

Module Performance Report

ATLAS Calorimeter Level-1 Trigger- Common Merger Module

B. M. Barnett, I. P. Brawn, C N P Gee

Version 1.0

23 February-2005

Atlas Level-1 Calorimeter Trigger

Table of Contents

1	Scope	3
2	Measured Performance.....	3
2.1	FPGA Occupancy	3
2.2	Device Execution Speed.....	4
2.3	Backplane Signal Timing Margins.....	5
2.4	Remote CMM Cable Timing Margins	6
2.5	Error-rate measurements between crates.....	7
2.6	Error-rate measurements on loaded backplane.	9
2.7	CMM Latency	11
2.7.1	Latency measurement by Simulation.....	11
2.7.2	Hardware latency measurement.....	12

List of Figures

Figure 1: Backplane Timing Margin	6
Figure 2: CMM cable input timings	7
Figure 3: One bit of front panel 40 MHz parallel LVDS output to CTP from CMM4..	8
Figure 4: Clock output on same connector as above.....	9
Figure 5: Input into CMM4, RTM from CMM5 via 3 meter cable.....	9
Figure 6: VME-- DS signal measured on the backplane.	10
Figure 7: Backplane input signals, measured at P10_4, from JEM in slot 10.	10
Figure 8: P3_0 signals for CPM in slot 3.	11
Figure 9: CMM latency.....	13

1 Scope

The Atlas Level-1 Calorimeter Trigger consists of three subsystems - the Pre-processor, the Cluster Processor (CP), and the Jet/Energy Processor (JEP). Both of the latter two subsystems use Common Merger Modules to perform intermediate and final result summing to prepare the hit flags sent to the CTP.

This document collects the device and module performance measurements available at the time of the PRR. The CMM and other related module specifications are available via the Level-1 trigger module web page, at <http://hepwww.rl.ac.uk//Atlas-L1/Modules/Modules.html>.

2 Measured Performance

2.1 FPGA Occupancy

Measurements by I.P.Brawn, January 2004

The Xilinx part numbers to be used for the production version of the CMM are:

Crate FPGA: XCV1000E-6FG860

System FPGA: XCV1600E-6FG860 (XCV1000E on current CMMs)

The XCV1000E, XCV16000E & XCV2000E devices all use the same FG860 footprint, so upgrades are possible. The following utilisation numbers are for the prototype CMMs, on which both the Crate & System CMMs are XCV1000E devices. They are:

	Cluster Crate Logic			Cluster System Logic		
	Used	Available	% Used	Used	Available	% Used
Flip Flops:	4,445	24,576	18%	949	24,576	3%
LUTs	4,371	24,576	17%	1,231	24,576	5%
IOBs:	441	660	66%	196	660	29%
Block RAMs:	60	96	62%	16	96	16%

	Jet Crate Logic			Jet System Logic		
	Used	Available	% Used	Used	Available	% Used
Flip Flops:	4,986	24,576	20%	762	24,576	3%
LUTs	4,762	24,576	19%	1,153	24,576	4%
IOBs:	508	660	76%	201	660	30%
Block RAMs:	68	96	70%	44	96	45%

Atlas Level-1 Calorimeter Trigger

	Energy Crate Logic			Energy System Logic		
	Used	Available	% Used	Used	Available	% Used
Flip Flops:	5,134/	24,576	20%	1255	24,576	5%
LUTs	6,052	24,576	24%	1,509	24,576	6%
IOBs:	539	660	81%	274	660	26%
Block RAMs:	73	96	76%	41	96	41%

2.2 Device Execution Speed

Measurements by I.P.Brawn, January 2004

The figures in this section are for the XCV1000E-6 devices on the current CMM, which are of the slowest speed grade available for that family.

Experience shows that figures from Precision Synthesis, which estimate the maximum frequency of a design before it has been placed & routed, are not particularly reliable.

The following numbers are taken from the Xilinx Static Timing Analyzer. They are calculated after place-and-route and are more reliable, and also use the same source information as the gate-level timing simulations.

Six items of CMM firmware have been used extensively in tests. For these, the Static Timing Analyzer gives the minimum clock periods as:

cluster crate-logic, crate-CMM 24.0 ns

cluster crate-logic, system-CMM 23.9 ns

cluster system-logic 22.6 ns

Jet-energy crate-logic, crate-CMM 23.9 ns

Jet-energy crate-logic, system-CMM 23.9 ns

Jet-energy system-logic 23.8 ns

The lack of headroom in these figures reflects two things. Firstly, the circuits were designed to run with a 24 ns clock with minimum latency. Each clock cycle has therefore been packed with as much logic as can be executed in 24 ns. The 24 ns clock is convenient to work with because it gives 1 ns of headroom, and has more factors than 25 so is easier to work with in simulation. Secondly, the Xilinx place-and-route tools don't do any more work than they have to, and stop optimising the design once the 24 ns clock specification has been met. This point is illustrated further below.

To test how much faster these designs could go, the cluster crate-logic for the crate-CMM was experimentally re-routed with a range of timing constraints and for all possible speed grades of device. The results are:

Atlas Level-1 Calorimeter Trigger

speed grade	clock period constraint	min clock period achieved
-6	24 ns	24.0 ns
-6	23 ns	22.8 ns
-6	22 ns	22.7 ns (constraint not met)
-7	24 ns	23.9 ns
-7	23 ns	22.9 ns
-8	23 ns	22.8 ns
-8	22 ns	21.4 ns
-8	21 ns	20.1 ns
-8	20 ns	19.9 ns

Thus, by keeping the current speed grade, tightening the constraints and re-routing the design, the minimum clock period could be reduced by ~1 ns. Further reductions could be achieved by using faster devices.

Regarding the interpretation of these figures, the Xilinx Static Timing Analyzer bases its calculations on "worst-case guaranteed delay" values. The documentation implies (but doesn't state explicitly) that designs will run ~25% faster than these speeds under normal operating conditions. However, an article in a Xilinx journal suggests adding 10% slack to the timing constraints.

Overall, using the current -6 speed grade devices, current firmware could probably be re-routed out to work with a 23 ns clock. To achieve a 22 ns clock, a speed-7 device would probably be needed.

2.3 Backplane Signal Timing Margins

Measurements by C N P Gee, February 2005.

Backplane input timings were measured with a single CPM in physical slot 6, sending data to the CMM in slot 20. The most sensitive indicator is the parity error counter, which was set to count errors from this source only. The error rate was measured by moving the deskew-1 clock timing over a complete clock period. The results are shown in Figure 1. Error rates are exactly zero for delays from 0 to 9.5 ns and from 19 to 25 ns.

The measurement thus illustrates an error-free region of approximately 15 ns of the 25 ns clock cycle.

Atlas Level-1 Calorimeter Trigger

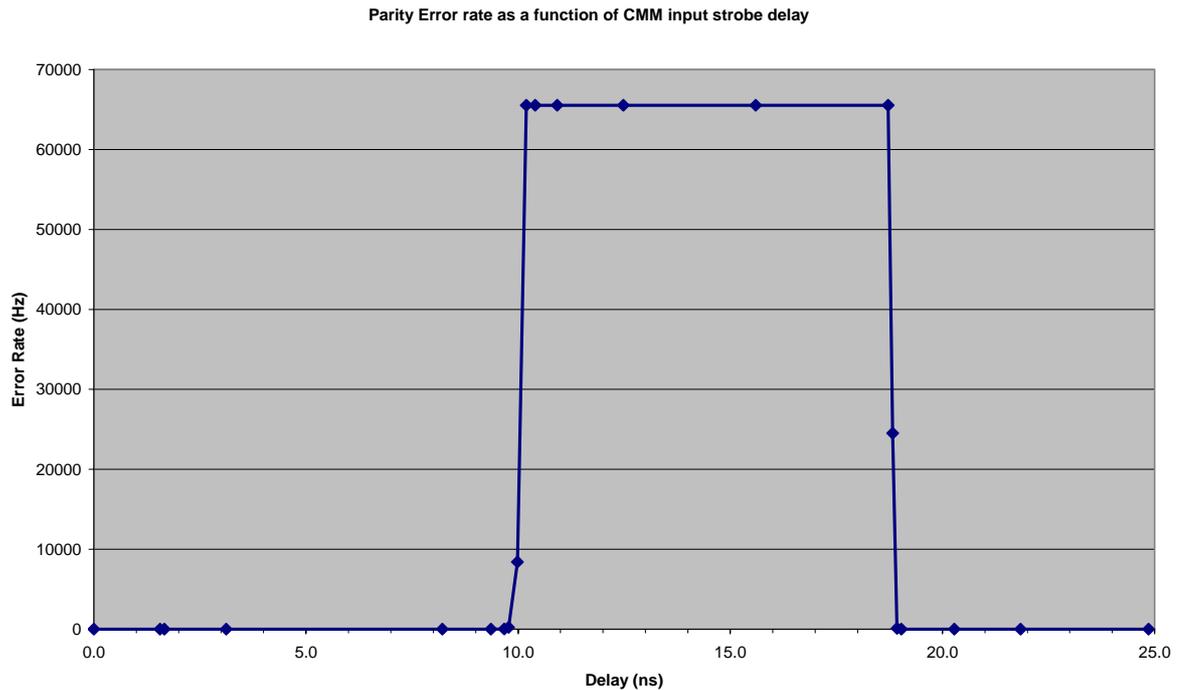


Figure 1: Backplane Timing Margin

2.4 Remote CMM Cable Timing Margins

Measurements by C N P Gee, February 2005.

Cable input timings were measured with a single c CPM in physical slot 6, sending data to a crate-summing CMM in slot 20 and then onwards to a second system-summing CMM in slot 0. The system-summing CMM parity error counter was set to count errors from the cable source only. The error rate was measured by moving the deskew-2 clock timing over a complete clock period. The results are shown in Figure 2. Error rates are exactly zero for delays from 0 to 2 ns and from 8 to 25 ns.

The measurement thus illustrates an error-free region of approximately 19 ns of the 25 ns clock cycle.

Atlas Level-1 Calorimeter Trigger

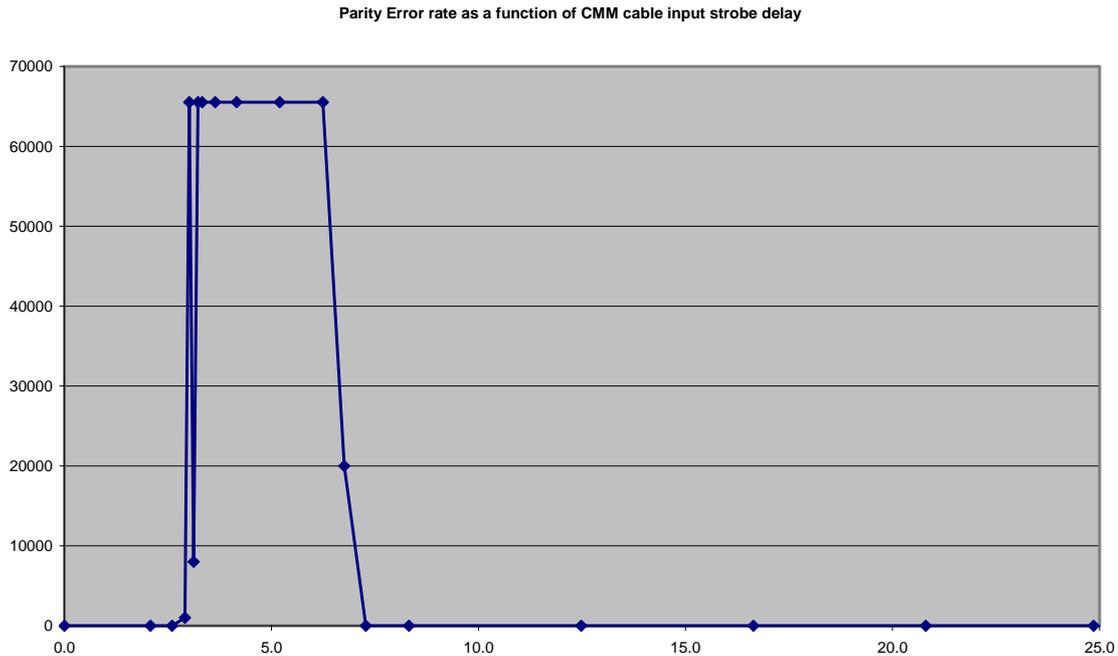


Figure 2: CMM cable input timings

2.5 Error-rate measurements between crates

Measurements by B M Barnett and C N P Gee, January 2005

Two Level-1 Calorimeter trigger algorithmic processor crates were set up in two separate racks with no particular care at inter-rack grounding having been taken – inter-crate grounding being provided only through the mains ground. The *system* crate (on the left) contained a single CMM in the right position, configured as a system merger. The other crate (on the right) hosted a CMM in both positions – configured as *crate* mergers. All CMMs were configured as e/gamma mergers.

The mergers were connected with a standard 3 meter SCSI-III cable running from each of the *crate* mergers to the *system* merger input (inputs 1 and 2 were used.)

The source (right) crate hosted 2 CPM modules (slots 3 and 4) and 2 JEM modules (slots 9 and 11) acting as pattern generators. The CPMs were loaded so as to generate a *stress* pattern (switching between AAA and 555 in alternate time-slices) whereas the JEMs were configured to generate ramps on all channels. L1A Traffic was generated at a high rate (around 130 KHz.)

The CMM backplane and cable timings were set up so that no parity errors were observed at the backplane inputs (crate mergers) or cable inputs (system merger.) The system was allowed to run for 2100 seconds (no error in $\sim 4 \cdot 10^{12}$ bits over the cables.) In a longer run, around 70 minutes involving only CPMs as data sources, neither were parity errors observed nor errors in the readout through a ROD when the results were compared with those of a simulation of the hardware (although only 500000 events were analysed.)

The signal qualities are illustrated in the following figures, and were measured with 1.7 GHz differential probe.

Atlas Level-1 Calorimeter Trigger

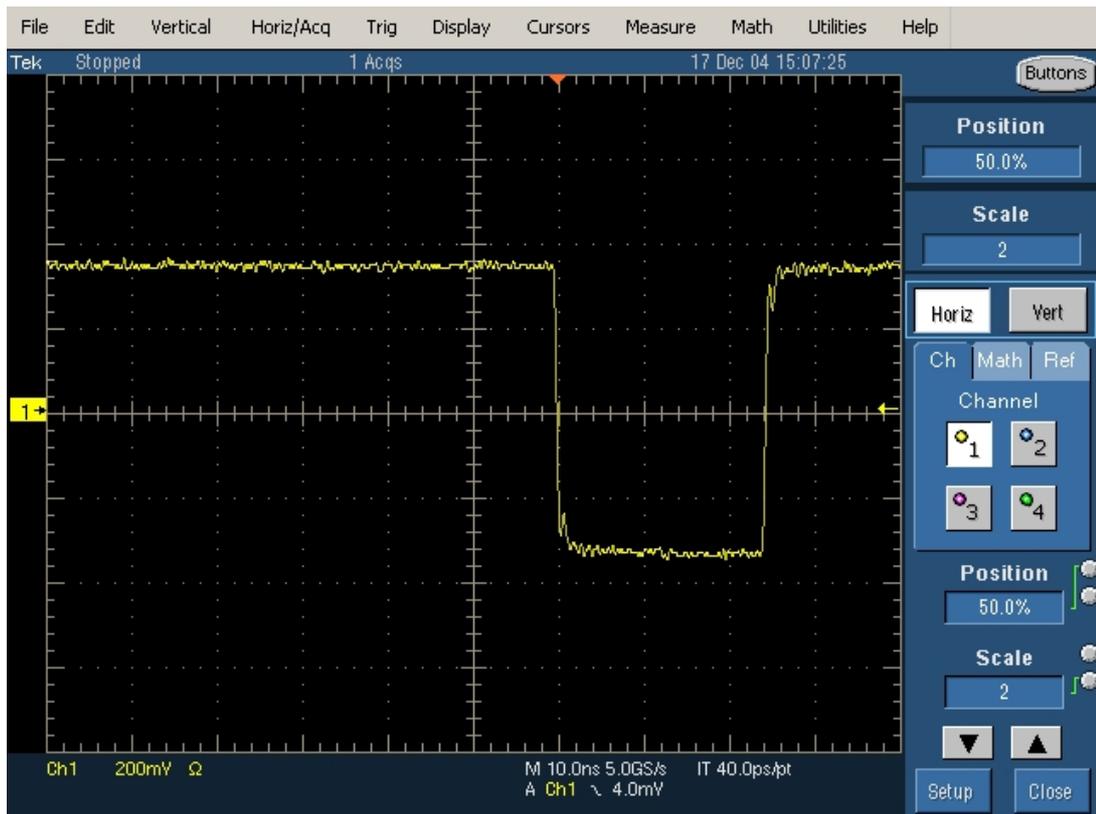
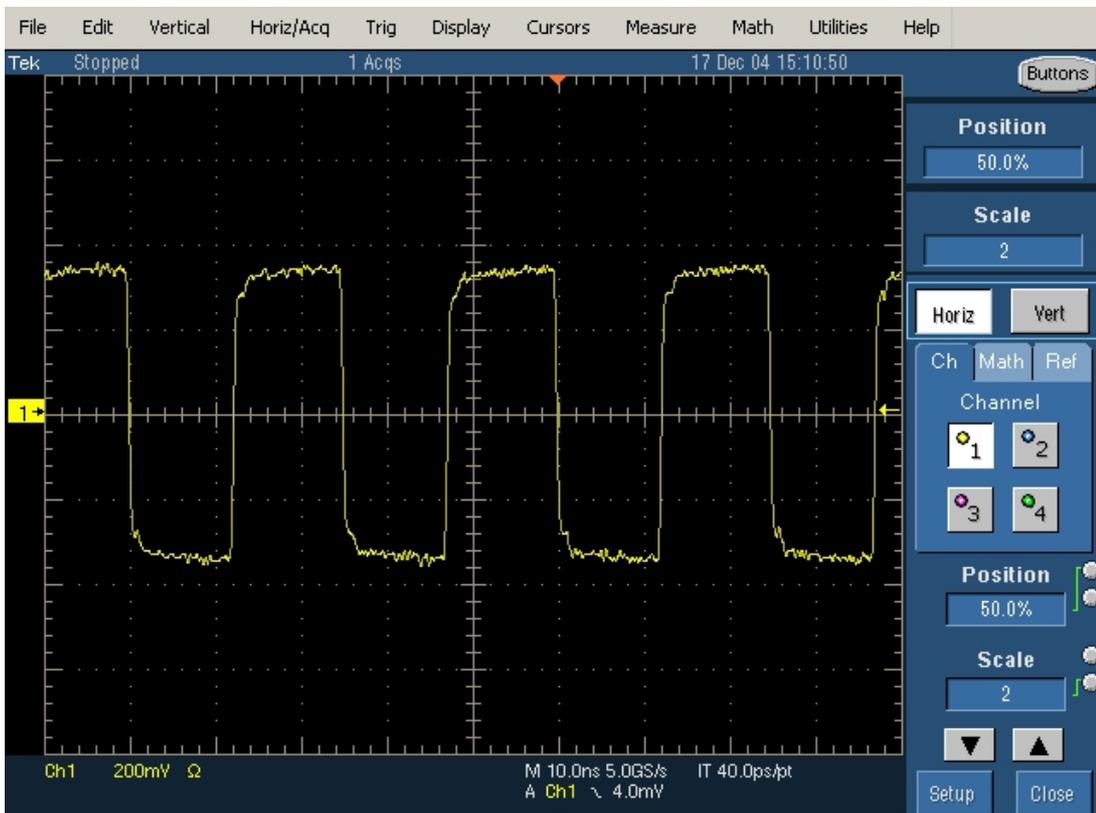


Figure 3: One bit of front panel 40 MHz parallel LVDS output to CTP from CMM4. Measured using differential probe with 100 ohm resistance between pins, measured at CMM front panel.



Atlas Level-1 Calorimeter Trigger

Figure 4: Clock output on same connector as above

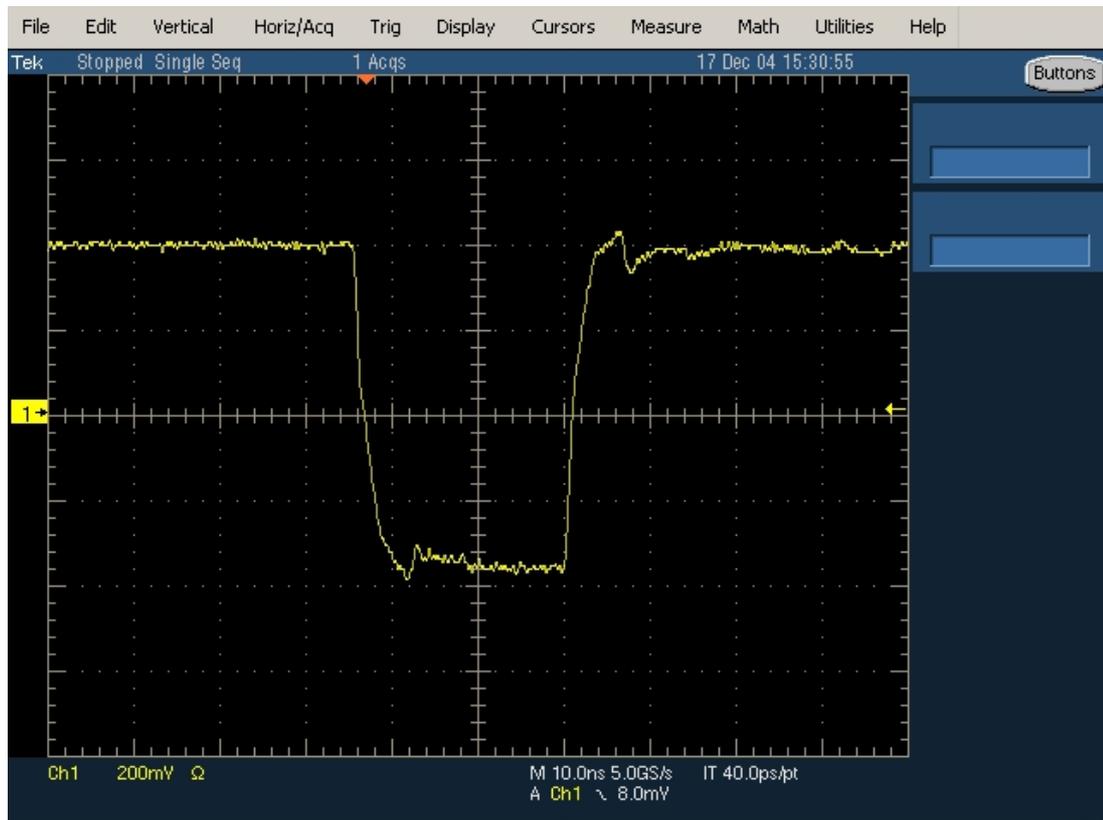


Figure 5: Input into CMM4, RTM from CMM5 via 3 meter cable, measured at connector on RTM (no extra resistor on probe.)

2.6 Error-rate measurements on loaded backplane.

There has been some concern about signal quality in the presence of much heavier backplane traffic, especially the 160 Mbit/s fan-in/fan-out data between CPMs. A single Level-1 Calorimeter custom processor crate was therefore populated with 6 CPMs (slots 1-6), 2 JEMs (slots 9-10) and 2 CMMs. The CMMs were configured as CP mergers, with the 2 JEMs driving a ramp into 1 channel each, to one CMM and the 6 CPMs, driving alternate 1s & 0s to all channels to both CMMs. The VME was exercised continuously in parallel by reading and writing a ramp to a 32 word memory on the crate TCM.

The CMM Timings were set up so that no parity errors were observed at the CMM backplane inputs and a run was started. A run of length 1.5 hours (5400 seconds) resulted in no parity errors at those inputs (no error in $\sim 10^{14}$ bits).

Signals were examined and found to be very clean:

Atlas Level-1 Calorimeter Trigger

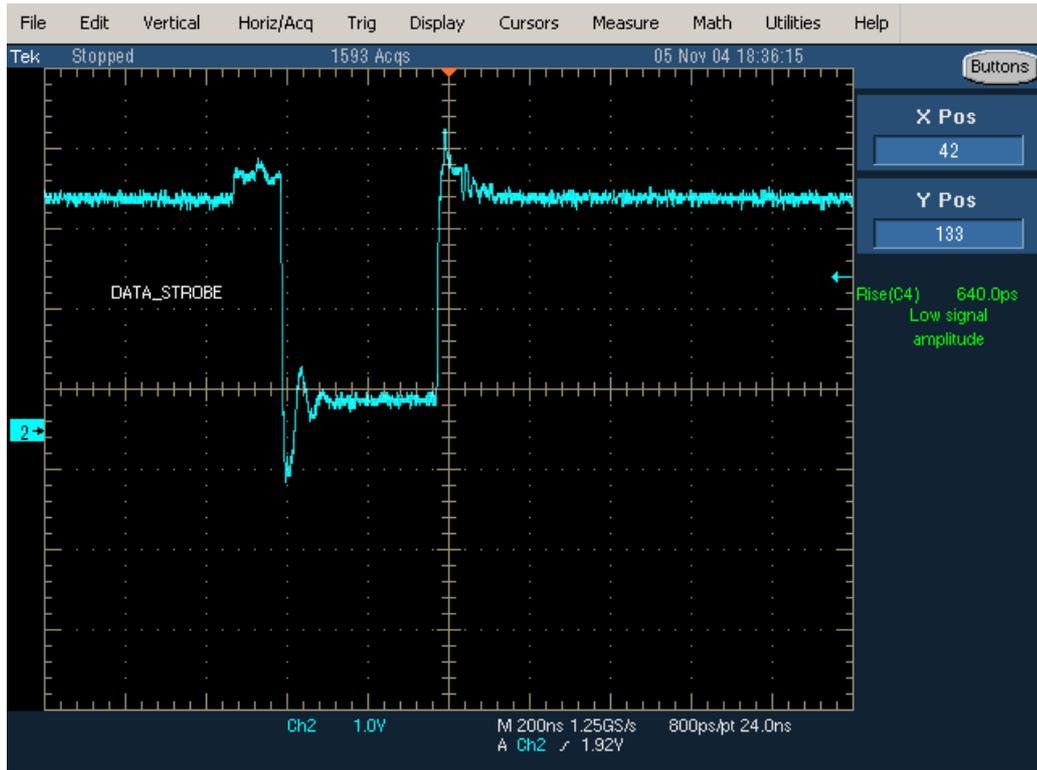


Figure 6: VME-- DS signal measured on the backplane.

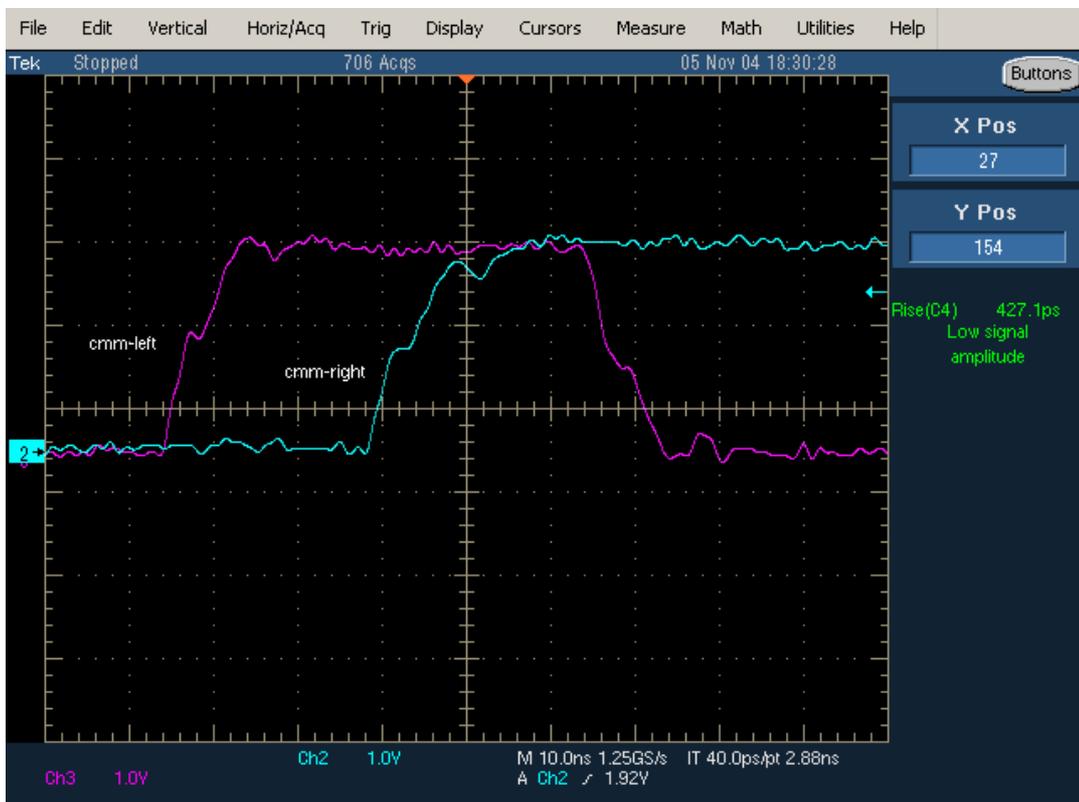


Figure 7: Backplane input signals, measured at P10_4, from JEM in slot 10.

Atlas Level-1 Calorimeter Trigger



Figure 8: P3_0 signals for CPM in slot 3.

2.7 CMM Latency

The latency in the CMM real-time path has been obtained both by simulation of the firmware and by lab measurement.

2.7.1 Latency measurement by Simulation

Measurements by I P Brawn, May 2003.

The total CMM (CP version) latency = 4 BC + 20 ns + cable delay. This breaks down as follows:

set-up time @ Crate FPGA:	1.5 ns		
Crate FPGA CP firmware:	2 BC	+	4.8 ns
CMM-CMM cable drivers:			2.6 ns
CMM-CMM cable:	?		
CMM-CMM cable receivers:			2.6 ns
set-up time @ system FPGA:			1.5 ns
System FPGA CP firmware:	2 BC	+	4.8 ns
CMM-CTP cable driver			1.7 ns

Atlas Level-1 Calorimeter Trigger

The latency for the firmware is defined to be the duration between the clock edge that captures the input data and the time when the output data becomes valid at the FPGA port. The firmware latency here is taken from timing simulation; the other numbers are taken from data sheets.

This calculation assumes that the skew of the TTC clocks to the CMM has been optimised so that a rising clock edge arrives at each FPGA as soon as the input data is ready to be clocked in. In practise a slightly longer value would be needed to ensure an adequate margin for error-free data capture.

The firmware numbers above are for the CMM CP logic. Work on the Energy firmware isn't quite finished is expected it to have a very similar latency (to within a few ns).

2.7.2 *Hardware latency measurement*

Measurements by N.Gee, February 2005.

Latency of the merger module system was measured using test data injected from a single CPM in physical slot 6 of a processor crate. The test data consisted of a single cycle of hits followed by a large number of empty cycles. The CMM clocks were adjusted so that CPM and the CMM system-summing logic were in phase, while deskew-1, which drives the crate-summing logic, was delayed by 12.5 ns. This meant that backplane data from the CPM was captured into the crate logic within a half clock cycle, and the system-summing logic could capture the CMM-CMM cable data again with less than a half-cycle delay. The CMM-CMM cable used was 3m in length. To compensate for the cable delay, a pipeline delay of 1 tick was required in the system-summing CMM.

The timing results are illustrated in Figure 9. The overall latency is approximately 156 ns, or a little over 6 bunch crossing periods. Given that latency of 1 bunch crossing has been introduced by the CMM-CMM cable, this compares well with the simulation value.

During test beam running in 2004, the behaviour of a single CMM was measured as 4.9 bunch crossing periods, providing both the crate and system-summing functions, with a pipeline delay value of zero. The present measurements also match this value well.

Atlas Level-1 Calorimeter Trigger



Figure 9: CMM latency. The upper trace is the CPM hit output, measured at the rear of the CPM. The middle trace is the CMM crate-summing output, measured at the rear of the CMM RTM. The bottom trace is the system-summing CMM output to the CTP. Traces 1 and 2 used single-ended probes, trace 3 used a differential probe.