SDC-93-427

SDC
SOLENOIDAL DETECTOR NOTES

LATENCY AT TRIGGER LEVEL 2

January 1993

B. Foster and G. P. Heath
University of Bristol, UK

R. Nickerson
Oxford University, UK

Latency at Trigger Level 2

B. Foster and G.P. Heath
H.H. Wills Laboratory, University of Bristol, UK
R. Nickerson
Nuclear and Particle Physics Lab, Oxford University, UK

As designs for level 2 trigger processors and readout systems become more concrete, it is worth reconsidering the question of the maximum Level 2 latency before chips are built and our decisions become set in concrete. There are clear and well-known trade-offs between making the latency very short, thereby decreasing storage required between L1 and L2 but leaving a short time for the L2 processing to arrive at a decision, and making it very long, thereby increasing storage requirements but simplifying the job of the L2 processors. We discuss these trade-offs in this note.

In the Technical Design Report our current maximum L2 latency is set at $50\mu s$, and the average time between L2 decisions is required to be $10\mu s$. Given that the maximum L2 trigger rate which we envisage is 1 kHz, there is at least scope for considering a considerable increase in this number. Below we list some of the pros and the cons of such an increase.

1. Increasing the maximum latency at L2 must make the design of L2 processors easier.

2. It increases the flexibility of the overall trigger system; if the processors at L2 are programmable, which is in general the case, then a larger latency will allow us to move processing around from say L3 to L2 or from L1 to L2 to react to unforseen problems in rates or bandwidths in the system.

3. There are possible qualitative changes of strategy possible as we increase the L2 latency. With the current value of $50\mu s$, the design of the L2 processors is constrained to rely substantially if not completely on special purpose hardware processors of one type or another. If we consider a relaxation of latency to say 1 ms, the number of available technologies is much larger. There are obvious advantages of flexibility and ease of design in the use of microprocessor type devices. It would also open up the possibilities of non-pipelined architectures at L2, and highly parallel systems. While we are not specifically advocating such systems, their availability for consideration may well pay dividends in the longer term.

4. Lengthening the L2 latency affects the necessary readout bandwidth for those devices, such as the silicon tracker, which transfer data off-detector after the L1 trigger decision. The current maximum transfer time of $10\mu s$ together with a maximum latency of $50\mu s$ requires a bandwidth much higher than the average transfer rate in order to cope with the statistical fluctuation of the arrival times of L1 accepted events. As the maximum latency is increased, the bandwidth can be moved towards the average transfer rate, or alternatively if the bandwidth is kept constant, the allowable maximum FLT rate increases. Such flexibility may prove to be very valuable in the highly uncertain conditions which will face us at SSC.

There are of course some cons:

1. The obvious problem is that an increase in L2 latency necessitates an increase in buffer size between L1 and L2. This increase in buffering applies both at the component level and at the global L2 decision processor. However, this should only be a problem for those components which keep their data on detector until after the L2 accept. The total amount of storage required with the maximum conceivable L2 latency extension would be less than 1 GByte for the whole experiment, which off detector is a negligible cost. This is not the case on detector however, where custom ASICs must be built in which memory space is likely to be at a premium. Nevertheless, broken down into the likely detector granularity, and in comparison with the storage required in any case at L1, such an increase is relatively modest.

2. Changing the L2 latency, even at this relatively early stage in SDC, may involve some redesign of essentially complete chips and boards and is therefore an additional cost to the experiment.

Given the above analysis, we believe there is a strong case to reexamine the maximum L2 latency with a view to extending it. As a parochial example of the importance of such a reexamination, we cite the ongoing design of the silicon tracker L2 processor. A attractive and flexible way of doing this is via a 2-stage algorithm, using a special purpose ASIC to filter the raw data and pass on potential high p_t patterns to a second stage based on a Digital Signal Processor. However, it is clear from preliminary investigations of possible devices and algorithms that the $50\mu s$ latency makes such a solution very marginal. It is interesting to note that the proposed LHC experiments seem to favour a rather long latency at their equivalent of our L2.

The magnitude of the L2 latency depends critically on the details of the various component trigger processor, buffer and readout schemes and we are not in a position to propose any numbers here. The components should enter into consultation on a likely optimum interval, with a view to making a final decision in the next few months. In such a consultation it will be necessary to specify the maximum latency allowed between various stages in the DAQ path, for instance between the component processors and global L2 processor. However, it should be realised that such a decision is likely to have to be made before all the required information is available. Our experience on the design of the ZEUS trigger scheme was that the originally proposed L2 latency was reexamined after processor design was well under way and that it was indeed lengthened substantially. That decision was however taken on rather general considerations and before it could be clear precisely what all the costs were likely to be. For definitiveness, we propose to increase the L2 latency to 1 ms. Given the timescales for SDC, we believe that we need to finalise this question in the next few months.