Architecture of the ATLAS High Level Trigger Event Selection Software

S. Armstrong, K. Assamagan, J.T. Baines, C.P. Bee, M. Biglietti, A. Bogaerts, V. Boisvert, M. Bosman, S. Brandt, B. Caron, P. Casado, G. Cataldi, D. Cavalli, M. Cervetto, G. Comune, A. Corso-Radu, A. Di Mattia, M. Diaz Gomez, A. dos Anjos, J. Drohan, N. Ellis, M. Elsing, B. Epp, F. Etienne, S. Falciano, A. Farilla, S. George, V. Ghete, S. González, M. Grothe, A. Kaczmarska, K. Karr, A. Khomich, N. Konstantinidis, W. Krasny, W. Li, A. Lowe, L. Luminari, H. Ma, C. Meessen, A.G. Mello, G. Merino, P. Morettini, E. Moyse, A. Nairz, A. Negri, N. Nikitin, A. Nisati, C. Padilla, F. Parodi, V. Perez-Reale, J.L. Pinfold, P. Pinto, G. Polesello, Z. Qian, S. Rajagopalan, S. Resconi, S. Rosati, D.A. Scannicchio, C. Schiavi, T. Schoerner-Sadenius, E. Segura, T. Shears, S. Sivoklokov, M. Smizanska, R. Soluk, C. Stanescu, S. Tapprogge, F. Touchard, V. Vercesi, A. Watson, T. Wengler, P. Werner, S. Wheeler, F.J. Wickens, W. Wiedenmann, M. Wielers, H. Zobernig

The ATLAS High Level Trigger Group[1]

Presented by Nikos Konstantinidis[a]

[a]Dept of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, U.K.

We present an overview of the strategy for Event Selection at the ATLAS High Level Trigger and describe the architecture and main components of the software developed for this purpose.

1. Introduction

The conditions at the LHC are expected to pose enormous challenges to the Triggering and Data Acquisition (DAQ) systems of all the experiments. With an expected event size of $\sim$ 1.5 MBytes in ATLAS and CMS (reflecting the large scale and high granularity of the various detector systems), the event recording capabilities are projected to be on the order of a few hundred Hz. Given the LHC bunch crossing rate of 40MHz, this means that the Trigger systems must achieve an event rate reduction of almost six orders of magnitude, selecting a handful of interesting events from each million of bunch crossings.

The ATLAS Trigger System features a (typical in HEP) three-level architecture. An overview of the system can be seen in figure 1. The first level (LVL1) trigger is based purely on hardware and uses coarse granularity information from the calorimeters and special muon trigger detectors in order to accept or reject events. The decision for an event must be reached within about $2\,\mu$s, the maximum latency of the LVL1 Trigger. The maximum output rate from LVL1 is 75 kHz, upgradeable to 100 kHz; the latter is the maximum rate that the readout electronics of the various subdetectors can sustain. When an event is accepted by LVL1, the data for this event are passed from the front-end electronics of the various subdetectors to specially designed Readout Buffers (ROBs), where they remain (and can be accessed) throughout the second level (LVL2) Trigger processing. Event selection at LVL2 is based on specialised software algorithms, running on dedicated farms of PCs. The average processing time per event is expected to be around 10 ms, implying a farm size of about 1000 processors. At this stage, the data from the inner tracking detectors become also available, as well as the full granularity data from the calorimeters and the precision muon chambers. It is also possible, at this stage, to combine/compare information from the different sub-detectors. As the ratio of interesting to non-interesting events is still very low at LVL2, the emphasis of the selection is on rejecting events as early as possible by accessing the bare minimum of data from the ROBs. LVL2 is expected to bring the event rate down to a $\sim 2\,\mathrm{kHz}$. If an event is accepted by LVL2, its data fragments from all ROBs are sent to the Event Builder, which builds the complete event and subsequently forwards it to the PC farm of the third Trigger selection stage, the Event Filter (EF). The average processing time at the EF is on the order of a few seconds. At this stage, full calibration and alignment information is available for the data, allowing (together with the increased time budget) the execution of precise, offline-like algorithms for the final online selection. The EF is also envisaged to perform event classification so that, for instance, potentially interesting events can be directed to a special, fast, offline analysis stream.

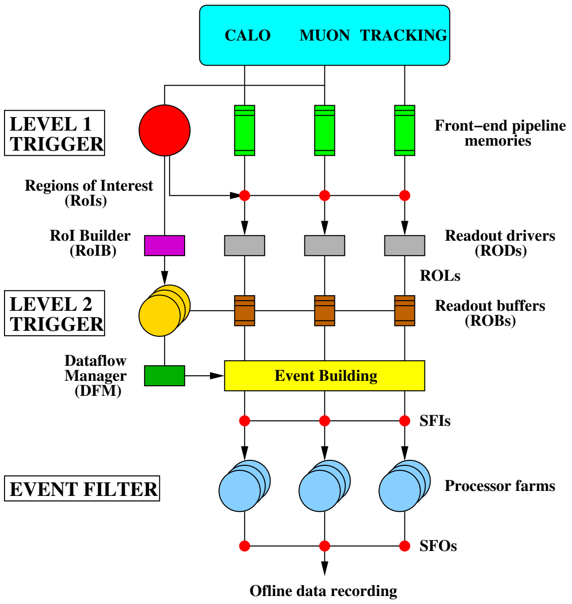

Figure 1. Schematic overview of the ATLAS Trigger and DAQ system.

The LVL2 and Event Filter together are referred to as the ATLAS High Level Trigger (HLT). The current hypothesis –based on projections for event recording capabilities, storage capacity and availability of offline processing power– is that the HLT event output rate (i.e. the event recording rate) will be around 200 Hz. Hence the HLT is expected to achieve an event rejection factor of about 500, similar to LVL1.

2. Event Selection Strategy at the HLT

The HLT event selection must make optimal use of the available CPU and network resources. This motivates the two main design features: (a) Processing in Regions of Interest, and (b) Step-by-step signature validation.

The Regions of Interest (RoIs) correspond to geometrical regions in the detector, where some interesting activity is detected at a previous processing step. When events are accepted by LVL1, information about the type and location of the RoI(s) that lead to the LVL1 decision is forwarded to LVL2. In this way, the LVL2 processors only need to access data fragments from a small fraction of ROBs and feature extraction algorithms need to run on just a small fraction of the full solid angle. For example, the typical size of LVL1 RoIs containing an electromagnetic (EM) cluster is (0.4×0.4) in (η,ϕ). The size and location of the RoIs can be refined after each processing step. For example, using the full granularity information from the calorimeter at LVL2, the size of EM RoIs can be reduced to (0.2×0.2), requiring access to (and processing of) an even smaller amount of Inner Detector data. Even at the Event Filter, where complete events are available, it is envisaged that reconstruction will be seeded by the (refined) LVL2 RoIs and full event processing may not always be necessary, leading to a reduction on the required CPU resources.

The second important feature of the HLT event selection strategy is a sequential approach to signature validation. Processing is performed in steps of feature extraction followed by hypothesis testing algorithms. This allows events to be rejected at the earliest possible step. In addition, this modularity facilitates the tuning and ordering of feature extraction so as to achieve the fastest event rejection. For example, for an event that is accepted by LVL1 due to an EM RoI, the first step at LVL2 is the feature extraction of the

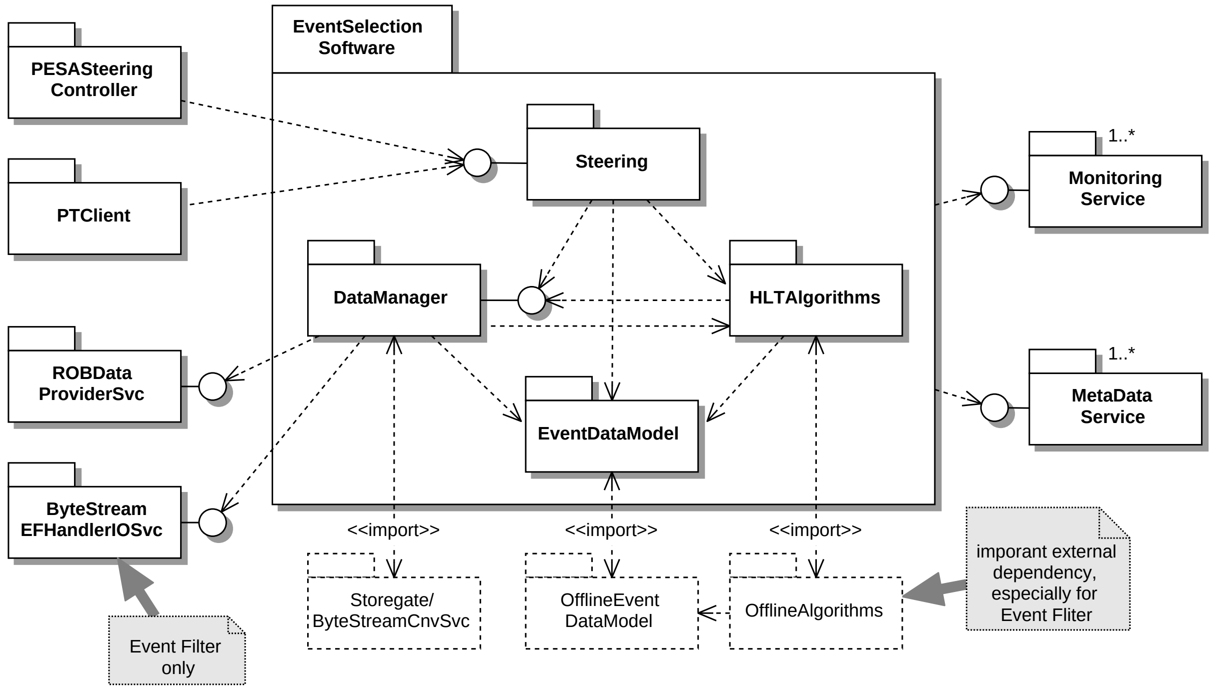

Figure 2. Package diagram and external dependencies of the HLTSSW. The central components are described in the text. The PESA Steering Controller and PTClient are external control components of the HLTSSW. The MetaData Service provides access to alignment and calibration information. StoreGate is the ATLAS offline Event Data Manager.

EM cluster shape. If this does not match the shape of a single electron (or photon cluster) the event is rejected before wasting network bandwidth and CPU to access the Inner Detector data and attempt to perform track reconstruction.

3. The HLT Event Selection Software

The design choices for event selection at the HLT, as described above, have been implemented in the HLT Event Selection Software (HLTSSW). The HLTSSW will be running on both the LVL2 and EF processing units, reinforcing the coherent approach for the two Trigger levels and providing the desired flexibility in defining the boundary between them. Any differences in terms of external dependences between the two Trigger levels are implemented through well-defined interfaces.

An overview of the main components and important external dependences of the HLTSSW is shown in figure 2. In broad terms, the HLTSSW comprises four components: (a) the Steering, (b) the Data Manager, (c) the Event Data Model, and (d) the HLT Algorithms.

The Steering [2] controls the order in which the various algorithms are executed, given the results at each processing step. Starting from the LVL1 result and having defined the complete set of final physics signatures that the HLT should accept (Trigger Menu), it is possible to draw up the full list of algorithm execution sequences that have to be followed in order to validate an event.

The Data Manager handles the event and trigger-related data during processing. It provides the HLT Algorithms with the event data they request and the Steering package with the algorithmic results so that the processing sequence can be determined. An important functionality of the Data Manager is to provide the data within an RoI. Given the location and size of an RoI, the Data Manager first finds the list of Detector Elements that are contained (fully or partly) within the RoI and subsequently, if the corresponding data are already in memory (as is the case in the EF) they are returned to the HLT algorithm, otherwise (as in the case of LVL2) a request is sent to the ROB Data Provider to fetch the data from the ROBs corresponding to the list of Detector Elements. Special emphasis was given on implementing this whole data access procedure in the most realistic way, just as it will be taking place in the running experiment, in order to diagnose any potential bottlenecks and provide solutions accordingly.This included a realistic representation of the data in bytestream format, and implementation of the software for converting the bytestream to raw event data objects.

The other two components of the HLTSSW, the Event Data Model and the HLT Algorithms, are naturally closely coupled to the offline software. Interfacing the algorithms through common event data entities (such as Track or Space Point) facilitates the migration between LVL2/EF/offline and makes comparisons and performance studies easier, leading to considerable savings in development and maintenance effort.

4. Conclusions - Outlook

It is beyond doubt that the HLT will have a central role in optimising the ATLAS capability for exploiting the physics of the LHC. The HLT group has been working towards developing a system which is (a) flexible, so that event selection can be tuned according to the physics requirements; (b) robust, in order to be able to cope with unforeseen difficulties such as increased machine background or detector inefficiencies; and (c) inclusive enough to allow the online selection of potentially unexpected new physics processes.

The aim of achieving the most efficient use of the networking and CPU resources is addressed with two crucial strategy choices for the event selection: (1) Processing will be restricted to Regions of Interest, i.e. geometrical regions in the detector where interesting activity was found at previous triggering steps; and (2) Processing will be performed in sequential steps, so that uninteresting events can be rejected at the earliest possible stage.

The HLT Event Selection Software was designed to accommodate the above requirements and strategy choices. Special emphasis was also placed in exploiting as much as possible the commonalities with the offline software, for obvious reasons of efficiency, portability and maintenance. It is currently under investigation how far we can go in using common software for the offline and HLT frameworks, while satisfying the stringent constraints (mostly in terms of timing) of the online environment, especially at the LVL2 Trigger. Finding the right balance at this early stage is crucial for ensuring an optimal use of human resources in software development for the future.

Altogether, the work has culminated in implementing a complete "vertical slice" of the Trigger-DAQ system for a specific Physics signature, including a realistic data access mechanism from the raw, bytestream format and realistic algorithms. Detailed results and conclusions from the current studies will appear in the HLT-DAQ Technical Design Report which is due to be submitted to the LHCC at the end of June 2003.

Acknowledgements

The authors would like to acknowledge the help and support of the Atlas Data Acquisition Group and the Atlas Offline and Detector software groups. N.K. would like to thank the local organising committee for the excellent organisation of the conference.

REFERENCES

1. http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/HLT/AUTHORLISTS/pisa2003.pdf
2. G.Commune et al, "The Algorithm Steering and Trigger Decision mechanism of the ATLAS High Level Trigger", presented at CHEP03, La Jolla, CA, USA, to appear in the proceedings.