

Chapter 2

The ATLAS Experiment

A substantial part of this chapter is based off of Ref. [1] for the LHC and Ref. [2] for the ATLAS detector.


2.1 The CERN Large Hadron Collider

The Centre for European Nuclear Research (CERN, *Conseil Européen pour la Recherche Nucléaire*), established in 1954, is the world's largest institution for particle physics research. Based in the vicinity of Geneva, Switzerland, on the border between Switzerland and France, the organization is currently comprised of 22 full member states and maintains differing levels of relations with numerous other nations from around the world. The primary focus of CERN is the development, construction, and maintenance of a series of high energy accelerator-based experiments, of which the Large Hadron Collider (LHC) is the current flagship.

The LHC, a circular proton-proton synchrotron collider, first became operational in November 2009. With a circumference of 27 km, and a centre of mass collision energy of 7 or 8 TeV, the A Toroidal LHC ApparatuS (ATLAS) is the currently the largest and highest energy particle accelerator in the world. The LHC is currently undergoing upgrades, and is expected to enable collisions to occur at a new energy frontier of 13 TeV in the spring of 2015, with a possible further upgrade to attain the design energy of 14 TeV. In addition to producing proton-proton collisions, the LHC is also used for proton-lead and lead-lead collisions. While interesting in their own right, these other types of collisions will not be considered further.

The LHC supports four main independent experimental collaborations, known as ALICE, ATLAS, CMS, and LHCb, which are spread along the circumference of the LHC as shown in Fig. 2.1. The layout of the ring itself is broken up into octants, numbered in clockwise order, with the main Meyrin site at the bottom of the ring occupying octant 1. The eight octants each have a primary purpose, with

Fig. 2.1 The locations of the four main LHC experimental collaborations shown along the circumference of the CERN LHC. The main Meyrin site at Point 1 and the Geneva airport at Point 8 are also shown [3]



four nominally representing collision points and associated experiments, and four representing purely LHC-related functionality [1]:

1. The ATLAS experiment, a collision point for proton-proton, proton-lead, and lead-lead collisions.
2. The ALICE experiment, a collision point for proton-lead and lead-lead collisions.
3. Beam cleaning for the LHC beams.
4. Radio Frequency (RF) cavities for the actual acceleration of the LHC beams.
5. The CMS experiment, a collision point for proton-proton, proton-lead, and lead-lead collisions.
6. Beam dumping, for the LHC beams.
7. Beam cleaning for the LHC beams.
8. The LHCb experiment, a collision point for proton-proton collisions.

The LHC is the final step of a multi-stage accelerator system, and is only one of the many experiments powered by the CERN accelerator complex, as shown in Fig. 2.2 and detailed in Ref. [4]. The injector chain for the LHC starts with a linear accelerator to bring the protons to an energy of 50 MeV. This linear accelerator is used to fill the 157 m circumference Proton Synchrotron Booster, which accelerates the protons to an energy of 1.4 GeV. These are then injected into the Proton Synchrotron, a 628 m circumference accelerator which brings the protons to an energy of 25 GeV. The next stage is the 7 km circumference Super Proton Synchrotron, which accelerates the protons to 450 GeV. These beams are then injected into the LHC, which accelerates the protons using the RF cavities in Point 4 to 3.5 TeV (2010 and 2011), 4 TeV (2012), 6.5 TeV (planned 2015), and potentially higher. Lead ions follow the same steps, with the exception of starting with a different linear accelerator, and replacing the Proton Synchrotron Booster with the 78 m Low Energy Ion Ring.

Within the LHC, two separate beams are run in opposite directions, which are then set on a collision course using quadrupole magnets at the four interactions points. These beams consist of *bunches* of protons, where the size and frequency of these

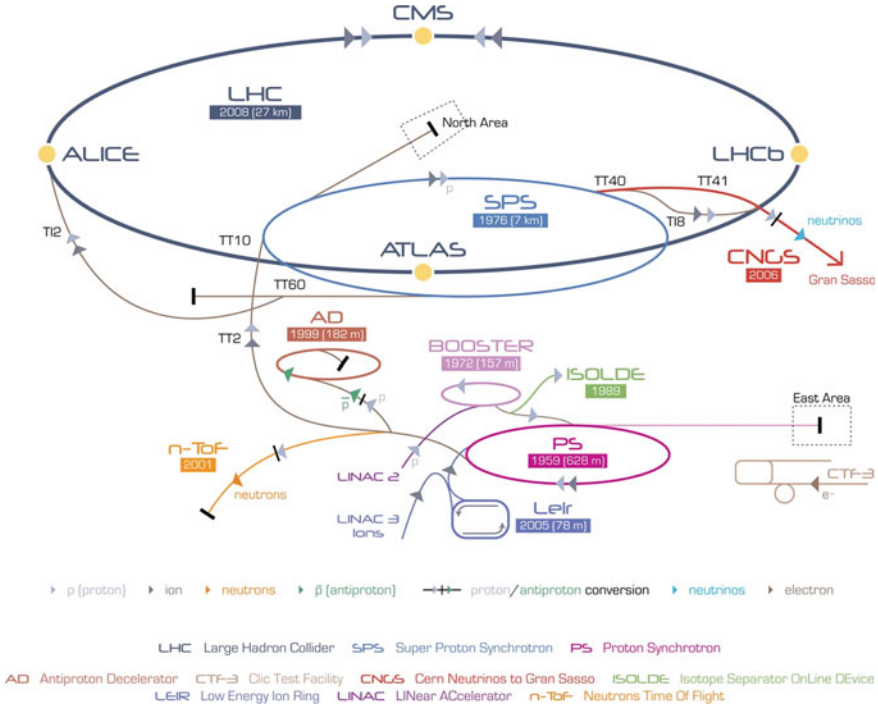


Fig. 2.2 The CERN accelerator complex, showing the multi-stage system of injectors which are used to fill the LHC. A subset of the many experiments supported by these accelerators is also shown [5]


bunches is directly relevant to the rate of interesting physics produced in collisions. This measure is known as the luminosity \mathcal{L} , which relates the number of events in a given amount of time to the cross-section of a desired process, as per Eq. 2.1. The luminosity thus depends on the size of the bunches, the frequency with which bunches interact (frequency of bunch crossings), and the beam cross-sectional area of overlap \mathcal{A} , as defined in Eq. 2.2.

$$\frac{dN_{\text{events}}}{dt} = \mathcal{L}\sigma \quad (2.1)$$

$$\mathcal{L} = \frac{f_{\text{crossing}} N_{\text{bunch1}} N_{\text{bunch2}}}{\mathcal{A}} \approx \frac{f_{\text{crossing}} N_{\text{bunch1}} N_{\text{bunch2}}}{4\epsilon\beta^*} \quad (2.2)$$

The beam cross-sectional area of overlap \mathcal{A} is related to the beam cross-sectional size at the interaction point β^* and the beam emittance ϵ . The envelope that can be formed around a particle's trajectory as it undergoes betatron oscillation while traversing the accelerator is known as the amplitude function β , and the value at the interaction point is β^* . The beam emittance is independent of the location of the

Fig. 2.3 The 2011 dataset switched the value for β^* partway through the year. A fit to the two values against the mean number of interactions per bunch crossing demonstrates the impact of squeezing beams on the mean number of collisions per crossing, and thus the luminosity [6]



particle in the accelerator, and is instead related to the area of the phase ellipse in which particles can circulate.

Given the need to maximize the sensitivity to new physics signals, it is important to increase the luminosity as much as possible, thus providing as much integrated luminosity $\int \mathcal{L} dt$ as possible for analysis use. This can be done by increasing the number of particles in a bunch, increasing the crossing frequency, or reducing the cross-sectional area of overlap between the beams. Two different values for β^* are shown in Fig. 2.3, demonstrating how reducing the amplitude function (squeezing the beam) leads to a higher number of average interactions per crossing. The overlap between the beams is unlikely to further change in a significant fashion, as this is limited by the quadrupole magnets, which are already used at near design values.

The time between bunches was typically 50 ns in Run-I of the LHC at both the ATLAS and Compact Muon Solenoid (CMS) experiments. This will be retained for the first data in Run-II, in order to ensure that the collider is stable with beams at 6.5 TeV, but is planned to be halved to 25 ns bunch spacing by the end of 2015.

Alternatively, the luminosity can be increased by adding more particles to each bunch. Given that the transverse area of each bunch is approximately fixed, this increases the probability that particles from opposing bunches will collide when crossed. This was the preferred technique for increasing the luminosity in the later stages of Run-I, most notably in 2012.

Increasing the luminosity from any of these techniques leads to the creation of pileup, which takes two forms. In-time pileup is the phenomenon by which multiple collisions happen in a single crossing, and thus the detector must be able to separate particles from multiple simultaneous interactions. Reducing the beam size or increasing the number of particles per bunch contributes to increasing the number of simultaneous collisions, and thus the in-time pileup. The ability of different types of detectors to differentiate between simultaneous collisions varies, as will be discussed later in this chapter.

Out-of-time pileup relates to the finite time required for the signals to develop in and be read from the detector. Residual energy from a previous bunch crossing or

energy from a future bunch crossing which is deposited before the read-out process is completed can contribute to the electrical signals associated to a given bunch crossing. This effect is primarily controlled by the amount of time between subsequent bunch crossings, and thus the amount of energy overlap. Some types of detectors can be read out at a rate faster than the bunch crossing frequency, rendering this irrelevant, but others take a longer amount of time. The effect on different types of detectors will be discussed later in this chapter.

Pileup is typically parametrized in terms of the number of reconstructed primary vertices N_{PV} for in-time contributions and the average number of interactions per bunch crossing $\langle\mu\rangle$ for out-of-time contributions. N_{PV} is a per-event quantity as calculated from the number of reconstructed vertices in the tracking detectors, as will be discussed in Sect. 2.3, and which comes coupled with a vertex identification efficiency. On the other hand, $\langle\mu\rangle$ is averaged over a block of events, and is calculated from the luminosity. As such, the two values provide useful independent estimates of the amount of pileup activity.

As the LHC operations team becomes more familiar with the accelerator, they have successively increased the luminosity, and thus the associated pileup. A summary of the beam conditions as delivered to the ATLAS experiment is provided in Table 2.1. Figure 2.4 shows the peak instantaneous luminosity and peak instantaneous interactions per crossing, demonstrating the strong correlation between the two variables. This is then integrated over time to obtain Fig. 2.5.

The ATLAS detector was designed to handle and excel under these high levels of pileup, and has been able to make full use of the data provided by the LHC.

Table 2.1 Luminosity-related LHC beam conditions in 2010, 2011, and 2012 as well as the design values for the LHC [6, 7]

Variable	2010	2011	2012	Design value
Per-beam energy (TeV)	3.5	3.5	4	7
β^* (amp. at int. point) (m)	3.5,2.0	1.5,1.0	0.6	0.55
Bunch spacing ($1/f_{\text{crossing}}$) (ns)	150	75,50	50	25
Max number of bunches	368	1380	1380	2808
Max number of protons (10^{11} /bunch)	1.2	1.45	1.7	1.15
Peak luminosity ($\text{cm}^{-2} \text{s}^{-1}$)	2.1×10^{32}	3.7×10^{33}	7.7×10^{33}	1×10^{34}
Max $\langle\mu\rangle$ (int/crossing)	4	17	37	19

The design conditions listed correspond to the LHC program itself, not the recently approved HL-LHC program. Max $\langle\mu\rangle$ corresponds to the peak value measured when averaging over a full luminosity block, as shown in Fig. 2.4.




Fig. 2.4 The peak of the number of interactions per bunch crossing averaged over the luminosity block (μ) (*top*) and instantaneous luminosity (*bottom*) delivered by the LHC as observed by ATLAS. The 2010, 2011, and 2012 data taking periods are shown as a function of time, demonstrating the significant increase in available data as time progresses [6]

2.2 The ATLAS Detector

The ATLAS detector is a large general-purpose experiment located at LHC Point 1, underneath the CERN main site located in Meyrin, Switzerland. The experiment is sensitive to many different areas of particles physics, from both proton-based or heavy ion-based collisions, including searches for new phenomena, investigations into the properties of the observed Higgs, and ultra high precision SM measurements.

Cylindrically shaped, the ATLAS detector is designed to be as close to hermetic as possible, thus enabling reconstruction and identification of nearly every energetic particle produced in the event. The detector has been designed to handle the expected




Fig. 2.5 The mean number of interactions per bunch crossing as calculated from the instantaneous luminosity (*left*) and the integrated luminosity (*right*) delivered by the LHC as observed by ATLAS, for the 2011 and 2012 datasets. The doubling of $\langle \mu \rangle$ and halving of β^* leads to approximately four times more integrated luminosity in 2012 [6]

high luminosity and high pileup conditions, with multiple optimizations designed to counteract the contribution of pileup, as will be discussed where relevant.

ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the centre of the LHC ring, and the y -axis points upward. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the beam pipe. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln \tan(\theta/2)$.¹

The ATLAS detector is made up of several sub-detectors. The innermost layers, immersed within a magnetic field from a 2 T solenoid, serve as tracking detectors with a coverage of up to $|\eta| < 2.5$. The middle layer of calorimeters are segmented into three main regions, with an ultimate coverage all the way out to $|\eta| < 4.9$. The outermost layer of the detector is the muon spectrometer, bathed in a toroidal magnetic field, which provides coverage for $|\eta| < 2.7$. These three levels of detectors will be discussed in Sects. 2.3, 2.4, and 2.5 respectively. A schematic diagram showing the detector and its many components is provided in Fig. 2.6.

Many sub-detectors involve two types of components, designed to maximize the precision of measurements over a wide range of angles. In the central $|\eta|$ region, where particles are approximately perpendicular to direction of the beam, long cylindrical detectors around the beam pipe are used. Once the particle trajectory becomes more aligned with the beam, wheel-shaped endcap detectors are used, where the axle of the wheel is the beam pipe. These two scenarios are associated with very different levels of particle flux, and thus materials with different properties are often required, hence the typical separation of sub-detectors based on $|\eta|$.

¹This paragraph is directly copied from the ATLAS common coordinate system definition, as recommended by the ATLAS publications committee: <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/CoordinateSystemFootnote>.




Fig. 2.6 A computer-generated schematic representation of the ATLAS detector, with the various sub-detectors indicated and people shown for scale [8]

This split between barrel and endcap sub-detectors applies to all three types of ATLAS detector classifications, where the exact $|\eta|$ separation between the two differs. Due to this detector scheme, the central $|\eta|$ region is also sometimes called the barrel region, and similarly for the endcap region. Only the calorimeters have significantly extended $|\eta|$ coverage, providing instrumentation in what is typically referred to simply as the forward region. Note that this classification excludes luminosity detectors and beam conditions monitors, where the former is in the very forward region and the latter exists in multiple regions.

2.3 Inner Detector

The ATLAS inner detector, comprised of three sub-detectors, is entirely enclosed by a superconducting solenoid magnet providing a 2 T magnetic field. The tracker is designed to provide excellent track and vertex reconstruction with superb momentum resolution within the full range of $|\eta| < 2.5$. The nominal lower track p_T threshold is thus 500 MeV, although in some cases lower thresholds are possible. Electron identification is also provided by the outermost sub-detector over a large kinematic range of 500 MeV–150 GeV within $|\eta| < 2.0$. A schematic diagram of only the ATLAS inner detector is provided in Fig. 2.7. A technical diagram of the inner detector is provided in Fig. 2.8.




Fig. 2.7 A computer-generated schematic representation of the ATLAS inner detector, with the various sub-detectors indicated [9]




Fig. 2.8 A quarter-sectional view of the ATLAS inner detector, showing each of the sub-detectors, their envelopes, and the geometrical space that they cover [2]

Pixel Detector

The innermost layer of the inner detector is the Pixel Detector, which contains three layers of high-precision and high-granularity semiconductor modules. These 1744

sensor modules each contain 47,232 pixels, which are individually $250\mu\text{m}$ thick. For technical reasons relating to the read-out system, a small number of pixels per module are ganged, thus the total number of independent read-out channels is 46,080 per module. This segmentation provides excellent position resolution, amounting to a per-layer resolution of $10\mu\text{m}$ in the transverse ($r - \phi$) plane and $115\mu\text{m}$ in the axial plane (z , barrel) or radial plane (r , endcaps).

The three layers within the barrel of the Pixel Detector are located at z of 50.5, 88.5, and 122.5 mm from the interaction point. The three layers within the endcap are located at r of 495, 580, and 650 mm from the interaction point. Due to the proximity of the lowest layer of the Pixel barrel to the interaction point, the lowest layer covers out to $|\eta| < 2.5$. At all points there are at least three layers of pixel active, but there is always one from the barrel and zero through two in the endcap.

SemiConductor Tracker

The second layer of the inner detector is the SemiConductor Tracker (SCT), which contains several layers of microstrip silicon wafers. These 15,912 sensors are arranged in four layers within the barrel, and nine in the endcap. The endcap layers are arranged such that there are always four levels of sensors for particles to pass through. This segmentation provides high quality position resolution, amounting to a per-layer resolution of $17\mu\text{m}$ in the transverse ($r - \phi$) plane and $580\mu\text{m}$ in the axial (z , barrel) or radial (r , endcaps).

The four layers within the barrel of the SCT are located at z of 299, 371, 443, and 514 mm from the interaction point. The nine layers within the endcap are located at r of 853.8, 934, 1091.5, 1299.9, 1399.7, 1771.4, 2115.2, 2505, and 2720.2 mm from the interaction point. Unlike the pixel detector, the barrel modules are only available up to $|\eta| \approx 1.5$, after which the endcap modules are used independently. Further details, such as the z extent of the endcap modules, is provided in Fig. 2.8.

Transition Radiation Tracker

The third and outermost layer of the inner detector is the Transition Radiation Tracker (TRT), which as the name implies is a combined transition radiation detector and tracking detector. The primary building-block of the TRT is a set of 73 (160) layers of polyimide straw tubes in the barrel (endcap). These straws, each 4 mm in diameter, contain a gas mixture of 70 % Xe, 27 % CO_2 , and 3 % O_2 . The resulting position resolution is significantly weaker than the Pixel Detector or SCT, with the entire detector providing a resolution of $130\mu\text{m}$.

Any charged particle passing through the straws ionizes the gas, thus producing ionization electrons. The outer wall of the straw, which is clad in a thin layer of metal, is held at a high negative voltage thus forcing the electrons toward the central anode wire. The electrons take up to 48 ns to reach the wire, where the arrival of the first ionization electron can be used to determine how close the particle passed to the centre of the straw.

The transition radiation part of the name comes from electrons, which are typically the only particles with a large enough Lorentz boost factor ($\gtrsim 1000$) emit X-rays as they pass through (transition between) the gas mixture and polyimide straws. The

inclusion of 3 % oxygen was found to reduce the ionization electron mobility, but increase the transiting electron identification power, and thus was retained. As the energy is increased, other particles (most notably charged pions) also begin to emit X-rays, hence the upper electron identification threshold of approximately 150 GeV.

The barrel of the TRT stretches from 563 to 1066 mm in z with respect to the interaction point, while the endcap covers 848 to 2710 mm in r . The barrel is only available within $|\eta| < 1.0$, and the endcap extent ends at $|\eta| \approx 2.0$. This is why the electron identification power of the inner detector is restricted to $|\eta| < 2.0$, rather than 2.5.

2.4 Calorimetry

After particles make it through the tracking detector, they reach the calorimeters. Calorimeters are designed to provide accurate energy measurements of all particles passing through them which undergo either electromagnetic or hadronic interactions. Only neutrinos and muons (which are Minimally Ionizing Particles (MIPs) at the LHC energy scale) pass through the calorimeters mostly unaffected.

As particle(s) from the hard-scatter collision interact with the calorimeter material, they create cascades of particles, also called showers. There are two types of cascades, depending on the nature of the source, and thus they are named electromagnetic and hadronic showers. These two types of showers have radically different properties, and thus require separate techniques for high-precision detection. Electromagnetic showers are primarily initiated by electrons or photons, while hadronic interactions are initiated by hadrons. Note that the hadrons which reach the calorimeter can come from very different sources. For example, hadronic showers can be the result of a single pion from tau decays, or they can be collections of hadrons from the fragmentation of quarks or gluons produced at the interaction point. These are but two examples of many possible processes, and demonstrate the large degree of variation that can be expected from hadronic showers.

The ATLAS calorimetry, comprised of several sub-detectors, is between the inner superconducting solenoid magnet and the outer toroidal magnets. The calorimeters are designed to provide excellent electron and photon reconstruction with superb energy resolution and containment, while also providing exceptional hadronic shower reconstruction and performance. In order to meet these needs, there are two levels of detectors throughout the full range of $|\eta| < 4.9$, with an inner layer of electromagnetic calorimeters and an outer layer of hadronic calorimeters. Within each layer, there are three or four sub-detectors, covering the barrel and/or extended barrel, endcap, and forward regions. A schematic diagram of the ATLAS calorimetry is provided in Fig. 2.9.

The ATLAS calorimeters are all sampling calorimeters, meaning that alternating layers of absorbing and active material are used. As particles traverse the calorimeters, the absorbing material is intended to force an interaction, thus producing either an electromagnetic or hadronic shower. The following active layers are designed




Fig. 2.9 A computer-generated schematic representation of the ATLAS calorimeters, with the various sub-detectors indicated [10]

to accurately measure the amount of energy in the particle cascade. This alternating layer design allows for creating compact calorimeters with excellent shower containment, but at the cost of energy lost in the absorption layers. The fraction of observed energy, known as the sampling fraction $f_{\text{samp}} = E_{\text{active}} / (E_{\text{active}} + E_{\text{passive}})$, provides a measure of quantifying this effect. If the sampling fraction for a given calorimeter is known, then the observed shower energy can be rescaled to account for the unobserved contribution.

2.4.1 Particle Showers

There are several critical differences between electromagnetic and hadronic showers which control the associated detector requirements. Electromagnetic showers can be broken up into two types of contributions, namely electron/positron and photon interactions. At high energy, electrons emit photons via Brehmsstrahlung radiation while photons produce electron-positron pairs, demonstrating how these two sectors are not distinct. At medium energies, photons can also Compton scatter to reduce their energy, while at low energies the photoelectric effect converts them into low energy electrons. Low energy electrons are subsequently absorbed to form ions, thus ending the showering process.

In the end, the vast majority of the energy from electromagnetic showers is absorbed in the calorimeters, so long as a sufficiently deep calorimeter is used. However, not all of the material that the shower passes through is an active (recording) medium. Energy is also deposited in the absorbing (passive) layers, thus it is better to have a larger fraction of active material for high precision measurements, so long as containment is not sacrificed. Typical sampling fractions for electromag-

netic sampling calorimeters are at the level of 10 %, such as the CDF lead-scintillator calorimeter with $f_{\text{samp}} = 12\%$ [11]. The ATLAS Collaboration decided that a higher-quality electromagnetic barrel calorimeter was worth the increased expense, and thus the sampling fraction is 18 % [12].

Hadronic showers, on the other hand, are much more complicated. Approximately 1/3 of hadronic interactions produce a π^0 meson, which then almost always decays to two photons, thus adding an electromagnetic component to the hadronic shower. The remaining 2/3 of hadronic interactions produce additional hadronic activity. This can be the production of charged mesons, such as π^\pm , or the fragments of nuclei, such as protons. The fraction of each process depends on the energy and target material, but these two processes account for approximately 20 and 30 % of the hadronic interactions respectively. Another approximately 10 % goes into neutrons, while the remaining 40 % results in invisible processes [13]. This is primarily due to the loss of nuclear binding energy when an atom is fragmented through a strong interaction, although other sources such as nuclear recoil also contribute.

In other words, while the vast majority of energy in electromagnetic showers is captured, a significant fraction of hadronic shower energy is undetectable. Some types of hadronic calorimeters have a means of correcting, or compensating, for this effect. The ATLAS calorimeter, however, is non-compensating, meaning that the response ratio of $e/\pi < 1$. Furthermore, the fraction of energy which is observed for contained hadronic showers is dependent on the energy, meaning that the response ratio of e/π is also non-linear. An explanation of this energy dependence, and the way in which hadronic systems are calibrated to account for these reconstruction differences, is presented in Sect. 4.3.

The quantity that describes the strength of the interaction between the detector material and the shower also varies between hadronic and electromagnetic showers, as is expected given that different forces are at work. The property which governs the average distance of material which must be traversed to reduce the energy of an electron to $1/e$ of the starting value (and $7/9$ of the distance for the average high-energy photon to produce an electron-positron pair) is called the radiation length (X_0). This is typically specified in units of g/cm^2 , where dividing by the density of the material gives a radiation length in cm. Typical electromagnetic calorimeters work with enough material to provide more than 20 radiation lengths of material in order to fully contain the electromagnetic showers.

Hadronic showers, on the other hand, are governed by the strong force. Atoms are mostly empty space, and thus a given particle can be expected to travel much further through a typical material before undergoing an hadronic interaction as compared to an electromagnetic interaction. The average distance travelled by a hadron before undergoing an inelastic collision with the surrounding material, also known as the nuclear interaction length (λ_I), should thus be much larger than X_0 . This is observed, with the nuclear interaction lengths for typically used materials 5–30 times longer than the corresponding radiation length [14]. Similarly to radiation lengths, this is typically quoted in g/cm^2 , thus dividing by the material density gives a nuclear interaction length in cm. Calorimeters typically have closer to 10 nuclear interaction lengths of active material, which contains the majority of hadronic showers.

This difference in the average distance a particle will travel before interacting makes it relatively easy to fully contain all electromagnetic showers up to reasonably high energies, while it is essentially impossible to do the same for hadronic showers. When a shower passes through the entirety of the calorimeter, this also reduces the calorimeter response, as a potentially large fraction of the energy is lost. This effect, known as punch-through, becomes more frequent as the energy scale of the shower is increased. This phenomenon and preliminary corrections for such scenarios is discussed in Sect. 4.4, where it is seen that punch-through becomes an important effect at high energies.

Note that while electromagnetic and hadronic showers have different properties, hadronic showers still show up in the electromagnetic calorimeters. The division between electromagnetic and hadronic calorimeters refers to the type of shower that they are intended to fully contain, not the type of shower which can be present in the detector.

2.4.2 Electromagnetic Calorimeters

The most forward of the electromagnetic calorimeters will be discussed in Sect. 2.4.4, as it is substantially different from the other electromagnetic calorimeters discussed here.

Despite being broken up into three separate sub-detectors, the ATLAS electromagnetic calorimetry has a predominantly unified form. All of these detectors make use of Liquid Argon (LAr) as the active material, where LAr was chosen due to its uniform nature, stability, and radiation-hardness. The passive material was chosen to be lead, which has a short radiation length of $X_0 = 0.561$ cm [14], thus enhancing the level of electromagnetic shower containment within the available space. Throughout the detectors, the typical number of radiation lengths of calorimeter material is typically larger than 22, as shown in Fig. 2.10.




Fig. 2.10 The number of interaction lengths X_0 of material present in the barrel (*left*) and endcap (*right*) electromagnetic calorimeters. The *crack region* where the barrel accordion meets the endcap accordion is not shown [2]

As particles pass through the calorimeters, they ionize the LAr. A high voltage is applied to the plates enclosing the LAr, causing the electrons to drift to the copper electrodes within a time of approximately 450 ns. The uniformity of the LAr results in a clean triangular signal shape, where the majority of the signal arrives within a short period before decreasing to zero linearly. Waiting for the full signal to arrive is not feasible, as 450 ns spans many bunch crossings. Instead, a signal shaper is applied to the raw triangular distributions, creating a bipolar pulse signal with a much shorter readout time. The shaper was designed to match the LHC design bunch spacing of 25 ns, thus the calorimeter performance is expected to improve in 2015 as the bunch spacing is switched from 50 to 25 ns. In order to ensure that the pulse is optimal for the data taking conditions, the constants used in the shaping function are frequently updated in dedicated calibration runs. These calibration runs determine the optimal amplitude, noise level (primarily pileup activity), and pulse shape [15].

In order to not bias future events where the charge is still being collected, the pulse shape is set to have a large negative tail after the peak, with the integral of the full pulse constructed to be zero. This tail balances the residual charge being collected, thus mitigating the contribution of residual energy from a given interaction affecting future events. Fluctuations in the pileup activity render this an average rather than absolute correction, as cells can have either a positive or negative energy when read out at a given time due to such fluctuations. Further suppression of pileup related effects in the calorimeter are discussed in Sects. 4.2.1 and 4.3.2.

The barrel and endcap sub-detectors both make use of a unique design, by which the various sampling and active layers are aligned in an accordion geometry. This geometry reduces the number of blind regions of the calorimeter, and provides full ϕ coverage, which is an improvement over previous calorimeter geometries. The accordion structure in use in the electromagnetic barrel can be seen in Fig. 2.11 along with the segmentation of the detector.

Electromagnetic Barrel Calorimeter

The ElectroMagnetic Barrel (EMB) calorimeter, as shown in Fig. 2.11, consists of three layers with differing granularities and depths, covering $|\eta| < 1.475$. The first layer of the EMB is special in that it provides extremely fine segmentation in the η dimension, with strip cells of 0.0031×0.098 in (η, ϕ) -space. This is useful for improving the η measurement precision, but is critical for discriminating photons from $\pi^0 \rightarrow \gamma\gamma$ decays. The second layer captures the large majority of the energy in the calorimeter, and thus a high level of segmentation is needed in both η and ϕ to achieve optimal resolution, resulting in cells of 0.025×0.0245 in (η, ϕ) -space. The final layer is much more coarse, and is intended to collect only the last parts of the electromagnetic shower, and thus high precision measurements are less critical. As such, the third layer is 0.05×0.0245 in (η, ϕ) -space. However, having a separate third layer rather than just extending the second layer has other uses, such as providing useful information on the nature of the shower, as exploited in Sect. 4.3.5.




Fig. 2.11 A view of a section of the electromagnetic barrel calorimeter, displaying the accordion geometry structure and detector segmentation. The three layers of calorimeter cells and their corresponding sizes and radiation lengths X_0 are shown [2]

Electromagnetic Endcap Calorimeter

Two coaxial wheels form the ElectroMagnetic EndCap (EMEC) calorimeter, providing coverage for the range of $1.375 < |\eta| < 3.2$. The EMEC has largely the same structure as the EMB, just geometrically rotated into the radial rather than axial direction, and with the internal accordion geometry rearranged for this orientation. The transition region between the EMB and the EMEC, within $1.375 < |\eta| < 1.52$, is traditionally known as the *crack* region. This region contains a significant quantity of material which provides necessary services to the inner detector. These services amount to several radiation lengths of inactive material, and thus are a source of energy loss and reduce performance. Analyses which require high precision electrons or photons typically remove this region from their analysis selection.

Electromagnetic Presamplers

To account for energy lost in material before the electromagnetic calorimeters, finely segmented LAr presamplers are placed in front of other material whenever possible, covering $|\eta| < 1.7$. This includes a presampler before the service material mentioned in the *crack* region, thus reducing but not removing the impact.

2.4.3 Hadronic Calorimeters

The most forward of the hadronic calorimeters will be discussed in Sect. 2.4.4, as it is substantially different from the other hadronic calorimeters discussed here.

Unlike the electromagnetic calorimeters, two radically different technologies are used for the three sub-detectors. The barrel and extended barrel calorimeters share one common type of technology, while the technology used for the endcap calorimeter is closer to what was used for its electromagnetic counterpart. The number of interaction lengths of calorimetry of either type varies throughout the ATLAS detector, from approximately 7–16, as shown in Fig. 2.12. This non-uniformity and lower value of λ_I compared to X_0 leads to many more situations in which hadronic showers punch-through the full calorimeter, as is studied in Sect. 4.4.


Barrel and Extended Barrel Hadronic Calorimeters

The barrel and extended barrel segments of the hadronic calorimeter are known as the Tile Barrel and Tile Extended Barrel due to their use of plastic scintillating tiles as the active medium. The absorbing layers for these calorimeters are comprised of steel, which has a moderate nuclear interaction length of $\lambda_I = 16.8$ cm [14]. The scintillating tiles are arranged radially in a fashion perpendicular to the beamline. The barrel covers out to $|\eta| < 0.9$, while the extended barrel covers $0.8 < |\eta| < 1.7$, as shown in Fig. 2.13.

As charged particles traverse the scintillating tiles, photons are produced in proportion to the amount of energy deposited. These photons are then collected with wavelength-shifting fibres within each calorimeter cell. Rows of cells are grouped together in η to form a module, which is then read out through PhotoMultiplier Tubes (PMTs) to measure the deposited signal. This procedure means that an entire row of cells is lost if any problems occur within the read-out services. The impact of such failures is discussed in Sect. 4.6.

The scintillator and steel layer design was chosen to keep the cost reasonable, a necessary requirement given that the detector is more than 10 m long. The detector

Fig. 2.12 The number of interaction lengths of material in the ATLAS detector as a function of the pseudorapidity $|\eta|$. The unlabelled cyan region is additional material after the calorimeters and before the muon spectrometer, and is only shown up to $|\eta| = 3.0$. The unlabelled gold area represents additional material before the calorimeters [2]






Fig. 2.13 A technical diagram of the Tile barrel (*left*) and Tile extended barrel (*right*) hadronic calorimeters, showing the module placement and pseudorapidity coverage [2]

consists of three layers in both the barrel and extended barrel, with additional sensors in the transition region between the two sub-detectors to partially recover the energy lost in the crack. The Tile calorimeter has more coarse granularity than the electromagnetic calorimeters, with cells of $0.1 \times \pi/32$ for the first layers in (η, ϕ) -space. This provides suitable precision for hadronic showers, as they are much less compact than electromagnetic showers, and thus higher granularity is less important.

Hadronic Endcap Calorimeter

The Hadronic End Cap (HEC) calorimeter technology is similar to that of the EMEC, with an active medium of LAr, but with a passive absorption medium of copper rather than lead. The geometrical arrangement of the layers also differs, where a simple design of flat copper plates is used rather than the accordion geometry used in the EMEC, as shown in Fig. 2.14.

The HEC is comprised of two wheels, each with two layers, covering the range of $1.5 < |\eta| < 3.2$. The first wheel has a higher granularity with 24 copper plates per module, while the second wheel uses only 16. Individual cells are roughly 0.1×0.1 in (η, ϕ) space for $|\eta| < 2.5$ and 0.2×0.2 for the rest of the detector.

2.4.4 Forward Calorimeters

The Forward CALorimeter (FCal) is substantially different from the rest of the ATLAS detector. While the inner detector and muon systems both have barrel and endcap detectors, only the calorimeters have a forward detector, which covers $3.1 < |\eta| < 4.9$. Placed 4.7 m from the interaction point, the FCal is placed in close contact with the endcap calorimeters, thus lessening the impact of the transition between sub-detectors. A diagram of the placement of the FCal and surrounding endcap material is provided in Fig. 2.15.




Fig. 2.14 A technical diagram of the Hadronic EndCap calorimeter, showing the flat copper plates and the larger number of plates in the first layer with respect to the second [2]




Fig. 2.15 A diagram showing the placement of the forward calorimeter with respect to the surrounding endcap calorimeters and other material [2]

The FCal consists of three separate layers, of which the first is an electromagnetic calorimeter, and the remaining two are hadronic calorimeters. The active medium for all three is LAr, while the passive medium is copper for the first layer and tungsten for the latter two. This very forward region experiences extreme particle flux at very high energy, and thus must be extremely radiation-hard and use materials which can stop very energetic particles. Very dense materials must be used to maximize

the number of interaction lengths, thus limiting both the transverse and longitudinal shower development. This partially motivates the use of copper rather than lead for the electromagnetic calorimeter (X_0 of 1.44 vs 0.561 cm), as copper has a shorter nuclear interaction length (λ_I of 15.3 cm vs 17.6 cm). However, the more important property of copper and the ultimate reason it was selected is its conductivity. The extreme particle flux in the forward region quickly heats the material, which quickly leads to problems if left unchecked. Using very conductive materials such as copper reduces the amount of cooling services required to counter this effect.

The second and third layers of the FCal are already sufficiently shielded by the first layer, and thus the cooling is less critical, albeit still important. At this point, the stronger requirement is the need to stop hadronic showers from exiting the calorimeter. For this reason, tungsten is used, as it has a very short nuclear interaction length of 9.94 cm.

This combination of materials results in 27.6 radiation lengths of electromagnetic calorimeter material, thus ensuring that no electrons or photons escape. Similarly, this configuration provides 9.94 nuclear interaction lengths of material in a very small space. This prevents all but the most energetic particles from escaping the FCal. In order to provide even more protection from escaping high-energy particles and shield the muon endcap region, a large uninstrumented brass plug is placed behind the FCal.

The very high particle flux can lead to ion buildup within the LAr. In order to avoid this, very narrow LAr gaps are used, thus reducing the drift time from the 450 ns in the EMB to 60 ns for the first module, and longer for the second and third (scaling with the LAr gap width). This is achieved by using gaps as small as 269 μm , 376 μm , and 508 μm for the first, second, and third modules respectively.

2.5 Muon Spectrometers

For particles to make it through the full depth of the calorimeters, they must either be non-interacting, a semi-stable MIP, or a shower that was not contained in the available material. The latter process is rare but important at high energies, and is discussed in Sect. 4.4. In the first case, which includes particles such as neutrinos, nothing can reasonably be done to contain them. At typical LHC energies, the only SM example of the second case is the muon. Muons thus provide very clean signatures for both searches and measurements, and a set of detectors designed to identify and measure them has been developed.

The ATLAS muon spectrometer consists of four sub-detectors, two in the barrel and two in the endcap. One sub-detector for each region provides a coarse, quick measurement, while the second provides a high precision particle momentum measurement. The coarse detectors cover the range of $|\eta| < 2.4$, and provide the ability to trigger on events with muon candidates. The much more precise detectors cover the range of $|\eta| < 2.7$, but are not fast enough to use in a trigger system. In this way, the two sets of two sub-detectors fulfil complementary roles. The full muon system is integrated with a series of toroidal magnets, placed outside of the calorimeter sys-




Fig. 2.16 A computer-generated schematic representation of the ATLAS muon spectrometer, with the various sub-detectors indicated [16]

tems, which provide the necessary magnetic field to measure particle momenta via their curvature. A schematic diagram of only the ATLAS muon system is provided in Fig. 2.16. A technical diagram of the muon system is provided in Fig. 2.17.

Resistive Plate Chambers and Thin Gap Chambers

The fast muon systems within the barrel region are made up of Resistive Plate Chambers (RPCs), while the endcap region is made of Thin Gap Chambers (TGCs). Both technologies produce a signal within 15–25 ns, which enables their use as triggering detectors. However, the cost of a very fast response appears in the measurement resolution. The RPC (TGC) provide an axial z (radial r) position resolution of 10 mm (2–6 mm) and an azimuthal ϕ resolution of 10 mm (3–7 mm).

The 606 RPC are split into three layers, and provide a total of 373×10^3 channels. They cover the range of $|\eta| < 1.05$, and are located approximately 7–10 m from the interaction point along the radial direction. The 3588 TGC are split into four layers, and consist of 318×10^3 channels. They cover $1.05 < |\eta| < 2.4$, and are located approximately 7, 13, 13.8, and 14.2 m from the interaction point along the axial direction.

Monitored Drift Tubes and Cathode Strip Chambers

The slower but much more precise muon detectors are the Monitored Drift Tubes (MDTs). The MDT are used in both the barrel and endcap regions, covering the full range of $|\eta| < 2.7$. The 1150 modules are split into three barrel layers and four endcap layers, forming a total of 354×10^3 channels. This structure provides an axial or radial (z or r) position resolution of $35 \mu\text{m}$. This is at least 57 times more precise than the trigger chambers, but the maximum MDT charge collection time




Fig. 2.17 A quarter-sectional view of the ATLAS muon spectrometer, showing each of the sub-detectors and the geometrical space that they occupy. Note that the MDT sub-detectors are named after where they are located, rather than the detector type. As such, BOL/BML/BIL is the barrel outer/middle/inner layer of the MDT, and similarly for the endcap EOL/EML/EIL (EEL is extra). The other sub-detectors are named after the technology, namely the RPCs, TGCs, and CSCs [2]

is 700 ns, thus it is also 28 times slower. This demonstrates the benefit of having separate trigger and precision sub-detectors, which together can provide much more impressive performance than would be possible independently.

Cathode Strip Chambers (CSCs) have a much more specific purpose, and are used only for the innermost module of the innermost layer of the endcap region, where $2.0 < |\eta| < 2.7$. As such, there are only 32 modules, with a total of 30.6×10^4 read-out channels. They provide a similar radial resolution of $40 \mu\text{m}$, but with a maximum collection time of 40 ns, providing a time resolution of approximately 7 ns per plane. This reduced collection time is necessary due to the increased particle flux in the forward region, which would overwhelm MDT if they were used.

2.6 Triggering System

The LHC bunch crossing rate as delivered to the ATLAS experiment for the 2012 dataset was every 50 ns, providing a collision frequency of 20 MHz. The full ATLAS detector cannot be read out and recorded at this frequency, necessitating the implementation of a triggering system to select only the most interesting events. When specifying the definition of what constitutes an interesting event, care must be taken to both cover all possible topologies required for physics analysis in addition to ensuring the trigger does not bias any particular process.

At a hadron collider, the relevant variable of interest in discriminating interesting inelastic collisions is the transverse momentum or transverse energy of objects in the event. As collisions become harder, the probability of a high transverse momentum object being observed grows, thus allowing for a selection of the most energetic processes which were beyond the reach of previous collider experiments. Typical objects considered are high p_T electrons, muons, taus, photons, jets, and MET. Definitions of all of these objects within ATLAS are provided in Chap. 3.

The probability for a given process to occur, and thus the frequency with which the trigger can be expected to observe such a process, varies dramatically between different final states. For example, the di-jet cross-section dominates all other hard-scatter processes by multiple orders of magnitude at hadron colliders, which in turn is a rare process compared to low-energy scattering events. As such, a naive definition of triggers based purely on the interesting object passing a specific p_T threshold will result in a dataset which is almost completely di-jet events. Defining a separate p_T threshold for each different object helps, but this is just the first step, as it also removes the sensitivity of the experiment to interesting physics which is below the threshold. In order to alleviate this problem, a *prescale* system is used. In this system, after the trigger decides that the event is interesting, a random number is used to decide whether or not to consider the event. This is governed by the prescale factor f_{prescale} , where $f_{\text{prescale}} = 100$ means that 1 in 100 events marked as interesting is recorded, and $f_{\text{prescale}} = 1$ is the special case of an unprescaled trigger.

The majority of ATLAS analyses make use of unprescaled triggers for simplicity, as they are purely interested in the high p_T regime and thus are willing to discard data from lower thresholds where more work is required to perform combinations of prescaled triggers. However, there are several very notable exceptions, such as the di-jet resonance analysis in Ref. [17]. This analysis makes use of a combination of many prescaled triggers with f_{prescale} up to 460,000 in order to extend its sensitivity to the lower di-jet mass (m_{jj}) regime. While this regime has already been covered by previous collider experiments, the enormous LHC luminosity allows for further improving the sensitivity at the lower energy regime. This is very important, as it is possible that new physics exists in the low energy regime, but that it is a very rare process, and thus it would otherwise be missed by relying purely on unprescaled triggers.

The ATLAS trigger system is comprised of three levels named Level 1 (L1), Level 2 (L2), and Event Filter (EF), and reduces the event rate through both prescaled and unprescaled triggers by a factor of 10^5 from 20 MHz to 200 Hz. In order to maintain this reduction factor, the exact trigger combination, thresholds, and prescale factors must all be updated on a regular basis as the LHC data conditions change, particularly when the centre of mass energy is increased.

If the trigger was infinitely fast, then a reduction of the rate from 20 MHz to 200 Hz would be a simple process. However, the trigger system must first wait for the detectors to complete their read-out of the event, and then must take time to build up a picture of the collision from the raw signals. Only once this is done can it make a decision about whether or not the event is interesting. This process takes much longer than the time between collisions, and thus substantial storage buffers are used.

This allows for retention of the complete detector read-out for many independent collisions, and thus is a fundamental component of the trigger system.

The first level of the trigger requires some level of activity in either the calorimeters or muon system, as these are the only parts of the detector which can perform a simple reconstruction and trigger decision within $2.5\,\mu\text{s}$ on an event-by-event basis, with the aim of reducing the event rate from 20 MHz to 75 kHz. The L1 system makes use of simple calorimeter towers (sums of energy of geometrical groups of calorimeter cells) for electron/photon and jet reconstruction, while muons make use of the trigger-specific chambers. If the L1 trigger decides an event was interesting, it defines one or more Regions of Interest (RoIs) where the η and ϕ coordinates of the interesting object(s) are located, and which specifies the specific trigger requirement that they passed in order to fire the trigger.

This information is then sent to the High-Level Trigger (HLT), which includes both the L2 and EF steps. The L2 trigger matches inner detector information to the RoI found by the L1 trigger, and makes a decision about whether the event still passes the appropriate trigger criteria within 40 ms, with the intent of reducing the event rate from 75 to 3.5 kHz. Successful events are then passed to the EF, which must reduce the trigger rate to the final level of 200 Hz. This is done by performing a full event reconstruction very similar to what is used for offline analysis, thus including calibrations, alignment corrections, and advanced algorithms. If an event is marked as interesting by the EF, it is recorded to disk in the RAW data format, which includes the information on all of the digitized energy deposits throughout the full ATLAS detector.

It is worth noting that the above description applies to the majority of triggers intended for use in physics analysis. There are also calibration triggers, minimum bias triggers, and other types of triggers which are typically highly prescaled, and select very different types of events. These types of triggers are not discussed further.

2.7 Simulation

The production of simulated events is a major requirement for ATLAS, given that almost all analyses make heavy use of Monte Carlo (MC) samples. MC can be used for many different purposes, ranging from estimating SM backgrounds to probing different signal hypotheses. MC can also be used to determine calibrations, such as for jets where a substantial amount of energy is not visible in the detector, but the MC retains information on the energy of the original particles. Providing a robust and efficient MC framework is thus a priority, and it is important that this simulation match data as closely as possible.

The ATLAS simulation framework is detailed in Ref. [18], and makes use of the GEANT toolkit from Ref. [19] for detector modelling. There are multiple steps in the generation processes, splitting the truth level step from the interaction between the truth particles and the detector. This allows for filters to be applied after the truth

level generation, thus enabling the selection of interesting events before performing the very computationally expensive digitization step.

2.7.1 Event Generation

The details of this step depend on the MC generator(s) used. In some cases, as will be discussed in Appendix B.1, the matrix element calculation and the partons showering are split into separate steps which are each handled in an independent generator. In other cases, both steps are performed within a single generator. The end product regardless of the specific generator choice is a set of truth-level events for a given process of interest, where the truth-level events are independent of the detector.

ATLAS makes use of many different Parton Distribution Function (PDF) sets to model the partons distributions within the colliding protons. This becomes particularly important for very high energy processes, where the LHC is operating in a new regime of parameter space. The Leading Order (LO) CTEQ6L1 [20] and MSTW2008LO [21] are particularly common for signal samples, while SM backgrounds are more likely to make use of the Next to Leading Order (NLO) CT10 [22] PDF set.

MC generators in ATLAS do not always provide a good description of data if taken as-is. In order to provide a better MC description, the partons shower model used by MC generator is often tuned, such as is detailed for Pythia6 [23] and Pythia8 [24] in Ref. [25]. In addition to tuning the specific generator showering model, the underlying event model used in the simulation is also adjusted. These tunes are typically derived with respect to a given PDF set, with CTEQ6L1 and CT10 again being the most common choices.

2.7.2 Detector Simulation and Digitization

This step propagates the stable truth particles produced in the previous step through a complete GEANT4 model of the ATLAS detector, where the current best understanding of interactions between relativistic particles and material is used. This includes the creation of electromagnetic and hadronic cascades from stable truth particles, and their interactions within the calorimeters. Electromagnetic showers are generally reasonably well modelled, but hadronic showers remain a topic of active investigation. As such, there are currently two showering algorithms widely used, with QGSP_BERT as the ATLAS default and FTFP_BERT as the alternative often used for systematic variation studies. The first is the quark gluon string model of nuclear fragmentation, described in Ref. [26], while the second is the FRITIOF model, as per Ref. [27]. In both cases, the Bertini cascade model is used, which is detailed in Ref. [28].

The full GEANT simulation is used wherever possible for MC production in ATLAS. However, this full simulation is very computationally expensive, especially when computing showers, and thus a fast simulation parametrization known as AFI I is often used where the highest level of precision is not necessary. This framework parametrizes the average GEANT response to a given type of particle in a given kinematic bin within the various calorimeters, thus providing look-up tables for energy deposits and interaction probabilities rather than performing full calculations for each step. The inner detector and muon spectrometer simulations are unchanged, and continue to use the GEANT model. This reduces the required computational time by an order of magnitude, at the cost of lower calorimeter precision. Many ATLAS analyses make use of MC created by a combination of both the full GEANT simulation and the fast AFI I parametrization.

The hard process dealt with so far is independent of pileup, which is a significant difference from real collision events. This is addressed by overlaying a set of minimum bias events created with a specially constructed Pythia8 LHC tune. These simulated minimum bias events add a pileup contribution, but do not encapsulate all of the possible detector effects, such as cavern background contributions to the muon spectrometer. Studies into the use of overlaid zero bias events from data are ongoing, as shown in Ref. [29], but have not yet become the standard. An example of the difference between overlaid MC minimum bias and data zero bias events is shown for muon variables in Fig. 4.14.

All of the particles, whether from the hard-scatter vertex or from the overlaid interactions, are passed through a complete model of the detector. Energy deposits (hits) are then assigned to each active portion of the detector which is affected. The digitization process described in Ref. [29] is then applied to convert the simulated energy deposits into electronic signals which would be observed in the actual detector. This is then reconstructed using the same procedure as data, giving rise to the Raw Digital Object (RDO) format, which primarily differs from the Raw Digital Object (RDO) format simply by the presence of additional truth information.

2.8 Computing and Data Management

The enormous amount of data produced either from LHC collisions or MC simulations requires a high-performance flexible distributed computing system. This requirement is met through the implementation of the Worldwide LHC Computing Grid (WLCG), which is composed of four levels of nodes, called Tiers. All of the events are initially recorded and reconstructed at the CERN Tier 0 (T0) data centre, which contains approximately 20 % of the computing power of the full WLCG.

The raw datasets and first reconstruction are then distributed around the world to a small number of Tier 1 (T1) nodes, such as TRIUMF in Canada, which represents 5 % of the full grid computing power. These datasets in turn can be requested by Tier 2 (T2) computing centres, which are typically hosted by universities or groups, and which provide additional computing resources as needed. The final level of Tier 3

(T3) services have no defined structure within the WLCG, and can vary from local clusters of computers to a single laptop connecting to the WLCG to download data or make processing requests.

Within ATLAS, there are a few main data formats in use, which facilitate easier management. The RAW and RDO formats are essentially restricted to T0 and T1 sites, with very little additional access available due to the very large computing resources and very complicated processing interface required to reconstruct events from raw events. The next level is the Event Summary Data (ESD), which contains all of the reconstructed data, and thus is a very large data format. These are typically not stored, but in some very specialized cases they are required, and thus the functionality exists for users to access and work with ESD.

An additional processing step by which the majority of the individual cell information is discarded and other space saving steps are taken forms the Analysis Object Data (AOD). This is the intended format for primary use within ATLAS, where all of the necessary information for reconstruction of interesting objects remains. For example, jets or the MET can be properly reconstructed at this level, as their necessary building blocks remain. In practice, this format is rarely used, although the development of the xAOD for Run-II looks set to change this by addressing many of the limitations of the AOD format.

Users who do work with ESD or AOD will typically also apply some additional analysis selection to reduce the size, thus producing Derived Event Summary Datas (DESDs) or Derived Analysis Object Datas (DAODs). However, the most common data format in use in Run-I applies a third level of filtering, earning it the name of the Derived³ Physics Data (D3PD), in addition to flattening the data format. This means that rather than storing the data in objects, every variable is stored as a simple type, such as an integer, vector (list) of integers, or so on. This format typically is created on a per-analysis or per-working group level, and applies specific selection cuts, triggers, and so on to further reduce the size of the file. Many groups found this necessary in order to reduce datasets to a manageable level. Even with this amount of reduction, a typical analysis will work with many TeraBytes (TBs) of D3PDs.

Many private formats where even further levels of derivation or with redesigned information management systems exist, but they are not officially supported. The above list is the officially supported set of ATLAS data formats. These supported data formats are typically stored across T0, T1, and T2 sites, while private data formats are typically restricted to T3 clusters or local use.

References

1. L. Evans, P. Bryant, LHC machine. J. Instrum. **3**(08), S08001 (2008), <http://stacks.iop.org/1748-0221/3/i=08/a=S08001>
2. ATLAS Collaboration, The ATLAS experiment at the CERN large hadron collider. J. Instrum. **3**, S08003 (2008)
3. AC Team. The four main LHC experiments. CERN-AC-9906026 (1999)
4. M. Benedikt et al., *LHC Design Report* (CERN, Geneva, 2004)

5. C. Lefevre, *LHC: the guide*, CERN-Brochure-2008-001-Eng (2008)
6. ATLAS Collaboration, Luminosity public results, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>
7. ATLAS Collaboration, Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C* **73** 2518 (2013). [arXiv:1302.4393](https://arxiv.org/abs/1302.4393) [hep-ex]
8. J. Pequeno, Computer generated image of the whole ATLAS detector (2008)
9. J. Pequeno, Computer generated image of the ATLAS inner detector (2008)
10. J. Pequeno, Computer generated image of the ATLAS calorimeter (2008)
11. A.V. Kotwal, C. Hays, Electromagnetic shower properties in a lead-scintillator sampling calorimeter. *Nucl. Instrum. Methods A* **729**, 25–35 (2013)
12. M. Aharrouché et al., Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam. *Nucl. Instrum. Methods A* **568**(2), 601–623 (2006)
13. C. Grupen, B. Shwartz, *Particle Detectors* (Cambridge University Press, Cambridge, 2008)
14. Particle Data Group Collaboration, K. Olive et al., Review of particle physics. *Chin. Phys. C* **38**, 090001 (2014)
15. N. Ilic, Performance of the ATLAS liquid argon calorimeter after three years of LHC operation and plans for a future upgrade. Tech. Rep. ATL-LARG-PROC-2013-016, CERN, Geneva (2013)
16. J. Pequeno, *Computer Generated Image of the ATLAS Muons Subsystem* (2008)
17. ATLAS Collaboration, Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector. [arXiv:1407.1376](https://arxiv.org/abs/1407.1376) [hep-ex]
18. ATLAS Collaboration, The ATLAS simulation infrastructure. *Eur. Phys. J. C* **70**(3), 823–874 (2010). doi:[10.1140/epjc/s10052-010-1429-9](https://doi.org/10.1140/epjc/s10052-010-1429-9)
19. GEANT4 Collaboration, GEANT4: a simulation toolkit. *Nucl. Instrum. Methods A* **506**, 250–303 (2003)
20. P.M. Nadolsky et al., Implications of CTEQ global analysis for collider observables. *Phys. Rev. D* **78**, 013004 (2008). [arXiv:0802.0007](https://arxiv.org/abs/0802.0007) [hep-ph]
21. A. Martin, W. Stirling, R. Thorne, G. Watt, Parton distributions for the LHC. *Eur. Phys. J. C* **63**, 189–285 (2009). [arXiv:0901.0002](https://arxiv.org/abs/0901.0002) [hep-ph]
22. H.-L. Lai et al., New parton distributions for collider physics. *Phys. Rev. D* **82**, 074024 (2010). [arXiv:1007.2241](https://arxiv.org/abs/1007.2241) [hep-ph]
23. T. Sjostrand, S. Mrenna, P.Z. Skands, PYTHIA 6.4 physics and manual. *J. High Energy Phys.* **0605**, 026 (2006). [arXiv:hep-ph/0603175](https://arxiv.org/abs/hep-ph/0603175) [hep-ph]
24. T. Sjostrand, S. Mrenna, P. Skands, PYTHIA 8.1. *Comput. Phys. Commun.* **178**, 852 (2008)
25. ATLAS Collaboration, ATLAS tunes of PYTHIA 6 and Pythia 8 for MC11. Tech. Rep. ATL-PHYS-PUB-2011-009, CERN, Geneva (2011)
26. G. Folger, J. Wellisch, String parton models in GEANT4, eConf C0303241 (2003) MOMT007. [arXiv:nucl-th/0306007](https://arxiv.org/abs/nucl-th/0306007) [nucl-th]
27. H. Pi, An event generator for interactions between hadrons and nuclei FRITIOF version 7.0. *Comput. Phys. Commun.* **71**(12), 173–192 (1992), <http://www.sciencedirect.com/science/article/pii/001046559290082A>
28. H.W. Bertini, Intranuclear-cascade calculation of the secondary nucleon spectra from nucleon-nucleus interactions in the energy range 340 to 2900 MeV and comparisons with experiment. *Phys. Rev.* **188**, 1711–1730 (1969). doi:[10.1103/PhysRev.188.1711](https://doi.org/10.1103/PhysRev.188.1711)
29. J.D. Chapman et al., The ATLAS detector digitization project for 2009 data taking. *J. Phys. Conf. Ser.* **219**(3), 032031 (2010), <http://stacks.iop.org/1742-6596/219/i=3/a=032031>

<http://www.springer.com/978-3-319-44452-9>

Searching for Dark Matter with the ATLAS Detector

Schramm, S.

2017, XXIII, 324 p. 145 illus., 135 illus. in color.,

Hardcover

ISBN: 978-3-319-44452-9