ATL-DA-ON-0003
Version 1.0
March 15, 2005

ATLAS/TDAQ

Proposal for a new online command – *Check System*

Author(s): T. Wengler
Contact(s): Thorsten.Wengler@cern.ch
Keywords: Synchronization, Timing signal distribution, L1A, L1ID, BCID
Status: Proposal released for comments

Abstract

A new command in the online software is proposed dedicated to the diagnostics of the L1A distribution tree during a run. The same mechanism may be used to check other aspects of the system in mid-run.

Change history

version	date	comments
0.1	28-Oct-04	– first draft
0.2	07-Dec-04	– comments from G. Lehmann, S. Hillier, N. Gee worked in (mainly on implementation - moving from a state change to a set of commands that uses existing functionality, and stressing the need to leave the door open for extensions to other checks)
1.0	15-Mar-05	– polished for release as proposal

1 Introduction

During the ATLAS combined test-beam periods in 2003/04 an increasing number of subsystems joint in combined runs. This not only required the integration of all run controllers into a common partition, but also common timing signals. In particular the L1A signal distributed to all subsystems and the ROD_BUSY signal send by each ROD to inhibit further L1As in case the ROD is not ready to receive them are crucial to keep the system synchronised. In normal operation a sub-system receives the L1A signal as a single pulse. The front-end system receiving the L1A then increments a local counter which forms the lower 24 bits of the L1ID. A further local counter counts the arrival of event counter reset signals (ECR), which resets the 24-bit L1A counter. This counter forms the upper eight bits of the L1ID [1].

2 Missing or additional L1A signals

A persistent problem in test-beam operation has been the occurrence of additional or lost L1A signals. This leads to an additional or missing event fragment from the sub-system in question. From this point on the synchronization of the read-out is lost and the full events are assembled from fragments belonging to different events.

The mechanism foreseen in ATLAS to detect such a loss of synchronization is the consistency of the BCID of each fragment. The BCID identifies the bunch crossing in which the data originated. As in case of the L1ID, the BCID is derived from a local counter. A separate timing signal, the BCR, resets the local BCID counters and is responsible for the correct synchronization of the BCID across the sub-systems and with respect to the LHC. Both the L1A and the BCR will need to be timed in for each timing zone receiving separate timing signals at the beginning of operation, and will have to be monitored continuously [2].

During test-beam operation only a fraction of the participating sub-systems was prepared to use the BCR in the manner foreseen for the final system. As a consequence the checking of the BCID in the read-out data had to be switched off and no check was possible. Detecting the problem was hence left to the examination of the recorded data, with considerable delay. An independent fast check of the number of L1As received by each sub-system would have been very helpful in detecting and diagnosing the problem.

Although one might expect the situation to be better during the commissioning and running of ATLAS, with BCID checking in place and online histograms reporting uncorrelated data from various sub-systems, the following points make the possibility to independently check the number of L1As received by each sub-system very desirable:

1. During the commissioning there are likely to be periods where the BCID mechanism cannot be used, because the signal is not easily available or because prototype modules are still used, which can not use it.

2. The timing-in of the BCR signal is a separate step in the timing-in procedure, which may not be undertaken by all sub-systems at the same time as the timing-in of the L1A signal.

3. Switching between running modes may cause the BCR signal to have different latency (i.e. global mode with the BCR signal coming from the CTP directly in LVDS via the CTP_OUT to LTP link vs. local mode with three or more partitions combined, where the BCR is delivered via local LEMO cables. This will lead to BCID mismatches unless the BCR signals are re-aligned or much care is taken to time-in local signals with respect to global signals, which could prove difficult.

4. The BCID errors which will be reported by the DAQ may be caused by failures of the BCR signal or failures of the L1A signal. It needs inspection of the data across sub-systems to determine the root of the problem. This requires the monitoring mechanisms to be in place and operational at the EB/EF level, which might not always be the case during the commissioning.

3 A quick and unambiguous way to check the L1A distribution tree

A very simple method to check the integrity of the L1A signal distribution starts by stopping to send L1A signals without resetting the local counters. One may then read the current value of the local counter via VME access and compare it both across sub-systems and with the count kept at the trigger source (e.g. the CTP). This method has been used in a manual fashion successfully several times during the test-beam.

This check should be carried out without stopping the run itself, since these counters may be reset at the end of a run (although they should not be), and since stopping and starting a run requires some time. The proposed check-point mechanism is very likely not suitable for this purpose either, as one will want to perform such a check without incrementing the run number. Instead the following mechanism is proposed:

- A new online command is created called *Check System*, represented by a corresponding button on the main run control panel.

  Note 1: This command may also be issued automatically at regular intervals.

- Pressing the *Check System* button causes the following sequence of events:

  1. The trigger source (e.g. the CTP) is signaled to stop the generation of triggers.

     Note 1: This may be achieved using the PAUSE transition already in place today.

     Note 2: If the *Check System* command is extended to include further checks it must be ascertained that actions taken by sub-system controllers in this transition do not interfere with the tests.

  2. The system waits a few seconds after the triggers have been stopped (e.g. the PAUSE state has been reached) to make sure all remaining events have left the system.

  3. The 32-bit L1ID is obtained from the appropriate local controllers (e.g. the ROD controllers, but also some FE controllers).

     Note 1: If an agreement is reached to publish this value each time the probe/getinfo command is issued, the values can be obtained from the IS entries thus published, and no extra work on the controller side is needed. One simply has to issue a probe/getinfo command, or wait until it is automatically executed (if the frequency is appropriate).

     Note 2: Should the method in Note 1 not be feasible, the controllers may be notified via IS, asking them to publish the numbers in return.

  4. In a similar fashion as for the L1ID the status of each BUSY line is obtained from the ROD_BUSY controllers.

     Note 1: After a few seconds of no triggers no sub-system should be busy.

     Note 2: The ROD_BUSY module features the necessary VME access to provide this information.

5. Once all answers have been received the L1ID numbers can be compared across the sub-systems and with the counter of the trigger source and the busy status of all lines can be checked.

- The result of the check may be displayed in an intuitive way, ranging from a simple "o.k." or green light in case all numbers are the same and no system is busy, to a detailed graphical display which indicates by color code which system is out of step (and in which direction) with respect to the reference value, i.e. the trigger source, and which system is busy.

Such a mechanism would in the matter of seconds reveal problems in the distribution tree of the L1A signals. In comparison during the test-beam each iteration of this operation took many minutes to complete, and needed the help of all relevant sub-system experts to provide the L1ID numbers while the system was in paused state. This manual version of the procedure will not be feasible any longer during full scale commissioning and operation.

One may be tempted in this context to utilize the TTCvi [3] capabilities. The TTCvi has an internal L1A counter, the value of which may be transmitted for checking purposes as asynchronous B-channel data to the FE receivers. However, the 24-bit counter implemented in the TTCvi is not reset by an ECR command, and can therefore not be used in this context. The LTP may provide such functionality (to be confirmed), and could hence be part of the test.

4 Prerequisites / Requirements on Sub-systems

The mechanism described above will be possible for the L1ID check only for sub-systems which can read their respective current L1ID values (all 32 bits) via VME on demand. This needs to be clarified with the sub-systems. Reading the busy status should be possible as all sub-systems will use the ROD_BUSY module, which provides the required functionality.

The necessary modifications on the run-control are not expected to be problematic, as this method does not appear to imply new functional requirements on the run-control, but rather another use-case for existing functionality. Changes are, however, necessary to the IGUI and the new use cases must be implemented, which needs to be agreed upon with the responsible group.

It should be noted that the proposed mechanism can easily be extended to other checks in mid-run. Further checks should however be light (fast in execution) and non-disruptive to the running state of the system.

Two such possible extensions concern the BCID and Trigger Type belonging to the L1ID (the last event before *Check System* is activated). As with the L1ID only those modules can participate which can read the last active BCID and Trigger Type via VME on demand.

5 Summary

A new command in the ATLAS online software is proposed. The new command *Check System* would allow to check in the matter of a few seconds, without stopping the run, if all sub-systems have received the same number of L1A signals, a functionality which has been found in test-beam running to greatly facilitate the system diagnosis. In addition the busy status (and other variables if appropriate) may be checked in mid-run by this mechanism.

The mechanism requires sub-systems to read and publish their current 32-bit L1ID counter value and busy status on demand. This functionality needs to be verified with the sub-systems. It furthermore requires the implementation of this command in the ATLAS online software.

References

[1] P. Farthouat, Definition of L1ID, EDMS ATL-D-ES-0014.

[2] N. Ellis, P. Farthouat, K. Nagano, T. Pauly, T. Wengler, Setting up the timing of ATLAS, EDMS ATL-DA-ON-0001.

[3] P. Farthouat, P. Gällnö, TTC VME bus Interface TTCvi, EDMS 110746.