

Study of Jet Substructure in the ATLAS Experiment using Distributed Analysis within Spanish Tier-2 Infrastructures

Elena Oliver García

Dirigido por:

José Francisco Salt Cairols and Santiago González de la Hoz

Tesis Doctoral

Octubre 2013



VNIVERSITAT
DE VALÈNCIA

Facultat de Física

Departament de Física Atòmica Molecular i Nuclear



Dr. José Francisco Salt Cairols.
Profesor de Investigación del Consejo Superior de Investigaciones Científicas.
Dr. Santiago González de la Hoz.
Contratado Doctor del Departamento de Física Atómica, Molecular y
Nuclear de la Universidad de Valencia.

CERTIFICAN:

Que la presente memoria '**Study of Jet Substructure in the ATLAS Experiment using Distributed Analysis within Spanish Tier-2 Infrastructures**' ha sido realizada bajo nuestra dirección en el Departamento de Física Atómica, Molecular y Nuclear de la Universidad de Valencia por Doña Elena Oliver García y constituye su tesis para optar al grado de Doctor en Física.

Y para que así conste, firmamos el presente Certificado.



Firmado
Dr. José Francisco Salt Cairols



Firmado
Dr. Santiago González de la Hoz

El trabajo descrito en esta tesis se ha llevado a cabo en el *Instituto de Física Corpuscular* (IFIC) en Valencia, España.



El IFIC es un centro mixto de la *Universitat de València* y el *Consejo Superior de Investigaciones Científicas* (CSIC)



Contents

Preface	3
1 Introduction	5
1.1 Standard Model and beyond	5
1.2 Large Hadron Collider	11
1.3 Grid Computing	15
1.4 Worldwide LHC Computing Grid project	17
2 ATLAS Experiment	21
2.1 Inner Detector	23
2.2 Calorimetry	25
2.3 Muon System	27
2.4 Trigger System	28
3 ATLAS Software and Computing Model	31
3.1 ATLAS Computing Model	31
3.2 Event Data Model	34
3.3 ATLAS Software	36
3.3.1 Athena Framework	36
3.3.2 Database	38
3.4 Computing testing and challenges before starting-up of LHC .	38
3.5 Conclusions	40
4 ATLAS Distributed Analysis System	43
4.1 Production and Distributed Analysis system	43
4.2 Distributed Analysis End-User Tools	46
4.2.1 Tools to access the Data	46
4.2.2 Tools to send analysis jobs to the Grid	48
4.3 Monte Carlo Production for ATLAS	50
4.3.1 Simulation Process	50
4.3.2 ATLAS Monte Carlo Production workflow on PanDA .	51

4.4	ATLAS Distributed Analysis Progress	52
4.5	Distributed Analysis Support Team	54
4.6	Real Distributed Analysis Workflow Example	56
5	The ATLAS Spanish Tier-2 Federation	61
5.1	ES-ATLAS-T2 Resources	63
5.2	Reliability and Availability of the Spanish ATLAS Tier-2 . . .	66
5.3	Data Transfers of the ES-ATLAS-T2	68
5.4	Data stored in ES-ATLAS-T2	84
5.5	ATLAS Production and Analysis jobs at ES-ATLAS-T2	94
5.5.1	Production Jobs	94
5.5.2	Analysis Jobs	97
5.6	User Support at ES-ATLAS-T2	102
5.7	Conclusions	103
6	Jet Substructure	105
6.1	Motivation: jets and boosted objects	105
6.1.1	Jet algorithms	107
6.1.2	Jet substructure variables	108
6.2	Dataset and reconstruction	109
6.3	Monte Carlo samples	111
6.4	Detector-level distributions	111
6.5	Pile-up dependence	113
6.6	Systematic uncertainties	113
6.7	Data correction	115
6.8	Results	115
6.9	Conclusions and perspectives	119
7	Summary/Resumen	121
7.1	Summary in English	121
7.1.1	Motivation	121
7.1.2	Introduction	121
7.1.3	Distributed Analysis framework	123
7.1.4	ATLAS Spanish Tier-2 Federation	124
7.1.5	Jet substructure	125
7.1.6	Conclusions	126
7.2	Resumen en español	127
7.2.1	Motivación	127
7.2.2	Introducción	128
7.2.3	Análisis distribuido	130
7.2.4	El Tier-2 español federado de ATLAS	131

7.2.5	Subestructura de jets	132
7.2.6	Conclusiones	133
Contributions and Impact		137
Agradecimientos		141
Bibliography		143

Preface

High Energy Physics has been developed for the last decades in huge accelerators. The present theoretical Model to explain the behaviour of the particle physics is the so-called Standard Model (SM). It is characterized by many parameters and observables. The experimental measurements of these parameters have shown a broad agreement with the SM predictions. Despite of this, some experimental evidences are not explained like neutrino mass or the dark matter. Other theories Beyond Standard Model (BSM), such as Supersymmetry or Extra Dimensions, predict these evidences and add new particles at high energies.

To explore this new Physics, the Large Hadron Collider (LHC) with a energy in TeV scale has been built. This collider is located at CERN (European Organization for Nuclear Research) in Geneva. LHC started in November 2009 and has just reached 8 TeV energy of center of mass. One of the four detector experiments, A Large Toroidal LHC Apparatus (ATLAS) was designed for general Physics purpose and it has taken a integrated luminosity of 27.03 fb^{-1} with around 1.80×10^{15} collisions for 3 years (2010-2012).

The data collected by ATLAS have been stored, processed and made available to thousands of people. These have been achieved thanks to the Grid paradigm. It consists of computing farms interconnected in several places around the world. With this computing effort, the LHC scientists have been able to access to data which are taken with a rate of 25 petabytes per year.

The Instituto de Física Corpuscular (IFIC) which has been evolved in ATLAS detectors, has contributed in this Grid computing effort by participating in a R&D project funded by the HEP National program: the development, maintenance and operation of the ATLAS Spanish Tier-2 (ES-ATLAS-T2) together with computing farms in Institut de Física d'Altes Energies (IFAE, Barcelona) and Universidad Autónoma de Madrid (UAM).

The studies reported in this thesis have been performed in that ATLAS framework. It can be divided in two different parts: the first part reflects

the work done in the comprehension and the use of the Distributed Analysis, the monitoring of the ES-ATLAS-T2, and the support for users in Grid computing at IFIC and for the ATLAS collaboration. The second part is the study of the Jet Substructure at LHC TeV scale. This measurement is interesting in its own right, although it is also a crucial commissioning step for the identification of boosted particles. All this was developed in the IFIC Exotics Physics Group in collaboration of several institutes.

This work starts first with a summary of the SM, adding a description of LHC and Grid in Chapter 1. Then, the ATLAS experiment is described in terms of detector layers in Chapter 2. Later, the computing and software that involves ATLAS is shown in Chapter 3. In Chapter 4, the Distributed Analysis in ATLAS is presented with its components, tools and help given for users including a real example. Next, the particular case of the ES-ATLAS-T2 is shown in Chapter 5, displaying its infrastructure, operations and performance. After that, the description of Jet Substructure and the ATLAS commissioning results are presented in Chapter 6. Finally, a summary of all this work is shown in Chapter 7.

Chapter 1

Introduction

High Energy Physics aims to understand the elementary particles of matter and their interactions. Huge accelerators are needed to observe these particles and study their properties. Nearly, a theoretical framework known as Standard Model (SM) describes all these observations. In this chapter the theory will be briefly introduced, along with some observations that were confirmed by the formalism. The focus of the second part of this chapter is on some of the issues of the theory and proposals that have been made to extend the SM. Then, the Large Hadron Collider will be described which has provided the collisions for this study in the ATLAS experiment that is discussed in the next chapter. Finally, the basis of Grid Computing will be shown, the LHC collaboration choice for this computing model, and the brief description of the Worldwide LHC Grid Computing project.

1.1 Standard Model and beyond

The current model of the most fundamental components of matter and the interactions among them is called Standard Model [1]. Table 1.1 represents the list of particles implied. This model is based on Quantum Field Theory: the union of Quantum Mechanics and Special Relativity developed in 20th century. The elementary particles that make up the matter have the inherent property called spin with half-integer and so-called fermions. They follow the Pauli exclusion principle that establishes that two identical fermions cannot be in the same quantum state at the same time.

Each fermion has its antiparticle with the same mass but different electric charge. Not taking into account these antiparticles, there are 12 types of matter's components that are grouped in 3 generations. Only the first generation is stable and forms the matter we are made of. Particles in the second

Table 1.1: Elementary particles properties

Family	Sub-family	Name	Symbol	Mass	Electric Charge	Spin
Fermions	Leptons	electron	e	0.511 MeV	-1	$1/2$
		muon	μ	105.7 MeV	-1	$1/2$
		tau	τ	1.777 GeV	-1	$1/2$
		electron neutrino	ν_e	< 2 eV	0	$1/2$
		muon neutrino	ν_μ	< 0.17 MeV	0	$1/2$
		tau neutrino	ν_τ	< 15.5 MeV	0	$1/2$
	quarks	up	u	1.5 to 3.3 MeV	$2/3$	$1/2$
		charm	c	1.27 GeV	$2/3$	$1/2$
		top	t	171.2 GeV	$2/3$	$1/2$
		down	d	3.5 to 6.0 MeV	$-1/3$	$1/2$
		strange	s	104 MeV	$-1/3$	$1/2$
		bottom	b	4.20 GeV	$-1/3$	$1/2$
Bosons		photon	γ	0	0	1
		W	W	80.398 GeV	± 1	1
		Z	Z	91.188MeV	0	1
		gluon	g	0	0	1
		Higgs	H	125 GeV	0	0

and third generations are more massive and unstable. For that reason, these particles decay in ground state leading to others through SM interactions.

Each particle has a group of properties or charges that determine if it is affected by a certain interaction or not. The electric charge is related to the electromagnetic interaction, the color charge to the strong interaction and all particles are sensitive to the weak interaction. According to this criterion, they can be classified in leptons and quarks. Leptons are the elementary particles without color charge. Three of these leptons have electric charge -1 (their antiparticles +1). The electron (e) was postulated in 1838 by Richard Laming to explain the chemical properties of atoms. However, the discovery is attributed to J.J. Thompson who indicated cathode rays really were unique particles, rather than waves. The muon (μ), in the second generation, was discovered during a study of cosmic radiation by Carl D. Anderson and Seth Neddermeyer in 1936 [2]. The tau lepton (τ) was produced in a e^+e^- collider called SPEAR at SLAC [3] between 1974 and 1977 by Martin Lewis Perl and his group.

The other three leptons only interact weakly without electric charge, the so-called neutrinos. They were postulated by Pauli in 1930 to explain the apparent violation of the energy conservation in β -decay [4]. It was confirmed the detection in 1956 [5]. These leptons are massless in the SM, but they have non-zero masses as it has been shown in the long baseline neutrino oscillation

experiments [6]. In a generation, the charged lepton is related to its neutrino by weak interaction. They are therefore known as electron, muon and tau neutrinos.

The elementary fermions with color charge are called quarks. They have fractional electric charge and in each generation there is a quark with $+\frac{2}{3}$ (up, charm, top) and another one with $-\frac{1}{3}$ (down, strange, bottom). Quarks come in three possible colors: red, green and blue. Because of the color charge they have, a property called color confinement is produced where quarks cannot be observed as free particles. Hadrons are the combination of quarks with no net color charge (a $q\bar{q}$ pair is called a meson, a qqq or $\bar{q}\bar{q}\bar{q}$ triplet is called a baryon). The existence of quarks was inferred by Murray Gell-Mann and George Zweig in 1964 [2] from regularities in the zoo of hadrons produced at colliders.

In the SM, the strong interaction is described by Quantum Chromodynamics and the electromagnetic and weak interactions by the electroweak interaction. The fourth fundamental interaction, gravity, is too weak at high energy scales. In Figure 1.1, a representation of all the interactions is shown. Also, in Table 1.2, the interactions properties are displayed comparing the strength with the electromagnetic force in the situation of two u quarks in the distances 10^{-18} m (upper limit for the size of quarks and electrons) and 3×10^{-17} m (over the range of the weak force). The strength of strong and weak interactions depends on the distance.

In the SM, interactions are produced by exchange of a particle with integer spin among leptons and quarks. These particles are called bosons. In the case of Quantum Chromodynamics, the exchange particle is the gluon. It doesn't have mass and there are 8 possible types, each one with a color and an anti-color. Thanks to the gluon's color charge, it can interact with itself. Gluons form the glue that holds together hadrons, specially protons and neutrons, the components of the atomic nucleus. Whenever a quark (or gluon) is ejected from its hadron, a spray of hadrons follows approximately the direction of this parton. These so-called jets are the Physics subject of this thesis.

The corresponding boson to the electromagnetic interaction is the photon that has no mass. There are three bosons for the weak interaction: W^+ , W^- and Z^0 discovered at CERN in 1983 [8, 9]. The interaction is much weaker than the strong interaction thanks to their large mass. The weak interaction produces flavour changes (the mentioned types of leptons and quarks), and




Figure 1.1: The current particles and interactions displayed in the nature.

Table 1.2: The interactions comparison. The strengths of the interactions are shown relative to the strength of the electromagnetic for two u quarks separated by the specified distances [7]. The distances are 10^{-18} m which is upper limit for the size of quarks and electrons, and 3×10^{-17} m which is over the range of the weak force. The strength of strong and weak interactions depends on the distance.

Interaction	Gravitational	Weak	Electromagnetic	Strong
Acts on	Mass-Energy	Flavour	Electric Charge	Color Charge
Particles experiencing	All	Quark, Lepton	Electrically Charged	Quarks, gluons
Particles mediating	Graviton	W^+, W^-, Z^0	Photons	Gluons
Strength at 10^{-18} m	10^{-41}	0.8	1	25
Strength at 3×10^{-17} m	10^{-41}	10^{-4}	1	60

thus allows transitions between the generations. In the nucleus scale, weak interaction is responsible for radioactivity.

According to this theoretical formalism, the electromagnetic and weak interactions could only accommodate different aspects of the same interaction. Initially, the formalism had massless W^+ , W^- and Z^0 bosons and fermions. This was in open contradiction with the experimental observations. To solve this contradiction, a new boson was introduced thanks to a new field added by Higgs to break the symmetry spontaneously [10]. This Higgs boson is one of the reasons for the construction of LHC.

The SM was a successful model for the agreement with the experimental results about the Z^0 boson discovery at CERN [11] in 1983 and the top quark at Tevatron [12] in 1995. Despite of this, the SM cannot explain some experimental evidences. For instance, it does not consider the neutrino mass (experimentally confirmed). As indicated before, the gravity is not included. Other exchanging boson should be existed hypothetically, called graviton, but it has not been observed at the moment. Also, SM does not give dark matter candidates. Dark matter is the name for the unobservable mass that galaxies should have according to their movements in General Relativity [13]. None of the particles in SM is a dark matter candidate to affect the gravity at macroscopic level. Another issue not explained in the SM is why matter and antimatter are not in the same proportion in the universe. The most of the part of the universe observed is matter, although particle-antiparticle pairs are produced from the photons without any preference for one or another.

The value of the Higgs mass in the SM gives rise to the so-called hierarchical problem. For the Higgs boson, the mass with quantum corrections tends to roughly the Planck Scale (10^{18} GeV), which is the energy where all interactions can be unified including the gravity. However, if the Higgs boson has this large mass, the mass contribution to the rest of particles will not be produced. Therefore, it is natural to expect a mass of the order of the electroweak scale (100 GeV). On 4th July 2012, a new Higgs-like boson has been announced by CMS and ATLAS experiments with a mass of 125 GeV [14]. An overview of the energy scales is shown in Figure 1.2 to get an idea of the energy differences.

Several extensions of the SM have been provided to solve the hierarchical problem. The most popular model is Supersymmetry (SUSY) [15]. It considers that there is a new particle (superparticle) for each matter particle (fermion), which is a boson with the same flavour. In a similar way, there is a new superparticle for each exchange boson, which is a fermion with the

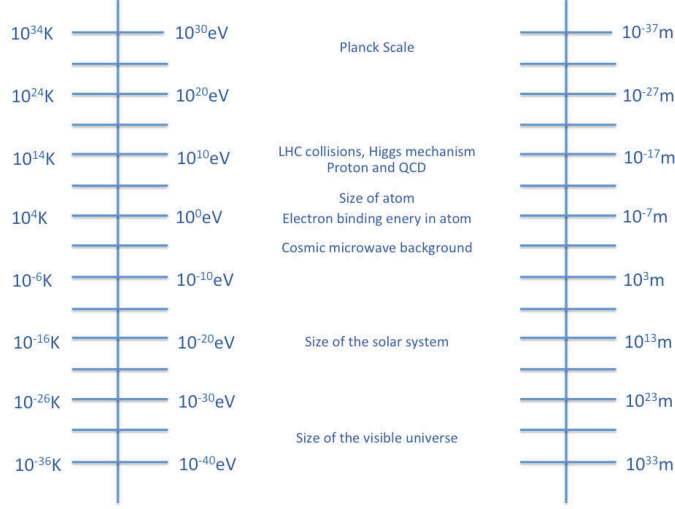


Figure 1.2: Energy scales in the nature. The great difference between Higgs mechanism and Planck scale is shown.

same flavour. This cancels the divergent contributions to the Higgs boson mass by the terms of the corresponding superpartners. The fact that none of the superparticles has been discovered indicates that, if SUSY exists, it is a broken symmetry. The breaking of the symmetry allows superpartners to acquire much larger masses (of the order of TeV) than those of their correspondent SM particles. The LHC experiments offer an excellent opportunity to discover squarks and gluinos with masses up to several TeVs. SUSY yields dark matter candidates. A graphical representation of the corresponding superparticles is shown in Figure 1.3.




Figure 1.3: Imaginary representation of particles and their correspondent superparticles at the bottom with higher mass.

Another popular Beyond Standard Model (BSM) proposal is called Extra Dimensions. These models consist in adding one or more space dimensions to the known three. These models have been studied in the beginnings of 20th century when Quantum Physics started. The so-called Kaluza-Klein [16] models attempt to unify gravity and electromagnetism in 5 dimensions. More recent proposals are Arkani-Hamed-Dimopoulos-Dvali (AAD) [17] and Randall-Sundrum (RS) model [18] which proposes the gravity can propagate between branes (i.e., 3+1 dimensions portions) and therefore this can explain why the gravity is so weak and would resolve the hierarchical problem.

All these extensions of the SM make testable predictions. The LHC experiments discussed in the next Section form an excellent opportunity to either confirm the existence of new particles or to severely constrain the parameters of such models.

1.2 Large Hadron Collider

The Large Hadron Collider (LHC) [19] is a hadron accelerator placed at CERN substituting the previous one, the Large Electron-Positron Collider (LEP) [20]. It consists in the collision of two proton beams that can rise a centre of mass energy maximum of 14 TeV, at a high luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

The LHC describes a circumference of 27 km, and 100 metres of depth in the region of France-Switzerland frontier. The energy of each beam depends on the collider length and the magnetic field from the dipolars magnets. Apart of this, there are magnets of higher orders to focus the beam and provide stability. The two proton beams share the mechanic and cryogenic structures but they circulate through separate vacuum tubes with different magnetic fields. 3000 bunches with 1.15×10^{11} protons can circulate each one is time separated in 25 ns. A schema of the different phases of the acceleration process is shown in Figure 1.4.

There are four impact points at LHC where the experiments are placed for detecting the collisions (see Figure 1.5): The Compact Muon Solenoid (CMS), A Large Ion Collider Experiment at CERN (ALICE), Large Hadron Collider beauty experiment (LHCb) and A Large Toroidal LHC Apparatus (ATLAS). This work has been developed in the framework of ATLAS and therefore it is described in Chapter 2. A brief description of the rest of LHC experiments is shown below:




Figure 1.4: The LHC process of the acceleration of protons is shown in this schema.

1.2. Large Hadron Collider

- CMS [21] is a detector for general purpose, that means it covers variety of subjects. It is similar to ATLAS, but the main difference is in the superconductor solenoidal magnet to determine the charged particles momentum and this makes the CMS more compact.
- LHCb [22] is focused in the study of the quark b to analyse symmetry violations and to determine elements of Cabibbo-Kobayashi-Maskawa matrix (CKM, [1]) which explain the particles flavour exchange. LHCb is similar to a spectrometer with one arm because of the low polar angle of the b-quark production.
- ALICE [23] is based on Heavy Ion collisions (as lead at 2.76 TeV and luminosity $2 \times 10^{27} \text{cm}^{-2} \text{s}^{-1}$) for the study of Quark-Gluon Plasma, a state of matter. This plasma supposedly existed 10^{-5} seconds after the Big Bang.




Figure 1.5: LHC draw with the experiments locations

The LHC started to work on 20th November 2009 with a center-of-mass energy of 900 GeV. Three days later, the first collisions were produced in the detectors. The energy was increased to 2.36 TeV on 30th November of the same year, establishing a new energy record in the history of colliders. In the beginning of 2010, energy of center of mass reached 7 TeV. During this time of period, a luminosity peak was achieved in $4.67 \times 10^{32} \text{cm}^{-2} \text{s}^{-1}$. The integrated luminosity reached 5fb^{-1} ($1 \text{fb} = 10^{39} \text{cm}^2$) in CMS and ATLAS in October 2011.

In April 2012, the energy of center of mass changed to 8 TeV. In March 2013, the LHC stopped its activity with a integrated luminosity over 27 fb^{-1} in the general purpose experiments. In Figure 1.6 the delivered luminosity over the time is observed. It is in downtime until the end of 2014 to refurbish the machine for operations at the design energy of 14 TeV.




Figure 1.6: Cumulative luminosity versus day delivered to ATLAS during stable beams and for proton-proton collisions in logarithmic y-scale in 2010 (green), 2011 (red) and 2012 (blue) running.

The LHC with a energy in center of mass at 8 TeV in 2012 was 4 times more than Tevatron [12] which closed in 2011. The LHC luminosity reached was 27 fb^{-1} , higher than Tevatron luminosity accumulated in 10 years.

In this new kinematic regime, known particles behave in different way than in previous experiments. The massive SM particles have mass around 100 GeV. At LHC, these particles can be produced with transverse momentum higher than their mass. Generally, heavy particles with mass $p_T \gg m$ are called boosted objects [24]. The LHC is the first accelerator where boosted objects are produced abundantly. The production quark tops rate at LHC is compared with the Tevatron production and the designed LHC energy [25] in table 1.3.

In 1995, 10000 top collisions were produced at Tevatron and only 23 were expected with an invariant mass of 1 TeV. At LHC (a real tops factory) in 2012, top quarks were produced 20 times more frequently than at Tevatron, and the boosted tops portion increased several orders of magnitude. 60000

1.3. Grid Computing

Table 1.3: Top pairs production evolution according to Tevatron, LHC and the future. The boosted tops number grows significant. From CERN Courier [25].

# events expected	Tevatron run II 10 fb^{-1} , 1.96 TeV	LHC 2012 20 fb^{-1} , 8 TeV	LHC desing 300 fb^{-1} , 13 TeV
Inclusive $t\bar{t}$ production	60000	4000000	200000000
Boosted production $m_{t\bar{t}} > 1$ TeV	23	60000	5200000
Highly boosted $m_{t\bar{t}} > 2$ TeV	0	480	110000

top pairs will be produced in all 2012 with $m_{t\bar{t}} > 1$ TeV. To study boosted objects, the Jet Substructure is a essential tool for their analysis and is explained in Chapter 6.

1.3 Grid Computing

During the LEP experiment performance [20], big computing challenges have been noticed in storage, processing and access to a considerable amount of information in events from the detectors. Essentially at the LEP era the computing model rely on a big centre at CERN plus several National Computer centres in the most important countries. This model was not enough for the LHC experiments.

By the end of 90's, it was realized that computing and storage challenges would be 100–1000 times more difficult in the LHC experiments. It was calculated that LHC would produced 15 PB of data per year in 15 years. Apart from this, 5000 physicists around the world would access to this information ¹. The analysis of all the data, including comparison with simulations, would require around 260,000 CPUs ² to process them. All the countries of the LHC collaboration were necessary to contribute in it. Since internet connections were improved regarding to the LEP beginnings, a new model was decided to use which allows the access of distant computing centres: a model based on Grid technologies, whose basis were developed by Ian Foster and Carl Kesselman [26].

The Grid is to share computer power and data storage capacity connected by internet. This technology connects computers that are spread over the world. Just as the World Wide Web enables access to information, the Grid enables access to computing resources.

¹At present, LHC experiments are producing 25 PB per year and more than 8000 physicists are working on them

²evaluated in 2012

The Grid provides advantages to the LHC computing [27]: The huge costs of maintaining and upgrading the resources are more easily sustainable in a distributed environment, where individual institutes and national organizations can finance local computing resources and share responsibility for these, while still contributing to the global objective. Multiple copies of data and automatic reassigning of tasks to available resources ensures a good improvement of resources and facilitates access to the data for all the scientists involved, independent of geographical location. Because all time zones are involved, it is possible round-the-clock monitoring and support.

The Grid implies to have adequate levels of network bandwidth among contributing resources, installing the correct software versions in various locations, dealing with heterogeneous hardware, managing and protecting the data over the lifetime of the LHC, and providing accounting mechanisms. It requires the development and deployment of middleware that enables the use of distributed resources as a virtual big computer. It will be defined with other concepts typically from Grid [28]:

- Middleware is the software that organises and integrates the different computing farms of Grid. Its main role is to automate all the “machine to machine” negotiations required to interlace the computing and storage resources and the network into a single computational factory. It lets the various elements (servers, storage, networks, etc.) to participate in a unified Grid environment. The mainly middleware components are:
 - Worker Node (WN) is the essential element which processes input data.
 - Computing Element (CE) receives job requests and delivers them to its assigned WNs. It provides an interface to the local batch queueing systems. A CE can manage one or more WNs. A Worker Node can also be installed on the same machine as the Computing Element.
 - Storage Element (SE) is the group of services responsible for a storage resource on the Grid. Services may include data access, quota management or space management. The SE is a sufficiently flexible interface to Grid storage units that allows interoperability of the Grid without forcing local sites to change their existing software stack.

- User Interface (UI) allows users to access the Grid facility and receives the input Grid jobs written in Job Description Language (JDL).
 - Workload Management System (WMS) is the module that receives users requests and queries the Information Index to find suitable resources. It essentially provides the facilities to manage jobs (submit, cancel, suspend/resume) and to ask about their status.
 - Information Service keeps information about the available resources to the Grid. One of the most used is the Berkeley Database Information Index (BDII) [29].
- Virtual Organization (VO) consists of people or institutions with a common objective as working in a scientific experiment. They need usually the same tools or the same data. Then, a certain Grid can have many different VOs that choose to share their resources, giving direct access to computers, programs, files, data, sensors and networks.
 - Grid Certificate is a digitally signed statement from one entity as a Certificate Authority (CA). A grid certificate gives to the users the authorization to run jobs on the Grid to access the data stored. User has to request a Grid Certificate to the CA of his region. This authority says that this public key of another entity like a Grid user has some specific value.
 - The Virtual Organization Management Service (VOMS) middleware package allows the definition of user groups and roles within the ATLAS VO. For the moment, there is a VOMS group for each Physics Working Group, Combined Performance group and Detector System, as well as a generic one, and another one for software testing, validation and central production activities.

Implementing Grid model implies an effective coordination among all LHC collaboration centres which was established in the Worldwide LHC Computing Grid project.

1.4 Worldwide LHC Computing Grid project

The Worldwide LHC Computing Grid project (WLCG) [27] was created to set up and operate the Grid computing infrastructure for the simulation, processing and analysis of the data coming from the LHC experiments. It

is a global collaboration of more than 140 computing centres in 35 countries for the 4 LHC experiments. In Figure 1.7 you can see the map of all centres of the project.

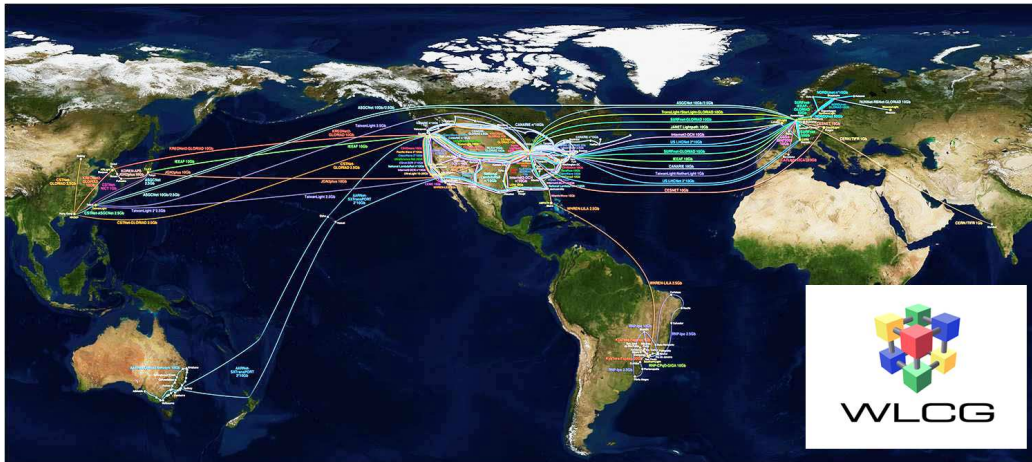


Figure 1.7: This map represents all the locations of the Grid infrastructure in the WLCG project around the World

A common computing architecture was developed to ensure the LHC requirements. It was in charge to design the schedule steps to achieve the Grid computing system for the several centres to be ready before the data-taking. Common software has been developed and new computing technologies have been studied to improve the system. Also, WLCG has been following the centres performance by using periodical tests.

It was defined the different organization groups assigning their corresponding tasks inside the project. The description of some components is presented:

- The Grid Deployment Board (GDB) is the forum within the Project where the computing managements of the experiments and the regional computing centres discuss and take, or prepare, the decisions necessary for planning, deploying, and operating the LHC Computing Grid.
- The Resources Review Board (RRB) representatives of each Experiment's Funding Agencies and the managements of CERN ,and of each Experiment's Collaboration. It is chaired by the CERN Director for Research and Computing.
- The Computing Resources Scrutiny Group (C-RSG) has to inform the decisions of the Computing Resources Review Board (C-RRB) for the LHC experiments.

The WLCG architecture consists of an concerted set of services and applications running on the Grid infrastructures. These infrastructures at the present are integrated in the European Grid Infrastructure (EGI) ³ [32] project in Europe and also in Asia and South America, the Open Science Grid (OSG) [33] project in the U.S.A. and the Nordic Data Grid Facility (called NorduGrid) [34] in the Scandinavian countries.

Apart from the WLCG project, every LHC experiment has designed specific computing characteristics and operations. This has allowed using different solutions and comparing performances. In Chapter 3, we will go deeper in the ATLAS case by showing the Computing Model. To be ready for the LHC start-up, an important program of testing and challenges had been run for six years. When the initial prototype grid system was put in place with only a small number of grid sites in 2003, the system was started to use for the production of simulated data, and the next years an extensive program of data and service challenges was employed. It will be also commented for the ATLAS case in Chapter 3.

³Previous projects to EGI started with the implementation of Grid computing in Europe: DataGrid (2001-2004) [30] and “Enabling Grids for E-science” (EGEE, 2004-2010)[31]

Chapter 2

ATLAS Experiment

As mentioned before, ATLAS is a High Energy Physics general purpose detector in the LHC scientific program [35]. It has a big size: 25 metres of diameter, 46 m of length and 7000 tones of weight.

Its structure is in concentric layers around the beam tube as it is shown in Figure 2.1. From inside out, the center consists in track detectors using the silicon technology in a multi-layer structure. Next, there is a sampling electromagnetic calorimeter of liquid argon and lead with an accordion structure. The following layer is a hadronic calorimeter which is also sampling with steel plates and with scintillating tiles. Finally, the last cover consists in muon chambers surrounded by a toroidal magnetic field. This allows to measure the muon momentum. There is also a solenoid magnetic field for the rest of layers.

Thanks to these characteristics of the detectors, a very precise measurement resolution can be achieved in the electron and photon energy, in the charged particles momentum (specially for muons), and a reasonable measurement resolution for the jets energy. Also, there is an excellent track reconstruction close to the interaction zone to detect long-lived particles (the case of b and c quarks and muons).

Taking into account the ATLAS shape, the cylindrical coordinates are the most appropriate choice. Because of collisions are not from elementary particles, the mass center system does not coincide with the laboratory system, then, the invariant Lorentz coordinates chosen in the z axis (the beam direction) are the azimuthal angle ϕ , the transverse momentum and the rapidity which is defined as:

$$y = \frac{1}{2} \ln \left[\frac{E + p_z}{E - p_z} \right] \quad (2.1)$$




Figure 2.1: The ATLAS experiment schema showing its parts.

where E is the particle energy and p_z is the momentum component in the z axis. In the practice, the pseudorapidity is used (i.e., the rapidity in the relativity approximation):

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right] \quad (2.2)$$

where θ is the polar angle between the particle track and the z axis as it is shown in Figure 2.2. In next Sections, the ATLAS components will be described.




Figure 2.2: Representation of the polar angle θ in the xz coordinates used for the pseudorapidity definition.

2.1 Inner Detector

The ATLAS tracker, also known as the Inner Detector (ID) [36], performs the pattern recognition, momentum and vertex measurements together with electron identification, providing a pseudorapidity coverage up to $|\eta| < 2.5$. These capabilities are achieved with a combination of discrete high-resolution semiconductor pixel and strip detectors in the inner part of the tracking volume, respectively the Pixel and the SemiConductor Tracker (SCT), and a straw-tube tracking detector, the Transition Radiation Tracker (TRT), with the capability to generate and detect transition radiation that enables the electron-pion identification, in its outer part. The ID operates embedded in a 2 T axial magnetic field generated by a solenoid [37]. This magnetic field is used for bending the charged particles and measure their charge and momentum. The tracker schema is shown in Figure 2.3.




Figure 2.3: The Inner detector schema

Because of the experimental conditions at the LHC, around 1500 charged particles will cross the ATLAS ID every 25 ns at high luminosity ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$). The ID electronics and all the sensor elements must be fast enough and of course radiation hard. In addition, a very fine granularity is needed to handle the particle fluxes and to reduce the influence of overlapping events. For this purpose the ID has 5832 individual silicon modules (with about 86 million of readout channels).

The Pixel subdetector is based on silicon pixel technology and it is arranged in three cylindrical barrels and three discs on each side of the central barrel. The pixel elements are $50 \times 400 \mu\text{m}^2$ resulting in an intrinsic resolution of

10 μm in the transversal direction with a direct 2D readout. This system is designed to provide a very high granularity (with 80.4 million channels) as well as high precision set of measurements as close as possible to the interaction point. It consists in three barrels at average radii of 5.05 cm, 8.85 cm and 12.25 cm, and three discs on each side at 49.5 cm, 58.0 cm and 65.0 cm from center of the ATLAS coordinates system, i.e. $Z=0$.

The SCT is a silicon microstrip based detector which is located just after the pixel detector. The SCT modules are arranged on four barrel layers and nine end-cap discs on each side. It has been designed to provide eight precision measurements via 4 layers of back-to-back silicon microstrip detector modules with a relative 40 mrad stereo angle. There are five sensor topologies, one for the barrel which has parallel strips with 80 μm pitch and 4 for the end-caps with fan-out structure (54.53–90.34 μm pitch). With 80 μm strip pitch on average a SCT module ensures a 17 μm precision in $r\phi$ and its stereo angle of 40 mrad allows 580 μm in z . The SCT has 4088 modules (2112 barrel and 1976 end-cap modules) which means 61 m^2 of silicon sensors with 6.3 million channels.

Finally, the TRT consists of about 300.000 gaseous straw tubes arranged in a barrel and two end-caps on each side of this barrel. It has 176 modules, 73 layers in 3 rings in the barrel region and 2 160 straw planes in 40 four-plane assembly units in the end-cap regions. The TRT gas mixture Xe/CF₄/CO₂ (70%/20%/10%) provides an efficient X-ray absorption, a fast charge collection and a stable operation over a sufficient high-voltage range even at high particle rates. Its technology allows to have an intrinsic resolution of 130 μm per straw (i.e. in the direction perpendicular to the wire) where each straw tube has a diameter of 4 mm. From the practical design point of view it has been built to provide, in the barrel region and on average, 36 TRT hits for tracks coming from the interaction point. In Figure 2.4, a graphical representation of the ID layers is shown.




Figure 2.4: The Inner detector layers.

2.2 Calorimetry

Outside the ID solenoid are the calorimeters which perform energy measurements and particle identification. The Electromagnetic Calorimeter (ECAL) uses liquid argon (LAr) [38] as an ionization medium (it is also known as LAr calorimeter), with the lead absorbers arranged in an accordion geometry. This kind of geometry provides complete ϕ symmetry without azimuthal cracks and the lead thickness in the absorber plates is optimized as a function of η in terms of performance in energy resolution. It allows an excellent performance in terms of energy and position resolution as well as in the identification of electrons and photons providing coverage up to $|\eta| < 3.2$. It is surrounded by cryostat as it needs very low temperatures to operate.

Surrounding the latter is the Hadronic Calorimeter (HCAL) with a coverage up to $|\eta| < 4.9$ which measures hadronic jets. A sampling technique with plastic scintillator plates (called tiles) embedded in an iron absorber is used for the hadronic barrel tile calorimeter (also known as TileCal [39]). The TileCal is separated into a large barrel and two smaller extended barrel cylinders, one on either side of the central barrel. In the end-caps ($|\eta| < 1.6$), LAr technology is also used for the hadronic calorimeters, matching the outer $|\eta|$ limits of end-cap electromagnetic calorimeters. The LAr forward calorimeters provide both electromagnetic and hadronic energy measurements, and extend the pseudorapidity coverage to $|\eta| < 4.9$. The calorimeter schema is displayed in Figure 2.5. The fractional energy resolution of the electromag-

netic calorimeter was measured with electrons for various energies between 10 and 245 GeV and at various η values [35]:

$$\frac{\sigma_E}{E} = \frac{10\%}{\sqrt{E}} \oplus 0.7\% \text{ (GeV)} \quad (2.3)$$

and the fractional energy resolution for the Tile calorimeter using hadrons [35]:

$$\frac{\sigma_E}{E} = \frac{50\%}{\sqrt{E}} \oplus 3\% \text{ (GeV)} \quad (2.4)$$

The fractional jet transverse momentum resolution, which considers both calorimeters, has the expression:

$$\frac{\sigma_{p_T}}{p_T} = \frac{N}{p_T} \oplus \frac{S}{\sqrt{p_T}} \oplus C \text{ (} p_T \text{ in GeV)} \quad (2.5)$$

where N is the effective noise term (the noise contributions and the pile-up), S is the stochastic term (statistical fluctuations), and the C is the constant term (fluctuations of the momentum). The N term is significant at low transverse momentum below 30 GeV. The C term is expected to dominate in the high momentum region over 400 GeV. The S term is the real limiting factor of the resolution for this measurement that dominates in the intermediate momentum interval. This resolution was measured with data from pp collisions at center-of-mass of 7 TeV and with integrated luminosity of 35 pb^{-1} [40]. For transverse momentum between 50 and 500 GeV the fractional resolution drops from 15% to 8% in the central detector ($|y| < 0.8$). The Monte Carlo simulation are in agreement with data.




Figure 2.5: The ATLAS Calorimeter Schema.

2.3 Muon System

The outermost detector is the muon spectrometer [41] which defines the overall dimensions of the ATLAS detector. Its layout can be seen in Figure 2.6. It consists in four technologies which are the Monitored Drift Tubes (MDT), the Cathode Strip Chambers (CSC), the Resistive Plate Chambers (RPC) and the Thin Gap Chambers (TGC). The former two detectors provide high precision momentum measurements for muons, needed to perform the tracking. The latter two detectors are used for triggering with timing resolution of the order of 1.5–4 ns and bunch crossing identification. The muon spectrometer is designed to achieve a transverse momentum resolution of $\frac{\Delta p_T}{p_T} < 10^{-4}$ for $p_T > 300 \text{ GeV}/c$. At smaller momenta, the resolution is limited to a few per cent by Multiple Coulomb Scattering (MCS) effects in the magnet and detector structures, and by energy loss fluctuations in the calorimeters.




Figure 2.6: The ATLAS Muon Chambers Schema.

An air-core toroid system [37] generates strong bending power in a large volume within a light and open structure. MCS effects are thereby minimised, and excellent muon momentum resolution is achieved with three layers of MDT chambers achieving a precision of approximately $50 \mu\text{m}$ in muon position measurements. This magnetic system as can be seen in Figure 2.7 has a barrel (25 m long, with an inner bore of 9.4 m and an outer diameter of 20.1 m) and two inserted end-cap magnets (with a length of 5.0 m, an inner bore of 1.65 m and an outer diameter of 10.7 m). The barrel toroid consists of eight flat coils assembled radially and symmetrically around the beam axis

and the magnetic field provides for typical bending powers of 3 Tm in the barrel and 6 Tm the end-caps. The end-cap toroid coils are rotated in azimuth by an angle of 22.5 degrees with respect to the barrel toroid coils to provide radial overlap, and to optimize the bending power in the transition region.




Figure 2.7: ATLAS Toroid Magnet - Barrel.

2.4 Trigger System

The LHC proton bunches will collide at a frequency of 40 MHz, i.e. at 25 ns and at the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ on average about 23 inelastic proton-proton collisions will be produced at each bunch crossing. Therefore, the trigger system needs to efficiently reject a large rate of background events and still select potentially interesting ones with high efficiency. To deal with this amount of data that these collisions will generate the ATLAS trigger is based on three levels of online event selection which is shown in Figure 2.8.

Each trigger level refines the event selection done by the previous level, applying new criteria. The level-1 trigger (LVL1) [42], which is hardware-based (i.e. it is implemented in custom electronics), is responsible for the first level of event selection, reducing the initial event rate to less than 75 kHz (limited by the bandwidth of the readout system, which is upgradeable to 100 kHz). It uses information from the calorimeters and muon trigger chambers to make a decision on whether or not to continue processing an event in about $2.5 \mu\text{s}$. The subsequent two levels are software-based and are collectively known

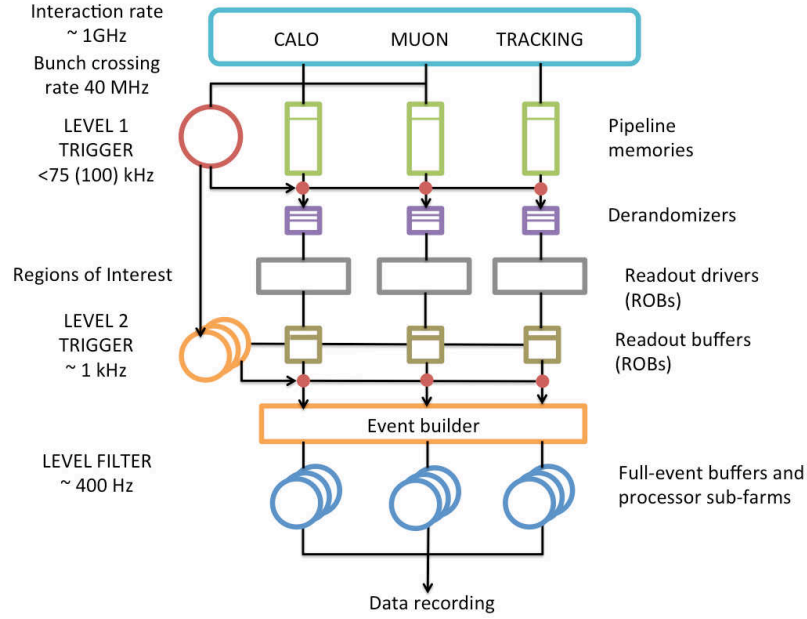


Figure 2.8: ATLAS Trigger Diagram.

as the High Level Trigger (HLT) [43]. One hand, the level-2 trigger (LVL2) decides in $O(10)$ ms if the event should be rejected making use of the Regions of Interest (RoI) provided by the LVL1 data with full granularity from all detectors. On the other hand, the Event Filter (EF) uses offline algorithms to perform its refine selection in $O(1)$ s. The HLT provides the reduction to a final data-taking rate of approximately 400 Hz where each selected event is estimated to have a total size of approximately 0.8 MB. This reduction is possible because the HLT uses seeded, step-wise and fast selection algorithms based on the reconstruction of potentially interesting physical objects like electrons, muons, jets, etc and which can provide the earliest possible rejection of background events.

In conclusion, ATLAS is a big experiment designed to discover new Physics. This task has been achieved thanks to the high granularity in the inner detector, a better energy precision in the calorimeters which allows an improvement is the distinction between electrons and photons. Also, a high precision in the jet energy is acquired in the hadronic calorimeter. Due to the muon chambers and the magnetic fields, a better position of muons is obtained. The signals are filtered by a 3 level trigger system to focus on the interesting events. After that, the next step is the storage and processing of data using the Grid which is described in the following chapter for the ATLAS case.

Chapter 3

ATLAS Software and Computing Model

The ATLAS experiment produces around 3.2 petabytes per year in real data. This data must be stored, processed and accessible for about 3000 ATLAS physicists spread around the world. As mentioned before in Chapter 1, to face up this challenge, Grid computing technology was chosen and the WLCG project was designed. In this chapter, the particular case of ATLAS is presented in the implementation of the WLCG guidelines and the changes thanks to the flexibly computing architecture. Next, the description of part of the ATLAS software: the Athena framework and Database. We devote the last section to the summary of the tests before the data-taking.

3.1 ATLAS Computing Model

According to WLCG, different types of computing centres have been defined in Tiers [44]:

- Tier-0 centre at CERN which records the raw data emerging from the data acquisition systems after the trigger filter. In addition, it executes the first-pass reconstruction, where a copy of the reconstructed data will be stored. It distributes a second copy of the raw data after the first processed across the Tier-1 centres associated with the experiment.

To cope with the increasing requirements for LHC computing, the Wigner Research Centre for Physics in Budapest, Hungary, will operate as an extension to the CERN Tier-0. That centre will act as a remote Tier-0, hosting CERN equipment to extend the Grid's capabilities. This will guaranty the well-working of the task in case of CERN problems [45].

- The Tier-1 centres have the prime responsibility for managing the permanent data storage (raw, simulated and processed data) and providing computational capacity for reprocessing and for analysis processes that require the access to large amounts of data. At present 10 Tier-1 centres are defined in ATLAS. Each Tier-1 has some Tier-2 centres associated to it. The group formed by these associated sites is called Cloud¹. For instance, there is a Tier-1 centre in Barcelona, Spain, which is called Port d'Informació Científica (PIC) [46] [47].
- The role of the Tier-2 centres is to provide computational capacity and appropriate storage services for Monte Carlo event simulation and for end-user analysis. The Tier-2 centres get data as required from Tier-1 centres, and the data generated are sent to Tier-1 centres for permanent storage. Also the Tier-2s can give capacity for calibrations. Around 80 centres have been identified in ATLAS. It can be defined a federated Tier-2 which consists of several locations or sites. One example of federation is the ATLAS Spanish Tier-2 which we see in Chapter 5.
- Other computing facilities in universities and laboratories take part in the processing and analysis of LHC data known as Tier-3 infrastructures [48]. These are not defined by the ATLAS group requirements, some are included in a Tier-2 centre and some are stand alone, although they must be provided with access to the data and analysis facilities. They are used for local needs, however, some of them contribute with CPU cycles for simulation and analysis when is possible. An example of Tier-3 is the one located at IFIC (Valencia, Spain) which has been used for the analysis presented in Chapter 4.

CERN also has a centre (size similar to a Tier-2) for analysis called CERN Analysis Facility (CAF). It is specially destined for calibration, alignment and algorithm developing for its direct access to RAW data from Tier-0.

This computing model has evolved for these years. At the beginning, the MONARC² [44] model has been used which consider the best connection is among the Tier-1 and associated Tier-2s centres. The Figure 3.1 represents this hierarchical Computing Model. According to this, a transfer between two Tier-2s from different Clouds has been done crossing the two Tier-1s. However, it was observed with transfer tests called 'sonar' [49] during the

¹Different term is Cloud Computing, another performance to deliver hosted services over the Internet

²This is the abbreviation of "Models of Networked Analysis at Regional Centres for LHC experiments".

LHC data-taking that some Tier-2 centres have good connection of a direct transfer.

According to the sonar test, the interconnection of some Tier-2s to other sites outside of their Clouds has been decided, specially to other Tier-1s centres. These Tier-2s are called Tier-2 Direct (T2D) [50] and they have to fulfil the requirements of a good connectivity and give a certain level of commitment and reliability. The validation criteria to be a T2D consists of the transfers of big files to or from Tier-1s should move at 5 MB/s during the last week and 3 out of the 5 last weeks. Currently, the interconnections are directly to any Tier and there is not Cloud distinctions. In fact, the whole infrastructure will behave like one Cloud, this is the so-called Mesh model. This model is represented in Figure 3.2.

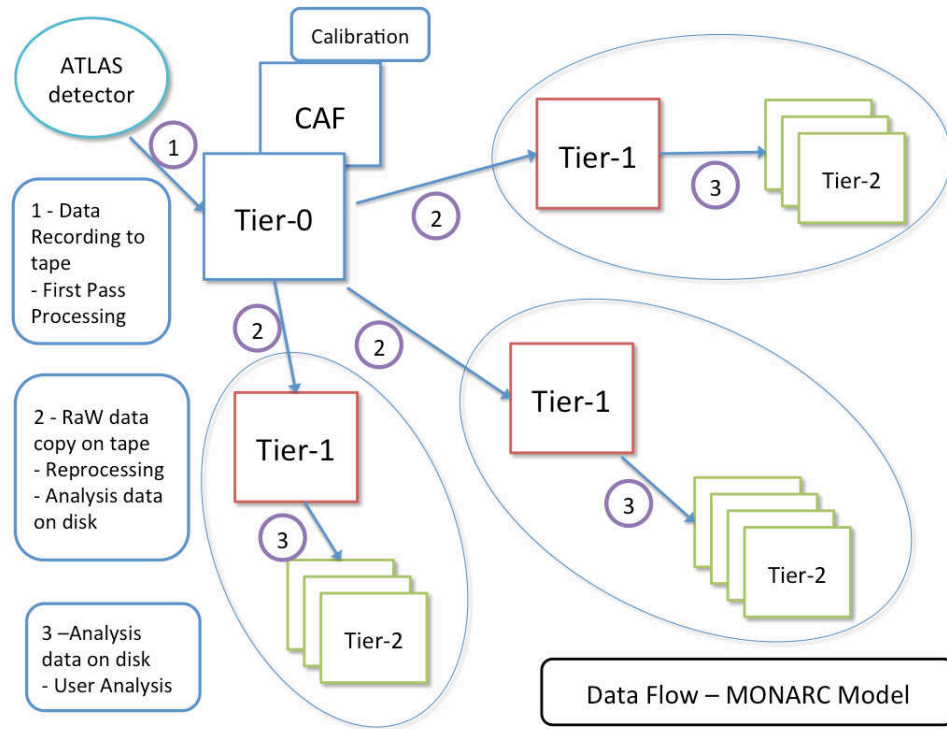


Figure 3.1: The ATLAS Data Flow keeping the hierarchical Computing Model.

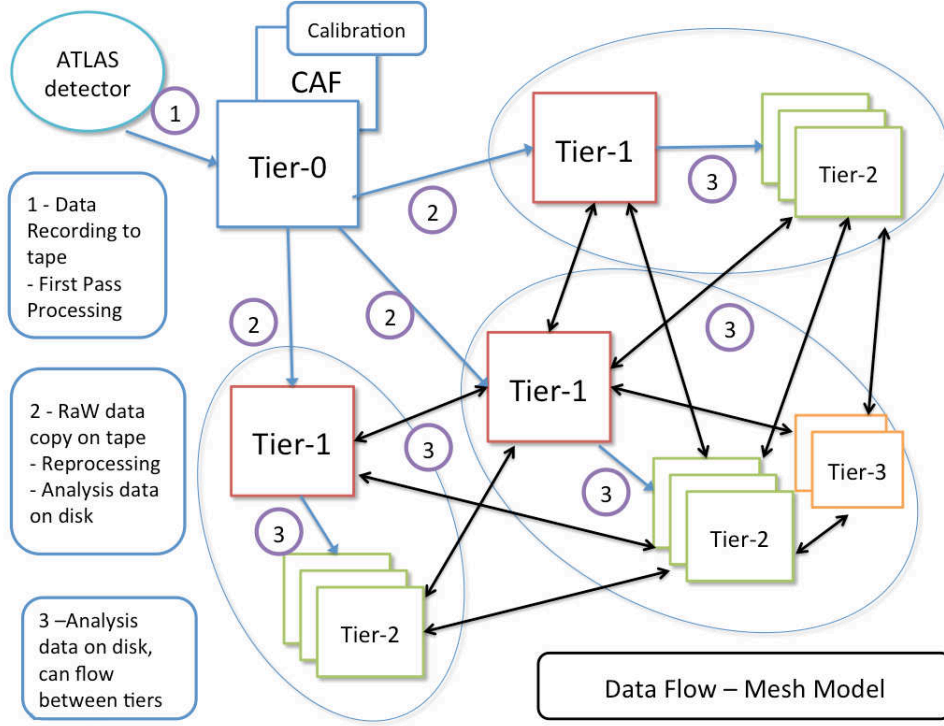


Figure 3.2: The ATLAS Data Flow with the Tier-2 centres connected with Tiers from other Clouds in the Mesh computing model.

3.2 Event Data Model

The Event Data Model in ATLAS defines a number of different formats [44] for the data from the detector and for the simulation produced that is described in the section 4.3. A schema is shown in the Figure 3.3. Firstly, the RAW is in “ByteStream” format and the size around 0.8 MB/event.

The Event Summary Data (ESD) is the full output of reconstruction in POOL/ROOT object format [44]. It includes tracks (and their hits), calorimeter clusters, calorimeter cells and combined reconstruction objects. The nominal size is 2.5 MB/event initially, to decrease as the understanding of the detector improves.

The Analysis Object Data (AOD) is the summary of event reconstruction with Physics objects, for example, electrons, muons and jets. The nominal size is 350 kB/event, however it can be the double of that.

The Derived Physics Data (DPD) is skimmed, slimmed and thinned events with other useful user’s data derived from AODs and conditions data if it is

3.2. Event Data Model

mainly skimmed ESD. The nominal size is 10 kB/event on average but there are large variations depending on Physics channels.

Finally, the TAG is a database or ROOT files used to quickly select events according to some specific criteria. The TAG is producible from AOD, though TAG databases contain or are linked to sufficient navigational information to allow retrieval of event data at all production processing stages, i.e., AOD, ESD, Raw Data (RDO).

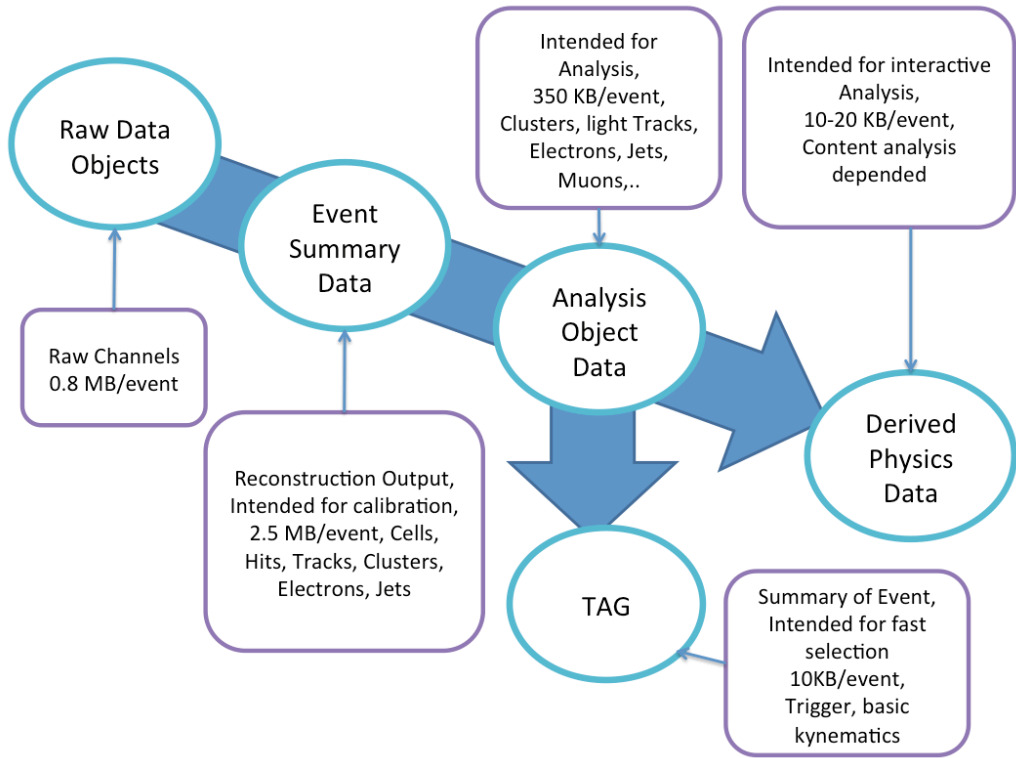


Figure 3.3: The Event Data Model defined for ATLAS

The use of the different data types has been changing during these last years according to the needs of the analysis. Before the LHC starting-up, simulated and cosmic radiation data had been analysed by physicists. The big part of the analysis processes has used the AOD format as input files to create ntuples with the histograms. It has been seen that format is too heavy for an analysis in a local machine, so DPD format was proposed to cope with this. The first definition of DPD was a ntuple customised by the end-user with only the variables needed for the local analysis. Then, official DPDs were also decided to produce from ESDs and AODs called dESD and dAODs.

When the data-taking started, the majority of the input files was ESDs for few months due to the first studies in calibration and alignment. The ATLAS policy related to the trigger frequency changed and the data rate has been increased more than it was expected. Because of this, RAW data required more size to store. Then it was decided that the ESD format should be stored in a short period of time, meanwhile, the calibration and the alignment are being executed and after that, the ESDs should be deleted to allow more space for RAW data in Tier-1 centres.

AODs were still popular, however the ATLAS physics groups started to define and to generate specifically DPDs. Currently, the DPD is the data type most used for analysis as input, specially the so-called D3PDs which are produced officially by ATLAS [50]³. The content of these D3PDs are defined by the Physics groups. The D3PDs require more space than a laptop hard disk (the necessary D3PDs can be rise 1 TB total) and this fact does not allow the original DPD idea of an interactive execution. Therefore, ntuples are created from D3PDs to do the last steps of the analysis locally.

3.3 ATLAS Software

ATLAS software has been developed for operational activities as trigger, calibration, alignment and simulation. It is also for all the variety of analysis of the different Physics Groups. During the design, it has taken into account the complexity and the scale of ATLAS, and its continuous updating. In this section, Athena, the principal software program is shown as well as Database.

3.3.1 Athena Framework

Athena is a control framework and it is a concrete implementation of an underlying architecture called Gaudi [51]. Gaudi is a kernel of software common to ATLAS and LHCb experiments. On the basis of this, Athena is the sum of this kernel plus ATLAS-specific enhancements. It is written in C++, and it is designed with a modular component architecture, consisting of core packages, such as the kernel, services and various tools, and supplemented by external libraries. The major components are:

- Application Manager: It is the overall driving intelligence that manages and coordinates the activity of all other components within the

³the number '3' is due to the old DPD process which generated 3 types of DPDs (basic, specific information and user custom). D1PD and D2PD have been discarded.

application. There is one instance of the application manager and it is common to all applications.

- **Algorithms and Sequencers:** They are algorithms share a common interface and provide the basic per-event processing capability of the framework. A Sequencer is a sequence of Algorithms, each of which might itself be another Sequencer, allowing for a tree structure of processing elements. A filter Algorithm is a event selection criteria. It can indicate that the event being processed fails to meet its filter criteria and inhibit the processing of downstream Algorithms.
- **Tools:** They are similar to an Algorithm but differs in that it can be executed multiple times per event.
- **Transient Data Stores:** They are the data objects accessed by Algorithms are organized in various transient data stores depending on their characteristics and lifetimes. The event data itself is managed by one store instance, detector conditions data, such as the geometry and alignment, by another store, etc.
- **Services:** A Service provides services needed by the Algorithms. In general these are high-level, designed to support the needs of the physicist. One example can be the random-number generators.
- **Selectors:** These components perform selection. For example, the Event Selector provides functionality to the Application Manager for selecting the input events that the application will process. Other types of selectors permit the selection of objects within the transient data stores.
- **Converters:** These are responsible for converting data objects from one representation to another. They implements the file input/output of the objects.
- **Properties:** All components of the architecture can have adjustable properties that modify the operation of the component. Typically these are basic types (integer, float, etc.), but can also be specified as having upper and lower bounds.
- **Utilities:** These are C++ classes that provide general support for other Athena components.

3.3.2 Database

The Athena architecture for persistent data access is independent of the implementation technology. As a result, without any changes in the sub-systems software code, Athena can access database-resident data through various technologies (Oracle [52], SQLite [53], Frontier/Squid [54]) by using the common LCG-AA software packages COOL/CORAL and a set of configuration switches [55].

ATLAS stores in the offline Oracle databases a wealth of information about the detector geometry and status, data-taking conditions, calibrations, alignments, data quality, luminosity, data management information and so on [56]. In addition, short event index records, the TAGs, can be stored in Oracle databases as well as in POOL/ROOT files, but their access pattern and related tools are completely different from the rest of database-resident information.

The master Oracle databases used for all offline computing purposes are placed at CERN. Part of the information (Conditions DB and Trigger DB) is replicated in realtime to Oracle database servers placed at 4 Tier-1 computing centres. In addition to the constants stored in the COOL Conditions Database, some subsystems need additional amounts of information that is stored in POOL files; these files are referenced from COOL.

3.4 Computing testing and challenges before starting-up of LHC

A great effort was made to implement the Grid computing and all the Tier centres for more than 5 years. Several tests coordinated by the WLCG project were executed. The goal was to get ready the Grid infrastructure before the data-taking of November 2009. These helped to achieve the well-working of all the sites after that [57]:

Data Challenges (DC) was started in 2002, with the goal of validating the Computing Model, the complete software suite, and the data model, and to ensure the correctness of the technical choices to be made for the final offline computing environment.

- Data Challenge 1 (DC1) (2002-2003) [58]: was run on conventional non-Grid infrastructure (PC clusters and farms) in Europe, Asia and North America and fully on Grid in the Nordic sites. It was first exercise on

worldwide scale. The goal was to set up the production infrastructure in a real worldwide collaborative effort to gain experience. It proved the necessity of the interactions between various different groups, for example groups involved in ATLAS Computing, Grid middleware developers, and CERN IT.

- Data Challenge 2 (DC2) (May-December 2004)[59]: It was deployed a new Production System able to submit jobs on three different Grids: LCG, OSG and NorduGrid. The three Grids in use in ATLAS have proven to be usable for large-scale productions despite of the short manpower. All involved components both on the Grid and the ATLAS side needed substantial improvements.

In 2004, the WLCG project proposed a series of Service Challenges (SC) that were intended to test various aspects of the grid service. Initially, these challenges focused on the data transfers. They expanded to the testing of data management, scaling of job burden, and key support processes including response to simulated security incidents. Finally basic services were tested and it was clear the necessary improvements in transfers [60]:

- Services Challenges (SC): SC1 (December 2004), SC2 (March 2005), SC3 (July 2005) These tests have consisted in checking with special emphasis on Data Management. The goals were largely exceeded for the Tier-2 sites in service reliability and sustained transfer rates.
- SC4 (June 2006) Offline data processing requirements can be handled by the Grid to the nominal LHC data rate. There were a large participation of Tier-2 sites and with all the Tier-1s sites. The transfer rates (disk-tape) planned were achieved and in some cases exceeded.

In addition to these organized tests, between 2008 and 2009, when the LHC restarted operations, significant amounts of cosmic-ray data were acquired and processed by the experiments. These data allowed full testing in real conditions of the experiments entire computing models, from data acquisition to analysis, and consequently validated the entire system. Often, by increasing the data-acquisition rates this testing could be performed at rates close to that anticipated for LHC running. The testing also allowed preliminary calibrations and alignment of the detectors in preparation for the LHC start-up. The testing program culminated in two major readiness challenges—one in 2008 (CCRC08) and one in 2009 (STEP09)—that demonstrated WLCGs preparation for the start of data taking:




Figure 3.4: Data rates achieved during STEP'09

- Common Computing Readiness Challenge (CCRC08) (March and June 2008) It was made the measurement of the readiness of the Grid services and operations before the real data taking expected that year. All the experiments simultaneously have been stressing the WLCG infrastructure in close to real conditions. Experiments have been running their Full Dress Rehearsals and scheduling key periods together with the CCRC08 challenge. At the end, the LHC data taking was planned for 2009. Therefore, another important test was done before it.
- Scale Testing for the Experiment Program 09 (STEP09) (May 2009) It was a stress and scale testing of all experiment workloads including tape recall and massive end user analysis. Concretely in ATLAS, the simulated data production was tested lasting millions of hours to create 12 millions of events and worked fine during the process. In the ATLAS data distribution, 4 PB were moved successfully in spite of full disks. In Figure 3.4, the data transferred from STEP09 to first data-taking is displayed [57].

3.5 Conclusions

ATLAS has followed the Tiers classification proposed by the WLCG project. Its computing model has been designed as interconnected Tier centres, which has been changing from hierarchical (MONARC) to flexible (Mesh) model. Thanks to these changes, the system performance has been improving according to the new experiment needs during the data-taking. The final objective is to increase the interconnections among Tier centres until all together became to one Cloud.

3.5. *Conclusions*

Data model was presented: the main data types are Raw data, ESD, AOD, DPD and TAG. A brief description of the data input type popularity was provided. Currently D3PD has displaced the AODs, with its own official ATLAS production.

The ATLAS software, Athena framework, was shown. It is using for important activities such as Monte Carlo Production. The Database is connected with Athena and giving the geometry information of the experiment and the data-taken conditions.

Finally, WLCG tests before data-taking have been enumerated. Thanks to these challenges exercises, Grid infrastructures and operations have been checked successfully and ATLAS data were available for the Physics Analysis, in particular for the analysis which allowed the Higgs-like particle discovery [14].

Chapter 4

ATLAS Distributed Analysis System

The LHC experiments need to analyse millions of events in weeks as it has been mentioned in Chapter 3. ATLAS has developed a system to perform the analysis in a distributed way using the Grid infrastructure as the one deployed by the Collaboration. This chapter is devoted to describe the Production and Distributed Analysis system including the components, the functionality and the tools to carry out these tasks. Special attention is given to the Monte Carlo Production workflow and the evolution of the Distributed Analysis system. The work of the Distributed Analysis Support Team is also discussed in this chapter. To finish, in the last section, the concrete case of the use of the Distributed Analysis for the Physics study reported in Chapter 6 is provided.

4.1 Production and Distributed Analysis system

ATLAS has a specific job management called Production and Distributed Analysis system (PanDA) [61]. It has been developed since summer 2005 to meet ATLAS requirements for production and Distributed Analysis operating at LHC data scale. ATLAS processing and analysis require challenges on throughput, dimensionality, robustness, efficient resource utilization, minimal operations manpower, and integration of data management. PanDA was initially created only for United States use in ATLAS activities. Since September 2006, PanDA has also been a principal component of the US Open Science Grid (OSG) [33] program in just-in-time [62] (pilot-based) workload management. In October 2007, PanDA was adopted by the ATLAS Collabo-

ration as the unique system for distributed processing production across the Collaboration. At present, PanDA unifies the different Grid environments (EGI [32], OSG [33] and NorduGrid [34]) in ATLAS sites for production and analysis. One of its characteristic is the useful monitoring web page which can be consulted by the users and experts. The PanDA components are the following:

- PanDA Server: central PanDA hub composed of several components that make up the core of PanDA. Implemented as a stateless REST web service over Apache mod-python and with a MySQL backend:
 - Task Buffer: the PanDA job queue manager, keeps track of all active jobs in the system.
 - Brokerage: matches job attributes with site and pilot attributes. Manages the dispatch of input data to processing sites, and implements PanDA's data pre-placement requirement.
 - Job Dispatcher: it receives requests for jobs from pilots and dispatches job payloads. It compares the capabilities of the site and worker node (data availability, disk space, memory etc.) with the jobs requirements. Jobs are assigned to sites according to this matching.
 - Data Service: data management services required by the PanDA server for dataset management, data dispatch to and retrieval from sites, etc. Implemented with the ATLAS Distributed Data Management (DDM) system (for the dataset and DDM definition see the subsection 4.2.1).
- PanDA DB: Oracle database for PanDA.
- AutoPilot: Pilot submission, management and monitoring system. Supersedes first generation PanDA JobScheduler.
- Pilot: the lightweight execution environment for PanDA jobs. Pilots request and receive job payloads from the dispatcher, perform setup and cleanup work surrounding the job, and run the jobs themselves, regularly reporting status to PanDA during execution. Pilot development history is maintained in the pilot blog.
- PanDA Schedconfig: Database table to configure resources, used by PanDA server and AutoPilot.

4.1. Production and Distributed Analysis system

- PanDA Monitoring: web based monitoring and browsing system that provides an interface to PanDA for operators and users. Currently, it is migrated to an updated platform.
- PanDA Logger: logging system allowing PanDA components to log incidents in a database via the Python logging module and HTTP Bamboo interface between PanDA and the ATLAS production database. Supersedes the PanDA Executor Interface.

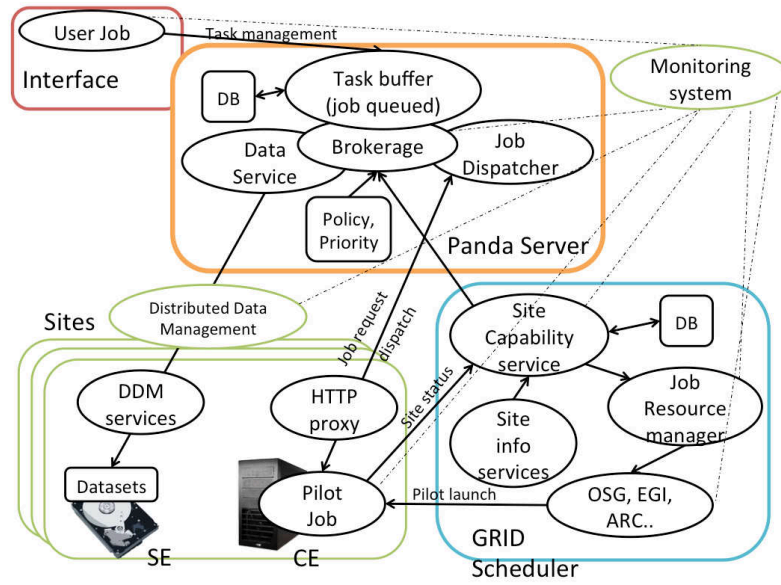


Figure 4.1: A schema of the PanDA architecture reflecting some of its components and its workflow.

The PanDA jobs workflow is described in Figure 4.1. Jobs are submitted to PanDA in a interface where production group or analysis users define job sets, their associated datasets and the input/output files within them. Job specifications are transmitted to the PanDA server via secure HTTP protocol (authenticated via a Grid certificate proxy). The PanDA server is the main component which provides a task queue and manages all job information centrally. The PanDA server receives jobs into the task buffer. At the same time a brokerage module operates to prioritize and to assign work on the basis of job type, priority, software availability, input data and its location, and available CPU resources. For production jobs, the allocation of job sets to sites is followed by the dispatch of corresponding input datasets to those sites, handled by a data service interacting with the Distributed Data Management system. Data pre-placement is a strict precondition for

job execution: jobs are not released for processing until their input data arrives at the processing site. When data dispatch completes, jobs are made available to a job dispatcher (more in deep in section 4.3). In analysis jobs, Distributed Data Management is consulted to find the input datasets and jobs are assigned to sites where data are.

In the next step, an independent subsystem manages the delivery of pilot jobs to worker nodes via a number of scheduling Grid systems. The pilot job once launched on a worker node contacts the job dispatcher and receives an available job appropriate to the site.

Pilots retrieve jobs from the PanDA server in order to run the jobs as soon as CPUs are ready. Pilots use resources efficiently; they exit immediately if no job is available and the submission rate is regulated according to workload. Each pilot executes a job on a worker node, detects zombie processes (they are processes that have completed execution but still have entries in the process table [63]), reports job status to the PanDA server, and recovers failed jobs. The pilot isolates the PanDA system proper from Grid heterogeneities, which are encapsulated in the pilot, so that at the PanDA level the Grid seems homogeneous.

The PanDA monitoring receives job information during the workflow from different PanDA and Grid components and ATLAS users can check the job status.

4.2 Distributed Analysis End–User Tools

The ATLAS Computing Team has developed specific tools for users to use the Grid infrastructure and to easy perform the Distributed Analysis tasks. These tools have been designed to get data information from the experiment, to submit the analysis job to the ATLAS Grid infrastructure and to retrieve the output [64]. They are user-friendly which allow users to focus more on their Physics analysis than understanding the complexity of the Distributed Analysis on Grid.

4.2.1 Tools to access the Data

The Distributed Data Management (DDM) [65] group is involved in developing a system to manage the access to ATLAS data that is distributed at sites around the world. This system is called Don Quixote 2 (DQ2) [65]. It consists of a bookkeeping system and a set of local site services to handle data

transfers, building upon Grid technologies. The data is grouped in datasets which are defined as versioned collection of Grid files, logical entity with a unique DSN (Dataset Name). They can be replicated (create an exact copy) to another Grid site. Every ATLAS dataset is catalogued and stored in the ATLAS Tier centres independently of the Grid type, this allows a easy data management. Users apply DQ2 to acquire data information about dataset name, number of files and sites where dataset is stored.

Another tool to get data information is the ATLAS Metadata Interface (AMI) which has all catalogued the metadata information of official ATLAS datasets. It can provide dataset name, number of files, number of events, the Athena version in the reprocessing, availability and generation information if data are Monte Carlo simulation (i.e., the cross section of the simulated channel). Powerful web interfaces are available, and a web service is provided [66].

Also, users can register their result files using only DQ2. That is the reason for what it is the most common tool for getting the output result files. However, the output dataset is stored in local devices and it would be full. For large-sized output files, one option is Data Transfer Request (DaTRI) [67]: a tool which allows ATLAS users can make a request for copying a dataset (usually their own output results) in another site under restrictions of the nationality, Physics group or activity (this point clarified in the space token section in Chapter 5).

DDM needs always to evolve to satisfy new user requirements and to integrate new technologies. The current DQ2 system has the issue to not be compatible with some of these requirements. Apart from this, its old architecture has limits as the datasets versioning. This implies pointer references between datasets and their files which complicates the database. Although DQ2 works, the operational burden is heavy. Therefore a new DDM system is being implemented, the so-called Rucio system [68].

Some improvements provided by this system are the definition of ‘accounts’ to assign them permissions and space quotes. Accounts can refer to users, groups or ATLAS activities. Files, datasets and containers have the same identifier system instead of these three types in DQ2. Also, Rucio allows to group containers into one container. In addition, the replicas control is improved and more metadata information is added and data searches can be according to it. Finally, the file address inside the SE (called Physical File Name) is unique and has additional protection to avoid a non-uniform distri-

bution of files to directories and access problems because of large directories names. All the system improvements are transparent for the user.

4.2.2 Tools to send analysis jobs to the Grid

Some tools have been developed for the easily use of the Grid for Distributed Analysis in ATLAS: PanDA-Client [69] and Ganga (Gaudi/Athena and Grid Alliance) [70]. Every user can choose one of them for sending the analysis code. Each tool has almost the same options and they access to all the Grid sites independently of their diverse configurations. Then, the differences in the structure and performance of the tools are not realized by users. The Figure 4.2 shows the access of the tools to the backends. Despite of both tools are similar, PanDA-Client is the most common tool used. However, the study of this thesis has been done using Ganga, thus this tool is described in more detail in this section.

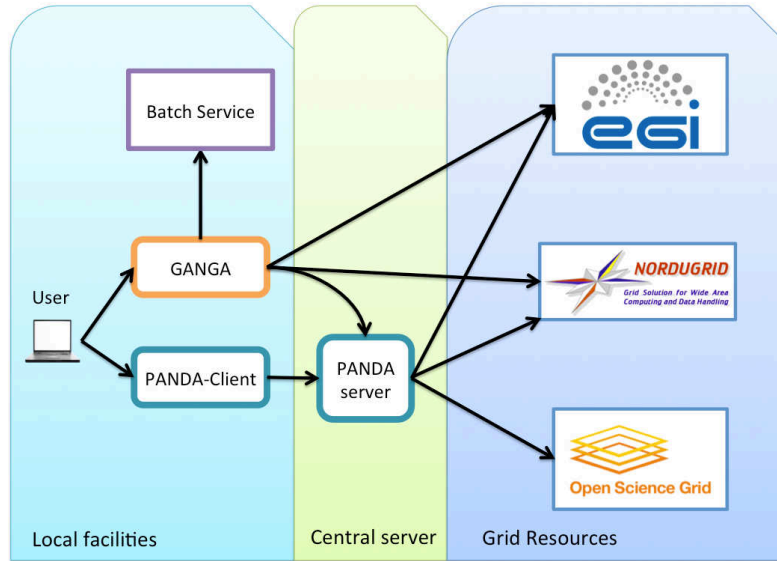


Figure 4.2: Relation among the users, Ganga and PanDA-Client and the different Grid flavours that can act as a backend of the job submission.

PanDA-Client

The same PanDA developers have elaborated the Grid Tool to submit jobs which is called PanDA-Client [69]. At the moment, it is the most popular Grid tool because of easy-to-use handling and its direct relation with PanDA

system. There are different programs to apply the several ATLAS software or Grid activities: The first is ‘pathena’ and lets user to submit Athena jobs. For the rest of applications like ROOT or simple Shell code the program is ‘prun’. ‘psequencer’ manage to perform sequential jobs or operations in PanDA based analysis. ‘pbook’ is the application for bookkeeping the PanDA analysis jobs and ‘puserinfo’ is the access control on PanDA analysis jobs.

Gaudi/Athena and Grid Alliance (Ganga)

Ganga is a user-friendly job definition and job management tool that allows simple switching between testing on a local batch system and large-scale data processing on distributed resources (Grid), implemented in Python [70]. It was developed for ATLAS and LHCb experiments to get a Grid user interface, and includes integrated support for configuring and running applications based on the Gaudi/Athena framework common to the two experiments, and, in the case of ATLAS, for Production System JobTransforms, and for DQ2 Data Management system.

A job in Ganga is constructed from a set of building blocks which is shown in the Figure 4.3. All jobs must specify the software to be run (application) and the processing system (backend) to be used. Many jobs will specify an input dataset to be read and/or an output dataset to be produced. Optionally, a job may also define functions (splitters and mergers) to divide a job into subjobs that can be processed in parallel, and to combine the resultant outputs. Ganga provides a framework for handling different types of application, backend, dataset, splitter and merger, implemented as plugin classes. Each of these has its own schema, which places in evidence the configurable properties. Because Ganga is based on a plugin system, it is readily extended and customised to meet the needs of different user communities. For instance, one of the activities where Ganga is used outside of ATLAS and LHCb is in simulation studies for medical Physics [71].

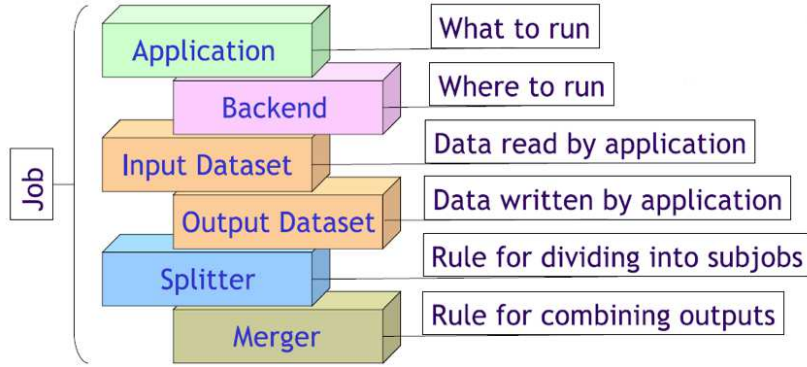


Figure 4.3: The Job Schema in Ganga showing the definition blocks [70].

4.3 Monte Carlo Production for ATLAS

In High Energy Physics, a simulation of particle events in the detector is necessary to foresee their behaviour and to design how to approach to the new physics with the comparison of the well-known physics background. This is a elaborated process that requires lots of calculations and computing size and is managed using the ATLAS Grid infrastructure. The simulation process and its workflow [44] inside the ATLAS Tier centres will be shown in the next subsection.




Figure 4.4: The full chain schema of the Monte Carlo Production process. The final chain after digitization is similar to real data processing.

4.3.1 Simulation Process

The target is to get Raw Data like real produced data in the detector, adding the Monte Carlo Truth information. Some steps are followed as in the Figure 4.4:

- **Generation:** is the production of particle four vectors from specified physics processes. Event generators are the tools for the modelling of the complex physics processes and can set detector requirements,

formulate analysis strategies, or calculate acceptance corrections. Generators can design the physics of hard processes, initial- and final-state radiation, multiple interactions and beam remnants, hadronization and decays. They run inside Athena and their output goes to a common format.

- **Simulation:** is to simulate the events moving through the detector. Geant4 (G4) [72] toolkit (inside Athena also) is used for this task which gives optimized solutions for geometry description and navigation through the geometry, the propagation of particles through detectors, the description of materials, the modelling of physics processes visualisation, etc.
- **Digitization:** this step translates the output of simulation process in Raw Data Object (RDO) that should be similar as the real detector data. It is of key importance that digitization is tuned by comparing the RDO output to real data in system tests to produce a realistic tuning of the detector response. The simulation of the pile-up, which is the overlaying of signal and background events, is processed optionally in this step.
- **Reconstruction:** is to derive from the stored raw data the relatively few particle parameters and auxiliary information necessary for physics analysis. The outputs are in that case ESDs and AODs that will be reprocessed to obtain DPDs.

4.3.2 ATLAS Monte Carlo Production workflow on PanDA

The simulation production workflow in PanDA is described in this subsection. It started when the physics groups define the jobs that have to be run on the Grid. These jobs are located in the central production database (ProdDB) and the executor (Bamboo) drives this information to the production engine (PanDA). The jobs are assigned to the clouds and every time the job status is updated in DB. Production jobs go to the PanDA server and the submission information returns to the client. The brokering module assigns work on the basis of job type, priority, input data and its location, available CPU resources and other brokerage criteria. The jobs arrive to the site where the input files are located. The DDM system stores the produced data at different sites in the PRODDISK space token (place dedicated to production activities, see Chapter 5) and registers them into the defined catalogues. In the Figure

4.5, the production jobs are a constant flow of algorithms previously tested and they allow to check the status of all sites. The ATLAS Distributed Computing (ADC) group designed shifts followed by ATLAS members to monitor the Monte Carlo production and alert site administrators about any problem sending a ticket to the Grid Global User Support (GGUS) [73] which is the main support access point to the Grid sites.

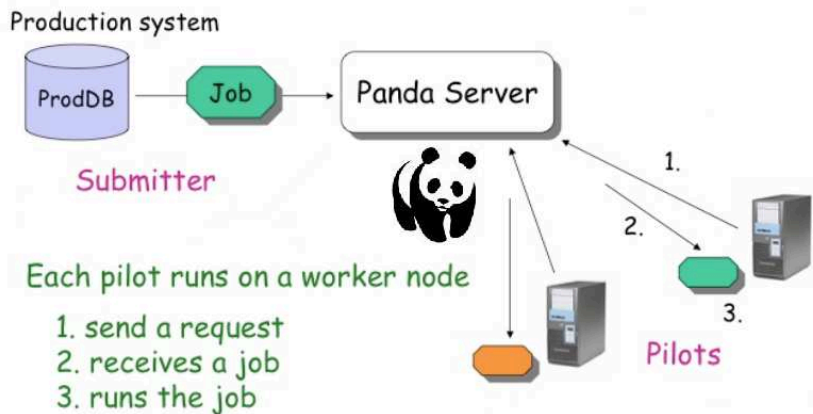


Figure 4.5: The Production workflow in the ATLAS Grid infrastructure [74]

The distribution of the production jobs is a clear example of the ATLAS computing evolution commented in Chapter 3. The MONARC model has been used favourably at the beginning of the WLCG project. However complications appeared during the data simulation production. Only one cloud was assigned to do a production task and sometimes more CPU was required meanwhile resources stayed unused in other clouds, thus they were not well-spent. With the ATLAS Computing Model changes, one Tier-1 can work with other Tier-2s from different Clouds to complete the production tasks. This involves these tasks are finished quicker for having more CPUs and there are more space to store the production outputs [54].

4.4 ATLAS Distributed Analysis Progress

The ATLAS Computing team has been improving the Distributed Analysis system since the LHC start-up. One of the most important improvements was in the HammerCloud (HC) [75]. The HC consists in an automatic analysis jobs tests system. HC sends jobs from analysis examples every day and tests the well-working of the sites. One of the tasks incorporated to HC

was the site exclusion. If HC shows problems after several tries, the sites will be excluded for analysis automatically and no more user jobs will go to them. Site administrators are alerted by e-mail with a summary of the errors presented to be able to fix them. After the site administrators provide a solution, HC continues doing tests until a maximum of jobs finish successfully. Because of the automatic site exclusion of the HC, failed jobs have been decreased and this saves time to the end-users analysis. Recently, HC is used also for simulation production.

Another aspect which has improved the Distributed Analysis is the so-called PanDA Dynamic Data Placement (PD2P) [76]. It was developed by the PanDA team. PD2P consists in doing data copies according to the popularity of use in the analysis jobs in an automatic way by PanDA. Apart from optimising the resources use because only the most required data are replicated, it allows other sites to run the jobs in case of downtimes or site complications.

Also it has changed the access of Athena software and ATLAS detector database. Both parts are necessary for ATLAS jobs, and for that reason, they have been distributed to all Grid sites (it had been inefficient to send Athena and database every time a job is submitted). Therefore, every Athena release version was installed in every site ready for the analysis jobs. The complete database release was only in Tier-1 centres and one part (the only information needed for simulation and analysis and not for reprocessing) was distributed as one file to the Tier-2s by DQ2. This database access mode has limited the assigned tasks of the Tier-2 centres to only analysis and production.

When the analysis jobs volume increased at the same level of the simulation jobs, a bottleneck problem was observed. Many requests for accessing the same Athena version in the same site or for accessing the same database file produced this problem. Therefore, the systems were changed to CernVM-FS (Cern Virtual Machines - File System) of Athena releases and Frontier/Squid for database [54]. Both systems have a infrastructure of cascading caches, which prevent direct access to the CERN serves by most of the jobs. The CernVM-FS system is based on HTTP protocol with 3 level-1 caches and allows the access of all Tier centres to Athena releases installed only at CERN. The Frontier/Squid is also based on HTTP protocol with caches at each site and all the Tier-2 centres can access to the whole database stored at CERN. This avoid a high latency if many sites want to request the access at the same time. Due to this system, Tier-2 centres can use all the database to

run every kind of ATLAS jobs including reprocessing. The real advantage of

Apart from all these improvements, for doing more them, a survey to Distributed Analysis end-users has been done in March 2011. This let know the user needs and their complications during the process. Interesting information was resulted, for instance, a better reliability is required from all Grid because around 10-20% of jobs were resubmitted for transient problems. That is why the auto-retry was developed. It consists in rerunning the failed job by PanDA system when these transient errors are detected. Another idea from this survey was to send automatic jobs only for merging the output files. It is called OutputMerging. This allows to save time to the user and makes the transfers go faster because the number of files decreases.

4.5 Distributed Analysis Support Team

The Distributed Analysis Support Team (DAST) [77] is a group of experts shifters that provide to user the first point of contact to addressing all Distributed Analysis questions, including analysis tools related problems, DDM related problems and in some cases offline software problems. First DAST shifts started in September 2008. DAST shifters must know the whole Distributed Analysis system, they have experience in the analysis tools and they identify the causes of the user issues. DAST shifters allow the understanding between developers and end-users. I am one of these DAST experts and I have assisted 442 cases in the period from 2nd February 2010 to 30th October 2013. This activity is one of my main contributions to the ATLAS Computing within the Spanish Tier-2 project.

The activity of DAST is considered one of ADC shift tasks to contribute in ATLAS. There is always a shifter in the Europe and in the North America time zones working from his/her institute. Another shifter can be incorporated like support, specially if the first shifter has not enough experience as DAST member. Shifters have a shared e-mail account where to label the user cases according to the status of the issue, the topic and the person assigned. Also, this account make easier a quick search in the cases history. The shifter can answer to the user about the solution directly or escalate the issue to an expert. Most of these experts are joined in DAST e-groups ¹ and reply users without DAST intervention. The shifters can create GGUS tickets if it is a site problem or savannah ticket in case of bug in one of the Grid tools.

¹E-groups is the interface to manage groups at CERN. The e-groups can be used as a mailing list or to grant access to different CERN resources.

One example of DAST procedure is explained below:

- some Distributed Analysis jobs are failed because an unknown reason for a user looking the PanDA monitoring.
- A simple test is done in the local computer with success, thus it is decided to ask for help to DAST.
- The e-groups hn-atlas-dist-analysis-help webpage is consulted, if there is another similar problem reported by some user before. Because nothing is found, then an e-mail is sent to the e-groups list.
- This e-mail explains briefly what it is going on, with the PanDA monitoring link of failed analysis jobs, the error messages and the expected result.
- This e-mail goes to a shared e-mail account and is seen by DAST member on shift. This new e-mail has the label 'ACTIVE' which indicates the status issue.
- The e-mail is also labelled by the shifter indicating he (or she) is looking after the issue. Thanks to this, other shifters will follow the rest of cases and the action is not repeated.
- Now, the user e-mail and the log files with the execution messages resulted from the PanDA monitor link are read carefully.
- The main error message is found in the log files and searched in the cases history. Looking the error in old similar cases, it is considered a site problem. Always, the main objective in every case is to identify the type of the problem. Basically, there are four types of problem: Grid tools bugs, site problems, global services failures, and user mistakes. According to this, the thread is labelled in function of the category.
- In the case of site problem, a GGUS ticket is created where the situation is described and it is sent to the site.
- An e-mail is sent to the user by the shifter explaining the reason of the error and noticing the GGUS ticket.
- When the problem is fixed by site administrators and/or the GGUS ticket is solved, then the user e-mail is archived and the issue is considered closed.

At the end of the week, the shifter writes a report with all the cases occurred. The most important information reported is the issues that are still open in order for the next shifter can know the work done and continue in this point. DAST members meet every Monday to discuss the cases during the week. In Figure 4.6 an example of e-mail list labelled is shown.

<input type="checkbox"/>	☆	Regina Mota Valls	Status/ACTIVE	Site has been blacklisted - Dear DAST experts, I have problems to re	4:25 pm
<input type="checkbox"/>	☆	Julien, ms. Christopher (2)	Assigned To/EU1	Status/ACTIVE Status/WAITING Re: A question about dataset cor	4:22 pm
<input type="checkbox"/>	☆	Fabien, Mlle (2)	Assigned To/PAT	Re: CLHEP and ROOTCOREDIR AFS vs Grid (FIXED) - Hi Fabien,	3:27 pm
<input type="checkbox"/>	☆	Feder . Feder (7)	Assigned To/PAT	Problem with usin Event Loop Grid Driver: unable to submit for :	2:41 pm
<input type="checkbox"/>	☆	Marcus, ms. Rodney (2)	Assigned To/EU1	Category/Site Problem Status/ACTIVE Status/USER2USER File copy	1:57 pm
<input type="checkbox"/>	☆	Katherine . Stephens (2)	Category/hammercloud	Status/ACTIVE Status/USER2USER Problems at Glasgow? - Ti	1:57 pm
<input type="checkbox"/>	☆	Giovanna . Rodney (2)	Assigned To/EU1	Category/hammercloud Category/Site Question Status/ACTIVE Status	11:39 am
<input type="checkbox"/>	☆	Tom, Jenni, Rob, Mike (2)	Assigned To/SW	reco-trf - still not working... - hi Tom Happy New Year to you too... Im i	10:36 am
<input type="checkbox"/>	☆	Sarah Williams	Status/ACTIVE	Re: Jobs stuck in pre-run - Hi Lauren, According to Panda, they're w	1:17 am
<input type="checkbox"/>	☆	Stephen, ms (2)	Assigned To/EU1	Catego.../Release Question Status/ACTIVE Status/WAITING Compiler	Jan 6
<input type="checkbox"/>	☆	Geoff, ms (2)	Assigned To/EU1	Category/DaTRI Status/ACTIVE Status/WAITING cancel datri reques	Jan 4
<input type="checkbox"/>	☆	Neil, ms, Arind (7)	Assigned To/NA1	Catego.../Release Question Category/Site Problem Escalated.../Cloud su	12/21/12

Figure 4.6: This is an example of the DAST list e-mails in a flash where the labelling procedure is shown. User names have been blurred to reserve the privacy.

The DAST activity is more fundamental because of the complicated identification of the problem. Normally, the user cannot recognise why his/her job or transfer has been failed and needs the intermediation of someone with more experience. Thanks to this group, around 40 issues are resolved per week and let to users discuss Distributed Analysis topics among other users.

4.6 Real Distributed Analysis Workflow Example

The Distributed Analysis system deployed on the ATLAS Grid infrastructure has been absolutely necessary to access to the data and process them. As an example, The work described in this section corresponds to the personal activity of the author for this PhD measurements in Jet Substructure which is reported in Chapter 6. ATLAS official AODs data were the input files for analysis jobs on Grid. In Table 4.1 the size and the events of the AOD input datasets used are shown.

The data size to analyse reached 7.66 TB for 15 millions simulation events in several types of generators and 25 TB for 158 millions of the real data with $\sqrt{s} = 7 \text{ TeV}$ and integrated luminosity of 35 pb^{-1} in 2010. It is expected the real data analysis would last more than the simulation data since it has the threefold size.

4.6. Real Distributed Analysis Workflow Example

Table 4.1: The input files information for Monte Carlo (MC) and real data (RD) studies. The number of datasets, the number of files, the size of the whole samples, the size per file, the number of events, the number of events per file and the size per event are displayed.

Inputs	Datasets	Files	Size (TB)	Size/file (GB)	Events	Events/file	Size/event (KB)
MC	56	3060	7.66	2.56	15227317	4976.25	514.44
RD	69	8158	25.23	3.09	157780181	19340.55	159.77

These jobs have taken the essential information which was stored in ROOT n-tuples with small size to allow the last analysis step locally. A workflow of the Distributed Analysis jobs is displayed in Figure 4.7. Ganga was the ATLAS Grid tool. The interface was used to send also the Athena package needed and to execute the suite tool to find the input data with DQ2 for the Monte Carlo case and AMI for real data using a Good Run List ². With this information, Ganga got the input files name and their number for the splitting in subjobs. Apart from that, the input datasets information gave the sites where they were stored and Ganga checked if there was at least one available site.

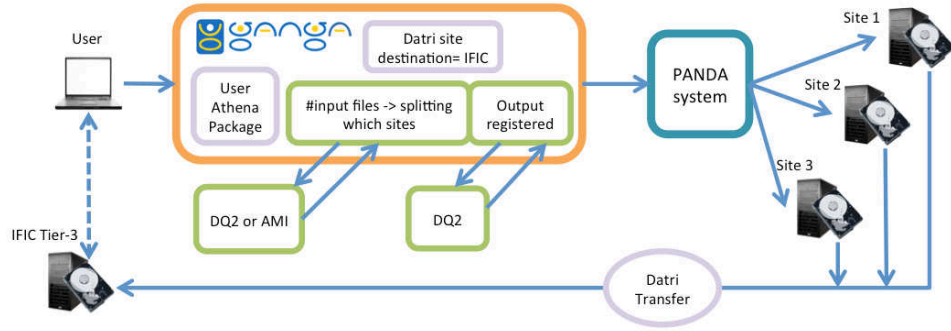


Figure 4.7: Distributed Analysis workflow for this example.

The output datasets were registered in DQ2 according to the user definition after Ganga had checked the name did not exist. Ganga also gave the instructions for a DaTRI request when output files would have finished defining the destination site which was IFIC Tier-3 centre [79] for this analysis. Grid jobs were split into 11373 subjobs according to the condition one file per subjob as it can be seen at Table 4.2. Therefore, PanDA server was contacted and it chose the suite Grid sites for running. Jobs were assigned to 36 sites where

²The Good Run List is a file in xml with the information from Data Quality [78]

Table 4.2: The subjobs of this Distributed Analysis Example: the number of subjobs with Monte Carlo (MC) and Real data AOD input files, the subjobs status and the resubmission times to achieve all the output ntuples.

Inputs	subjobs	success	failed	cancelled	resubmissions
MC	3205	3060	145	27	1
RD	8168	8157	11	1	6
Total	11373	11217	156	28	7

the 2.75% of the total subjobs run into the ATLAS Spanish Tier-2 [80]. One of the main subjects of this work is to report the behaviour of this Tier-2 as well as the use of its User Interfaces, the local storage and the User Support help (an extended description in Chapter 5). Only 1.37% of all jobs have been failed and its subjobs have been resubmitted successfully except one subjob. The problem was in a corruption file from real data. This subjob was sent to different sites at least 6 times but the result was the same: an error because of some variables were not found in the Athena running. Despite of this, results have not been affected because the broken file represented the 0.02% of all events.

The failure of the pilot-job has happened twice, in this case the subjobs to a specific site have been cancelled. To solve this situation, same subjobs have been sent to other sites without problems. In the Figure 4.8 the subjobs distribution in ATLAS PanDA queues are shown. Although the Monte Carlo and Real data analysis have been done in different moments separated in months, the biggest part of the subjobs went to the same site: the ANALY_MWT2 PanDA queue which belongs to the MidWest Tier-2 centre of United States of America. This is because that site was assigned by the system to run jobs for input datasets with the most number of files in both cases.

It is observed in Table 4.3 the CPU time average per event was 0.0043 seconds for Monte Carlo data and 0.0015 seconds for real data. According to these times, if this study had been done in a local computer with one CPU and subjobs had been run without splitting, the CPU time for all events in Monte Carlo would have been 18 hours and for all real data events 64 hours. Although the processing time of the Monte Carlo AODs on Grid is 13 hours higher than the hypothetical local running (because of the extra time for the pilot job, the PanDA priority assignment, the output datasets registration in DQ2 and the resubmissions), the time for the data jobs has been reduced in 28 hours.

4.6. Real Distributed Analysis Workflow Example




Figure 4.8: Subjobs distribution according to PanDA queues which represents the sites where analysis jobs have been run for the Jet Substructure analysis for Monte Carlo (top) and for Real data (bottom).

Table 4.3: Jobs time statistics with Monte Carlo (MC) and real data (RD) AOD input files: the time interval where the job submissions have been done, the walltime of all jobs, the events processed per second, the walltime divided by the number of events, the CPU time per event on Grid and the CPU time if all the events are analysed in a local computer with one CPU.

Inputs	Submissions interval	Walltime (hh:mm)	events/ Walltime (1/sec)	Walltime /events (sec)	CPU time /event (sec)	CPU time locally (hh:mm)
MC	9 days	31:02	54.53	0.0073	0.0043	18:10
RD	3 days	36:16	63.61	0.0014	0.0015	64:46

Table 4.4: The Output Files information for Monte Carlo (MC) and real data (RD) studies. The number of datasets, the number of files, the size of the whole samples, the size per file, the number of events, the number of events per file and the size per event are displayed.

Outputs	Datasets	Files	Size (GB)	Size/file (MB)	Events	Events/file	Size/event (KB)
MC	59	3060	89.41	29.22	15227317	4976.25	5.83
RD	77	8157	185.87	22.79	157718929	19335.41	1.18

The output ntuples size has been reduced in 98.8% with 89 GB for Monte Carlo and in 99.2% with 186 GB for real data which is shown in Table 4.4. These n-tuples have been registered in DQ2 and transferred to the IFIC Tier-3 centre [79] by DaTRI thanks to the automatic request. Thus, the access to this Tier-3 storage was easy and direct because of the Storm+Lustre³ to do selections and final results.

We should take into account that this study was processed before most of the applications of Distributed Analysis improvements were implemented. For instance, the auto-retry commented in section 4.4 could have been useful to avoid the resubmissions and this Physics study would have been faster, specially for Monte Carlo data which they were analysed at the end of 2010 meanwhile the real data were processed in the middle of 2011.

³the Lustre file system using Storm as Storage Resource Manager [81]

Chapter 5

The ATLAS Spanish Tier-2 Federation

The ATLAS Computing model according to WLCG project and the Distributed Analysis have been presented in Chapters 3 and 4. The main goal of this chapter is the description and performance of the ATLAS Spanish Tier-2 (ES-ATLAS-T2). The reason is two-fold: to give an approach to the daily work from the perspective of an ATLAS tier-2 and it serves as an illustration of how ATLAS evaluates a site based on some monitoring metrics [82]. In addition, I have contributed to the operation of the ATLAS Spanish tier-2.

ES-ATLAS-T2 is a federated Tier-2 formed by three institutions: the Institut de Física d'Altes Energies of Barcelona (IFAE), the Universidad Autónoma of Madrid (UAM), and the Instituto de Física Corpuscular of Valencia (IFIC) [80]. The location of these three sites can be seen in Figure 5.1. It is associated to the Iberian Cloud, in the South West of Europe. The Tier-1 assigned is the Port d'Informació Científica (PIC) in Barcelona. There is another Iberian federated tier-2 in Portugal formed by LIP (Laboratório de Instrumentação e Física Experimental de Partículas at Coimbra) and LI-PLNEC/FCCN (LIP, Laboratório Nacional de Engenharia Civil and Fundação para a Computação Científica Nacional) and outside the European zone there is a Tier-3 in Valparaíso, Chile as part of the E-science infrastructure shared between Europe and Latin America (EELA) [83].

The history of the ES-ATLAS-T2 started in 2002 when the LHC Computing Grid (LCG) was created, the previous project to WLCG, coordinated by CERN [27]. The purpose was to apply a Grid computing system to allow physicists from all the world to analyse the huge quantity of data from the



Figure 5.1: Logo and map with the location of the ATLAS Spanish Tier-2 sites.

LHC. Thus, the computation changed from being centralized at CERN to be distributed between computer centers all over the world. These centers belonged to institutes involved in the LHC, which began to share computing responsibilities for the Experiments. Then, Spanish institutes that have been working on detectors or Physics studies related to LHC joined that Grid project. The institutes were IFCA (Santander), IFIC (Valencia), CIEMAT (Madrid), UAM (Madrid), UB (Barcelona), USC (Santiago de Compostela) and IFAE (Barcelona) [84].

After some national projects to establish the base of a Grid infrastructure in Spain (Preparation of Local Infrastructure for participating in Grid for ATLAS 2000-2002 and European Data Grid 2002-2005), the ES-ATLAS-T2 started to be developed in 2005 with the objective of doing an installation, performance and maintenance of a Tier-2 infrastructure for Distributed Analysis and Monte Carlo Production in ATLAS. This infrastructure must fulfil with the ATLAS collaboration requirements and to have an optimal stable performance for serving the ATLAS physicists in their studies.

The size of the ES-ATLAS-T2 in what refers to resources was set to 5% of all ATLAS Tier-2 centers.. Because it is a federated Tier-2, resources are shared in three sites. Thanks to this, there are experts from the three institutes who share responsibilities, exchange ideas, and solve problems together in shorter time. These experts meet every two weeks by videoconference and face to face meetings every 6 months. IFIC assumes the coordination of all the ES-ATLAS-T2 activities. IFIC provides 50% of the resources while UAM and IFAE 25% each. From 2006 till 2013 the ES-ATLAS-T2 has produced many million of Monte Carlo events for the ATLAS Collaboration and has provided a very good service to users from all over the world to analyze data. In particular, part of the data for the Higgs discovery was recorded on the

ES-ATLAS-T2 disk servers. The fast access and good performance of these servers helped to finish the analysis in time, as a consequence of an excellent operation at the three sites of the ES-ATLAS-T2.

5.1 ES-ATLAS-T2 Resources

The resources assigned for the ATLAS experiment are fixed every year by WLCG project in the Memorandum of Understanding (MoU) [85]. The ES-ATLAS-T2 infrastructure has been increased in terms of these pledges. In Table 5.1 ES-ATLAS-T2 pledges are shown and one can observe the rising evolution. It is observed that disk are accounted in Terabytes (TB) and CPUs in HEP-SPEC06 units. The HEP-SPEC06 [86] is a HEP-wide benchmark for measuring CPU performance and, in this way, different CPU models can be compared. Currently, a Intel Xeon X5690 processor has around 247.48 HEP-SPEC06.

We could highlight the important ES-ATLAS-T2 hardware increments, particularly the first years before the data taking 2006-2009. The most significant increment was in 2008 where the resources have been extended at 600% regarding to 2007. In the rest of the years 2010-2014, the increase is smoother around 80-20% where in 2012 drops a little the required CPU.

The CPU and disk values in June 2013 in ES-ATLAS-T2 are shown in the table 5.2. It is observed the pledge is carried out roughly. Also, IFIC contributes around 50% of the infrastructure of ES-ATLAS-T2, as was established. In Figure 5.2 some views of the Valencia, Madrid and Barcelona infrastructures can be seen.

As mentioned before, these resources are used for processing the jobs from users' analysis and from official Monte Carlo production. According to the ATLAS collaboration, a Tier-2 centre must divide its resources in these two tasks to 50% each other.

The computing resources are accessed via 6 CEs and 4 SEs. In table 5.3 these Grid elements are shown, they are distributed in the different ES-ATLAS-T2 sites. To complete the ES-ATLAS-T2 services, each site has also a BDII

Table 5.1: The evolution of the pledges in CPU and disk for the ES-ATLAS-T2.


ES-ATLAS-T2	2006	2007	2008	2009	2010	2011	2012	2013	2014
CPU(HEP-SPEC06)	92	243	1750	5390	10308	13900	13300	18000	20600
DISK(TB)	14	63	387	656	1107	1880	2350	2550	2800



(a)



(b)



(c)

Figure 5.2: Some pictures of the ES-ATLAS-T2 infrastructures: IFIC Storage at Valencia (a) and UAM Data Center at Madrid (b) and IFAE Worker Nodes at Barcelona (c).

5.1. ES-ATLAS-T2 Resources

Table 5.2: The ATLAS Spanish Tier-2 resources in June 2013.

ES-ATLAS-T2	CPU(HEP-SPEC06)	DISK(TB)
IFIC	14573	1316
UAM	3780	618
IFAE	4599	679
Total	22952	2613

Table 5.3: CEs and SEs services of ES-ATLAS-T2 distributed in its three sites.

	CE	SE
IFIC	IFIC-LCG2-CE-ce03.ific.uv.es IFIC-LCG2-CE-ce05.ific.uv.es	IFIC-LCG2-SRM-srmv2.ific.uv.es
IFAE	ifae-CE-ifaece02.pic.es ifae-CE-ifaece03.pic.es	ifae-SRM-srm-disk.pic.es ifae-SRM-srmifae.pic.es
UAM	UAM-LCG2-CE-grid001.ft.uam.es UAM-LCG2-CE-grid003.ft.uam.es	UAM-LCG2-SRM-grid002.ft.uam.es

server and 11 UIs in total. These Grid elements have been defined in Chapter 1.

The file systems to manage the SEs space are different in the ES-ATLAS-T2 sites. In UAM and IFAE (disk+SRMposix) dCache is used [87]. This system was specially developed for High Energy Physics to manage a huge amount of data. The file system chosen by IFIC is Lustre [81] which lets a easy access to the data stored in SE for the local end-users.

The network is provided by the Spanish NREN (National Research and Education Network) RedIRIS [88] and local providers for the last mile. The links are 10 Gbps between POPs (Point of Presence) with alternate paths for backup. Connectivity from the sites (IFIC, IFAE, UAM) to network is at 10 Gbps. In addition, these sites are interconnected by a triangular backbone link provided by RedIRIS.

A monitoring system is important for detecting quickly any possible failure of the services or the hardware. Two modes are being used, one from the global LHC Grid for instance the Service Available Monitoring (SAM) [89] and the other is an internal monitoring with tools like Nagios or Cacti ([90], [91]).

The software installation for many machines requires automatic tools, in particular to install and configure the operating system, Grid middleware and the storage system. PUPPET [92] is an example, used by the Spanish three sites. The ATLAS software is accessed for the sites machines using the CernVM-FS ([50]) which is detailed in Chapter 4.

5.2 Reliability and Availability of the Spanish ATLAS Tier-2

WLCG Collaboration considers essential to control the performance of the Tier centres. Besides confirmation the resources pledges, a Tier-2 infrastructure must be available the longest possible time for the end-user. The way to do that is with reliability and availability metrics. They are defined by the Service Availability Monitoring system (SAM) [89] as follows:

$$Reliability = \frac{Uptime}{Total\ time - Scheduled\ Downtime - Time\ status\ unknown} \quad (5.1)$$

$$Availability = \frac{Uptime}{Total\ time - Time\ status\ unknown} \quad (5.2)$$

Availability is the ratio of the time when the site is available over the total time and reliability is the same ratio corrected the total time by taking away the scheduled downtime which is the period in a known technical stop of the site usually for system updates or for adding new hardware. To get this metrics, SAM system runs a range of critical tests at regular intervals of time throughout the day over the Tiers. These Tier centres are considered to be reliable and/or available when these tests are completed successfully.

Applying these metrics to ES-ATLAS-T2, Figure 5.3 and 5.4 show that the Spanish Tier-2 federation has been working with good levels of reliability and availability from January 2010 to February 2013.

In some punctual moments, the value drops considerably in one of the sites, but the ES-ATLAS-T2 average has been keeping between 90-100% in both metrics because it is federated. In particular, the availability has more falls because of considering the scheduler downtimes in the statistics. These periods of time do not reflect a problematic situation however they are not accessible moments for the end-user. Therefore, both metrics are necessary to find out the cause of a inaccessibility of a Tier centre.

Another availability to take into account is evaluated by ATLAS Computing operations. It consists in the time when the site is online over all the time considering also the Hammercloud exclusions (see Chapter 4). In Figure 5.5 the ATLAS availability for the ES-ATLAS-T2 queues for analysis and production from September 2011 to March 2013 is shown.

5.2. Reliability and Availability of the Spanish ATLAS Tier-2




Figure 5.3: Reliability of ES-ATLAS-T2 for the last 3 years. The average is always in the interval 90-100%




Figure 5.4: Availability of ES-ATLAS-T2 for the last 3 years. The average is always in the interval 90-100%

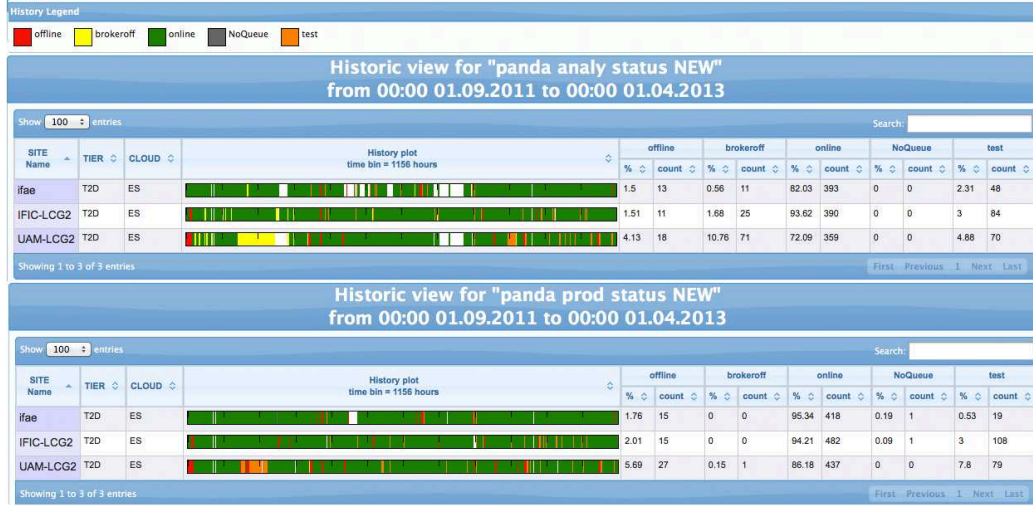


Figure 5.5: The ATLAS availability of the ES-ATLAS-T2.

It can be seen that the analysis availability is between 70–90%. In the case of the production availability, IFAE queues have the best values around 95% meanwhile UAM production queues are around than 86%. If the average is made, ATLAS analysis availability has a value of 82.58% and production 91.91% for the whole Spanish Tier-2 federation. In that case, the availabilities can be considered acceptable.

5.3 Data Transfers of the ES-ATLAS-T2

As commented before, in ATLAS Computing, data is being transferred among Tier centres after the reprocessing and after the ATLAS official production. Data replicas are transferred also with PD2P (see Chapter 3) according to the data use in an automatic way and with DaTRI when the user request it. In addition, there are data transfer tests every day and sometimes specific checking such as Sonar tests to measure the connections between sites [49].

Transfer can be considered as arrows where some parameters are taken into account to control their performance: source site, destination site, number of transfers, transfer speed of the data volume and if the transfer was successful or not. All this can be observed in the next set of Figures for the ES-ATLAS-T2 comparing with all ATLAS transfers from January 2010 to 1st April 2013. This information are acquired from the ATLAS Dashboard [93].

The represented parameters are the throughput (amount of data size trans-

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.6: Throughput of all ATLAS transfers in function of different Clouds sources from January 2010 to 1st April 2013. The behaviour is growing in average (it follows a raising trend).

ferred in function of time), the transfer efficiency (transfer successes regarding the total number of transfers) and the number of transfers with success according to the sources, the destinations and the transfer activities.

First, the throughput of the all ATLAS transfers is displayed in Figure 5.6. It has been increasing for the last three years from 782 MB/s in January 2010 to the maximum of 9426 MB/s in January 2013. Around 398 PB has been transferred as it is shown in Figure 5.7.

Then, the same parameter for only the transfers to ES-ATLAS-T2 in Figure 5.8. The most important effect to be drawn from this Figure is the change in the model of computation. Before 2011, hierarchical model has still working in the main activities and all transfers to ES-ATLAS-T2 passed through its Tier-1 centre (PIC). After that, steadily transfers come from every possible Cloud, although the most of them are from Iberian Cloud (ES in Figures). The Throughput increased reaching a maximum of 243 MB/s in February 2012. According to Figure 5.9, 9.60 PB has been transferred to ES-ATLAS-T2.

The throughput of the transfers from ES-ATLAS-T2 to the rest of ATLAS has been also increasing which is shown in Figure 5.10. The peak is 127 MB/s in February 2013 whose the highest contribution is from IF AE. Around 3547 TB has been transferred from the Spanish Tier-2 federation according to the Figure 5.11.




Figure 5.7: Cumulate volume of all ATLAS transfers in function of sources grouped in Clouds from January 2010 to 1st April 2013 in constant raise.




Figure 5.8: Throughput of ATLAS transfers to ES-ATLAS-T2 in function of sources grouped in Clouds from January 2010 to 1st April 2013. It can be observed the change in the computing model when the transfers have only one colour (Iberian Tier-1, PIC) to different colours in the histograms.

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.9: Cumulate volume of ATLAS transfers to ES-ATLAS-T2 in function of sources grouped in Clouds from January 2010 to 1st April 2013.

The efficiency of all ATLAS transfers is shown in Figure 5.12. It is observed the transfer efficiency is between 80-100% with punctual drops lower than 75% from some Clouds, inclusively decreasing less than 30% in February 2013. Taking into account all the transfers for these three years, the total transfer efficiency was 92%.

The efficiency for the transfers to ES-ATLAS-T2 is shown in Figure 5.13. Its values are between 80-100% in the beginning of 2010 when the transfers were from the Tier-1 PIC. Then in October 2010, transfer tests started where other Clouds came in as sources and the efficiency range was widened to 30-100%, dropping less than 5% in November 2010. At the end, the new computing model was applied at the beginning of 2011. The efficiency has been improving in that moment, specially in the final months, for instance the value rate is 90–100% in November 2012. Accounting all the transfers to the Spanish Tier-2 federation, the efficiency is about 88%.

The efficiency of the transfers from the ES-ATLAS-T2 is displayed in Figure 5.14. It can be seen that the efficiency is better than the opposite direction, in the range of 100-80%. There are only four moments when efficiency became out of this range. Considering all the transfers from this Tier-2 centre, the total efficiency is 94%, a little higher than the all ATLAS transfers efficiency.

The transfer efficiencies have to be evaluated taking into account failures can be possible for three situations: problem in the source site, error from the




Figure 5.10: Throughput of ES-ATLAS-T2 transfers to all ATLAS in function of the three sites as sources from January 2010 to 1st April 2013.




Figure 5.11: Cumulate volume of ES-ATLAS-T2 transfers to all ATLAS in function of the three sites as sources from January 2010 to 1st April 2013.

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.12: Efficiency of all ATLAS transfers in function of sources Clouds from January 2010 to 1st April 2013. It almost stable in 80-100%.




Figure 5.13: Efficiency of ATLAS transfers to ES-ATLAS-T2 in function of sources Clouds from January 2010 to 1st April 2013. The new computing model implementation is observed.




Figure 5.14: Efficiency of ES-ATLAS-T2 transfers to all ATLAS in function of the three sites as sources from January 2010 to 1st April 2013.

destination or a lost connection. Therefore, the low transfer efficiencies related to ES-ATLAS-T2 could be for external reasons. For instance a problem in other sites caused the value of 3% in Figure 5.13 due to long time waiting the action of the source site during a functional test.

The number of successful transfers in ATLAS is displayed in Figure 5.15. The transfers has been increasing with a peak of 41.7 million in November 2012. Around 1007 million of transfers have been done for last three years according to Figure 5.16.

The number of successful transfers to ES-ATLAS-T2 is shown in Figure 5.17. Accounting all the transfers to this Tier-2 federation, 14.6 million has been transferred in this period which is displayed in Figure 5.18. The tend is to grow up reaching 810000 transfers in November 2012. Also, the computing model change is noticed in the beginning of 2011 , as we have seen before in the throughput and the efficiency.

Next, the successful transfers from the ES-ATLAS-T2 is displayed in Figure 5.19. An upward tendency is noticed too, however it is not constant with a peak of 633 thousands transfers in February 2013. Around 14.2 million of transfer successes have been sent by the ES-ATLAS-T2 which is shown in Figure 5.20.

The Figure 5.21 shows the same number of transfers to ES-ATLAS-T2 seen in Figure 5.17, but this histogram is in function of the destination site, therefore,

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.15: Transfer Successes of all ATLAS transfers in function of sources grouped in Clouds from January 2010 to 1st April 2013.




Figure 5.16: Cumulate transfer Successes of all ATLAS transfers in function of sources grouped in Clouds from January 2010 to 1st April 2013.




Figure 5.17: Transfer Successes of ATLAS transfers to ES-ATLAS-T2 in function of sources Clouds from January 2010 to 1st April 2013. The change of the computing model is significant in March 2011.




Figure 5.18: Cumulate transfer successes of ATLAS transfers to ES-ATLAS-T2 in function of sources Clouds from January 2010 to 1st April 2013.

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.19: Transfer successes of ES-ATLAS-T2 transfers to all ATLAS in function of the three source sites from January 2010 to 1st April 2013. An upward tendency is noticed.




Figure 5.20: Cumulate transfer successes of ES-ATLAS-T2 transfers to all ATLAS in function of the three source sites from January 2010 to 1st April 2013.




Figure 5.21: Transfer Successes of ATLAS transfers to ES-ATLAS-T2 in function of the three sites as destinations from January 2010 to 1st April 2013.

the three Spanish sites can be distinguished. In this way, another point of view can be noticed. It is observed that the transfers received increase almost twice after July 2011 (from 300000 to 600000). Transfers tend to be proportional to the resources percent until the date mentioned, where IFAE are receiving more than IFIC.

The Figure 5.22 displays the number of successful transfers from ES-ATLAS-T2 in function of the destination. In this point of view, the computing model changing is discerned. Previously, all the Tier-2 transfers run through its assigned Tier-1, therefore, the only destination at the beginning is the Iberian Cloud. Then, every Cloud receives transfers from ES-ATLAS-T2, although Iberian Cloud transfers the majority contribution.

The number of transfer successes in all ATLAS in function of the activities is shown in Figure 5.23 and in this way, transfers behaviours can be displayed. The activities are labelled in the plot as:

- 'T0 Export' are the transfers of ATLAS data from Tier-0 to Tier-1 centres and Tier-2 with calibration tasks.
- 'Data consolidation' are the daily transfers of the ATLAS official data as AOD from Monte Carlo or reprocessing.
- 'Data brokering' are the transfers of data replicas generated by PD2P exclusively.

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.22: Transfer Successes of ES-ATLAS-T2 transfers to all ATLAS in function of the destination grouped in Clouds from January 2010 to 1st April 2013. The Iberian Cloud is still the destination of the most of transfers despite the computing model change.

- 'Functional Test' are transfers only for daily testing.
- 'Production' is referring to all the transfers during the ATLAS official production.
- 'Group Subscription' are the transfers requested by Physics Groups using DaTRI.
- 'User Subscription' are the DaTRI transfers requested by users.

Two activities have been added for two years thanks to the Distributed Analysis improvements and new analysis tendency to use D3PD as input. They are the PD2P activity and the DaTRI subscriptions for the Physics groups which appear in August 2010. The highest percentages are from 'Production' with 35%, the next is 'Data consolidation' with 19% and 'User Subscription' with 29% which reflects the Distributed Analysis usage.

The next Figure 5.24 shows the transfers to ES-ATLAS-T2 according to the ATLAS activities. The PD2P transfers appear in November 2010, two months like in the previous figure. It is observed a lot of user subscriptions with an almost constant flow around 200000 per month. Although new activities have increased the transfers volume, the activities with the most contribution are Data Consolidation 29%, User Subscription 27%, Production 16% and Data Brokering 14%. The great percentage in the user data transfers reflects the




Figure 5.23: Transfer Successes of all ATLAS transfers in function of the transfer activities (top) and its percentages (bottom) from January 2010 to 1st April 2013. The most of transfers are from 'Production', 'Data Consolidation' and 'User Activity'.

Spanish users analysis activity. The small contribution of 'T0- Export' 2% are due to the transfers to IFIC which is a site assigned to do ATLAS calibration tasks, which detector components variables are needed to the correct Physics analysis, since April 2011.

Finally, the Figure 5.25 displays the successful transfers from the ES-ATLAS-T2 in function of the activities. We note that the majority of the transfers for the ES-ATLAS-T2 are those for Production 65%. The next significant activity is the user subscription (22%) since the users from every part of ATLAS centres execute DaTRI to move their output files stored in ES-ATLAS-T2 to Tier-3 centres or similar places of their institutes.

In conclusion, hundred of million of transfers have moved hundred of PBs for the last three years in ATLAS. A part of them were related to the ES-ATLAS-T2 with a good efficiency average reflecting the computing model change and the ATLAS activities, especially the user actions.




Figure 5.24: Transfer Successes from ATLAS to ES-ATLAS-T2 in function of the activities (top) and its percentages (bottom) from January 2010 to 1st April 2013. 'Data Consolidation' is the activity with more transfers, and the next is the 'User Subscription' due to the Spanish users.

5.3. Data Transfers of the ES-ATLAS-T2




Figure 5.25: Transfer Successes from ES-ATLAS-T2 to all ATLAS in function of the transfer activities (top) and its percentages (bottom) from January 2010 to 1st April 2013. The most part of the transfers in a Tier-2 centre is due to the 'Production' activity.

5.4 Data stored in ES-ATLAS-T2

Data transfers have been seen before, the data stored in ES-ATLAS-T2 after these transfers is shown in this section. Independently the file system used, every ATLAS Tier-2 centre has distributed data in space tokens. They are reserved storage space for entity or purpose [94]. In particular, permissions are defined to users or activities for reading or adding or removing data in a space according to the Grid certificate role and the approval of people in charge.

The ATLAS Tier-2 space tokens are the following:

- DATADISK: where the official ATLAS data are stored, from the experiment or from the Monte Carlo production. Only users with production role can store in.
- HOTDISK: files with frequent use are stored in, like database files. Currently, because of the database access has been changed by Squid system (see Chapter 4), its use is the shortest of all the space tokens and this space token is removing in every Tier.
- CALIBDISK: is the place booked for calibration operations of the ATLAS detector, only in some sites of Tier-2 federation, for instance the IFIC from ES-ATLAS-T2.
- PRODDISK: is the specific space for the files needed during the process of the ATLAS simulation production (commented in Chapter 4).
- GROUPODISK: a space for ATLAS Physics groups. Only users working in the Physics group can store replicas under the approval of their conveners.
- SCRATCHDISK: where all output analysis job files are stored when the process has finished. It is a temporal space, user must move the output files in less than one month. Every user has permission to use it.
- LOCALGROUPODISK: is the space for local users of an institute or a Cloud. It is frequently a Tier-3 storage not subject to ATLAS pledges.

The Figure 5.26 shows the occupancy in function of the space token of all ATLAS Tier-2 centres. The greatest space is the DATADISK with 62% which is followed by the GROUPODISK if the different related space token




Figure 5.26: Occupancy percent of the space tokens at all ATLAS Tier-2 centres.

are added with a percentage of 19%. The next significant space token is the LOCALGROUPDISK with 17%.

The next Figure 5.27 represents all the ES-ATLAS-T2 distinguishing the great domain of the space for real and final data production around 57%. The second most important contribution is the space for Physics groups in 31%.

Figure 5.28 corresponds to each particular site (IFIC, IFAE and UAM). In each pie graph, it is discerned the high occupancy of GROUPDISK, the highest is SUSY at IFIC and the next is Higgs at UAM. The assigned space token to the top physics group at IFAE is also high with 149 TB but its LOCALGROUPDISK has also occupied space due to the frequently use of its Tier-3 centre by the local users in 87 TB.

In Figure 5.29, the total occupancy in function of ES-ATLAS-T2 reaches 2 PB where it is noticed that the sites percent is proportional to the resources. The little excess of IFAE occupancy can be explained for the LOCALGROUPDISK influence. If we compare the ES-ATLAS-T2 occupancy regarding all data stored in the ATLAS Tier-2 centres, the contribution of the Spanish Tier-2 federation is 5.1%.




Figure 5.27: Occupancy percent of the space tokens at ES-ATLAS-T2.

The ATLAS data occupancy evolution for four years is shown in Figure 5.30. An increasing of the space every year reaching 140 PBs. This disk space accounts for the data from the detector, but also the replica, the simulation data and the ATLAS user results (ntuples).

The next group of figures show the space tokens evolution of the ES-ATLAS-T2 sites from the token creation to February 2013. The green line represents the available space, the blue one is the size of the registered files in DQ2 and the red line is related to the real Physical space. First, the DATADISK evolution of IFAE is shown in Figure 5.31. It is discerned in all plots the occupancy has increased until November 2009, when the LHC was starting, the values are the highest and trend to rise.

Figure 5.32 is the CALIBDISK size evolution. Only IFIC owns this space token for the ATLAS calibration jobs which started in March 2011. Calibration files occupy until 17 TB, and it is growing smoothly.

Figure 5.33 shows the UAM HOTDISK space token. It has few TBs assigned and they are wasted. That is it surely for the database access system change as it has been said in Chapter 4.

5.4. Data stored in ES-ATLAS-T2




Figure 5.28: Occupancy percent of the space tokens at IFIC, IFAE and UAM.




Figure 5.29: The occupancy of ES-ATLAS-T2 in function of sites.




Figure 5.30: Occupancy evolution of data in all ATLAS Grid sites.




Figure 5.31: Occupancy evolution of the DATADISK at IFAE.




Figure 5.32: Occupancy evolution of the CALIBDISK at IFIC.




Figure 5.33: Occupancy evolution of the HOTDISK at UAM.




Figure 5.34: Occupancy evolution of the PRODDISK at IFIC.

The next space token shown is PRODDISK in Figure 5.34 for the IFIC case. There are occupancy fluctuations between 5 and 10 TB because of the creation and deletion frequently of the Monte Carlo product files.

Figure 5.35 shows the UAM SCRATCHDISK evolution where output analysis jobs files are stored. As same as PRODDISK, data are kept temporally, SCRATCHDISK files are deleted after a month. This is noted in the plot profile since there are always free space at the end. However, the occupancy has been increasing in each site reaching 20 TB.

Figure 5.36 represents the space token GROUPDISK evolution for the ES-ATLAS-T2 sites. In this case, the time profile is different in each site since there are not the same Physics groups assigned for the space. At IFIC, space are reserved for Exotics and SUSY Physics group, it has been used greatly reaching 100 TB at the beginning of 2012 and 300 TB in 2013. Top Physics group has reserved its space at IFAE, the occupancy has increased also in a smooth way, reaching a peak of 150 TB. The abrupt increment is at UAM, where GROUPDISK is reserved for Physics related to Higgs boson. In October 2011, occupancy was from less than 10 TB to 70 TB (around 190 TB in March 2013).

Finally, Figure 5.37 represents the LOCALGROUPDISK evolution. It can be seen the Tier-3 use difference of each institute for local physicists. In the IFIC case, the majority use of its Tier-3 centre has been for output storage.




Figure 5.35: Occupancy evolution of the SCRATCHDISK at UAM

Therefore, the increase started more after data-taking, in November 2010 and a high occupancy reaching 75 TB. The IFÁE LOCALGROUPDISK has been filled since its Tier-3 centre is used for analysis jobs and it implies to store the D3PDs inputs too. This brought an abrupt increment in April 2011 from less than 10 TB to 50 TB and it reached 100 TB in January 2013. The LOCALGROUPDISK of UAM is the smallest, with an occupancy of 2 TB, because of a Physics local community of few users.

In conclusion, space tokens achieve each ATLAS activity owns its space while users dispose resources for Physics analysis. Observation of stored data in ES-ATLAS-T2 indicates the ATLAS computing changes and new user tendencies. The ES-ATLAS-T2 occupancy regarding all Tier-2 centres has the 5% of the data destined to be in these Tier centres which reflects a good data assignation to it in function of the resources provided by the Spanish computing community. Space tokens evolution profiles reflect characteristics of ATLAS activities, Physics analysis and local users.




Figure 5.36: Occupancy evolution of the GROUPDISK at IFIC, IFAE and UAM.

5.4. Data stored in ES-ATLAS-T2




Figure 5.37: Occupancy evolution of the LOCALGROUPDISK at IFIC, IFAE and UAM.

5.5 ATLAS Production and Analysis jobs at ES-ATLAS-T2

Production and Analysis are the most important ATLAS computing activities in a Tier-2 centre. Therefore, the performance checking has been taken into account since the Spanish Tier-2 beginnings [95].

In this section how many jobs has been sent to ES-ATLAS-T2 is shown, and the same parameter for all ATLAS to compare with the general tendency. In addition, the efficiency is displayed which was resulted on the basis of the successful jobs over all the completed jobs. The ATLAS Dashboard monitoring webpage has been used to achieve these histograms and statistics data [93], in the period of time from January 2010 to first day of April 2013. This tool was developed by ATLAS to monitoring the several operations.

5.5.1 Production Jobs

ATLAS production executes the official simulation data creation as commented in section 4.3 and recently the official DPDs for the Physics groups. These production jobs are useful to test any time the ES-ATLAS-T2 because a reliable code is executed and jobs flow is present.

The performance of this jobs flow in ATLAS and ES-ATLAS-T2 is shown in Figure 5.38. In the histogram at the top, the ATLAS completed jobs are displayed and the production is not totally regular per month. For instance, there were low activity periods like June-July 2010, and the opposite in September 2011 and August 2012.

Around 7 million of jobs per month has been processed in ATLAS reaching the 10 million in September 2011. While 120 thousand jobs per month in ES-ATLAS-T2 were executed with a maximum over 250 thousands of jobs. During all this period of time of three years, 265 million of production jobs have been run, and 4.7 million in ES-ATLAS-T2 which is discerned in Figure 5.39. The main contribution is from IFIC with 2.3 million of jobs, this is expected because this site is providing the 50% of the ES-ATLAS-T2 resources. The next contribution are from IFAE with 1.4 million and the last UAM with 958 thousand. The ES-ATLAS-T2 contribution percentage regarding ATLAS is 1.78% in production.

We follow with the jobs efficiency comparison which is shown in Figure 5.40. In the ATLAS histogram (top), the average efficiency in this period is around

5.5. *ATLAS Production and Analysis jobs at ES-ATLAS-T2*




Figure 5.38: The evolution of completed jobs in production for all ATLAS (top) and the ES-ATLAS-T2 (bottom).




Figure 5.39: The cumulation of completed jobs in production for all ATLAS (top) and the ES-ATLAS-T2 (bottom). The Spanish Tier-2 contribution is 1.78%.

95.7%. In the histogram at the bottom, the ES-ATLAS-T2 sites efficiencies are displayed, only three times the efficiencies drop less than 80% (January 2010 and September 2011 at IFIC, and March 2011 at UAM). However, the efficiency average is 93%.

5.5.2 Analysis Jobs

The activity called ‘Analysis’ consists of the jobs which are sent by ATLAS physicists to their studies. In the beginning of the ATLAS Grid use, the analysis jobs were low, however this tendency changed completely when data-taking started.

The rise began in April 2010 as it can be seen in Figure 5.41. Up and down are observed in the ATLAS histogram (top) of completed analysis jobs. These peaks can be explained if an important conference is close in time like July (iCHEP). In August 2012 case, the increment of 4.6 million could be on account of the intensive activity of the Higgs Physics group after the new particle discovered announcement to confirm if it has exactly the Higgs properties.

In the ES-ATLAS-T2 histogram at the bottom, a progressive increment is better observed, with a marked peak in February 2013 of over 400 thousands jobs. This value is higher than production jobs in the same month.

In Figure 5.42, the total number of completed analysis jobs in ATLAS in this period is 407 million whose 7.2 million went to ES-ATLAS-T2, the 1.76% percent of ATLAS jobs.

In Figure 5.43 the analysis efficiency is shown. It is observed as much as for all ATLAS as for only ES-ATLAS-T2 the analysis jobs efficiencies vary in a irregular way. The ATLAS efficiency average is 87.9% and ES-ATLAS-T2 has 89.3% for what this last value exceeds a little the general efficiency average. Both values are smaller regarding the production efficiencies although this is logical since the production jobs are stabler and more reliable.

The case of analysis jobs is the errors could not be from computing resources, they can be due to wrong user algorithms or input data not found. Therefore, the analysis jobs efficiency should not be considered a site status parameter. For that reason the HammerCloud has been created (see section 4.4), to have a constant analysis jobs flow and thus to test the sites. Its jobs have the




Figure 5.40: The success/total efficiency of completed jobs in production for all ATLAS (top) and the ES-ATLAS-T2 (bottom). The Spanish Tier-2 efficiency average in production is 93%.

5.5. *ATLAS Production and Analysis jobs
at ES-ATLAS-T2*




Figure 5.41: The evolution of completed jobs in analysis for all ATLAS (top) and the ES-ATLAS-T2 (bottom).




Figure 5.42: The cumulation of completed jobs in analysis for all ATLAS (top) and the ES-ATLAS-T2 (bottom). The Spanish Tier-2 contribution in ATLAS is 1.76%

5.5. *ATLAS Production and Analysis jobs
at ES-ATLAS-T2*




Figure 5.43: The success/total efficiency of completed jobs in analysis for all ATLAS (top) and the ES-ATLAS-T2 (bottom). The Spanish Tier-2 efficiency average is 89%.

analysis characteristics, many input/output data different from production jobs with more CPU waste.

Production and analysis are significant ATLAS computing activities in a Tier-2 centre. The monitoring of them is vital to know a Tier-2 performance. It can be consider ES-ATLAS-T2 demonstrates to be efficient during the data taking in the Grid jobs execution as same as in the optimal use of the resources.

5.6 User Support at ES-ATLAS-T2

One of the ES-ATLAS-T2 activities is the support of the users of each center in the implementation of the Distributed Analysis focused on ATLAS Grid. Each institute has its own User Support team who takes the control in creating and updating web pages with the essential information to start and to access to the Grid in each center. In addition, the team attends questions and problems of these users by e-mail, phone or face to face. Specially, I am a member of the User Support at IFIC since 2008.

The most frequent topics are how to get the Grid certificate and join in the ATLAS Virtual Organization, where the storage places are to save their own user data and how to use the ATLAS tools in the institute resources. User support are also a direct link from users to ES-ATLAS-T2 administrators in case of local problems. The team can notify Grid news as downtimes of sites or services because of a local or ATLAS global issue.

Sometime, external links are required to solve the user problems. In this case, the DAST list is consulted (see Section 4.5) or a savannah ticket is created when is a tool bug [96] or a GGUS ticket made because of a external site error. Because of the Grid experience of the User Support team, they achieve to identify better the cause of the problem and they know the steps to solve the problem.

Although there are one person or more people of User Support in each centre and take care of their own users, they exchange in every ES-ATLAS-T2 meeting. This helps to improve and collaborate like it happened in October 2009 when a Distributed Analysis tools Tutorial has been done before the data-taking. This has as effect the decreasing of user doubts in that moment. This helped to decrease the user doubts in that moment. Over time, the three User Support teams have been adapted to their user analysis styles. For

instance, IFAE team is focused on the Tier-3 activities, IFIC in the Ganga use and UAM in pathena.

ATLAS users from Madrid, Barcelona and Valencia have the extra support from local experts apart from the global established ways by the experiment who give a friendly help for their difficulties and a Grid knowledge to improve their Physics activities.

Finally, the work in the IFIC User Support by me has been intensive attending 222 cases since her incorporation as a member in 2008 until now. This experience has been very productive for the execution of the Physics Analysis and Jet Substructure and helping other studies as top-antitop resonances, tau lepton searches, SUSY, asymmetries and Higgs searches. A twiki has been done by the IFIC User Support that shows the basic knowledge to start in Distributed Analysis using the IFIC infrastructure in Reference [97].

5.7 Conclusions

The ES-ATLAS-T2 description as a federated Tier-2 assigned to the Iberian Cloud has been presented. Its history since 2002, its infrastructure and its coordination among the three Spanish sites have been included. The great effort to implement them in this Spanish collaboration has been illustrated. Its ATLAS contribution is 5% satisfying all the requirements according to a Tier-2 centre.

The current WLCG availability and reliability of the Spanish Tier-2 since 2010 are around 90–100% and its ATLAS availabilities of 82.64% for analysis and 92.58% for production. Thus, it responds to the experiment requirements.

Transfers related to this Tier-2 federation have been observed since 2010 as much received transfers (14.6 million) as sent transfers (14.2 million). Throughput has been evaluated with a peak of 243 MB/s in February 2013 and 9.60 PB total size received for last three years. For the sent transfers, the throughput reached a maximum of 127MB/s in February 2013, with a total of 3.5 PB to all ATLAS. Transfers efficiency has been 88% when reception and 94% in dispatch. This last efficiency exceed the ATLAS efficiency average. With all, Computing Model changes to pass from one transfer source to all sites have been observed. Also, it was noticed the ATLAS activities behaviour looking transfers, including the recent PD2P.

Total stored data at ES-ATLAS-T2 is a little over 2 PB (IFIC has 1007 TB, IFAE 602 TB and UAM 475 TB), the 5% over all Tier-2 centres. The storage distribution according to the ATLAS activities by space tokens has been explained and demonstrated when the Spanish Tier-2 is observed since 2010. One example of that is the huge contribution of the experiment data and their derived files. In addition, the space tokens evolution reflects the use of each activity and each Spanish user.

Jobs of the most two important ATLAS activities have been studied: analysis and production. 4.7 million of production jobs have arrived to this Tier-2 federation for last three years with a high efficiency of 93% which contributes 1.78% regarding all ATLAS. A little better results in analysis part where 7.2 million of jobs have run, contributing in 1.76% regarding all the sites with an efficiency of 87.9% higher than global average.

Finally, ES-ATLAS-T2 has provided a group of experts in computing know-how to help the Spanish physicists from the three centres and part of them have collaborated in global ATLAS tasks as DAST (the author included). Every of these aspects puts the Spanish ATLAS Tier-2 in a good position for the future challenges.

Chapter 6

Jet Substructure

6.1 Motivation: jets and boosted objects

With the LHC a new era has started for particle Physics. In 2012, the LHC was operated with a energy in center of mass at 8 TeV; four times more than Tevatron with the data collected in Run I [12]. The integrated luminosity reached was 27 fb^{-1} .

The LHC thus accesses a new kinematic regime, in which known particles behave in different way than previous experiments. The most massive SM particles have a invariant mass of order 100 GeV. When a heavy particle decays at rest, its decay products take opposite directions (back-to-back). At the LHC, these particles can be produced with a transverse momentum that is much larger.

For particles with transverse momentum higher than its mass, the decay products are emitted in small angles in the parent particle direction. Heavy particles with mass $p_T \gg m$ are called boosted objects [24]. The topologies of a top quarks decaying of at rest and a highly energetic top quark are shown in Figure 6.1.

Reconstruction of boosted objects with hadronic decays presents an experimental challenge. The decay jets are collimated in a small detector region. For large boost, the jets cannot be resolved. An alternative was proposed by Seymour in Reference [98], which considers to reconstruct a single massive jet ("fat" jet) with all the boosted object energy. In practice, this alternative offers more advantages when transverse momentum of boosted object exceeds its mass more than a factor of 3. Therefore, the boosted object regime starts at $p_T > 250 \text{ GeV}$ for W bosons and at $p_T > 500 \text{ GeV}$ for quarks tops.




Figure 6.1: Difference observed when there is (right) or not (left) boost in a heavy particle like the top quark. In the latter case, decay products are not easily separated and it can be complicated to identify each.

We can distinguish the "fat" jets that contain the decay of boosted object from jets initiated by light quarks or gluons (which are produced plentifully at LHC) with a jet substructure analysis [99]. New techniques for jet reconstruction (jet grooming) have been developed to evaluate jet substructure better. Several groups have defined packages to identify boosted objects (top taggers). Many different observables have been proposed. The most important variables for this thesis are described in subsection 6.1.2. A complete summary is found in Reference [24], [100].

Jet substructure has not received much attention in previous experiments [101, 102, 103, 104, 105, 106]. Therefore, there is a relative high uncertainty about the modelling of the development of the jet (parton shower) in the Monte Carlo (MC) tools used. Also, there is a significant uncertainty in the description of the detector response. The success of the study of boosted objects depends on the control of these theoretical and experimental uncertainties. For the commissioning of these techniques, ATLAS has studied the first high p_T jets registered by the experiment in 2010. The jet mass measurement and some substructure variables [107] are the main topic of this chapter.

6.1.1 Jet algorithms

There is no unique way to define a jet. Existing jet algorithms cluster the detected particles according to the transverse momentum and some distance metric. The algorithms used in ATLAS satisfy infrared and collinear-safe properties [108]. At each step, particles are clustered based on the distance between them and on the distance between the softest particle and the beam. Then, the distance metrics used in the most popular algorithms can be expressed:

$$d_{ij} = \min(k_{ti}^{2p}, k_{tj}^{2p}) \Delta R_{ij}^2 / R^2 \quad (6.1)$$

$$d_{iB} = k_{ti}^{2p} \quad (6.2)$$

$$\Delta R_{ij} = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2} \quad (6.3)$$

where k_{ti} represents the transverse momentum, ΔR_{ij} is the angular distance and R is a parameter that governs the maximal size of a jet. As long as the smallest distance d_{ij} is smaller than d_{iB} , particles are merged. When the minimum distance is between i and the beam, the jet is finished. The parameter p determines the relative importance of k_t and ΔR .

- If $p=0$, the particles are clustered without momentum dependence. This geometrical distance criterion yields the Cambridge-Aachen algorithm.
- If $p=1$, distance depends on the transverse momentum of the softest particle. The resulting algorithm is called k_t algorithm.
- If $p=-1$, the algorithm is called anti- k_t . It acts by iteratively merging the nearest objects in the event starting with the highest transverse momentum and finishing with the softest. This jet algorithm is default in ATLAS.

In Figure 6.2, the result of the three algorithms are shown for a simulated event. A few differences are obvious. Cambridge-Aachen and k_t jets have an irregular footprint while anti- k_t jets are nearly circular. In this study, jets are constructed using the jet algorithm anti- k_t which is implemented in the FASTJET package [109]. Contrary to LEP, at the LHC this jet reconstruction is inclusive that means the number of resulted jets is not defined. The

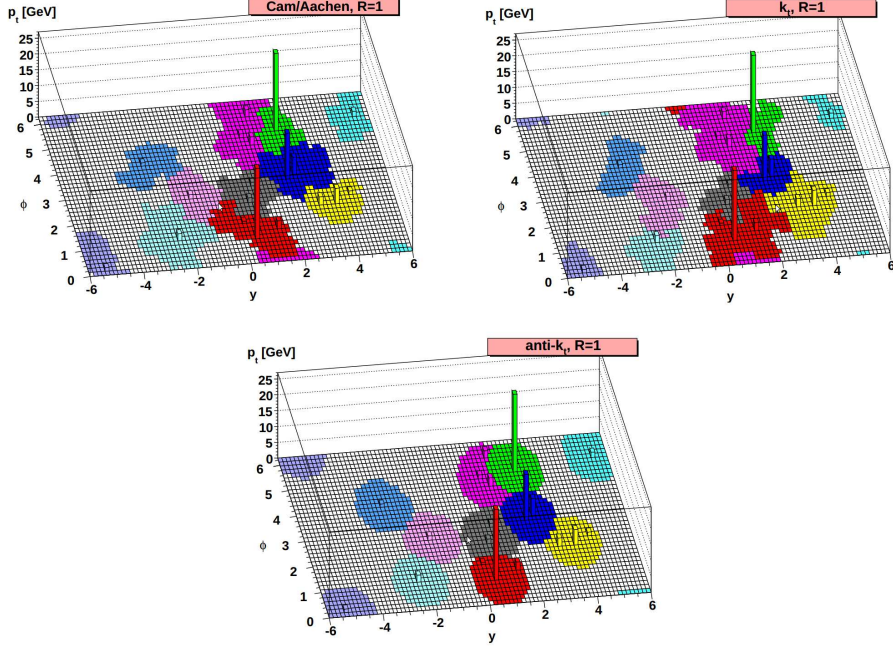


Figure 6.2: One ATLAS simulated event clustered with different jet algorithms [108].

parameter R is fixed to 1.0, a radius larger than the ATLAS standard radii of 0.4 and 0.6 to identify the fat jet and according to previous studies of heavy boosted objects [110].

6.1.2 Jet substructure variables

In the last years, a huge number of observables have been proposed. The most straightforward observable is the jet invariant mass. It is the sum over the mass-less four-vectors of the jet components:

$$m_{jet}^2 = (\sum E_i)^2 - (\sum \vec{P}_i)^2 \quad (6.4)$$

where E_i and P_i are the energy and momentum components of the jet four-vectors. For boosted objects, the jet mass reflects the mass of the boosted particle, while jets from light quarks and gluons typically have a smaller mass. The jet mass is proportional to $m \sim \alpha_s p_T R$ [99], i.e. a TeV jet with $R=1$ yields a 100 GeV jet mass.

The k_t splitting scales are defined by reclustering the constituents of the jet




Figure 6.3: Calorimeters η - ϕ displays of two collimated jets from a bottom quark without substructure (left) and from a boosted top quark (right). These images belong to a event from a $t\bar{t}$ resonances study of pp collisions at $\sqrt{s} = 8$ TeV and integrated luminosity of 14 fb^{-1} . [114].

with the k_t recombination algorithm [111, 112]. The k_t -distance of the final clustering step can be used to define a splitting scale variable $\sqrt{d_{12}}$:

$$\sqrt{d_{12}} = \min(p_{Tj1}, p_{Tj2}) \times \delta R_{j1,j2} \quad (6.5)$$

where 1 and 2 are the two jets before the final clustering step [113]. In a jet containing a boosted object, the splitting scale reflects the mass of the decaying objects. For jets initiated by light quarks or gluons, the splitting is generally asymmetric. The splitting scale therefore tends to yield smaller values. In the case of boosted top quarks, the $\sqrt{d_{12}}$ value is roughly of order m_W meanwhile the background jets have a value $\leq 20 \text{ GeV}$. The differences between heavy boosted particles and QCD jets can be visualised in Figure 6.3.

The variable $\sqrt{d_{23}}$ is defined in a similar manner, it refers to the next-to-last clustering step. Again, top jets present larger values than background.

6.2 Dataset and reconstruction

The data corresponds to pp collisions with center-of-mass energy of $\sqrt{s} = 7$ TeV. Data were registered in ATLAS in 2010 when the detector conditions and LHC beam were stable, the luminosity was correctly monitored and trigger was working. The integrated luminosity is $35.0 \pm 1.1 \text{ pb}^{-1}$ [115].

The events were selected first by L1 Calorimeter Trigger system efficiently and without any significant bias. To discard events with many detector noise

or non-collision backgrounds, a primary vertex should be coherent with the LHC beamspot and it should be reconstructed from 5 tracks of $p_T > 150$ MeV. The 3% of the events are discarded with this selection.

Also, jet mass and substructure variables are affected by additional pp collisions, so-called pile-up [24]. To suppress this effect, the number of primary vertices is fixed to one ($N_{PV} = 1$). This requirement reduces the events to 22%. This process is efficiency to reject pile-up and does not add extra systematic uncertainties (this part is comment in Section 6.5).

A 3-dimensional topological algorithm is used to cluster calorimeter cells [116]. Clusters are classified in hadronic or electromagnetic based in shape, depth and energy density. Their energy is corrected by calibration constants which depend on this cluster classification. Thus calorimeter non-compensation effects are partially recovered.

A further calibration is applied once jets are reconstructed. Jet mass, energy and η , are corrected using constants based on PYTHIA [117]) MC. Corrections for the three variables are defined with the matched pairs of hadron-level and reconstructed jets which depend on the energy and η . This jet-level correction is around 10–20% for mass and energy, and 0.01 for η .

The jets from tracks are used to calculate systematic uncertainties. To get these jets (track-jets), the same algorithm is applied that is also used for the calorimeter jets. The tracks from the inner detector are introduced into the algorithm if they have $p_T > 500$ MeV, $|\eta| < 2.5$, $|z_0| < 5$ mm and $|d_0| < 1.5$ mm, where z_0 and d_0 are the longitudinal and transversal impact parameter of the track. The jets are with $|y| < 2$ in four 100 GeV p_T bins from 200 to 600 GeV. This is not biased by the trigger effect and the selected jets are contained entirely into the barrels and endcaps of the sub-detectors.

6.3 Monte Carlo samples

The samples with jet inclusive events are produced with different MC generators. PYTHIA 6.423 [117] and HERWIG++ 2.4 [118] were used to calculate matrix elements in leading-order (LO) perturbative QCD (pQCD) for $2 \rightarrow 2$ processes. ALPGEN 2.13 [119] and SHERPA 1.2.3 [120] for some cross-checks implemented in $2 \rightarrow n$ processes like QCD multijets productions.

Parton shower is calculated in leading-logarithm approximation. If PYTHIA generator is used, showers are p_T ordered and in the case of HERWIG++ is angular ordered. Fragmentation into particles, PYTHIA is based in string model [121] and HERWIG++ in cluster model [122]. ALPGEN is interfaced with HERWIG [123, 124] for parton shower and fragmentation, and JIMMY [125] for underlying event model.

PYTHIA use the AMBT1 tune [126] if other options are not defined. The parton density functions at leading-order are extracted from the MRST2007 LO* set [127, 128]. Pile-up is not included in the samples.

All the MC samples have been made by ATLAS full simulation with GEANT4 [72] using the distributed production described in Chapter 4.

6.4 Detector-level distributions

Distributions are shown in Figures 6.4 and 6.5 for jet p_T , mass, $\sqrt{d_{12}}$ and $\sqrt{d_{23}}$. The statistical uncertainty presented in the ratios at the bottom of the plots is from MC and the added data in quadrature. The substructure variables are shown only for p_T interval of 300–400 GeVs. The MC is normalized in relation to data separately in each plot. It is observed that the jet substructure is adequately represented with the MC leading-order parton shower mode.

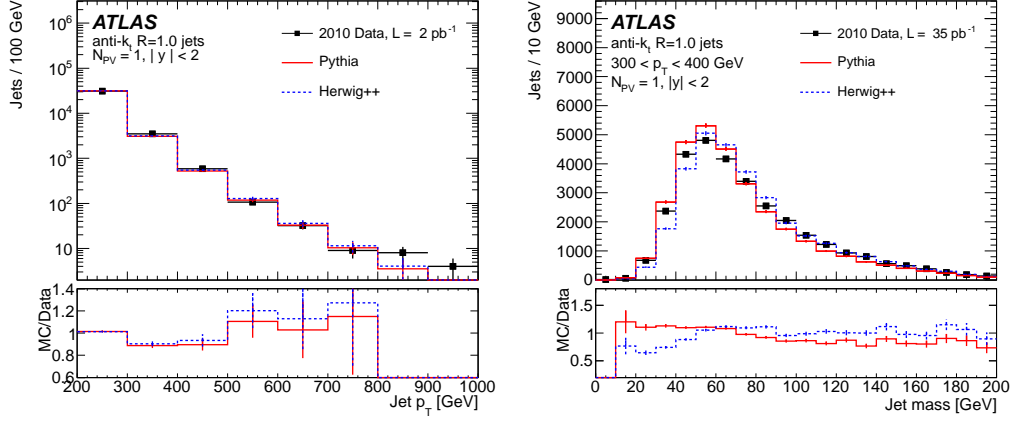


Figure 6.4: p_T distribution of anti- k_t $R = 1.0$ jets with $p_T > 200$ GeV (left) and mass distributions for anti- k_t jets with $|y| < 2.0$ in the 300–400 GeV p_T bin.




Figure 6.5: Distributions for $\sqrt{d_{12}}$ (left) and $\sqrt{d_{23}}$ (right) of anti- k_t $R = 1.0$ jets with $|y| < 2.0$ in the 300–400 GeV p_T bin.

6.5 Pile-up dependence

The data analysed contain only events with single pp interactions. The impact of multiple simultaneous pp interactions that produce extra events is evaluated in this section[129]. Pile-up adds a background of soft diffuse radiation that affects to the energy of the jets measured. Some substructure observables are expected to be especially sensitive to pile-up [24], particularly for the invariant mass of large-size jets.

The correlation of the mean jet mass of anti- k_t jets with $R=0.4$, 0.6 and 1.0 with the number of reconstructed primary vertices is shown in Figure 6.6. All jets with $p_T > 300$ GeV and rapidity in the range $|y| < 2$ were chosen. The jet mass is found to increase linearly with the number of pile-up vertices. This progression is stronger for long R , as expected.




Figure 6.6: The mean mass for jets with $p_T > 300$ GeV as a function of the number of primary vertices identified in the event. Comparisons show the effect for anti- k_t jets with different R -parameters.

6.6 Systematic uncertainties

The most important systematic uncertainty belongs to the modelling of the calorimeter response. Because of this, the jets calibration should be validated with MC. Due to the use of a big R -parameter of 1.0 in the anti- k_t algorithm, the ATLAS jets energy scales uncertainties cannot be implemented which are prepared for $R=0.4$ and 0.6 . The main systematic effects stem from the uncertainty in the jet p_T scale (JES) and jet p_T resolution (JER). Also, there are uncertainties per each substructure variable such as jet mass scale (JMS) and jet mass resolution (JMR).

The scale uncertainties are validated using track-jets. Uncertainties of the inner detector and the calorimeter are non correlated. Comparing the measurements obtained in each sub-detector allows to separate their physics effects. There is a limited precision of 3–5% in this method because of the efficiency of the inner-detector tracking and the MC modelling of the charged and neutral particles of the jets. Track-jets are matched with calorimeter-jets for $\delta R < 0.3$. The ratios are defined:

$$r^X = \frac{X_{\text{calorimeter-jet}}}{X_{\text{track-jet}}} \quad (6.6)$$

and the double ratio:

$$\rho^X = \frac{r_{\text{data}}^X}{r_{\text{MC}}^X} \quad (6.7)$$

where x can be p_T , mass or a substructure variable. The double ratio ρ^X forms a data-driven constraint on the scale. Deviations in the double ratio indicate that either the tracker or the calorimeter response is mismodelled.

Finally, the estimated uncertainty on the inner-detector measurement with the observed derivation in the double ratios is added in quadrature. The scale uncertainties obtained are around 3–6% and dominate the systematic uncertainties.

Another cross-check done was using MC test to determine how the detector responds in function of several variables. The samples used are different from the generator type, and also from various hadron physics models and detector geometry. The test results are similar in order of magnitude to the in-situ studies. The track-jet in-situ study is constricted because of the inner detector acceptance and only ranges over $|\eta| < 1.0$, which is the $\simeq 75\%$ of jets in the presented distributions. Despite this, the tests also show the η dependence is weak for every studied mismodelling. Therefore, this systematic uncertainty has been included.

The test of JER for anti- k_t in $R=0.4$ and $R=0.6$ [130] indicates jet p_T resolution by simulation is in agreement with data. The resolution uncertainties have been taken from this test because the resolution of the mass and substructure variables are difficult to validate in-situ with the data used. It was observed the resolution uncertainties are smaller than 20%.

6.7 Data correction

It is necessary to correct the effect of the detector resolution and acceptance to compare the results with the theoretical predictions. To do that, a matrix unfolding technique is used called Iterative Dynamically Stabilised (IDS) unfolding [131, 132].

It consists in a matching of the truth and the reconstructed jets of the MC samples if $\delta R < 0.2$. A match of $> 99\%$ of the jets is achieved. With the resulted pairs, a transfer matrix is obtained which is used to correct the effects of the detector. It can be possible a formed pair has one of its components that fails in the p_T cut. For that reason, each matching is done per each p_T bin separately. Data are scaled by the reconstructed matching efficiency, multiplied by the transfer matrix and divided by the truth matching efficiency. In an iterative process, the MC spectra are normalized by a regularization function.

6.8 Results

Finally, the normalized cross sections are presented in Figures 6.7, 6.8 and 6.9. In the ratio part of the plots, the data statistical uncertainty is not included. PYTHIA is found to provide a good description of the data, even if it tends to yield slightly too soft distributions. The opposite behaviour is for HERWIG++ distributions. The splitting scales show better agreement with MC predictions than mass. For high p_T bins, the measurements precision are limited by statistical fluctuations. However the level of agreement is roughly constant between bins in every variable.

Correlations between bins are introduced since the unfolding method was applied. Because the statistical uncertainty represents the diagonal element of the covariance matrix only, the full covariance matrices are needed to compare with alternative predictions. Then, the matrices and the results can be consulted in HEPDATA [133].

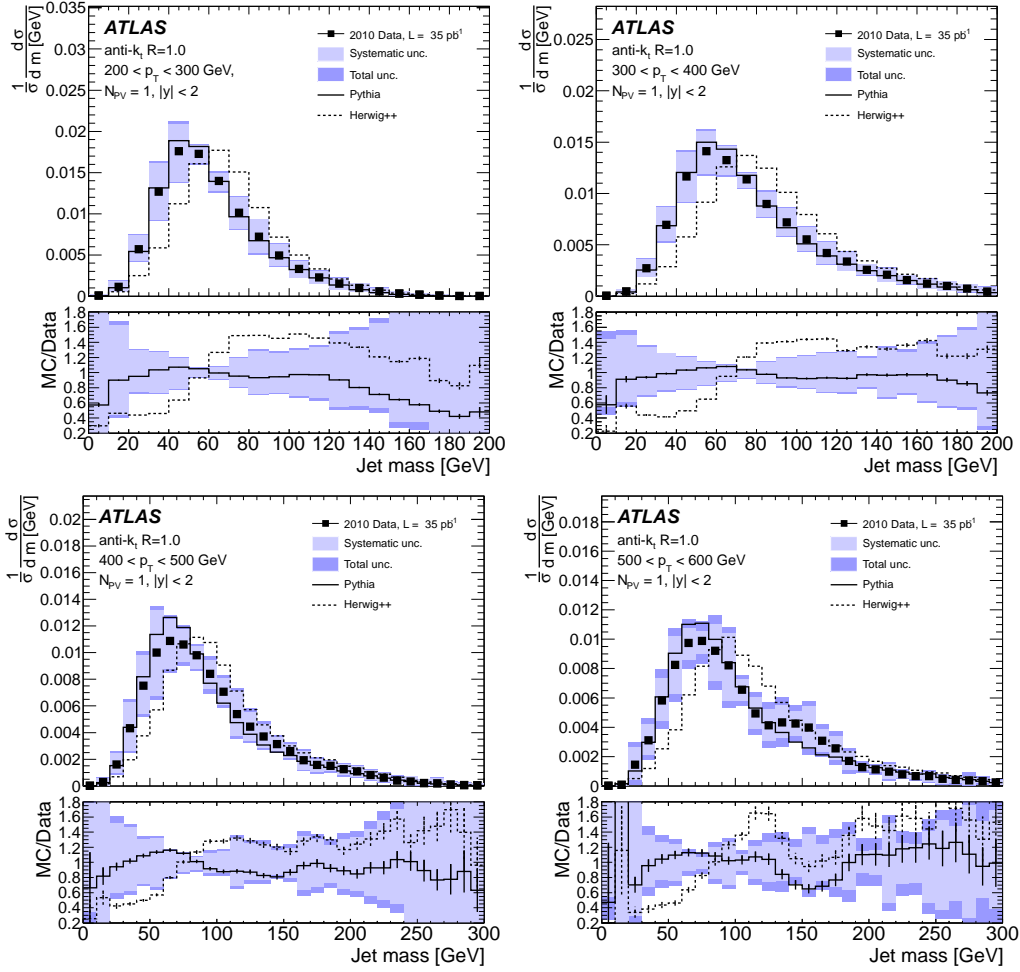


Figure 6.7: Normalised cross-sections as a function of mass of anti- k_t jets with $R = 1.0$ in four different p_T bins.

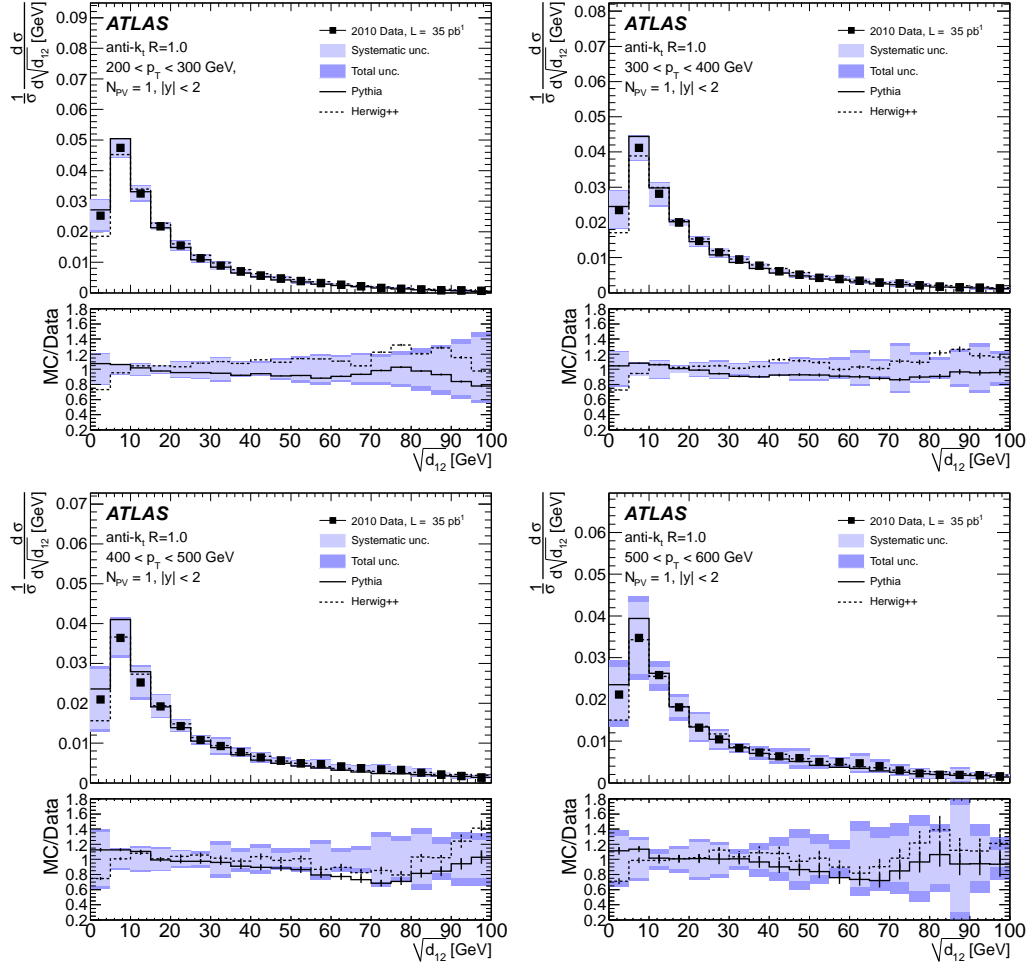


Figure 6.8: Normalised cross-sections as a function of $\sqrt{d_{12}}$ of anti- k_t jets with $R = 1.0$ in four different p_T bins.

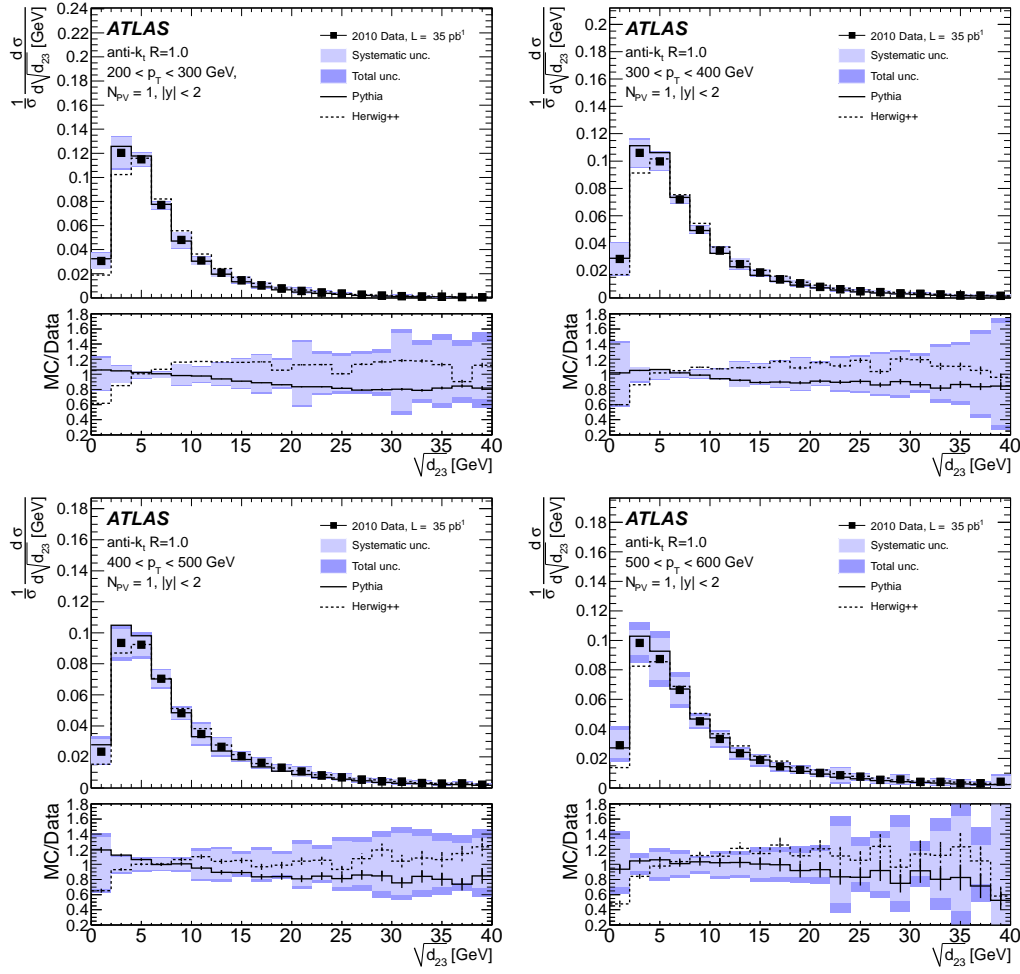


Figure 6.9: Normalised cross-sections as a function of $\sqrt{d_{23}}$ of anti- k_t jets with $R = 1.0$ in four different p_T bins.

6.9 Conclusions and perspectives

Jet substructure variables have been measured: jet mass and splitting scales. This was the first measurement of these variables at the LHC [107].

Data are in agreement with MC predictions from PYTHIA. The variable with the worst agreement with MC is jet mass, and the substructure variables $\sqrt{d_{12}}$ and $\sqrt{d_{23}}$ have smaller systematic uncertainties. HERWIG++ does not reproduce the jet mass distribution but it is adapted for splitting scales. These measurements show that the modelling parton-shower in common MC can reproduce well these variables.

The pile-up dependence of substructure variables have been studied, confirming the sensitivity of jet substructure to the effect of multiple proton-proton interaction, $N_{pv} = 1$ have been tested to reduce it. Elsewhere, methods to mitigate this dependence has been developed and used with some success [134].

Fully corrected distributions are available for comparison for physics analysis including the unfolding of detector effects. In fact, several groups are making an effort in obtaining a theoretical expression for the jet mass [135, 136, 137, 138].

Since the publication of this result, the first resonance searches geared to boosted top quarks have been performed by ATLAS that rely on jet substructure [139]. These first results demonstrate that novel techniques can enhance the sensitivity to BSM physics considerably.

Chapter 7

Summary/Resumen

7.1 Summary in English

7.1.1 Motivation

The first study of jet substructure on LHC data was performed by the ATLAS experiment. The jet algorithm chosen was AntiKt with R-parameter=1.0. This study has been important to check the working of the substructure variables which allow to distinguish boosted objects from background. These objects are produced abundantly at LHC and are expected to gain relevance in future operation at large energy.

In this study, the computing part has had a great importance because the work done into the ATLAS Spanish Tier-2 federation on understanding its performance and its operations. This has allowed the access of hundred of million of events to obtain the results using Grid technologies for Distributed Analysis. Also, this activity helped in other physics studies of ATLAS experiment.

7.1.2 Introduction

This work is framed within the context of the Standard Model (SM). The SM explains the interactions between the smallest constituents of matter [1]. The SM classifies the constituents into fermions (half-integer spin) and bosons (integer spin). In Table 1.1 elementary particles are listed. Fermions can be divided in leptons and quarks grouped in three generations which are different from mass. Quarks have fractional electric charge and also color charge.

The SM describes the strong, weak and electromagnetic interactions. The last two are unified by the electroweak formalism whose bosons are photon, W^{+-} and Z^0 . This model predicts all bosons are massless which is in disagreement with experiments. To solve that, the Higgs mechanism was formulated adding an extra boson. The strong interaction takes over the quarks confinement by the color charge and the exchanged particles are gluons.

The Standard Model is in agreement with the most of the experiments results. The last recent discovery of the Higgs boson [14] puts the SM on an even more solid empirical basis. However, there is some evidence of effects that the Standard Model cannot explain such as the neutrino mass and the presence of dark matter [13]. Extensions have been proposed among which the most popular are Supersymmetry [15] and Extra Dimensions [18] which incorporates more spatial dimensions. They are grouped under the expression “Beyond Standard Model” (BSM).

The Large Hadron Collider (LHC) [19] is located at CERN [11], with a circumference of 27 kilometres, at 100 metres of depth. It can reach a center-of-mass energy of 14 TeV, at a luminosity of $10^{34} cm^{-2} s^{-1}$. The LHC has four impact points where experiments are situated to detect collisions. A schema of these points is presented in Figure 1.5.

One of these experiments is ATLAS (A Large Toroidal LHC Apparatus) [140], the detector which has produced the data studied in this work. It is a general purpose detector, i.e., it covers many different physics objectives. Its shape is cylindrical with a diameter of 25 metres and 46 metres of length. Its weight reaches 7000 tones. A schema of the ATLAS detector is shown in Figure 2.1.

The ATLAS subdetectors are distributed in concentric layers around the beam. The deepest subdetector can detect the particle tracks, the so-called Inner Detector. The following layer is the calorimeter that measures the energy of the particles and allows to identify them. The first component of calorimetry is the electromagnetic calorimeter (Liquid Argon) which detects energy from electric charged particles and photons. It is followed by the hadronic calorimeter (TileCal) that measures energy of hadrons. The most external layer is formed by the muon chambers. They measure the momentum of the muons and allow for a better identification. There are many events without interest produced in ATLAS. To filter these events, there is a three-level trigger system which is shown in Figure 2.8.

The LHC experiments have strong challenges in computing aspects: a huge storage to save 25 PB per year, the processing of 100,000 CPUs and interna-

tional access for 8000 physicists. The data of all LHC experiments are processed and analysed in a worldwide distributed computing environment based on Grid computing concept [26]. Grid computing consists of sharing resources by centres spread around the world. In the case of LHC, a Grid guideline has been established for the four experiments in the Worldwide LHC Computing Grid (WLCG) [27] project. All LHC experiments have chosen a hierarchical architecture that includes several “Tier” centres. A Tier centre is a cluster of computers capable of running the totality of the experiment-specific data processing software. In addition, the centre has to be Grid-enabled, which means that the general Grid software must be installed and maintained so that each centre can act as a component of a unique worldwide computing network. The adapted tiered hierarchy for computing model at the LHC is: one Tier-0 centre (CERN), 10 Tier-1 centres (National centres), around 80 Tier-2 centres (Regional centres) and several Tier-3 centres (institution level resources).

In ATLAS, other guidelines have been added according to the requirements of the experiment since the data-taking. The MONARC model changed to ‘Mesh’ where transfers do not need to pass through intermediate centres among Tier-2s that are certificated as Tier-2 Direct (T2D) [50] due to its excellent connection and availability. Specific software has been developed by the ATLAS collaboration such as Athena [51] for all the activities of the experiment (calibration, alignment, simulation and reprocessing). Another software example is Database [55] which gives all the detector information in geometry and other parameters. All Grid computing resources have been tested successfully before data-taking in November 2009 by WLCG to be ready and functional [60].

7.1.3 Distributed Analysis framework

The ATLAS analysis is distributed because the data are distributed around the world, and PanDA (Production and Distributed Analysis) [61] system was developed to control all processes efficiently. This system achieves the compatibility of all sites with the ATLAS requirements. Some tools have been designed to make easier the access of Distributed Analysis for the ATLAS users such as Don Quixote 2 (DQ2) [65] and ATLAS Metadata Interface (AMI) [66] to manage data, and PanDA-Client [69] and Ganga [70] to send Grid jobs. The latter has been used for this study.

The ATLAS collaboration has also done innovations to improve the Distributed Analysis such as implementing analysis and production tests called

HammerCloud [75] which excludes inefficient sites. Another example is to do automatic official data replicas to other sites in function of the user demand, the so-called PanDA Dynamic Data Placement (P2DP) [76]. In addition, there is a big task which consists of attending doubts and problems from ATLAS users in Distributed Analysis by the Distributed Analysis Support Team (DAST) [77], where I participate since 2008.

An example of the operations of the Distributed Analysis System in ATLAS has been shown with the study of Jet Substructure. Ganga has been used to send the Athena package needed and to choose the accurate tool to find the input data for analysis. If we were processing simulated data, then DQ2 was used but in the case of real data, AMI was the tool since it can read the list with the best data quality called Good Run List [78].

When the input data information was obtained, job was split into subjobs according to the number of input files. Ganga looked for the available sites with the input data. It also gave the instructions to request a replica of the output data to another centre which was the IFIC Tier-3 centre [79] in this case. The workflow is presented in Figure 4.7.

These Grid jobs were distributed in 36 sites. In a period of time of 31 hours, 15 million of simulated events were analysed. In the case of real data, 158 million of events, which were recollected in 2010 with a center-of-mass energy of 7 TeV, were processed in 36 hours .

7.1.4 ATLAS Spanish Tier-2 Federation

To illustrate the operations of a Tier-2 centre, we focused on the ATLAS Spanish Tier-2 [80]. It is a federation of three sites: IFIC (Valencia), IFAE (Barcelona) and UAM (Madrid). It contributes the 5% of all the ATLAS Tier-2 centres computing resources. Its availability, data transfers, storage and processing were observed to check its operations. The Spanish Tier-2 functions changed due to the computing model evolution. Because it is considered as a T2D, it has sent and received transfers which did not need to pass through its assigned Tier-1 centre (PIC). Other changes were noticed such as the possibility to make calibration tasks because the access of all the Database through the Frontier/Squid [54].

According to the Service Availability Monitoring (SAM) [89] test for three years, the reliability and availability are between 90% and 100%. The specific ATLAS availability average was 83% for Hammercloud analysis jobs and 92% for production from September 2011 to April 2013. In Figure 5.8 the transfers

sources were shown for three years. Transfers were 14 million as receptor and the same as emitter. The data volume transferred were 9 PB and 3.5 PB respectively. The total storage occupancy reached the 2 PB. Observing the occupancy evolution of the space tokens [94], the different uses of the ATLAS activities were reflected as in Figures 5.31 –5.37. Production and analysis are the most important activities in an ATLAS Tier-2 centre. 4.7 million of production jobs arrived at the Spanish Tier-2 in three years with a high efficiency (93%) which contributes the 1.78% of all ATLAS. In the analysis activity, 7.2 million of jobs run, contributing the 1.76% with respect to all ATLAS, with the efficiency of 87.9%.

The Spanish Tier-2 federation is a reality and has a high efficiency thanks to a group of experts who has participated in the installation, maintenance and operations. Some of them have taken part in a team to help the physicist community from the three Spanish centres. They are called User Support and attend the topics about Distributed Analysis and Grid computing in ATLAS, where I am a member in the IFIC centre.

7.1.5 Jet substructure

Finally, the first study of jet substructure was performed by ATLAS experiment [107]. Substructure analysis is a crucial tool to identify boosted objects that may prove to be a powerful signature for BSM Physics. The jet algorithm used in this measurement is anti- k_t [108] with R-parameter value 1.0. The jet substructure that were measured are jet mass and splitting scales. The latter consists of the scale of the k_t jet algorithm from two to one $\sqrt{d_{12}}$ and from three to two components $\sqrt{d_{23}}$.

QCD simulated samples from different Monte Carlo generators and data from ATLAS detector taken in 2010 with a center-of-mass energy of 7 TeV and integrated luminosity of $35.0 \pm 1.1 \text{ pb}^{-1}$ were analysed. Substructure variables are affected by the pile-up collisions for jet mass. The election of the first ATLAS data and selecting events with a single vertex allowed to suppress the pile-up effect.

The calibration used was obtained by Monte Carlo (PYTHIA generator [117]) comparing the variables values at hadron-level with the values from the reconstructed jets which depend on the mass, the energy and the pseudorapidity η .

The most important contributions in the systematic uncertainties were on scale and resolution of the jet transverse momentum. In addition, there were

resolution and scale systematic uncertainties for each substructure variable. To validate the scale uncertainties, track-jets were used. They were compared with calorimeter jets. To calculate the resolution systematic uncertainties several Monte Carlo models were used. These models were different in generators, physics models to form hadrons or detector geometries.

The effect in resolution and acceptance of the ATLAS detector was corrected to compare results with theoretical predictions. The unfolding matrix technique used is called Iterative Dynamically Stabilised unfolding [131, 132]. It consist of the creation of a transfer matrix matching the simulated data at hadron-level and the detector-level data. Each matching was done per each p_T bin separately. Data were scaled by the reconstructed matching efficiency, multiplied by the transfer matrix and divided by the truth matching efficiency.

The normalized cross section distributions in function of the substructure variables were resulted. They were shown in Figures 6.7, 6.8 and 6.9. The generator PYTHIA was in good agreement with data. The opposite behaviour happens with HERWIG++ generator [118]. The splitting scales were the best substructure variables which agree with data.

7.1.6 Conclusions

The ATLAS computing model has evolved from hierarchical to Mesh model which has affected successfully to the T2Ds because this increases the running jobs because more data is received. The final target is to interconnect all kind of Tiers until to achieve that all work as one Cloud.

The ATLAS Distributed Analysis is being done successfully since the tools developed by the collaboration to make easier the Grid use. The tools to manage data and to send jobs make user-friendly the Grid for ATLAS physicists like Ganga which was used for this thesis. During the data-taking, the analysis jobs have being increased. Meanwhile the Distributed Analysis have being improved applying new ideas. An ATLAS user support is being given by DAST which attend problems and clarifies doubts about Distributed Analysis.

A Distributed Analysis example was shown which represents the data getting of the jet substructure study. A total of 11373 subjobs were sent (3205 for input simulated data and 8168 for real data). They were distributed in 36 sites and have lasted 31 hours for 15 million of simulated events and 36 hours for 158 million of real data events.

The ATLAS Spanish Tier-2 was presented. Its contribution of 5% has been in a good agreement with the operations like availability (between 90 and 100%), transfers (14 million of transfers), storage (over 2PB of occupancy) and processing (4.7 million of production jobs and 7.2 million of analysis jobs). The observation of these parameters for three years showed how the computing model evolution affected satisfactorily a Tier-2 centre and also the different uses of the technology by each ATLAS activity. Another computing operation was presented which is the User support to help the users from the three Spanish sites.

Jet substructure variables were measured: jet mass and splitting scales [107]. The Monte Carlo predictions with PYTHIA was in agreement with data meanwhile HERWIG++ did not reproduce well the jet mass distribution but it was adapted for the splitting scales. The splitting scales had also smaller systematic uncertainties.

A pile-up dependence was observed which was reduced fixing the number of primary vertices to one. Elsewhere, methods have been developed to decrease this dependence with some of success [134]. The distributions were corrected to compare them with theoretical models. In fact, some studies have tried to achieve a theoretical expression of jet mass [135, 136, 137, 138].

Since the publication of these results, top-antitop resonances studies have used this study to search boosted top quarks [139]. These techniques increase the sensitivity of new physics searches.

7.2 Resumen en español

7.2.1 Motivación

En el Large Hadron Collider (LHC) se producen partículas consideradas objetos ‘boosted’ donde sus productos de desintegración se concentran en una pequeña parte del detector. El estudio de estos objetos pueden dar pistas de nueva física como en dimensiones extra y supersimetría, así como el estudio del bosón de Higgs. Con el fin de caracterizar estos objetos se han introducido variables de subestructura de jets. Este trabajo recoge el primer estudio de subestructura de jets con datos reales usando el algoritmo *antiK_t* con parámetro $R=1.0$. Todo ello se ha realizado en el marco del Modelo Estándar y a partir de los sucesos reales del experimento ATLAS.

El análisis de datos de ATLAS se realiza utilizando las tecnologías Grid y el análisis distribuido. En este estudio de subestructura de jets ha tenido

una gran importancia, ya que el trabajo realizado dentro del Tier-2 federado español en el entendimiento de su funcionamiento y la realización de algunas de sus operaciones ha permitido el acceso a cientos de millones de eventos para obtener los resultados mostrados. También ha permitido ayudar en otros estudios de física del experimento ATLAS donde se requería tal experiencia.

7.2.2 Introducción

El marco teórico en el que se enmarca este trabajo es el Modelo Estándar, que es la interpretación actual del comportamiento de las partículas elementales [1]. Separa las partículas en fermiones (spin semientero) encargadas de constituir la materia, y en bosones (spin entero) que son intercambiadas por los propios fermiones para interactuar.

Los tipos de interacción son fuerte, débil y electromagnética. Éstas dos últimas se unen mediante el formalismo electrodébil cuyos bosones responsables son el fotón, W^{+-} y Z^0 . Este formalismo pronosticaba que todos los bosones no tenían masa lo que no concordaba con los resultados de los experimentos así que se formuló el mecanismo de Higgs que aportaba otro bosón extra.

Este modelo está en completo acuerdo con los resultados de los experimentos de aceleradores de partículas, como por ejemplo el último descubrimiento del bosón de Higgs [14]. Sin embargo, siguen varias discrepancias sin resolver como que no se incluye la gravedad, ni se considera que los neutrinos tengan masa o que no hayan candidatos para la materia oscura [13]. Para incluirlas, otros modelos teóricos se han desarrollado denominados en general 'Más allá del modelo Estándar'. Los más populares son Supersimetría [15], que propone partículas nuevas llamadas 'supersimétricas', y Dimensiones Extra [18], que incorporan más dimensiones espaciales a las tres ya conocidas.

Estos retos que surgen de una física nueva han sido la razón de la contrucción del acelerador de protones Large Hadron Collider [19]. Está localizado en el CERN [11], tiene una circunferencia de 27 kilómetros y situado a 100 metros de profundidad. Puede alcanzar una energía en el centro de masas de 14 TeV, a una luminosidad de $10^{34} cm^{-2} s^{-1}$. Tiene cuatro puntos de impacto donde hay experimentos situados para detectar las colisiones. Un esquema de esos puntos se encuentra en la Figura 1.5.

Uno de esos experimentos es ATLAS (A Large Toroidal LHC Apparatus) [140], cuyos datos han sido analizados para este estudio. Su propósito es de carácter general, por tanto cubre una gran variedad de objetivos de la física.

Tiene forma cilíndrica con 25 metros de diámetro y 46 metros de longitud con un peso de 7000 toneladas. Un esquema del detector se encuentra en la Figura 2.1.

Sus componentes o subdetectores se distribuyen de forma concéntrica alrededor del haz de protones. El primer subdetector más interior es el encargado de detectar las trazas, también llamado ‘Inner Detector’. El siguiente subdetector es el calorímetro que mide la energía de las partículas, lo que permite identificarlas. El primer componente es el calorímetro electromagnético (‘Liquid Argon’), que detecta la energía de las partículas cargadas y de los fotones. Después se sitúa el calorímetro hadrónico (‘TileCal’), que detecta la energía de los hadrones. La capa más exterior de subdetectores la forma las cámaras de muones que miden el momento de estas partículas y permiten una mejor identificación gracias a estar rodeados de un campo magnético toroidal. En ATLAS se producen muchos eventos sin interés para la física actual. Para filtrar los eventos no deseados y guardar aquellos eventos que nos aportan información se usa un sistema de ‘trigger’ a tres niveles como se puede ver en la Figura 2.8.

El LHC y sus experimentos está suponiendo un desafío computacional como es la necesidad de una enorme capacidad para almacenar 25 PB por año, de gran procesamiento entorno a 100,000 CPUs y de que haya acceso internacional a los datos distribuidos por todo el mundo para 8000 físicos. Los datos procedentes de todos los experimentos del LHC son procesados y analizados por medio de una computación distribuida basada en el modelo Grid [26]. El Grid consiste en compartir recursos computacionales por todos los centros repartidos por todo el mundo que trabajan en el experimento. Para el caso del LHC, se establecieron unos parámetros comunes mediante el proyecto Worldwide LHC Computing Grid (WLCG)[27]. Los experimentos del LHC han establecido una arquitectura jerárquica basada en ‘Tiers’. Un Tier es un centro computacional capaz de ejecutar la totalidad del software específico para procesar los datos del experimento. Además, el centro ha de ser accesible en el Grid, lo que implica que el software del Grid (middleware) ha de ser instalado y mantenido de tal forma que cada centro pueda actuar como un componente de una red única y mundial. El modelo computacional jerarquizado en el LHC es [44]: un Tier-0 (CERN), 10 Tier-1s (centros nacionales), unos 80 Tier-2s (centros regionales) y varios Tier-3s (recursos a nivel institucional).

Dentro de ATLAS, se han ido añadiendo otros requisitos para la computación donde se han adaptado las exigencias del experimento desde que comenzó la

toma de datos. Se cambió del modelo jerárquico al modelo ‘Mesh’ (de rejilla) donde las transferencias no necesitan pasar por centros intermedios al calificarse centros de destino como Tier-2 directos (T2D) [50] por su excelente conexión y disponibilidad. Se ha desarrollado software específico por la colaboración ATLAS como por ejemplo Athena [51] para el análisis y procesamiento de los datos. Otro ejemplo de software es la base de datos (‘Database’ [55]) que actualiza la información del detector ATLAS en geometría y otros parámetros. Con el fin de tener todo preparado y operativo para la toma de datos, se procedió a la verificación (test) de todo el modelo Grid obteniéndose buenos resultados dentro del contexto general del proyecto WLCG [60].

7.2.3 Análisis distribuido

El análisis en ATLAS se hace de manera distribuida dentro del Grid. Se desarrolló un sistema para controlar con eficacia todos los procesos relacionados con los trabajos Grid para ATLAS, este sistema se llamó PanDA [61]. Consigue la compatibilidad de todas las granjas computacionales con los requisitos de ATLAS. La colaboración ATLAS ha creado herramientas para facilitar el acceso al análisis distribuido a los usuarios, como son Don Quixote 2 (DQ2) [65] y ATLAS Metadata Interface (AMI) [66] para la gestión de datos, y PanDA-Client [69] y Ganga [70] para el envío de los trabajos al Grid. Concretamente, hemos utilizado Ganga para el estudio de subestructura de jets.

La colaboración ATLAS, para mejorar el análisis distribuido, ha implementado una serie de tests basados en modelos de análisis y producción llamados HammerCloud [75], que excluyen aquellos centros de participar en esas actividades. Otro ejemplo de estas mejoras son las copias automáticas de datos oficiales a otras granjas según la demanda cuya actividad se denomina PanDA Dynamic Data Placement (P2DP) [76]. También hay que mencionar la gran tarea de atender las dudas y problemas de los usuarios en análisis distribuido por parte del Distributed Analysis Support Team (DAST) [77].

Un ejemplo de cómo se ha usado el análisis distribuido ha sido el realizado para el estudio de subestructura de jets. Ganga se ha usado para enviar el paquete de Athena necesario y elegir la herramienta adecuada para encontrar los datos de entrada que analizar. Si se escogía datos simulados como datos de entrada entonces se usaba DQ2 pero si se trataba de datos reales entonces se elegía AMI porque podía interpretar una lista de mejores eventos bajo criterios de calidad que se denomina Good Run List [78].

Una vez obtenida la información de los datos de entrada, se dividió el trabajo principal en sub-trabajos en base al número de ficheros. Buscó los centros de computación disponibles con los datos. Ganga también aportó instrucciones para que se pidiera una réplica automática de los datos de salida a otro centro, que en este caso se guardó en el Tier-3 del IFIC [79]. Se puede ver en la Figura 4.7 el flujo de datos de trabajo descrito.

Estos trabajos se han distribuido en 36 centros, y el tiempo empleado ha sido de 31 horas para 15 millones de eventos simulados y 36 horas para 158 millones de eventos de datos reales. Estos datos corresponden al año 2010 con energía del centro de masas de 7 TeV.

7.2.4 El Tier-2 español federado de ATLAS

Con el fin de ilustrar el funcionamiento de un centro Tier-2 de ATLAS, nos hemos centrado en el Tier-2 español [80]. Es una federación de tres centros: IFIC (Valencia), IFAE (Barcelona) y UAM (Madrid). Tiene una contribución del 5% respecto de todos los Tier-2 de ATLAS en recursos computacionales. Sus funciones han cambiado por pasar a otro modelo de computación donde ha recibido y enviado transferencias sin necesidad de ir a través del Tier-1 español, PIC. En la Figura 5.8 se puede observar cómo se incorporaron otros centros como origen de transferencias al Tier-2 español. También los cambios se han notado al poder hacer tareas de calibración, debido a un acceso más flexible de toda la base de datos mediante el Frontier/Squid [54] por parte de los Tier-2s.

Según los tests de Service Availability Monitoring (SAM) [89] en el periodo de tres años, los parámetros de fiabilidad y accesibilidad fueron entre 90–100%, y la accesibilidad específica media de ATLAS se ha evaluado en un 83% para análisis y 92% para producción en el periodo desde septiembre de 2011 hasta abril de 2013. Las transferencias han sido de 14 millones tanto las recibidas como las enviadas, donde se ha transferido un volumen de datos de 9 PB y 3.5 PB respectivamente. Se ha observado que la ocupación del espacio supera 2 PB. Mediante la observación de la evolución del espacio de los space tokens (espacios reservados) [94], se ha reflejado los diferentes usos por cada una de las actividades de ATLAS como se muestra en la Figura 5.37. La producción de datos simulados y el análisis son las actividades más importantes en un Tier-2 de ATLAS. 4.7 millones de trabajos de producción llegaron a este Tier-2 estos últimos tres años con una eficiencia alta de 93% que contribuye el 1.78% respecto de todos los trabajos en ATLAS. En la

actividad de análisis se ejecutaron 7.2 millones de trabajos, contribuyendo en un 1.76% respecto al total de ATLAS, con una eficiencia del 87.9%.

En la instalación, mantenimiento, operación y explotación del Tier-2 federado español han participado un grupo de expertos que lo han hecho posible y eficaz. Algunos de ellos han formado parte del equipo de ayuda para los físicos que realizan los análisis en los tres centros españoles. Este equipo es el denominado User Support en el que yo formo parte activa.

7.2.5 Subestructura de jets

Finalmente, destacar que se ha colaborado en el estudio de física pionero en ATLAS de la subestructura de jets [107], que está sirviendo para determinar objetos ‘boosted’ (empujados). Éstos objetos son de interés para la física más allá del Modelo Estándar dentro de ATLAS y LHC. Se caracterizan porque a momentos transversos mucho más grandes que su masa, los productos de desintegración se concentran en una pequeña parte del detector y no pueden diferenciarse [24]. Entonces se aplica la técnica de considerar un único jet que englobe todos los productos de desintegración, para luego estudiar su estructura interna y así poder identificar las partículas pesadas ‘boosted’ y descartar gluones y quarks ligeros.

El algoritmo de jet estudiado ha sido el denominado AntiKt [108], que es el usado por defecto en el detector, aunque el parámetro R se ha aumentado hasta 1.0 para poder englobar todos los productos de la desintegración. Se han escogido como variables de subestructura de jets las más fundamentales, que son la masa intrínseca del jet y las ‘splitting scales’ (escalas de división). Estas últimas se obtienen deshaciendo los últimos pasos del algoritmo Kt para pasar de un jet a dos partes ($\sqrt{d_{12}}$) y de dos a tres partes ($\sqrt{d_{23}}$).

Se han analizado muestras simuladas de procesos de interacción fuerte con diferentes generadores de Monte Carlo y datos reales procedentes del detector ATLAS obtenidas en el año 2010 con una energía en el centro de masas de 7 TeV y con una luminosidad de $35.0 \pm 1.1 \text{ pb}^{-1}$.

Las variables de subestructura se ven afectadas por las colisiones adicionales que se originan (‘pile-up’) como se muestra en la Figura 6.6 para masa de jets. Al estudiar los primeros datos y fijar que el número de vértices primarios sea uno, hemos conseguido evitar el efecto que el ‘pile-up’ hace en estas variables.

La calibración usada se ha obtenido mediante Monte Carlo (generador usado es PYTHIA [117]), comparando el valor a nivel hadrónico con el de los jets

reconstruidos de las variables de subestructura, los cuales dependen de la masa, la energía y la ‘pseudorapidity’ η .

Las contribuciones más importantes en los errores sistemáticos han sido en la escala y resolución del momento transversal del jet. Se han incluido además errores sistemáticos de la escala y resolución de la masa del jet y las ‘splitting scales’. Para validar los errores sistemáticos de escala se han usado los jets sacados del detector de trazas, que se han comparado con los jets del calorímetro usando el mismo algoritmo de jets. Para calcular los errores sistemáticos en la resolución se ha usado varios modelos de Monte Carlo con parámetros variados como el tipo de generador, el modelo teórico de la formación de hadrones e incluso cambiando la geometría del detector.

Se ha corregido el efecto de la resolución y la aceptación del detector, y así conseguir que pueda ser comparado con predicciones teóricas. Para ello se ha usado una técnica de desdoblamiento con matrices denominado Iterative Dynamically Stabilised unfolding [131, 132]. La técnica consiste en crear una matriz de conversión con datos simulados comparando los valores a nivel de generación y a nivel detector, de donde se obtienen las parejas de jets que coinciden. La comparación se hace para cada intervalo de momento transversal. Para obtener los datos corregidos se ha multiplicado la eficiencia de la coincidencia en la parte reconstruida luego se ha multiplicado por la matriz de conversión y dividido por la eficiencia de la coincidencia en la parte a nivel de generación.

Se han obtenido como resultados las distribuciones de la sección eficaz normalizada en función de las variables de subestructura como se muestran en la Figuras 6.7, 6.8 y 6.9. El generador PYTHIA ha descrito bien los datos reales, sin embargo, no ha pasado lo mismo con el generador HERWIG++ [118]. De entre las variables de subestructura, las ‘splitting scales’ han concordado mejor con los datos que la masa del jet.

7.2.6 Conclusiones

Este estudio ha mostrado los primeros resultados de subestructura de jets para el algoritmo AntiKt con $R=1.0$ en el experimento ATLAS y a través de un análisis distribuido mediante computación Grid. El modelo computacional de ATLAS ha evolucionado de un modelo MONARC a otro ‘Mesh’ ante las mejoras de conexión a internet que ha afectado positivamente a las operaciones de los Tier-2s. El objetivo final es interconectar todos los Tiers hasta que todo se comporte como una única Cloud.

El análisis distribuido en ATLAS se está realizando con éxito gracias a las herramientas desarrolladas por la colaboración, que facilitan el uso del Grid. Las herramientas para la gestión de datos y enviar los trabajos de análisis permiten un manejo más fácil del Grid para los físicos de ATLAS, como es el caso de Ganga que se ha usado para este trabajo.

A medida que se han tomado los datos en el experimento y ha aumentado el uso de las granjas que conforman la computación de ATLAS en la actividad de análisis, se ha ido mejorando el análisis distribuido aplicando nuevas ideas. También se está dando un servicio de apoyo a los usuarios de ATLAS por medio de DAST, que facilita la resolución de problemas y aclaración de dudas en el tema de análisis distribuido.

Se ha presentado un caso de análisis distribuido que representa el proceso utilizado para la obtención de los datos del estudio de subestructura de jets. Se han enviado un total de 11373 sub-trabajos (3205 procesaron datos simulados y 8168 datos reales). Se han distribuido a 36 granjas y ejecutado en un periodo de tiempo de 31 horas para 15 millones de eventos simulados y 36 horas para 158 millones de eventos de datos reales siendo sólo los pertenecientes al 2010 con energía del centro de masas de 7 TeV.

Hemos mostrado como ejemplo de Tier, el Tier-2 federado español de ATLAS. Su contribución del 5% satisface todos los requisitos de un Tier-2 en accesibilidad (entre 90 y 100%), transferencias (14 millones de transferencias realizadas), almacenamiento (ocupación de más de 2PB) y procesamiento (4.7 millones de trabajos en producción y 7.2 millones en análisis). Observando estas operaciones durante tres años se ha apreciado cómo ha influido el cambio de modelo y las propias actividades de ATLAS al Tier-2. Además el trabajo realizado ha dado lugar a el User Support que presta ayuda a usuarios de los tres centros en temas de análisis distribuido y computación Grid.

En cuanto a la parte de análisis de física: Las variables de subestructura de jets que se han medido son la masa del jet y las ‘splitting scales’. Han sido las primeras medidas en el LHC [107]. Las predicciones de Monte Carlo con generador PYTHIA han estado en concordancia con los datos del experimento. El generador HERWIG++ no ha reproducido bien la distribución de la masa del jet, pero se ha adaptado para las splitting scales. La variable de subestructura con peor concordancia con Monte Carlo ha sido la masa del jet. Las splitting scales han aportado errores sistemáticos pequeños.

Existe una dependencia de las variables de subestructura con el ‘pile-up’ y para reducir ese efecto hemos fijado en uno el número de vértices primarios.

En estudios posteriores se han desarrollado métodos que han disminuido en parte esa dependencia [134]. Se han corregido las distribuciones para que pueda ser comparada con modelos teóricos. De hecho, se pueden consultar estudios teóricos en los que se ha pretendido obtener la masa de los jets [135, 136, 137, 138].

Desde la publicación de este resultado, grupos de física que realizan estudios de resonancias top-antitop han usado este trabajo para buscar quarks top ‘boosted’ del alto momento [139]. Estas técnicas aumentan la sensibilidad en búsquedas de nueva física.

Contributions and Impact

This ATLAS physics analysis of Jet Substructure has been elaborated by several physics groups in different universities and institutes. IFIC was the institute that initiated the idea and was the responsibility for the edition. The main contribution belongs to the institutes IFIC, UCL (University College London) and SLAC National Accelerator Laboratory. Other contributions are from University of Oxford, DESY (Deutsches Elektronen-Synchrotron) and McGill university. Also, a fundamental part was played by the ES-ATLAS-T2 and the IFIC User Support for the Grid access to get the ATLAS data and the knowledge of the ATLAS Grid Tools for Distributed Analysis.

The results of this analysis have been presented in several international conferences, this is a sample of them:

- “Measurements of jet substructure in ATLAS”, Adam Davison, Boost2011, Princeton (USA), 22-26 May 2011.
- “Where we stand with sub jet analyses” Marcel Vos, Internationales Wissenschaftsforum Heidelberg LHC New Physics Forum, Heidelberg (Germany), December 19-21, 2011.
- “BSM physics with top quarks in the ATLAS experiment” Elena Oliver, Physics in Collision 2011, Vancouver (Canada), 28 Aug - 1 Sep 2011, ATL-PHYS-SLIDE-2011-604.
- “Operations of the ATLAS Spanish Tier-2 during the LHC Run I”. Elena Oliver et al, IBERGRID 2013, 7th Iberian Grid Infrastructure Conference. Madrid, Spain. 19-20 September 2013.
- “Evolution of the Atlas data and computing model for a Tier2 in the EGI infrastructure”. Alvaro Fernandez, The International Symposium on Grids and Clouds (ISGC) 2012 Conference, March 2012, Taipei. ATL-SOFT-PROC-2012-004.

- “Response of the Iberian Grid Computing Resources to the ATLAS activities during the LHC data-taking”, M. Kaci, IBERGRID 2012, 6th Iberian Grid Infrastructure Conference. Lisbon, Portugal. 7-9 November, 2012. ATL-COM-SOFT-2012-154.
- “Illustrative Example of Distributed Analysis in ATLAS Spanish Tier-2 and Tier-3 centers at IFIC”. Elena Oliver, Physics in Collision, 28 August – 1 September 2011, Vancouver, Canada. ATL-SOFT-SLIDE-2011-449.
- “Data analysis on the ATLAS Spanish Tier2” Elena Oliver, 4th IBER-GRID Conference. Braga (Portugal), May 24-27th, 2010.
- “Readiness of the ATLAS Spanish Federated Tier-2 for the Physics Analysis of the early collision events at the LHC”, Elena Oliver, International Conference on Computing in High Energy and Nuclear Physics (CHEP’09) 21-27 March 2009, Prague, Czech Republic.

The list of publications based on this work are:

- “A search for $t\bar{t}$ resonances in lepton+jets events with highly boosted top quarks collected in pp collisions at $\sqrt{s} = 7\text{TeV}$ with the ATLAS detector”, ATLAS Collaboration, Journal-ref: JHEP 1209 (2012) 041, arXiv:1207.2409.
- “ATLAS measurements of the properties of jets for boosted particle searches”, ATLAS Collaboration, arXiv:1206.5369.
- “A search for $t\bar{t}$ resonances with the ATLAS detector in 2.05fb^{-1} of proton-proton collisions at $\sqrt{s} = 7\text{TeV}$ ”, ATLAS Collaboration, Journal-ref: Eur.Phys.J. C72 (2012) 2083, arXiv:1205.5371.

Several authors have responded to the challenge of calculating the jet mass distribution from this study:

- “On jet mass distributions in Z+jet and dijet processes at the LHC”, Mrinal Dasgupta, Kamel Khelifa-Kerfa (Manchester U.), Simone Marzani, Michael Spannowsky (Durham U.), arXiv:1207.1640.
- “Investigation of Monte Carlo Uncertainties on Higgs Boson searches using Jet Substructure”, Peter Richardson, David Winn, arXiv:1207.0380.
- “Probing colour flow with jet vetoes”, Simone Marzani (Durham U.), arXiv:1205.6808.

- "Energy flow observables in hadronic collisions", F. Hautmann, arXiv:1205.5411.
- "Precision Jet Substructure from Boosted Event Shapes", Ilya Feige, Matthew D. Schwartz, Iain W. Stewart, Jesse Thaler, Journal-ref: Phys. Rev. Lett. 109, 092001 (2012), arXiv:1204.3898.

Agradecimientos

Han sido muchas las personas que me han ayudado en estos años. Nombro en especial a mis directores José Francisco Salt Cairols y Santiago González de la Hoz que he han dedicado todo su tiempo y aconsejado en base a su experiencia. También le doy las gracias a Marcel Vos y a Miguel Villaplana por prestarme atención y responderme a todas mis dudas.

Agradezco el trabajo realizado por Luis March y Juan López que me sirvió para empezar en esta aventura. También a los miembros del grupo de Grid Computing del IFIC: Azzeddine, Farida, Álvaro, Gabriel, Alex, Javier N., Javier S., Víctor L. y Víctor M. que estuvieron siempre dispuestos a ayudarme. Estoy muy agradecida a los miembros del Tier-2 español con los que he trabajado como Jordi, Xavi, los dos Carlos, Juanjo, Andreu, Alexei y Jose que ha sido un placer colaborar con ellos.

Quiero agradecer los consejos e ideas de los componentes del grupo aTOPe, tanto seniors como estudiantes: a Susana, Salva, Maria José, Vasia, Juan, Carmen, Teresa, María, Emma, Regina, Vicente, Adrián, Carlos, Viki, Elena, Sebas y Patricia.

Ha sido una experiencia muy gratificante trabajar en DAST donde he conocido profesionales como Dan, Nurcan, Mark, Manoj, Alden, Hurng, Wensheng, Kamile, Jarka, Frederic, Daniel, Marya, Nils, Bill, Mike, Jack, Philipp, Sergey y demás expertos de la colaboración ATLAS. También, ha sido un placer colaborar con universidades de Marruecos y así disfrutar de la compañía de Rajaa y Naima.

Quiero dar las gracias a mi familia y en especial a mi marido Paco que me ha ayudado a sobrellevar el día a día en el tiempo que ha llevado realizar esta tesis.

A todas estas personas les agradezco su dedicación, profesionalidad y su paciencia para guiarme todo este tiempo.

Bibliography

- [1] A. Ferrer and E. Ros. *Física de partículas y de astropartículas*. Educació. Sèrie Materials. Publicacions de la Universitat de València, 2005.
- [2] Cahn, R.N. and Goldhaber, G. *The Experimental Foundations of Particle Physics*. Cambridge books online. Cambridge University Press, 2009.
- [3] '<http://www.slac.stanford.edu/>'.
- [4] Anita Hollier. Wolfgang Pauli and Modern Physics. Technical Report CERN-OPEN-2012-001, CERN, Geneva, Sep 2000.
- [5] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [6] J. N. Abdurashitov et al. Measurement of the Solar Neutrino Capture Rate by SAGE and Implications for Neutrino Oscillations in Vacuum. *Phys. Rev. Lett.*, 83:4686–4689, Dec 1999.
- [7] '<http://www.cpepphysics.org/particle-chart.html>'.
- [8] M. Banner et al. Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider. *Physics Letters B*, 122(56):476 – 485, 1983.
- [9] G. Arnison et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s}=540$ GeV. *Physics Letters B*, 122(1):103 – 116, 1983.
- [10] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, Oct 1964.
- [11] Pol, Andri. *CERN: European Organization for Nuclear Research*. Lars Muller, Zurich, 2013. '<http://home.web.cern.ch/>'.

- [12] TeVI Group. Design Report Tevatron 1 project. 1984.
- [13] Raine, D.J. and Thomas, E.G. *An introduction to the science of cosmologyd*. Series in Astronomy and Astrophysics - Institute of Physics. Institute of Physics Publishing, 2001.
- [14] The ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1 – 29, 2012.
- [15] Yao, W. -M. and others. Review of particle physics. *J. Phys.*, G33:1–1232, 2006.
- [16] T. Appelquist, A. Chodos, and P.G.O. Freund. *Modern Kaluza-Klein theories*. Frontiers in physics. Addison-Wesley Pub. Co., 1987.
- [17] Antoniadis, I. and Arkani-Hamed, N. and Dimopoulos, S. and Dvali, G.R. New dimensions at a millimeter to a Fermi and superstrings at a TeV. *Phys.Lett.*, B436:257–263, 1998.
- [18] Ben Lillie, Lisa Randall, and Lian-Tao Wang. The Bulk RS KK-gluon at the LHC. *JHEP*, 09:074, 2007.
- [19] Lyndon Evans and Philip Bryant. Lhc machine. *Journal of Instrumentation*, 3(08):S08001, 2008.
- [20] S. Myers and E. Picasso. The design, construction and commissioning of the cern large electronpositron collider. *Contemporary Physics*, 31(6):387–403, 1990.
- [21] The CMS Collaboration. The CMS experiment at the CERN LHC. *JINST*, 3:S08004, 2008.
- [22] A. Augusto Alves Jr. et al. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3(08):S08005, 2008.
- [23] K Aamodt et al. The ALICE experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08002, 2008.
- [24] A. Abdessalam et al. Boosted objects as a probe of beyond the standard model physics. *J. Phys. G.*, 2010.
- [25] M. Vos. Boosting sensitivity to new physics - CERN Courier.
- [26] Ian Foster and Carl Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 2003.

- [27] Christoph Eck et al. *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*. Technical Design Report LCG. CERN, Geneva, 2005.
- [28] F. Magoulès. *Introduction to Grid Computing*. Number v. 10 in Chapman & Hall/CRC numerical analysis and scientific computing. CRC Press, 2009.
- [29] Yang, X. and Wang, L. and Jie, W. *Guide to e-Science: Next Generation Scientific Research and Discovery*. Computer communications and networks. Springer, 2011.
- [30] Robert J. Wilson. The European DataGrid. In *Snowmass 2001*, 2001.
- [31] Erwin Laure and Bob Jones. Enabling Grids for e-Science: The EGEE Project. Technical Report EGEE-PUB-2009-001. 1, Sep 2008.
- [32] Candiello, A. and Cresti, D. and Ferrari, T. and Karagiannis, F. and Kranzlmuller, D. and others. A business model for the establishment of the European grid infrastructure. *J.Phys.Conf.Ser.*, 219:062011, 2010.
- [33] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Wurthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick. The Open Science Grid. *Journal of Physics: Conference Series*, 78(1):012057, 2007.
- [34] M. Ellert et al. The NorduGrid project: using Globus toolkit for building GRID infrastructure. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 502:407 – 410, 2003. Proceedings of the VIII International Workshop on Advanced Computing and Analysis Techniques in Physics Research.
- [35] The ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.
- [36] The ATLAS Collaboration. *ATLAS inner detector: Technical Design Report, 1*. Technical Design Report ATLAS. CERN, Geneva, 1997. CERN-LHCC-97-016.
- [37] The ATLAS Collaboration. *ATLAS magnet system: Technical Design Report, 1*. Technical Design Report ATLAS. CERN, Geneva, 1997. CERN-LHCC-97-018.

-
- [38] The ATLAS Collaboration. *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1996. CERN-LHCC-96-041.
 - [39] The ATLAS Collaboration. *ATLAS tile calorimeter: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1996. CERN-LHCC-96-042.
 - [40] The ATLAS Collaboration. Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector. *Eur.Phys.J.*, C73:2306, 2013.
 - [41] The ATLAS Collaboration. *ATLAS muon spectrometer: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1997. CERN-LHCC-97-022.
 - [42] The ATLAS Collaboration. *ATLAS level-1 trigger: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 1998. CERN-LHCC-98-014.
 - [43] P. Jenni, M. Nelli, M. Nordberg and K. Smith. *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*. Technical Design Report ATLAS. CERN, Geneva, 2003. CERN-LHCC-2003-022.
 - [44] *Atlas Computing: technical design report*. CERN, Geneva, 2005.
 - [45] '<http://home.web.cern.ch/about/computing>'.
 - [46] www.pic.es.
 - [47] Xavier Espinal and et al. Iberian ATLAS computing: Facing data taking. In *IBERGRID 2009*, pages 187–209, 2009.
 - [48] R. Brock et al. U.S. ATLAS Tier 3 Task Force . Preprint U.S. ATLAS 2009.
 - [49] <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/DDMSonarTests>.
 - [50] I Ueda. ATLAS Distributed Computing Operations in the First Two Years of Data Taking. Technical Report ATL-SOFT-PROC-2012-003, CERN, Geneva, Mar 2012.
 - [51] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino and D. Quarrie. The athena control framework in production, new developments and lessons learned. pages 456–458, 2005.

- [52] 'http://en.wikipedia.org/wiki/Oracle_Database'.
- [53] '<http://www.sqlite.org>'.
- [54] S. Gonzalez de la Hoz. The evolving role of Tier2s in ATLAS with the new Computing and Data Distribution model. Technical Report ATL-SOFT-PROC-2012-010, CERN, Geneva, May 2012.
- [55] R. Basset, L. Canali, G. Dimitrov, M. Girone, R. Hawkings, P. Nevski, A. Valassi, A. Vaniachine, F. Viegas, R. Walker, and A. Wong. Advanced technologies for scalable atlas conditions database access on the grid. *Journal of Physics: Conference Series*, 219(4):042025, 2010.
- [56] A. V. Vaniachine and J. G. von der Schmitt. Development, deployment and operations of atlas databases. *Journal of Physics: Conference Series*, 119:072031, 2008.
- [57] Ian Bird. Computing for the large hadron collider. *Annual Review of Nuclear and Particle Science*, 61(1):99–118, 2011.
- [58] The ATLAS Collaboration. ATLAS Data Challenge 1. Technical Report ATL-SOFT-2003-012, CERN, Geneva, Nov 2003. revised version number 1 submitted on 2003-11-13 17:27:28.
- [59] G. Poulard L. Goossens and D. Barberis. ATLAS Data Challenges: large-scale productions on the Grid. *PoS*, HEP2005:394, 2006.
- [60] '<http://lcg-archive.web.cern.ch/lcg-archive/peb/bs/BSReport-v1.0.pdf>'.
- [61] T. Maeno. PanDA: distributed production and distributed analysis system for ATLAS. *Journal of Physics: Conference Series*, 119(6):062036, 2008.
- [62] 'http://en.wikipedia.org/wiki/Just-in-time_compilation'.
- [63] 'http://en.wikipedia.org/wiki/Zombie_process'.
- [64] E. Oliver, A. Fernández, S. González de la Hoz, M. Kaci, A. Lamas, J. Salt, J. Sánchez, and M. Villaplana. Illustrative Example of Distributed Analysis in ATLAS Spanish Tier-2 and Tier-3 centers. Aug 2011.
- [65] M. Branco, D. Cameron, B. Gaidioz, V. Garonne, B. Koblitiz, M. Lassinig, R. Rocha, P. Salgado and T. Wenaus. Managing ATLAS data on a petabyte-scale with DQ2. *Journal of Physics: Conference Series*, 119(6):062017, 2008.

-
- [66] Solveig Albrand, Jerome Fulachier and Fabian Lambert. The ATLAS metadata interface. *Journal of Physics: Conference Series*, 219(4):042030, 2010.
- [67] 'twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/DataTransferRequestInterface'.
- [68] Vincent Garonne, Graeme A Stewart, Mario Lassnig, Angelos Molfetas, Martin Barisits, Thomas Beermann, Armin Nairz, Luc Goossens, Fernando Barreiro Megino, Cedric Serfon, Danila Oleynik, and Artem Petrosyan. The atlas distributed data management project: Past and future. *Journal of Physics: Conference Series*, 396(3):032045, 2012.
- [69] P. Nilsson. Distributed Data Analysis in ATLAS. (ATL-SOFT-PROC-2009-007), Nov 2009.
- [70] Frederic Brochu, Ulrik Egede, J. Elmsheuser, K. Harrison, R. W. L. Jones, H. C. Lee, Dietrich Liko, A. Maier, Jakub T. Moscicki, A. Muraru, Glen N. Patrick, Katarina Pajchel, W. Reece, B. H. Samset, M. W. Slater, A. Soroko, C. L. Tan, Daniel C. Van der Ster. Ganga: a tool for computational-task management and easy access to Grid resources. *CoRR*, abs/0902.2685, 2009.
- [71] Jakub T. Moscicki, Susanna Guatelli, Alfonso Mantero, and M.G. Pia. Distributed geant4 simulation in medical and space science applications using DIANE framework and the GRID. *Nuclear Physics B - Proceedings Supplements*, 125(0):327 – 331, 2003. Innovative Particle and Radiation Detectors.
- [72] GEANT4: A simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [73] 'https://ggus.eu/pages/home.php'.
- [74] T. Maeno, K. De, T. Wenaus, P. Nilsson, R. Walker, et al. Evolution of the ATLAS PanDA Production and Distributed Analysis System. 2012.
- [75] Daniel C van der Ster, Johannes Elmsheuser, Mario Ubeda Garcia, and Massimo Paladin. Hammercloud: A stress testing system for distributed analysis. *Journal of Physics: Conference Series*, 331(7):072036, 2011.
- [76] T Maeno, K De, and S Panitkin. PD2P : PanDA Dynamic Data Placement for ATLAS. Technical Report ATL-SOFT-PROC-2012-016, CERN, Geneva, May 2012.

- [77] <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/AtlasDAST>.
- [78] "Max Baak". "good run lists, what about them?".
- [79] M Villaplana, S Gonzalez de la Hoz, A Fernandez, J Salt, A Lamas, F Fassi, M Kaci, E Oliver, J Sanchez, and V Sanchez-Martinez. ATLAS Tier-3 within IFIC-Valencia analysis facility. Technical Report ATL-SOFT-PROC-2012-017, CERN, Geneva, May 2012.
- [80] E. Oliver, J. Nadal, J. Pardo, G. Amorós, C. Borrego, M. Campos, L. Del Cano, J. Del Peso, X. Espinal, F. Fassi, A. Fernández, P. Fernández, S. González, M. Kaci, A. Lamas, L. March, L. Muñoz, A. Pacheco, J. Salt, J. Sánchez, M. Villaplana and R. Vives. Readiness of the ATLAS Spanish Federated Tier-2 for the Physics Analysis of the early collision events at the LHC. *Journal of Physics Conference Series*, 219(7):072046, apr 2010.
- [81] '<http://twiki.ific.uv.es/twiki/bin/view/Atlas/LustreStoRM>'.
- [82] Elena Oliver and et al. . Operations of the ATLAS Spanish Tier-2 during the LHC Run I'. In *IBERGRID 2013*, pages 227–229, 2013.
- [83] Gabriel Amorós and et al. . Data analysis on the ATLAS Spanish Tier2. In *IBERGRID 2010*, pages 221–222, 2010.
- [84] Jose Salt Cairols. Final inform of the project FPA2005-07688-C03-01. Technical report, 2008.
- [85] 'https://espace.cern.ch/WLCG-document-repository/MoU/signed/CERN-C-RRB-2005-01_rev_Spain.pdf'.
- [86] <http://w3.hepux.org/benchmarks/doku.php/>.
- [87] <http://www.dcache.org/manuals/dcache-whitepaper-light.pdf>.
- [88] '<http://www.rediris.es/>'.
- [89] '<https://wiki.egi.eu/wiki/SAM>'.
- [90] <http://www.nagios.org/>.
- [91] <http://www.cacti.net/>.
- [92] '<http://puppetlabs.com/solutions/configuration-management/>'.
- [93] '<http://dashboard.cern.ch/atlas/>'.

-
- [94] <https://indico.cern.ch/materialDisplay.py?contribId=2&materialId=0&confId=20248>.
- [95] M. Kaci, G. Borges, M. David, N. Dias, F. Fassi, A. Fernandez, H. Gomes, J. Gomes, S. Gonzalez de la Hoz, V. Lacort, J.P. Martins, M. Oliveira, E. Oliver, A. Pacheco Pages, J. del Peso, J. Pina, J. Salt, J. Sanchez, V. Sanchez, S. Sayzhenkova, A. Sedov, M. Villaplana, and H. Wolters. Response of the Iberian Grid Computing Resources to the ATLAS activities during the LHC data-taking. Technical Report ATL-COM-SOFT-2012-154, CERN, Geneva, Aug 2012. IBERGRID 2012, 6th Iberian Grid Infrastructure Conference Lisbon, Portugal, November 7-9, 2012 <http://www.ibergrid.eu/2012/index.php>.
- [96] <https://savannah.cern.ch/>.
- [97] <http://twiki.ific.uv.es/twiki/bin/view/Atlas/AtlasDataProcessingAtIFIC>.
- [98] Michael H. Seymour. Searches for new particles using cone and cluster jet algorithms: A Comparative study. *Z. Phys.*, C62:127–138, 1994.
- [99] Jesse Thaler and Lian-Tao Wang. Strategies to Identify Boosted Tops. *JHEP*, 07:092, 2008.
- [100] Proceedings of the BOOST 2011 conference. 2012.
- [101] Chekanov, S. and others. Measurement of subjet multiplicities in neutral current deep inelastic scattering at HERA and determination of α_s . *Phys. Lett. B*, 558:41, 2003.
- [102] Chekanov, S. and others. Substructure dependence of jet cross sections at HERA and determination of α_s . *Nucl. Phys. B*, 700:3, 2004.
- [103] Chekanov, S. and others. Subjet Distributions in Deep Inelastic Scattering at HERA. *Eur. Phys. J*, 63:527, 2009.
- [104] D. Acosta et al. Study of Jet Shapes in Inclusive Jet Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV. *Phys. Rev. D*, 71:112002, 2005.
- [105] T. Aaltonen et al. The Substructure of High Transverse Momentum Jets Observed by CDF II. *CDF Note*, 10199, 2010.
- [106] T. Aaltonen et al. Preliminary Results of a Search for Boosted Top Quarks by CDF II. *CDF Note*, 10234, 2010.
- [107] The ATLAS Collaboration. Jet mass and substructure of inclusive jets in $\sqrt{s} = 7$ TeV pp collisions with the ATLAS experiment. *JHEP*, 1205:128, 2012.

- [108] Gavin P. Salam. Towards Jetography. *Eur. Phys. J.*, C67:637–686, 2010.
- [109] Matteo Cacciari and Gavin P. Salam. Dispelling the N^3 myth for the k_t jet-finder. *Phys. Lett.*, B641:57–61, 2006.
- [110] The ATLAS Collaboration. Prospects for top anti-top resonance searches using early ATLAS data. ATL-PHYS-PUB-2010-008, Jul 2010.
- [111] S. Catani, Yuri L. Dokshitzer, M. H. Seymour, and B. R. Webber. Longitudinally invariant k_\perp clustering algorithms for hadron hadron collisions. *Nucl. Phys. B*, 406:187–224, 1993.
- [112] Stephen D. Ellis and Davison E. Soper. Successive combination jet algorithm for hadron collisions. *Phys.Rev.*, D48:3160–3166, 1993.
- [113] J. M. Butterworth, B. E. Cox, and Jeffrey R. Forshaw. WW scattering at the CERN LHC. *Phys. Rev.*, D65:096014, 2002.
- [114] The ATLAS Collaboration. A search for $t\bar{t}$ resonances in the lepton plus jets final state with ATLAS using 14 fb¹ of pp collisions at $\sqrt{s} = 8$ TeV. Technical Report ATLAS-CONF-2013-052, CERN, Geneva, May 2013.
- [115] The ATLAS Collaboration. Updated Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector. ATL-CONF-2011-011, 2011.
- [116] The ATLAS Collaboration. Calorimeter clustering algorithms: description and performance. Technical Report ATL-LARG-PUB-2008-002, CERN, Geneva, 2008.
- [117] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 05:026, 2006.
- [118] M. Bahr et al. Herwig++ Physics and Manual. *Eur. Phys. J.*, C58:639–707, 2008.
- [119] M. L. Mangano et al. ALPGEN, a generator for hard multiparton processes in hadronic collisions. *JHEP*, 07:001, 2003.
- [120] T. Gleisberg et al. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009.

-
- [121] Bo Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand. Parton Fragmentation and String Dynamics. *Phys. Rept.*, 97:31–145, 1983.
- [122] Webber, B. R. A QCD Model for Jet Fragmentation Including Soft Gluon Interference. *Nucl. Phys.*, B238:492, 1984.
- [123] G. Corcella et al. HERWIG 6.5 release note. 2002.
- [124] G. Corcella et al. HERWIG 6: An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *JHEP*, 01:010, 2001.
- [125] J. M. Butterworth, Jeffrey R. Forshaw, and M. H. Seymour. Multiparton interactions in photoproduction at HERA. *Z. Phys.*, C72:637–646, 1996.
- [126] The ATLAS Collaboration. Charged-particle multiplicities in pp interactions at $\sqrt{s} = 900$ GeV measured with the ATLAS detector at the LHC. *Phys.Lett.*, B688:21–42, 2010.
- [127] Martin, A.D. and Stirling, W. J. and Thorne, R.S. and Watt, G. Parton distributions for the LHC. *Eur. Phys. J.*, C63:189–285, 2009.
- [128] Sherstnev, A. and Thorne, R.S. Parton Distributions for LO Generators. *Eur. Phys. J.*, C55:553–575, 2008.
- [129] The ATLAS Collaboration. In-situ jet energy scale and jet shape corrections for multiple interactions in the first ATLAS data at the LHC $\sqrt{s} = 7$ TeV. ATLAS-CONF-2011-030, 2011.
- [130] The ATLAS Collaboration. Jet energy resolution and reconstruction efficiencies from in-situ techniques with the ATLAS Detector Using Proton-Proton Collisions at a Center of Mass Energy $\sqrt{s} = 7$ TeV. ATLAS-CONF-2010-054, 2010.
- [131] B. Malaescu. An Iterative, Dynamically Stabilized(IDS) Method of Data Unfolding. *ArXiv e-prints*, 2011.
- [132] B. Malaescu. An iterative, dynamically stabilized method of data unfolding. *ArXiv e-prints*, 2009.
- [133] The ATLAS Collaboration. HepData relating to the article 'Jet mass and substructure of inclusive jets in $\sqrt{s} = 7$ TeV pp collisions with the ATLAS experiment', 2012. <http://hepdata.cedar.ac.uk/view/red4953>.

- [134] A. Altheimer, S. Arora, L. Asquith, G. Brooijmans, J. Butterworth, et al. Jet Substructure at the Tevatron and LHC: New results, new tools, new benchmarks. *J.Phys.*, G39:063001, 2012.
- [135] Yang-Ting Chien, Randall Kelley, Matthew D. Schwartz, and Hua Xing Zhu. Resummation of Jet Mass at Hadron Colliders. *Phys.Rev.*, D87:014010, 2013.
- [136] Mrinal Dasgupta, Kamel Khelifa-Kerfa, Simone Marzani, and Michael Spannowsky. On jet mass distributions in Z+jet and dijet processes at the LHC. *JHEP*, 1210:126, 2012.
- [137] Ilya Feige, Matthew D. Schwartz, Iain W. Stewart, and Jesse Thaler. Precision Jet Substructure from Boosted Event Shapes. *Phys.Rev.Lett.*, 109:092001, 2012.
- [138] Teppo Jouttenus, Iain Stewart, Frank Tackmann, and Wouter Waalewijn. Jet Mass Spectra in Higgs + One Jet at NNLL. 2013.
- [139] ATLAS Collaboration. A search for $t\bar{t}$ resonances in lepton+jets events with highly boosted top quarks collected in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. oai:cds.cern.ch:1461455. Technical Report arXiv:1207.2409. CERN-PH-EP-2012-158, CERN, Geneva, Jul 2012. Comments: 25 pages plus author list (47 pages total), 7 figures, 4 tables, submitted to Journal of High Energy Physics.
- [140] The ATLAS Collaboration. ATLAS detector and physics performance. Technical design report. Volume 1 and 2. CERN-LHCC-99-14/15.