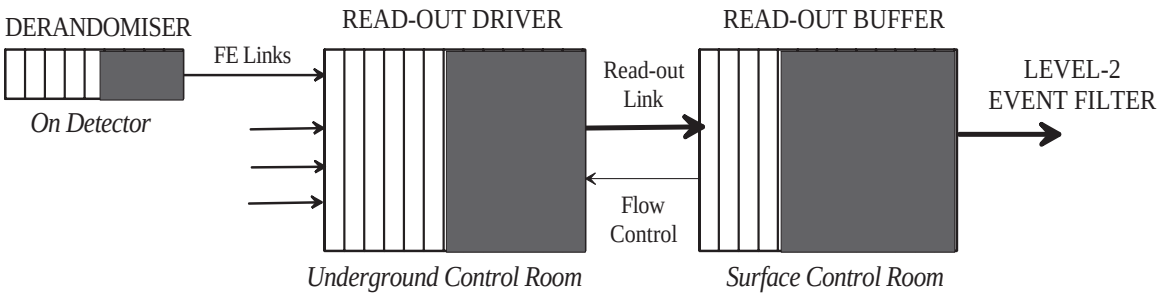

Figure 20-1 Buffers in the readout chain.

20.2 Systematic deadtime

Most of the subdetector front-end systems transmit the data of several consecutive bunch crossings on receipt of a L1A signal. This means that the data of a given bunch crossing could belong to more than one event and hence have to be read out several times. This complicates the readout control part of the front-end electronics when the level-1 pipeline is digital, and is almost impossible to achieve when the level-1 pipeline is analogue. The most demanding sub-detector in this respect is the liquid-argon electromagnetic calorimeter which has an analogue level-1 pipeline and reads out five consecutive bunch crossings per event. As a consequence, a systematic deadtime of four bunch crossings is introduced in the CTP after each L1A signal is issued.

20.3 Derandomizers

The derandomizers are part of the front-end electronics and must be large enough that no more that 1% of the events are lost due to deadtime when the L1A rate is 75 kHz [20-1]. There are several hundreds of thousands of such devices in ATLAS and their characteristics are sub-detector dependent. In some cases, zero-suppression has been applied on the data before they are stored in the derandomizer (e.g. SCT, MDT), while in other cases the full event is stored (e.g. TRT, LAr calorimeter). Some sub-detectors have small derandomizer size but high-speed front-end links, while others have large derandomizer size and low-bandwidth front-end links.

The level-1 trigger electronics must be throttled when the derandomizing buffers are nearly full, in order to avoid the loss of events in some places and hence the loss of synchronization between different parts of the detector. The classical scheme, where each buffer provides a signal to warn it is filling up dangerously, is not practical for two reasons. First, it would lead to a huge number of cables from the detector to the underground counting room where the CTP is located. Second, the time for a signal to transit from the on-detector electronics to the CTP is of the order of 500 ns (100 m of cable), and during that time up to four events can occur which corresponds to 30–50% of the derandomizer capacity.

The occupancy of the front-end derandomizers will be controlled by the CTP itself where an algorithm will be applied to make sure that the derandomizers never fill. This algorithm guarantees that there are no more than N L1A signals within any time window of M microseconds (see Chapter 15). The algorithm used is the ‘leaking bucket’ algorithm: each time an L1A occurs the content of the bucket is increased by a given value (X) and the bucket is

leaking at a constant rate (R). When the bucket is full, deadtime is introduced. The two programmable parameters (X and R) are used to control the deadtime. At the time of writing, the liquid-argon calorimeter electronics has the most stringent running conditions: the derandomizer contains up to 8 events, and 10 μs are necessary to transport an event on the front-end link. The CTP will be programmed so that there are no more than 8 events within any 80-μs time window.

This algorithm does not work when there is zero-suppression in between the level-1 pipeline and the derandomizer. In this case the event size is variable and the derandomizer occupancy is unpredictable. For these configurations, the subdetector electronics have to always have data for all the events and to produce a flag in the data to warn that there are data lost.

20.4 Readout drivers

In the current estimate, there are a total of about 1700 RODs in ATLAS. Most of these modules are located in crates in the underground control room (USA15). The RODs contain buffers which can fill up. These modules have to produce a signal (ROD_BUSY) when their buffer is close to being full in which case deadtime must be introduced. The logical OR of all these ROD_BUSY signals will be used to veto the CTP.

In order to avoid having to OR and monitor up to 1700 at the CTP level, a module will be provided to subdetector groups as described in Section 20.6.

20.5 Readout buffers

For the ROBs, an identical strategy to that of the RODs could be applied. However, as the ROBs are located in the surface counting room and the CTP is in the underground counting room, it is not deemed a very practical solution. Furthermore, deadtime could be introduced too early as a given ROB which is filling up does not know about the availability of buffer space in the ROD to which it is connected.

It is planned to use the flow control mechanism of the readout link to apply back pressure on the ROD. If a ROB has no available space in its buffer it disables the data transmission from the ROD. The buffer of the ROD may then fill up in which case the ROD_BUSY will be asserted in due course (if necessary).

20.6 Busy handling

20.6.1 BUSY module

As mentioned earlier, there will be about 1700 RODs in ATLAS, each of them providing a ROD_BUSY signal and capable of introducing deadtime. It is therefore very important to:

- monitor and control all of the ROD_BUSY signals so that any pathological ROD can be prevented from introducing deadtime;

- monitor the duration of these signals;
- gather the ROD_BUSY signals in a tree structure so that at the central-trigger-logic level only one BUSY signal per subdetector appears. An additional module will handle these signals and provide a VETO signal to the CTP.

A VME module handling up to 20 ROD_BUSY signals (i.e. corresponding to one crate full of RODs) will be made available. Its block diagram is shown in Figure 20-2.

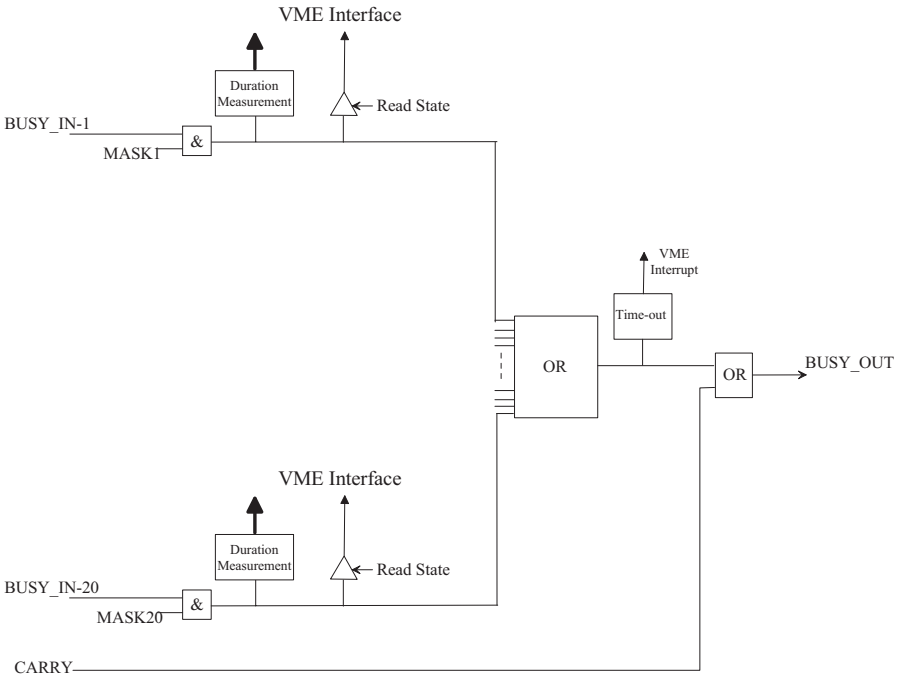

Figure 20-2 BUSY module block diagram.

Each input BUSY_IN signal can be individually masked and its integrated duration measured. A 16-bit counter clocked with the BC clock is used for this purpose. The duration measurement is available both in time and in number of BCs. It requires a readout of the counters every 1.6 ms, which is easily done by the processor controlling this module.

The logical OR of the unmasked BUSY_IN signals is made available as an output. The state of the BUSY_IN inputs can be read out through the VME interface for debugging purposes. A VME interrupt can be issued if the duration of one of the BUSY_IN signals exceeds a time limit in order to warn the processor controlling this module that an action has to be effected.

In order to be able to chain the modules in a tree structure (as shown in Figure 20-3), a CARRY input is made available. This input is connected to the BUSY_OUT of the previous module.

At the central-trigger-logic level, an additional module of the same kind is used to handle one ROD_BUSY signal per subdetector, and its output is used to veto the CTP.

Such a structure allows an efficient control of the deadtime in the experiment as well as an easy way to detect a faulty module introducing deadtime. It also allows the partitioning of the readout as a subdetector BUSY contribution can easily be removed from the CTP VETO. In the same way, additional subpartitioning within subdetectors can be implemented.

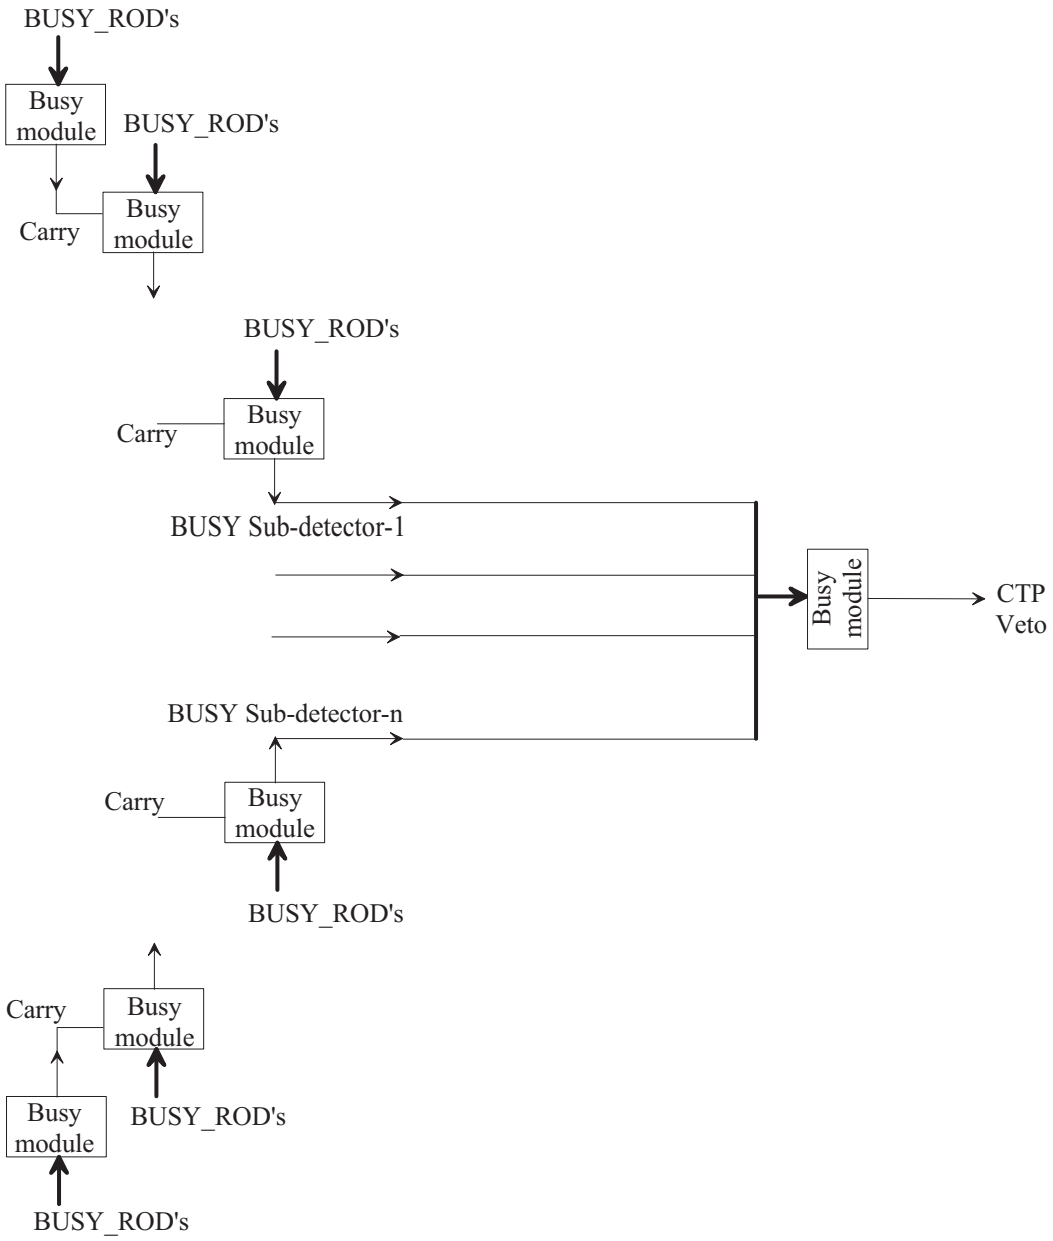

Figure 20-3 Chaining the BUSY modules.

20.6.2 Quality assurance and review procedure

Figure 20-4 shows the schedule for the design and production of the BUSY module.

After the specifications are defined, a preliminary design review (PDR) will take place. It will examine and assess the full requirements and specifications for the module, identify any missing functionality, ensure full compatibility with all connecting subsystems and determine overall feasibility.

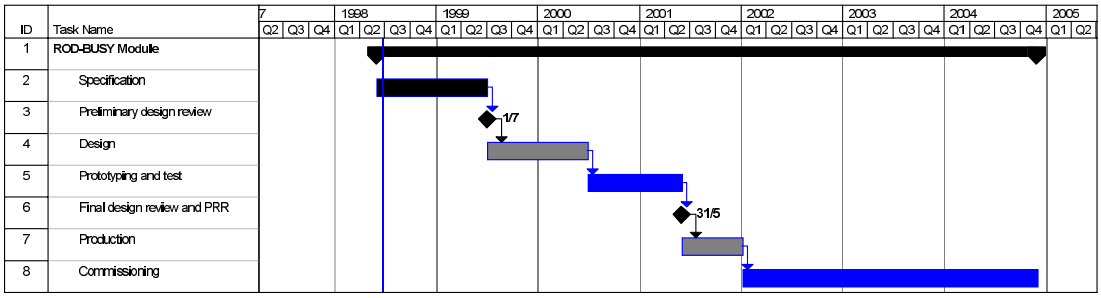

Figure 20-4 Schedule for the BUSY module production.

After the design is finished and before the production is launched, a final design review (FDR) will take place and be used as an ATLAS production readiness review (PRR).

The production of the boards will follow usual rules, namely:

- Electrical tests of the PCB before assembly.
- Burn in of the boards after assembly by leaving them under power without cooling for two days.
- Test of the full boards.

Each element will receive a serial number and a database will be set up to store the information relative to each board: origin of the components which populate the board, results of tests, history of failure and location.

In order to facilitate the test and the maintenance of the boards, the JTAG test bus will be used as much as possible.

20.7 References

20-1 *Trigger and DAQ Interfaces with Front-End Systems: Requirement Document (version 2.0)*, ATLAS note DAQ–NO–103, June 1998.
http://www.cern.ch/Atlas/GROUPS/FRONTEND/FEreq980310.ps