AGH UST Faculty of Physics and Applied Computer Science
Wydział Fizyki i Informatyki Stosowanej AGH
Al. Mickiewicza 30
30-059 Kraków, Poland
www.agh.edu.pl
Kraków 2016

High Level Trigger role in extending physics reach of the ATLAS experiment at the LHC

System design and performance

Tomasz Bołd
Habilitation Thesis

This work was supported in part by Ministry of Science and Higher Education grant 665/N-CERN-ATASL/2010/0 and the National Science Center grants UMO-2013/08/M/ST2/00320, DEC-2011/03/B/ST/02631 and UMO-2015/18/M/ST2/00087.

Wydrukowano nakładem Wydawnictwa EXPOL
Recenzent: prof. dr hab. Danuta Kisielewska
ISBN 978-83-65256-07-2
Kraków 2016

This book is dedicated to my wife Iwona and children, Szymon & Milena.

Abstract

The trigger system is an indispensable ingredient of every high energy physics experiment. To study rare processes high luminosities are required which in turn result in a high rate of background. Only the early filtering of events containing interesting signatures makes it possible to carry measurements without unsurmountable burden of analysing already known processes. In the proton-proton, proton-lead and lead-lead collisions at the LHC the ratio of signal to background is about $10^{-7} - 10^{-9}$ depending on the process and centre of mass energy. At the LHC experiments, the trigger system is responsible for implementing bulk of this rejection, preserving about 10^{-5} (p-p) of collision events. In the filtered sample the interesting signal events are preserved with a high efficiency and backgrounds suppressed significantly.

The ATLAS trigger system consists of the hardware electronics system, the Level 1, capable of handling 40 MHz of collisions rate of the LHC and the software High Level Trigger part that is responsible for a precise events selection. For every collision the trigger system is responsible for selecting events from many hundreds of channels thus maximising physics output of the experiment. Due to finite capacity of the detector readout buffers the hardware system is designed to take the decision about an event in 2.5 μs. It uses for that purpose coarse granularity information from the Calorimeters and fast detectors of Muon Spectrometer to reduce the rate to about 100 kHz. The High Level Trigger acts on preselected events and uses full granularity information from Calorimeters, Muon Spectrometer and Inner Detector trackers. For the filtering a resource saving reconstruction in regions indicated by hardware is performed. The flexible system allows also, where justified by efficiency and/or purity, to perform complete event reconstruction. About a thousand of fast and precise algorithms orchestrated by the High Level Trigger steering component are engaged in the decision taking process. As of 2015 the High Level Trigger reduces rate of events which are saved for analysis to 1000 kHz in average processing time of 300 ms per event.

In this monograph the High Level Trigger is described. However for the sake of completness also the hardware trigger and surrounding data acquisition components are covered briefly. The purpose and constraint of the system are discussed first and referred to later in the part describing the design and implementation. Bulk of the content contains discussion of performance of triggering strategies from the perspective of the how they are implemented online. Remarks on the steady improvements and planned upgrades of the hardware trigger system and the High Level Trigger are also presented.

System wyzwalania, tryger, jest nieodzownym elementem każdego eksperymentu fizyki wysokich energii. Obserwacja rzadkich procesów wymaga wysokich świetlności, a co za tym idzie częstość obserwowanych procesów tła rośnie. Tylko wczesna selekcja na etapie zbierania danych

pozwala na przygotowanie próbki, która nie jest zdominowana przez tło, a przez co jest technicznie możliwa do analizy. W zderzeniach proton-proton, proton-ołów i ołów-ołów stosunek sygnału do tła wynosić może $10^{-7} - 10^{-9}$ w zależności od procesu i energii w centrum masy. Na LHC, systemy wyzwalania są odpowiedzialne za znaczną część filtracji pozostawiając do analizy około 10^{-5} ze wszystkich przypadków zderzeń. W zachowanej próbce interesujące procesy zachowywane są z wysoką efektywnością i czystością.

System wyzwalania eksperymentu ATLAS składa się z części sprzętowej, która jest zdolna do przetwarzania synchronicznie z częstością zderzeń wynoszącą 40 MHz na LHC i częścią wysokiego poziomu odpowiedzialną za precyzyjną selekcję. Dla każdego przypadku system weryfikuje hipotezę sygnału w kilkuset kanałach maksymalizując tym samym zdolność badawczą eksperymenu. Ze względu na ograniczenia w buforowaniu danych poziom pierwszy trygera jest tak zaprojektowany, aby podjąć decyzję o akceptacji bądź odrzuceniu przypadku w 2.5 μs. W tym celu używane są informacje o zgrubnej rozdzielczości z kalorymetrów i szybkich detektorów spektrometru mionów. Nominalna częstość akceptowanych przypadków przez poziom pierwszy wynosi 100 kHz. Zaakceptowane przypadki filtrowane są następnie w programowalnym filtrze wysokiego poziomu, gdzie pełna granulacja informacji z kalorymetrów, spektrometru mionów i detektora śladów jest używana. Zoptymalizowana rekonstrukcja wykonywana jest w rejonach detektora wskazanych przez poziom sprzętowy. W systemie możliwa jest także rekonstrukcja pełnego przypadku jeśli jest to uzasadnione zwiększeniem wydajności czy czystości filtrowania. W prcesie filtracji bierze udział około tysiąca szybkich i precyzyjnych algorytmów, których praca jest koordynowana przez komponent sterujący. W roku 2015tym filtr wysokiego poziomu redukował ilość przypadków zapisanych średnio do 1000 na sekundę w czasie ok. 300 ms na przypadek.

Niniejsza monografia opisuje system filtra wysokiego poziomu. Dla kompletności opisu zawarte są w niej także informacje o filtrze sprzętowym i systemie akwizycji danych. Na wstępie dyskutowane są cele i ogranicznia systemu filtrowania, jego projekt i konstrukcja. Dużą część zajmuje opis strategii z perspektywy implementacyjnej, w którym opisano obecnie już historyczne rozwiązania, bieżący system, a także planowane usprawnienia. Na zakończenie, krótko opisane zostały plany unowocześnienia całego systemu wyzwalania wraz z filtrem wysokiego poziomu.

Acknowledgements

First of all I would like to thank prof. Danuta Kisielewska. Without her none of my involvements in high energy physics would have ever happened. It was a pleasure to work under her wings for so many years.

I would like to thank the deans of Faculty of Physics and Applied Computer Science prof. Janusz Wolny, prof. Wojciech Łużny and prof. Bartłomiej Szafran for their support of High Energy Physics. Also big thanks to prof. Władysław Dąbrowski for leading the Department of Particle Interactions and Detection Techniques in an efficient manner for so many years. I profited for all this years from a friendly atmosphere created by the staff of the faculty. I am grateful for that.

I would like to thank prof. Andrew Lankford for a friendly atmosphere in the UC Irvine group in which I worked during my long stay at CERN since 2005 til 2011. A friendly, yet productivity oriented spirit in the trigger project created over the years by the ATLAS trigger and DAQ managers: Nick Ellis, Chris Bee, Srini Rajagopalan, Brian Petersen, Frank Winklmeier, Anna Sfyrla, Kunihiro Nagano, David Francis and David Strom is much appreciated. It is pleasing to add a piece to a well organised project. My, equally warm, thoughts go to the entire ATLAS Collaboration.

My contribution to the ATLAS HLT would not be possible without collaboration with many other colleagues, whom all I will not be able to mention in this short address as I am pretty sure, I would forget someone (and there is a page limit). However I would like to specifically thank a few who were at the cornerstones. First thanks go to Krzysztof Korcyl for telling me about the ATLAS trigger for the first time and explaining how challenging it is. Youngsters rarely need more than this to be ignited. He knew that well. While I was a novice at CERN, Werner Wiedenmann was an all-knowing guru who I was always able to ask questions being sure that the answer will be thorough and enlightening. For that I would like to thank him heartily. My great thanks go to Nicolas Berger and Till Eifert with whom I had a great fun to work on essential parts of HLT which have been functioning till now. For many years I had a pleasure to coordinate the core HLT software group with Joerg Stelzer. It was an inspiring time when I learned that impossible usually means difficult, but in reality means that “we need a bit of time”. When I recall more recent times I would like to thank John Baines with whom we look into the future of ATLAS HLT, a bright future I am sure, given John’s persistence on never losing the goal from the sight.

I would like to thank John Baines for reading this manuscript and providing useful comments on the style, english and content. I would also like to thank Szymon Bołd for making the nice drawing on the front page.

Statement of author contribution

My involvement in the High Level Trigger (HLT) dates back to 2005. At that time I joined Univ. of California, Irvine group at CERN. At first, I was involved in development of the infrastructure beyond the HLT itself responsible for monitoring. It provided means of handling monitoring information from HLT at runtime which functions till today. In 2006 the HLT passed from the stage of prototype to the phase where a functional system had to be developed. I took a leading role in this effort contributing to development of all aspects of the system starting from the core components, migration of the algorithms to the infrastructure facilitating construction of complex trigger menus. After the trigger organisational structure was established I was co-leading the Trigger Core software group which responsibility was to maintain robustness of the whole system. Several additional developments were performed under the umbrella of this group to which I contributed significantly: for example the analysis tools or reliable online rates metering subsystem for the HLT. Development of trigger strategies continued at that time and all non-trivial approaches were consulted by the Trigger Core group and optimal strategies were developed. I was privileged to present some of related results on behalf of the ATLAS at ICHEP 2008 [1] and CHEP 2013 conferences and contributed to several conference presentations and proceedings for instance [2, 3].

At the time of the LHC startup, I was heavily involved in the online system deployment being responsible for the overall trigger operations as a so called "trigger online expert on call" from 2009 till 2011. My involvement in the TDAQ system in Run 1 was recognised by the ATLAS collaboration by nominating me and then awarding the ATLAS Outstanding Achievements Award.

After re-joining Krakow AGH UST group I still keep contributing to the HLT serving with consultancy (i.e. to the currently used *b*-jet trigger strategy), maintaining the software and performing some expert-level shifts.

I also look into the future of the ATLAS HLT. Namely my current responsibility is to co-lead R&D of the future HLT framework for Run 3 (2020 and beyond) and beyond and an involvement in the upgrade activities by contributing to the design discussions [4, 5].

I also serve in various bodies of ATLAS Trigger/DAQ like the Steering Committee responsible for overseeing current and future of this ATLAS component. In years 2015 and 2016 I serve in the Trigger DAQ Speaker Committee responsible for organising TDAQ community contributions to many conferences. Last but not least I represent AGH UST & Jagiellonian Univ. in the Trigger DAQ Institutes Board.

Before the first Heavy Ion (HI) run in 2010 I was responsible for preparations of the trigger system to running in new conditions. That attracted my interest in the relativistic HI physics. In 2011-2012 I performed an analysis of hydrodynamic flow [6] and to that day I am involved in the HI programme preparing new results, serving as a member of editorial boards and coordinating trigger preparations for HI runs as a coordinator of Heavy Ion Trigger Menu Forum as well as by preparing specific trigger selection algorithms. I have presented HI results on behalf of the ATLAS Collaboration at the PLHC in 2011 in Perugia [7], Quark Matter in 2014 in Washington [8], Epiphany conference 2014 in Krakow [9] and recently at Rencontres de Moriond in 2016 [10].

Contents

List of abbreviations

Pixel - Pixel detector
Athena - Athena, ATLAS Offline Framework
BDT - Boosted Decision Tree
BPTX - Beam Pickup
BS - Beam Spot
CSC - Catode Strip Chambers
DAQ - Data Acquisition
DF - Data Flow system
FCal - Forward Calorimeter
EF - Event Filter
EMC - Electromagnetic Calorimeter
FEX - Feature Extraction Algorithms
HCAL - Hadronic Calorimeter
HLT - High Level Trigger
Hypo - Hypothesis Algorithms
L1 - First Level Trigger
L1CT - Level 1 Central Trigger
L1Calo - Level 1 Calorimeter Trigger
L2 - Level-2 trigger
LAr - Liquid Argon
LHC - Large Hadron Collider
LL - Log-Likelihood
MBTS - Minimum Bias Trigger Scintillators
MDT - Monitored Drift Tubes
MS - Muon Spectrometer
MV - Multivariate
RoI - Region of Interest
RoIB - Region of Interest Builder
RoS - Readout System
RoL - Redout Link
SCT - SemiConductor Tracker
SFI - Sub-Farm Input
SFO - Sub-Farm Output
TE - Trigger Element
TGC - Thin Gap Chambers
TRT - Transition Radiation Tracker
RPC - Resistive Plate Chambers
ZDC - Zero Degree Calorimeter

1

Introduction - online filtering of the experimental data

LHC experiments have been constructed to search for extremely rare processes. A quantity characterising the frequency of given process is called a cross-section, σ measured in barns[b]. It is a characteristic quantity for a process and is an indirect measure of probability of the process to take place. Production and decay both contribute to the observed cross-section. Each of stages is characterised by its probability (production probability and probability of the decay into an interesting channel called branching ratio) which are to be multiplied for the whole process. Experimentally, there is always additional reduction of the probability of observing the whole process related to detector effects and selectivity of the final signatures.

At the LHC main experimental difficulty as compared to experiments at LEP or HERA is presence of high QCD background for which cross-sections are shown in Fig. 1.1 (inclusive jets cross-section). For instance a total QCD cross-section $\sigma_{tot} = 98.3 \pm 0.2(\text{stat})^{+2.8}_{-2.7}(\text{sys})$ mb [12] is about ten orders of magnitude higher than total cross-section for Higgs boson production of $\sigma_{pp\to H} = 33.0 \pm 5.3(\text{stat}) \pm 1.6(\text{sys})$ pb [13]. For that very reason sophisticated selection strategies have to be applied on-line in order to drastically reduce rate of background events while preserving rare signal of interest.

ATLAS [14] experiment consists of about $140 \cdot 10^6$ electronic readout channels. After a commpression (mostly zeros cancellation) the information recorded for typical collision amounts to about 1.5 MB. With a slightly simplifying assumption that collisions take place at a rate of 40 MHz a data stream of 60 TB/s is generated by running experiment (without zero cancellation about 1PB/s would be generated). With the experiment usually taking data for $10^7 s/year$ about 500 EB (exabytes) of data would need to be recorded. That number significantly exceeds technical capabilities for the transfer, storage and post-processing of contemporary computer systems. The figures for the ATLAS experiment are quoted here however the situation is nearly

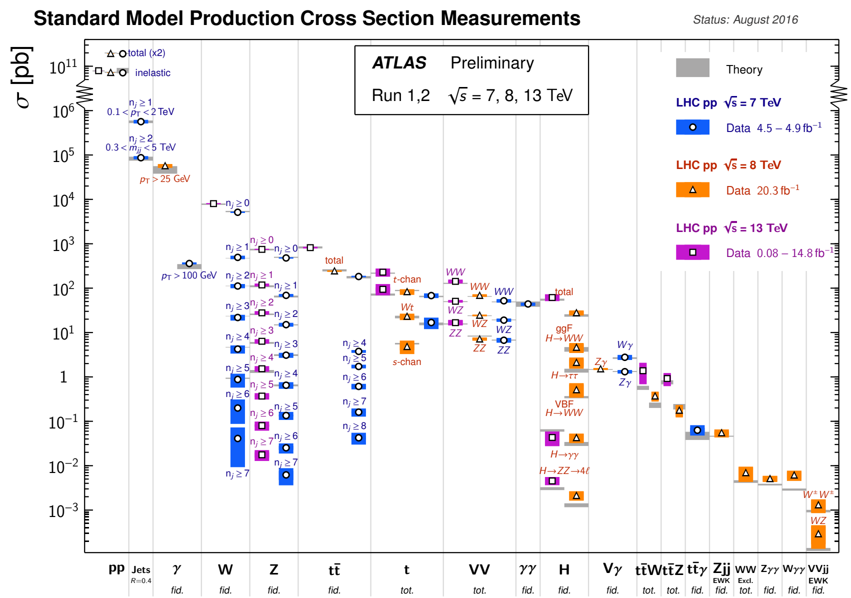

FIGURE 1.1: A recent summary of several Standard Model total and fiducial production cross-section measurements illustrating relative scales of signals and backgrounds [11].

identical in CMS [15] detector and to a lesser extent in LHCb [16] and ALICE [17] experiments.

However, as discussed in previous paragraph small cross-sections of interesting processes provides a means to stay within the limitations of contemporary technology by constraining the output data to contain only fraction of events of interest. An infrastructure devoted to this task is called trigger system. It is tightly coupled with the data acquisition system (DAQ) responsible for collection of detector data, synchronisation and recording. Elements of trigger system provide filtering components used by DAQ at subsequent stages of data collection. Therefore trigger systems are not single uniform component. Depending on the place in the DAQ chain where the trigger component is located technologies vary significantly (from custom electronic chips to high level software). On the other hand it provides consistent and uniform filtering functionality. Trigger systems of LHC size experiments have capability of applying concurrently multiple filtering strategies in order to be able to collect data for multiple physics measurements at the same time thereby maximising scientific outcome of the data collected with multipurpose detectors. From that perspective yet another level of modularity has to be achieved in the design and construction trigger system. In summary, for the reasons described above, trigger systems are indispensable elements of ATLAS and similar experiments.

1.1 Signatures for triggering

The trigger implements strategy which reflects the main physics goals of the experiment. More details are given in chapter 5 with discussion of implementation of tiggers for specific physics signatures. Due to a high QCD backgrounds at hadron machine most derirable signatures have significantly distinct from the features of QCD jets. The Higgs boson discovery is a primary goal of the experiment and trigger. In particular triggering on $W^{\pm}, Z$ bosons decays is a important tool to collect statistics needed for this and many other studies. Therefore the trigger system is built for triggering on muons and electrons with a high efficiency. For the same reason the hadronic tau trigger, more difficult experimentally, is also fully supported by the ATLAS trigger system. To enhance signal efficiencies for $W^{\pm}$ missing momentum is also available for online selection. Due to the clean signature of the Higgs boson decay to two photons single and di-photon triggers are also key ingredients of the trigger menu. Collecting event samples with QCD jets is not difficult at the LHC, however recording a high purity samples with high p_{T} or complex topologies is by no means straightforward. In this case high selectivity has to be achieved already at the hardware trigger stage where compute resources are limited. A filter capable of discriminating light from heavy flavour jets is highly desirable when hadronic branching ratios significantly exceed these with leptons in the final state despite the high purity and efficiency of triggers for the latter. All the above mentioned signatures with high p_{T} thresholds are explored in searches for new particles. Finally, a trigger that efficiently identifies collisions is critically important during the commissioning phase and heavy ion data taking and needs to be supported for a general purpose detector like ATLAS.

1.2 Trigger system figure of merit

1.2.1 Efficiency

The quality of the trigger system can be assessed by quoting a few statistical properties. The first one, an efficiency, is the probability of a positive trigger decision under the condition of the presence of a signal in the event. A good trigger system has an efficiency close to 1. Precise knowledge of this parameter is needed in the data analysis as it allows a translation from the number of registered events to the actual number of events. Such correction procedure is called a correction for the trigger efficiency. In practice, efficiencies are characterised as a function (histogrammed or parametrised analytically) of a few key phase space parameters. They are usually plotted as a marginal distributions, an example of which is shown in Fig. 1.2.

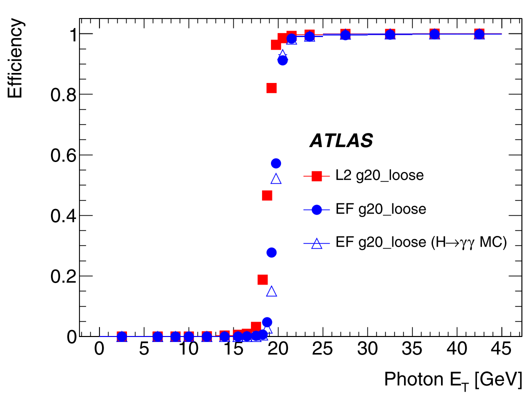

Figure 1.2: Photon trigger efficiency as a function of the offline reconstructed photon transverse energy, E_{T} [18]. The aim of this is to select photons of $E_{\mathrm{T}} > 20$ GeV (EF g20_loose). The efficiencies at the early stages of triggering are also shown (L2 g20_loose). The efficiency has been measured using MC simulation with the Higgs boson decaying into two energetic photons (ATLAS golden channel).

As illustrated, the efficiency as a function of the most important parameter (energy, momentum) can be divided into the three regions. The region of A) inefficiency, B) the rise region (turn-on), and C) plateau. In the construction of the trigger, the ideal is to reach zero efficiency in the region A and close to one in the region C. The narrower the region B is the better.

The efficiency estimator ϵ in the range i of variable x simply is a relative frequency f:

$$f_i(x) = \frac{k_i(T)}{n_i}$$

where $k_i(T)$ denotes number of positive trigger decisions and n_i number of events with signal. In the case where k approaches n, and n is close to zero more rigorous estimations have to be used [19].

The reasons for inefficiency can be, among others, an incomplete phase space coverage by the triggering detectors. This is the case in the ATLAS experiment where the fast muon detectors [20] do not cover the full azimuthal angle due to constructional constraints.

Also due to the time contraints for trigger reconstruction simplified algorithms (or even hardware implementable approximations) are used instead of precise ones. Therefore the variables used in the trigger selection are only roughly correlated with the offline references. The worse the correlation between the variables the wider the turn-on region is. A similar incompatibility is caused by inexact detector calibrations known during the trigger reconstruction.

It has to be stated that above mentioned factors cannot be eliminated. For instance, the exact calibration of detector elements changes with time and is often assessed using the data for which the calibration has be made. In consequence the trigger algorithms always work with slightly outdated calibration of the detectors.

Also, reference quantities are often defined after data taking. An example is the estimation of the jet energy scale in presence of the pile-up [1] which is known only months after careful analysis of the data taken. Moreover that incompatibility can be also analysis-specific. For instance additional kinematic cuts may influence the measured efficiency.

1.2.2 Purity

An equally important parameter of the trigger system is purity. It can be understood as the conditional probability of a positive decision in the absence of a signal. In practice, the inverse quantity is used, a frequency/rate of positive trigger decisions. The efficiency and purity are negatively correlated. That means that an increase of purity (drop of the rate) reduces the efficiency, usually in a nonlinear manner. The trigger rate is a parameter which does not play any role in a physics analysis. It is mainly an operational concern. Failing to control the trigger rate results in a sample of low purity which requires more data handling effort. In the case of many selections preformed concurrently an impure selection reduces the bandwidth that could otherwise be used by other triggers.

About

The rest of this book presents the design, implementation and performance of the ATLAS High Level trigger. In the introductory material key aspects considered in the design of the system are recalled. Then the design and implementation during Run 1 are described and ongoing developments during Run 2 of the LHC are presented. In chapter 5 the perfomance of the trigger during the LHC Run-1 is reviewed with emphasise on the performance of the triggers essential for key measurements but with discussion of less standard applications like triggering in Pb-Pb runs. In this chapter upgrades for Run 2 will be briefly presented. The reader's

[1] Several collisions per bunch crossing are referred to as the pile-up. The average number of collision per bunch crossing is denoted as $\langle\mu\rangle$.

knowledge of experimental particle physics at a higher than elementary level is assumed in the text, however a brief introduction to trigger-specific concepts and terminology is given.

2

The ATLAS experiment and Trigger Architecture

2.1 The ATLAS experiment

The ATLAS detector surrounds the LHC collision point closest to the main CERN site in Meyrin, Switzerland. It is located in the experimental cavern about 100m under the surface. It is a large piece of equipement, 40 meters long and 25 meters tall. In addition to the experimental facilities; cooling, gas supplies and powering have to be added which reside in an adjacent cavern and on the surface. To the whole picture, globally distributed data centers have to be added for completness.

About 10 thousand people participated in the construction and commissioning of the ATLAS experiment over a period of 20 years. At the moment about 3000 authors sign publications of the experiment. In addition to maintaining and operating subsystems, upgrades are planned and studies performed to prepare the detector for high luminosity LHC. All activities in the experiment are carried out collaboratively, as was the development of the systems described in this monograph. A fusion of many backgrounds, skills and knowledge and the dedication of the researchers involved allows enterprises like the ATLAS experiment to be brought to the stage of succesfull scientific project.

2.2 ATLAS detector

The detector surrounds the luminous region where the beams cros that is about 1 m long and 1 mm wide. Thus the detector has a cylindrical geometry with complete coverage in azimuthal angle and near complete pseudorapidity range[1]. A drawing of the detector is shown in Fig. 2.1. Only the main sub-detectors are briefly introduced here since the smaller ones are mostly for the triggering purposes and will be described in the text where appropriate.

2.2.1 Inner Detector

The innermost volume is occupied by the Inner Detector (ID) [21, 22], a tracking detectors immersed in the solenoidal 2 Tesla field aligned with z axis generated by a superconducting magnet. It consists of three sub-detectors utilising various technologies to provide complementary measurements. They are arranged in a cylindrical geometry with the barrel and end-cap parts as shown in Fig. 2.2. The barrel and end-cap sub-detectors are labeled and dimensions are given.

Closest to the interaction point are three layers of silicon Pixel detector (Pixel). An additional layer in between the beam pipe and innermost Pixel layer was added during the maintenance period 2012-2015 [23]. There are 1744 identical detector sensors of size 19×63 mm^2 having 47234 50×400 μm^2 pixels used in the barrel layers (1456 modules) and the end-caps (288 modules). The number of readout channels of the Pixel detector amounts to ∼80 million. The resolution provided by Pixel detector is 10 μm in ϕ and 115 μm along z (in the barrel) and R (end-caps).

Further from the centre, 4 in the barrel, 9 in each end-cap double-sided layers of SemiConductor Tracker (SCT) are mounted. It uses two kinds of modules, rectangular in the barrel and trapezoidal in end-caps. Each of the 4088 two-sided modules is made up of two silicon wafers with strip shaped readout. Sensor strips on both sides of the module have a small stereo angle of 40 mrad in order to provide z/R position reconstruction. The SCT has binary readout with ∼6 million of channels. The hit position resolution in the SCT is 17 μm in ϕ and 580 μm along z (barrel) and R (end-caps). Both, Pixel and SCT detectors, cover pseudorapidity $|\eta| < 2.5$.

The Transition Radiation Tracker (TRT) is located outside of the SCT. It is made out of many layers of gass filled straw tubes (4 mm of diameter) interspaced by plastic fibres (barrel) and foils (end-caps) as transition radiation sources. Straws in the barrel are 144 cm long and are arranged along the z axis. In the forward regions they are 39 cm long and are arranged radially. The gas mixture, Ar + Xe, allows for the electron-hadron (mostly $\pi^{\pm}$) identification through a high absorption of low energy transition radiation photons by the Xe and thus large amplitude signals induced as compared to minimum ionising particles. In total, the TRT consists of 50000 straws in the barrel and 25000 in the end-caps leading to 350k readout channels. The TRT provides 130 μm position resolution thanks to the many hits left by the track and covers a pseudorapidity range, $|\eta| < 2$.

The ID tracking detectors do not provide fast enough readout, nor is track reconstruction possible quickly enough that tracking information could be made available in the earliest, hardware, stage of the trigger system. Information from the ID is an indispensable ingredient in the later stages of online events filtering. Most of the rejection power is achieved by the use of the ID tracking.

[1] ATLAS has a right-handed coordinate system that is centered on the nominal interaction point and oriented so that x, y, z axes point to the center of the LHC ring, upwards and along the beam axis respectively. Pseudorapidity, defined as $\eta = -\ln(\tan\theta/2)$ is a function of polar angle, θ between the object and the azimuthal plane. The azimuthal angle, ϕ is measured from the x axis and has range from $-\pi$ to π.

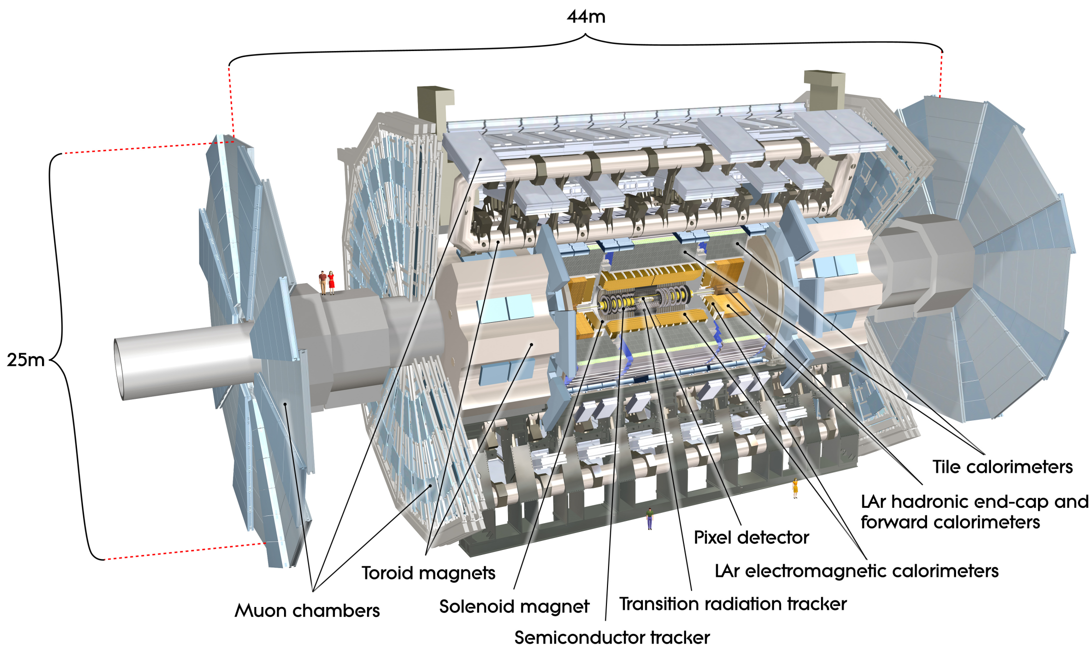

FIGURE 2.1: A sketch of the ATLAS detector. The main elements are indicated in the legend. For a size comparison human silhouettes are drawn in a few places [14].

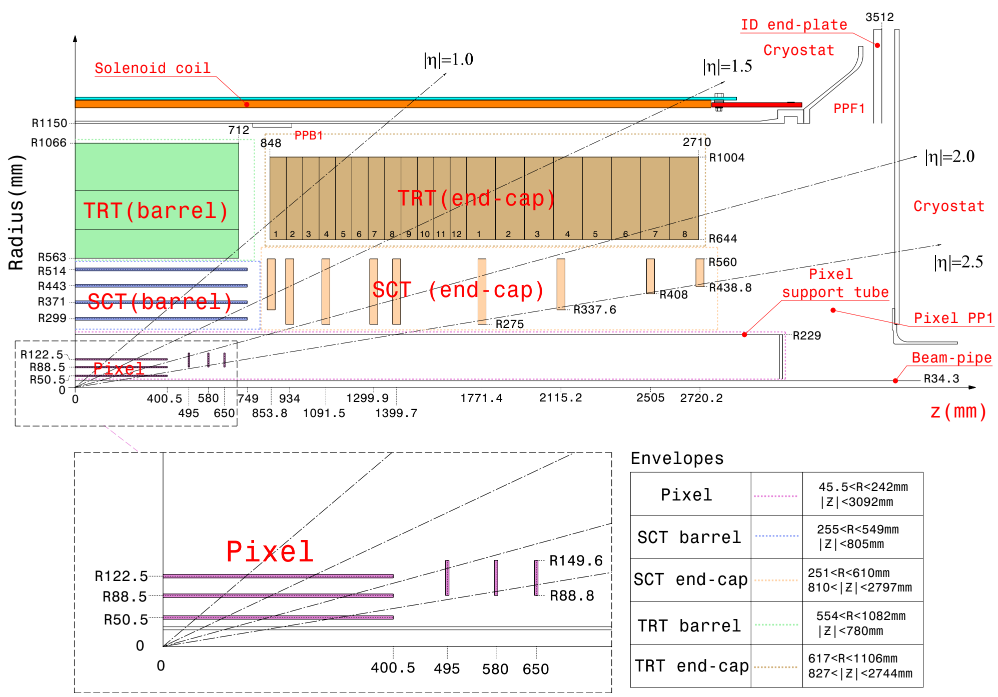

FIGURE 2.2: Detailed layout of the ATLAS Inner Detectors [14].

2.2.2 Calorimeters

Outside of the ID and solenoid the calorimeters are located. First the Electromagnetic Calorimeter (EMC) and then Hadronic Calorimeter (HCAL) covering range of $|\eta| < 4.9$ and the full azimuth. The sketch of the arrangement of the various sub-detectors is shown in Fig. 2.3. The central calorimeter is referred to as the barrel, the hadronic calorimeter of the same radii as the barrel is called the extended barrel. At $\sim 1.7 < |\eta| <\sim 2.5$ the hadronic and electromagnetic end-caps are located and in the even more forward region $3.1 < |\eta| < 4.9$ the Forward Calorimeter (FCal) is located. The electromagnetic barrel, EMC, HCAL end-caps and FCal use Liquid Argon (LAr) as the active medium and therefore are wrapped in a cryostat and operate at 90 K.

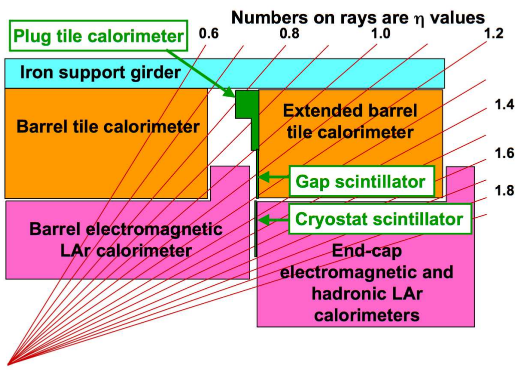

Figure 2.3: Sketch of ATLAS Calorimeters arrangement [14].

In the barrel the EMC uses lead as an absorber shaped in accordion wafers to avoid cracks between cells. Cells are arranged in a projective manner in order to improve particles impact position resolution. There are three layers of cells in the EMC in $|\eta| < 3.1$ and one in the FCal. The cell geometry of the barrel EMC at $\eta \sim 0$ is shown in Fig. 2.4. The granularity is reduced and the number of layers varies with growing pseudorapidity as described in Ref. [14](Table 1.3). In the electromagnetic layer in the FCal copper is used as the absorber and plastic fibres surrounding absorber rods to collect the light. EMC LAr signals are sampled with 12 bit precision and are integrated for about 600 ns. In total the LAr has 115.5k and 69.5k readout channels in the barrel and end-caps respectively. The FCal part has an additional 1.7k channels.

The HCAL in the barrel and extended barrel has is made of plastic and steel tiles and therefore operates at room temperature. The detector is assembled from wedge shaped modules subdivided in the three cells at various interaction depths read out by photomultipliers (two per cell). 64 Tile modules in ϕ and projective compartmentalisation along z results in approximate $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ granularity (0.2×0.1 in last layer). The Tile calorimeter wedge is shown in Fig. 2.5. The HCAL in the end-cap part has two layers constructed identically to the first electromagnetic layer, but with tungsten as the absorber. There are in total about 10 thousand 10 bit resolution channels in the Tile HCAL, 6.6k in the hadronic end-cap and 3.5k in the FCal.

The EMC and HCAL have sufficient depth for containment of the electromagnetic and hadronic cascades. They provide more than 22 radiation lengths and the HCAL more that 11 interaction lengths (including the support structure) which allows for precise measurements as well as for elimination of a leakage (punch-through) to the Muon Spectrometer (MS). The

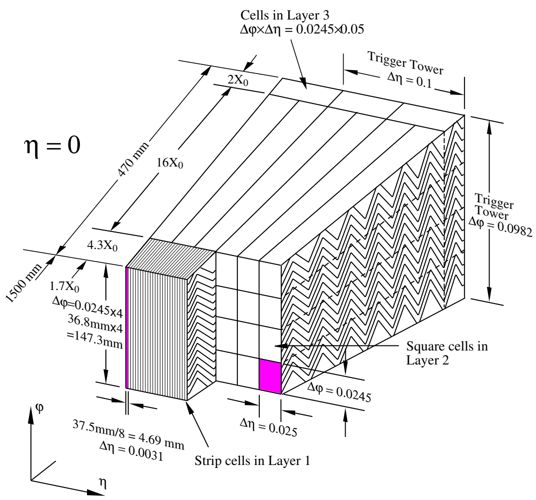

FIGURE 2.4: Drawing of the Electromagnetic Calorimeter cell arrangement in the barrel part [14].

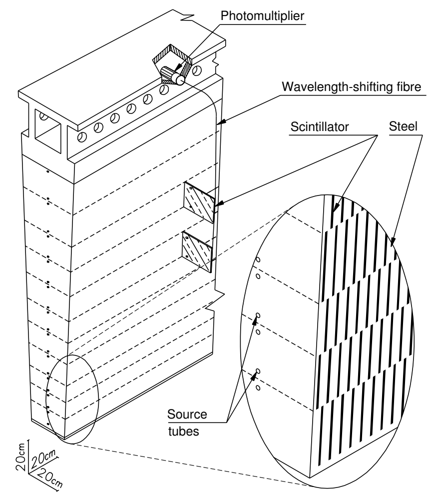

FIGURE 2.5: Drawing of the ATLAS Hadronic Tile Calorimeter module [14].

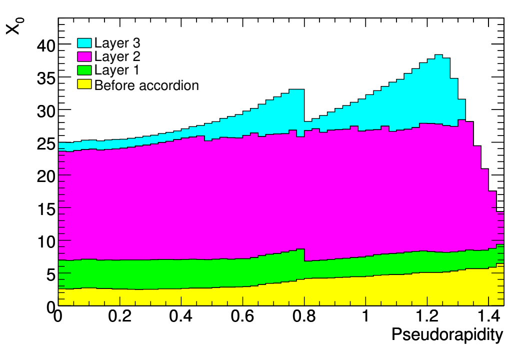

FIGURE 2.6: The material of the layers of the LAr EMC in the barrel section and the ID material in front of it in radiation lengths, X_0. More details, including distributions in the gap and the forward regions can be found in Ref. [14].

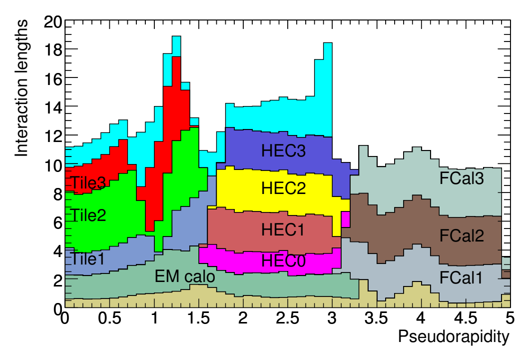

FIGURE 2.7: Material in front and inside of the hadronic calorimeter layers in interaction lengths. The support structures of the MS are drawn in cyan [14].

material distribution as a function of pseudorapidity is shown in Fig. 2.6 and Fig. 2.7.

In addition to the full granularity readout from each cell, the calorimeters are quipped with a coarse granularity readout for triggering purpose. Signals from groups of cells, trigger towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ (in the central region and bigger in the forward), formed in the front-end electronics are transmitted to the first stage hardware trigger within 1 μs. In further selection stages and in offline reconstruction the full cells information is used for precise calorimetric reconstruction.

2.2.3 Muon Spectrometer

Outermost of the ATLAS detector system, is the MS providing reconstruction of muon tracks in a magnetic field produced by large, air-core, superconducting toroidal magnets (in Fig 2.1 the coils and MS chambers are shown). A field of 0.5 T is generated in the barrel and a less uniform, but stronger field of 1 T in the forward regions. The large dimensions of the spectrometer are dictated by the requirement of about 10% p_{T} resolution for very high p_{T} 1 (TeV) muons.

Main detectors of the MS are the Monitored Drift Tubes (MDT) , covering $|\eta| < 2.7$ in the barrel and end-caps, designed to provide precision measurements. The MDT has three layers except in the very forward region $2 < |\eta| < 2.7$ where the innermost MDT chambers are replaced by Cathode Strip Chambers (CSC). There are 80 variants of the MDT modules as the coverage of this detector is huge and 1088 chambers have to adapt to various geometrical constraints. In each, drift tubes of 3cm diameter are filled with an Ar/CO_2 mixture and are used to collect ionisation signals which measure timing that translates to a position, provides 80 μm resolution in the bending plane. Combing information from layers of tubes in the chamber a 35 μm resolution is obtained. The construction of MDT chamber is shown in Fig. 2.8.

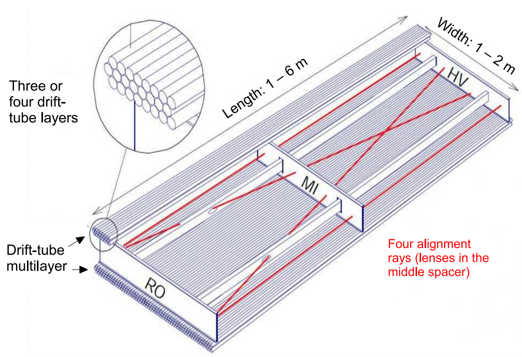

Figure 2.8: Illustration of the MDT tube (left) and chamber (right) with the number of tubes and support structure [14]. The HV and RO symbols denote the locations of the high voltage supply and readout electronics respectively.

The CSC are used in the forward region as they are capable of handling higher muon rates (multiple tracks per module) than the MDT. They are multiwire proportional chambers with anode wires aligned radially and perpendicular to them cathode strips housed in 8 small and 8 larger modules with 4 layers each. The CSC has 60 μm resolution in the bending plane, 5 mm in azimuthal direction and a single plane has time resolution of 7 ns.

For triggering purposes the MDT chambers are surrounded by Resistive Plate Chambers (RPC) at $|\eta| < 1.05$ and Thin Gap Chambers (TGC) $1.05 < |\eta| < 2.4$. These detectors are capable of delivering signals within tens of ns after the passage of a muon and have a time resolution of 15-25 ns thus allowing for the association of the signals to a bunch crossing. The RPC chambers are attached to the barrel MDT modules, two at smaller radii and one in the outermost layer. With this arrangement, the two inner chambers give a good efficiency for low p_{T} muons while, when combined with outer stations a good p_{T} resolution is achieved with moderate readout segmentation. Each RPC station has two independently read out layers so that a single muon track leaves up to 6 hits (typically 3 are needed for high p_{T} tracks). With orthogonally arranged readout strips, the resolutions along z and ϕ are 10mm and complement MDT measurement in non bending plane.

The TGC is a very similar technology to the CSC and was chosen due to a good position (2-6 mm in R, 3-7 mm in ϕ) and time resolution (4 ns) . In addition to the triggering function it provides a second coordinate measurement to complement the MDT. The TGC has 4 stations, one at small z and 3 surrounding the middle MDT module.

The arrangement of MS elements is shown in Fig. 2.9. The MS readout provides 345k MDT, 30.7k CSC, 373k RPC and 318k TGC channels.

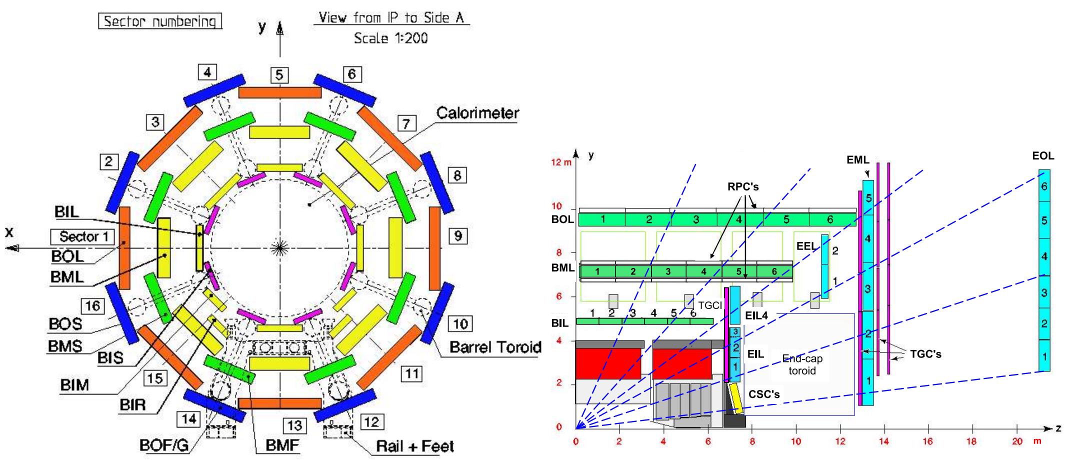

Figure 2.9: A cross-section of the MS illustrating the arrangement of chambers in the azimuthal plane in the barrel (left) and in the bending plane (right) [14]. In the drawing on the left various colours are used to represent the different types of chambers.

2.3 Overview of Trigger and DAQ in Run 1

The architecture of the atlas Trigger and Data Acquisition (DAQ) system is show in Fig. 2.10 for Run 1. Two major sub-detectors, the calorimeters and fast muon detectors [25] are designed to be able to deliver signals sufficiently fast so that the decisions based on that signals can be made at the LHC bunch crossing frequency of 40 MHz and with latency small enough that the on-detector readout buffers do not overflow (2.5 μs). Signals from these detectors are delivered to specialised systems capable of finding the signatures of physics objects in a short time. These are the Level 1 Calorimeter Trigger (L1Calo) and Level 1 Muon Trigger (L1Muon) systems. In addition signals from smaller detectors like the Minimum Bias Trigger Scintillators (MBST) or Absolute Luminosity for ATLAS detectors (ALFA)[2] are made available to the Level 1 Central Trigger (L1CT) and can contribute to the event acceptance. This systems are collectively

[2]The signals from this detector is used only in a runs with special, low luminosity, runs.

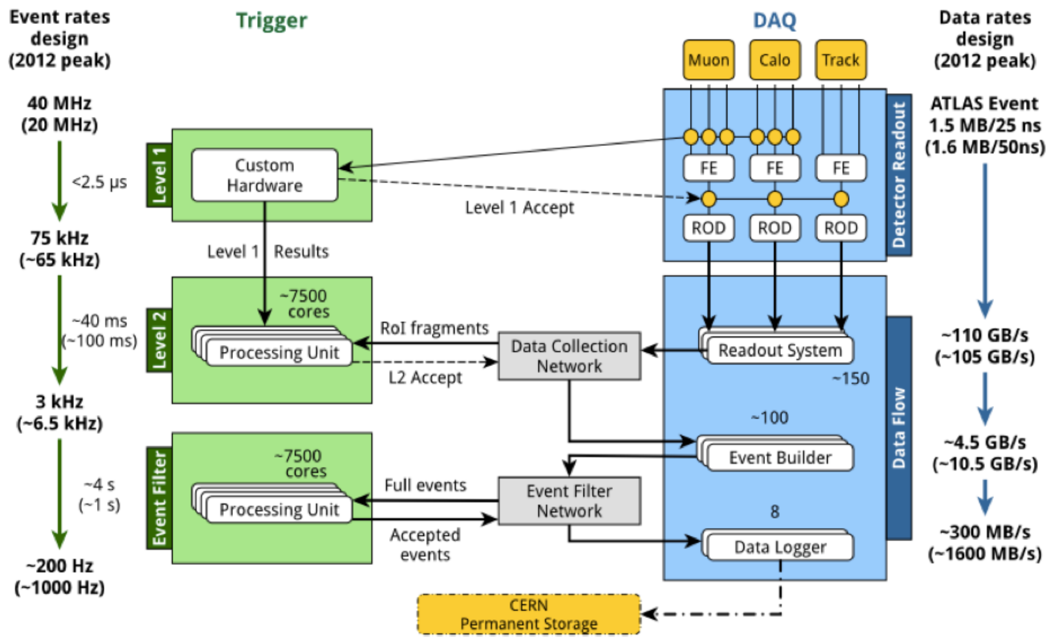

FIGURE 2.10: Schematic view of the ATLAS Trigger and DAQ architecture used during the Run 1 with nominal (upper ones) and average operational parameter as of the end of the run in 2012 (bottom ones) [24]. In example 40 MHz was the designed L1 input while in the Run 1 the LHC bunch filling scheme with 50 ns spacing instead of nominal 25 ns provided colliding bunches at the rate of 20 MHz.

referred to as the First Level Trigger (L1). For accepted events the detector data is transfered via optical links to the Readout System (RoS) to await the High Level Trigger (HLT) decision. The RoS system consists of about 150 computers, each receiving data through, up to four, input cards (ROBINs) which can serve up to three links. The arrangement of the data links is such that adjacent regions of the detector are served by the same RoS PC and the data "density" (amount and frequency at which it is requested) is optimised such that typical operation exerts an even load on the system, as much as possible.

A record containing information from the L1 calorimeter, L1 muon systems and L1CT is assembled concurrently by Region of Interest Builder (RoIB) component and passed through the L2 supervising machine to an L2 node. The mean allotted time to take the decision for a particular event at L2 depends on the size of the farm and increased from 40 ms in the design to 75 ms at the end of the Run 1. An additional limitation on the maximum time is also applied in order to guarantee safe system operation and is set as high as 15 seconds[3]. In a normal operational conditions L2 accepts about 10% of events. For rejected events, the RoS system is instructed to drop the data related to that event from the buffers. On the contrary, for accepted events the detector data fragments are assembled by the event building component, called the Sub-Farm Input, into a complete record including the L2 decisions, and are made available to the third, Event Filter (EF) , stage. There, the decision taking process is continued and as for L2 the reject/accept decisions dictate the fate of the event. Accepted events are

[3]That is one particular setting optimal for running conditions at a particular time, it can vary depending on the trigger menu or beam types. For instance in Pb-Pb runs it was usually increased to accomodate the ultra-central events with very high detector activity.

transfered to the Sub-Farm Output to be recorded for offline reconstruction and analyses while rejected events are simply abandoned. Details of the decision taking process at the L2 and EF, collectively called High level Trigger (HLT), is the subject of chapter 4.

2.4 Trigger and DAQ in Run 2

A slightly modified approach was used for Run 2 as shown in Fig. 2.11. Namely the two levels of the HLT and event building were merged into one component. The main motivation for this was reduction of the RoS requestes load occurring in the Run 1 RoS system that had to handle requestes for both L2 and event building. In the Run 2 system, the detector data fragments are initially retrieved as for L2 as in Run 1 and if a certain stage of processing is reached, all remaining fragments are retrieved and the event built. For accepted events such, the assembled event fragments are returned to the data flow component that is responsible for recording to permanent storage. In the case of rejection, a signal is sent back to the RoS to drop buffered data for the current event. This architecture also provides a significant simplification as compared to the previous one in terms of number of components and the networking needed to support it. Also, more refined triggering strategies can be implemented in HLT in Run 2 as will be discussed in chapter 5.

An additional hardware component of relevance for triggering is included between L1 and HLT, the Fast TracKer (FTK) [26]. FTK obtains signals from silicon tracking detectors for every event accepted by L1 (100 kHz) and provides tracking information for the entire event to the HLT right at the beginning of the reconstruction. The early use of the tracking information in the HLT selection allows triggering on demanding signatures involving hadronically decaying τ and b-jets. In addition the ability to distinguish between the primary interactions and contributions from the pile-up of other p-p interactions enables high trigger purities and thus a linear scaling of trigger rate with luminosity to be maintained. At time of writing, FTK was in the final commissioning stages and so no results from this system can be presented besides those from simulation.

2.5 Level 1 calorimeter trigger

The hardware calorimeter trigger (L1Calo) [25] receives signals from calorimeters cells from the front-end electronics, grouped into the trigger towers 0.1×0.1 in $\Delta\eta \times \Delta\phi$ (bigger in the forward region $\eta > 2.4$). In the L1Calo PreProcessor these inputs are timed, digitised and fed through lookup tables in order to subtract pedestals, suppress noise, calibrate the energy and if needed mask out problematic towers. Two clusters processors then operate on the towers to locate regions of interesting activity. The Cluster Processor Module (CPM) attempts to find clusters of energy corresponding to e/γ and τ's. A sliding window algorithm is used for that purpose. In each 4×4 windows the four two-tower sums in the center towers are compared with a threshold. A window is considered to contain local maxima if any of the sums passes the threshold, as illustrated in Fig. 2.12. For the e/γ trigger it is required that there is no signal in the hadronic calorimeter (hadronic isolation), while for cluster originating from the decay of the τ lepton the transverse energy in the hadronic calorimeter is used. Transverse energy deposited in towers around the 2×2 core define isolation which enables the selection to be discriminative agains larger deposits from jets and background activity caused by pile-up. A similar algorithm is implemented for finding jets in the Jet/Energy Module (JEM). The main difference in the

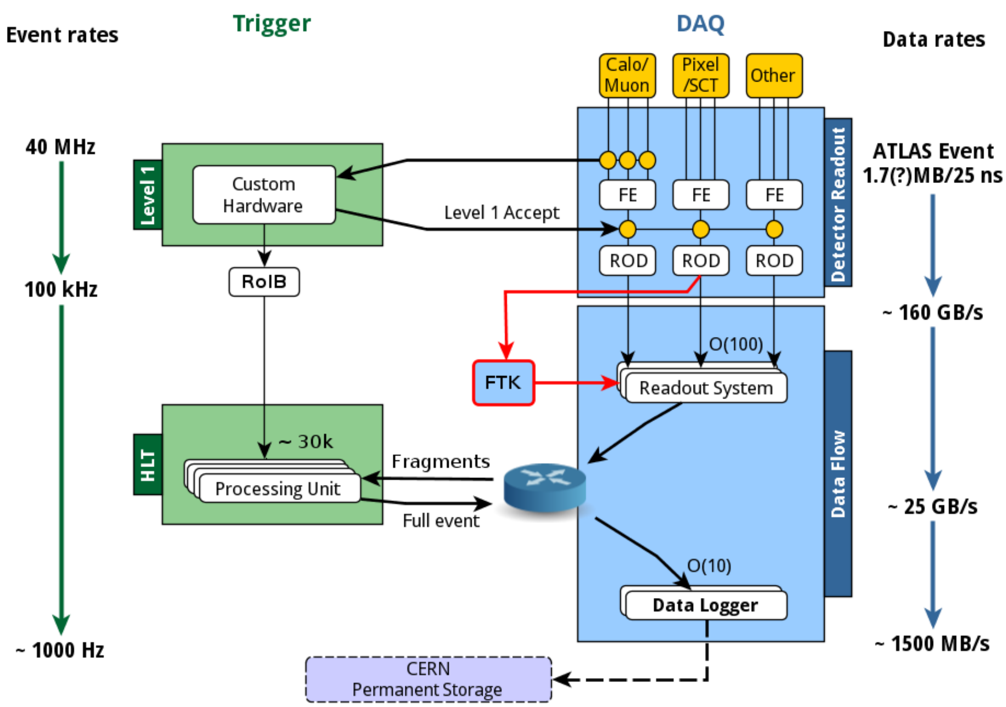

Figure 2.11: Schematic view of the ATLAS Trigger and DAQ architecture used during Run 2 with the average expected operational parameters.

JEM is that towers from electromagnetic and hadronic calorimeters are summed up implicitly and the size of the core of the window is configurable.

The summation windows are arranged so that they overlap by one row of towers. As a result the local maxima can be exceeded in two adjacent widows and therefore the merging module (CMM) is added in the L1 calorimeter system to disambiguate the position of deposits. The position of the deposits and the thresholds passed are propagated to HLT to seed the reconstruction.

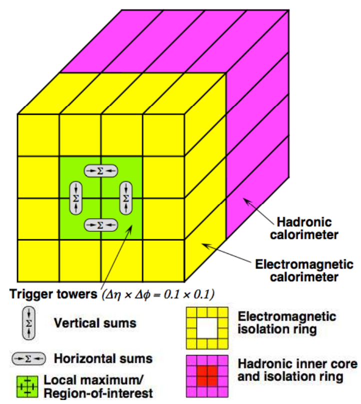

FIGURE 2.12: Elements of the hardware algorithm used to find signals from e/γ and τ in the ATLAS calorimeter [27]. The labels "vertical-sums" and "horizontal-sums" refer to the summation along the η and ϕ directions in cylindrical coordinates respectively.

2.5.1 Commissioning and performance

The performance of the L1 calo trigger is quantified mostly in terms of the resolution or bias with respect to the energies reconstructed by the offline algorithms. The first parameter tuned in order to obtain high precision was the timing of the trigger towers signals so that the digitisation is optimal in terms of the linearity and strength of the signal. A prominent example of the effect of better timing is visible in the efficiencies of a L1 electromagnetic trigger shown in Fig. 2.13 for three different running periods for which the timing adjustments were gradually improved [28, 29].

Besides the improvements of timing the extraction of energy value from the signal shape has been improved over time by the use of a fine calibration function encoded in a LUT. Resolutions presented in Fig. 2.14 show optimisations resulting in the disappearance of the energy bias and improved resolution.

2.5.2 Upgrades

For Run 2 the L1 calo system was upgraded to maintain high performance during LHC operation with a 25 ns bunch spacing and increased luminosity and thus pile-up. The most significant upgrade was carried out in the L1 calo preprocessor system which times and digitises the detector signal pulses. The timing resolution improvement was achieved thanks to a replacement of the chips sampling the detector signals with new version with higher frequency (80 MHz) and implementing advanced FIR filters correcting for pedestal variations and effectively stabilising the calibration as a function of pile-up [31, 32]. The latter was of a special importance for the trigger signatures that use large areas of the calorimeter such as jets or missing transverse energy.

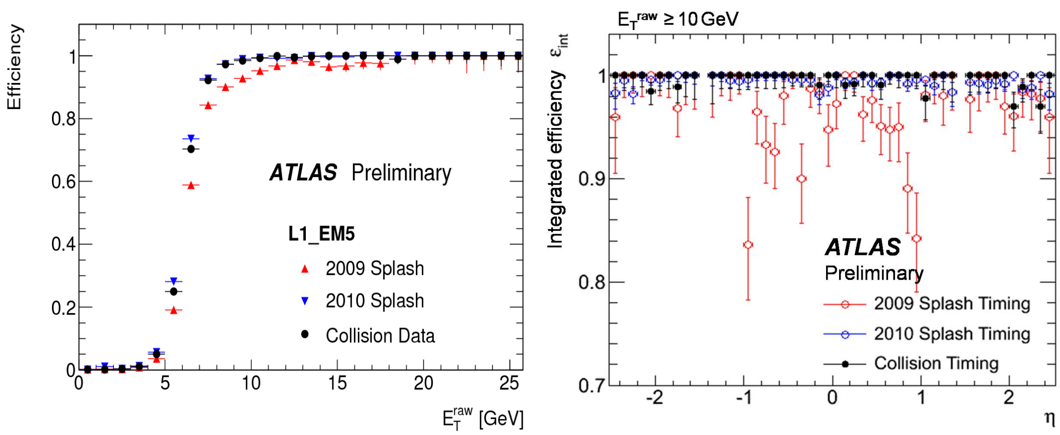

FIGURE 2.13: Efficiency of L1 calorimeter trigger EM5 as a function of raw transverse energy, E_{T} (left), and pseudorapidity, η (right), for three running periods with improving trigger towers timing adjustments [28].

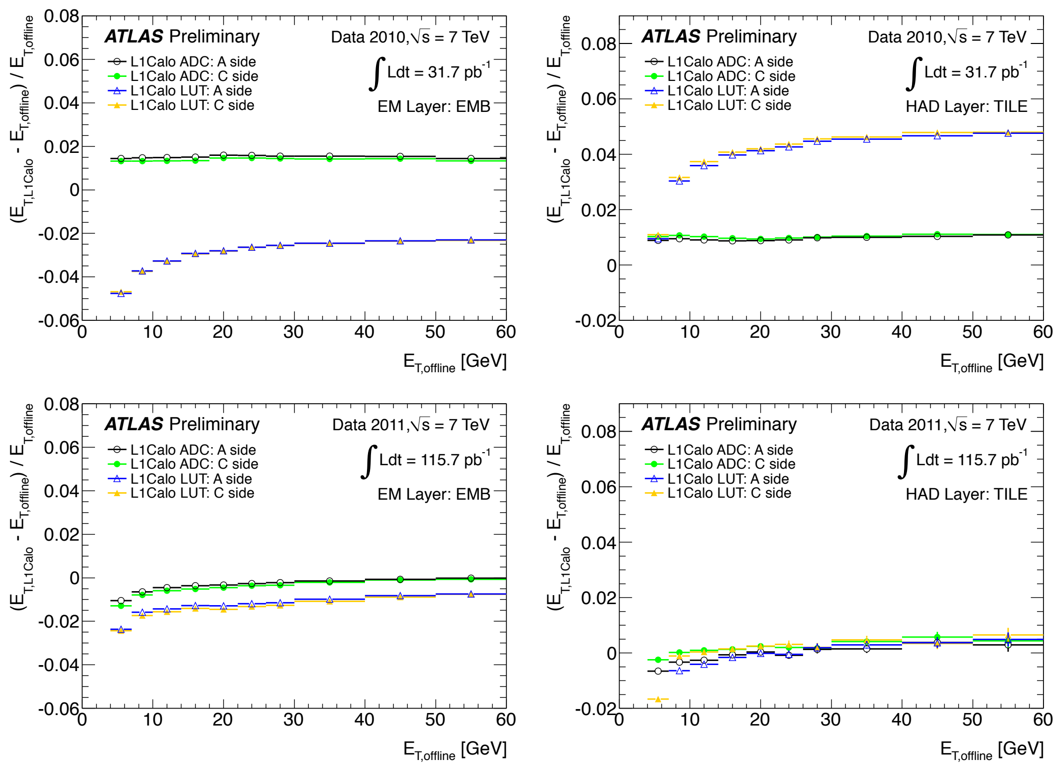

FIGURE 2.14: Energy resolution of the L1 calorimeter trigger in 2010 (top) and 2011 (bottom) for electromagnetic barrel calorimeter (left) and tiles calorimeter (right) [30].

An effect of the improvement, simulated by disabling pedestal filters, is shown in Fig. 2.15. For the EM triggers the performance is not deteriorated while for the $E_{\mathrm{T}}^{\mathrm{miss}}$ significance the highly non-linear behaviour due to the fake triggers is improved so as to be as good as in the case of

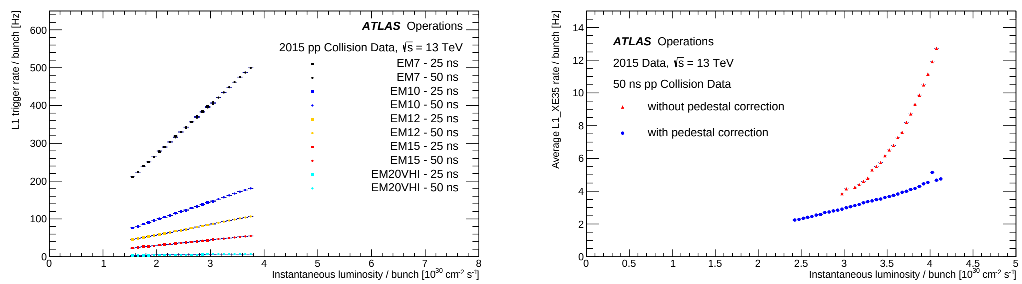

Figure 2.15: Rate of the triggers per bunch crossing as a function of instantaneous luminosity per bunch crossing for electromagnetic triggers (left) and missing transverse energy significance (right) [30].

the EM triggers.

Another significant addition to the L1 calo system for the Run 2 is the topology processor (L1Topo) allowing the combination of trigger primitives (trigger objects TOBs) in various ways. This system obtains TOBs from the L1 calo merger modules and L1 muon processor and is able to perform simple calculations on these objects. Quantities like invariant mass, M_{inv}, radial distance, ΔR and similar can be calculated within the time constraint of the L1 system. By applying cuts on such quantities a significant reduction of the background can be achieved or a reduction of the E_{T} thresholds, as compared to the combined triggers without any topology requirement, and thus improved signal efficiencies. Several case studies are presented in Ref. [24]. The L1Topo system algorithms are part of the firmware and thus their logic (e.g. formula) and parameters can be reprogrammed and so optimised to the running conditions.

For the operations after the LHC Run 2 at yet higher luminosities ($\mathcal{L} = 3 \times 10^{34} cm^{-2}\text{s}^{-1}$) an upgrade to the cluster processors is planned. The CP and JEP will be replaced by eFEX and jFEX processors correspondingly providing higher resolution. The first will be fed with more data (finer granularity towers from electromagnetic calorimeter) in order to perform shower shape quantities estimation which will help to reduce jet backgrounds. The jFEX will implement jet, hadronic taus and $E_{\text{T}}^{\text{miss}}$ like algorithms that will be more robust against the pile-up. An additional component responsible for the extraction of large area objects, the gFEX, is also planned [24, 33].

2.6 Level 1 muon trigger

The Level 1 muon trigger is a large-scale coincidence logic processor. The MS RPC in the barrel $|\eta| < 1.05$ and TGC in the endcap ($1.05 < |\eta| < 2.4$) provide inputs to the trigger system [20]. The RPC system consists of three double layered planes while the TGC has four double wheels (out of which only three were used initially). Readout from the detector layers provides both η and ϕ information. The geometrical acceptance of the TGC is nearly complete while in the barrel only about 80% of the surface is covered by the RPC planes because of the space needed for calorimeter services (at $|\eta|$ around 0), elevators at the bottom of the detector and the toroid magnet support structures.

Six threshold configurations are subdivided in the three low and three high-p_{T} ones. For each threshold roads in which a muon originating from the interaction point appears can be determined from simulation. A coincidence of the signals from the detectors lying within the bounds of the roads is considered as a signature of a muon. For low p_{T} muons the coincidence of the three signals in the inner RPC planes out of four is required. In the forward region the low-p_T trigger logic require three hits in the outer TGC planes out of four. The high-p_T coincidence logic require an additional hit in the outer RPC plane and signals in the inner TGC planes. The sketch in Fig. 2.16 shows the arrangement of the RPC and TGC chambers as well as the trigger roads. The geometry of the trigger road defines the p_T threshold. The narrower the road the higher the p_T of the muon track needs to be in order to leave hits in a single road.

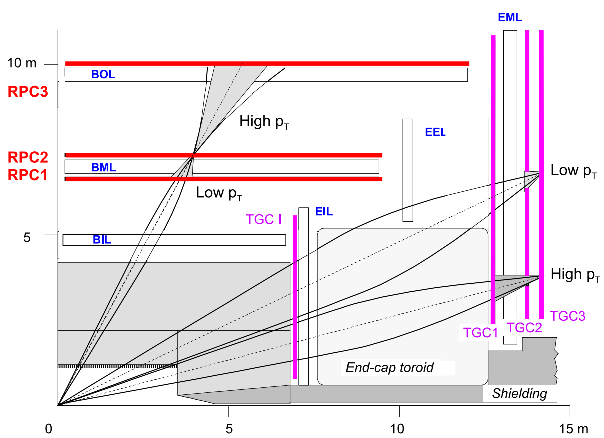

FIGURE 2.16: A sketch of the RPC, TGC MS trigger chambers and the visualisation of the trigger roads [18].

2.6.1 Commissioning and performance

The MS, and L1 muon trigger started operation with cosmic rays well before the first beams appeared in the LHC. For that purpose the coincidence and timing has been defined so that the muon tracks originating outside of the detector cause the trigger to fire. Thanks to that and the use of test pulses the relative timing of signals from the detector chambers were established in advance so that when the first beams were seen about 70% of the RPC triggers were properly timed. The timing for TGC was established mostly using test pulses [18, 34].

In early running the L1 muon trigger also used signals not constrained to roads. The efficiency of this trigger with respect to offline reconstructed muons is shown in Fig. 2.17. Due to the geometric acceptance of the RPC a plateau value of about 70% is reached while for the TGC about 90% of muons are found for $p_T \gtrsim 5$ GeV.

Two configurations with no trigger road contraint were used initially, low (MU0) and high-p_T (MU0_COMM) and four higher threshold intended to be efficient at 6, 10, 15 and 20 GeV. In 2011, dictated by an increase of luminosity, the lowest thresholds were replaced by higher ones and the low p_T =10 GeV and later high p_T =11 GeV thresholds became the primary trigger thresholds. Their performance is shown in Fig. 2.18. The merit of using the low p_T logic (denoted as 2 station coincidence) is visible in the efficiency curves for the RPC. Also, slightly better performance in the plateau regions can be observed as a result of the improved detector timing. In later running the lowest unprescaled trigger was MU15.

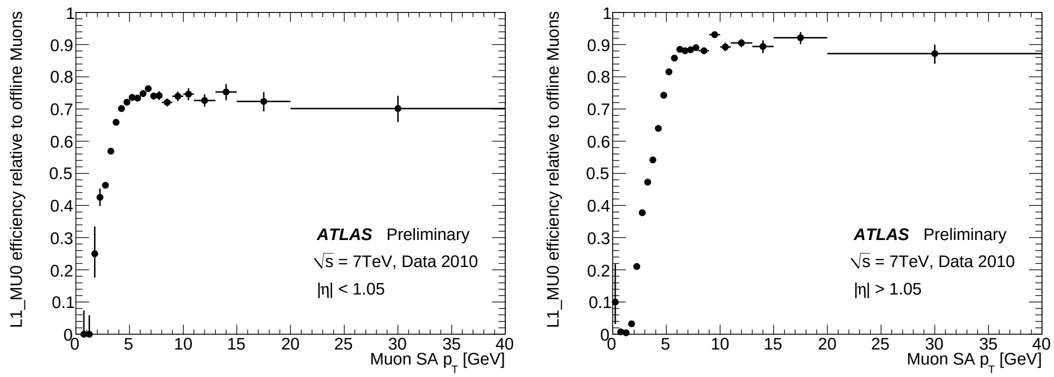

FIGURE 2.17: Efficiency of the L1 MU0 trigger as a function of reference offline muon transverse momentum, p_T, in the barrel region covered by the RPC (left) and the endcaps covered by the TGC (right). The efficiency was measured using a minimum bias data sample [34].

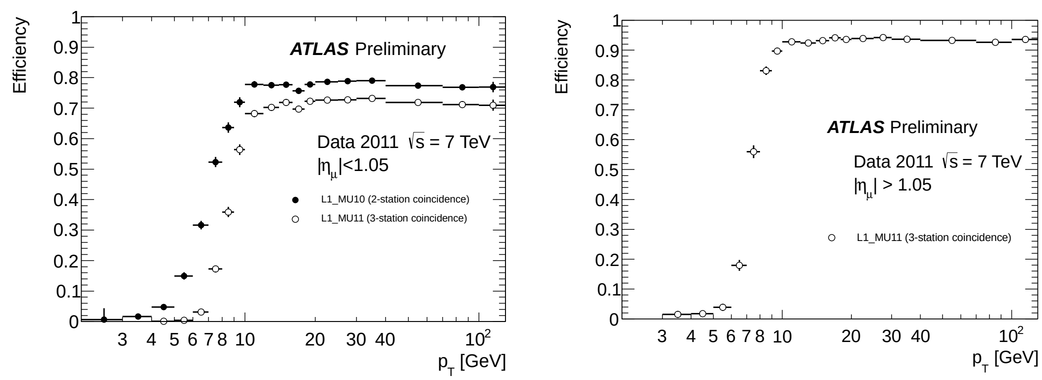

FIGURE 2.18: The efficiency of lowest unprescaled triggers in 2011 [35]. The RPC efficiency is shown for low-p_T MU10 and high-p_T MU11 (left) and the TGC high-p_T MU11 on the right.

2.6.2 UPGRADES

The rate of L1 muons at the end of Run 1 was dominated by the fake signals at $\eta > 1$ from secondary particles generated in the magnet and shielding as well as due to the muons from beam halo. Figure 2.19 shows the rates of muon TOBs corresponding to reconstructed offline muons as well as fake ones. In the RPC region both rates are the same which means that the trigger has high purity while the forward region, handled by the TGC shows a significant excess due to fakes. An increase of the threshold would mean a reduction of the physics reach of the experiment in this phase space. Therefore a coincidence between the inner stations of the TGC and the signals in additional inner layers called Forward Inner was added. The resulting fake-muon rate reduction in shown in Fig. 2.20. About 50% of fake muons are rejected in the areas where this additional coincidence can be exploited. The efficiency for finding real muons was not degraded, as verified from data collected at the start of Run 2 and is shown in Fig. 2.20. Coincidences with Tile calorimeter signals will be added in the future for the same purpose.

Further improvements aimed at fake reduction are planned in the shutdown between Run 2

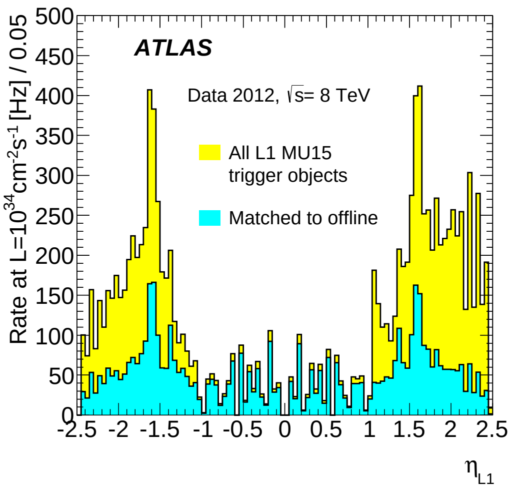

FIGURE 2.19: The rate of muons found by MU15 trigger as a function of pseudorapidity, η, at the end of Run 1. All, real and fake, MU15 candidates (yellow, upper histogram) and the real, matched to offline reconstructed muons, (cyan, lower histogram) are shown [36].

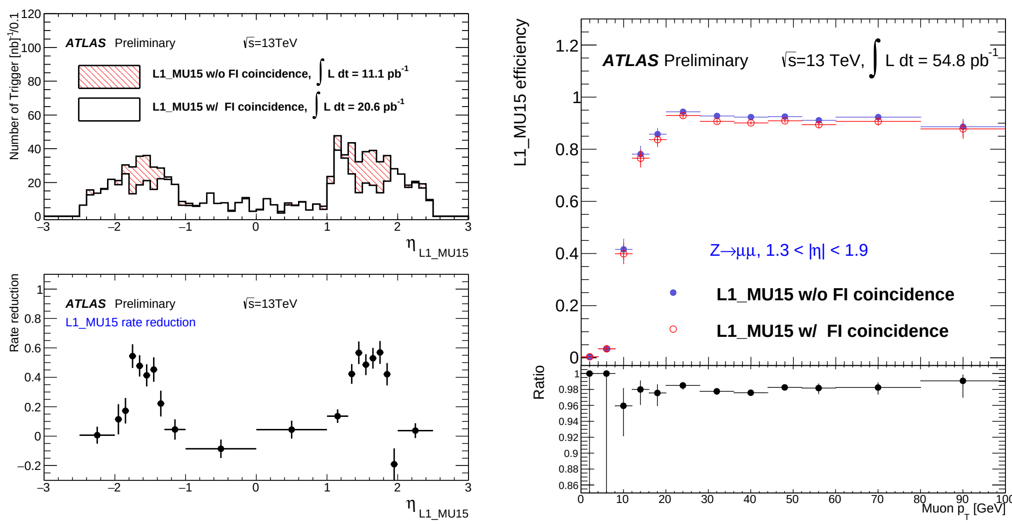

FIGURE 2.20: (left)Measurements of the rate with additional coincidence signals incorporated in the MU15 logic (w/ FI) compared to the rates measured in data taken without this coincidences during the Run 2 (w/o FI) at the beginning of Run 2. Reduction of the rates is shown at the bottom. (right) The efficiency of the MU15 trigger as function of transverse momenta, p_{T}, of the matched offline muons in the area of TGC with additional coincidences and without them. Ratio of the two is shown at the bottom. The efficiency was measured using tag-and-probe method using sample of events with $Z \to \mu\mu$ decay [37].

and 3. A detector system, called the New Small Wheel, is planned to be installed to cover the region $1.3 < |\eta| < 2.7$ extending the detector acceptance as well as adding discrimination power against the expected high cavern background of photons and neutrons [24]. The expected L1 muon rates reduction from about 60 kHz to about 30 kHz for ~ 20 GeV muons will allow the experiment to retain acceptance for EW bosons decaying to muons even at a luminosity as high as $\mathcal{L} \simeq 3 \times 10^{34} cm^{-2}\mathrm{s}^{-1}$.

2.7 Level 1 Central Trigger in Run 1

Signals from the L1 calorimeter, L1 muon system and several detectors in the forward region such as Minimum Bias Trigger Scintillators (MBTS), Zero Degree Calorimeter (ZDC), Beam Pickup (BPTX) are delivered to the L1CT [38]. For muons additional processing is performed by a Muon to CTP Interface (MUCTPI) which extracts the object multiplicities and removes duplicates in the overlapping regions found by L1 muon system. The 372 signals are then fed to the L1CTinput cards whose role is to time-align the inputs as well as to select 160 of them to be further passed to the core decision-making board. In the core L1CT card these signals are combined with information about bunch patterns in a large lookup table.

The bunch patterns are configured externally and define in which bunch crossings triggers are allowed to fire. The majority of items are designed to be active in the bunches where both beams are filled, but triggers designed to collect statistics for beam background studies are also gated with various unpaired and empty bunches. There are eight possible patterns which can be used. An example bunch pattern is shown in Fig. 2.21. Groups presented in this figure are designated to: 0 - Bunch Counter Reset Veto - synchronisation, 1 - primary triggering, 2 - calibration pulses, 3 - 8 - background studies with varying proximity to the colliding bunches. In this example filled bunches are structured in so called trains with spacing between each colliding bunch of 75 ns (two empty bunches in between i.e. the filled bunches are 7,10,13,16 ...). The configuration of bunches is established for every run by measuring signals in the BPTX detectors located at $z = \pm 175$ from the nominal interaction point.

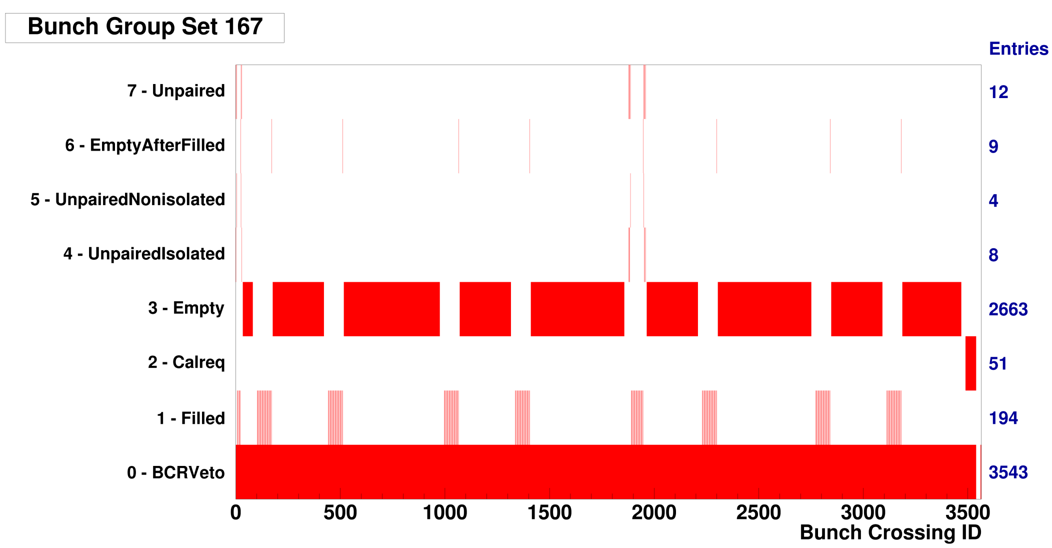

Figure 2.21: Example bunch pattern, identified as 167 in the database, configured during a stable run in March 23, 2011. The filled vertical lines mean that particular bunch crossing (horizontal axis) is considered to be in the group labeled on the left. On the right the number of bunches in each group is shown. For the description of the bunch groups see the text.

As a result of processing in the core L1CT module the 256 trigger items are produced. They are named according to the convention: L1_[multiplicity][_ trigger object and threshold ... bunch group . For instance L1_EM30 denotes requirement of presence of a L1 calorimetric deposit of energy above the 30 GeV threshold. A few examples of logic behind the trigger items are given in the table 2.1.

#	item	condition
1	L1_EM30	EM30 & BGRP0 & BGRP1
2	L1_EM5_2MU6	EM5 & 2 × MU6 & BGRP0 & BGRP1
3	L1_J10_UNPAIRED_ISO	J10 & BGRP0 & BGRP4
4	L1_J10_FIRSTEMPTY	J10 & BGRP0 & BGRP6
5	L1_J10_EMPTY	J10 & BGRP0 & BGRP3

Table 2.1: Examples of L1 items and the logic they implement. Items 1 and 2 are designed for physics, 3,4 to study backgrounds related to beams while 5 a beam unrelated one like detector noise.

Decisions for each of the item are then additionally masked according to the scalers in order to control the L1 output rates. An OR operation over the decisions after prescaling, when positive, is subjected to a dead-time algorithm. The purpose of the dead-time algorithm is to avoid bursts of trigger-accept signals which would cause instabilities in the detector readout. There are two components of that algorithm called simple and a complex dead-time. The former disallows a trigger in a given number of bunches after the event acceptance. The size of the inactivity window is a configurable parameter and is set typically to about 3-5 bunch crossings. The complex dead-time algorithm has the behaviour of a "leaky bucket". That is, when the "bucket" is full, accept decisions are disallowed and there is dead-time. However with the progression of the time the "bucket" gets emptied and new triggers are allowed. The complex dead-time algorithm is controlled by two parameters, the depth of the bucket and time it takes to leak one entry. For instance, setting the depth to 5 and the leak rate to 200 it means that it would take 1000 bunch crossings to empty a full bucket, and 200 bunch crossings before one new entry, the next trigger, is allowed.

The interplay of the dead-time algorithms with the bunch pattern results in different dead-times for various trigger items and makes the dead time also dependent of the bunch crossing as shown in Fig. 2.22. The simple dead-time algorithm does not influence the first two bunches therefore they have higher rate. The calorimetric trigger L2_TE180 is significantly affected by pile-up due to the long length of the calorimeter signal pulse shape and therefore the variation is more visible. For L1_MU10 the increase in rate comes from late arriving muons [39]. In addition to L1CT internal algorithms, a feedback "busy" signal from detectors and DAQ also contributes to the dead-time. An example of this contribution over time is shown in Fig. 2.23. The dead-time fraction is typically below 1%.

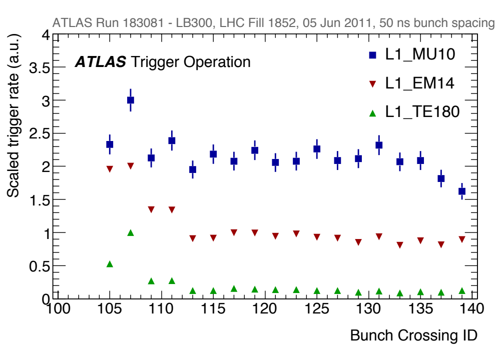

Figure 2.22: The influence of dead time on the trigger items rates as a function of the bunch in the train [39].

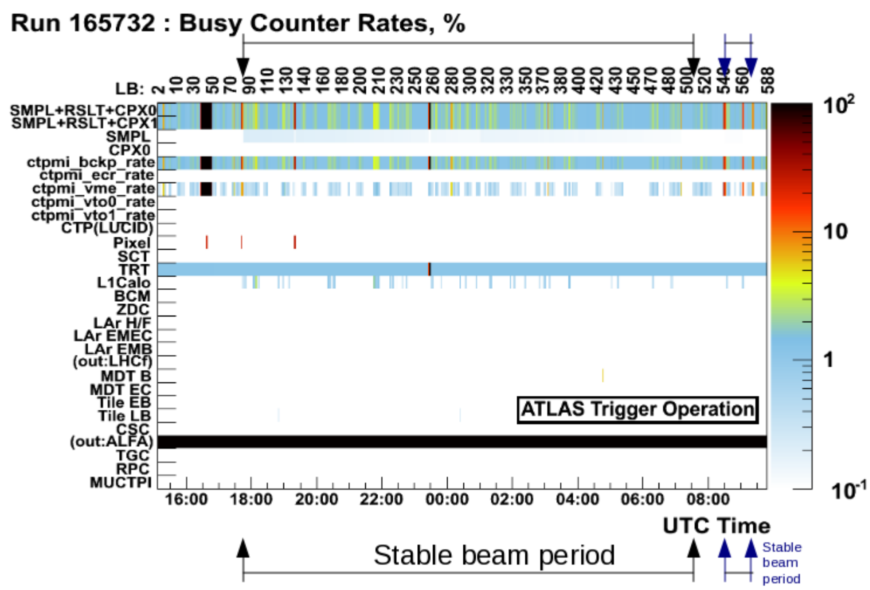

Figure 2.23: The dead time fraction caused by the feedback "busy" of various sources of CTP signals as a function of time [39]. On top the internal CTP sources are shown. The signal from ALFA was ignored as this detector was not taking part in that run.

In summary, the core L1CT module produces, for accepted events three sets of decisions, Trigger Before Prescale (TBP), Trigger After Prescale (TAP) and Trigger After Veto (TAV). Any positive decision in the TAV set causes the L1 accept signal (L1A) to be transfered to the detectors via the Timing, Trigger and Control (TTC) network and causes the movement of event data to the RoS. The L1CT decision bits and internal information are also transfered to a dedicated RoS in this case.

The L1CT has also hardware module providing precise monitoring of rates, dead time and software infrastructure for the collection of this information and the configuration of entire system.

2.7.1 Level 1 Central Trigger upgrade

The main motivation for the upgrade of the L1CT was a need to accomodate more trigger items needed at the higher energies and luminosities of Run 2. These major changes made the system more flexible and better monitored. The capacity of the input cards was increased and an additional 192 low-latency direct inputs are available in Run 2. They serve the L1 topological outputs and the ALFA detector input. The number of signals transfered to the core L1CT board is doubled and so is the number of trigger items ($256 \rightarrow 512$). The number of bunch groups which can be used to pick particular bunches is increased to 16 and there are four complex dead-time algorithms operating concurrently in order to better support the variance in the detector readout buffers. A new monitoring module allows the measurement of per-bunch-crossing dead-times for half of the triggers instead of only 12 in Run 1. The associated software suite has been upgraded to improve reliability of the system [40].

3

Requirements and design

In this chapter detailed requirements for the HLT are given. First the general requirement applicable to any trigger system and secondly more detailed requirements specific to various uses of the HLT in ATLAS are given. Later the design principles are presented which allow the set of requirements to be satisfied.

3.1 Physics requirements

The primary set of requirements is motivated by the physics goals of the ATLAS experiment. They are, in general, to collect a valuable data with high efficiency. The data has to be useable for the wide range of studies possible with a multipurpose detector. The scientific quality of the planed results depends on a near-perfect functioning of all the detector elements, including the HLT.

3.1.1 High Efficiency

The most important requirement of the trigger system is a very high efficiency. In addition, the trigger system, including HLT, has to allow for collection of auxiliary information for the efficiency to be measured on data. Efficiency uncertainties, including systematic ones, contribute to the overall uncertainty of the measurement and need to be minimised. In all these methods the existence of reference samples or auxiliary data is necessary. Their acquisition has to be planned prior to data taking and the trigger system has to be specially equipped for that purpose. A commonly used methods of estimating efficiency from the data are recalled.

Orthogonal trigger method

The assessment of trigger efficiency can be performed with samples of events in which the signal for which the trigger is designed is available. An important condition to meet, however, is that the sample cannot be biased towards the trigger under study. An extreme example of such a bias would be in study of the performance of the trigger aiming to select 20 GeV photons in a sample recorded with another trigger aiming at the selection of 25 GeV photons. The efficiency obtained using this sample will be unity above 25 GeV and there sill be no statistics below that threshold.

An ideal candidate for an orthogonal selection is a random trigger. However, due to a minimal statistical overlap of events taken at random and events containing signal of small cross-section the random orthogonal trigger is rarely useful and has to be replaced by one improving the overlap with the signal while avoiding a bias. A family of these triggers is collectively called minimum-bias triggers (min-bias) and they are superior to random triggers because events with no collisions are eliminated from the sample. An example of the evaluation the trigger efficiency with respect to the minimum bias triggers is shown in Ref. [41] and many early ATLAS publications.

The choice of orthogonal trigger is not limited to random triggers. Yet better option is a trigger enhancing a set of events but which is independent in terms of detectors, trigger hardware or software used. An example of a good orthogonal trigger could be a calorimetric missing transverse energy (signature of a neutrino) trigger selecting events with W boson decays which are used to study the efficiency of triggers aiming at the second decay product of the W, namely electrons, muons or τ leptons.

Bootstrap method

For the application of the bootstrap method a range of triggers has to be planned in advance. As an illustration, assume a set of triggers applying a range of cuts on the total transverse energy. The lowest threshold has to be set such that efficiency of this trigger can be measured using the statistics available in the min-bias sample. The efficiency of higher threshold(s) is then established by the fraction of coincidences with lower threshold triggers. One obtains a set of relative efficiencies which are multiplied in order to obtain the efficiency of the trigger in question. In practice, the intermediate triggers need to be prescaled, which significantly reduces the statistical precision of the efficiency curve estimation. In order to avoid such an effect the "decision to be" has to be known for intermediate triggers [42]. The HLT system has to allow for a collection of this information.

Tag and probe

In the cases where the signal appears twice in the event (e.g. $Z \to ee$ or $Z \to \mu\mu$) the efficiency of the single-object trigger can be found using a tag-and-probe method [43]. In this case the procedure is as follows: from all events which were recorded by this single-object trigger, those containing the Z bossons (or J/ψ) are selected. In these events the trigger decision for the second object (probe object) is checked and the efficiency is estimated as

$$\epsilon_{t\&p} = \frac{n_{\text{probe+trigger}}}{n_{\text{probe}}},$$

where n_{probe} is the number of objects selected and $n_{\text{probe+trigger}}$ is number of selected objects passing the trigger in addition. The trigger pass/reject information for all objects in the event has to be present for that method to be usable.

3.1.2 Rejection power

The trigger system, and the HLT in particular should provide high rejection power. The rejection power is insufficient or signal rate is to high an additional rate control mechanism should be provided.

3.1.3 High availability

In order to achieve high efficiency the fraction of the time the trigger and DAQ systems are inactivated after an event is inspected has to be minimised. The design allows for small fraction (percent level) of collisions to be lost due to the detectors inability to cope with data readout from adjacent crossings. Besides that the design of the HLT and the overall system working point should minimise deadtime. In case that there is a problem with online data processing a procedure is required to recover the events for analysis.

3.1.4 Multichannel selection

Even though, for first period of exploitation of the LHC data, the primary goal was the discovery of the Higgs bosson, the number of studied subjects covers nearly all areas of modern high energy particle and nuclear physics. In order to allow this program to be effective "multichannel" selection has to be implemented. The number of independent signatures which the trigger system is able to process is defined to be 256 at L1 and 8012 at the HLT. For L1 the number was constrained during LHC Run 1 and as increased by factor of two for Run 2 (see sec. 2.7). Limitation of the HLT system were never approached. All signatures are required to function independently of each other unless specifically configured otherwise. Adhering to this requirement lead to the design of the trigger system where each of the signatures is configured independently. Evaluation of the signatures is also guaranteed independent, while the resources to perform that evaluation are shared.

3.1.5 Flexible construction of complex signatures

The main physics goals of the experiment drive the design of the detectors and systems. For instance, a strength of the ATLAS experiment in the search for Higgs decay products is very good electron and photon identification. Consequently the trigger system has to reach the utmost performance to support such a channels. However, at the same time it was required that the trigger system shall allow for flexible construction of complex signatures. Amongst the main requirements were multi-object triggers, for instance the di-electron trigger. Also needed were triggers in which information derived from combined objects are used for selection, i.e. di-muon with invariant mass in the mass window of the J/Ψ meson.

3.1.6 Decision introspection

As has been discussed, data analysis incorporates, among other aspects, corrections for trigger inefficiency. Therefore a record of the decision taking process is made available to the analysts. In the record of the complex decision taking process, both the decisions and the objects description (physics objects such as electrons) which were reconstructed must be present. The trigger decision may need to be inspected for other reasons. One of them is dividing sets of events into categories, so called streams. That process happened online during the Run 1 but may (as is planed for Run 2) also be used to perform fast event categorisation in offline processing.

For the trigger configuration system, an important requirement is that the knowledge of configuration of the system used for data taking is preserved. This is to ensure that trigger efficiencies are reproducible in MC simulation. Key aspects of the configuration need to be made available during data analysis (essentially those properties which allow to the efficiency of the trigger to be determined).

3.1.7 Single event reproducibility

The selection process has to be deterministic. The reproducibility of the algorithmic code is a guarantee of its correctness but there are other more important and very positive practical implications of reproducibility. Tuning of the system parameters involves an iterative process which converges more easily fluctuations stemming from irreproducibility are absent. In the HLT system there is a need to pre-scale high rate signatures and it is desirable to do it in a way that assures randomness. The random number generator (RNG) used for pre-scaling needs to be in a predetermined state for each event that is independent of the number of events processed before or the order of the processing.

Single event reproducibility excludes an important class of algorithms from use. Namely all algorithms which would rely on local machine learning during data processing.

3.2 Operational requirements

3.2.1 Processing time requirements

The interesting collision events occur at a certain rate defined by the cross section and luminosity. The iime span between two consecutive events is given by a Poisson distribution. Therefore, it is possible that the interesting signatures appear in adjacent bunch crossings and the readout and trigger system have to reliably select such events. The method, commonly used in particle detectors is to use derandomizing buffers rather than hold the readout until the event is processed. With the derandomizing infrastructure available the temporal structure of the processing plays no role and the required event processing time can be calculated from the average rate. For instance for a distributed system to process an average input rate R using N processors results in requirement for the average time to process the event of

$$T_{\text{avg}} = N/R.$$

This simple relation can also be interpreted in the spirit of Gustafson law for parallel computation for the so called embarrassingly parallel problem [44] when the rate, R, is understood as the size of the problem. An increase of the number in processing units allows higher rate to be processed without reducing the time T to process an event. Conversely, with a predetermined number of processors, N, the demand to process higher rates can be realised only if the average processing time is reduced. In summary, the increased system throughput (rate), which is of primary interest form the physics perspective, can be obtained by shortening processing time or increase of the number of processing units.

A derandomizing infrastructure of any sort usually has limited storage capacity. That capacity defines the maximum time of processing, T_{max} which is allowed without overflowing buffers. The larger the capacity of the readout system the bigger T_{max} that can be allocated to the decision taking step. Here, however the relation is not as clear as for averages. It depends, for instance, on the method of addressing data in the derandomizers and other logistical constraints. In the case of ATLAS the constraint on the maximal time allowed for the HLT processing is dictated by the logistics of event data handling in the data flow system.

An interesting approach for accelerator experiments would be to significantly extend the buffering capability to allow processing during periods of accelerator inactivity, thus allowing for much bigger processing times and as a result better filtering. An implementation of this strategy is the trigger and DAQ system of the LHCb experiment where the event rate is only slightly reduced by a prompt pre-selection during collisions and such a pre-filtered data are stored on discs to be further processed at a latter time, when collisions are not taking place [45].

3.2.2 Handling excessive processing time

In order to assure system stability timeouts handling is implemented in the HLT. For events in which the HLT is not able to work out the decision within $T_{\max}$ the process should be interrupted and as much of the event information as is available needs to be saved. Events of this type need to be directed outside of the main data stream and the decision taking proces has to be repeated offline without constraining the processing time. If the case when gentle timeout handling in HLT process fails, the data flow system is allowed to force accept the event without the trigger decision. The trigger decision taking process needs to be repeated offline as in the former case.

3.2.3 Online monitoring

The efficiency of complex system like the ATLAS experiment's trigger, which is exposed to variable conditions (physics and detector), can only be assured if the health of the system is constantly monitored. For the trigger the primary monitored quantities are the rates at various stages of the selection process. This allows pathologies to be spot, like excessive or zero rate, and to either disable malfunctioning signatures or tune the rates so that the stability of data collection is not endangered. Also the variables used in the selection are monitored. It should be nothed that the monitored information concerns mostly the events which are rejected and this invaluable information can not be recreated from the recorded data.

3.2.4 Input and output data bandwidth

In order to process the event in the HLT to decide if the event shall be accepted for recording or rejected, the information has to be collected from the readout system of the detectors. An important contraint arising from the choice of working point is that not all the data can be requested in each processed event. This is the result of a balance between the rate of events which need to be inspected to allow recording of interesting physics (highest possible) and the construction costs of the readout system (lowest possible). In case of the ATLAS experiment the design HLT input rate is 100 kHz while the readout system allows access to the full event information for about 1/5[1] of this rate.

A possible solution was identified to reduce the amount of detector data needed for the trigger decision, especially to reject the events as they constitute majority of the input rate. For accepted events a higher expense in processing time and data bandwidth can be afforded in order to benefit from an improved data usability in the analyses (no early accept i.e. the full HLT processing is run to completion rather than halted on the first trigger accepted).

Limitations on the output rate arise from the offline data storage and handling costs. Therefore the HLT process has to have the flexibility to use available offline output bandwidth optimally in changing experimental conditions.

[1] For LHC run 2 the capacity of the readout system is increased to about 1/3 of the HLT input rate.

3.2.5 Deployment environments

In the design the following possible use cases of the HLT needed to be considered as they require the introduction of a certain degrees of flexibility within the framework.

1. **Online selection of collision events.** This is the primary environment and most of the operational requirements (i.e. timing, reliability) are dictated by the online data taking conditions. Also the requirements of the ATLAS physics programme have to be implemented in this environment.

2. **Part of MC production.** The HLT process is used evaluate the trigger decisions in MC simulations (however not to decide on rejection of the event). The trigger result contained in the MC data allows for a direct performance studies of trigger efficiency in signal samples. An important difference, compared to the online environment is the fact that the resulting decision is made available for each processed event.

3. **Recovery.** In the rare occasions when the online system running constraints disallow the decision to be taken online or in the case of malfunction the events are directed to, so called, debug streams. For these events the decision has to be reevaluated. Normally unlimited time is allotted to process events from debug streams in order to evaluate the decision as if it would be taken online.

4. **Reprocessing.** The deployment of new trigger algorithms and configurations is certified by re-evaluation of the entire trigger decision on test data samples.

5. **Development.** The development and refinements of the HLT process is made as developer friendly as possible. The Athena (ATLAS Offline Framework) [46], is used as the basis of the HLT process and where possible the infrastructure (i.e. part of configuration system) is reused.

3.3 Key design principles

The above described requirements constitute the corner stones of the design of ATLAS trigger system and the HLT in particular. They are listed below without the technical details which are deferred to chapter 4.

In order to be able to inspect the event data at a rate higher than the average readout rate two key ideas have been implemented. In addition they allow a reduction of the CPU resources needed for event processing. That aspect is not emphasised in the following discussion except when specific CPU savings are independent of the detector data bandwidth.

These key design principles are implemented in the HLT framework in a way that is transparent to the HLT algorithms.

3.3.1 Regions of interest and abstracted data access

The L1 trigger system analyses signals from entire detector. Simplified algorithms are employed, that are readily implementable in the hardware and the data used as an input are compressed. These are algorithms such as sliding window over trigger towers (groups of cells) for calorimeter cluster finding or lookup table matching in the case of muons. As a result the geometrical positions of the reconstructed objects are known together with some coarse characteristics of them. This positional information takes the for of Region of Interest (RoI) that is fundamental to reconstruction at the HLT. Only data from detector regions pertaining to the RoI(s) are read out during the HLT process. With an appropriate mapping of the readout channels to the RoS layer (cabling) only a fraction of the RoS infrastructure is needed to provide the data for a particular event. Notice that other cabling schemes would be possible i.e. in order to increase the fault tolerance of the RoS adjacent detector elements could be read by independent elements. In such an arrangement reading RoI data would involve many elements of RoS and readout rate would not be significantly reduced by the RoI approach. Therefore the mapping of the detector readout elements to the RoS elements is instrumental to a successful implementation of RoI concept.

The detector data access in the algorithms has to be abstracted so that an algorithm in test environment (when the data is readily available in the memory) and in the deployment environment (when the data are stored in RoS system buffers) can access the data without modification.

3.3.2 Early rejection

A further reduction of the data readout rate is achieved by the application of partial selection requirements as early in the reconstruction process as possible. In the HLT the reconstruction process is divided to a sequence of small algorithms. The output of the reconstruction algorithms are then input to the selection hypotheses algorithms and failure to pass the selection requirements results in abandoning further reconstruction for that RoI. This procedure is applied to each RoI. In addition, the most time consuming and data hungry reconstruction algorithms are deferred to the later stages. In this way the amount of the data which be needs to be delivered to reconstruction is significantly reduced.

3.3.3 Step-wise validation of decisions

In multi-object signatures further optimisation is achieved by grouping the reconstruction steps and requiring coherent validation of the hypotheses. For instance, in the di-muon trigger an early application of a hypotheses for each individual muon RoI may result in only a single muon remaining valid in the event. In such a case reconstruction is abandoned and detector readout bandwidth is reduced.

3.3.4 Avoidance of repeated reconstruction

The requirement of independent selection chains implies the presence of many isolated sequences of algorithms. On the other hand, reconstruction algorithms are often common to several sequences. It is therefore desirable that duplicated reconstruction of the same detector regions is avoided. This is achieved by caching the reconstruction results and reusing them for identical sequences. This optimisation specifically allows for a reduction of CPU resources. It does not reduce data requests as the detector data in repeated reconstruction would not need to be fetched thanks to an additional buffering layer.

4

HLT process implementation

4.1 Components of HLT process

The events accepted by the L1 are taken over by the data flow infrastructure and assigned to a machine running the Level-2 trigger (L2) process. The RoI information and trigger decisions are unpacked and used to prepare the seeding information for the *HLT algorithms*. The goal of the processing at the HLT is the confirmation of the coarse hypotheses from the L1 hardware. The overall decision evaluation process is therefore decomposed into the evaluation of partial decisions in *HLT trigger chains*. Once all evaluations are finished the event is accepted if the decision of at least one HLT chain is positive. The design principles described in section 3.3 are respected during the decision evaluation process. For the accepted events the decision summary is constructed and the event is handed by the Data Flow system (DF) infrastructure to the next level of the HLT, the EF, where decisions are further refined with more precise algorithms. Upon a positive evaluation of at least one HLT trigger chains at the EF the event is recorded for use in the analyses. A negative decision at L2 or EF results in the rejection of the event. In order to provide the data for a wide spectra of physics analyses performed with the ATLAS detector, about 1000 trigger chains are used in the selection process.

The complex task of decision-taking is decomposed into the components that are enumerated below. The interplay between them in the decision taking process will be clarified in the subsequent text.

1. *HLT steering* - is the top level component which orchestrates the decision taking process. More elementary tasks are delegated to more specific components.

2. *L1 and L2 converters* - these are the components responsible for initialisation of the HLT decision taking process. The task of the L1 converter is the unpacking of information from the hardware, namely the CTP decisions and RoI information. The L2 converter

serves the same purpose in the EF but in this case the information from L2 is used as the seed.

3. *HLT Feature Extraction Algorithms (FEX)* - the algorithms responsible for reconstruction of physical objects (tracks, energy clusters, etc.) required to confirm hypothesis. The reconstruction is typically, though not exclusively, performed in single RoI.

4. *HLT Hypothesis Algorithms (Hypo)* - the algorithms implementing the task of hypothesis verification. A trivial example if the implementation is the requirement of a value of object property (e.g. electron E_{T}) above a threshold.

5. *HLT sequence* - is a list of HLT algorithms with the specification of the input and the output. A typical sequence consists of set of FEX algorithms delivering an input to Hypo algorithm.

6. *HLT trigger chain* - the decision-taking unit. The task of the chain is decomposed into the steps. Rejection by the hypotheses, tested at each step, results in the abandonment of the evaluation of later sequences thus is an implementation of the early rejection principle.

7. *HLT navigation* - a is structure allowing a flexible interconnection between the algorithms by providing the context of the execution.

8. *Result Builder* - is a component responsible for packaging of the decisions and intermediate data into the decision record. An additional task of the Result Builder is the formation of the event streaming directives which are used by DF to assign the event to an output file.

9. *HLT monitoring* - is a set of components responsible for the acquisition of statistical information (histograms) concerning the decision taking process. The most important for online operation are the rate and error counters [2]. Also, numerous physical quantities used in the selection process are collected and made available in online system for inspection by experts or expert systems to assure healthy functioning of the HLT. In addition, operational information, such as the amount of data used by the HLT or algorithm execution timing are monitored.

10. *Trigger Cost* - is a set of components responsible for the collection of very detailed logging information about the HLT execution. Postprocessing of this data allows the emulation of the behaviour of a new configuration in terms of the rates for particular trigge as well as the overall rates (accounting for overlaps). Additionally the operational performance can be assessed with the help of the information collected by these components [47].

11. *Trigger Configuration Services* - is set of tools providing the configuration of the trigger chains, and sequences to the HLT steering [48]. In the online environment the Trigger Configuration Services are also responsible for configuring the algorithms and other components of the system. In offline conditions the latter is provided by infrastructure of Athena framework.

4.2 Decision evaluation

4.2.1 Initialisation of event processing

RoI information at L2

The RoI information is used to guide regional reconstruction. The description of the reconstruction region is a set of parameters for the location in the detector, η, ϕ, size $\Delta\eta$, $\Delta\phi$ and a position and span along the beam direction z, Δz. The location parameters are contained in the L1 data while the predefined sizes are associated to RoI according to the type of the region, taking into account geometrical properties typical of the physics process. For instance for jet RoIs the size is $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$ while for electrons only 0.1×0.1. The RoI information is updated to be more precise during processing in the HLT, where also the z and Δz are used to limit reconstruction only to interesting regions. Unlimited number of RoIs is allowed by the HLT system. In high density events the RoI may overlap. In the initialisation stage there is no attempt to mitigate this problem by merging or elimination. Global event information such as missing or total transverse energy is also represented by a variant of the RoI object. The RoI information is associated to the HLT navigation structure which is described in one of the next sections.

HLT chains initialisation at L2

The other part of information used at the start of processing in the HLT are the L1CT bits. They are used to decide for which of the HLT chains the decision needs to be evaluated in that event. The prescaling procedure (unconditional rejection of a fraction of decisions) is then applied in order to further reduce the number of chains for which the decision and thus processing is needed. The level of rejection is determined by one floating point value, a prescale factor, or prescale for short. The value of the parameter p signifies that $1 - p$ out of p decisions are discarded. Each HLT chain has an independent [1] prescale factor. The set of prescales constitue important meta information about the data taking conditions and will be also described in more details later in the text.

In most of the cases the prescaling procedure is applied at L1 as it allows biggest savings of resources; that is, the events rejected at L1 are then not even read out from the detectors. Prescaling at the HLT is necessary when several HLT chains are based on the same L1 item. In this case the prescale factor used at L1 can be as high as the most rejecting HLT chain requires. In order to control the output the least rejecting HLT chains need to be additionally prescaled at the HLT. A RNG is used to decide about the prescaling decision to assure independence of samples collected by trigger chains. However, in order to assure the reproducibility of this process the RNG is restarted for each event with the semi-random seed present in the event data. Such an implementation allows the requirement of the reproducibility of the HLT decision to be respected as the RNG state is independent of number and properties of previously processed events (i.e. the prescaling decisions for instance).

If the efficiency of the prescaled chains is to be measured using the bootstrap method then, in order to improve statistical significance of such a measurement, coherent prescaling may be applied. For example if the chain of interest has prescale factor of 100 and its bootstrapping chain prescale factor is 10^3 then only in one event out of 10^5 are both decisions known and are useful in the measurement. If the decision for the chain of interest is evaluated for every event for which the bootstrapping chain caused the event to be recorded, the number of usable events

[1] See the paragraph about coherent prescales.

is increased to one out of 10^3. It is a significant improvement in measurement quality while at the same time the resources involved (number of recorded events and CPU cycles for decision taking) are actually reduced.

For coherent prescaling the chains need to be assembled into a coherence group and the prescales are applied as follows. The chains in the group are ordered according to the value of prescale factor. The prescaling decision for the least prescaled (lowest prescale factor value) is evaluated accordingly to the prescale factor, p_1, for that chain. When the chain is rejected from further consideration in the decision taking process the remaining chains are also rejected. When the chain is designated to be further processed the cumulative prescale factor ($p_1 \times p_2$) is applied to the second chain and the procedure repeated recursively for all chains in the group. The decision taking process is then performed for all chains independently. As a result for every event where the chain with the highest prescale resulted in recording of the event, the decision for the chain with the lower prescale is also known.

The coherent prescaling method can be applied only to a group of chains seeded by the same L1 item. In cases where the bootstrap method is intended to be applied to the HLT chains seeded by independent L1 items re-running of the HLT chains for accepted events can be selectively applied. The re-running of selected HLT is initiated in all events for which the seeding L1 passed the event before the prescaling step.

For detailed trigger studies it is often desirable to record some fraction of events irrespective of the selection decision. A way to achieve that is by adding a parallel, never-rejecting chain with the desired prescale and group it in a coherent prescaling set with the chain to be monitored. However, in order to reduce the number of chains to maintain it the configuration an additional scaler, a so called pass-through prescale, is configurable for a chain. A fraction of events are recorded according to this scaler, irrespective of the actual trigger chain decision.

Auxiliary chains, used to collect monitoring data, can be seeded by the presence of any (or any from a predefined set) positive L1 decision.

Sanity checks

Additional checks are performed on the information from L1. The L1CT decision is reverse engineered to determine passing items and a check is performed to see appropriate number of RoIs are in fact present in the events. As an example, if a L1 item requires the presence of two electron objects above some threshold and a muon of transverse moment above a certain threshold then the presence of these objects is checked in the part of the L1 decision encoding RoI information. Such inconsistencies may appear when the L1CT and RoI information from distinct events are erroneously combined. Operational experience has shown that this is a very rare condition but early flagging of it at the HLT allows corrective measures to be applied quickly.

Another check performed on L1 data at the start of the HLT concerns the presence of overflow bits. In very busy events the number of RoIs may exceed the capabilities of the electronics and data links. This case is signalled by the presence of flags in the L1 information. In response to these condition the HLT process provides monitoring information (frequency histograms) to the DAQ monitoring system. Additional actions can be configured such as to abandon processing of the events to different stream. Operationally these conditions are rare and redirection to a special stream feature is not routinely used.

Initialisation in the Event Filter

The initialisation procedure at the EF ressembles that of the L2 with the following difference. The L1 RoI information is replaced by more precise refined L2 information and instead of L1 items the results of the chains from L2 are used to seed the chains at the EF.

4.2.2 Validation of chains decision

The HLT selection is configured in terms of chains. The chain consists of a set of requirements that are imposed on the event in order to pass the selection. These requirements are static, this is the definition stays unchanged during data taking. Even though it is possible to modify the definitions from one data taking session/run to another it is preferred, for the sake of simplicity of bookkeeping in the analyses, not to do it. In practice, families of chains with selection criteria optimised for different luminosity conditions are present in the configuration and are activated when the conditions change (e.g. luminosity falling during the data taking run).

For optimisation purposes the selection within the HLT chains is divided into steps. At each step part of the chain definition is evaluated and a negative decision results in the abandonment of further processing of that chain, thus saving of resources.

A subset of all configured chains is activated for a given events. The rules determining this activation were described in section 4.2.1. Evaluation of the chain decisions is performed in a step wise manner, that is step one is evaluated first for all chains. If the features reconstructed in the event do not meet the requirements of the chain, that chain is deactivated and it is not considered in the next step. If that happens for all activated chains then the event is early-rejected.

The definition of the chain step is expressed in an abstract manner through the requirement for the presence of a Trigger Element (TE) together with required multiplicity. Each TE corresponds to the result of sequence of algorithms as will be explained in details further in the text. For readability TE is labeled with an acronym describing briefly that sequence. Tables 4.1 and 4.2 present the definitions of two HLT chains, a relatively simple muon selection and more complex one involving multiple TEs at each step. A brief explanation is provided for each TE. Thanks to this abstract approach the evaluation of the chain decision is reduced to counting of the TEs. The entire selection complexity is hidden within the HLT algorithms that implement reconstruction and hypothesis verification.

If the decision evaluated in any chain is positive[2] then the event is accepted for further processing, in case of L2 and permanent storage in case of the EF. Based on the decisions of chains that have passed additional directives concerning the destination of the event are evaluated. This is the name of the stream to which the event should be written out and, if only a fragment of the event is interesting, the identifiers of the detector readout fragments are made known to the DF system.

For accepted event the second round of the decision evaluation is performed for chains which are set to be re-run. The feature objects and decisions are saved in the event record. Similarly as for the pass-through feature the re-running once enabled for a chain at L2 it is enabled at EF as well. The decisions worked out in the second reconstruction round do not influence the streaming information concerning the event.

4.2.3 HLT algorithms and sequences of thereof

The reconstruction and hypothesis validation task is outsourced to the external HLT components, the HLT algorithms. From the decision taking point of view the algorithms are never used alone but are grouped into sequences. A typical sequence consists of a FEX and Hypo algorithm however sequences with tens of algorithms are also possible and are present in the

[2] Including positive decions due to the pass-through mechanism.

Step	TE	meaning
L2		
1	L2_muon_standalone _mu6	At least one muon was found in the MS with $p_{\mathrm{T}} > 6$ GeV
2	L2_mu6	A muon track was found in the ID and could be combined with track from the MS. The better estimated p_{T} is above 6 GeV.
EF		
1	EFID_mu6	ID tracks in the region pointed to by L2 are reconstructed with improved algorithms.
2	EF_mu6	A muon reconstructed in the MS and combined with ID track satisfies the hypothesis of muon of $p_{\mathrm{T}} > 6$ GeV.
3	EF_muI_efid_mu6	Charge particle tracks in a wider RoI are around the track identified as muon are present.
4	EF_mu6i	A muon in isolated cone is found.

Table 4.1: Definitions of relatively simple L2 and subsequent EF chains designed to select events containing at least one muon of $p_{\mathrm{T}} > 6$ GeV and without additional charged particles in its vicinity.

configuration. The input and output to the HLT sequences are sets of TEs. In the output only one type (label/name) of TEs can be specified, that is, the sequence provides only one unit of reconstruction and hypothesis validation. In the input several sets of TEs are allowed and, depending of the algorithms in the sequence, various combinations of the information delivered via the TEs can be created. Most typically though, the HLT sequence is defined by only one set of TEs at the input and output (e.g. EFID_mu6 and EF_mu6 respectively i.e. the consecutive requirements for chain EF_mu6i in table 4.1). Through the specification of the inputs and outputs the sequences form a complete reconstruction and verification stream for each type of object in the event. A typical configuration consists of about one thousand sequences. Some sequences, used by the chains shown before as examples, are presented in the table 4.3.

Step	TE	meaning
L2		
1	L2_muon_standalone_mu6 × 2 & L2_e7T_ medium1cl	At least two muons were found in the MS with $p_T > 6$ GeV and an electromagnetic cluster with a transverse energy above 7 GeV. A medium cluster quality is required.
2	L2_muon_standalone_ wOvlpRm_mu6 × 2 & L2_e7T_medium1id	Muons are required not to be close by and an additional requirement for the charge of muons to be the same is applied. A track is present in the region where the cluster is found.
3	L2_mu6_wOvlpRm × 2 & L2_e7T_medium1	An ID track compatible with muons is present and matches muons from the MS. A track found the vicinity of the calorimetric cluster, when combined with the cluster information, is compatible with an electron.
4	L2_muon_combined_ wOvlpRm_mu6 × 2 & L2_e7T_medium1	The overlap of the muon track with the track identified in ID is required. The electron requirement is confirmed.

Table 4.2: Definition of a complex L2 HLT chain designed to select events containing a pair of non-overlapping muons and an electron. For brevity the EF part is skipped as the implementation is similar but with improved algorithms.

Input TEs	Algorithms	Output TE
MU6 (from L1)	MS Fast muon reconstruction, muon hypothesis $p_T > 6$ GeV	L2_muon_ standalone_mu6
L2_muon_ standalone_mu6	ID track finding in silicon detectors, MS and ID combiner, Hypothesis verification on combined muon track	L2_mu6
L2_mu6	Pixel clustering, SCT clustering, TRT drift-circles making, silicon hits finding, silicon tracking, tracks ambiguity resolving, extension to the TRT, extension processing (confirmation), tracks slimmer, vertex finding	EFID_mu6
L2_muon_ standalone_mu6	Overlap removal algorithm	L2_muon_ standalone_ wOvlpRm_mu6

Table 4.3: Example sequences used by the HLT chains shown in tables 4.1 (first three) and 4.2 (all of them and the first three are also used by this chain). Instead of the algorithm names brief description of their function is presented.

In the event processing the number of TEs in the input depends on the presence of features in the event. In particular it can be zero when objects are absent or when the objects found do not satisfy the selection criteria. The starting set of TEs used by the initial HLT sequences is prepared in per-event initialisation from the RoIs found by L1 (cf. 4.2.1). A separate TE is made for each RoI and another sets of TEs are made, one for each threshold passed at L1 by this object as shown in Fig. 4.1. These TEs are then picked up by the HLT sequences.

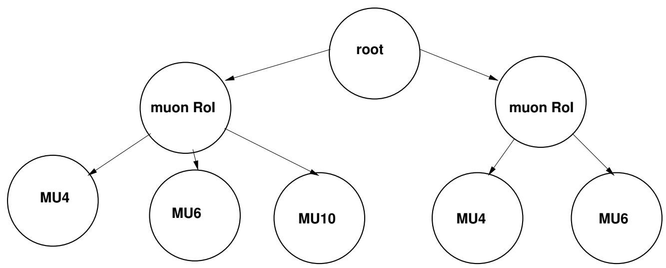

Figure 4.1: A sketch of the structure of TEs (circles) with the seeding relation (arrows) after the initialisation of event processing. For each threshold passed a separate TE is created.

The invocation of the sequences is driven by HLT chains. Validating the partial chain decision at a given step is in fact a verification of the presence of TEs. Since the TE uniquely corresponds to the HLT sequence the algorithms in that sequence are invoked. If the input TEs to that sequence have already been created, then the algorithms in that sequence retrieve data using the input TEs and in the end produce the output TEs. If the input TEs are missing the appropriate sequences are invoked recursively. A sequence is executed only once per event and the resulting set of TEs is cached. With this one invocation, reconstruction and hypothesis validation is performed on all inputs. For example the aforementioned sequence: EFID_mu6 $\rightarrow$ EF_mu6, when invoked, performs reconstruction of all muons in the event which passed the selection so far.

The strategy with which the input TEs are used to make the output TEs depends on the first algorithm in the sequence. This, precursor algorithm is given the freedom to merge, filter, combine and above all to create output TEs. Depending of their functionality the HLT FEX algorithms are categorised as follows.

1. Single RoI FEX algorithm which for every input TE produces an output TE. This algorithm is equipped with a filtering possibility to prevent creation of the output TE for some inputs (rarely used).

2. Combining FEX algorithm which needs to be seeded by at least two sets of TEs and creates an output TE for each input combination. This algorithm is also quipped with filtering functionality.

3. Full-scan FEX algorithm (or all TE algorithm) which obtains complete input sets and performs arbitrary combination of them. Typical cases are:

    - Information extraction from the joint input. An example is the algorithm calculating missing transverse including information from all muons present in the event or verification of combined information i.e. the last example in table 4.3.
    - Creation of the seeds. For instance an algorithm creating an RoI object encompassing the entire detector which is then used by a single RoI tracking algorithm.

After the execution of the precursor algorithm the output TE set is frozen and the remaining algorithms operate on it. Therefore their only function is to add more information to the output TE.

The last[3] algorithm in a sequence is typically a Hypo algorithm. It always functions only on a single TE and its ultimate goal is to asses the validity of the hypothesis for the objects associated to this TE. The TEs (RoIs) for which the hypothesis is invalid are set to be inactive and they are excluded from further inspection thus reducing the number of regions for which the reconstruction is performed.

The result of the execution of the sequence is returned to the HLT chain where the partial decision for this chain is evaluated by counting only active TEs.

4.2.4 Information exchange between HLT algorithms

The HLT algorithms perform progressive reconstruction of the event properties. Since most of the reconstruction happens within RoIs the data exchange infrastructure has to assure that each algorithm can quickly obtain access to the data in a RoI which was reconstructed by preceding algorithms. For the FEX algorithms this input data serves as the seed for next reconstruction step while for the Hypo algorithms the input data is used to apply the selection algorithms.

A structure used for the intercommunication is formed as a directed acyclic graph resembling a tree. The nodes of the graph are TEs while edges are the connections corresponding to the reconstruction direction (seeding). The node representing the event is made the root. Next two layers are created by the L1 conversion process that represent RoIs and the L1 decisions for each RoI. A linear continuation of these initial branches is typical for inclusive triggering. Referring to the earlier example, two TEs, EFID_mu6 and EF_mu6, are the nodes and the reconstruction/seeding relation between them points from first to the second TE. The topological requirements result in joins between the nodes. New TE can also be created outside of the main tree and are used to represent global event quantities such as the primary vertex or missing transverse energy information.

As stated before the TE is an abstract object and does not contain any attributes describing physical quantities. However each TE is capable of storing a link to any arbitrary object (e.g. tracks, clusters) which appropriate attributes. A general term used for these objects is "features" and hence the name FEX for the class of algorithms in which they are created. The directed graph of TEs and the links to the auxiliary data objects are called the HLT navigation. A sketch of the navigation is show in Fig. 4.2 while a visualisation of the actual navigation structure within an event is shown in Fig. 4.3. Only the structure resulting from the processing at L2 is shown. The entire structure for the event is roughly doubled in size once the event is processed at the EF level.

The navigation structure for events accepted at L2 is streamed to a flat bytes array and is transported to the EF where the di-graph structure is recreated and the reconstruction starts from the final nodes of L2. For accepted events the entire structure (containing the TEs originating from the processing at L2 and the EF) is saved in the event record. This structure is then used in analysis in order to discover which objects reconstructed during triggering contributed to the event passing the selection and hence this structure is instrumental in efficiency measurements and so crucial to the analyses.

[3]When placed first in the sequence it acts as if it were preceded by a single RoI FEX which does nothing but create one output per input.

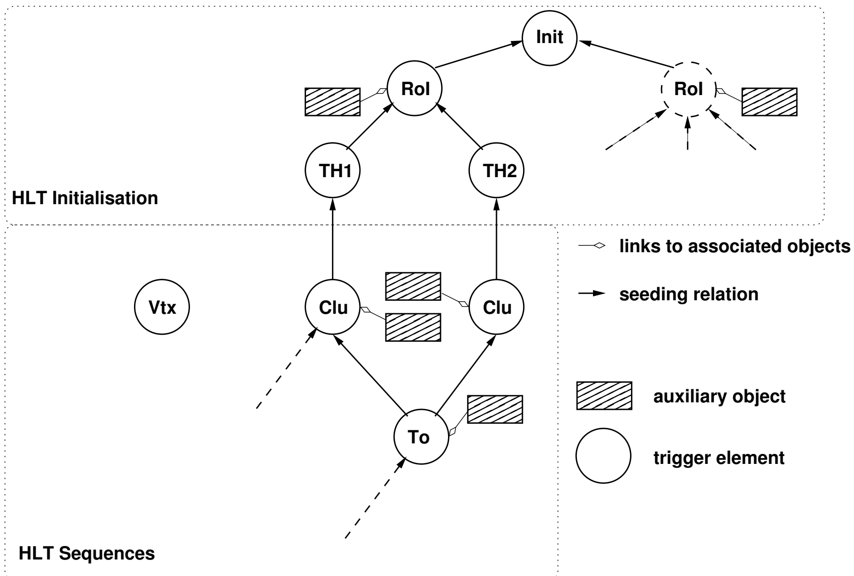

Figure 4.2: A sketch of the navigation structure illustrating the relations between the TEs as well as between the TE and auxiliary objects. The L1 RoIs encoded in the structure are shown at the top together with the event root node while the remaning fragment is created by the HLT algorithms operating in sequences. The standalone element "Vtx" shows how the unseeded reconstruction are represented in the structure. The element "To" shows a topological connection between the objects. Dashed lines are meant to illustrate that the structure extends.

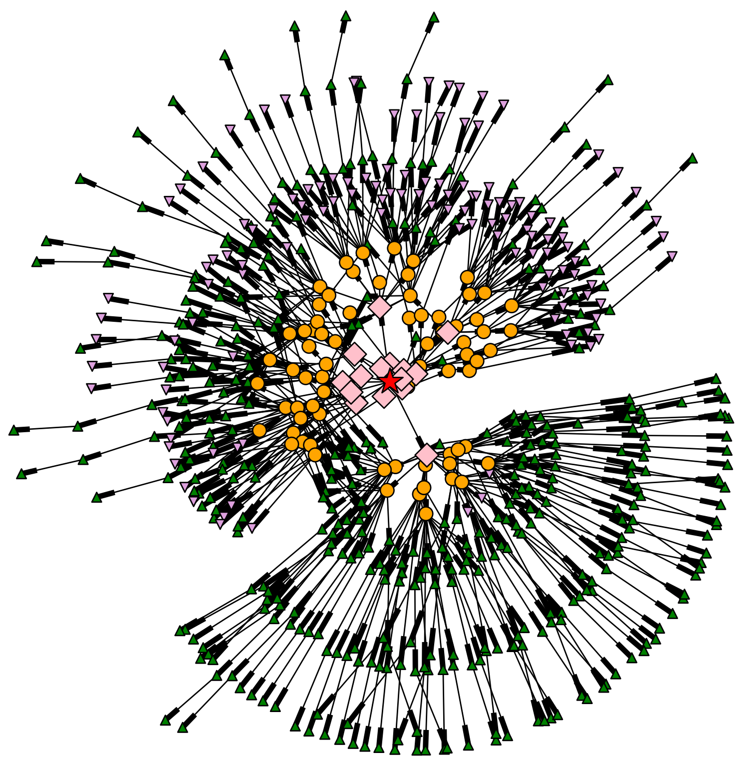

FIGURE 4.3: The L2 navigation structure for an example event processed with a representative trigger configuration. The root node (red star) is drawn in the centre, outward are drawn L1 RoI TEs (pink diamonds), L1 threshold nodes (orange circles) and regular TEs which were created by the HLT sequences and passed the selections (upward pointing green triangles) or were rejected (downward pointing violet triangles). The graph layout drawing has been made with the neato algorithm [49].

4.2.5 Structural information encoding

The use of a di-graph in the encoding of seeding information has several advantages. For a typical algorithm, FEX or Hypo, a single[4] TE is the entry point to the structure allowing the possibility to quickly look back at the history of reconstruction in the context of the TE (RoI) without resorting to the global event data repository.

Another significant advantage is present in the validation of the common origin of two objects. For instance if a trigger signature requires two electrons reconstructed by different algorithms (e.g. by different tracking algorithms or simply when the selections applied are different, e.g. 10 and 15 GeV) the condition of them being two separate objects is conveniently checked by verifying that the branches are not rooted in the same RoI node. This method is extended to selection chains of higher object multiplicity.

The fact that the TEs can be merged allows for a convenient encoding of the origin of the combined object (see join node "To" in Fig. 4.2). When requirement on combined objects are made in succesive selection steps, the join in nodes play a crucial role in ensuring that checks are only applied to combinations that were validated in the previous steps.

The functionality mentioned above can be implemented in the HLT thanks to the di-graph representing the navigation structure without the need for the HLT framework to rely on physical quantities like η and ϕ. In addition to improving the performance it allows the HLT steering to remain physics agnostic and all the decisions are implemented in the HLT algorithms.

4.2.6 Optimisations

The working point of the trigger system, in particular the HLT, demands specific optimisations. Because most of the events are rejected it is more important to be able to arrive at a negative decision with the least possible resource utilisation rather than to optimise the system for the events which are accepted. Other optimisations were already mentioned:

- RoI based reconstruction.
- Step wise reconstruction in chains, allowing the rejection of combinations as well as the grouping of data requests.
- Prescaling of chains at the beginning of the event. Features like coherent prescaling and re-running of chains were added in order to prevent loss of information.
- Buffering and combining of detector data requests.
- Caching of the results of sequences execution.

Caching of the results of FEX algorithms execution, complements the set of optimisations. The configuration of the trigger selection contains many sequences that have an identical set of FEX but distinct Hypo algorithm as these arrays of trigger chains are needed to cover varying luminosity conditions. The execution of the FEX algorithms does not depend on the decisions taken by the Hypo algorithms and therefore it would be suboptimal to re-execute them for every sequence. On the other hand the execution depends on the feature objects reconstructed in the earlier stages (e.g. a different clustering algorithm may result in different RoI position estimate and therefore the tracking algorithm needs to be evaluated for each of these positions.) An automated mechanism has been added that decides wether to execute the

[4] In fact, for optimisation purposes, the sequence input TE and the output TE are passed to the algorithm. This is however only an optimisation allowing the path to the input data to be shortened in some cases. Functionally the output TE on its own is sufficient.

FEX algorithm or use the results already reconstructed in the RoI. This mechanism provides significant resources savings. This problem is similar to the memorisation approach to resource consuming computations in which the input and output of the execution are stored pairwise and for each new execution the result is shipped from the memory if the inputs were already processed by the algorithm. A complication of the HLT case is that the input to the HLT algorithm is a TE. By design each execution is performed on a different TE as the algorithms are parts of distinct sequences and some other similarity of the input has to be found. The structural similarity of TEs (i.e. that they originate from the same RoI) can not be exploited as it is not possible to distinguish the cases when the reconstruction paths are distinct and the execution should be preformed again on the RoI or conversely, when the algorithm execution is not influenced by the different paths and so no repeated execution is needed.

Therefore the following solution, assuming only the reproducibility of the algorithm execution given the same input, is implemented. In the first execution of the algorithm the set of input data it requested and set of output data are recorded. At the next execution the input data requests are preformed on the input TE and if the resulting set of objects is the same as in one of previous executions the repeated execution is avoided and the output feature objects are linked to the output TE for which the execution was scheduled.

Figures 4.4 shows the benefits this implementation as measured for the L2 configuration on data taken in October 2012 with a range of luminosities $\sim 1.5-6\cdot10^{33}cm^{-2}s^{-1}$. As can be seen the optimisation reduces the number of executions which has a significant effect on accepted events in terms of the size of the result. As a consequence the execution times per event are reduced. The most significant reduction is visible for accepted events (factor of $\simeq$2) but the figure of merit, the reduction of execution time for all events, is also reduced significantly. Reduction of the tails in the distributions indicates that this optimisation is most relevant to busy events.

4.2.7 Abnormal execution

The instabilities in the behaviour of the selection algorithms and the system overall in response to unexpected conditions were anticipated and factored into the design of the DAQ and HLT. The requirement states that the event must be recorded for detailed inspection in the case that any problem is encountered in the decision evaluation online. These events are used to improve the algorithms robustness. In addition these events could have special properties resulting from yet unknown physics [50, 51]. Last but not least, the HLT system has no dead time; so all events leaving the L1 have to be either rejected or accepted depending solely on physics properties of the event. This requirement has to be confronted with the need for stable operation of the HLT system. In order to give a better understanding the reasons for malfunctioning of the HLT system a few possible conditions are listed:

- missing detector information,
- unexpected input in the detector data, e.g. very high multiplicity,
- misconfiguration of the algorithm,
- a defect in the algorithm.

In order to provide sufficient flexibility to handle this nonuniform set of conditions several errors-reporting and handling components are added to the HLT.

The HLT algorithm are capable of reporting to the HLT framework error conditions. That report encodes the severity of the error in the form of the desired action of the framework to take in response. These possibilities are as follows:

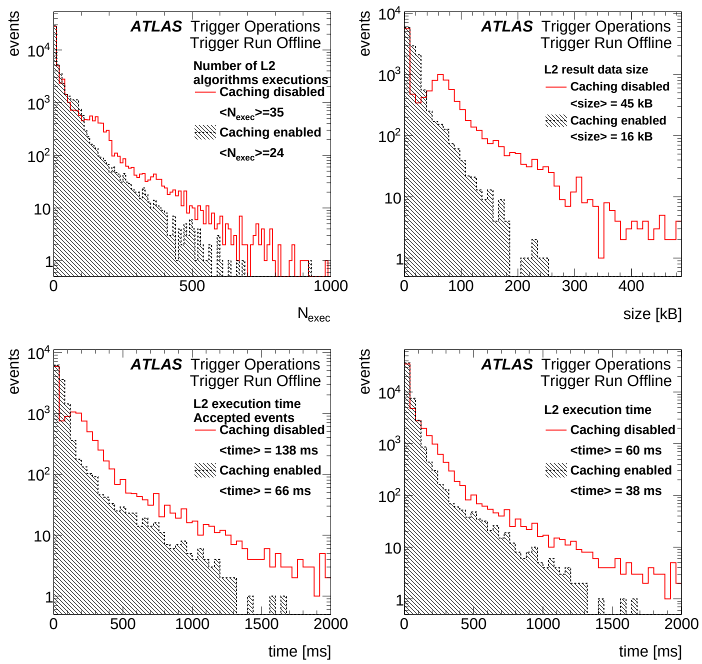

FIGURE 4.4: The effect of enabling FEX algorithms caching on the performance of the L2 system. The reduction in the number of executions of the algorithms with caching enabled (grey histogram) as compared to the case with no FEX algorithms caching is shown on the top right panel. The difference in the result size is shown on the top right. The execution times for accepted (left) and all events (rights) with caching enabled and disabled is shown at the bottom [39].

1. Request to abort execution of the process. It is usually used when misconfiguration of the algorithm is encountered. This type of error condition is found during the process of configuration testing [52] and never occurred during the routine data taking,

2. Request to abort execution of that particular event. A typical condition leading to this type of algorithm behaviour is unexpected input data which results in excessive execution time.

3. Request to abort chain. This is used to flag an interim problem, usually in the input data.

The HLT steering then redirects the events to special, so called *debug*, streams for the recovered. In cases when the HLT process is unable to respond with the decision in a predefined time the DF enables a force-acceptance of that event. These events also end up in the debug stream.

The process of recovery is permitted to consume as much time as needed and the events with non-typical inputs are successfully processed. The bulk of the debug stream content is recovered in this way. In case of a genuine defect in the algorithm an updated version of the code is prepared and used to reprocess the event.

Recovered events are integrated in to the physics data and used in the analyses. It shall be added that this procedure must be fully reproducible including the prescaling which uses an RNG as described in sec. 4.2.1. The aforementioned severe conditions are rarely present in data taking.

Other minor problems occur online at a slightly higher rate and for those cases distinct monitoring system is needed. While reporting the execution status the HLT algorithms in addition action can append information of the cause of issue. Examples are, a missing feature object, a timeout, missing detector data etc. This information is monitored in the form of frequency histograms and reported via the online monitoring system. It is left to the judgement of the operators wether the rate of an error is a new unusual condition or if is within the tolerable margin.

4.3 Monitoring

The HLT system is equipped with an extensive monitoring infrastructure [2, 53] allowing it to verify correct functioning during running. The main purpose is to be able to notice malfunctioning as early as possible and apply remedies. Key aspects are monitored by shifters and by automated systems during running. The detailed performance of specific trigger chains is scrutinised immediately after the data recording. The trigger monitoring is therefore split into online and offline parts. Elements monitored online are:

- rates of trigger chains (and L1 seeds),
- overall recording rates and rates of events directed to particular streams (including debug streams),
- counts of errors (discussed in the previous section),
- several critical quantities per algorithm which have the potential to influence data taking e.g. $\eta \times \phi$ maps of RoI rates which can indicate dead or inappropriately active regions,
- consistency of configuration of entire HLT farm.

The quantities of interest are exported from the algorithms and used to fill the histograms. Mostly plain 1D histograms and correlation plots are used. One HLT process publishes about 10 000 histograms. This histograms are published periodically to the online information system where the partial information is hierarchically gathered and published back to the information system. From there the monitoring information is taken to the online display programs used by shifters and to the automated data quality monitoring framework (DQMF) [54]. The role of the DQMF is to apply test programmed in advance (i.e. comparison to a reference) to the histograms and flag potential missbehaviour in a condensed form to the system operators.

The rate information is also published in the form of the histograms. I order to eliminate the potential for partial summation, and as consequence missed entries, a more robust, yet more resources-consuming procedure is used to collect the rate histograms. The rate information summed over an entire farm is available every 20s[5]. An example of a rate monitoring display is presented in Fig. 4.5. All online monitoring information is recorded for future reference. Summed counts from the entire data taking run is available for all monitoring information and in addition there is per luminosity block information for selected histograms. The rate histograms are recorded with the original publication frequency.

4.4 Configuration of the HLT trigger

The configuration of the trigger system, a menu, is defined by a coherent set of settings for the L1 and the HLT. Constituents of the HLT configuration are the set of chains, sequences

[5] The interval is configurable.

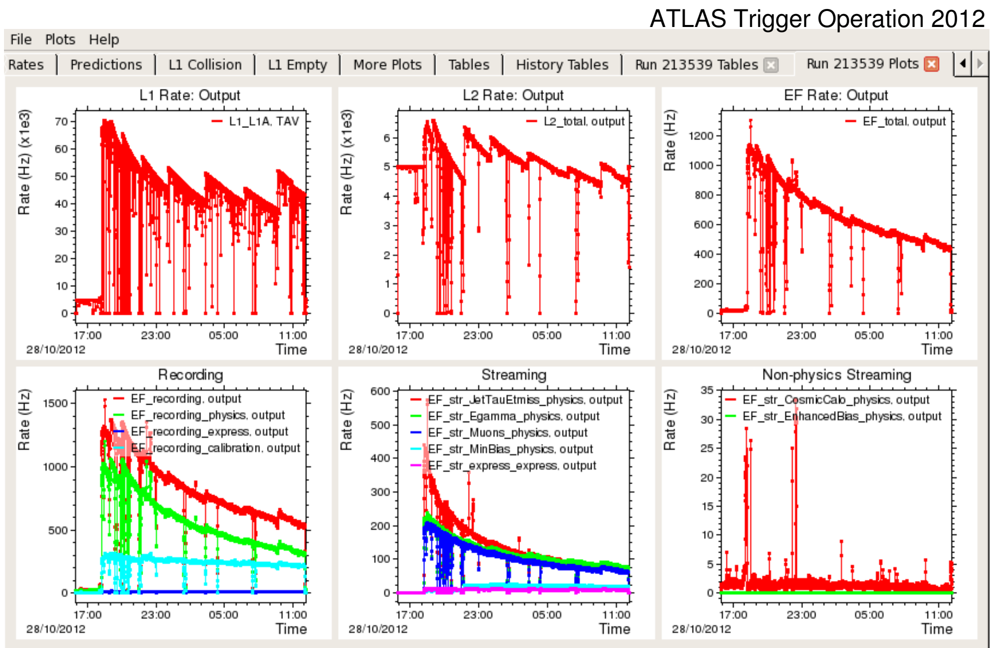

FIGURE 4.5: A screen capture of the online rates monitoring display. Overall rates of L1 L2 and EF are shown in the top row while EF output rates for the streams at the bottom. Spikes and dips, present in all plots are related to prescale changes and dead time respectively [39].

and algorithms and the configuration of each of them. Typical configuration contains about one thousand chains, sequences and related algorithms. Important additional elements of the configuration are prescales for the HLT chains. They can be updated during running online [55] allowing adaptation of the trigger to the data taking conditions.

The configuration is developed in the offline framework using the Python scripting language. Once the development phase ends the configuration is converted to an intermediate XML format and then copied to a relational database, the TriggerDB [56, 57]. Each configuration used for data taking ends up in that database and from there the HLT selection process is configured[6]. Configurations are never erased from the TriggerDB and that allows the inspection of the data taking conditions for all recorded data. Once the new configuration is uploaded to the database each configuration entity (e.g. configuration of the algorithm) is inserted only if the new configuration of that component is different from any configuration used in the past. This approach allows the size of the database to be limited as well as to facilitate the understanding of the changes in the configurations.

Complete configurations are tagged by the database keys assigned to the top level tables which, via relational dependencies, point to the configuration of every element of the L1 and HLT system. There are three keys denoting uniquely the entire configuration, a super-master-key (SMK) which labels the configuration of the L1 and HLT components and separate prescale set keys, separate for the L1 and HLT that define the values of the prescale factors for each item and chain of the respective system. Usually one SMK is associated with many L1 and HLTprescale keys.

A summary of the configuration is extracted from the TriggerDB and saved in the data taking conditions database. It is used during analysis to translate between the compact integer

[6] The HLT farm processes and their connection to DF are configured using different technology [58].

identifiers used internally in the decision record and the human-readable names of the HLT chains, L1 items and their associated information.

In addition to storing new configurations the TriggerDB can be modified directly using a front end tool which allows the modification of every setting of the system, however it is mostly used to hand-tune the prescale factor values during untypical running conditions. Configuration browsing via the web pages is also provided.

4.5 Data analysis

Analysis involving estimation of the trigger efficiencies is supported by a set of tools that allow the analyst to obtain answers to the most typical query, namely if the trigger positive decision is caused by an object of question without hurdle of the internal organisation of trigger information. The analyst is supported by two main tools. A tool allowing the decision of a particular chain to be obtained as wel as the the prescale factors for that chain and the feature objects that lead to that trigger decision, the *TrigDecisionTool*. Another tool automates the procedure match trigger objects to reference objects (i.e. objects found in offline reconstruction), the *TrigMatchingTool*. At the same time the tools support expert analysis, for instance inspection of intermediate feature objects irrespectively of the trigger chain decision.

In order to provide these two functionalities small aspects of the configuration are recorded in the analysis files. Because the configuration is constant for a period of time and can be linked to the event using the run number and luminosity block number as an index it can be stored in the metadata section of analysis files.

The opposite is the case for the HLT navigation structure which is changing for every event. As suggested by the Fig. 4.3 the navigation structure can become quite sizeable. Therefore algorithms were written that allow this size to be reduced by removing navigation tree fragments of no interest for certain analyses. For instance, in files for the analyses that rely solely on muons the branches related to jets, electrons, taus etc. can be removed without loss of information. Branches rejected during processing as well as the intermediate steps are also eliminated as, typically, only a final decision per RoI is sufficient in analysis.

4.6 Revised design in Run 2

Following the successful operation of the trigger during the Run 1 (2009-2013) several upgrades were needed to ensure equally performant triggering in the new experimental conditions at the LHC in Run 2. The expected luminosity increase from $8 \cdot 10^{33} cm^{-2}s^{-1}$ to $1.2 \cdot 10^{34} cm^{-2}s^{-1}$ and centre of mass energy increase from 8 TeV to 13 TeV results in a 5-6 fold increase in the rate of signatures of interest. At the same time the efficiency for specific signatures (i.e. inclusive lepton triggers, b-jets, τ-leptons decaying hadronically) should not deteriorate. The rise of signal rates comes at the expense of raised backgrounds which need to be efficiently rejected in the more difficult pile-up conditions. For completness the conditions expected in the future LHC runs listed in Table 4.4.

In te case of the HLT the main limitation from Run 1 was alleviated, that is the input rate limit to the more advanced offline-type algorithms. This has been achieved by merging the L2 and EF into one system. That is, in Run 2, there is only one process performing the entire selection starting from fast (L2-like) algorithms followed by advanced (EF-like) algorithms in one pass. In most of the cases the partial configurations of the L2 and EF chains were

LHC run	Years	Instantaneous luminosity (cm^2 s^{-1})	Integrated luminosity (fb^{-1})	Average pile-up, $\langle\mu\rangle$
1	2011-2012	8×10^{33}	25	10-20
2	2015-2017	1.2×10^{34}	100	∼27
3	2019-2021	$\sim 2 \times 10^{34}$	300	∼55-80
4	2023-	$\sim 5 \times 10^{34}$	3000	∼140

Table 4.4: LHC luminosity conditions for past, ongoing (Run 2), and future runs.

concatenated into a longer chain of requirements. In some cases it was found more useful to abandon the simple and fast part and use the advanced offline or offline type algorithms from the beginning. For instance in case of the jet trigger the RoI approach used during Run 1 has changed in Run 2 to the scan of all the deposits from the merged RoIs or even the entire calorimeter [59]. A general trend to use more offline-type algorithmic code adapted to the HLT was observed [60] while a the same time the operational performance was improved.

Only a minimal adaptation of the HLT framework were required before the Run 2. It required mostly new configurations of the existing framework core components. The same initialisation procedure as was used for L2 in Run 1. Decoding was expanded for L1 information from L1Topo system and for a new information of cluster energies from L1Calo. The evaluation of the decision of the chains has been slightly modified in order to take from the merger of the L1 and EF systems, namely the retrieval of complete event information is associated to a chain reaching a certain step. That event building step became a configuration parameter. The trigger chains and event-building points are arranged so that the event-building happens at an optimal moment allowing expensive reconstruction algorithms requiring full event information to run when detector data ia already available (e.g. missing transverse E_{T} algorithm).

5

Overview of triggering strategies and their evolution

5.1 Early commissioning of the HLT

5.1.1 Commissioning prior to the first collisions

The commissioning of the HLT triggers started before the first data was observed. It was performed on the MC simulated events containing signals for efficiency measurements and min-bias sample for rates estimate. The summary of the expected performance of the entire trigger and specifically the HLT is summarised in Ref. [61] (chapter 8). The goal of pre-data validation was to ensure excellent physics performance (signal efficiencies) as well as the operational health of the trigger menu composed of many FEX and Hypo algorithms - that is that the execution times were within the constraints of the DAQ system design.

The overall strategy for early running is summarised in Fig. 5.1 where the rates for groups of triggers is shown.

As an example of the HLT performance the tracking efficiency as a function of η is shown in Fig. 5.2 for simulation together with the results obtained on data in 2010. As can be seen from this comparison several improvements were implemented before Run 1 to improve algorithms performance. This was the case for nearly all aspects of the HLT.

The timing tests performed on MC did not take into account the detector data retrieval time and therefore once the TDAQ system became available for tests in 2008 the data were pre-loaded into the RoS system and fed to the HLT in order to obtain measurements including the time spent in processing as well as the time waiting for data retrieval. The result of such

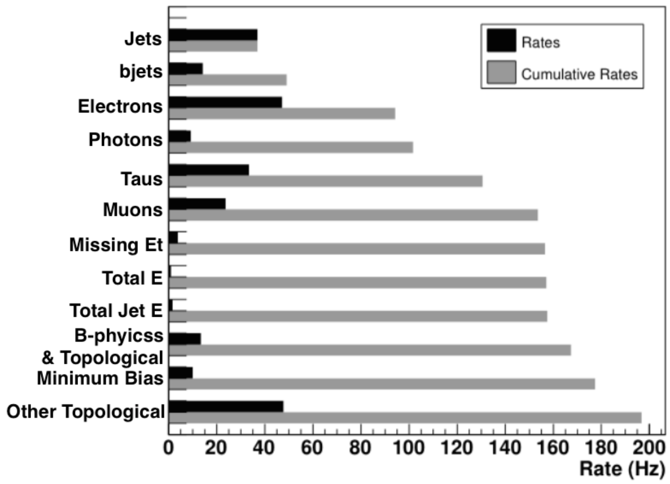

FIGURE 5.1: Rate allocation plan established before Run 1 using simulated min-bias events and optimising for an overall recording bandwidth of 200 Hz at the luminosity of $\mathcal{L} = 10^{31}\mathrm{cm}^{-2}s^{-1}$ [61].

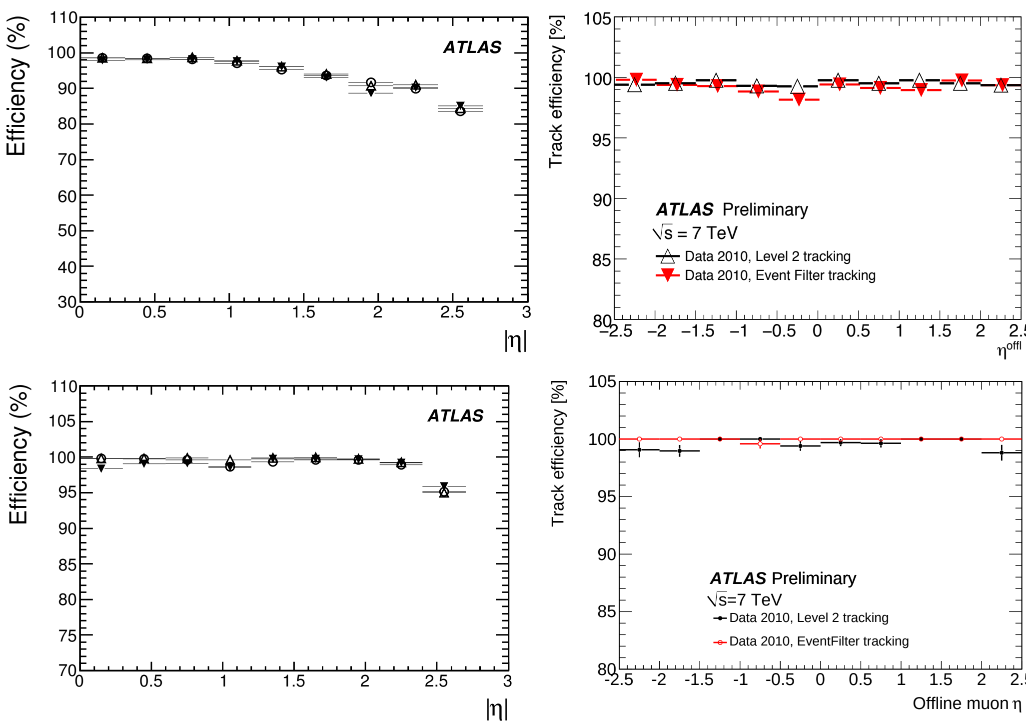

FIGURE 5.2: Comparison of tracking efficiencies obtained from MC in preparation to Run 1 [61](left) and the efficiencies measured on data during year 2010 (early Run 1) [62] (right) for electron (top) and muon (bottom) tracks are shown. On the left efficiencies for two tracking algorithms prepared for L2, *SiTrack* - combinatorial pattern recognition (full triangles), *IdScan* - histogramming algorithm (empty triangles) and EF - offline tracking in RoIs (circles) are presented. The L2 algorithm on the data is the *IdScan* one. In all cases the efficiencies are derived independently of each other are with respect to offline reference.

a comparison obtained back then is shown in Fig. 5.3. An approximate performance scaling between the tests and the HLT system operation at P1 was found using these measurements.

An important next step of the commissioning was performed with cosmic rays which allowed

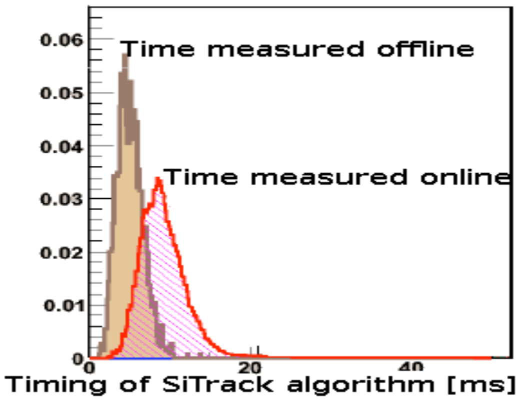

Figure 5.3: Comparison of the algorithm execution time as measured online in tests involving data preloaded to the RoS system and offline [1].

the HLT algorithms to be tested on the data from ATLAS sub-detectors. After that the first collisions[1] in 2008 occurred which allowed validation of mostly the L1 trigger system i.e. the signal timings. After the LHC accident in 2008 operation resumed in the fall of 2009 with a very low luminosity of $\mathcal{L} \simeq 1 \cdot 10^{27} cm^{-2} s^{-1}$.

5.1.2 Low luminosity running in 2010

The general strategy for commissioning the HLT in 2010 was to defer active rejection at the HLT until it had been validated and run demonstrated to run safely on data taken solely using the L1. For that purpose the first menu deployed on data contained HLT chains which did not contain any algorithmic steps. Each L1 items had an associated HLT chain and the presence of a L1 trigger resulted in event acceptance. The HLT framework, was thus utilised from the first day of data taking providing the important functionality of event streaming. The streams in this commissioning configuration were used to separate events with minimum bias signatures (mostly events with a signal in the MBTS detectors), calorimetric signatures with electromagnetic and hadronic objects and muon triggered events. A fraction of events were also directed by the HLT to the so called express stream which is used in ATLAS to provide input to data driven calibrations, detector checks and data quality assessments. Initially these were a subset of events from all L1 items but that configuration was soon refined to fill the express stream more selectively. That was done by means of adjusting the fractions of events from HLT chain which cause an event so be directed to this stream.

In the next stage HLT chains containing algorithms were added to the configuration but without actually enabling them. Events recoded with the help of L1 only triggers were then reprocessed offline with the HLT algorithms enabled. After confirmation of the correct functioning when the algorithms were enabled they were used online in a limited scope i.e. the Hypo algorithms were set into always-accept mode. This mode of operation allowed the monitoring infrastructure to be tested with real data as well as to verify of the data taking stability with HLT farm requesting real detector data for the first time. Once the luminosity grew it was necessary to enable HLT rejection for some of the trigger chains. The first trigger for which this was done was a minimum bias performance trigger, *mbSpTrk*, in which the L1 was based on all filled bunches and the HLT was used to select collision events reling on counting of hits in the ID at L2 and track finding at the EF. Figure 5.4 shows the very first data taking run where that

[1] First beams provided so called beam-splashes where low intensity proton bunches were steered to tertiary collimators and showers of the size of entire ATLAS detector resulted in signals that were triggered on.

HLT chain was enabled. At the start of the run the rates were mostly due to random triggers and cosmic rays (a non zero rate of muons caused by muon trigger from cosmic showers). Once stable beam conditions were reached the data taking lasted until 15:30 when the L1 random trigger seeding the mbSpTrk chain was enabled in several steps giving a significantly higher rate which allowed the HLT to obtain more and more events. The observed EF output rate did not change which is a sign that a significant fraction of events were empty and only some (~1/1000) contained a track and were retained by the HLT.

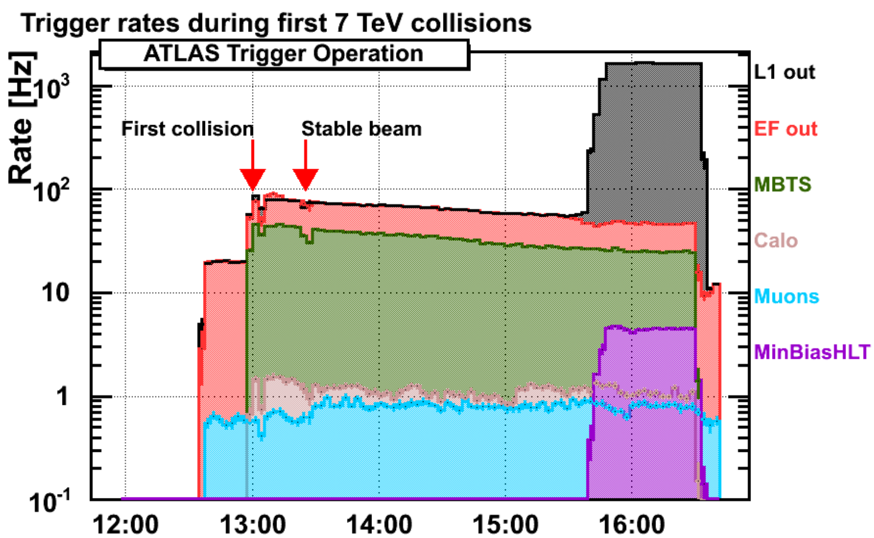

FIGURE 5.4: Selected online monitored rates as a function of wall clock time in the first run when active rejection at the HLT took place in 2010 at $\sqrt{s} = 7$ TeV. The L1 and EF rates shown are total ones. Rates labeled as Muons and Calo are for the respective streams seeded by the L1 items while the MBTS is for an or of L1 trigger rates is shown. The MinBiasHLT denotes the rate of events triggered by the mbSpTrk chain [39].

The example of the mbSpTrk commissioning chain illustrates that the independence of HLT chains served well in the initial commissioning phase. Gradually, with the rising luminosity all HLT chains were functioning online and were used to reject events. At $\mathcal{L} \simeq 1 \cdot 10^{30} cm^{-2}\mathrm{s}^{-1}$, in the middle of 2010, the commissioning menu was replaced by the physics menu in which all selections relied on the HLT, at least for low p_{T} thresholds. The menu reflected the ATLAS physics strategy expressed in torms of a guidance on the overall rates for type of trigger as well as a specification of the so called primary physics triggers. The requirement for these triggers was that they should stay unprescaled (provide loss-less selection) throughout a significant part of the data taking period. The raise of the thresholds as well the tightening of selection requirements were exercised in adjusting the menu to increasing luminosity. Table 5.1 summarises the most important primary triggers in the physics menu from 2010 as a function of instantaneous luminosity averaged over the duration of the LHC fill. The luminosity values given in this summary are only approximate, for instance the 10 GeV muon trigger was prescaled for the first time during a run at a luminosity slightly above $\mathcal{L} \sim 3.5 \cdot 10^{31} cm^{-2}\mathrm{s}^{-1}$ on the 1st of October 2010. More details on the trigger thresholds and rates will be given in corresponding sections. The time evolution of the peak instantaneous luminosity delivered by the LHC in 2010 is shown in Fig. 5.5. Trigger chains with lower p_{T} thresholds were prescaled dynamically in order to provide sufficient statistics for auxiliary measurements (for instance efficiency estimations) while

Luminosity [$cm^{-2}s^{-1}$]	$3 \cdot 10^{30}$	$2 \cdot 10^{31}$	$2 \cdot 10^{32}$
Category	p_T threshold [GeV],selection		
	Inclusive		
muon	4	10	13,tight
electron	10,medium	15, medium	15, medium
photon	15,loose	30,loose	40,loose
tau	20,loose	50,loose	84,loose
jet	30	75,loose	95,loose
	Di-objects and other		
di-muon	4	6	6,loose
di-electron	3,loose	5,medium	10,loose
di-photon	5,loose	15,loose	15,loose
E_T^{miss}	25,tight	30,loose	40,loose
B-physics (di-muon)	4	4	2 + 4

Table 5.1: Unprescaled triggers of the lowest p_T thresholds in the initial physics menu in 2010 [18]. The B-physics triggers were required two muons with additional invariant mass cut around the J/ψ and Υ resonances.

maintaining overall recording rates within the limits.

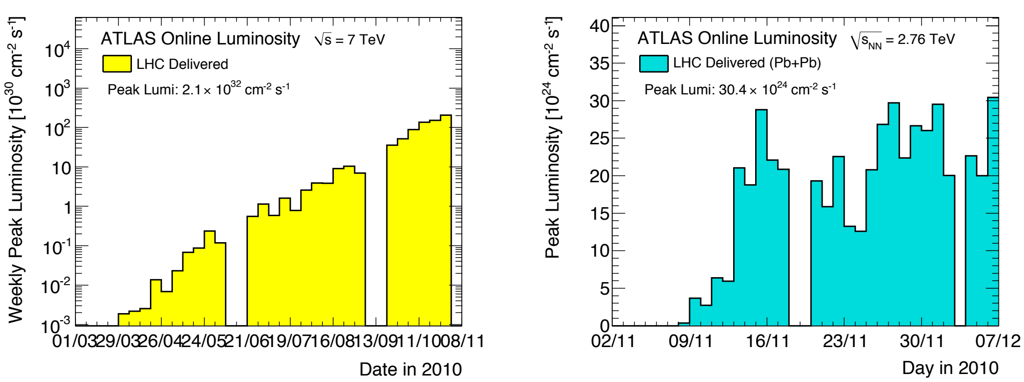

Figure 5.5: The peak instantaneous luminosity delivered by LHC as a function of time progression in 2010 during pp - one week granularity (left) and Pb-Pb - single day granularity (right) data taking periods [63].

Commissioning in the 2010 concluded with the first Pb-Pb run at a centre of mass energy per nucleon pair $\sqrt{s_{NN}} = 2.76$ TeV. Luminosities during that runs reached $\mathcal{L} \sim 30 \cdot 10^{24} cm^{-2}s^{-1}$ and rates were so low that the L1 alone could be used for data taking. The HLT was used to stream the data according to the L1 selections with most events directed to a single *bulk* stream. A small fraction of events without any bias ($\simeq 0.5\%$) and background ($\simeq 2\%$) were also written out separately. Similarly to the early commissioning during the pp runs the only actively functioning trigger at the HLT was used to select events based on the activity in the ID, the *mbSp* trigger which was seeded by a random subset of colliding bunch crossings and then counted number of detector signals in the ID at the HLT. That trigger was used to measure the efficiency of the MBTS based L1 selection.

5.2 Triggers relying on reconstruction in RoIs

Among all the strategies implemented in the HLT the strategies in which algorithms perform reconstruction in RoIs are dominant. The reconstruction of calorimeter clusters, tracks in MS and ID are performed in this way. Accordingly, most of the Hypo algorithms function in a single RoI. A brief description of the main algorithms will be given below concluding with an overview of combined strategies using several detectors.

The L1 delivers the positions of RoIs in the calorimeters and MS. An extrapolation to the ID has to be performed in order to associate the ID tracking information. In case of calorimeter information, to a good approximation, the position (η, ϕ) of the RoI in the ID is the same as in calorimeter however the shape of the region needs to take into account all possible projections from the luminous region to the cluster. In the muon trigger the track reconstructed in MS has to be extrapolated to the ID taking into account the magnetic field, muon position and p_{T}. The RoIs relations are visualised in the sketch in Fig. 5.6.

Various optimisations are implemented on top of this basic scheme in order to reduce volume in which reconstruction is performed. For instance, taking into account cluster orientation or simple vertexing reduces the volume in which track reconstruction is performed.

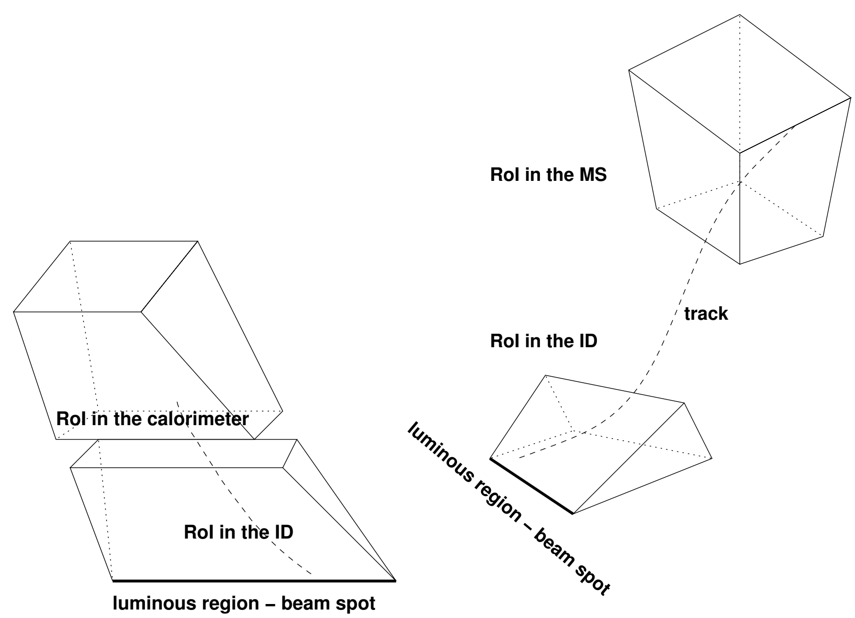

Figure 5.6: A sketch of the RoI in the calorimeter and the RoI associated to it in the ID (left) and in the MS and the corresponding one in the ID (right). The luminous region, a potential origin of the particle (dashed line), is drawn at the bottom (thicker line).

The inclusive triggers based on an RoI lead to the reconstruction of all regions identified by the L1 system. For more complex signatures e.g. di-object triggers, the HLT steering assures the correctness of the decision evaluation which simplifies trigger menu maintenance. For example a di-electron trigger has all the FEX and Hypo algorithms and sequences configured as a single electron trigger but with the requirement for two objects at each chain step. For symmetric triggers (e.g. L2_2e15 which is 2×15 GeV electron) simply trigger element counting is performed. For asymmetric triggers (e.g. L2_e15_e20 which is requirement of two electrons of which one is above 15 GeV and other above 20 GeV) an additional uniqueness check is performed. This overlap check relies on a fast introspection of the seeding relation in the HLT navigation structure as it has to be performed very often. The overlaps are removed by the HLT steering for every type of signatures and for combination of e/γ and τ triggers as they are initiated by the same L1 hardware. In cases where overlap removal is needed among other objects either a general purpose HLT overlap removal algorithm, which relies on geometric information extracted from the RoI descriptor object is used.

5.2.1 Electron and photon triggers

The main motivation for the electron triggers is to be able to record events with electroweak bosons through the decays $W \to e^{\pm}\nu$ and $Z \to e^{+}e^{-}$. This events are useful because they allow for many measurements within the SM and can also be used to test detector, trigger and reconstruction performance. In the initial running the decay $J/\psi \to e^{+}e^{-}$ was also targeted for the same reasons. However the most important purpose is the possibility to record events containing Higgs boson decays via electroweak boson channels and searches for new gauge bossons with decays containing the electron in the final state. Therefore an effort is made to keep the E_{T} thresholds above which the electron trigger is nearly 100% efficient as low as $E_{\mathrm{T}} \sim 20$ GeV, in order be fully efficient for of Z and W decays. It is achieved by employing sophisticated selection criteria and in general closely following developments of the techniques used in the offline electron identification.

Similarly for photon triggers the main physics goal is to collect events in the golden Higgs decay channel $H \to \gamma\gamma$ [2] [64]. Signatures with photons are important in many searches for new phenomena [65]. Photons produced in association with jets via $qg \to \gamma q$ are an important tests of QCD and can also be used to calibrate jet energy resolution and biases [66] or study jet quenching in the QGP. The E_{T} thresholds for di-photon triggers were kept at the lowest affordable values, around 20-30 GeV while single photon E_{T} threshold were significantly higher, reaching 120 GeV for an unprescaled trigger in 2015 p-p running and around 20 GeV in Pb-Pb (a summary of the thresholds will be given later).

The electron and photon trigger reconstruction uses identical calorimetric FEX algorithms and nearly identical Hypo (for the calorimetric selection). They differ in the tracking which is performed for the electron candidates only. The logical sequence of the steps performed in the electron trigger is shown in the diagram in Fig. 5.7. First the calorimeter clustering and then tracking will be described. At the end of the section the Hypo algorithms for photons and electrons and their evolution will be covered.

Reconstruction of the electron or photon is initiated by the presence of an EM cluster found by the L1 e/γ trigger[3]. The position of the RoIs (a set of tuples containing η, ϕ and energy thresholds) are made available to the clustering algorithm by passing the TE associated to each EM RoI. The RoI position is translated to a set of readout buffer identifiers by a *Region Selector* service component. Then the L2 HLT calorimeter algorithm initiates retrieval of the detector data from the RoS only for fragments of relevance to processing a given RoI. Cluster size is assumed to be within the area of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ around each region indicated by the L1Calo system.

[2] ATLAS detector is particularly suited for this measurement due to the construction of the electromagnetic calorimeter.

[3] Photons and electrons are identified at the L1 using algorithm operating on trigger towers in the calorimeter and therefore they are indistinguishable at this stage (see section 2.5).

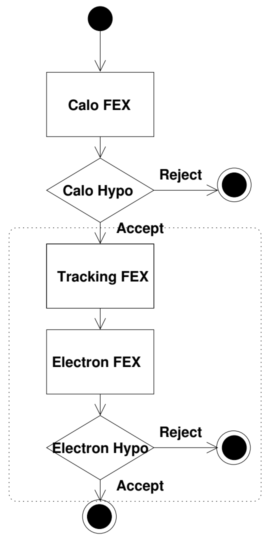

Figure 5.7: The activity diagram of a fragment of the electron trigger chain. The part surrounded in the dotted rectangle is only present in the electron trigger chains configuration. Rectangles denote the FEX algorithms of the sequences while the romboidal shapes denote Hypo algorithm instances.

The longitudinal development of the electromagnetic cascade results in the largest deposits in the second calorimeter layer. Therefore the seed, cell with the largest energy, is sought in this layer. Around that cell the cluster is formed by including all cells within the region in the $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ and calculating an energy weighted position. Similarly cells in other calorimeter layers are added to the cluster. The energy of the cluster, after correction for leakage outside of the defined region is assumed to be the energy of the photon or electron. In addition a shower shape variables are calculated and used to discriminate between the signal electron and photon and hadronic background [43, 67]. In order to improve pile-up robustness an alternative technique based on an artificial neural network is being studied for use in Run 2.

The fast clustering stage is retained even in the merged HLT system used in Run 2 as it offers a good rate reduction before, the fast, but more expensive, tracking algorithm.

At the EF stage the clustering is performed by offline algorithms and tools adapted for the trigger case, namely to be capable of running within an RoI. The energy in cells is grouped in so called towers in order to avoid false seeds due to a single noisy cell. Similarly to the L2, after this additional cells grouping, seed finding and then clustering are performed [68]. Furthermore, at the EF stage a better cells energy calibration is used.

A figure of merit for clustering algorithms, energy resolution with respect to the energy measured in the offline, is shown in Fig. 5.8 for several running periods. As can be seen the fast L2 algorithm is delivers a poorer resolution compared to the EF like algorithm and that progress in the development of the algorithms has resulted in a, significant improvement in the resolution, even though the pile-up conditions have become harsher.

The Pb-Pb running conditions necessitated yet another development in the clustering algorithms of the e/γ triggers. This is because soft particles, produced copiously in the collisions, contribute to the energies in calorimeter cells and distort the energy measurement. Also the shower shape variables used in the e/γ identification are significantly changed. Correction techniques used for pile-up corrections in p-p runs can not be used in Pb-Pb run because the level of the activity changes from collision to collision due to stochastic nature of collision centralities. Therefore additional algorithms are inserted in the sequence which are used to correct the energy of each individual cell before the EF clustering stage. The first algorithm calculates

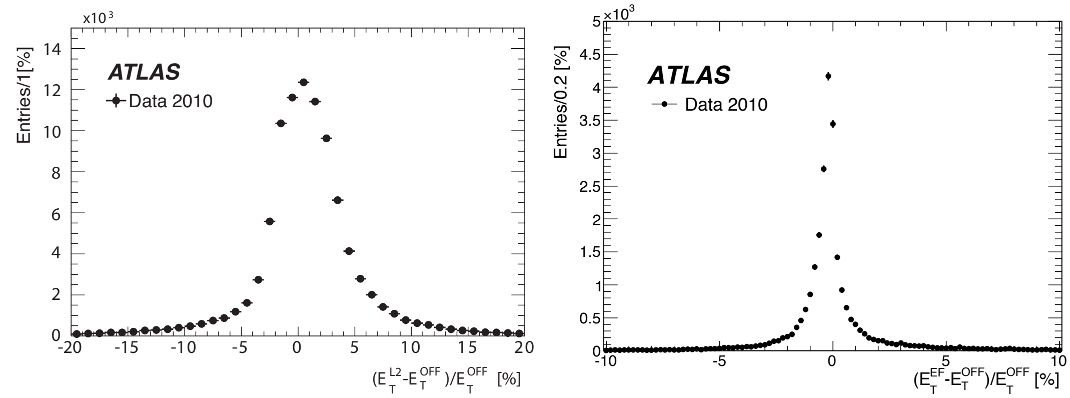

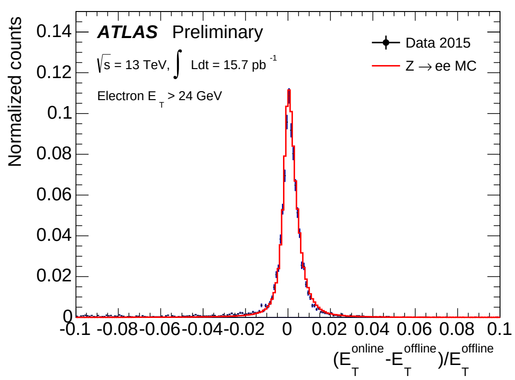

FIGURE 5.8: Transverse energy resolution of the HLT reconstruction algorithms with respect to the energy reconstructed by an offline algorithm. Performance from very beginning of the LHC run 2010 is shown for L2 (top left) and EF(top right) [28] and recent from 2015 (bottom) [69, 70].

the average energy in slices of $\Delta\eta = 0.1$ for each calorimeter layer. This algorithm does not operate in RoIs but reads information from entire calorimeter. The second algorithm, using these averages, provides a set of cells with corrected energies to the p-p like clustering stage. After that modification the cluster energy and shower shape become almost independent of centrality and similar to those observed in p-p running conditions which allows reuse of the p-p selection cuts for Pb-Pb running.

For electron triggers the ID tracking is performed in two stages, fast but less precise algorithm after the L2 clustering and an offline-like algorithm after the EF clustering. Both start from a conversion of the encoded detector data into the *Space Points*, that are 3D positions of the hits.

The L2 tracking algorithm, *IdScan*, uses a histogramming technique. Identical steps are performed inside each RoI that contains an electron candidate. Initially the approximate vertex z position is found by identifying a peak in a histogram obtained by straight line extrapolation to the luminous region of all possible track segments formed from pairs of space points. The domain for the pairs is reduced by dividing the RoI into slices in ϕ in order to avoid combinatorial complexity as well as to give a preference to tracks of higher p_{T}. Once the vertex z coordinate is established the relative η, ϕ positions of the hits are used to fill a 2D histogram. Peaks in this histogram indicate hits that belong to the same track. Hits forming the peaks are taken to yet another 2D histogram $1/p_{\mathrm{T}} \times \phi$ that filled using all hit triplets from distinct detector layers. This tracking technique gives very high efficiency at the expense of a higher probability of finding a fake track. It is also applicable to conditions of low pile-up due to an assumption a limited number number of peaks in the first (vertex z position) histogram.

Track candidates are then fed into a Kalman filter algorithm and subsequently extrapolated to the TRT ti include more hit information. Together with the track parameters an information on the number of high threshold hits from the TRT associated to that track is recorded in order to enhance discrimination of e from $\pi^{\pm}$ particle tracks.

At the EF a set of offline tools is used. It happens that the choices made for wrapping the offline tools has been such that EF tracking is the longest sequence in the HLT consisting of more than 10 algorithms doing data preparation, triplets formation from hist in the Pixel and SCT layers and tracking using a Kalman filter. They are followed by the track candidates filtering to retain only the high quality tracks, ambiguity solving and extrapolation to the TRT.

The performance of the tracking part of the electron trigger chain is measured using a supporting electron trigger in which the track requirement was not applied and therefore a possible inefficiency in the tracking does not cause event rejection. After offline reconstruction of the sample collected using this supporting trigger the HLT reconstructed track close (min $\Delta R < 0.1$) to each offline track associated to an electron candidate is sought. Absence of such a track is a sign of inefficiency. The same technique is used to verify the performance of the L2 and EF parts.

The overall excellent efficiency of the HLT tracking algorithms in Run 1 is shown in Fig. 5.9. As can be seen efficiency is better at the EF than at the L2. Nonetheless both remain very high despite the significant increase in luminosity and therefore pile-up during Run 1.

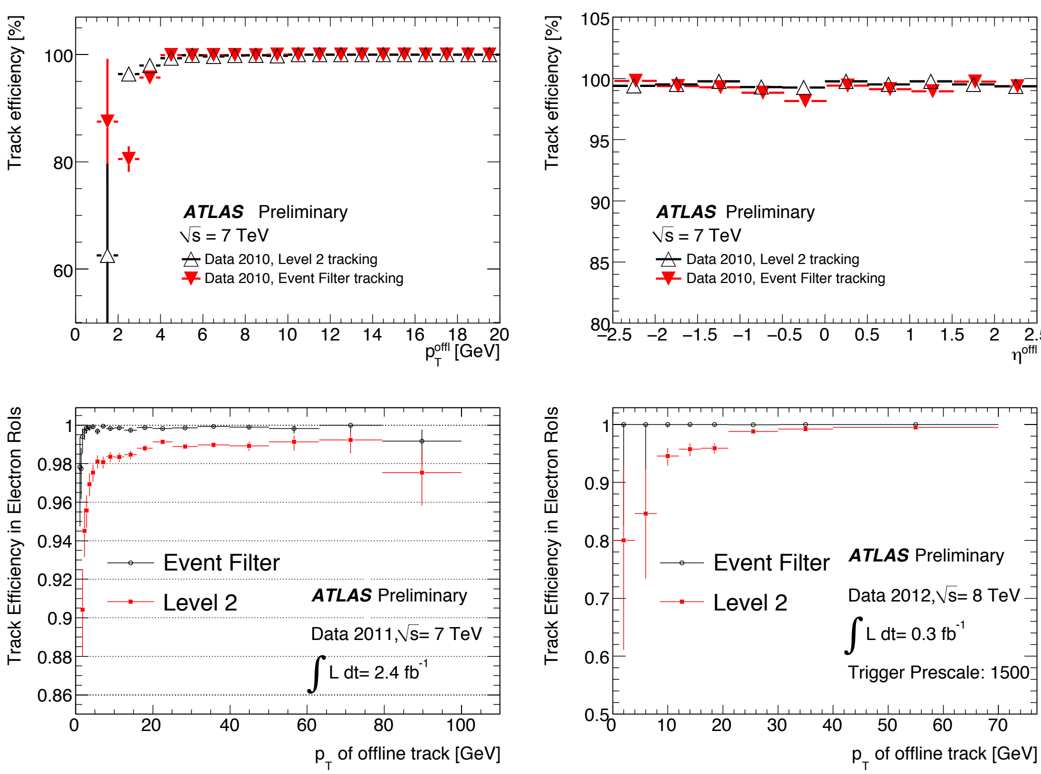

FIGURE 5.9: Performance of the L2 and EF tracking algorithms in electron trigger chains as a function of offline track transverse momenta, p_{T}, and pseudorapidity, η, in 2010 (top) 2011(bottom left) 2012 (bottom right) [71]. Electrons used to obtain the performance plots in 2010 were required to be of medium quality and above 5 GeV, triggered by a 22 GeV electron tracking monitoring trigger in 2011 and 24 GeV in 2012.

HLT tracking algorithms underwent upgrades before the Run 2. The main reason for the improvements was to profit from both the fast and the precise tracking algorithms running within the same process in the merged HLT system. For instance the data preparation and early track forming steps could be performed only once per RoI. In addition the incorporation of the IBL layer resulted in an improved resolution of the tracks impact parameters which is important in high pile-up conditions. Together with the refurbishment of the tracking code components a significant effort was invested in optimisation which in turn resulted in equally good tracking algorithm performance for electron triggers as in previous years, but, running significantly faster [72, 73]. The performance measured on data as well as the improved execution times are shown in Fig. 5.10 and 5.11

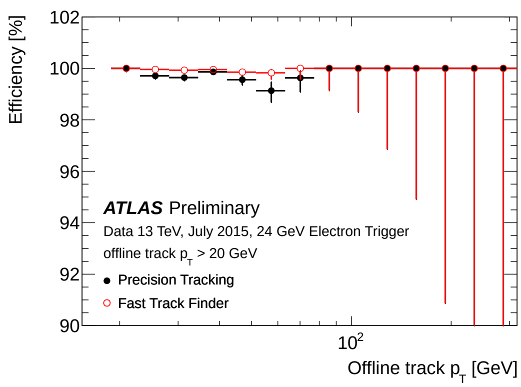

Figure 5.10: The efficiency of the fast (L2-like) and precision (EF-like) tracking algorithms in electron trigger chains as a function of the offline track transverse momenta, p_{T}, in 2015 [73].

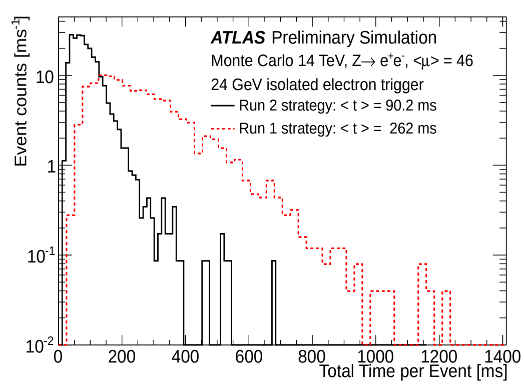

Figure 5.11: A comparison of the tracking algorithm execution times with the strategy used in Run 1 ported to a single stage HLT system with the improved strategy used in Run 2 [73].

Selection algorithms in electron and photon chains

The trigger menu consists of numerous electron and photon chains designed to probe a spectrum of E_{T} ranges. The reconstruction algorithms are the same for all of them and the differentiation is implemented by Hypo algorithms, two in photon and four in electron chains. They do selection on quantities prepared by the proceeding FEXe and therefore use either calorimetric only or combined electron objects.

The former are shower shapes listed in the table 5.2 and E_{T}, while electron Hypo algorithm apply selection on the track to cluster matching and track quality as documented in table 5.3. Figure 5.12 shows MC predictions for one of the shower shape variables before data taking as well as the comparison of the signal distributions as seen in the data. The MC models the data reasonably well.

Three sets of criteria are prepared in order to achieve balance between the rates (purity) and efficiency with increasing luminosity. They are traditionally called selection tightness and are named *loose, medium* and *tight* which denote higher background rejections respectively. It is achieved by including more discriminating variables and, in cases, applying a more constraining cuts on the variables. Thresholds are optimised as a function of η (significant variation for some variables) as well as a function of E_{T} (small tuning in steps of $\sim$ 10 GeV). At the L2,

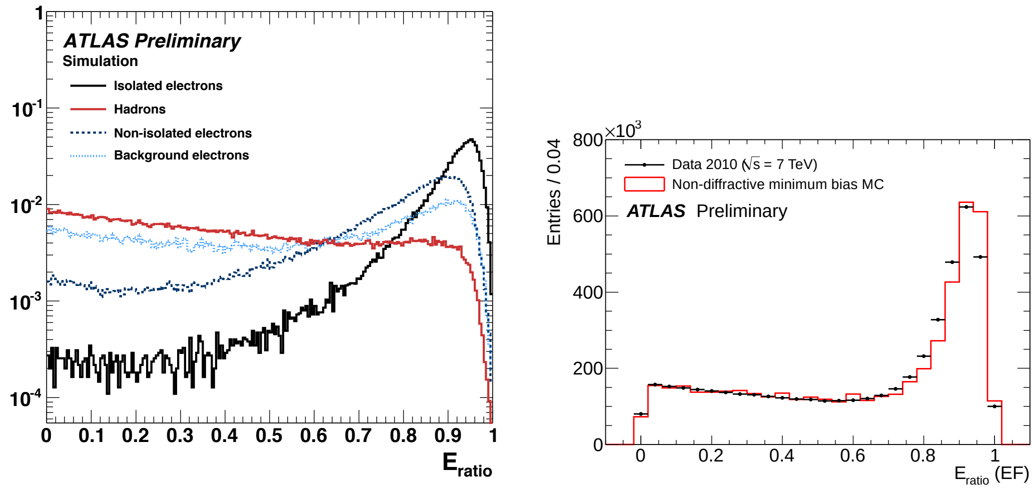

Figure 5.12: Electron identification variable, E_{ratio} for signal and backgrounds in MC prior to the data taking (left) [74] and data to MC comparison from 2010 (right) [75].

in addition to E_{T}, only two cuts, R_{had1} and R_η were used throughout the Run 1 and so is the case in Run 2. Cuts for *loose* photons are the same as for electrons.

The efficiency of the electron and photon triggers in early runs are shown in Fig. 5.13. A very steep turn-on as a function of E_{T} can be observed as well as uniform, and nearly 100%, efficiency as a function of η. For electrons the measurement was performed using tag-and-probe and orthogonal trigger methods in events with Z and $W^{\pm}$respectively. For photons the bootstrap techniques was used and a cross check with $H \to \gamma\gamma$ MC performed.

In 2011 at luminosity $\mathcal{L} \sim 3 \cdot 10^{33} cm^{-2}\mathrm{s}^{-1}$ the selection criteria were redefined (labeled as *loose1, medium1, tight1*) in order to adhere to changed offline identification criteria wich in turn was a response to higher pile-up conditions. Also at the L1 number the triggers were redefined so that the E_{T} requirements became a function of η. Also the hadronic isolation/veto requirement has been applied at the L1 at $\mathcal{L} > 2.3 \cdot 10^{33} cm^{-2}\mathrm{s}^{-1}$ [75]. Comparison of efficiencies for the primary electrons and photons in 2011 is shown in Fig. 5.14. As compared to the 2010 the thresholds are higher. The turn-on as a function of p_{T} for electron triggers was measured with respect offline *medium++* and has gentler raise and plateau for the *medium1* selection reaches only 98% at ∼35 GeV. A small drop in efficiency of *e22vh_medium1* as compared to *e22_medium* is justified by about two-fold decrease of the rates [75]. The influence of the pile-up on the electron trigger performance is also shown in Fig. 5.14. The efficiency deterioration by about 3% is caused by the L2 selection.

Photon triggers efficiency was measured with bootstrap method for L1 with respect to minimum-bias and then the HLT reative to L1. The latter is shown in Fig. 5.14. The L1 efficiency (not shown) reaches nearly 100% [75] and the HLT remains also very high.

For the Run 2 the discriminating quantities used in electron chains are combined using likelihood method and form one discriminating parameter. Main reason for following this approach was that the offline identification relied on it and that it performs better after the change of gas-mixture in TRT which made f_{HT} (see Table 5.3 less discriminating. The performance of

variable	definition	2010/2011 *loose*	2010/2011 *medium*	2010/2011 *tight*	2011/2012 *loose1*	2011/2012 *medium1*	2011/2012 *tight1*
Hadronic calorimeter							
Hadronic leakage, R_{had1}	Ratio of the E_{T} in the first sampling og the hadronic Tile calorimeter to the cluster E_{T}	✓	✓	✓	✓	✓	✓
Hadronic leakage, R_{had}	As above but hadronic deposites from all samplings is used	✓	✓	✓	✓	✓	✓
Middle layer of the EM calorimeter							
R_η	Ratio of E_{T} in $\eta \times \phi$ window of 3×7 to the E_{T} in 7×7	✓	✓	✓	✓	✓	✓
R_ϕ	Ratio of E_{T} in $\eta \times \phi$ window of 3×3 to the E_{T} in 3×7						✓
Lateral shower width, $w_{\eta 2}$	$\sqrt{\Sigma E_i \eta_i^2 / \Sigma E_i - (\Sigma E_i \eta_i / \Sigma E_i)^2}$ where sum goes over cells in a window 5×3	✓	✓	✓	✓	✓	✓
Strip layer of the EM calorimeter							
E_{ratio}	Energy difference of two highest deposits over their sum		✓	✓		✓	✓
Total shower width, w_{stot}	$\sqrt{(\sum E_i (i - i_{max})^2)(\sum E_i)}$ where i is the strip number and sum goes over $\sim$20 strips		✓	✓		✓	✓
w_{s3}	Shower width with the three strips around maximum						✓
F_{side}	$(E_7 - E_3)/E_7$ where E_n stands for sum over n strips around the maximum.						✓
ΔE	Energy difference between second largest and smalles one in between two maximal deposits.						✓

Table 5.2: Definition and use of shower shape variables in electron and photon identification [76, 75]. Additional constraint of that is cluster is not in the region of the detector gap $1.37 < |\eta| < 1.52$ and within $|\eta| < 2.47$ and < 2.37 for electrons and photons respectively.

		2010/2011		2011/2012		
variable	definition	*medium*	*tight*	*loose1*	*medium1*	*tight1*
Track quality						
Pixel hits, N_{pix}	Number of hits in the pixel detector	✓	✓	✓	✓	✓
B-layer hits, N_{BL}	Number of hits in the pixel detector b-layer		✓		✓	✓
SCT hits, N_{Si}	Number of hits in the SCR detector ≥ 7		✓	✓	✓	✓
Transverse impact parameter, d_0	Distance of closest approach between the track and primary vertex, < 5 mm (< 1 mm in tight selections)	✓	✓	✓	✓	✓
Track matching						
Matching, $\Delta\eta_1$	Longitudinal distance between the cluster and extrapolated track impact position in the strip calorimeter layer, < 0.01 (< 0.005 for medium1 section)	✓	✓	✓		✓
E/p	Energy ratio of the cluster to the track momenta		✓			✓
$\Delta\phi_2$	Azimuthal distance between the track extrapolated to second calorimeter layer and the cluster centre in that layer		✓			✓
TRT information (for tracks $\|\eta\| < 2$)						
N_{TRT}	Number of hits in TRT for the track		✓			✓
f_{HT}	Ratio of number of high threshold hits to all hits in TRT for the track		✓		✓	✓

Table 5.3: Description of tracks related discriminating variables used in the electron triggers through Run 1 [76, 75].

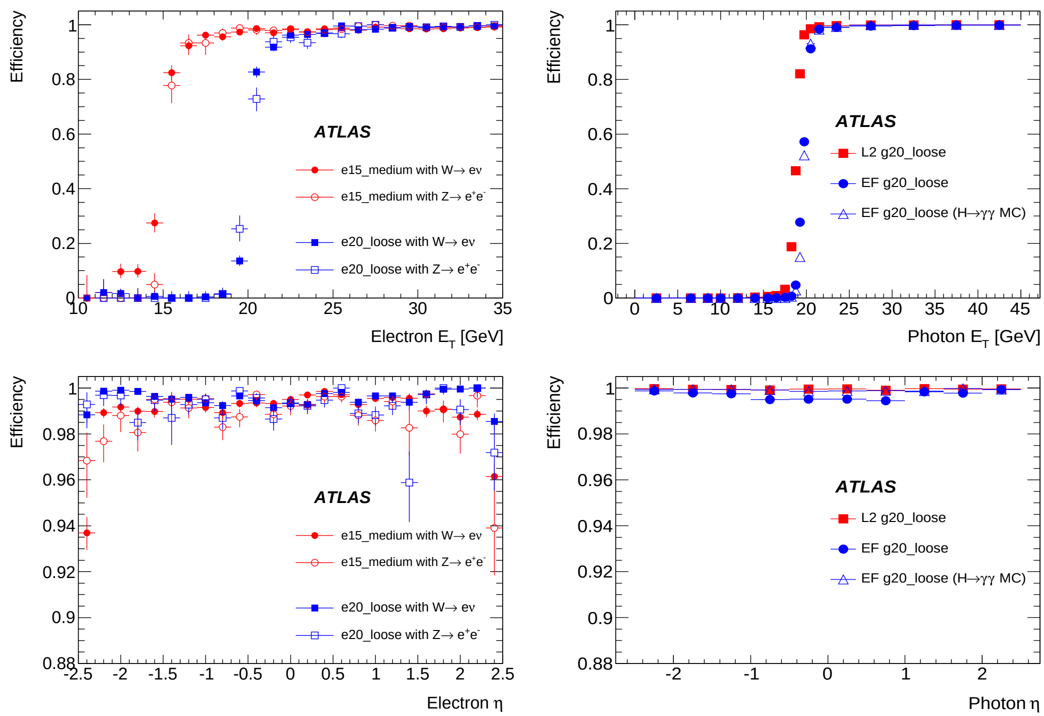

FIGURE 5.13: Efficiencies of the electron (left) and photon (right) triggers in 2010 as a function of transverse momenta, p_{T}, and pseudorapidity, η, of the offline reconstructed objects [18].

the electrons and photon trigger in 2015 is shown in Fig. 5.15. Introduction of the likelihood based triggers did not affect efficiencies while a small rates reduction was achieved. The photon trigger efficiencies remain very high for low and high E_{T} regions as also shown in this figure.

The aforementioned modification to electron and photon chains for Pb-Pb run resulted in a significant performance improvement as compared to an unaltered p-p configurations. The efficiency of the photon trigger as a function of collision centrality (which translates directly to the occupancy int the calorimeters and thus shower shapes modification) is shown in Fig. 5.16. The impact of applying the correction to the cell energies can be observed by comparing corrected and uncorrected chain efficiencies in the most central events (highest $\Sigma E_{\mathrm{T}}^{\mathrm{FCal}}$). At the peripheral events both efficiencies become similar as the activity level from the underlying event becomes similar to the p-p running conditions. A similar approach, in which full calorimeter scan is employed in the electron and photon chains in calculation of hadronic isolation is planned for high pile-up p-p running.

Selection of primary e/γ trigger rates is shown in the table 5.4 for selected runs along the Run 1 and beginning of the Run 2.

	$\mathcal{L}$ × 10^{30} [$cm^{-2}s^{-1}$]	single-e [Hz]	di-e [Hz]	single-γ [Hz]	di-γ [Hz]
2010					
begin	∼1.7 10^{-2}	8.5 [3,medium]	0.3 [2×3,loose]	11 [9,etcut]	4 [2×3,loose]
end	∼1.9 10^{2}	42 [15,medium]	2.5 [2×10,loose]	9 [40,loose]	2.5 [2×15,loose]
2011					
begin	∼3.5 10^{2}	18 [15,tight]	0.4 [2×12,med]	6 [40,tight]	5.5 [2×15,loose]
end	∼3.1 10^{3}	60 [22,med1+ had. iso.]	3 [2×12,med+ had. iso.]	10 [80,loose]	6 [2×20]
2012					
begin	∼2.5 10^{3}	32 [24,iso.,med1+ had. iso.]	1.6 [24,med1+ 7,med1]	4 [120,loose]	5.5 [35,loose+ 25,loose]
end	∼6.6 10^{3}	90 [24,iso.,med1+ had. iso.]	4 [24,med1+ 7,med1]	10 [120,loose]	9 [35,loose+ 30,loose]
2013					
p-Pb	∼4.5 10^{-1}	16 [9,tight1]	1.6 [9,tight1+ 5,E_{T} only]	14 [20,loose]	3 [2×10,loose]
2015 (Run 2)					
end	∼4.3 10^{3}	55 [24,LH med, iso.]	12 [2× 12,loose]	18 [120,loose]	8.5 [35,loose+ 25,loose]

Table 5.4: Evolution of the rates and selection criteria of the lowest p_{T} primary (unprescalled) e/γ triggers for selected example data taking runs during Run 1 and beginning of Run 2. A peak luminosity and the rates in Hz at the beginning of the run are quoted. Threshold in unit of GeV and selection working points are described in square brackets. Typically, out of many triggers with slightly varying selections the simples are presented.

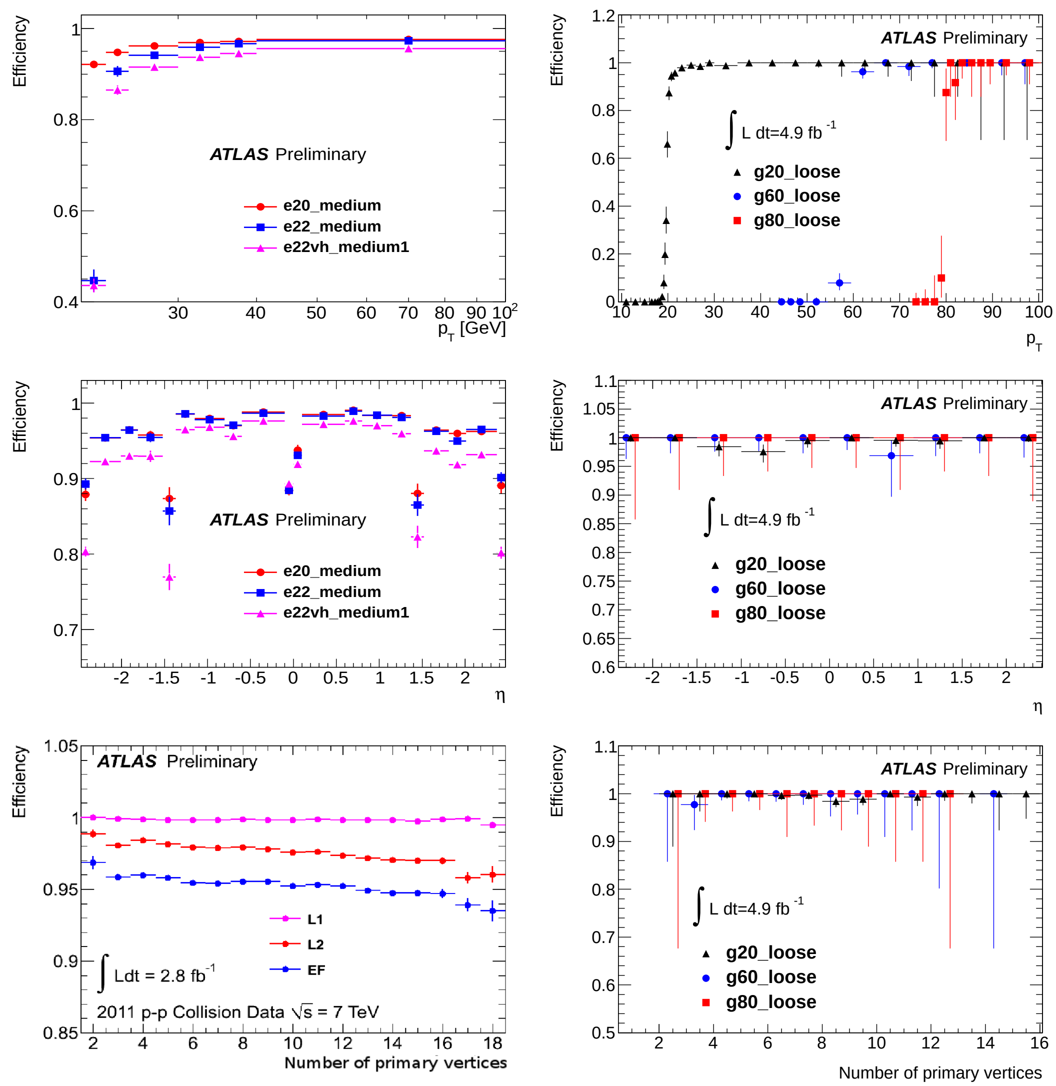

FIGURE 5.14: Efficiencies of the electron (left) and photon (right) triggers in 2011 as a function of transverse momenta, p_{T}, and pseudorapidity, η, of the offline reconstructed objects and number of vertices (pile-up) [75].

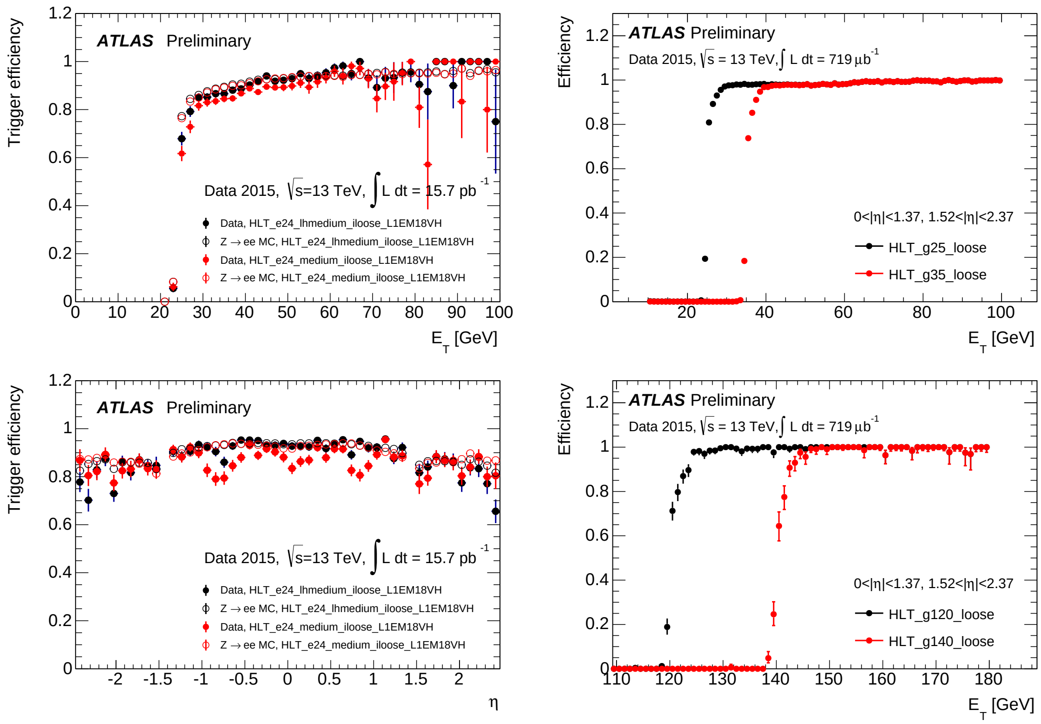

FIGURE 5.15: Efficiencies of the electron (left) and photon (right) triggers in 2015 as a function of transverse energy, E_{T}, and pseudorapidity, η, of the offline reconstructed objects and number of vertices (pile-up) [70, 69].

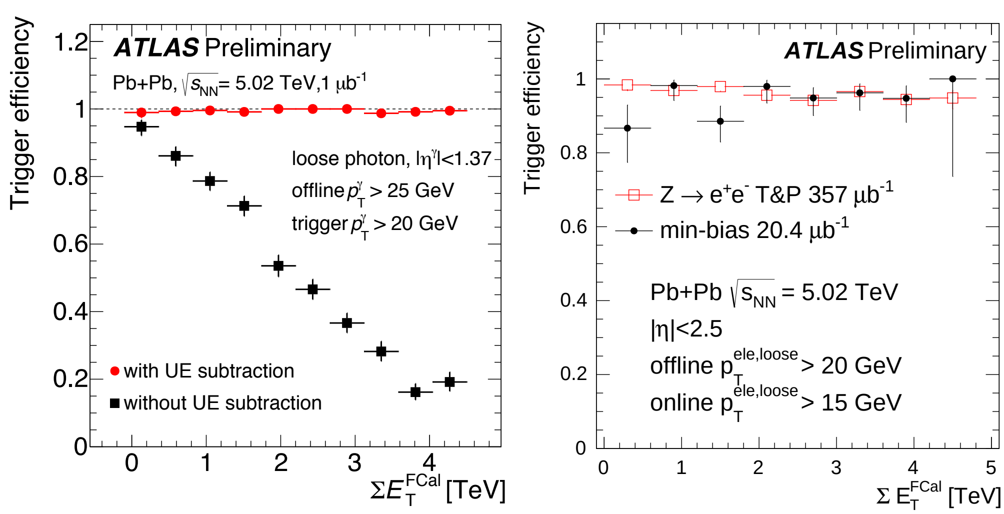

FIGURE 5.16: The efficiency of the primary photon (left) and electron (right) triggers during 2015 Pb-Pb data taking run as a function of the collision centrality estimated with the transverse energy deposition in the Forward calorimeter, $\Sigma E_{\mathrm{T}}^{\mathrm{FCal}}$. On the left the p-p photon trigger configuration (without UE subtraction) and Pb-Pb configuration (with UE subtraction) are shown while in case of electron only the Pb-Pb trigger is presented [69].

5.2.2 Muon triggers

Muons are important signatures of many processes of interest for ATLAS experiment. Similarly to electrons, muons very clean signature allows for triggering on EW bossons. In harsh conditions alike in Pb-Pb collisions muons provide superior quality signal compared to electrons as they remain unaffected by the underlying event [77]. Efficient triggering on muons in wide p_{T} range allows for a selection of weak bosons via $W \to \mu\nu$, $Z \to \mu\mu$ or indirectly $Z \to \tau\tau$ [78], top quark [79] decays as well as direct searches for a new physics [80]. At low values of p_{T} efficient μ trigger allows for selection of events sample used for B-physics sector, also at high luminosities when the di-object trigger were used to trigger for events with J/ψ or Υ. Pure muon detector signature, namely the track outside of the hermetic calorimeter, allows for measurements with no misidentification biases. Similarly to electrons muons are used to cross calibrate the detector and reconstruction procedures. For example decays into two muons $J/\psi \to \mu\mu$ and $Z \to \mu\mu$ benchmark processes can be used to validate muon p_{T} calibration [81, 82] while $W \to \mu\nu$ can be used to estimate biases and resolution of missing transverse energy measurement [83].

Muon triggers at HLT are based on tracking in MS and ID. Since the L1 RoIs contain an information on the muon track position in the MS the reconstruction starts from finding track segment in the spectrometer. With p_{T} estimate, track coordinates a back extrapolation to the luminous region results in an estimate of possible location of the muon track in ID, an ID RoI of muon. The tracks reconstructed in ID and MS are then combined in order to obtain a improved p_{T} measurement, better η, ϕ and impact parameters estimates as exemplified in Fig. 5.17 and 5.18.

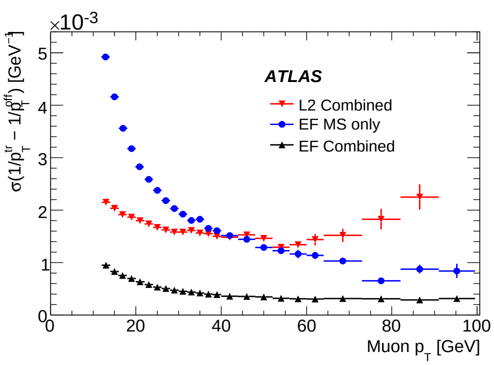

FIGURE 5.17: Resolution of inverse of transverse momenta, $1/p_{\mathrm{T}}$, residuals as a function of offline muon p_{T} for estimates at the L2 and standalone and combined with ID EF [36].

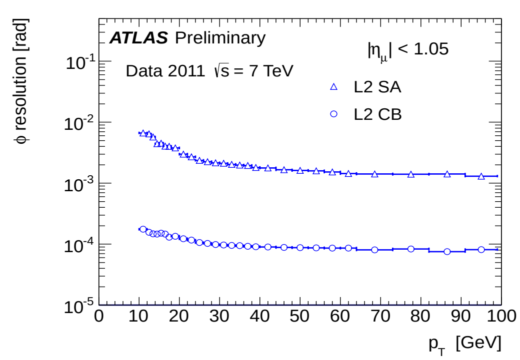

FIGURE 5.18: Resolution of ϕ estimate as a function offline muon track p_{T} for standalone and combined L2 reconstruction of muons [35].

The MS tracking at L2, also referred to as standalone MS tracking, opens a window around the L1 RoI and uses the hits from the trigger detectors (RPC,TGC) and precise detectors (MDT, CSC[4]). The track is then fit using the hit information and a p_{T} estimate is produced from a predefined look-up table. The use of look-up tables results in worse resolution at higher muon p_{T} (Fig. 5.17) however it does not cause any inefficiency as the thresholds used for muon triggers are well below that region. The L2 ID tracking algorithm of the same type as for electrons is used for muons, namely IdScan. The two algorithms are followed by the

[4]The CSC information was used from 2011 on.

one combing two tracks. The absence of track in the ID allows to reject muons from cosmic radiation and from in-flight decays. If compatible ID and MS tracks are found the combined track is formed with the p_{T} estimate being a weighted average of p_{T} estimated from fits in MS only and ID algorithm. Functionally identical set of algorithms is used at the EF but the offline tools are used for both MS and ID tracking. In addition to the *outside-in* approach an alternative arrangement is available at the EF in which tracks are reconstructed firstly in the ID and then extrapolated to the MS.

The di-muon triggers for B-physics measurements use two approaches. Either with two muons identified at L1 (topological trigger) or with one muon identified at L1 and the second at the HLT. The drawback of the former approach is that the low L1 efficiency in the barrel region reduces the overall efficiency for this trigger, particularly when both muons fall in the barrel region. In the second approach, upon confirmation of the muons at the HLT, search is made for the second mun in a wide cone ($\Delta\eta \times \Delta\phi = 1.5 \times 1.5$) around the L1 RoI position or even in entire detector. First tracks are reconstructed in the ID and then extrapolated to the MS in order to form μ candidates. As a consequence the L1 inefficiency imparts only one of the muons. On the other hand that approach is significantly more resource consuming due to the larger detector volume processed. In each case the invariant mass of the di-μ system is calculated and cuts applied to select values within a window of J/ψ or Υ masses [84]. The HLT provides support for forming the combinations used in both approaches at the stage where the mass cuts are applied (see section: 4.2.3).

In order to reduce the p_{T} threshold of the single muon triggers an isolation requirement may be applied. Calorimetric and track based information is used in the muon isolation algorithm. In the calorimeter, the energy in annuli of $0.07 < \Delta R < 0.2$ and $0.1 < \Delta R < 0.2$ in electromagnetic and hadronic calorimeter respectively and the sum of track p_{T} in a cone of $R = 0.2$ are used to assess isolation. The muon track is excluded from the sum and a minimum track p_{T} of 1 GeV is required. The energy and p_{T} sums as well as the ratios $\sum E_{\mathrm{T}}/p_{\mathrm{T}}(\mu)$ and $\sum p_{\mathrm{T}}/p_{\mathrm{T}}(\mu)$ are used to define the isolation criteria. After optimisation it was found that the condition $\sum p_{\mathrm{T}}/p_{\mathrm{T}}(\mu) < 0.12$ retains the full efficiency for muons from Z decays while rejecting a significant fraction of the background from heavy flavour, pion and kaon decays [35, 36].

Muon trigger performance

The efficiency of the inclusive muon trigger is studied with data using a tag-and-probe method exploiting the J/ψ and Z decays into $\mu^+\mu^-$ pairs. At very high p_{T} muons from top quark decays and $W+$ jet (with W decaying into μ) are used. Typically efficiency is measured as a function of p_{T}. In addition, due to the RPC detector non-uniformities, η, ϕ maps are made.

The efficiency of the L1 muon trigger was discussed in section 2.6. The efficiencies are different in the forward region covered by the TGC and barrel region covered by the RPC due to different geometrical acceptances that reach ∼99% and ∼80% respectively. In addition a reduction in both the RPC and TGC muon efficiency is observed for high-p_{T} thresholds (that are a coincidence of signals from three layers) due to detector inefficiencies and because the additional RPC layer used in the coincidence has an even smaller acceptance.

In the evaluation of the HLT performance on early data, the events used for the measurement were required to contain a vertex with at least three tracks. Good quality reconstructed muons formed from ID and MS track segments were used as the offline reference. The muons were required to have opposite sign charge and the invariant mass of the pair was required to lie within a window around the respective resonance masses. The efficiencies obtained in 2010 are shown in Fig. 5.19. The low-p_{T} 6 GeV trigger efficiency turn-on is rapid and the plateau for barrel and endcap regions are as expected from geometrical constraints. The agreement

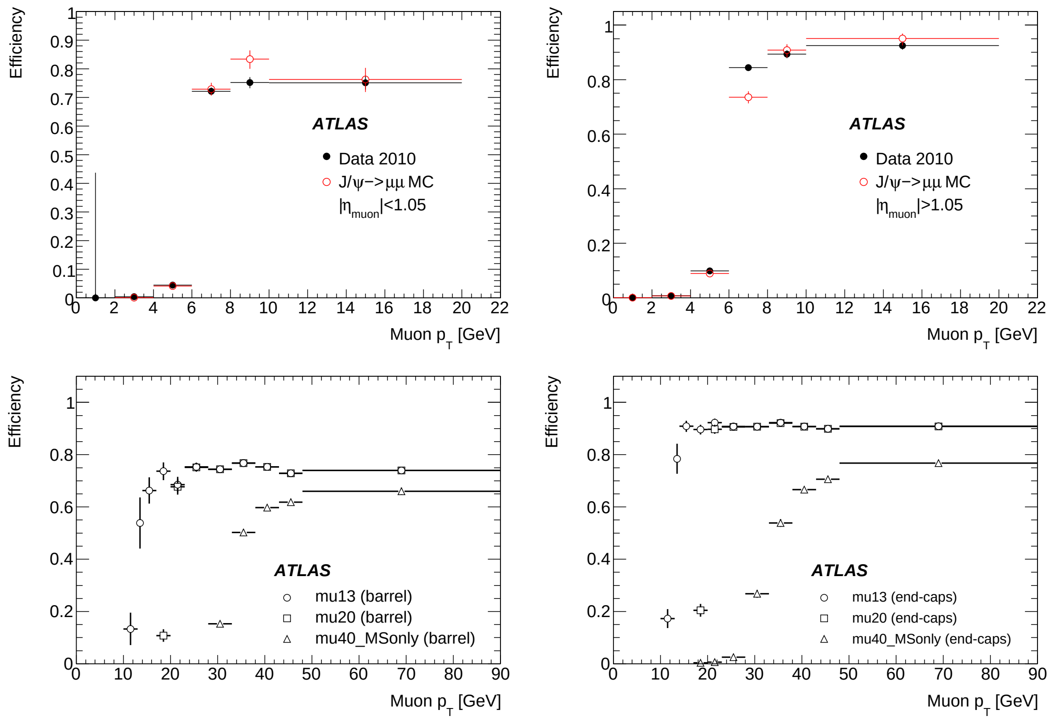

FIGURE 5.19: Efficiency of the muon triggers (L1 + L2 + EF) in 2010 for 6 GeV (top) and 13 and 20 GeV (bottom) triggers. The barrel (left) and endcap (right) are presented separately [18].

in modeling the efficiency in MC is also satisfactory. The high-p_{T} thresholds have a similar efficiency. The drop in efficiency at the very high p_{T} is due to the lack of statistics from $Z \to \mu\mu$ decays and that the wide bin has dominant contribution from muons at the lower end of the p_{T} spectra where the plateau efficiency is not yet attained.

The efficiency of the di-muon trigger was measured using a Bayesian approach and is given by

$$\epsilon_{2\mu} = \epsilon_{1\mu}\epsilon_{2\mu|1\mu}/\epsilon_{1\mu|2\mu}.$$

The parameter $\epsilon_{2\mu|1\mu}$ and $\epsilon_{1\mu|2\mu}$ are the conditional probabilities that the di-muon trigger selected an event under the condition that a single muon trigger did so and vice versa respectively. For the topological triggers the second term is, by definition, equal to unity while in single muon seeded trigger it is close to one. All terms can be measured from data. As shown in Fig. 5.20 the topological trigger chains (EF_2mu*) exhibit a lower efficiency as compared to the ones seeded from single L1 RoI (EF_mu*) as expected.

The enhancement in the data collected in of the regions of the di-muon resonances in mass spectra thanks to these triggers is shown in Fig. 5.21. The dataset collected using dedicated topological and single L1 seeded di-muon triggers is an order of magnitude larger than that which would have been recorded using single muon trigger alone due to the prescaling incurred by recording rate limitations.

The primary unprescaled muon trigger in 2011 was *mu18_medium*. Initially, this trigger was seeded by a L1 MU10 threshold (2 stations coincidence in the barrel and 3 in endcap) until mid 2011 when luminosity reached $\mathcal{L} \sim 2 \cdot 10^{33} cm^{-2}\mathrm{s}^{-1}$ and the MU11 L1 threshold was

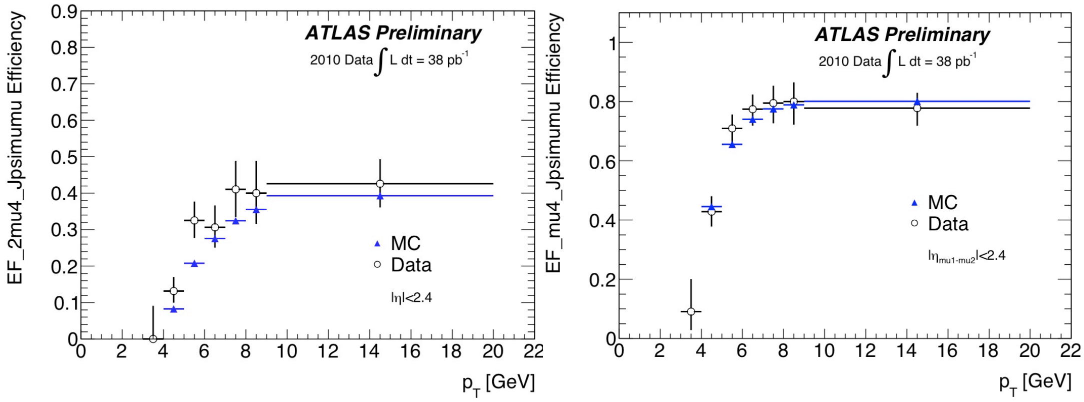

FIGURE 5.20: Efficiency of the di-muon J/Ψ as a function of the higher of the two muons p_{T} for topological trigger (left) and trigger relying on HLT for finding the second muon (right) [84].

used. The MU11 threshold is a high-p_{T} L1 threshold in barrel and end-caps. The efficiency of this trigger is shown in Fig. 5.23. A steep turn-on and flat plateau can be observed. A similar performance was observed for the inside-out configuration, however the algorithms in this configuration performed worse in terms of the required compute resources, as shown in 5.22. The primary trigger was found to only slightly degrade in performance with increasing pile-up. A slight discrepancy between MC and data visible in the barrel can be attributed to a missmodeling of the RPC detector efficiencies in the MC. In order to account for this weights that scale the number of events in bins in η and ϕ are established [35].

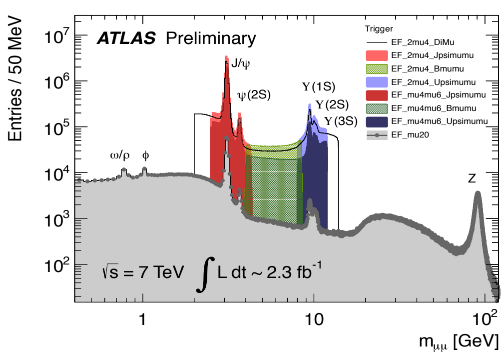

FIGURE 5.21: Mass spectra of oppositely charged muons collected by the various triggers denoted in the legend used in 2011 [85].

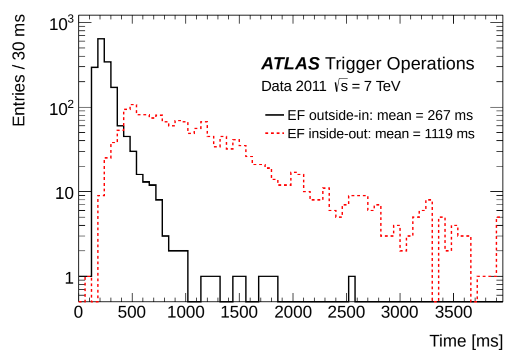

FIGURE 5.22: A comparison of the reconstruction time for the EF outside-in (from MS to ID) and inside-out (from ID to MS) strategies [35].

In 2012, due to the increase of luminosity to up to $\mathcal{L} \simeq 7.7 \cdot 10^{33} cm^{-2} \mathrm{s}^{-1}$ the p_{T} threshold of the primary inclusive muon chain had to be increased. The 4,6, and 10 GeV thresholds used a low-p_{T} configuration in the barrel and a high-p_{T} in the endcaps. A high-p_{T} threshold were 11, 15, 20 GeV. From the 15 GeV L1 item the 24 GeV isolated, 36 GeV un-isolated and MS only 40 GeV chains were seeded. Lower p_{T} thresholds were either prescaled or combined in multi-object triggers such as 2×13 GeV or 18 GeV + 8 GeV where the second 8 GeV threshold is obtained

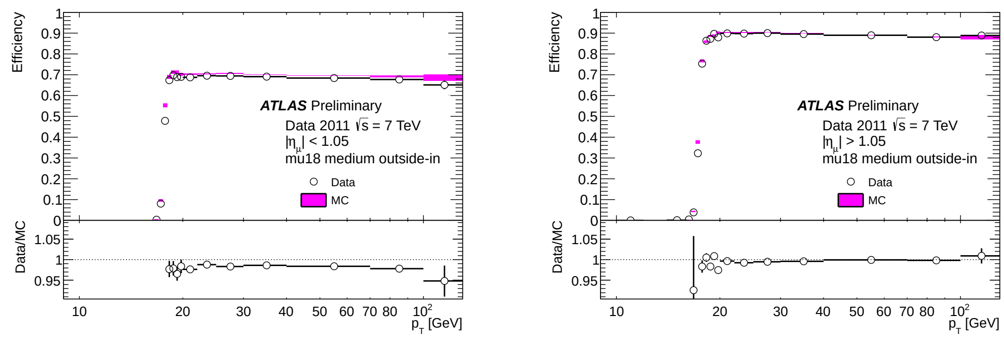

Figure 5.23: Efficiency of the *mu18_ medium* trigger (seeded of MU10) as a function of offline muon p_{T} for barrel (left) and endcap (right) [35]. Efficiency ratio between MS and data is shown at the bottom.

using MS detector scan at the HLT. The efficiencies for the high-p_{T} inclusive triggers are shown in Fig. 5.24. The efficiency rapidly raises near the threshold and remains high and flat at high

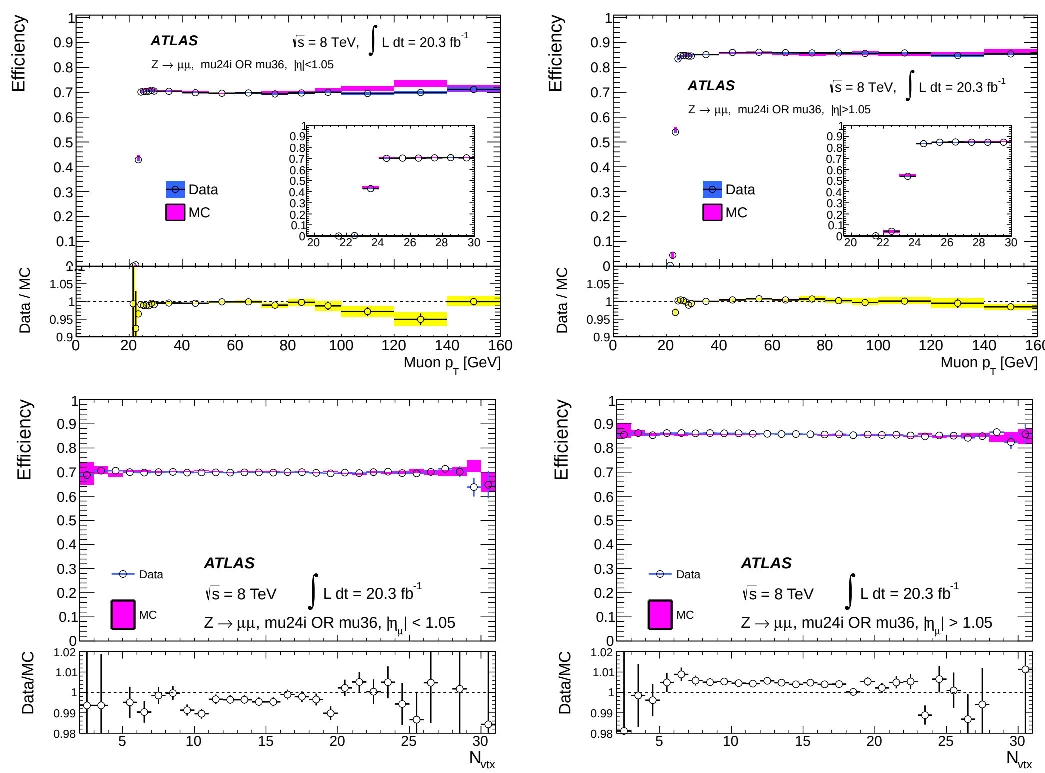

Figure 5.24: (top) Efficiency of the *mu24i* OR *mu36* trigger (seeded of MU15) as a function of offline muon p_{T} for barrel (left) and endcap (right). (bottom) The efficiency as a function of number of vertices in the event (pile-up) for the η regions corresponding to the barrel and endcaps [36].

p_{T}. The isolation cut $\sum p_{\mathrm{T}}/p_{\mathrm{T}}(\mu) < 0.12$ applied in the 24 GeV trigger does not introduce any pile-up influence, either in the barrel or in the endcap. At medium p_{T} the di-muon chains used in the 2012 were based on a single L1 muon trigger with the second muon reconstructed at the HLT. As a result the efficiency for the second muon was nearly 100% as shown in Fig. 5.25 together with the efficiencies for the lower-p_{T} muons.

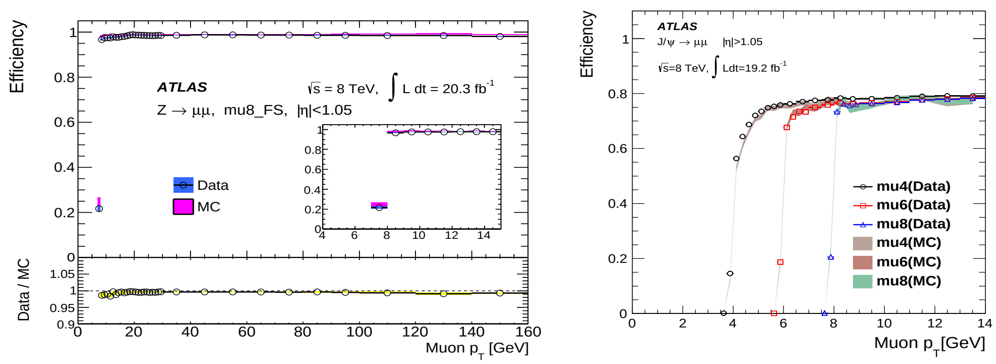

FIGURE 5.25: (left) Efficiency of the *mu8FS* and low-p_{T} (right) threshold chains [36].

In Run 2 additional efficiency drop is expected due to the 3-station coincidence requirement in the forward region needed to limit trigger rate. Besides this an efficiency is not expected to deteriorate in Run 2. The efficiency in the region covered by the RPC is also not expected to change. Both efficiencies are shown as a function of muon p_{T} in Fig. 5.26. For clarity the excellent HLT efficiency relative to the L1 15 GeV seed is also shown. The efficiency of the logical OR of the triggers chains are presented. This strategy is used in analyses as together the two chains provide good performance over a wide range of p_{T}. At lower p_{T} the isolation criteria is required while an inclusive trigger turning on at 50 GeV is used to efficiently cover higher p_{T}.

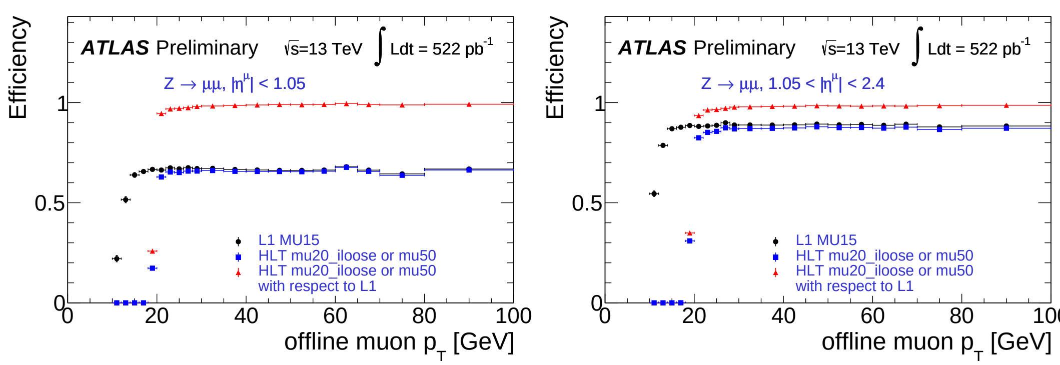

FIGURE 5.26: Overall and relative efficiencies of the muon barrel (left) and endcap (right) triggers as a function offline muon p_{T} in 2015 [37].

The muon trigger rates scale linearly with luminosity and no sign of pile-up dependence is observed at the L1 or HLT as shown in Fig. 5.27. At low luminosities (2010) the effect of an increase in the of number of colliding bunches can be seen as jumps in the line fitted to the measured rates. This is because muons from cosmic rays that reach the MS are sampled for a longer fractions of time as the number of bunch-crossings increases. At high luminosities a slight nonlinearity can be observed for the L1_3MU4 trigger caused mostly by the rise in the number of fake muons in the forward region. At the HLT the isolation requirement included in the 24 GeV trigger does not introduce any pile-up dependence.

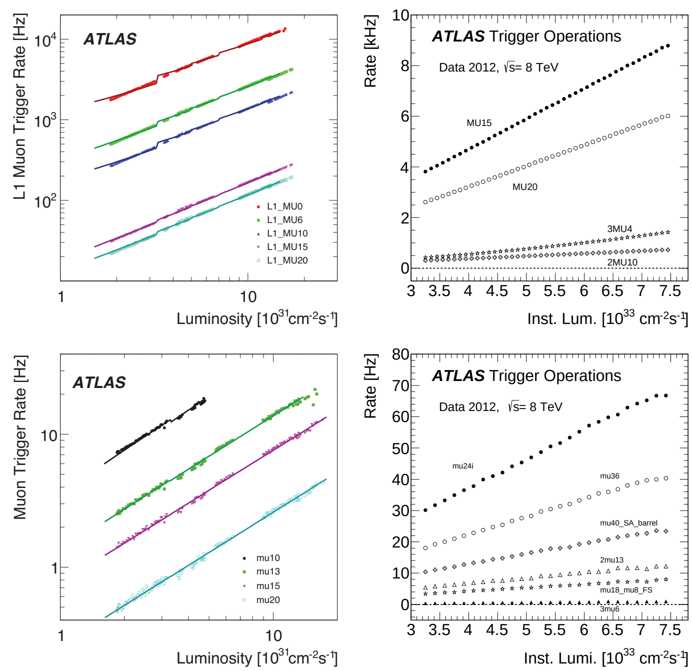

Figure 5.27: Rates of the muon triggers as a function luminosity for L1 (top) and HLT (bottom) in two luminosity and thus pile-up regimes in 2010 (left) and 2012 (right) [18, 36].

Table 5.5 shows an overview of the example primary muon triggers throughout Run 1 and at the beginning of Run 2. Also the triggers used in the Pb-Pb data taking periods are shown.

	$\mathcal{L}$ $\times 10^{30}$ $[cm^{-2}s^{-1}]$	single-μ [Hz]	di-μ [Hz]	B-phys [Hz]
2010				
begin	~1.7 10^{-2}	1.5 [4]	0.02 [2×4 MSonly]	0.01 [4 (L1 MU0, HLT FS)]
end	~1.9 10^{2}	25 [13,tight]	30 [2×6]	1 [4+6, J/ψ mass cut†]
2011				
begin	~3.5 10^{2}	5 [15,isol]	0.6 [2× 10,loose]	2.5 [4×2, J/ψ mass cut†]
end	~3.1 10^{3}	110 [18,med]	3.5 [15+10 tight EFFS]	20 [4T×2, J/ψ mass cut†*]
Pb-Pb	~4.5 10^{-4}	25 [4,loose]	-	-
2012				
begin	~2.5 10^{3}	18 [20,isol]	4 [18,tight+ 8,EFFS]	12 [6+4, J/ψ mass cut†]
begin	~6.6 10^{3}	60 [24,tight iso.]	10 [18,tight+ 8,EFFS]	35 [6+4, J/ψ mass cut†]
2013				
p-Pb	~4.5 10^{-1}	6 [10]	4 [4+ 2,EFFS]	-
2015 (Run 2)				
end	~4.3 10^{3}	53 [24,med iso.]	17 [18+ 8,EFFS]	4.5 [18+ MSonly muon and J/ψ mass cut†]

Table 5.5: The rates and selection criteria for μ triggers. Rest as in 5.4. †While there was an array of triggers used for J/ψ selection these are chosen to indicate typical rates devoted to this type of triggers. *The suffix "T" in the trigger name denotes higher threshold at L1 (4GeV as compared to earlier used requirement of hits without a muon track characteristics).

5.2.3 Tau triggers

The ability to select events containing a τ lepton allows measurements in the SM sector at the EW scale [86]. The Higgs boson coupling to the τ is an example of such an important measurement [87]. Besides that, the final states with τ can be used in searches beyond the standard model [88, 89, 90].

The most straight-forward τ lepton decay channel for triggering is $\tau \rightarrow l\bar{\nu}_l\nu_\tau$ (l is either μ or e) which occurs in about 35% of cases while the most abundant decay channel into hadrons (mostly pions) has a BR of 65%. The τ trigger aims to achieve high efficiency and purity for the hadronic channel as the lepton decay modes are well covered by the low-p_{T} muon and electron triggers. Characteristics of the hadronic τ decay in the detector are a narrow hadronic cluster and significantly fewer charged particle tracks as compared to background QCD jets, as shown schematically in Fig. 5.28. The τ lepton decays into a single π hadron in $\sim 45\%$ (called one-prong) and the decay to a higher number of charged hadrons in $\sim 12\%$ of cases (multi-prong) have different characteristics in terms of possible efficiency and rejection power.

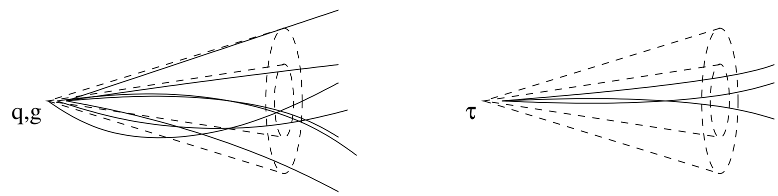

FIGURE 5.28: Schematic illustration of the QCD jet (left) and τ (right) decay topology exploited in the hadronic τ triggers. The core (inner) cone and outer (isolation) cones are shown with dashed lines.

Due to the copious background and a limited discriminating power, the τ hadronic trigger is usually combined with other signatures allowing to use of lower p_{T} thresholds for the τ while at the same time remaining within recording rate limitations. A combination of hadronic τ and missing transverse energy, $E_{\mathrm{T}}^{\mathrm{miss}}$, muon and electron triggers are used to construct very targeted triggers [88, 90].

The reconstruction at L2 begins, similarly to electrons, with the calorimeter clustering algorithm after which the E_{T} and shower shape quantities are used to reject a fraction of predominantly background candidates. In the L2 calorimeter algorithm the seed cell with the highest E_{T} is found first and the cluster is built around that cell. The RoI at the L2 used to search for the cluster had dimension $\Delta\eta \times \Delta\phi = 0.6 \times 0.6$ in 2010 and was enlarged to 0.8×0.8 from 2011 on. For candidates passing calorimeter-based selection the tracking is performed and track-based and combined track-cluster discriminants are used in τ hypothesis verification. At the EF the offline topological clustering algorithm, adapted to the HLT environment, is used in a region $\Delta\eta \times \Delta\phi = 0.8 \times 0.8$. The cell with the highest signal/noise ratio becomes a seed, i.e. first cell of the cluster, and the cluster is enlarged by adding geometrically neighbouring cells until all cells of the signal to noise RM ratio >2 are incorporated. Tracking is then performed using offline-type algorithms for candidates passing. The Hypo algorithms in the last stage reject candidates based on the same discriminating variables as used offline.

Since tracking in relatively large detector volume is required in the τ trigger, it is configured to use only silicon detector information (no TRT information). Also, the track minimum p_{T} requirement is set as high at 1.5 GeV. In Run-2 a sophisticated tracking strategy was deployed in order to further reduce the processing time. First tracking is performed in a region that is narrow in a $\Delta\phi$ and $\Delta\eta$ but has a wide span along z. In the second stage, tracking is performed in larger area in both η and ϕ but with a narrow range in z around the leading track.This strategy allows for saving reconstruction time due to the reduced detector volume in which the processing is performed, as shown in 5.29.

The main discriminating variables used in the L2 and EF parts of the τ trigger are similar and rely on the cluster sizes and track multiplicities. They are summarised in the table 5.6.

The distribution of the shape parameters and track multiplicity in the core cone are shown in 5.30. The discriminating power is in general higher for the one-prong as compared to the multi-prong decays. Therefore distinct selection cuts are applied for these two modes.

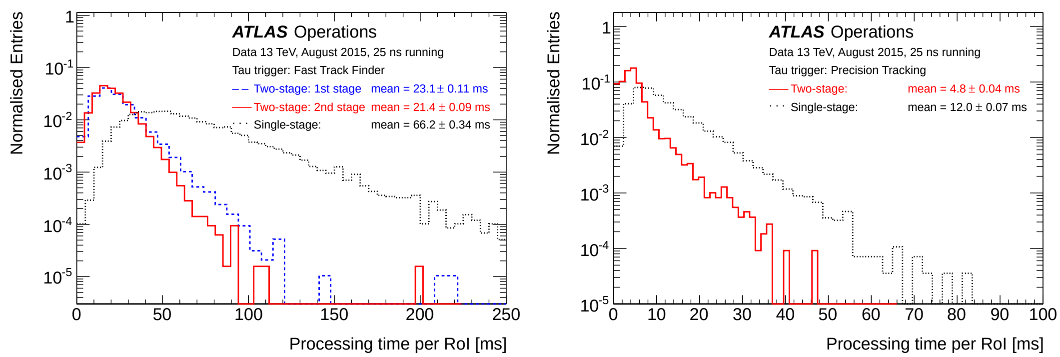

FIGURE 5.29: Comparison of the fast (left) and precision (right) tracking algorithm timings for τ triggers in the single stage strategy from Run 1 and the two stage approach used in Run 2 (see text) [91].

The optimisation of the parameters used for the discrimination of signal and background and the optimal combination of these parameters is the subject of contining improvements. For instance, as shown in table 5.6, the definition of R_{EM} at L2 has been optimised ($(\Delta R)^2 \rightarrow \Delta R$) in 2011 to obtain a better match between offline and EF selection quantities and therefore an optimal efficiency of the trigger with respect to offline. In Run 1 the sequence of cuts applied starts with a cut on the L2 R_{RM} at a value optimised for multi-prong decays. At a later stage, after the tracking information is available, a cut more optimal for one-prong decays is applied if only one track was found [94, 95]. These cuts are E_{T} dependent and optimised for three working points *loose*, *medium* and *tight* which correspond to a different balance between background rejection and efficiency. For the data taking at end of Run 1 (in 2012) the pile-up dependence was mitigated by reducing size of the region used in L2 clustering to $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$ where search for the cone of $R < 0.2$ was performed. Tracks taken into account were required to share the origin with the leading p_{T} track ($|\Delta z| < 2mm$). Similar fine tuning was applied at the EF. The cuts on discriminating variables were replaced with the use of multivariate (MV) techniques. The variables were combined using Boosted Decision Tree (BDT) and Log-Likelihood (LL) methods in order to obtain one discriminant [96]. The performance of the MV selection for one- and multi-prong τ at the EF is shown in Fig. 5.31 and demonstrates the superiority of the more advanced BDT discriminant in the multi-prong decay case [97].

Tau trigger performance

The τ hadronic triggers started to be commissioned on the first data in 2010. The trigger chains were kept in monitoring mode, initially, in order to verify reconstruction robustness and performance. At $\mathcal{L} \sim 1.2 \times 10^{29} cm^{-2}\mathrm{s}^{-1}$ the τ triggers were enabled and provided rejection. Throughout 2010 the primary inclusive trigger threshold was increased progressively from 20 to 50 and then 84 GeV while lower thresholds were used in triggers combining τ's with other signatures [18].

For the early data the trigger efficiency was studied using a QCD di-jet samples due to the low statistics for the τ signal in the data. As a result, the efficiency with respect to the offline reference only indicates the degree of agreement between offline identification and online selection. Figure 5.32 shows a comparison of the efficiency between data and MC for the background sample as well as a comparison for the signal MC. Both agree at the plateau while the turn on is less steep for the di-jet sample due to the L1 trigger used to collect the di-jet

variable	definition	2010 L2	2010 EF	2011/2012 L2	2011/2012 EF
EM radius $R_{L2,EM}$	$\frac{\sum E\cdot(\Delta R)^2}{\sum E}$ where sum runs over second layer of ECal	✓		✓	
EM radius $R_{EC2,EM}$	$\frac{\sum E_{\mathrm{T}}\cdot(\Delta R)}{\sum E}$ where sum runs over second layer of ECal		✓		
EM radius R_{EM}	-||- where sum runs over first three ECal layers			✓	✓
Track multiplicity	Number of tracks in cone of $R < 0.1$	✓	✓		
p_{T} of the leading track R_{EM}	$\frac{\sum E_{\mathrm{T}}\cdot\Delta R}{\sum E}$ where sum runs over second layer of ECal		✓		
Tracks invariant mass m_{tracks}	$\sqrt{(\sum E)^2 - (\sum \vec{p})^2}$ where the sum runs over the tracks		✓		
Core isolation	$\sum_{\text{iso}} p_{\mathrm{T}} / \sum_{\text{core}} p_{\mathrm{T}}$ where "core" sum run over the tracks within the narrow cone (R<0.1) and "iso" sum over tracks in with $0.1 < R < 0.3$ from the leading track	✓	✓	✓	
Track radius R_{Track}	$\frac{\sum p_{\mathrm{T}}\cdot\Delta R}{\sum p_{\mathrm{T}}}$ where the ΔR is calculated from τ candidate direction and sum runs over the tracks within $\Delta R < 0.2$				✓
Leading track momentum fraction f_{track}	$\frac{p_{\mathrm{T}}^{\text{lead.-track}}}{p_{\mathrm{T}}^{\tau}}$				✓

Table 5.6: Description of τ trigger discriminating variables used through Run 1 [92, 93].

sample.

The pile-up dependence was studied already in 2010 as τ trigger relies on clustering and tracking in a relatively large detector volume. A comparison of the distribution of the discriminating variables for two data taking periods differing by three orders of magnitude in luminosity is shown in Fig. 5.33. Also shown is the scaling of rates with luminosity for both the inclusive and combined triggers. No significant pile-up dependence was observed in 2010.

In 2011, when luminosities increased to $\mathcal{L} \approx 3.5 \times 10^{33} cm^{-2}\mathrm{s}^{-1}$, the hadronic τ triggers had to be adapted due to the pile-up effects and the thresholds had to be increased in to limit the increase in trigger rates. The pile-up dependence of the di-τ triggers is shown in

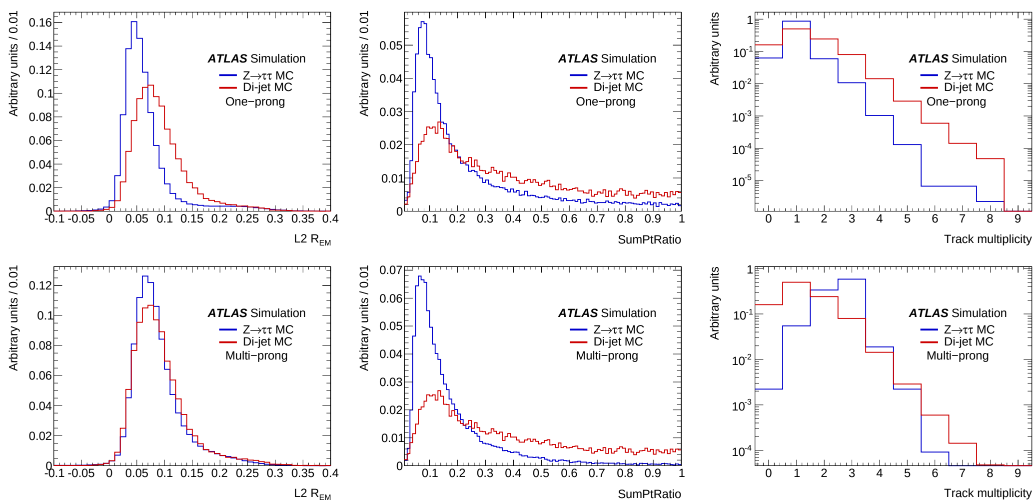

Figure 5.30: The τ discriminating variables for signal (blue) and background (red) candidates. Sample of RoIs when τ decays into one hadron (one-prong) (top) and into multiple hadrons (multi-prong) (bottom) are separated while the di-jet background shapes are the same. From left to right the R_{EM}, $\sum p_{\mathrm{T}}^{\mathrm{iso}} / \sum p_{\mathrm{T}}^{\mathrm{core}}$, and track multiplicity in cone of $R = 0.1$ around the leading track are shown [92].

Fig. 5.34. The nonlinear behaviour (flattening at higher luminosity values) was caused by a significant modification of shower shape variables on which the cuts were applied at L2. In order to reduce the degradation of performance due to the pile-up, in 2012 the cone size in which clustering was performed was reduced to $R = 0.2$ and the tracks used in tagging were required to originate from the same vertex, as mentioned earlier. In addition the τ identification started to use multivariate techniques which included some pile-up robust variables [96, 97, 98]. The improvement in the linearity of the scaling of rates with luminosity can be seen in Fig. 5.34 where a linear scaling of di-τ trigger rates was recovered in 2012.

Another challenge addressed in an τ trigger was related to the use of hadronic and electromagnetic isolation criteria. These were used to reduce the rates at low p_{T} however at high p_{T} a leakage of the τ cluster energy added to energy depositions due to the pile-up caused a loss of isolation criteria efficiency. Because of this the high-p_{T} triggers were seeded from L1 items which did not require isolation while at low p_{T} isolation was required. Consequently in the inclusive analyses both of these triggers were used [92]. After the upgrades for Run 2 it is possible to calculate L1 isolation relative to the energy in the core of the cone and therefore the isolation-based triggers retain high efficiency at high p_{T}, as shown in Fig. 5.35 [99]. Additionally topological requirements implemented in the L1 topo system will be used at high luminosity. Improvements in the HLT reconstruction such as better cells calibration and the use of offline reconstruction software earlier in the selection chains are also expected also to give improved performance. The expected performance is shown in Fig. 5.35. The efficiency at high p_{T} is significantly better.

In Table 5.7 typical rates are given for the primary τ triggers.

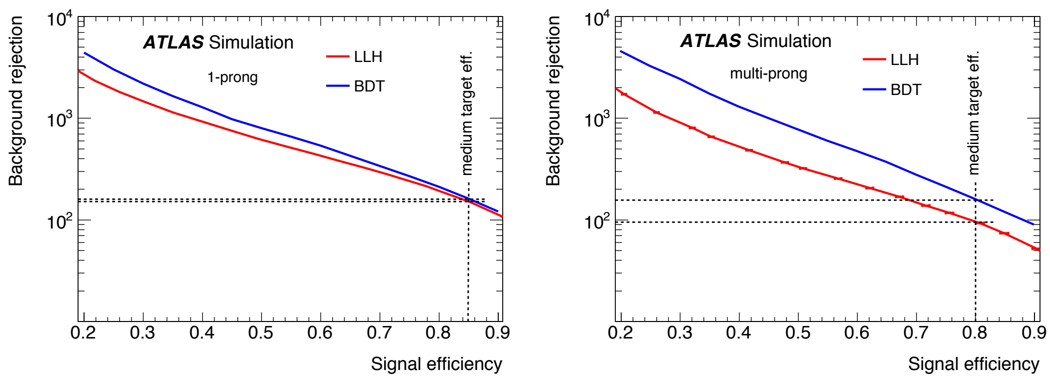

FIGURE 5.31: Performance of the τ Boosted Decision Tree and Log-Likelihood discriminants for a MC ($Z \to \tau\tau$) simulation. The jet background rejection as a function of signal efficiency measure with respect of the offline reconstructed and identified τ decaying in one-prong mode (left) and multi-prong mode (right) [97].

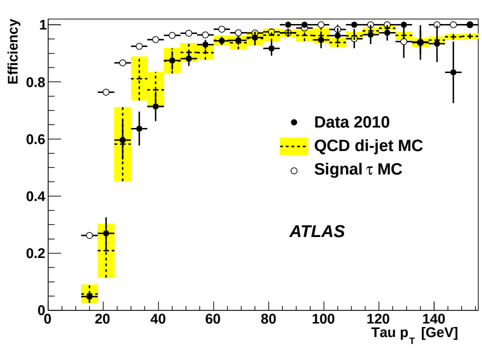

FIGURE 5.32: The τ trigger (tau16_loose) efficiency measured on data and MC for background resembling taus (QCD di-jets) and genuine τs in MC.

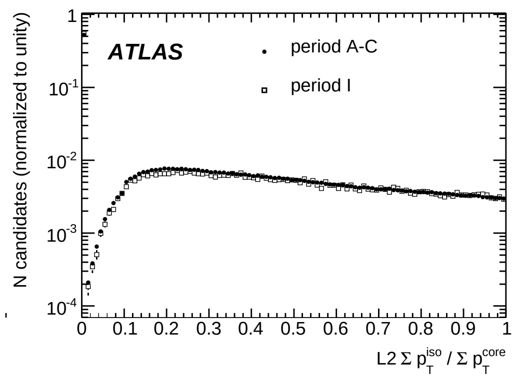

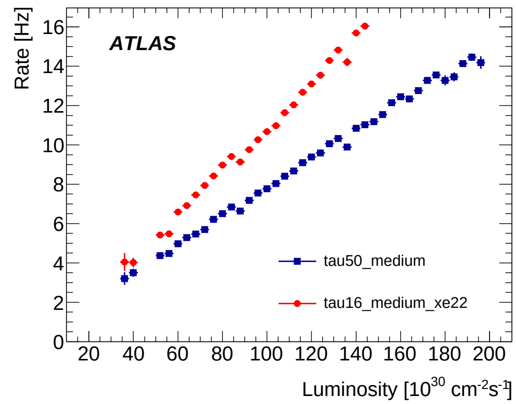

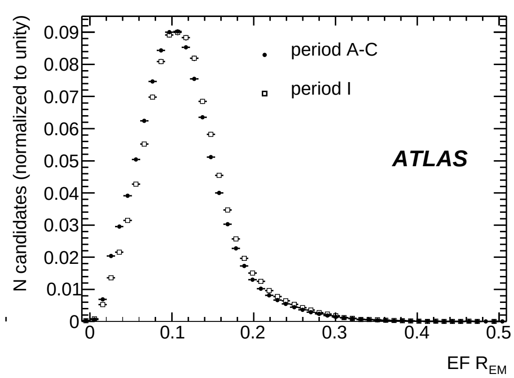

Figure 5.33: (left) The variation of the discriminating variables $\sum p_{\mathrm{T}}^{\mathrm{iso}}/\sum p_{\mathrm{T}}^{\mathrm{core}}$ and R_{EM} with between two periods of data taking with luminosity $\mathcal{L} < 2.4 \times 10^{29} cm^{-2}\mathrm{s}^{-1}$ (labeled as period A-C) and $\mathcal{L} \approx 1.5 - 2.1 \times 10^{32} \mathrm{cm}^{-2}\mathrm{s}^{-1}$ (labeled as period I). (right) The τ trigger rates scaling with instantaneous luminosity in 2010 [18].

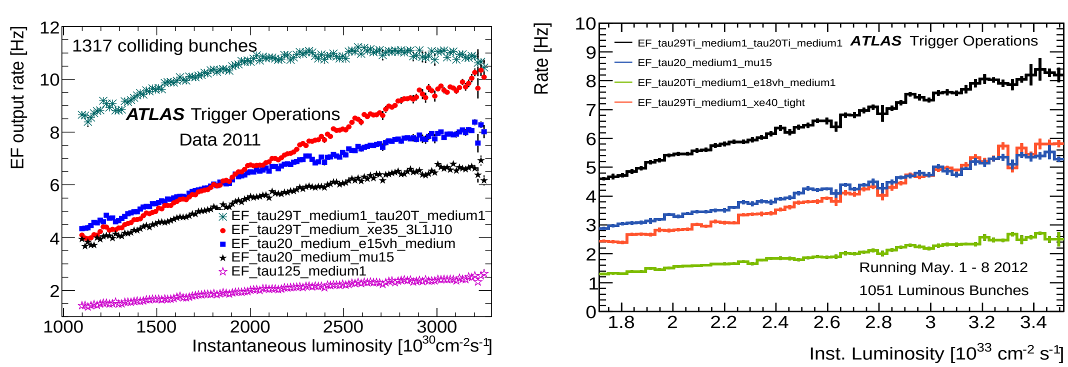

Figure 5.34: (left) As indicated in the legend the di-τ and combined τ trigger rates as a function the instantaneous luminosity [92]. (right) The rates of the τ triggers in 2012 [96, 97, 98]

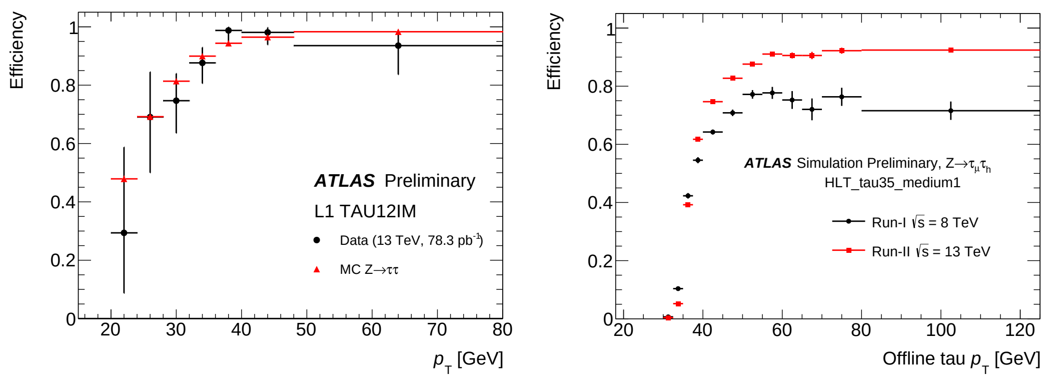

FIGURE 5.35: (left) Efficiency in data and MC of the L1 τ trigger requiring isolation at EM layers after the upgrade for the Run 2. (right) Comparison of the trigger efficiency from Run 1 and expected one Run 2 for non-isolated inclusive τ trigger [99, 91].

	$\mathcal{L}\ \times 10^{30}$ $[cm^{-2}s^{-1}]$	single-τ [Hz]	di-τ [Hz]	$\tau + e$ [Hz]	$\tau + \mu$ [Hz]
2010					
begin	$\sim$1.7 10^{-2}	3.5 [12,loose]	0.05 [2× 12,loose]	0.01 [12,loose+ e 10,loose]	<0.01 [12,loose+ μ 10]
end	$\sim$1.9 10^{2}	7 [84,loose]	3 [2× 38,loose]	7 [12,loose+ e 10,med]	6 [12,loose+ μ 10]
2011					
begin	$\sim$3.5 10^{2}	2.5 [100,med]	2.5 [20+ 29,med1]	1.2 [16,loose+ e15,med]	2.5 [16,med+ μ 10]
end	$\sim$3.1 10^{3}	2.5 [125,med]	11 [20+ 29,med1 isol., tight L1]	8 [20,med+ e15,med had. iso.]	7 [20,med+ μ 15]
2012					
begin	$\sim$2.5 10^{3}	0.8 [125,med1]	6 [29,med1 iso.+ 20,med1 iso.]	2 [20,med1 iso.+ e18,med1 had. iso.]	5 [20,med1+ μ 15]
end	$\sim$6.6 10^{3}	2 [115,med1]	18 [29T,med1 iso.+ 20T,med1 iso.]†	6 [20T,med1 iso.+ e18,med1 had. iso.]†	12 [20,med1+ μ 15]
2015 (Run 2)					
begin	$\sim$4.3 10^{3}	12 [125,med1]	19 [35,loose1+ 25,loose1]	3.5 [25,med1+ e17,LH med iso.]*	9 [25,med1+ μ 14]

Table 5.7: The rates and selection criteria for τ triggers. †The "T" in the threshold denotes modification of the L1 seed to be closer to the HLT threshold. *An array of the triggers with varying electron selection requirements coexist in the menu. Rest as in 5.4.

5.3 Global event triggers

The HLT also supports reconstruction and triggering for cases when the processing is not restricted to the RoIs. A typical case is triggering on global calorimetric quantities like missing or total transverse energy. These kinds of algorithms are bound to consume a significant fraction of HLT (CPU) and DAQ (ROS and network bandwidth) resources and therefore are deferred to late stages in the chains. In Run 1 this algorithms were usually run in the EF and in Run 2 after event-building step. This significantly reduces the ROS bandwidth consumption as the event data is already retrieved and readily available in the HLT processors. To minimise the CPU consumption in processing data from a large detector volume only a minimal reconstruction is performed first and some rejection achieved before the next a more CPU demanding steps. For instance, in some global triggers selecting events depending on the tracks multiplicity the tracking stage is preceded by a significantly simpler algorithm that counts the hits in the ID allowing a rejection before the more expensive tracking steps.

The HLT steering evaluates the FEX algorithms in different manner as compared to a more typical RoI reconstruction. The FEX algorithms which need to reconstruct the entire sub-detector (or event) are called full-scan algorithms and are given more freedom in using the input and in producing the output TEs. That is, the multiplicity of the algorithm output is not fixed to a single TE per invocation. Typically the input TE is connected to the output however this is not a requirement, for instance TEs pointing to RoIs with muons, used in the reconstruction of $E_{\mathrm{T}}^{\mathrm{miss}}$, are not linked to the output. The input to the full-scan FEX algorithms is also treated in a custom manner, including the possibility of ignoring the input. Consequentially, full-scan FEX algorithms can also function without any input TEs, which is impossible in case of RoI FEX algorithms. Last but not least, execution of full-scan FEX algorithms can result in the creation of RoI(s) which are then processed in the same way as those originating from L1.

Also, since the full-scan FEX algorithms perform non-typical data processing the HLT built in monitoring and caching are supported to a reduced extent. In practice the algorithm has to implement custom monitoring and caching exploiting building blocks provided by the HLT steering framework.

In the following the min-bias and missing E_{T} signatures, their implementation and performance will be described. Also specific full-scan approaches of particular interest will be discussed. The interesting use-case of the jet triggers which moved from a RoI based to a full-scan approach will be presented later.

5.3.1 Minimum bias triggers

The minimum bias trigger is designed to record an inclusive sample of p-p (p-Pb or Pb-Pb) collisions irrespectively of the decay channel. The collection of a sample of these events played an essential role in the early running period when the most basic measurement of charged particle multiplicities allowed to the p-p cross section to be constrained to establish tuning of the MC generators with the first LHC data at 0.9 and 7 TeV [100, 101, 102, 103]. During the p-Pb and Pb-Pb data taking the minimum bias trigger plays an essential role and is used to record about half of the dataset [104, 105].

In pp running at 900 GeV and 7 TeV the MBTS detector was used to seed the minimum-bias triggers. This detector delivers signals from all of its 32 pads directly to the CTP which allows the configuration of the various requirements on the number and arrangement of the hits necessary to select the event. In particular the L1_MBTS_N is a requirement for N hits on either side while the notions L1_MBTS_N_N denotes the requirement for a coincidence of N signals from both detector sides. The trigger with a single hit requirement was run as long as the rate permitted.

In order to estimate the efficiency of the MBTS for selecting the inclusive event sample an orthogonal trigger was setup. This trigger, called *mbSpTrk*, relies on the detection of charged particle tracks in the ID in the randomly sampled set of events. The reconstruction sequence for this trigger involved counting the space points[5] at the L2 followed by tracking at the EF. A two stage selection was used based on the number of space points above the typical noise level and requirement of the presence of at least one track at the EF stage. The former requirement allows for the rejection of empty events. In the track counting step tracks are required to originate from the Beam Spot (BS) ($|z_0| < 200mm$) in order to reduce the contribution from beam related backgrounds.

The seed for the mbSpTrk trigger was provided by the Beam Pickup (BPTX) and later during 7 TeV operation by a random trigger in a filled bunch crossing. Efficiencies at 900 GeV and 7 TeV measured as a function of the number of tracks originating from the luminous region are shown in Fig. 5.36 [106]. The MBTS trigger requiring a single hit on either side of the detector reaches 100% efficiency for about three tracks. The MBTS_2 trigger would be next good candidate for the primary minimum bias trigger while the requirement for signals on both sides significantly shifts turn-on region.

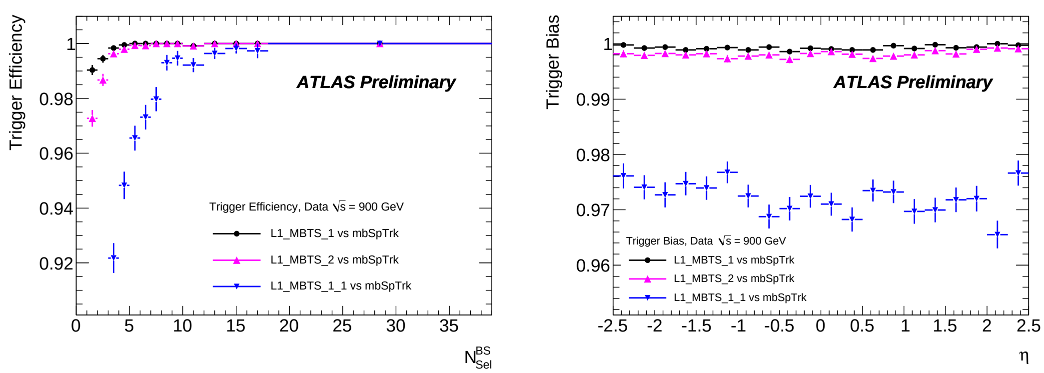

FIGURE 5.36: The efficiency of various MBTS triggers in 2009 as a function of the, offline reconstructed, tracks,N^{BS}_{Sel} (left) and bias as function of η [106].

At 7 TeV the single arm trigger was used initially. The efficiency of the primary L1_MBTS_1 trigger as well as other triggers efficiencies measured using the same supporting mbSpTrk trigger are shown in Fig. 5.37. During 2010 running, similarly to 2009 ,the single and two hits triggers perform very well while the requirement of a two-side coincidence resulted in loss of a significant fraction of the inclusive p-p cross-section at low charged-particles multiplicities. A cross-check of the efficiency of the ID trigger was made using a sample of events selected by the L1_MBTS_2 trigger. As shown in Fig. 5.37 this trigger can be considered to be fully efficient [107].

The minimun-bias triggers in Pb-Pb data taking runs in 2011 played an essential role as about half of the data samples was selected by them as shown in Fig. 5.38. A set of measurements in relativistic heavy-ion physics aim at measurements across the whole spectra of collision centralities. Therefore a very good understanding of any possible trigger bias of the

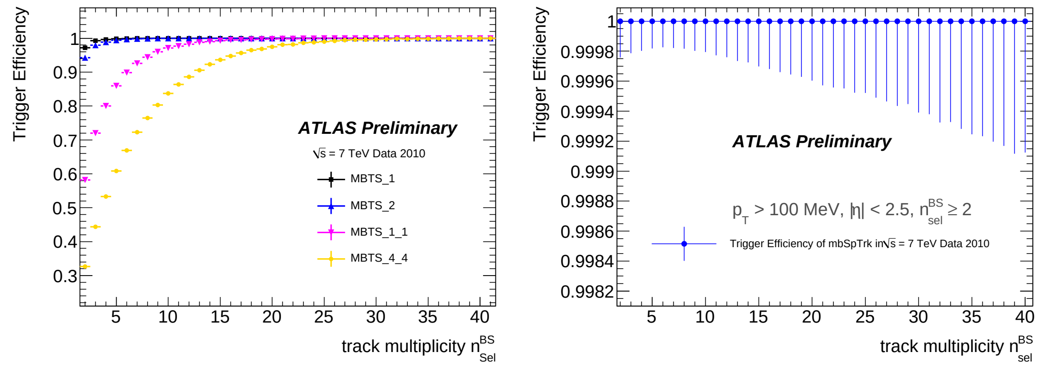

FIGURE 5.37: (left) The efficiency of MBTS triggers in early 2010 as a function of number of, offline reconstructed, tracks, N_{Sel}^{BS}. (right) The efficiency of the ID *mbSpTRk* trigger [107].

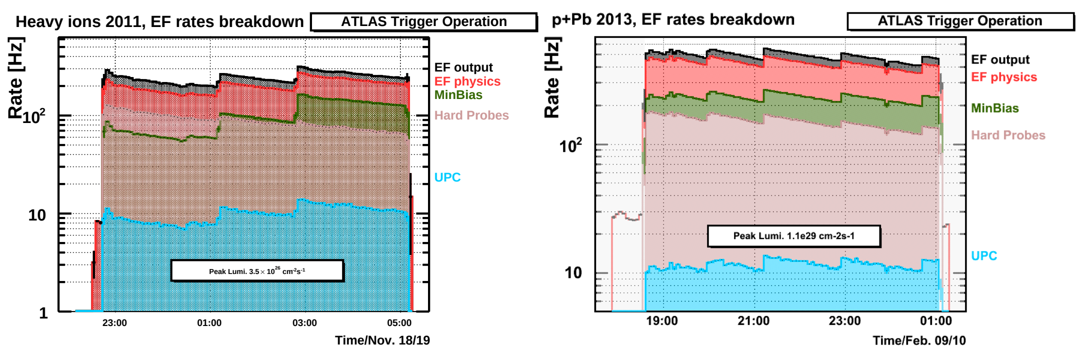

FIGURE 5.38: Rates of the main physics streams during Pb-Pb (left) and p-Pb (right) runs [39].

data sample introduced by the minimum-bias trigger.

For 2011, due to radiation damage of the MBTS [105] the minimum-bias trigger relied on the ZDC[6] and the total calorimeter energy (TE) triggers at L1. The former covers the low multiplicity region of the centralities spectrum while the later is fully efficient and has high purity in central collisions. The relatively low purity of the L1 ZDC trigger led to the requirement for additional filtering at HLT. The HLT selection requiring presence of the few tracks was used in order to reject empty events. Because track reconstruction in more central events, where multiplicity is high, would cause significant operational contraints on the number of events that could be processed or could even result in timeouts in the processing, a unique arrangement of triggers at L1 was used. Namely the total energy trigger $E_{\mathrm{T}} > 50$ GeV without subsequent reconstruction at the HLT was used to select central events, while at low multiplicity the L1 ZDC signals on both sides of the calorimeter in coincidence with the L1 $E_{\mathrm{T}} \leq 50$ GeV total energy trigger is used for peripheral events. The efficiency achieved by the combination of the two triggers is shown in Fig. 5.39. In more peripheral events (low $\Sigma E_{\mathrm{T}}^{\mathrm{FCal}}$ and multiplicity) the total energy trigger is inefficient and the events are saved thanks to the ZDC-based trigger. The latter trigger becomes inefficient at certain value of energy and at this point the other trigger based on total energy picks up the events.

[5] Hits in the ID

[6] The ZDC is a very forward ($|\eta| > 8.3$) calorimeter at $z = \pm 140m$ aimed mainly at the detection of neutrons

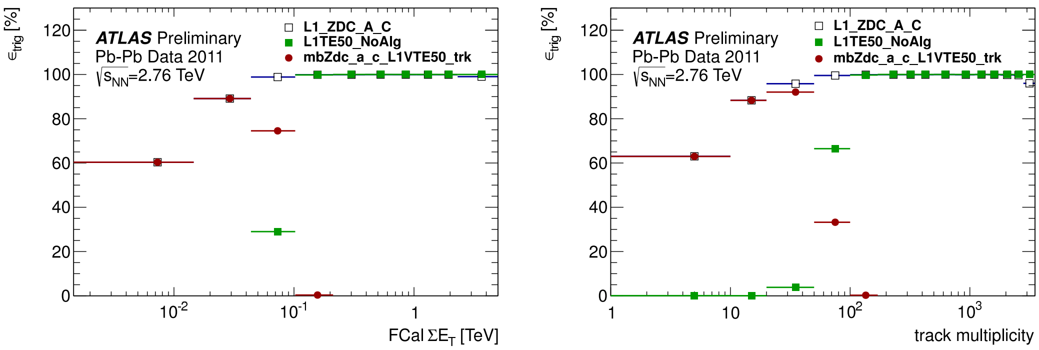

FIGURE 5.39: The efficiency of primary minimum-bias triggers *mbZdc_a_c_L1VTE50_trk* and *L1TE50_NoAlg* in Pb-Pb run in 2011 as a function (left) and (right) [105]

The minimum-bias trigger used in p-Pb data-taking in 2013 was configured in a similar way to that used during Pb-Pb running however a bigger range of triggers were prepared in anticipation of a change in the beam configuration in the middle of the run (p-Pb $\rightarrow$ Pb-p). In total about half of the statistics was recorded by a set of minimum-bias triggers relying on MBTS at L1 and a timing difference constraint at the HLT in combination with ZDC and TE triggers as shown in Fig. 5.38 [104].

High multiplicity triggers (HMT) were used to select events from the high end of the multiplicity spectra. For these trigger the track reconstruction and counting is performed at the EF. It is preceded by a L1 TE requirement and space-point counting at L2. The HMT allowed the data sample to be increased by about three orders of magnitude for the events of interest. Due to the constant rise in the luminosity during the short run, numerous triggers were enabled and controlled during data taking using prescaling at the HLT. As a result several triggers had to be combined in the analyses as shown in Fig. 5.40 [41]. An appropriate treatment of the prescaling information has to be applied to the sample in order to avoid double counting as described in [41].

originating from Pb nuclei dissociation.

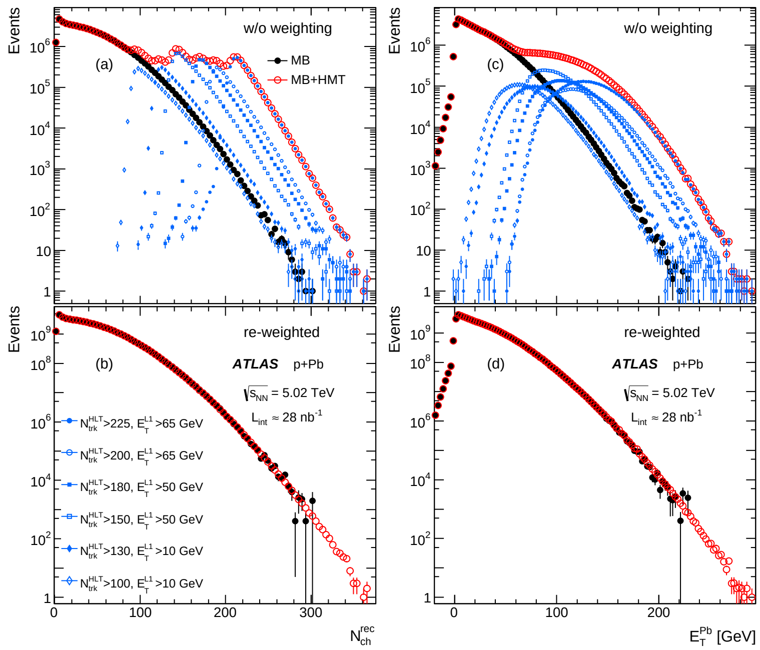

Figure 5.40: Distributions of the number of offline reconstructed charged particles, $N_{\rm ch}$, (left) and the sum of transverse energy in the forward calorimeter in the Pb-going direction, $E_{\rm T}^{\rm Pb}$, (right) for the minimum-bias trigger (black) and the number of HMT triggers (blue). Distributions, before a weighting which takes into account the prescales, are shown at the top and distributions taking prescales and correlations between the triggers into account at the bottom [41].

5.3.2 Transverse momentum triggers

In ATLAS two types of total transverse momentum triggers are implemented. The imbalanced (missing) transverse momentum trigger for the selection of events in which particles traverse the detector without leaving any signals. This signature is useful for triggering on SM processes with neutrinos involved [108, 109] but is also very valuable in searches for new physics when the postulated particles are weakly interacting and leave no signal in the detector [110]. Samples collected using the missing transverse momentum triggers are also used to estimate the trigger efficiencies of associated leptons.

The total transverse momentum of the event is an indication of whether there is a hard process in the collision and therefore can be used in minimum-bias triggering when a very low threshold is required, as mentioned in section 5.3.1. It can also be used with high threshold for new physics searches however at higher luminosities a strong pile-up dependence makes this approach impractical.

Due to the high CPU cost of full detector tracking at high frequency in the HLT and the complete infeasibility of using track information in the current L1, transverse energies measured in the calorimeter, are used instead.

The total $E_{\rm T}$ triggers (denoted as TE) apply a cut on the scalar $\sum E_{\rm T}$. The missing $E_{\rm T}$, $E_{\rm T}^{\rm miss}$, is calculated as a vector sum of energies $|E_{\rm T}^{\rm miss}| = \sqrt{(E_{{\rm T}x})^2 + (E_{{\rm T}y})^2}$ where the $E_{{\rm T}x} = -\sum E_i \cos(\phi_i)\sin(\theta_i)$ and $E_{{\rm T}y} = -\sum E_i \sin(\phi_i)\sin(\theta_i)$ are sums over the whole calorimeter [111].

At L1 the trigger towers are used to construct this quantities while at the EF fully built event information allows for a more precise calculation using cells. The towers at L1 are filtered for the presence of noise by applying a threshold (1 GeV in early running) and the energy sums are corrected to obtain a more linear correspondence to the quantities calculated offline. At the EF cells are also rejected if their signal deposition is lower than three standard deviations of the noise level of the cell.

At the HLT the momenta of the muons can be added to both sums. In order to be able to do that the muon reconstruction sequences have to be invoked in the MET and TE chains without actually requiring the presence of a muon. This is achieved thanks to the recursive mechanism for execution of HLT sequences (see section 4.2.3). In practice the MET and TE sequences demand as an input TEs that are the a result of full or partial muon reconstruction (MS only) chain and use this information if it is available as illustrated in the Fig. 5.41. This information can is available already at L2 however in early running the correction was found to be small and only calorimetric energies were used. Besides this complication the HLT reconstruction for $E_{\mathrm{T}}^{\mathrm{miss}}$ is straightforward from the perspective of the framework. However a careful treatment of the detector condition information, such as an inactive regions and precise noise levels, is needed in order to achieve operational stability as all such irregularities lead to artificial $E_{\mathrm{T}}^{\mathrm{miss}}$.

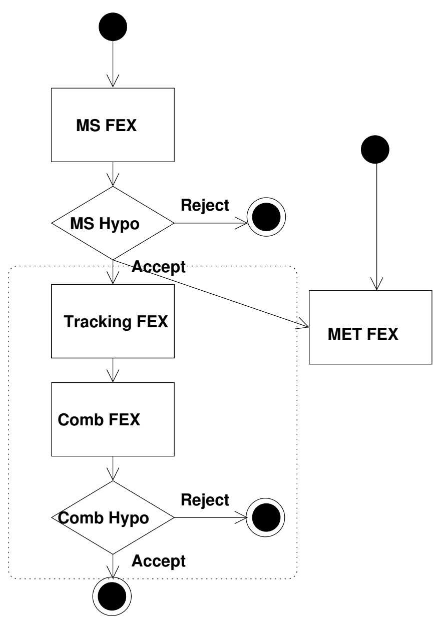

FIGURE 5.41: An illustration of the MET/TE and muon reconstruction sequences in HLT. The former depends on the presence of a confirmed muon while reuse of the same sequences for the MET trigger avoids repeated execution. In the muon sequences the MS stands for Muon Spectrometer only part of the chain while "Combo" indicates a reconstruction combining the ID and MS information as described in sec. 5.2.2.

Initially the L2 reconstruction was limited to a recalibration of the L1 $E_{\mathrm{T}}^{\mathrm{miss}}$ estimate using L1 E_x and E_y values to calculate a more precise $E_{\mathrm{T}}^{\mathrm{miss}} = \sqrt{E_x^2 + E_y^2}$ as at L1 the square root function was only estimated using a LUT.

Due to the fact that the transverse energy triggers use information from the entire calorimeter in-time and out-of-time pile-up effects have a stronger influence on them as compared to RoI based approaches. This aspect of the $E_{\mathrm{T}}^{\mathrm{miss}}$ and $\sum E_{\mathrm{T}}$ triggers has driven an evolution during Run 1 in order to retain low thresholds suitable for selecting physics at the EW scale while at the same time keeping rates under control. Advancements in the reconstruction approaches will be presented together with an overview of performance from early running through to the beginning of the Run 2.

PERFORMANCE TRANSVERSE ENERGY TRIGGERS

The $E_{\mathrm{T}}^{\mathrm{miss}}$ and $\sum E_{\mathrm{T}}$ triggers were included in data-taking as early as in 2010, both to select physics and also because, by performing full calorimeter scan, they provided detector condition monitoring. There is a big difference between the resolution of $E_{\mathrm{T}}^{\mathrm{miss}}$ at L1 and the EF as shown in 5.42. It presents the correlation between the L1 and EF $E_{\mathrm{T}}^{\mathrm{miss}}$ and offline values calculated using an algorithm based on topological clustering [68]. The resulting L1 efficiencies have a broad turn-on region. The EF (and L2) thresholds are set significantly higher to compensate, at least partly for the broad L1 turn-on. However this was only possible as long as the L1 rates were acceptable. Triggers with a reduced difference between the L1 and EF thresholds were named tight and their turn-on curves have a characteristic double slope structure as shown in Fig. 5.42. These efficiency measurements were obtained using a sample of offline selected W events collected using an orthogonal inclusive 13 GeV μ trigger.

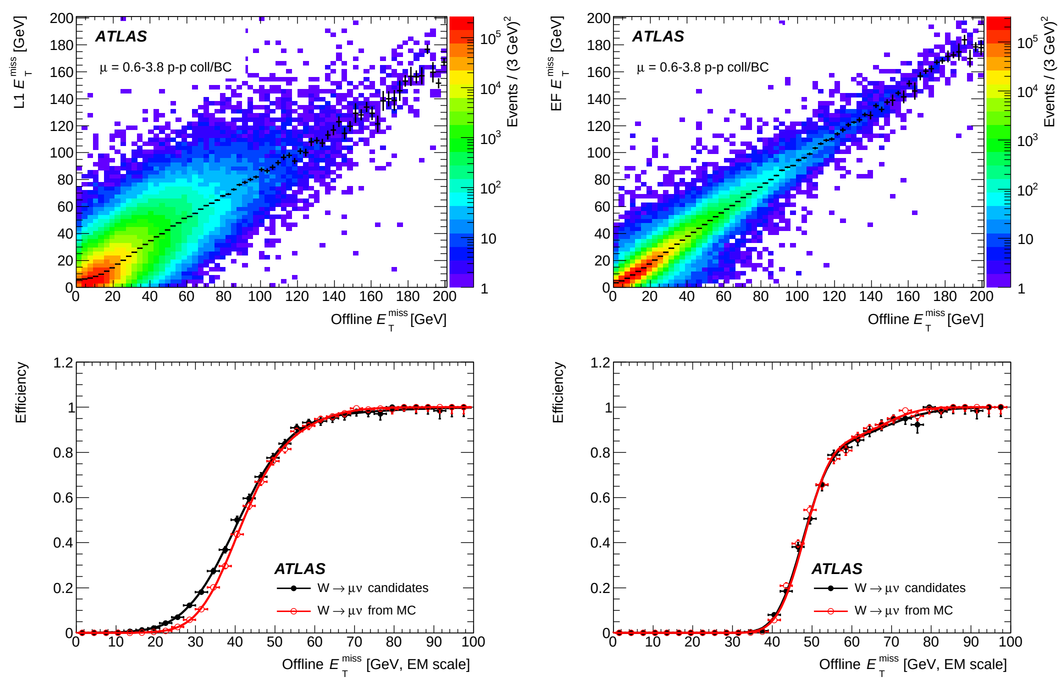

FIGURE 5.42: (top) Resolution of the $E_{\mathrm{T}}^{\mathrm{miss}}$ variable at L1 and EF with respect to offline reconstructed quantity. (bottom) The efficiency of the L1 and EF triggers for data and MC. A single, in case of L1 and double, in case of EF, gaussian distribuant functions are fit to the points [18, 112].

As early as in 2010, with a negligible pile-up of 0.6-2 collisions per crossing, the $E_{\mathrm{T}}^{\mathrm{miss}}$ and $\sum E_{\mathrm{T}}$ rates dependence on the luminosity started to deviate from linearity as shown in Fig. 5.43 [18]. This is due to the fact that the distribution of $E_{\mathrm{T}}^{\mathrm{miss}}$ changes due to the additional energy in the calorimeters and so the larger magnitude of fluctuations mimics genuine signals. The change of the $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution and rates at L1 and EF are shown in Fig. 5.43. The effect is bigger at L1 where the resolution is worse. It also demonstrates the necessity of reducing the pile-up dependence already at L1.

Running in 2011 with a pile-up of 7-15 necessitated an upgrade of the selection at L1. Instead of plain $E_{\mathrm{T}}^{\mathrm{miss}}$ the $E_{\mathrm{T}}^{\mathrm{miss}}$ significance, denoted as XS, was used for some items with

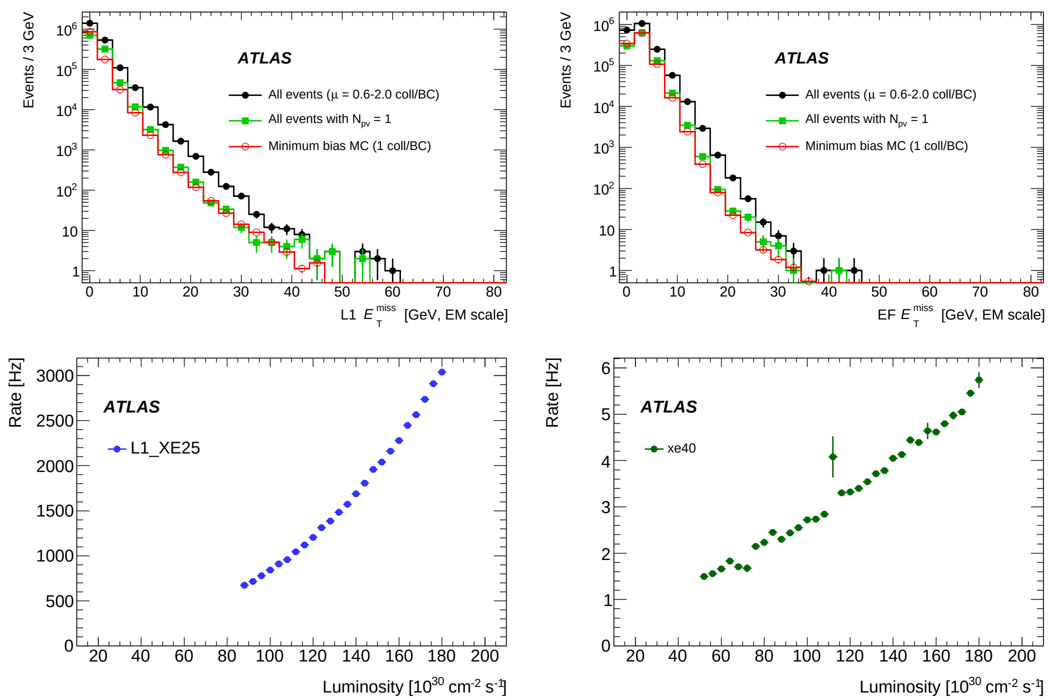

FIGURE 5.43: (top) The L1 (left) and EF (right) E_T^{miss} measured in with requirement of exactly single vertex and the quantities obtained for all events with pile-up, μ from 0.6 to 2. (bottom) L1 (left) anf EF (right) rates of E_T^{miss} triggers as a function of instantaneous luminosity coresponsding to variation of pile-up from 0.6 to 2 collisions/BC [18]

low thresholds. The E_T^{miss} significance is defined as $E_T^{miss}/f(\sqrt{\sum E_T})$, where f is a linear function [113]. For the sake of increased trigger robustness against detector operational issues an additional cuts of $E_T^{miss} > 10$ GeV and 16 GeV $<$ XS $<$4 TeV were applied. Also, because the pile-up contribution affects mostly the forward region a higher noise threshold was applied when summing cells from that region at the EF level. The rate dependence on pile-up is shown in Fig. 5.44 for the E_T^{miss} and XS triggers. This filtering of cell energies in the forward calorimeters (illustrated by the difference between rates in data taking periods E,F without filtering and G,H with filtering) resulted in a significant rate reduction. Nevertheless the rates grow exponentially with the pile-up probability. The XS triggers, for which rates are also presented in Fig. 5.44, show a weaker dependence of the rate on pile-up for low XS thresholds. For higher values of XS threshold the rate even decreases with luminosity. This is at the expense of a reduced efficiency of this trigger. As a result a balance between the E_T^{miss} and XS triggers was chosen so that the first was used as an inclusive trigger and with increased luminosity the thresholds were increased (the primary E_T^{miss} tigger ended up with a threshold of 80 GeV) or the prescales increased. The XS triggers were used in combination with other signatures.

In 2011 the L2 played a minimal role in the E_T^{miss} selection. In 2012 improvements were focused on the use of fine granularity information at L2. For this reason the readout system was modified to sum the information delivered by the front end boards of the calorimeter readout to each Readout Link (RoL) (128 cells for most of the detector) and present this information as a virtual sub-detector. Since these sums have a relatively small data-size as compared to the full cells data, they can be retrieved at L2 and need to calculate E_T^{miss} resulting in significantly

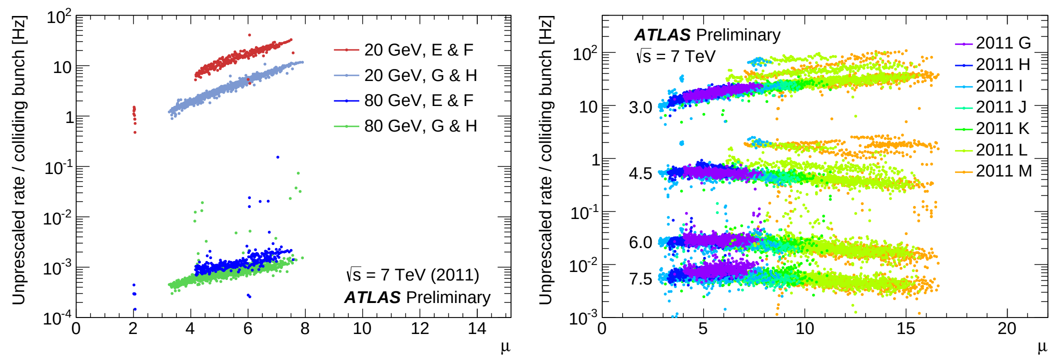

FIGURE 5.44: (left) L1 $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger rates per bunch crossing as a function of instantaneous luminosity expressed in terms of pile-up probability, μ. Two periods are compare in which the more stringent cell noise cuts were applied to the cells in the forward region (G,H) as compared to earlier conditions (E,F). (right) As on the left the rate of the L1 $E_{\mathrm{T}}^{\mathrm{miss}}$ significance (XS) triggers with various thresholds and for several running periods (G-M) characterised by varying luminosity [113].

improved resolution as compared to L1. This approach allowed to apply the selection at the L2 and the output rate to be reduced. Therefore at the EF a more time is allowed for time consuming algorithms such as offline-like topological clustering giving better performance at the higher levels of pile-up of about 40 collisions per crossing in 2012. Figure 5.45 shows a comparison of the resolutions of L2 and L1 $E_{\mathrm{T}}^{\mathrm{miss}}$ with respect to the EF value and the resolution of two different EF algorithms with respect to offline calculated $E_{\mathrm{T}}^{\mathrm{miss}}$. Each of these improvements allowed the turn-on curve to be sharpened, remaining fully efficient in plateau, while keeping the rates under control.

A set of improvements were developed in preparation for the Run 2, where luminosity and pile-up will increase further, especially the out-of-time pile-up in 25 ns bunch spacing. At L1 it is expected that the simple sum over cells will be replaced by a more refined quantity $E_{\mathrm{T}}^{\mathrm{miss}} - \sum_i w_i p_{\mathrm{T},i}^{\mathrm{jet}}$ in which the appropriately weighted (w_i) jet contribution will be subtracted. The weights will be found using a Kalman filter algorithm. The expected performance is shown in Fig. 5.46. At the HLT more advanced algorithms are to be applied earlier in the selection. Several of these have already been prepared and the best performing will be chosen. They are: simple summation of the cells with a symmetric 2σ noise cut, topo-cluster based algorithms with and without pile-up correction, a pile-up fitting algorithm and an algorithm based on jet momenta. Their preliminary efficiencies are shown in Fig. 5.47.

The Table 5.8 presents rates and selection cuts of few out of many $E_{\mathrm{T}}^{\mathrm{miss}}$ triggers during Run 1 and beginning of Run 2.

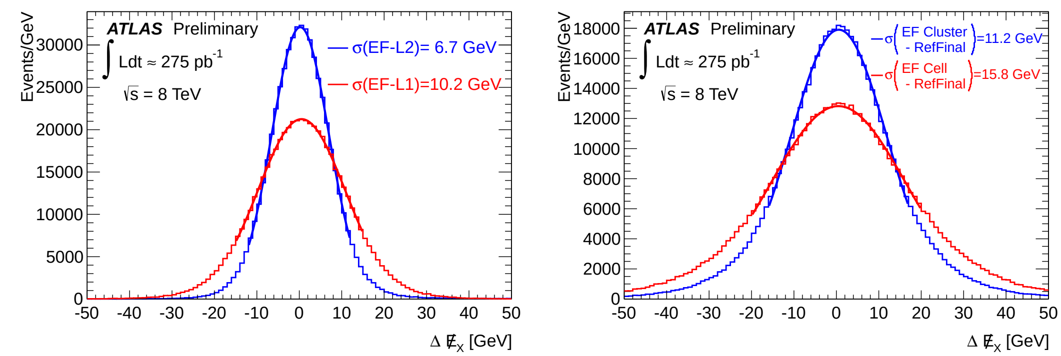

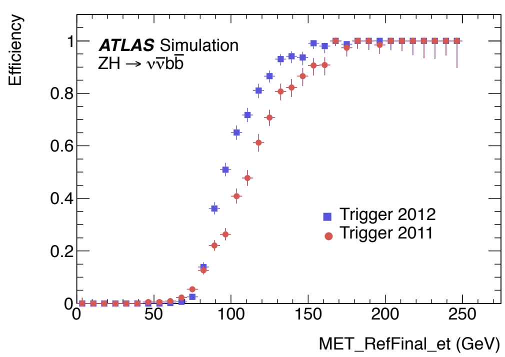

FIGURE 5.45: (top left) Resolution of L2 FEB $E_{\mathrm{T}}^{\mathrm{miss}}$ with respect to $E_{\mathrm{T}}^{\mathrm{miss}}$ estimated at the EF compared to the resolution of the L1 quantity. (top right) Resolution of $E_{\mathrm{T}}^{\mathrm{miss}}$ obtained with sum over all cells, filtered for the noise with respect to the offline reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ compared to the EF quantity obtained from cells grouped by a topo-clustering algorithm prior to the summation [114].

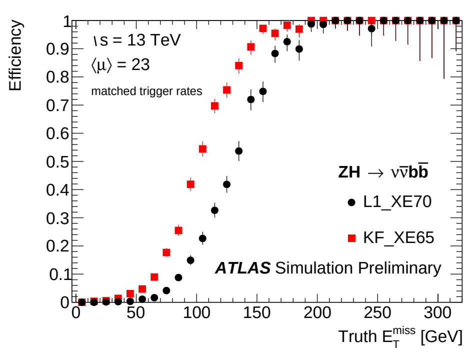

FIGURE 5.46: Expected efficiency of the L1_XE70 $E_{\mathrm{T}}^{\mathrm{miss}}$ and newly designed topological trigger KF_XE65 for the process $HZ \to \nu\bar{\nu}b\bar{b}$ [115, 116].

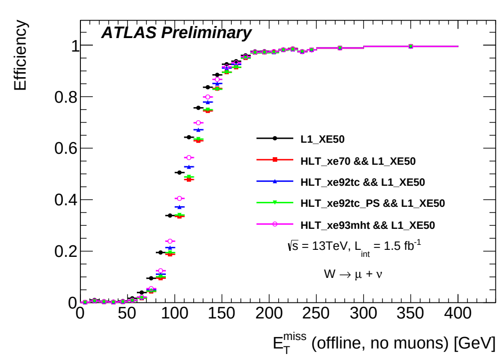

FIGURE 5.47: Performance of variants of $E_{\mathrm{T}}^{\mathrm{miss}}$ configuration prepared for the Run 2, HLT_xe70 based on plain sums over cells, HLT_xe92tc using topological clustering, HLT_xe92tc_PS topo-clustering with η-ring pileup subtraction, HLT_xe93mth based on sums of jets momenta [114, 117, 116].

	$\mathcal{L}\ \times 10^{30}$ $[cm^{-2}s^{-1}]$	$E_{\mathrm{T}}^{\mathrm{miss}}$ [Hz]	$E_{\mathrm{T}}^{\mathrm{miss}}+\mu$ [Hz]	$E_{\mathrm{T}}^{\mathrm{miss}}+e$ [Hz]	$E_{\mathrm{T}}^{\mathrm{miss}}+\tau$ [Hz]
2010					
begin	$\sim$1.7 10^{-1}	0.25 [20]			0.04 [20+ τ12,loose]
begin	$\sim$1.9 10^{2}	6 [40]		6 [20+ e20,loose]	24 [22+ τ16,med]
2011					
begin	$\sim$3.5 10^{2}	1.5 [60]		2 [30+ e15,med]	4 [35+ τ29,med]
end	$\sim$3.1 10^{3}	5 [70]	15 [30+ μ15]		$\sim$0 [35+ τ29,med]
2012					
begin	$\sim$2.5 10^{3}	4 [70]		0.2 [60XS +e20E_{T} cut]	4 [40+ τ29 iso.]
end	$\sim$6.6 10^{3}	55 [80]‡		4 [45XS+ e13 E_{T} cut]	14 [40,tight+ τ29T iso.]
2015 (Run 2)					
begin	$\sim$4.5 10^{3}	120 [70]*		1 [60+ e60LH,loose]	1.2 [70+ τ35,med1]

Table 5.8: The rates and selection criteria for a typical $E_{\mathrm{T}}^{\mathrm{miss}}$ triggers. Triggers combining only μ and the $E_{\mathrm{T}}^{\mathrm{miss}}$ were only used temporarily, in latter data taking a more sophisticated combinations with other signatures were used. †The $E_{\mathrm{T}}^{\mathrm{miss}}$ significance triggers were in preparation and were not used unprescaled, lowest threshold is quoted. ‡ The 60 GeV was present in the menu but events triggered by it were directed to a data stream processed with a delay as compared to the main physics stream. Also many triggers at 80 GeV were set up in this period varying in L1 seeds, L2 cuts, noise cuts and pile-up suppression strategy. The rate of the variant delivering highest rate is quoted. *About a dozen of $E_{\mathrm{T}}^{\mathrm{miss}} > 70$ GeV triggers were configured varying by the approach to pile-up suppression. The rate of the one delivering highest rate is quoted. Rest as in 5.4.

5.4 Hybrid approaches

Another class of triggers that require significant support from the HLT framework in order to be implemented efficiently use a hybrid approach using RoI and full scan reconstruction. This can be justified by better physics performance or resources consumption. When combining this two approaches the RoIs need to be either merged into larger regions or they have to be produced as the result of full scan reconstruction. In some cases the RoIs are not restored after full scan reconstruction.

Two types of signatures are given as examples of these strategies. However these complex processing strategies are beginning starts to be used also in signatures described above. For instance the pile-up corrections typically use global reconstruction in order to associate tracks to vertices or identify clusters from the secondary interactions.

5.4.1 Light jet triggers

The jet trigger is used to collect events for studies of QCD [118], fully hadronic top quark decays [119] as well as many searches such as the first search by ATLAS for new particles decaying in two quarks [120] or contact quark interactions [121]. Due to the very high di-jet production cross-section misidentified jets constitute a background for e/γ and τ triggers. Control jet samples are therefore recorded in order to study these effects. A complete set of jet signatures consists of single and multi-jet triggers. Also, topological requirements are applied in order to acquire events with certain characteristics. In combination with other signatures jets constitute a significant part of the trigger menu.

The L1 jet trigger uses a sliding window algorithm that is optimal for a hardware implementation. The state-of-the-art jet-finding algorithm, anti-k_T [122] used in offline reconstruction, is inherently different from the sliding window algorithm resulting in broad L1 turn-on curves. This is compensated by setting lower L1 thresholds and then performing reconstruction using close to offline at the EF preceded by a cone algorithm [61, 123] executed inside of RoI at L2 with an enforced limit on the number of cone finding iterations (three in 2010). In general the L2 and EF thresholds are set so that the single jet trigger efficiency reaches maximum at the same point that the L1 trigger does, that is they affect only the steepness of the turn-on curves.

However an undesired feature of the distinct L1 algorithm is the behaviour for large area jets and multi-jet triggers. The former are affected because the wide-spread deposits appear as many low energy jets at L1 while in the case of multi-jets the square shape of the trigger towers and poor L1 resolution results in a miscounting of jets that can cause failure of the multi-jet trigger. Of special concern was a nonlinear rate dependence on the luminosity of this triggers [124]. For this reason a clustering algorithm closer to the offline algorithm is desired early in the reconstruction stage. In 2011 the *L1.5 jet* algorithm using trigger tower information from the L1 calorimeter trigger readout system was developed. Towers were input to the anti-k_T algorithm which scanned the entire calorimeter detector without requiring any L1 seed nor was it constrained to an RoI. That was possible because none of the high rate L2 reconstruction used the L1CaloRoS and so an additional readout rate at the level of 7 kHz (15 kHz after RoS upgrade for run in 2012) could be afforded. This approach allowed multi-jet triggers efficiency to be recovered and permits the evaluation of an anti-k_T algorithm with a wide radius of $R = 1$. However the resolution of this algoritm was found to be inferior in comparison to the L2 cone algorithm using cell information. Therefore the L1 and L1.5 jet algorithms were used to find RoIs while the cell-based cone algorithm performed precise E_{T} reconstruction. For multi- or large area- jet triggers the L1.5 jet information was used to accept event at L2 that were then processed by the anti-k_T algorithm using cell information at EF.

Due to the poor energy resolution of the algorithm based on trigger towers, it was clear that an optimal algorithm should perform an anti-k_T reconstruction using cells as early in the chain as possible. Retrieving full calorimeter information and performing topological clustering and jet finding is not feasible at high rate. A feasible approach is to implement an algorithm performing cell-based reconstruction in all the regions pointed to by L1 jet RoIs in one pass, a partial-scan algorithm. Figure 5.48 shows the construction of the partial-scan region as well as its operational performance showing the superiority of this approach. The partial scan reconstruction does not use L1 information as seeds and therefore the overlaps are eliminated by construction [59]. As can be seen, the number of retrieved cells and therefore their readout time and the clusters formation time are significantly lower when compared to a full scan approach. That development was done in preparation for the Run 2 data taking period. Even though the L2 and EF have been combined the Run 2 system a fast rejection algorithm, prior to the event-building stage allows for a greater flexibility.

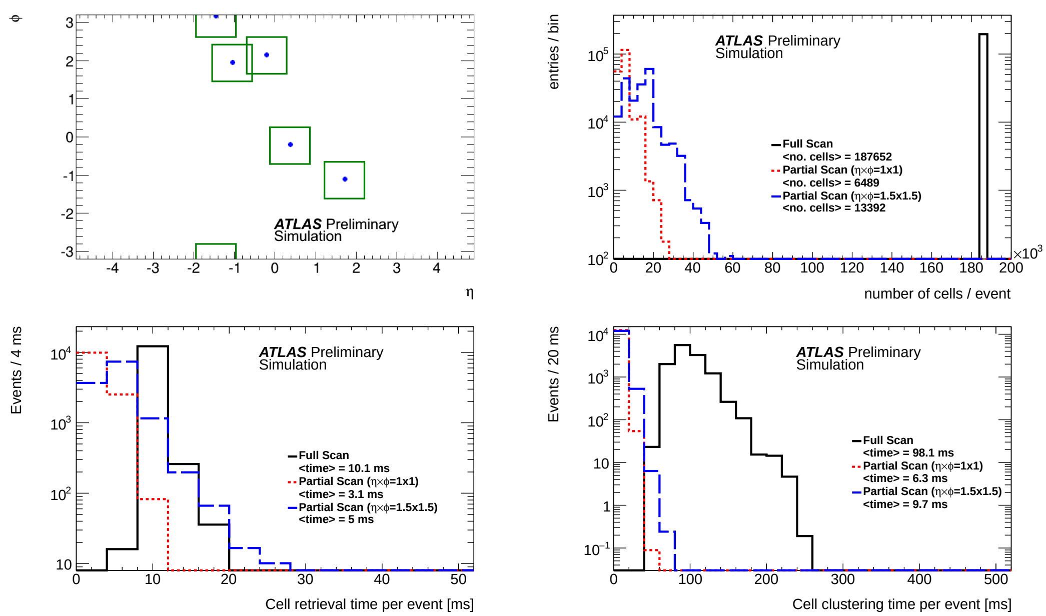

FIGURE 5.48: (top left) Event with multiple L1 jet RoIs (blue dots) and the area read of in the partial-scan algorithm (green squares) [125]. (top right) Number of cells retrieved for the purpose of jets reconstruction in partial scan (two region sizes) and a full scan approach in which for every event entire detector is red out. (bottom left) Time to retrieve cells in aforementioned approaches. The difference is not as significant because of the other calorimetric reconstruction and because of a constant overhead of the request. (bottom right) Time to reconstruct clusters from the retrieved cells [59].

The jet reconstruction algorithm used at the EF is an offline anti-k_T algorithm wrapped for the HLT. It is preceded by preparation of the cell data (decoding and calibration) and a noise robust, clustering algorithm known as topological clustering [68]. Initially a RoI-based approach was used but for higher luminosities it was found that a more optimal approach was to abandon RoI-based reconstruction and perform the reconstruction using the entire calorimeter information. The main reason was to avoid duplicated reconstruction of the same jets as well as a significant simplification of the adaptation of the offline algorithm to run in the HLT. It also opened up the possibility of using an independent seeding scenario i.e. L1 random or total E_{T} triggers. In the final stage of the jet trigger the RoIs were not re-constituted and the hypotheses algorithms worked instead on one set of jets per event rather that on RoIs. Figure 5.49 shows graphically all the discussed options.

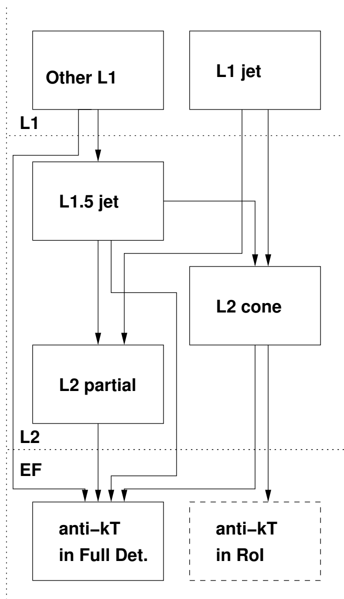

Figure 5.49: A sketch of the possible jet reconstruction algorithmic blocks (boxes) and their interrelation. Details involved in each block, like data preparation or clustering are not shown for simplicity. With dashed line EF RoI based reconstruction used in early Run 1 (2010) is marked.

For the purpose of triggering during the Pb-Pb collisions an additional algorithm was developed. It proceeded the cluster formation step in order to correct cell energies for the underlying event. In this algorithm the average energy deposition in slices of pseudorapidity was calculated and subtracted from the cluster energies used to form jets. In contrast to the offline algorithm the area of the jet was included in the estimate which resulted in a slight bias with respect to the offline where, in iterative process, contribution from jets to the underlying event as well as the underlying event azimuthal variation are eliminated [126]. Nevertheless this enabled collision centrality independent jet algorithm performance to be obtained.

For Run 2 a novel approach was introduced in order to access the jet spectra at very low E_{T} where measurements are statistically limited. Namely the jets reconstructed by the HLT algorithms are recorded without the detector data for that event. This effectively reduces the event size and allows a much larger number of events to be recorded.

Jet trigger performance

The jet HLT trigger was used to actively reject events already in 2010. Only the L2 selection was used as it was sufficient to reduce the rate to the desired level. Clusters used by the cone algorithm were calibrated at the EM scale and as a result actual E_{T} threshold cut that was applied was lower than the offline value. The residual E_{T} bias due to the diferences in cone and anti-k_T algorithms is shown in Fig. 5.50 as function of E_{T} and η. As can be seen the L2 algorithm overestimates the energy at low E_{T} but is compatible with the offline at high E_{T}. The effect is well reproduced in MC simulation. No significant deterioration of the performance was observed in the forward region. Both L2 and L1 jets were found to be nearby to the offline counterparts [127].

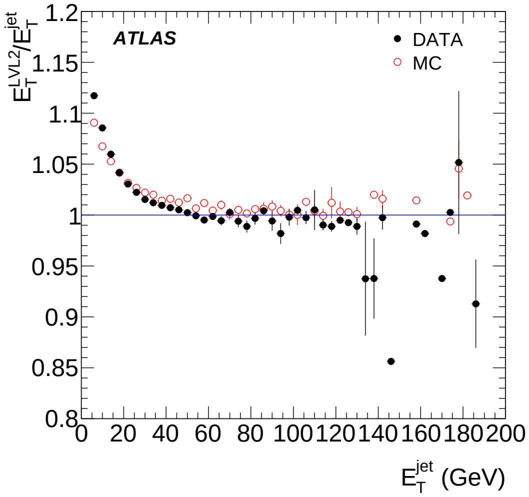

FIGURE 5.50: Bias of the E_T estimation as a function of offline jet E_T (left) and η (right) of the L2 jet algorithm from 2010 in data and MC [18, 127]. Results of the offline algorithm was re-calibrated to EM scale for comparison.

For optimal efficiency the thresholds at L1 and L2 were aligned so that the L2 only removes events from the L1 turn-on region as shown in Fig. 5.51. The plateau is reached at around 80 GeV for 30 and 45 GeV cuts at L1 and L2 respectively. A slight variation of the L1 efficiency is a result of calorimeter cell miss-calibration in the running conditions with bunch trains. Figure 5.52 shows the scaling of rates with luminosity for the entire range of luminosities from 2010. The scaling is linear as pile-up levels were very modest.

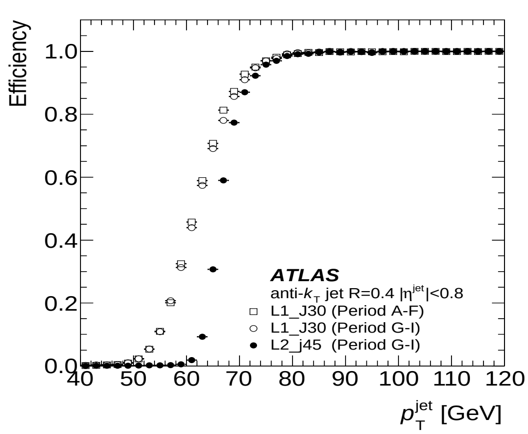

FIGURE 5.51: L1 and L2 jet trigger efficiencies as a function of offline reconstructed jet E_T [18].

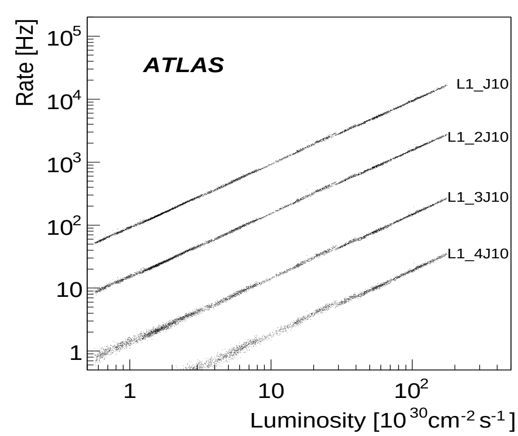

FIGURE 5.52: (right) Rates of single and multi jet triggers as a function of instantaneous luminosity in 2010 [18].

In 2011 the EF selection was enabled for jets. The hadronic calibration was not yet applied resulting in a reproducible bias as shown in Fig. 5.53. The resolution (RMS of the E_T ratio of EF to offline reconstruction) was stable with E_T and amounted to about 2.5% as shown in Fig. 5.53. The bias becomes smaller at high η values while the resolution stays approximately contant as it depends on the amount of inactive material in the calorimeters. A missmatch between data and MC was attributed to a missmodeling of pile-up in the MC simulation.

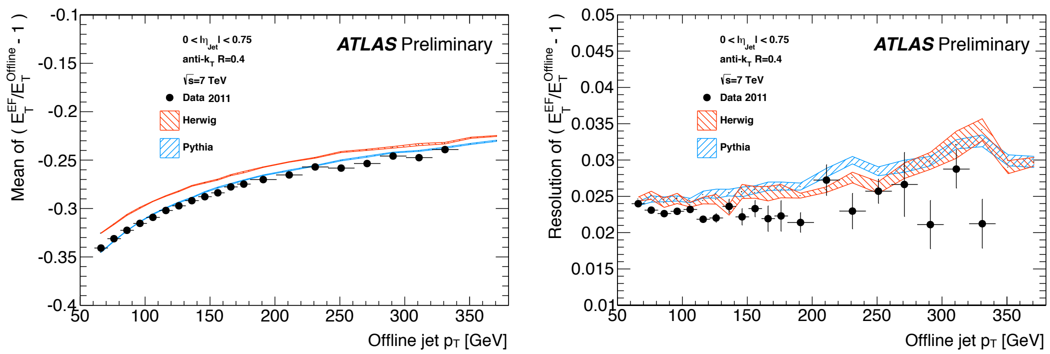

FIGURE 5.53: (left) Bias of the jet p_T estimation of the EF jet algorithm and the one used in offline reconstruction in data and MC [125]. No recalibration to the EM scale was performed for the offline measured jets. (right) Resolution of EF jet E_T estimation with respect to offline reconstructed jets.

A simple pile-up and noise correction was implemented for 2011 running. Cells in which the signal deposition was below a configurable η and noise dependent threshold were not taken into account in the clustering step. As a result a fraction of the energy from soft particles was eliminated from the jet which sharpened the turn on curves as shown in Fig. 5.54 [128]. As can be seen the most significant effect occurred at L2 as the EF topo clustering algorithm already reduced significantly the chance that isolated cell signals due to noise are included in clusters.

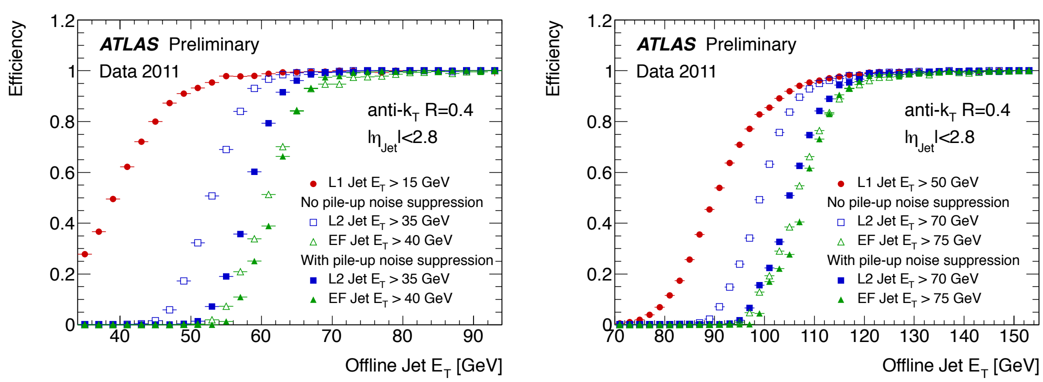

FIGURE 5.54: Illustration of the change of the efficiencies at L2 and EF with pile-up correction introduced in 2011 for low (left) and high (right) E_T jet triggers [125, 128].

Figure 5.55 shows a summary of the L1 L2 and EF efficiency turn-on curves in 2011 [127]. Low E_T triggers use a random L1 trigger to seed directly the EF jet finding algorithm and are efficient at energies as low as 20 GeV. Efficiencies were measured using a min-bias trigger at low E_T and a bootstrap method for high E_T. Discrepancies between the data and MC in the turn-on region had no effect on the analyses as the wide array of E_T thresholds allowed the use of several triggers, always at the efficiency plateau which, as indicated by both MC and data, is nearly 100%.

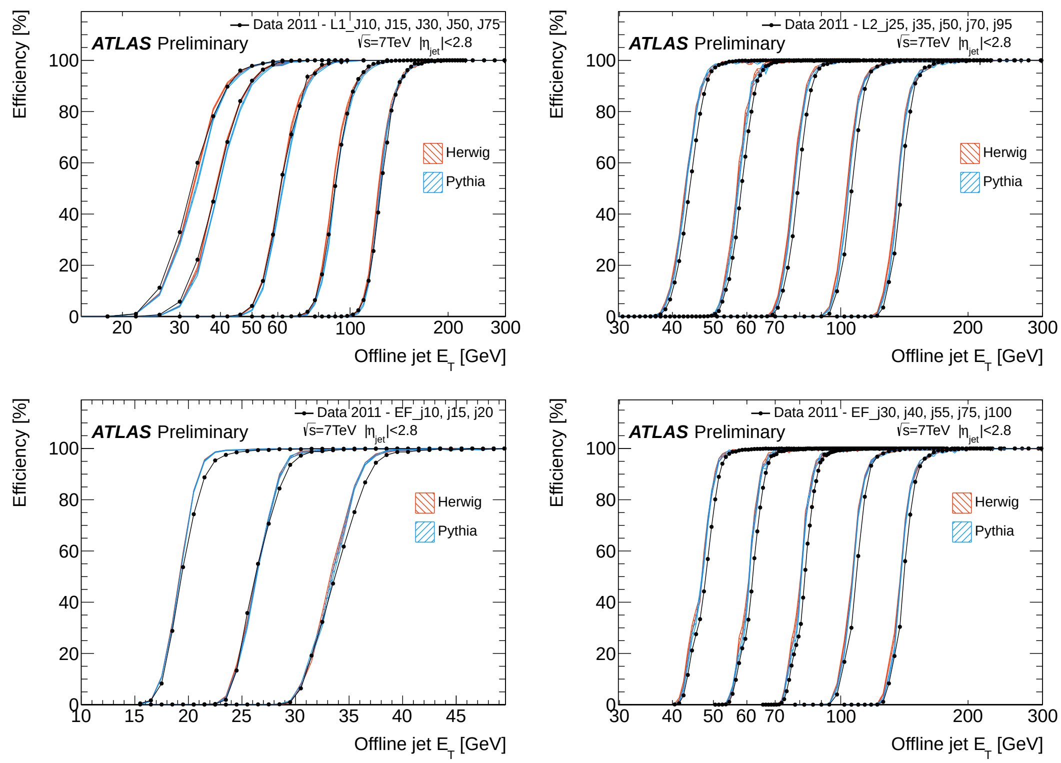

FIGURE 5.55: Efficiencies of the L1 (top left) L2 (top right) and EF jet triggers seeded by the L1 random (bottom left) and passing full chain of selection (bottom right) jet triggers in 2011 as a function offline jet E_{T} [125, 127].

In 2011 the L1.5 algorithm was developed and started operation in 2012. This algorithm was used in multi-jet signatures which were seeded at L1 by the requirement of a fewer number of jets. The advantage of this approach is that is gave a significant improvement in angular resolution as shown in Fig. 5.56. As can be seen the resolution obtained with the anti-k_T algorithm is much improved as compared to L1, even with the same size of towers of 0.2×0.2 used at L1 and the HLT. With a tower size of 0.1×0.1 a resolution approaches that of the L2 algorithm that uses cell information. In practice the L1.5 efficiency is superior in events with nearby jets. It can be quantified as a relative efficiency increase as a function of the jet separation distance as is shown in Fig. 5.57. An increase in efficiency of up to 20% is observed for jets separated by $R < 0.5$. No advantage of L1.5 was observed when the distance between the jets is larger than this.

The commissioned L1.5 algorithm was used in entire operation in 2012 and the multi-jet triggers used this approach by default. A clear efficiency increase with respect to the offline selection was obtained as shown in Fig. 5.58. The trigger turns on is more rapid and the plateau efficiency reached almost 100% in contrast to the L1 efficiency that was the same as in 2011.

In 2011 the HLT jets trigger had to be used in the HI run as the instantaneous luminos-

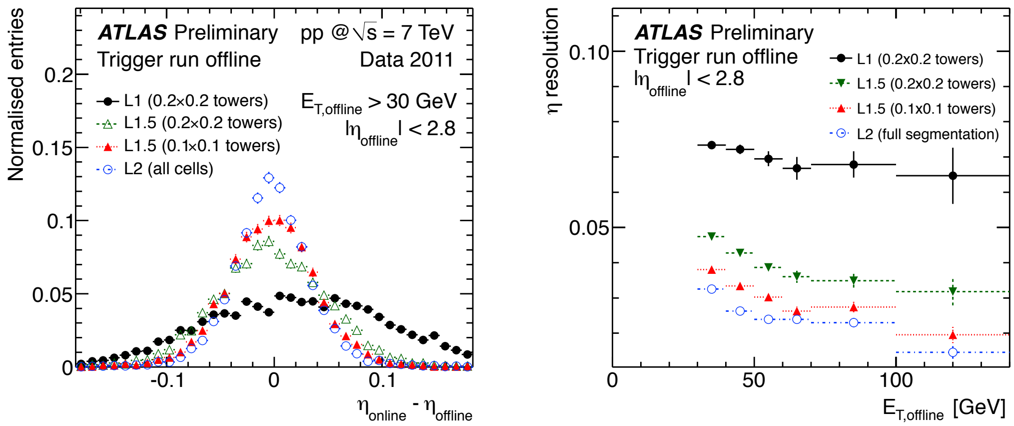

Figure 5.56: (left) Difference between the pseudorapidity, η between the offline reconstructed jet and the one reconstructed by the L1 and three L2 jet finding algorithms for all jets of $E_{\mathrm{T}} > 30$ GeV. (right) Resolution of the L1 and three L2 algorithms as a function of E_{T} [125].

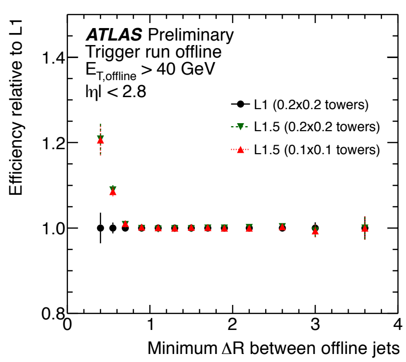

Figure 5.57: Relative efficiency increase of the L1.5 jet trigger with respect to the L1 trigger [125].

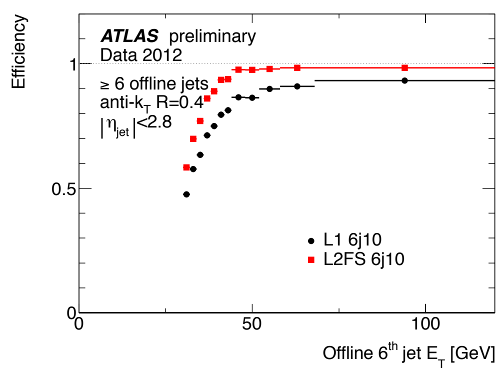

Figure 5.58: Comparison of efficiencies of L1 and L1.5 jet trigger (renamed to L2 full-scan or for short L2FS) based triggers in 2012 [125, 124].

ity was to high too collect data using only the L1 minimum-bias selection. The L2 step was skipped in the configuration and L1 jet seeds were used for high E_{T} while a total E_{T}, prescaled appropriately, was used to seed EF jets reconstruction at low E_{T}. The aforementioned underlying event subtraction algorithm was applied and this allowed for the stable functioning of the trigger across collision centralities as shown in Fig. 5.59.

The performance of the L2 partial scan approach was validated on simulated data. A comparison was made with respect to the results obtained when the entire calorimeter information was used by the L2 algorithm. Two sizes of scanned regions were studied, $\eta \times \phi = 1$ and 1.5×1.5. No significant difference was observed in the number of jets, as shown in Fig. 5.60. However, as expected, a better energy estimate was obtained when larger region size was used at the expense of higher cell data retrieval and clustering times as shown earlier in Fig. 5.48.

Functionality of recording solely HLT reconstructed jets without an event information, called

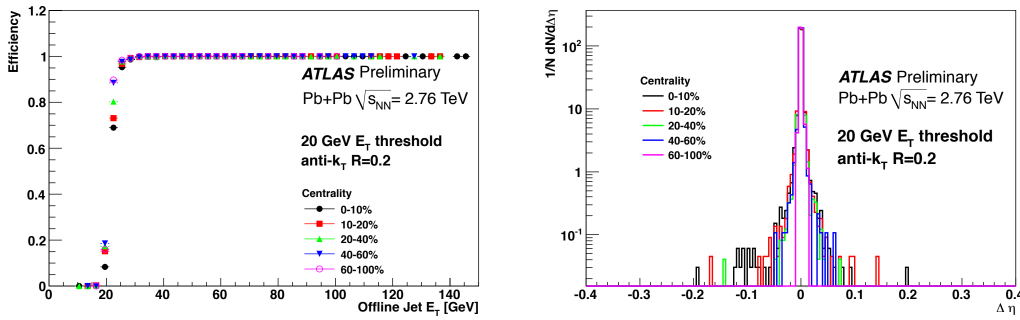

FIGURE 5.59: (left) Efficiency of the low E_T jet trigger (seed by $\sum E_T$ L1) as a function of offline reconstructed jet E_T for 5 bins of collisions centrality (0-10% are the most central collisions with highest underlying event contribution). (right) Angular resolution of the HLT jet reconstruction. Rest as on the left [125].

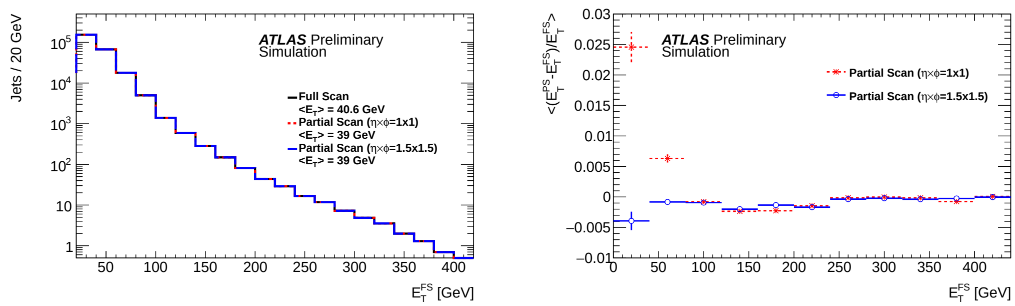

FIGURE 5.60: (left) Spectra reconstructed by L2 partial scan algorithms differing in scanned area sieze compared to a full-scan approach. (right) Energy bias of the L2 partial scan algorithms with respect to the full scan algorithm [125, 59, 128].

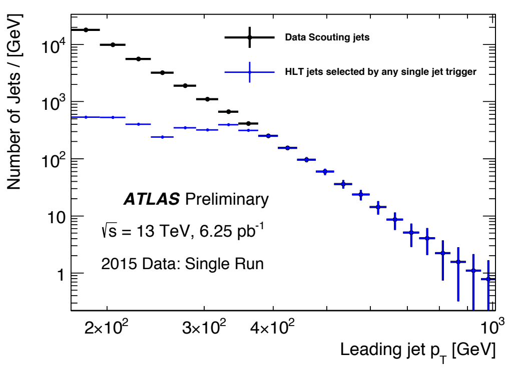

FIGURE 5.61: Leading jet transverse momenta in physics stram and in the stream when recording of only the HLT jets recording is performed (labeled as "Data scouting jets") [125].

data scouting, was commissioned for Run 2. Figure 5.61 shows the yield as a function of the leading jet p_T obtained by regular, prescalled at low E_T, triggers and the yield of all jets reconstructed at HLT and recorded without a loss. At high E_T both triggers provide the same statistics while in the jets in scouted data stream deliver an order of magnitude higher statistics at $E_T < 400$ GeV. With increasing luminosity, when unprescaled jet triggers will need to have yet higher E_T thresholds the advantage of this method will expand to higher E_T and will be even more pronounced at lower E_T. The E_T resolution at the HLT, which is normally a concern with this approach plays even more important role as a too biased E_T measurement or too poor resolution of the E_T estimation would directly affect the analysis. The resolutions with respect to MC simulation and offline reconstructed jets are shown in Fig. 5.62. A bias at the level of 1% at lowest E_T is observed while resolution is at the level of few % assures usability of the scouted jets data for analyses.

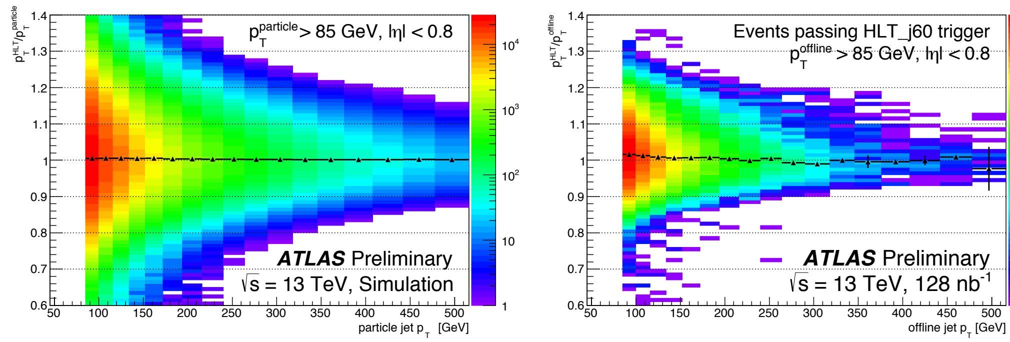

FIGURE 5.62: Comparison of E_T reconstructed in HLT to the truth particle p_T (left) and jet E_T reconstructed offline (right) [125].

Example of primary jet triggers during Run 1 and beginning of Run 2 are show in the Table 5.9 together with the rates at specific luminosities.

	$\mathcal{L}\ \times 10^{30}$ $[cm^{-2}s^{-1}]$	single-jet [Hz]	di-jet [Hz]	multi-jet [Hz]
2010				
begin	~1.7 10^{-2}	6 [40]	8 [2×10]	
end	~1.9 10^{2}	0.6 [95]†	0.4 [2×35]†	4 [4×35]*
2010				
begin	~3.5 10^{2}	3 [180, anti-kT R=0.4]	2 [180,30 $+\Delta\phi >$ 0.4 cut]	1.5 [4×45]*
end	~3.1 10^{3}	1.5 [320, anti-kT R=0.4]	1 [2×135 $+\Delta\phi >$ 3.5 cut]	11 [4×45]*
Pb-Pb	~4.5 10^{-4}	35 [20, anti-kT R=0.2, L1 TE10 seed]	-	-
2012				
begin	~2.5 10^{3}	3 [250, anti-kT R=0.4]	1.2 [2×200 $+\Delta\phi >$ 0.4 cut]	4 [4×80 had. calib.]*
end	~6.6 10^{3}	7 [360, anti-kT R=0.4 had. scale]‡	0.6 [180+145 $+\Delta\eta > 3.4$ $m_{jj} > 2TeV$]	8.5 [4×80 L2FS]*
2013				
p-Pb	~4.5 10^{-1}	9 [90, anti-kT R=0.4, EFFS]◇		
2015 (Run 2)				
end	~4.3 10^{3}	17 [360]	△	18 [4×85]

Table 5.9: The rates and selection criteria for example primary jet triggers. †This triggers are prescaled in that particular run and are presented as an example of rates allocation given to this trigger. *Many multi-jet triggers were used in this runs (including 6-jet and with changing L1 seeding and HLT approach optimising multi-jet performance described earlier), a typical examples are given. ‡Lower unprescalled thresholds, down to 220 GeV, were used to select events to the stream processed with a delay as compared to the main physics data streams. ◇A number of prescaled triggers based on min-bias L1 seeds were used in p-Pb run to record events with jets of lower E_{T} which were in fact of primary use for physics analyses. △The di-jet triggers were mostly devoted to heavy flavour jets and contained an additional b-tag requirements. Rest as in 5.4.

5.4.2 Heavy flavour jet triggers

Heavy flavour jets from b and t quarks are important for many hadronic channels. Additional jet flavour tagging leads to a rate reduction and thus is beneficial at low p_{T} where light jets triggers would otherwise to be significantly prescaled. In early measurements when the minimum p_{T} threshold for the jet was sufficiently low so that inclusive jet triggers were used to provide datasets for heavy flavour the analyses. The b-jet identification performed offline was sufficient [129]. At higher luminosities, only b-jet triggers can offer the possibility to select hadronic-only decays with heavy flavour jets, such as $hh \to 4b$ where there is no additional lepton that can be used for triggering and plain jet triggers are excluded due to the recording rate limit [130, 131].

A b-jet can be identified through the property of the b-hadron decay within the jet. The presence a the muon inside the jet cone is the simplest possibility. Another technique exploits the presence and properties of the secondary vertex of b-hadron decay. The first method was used initially to collect statistics to calibrate the secondary vertex method. The muon-based trigger was also used during the Pb-Pb run. The secondary vertex method uses the fact that the vertex from which the tracks in the b-jet cone originate is distinct from the primary interaction vertex. This approach provides a much higher identification efficiency as the branching-ratio for direct and casacade muon production is only about 20%.

In offline reconstruction both techniques are used however the primary b-jet tagger relies on secondary vertex information [132]. Offline identification relies on highly efficient tracking and vertexing in addition to jet reconstruction. In the HLT due to the resource limitations, tracking for b-jet triggering can only be performed within a jet cone and so the primary vertex position is only approximated. Therefore the resolution of secondary vertex tagging is larger than in the offline. However this aspect was significantly improved during Run 1. The primary vertex is assumed to be contained within the luminous region determined by conditions data. Additional chains were present in the trigger configuration to provide measurements of the transverse and longitudinal position and extent of the beam-beam interactions region (see section 5.5). At the start of the data taking run, during the first few luminosity blocks, when the beam position is not yet known, the b-jet tagging algorithms reject every event to avoid a rate increase by triggering on light jets. The difference in the distribution of the b-tagging discriminant variable before and after the beam-spot information update is shown in Fig. 5.63 [18]. A significant shape modification, most notably a reduction of counts in the peak around zero, would result in a purer b-jets selection, improved efficiency and reduced rate.

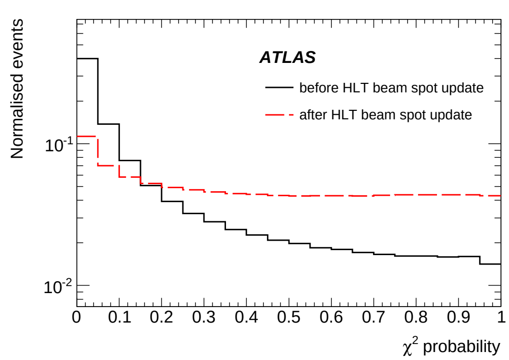

Figure 5.63: The b-jet discriminant variable used at L2 in 2010, χ^2 probability, before and after beam-spot position update during the run [18].

The z coordinate of the primary vertex was measured in every event. For that purpose track trajectories were projected back to the luminous region and the intercepts were histogrammed using all tracks in the RoI. A sliding window algorithm was used to find peaks in this histogram and the one with most tracks associated to it was considered to be a primary vetex.

Over the years various multivariate techniques were used to obtain b-tagging information. Initially the most robust technique was based on the signed impact parameter significances (transverse and longitudinal) with respect to the primary vertex, translated to a χ^2 probability. Impact parameter is defined as the distance of closest approach of a track trajectory to the primary vertex signed with a track intercept with jet axis so that the track originating from a position along the jet axis direction will have positive sign as illustrated in Fig. 5.64. For light quarks, the χ^2 probability is expected to be uniform (all tracks originate from the vertex with the spread determined by track and vertex resolutions) and a peak at zero is expected for b-jets [18].

A change of the approach used for primary vetex finding was implemented for 2012. All candidate RoIs passing the jet trigger selection were combined as input to the primary vertex finding algorithm, greatly improving resolution as it linearly depends on the number of tracks involved in vertexing. As a consequence the 2D χ^2 probability tagging method was used. It improved efficiency and reduced the misstag rate. Other taggers explored the other b-hadron characteristics such as the mass, the number of two track vertices and the ratio of energy associated with the secondary vertex to the total jet E_{T} [133, 134].

For the Run 2 more sophisticated configurations are prepared in which yet better primary vertex estimation is available. This is achieved by implementing a two-stage tracking involved in the selection chain. In the b-jet chains early rejection is achieved by using the light jet algorithms which apply a cut on E_{T}. If the event passes that stage, the lower E_{T} jets are grouped in one super-RoI and tracking and vertex finding is performed in the super-RoI providing an improved estimation of the primary vertex position. Another stage of jets selection is then performed. For the jets passing this selection further precise tracking is repeated inside of each RoI the results of which are delivered to the b-tagging tools together with the primary vertex information. The arrangement, with significant simplifications, of the b-jet chain components is shown in Fig. 5.65.

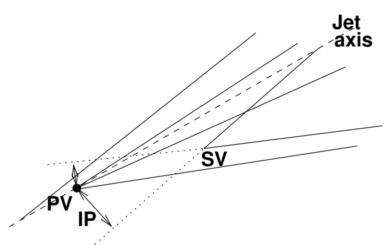

Figure 5.64: Illustration of impact parameter (IP) definition which is a distance of track trajectory from primary vertex (PV) and signed according to the intercept with jet axis. The secondary vertex (SV) is also shown in the drawing.

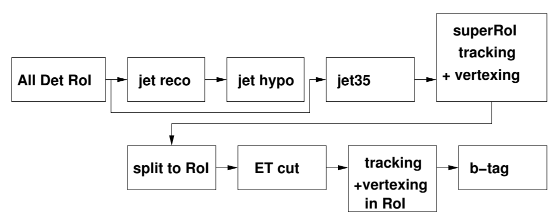

Figure 5.65: Sketch of the selection flow in the b-jet triggers in Run 2. See text for description.

Performance of the b-jet triggers

In 2010 the b-jet triggers based on vertex properties were not used to actively reject events. Instead the jet+muon events were used to collect samples enriched in b-jets for performance studies of the offline and trigger secondary vertex taggers. Studies of impact parameter resolution were performed using this data. Lifetime based triggers were run parasitically so that trigger level quantities were recorded in the data for offline performance studies. An example is the impact parameter distribution shown in Fig. 5.66 for L2 tracking where a slight asymmetry due to b-enrichment can be observed [18]. The left side of the impact parameter significance

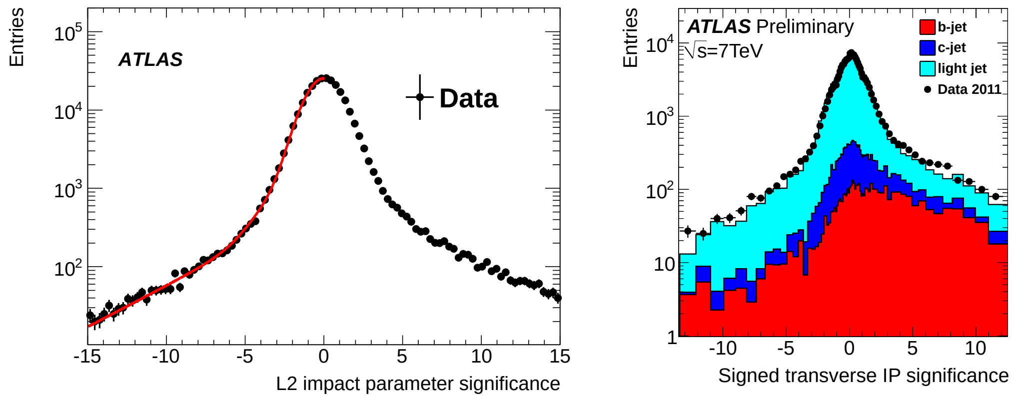

FIGURE 5.66: Distribution of transverse impact parameter significance in candidate b-jets at L2 in 2010 (left) and EF in 2011 (right) [18, 135].

distribution, after reflection to the right side, is used to study to estimate the efficiencies and fake rates of the taggers for running in 2011. Also shown is the same distribution measured on data in 2011 and a comparison to the MC with mix of light and b- and c-jets [135]. The good data-MC agreement allowed the desired identification efficiencies to be attained. The resulting performance with respect to the offline b-jet tagging is shown in Fig. 5.67 where the distribution of the negative logarithm of the probability that the jet originates from a b-hadron is plotted. Near zero, the light-jet region, use of b-jet trigger removed most of the background events while

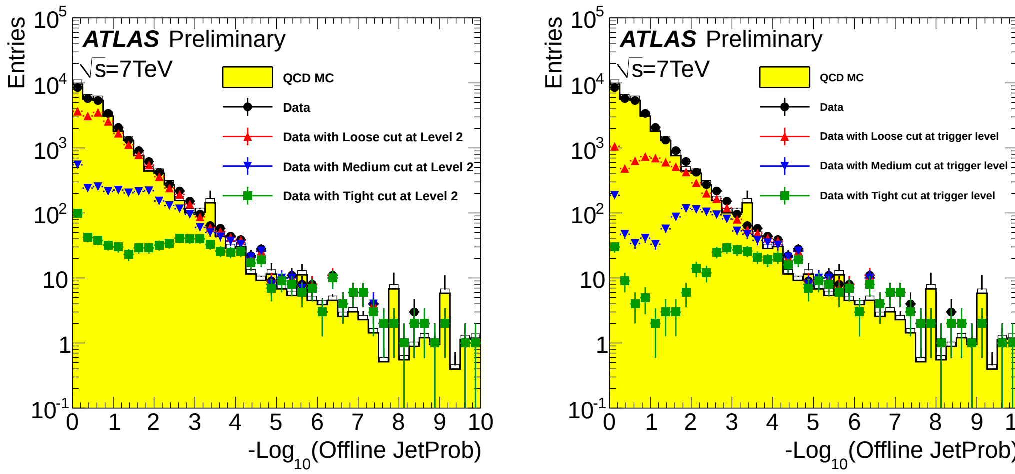

FIGURE 5.67: A comparison of online b-tagging performance of three selections (loose, medium, tight) for unbiased jet trigger and enriching b-jet sample at L2 (left) and combined L2+EF (right) [136].

in the region of probability approaching zero (see Fig. 5.63) where the negated logarithm grows nearly all offline b-tagged jets are also tagged by the trigger. The working points of the HLT b-jet triggers providing 50% (tight), 60% (medium) and 75% (loose) efficiencies calibrated on a simulated sample of $t\bar{t}$ were used in 2011. Most of the dataset was selected using the medium selection while the tight selection served as a backup in case of a significant rate underestimation. The b-jet triggers did not exhibit a nonlinearity in rate dependence a a function of instantaneous luminosity as they involve several pile-up robust reconstruction techniques and the sophisticated tagging itself involves association of tracks to the primary vertex. A set of calibration jet+μ triggers were maintained in 2011 and delivered sample of b-jets up to a very high jet p_{T}. The rates of this trigger were tuned by means of prescaling so that they deliver a flat coverage for a wide range of p_{T} as shown in Fig. 5.68. These triggers were actively used in the Pb-Pb data taking period.

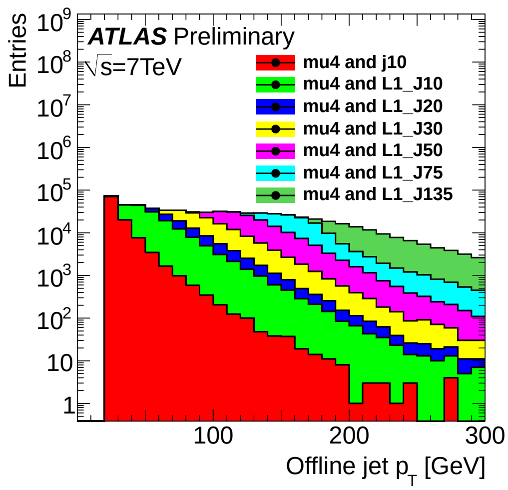

Figure 5.68: Spectra of jets containing muon collected by a various jet+μ triggers in 2011 [136, 135].

As mentioned before, in 2012 a more sophisticated tracking approach allowed for a better secondary vertex estimation and therefore more sophisticated taggers were used exploiting the impact parameters of tracks with respect to the primary vertex combined with the properties of the tracks associated with the secondary vertex. Their performance can be illustrated by the efficiency of a combined b-tagger at L2 and EF. In Fig. 5.69 the tagger weight, a quantity obtained from addition of the output of the two taggers and upon which the cut was applied, is shown. Working points, loose, medium and tight achieving 60%, 50% and 40% identifications efficiencies were set by applying cuts on the combined tag quantity at L2/EF or around 1.3/1.1, 3/3.1 and 3.5/4.5 respectively.

During the first year of Run 2 a setup identical to Run 1 was used after optimisation for the new conditions with efficiency points of 79%, 72% and 62% for loose, medium and tight. It is expected that a complete set of offline taggers combined with the multivariate technique (BDT) will be used in the future. The expected performance in comparison to Run 1 is shown in Fig. 5.70. These taggers will rely on secondary vertex tagging, secondary vertex properties (as in Run 1) and b- and c-hadron decays topologies. As can be seen a significant increase in identification efficiency is expected for a similar miss-tag rate. Further improvements are expected after incorporating the FTK tracking information [133].

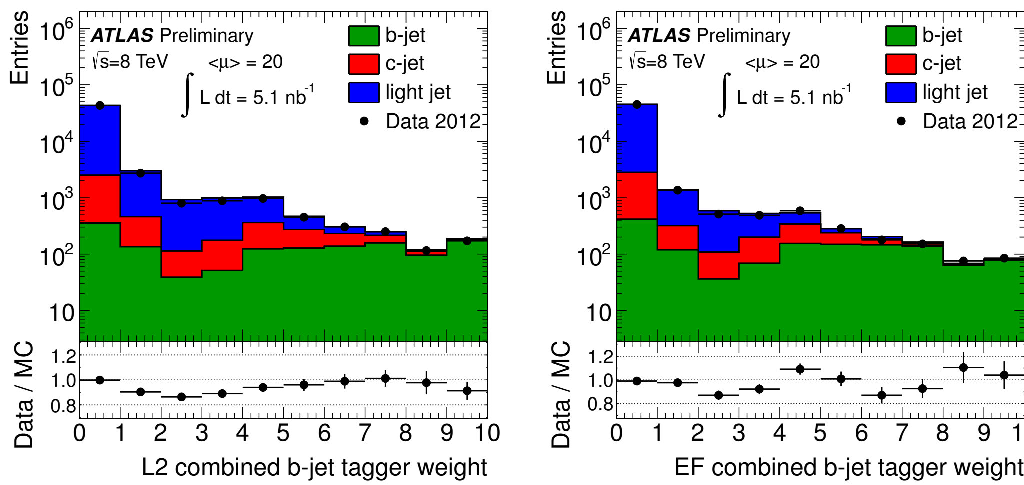

Figure 5.69: Performance of the combined b-jet taggers at L2(left) and EF(right) in 2012 [136, 137].

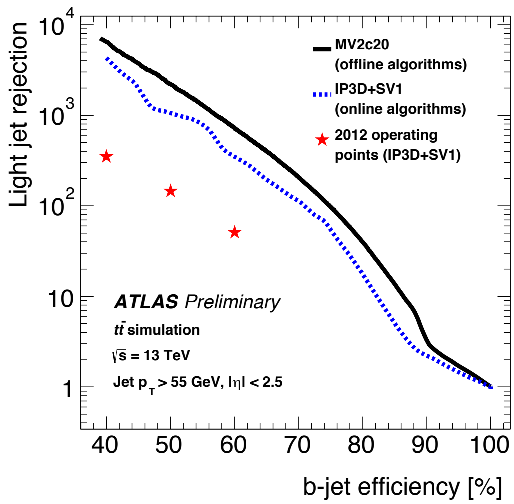

Figure 5.70: Comparison of 2012 performance of b-jet tagging online algorithms (stars) and expected performance in Run 2 for BDT combined Run 1 taggers (dotted) and the Run 2 taggers (trained with background of 80%, 20% light-,c-jets respectively). Performance is measured in $t\bar{t}$ sample for jets of $p_{\mathrm{T}} > 55$ GeV [133, 136].

5.5 Auxiliary functions of HLT

In addition to the primary role of HLT, that is to provide event filtering, it performs few additional tasks beneficial to the detector calibrations and overall performance of the experiment. This is predominantly due to the much higher number of events it sees as compared to the analyses performed offline. For that reason, typically the L2 is system, or early stages of HLT in Run 2 are charged with this task. Monitoring of the detector condition is one of these auxiliary tasks. The high rate of input events allow the HLT to see rarely occurring noise bursts or rarely occurring anomalies. The HLT system is capable of taking special actions in such occasions.

The DF system allows for event to be partially built. That is, only some of the detector data fragments are saved in the event record. This information, called Partial Event Building (PEB), is accompanied by a stream tag. In the HLT configuration the stream tag is assigned to an event if a related chain causes the event to be accepted. In order to combine both sets of information the PEB collection per chain, and so per stream is facilitated by the steering component. The PEB information is entered into a per-stream container and then shipped with the streaming decision. If partially built event is requested by the L2 the DF collects only detector fragments listed in PEB lists from the RoS system while in case of the EF the DF system role is to remove unnecessary parts from already built event. This actions are performed per output stream as it is possible that event is decided to be useful for physics (full event record) and calibration or monitoring at the same time.

5.5.1 Detectors calibration and monitoring

Calibration of the muon detector timing is a vital ingredient necessary for high precision data reconstruction. In order to achieve precision for each of the drift tubes a significantly higher number of muons as compared to those available in the main physics stream is necessary. For that reason a special collection system was developed to which the L2 MS algorithms communicate directly and it was filled with about a 10 times higher rate of detector data solely from the tubes in which muons were identified. In this approach the stream of data was significantly smaller and was handled by the control network of the DAQ system. The collected data was made available at a few calibration computing centers and up-to-date drift times were made available within 36 hours [138]. For Run 2 an approach relying more on the DF system was devised based on the partial event building mechanism mentioned above.

Another aspect, detector noise bursts, is explored in the calorimeter noise streams. Events in which the signal is recognised algorithmically as unphysical i.e. a significant deposit in a single isolated cell, would normally be filtered out. However, in order to enhance statistics for the noise studies it is beneficial to record events with such occurrences with a steady rate. Only detector fragments associated to the region where the noise was observed are recorded exploiting the PEB mechanism.

5.5.2 Online beam spot measurement

Last discussed here are the set of tools used for the online estimation of the beam-spot parameters. As mentioned in section 5.4.2, it is important to have an updated beam-spot parameter for online triggering. Another client of this information is the LHC operation team as the quick and precise beam measurements at experiments allows for optimisations of the beam optics and thus increased luminosities. An example of the beam-spot properties extracted online are shown in Fig. 5.71. The HLT configuration consists of several chains devoted to the task of extraction of the BS parameters. The tracking is performed first followed by vertex finding in this chains. Position of collision vertices averaged over many events provides an information on the luminous region. A fraction of most suitable events for vertex position finding is recorded using PEB mechanism. However, most importantly, the vertex information is entered in to the online monitoring system. After gathering the information from the entire HLT farm a precise information about the beam movements is available. In case of significant movements the up-to-date parameters are entered in to conditions database and fed back to be used by the HLT. Figure 5.71 shows updates occurring during the run due to either, change of position or widths. Most of the updates were caused by beam-spot movements along the y coordinate. Typically updates occur when positions change by 10% as compared to the respective width or widths change by 10%.

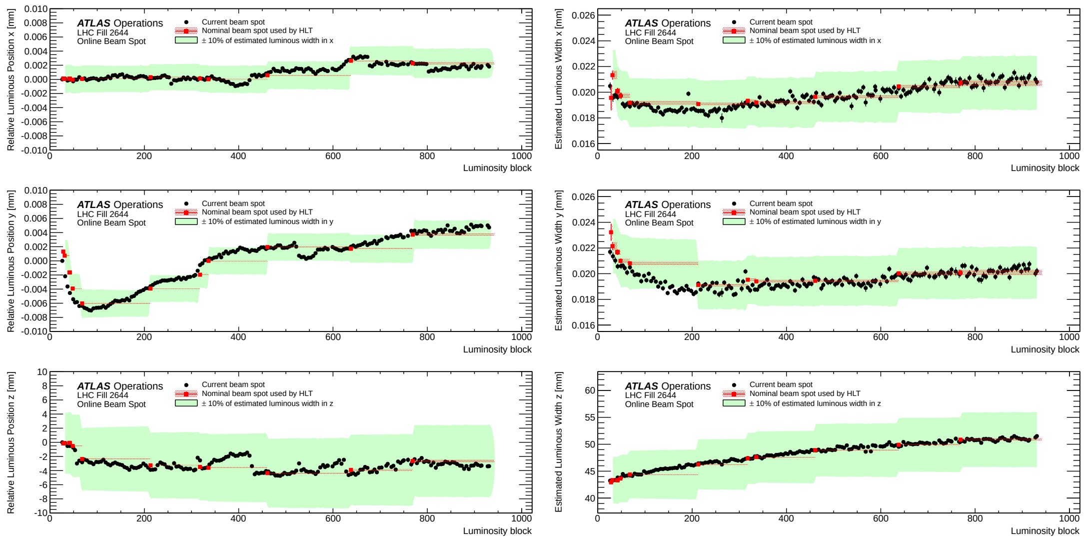

FIGURE 5.71: Beam-spot x, y, z (from top to bottom) positions (left) and widths (right) as a function of, about one minute long, luminosity block in one data taking run. With black points online measurements are shown and with red moments at which update of values used by HLT algorithms. The band represents condition of the update [139].

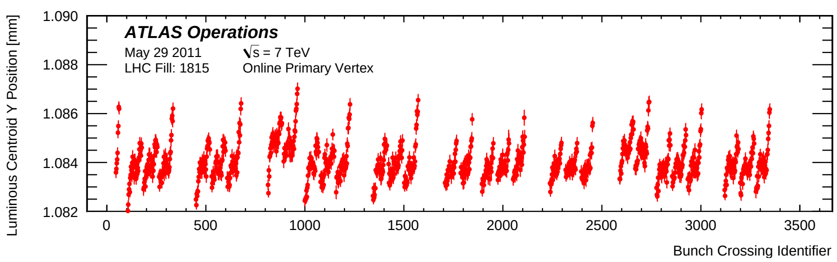

FIGURE 5.72: Position of the beam-spot along the y direction as a function ob bunch position along the beam orbit [140].

The BS parameters update is performed by signalling that new information is available to be used since a predefined luminosity block number. Data to signal that information is inserted in a CTP data fragment and reaches the HLT through the data path. As a consequence the HLT process, while unpacking L1 data, notices the need to re-read the beam-spot properties from the conditions database which is queried for the updated positions and widths of the BS. A tier like database, CORAL proxies, is put in place to assure that the update is completed within few milliseconds, sufficient not to cause a L2 processing timeout.

With high statistics, uniquely available at L2, the positions of the beam spot can be measured for specific bunch crossings. Figure 5.72 shows the result of that measurement. Small displacements, resulting from bunch-bunch interactions are observed thanks to the online beam spot infrastructure at L2 [140]. ‘

6

ATLAS HLT evolution beyond LHC Run 2

6.1 Motivation

The reason to upgrade the ATLAS trigger system for the high luminosity LHC (HL-LHC), starting in run 4, is to be able to fully profit from the increase of the luminosity delivered by the machine, as indicated in Table 4.4 in section 4.6. In particular, an increased luminosity, even neglecting pile-up effects, would lead to rates above the capability of the existing system. The most restrictive limit of the L1 output rate could only be met by increasing thresholds which would in turn mean a reduction of signal acceptances. In particular it was demonstrated that, for the flagship single lepton signatures, thresholds would need to increase by about 15 GeV. Depending on the analysis, that would mean drop of signal efficiency by approximately a factor of two. The effect can be even more pronounced for multi-lepton triggers depending on the studied channels. For instance the di-muon trigger efficiency for $h \rightarrow \tau\tau$[1] with both τ decaying to μ, would be reduced by a factor of about three. For photons, the need to preserve an inclusive photon trigger at the level of 200 GeV would consume a significant fraction of the 100kHz L1 bandwidth, not leaving room for other triggers. The hadronic τ trigger could only be used in combinations with the lepton triggers or forming di-τ signatures. An unmodified trigger system would require the increase of the di-τ thresholds from 40+30 GeV used in Run 1 and 2 to ~60+50 GeV for the HL-LHC thus significantly reducing signal acceptances. For instance in the h decay to hadronic τ this reduction would amount to 30%. Jets are not difficult to trigger on, however signatures involving many jets are not as straightforward because of the complex topologies involved and the rapid increase of the rates with pile-up. Efficient selection of b-tagged jets

[1] h is an additional Higgs bosson proposed as an extension ot the SM.

are useful to acquire datasets of t or h decays into $b\bar{b}$ ($h \to b\bar{b}$ BR =∼57% [141, 142]). The latter would allow for a new categories of searches exploring new particles coupling to the Higgs field [143]. In an unmodified trigger system the lowest b-jet threshold would need to increase from 50 GeV to about 100 GeV, thus limiting acceptance by about a factor two for such signals. It is also of critical importance to be able to trigger efficiently on $E_{\mathrm{T}}^{\mathrm{miss}}$ as this allows a unique possibility to select events containing weakly-interacting stable particles such as various dark matter candidates. The optimal $E_{\mathrm{T}}^{\mathrm{miss}}$ threshold varies depending on the scenario, however a reasonable value covering the needs of the various analyses lies around 150 GeV. Again, the limitation of the current system would push this threshold to around 300 GeV resulting in a significant reduction of signal efficiency [5].

6.2 Upgrade options

A general solution to all the issues mentioned above is to increase the L1 selectivity and to allow an increased output rate to the HLT[2] which in turn must also provide highly selective processing. For this reason, a major upgrade of the TDAQ system is planned after Run 3. It would involve all TDAQ system components and the detector readout system will be replaced. At the same time the ATLAS detector will undergo a major upgrade and the on-detector electronics will be replaced most of the detectors. As for the HLT, some aspects of upgrade will be implemented before Run 3. The selection algorithms and HLT framework interfaces are currently being re-worked, while throughout Run 3 and in the shutdown period before Run 4 effort will be concentrated on optimising the performance of the system.

For the hardware trigger two variants are being considered depending of the maximum readout rate that would be handled by the upgraded detectors [3]. One option is a two-stage system with a L0 that uses calorimeter and muon detector data in order to reduce the rate to 1 MHz. At this rate the Inner Tracker (ITk) data could be read out and hardware-based tracks reconstruction performed using a system similar to FTK. The next trigger level, L1 would be responsible for the formation of combined signatures and would thus be capable of reducing the event rate to 400 kHz. An alternative single staged system would consist of only the L0 and its output rate to HLT would be 1 MHz which is a 10-fold increase as compared to the current system.

Events selected by the hardware systems will be handled by the TDAQ data flow components shipping 2 TB/s (5 TB/s in case of single staged system) of data to the HLT. Another possibility being considered is to build in a significantly large intermediate buffering system (of size of up to ∼50 PB) which would decouple the HLT filtering from the active LHC cycle, thus allowing to optimal use of the HLT farm by extending processing into the inter-fill periods [5].

6.3 HLT upgrade

The HLT is intended to reduce the rate to 10 kHz from the input rate of 400 kHz (1 MHz) using a combination of trigger-specific fast algorithms and offline algorithms in the final stages of the selection.

[2]The HLT is renamed to Event Filter in the upgraded system, but for the sake of consistency it is still called the HLT in this text.

[3]It is not yet decided if the muon MDT and NSW on detector electronics will be upgraded to allow for 1 MHz readout assumed in one of the options.

Both, the jets and $E_{\mathrm{T}}^{\mathrm{miss}}$ signatures would require full-event tracking information in order to suppress the effects of pile-up by identifying the primary vertex and particles originating from it. Accordingly it is expected that full-event tracking will run with about 100 kHz and an even greater rate will be required for regional tracking. Therefore a hardware assisted tracking is planned to be integrated with the selection system. It might be the next generation of the FTK system but another upgrade option being studied is the possibility to use external accelerators such as the Graphical cards Processing Units (GPU) [144]. A demonstrator project was launched to implement tracking algorithms on GPUT. In addition GPU implementation of other time-consuming algorithms are being studied; tracking in the MS, calorimeter topological clustering and jet-finding [3]. In the approach taken in this project the algorithms scheduled by the framework are replaced by proxies which offload the computation to the GPU. Together with a feasibility assessment, the optimal balance between the number of CPUs and GPUs (in terms of throughput per Watt of consumed power) will be determined from the study.

The next generation HLT farm will be composed of processors with more compute power while the amount of operational memory is not expected to grow so significantly. The development at direction followed by the high performance computing industry suggests that the current ATLAS processing approach in which one (or 2 in case of hyper-threading enabled) process occupying a CPU core reconstructs a single event at a time will saturate and throughput increase could only be achieved by moving to a concurrent execution of algorithms in a multithreaded framework as a consequence allows to reduce need for the operational memory.

Therefore it is planned to redesign the HLT framework to allow for a straightforward integration of unmodified offline algorithms as well as to allow for concurrent processing of multiple events or multiple RoIs. The achievement of this goals is expected within a common framework for offline reconstruction and the HLT [4]. A feature allowing seamless use of offline algorithms at HLT is the *views* component which will allow either full or partial event data to be provided transparently as input to an offline algorithm thus removing the data access pattern dependence. By decomposition of the event into several views the concurrent reconstruction of a single event will be possible.

7

SUMMARY

This monograph describes the purpose, design, implementation and performance of the ATLAS High Level Trigger System. At the LHC, the full collision rate of 40 MHz can only be sampled by a hardware trigger, as is the case for ATLAS. The First Level Trigger output rate to the HLT is maximum of 100 kHz. The HLT provides further rate reduction to 200 Hz in 2010 growing gradually up to 1 kHz in 2015. The HLT selection process is composed of components written in C++ and executed on a farm of commodity computers.

A unique feature of the HLT is that it allows for about 3-5 times higher events processing rate than the bandwidth allowed by the detector readout system. It is possible, mainly thanks to two principles; RoI based reconstruction and early event rejection. In order to implement this functionality a sophisticated reconstruction approach was devised with a clear separation of algorithmic functions into reconstruction and decision-taking algorithms and a HLT steering framework to to control them. The early reconstruction steps are performed by specialised and fast trigger algorithms performing rejection before more precise, thus slower, offline or offline-like algorithms are used. Typically the HLT framework controls decision-taking for about a thousand (two thousand in Run 2) independent signatures, insulating the algorithms from the complexities of these signatures and the interrelations between them.

The HLT is equipped with a monitoring system necessary for online verification of the correctness of the selection process. For the saved events a record of the decision taking process is available allowing the extraction of trigger efficiencies from the data, thus contributing to the high precision results from analyses.

In order to optimise physics reach of the experiment the HLT allows for various reconstruction strategies. In particular a RoI based approach is implemented which is suitable for electrons, photons, muons and taus. There, the main role of the HLT framework is to handle regional reconstruction and the decision taking steps. In signatures where full-detector access is indispensable, such as $E_{\mathrm{T}}^{\mathrm{miss}}$ or min-bias the HLT only manages decision taking process. The most complex signatures, such as jets and b-jets combine both approaches to maximise rejection allowing for topological selections or a more precise extraction of event characteristics as compared to RoI-based reconstruction. All these use-cases are seamlessly integrated within the HLT framework that even provides for yet more specialised signatures.

The conditions of the HLT system operations will change significantly for LHC Run 4 starting in 2023 with input rates of 400 kHz or even 1 MHz depending on the upgrade scenario chosen. About a 10-fold rate reduction is still expected to be achieved at the HLT but, with a shorter processing time allotted per event and harsher pile-up conditions. It is clear that only precise offline algorithms are able to offer sufficient selectivity in these conditions. Therefore a significant upgrade is planned in order to prepare the system for high luminosity running. A key ingredient is the seamless integration of offline components while retaining the key trigger-specific features of RoI based reconstruction and early event rejection. Some aspects of the upgraded system are already under development and are expected to be ready for Run 3. The use of external co-processors dedicated to specific tasks is also being considered.

Bibliography

[1] T. Bold. Commissioning ATLAS Trigger. Number arXiv:0810.3269 in 34th International Conference on High Energy Physics Proceedings, 2008.

[2] M. zur Nedden, A. Sidoti, and T. Bold. The ATLAS Trigger Monitoring and Operation in proton-proton collisions. Technical Report ATL-DAQ-PROC-2010-006, CERN, 2010.

[3] A. Tavares Delgado, P. Conde Muiño, T. Bold, J. Augusto, S. Kama, M. Negrini, A. Sidoti, L. Rinaldi, S. Tupputi, J. Baines, M. Bauce, A. Messina, D. Emeliyanov, A. K. Elliott, Z. D. Greenwood, and S. Laosooksathit. An evaluation of GPUs for use in an upgraded ATLAS High Level Trigger. Technical Report ATL-DAQ-PROC-2015-061, CERN, 2015.

[4] J. Baines, T. Bold, P. Calafiura, S. Kama, C. Leggett, D. Malon, G. A. Stewart, and Wynne B. M. ATLAS Future Framework Requirements Group Report. Technical Report ATL-SOFT-PUB-2016-001, CERN, 2016.

[5] S. George, B. Petersen, E. Lipeles, M. Ishino, N. Konstantinidis, J. Adelman, D. Sankey, T. Pauly, J. Zhang, J. Baines, T. Bold, G. Lehmann Miotto, N. Gee, C. Bee, D. Francis, and D. Strom. Initial Design Review for the Phase-II Upgrade of the ATLAS TDAQ System. Technical Report ATL-COM-DAQ-2016-028, CERN, 2016. (ATLAS internal documentation).

[6] ATLAS Collaboration. Measurement of the centrality and pseudorapidity dependence of the integrated elliptic flow in lead-lead collisions at $\sqrt{s_{\mathrm{NN}}} = 2.76$ TeV with the ATLAS detector. *Eur. Phys. J.*, C74(8):2982, 2014.

[7] T. Bold. Results from lead-lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV with ATLAS at the LHC. In G. M. Bilei, S. Ciprini, and L. Fanò, editors, *Proceedings of the Physics at LHC 2011, Perugia, Italy, 2011*, volume eConf/C1106061, pages 59–62, 2012.

[8] T. Bold. Measurements of Elliptic Flow with the ATLAS Detector. *Nuclear Physics A*, 904:491c – 494c, 2013. Proceedings of Quark Matter Conference 2012.

[9] T. Bold. Soft Physics of p+Pb and Pb+Pb Collisions from the ATLAS Experiment at the LHC. *Acta. Phys. Pol. B*, 45(7):1331–1324, 2014.

[10] T. Bold. Heavy-ion measurements at ATLAS. In *Proceedings of the 51st Rencontres de Moriond on QCD and High Energy Interactions, La Thuile, Italy, 2016 (in preparation)*, 2016. Also available as tech. report ATL-PHYS-PROC-2016-043.

[11] Summary plots from the ATLAS Standard Model physics group. https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CombinedSummaryPlots/SM/.

[12] G. Antchev, P. Aspell, I. Atanassov, V. Avati, J. Baechler, et al. First measurement of the total proton-proton cross section at the LHC energy of $\sqrt{}$ s =7 TeV. *Europhys. Lett.*, 96:21002, 2011.

[13] ATLAS Collaboration. Measurements of the Total and Differential Higgs Boson Production Cross Sections Combining the $H \to \gamma\gamma$ and $H \to ZZ* \to 4l$ Decay Channels at $\sqrt{s}$=8 TeV with the ATLAS Detector. *Phys. Rev. Lett.*, 115(9):091801, 2015.

[14] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.

[15] S. Dasu et al. CMS. The TriDAS project. Technical design report, vol. 1: The trigger systems. Technical Report CERN-LHCC-2000-038, CERN, 2000.

[16] Jr. Alves, A. Augusto et al. The LHCb Detector at the LHC. *JINST*, 3:S08005, 2008.

[17] ALICE Collaboration. Performance of the ALICE Experiment at the CERN LHC. *Int.J.Mod.Phys.*, A29:1430044, 2014.

[18] ATLAS Collaboration. Performance of the atlas trigger system in 2010. *Eur. Phys. J. C*, 72(1):1–61, 2012.

[19] D. Casadei. Estimating the selection efficiency. *J. of Instr.*, 7(08):P08021, 2012.

[20] ATLAS Muon Collaboration. Atlas muon spectrometer, technical design report. *CERN/LHCC 97-22*, 1997.

[21] ATLAS Collaboration. ATLAS inner detector: Technical design report. Vol. 1. Technical Report CERN-LHCC-2010-013. ATLAS-TDR-19, CERN, 1997.

[22] ATLAS Collaboration. ATLAS inner detector: Technical design report. Vol. 2. Technical Report CERN-LHCC-2010-013. ATLAS-TDR-19, CERN, 1997.

[23] M. Capeans, G. Darbo, K. Einsweiller, M. Elsing, T. Flick, M. Garcia-Sciveres, C. Gemme, H. Pernegger, O. Rohne, and R. Vuillermet. ATLAS Insertable B-Layer Technical Design Report. Technical Report CERN-LHCC-2010-013. ATLAS-TDR-19, CERN, 2010.

[24] R. Bartoldus et al. Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System. Technical Report CERN-LHCC-2013-018. ATLAS-TDR-023, CERN, 2013.

[25] ATLAS TDAQ Collaboration. ATLAS level-1 trigger: Technical Design Report. Technical Report ATLAS-TDR-12 ; CERN-LHCC-98-014, CERN, 1998.

[26] M. Shochet, L. Tompkins, V. Cavaliere, P. Giannetti, A. Annovi, and G. Volpi. Fast TracKer (FTK) Technical Design Report. Technical Report CERN-LHCC-2013-007. ATLAS-TDR-021, CERN, 2013.

[27] R. Achenbach et al. The ATLAS Level-1 Calorimeter Trigger. *J. of Instr.*, 3(03):P03001, 2008.

[28] V. Lang and R. Stamen. Precision Synchronization of the ATLAS Level-1 Calorimeter Trigger with Collision Data in 2010 and 2011. Technical Report ATL-DAQ-PUB-2012-001, CERN, 2012.

[29] J. Bracinik. The performance of the atlas level-1 calorimeter trigger with lhc collision data. *J. of Instr.*, 6(01):C01078, 2011.

[30] ATLAS L1 Calo Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults.

[31] M. Palka. Operation of the enhanced ATLAS First Level Calorimeter Trigger at the start of LHC Run-2. Technical Report ATL-DAQ-PROC-2015-049, CERN, 2015.

[32] A. Khomich. Upgrade of the preprocessor system for the ATLAS level-1 calorimeter trigger. *JINST*, 5:C12047, 2010.

[33] M. Wessels. Upgrade of the ATLAS Level-1 Calorimeter Trigger. Technical Report ATL-DAQ-PROC-2014-003, CERN, 2014.

[34] ATLAS Collaboration. Performance of the ATLAS Muon Trigger in p-p collisions at $\sqrt{s} = 7TeV$. Technical Report ATLAS-CONF-2010-095, CERN, 2010.

[35] ATLAS Collaboration. Performance of the ATLAS muon trigger in 2011. Technical Report ATLAS-CONF-2012-099, CERN, 2012.

[36] ATLAS Collaboration. Performance of the ATLAS muon trigger in pp collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J.*, C75:120, 2015.

[37] ATLAS Muon Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MuonTriggerPublicResults.

[38] M. Stockton. The ATLAS level-1 central trigger. *J. Phys. Conf. Ser.*, 331:022041, 2011.

[39] ATLAS Trigger Operation Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults.

[40] J. Glatzer. Operation of the Upgraded ATLAS Level-1 Central Trigger System. *J. of Phys. Conf. Ser.*, 664(8):082013, 2015.

[41] ATLAS Collaboration. Measurement of long-range pseudorapidity correlations and azimuthal harmonics in $\sqrt{s_{NN}} = 5.02$ TeV proton-lead collisions with the ATLAS detector. *Phys. Rev.*, C90(4):044906, 2014.

[42] V. Lendermann, J. Haller, M. Herbst, K. Krüger, H-C. Schultz-Coulon, and R. Stamen. Combining Triggers in HEP data analysis. *Nucl. Instr. and Met. A*, 604(3):707 – 718, 2009.

[43] ATLAS Collaboration. Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data. *Eur. Phys. J.*, C72:1909, 2012.

[44] J. L. Gustafson. Reevaluating Amdahl's Law. *Commun. ACM*, 31(5):532–533, 1988.

[45] M. Frank, C. Gaspar, E. v Herwijnen, B. Jost, and N. Neufeld. Deferred High Level Trigger in LHCb: A Boost to CPU Resource Utilization. *J. of Phys.: Conf. Ser.*, 513(1):012006, 2014.

[46] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino, and D. Quarrie. The athena control framework in production, new developments and lessons learned. *Computing in high energy physics and nuclear physics. Proceedings, Conference, CHEP'04, Interlaken, Switzerland*, pages 456–458, 2005.

[47] R. Ospanov. Resource Utilization by the ATLAS High Level Trigger during 2010 and 2011 LHC Running. *Physics Procedia*, 37(0):1900 – 1907, 2012. Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011).

[48] J. Stelzer. The ATLAS High Level Trigger Configuration and Steering: Experience with the First 7 TeV Collision Data. *Journal of Physics: Conference Series*, 331(2):022026, 2011.

[49] E.R. Gansner, E. Koutsofios, S.C. North, and K.-P. Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–230, 1993.

[50] ATLAS Collaboration. Search for massive long-lived highly ionising particles with the ATLAS detector at the LHC. *Phys. Lett. B*, 698(5):353 – 370, 2011.

[51] ATLAS Collaboration. Search for Magnetic Monopoles in $\sqrt{s}$=7 TeV pp Collisions with the ATLAS Detector. *Phys. Rev. Lett.*, 109:261803, 2012.

[52] V. Bartsch. Experience with the custom-developed ATLAS Offline Trigger Monitoring Framework and Reprocessing Infrastructure. Technical Report ATL-DAQ-PROC-2012-040, CERN, 2012.

[53] D. Oliveira Damazio. Data Quality Monitoring for the ATLAS trigger System during the first data taking period of the Large Hadron Collider. Technical Report ATL-DAQ-PROC-2013-016, CERN, 2013.

[54] C Cuenca Almenar. ATLAS Online Data Quality Monitoring. Technical Report ATL-DAQ-PROC-2010-015, CERN, 2010.

[55] F. Winklmeier. Real-time configuration changes of the ATLAS High Level Trigger. Technical Report ATL-DAQ-PROC-2010-009, CERN, 2010.

[56] P. Bell, D. Berge, S. Brunet, G. Fischer, M. Goebel, J. Haller, S. Head, A. Höcker, T. Kono, A. Martyniuk, M. Nozicka, M. Owen, R. Spiwoks, J. Stelzer, T. Wengler, and W. Wiedenmann. The Configuration System of the ATLAS Trigger. Technical Report ATL-DAQ-PROC-2009-047, CERN, 2009.

[57] T. Perez Cavalcanti. The ATLAS High Level Trigger Steering Framework and the Trigger Configuration System. Technical Report ATL-DAQ-PROC-2011-016, CERN, 2011.

[58] R. M. Bianchi. Recent updates of the Control and Configuration of the ATLAS Trigger and Data Acquisition System. Technical Report ATL-DAQ-PROC-2011-031, CERN, 2011.

[59] A. Tavares Delgado. Design and performance of the ATLAS jet trigger system. Technical Report ATL-DAQ-PROC-2015-019, CERN, 2015.

[60] C. Schiavi. ATLAS High-Level Trigger algorithms for Run-2 data-taking. Technical Report ATL-DAQ-PROC-2015-018, CERN, 2015.

[61] ATLAS Collaboration. Expected performance of the ATLAS experiment: detector, trigger and physics. Technical Report CERN-OPEN-2008-020, CERN, 2009.

[62] ATLAS HLT Tracking Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/HLTTrackingPublicResults.

[63] ATLAS luminosity Public plots.
https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults.

[64] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys.Lett.*, B716:1–29, 2012.

[65] ATLAS Collaboration. Search for high-mass diphoton resonances in pp collisions at $\sqrt{s}$ = 8 TeV with the ATLAS detector. *Phys. Rev.*, D92(3):032004, 2015.

[66] ATLAS Collaboration. Measurement of the production cross section of an isolated photon associated with jets in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Phys. Rev.*, D85:092014, 2012.

[67] ATLAS Collaboration. Expected photon performance in the ATLAS experiment. Technical Report ATL-PHYS-PUB-2011-007, CERN, 2011.

[68] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, S. Menke, S. Rajagopalan, D. Rousseau, S. Snyder, and G. Unal. Calorimeter clustering algorithms: Description and performance. Technical Report ATL-LARG-PUB-2008-002, CERN, 2008.

[69] ATLAS Egamma Trigger Public Results.
https://twiki.cern.ch/twiki/bin/view/AtlasPublic/EgammaTriggerPublicResults.

[70] G. Pasztor. The Upgrade of the ATLAS Electron and Photon Triggers towards LHC Run 2 and their Performance. Technical Report ATL-DAQ-PROC-2015-053. arXiv:1511.00334, CERN, 2015.

[71] HLT Tracking Public Results.
https://twiki.cern.ch/twiki/bin/view/AtlasPublic/HLTTrackingPublicResults, 2015.

[72] A. J. Washbrook. The Performance and Development of the ATLAS Inner Detector Trigger. Technical Report ATL-DAQ-PROC-2013-045, CERN, 2013.

[73] O. Penc. The design and performance of the ATLAS Inner Detector trigger for Run 2. Technical Report ATL-DAQ-PROC-2015-058, CERN, 2015.

[74] ATLAS Collaboration. Expected electron performance in the ATLAS experiment. Technical Report ATL-PHYS-PUB-2011-006, ATL-COM-PHYS-2010-990, CERN, 2011.

[75] ATLAS Collaboration. Performance of the ATLAS Electron and Photon Trigger in p-p Collisions at $\sqrt{s} = 7$ TeV in 2011. Technical Report ATLAS-CONF-2012-048, CERN, 2012.

[76] ATLAS Collaboration. Performance of the Electron and Photon Trigger in p-p Collisions at $\sqrt{s} = 7$ TeV. Technical report, CERN, 2011.

[77] ATLAS Collaboration. Measurement of the centrality dependence of J/ψ yields and observation of Z production in lead–lead collisions with the ATLAS detector at the LHC. *Phys. Lett.*, B697:294–312, 2011.

[78] ATLAS Collaboration. Measurement of the Z to $\tau\tau$ Cross Section with the ATLAS Detector. *Phys. Rev.*, D84:112006, 2011.

[79] ATLAS Collaboration. Measurement of the top quark pair cross section with ATLAS in pp collisions at $\sqrt{s} = 7$ TeV using final states with an electron or a muon and a hadronically decaying τ lepton. *Phys. Lett.*, B717:89–108, 2012.

[80] ATLAS Collaboration. Search for microscopic black holes in a like-sign dimuon final state using large track multiplicity with the ATLAS detector. *Phys. Rev.*, D88(7):072001, 2013.

[81] ATLAS Collaboration. J/ψ Performance of the ATLAS Inner Detector. Technical Report ATLAS-CONF-2010-078, CERN, 2010.

[82] ATLAS Collaboration. Alignment of the ATLAS Inner Detector and its Performance in 2012. Technical Report ATLAS-CONF-2014-047, CERN, 2014.

[83] ATLAS Collaboration. Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS. *Eur. Phys. J.*, C72:1844, 2012.

[84] F. A. Conventi. A measurement of the ATLAS di-muon trigger efficiency in proton-proton collisions at $\sqrt{s}$= 7 TeV. Technical Report ATL-DAQ-PROC-2011-044, CERN, 2011.

[85] ATLAS B Physics Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BPhysicsTriggerPublicResults.

[86] ATLAS Collaboration. Measurement of τ polarization in $W- > \tau\nu$ decays with the ATLAS detector in pp collisions at sqrt(s) = 7 TeV. *Eur. Phys. J.*, C72:2062, 2012.

[87] ATLAS Collaboration. Evidence for the Higgs-boson Yukawa coupling to tau leptons with the ATLAS detector. *JHEP*, 04:117, 2015.

[88] ATLAS Collaboration. Search for charged Higgs bosons decaying via $H^{\pm} \to \tau^{\pm}\nu$ in fully hadronic final states using pp collision data at $\sqrt{s}$ = 8 TeV with the ATLAS detector. *JHEP*, 03:088, 2015.

[89] ATLAS Collaboration. Search for supersymmetry in events with large missing transverse momentum, jets, and at least one τ lepton in 20 fb^{-1} of $\sqrt{s}$ = 8 TeV proton-proton collision data with the ATLAS detector. *JHEP*, 09:103, 2014.

[90] ATLAS Collaboration. Search for the direct production of charginos, neutralinos and staus in final states with at least two hadronically decaying taus and missing transverse momentum in pp collisions at $\sqrt{s}$ = 8 TeV with the ATLAS detector. *JHEP*, 10:096, 2014.

[91] ATLAS Tau Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TauTriggerPublicResults.

[92] ATLAS Collaboration. Performance of the ATLAS tau trigger in 2011. Technical Report ATLAS-CONF-2013-006, CERN, 2013.

[93] ATLAS Collaboration. Performance of the Reconstruction and Identification of Hadronic Tau Decays with ATLAS. Technical Report ATLAS-CONF-2011-152, CERN, 2011.

[94] M. P. Casado. Performance of the ATLAS Tau Trigger system with 7 TeV pp collisions at the LHC. Technical Report ATL-DAQ-PROC-2011-037, CERN, 2011.

[95] C. Cuenca Almenar. Triggering on hadronic tau decays in ATLAS: semiconductor tracking detectors in action. Technical Report ATL-DAQ-PROC-2012-003, CERN, 2012.

[96] A J Tanasijczuk. The ATLAS hadronic tau trigger. Technical Report ATL-DAQ-PROC-2012-017, CERN, Geneva, Jun 2012.

[97] P. Czodrowski. Triggering on hadronic Tau Decays in ATLAS: Algorithms and Performance. Technical Report ATL-DAQ-PROC-2012-039, CERN, 2012.

[98] J. Mahlstedt. The ATLAS Hadronic Tau Trigger. Technical Report ATL-DAQ-PROC-2013-027, CERN, 2013.

[99] A. Karamaoun. The Upgrade and Performance of the ATLAS Tau Trigger towards Run 2. Technical Report ATL-DAQ-PROC-2015-040, CERN, 2015.

[100] ATLAS Collaboration. Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC. *New J. Phys.*, 13:053033, 2011.

[101] ATLAS Collaboration. Charged particle multiplicities in pp interactions at sqrt(s) = 7 TeV measured with the ATLAS detector at the LHC. Technical Report ATLAS-CONF-2010-024, CERN, 2010.

[102] ATLAS Collaboration. Charged particle multiplicities in pp interactions for track $p_T >$ 100 MeV at sqrts = 0.9 and 7 TeV measured with the ATLAS detector at the LHC. Technical Report ATLAS-CONF-2010-046, CERN, 2010.

[103] ATLAS Collaboration. Charged-particle multiplicities in pp interactions at $\sqrt{s} = 900$ GeV measured with the ATLAS detector at the LHC. *Phys. Lett.*, B688:21–42, 2010.

[104] ATLAS Collaboration. Performance of the ATLAS Minimum Bias and Forward Detector Triggers in pPb collisions. Technical Report ATLAS-CONF-2013-104, CERN, 2013.

[105] ATLAS Collaboration. Performance of the ATLAS Minimum Bias and Forward Detector Triggers in 2011 Heavy Ion Run. Technical Report ATLAS-CONF-2012-122, CERN, 2012.

[106] ATLAS Collaboration. Performance of the Minimum Bias Trigger in p-p Collisions at Sqrt(s) = 900 GeV. Technical Report ATLAS-CONF-2010-025, CERN, 2010.

[107] ATLAS Collaboration. Performance of the Minimum Bias Trigger in p-p Collisions at Sqrt(S) = 7 TeV. Technical Report ATLAS-CONF-2010-068, CERN, 2010.

[108] ATLAS Collaboration. Measurement of the $W \to \ell\nu$ and $Z/\gamma^* \to \ell\ell$ production cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *JHEP*, 12:060, 2010.

[109] ATLAS Collaboration. Measurement of the production cross section for W-bosons in association with jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Phys. Lett.*, B698:325–345, 2011.

[110] ATLAS Collaboration. Search for diphoton events with large missing transverse momentum in 7 TeV proton-proton collision data with the ATLAS detector. *Phys. Lett.*, B718:411–430, 2012.

[111] D. et al Casadei. The implementation of the ATLAS missing E_T triggers for the initial LHC operation. Technical Report ATL-DAQ-PUB-2011-001, CERN, 2011.

[112] ATLAS Collaboration. Performance of the ATLAS transverse energy triggers with initial LHC runs at sqrt(s) = 7 TeV. Technical Report ATLAS-CONF-2011-072, CERN, 2011.

[113] ATLAS Collaboration. The ATLAS transverse-momentum trigger performance at the LHC in 2011. Technical Report ATLAS-CONF-2014-002, CERN, 2014.

[114] ATLAS Missing Et Trigger Public Results.
https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MissingEtTriggerPublicResults.

[115] ATLAS Collaboration. Search for a CP-odd Higgs boson decaying to Zh in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2016-015, CERN, 2016.

[116] A. Mincer, F. Urs Bernlochner, A. Struebig, J. Schouwenberg, and J. B. Beacham. The ATLAS Transverse Momentum Trigger at the LHC. Technical Report ATL-DAQ-PROC-2015-028, CERN, 2015.

[117] A. Strubig. The ATLAS transverse momentum trigger evolution at the LHC towards Run II. Technical Report ATL-DAQ-PROC-2015-031, CERN, Geneva, Oct 2015.

[118] ATLAS Collaboration. Measurement of inclusive jet and dijet cross sections in proton-proton collisions at 7 TeV centre-of-mass energy with the ATLAS detector. *Eur. Phys. J.*, C71:1512, 2011.

[119] ATLAS Collaboration. Measurement of the top-quark mass in the fully hadronic decay channel from ATLAS data at $\sqrt{s} = 7\,\text{TeV}$. *Eur. Phys. J.*, C75(4):158, 2015.

[120] ATLAS Collaboration. Search for New Particles in Two-Jet Final States in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC. *Phys. Rev. Lett.*, 105:161801, 2010.

[121] ATLAS Collaboration. Search for quark contact interactions in dijet angular distributions in pp collisions at $\sqrt{s} = 7$ TeV measured with the ATLAS detector. *Phys. Lett.*, B694:327–345, 2011.

[122] M. Cacciari, G. P. Salam, and G. Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 04:063, 2008.

[123] G. P. Salam and G. Soyez. A Practical Seedless Infrared-Safe Cone jet algorithm. *JHEP*, 05:086, 2007.

[124] M. Campanelli. Evolution and performance of the ATLAS jet trigger. Technical Report ATL-DAQ-PROC-2013-046, CERN, 2013.

[125] ATLAS Jet Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/JetTriggerPublicResults.

[126] ATLAS Collaboration. Measurement of the jet radius and transverse momentum dependence of inclusive jet suppression in lead-lead collisions at $\sqrt{s_{NN}}$= 2.76 TeV with the ATLAS detector. *Phys. Lett.*, B719:220–241, 2013.

[127] ATLAS Collaboration. Performance of the ATLAS Jet Trigger in the Early sqrts=7 TeV Data. Technical Report ATLAS-CONF-2010-094, CERN, 2010.

[128] S. Shimizu. The ATLAS jet trigger performance in LHC Run I and Run II updates. *J. of Phys. Conf. Ser.*, 664(8):082051, 2015.

[129] ATLAS Collaboration. Measurement of the inclusive and dijet cross-sections of b^- jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Eur. Phys. J.*, C71:1846, 2011.

[130] ATLAS Collaboration. Searches for Higgs boson pair production in the $hh \rightarrow bb\tau\tau, \gamma\gamma WW^*, \gamma\gamma bb, bbbb$ channels with the ATLAS detector. *Phys. Rev.*, D92:092004, 2015.

[131] ATLAS Collaboration. Search for Higgs boson pair production in the $b\bar{b}b\bar{b}$ final state from pp collisions at $\sqrt{s} = 8$ TeVwith the ATLAS detector. *Eur. Phys. J.*, C75(9):412, 2015.

[132] ATLAS Collaboration. Performance of b-Jet Identification in the ATLAS Experiment. *JINST*, 11(04):P04008, 2016.

[133] L. Živkovič. Real-time flavour tagging selection in ATLAS. Technical Report ATL-DAQ-PROC-2015-048, CERN, 2015.

[134] D. Madaffari. Real-time flavor tagging selection in ATLAS. Technical Report ATL-DAQ-PROC-2014-036, CERN, 2014.

[135] A. Coccaro. Track Reconstruction and b-Jet Identification for the ATLAS Trigger System. Technical Report arXiv:1112.0180. ATL-DAQ-PROC-2011-051, CERN, 2011.

[136] ATLAS BJet Trigger Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BJetTriggerPublicResults.

[137] C. Bertella. Real-time flavour tagging selection in ATLAS. *J. Phys. Conf. Ser.*, 664(8):082005, 2015.

[138] E. Diehl. The ATLAS Muon Calibration Stream: Design and Performance. Technical Report ATL-DAQ-PROC-2012-059, CERN, 2012.

[139] ATLAS Online Beam Spot Public Results. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/OnlineBeamSpotPublicResults.

[140] R. Bartoldus. Online determination of the lhc luminous region with the atlas high-level trigger. *Physics Procedia*, 37:2080 – 2088, 2012. Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011).

[141] ATLAS Collaboration CMS Collaboration. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. *Submitted to JHEP [arXiv: 1606.02266]*, 2016.

[142] ATLAS Collaboration. Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s}$=7 and 8 TeV in the ATLAS experiment. *Eur. Phys. J.*, C76(1):6, 2016.

[143] ATLAS Collaboration. Search for pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Submitted to Phys. Rev. D [arXiv: 1606.04782]*, 2016.

[144] S. Kama. Triggering events with GPUs at ATLAS. Technical Report ATL-DAQ-PROC-2015-013, CERN, 2015.