Brief Overview of the SCT L1A Latency Budget Issue

Alex Grillo
27-Jun-1999

There has been considerable discussion for the past few months concerning a problem with the SCT L1A latency and its relationship to the length of fibre used to distribute the L1A signals to the SCT Front-End Electronics (FEE). Apparently, we have not distributed information concerning this issue to a wide enough audience or at least to the correct audience. There also seems to be some concern that the SCT must be distributing its signals inefficiently to have a problem since the TRT does not have a similar problem. While the TRT is planning to use the same protocol as the SCT, the existence or non-existence of a latency problem will still be critically dependent upon the distance that the signals must travel and the length of the respective pipelines for the two sub-systems, and these are not necessarily similar. However, it should be instructive to compare the latency budgets of the SCT and the TRT to see how efficient they are in distributing their signals given that they use the same protocol. This brief note is to provide a brief overview of the latency issue for the SCT and a comparison of the same latency budget for the TRT.

The table below lists the present understanding of the L1A latency for the SCT as tabulated by John Lane and for the TRT received from Peter Lichard. More details of the SCT latency are contained in John Lane's web page (http://www.hep.ucl.ac.uk/~jbl/SCT/SCT_latency.txt). The TRT lists two possible latency budgets depending upon whether they use the DORIC or not to de-multiplex the clock and command streams on the same link. Both are shown in the table. Since the breakdown of items is different for the two systems, a column is included at the far right to show the SCT times for equivalent items as in the TRT list.

Comparing the SCT to the TRT lists up to the ROD crate shows reasonably good agreement. There are small differences in the estimates which tend to balance out and the totals are within 10 ns of each other, the longer SCT estimate due to the CTP fanout included. A comparison of the SCT to the TRT with the DORIC option for timing from the ROD crates onward also shows good agreement. SCT estimates 192.5 ns to put the commands onto the fibre or cable while the TRT estimates 250 ns. After the fibre/cable going to the Front-End Electronics, the SCT estimates 200 ns and the TRT estimates 175 ns. So excluding the fibre/cable carrying the trigger command to the FEE, the two systems are within 25 ns of each other. It should be noted that these SCT estimates are slightly less than previously quoted due to some contingency previously included in individual items now collected in the 162.5 ns shown as a total SCT Contingency. We have not compared the SCT estimates to the TRT option without DORIC since the SCT is committed to the DORIC architecture.

Looking at the fibre/cable estimates for the run from ROD crate to the FEE, both estimates use 5 ns/meter for the propagation time even though one is optical and the other electrical. The SCT estimate of 97 m dates from roughly March-99 and assumes a "worst case" position of the ROD racks. That is, it assumes a position for the ROD racks which is most distant from the wall of USA15 nearest the ATLAS detector since the actual position of the ROD racks had not yet been fixed. To achieve this length of 97 m, a trial "special" path which breaks some of the phi-symmetry of the services layout inside the Inner Detector was worked out. This includes approximately 7 m contingency. A "normal" path for the SCT fibre, which preserved the phi-symmetry of the routing along with the worst case ROD rack position, would result in a length of about 122 m and add approximately 125 ns onto the latency. We believe that the TRT estimate of 100 m assumes a "normal" phi-symmetric routing through the detector but we are not sure what rack position is assumed for the RODs. It may be possible to significantly shorten the fibre paths by a positioning of the ROD racks closer to the ATLAS detector side of USA15 but that would require fixing the positions of the ROD racks.

The totals of the estimates for the two sub-systems, not counting the contingencies, are within 38 ns of each other. Of course, the critical comparison is how the latency totals match to the length of each pipeline. The SCT now has a pipeline length of 132 beam crossings (BCs) and the TRT has 134 BCs. This gives the TRT an extra 50 ns of margin in their latency budget. One can review how the SCT chose 132 and what assumptions were made years ago when that was chosen, but that is pointless now. The fact is that the chips have been designed with that length. To add to the pipeline length would require another chip submission and also enlarge the size of the chip which will cost more in production. One of the two SCT chips being considered must increase its pipeline in multiples of 12 so we would have to increase the pipeline length by about 10% to increase it at all.

In summary, the SCT latency budget looks reasonable with the 97 m fibre running from the ROD crates to the FEE. This leaves an overall safety factor or contingency of 6.5 BCs (162.5 ns). We think this is acceptable given the maturity of the design at this time but we would be uncomfortable to reduce this contingency until the designs are complete. A change to a significantly longer fibre from ROD crates to FEE would significantly reduce this safety. A comparison of the SCT estimates with the TRT shows no large differences. It appears that, with the exception of the TRT option that does not use the DORIC, the time required to fan-out and distribute the TTC signals for the TRT and the SCT are comparable. With the current estimates of 97 m fibre for SCT and 100 m cable for TRT running from ROD crates to FEE, both estimates fit comfortably within the maximums dictated by the pipeline lengths. It is important that these fibre/cable lengths be maintained in the final layout of ROD rack positions and services layout inside the detector. If a more optimal rack position can be selected which is both closer to the detector and closer to the CTP racks, a further reduction in latency budget can be realized by reducing the 32-40 m of fibre now budgeted to run from TTCvi to TTCrx. In other words, the total of 32 m and 97 m of fibre must be held fixed in order for the budget to work. Any reduction in the 32 m by better rack position can be re-allocated to the fibre run to the FEE in order to optimize its routing inside the detector.

Also, the Trigger Group is holding 17.8 BCs as contingency in the latency of the L1 Trigger. In this analysis, we have respected their hold on that contingency. If, after further refinements of their design, they can release some of their contingency, the total SCT fibre path could be lengthened. However, we doubt such a release could happen before the fibre path must be fixed.

	SCT			TRT w/o DORIC		TRT w/ DORIC		SCT Equiv.
	# BC	ns		# BC	ns	# BC	ns	ns
CTP Output (17.8 BC contingency)	77.9	1947.5	CTP Output (17.8 BC contingency)	78.0	1950.0	78.0	1950.0	1947.5
Cable (8m) & Fanout	2.0	50.0						
TTCvi	0.1	3.0						
Cable (0.5m)	0.1	2.5						
TTC Crate or TTCvx	0.9	22.0	TTC Elect (TTCvi+Encoder+TTCrx)	4.0	100.0	4.0	100.0	112.5
Fibre TTCvx to TTCrx (32m)	6.4	160.0	Cable CTP to ROD Crate (40m)	8.0	200.0	8.0	200.0	200.0
TTCrx	3.0	75.0						
TIM	2.0	50.0						
ROD Backplane	0.2	5.0						
ROD	3.0	75.0						
BOC	0.5	12.5						
BPM chip	2.0	50.0	TRT-TTC	5.0	125.0	10.0	250.0	192.5
Fibre BOC to FEE (97m)	19.4	485.0	Cable ROD to FEE (100m)	20.0	500.0	20.0	500.0	485.0
			Patch Panel	1.0	25.0	1.0	25.0	
DORIC	1.0	25.0	DORIC			1.0	25.0	25.0
ABC or ABCD	7.0	175.0	DTM ROC	5.0	125.0	5.0	125.0	175.0
SCT Contingency	6.5	162.5	TRT Contingency	13.0	325.0	7.0	175.0	162.5
Total Latency	132.0	3300.0	Total Latency	134.0	3350.0	134	3350.0	3300.0
Pipeline Length	132.0	3300.0	Pipeline Length	134.0	3350.0	134.0	3350.0	3300.0