The ATLAS detector at the LHC

Andrée Robichaud-Véronneau*
on behalf of the ATLAS collaboration
Université de Genève

July 17th, 2009

Abstract

The world's largest multi-purpose particle detector, ATLAS, is fully installed underground and in operation at CERN. Details of the different subdetectors are presented, together with expected performance aspects. The performance achievements obtained during the two real data-taking periods of ATLAS, cosmic rays and single beam runs of the LHC in fall 2008, are outlined.

1 Introduction

The ATLAS (A Toroidal Lhc ApparatuS) detector is one of the four experiments which will measure p-p (and Pb ion) collisions at the Large Hadron Collider (LHC). The LHC is a 26.7 km ring which will provide collisions at a nominal center-of-mass energy of 14 TeV for protons (1150 TeV for Pb ions) with a design luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$. 2808 bunches, each containing $\text{O}(10^{11})$ protons will collide every 25 ns (bunch crossing frequency of 40 MHz), resulting in a pile-up of $\sim$20 events at peak luminosity. This is an important challenge for the hardware and data acquistion of the four LHC experiments. With this unprecedented energy, the LHC detectors will be used to probe for new physics beyond the Standard Model.

ATLAS is a multi-purpose particle physics detector, which means it should be able to measure the signatures of all the possible final states

*E-mail:andree.robichaud-veronneau@cern.ch

we expect to observe from proton-proton (or heavy ion) collisions. They include mainly: tracks from charged particles (hadrons, electrons, muons), energy deposition of electromagnetic and hadronic matter in the calorimeter system (through showering of the particles in many interaction length of material) and tracks from muons escaping the detector. The two main tracking devices, the Inner Detector and the Muon Spectrometer lie inside magnetic fields to allow the measurement of charge and momentum using the tracks. The passage of neutrinos through the detector can be measured using the missing transverse energy in the hermetically closed detector volume. A detailed description of the detector and its performance is presented here.

2 Description of the detector

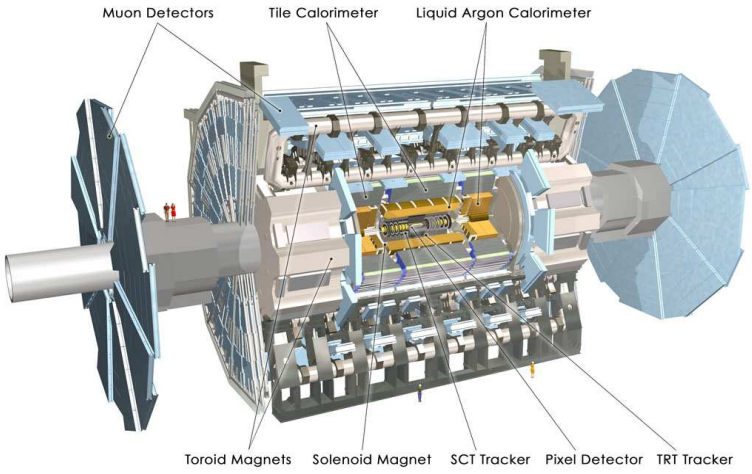

Figure 1: The ATLAS detector

The ATLAS detector [1] uses a cylindrical geometry with one end-cap on each side to ensure full coverage in solid angle, as can be seen in Fig. 1. It is divided into three main parts, each of them comprising several subdetectors. The Inner Detector uses silicon and straw tube technologies for tracking; the energy in the Calorimeter is measured with Liquid Argon (LAr) and Tile subdetectors, and finally, the Muon spectrometer is made of various drift chambers to record the muon tracks. The details for each subdetector are given below.

2.1 Inner Detector

The Inner Detector (ID) is the closest to the beampipe. It spans 2.3 m in height and 7.0 m in length. It is made of three subdetectors: the Pixel detector, the SemiConductor Tracker (SCT) and the Transition Radiation Tracker (TRT).

The Pixel detector is the innermost detector in the ID. It is made of three layers of silicon modules in the barrel and three wheels in each end-cap. The innermost layer is located directly on the beampipe, at 5.05 cm from the Interaction Point (IP) and the outermost layer, at 12.25 cm. It is made of 1744 identical modules. Each module is built from a 250 μm silicon sensor divided into pixels of 50 x 400 μm in size (50 x 600 μm on the edge) which is bump-bonded to the chips. The modules are arranged in ladders of 13 modules each in the barrel region (module plane parallel to the beam) and in wheels of 48 modules each (module plane perpendicular to the beam). About 80 million channels are necessary to read out the full detector. The Pixel then provides three absolute position measurements along the charged tracks and is particularly useful to determine the position of secondary decay vertices in the event.

The SCT is located around the Pixel detector, from R=25.5 cm to R=54.9 cm in the barrel and R=25.1 cm to R=61.0 cm in the end-cap. It uses silicon strips with 80 μm pitch assembled in double-sided modules with a stereo angle of 40 mrad between the two sides. Four different geometries for the 4088 modules are used: rectangular module in the barrel and three wedge-shaped module geometries in the end-caps, suited for the three different rings of modules on the end-cap disk: inner, middle and outer. SCT disks are not identical: the number of rings varies between one to three, starting from the outer ring. For all modules, strips are wire-bonded to the chips. The thickness of the silicon sensor in the SCT is 285$\pm$15 μm. This accounts for about 6 million read-out channels in the detector and covers a surface of silicon of 63 m^2, making the SCT one of the largest existing silicon detectors.

As both detectors lie close to the beamline and make use of semiconductor technology, they will suffer from radiation damage. The innermost layer of the Pixel detector, called *b layer*, is planned to be replaced after three years of LHC operation at nominal luminosity. The expected neutron equivalent fluence F_{neq} for the two other layers is about 8x10^{14} cm^{-2}, so they will be used for the full lifetime of the LHC program. The radiation effects are not limited to the Pixel, and the innermost layer of

the SCT is expected to receive a fluence F_{neq}=$\sim$2x10^{14} cm^{-2}. As a result, the silicon sensors were chosen for their radiation-hardness properties. In the Pixel detector, oxygenated n-type silicon bulk with n^+ implants for pixel elements was used, while for the SCT, p-in-n bulk with AC coupled STRIPS were used. In addition, the sensors have to be kept cold (-7 ° C for the SCT and -20° C for the Pixel nominally) to reduce leakage current after irradiation. Furthermore, their bias voltage will increase with time (from 150 V to 500 V [600 V] for the SCT [Pixel]) to compensate for type inversion in the bulk.

The ID is equipped with a complex environment monitoring system, which controls the gas atmosphere inside the detector volumes, and an evaporative cooling system. To avoid condensation and risk of fire, the two silicon detectors are filled with N_2 while the TRT is filled with CO_2 to avoid contamination from the N_2. A two-phase evaporative cooling system is used in common for both Pixel and SCT (the TRT operates at room temperature) to remove about 85 kW of heat from the ID volume. It uses C_3F_8 as a coolant, which undergoes pressure drops along the pipe and is then brought to boil inside the detector. Counter-flow heat exchangers are installed at the input pipes to cool down the incoming coolant with the outgoing coolant (subcooling). Finally, a heater system is installed to boil the remaining liquid in the output pipes to avoid condensation along the calorimeter on their way back to the cooling plant.

The last layer of the ID consists of the Transition Radiation Tracker (TRT). It is a straw tube detector interleaved with thin foil which provides transition radiation photons to be detected by the straw tubes. In the barrel, the detector is divided into 32 modules composed of three rings and the straws are parallel to the beam axis. In each end-cap, 20 wheels are assembled together, in which the straws are perpendicular to the beam axis. The 4 mm straw tube forms the cathode and a 31 μm gold-plated tungsten wire forms the grounded anode. The cathode is kept at -1530 V. Inside the straws, a mixture of 70% Xe, 27% CO_2 and 3%O_2 is used to be ionized by the transition radiation photon. 351 000 channels are read-out for the entire detector. The TRT provides R-ϕ information and adds $\sim$36 space points to the track (22 in the end-cap).

The ID is enclosed in a superconducting Solenoid to provide the 2 T magnetic field needed for tracking. In order to reduce the material in front of the calorimeter, the solenoid is enclosed in the same cryostat as the Electromagnetic Calorimeter (see section 2.2).

2.2 Calorimeter

Once they have crossed the tracking system, the neutral and charged particles will reach the calorimeters, where they will shower and hence their energy will be deposited and measured in the detector. In this process, electromagnetic and hadronic matter behaves differently and needs to be treated by two separate calorimeter systems.

The Electromagnetic (EM) Calorimeter is made of lead as absorber material and copper-kapton electrodes arranged in an accordion geometry. The accordion is kept in a cold Liquid Argon (LAr) vessel, which serves as the active material. Four individual vessels form the detector: two half-barrels and two end-caps, covering up to pseudorapidity[1] $|\eta| < 3.2$. There is a small gap at $|\eta| = 0$ and one at $1.37 < |\eta| < 1.52$. The calorimeter is divided into three radial layers (two in certain regions of pseudorapidity) and has a decreasing granularity as a function of distance from the beampipe. The cell granularity varies also in $|\eta|$ from $\Delta\eta$ x $\Delta\phi =$ 0.0031 x 0.1 to 0.1 x 0.1. A presampler is installed to determine the energy lost in the material in front of the EM calorimeter in the range $|\eta| < 1.8$.

The Hadronic Calorimeter (HCAL) uses two different technologies to absorb the hadrons from the collisions. At $1.5 < |\eta| < 3.2$, the Hadronic End-Caps (HEC) use LAr technology and are located inside the same cold vessels as the EM Calorimeter. The HEC is composed of two wheels, each equipped with 32 wedge-shaped modules, using copper plates as absorber material. Surrounding the EM and HEC calorimeters, the Tile Calorimeter (TileCal) uses steel plates as absorber material interleaved with scintillator tiles as active material. The light from the scintillating tiles is read out by wavelength shifting fibres and 9852 photomultiplier tubes (PMTs). It is divided into two parts: the Barrel covering the range $0 < |\eta| < 1.0$ and the Extended Barrel covering $0.8 < |\eta| < 1.7$. They are built of 64 wedge-shaped modules of size $\Delta\phi \sim 0.1$ rad. The gap between the Barrel and Extended Barrel is equipped with plug calorimeters to determine the amount of energy lost in this region.

Finally, a forward calorimeter (FCal) system is needed to cover the pseudorapidity range up to 4.9. The FCal is made of three layers: one for EM particle detection using copper plates as absorber and two more for hadronic particle detection using Tungsten plates as absorber. The detector is housed in the LAr vessel. It uses smaller LAr gaps between the

[1] Pseudorapidity is defined as $\eta = -\ln(\tan(\theta/2))$ for massless objects

plates to limit ion formation due to the high particle flux in this region.

2.3 Muon Spectrometer

Due to the high particle multiplicity in ATLAS, muons are a very important handle for the trigger system (see section 2.5). The Muon Spectrometer is then designed with two separate sets of detectors: trigger chambers (fast) and precision chambers (slow). They are enclosed in the toroidal magnet system of ATLAS which provides strong bending power in a large empty volume, allowing for a minimization of multiple scattering effects and a good charge identification for high p_T muons. The precision chambers allow one to have good momentum resolution without loss of trigger efficiency.

The toroidal magnet system is the skeleton of ATLAS, hence determines the size of the detector. It is composed of three toroids: 1 barrel and 2 end-caps. Each toroid consists of 8 superconducting coils, equally separated in azimuth. The end-caps have a 22.5° angle with the barrel toroid to provide radial overlap and optimise the bending power in the transition region between barrel and end-cap. The magnetic field strength is approximately 0.5 T in the barrel and 1 T in the end-cap.

The precision tracking muon system is composed of two subdetectors: the Monitored Drift Tubes (MDT) and the Cathode Strip Chambers (CSC). The MDT occupies most of the solid angle, with three concentric layers in the barrel and up to $|\eta|$=2.7 in the end-cap small and big wheels (located on each side of the end-cap toroid). A chamber consists of 3 to 8 layers of drift tube filled with gas. It has a maximum counting rate of 500 Hz/cm^2. The CSC covers the range $2 < |\eta| < 2.7$ only for the innermost small wheel. This is needed due to the high particle flux in this region of high pseudorapidity. The CSC is a multiwire proportional chamber which uses two perpendicularly segmented cathode planes to locate particles. Its maximum counting rate is 1000 Hz/cm^2. Both detectors use a mixture of Ar and CO_2 gaz inside the chambers for particle detection.

The trigger chambers are also divided into two subdetectors: the Resistive Plate Chambers (RPC) in the barrel and the Thin Gap Chambers (TGC) in the end-cap. The RPC extends to $|\eta|$=1.05 and uses the same detection principle as a spark chamber. The TGC covers the range $1.05 < |\eta| < 2.4$ and is a multiwire proportional chamber with a smaller wire-cathode gap than the wire-wire gap.

2.4 Forward Detectors

On top of the subdetectors presented above, other detectors are needed along the beampipe (at high pseudorapidity) for specific purposes. The LUminosity measurement using Cerenkov Integrating Detector (LUCID) is designed to measure relative luminosity using inelastic $p-p$ collisions. It is located at 17 m from the IP. Further down the beampipe, at 237 m, the Absolute Luminosity For ATLAS (ALFA) detector, as its name states, will measure absolute luminosity by detecting elastic $p-p$ collisions. Finally, between LUCID and ALFA, at 140 m, the Zero Degree Calorimeter (ZDC) aims to detects very forward neutrons from heavy ion collisions and also from minimum bias events. The challenge for these detectors is the trigger system, as the time required for their signals to reach the Central Trigger Processor is much greater than that for the rest of ATLAS.

2.5 Online Operation

In order to operate the detector, three main systems are needed: the Trigger system, to filter the interesting data, the Data AcQuisition (DAQ) system to record this data, and last but not least, the Detector Control System, to power and monitor the detector during operation.

The amount of data provided by the LHC will be enormous and not all of these collisions will be worth recording. Only a small portion of them will contain interesting physics. Also, the hardware used for recording the data has intrinsic limitations in bandwidth. From 40 MHz of collision rate, we need to filter it down to a rate of $\sim$200 Hz. To achieve this, a three-level trigger system is implemented in ATLAS. The Level 1 (L1) is a hardware layer which needs to take a decision in less than 2.5 μs. It introduces a new idea called Regions of Interest (RoI) where one triggered object (e.g. EM cluster, muon, photon) can be found in a $\Delta\eta$ x $\Delta\phi = 0.1$ x 0.1 region in the detector. The RoIs are then given to the Level 2 (L2) which have access to the full detector granularity *within the RoI* to make a decision in 40 ms on average. If L2 accepts the event, it is passed to the Event Filter (EF) which uses the full event information to take a decision and classify the event in 4 s on average. L2 and EF form the High Level Trigger (HLT) and are purely software based. The HLT algorithms work in steps and hence can reject an event as soon as possible.

A DAQ system is needed to collect data from the experiment and also to interact very closely with the trigger system. In ATLAS, the DAQ and

trigger are designed with a common data flow system. At L1, front-end electronics store the data in pipelines, waiting for the L1 decision. If L1 accepts the event, the data are transferred to a Read-Out Driver (ROD) in the standard ATLAS format. These data are then sent to the Read-Out System (ROS), which in turn sends it into Read-Out Buffers (ROB). The ROB can send data to the L2 processors on request. Finally, when the event is accepted by L2, the data are sent to the Event Builder (EB) and to the EF for one last trigger decision.

In order for the two previous systems to be operational, the detector hardware needs to be fully functional and well monitored. This is where the Detector Control System (DCS) comes into play. The DCS task is to control and monitor all hardware systems in ATLAS and to ensure its safe operation. This means controlling and monitoring of power supply parameters, environmental parameters and gas supplies. The DCS can also enable manual and automatic safety actions in case of incidents. The DCS and DAQ system can exchange information using the DAQ-DCS Communication (DDC) protocol. The DCS data are stored online and a small portion is also sent to the offline database for reconstruction. A common hardware device, the Embedded Local Monitor Board (ELMB), is used throughout ATLAS for the DCS systems of individual subdetectors.

3 Expected performance

Establishing benchmarks is essential to be able to evaluate the performance in real data-taking. In ATLAS, this was done in two steps, first in the Technical Design Report [2] and then updated in the ATLAS Technical Paper [1] and the full description of the expected performance [3]. The latest performance of the detector is outlined here.

In order to establish the expected performance of the detector, two aspects have to be considered simultaneously: the hardware commissioning, or how well the detector parts are working with respect to their specifications, and the physics commissioning, or how well we are able to reconstruct the physics objects we want to measure. Several tools are available to achieve this. The usage of simulated data (using Monte Carlo techniques) is necessary, covering the various scenarios of data-taking with different center-of-mass energies, with or without pile-up and misalignement, etc. Also, data from test beams, where a beam of pions with a fixed energy is sent to a slice of the final detector, are useful to determine the

detector parameters. Finally, acquisition of cosmic ray and single beam data is described below (see section 4).

In the hardware commissioning phase, we need to determine the intrinsic accuracies and resolutions of the subdectectors using test beams. For the inner detector, the value for the accuracy should be sufficiently small to be able to distinguish two close-by tracks. The pixel module is expected to measure hits with a precision of 10 μm in the transverse direction ($R-\phi$) and 115 μm in the longitudinal direction (z). In the SCT, due to the strip technology, the longitudinal precision is reduced, leading to 17 μm in $R-\phi$ and 580 μm in z. The TRT straw drift time accuracy is 130 μm, since it doesn't record the z position. In the calorimeters, the resolution in energy and the linearity in the response are the important parameters. For the LAr calorimeter, these parameters are shown in Equation 1 for the barrel and in Equation 2 for the Tile calorimeter for $\eta = 0.35$. The first term is the sampling term, measuring the fluctuations of the electromagnetic shower and the second term is the constant term, which measures linearity.

$$\frac{\sigma(E)}{E} = \frac{(10.1 \pm 0.4)\%}{\sqrt{(E)(GeV)}} \oplus (0.2 \pm 0.1)\% \tag{1}$$

$$\frac{\sigma(E)}{E} = \frac{(56.4 \pm 0.4)\%}{\sqrt{(E)(GeV)}} \oplus (5.5 \pm 0.1)\% \tag{2}$$

For the muon chambers, the time resolution is important for the trigger chambers, while the position accuracy is more relevant for precision chambers. The RPC has a 10 mm accuracy in z and ϕ with a 1.5 ns time resolution and the TGC, in the forward region, reaches 2-6 mm in z, 3-7 mm in ϕ and a 4 ns time resolution. The precision CSC chambers in comparison, have a 40 μm z accuracy, a 5 mm ϕ accuracy and a 7 ns time resolution, while the MDT reaches an average resolution of 35 μm per chamber.

The second part of the commissioning concerns the physics objects we wish to reconstruct using the data from the different subdetectors. Five main types of objects are of interest for physics. Low-p$_T$ charged particles are measured by the track they leave in the ID without escaping it. Electromagnetic objects are identified by their energy deposit (cluster) in the EM calorimeter. If a track is matched to this cluster, the object is an electron. If not, the cluster is identified as a photon. A jet is the experimental

signature of a quark or a gluon in the final state of the primary interaction. It is detected by the energy (cluster) it leaves in the hadronic calorimeter. Finally, a muon is identified by a track in the muon chambers matched to a track in the ID. For these objects to be correctly reconstructed, it is then important to identify the potential sources of inefficiency and correct for them.

In the ID, the choice of silicon as active material is beneficial for the position accuracy, but disadvantageous for the material budget, which should be made as low as possible. The services between the barrel and the endcap almost double the amount of material in front of the calorimeter. This leads to higher bremsstrahlung rate for electrons and higher rate of conversions for photons into electron-positron pairs. The particles then lose part of their energy before reaching the calorimeter. The tracking efficiency for electrons is affected by this, especially for low energy and forward electrons. Muons are not significantly affected by material in the ID. The TRT allows for a good separation between electrons and pions as measured with test beam data in the range from 2 to 350 GeV for incoming particles.

Jets in ATLAS can be reconstructed using two different methods. The tower jets are built from the sum of energy in cells in a given geometrical region ($\Delta\eta$ x $\Delta\phi$ = 0.1 x 0.1). The topological cluster jets are obtain by building a 3D cluster around a seed using only cells that are above the noise level. It was shown that, for jets from Vector Boson Fusion Higgs decays, the topological jets have a better recontruction efficiency and purity at lower p_T while the two types of jets have similar performance at higher energies.

Muon reconstruction efficiency was obtained using in-situ methods on simulated data. Over the full $|\eta|$ (up to 2.7) and ϕ range, muons with p_T >5 GeV can be measure with an efficiency ϵ >80%. Another issue for high energy muons is charge misidentification. The charge of particles is measured by the direction of their bending in the magnetic field. As the bending is proportional to the particle transverse momentum, high p_T eletrons or muons have a probability that their charge is wrongly identified. At 1 TeV, tracks are almost straight and this probability is greater than 1%. Charge misidentification is also dependent on η and can reach 30% for 2 TeV muons at $|\eta|$ >2.

Finally, performance benchmarks need to be established in order to trigger on all the physics objects described above. This is particularly important since the trigger decision online is irreversible and it needs to be

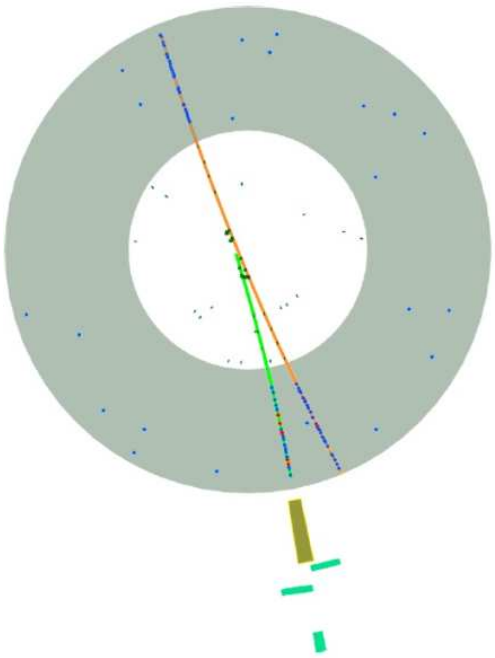

Figure 2: Event display of a cosmic ray muon inducing electron ionisation.

efficient in order not to miss any interesting physics events. In particular, in-situ methods for electron efficiency determination were compared to purely Monte-Carlo based methods and a good agreement was found between them.

4 Data-taking performance results

Other tools involving real data can be used to determine the performance of the ATLAS detector before the LHC delivers proton collisions. These are cosmic ray and single beam runs. The experience gained with them is outlined below.

4.1 Cosmic rays in 2008

Cosmic ray data have been used since 2005 in ATLAS for various detector commissioning activities on the surface and underground. They are particularly useful to determine dead or noisy channels, to perform calibrations

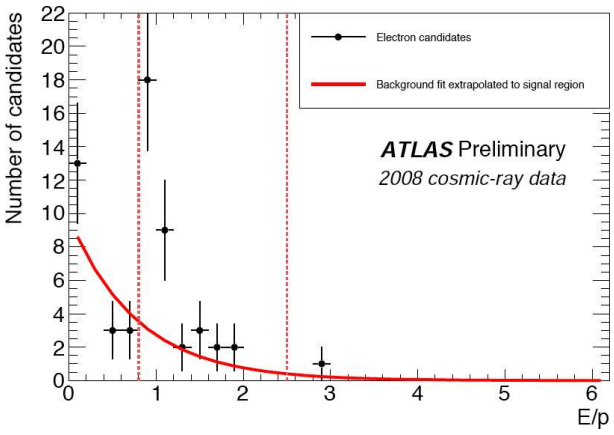

Figure 3: E/p for electron candidates

and alignment, for timing, and more generally, to gain experience operating the ATLAS detector as a whole. During the fall 2008 period, over 200 million events were collected using various L1 triggers, from muon chambers, calorimeters and minimum bias trigger scintillators (MBTS). Out of those, $\sim$ 7 million events were triggered by the L2 ID tracking trigger, the first time the HLT was operated online with real data.

The SCT efficiency and noise performance could be determined with these events. The noise occupancy was measured at 150 V bias voltage to be well below the specifications ($5\text{x}10^{-4}$) and the hit efficiency was greater than 99% (inactive modules are not considered here). Alignment of the ID was performed using tracks collected with the solenoid field turned on. For the barrel in each tracking subdetector, track residuals were in good agreement with the perfectly simulated geometry. In the end-cap, more statistics are needed for a similar measurement.

Noise measurements were done in all the calorimeter layers using random triggers and were found to be consistent in between runs. Also, a good uniformity over time (four months) is observed for the noise. Using the EM calorimeter, a search for electrons in cosmic ray events was performed. These electrons come from muon bremsstrahlung and electron ionisation on the trajectory of the muon, as shown in Fig. 2. After requiring 2 tracks in the ID, a cut on high to low threshold in the TRT and medium quality cuts for the electron selection, the first observation of real electrons in ATLAS can be seen in Fig. 3. Thirty-six events are selected in the signal region.

Good muon tracks recorded in all subdetectors were used to determine the alignment between the ID and the muon system. A good correlation between ID and muon coordinates and a good agreement with the simulated data were observed. As expected from simulation, the average energy deposited by muons in the calorimeters was ∼3 GeV.

4.2 Single beam runs

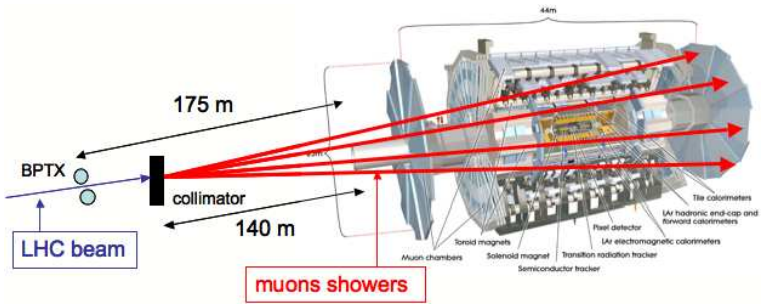

Figure 4: Experimental setup for beam splash events

On September 10th, 2008, the LHC was officially started and single beams were running in both directions (not simultaneously). This great success was followed by a few days of single beam running until a transformer malfunction forced the LHC operators to stop the beam and start training the magnets. On September 19th, an accident happened and the single beam running was fully stopped. During the few days of operation, two types of events were collected: beam splashes and single beam events. A beam splash event (see Fig. 4) is obtained when the beam in one direction hits a collimator placed 140 m before the detector and creating thousands of particles which cross the detector from one side. This generates high energy deposits in the calorimeters (up to ∼ 1000 TeV). Collimator "shots" were done every 42 s. A single beam event is obtained when the beam circulates for ∼100 turns; halo muons can be measured at very high pseudorapidity with lower deposited energy. The beam was made of bunches with $2\text{x}10^9$ protons each.

The ATLAS detector was ready to receive beam on September 10th. For safety reasons, the subdetectors close to the beam pipe were operated at lower voltages, like the SCT end-caps, FCal and muon chambers at high pseudorapidity, and some others were switched off, like the Pixel and SCT

barrel. The TRT was not using Xe as an active gas. The HLT was not used online, while L1 was used for streaming data.

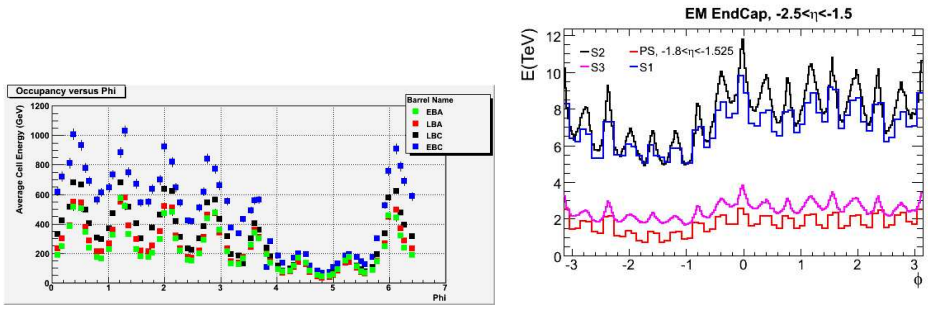

Figure 5: Left: 8-fold structure in Tile Calorimeter energy distribution in azimuth. Note the two reduced peaks due to the support pillars of ATLAS. Right: 16-fold structure seen in LAr End-Cap calorimeter energy distribution in azimuth.

In Fig. 5, the energy deposited in the TileCal from a beam splash event shows an eightfold peak structure in azimuth, which is explained by the magnet geometry of ATLAS. A similar structure was observed in the LAr calorimeter and especially in the EM end-cap, where a 16-fold structure is explained by the shielding components near the beam pipe. The SCT end-cap disks show a higher module occupancy on the side of the incoming beam. Single beam events were useful to do trigger timing. An improvement in the overlap of the different L1 triggers was visible after 2 days of running.

5 Conclusion

The ATLAS detector, a giant multipurpose particle experiment, is now fully installed and operational in a cavern underground at CERN, thanks to the work of all the people involved in this project. Its different subdetectors are operational and have been tested with cosmic ray and single beam data. The expected performance of its hardware and software components was determined and cross-checked with data. The detector is ready to resume beam operation as expected for 2009-2010.

6 Acknowledgements

We are greatly indebted to all CERN's departments and to the LHC project for their immense efforts not only in building the LHC, but also for their direct contributions to the construction and installation of the ATLAS detector and its infrastructure. We acknowledge equally warmly all our technical colleagues in the collaborating Institutions without whom the ATLAS detector could not have been built. Furthermore we are grateful to all the funding agencies which supported generously the construction and the commissioning of the ATLAS detector and also provided the computing infrastructure.

The ATLAS detector design and construction has taken about fifteen years, and our thoughts are with all our colleagues who sadly could not see its final realisation.

References

[1] The ATLAS Collaboration, G. Aad *et al.*, The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003.

[2] The ATLAS Collaboration, Detector and Physics Performance Technical Design Report, CERN/LHCC/99-15 (1999).

[3] The ATLAS Collaboration, G. Aad *et al.*, Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics, CERN-OPEN-2008-020, arXiv:0901.0512 [hep-ex].