Performance requirements of proposed ATLAS second level trigger architectures from simple models.

D. Calvet [d], R. J. Dankers [c], M. Dobson [a], S. George [a,1], J.R. Hubbard [d], A. Kugel [b], P. Le Dû [d], I. Mandjavidze [d] R. Männer [b], K.-H. Noffz [b], J. A. Strong [a], J. C. Vermeulen [c], Q. Zhu [a,3],

[a] *Royal Holloway and Bedford New College, University of London, UK* [2]

[b] *Universität Mannheim, Germany*

[c] *NIKHEF, Amsterdam, Netherlands*

[d] *CEA, DSM/DAPNIA, Centre d'Etudes de Saclay, Gif-sur-Yvette, France*

The second level trigger for the ATLAS experiment will require higher communication bandwidths and processing power than existing experiments. Simple models are used to estimate the rates and loads inside the LVL2 trigger system and the latency for different trigger architectures and processing strategies.

Key words: Trigger architectures; Large systems; LHC; ATLAS

1 Introduction

ATLAS [1] is a general purpose detector for the LHC. The LVL2 trigger uses full precision data from the detectors to examine regions of interest (RoI) identified by the LVL1 trigger. As an initial step in the ATLAS demonstrator programme [4], "paper models" [2] based on simple calculations are used to estimate the latency, rates and loads in the LVL2 system, and indicate potential bottlenecks and critical parameters for different trigger architectures. Since the calculations are all made with average numbers and do not take into account queueing or congestion, results must be interpreted as a "best case".

1 email: S.George@rhbnc.ac.uk

2 Supported by the UK Particle Physics and Astronomy Research Council.

3 Supported by the joint UK Royal Society/CAS research project.

These limitations are being addressed by detailed modelling, emulation and lab measurements.

2 Model description

After a LVL1 accept, data from the muon detectors, calorimeters, transition radiation tracker (TRT) and silicon tracker are sent to readout buffers (ROBs) with limited processing capability. The average number of ROBs for each type of RoI is taken from [5], and varies from 1 to 21. The rates and patterns of LVL1 RoIs are taken from the trigger menu [3]. RoIs can be either "trigger" RoIs which contributed to the LVL1 decision, or "secondary" RoIs with lower thresholds. The aggregate rate of each RoI type affects the number of processors and network bandwidth required. The relative rates of the menu items affects the average latency. The LVL2 trigger must be designed to accept a total event rate of up to 100 kHz.

Several possible selection strategies are considered, ranging from fully parallel to fully sequential. For minimal selection, only trigger RoIs are processed. In the parallel scheme, data from all relevant ROBs are read out and features extracted in parallel. The event decision is made when all features are available, and the output rate is a few kHz. In the sequential scheme [2], trigger algorithms are executed in a number of steps, and the majority of events are rejected at each step. This allows the scope of LVL2 to be extended to include some LVL3 algorithms such as b-jet tagging and missing E_T, which reduces the LVL2 output rate to a few 100 Hz ("extended menu"). The comparative physics performance of parallel and sequential selection has yet to be studied.

Feature extraction (FEX) algorithms have been simulated and timed [6,7]. Extrapolated to a 500 MIPS general purpose processor, calorimeter FEX code takes about 100μs per RoI; tracking takes about 1ms. Execution times are typically around 10μs on FPGA processors (with today's technology). Electron confirmation in the calorimeter alone rejects background by a factor of 3. This is increased to 25 when the tracking detectors are used. TRT scan and b-jet tagging algorithms are estimated to take 50ms and 250ms respectively. Processors are assumed to have i/o co-processors which handle protocols and set up messages. A context switch and i/o time of 50μs is added to the CPU load for each message in or out.

3 Results

These results are for low luminosity running, characterized by parallel feature-extraction and a scan of the full TRT for events with validated muons, at a LVL1 accept rate of 38 kHz. High luminosity running differs mainly by the absence of the TRT scan and a LVL1 rate of at maximum 100 kHz.

Processor i/o is critical in all models, accounting for around half the latency, ROB load and required processing power. The calorimeter and TRT ROBs are saturated with requests in most configurations. This problem can be alleviated by processing only trigger RoIs or by sequential selection.

The total number of processors required for fragment building, feature extraction, event building and analysis is in the range 600 to 900 for architectures B and C [4]. The majority of these are used for feature extraction. For architecture A [4], it is found that 15 FPGA processor boards (each with 24 FPGAs) are sufficient to perform feature extraction at 100kHz. The average latencies are in the range of 5 to 10ms, except architecture A which is about 1ms and the extended processing at LVL2 which takes about 26ms. Network traffic is dominated by calorimeter and TRT data. For low luminosity the bandwidth is less than 1.5 GByte/s and for high luminosity running at 100 kHz the bandwidth is below 3 GByte/s.

Sequential selection increases the latency, but also reduces the ROB loads, farm sizes and network bandwidth. The cost of more processors for extended selection is compensated by the reduction in the input bandwidth to LVL3 and the reduction in the number of processors needed in the LVL3 farm.

The use of FPGAs for local feature extraction makes it possible to continue using slow tracking algorithms (TRT scans and b-jet tags) at intermediate luminosity (3×10^{33} cm^{-2}s^{-1}).

4 Conclusions

These preliminary models demonstrate that there are potentially several solutions for the ATLAS second level trigger. All the models investigated have affordable latencies. We have given a first indication of the critical parts of the system and what sort of performance will be required of the technology. This work will continue with closer comparisons between the different models.

Acknowledgement

The authors wish to acknowledge the valuable contributions that many members of the ATLAS collaboration made to this work.

References

[1] The ATLAS collaboration, Technical Proposal, CERN/LHCC/94-43.

[2] J. Bystricky et al, A Sequential Processing Strategy for the ATLAS Event Selection, to be published in IEEE Trans. on Nucl. Sci. Vol 44, 3 June 1996.

[3] J. Bystricky et al, Trigger Menus at Luminosity 10^{33} $cm^{-2}s^{-1}$ ATLAS DAQ-NO-54, June 96.

[4] ATLAS Level-2 Trigger Groups, Options for the ATLAS Level-2 Trigger, presented at CHEP97 by J. R. Hubbard.

[5] R. Bock, P. Le Dû, Detector and Readout specifications, and buffer RoI relations, for the Level-2 trigger demonstrator program, ATLAS DAQ-NO-63, January 1997.

[6] R. Hauser, I. Legrand, Algorithms in second-level triggers for ATLAS and benchmark results, ATLAS DAQ-NO-27, December 1994.

[7] S. Sivoklokov, R. Dankers, J. Baines, Second Level Trigger in the Forward Region of the ATLAS Inner Tracker, ATLAS INDET-NO-111, December 1994.