Configuration of the ATLAS Trigger

A. dos Anjos, N. Ellis, J. Haller, M. Landon, R. Spiwoks, T. Wengler, W. Wiedenmann, and H. Zobernig

***Abstract*—The ATLAS detector at CERN's LHC will be exposed to proton-proton collisions at a nominal rate of 1 GHz from beams crossing at 40 MHz. In order to reduce the data rate to about 200 Hz, only potentially interesting events are selected by a three-level trigger system. Its first level is implemented in electronics and firmware whereas the higher trigger levels are based on software. To prepare the full trigger chain for the online event selection according to a certain strategy, a system is being set up that provides the relevant configuration information, e.g., values for hardware registers in level-1 or parameters of high-level trigger algorithms—and stores the corresponding history. The same information is used to configure the offline trigger simulation. In this presentation an overview of the ATLAS trigger system is given concentrating on the event selection strategy and its description. The technical implementation of the configuration system is summarized.**

***Index Terms*—ATLAS, configuration, LHC, online selection, trigger.**

I. Introduction

The Large Hadron Collider (LHC) [1] is currently being built at the European Organization for Nuclear Research (CERN) in Geneva. It is scheduled to start data taking in 2007. Proton beams will then be collided at a centre-of-mass energy of 14 TeV with luminosities of up to 10^{34} cm^{-2}s^{-1}. Such energies and luminosities allow a large physics programme to be carried out.

ATLAS [2] is one of the detectors at the LHC being built to record the pp interactions. Its design is similar to that of past and present collider experiments. Starting the description from the interaction vertex, the inner detector allows the identification and momentum measurement of charged particles in a magnetic field of 2 T which is provided by a superconducting solenoid. The inner detector is composed of silicon detectors with pixel and microstrip technology and a straw tube tracker. Calorimetry is provided by lead and copper liquid-argon sampling calorimeters and a iron scintillator-tile sampling calorimeter for hadronic calorimetry in the barrel. The calorimeters are surrounded by the muon spectrometer consisting of monitored drift tube chambers (MDT) and dedicated fast muon chambers for triggering—resistive-plate chambers (RPC) in the barrel and thin gap-chambers (TGC) in the forward region. The magnetic field in the muon system is provided by external air-core toroids. The ATLAS detector has a diameter of 22 m and a length of 46 m. The total weight is about 7000 tons.

At the LHC bunches of protons will cross with a nominal rate of 40 MHz, corresponding to a time interval between crossings of 25 ns. A total interaction rate of about 1 GHz is expected at nominal luminosity. The trigger of the ATLAS experiment must be able to reduce the rate to below the maximum rate that can be processed by the offline computing facilities, about 200 Hz, while retaining the capability of identifying previously undetected and rare physics processes. For example a Standard Model Higgs particle with a mass of 120 GeV, decaying into two photons, is expected to occur in one out of 10^{13} interactions.

II. Overview of the ATLAS Trigger System

The ATLAS trigger is designed as a three-level system. Fig. 1 gives a schematic overview of the trigger system. In the following only a short description of the three system levels is given; more details on the first trigger level can be found in [3] and [4]. The higher trigger levels are described in [5].

- The first-trigger level (LVL1) is a hardware based system which has to reduce the event rate to below 75 kHz (upgradeable to 100 kHz) within a latency of 2.5 microseconds. During that time the detector data are stored in pipelines. LVL1 makes its decision based on comparatively coarse information from the calorimeters and the muon trigger-chamber system. Track information from the inner detector is not included in the LVL1 decision. The LVL1 can be viewed in three parts (cf. Fig. 2): the calorimeter trigger (L1Calo), which receives and prepares calorimeter information, the muon trigger (L1Muon), which does the same for the information from the muon trigger chambers, and the LVL1 event-decision part implemented in the Central Trigger Processor (CTP).
- The Level-2 trigger (LVL2) is based on software selection algorithms running in processor farms. LVL2 can access data from all sub-detectors of ATLAS in so-called "Regions-of-Interest" (RoI) that were identified by LVL1. LVL2 has to reduce the event rate to about 1 kHz within an average LVL2 time budget of about 10 milliseconds. Hence a fast rejection is needed using specialized trigger algorithms.
- The Event Filter (EF) is also based on software selection algorithms and runs on large processor farms. In contrast to LVL2 it runs after the event building, such that the complete event information is available to the EF. Event selection and classification will be done within a few seconds. Events accepted by the EF are written to mass storage.

III. Level-1 Selection Strategy

The task of the LVL1 calorimeter trigger is to search for localized energy depositions that are signals for electrons/pho-

Manuscript received June 16, 2005; revised October 17, 2005.

A. dos Anjos, W. Wiedenmann, and H. Zobernig are with University of Wisconsin, Madison, WI 53706 USA (e-mail: Andre.Dos.Anjos@cern.ch).

N. Ellis, J. Haller, R. Spiwoks and T. Wengler are with CERN, Geneva 23, Switzerland (e-mail: Johannes.Haller@cern.ch).

M. Landon is with Queen Mary University of London, London E1 4NS, U.K. (e-mail: m.p.j.landon@qmul.ac.uk).

Digital Object Identifier 10.1109/TNS.2006.873307

[1]At LVL1, electrons and photons cannot be distinguished, nor can hadrons and hadronic decays of tau leptons into narrow jets.

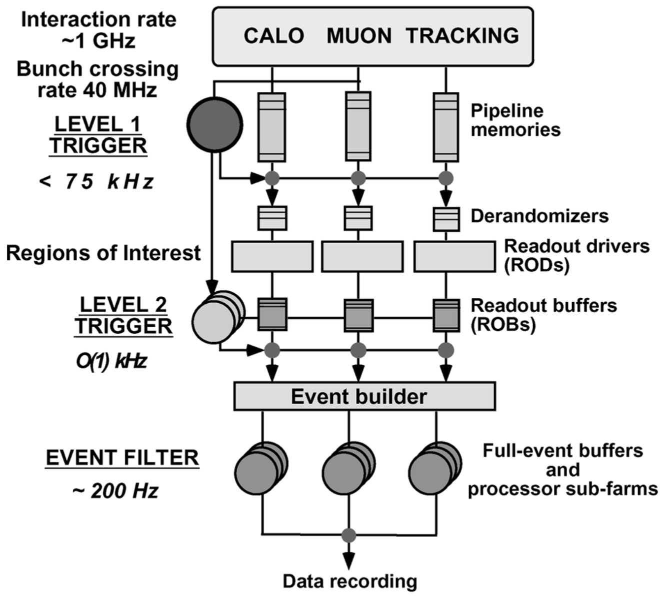

Fig. 1. A schematic view of the ATLAS trigger system.

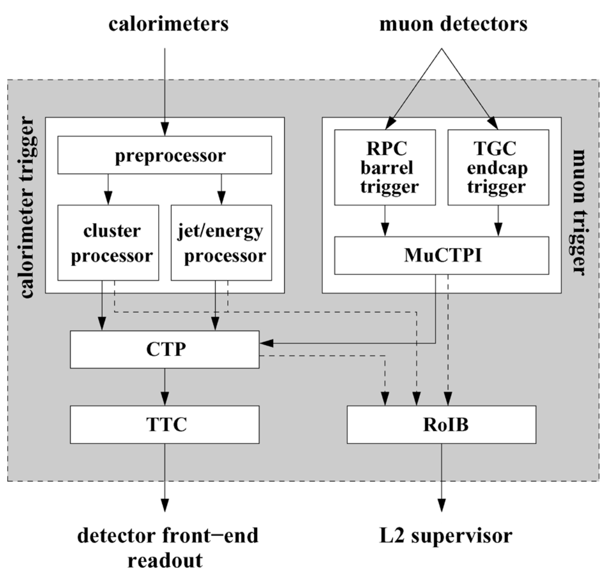

Fig. 2. Overview of the ATLAS LVL1 trigger system.

tons[1] (EM), tau leptons/hadrons (HA), jets (JT) and forward jets (FJ). The candidate objects are compared to programmable transverse-energy thresholds and the multiplicities of the objects passing the thresholds are counted. The 3-bit multiplicity values of 16 transverse energy thresholds for EM or TAU candidates (8 of these 16 thresholds are reserved for EM candidates, the other 8 can be configured to be either EM or TAU), 8 thresholds for JT candidates and 8 thresholds for FJ candidates are sent to the CTP for each event. The corresponding threshold values on the transverse energy as well as other selection variables—like isolation criteria or jet cone sizes—can be separately specified for various angular regions of the calorimeter (η, ϕ). In addition, global energy sums—total transverse energy (TE), missing transverse energy (XE) and transverse energy calculated from jets (JE)—can be derived and the bit-coded values of 4 TE thresholds, 4 JE thresholds and 8 XE thresholds are sent to the CTP.

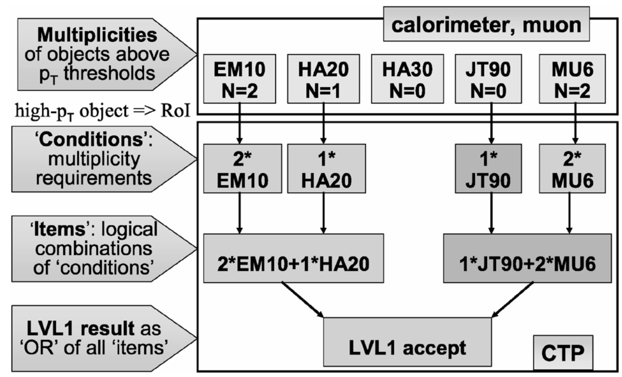

Fig. 3. Schematic example illustrating the LVL1 event selection strategy.

In analogy to the calorimeter trigger, the LVL1 muon trigger derives the multiplicities for muon candidates (MU) passing six programmable thresholds. The muon candidates are identified in the RPC and TGC detectors by requiring hit coincidences in various chamber layers. The momentum of the candidates is discriminated against the six thresholds by using coincidence windows of varying width exploiting the bending of muon tracks in the magnetic field. Contrary to the calorimeter trigger, the muon thresholds are common for the full angular range of the muon system. The 3-bit multiplicities are again sent to the CTP.

The Central Trigger Processor receives the multiplicities from L1Calo and L1Muon. In addition to these standard inputs, the CTP receives inputs such as beam-pickup signals, cosmic ray triggers, minimum bias and luminosity triggers. A total of 160 internal input bits can be taken into account by the CTP at any given time. Finally the CTP provides internal triggers from 2 random generators, for 8 bunch-crossing groups and for 2 prescaled clocks. The CTP combines the various input trigger bits and generates a level-1 accept signal (L1A) according to the LVL1 trigger menu consisting of up to 256 LVL1 trigger items each of which is a logical combination of one or more conditions on trigger inputs. For example, if "EM10" symbolizes the trigger input for electrons/ photons with a transverse energy of at least 10 GeV, then "2EM10" symbolizes the condition of there being at least two electron/photon candidates above that threshold. Several of these conditions can be logically combined to form a trigger item, e.g., "2EM10 and HA20". In addition, each trigger item has a mask and a prescale factor. A priority w.r.t. the dead-time generated after each L1A can be specified for each item.

Fig. 3 illustrates the LVL1 selection strategy for an example LVL1 trigger menu with two items, "2EM10 and HA20" and "JT90 and 2MU6." In this example, the event leads to a L1A as it fulfills the "2EM10 and HA20" signature. For simplicity, possible masks, prescale factors and active dead-time are not illustrated in the figure.

The full trigger menu configuration of the CTP can be changed via VME registers. The same is true for the L1Calo thresholds. A change of the TGC thresholds requires a firmware change. RPC thresholds can be changed by loading prepared coincidence matrix files into the on-detector electronics.

At every L1A, the CTP, the calorimeter trigger and the muon trigger send information to the level-2 trigger system via the RoI-Builder (RoIB). In case of the CTP this information contains among others, an 8-bit trigger type, the 160 input bits, the bits of the internal triggers and trigger bits for the 256 trigger items. The calorimeter trigger and the muon trigger send Region-of-Interest information of each candidate object found with their algorithms. These RoIs contain the geometrical region in the detector (η, ϕ) as well as the transverse energy thresholds passed. The corresponding thresholds of these RoIs are not necessarily required to be part of any LVL1 item (secondary RoI).

IV. High-Level Trigger Selection Strategy

The ATLAS High-Level Trigger (HLT) consists of LVL2 and the EF which are both implemented as pure software triggers running in processor farms. The LVL2 supervisor (L2SV) computers, about ten of which are envisioned for the final system, receive the LVL1 RoI information from the RoIB and send it to processors in the LVL2 farms. In the LVL2 farm processors, the LVL2 processing unit (L2PU) forms the interface between the L2SV, the read-out subsystems (ROS) and the true HLT selection software. During the selection procedure, information from various sub-detectors can be retrieved from the ROS. In general LVL2 is guided ("seeded") by the geometrical RoI information from LVL1 and only data in these regions are requested, reducing the amount of data to be moved and analyzed. Secondary RoI can be used to give (guided) access to lower p_T objects. The decision of the LVL2 selection is sent back to the L2SV which, in case of a positive LVL2 decision, passes it to the event building.

Event building is the data-acquisition step in which all event fragments from all ATLAS sub-detectors are requested from the ROS and assembled to give a full ATLAS event. The event is then sent to the EF where events are distributed to EF processors which run the event selection (and classification). In general a seeded (from information sent by LVL2) reconstruction is most suitable, but in principle a full event reconstruction is possible within the available resources. In case of a positive EF result, the event will be written to storage.

The HLT selection software provides a common framework to implement the LVL2 and EF selection. The concept of "step processing" subdivides the sequence of HLT algorithms to verify a given LVL1 RoI into several logical steps as illustrated in Fig. 4. "Trigger elements" are used at each step to abstract physics event properties. For example the trigger element "e30" indicates the presence of an electron candidate with E_T greater than 30 GeV. From step to step the trigger elements are refined by demanding information from more and more sub-detectors and combining it. This allows one to execute first those algorithms which give large rejection power for little data movement and processing time. This refinement is governed by so-called trigger sequences which specify the trigger elements required as input, the refinement algorithm to be run and the trigger element to be created if the algorithm completes its work successfully.

Requested trigger elements are logically combined in requested "trigger signatures," e.g., "e30 and e30." After each step the fulfilled signatures are discriminated against a menu table of requested signatures. Events fulfilling at least one

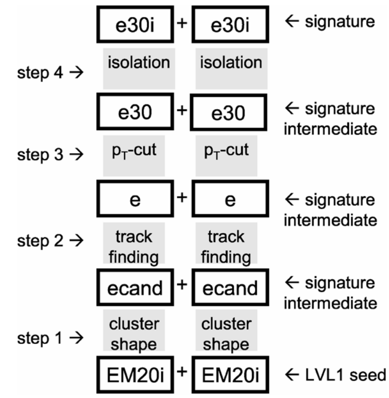

Fig. 4. Example illustrating the stepwise selection strategy of the ATLAS HLT. The LVL2-EF hardware boundary is not shown.

signature are passed to the next step. In case no signature of a menu table is fulfilled, the event processing is interrupted and the system is free to accept the next event. Random triggers are implemented in LVL1. The corresponding events can be accepted in the HLT with the help of dummy algorithms that do not refine detector data but simply copy trigger elements from one step to the next. Events accepted after the final step are written to mass storage.

The common software framework of LVL2 and EF allows moving decision-making (and possibly some refinement algorithms) across the LVL2-EF boundary. The HLT selection is driven by the LVL1 information, while the list of possible HLT algorithm sequences and required signatures will be configured at run start.

V. Trigger Configuration

The physics-inspired list of event signatures that the experiment should record needs to be translated into the dedicated description needed by the hardware and the software of the trigger. The configuration system stores this information and makes it available to the various users.

A. Use Cases and Constraints

Not only the online event selection needs access to the configuration information, but also the offline simulation of the trigger system and the later offline reconstruction of recorded events has to read these data for each run. The changing conditions of online data-taking will lead to frequent changes of the online event selection (e.g., changes of prescale factors) requiring a flexible and fast configuration system. In addition, it is mandatory to keep the correct history for later offline data analysis.

It has to be taken into account that many configuration parts are strongly coupled, e.g., the final HLT signatures are based on the LVL1 items and RoI (i.e., configured LVL1 thresholds) via the seeded mechanism. Hence consistency of the configuration is a key issue. In addition, the description of the trigger configuration must be generic, i.e., it must be able to describe all reasonable physics selection strategies.

B. The Description of the Trigger Menu

In this section the description of the LVL1 and HLT event selection are briefly introduced. A detailed description is beyond

TABLE I

th_name="EM10"	cab=2 s_bit=15 e_bit=17			
E_T	$\eta_{\min}$	$\eta_{\max}$	em. iso.	had. veto
8	-5.0	-2.5	1.0	3.0
10	-2.5	0.0	1.0	2.0
10	0.0	2.5	1.0	3.0
8	2.5	5.0	1.0	3.0

TABLE II

LVL1_item_name="DiElecHadTrigger", definition= (2EM10 and HA20), prescale=10, mask=ON, priority=HIGH

TABLE III

input TE	algorithm	output TE
"EM20i"	cluster shape	"ecand"
"ecand"	track finding	"e"
"e"	pt.gt.30	"e30"
"e30"	isolation	"e30i"

the scope of this presentation, only simplified examples can be given here.

A LVL1 configuration consists of:

- The definition of thresholds for which L1Calo and L1Muon deliver multiplicities to the CTP. As detailed above, in case of calorimeter objects these thresholds can be separately specified for various angular regions of the detector. Therefore a threshold is a list of so-called threshold values. This is demonstrated in Table I, where the EM Threshold "EM10" is composed of 4 threshold values covering the angular region from $\eta = -5$ to $\eta = 5$. In this simplified EM case, the transverse-energy threshold, the electro-magnetic isolation, as well as an hadronic veto can be specified for each threshold value. In addition, the physical cables over which the multiplicities are sent to the CTP need to be known by L1Calo and CTP. Similarly the other thresholds on calorimeter objects and muon candidates can be described.
- The definition of LVL1 items, i.e., the logical combination of requested L1Calo and L1Muon multiplicities derived in the CTP. An item includes a prescale factor, a mask and a priority w.r.t to the dead-time as detailed above. An illustrative example is given by Table II, where "EM10" and "HA20" are thresholds that are configured as detailed previously.
- Other more technical parameters like dead-time parameters, parameters for technical internal triggers like random generator rates, etc.

The HLT configuration is composed of:

- The definition of algorithm sequences that define which algorithms are used at each step to refine trigger elements (TE). A hypothetical sequence, corresponding to the example of Fig. 4, is given in Table III. Starting from the LVL1 RoI "EM20i," cluster shape and track finding algorithms, as well as cuts on the p_T and isolation criteria are used to reconstruct isolated electron candidates labeled with the trigger element "e30i."
- The definition of requested signatures, i.e., the logical combination of requested trigger elements that events have to fulfill in order to be accepted. This definition includes a prescale factor and a forced accept rate that is needed for monitoring purposes. An illustrative example is given by Table IV, where "e30i" is a trigger element as defined above.
- Specific parameters of the algorithms for feature extraction and event selection. Hypothetical examples are listed in Table V.

TABLE IV

HLT_signature_name="DiElectron", definition= (e30i and e30i), prescale=10, fa_rate=0.01

TABLE V

algorithm	parameter name	value
track finding	$\Delta\phi$	0.2
track finding	$p_{T,\min}$	5.
cluster shape	isolation	2.1
cluster shape	hot core	0.9
...	...	...

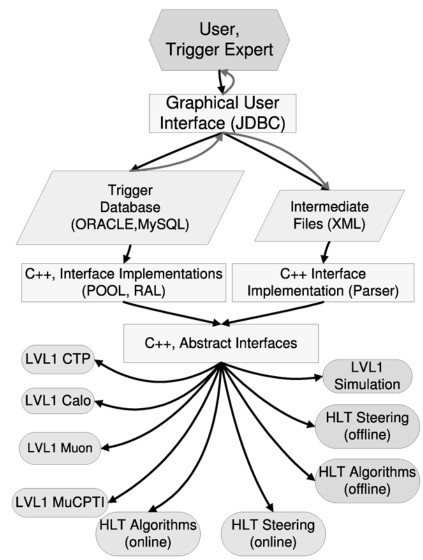

Fig. 5. Schematic overview of the dataflow in the configuration working model.

In order to allow an easy trigger operation the various configuration parts discussed previously must be available in many versions. This allows a re-use of parts of already used configurations.

C. Working Model and Technical Implementation

A configuration system is being built that provides the configuration information and stores the corresponding history. A schematic overview of the dataflow in the configuration system is given in Fig. 5. This system is based on and extends a static trigger configuration description [6] used so far. In the following an overview of the man components is given.

1) Trigger Database: The new system stores all versions of configurations in external relational tables ("Trigger Database") of the conditions database of the experiment. All parts of the trigger configuration are identifiable by unique keys. The full configuration, a relational tree, can be identified by a unique "master-key." This master-key will be stored in the proper interval-of-validity (IoV) table of the conditions database. This concept requires that configurations that were used already must not be changed afterwards in order to keep the correct history. However, new versions of configuration items must be introduced for each change. Thus, the database will grow with every change, not with every run (except the IoV table of the configuration masterkey). Replicas of the conditions database are foreseen to be available for reconstruction and simulation at several physical locations (outside institutes and GRID).

2) Graphical User Interface: The database is populated via a graphical user interface ("Trigger Tool") based on JAVA (JDBC). This interface allows users to browse the content of the database, to introduce new configurations (or parts of configurations) and to perform consistency checks of configurations. Various type of users with different access rights are foreseen. This tool will be essential for trigger operation during regular data-taking.

3) Access to DB Information: The trigger subsystems are given access to the configuration data via two independent paths. The first, foreseen for online running and regular simulation and reconstruction jobs, uses interfaces implemented in C++ based on the POOL/RAL package [7] that allows an DBMS independent implementation. The interfaces allow the direct access of the online state machines of the trigger subsystems to the relevant information in the Trigger Database tables at the start of a data-taking run. For online data-taking, the master-key that is needed to get the correct configuration information is made available to the subsystems via the RunControl [5] of the experiment and it is distributed by the Information Service (IS) [5]. The second path are intermediate text files (XML) which can be extracted from the DB using the Trigger Tool. They can be read by the subsystems via another, parser based, implementation of the access interfaces. This allows stand-alone tests (of the simulation) and debugging sessions without database access.

VI. Conclusion

An overview of the ATLAS trigger system has been given putting emphasis to the LVL1 and HLT selection strategies and the configuration of the full trigger system. A system is being built that configures all trigger levels consistently. It is based on a relational database back-end that holds all configurations and allows the correct configuration history to be retrieved via master-keys stored in the proper conditions database of the experiment. The trigger subsystems are given direct access to the database at run start. In addition, intermediate files can be extracted for stand-alone tests and debugging sessions. The system is expected to be developed further and see continuous improvements before the start of data-taking in mid 2007 and even beyond.

Acknowledgment

The authors would like to thank the ATLAS LVL1 Trigger Group, the ATLAS HLT Group and the ATLAS Data Acquisition Group for their help.

References

[1] M. Benedikt, LHC Design Report CERN, Geneva, Switzerland, CERN-2004-003-V1, CERN-2004-003-V2, CERN-2004-003-V3, 2004.
[2] ATLAS Detector and Physics—Technical Design Report (in The ATLAS Collaboration) CERN, Geneva, Switzerland, CERN-LHCC-99-14/15, 1999.
[3] ATLAS, First-Level Trigger—Technical Design Report (in The ATLAS Collaboration) CERN, Geneva, Switzerland, CERN-LHCC-98-14, 1998.
[4] P. Borrego-Amaral, "The ATLAS Central Trigger Processor (CTP)," presented at the 14th IEEE-NPSS Real Time Conference, Stockholm, Sweden, 2005.
[5] ATLAS, High-Level Trigger Data Acquisition and Controls—Technical Design Report (in The ATLAS Collaboration) CERN, Geneva, Switzerland, CERN-LHCC-2003-022, 2003.
[6] M. Elsing and T. Schörner-Sadenius, "Configuration of the ATLAS Trigger System," presented at the Conference for Computing in High Energy and Nuclear Physics, La Jolla, WI, 2003.
[7] I. Papadopoulos, "POOL Development Status and Plans," presented at the Conference for Computing in High Energy and Nuclear Physics, Interlaken, Switzerland, 2004.