

Trigger Timing Plots

John Nelson

14 November 2007

Chris Perkins introduced into L2 code, a set of time measurements which illustrate what is happening when L2 is processing events. This note adds explanations to the headings on the plots.

Following a trigger, L2 will receive data from all the DSM and QT crates as well as from the BTOW and ETOW data collectors. Data does not arrive in any particular order. For the purpose of the plots, all data sources are referred to as "DSMs"

Once all the 'DSM' data from a given event has arrived, L2 will cause another part of the code to begin handling this event, while remaining active in waiting for data to arrive from subsequent events.

Page 1: first DSM->lastDSM

This is time taken for data from the last 'DSM' source to arrive at L2 relative to the arrival of data from the first source.

lastDSM->trg

After the arrival of data from the last source, the main part of L2 begins processing the full event. The time taken from the arrival of data from the last 'DSM' source to the beginning of the main processing is given by this quantity.

Page 2: L2 algo time (trg->evtacc)

The processing of an event involves a sequence of algorithms examining data and finally providing a decision to accept (or abort) the event. The time is given by this quantity.

Accept (evtacc->accsent)

This is the time taken from the accept decision to the completion of sending the ACCEPT message to L0. This includes the construction of the message.

Page 3: Release (accsent->relsent)

If DAQ is enabled, all the data provided by the 'DSM' sources must be assembled into a single data block to be sent to DAQ. The time taken for this purpose is given by this quantity.

If DAQ is not enabled, then L2 will send a RELEASE message to L0 without building the event block. This time will therefore be much shorter than if DAQ is enabled.

Page 4: firstDSM->accept sent

This should be self-explanatory. It is the time taken from the arrival of data from the first 'DSM' source to the ACCEPT message being sent to L0. This is a critical time and should be less than 10 millisecs otherwise L0 will report a late accept.

lastDSM->accept

This time provides the complete processing time from the arrival of data from the last 'DSM' source, through all algorithm executions, to the ACCEPT decision being sent to L0.

Pages 5 to 8: 'DSM' times relative to L1

Data from 'DSM' sources arrives in no particular order. These plots show times relative to the arrival of data from L1. L1 is chosen since it must always be enabled. It is possible to run without any 'DSM' sources in the run.

Please note, that in the absence of a 'DSM' source, plots will still be made but entries should obviously be undefined.

Page 9: L1 Times

Three measurements are plotted. These times are determined by the L1 VME CPU.

L1 Time 1: this is only relevant if myriNet is being used by L1 to despatch the Build Event command to DSM crates. This despatch is done in a tree-arrangement and this plot shows how long L1 takes to send the Build Event command to the first 3 'DSM' sources in the tree.

L1 Time 2: following the despatch of the Build Event command (or, if STP is enabled, immediately) L1 must download data from its own set of DSMs. Time 2 is the time taken for the DMA readout.

L1 Time 3: L1 then examines all the triggers and produces a summary and, in principle, can execute algorithms to determine whether an event is to be aborted. (It does not do so at present.) 'L1 Time 3' is the time taken for this processing.

Pages 10 and 11: Algorithm times

These are time taken for the execution of all algorithms that have been enabled for this run.