Universitat de València (Estudi General, UVEG)

DEPARTAMENT DE FÍSICA ATÒMICA,

MOL·LECULAR I NUCLEAR

VNIVERSITAT DE VALÈNCIA
(ESTUDI GENERAL)

Application of Grid technologies and search for exotics physics with the ATLAS experiment at the LHC

Tesi doctoral

Luis March Ruiz

novembre 2007

D. EDUARDO ROS MARTÍNEZ,
Científico Titular del CSIC, y
D. SANTIAGO GONZÁLEZ DE LA HOZ,
Titulado Superior de Investigación y Laboratorio del CSIC,

CERTIFICAN :

Que la presente memoria que tiene por título: '*Application of Grid technologies and search for exotics physics with the ATLAS experiment at the LHC*', ha sido realizada bajo su dirección en el Instituto de Física Corpuscular (Centro Mixto Universitat de València - CSIC) por D. Luis March Ruiz y constituye su trabajo de tesis doctoral en el Departament de Física Atòmica, Mol·lecular i Nuclear de la Universitat de València para optar al grado de Doctor en Física.

Y para que conste, en cumplimiento de la legislación vigente, firmamos el presente certificado en Burjassot a 23 de noviembre de 2007.

Fdo. Eduardo Ros Martínez Fdo. Santiago González de la Hoz

V^oB^o de la tutora:
Da. $M^{\underline{a}}$ Victoria Castillo Giménez

A mis padres

Contents

Introduction

The work presented in this thesis has been performed within the ATLAS (A Toroidal LHC ApparatuS) collaboration. Two subjects have been investigated. One subject is the Computing System Commissioning (CSC) production using an instance of the Production System (ProdSys), called Lexor, and the test of the ATLAS Distributed Analysis (ADA) using ProdSys. The other subject is the simulation and subsequent analysis of processes involving new particles predicted by the Little Higgs model within the ATLAS detector.

An introduction to the Standard Model (SM), the Large Hadron Collider (LHC) and the ATLAS experiment, software and computing is given in chapter 1. The problems of the SM are discussed and some proposed solutions are reviewed. The SM introduction is followed by an overview of LHC and ATLAS. The main ATLAS subsystems are described and the ATLAS software and computing model is discussed.

Many physics processes within and beyond the Standard Model involve b-quark decays. New heavy particles, expected in models like the Little Higgs model, produce high p_T b-jets. For this reason, a detailed b-tagging study of high and very high-p_T jets, using the ATLAS full detector simulation, is presented in chapter 2.

The Little Higgs model is a novel theory introduced to solve the hierarchy problem of the Standard Model without Supersymmetry (SUSY). In this model, new heavy particles cancel the radiative corrections of the Standard Model to the Higgs mass, allowing a low mass Higgs boson. These new particles are within the reach of the LHC. The simplest model based on the Little Higgs ideas is the so-called Littlest Higgs model. This model has been used as reference for the simulation studies presented in this thesis. In chapter 3, the principles and phenomenology of this model are discussed. The possibility to discover the new predicted particles using the ATLAS detector has been investigated. More precisely, the hadronic decay channels of heavy gauge bosons are analysed using the ATLAS fast detector simulation (ATLFAST).

The ATLAS collaboration is preparing for data taking, scheduled to start in 2008. After data collection by the online trigger, the expected data volume recorded for offline reconstruction and analysis is of the order of few Petabytes per year, to be processed by institutions distributed worldwide. Therefore the required comput-

ing resources are two orders of magnitude larger than for previous experiments in particle physics. The LHC experiments need a worldwide distributed data management and computing system to handle this huge amount of data. The LHC Computing Review in 2001 recommended that LHC experiments should carry out Data Challenges (DCs) of increasing size and complexity to be prepared for LHC data taking. The goal of the ATLAS Data Challenges is the validation of the ATLAS computing model. A description of the ATLAS Production System, that allows to manage massive productions of Monte Carlo data using Grid resources, and the main results obtained for ATLAS DCs are discussed in chapter 4. A description of the IFIC contribution to the ATLAS CSC simulation production running an ATLAS ProdSys instance is given also.

Finally, as part of the preparation for data taking and physics analysis, the ATLAS Distributed Analysis model and tools are described in chapter 5, including an example of application using the ADA ProdSys. A description of a Tier-3 prototype for IFIC is presented as well.

Introducción

El trabajo presentado en esta tesis ha sido realizado dentro de la colaboración ATLAS (A Toroidal LHC ApparatuS). Dos temas se han investigado. El primero es la producción Monte Carlo denominada *Computing System Commissioning (CSC)* utilizando una instancia del sistema de producción (ProdSys) llamada Lexor y el test de Análisis Distribuido (ADA) utilizando ProdSys. El segundo tema es la simulación y análisis de procesos de ATLAS incluyendo nuevas partículas postuladas por el modelo *Little Higgs*.

Una introducción al Modelo Estándar (ME), al acelerador *Large Hadron Collider (LHC)* y al experimento ATLAS, su *software* y modelo de computación se incluye en el capítulo 1. Los problemas del ME y algunas de sus posibles soluciones se discuten también en este capítulo. La introducción al ME se completa con una descripción del LHC y ATLAS. Los principales subsistemas del experimento ATLAS se describen también, así como su modelo de computación y *software*.

Varios procesos en el ME y sus posibles extensiones incluyen desintegraciones de *quarks b*. Nuevas partículas pesadas, esperadas en modelos como el llamado *Little Higgs*, producen *b-jets* de alto p_T. Por este motivo, en el capítulo 2 se presenta un estudio detallado de identificación de *b-jets* de alto y muy alto p_T, utilizando una simulación detallada del detector ATLAS.

El modelo *Little Higgs* es una teoría reciente introducida para solucionar el problema de jerarquía del Modelo Estándar sin introducir SuperSimetría (SUSY). En este modelo, nuevas partículas pesadas cancelan las correcciones radiativas del Modelo Estándard en la masa del *Higgs*, permitiendo una masa ligera para dicha partícula. Estas nuevas partículas podrían ser descubiertas en el rango de energía del LHC. El modelo más simple basado en las ideas del modelo *Little Higgs* se denomina modelo *Littlest Higgs*. Este modelo se ha utilizado como referencia para los estudios de simulación que se presentan en esta tesis. En el capítulo 3, se discuten los principios y la fenomenología de este modelo. Se ha investigado la posibilidad de descubrir estas nuevas partículas utilizando el detector ATLAS.

La colaboración ATLAS se está preparando para la toma de datos del LHC, que debería de iniciarse en el año 2008. El volumen de datos esperados para su posterior reconstrucción y análisis es del orden de algunos *Petabytes* por año, para ser procesados por los diversos institutos que participan en ATLAS distribuidos por todo el mundo. Los recursos de computación necesarios son dos órdenes de

magnitud superiores a los recursos disponibles en experimentos anteriores de física de partículas. Los experimentos del LHC necesitan una gestión de datos y un sistema de computación distribuidos mundialmente para poder manipular esta enorme cantidad de datos. El *LHC Computing Review* de 2001 recomendó que los experimentos del LHC deberían realizar una serie de *Data Challenges (DCs)* de tamaño y complejidad crecientes para preparar la toma de datos del LHC. La meta de los *ATLAS Data Challenges* es la validación del modelo de computación del experimento ATLAS. En el capítulo 4, se describe el sistema de producción del experimento ATLAS, el cual permite la administración de producciones masivas de datos tipo Monte Carlo utilizando los recursos Grid y los principales resultados más destacados obtenidos en los ATLAS DCs. Además, se detalla la contribución del IFIC en la producción CSC, en la que se ha utilizado una instancia del ProdSys.

Finalmente, como parte de la preparación para la toma de datos y análisis de física, se describe en el capítulo 5 el modelo y las herramientas de Análisis Distribuido del experimento ATLAS, incluyendo un ejemplo de aplicación utilizando ADA con ProdSys. Además, se presenta también una descripción de un prototipo Tier-3 para el IFIC.

Chapter 1

LHC, ATLAS, physics and computing

The Standard Model (SM) is the only existing mathematically consistent description of fundamental particles and their interactions. It has been enormously successful in explaining and predicting the results of many high energy physics experiments. Despite the success of the Standard Model, there are still opened questions that the Standard Model can not answer.

In this chapter, an introduction to the Standard Model, its problems and some proposed solutions beyond the Standard Model are discussed in section 1.1.

The Large Hadron Collider (LHC) is a proton-proton and heavy-ion collider with an unprecedented center-of-mass energy of 14 TeV, located at CERN (Geneva). The accelerator is currently being installed in the LEP tunnel.

A Toroidal LHC ApparatuS (ATLAS) detector is currently under construction in one of the experimental areas of the LHC.

An overview of the LHC accelerator, the ATLAS experiment and its main detector subsystems are described in section 1.2.

Finally, the ATLAS software and computing model is discussed in section 1.3.

1.1 The Standard Model and beyond

The Standard Model (SM) is currently the only satisfactory description of the interactions between fundamental particles, with the exception of gravity. This model explains at present the results obtained by a large number of High Energy Physics experiments.

The Standard Model has three main components, associated with three fundamental interactions. Quantum Electrodynamics (QED) describes the electromagnetic interaction. The Glashow-Weinberg-Salam model describes the weak interaction unified with QED, called electroweak unification. Finally, Quantum Chromody-

namics (QCD) describes the strong interaction between quarks. Up to now, the integration of the fourth fundamental interaction, gravity, into the SM has not succeeded yet. The Standard Model describes the interactions between fundamental particles. An important property of fundamental particles is the intrinsic angular momentum or spin. The matter is made of fermions, particles with spin s=1/2. The interaction between fermions is explained as an exchange of intermediate particles or gauge bosons with spin s=1.

The SM is a gauge theory, based on the symmetry group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$. It describes strong, weak and electromagnetic interactions, as an exchange of the corresponding spin-1 gauge fields: 8 massless gluons and 1 massless photon for the strong and electromagnetic interactions, respectively, and 3 massive bosons, $W^{\pm}$ and Z, for the weak interaction. The fermionic matter consists of the known leptons and quarks, which appear in a 3-fold family structure:

$$\begin{bmatrix} \nu_e & u \\ e^- & d' \end{bmatrix}, \qquad \begin{bmatrix} \nu_\mu & c \\ \mu^- & s' \end{bmatrix}, \qquad \begin{bmatrix} \nu_\tau & t \\ \tau^- & b' \end{bmatrix} \tag{1.1}$$

Each quark appears in 3 different states called 'colours'. Within each family, quarks and leptons can be arranged into doublets and singlets:

$$\begin{bmatrix} \nu_l & q_u \\ \ell^- & q_d \end{bmatrix} \equiv \begin{pmatrix} \nu_l \\ \ell^- \end{pmatrix}_L, \quad \begin{pmatrix} q_u \\ q_d \end{pmatrix}_L, \quad \ell^-_R\,, \quad (q_u)_R\,, \quad (q_d)_R \tag{1.2}$$

Thus, the left-handed fields are $SU(2)_L$ doublets, while their right-handed partners transform as $SU(2)_L$ singlets. The 3 fermionic families in equation (1.1) appear to have identical properties (gauge interactions). They only differ by their mass and flavour quantum number.

The main properties of known fermions are listed in table 1.1. Fermions can be classified according to the the strong interaction. The fermions with strong interactions, i.e. those with colour charge, are called quarks. Quarks are also sensitive to the weak and electromagnetic interactions. There are six types of quarks (also called quark flavours). These flavours are *up*, *down*, *strange*, *charm*, *bottom* and *top*. Ordinary matter is made of u and d quarks. The combination of two u quarks and one d quark makes up the proton, while the neutron consists of two d quarks and one u quark. The heaviest quarks can only be produced in experiments with high energy accelerators. Quarks cannot been observed as free particles, but their existence is confirmed by a large number of experiments.

The remaining fermions have only weak and electromagnetic interactions, and are called leptons. Electron, muon and tau leptons have electric charge and weak isospin (i.e. weak interaction). Only electrons exist in ordinary matter. The corresponding neutrinos are electrically neutral and interact only weakly. The electronic neutrino ν_e was discovered originally in neutron decays: $n \rightarrow p + e^- + \bar{\nu}_e$. Neutrinos have a very small mass, but non-zero as discovered recently. Only imprecise values for these neutrino masses are currently available [1].

quarks	mass (GeV)	Q (q_e)		leptons	mass (GeV)	Q (q_e)
d	0.01	-1/3		e	0.000511	-1
u	0.01	+2/3		ν_e	~ 0	0
s	0.1	-1/3		μ	0.106	-1
c	1.2	+2/3		ν_μ	~ 0	0
b	4.3	-1/3		τ	1.8	-1
t	165	+2/3		ν_τ	~ 0	0

Table 1.1: *Fundamental particles of the Standard Model with spin 1/2. Charged lepton masses are known with high accuracy, but quark masses only in an approximate way. Each fermion is associated to a corresponding anti-fermion, with same mass but opposite charge.*

Fermions from the Standard Model are classified into three families or generations. Lower limits for the mass of hypothetical fourth generation fermions are available. On the other hand, the analysis of the Z width made at LEP [1] shows that there are exactly three light neutrino families. This result disfavours the existence of a fourth generation. The family structure and the different fermion masses are not explained by the Standard Model itself, but they are essential elements in its construction.

Fermions interact via intermediate particles, the gauge bosons. Each known interaction is associated to a particle with spin 1. These particles are listed in table 1.2. The gauge bosons of the weak and electromagnetic interactions are well determined experimentally. The carrier of the electromagnetic interaction is the photon, observed in daily light. The photon is massless (the upper experimental limit is $m_\gamma < 10^{-6}$ eV). The gauge bosons of the weak interaction, W and Z, have been observed in accelerator experiments. Their most important properties (mass and width) are known with high precision. Gluons, responsible for the strong interaction, cannot be observed as free particles, the same as quarks. There is however sufficient experimental evidence concerning their existence. At present gravity is not included in the Standard Model. It is believed however that an intermediate particle, called graviton, is responsible for the gravitational interaction.

The gauge symmetry is broken by the vacuum. The mechanism is known as Spontaneous Symmetry Breaking (SSB) of the electroweak group to the electromagnetic subgroup:

$$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y \stackrel{SSB}{\rightarrow} SU(3)_C \otimes U(1)_{QED}$$

The SSB mechanism generates a mass for the weak gauge bosons and for fermions as well. It generates an additional scalar particle in the model, the so-called 'Higgs boson'. The Higgs boson is not associated to any fundamental interaction, but is also included in table 1.2. This particle is neutral but the mass is unknown, since no experimental evidence is yet available. Therefore, the SM Higgs boson is a hypothetical particle and it's the simplest solution.

interaction	name	symbol	spin	mass (GeV/c^2)	Q (q_e)
electromagnetic	photon	γ	1	0	0
weak	intermediate bosons	Z W	1 1	91.2 80.4	0 ± 1
strong	gluons	g	1	0	0
gravitational	graviton	G	2	0	0
	Higgs	H	0	$m_H > 114$	0

Table 1.2: *Elementary bosons of the Standard Model. These particles with integer spin are responsible for the fundamental interactions between fermions.*

1.1.1 Quantum Electrodynamics

The Standard Model is a gauge theory, as explained before. This means that the theory is invariant under certain local transformations. The Lagrangian describing a free Dirac fermion is:

$$\mathcal{L}_0 = i\overline{\psi}(x)\gamma^\mu\partial_\mu\psi(x) - m\overline{\psi}(x)\psi(x) \tag{1.3}$$

where ψ is the wave function of the fermion and γ^μ are the Dirac matrices. There are two terms in the Lagrangian: the kinetic term and the mass term. $\mathcal{L}_0$ is invariant under global $U(1)$ transformations:

$$\psi(x) \overset{U(1)}{\rightarrow} \psi'(x) \equiv exp\{iQ\theta\}\psi(x)$$

where $Q\theta$ is an arbitrary real constant. The phase of $\psi(x)$ is therefore a conventional quantity without physical meaning. However, the free Lagrangian is no longer invariant if one allows the phase transformation to depend on the space-time coordinate, i.e. $\theta = \theta(x)$. In this case, we obtain:

$$\partial_\mu\psi(x) \overset{U(1)}{\longrightarrow} exp\{iQ\theta\}(\partial_\mu + iQ\partial_\mu\theta)\psi(x)$$

The 'Gauge Principle' is the requirement that the U(1) phase invariance should hold also locally. For this purpose, an additional piece has to be added to the Lagrangian, transforming in a way to cancel the $\partial_\mu\theta$ term in equation (1.3). This cancellation is achieved by introducing a new spin-1 field $A_\mu(x)$, transforming as

$$A_\mu(x) \overset{U(1)}{\longrightarrow} A'_\mu \equiv A_\mu(x) + \frac{1}{e}\partial_\mu\theta$$

The covariant derivative is defined in the following way:

$$D_\mu\psi(x) \equiv [\partial_\mu - ieQA_\mu(x)]\psi(x)$$

and has the required property of transforming like the field itself:

$$D_\mu\psi(x) \overset{U(1)}{\longrightarrow} (D_\mu\psi)'(x) \equiv exp\{iQ\theta\}D_\mu\psi(x)$$

The complete Lagrangian is therefore

$$\mathcal{L} \equiv i\overline{\psi}(x)\gamma^\mu D_\mu \psi(x) = \mathcal{L}_0 + eQA_\mu(x)\overline{\psi}(x)\gamma^\mu\psi(x) \tag{1.4}$$

and this Lagrangian is invariant under local $U(1)$ transformations.

The gauge principle has generated an interaction between the Dirac spinor and the gauge field A_μ, which is just the familiar interaction of Quantum Electrodynamics (QED). The corresponding electromagnetic charge eQ is completely arbitrary. A gauge-invariant kinetic term for the A_μ field is required as well. This term reads:

$$\mathcal{L}_{kin} \equiv -\frac{1}{4}F_{\mu\nu}(x)F^{\mu\nu}(x) \tag{1.5}$$

where $F_{\mu\nu} \equiv \partial_\mu A_\nu - \partial_\nu A_\mu$ is the usual electromagnetic field strength. A possible mass term for the gauge field, $\mathcal{L}_m = \frac{1}{2}m^2 A^\mu A_\mu$, is forbidden to preserve gauge invariance. Therefore, the photon field is predicted to be massless, as experimentally observed.

The complete Lagrangian in (1.4) and (1.5) yield the well-known Maxwell equations:

$$\partial_\mu F^{\mu\nu} = J^\nu \qquad\qquad J^\nu = -eQ\overline{\psi}\gamma^\nu\psi$$

where J^ν is the fermion electromagnetic current. The QED Lagrangian, that can therefore be obtained from a simple gauge invariance principle, leads to a very successful quantum field theory.

1.1.2 Quantum Chromodynamics

The large number of known mesonic and baryonic states implies the existence of a deeper level of elementary constituents for matter. These constituents are the quarks. Assuming that mesons are $M \equiv q\overline{q}$ states and baryons have three quark constituents, $B \equiv qqq$, one can classify the entire hadronic spectrum. However, in order to satisfy Fermi-Dirac statistics, it is necessary to assume the existence of a new quantum number, colour. Each quark type may have $N_c = 3$ different colours: q^α, α= 1,2,3 (red, green, blue). Baryons and mesons are then described by the colour-singlet combinations

$$B = \frac{1}{\sqrt{6}}\epsilon^{\alpha\beta\gamma}|q_\alpha q_\beta q_\gamma > \qquad\qquad M = \frac{1}{\sqrt{3}}\delta^{\alpha\beta}|q_\alpha \overline{q}_\beta >$$

In order to avoid non-observed states with non-zero colour, it is necessary to postulate that all physical states are colourless, i.e. singlets under rotations in colour space. This assumption is known as the confinement hypothesis.

A direct test of the colour hypothesis can be obtained from the ratio:

$$R_{e^+e^-} \equiv \frac{\sigma(e^+e^- \to hadrons)}{\sigma(e^+e^- \to \mu^+\mu^-)} \approx \sum_q N_c c_q^2$$

where the experimental value of $R_{e^+e^-}$ implies $N_c = 3$.

The hadronic production occurs through $e^+e^- \rightarrow \gamma^* / \ Z^* \rightarrow q\overline{q} \rightarrow hadrons$. Since quarks are confined, they hadronize and therefore the sum over all possible quarks is equivalent to the inclusive cross-section into hadrons.

Bjorken scaling, observed in deep inelastic scattering experiments, suggested that the strong interaction must have the property of asymptotic freedom. This means that the interaction is weaker at short distances, or equivalently that quarks behave as free particles for $Q^2 \rightarrow \infty$. Thus, the interaction between a $q\overline{q}$ pair is of a string type. If we try to separate the quark from the antiquark, the force increases. At some point, the energy on the elastic string is larger than $2m_{q'}$, and an additional pair $q'\overline{q'}$ is created. As a result, a mesonic system $q\overline{q'}+ q'\overline{q}$ is obtained. Increasing further the energy can only result in the production of additional mesons, but quarks remain confined within colour-singlet bound states. Conversely, if one tries to approximate two quarks into a very short-distance, the string loses its energy and quarks behave as free particles.

Electromagnetic interactions are associated with the fermion electric charge, and quark flavours (*up*, *down*, *strange*, *charm*, *bottom*, *top*) are related to electroweak phenomena. The strong forces are both flavour-conserving and flavour-independent. On the other side, the carriers of the electroweak interaction (γ,Z,$W^{\pm}$) do not couple to the quark colour. Colour is the only charge associated to the strong force and the quantum field theory describing this force is based on this property of quarks.

The theory of colour interaction is based on the $SU(3)_c$ group. It remains to be demonstrated that such a theory is able to explain both confinement and asymptotic freedom. In the following we denote q_f^{α} a quark field of colour α and flavour f and we introduce a vector notation in colour space: $q_f \equiv (q_f^1, q_f^2, q_f^3)$. The free Lagrangian reads therefore

$$\mathcal{L}_0 = \sum_f \overline{q}_f (i\gamma^{\mu}\partial_{\mu} - m_f) q_f$$

and is invariant under an arbitrary global $SU(3)_c$ transformation in colour space. As in QED, the Lagrangian is assumed to be also invariant under local $SU(3)_c$ transformations, $\theta_a = \theta_a(x)$. This requirement is satisfied replacing the quark derivatives by covariant objects. Since we have eight independent gauge parameters, eight new gauge bosons $G_a^{\mu}(x)$, the so-called gluons, are needed, and the covariant derivative reads:

$$D^{\mu} q_f \equiv [\partial^{\mu} - i g_s \frac{\lambda^a}{2} G_a^{\mu}(x)] q_f \equiv [\partial^{\mu} - i g_s G^{\mu}(x)] q_f$$

where the compact matrix notation

$$[G^{\mu}(x)]_{\alpha\beta} \equiv (\frac{\lambda^a}{2})_{\alpha\beta} G_a^{\mu}(x)$$

has been used. The matrices λ^a (a=1,2,...,8) are the generators of the fundamental representation of the $SU(3)_c$ algebra, also called the Gell-Mann (3x3) matrices.

The complete Lagrangian of Quantum Chromodynamics (QCD) is:

$$\mathcal{L}_{QCD} = \mathcal{L}_{kin} + \mathcal{L}_{fermionic} = -\frac{1}{2}Tr(G^{\mu\nu}G_{\mu\nu}) + \sum_f \overline{q}_f(i\gamma D_\mu - m_q)q_f$$

where the following terms have been introduced to build the gauge-invariant kinetic Lagrangian for the gluon fields:

$$G^{\mu\nu}(x) \equiv \frac{i}{g_s}[D^\mu, D^\nu] = \partial^\mu G^\nu - \partial^\nu G^\mu - ig_s[G^\mu, G^\nu]$$

All interactions are given in terms of a single universal coupling g_s, called the strong coupling constant. The existence of self-interactions among the gauge fields is a new feature that was absent in QED. It can be shown that these gauge self-interactions explain asymptotic freedom and it is expected (but not proved) that they also explain confinement.

1.1.3 Electroweak Unification

The symmetry group of the electroweak interaction is:

$$G \equiv SU(2)_L \otimes U(1)_Y$$

where L refers to left-handed fields and Y refers to hyper-charges. We consider first the free Lagrangian

$$\mathcal{L}_0 = \sum_{j=1}^{3} i\overline{\psi}_j(x)\gamma^\mu\partial_\mu\psi_j(x)$$

where the following notation for quarks is used:

$$\psi_1(x) = \begin{pmatrix} u \\ d \end{pmatrix}_L \qquad \psi_2(x) = u_R \qquad \psi_3(x) = d_R$$

and in the same way for leptons:

$$\psi_1(x) = \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L \qquad \psi_2(x) = \nu_{eR} \qquad \psi_3(x) = e^-_R$$

As in QED, $\mathcal{L}_0$ is invariant under global transformations in flavour space and should also be invariant under local $SU(2)_L \otimes U(1)_Y$ gauge transformations. The symmetry group introduces 4 different gauge bosons, $W^\pm$, Z and γ. The covariant derivative is defined as follows:

$$D_\mu\psi_j \equiv [\partial_\mu ig\frac{\vec{\sigma}}{2}\vec{W}_\mu - ig'\frac{y_j}{2}B_\mu]\psi_j(x)$$

where $\vec{\sigma}$ are the usual Pauli matrices (generators of the $SU(2)$ group), $\vec{W}_\mu$ and B_μ are the gauge fields identified with $W^\pm$, Z and γ. The constants g and g'

are the interaction couplings and y_j are the weak hyper-charges. The complete Lagrangian is:

$$\mathcal{L} = \mathcal{L}_{fermionic} + \mathcal{L}_{kin} = \sum_{j=1}^{3} i\overline{\psi}_j(x)\gamma^\mu D_\mu \psi_j(x) - (\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W^i_{\mu\nu}W_i^{\mu\nu})$$

with the following notation for vector fields

$$B_{\mu\nu} \equiv \partial_\mu B_\nu - \partial_\nu B_\mu \qquad \vec{W_{\mu\nu}} = \partial_\mu \vec{W_\nu} - \partial_\nu \vec{W_\mu} + g\vec{W_\mu} \times \vec{W_\nu}$$

The coupling constant g appears both in the interaction part of the Lagrangian and in the kinetic term. The gauge symmetry forbids mass terms for the gauge bosons. Fermionic masses would also produce an explicit breaking of the gauge symmetry. Therefore, the $SU(2)_L \otimes U(1)_Y$ Lagrangian can only describe massless fields.

The term $\mathcal{L}_{fermionic}$ in the Lagrangian contains interactions of fermion fields with gauge bosons. This term can be divided into two parts: charged currents (involving $W^\pm$ bosons) and neutral currents (involving Z and γ bosons). In the case of charged currents, the charged boson fields are defined as follows:

$$W^\pm_\mu \equiv (W^1_\mu \pm iW^2_\mu)$$

In the case of neutral currents, i.e. interactions involving the neutral gauge fields W^3_μ and B_μ, identified with Z and γ bosons, the definition is:

$$\begin{pmatrix} W^3_\mu \\ B_\mu \end{pmatrix} \equiv \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix}$$

therefore the neutral gauge fields in the Lagrangian are a combination of Z and γ fields, involving a rotation angle θ_W, called the Weinberg angle. In order to recover QED, the following relations are obtained:

$$g\sin\theta_W = g'\cos\theta_W = e \qquad \frac{Y}{2} = Q - T_3$$

where T_3 is the third component of the weak isospin, Q the electric charge and Y the weak hyper-charge. The first equation relates the $SU(2)_L$ and $U(1)_Y$ couplings to the electromagnetic coupling, providing the desired unification of weak and electromagnetic interactions. The second identity relates fermion hyper-charges with the electric charge and the weak isospin, fixing in this way all couplings between fermions and gauge fields. Neutral currents were first observed at CERN in 1973 [2]. The ratio of neutral to charged current was measured, providing an estimation for the Weinberg angle. The value of this parameter is $sin\theta^2_W \approx 0.22$. Some years later, in 1982 and 1983, an additional success of the SM was the discovery of the gauge bosons W^+ and Z at the CERN SPS collider [3]. Table 1.3 shows the electroweak quantum numbers of fermions from the first generation.

leptons	T	T_3	Q	Y		quarks	T	T_3	Q	Y
ν_e	1/2	1/2	0	-1		u_L	1/2	1/2	2/3	1/3
e_L^-	1/2	-1/2	-1	-1		d_L	1/2	-1/2	-1/3	1/3
e_R^-	0	0	-1	-2		u_R	0	0	2/3	4/3
						d_R	0	0	-1/3	-2/3

Table 1.3: *Electroweak quantum numbers of leptons and quarks.*

1.1.4 Spontaneous Symmetry Breaking

As discussed before, the electroweak Lagrangian only allows massless gauge particles. This is the case of the photon field, but the gauge bosons $W^{\pm}$ and Z are on the contrary heavy particles. In order to generate masses, the gauge symmetry has to be broken in some way. However, a fully symmetric Lagrangian is also required to preserve renormalizability. A solution to this dilemma is provided by the Spontaneous Symmetry Breaking (SSB) mechanism discussed below.

In a Quantum Field Theory, the ground state is the vacuum. The SSB mechanism keeps a symmetric Lagrangian, but introduces a non-symmetric vacuum. For general continuous symmetries, the SSB mechanism implies however massless degrees of freedom (scalar particles), that are not observed experimentally. This result is known as the Goldstone theorem, and can be deduced using a toy model consisting in the following Lagrangian for a complex scalar field $\phi(x)$:

$$\mathcal{L} = \partial_\mu \phi^\dagger \partial^\mu \phi - V(\phi) \qquad V(\phi) = \mu^2 \phi^\dagger \phi + h(\phi^\dagger \phi)^2$$

that is invariant under global phase transformations of the scalar field. In order to have a stable ground state, the potential should be bounded from below, i.e. $h > 0$. The quadratic term $\mu^2 < 0$ implies that the minimum energy state is obtained for those field configurations satisfying

$$|\phi_0| \equiv \sqrt{\frac{-\mu^2}{2h}} \equiv \frac{v}{\sqrt{2}} > 0 \qquad V(\phi_0) = -\frac{h}{4} v^4$$

and as a result, there is an infinite number of degenerate states with minimum energy, $\phi_0(x) = \frac{v}{\sqrt{2}} e^{i\theta}$. By choosing a particular solution, $\theta = 0$ for example, the symmetry is spontaneously broken. If the excitations over the ground state are written as

$$\phi(x) = \frac{1}{\sqrt{2}}[v + \varphi_1(x) + i\varphi_2(x)]$$

where φ_1 and φ_2 are real fields, the potential transforms to

$$V(\phi) = V(\phi_0) - \mu^2 \varphi_1^2 + hv\varphi_1(\varphi_1^2 + \varphi_2^2) + \frac{h}{4}(\varphi_1^2 + \varphi_2^2)^2$$

Thus φ_1 describes a massive field of mass $m_{\varphi_1}^2 = -2\mu^2$, but φ_2 is a massless field, i.e. the Goldstone boson. These massless states can however be avoided in the case of local gauge symmetries, as discussed below.

We consider in the following a $SU(2)$ doublet of complex scalar fields:

$$\phi(x) = \begin{pmatrix} \phi^{(\dagger)}(x) \\ \phi^{(0)}(x) \end{pmatrix}$$

The gauged Lagrangian of the scalar fields is:

$$\mathcal{L}_S \equiv (D_\mu \phi)^\dagger D^\mu \phi - \mu^2 \phi^\dagger \phi - h(\phi^\dagger \phi)^2 \qquad h > 0, \mu^2 < 0$$

and this Lagrangian is invariant under local $SU(2)_L \otimes U(1)_Y$ transformations. As in the toy model discussed before, there is an infinite set of degenerate states with minimum energy, satisfying

$$< 0|\phi^{(0)}|0 >= \sqrt{\frac{-\mu^2}{2h}} = \frac{v}{\sqrt{2}}$$

Since the electric charge is a conserved quantity, only the neutral scalar field can acquire a vacuum expectation value, so the $SU(2)_L \otimes U(1)_Y$ symmetry is spontaneously broken to the electromagnetic subgroup $U(1)_{QED}$. This subgroup remains the only exact symmetry of the vacuum. According to the Goldstone theorem, 3 massless states should then appear, but these states can be reabsorbed by the gauge fields. Indeed, the scalar doublet can be written as

$$\phi(x) = exp\{i\frac{\sigma_i}{2}\theta^i(x)\}\frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}$$

with 4 real fields, $\theta^i(x)$ and $H(x)$. The crucial point is that the local $SU(2)_L$ invariance of the Lagrangian allows to rotate away any dependence on $\theta^i(x)$, and as a result the massless excitations are unphysical states. The scalar field $H(x)$ is the only remaining physical scalar state. The associated particle is called the Higgs Boson and is discussed in the next section. Besides this, the vacuum expectation value of the neutral scalar fields generates quadratic terms for $W^\pm$ and Z, i.e. these gauge bosons acquire a mass:

$$M_Z \cos\theta_W = M_W = \frac{1}{2} v g$$

The photon, on the contrary, remains massless, because $U(1)_{QED}$ is an unbroken symmetry of the model.

The SM including the SSB mechanism described before provides a very large number of theoretical predictions, in perfect agreement with experimental data. M_Z is predicted to be larger than M_W, in agreement with the measured masses

$$M_Z = 91.1875 \pm 0.0021\ GeV \qquad M_W = 80.425 \pm 0.034\ GeV$$

From these experimental values, the electroweak mixing angle is deduced:

$$\sin^2\theta_W = 1 - \frac{M_W^2}{M_Z^2} = 0.222$$

An independent estimation of $\sin^2\theta_W$ can be obtained from the decay $\mu^- \to e^-\overline{\nu}_e\nu_\mu$:

$$\frac{g^2}{M_W^2 - q^2} \approx \frac{g^2}{M_W^2} = \frac{4\pi\alpha}{\sin^2\theta_W M_W^2} \equiv 4\sqrt{2}G_F$$

The measured muon lifetime, $\tau_\mu = (2.19703 \pm 0.00004)10^{-6}\ s$ provides a very precise determination of the Fermi coupling constant G_F:

$$G_F = (1.16637 \pm 0.00001)10^{-5} GeV^{-2}$$

The measured values of $\alpha^{-1} = 137.03599911(46)$, M_W and G_F imply

$$\sin^2\theta_W = 0.215,$$

in good agreement with the value obtained from the boson masses. The agreement can be improved further by introducing radiative corrections. The Fermi constant provides also a direct determination of the electroweak scale, i.e. the scalar vacuum expectation value (also known as Fermi scale):

$$v = (\sqrt{2}G_F)^{-\frac{1}{2}} = 246\ GeV$$

1.1.5 The Higgs Boson

The Higgs boson is the only particle predicted by the SM yet undiscovered. As discussed before, the Higgs boson provides a mechanism to generate masses for the fundamental particles. Its discovery is difficult for the following reasons:

- The mass m_H is not predicted by the SM. The only restrictions from theory (unitarity, triviality, and vacuum stability) do not provide stringent bounds. There are also limits from experimental data. Direct searches have been performed at LEP [4] using the production channel:

  $$e^+e^- \to Z^* \to ZH$$

  The absence of any significant signal provides a lower bound for the Higgs mass:

  $$m_H > 114\ GeV$$

  Experimental upper bounds are also derived from the measurement of observables sensitive to quantum loops. These observables depend logarithmically on m_H via radiative corrections. The following upper limit is obtained:

  $$m_H < 219\ GeV$$

- Production cross-sections are very small and highly dependent on the mass. The main production channels at LHC are shown in figure 1.1:
  (a) gluon-gluon fusion $(gg \to H)$,

(b) W and Z fusion ($qq \to qqH$),
(c) $t\overline{t}$ and $b\overline{b}$ fusion ($gg \to q\overline{q}H$) and
(d) Higgsstrahlung ($q\overline{q} \to Z(W)H$).

In figure 1.2(left) cross sections of all these mechanisms as a function of the Higgs mass are shown, for a center-of-mass energy of $\sqrt{s} = 14\ TeV$. In the mass range allowed by experimental data, the dominant production channel is gluon-gluon fusion. The associated production with W, Z and $t\overline{t}$ pairs provides also an important contribution for $m_H < 200\ GeV$.

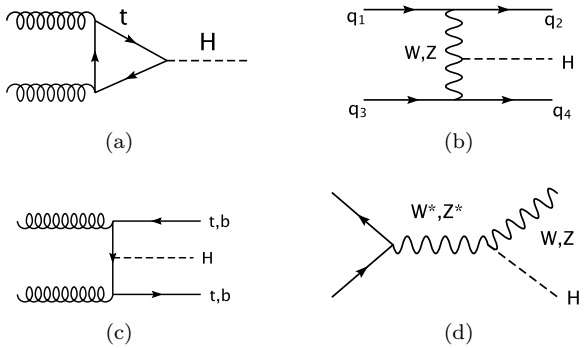

Figure 1.1: *The most important Higgs production mechanisms at LHC (pp collisions): (a) gluon-gluon fusion ($gg \to H$), (b) W and Z fusion ($qq \to qqH$), (c) $t\overline{t}$ and $b\overline{b}$ fusion ($gg \to q\overline{q}H$) and (d) Higgsstrahlung ($q\overline{q} \to Z(W)H$).*

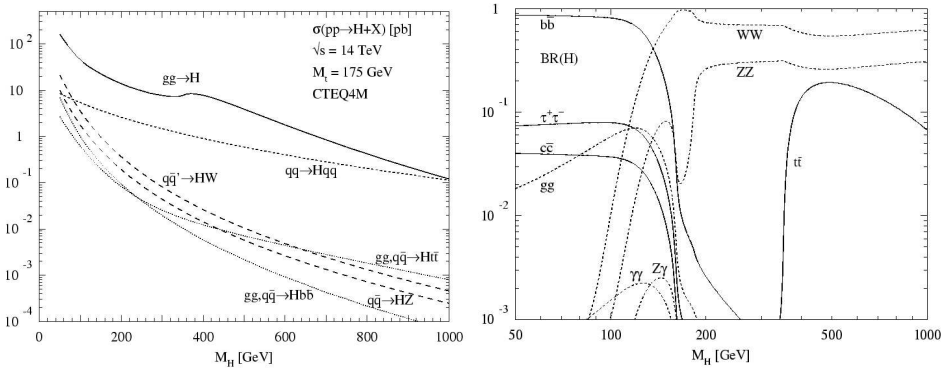

Figure 1.2: *Higgs production cross sections at the LHC ($\sqrt{s} = 14\ TeV$) for the various production mechanisms as a function of the Higgs mass (left) and branching ratios for the dominant decay modes of the SM Higgs particle as function of Higgs mass (right).*

In figure 1.2 (right) the branching ratios (BR) of the Higgs boson are plotted as a function of Higgs mass. They are again strongly dependent on the Higgs mass [5]. The various mass regions are discussed below.

- ***Low mass region*** ($114\ GeV < m_H < 130\ GeV$):
  the dominant decay channel in this range is $H \to b\overline{b}$, with nearly 100% rate. In hadronic collisions, the $H \to b\overline{b}$ decay suffers however from a very large background and the discovery is nearly impossible. the Higgs may also be produced associated with other particles ($t\overline{t}H$, WH, ZH) with a smaller cross section but also reduced backgrounds, so the search is possible. Another relevant channel for the Higgs discovery is $H \to \gamma\gamma$. The signature of this channel is very simple, two high-p_T photons, but this channel requires a very performing electromagnetic calorimeter in order to reject the very large background.

- ***Intermediate mass region*** ($130\ GeV < m_H < 180\ GeV$):
  the channels $H \to ZZ^{(*)}$ and $H \to WW^{(*)}$ have large branching ratios in this mass region. In addition, the gauge bosons decaying leptonically provide signatures with low background. Therefore these are the most advantageous channels in this mass range.

- ***High mass region*** ($180\ GeV < m_H < 219\ GeV$):
  this is the most favourable region to discover a Higgs boson signal, since the channel $H \to ZZ \to 4l$ is almost background free. Each lepton pair has an invariant mass equal to the Z mass and the irreducible background $pp \to ZZ \to 4l$ is much smaller than the signal.

1.1.6 Beyond the Standard Model

The SM is a very successful theory, as discussed before, but unfortunately incomplete. There are both conceptual problems and phenomenological indications for physics Beyond the SM (BSM). On the conceptual side the most obvious problems are quantum gravity and the related hierarchy problem. Among the main phenomenological hints for new physics we can mention coupling unification, dark matter, neutrino masses, baryogenesis and the cosmological constant. Some of these problems are discussed below [6, 7].

- **Problems outside the SM**

  Many gravity/cosmology results cannot be explained within the SM. Grand Unification theory is not included in the SM, since couplings do not converge at high energies. A non-zero cosmological constant, Λ, has been measured using a large number of experimental data. The cosmological constant accounts for a mysterious *dark energy* and represents more than 70 % of the total energy of the universe. The matter energy is the remaining 30%, and the SM cannot explain this large amount of matter either. Other important questions from cosmology are: the observed matter-antimatter asymmetry related to CP-violation, inflation, etc.

- **Problems within the SM**

  Within the SM, neutrinos are massless. In 1998, neutrino oscillations were observed, implying non-zero neutrino masses. Neither masses nor oscillations are predicted by the SM.

  There is a large number of constants in the SM and the origin of them is unexplained. Most masses are too small compared with the vacuum expectation value (v) or Fermi Scale. The origin of Cabbibo-Kobayashi-Maskawa (CKM) mixing angles is also unexplained.

  One of the main problems of the SM is the so-called 'hierarchy problem'. The SM has an intrinsic scale v which is insignificant compared to Planck Mass ($M_{Pl} \sim 10^{19}\ GeV$). There is also no reason for the Higgs boson to have such a low mass. The most important argument against a low Higgs mass are radiative corrections. These corrections arise, in order of relevance, from *top*, gauge and Higgs loops. Radiative corrections to the Higgs propagator are shown in figure 1.3(left).

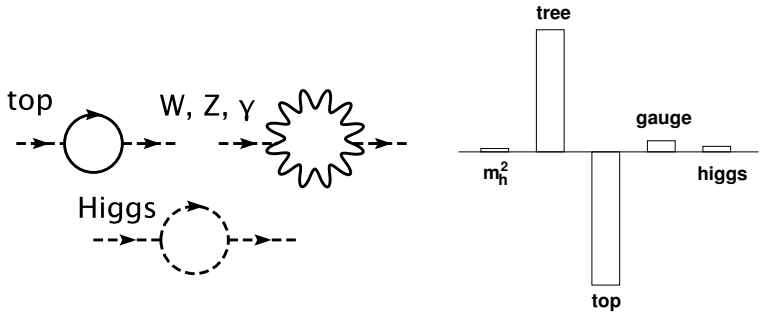

Figure 1.3: *The most significant loop contributions to the Higgs mass in the SM (left) and fine tuning required to obtain a low Higgs mass in the Standard Model with a cut-off of 10 TeV (right).*

Assuming that the SM is valid up to 10 TeV (which is the maximum energy range for LHC), loop corrections to the Higgs mass m_H^2 have the following contributions:

$$-\frac{3}{8\pi^2}\lambda_t^2\Lambda^2 \sim -(2\ TeV)^2$$

$$\frac{1}{16\pi^2}g^2\Lambda^2 \sim (700\ GeV)^2$$

$$\frac{1}{16\pi^2}\lambda^2\Lambda^2 \sim (500\ GeV)^2$$

from *top* loop, gauge loop and Higgs loop, respectively. The expected radiatively corrected mass of the Higgs boson is therefore:

$$m_h^2 = m_{tree}^2 - [100 - 10 - 5](200\ GeV)^2$$

In order to obtain a Higgs mass of a few hundred GeV, as required by data, fine tuning of parameters is needed to the 1% level. This is the hierarchy problem. Figure 1.3 (right) shows the level of tuning necessary to obtain a low Higgs mass. Since no fine tuning is required if $\Lambda \leq 2\ TeV$, the SM is valid up to this limit, implying new physics should appear in the range accessible to LHC.

In the following we discuss briefly a few theories proposed as possible extensions of the SM. These theories try to address in particular the hierarchy problem presented before.

Supersymmetry (SUSY)

The basic idea of SUSY is to introduce a new symmetry between fermions and bosons. SUSY predicts that every particle has a supersymmetric partner or *s-particle*. This partner has the opposite spin-statistics, but the same quantum numbers as the SM particle. In the limit of exact boson-fermion symmetry, the quadratic divergences of bosons loops are cancelled by fermion loops, so only logarithmic divergences remain. In particular, the top loop contributing to the Higgs mass is quenched by a partial cancellation with s-top exchange (the s-top being the scalar partner of the top quark). However, as no SUSY particles have been observed yet, the exact model must be broken in some way. The simplest SUSY theory including the SM is the MSSM or Minimal Supersymmetric Standard Model. MSSM introduces more than 100 extra parameters, 5 Higgs bosons and a super-partner for each SM particle. In SUSY, the gauge couplings of QCD and electroweak interactions converge at an energy scale of about $10^{16}\ GeV$, solving in this way the SM problem of gauge coupling unification. Finally, the SUSY particle with smallest energy (LSP) is expected to be neutral and stable. This particle is therefore an excellent candidate to solve the dark matter problem observed in cosmology.

Technicolour

According to this theory, the Higgs boson is a condensate of new fermions. There is no fundamental scalar Higgs sector, hence no quadratic divergences. This mechanism needs a very strong binding force, $\Lambda_{TC} \sim 10^3 \Lambda_{QCD}$. It is difficult however to explain why such a strong force is not showing up in precision electroweak tests. Hence Technicolour is disfavoured by LEP data, although not all models can be totally discarded.

Large compactified extra dimensions

The idea is that the reason why M_{Pl} appears very large, so gravity is very weak, is simply due to hidden extra dimensions. This possibility is very exciting and has attracted recently a lot of interest. It remains to be seen if these theories are able to solve the problems of the SM, still being compatible with experimental data.

Little Higgs models

These theories solve the hierarchy problem and the fine tuning by cancelling the corrections to the Higgs mass with new particles of the same spin. This cancellation allows Little Higgs models to be valid up to a scale of 10 TeV without fine tuning. In these models, the SM Higgs is a pseudo-Goldstone boson resulting from the breaking of an extra global symmetry. They are described in more detail in chapter 3.

1.2 The LHC and the ATLAS experiment

The Large Hadron Collider (LHC) is currently being installed in the LEP tunnel at the European Centre for Nuclear Research (CERN) in Geneva, Switzerland. It will accelerate protons to an unprecedented energy of 7 TeV. Two beams of opposite orientation will be made to collide at four interaction points, with a center-of-mass energy of 14 TeV, each of which will be equipped with a large experimental apparatus: ATLAS, CMS, LHCb and ALICE.

The LHC machine allows for a broad and ambitious physics program, such as physics in the Standard and beyond the Standard Model.

ATLAS (A Toroidal LHC ApparatuS) is a general purpose experiment, currently being installed at the LHC, designed to search for new physics and to perform precision measurements.

In this section, an overview of the LHC accelerator and the ATLAS experiment are given in sections 1.2.1 and 1.2.2, respectively.

A description of the ATLAS main detector subsystems: The calorimetry, the Inner Detector and the Muon system are shown in sections 1.2.3, 1.2.4 and 1.2.5, respectively.

Finally, the Magnet system is also discussed in section 1.2.6.

1.2.1 The Large Hadron Collider (LHC)

The Large Hadron Collider (LHC) is a proton-proton and heavy-ion collider with a center-of-mass energy of 14 TeV, located at CERN (Geneva). The accelerator is currently being installed in the LEP tunnel. The LEP accelerator reached a center-of-mass energy of 209 GeV using electron-positron collisions. Proton-proton collisions can reach more energy than electron-positron collisions because a much smaller fraction of the beam energy is lost into synchrotron radiation.

Other accelerators already existing at CERN are used to inject protons into the LHC, in order to reach the nominal center-of-mass energy of 14 TeV. As figure 1.4 (left) shows, the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS) are used to accelerate protons up to 450 GeV. Protons are then injected into the LHC, where they are are accelerated up to 7 TeV. Extremely large magnetic fields are needed to keep protons in orbit with an energy of 7 TeV inside the LEP/LHC circumference (27 km long). The 1232 superconducting dipoles installed inside the ring provide the 8.4 Tesla magnetic field needed for this purpose.

At the LHC, bunches of 10^{11} protons collide at four different interaction points, where the experiments are placed (see figure 1.4, right). These collisions have a frequency of 25 ns. The four experiments under construction at the LHC are: ATLAS, CMS, LHCb and ALICE. ATLAS and CMS are general purpose experiments, designed to search for new physics and to perform precision measurements. LHCb is a B physics and CP violation dedicated detector, and finally ALICE is a

heavy ion experiment designed to investigate nuclear matter at very high energies and densities. The research study presented in this thesis has been performed in the context of the ATLAS detector and simulation tools. A description of ATLAS is presented in the next section. More detailed descriptions of ATLAS and CMS can be found in [8] and [9], respectively.

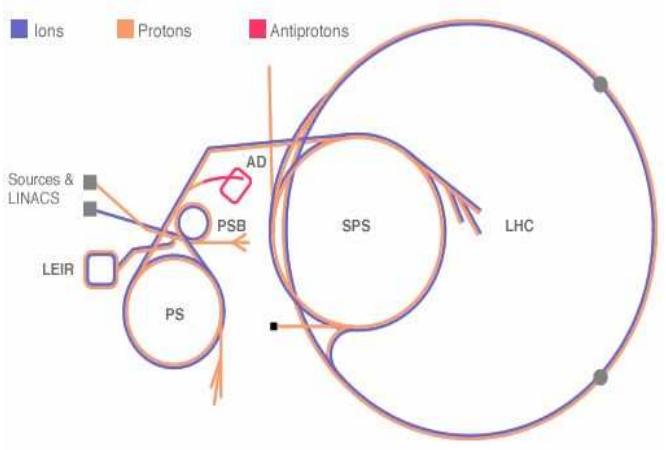

(a)

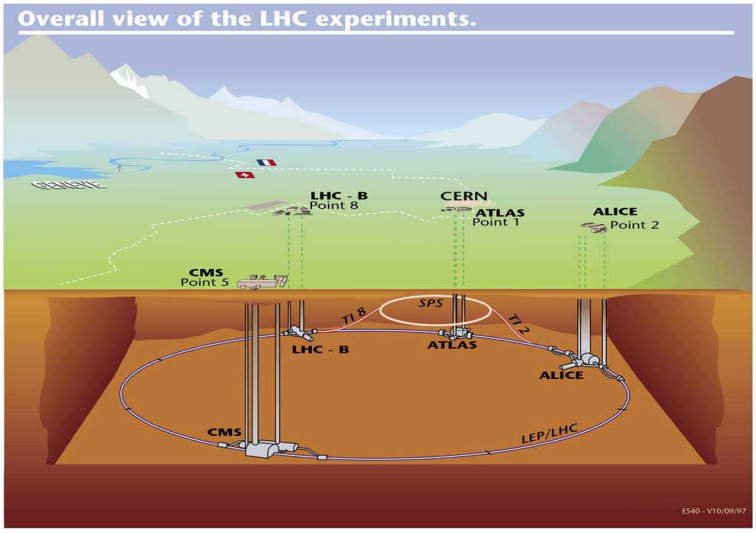

(b)

Figure 1.4: *(a) The LHC injection system and (b) The LHC experiments.*

Two operation phases are foreseen for the LHC. The nominal luminosity is expected to be $10^{33}\ cm^{-2}s^{-1}$ (low luminosity phase) during the first 3 years of operation, collecting an integrated luminosity of 10 fb^{-1}/yr for both ATLAS and

CMS. This integrated luminosity is comparable to the total integrated luminosity collected by the Tevatron in Run 2. Then, the luminosity should increase up to $10^{34}\ cm^{-2}s^{-1}$ (high luminosity phase) for the next years of operation, collecting 100 fb^{-1}/yr per experiment. The luminosity values quoted are just reference values currently used for ATLAS physics analysis.

The LHC machine allows for a broad and ambitious physics program. The main topics are briefly summarised below.

- Search for the Standard Model Higgs boson over the full allowed mass range, from the experimental LEP limit up to the theoretical upper bound of about 1 TeV. If the Higgs boson is found, it should be possible to measure its mass and some of its couplings with sufficient precision.
- Search for Supersymmetry and other physics models beyond the Standard Model, such as Extra Dimensions and Little Higgs models, Leptoquarks or models with additional leptons, quarks or gauge bosons, up to masses of the order of 5 TeV.
- Precise measurements of the W mass, of $WW\gamma$ and WWZ Triple Gauge Couplings, and of the mass, couplings and decay properties of the top quark. Several QCD measurements will also be possible. It will be possible for example to test the running of the strong coupling constant α_s over an unprecedented range of Q^2 values.
- Detailed studies of B-mesons and CP-violation. These measurements will also be performed (with better sensitivity) by the dedicated LHCb experiment [10]. In addition to the ATLAS and CMS studies (general purpose detectors).
- Studies of the phase transition from hadronic matter to a plasma of unconfined quarks and gluons. These studies will be performed by the dedicated ALICE experiment [11]. In addition to the ATLAS and CMS studies (general purpose detectors).

At high luminosity the expected event rate is $\simeq 10^9\ events/s$. The observed events can be classified as follows:

- *soft collisions*: they are due to long-distance interactions between the two incoming protons. The out-coming particles from soft collisions have large longitudinal momentum but a small transverse momentum, with average p_T of a few MeV. These events, also called *minimum bias* events, represent by far the majority of the collisions.
- *hard collisions*: they are due to short range interactions between two partons from the incoming protons. In these interactions the momentum transfer is large, allowing the production of final particles with high-p_T and the production of massive new particles as well.

At the LHC, bunches of about 10^{11} protons collide at each interaction point every 25 ns. several soft collisions occur at each bunch crossing, producing on the average a total number of $1,000$ charged particles in the region of $|\eta| < 2.5$. As a result, interesting high-p_T events are overlapped by some 25 soft interactions, forming what is usually called the *pile-up*. The detector parameters have been carefully adjusted in order to reduce the impact of pile-up on the various physics analyses.

1.2.2 The ATLAS detector

A Toroidal LHC ApparatuS (ATLAS) detector is currently under construction in one of the experimental areas of the LHC. The commissioning phase for some of the components has already started using cosmics. The overall ATLAS detector layout, shown in figure 1.5, is extremely complex and described in detail in the Technical Design Report [12].

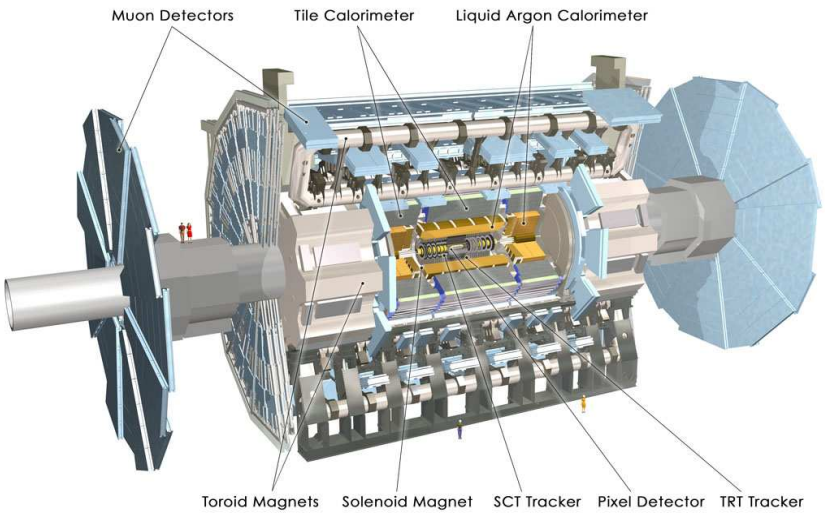

Figure 1.5: *The ATLAS detector.*

The basic design criteria of the detector are summarized below.

- Very good electromagnetic calorimetry for electron and photon identification and measurements, complemented by a full-coverage hadronic calorimetry for accurate jet and missing transverse energy (E_T^{miss}) measurements.
- High-precision muon momentum measurements, with the capability of accurate measurements at the highest luminosity using the external muon spectrometer alone.
- Efficient tracking at high luminosity for high-p_T lepton-momentum measurements, electron and photon identification, τ-lepton and heavy-flavour identification, and full event reconstruction capability at lower luminosity.

- Large acceptance in pseudo-rapidity (η) with almost full azimuthal angle (ϕ) coverage everywhere. The azimuthal angle is measured around the beam axis, whereas pseudo-rapidity is related to the polar angle θ, i.e. the angle from the beam axis.

The ATLAS detector is contained within a cylinder of about 44 m length and 22 m in diameter. The overall weight of the detector is $\sim 7,000$ tons. Like most collider detectors, ATLAS is designed as an onion-like structure. It can be split into three major sub-systems. Starting from the interaction point, these subsystems are: the inner detector, the calorimeters and the muon spectrometers. The trigger and data acquisition system can be considered as a fourth component of the detector.

1.2.3 Calorimetry

Calorimeters are usually divided into electromagnetic and hadronic sections. The ATLAS calorimeter consists of an electromagnetic (EM) calorimeter covering the pseudo-rapidity region $|\eta| < 3.2$, a hadronic barrel calorimeter covering $|\eta| < 1.7$, hadronic end-cap calorimeters covering $1.5 < |\eta| < 3.2$, and finally forward calorimeters covering $3.1 < |\eta| < 4.9$. Identification of electrons and photons are very important issues in calorimetry. The rate of QCD jets is very high and in order to be able to obtain clean samples of electrons with transverse momentum above 20 GeV, a rejection of QCD jets at the level of 10^6 is required. Such a rejection can be achieved with very good granularity in both the EM and hadronic part, in order to identify isolated energy deposits from electrons/photons and to allow a veto on hadronic energy collected behind the electromagnetic clusters.

The design of the EM calorimeter has been optimized for the analysis of Higgs processes involving decays to electrons and photons, $H \to ZZ^* \to 4e$ and $H \to \gamma\gamma$. The latter decay has a very large background and a mass resolution of 1% is required [13]. This implies that the constant term of the energy resolution remains smaller than 1%. The sampling term, on the other hand, is kept at the level of $0.1/\sqrt{E(GeV)}$. The dynamic energy range covered by the calorimeter is from 1 GeV electrons from B-meson decays, up to a few TeV for the decay of heavy new vector bosons.

The reconstruction of jets is an important issue as well, setting stringent limits on the linearity of the hadronic calorimeter response at high energies. The reconstruction of W bosons in hadronic decay modes requires in addition a good jet energy resolution. The discovery of a high mass Higgs boson decaying into high p_T bosons, $W \to jj$, requires a good jet-jet invariant mass reconstruction as well. Finally, forward jet tagging and energy measurements are important in order to calculate the missing transverse energy.

A measurement of the missing transverse energy is a way to detect particles escaping the detector without interactions. These can be either neutrinos or stable supersymmetric particles. For example, MSSM decays of $H_0 \to \tau\tau$ require good p_T^{miss} resolution and E_T^{miss} reconstruction. To identify missing transverse energy, the calorimeter needs a good hermicity which means a pseudo-rapidity coverage

up to $|\eta| = 5$ and no cracks in the detector for cables or any other purpose. Finally, the hadronic calorimeter needs to be thick enough to contain the total energy of hadrons. A good energy containment and resolution is achieved with 11 absorption lengths of material in front of the muon system.

Electromagnetic calorimeter

The EM calorimeter is a lead liquid argon (LAr) sampling calorimeter. Layers of lead/stainless steel absorber plates and LAr are inter-spaced. The lead ensures the shower development due to its short radiation length and the secondary electron ionization is measured in the narrow gaps of liquid argon. An inductive signal from the ionisation electrons drifting in the electric field across the gaps is registered by copper electrodes. To achieve a low capacitance of the detecting elements and thereby a fast signal, the lead plates have an accordion shape as shown in figure 1.6(a).

Hadronic calorimeter

The hadronic calorimeter of ATLAS is also of sampling type and it is divided in barrel and end-cap. The barrel, also called tile calorimeter, uses iron as absorber and scintillating plastic tiles as active material. The end-cap uses copper as absorber and LAr as active material. A module from the barrel part is shown in figure 1.6(b). The calorimeter in the central pseudo-rapidity region (barrel region) covers up to $|\eta| = 1.7$. The resolution of the hadronic calorimeter is $0.5/\sqrt{E} + 3\%$ for $|\eta| < 3$ and $1/\sqrt{E} + 10\%$ for $3 < |\eta| < 5$, where E is the energy in GeV. The scintillator tiles of the hadron calorimeter are placed in the direction of the incoming particles. The light created in the scintillators is read out via wavelength shifting fibres by photo-multipliers placed on the outside of the calorimeter.

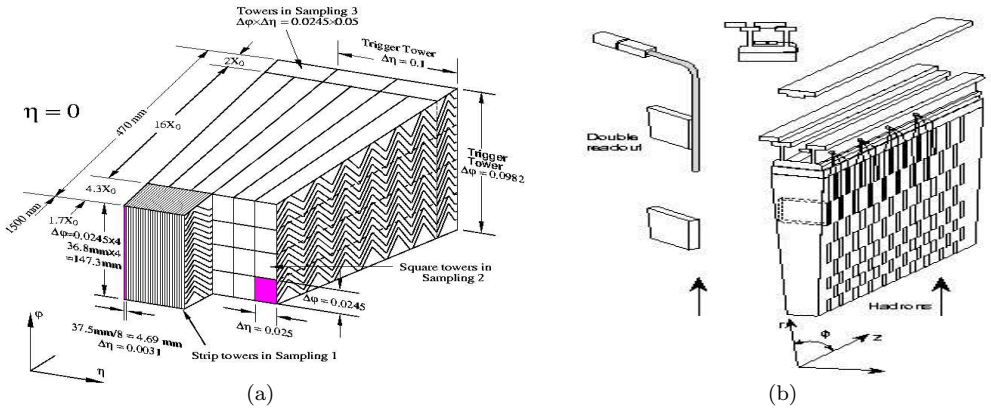

Figure 1.6: *Electromagnetic (a) and hadronic (b) calorimeter structures.*

1.2.4 The Inner Detector

The Inner Detector (ID) is the central part of ATLAS (see figure 1.7). With a combination of high-precision and high-granularity layers in the inner part, and straw tubes in the outer part, it can reconstruct charged particle tracks inside a solenoidal magnetic field of 2 T. The coverage of the ID extends up to pseudo-rapidity values of $\eta = 2.5$.

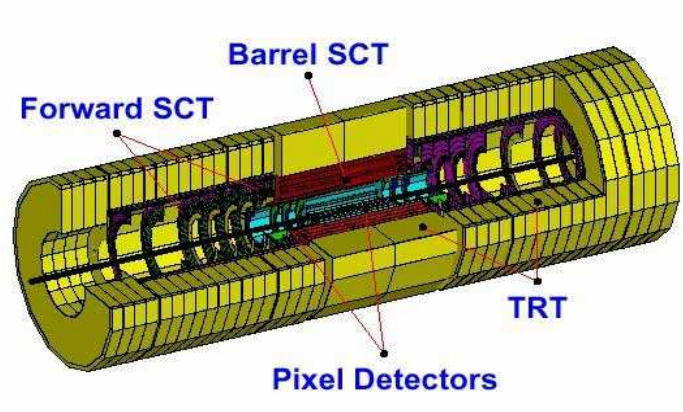

Figure 1.7: *ATLAS Inner Detector.*

The main requirements for the Inner Detector are:

- Tracking efficiency of more than 95% over the full coverage for isolated tracks with $p_T > 5$ GeV, with fake-track rates smaller than 1%.
- Identification of individual particles inside jets, wherever the calorimeter cannot resolve isolated tracks.
- Momentum measurement over a large range. Below $p_T = 0.5$ GeV the particles spiral inside the ID in the magnetic field and reconstruction is difficult. This lower limit affects the reconstruction of electrons from converted photons as well.
- Separate electrons and photons that create identical clusters in the EM calorimeter.
- Charge identification of particles with very large transverse momentum, originating from the decay of massive resonances.
- Decay length reconstruction for CP-violation studies in the B-system and B^0 mixing measurements.
- Tagging of jets originating from high energy b-quarks. The tagging is performed using impact parameter or secondary vertex techniques.

- Momentum measurement of low energy muons with large scattering angles in the hadronic calorimeter.
- Electron/jet separation, in addition to the separation provided by the calorimeter.
- Identification of the primary event vertex, overlapping with the many other vertexes from minimum bias events.

The momentum and vertex resolution requirements call for high-precision measurements to be made with fine-granularity detectors. Semiconductor tracking detectors, using silicon micro-strip (SCT) and pixel technologies are used. The highest granularity is achieved around the vertex region using semi-conductor pixel detectors (PIX). The total number of precision layers must be limited to avoid dead material and a high cost. Photon conversions, bremsstrahlung from electrons and pions from nuclear interactions with the detector material cause a degraded ID and calorimeter performance. ATLAS is designed in a way that three pixel layers and eight silicon strip layers are crossed by most tracks. A large number of tracking points (typically 36 per track) is provided by the straw tube tracker (TRT) in the outer part. The TRT provides almost continuous tracking with much less material per point and at lower cost. The TRT provides in addition electron identification. The combination of the various techniques gives very robust pattern recognition and high precision in both η and ϕ. In the barrel part of the Inner Detector, where $|\eta| < 1$, all of the detecting elements are arranged in cylindrical structures, while the two end-caps have the detecting elements arranged in wheels.

The Silicon Detectors

Two technologies are used for the ATLAS tracker: pixel detectors (PIX) placed close to the beam pipe and silicon strip detectors (SCT) placed further away. The SCT system is designed to provide eight precision measurements per track, contributing to the measurement of momentum, impact parameter and vertex position. It allows in addition a good pattern recognition due to its high granularity. The silicon strip detectors have an n-type bulk with a single sided readout of n+ strips. This choice has been made because of its radiation hardness to be very radiation resistant. In each layer, two single sided detectors are placed back to back. The micro-strips are parallel to the beam pipe, thus measuring the ϕ coordinate directly. The strips on the back side reconstruct the z(r) coordinate in the barrel (end-cap). Each strip has a length of 12 cm and a width of 80 μm.

The PIX system is designed to provide a very high-granularity, and thus high-precision hit measurements close to the interaction point. The system provides three precision measurements over the full acceptance, and determines mostly the impact parameter resolution. The system includes 140 million square detecting elements, each 50 μm in the $R\phi$ direction and 400 μm in z, thus providing a 3-dimensional coordinate on each layer. The best resolution is in the ϕ direction. The most important task of the pixel detectors is the determination of secondary

vertexes for the identification of B-mesons and b-tagging. It is also important for pattern recognition due to its low occupancy resulting from the high granularity. The first pixel layer is placed as close as possible to the beam pipe to optimize impact parameter measurements.

The Transition Radiation Tracker

The transition radiation tracker (TRT) is based on the use of straw detectors, operating at the very high rates expected at the LHC. This is possible by virtue of their small diameter and the the isolation of individual wires. Electron identification is ensured by employing xenon gas to detect transition-radiation photons created in a radiator between the straws. This technique is intrinsically radiation hard, and allows a large number of measurements, typically 36, to be made on every track at a modest cost. However, the detector must cope with a high occupancy and high counting rates at the LHC design luminosity. The large number of hits on the tracks is useful for pattern recognition as well. In total there are around $370,000$ straws with 4 mm diameter. They are placed radially in the end-caps and along the beam axis in the barrel region. These orientations are chosen to maximise the number of straws crossed by tracks originated in the interaction point. A 30 μm gold covered tungsten wire is placed in the centre of each straw.

1.2.5 The Muon System

The muon system serves a double purpose: a trigger to select events with high energy muons and a precision muon spectrometer. The general layout of the muon system is shown in figure 1.8.

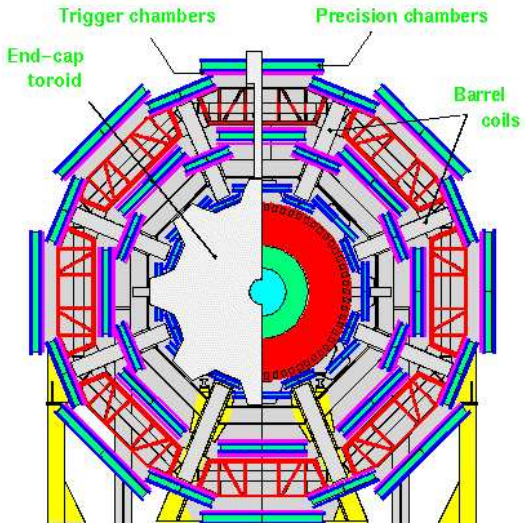

Figure 1.8: *ATLAS Muon spectrometer.*

The muon system has to fulfill the following requirements:

- A good transverse momentum resolution in the low p_T region. The limit is defined by the ability to detect $H \rightarrow ZZ^*$ decays into muons with a high suppression of the background. A resolution of around 1% is required.
- Sufficient resolution at high p_T for good charge identification of muons from eventual decays $Z' \rightarrow \mu\mu$, where Z' is a very massive resonance.
- A pseudo-rapidity coverage of $\eta < 3$ and an hermetic system to prevent particles to escape through holes.
- Measurement of spatial coordinates in two dimensions to provide good invariant mass resolution.
- A low rate of both punch through hadrons and fake tracks.

Particles from the primary vertex traverse three sets of muon chambers. In the barrel part of the detector the chambers are mounted on the vacuum vessels (cryostats) of the Barrel Toroid magnet system. The momentum is measured from the sagitta of the tracks. In the end-caps, the chambers are placed perpendicular to the beam axis and the muon deflection is provided by the magnetic field of the End-Cap Toroid (ECT) magnet systems.

The Monitored Drift Tube (MDT) chambers are proportional chambers made out of aluminum tubes of 30 mm diameter and length varying from 70 cm to 630 cm. To measure the coordinate in the bending plane of the magnet, the tubes are placed transverse to the beam axis. In order to reduce the level of fake tracks reconstructed from random associations of background hits, each set of MDTs contains 2 super-layers with 3 or 4 layers of tubes. Each MDT has a resolution of 80 μm which results in a momentum resolution $\Delta p_T/p_T < 10^{-4} \cdot p$ for tracks with $p_T < 300$ GeV. Cathode Strip Chambers (CSC) with finer granularity compared to the MDT, are used in the forward region, where the track density is higher. The CSCs are multi-wire proportional chambers with a wire spacing of 2.5 mm.

An important design criterion of the muon system was the capability to trigger on muon tracks. The precision chambers had therefore to be complemented by a system of fast chambers, capable to deliver track information a few tens of nanoseconds after the passage of the particle. In the barrel region Resistive Plate Chambers (RPC) were selected for this purpose, while in the end-cap Thin Gap Chambers (TGC) are used. Both chamber types deliver signals inside 15-20 ns, being able to tag the beam crossing.

The drift tubes have a large diameter which results in a maximum drift-time of 480 ns, much longer than the 25 ns spacing between the bunch crossings. For this reason additional layers of chambers are needed for the trigger. Only the trigger chambers are read out by the Level-1 trigger. The lower spacial resolution of the trigger chambers provide a reduced but sufficient momentum resolution. The MDTs provide only the η coordinate, but the trigger chambers also provide the ϕ coordinate of muon tracks.

1.2.6 The Magnet System

The ATLAS superconducting magnet system is an arrangement of a central solenoid providing a magnetic field for the ID, surrounded by a system of large air-core toroids generating the magnetic field for the muon spectrometer.

The solenoid magnet is placed before the electromagnetic calorimeter. The major problem is the increased amount of material in front of the calorimeter which causes many particles to start showering before they reach the active part of the calorimeter. The central solenoid is a superconducting magnet cooled by helium at 4.5 K. To reduce the material, the magnet shares the cryostat with the liquid argon calorimeter. The length is considerably shorter than the inner tracking system. This is the result of a compromise: a short coil reduces the material in front of the calorimeter and a long coil makes the magnetic field more uniform in the Inner Detector. The magnetic field along the z-direction is 2 T at the interaction point.

The toroid magnet is divided into one barrel part and two forward regions. There are three toroids, each consisting of eight coils assembled radially and symmetrically around the beam axis. The toroid coils are of a flat racetrack type with two double-pancake windings made of 20.5 kA aluminum-stabilized NbTi superconductor. In each forward region the magnetic field is provided by eight superconducting coils placed in a huge, gear-wheel shaped vacuum vessel (cryostat). With this toroidal arrangement, tracks are almost perpendicular to the magnetic field, over the complete pseudo-rapidity range. The low number of coils used to form the toroid field results in a field strength that varies strongly with the ϕ coordinate. The field is 4 T at the superconductors for both barrel and end-caps toroids.

1.3 ATLAS software and computing

The Large Hadron Collider (LHC) starts to operate in 2008 and will produce roughly 15 Peta-bytes of data annually, that thousands of scientists around the world will access and analyse. The mission of the Worldwide LHC Computing Grid (WLCG) project is to develop, build and maintain a distributed computing infrastructure for the storage and analysis of data from the four LHC experiments.

The ATLAS Computing Model defines the requirements on the architecture and performance of the software and the overall computing system.

In this section, an overview of the ATLAS software: Event generators, full simulation, reconstruction and fast simulation are given in sections 1.3.1, 1.3.2, 1.3.3 and 1.3.4, respectively.

Finally, a brief description of the WLCG project and the ATLAS Computing Model are shown in sections 1.3.5 and 1.3.6, respectively.

1.3.1 Event generators

There are several available Monte Carlo event generators for pp collisions, such as HERWIG [14], ISAJET [15] and PYTHIA [16]. Each of these simulates a hadronic final state corresponding to some particular model of the underlying physics. The details of the implementation of the physics are different in each of these generators, however the underlying philosophy of the generators is the same.

The basic process is a parton interaction involving a quark or gluon from each of the incoming protons. Elementary particles in the final state, such as quarks, gluons or W/Z bosons, emerge from the interaction. The fundamental process is calculated in perturbative QCD, and the initial momentum of the quarks or gluons is given by structure functions.

Additional QCD (gluon) radiation takes place from the quarks and gluons that participate in the basic scattering process. These parton showers are based on the expansions around the soft and collinear limits and can be ascribed to either the initial or final state. The algorithm used by HERWIG includes some effects due to quantum interference and generally produces better agreement with the data when detailed jet properties are studied. The showering continues down to some low energy cut-off. For some particular cases the matrix element calculations involving higher-order QCD processes are used. The events that have more energy in the parton process have more showering, and consequently more jet activity.

The collection of quarks and gluons must then be hadronized into mesons and baryons. This is done differently in each of the event generators, but is described by a set of fragmentation parameters that must be adjusted to agree with experimental results. HERWIG looks for colour singlet collections of quarks and gluons with low invariant mass and groups them together. This set then turns into hadrons. PYTHIA splits gluons into quark-antiquark pairs and turns the resulting set of colour singlet quark-antiquark pairs into hadrons via a string model.

ISAJET simply fragments each quark independently paying no attention to the colour flow.

Matrix elements are likely to provide a better description of the main character of the events, i.e. the topology of well separated jets, while parton showers should be better at describing the internal structure of these jets. The above models describe events where there is a hard-scattering of the incoming partons: either a heavy particle is produced or the outgoing partons have large transverse momentum. While these are the processes that are of most interest, the dominant cross-section at the LHC consists of events with no hard scattering. There is little detailed theoretical understanding of these minimum-bias events and the event generators must rely on data at current energies.These minimum-bias events are important at LHC, particularly at high luminosity, as they overlap interesting hard-scattering events such as the production of new particles. The generators use a different approach in this case. ISAJET uses a pomeron model that has some theoretical basis. HERWIG uses a parametrization of experimental data from proton colliders. PYTHIA uses a mini-jet model where the jet cross-section is used at very low transverse momenta, i.e the hard scattering process is extrapolated until it saturates the total cross-section. Whenever relevant, ATLAS has used the PYTHIA approach with dedicated modifications that agree with present data from the Tevatron. The multiplicity in minimum-bias events predicted by this approach is larger than that predicted by ISAJET or HERWIG.

Despite the huge efforts which has been put into developing physics generators for hadron colliders over the last years, the precision with which data from currently running colliders can be reproduced is not better than $10-30\%$, and in some cases is not better than a factor of two. In the work presented here, cross-sections are normally computed using leading-order QCD implemented in PYTHIA 6.2, using the CTEQ5L [17, 18] set of structure functions as a reference. Whenever better or more appropriate calculations are available, the production cross-section from PYTHIA is suitably normalized, or a different Monte Carlo generator is used.

1.3.2 Full Simulation

The complexity of the physics events to be analyzed at the LHC and the diversity of the detectors to be integrated into ATLAS imply an absolute necessity to provide an accurate detector simulation, in order to evaluate the detector and physics performance in detail. The ATLAS simulation program has been developed continuously since 1990, using initially the GEANT-3 package [19], written in FORTRAN, and since 2,000 the GEANT-4 package [20] written in C++. All simulated data up to the so-called DC1 (Data Challenge 1) [21] were produced using the GEANT-3 package, then the switch over to the GEANT-4 package took place. The ATLAS simulation program can be logically divided into three separate modules: event generation, detector simulation and digitization. These three parts communicate through a set of data structures (ZEBRA banks for the FORTRAN version, and Object Oriented structures for the C++ version).

In addition to describing the detector geometry and tracking particles through it, the GEANT framework is used to describe the materials constituting the detector, to visualize the detector components, and to simulate and record the response of the sensitive elements of the various systems. When simulating a very complex detector such as ATLAS, the simulation software has to take into account as many physics processes as possible, covering the broadest possible range of energies. Ideally, the program should be able to simulate physics processes with energies as low as 10 eV (*e.g.* ionization potential in the active gas of various detectors) and as high as a few TeV (*e.g.* catastrophic energy losses of muons traversing the calorimeters). The most challenging task in terms of consumption of CPU time is an accurate simulation of showers in the calorimeters.

The full simulation studies performed in this thesis are in chapter 2. The same physics channel has been simulated twice, the first one using GEANT-3/DICE for an official sample from DC1 production and the second one using GEANT-4 for a private sample from Rome production [22]. Both simulations have been used in the various analysis described later.

The description of the ATLAS geometry in GEANT is probably the most critical issue for the detector simulation program, since it must represent the right compromise between accuracy and performance. All the ATLAS sub-detectors have been described to a very high level of detail in DICE [23]. Inactive material (cryostats, support structures, services, etc) has been given a great emphasis, and an accurate description of both its layout and material distribution has been introduced wherever possible. These inactive parts of the detector have often been shown to have a direct impact on the physics performance of the experiment and therefore have to be evaluated very carefully. The numbers of GEANT volumes used to describe the geometry of the various ATLAS systems are shown in table 1.4, where they can be compared to the total number of active detector elements or equivalently to the number of independent readout channels and to the total number of modules or chambers. In many systems each active cell is described explicitly, using sometimes up to ten volumes or more per cell. In contrast, details like the pixel or micro-strip structure of the silicon detectors, or the cell structure of the barrel accordion calorimeter, have not been described in the geometry itself, but only introduced afterward at the digitization level in order to increase the flexibility and the performance of the simulation chain. However, GEANT-4 has a very detailed geometry description for all these detectors.

The event-generation phase is normally run separately in order to have a consistent input stream which can be used many times. The most widely used event generators (PYTHIA, HERWIG, ISAJET) are interfaced to GEANT. Ad hoc single-particle generators and test-beam geometries have also been developed for detector-specific studies. Particle filter algorithms can be applied, in order not to track particles outside the geometrical acceptance of the detector or with an energy below a certain threshold, thus achieving a significant gain in CPU time.

The digitization step is a second level of detector simulation, placed just at the interface with the reconstruction program, where the physical information is reg-

Detector system	Active elements	Modules or chambers	GEANT volumes
Pixels	140 000 000	2 200	26 000
Silicon micro-strips	6 280 000	4 100	50 000
Transition Radiation Tracker	420 000	240	2 260 000
LAr accordion calorimeters	170 000	48	9 960 000
LAr hadronic end-cap and forward calorimeters	9 000	134	890 000
Tile Calorimeters	10 000	192	900 000
Muon System	1 230 000	2 000	1 850 000

Table 1.4: *Number of active detector elements, number of modules or chambers, and number of GEANT volumes defined for the detailed simulation of each of the various ATLAS detector systems (extracted from reference [12]).*

istered and reprocessed in order to simulate the detector output, and eventually written out to be used by the reconstruction programs. The output from the digitization is obtained in a form similar to the expected one from the readout electronics in the actual experiment. This step was in fact originally conceived to give the user the possibility of changing the readout characteristics (for instance, the strip pitch in the silicon detectors or the cell granularity in the EM calorimeter) immediately after the detector simulation step, thus gaining considerably in the amount of CPU time needed during the phase of detector optimization. Detailed signal treatment (for instance, the most accurate possible treatment of $\mathrm{d}E/\mathrm{d}x$ in the TRT or digital filtering in the calorimeters), simulation of the front-end electronics behavior, noise injection, *etc.*, can also be performed at this level. The digitization step is very fast, except when pile-up at high luminosity is included. It is often therefore rerun as the first step in the reconstruction chain, if noise levels or single-channel efficiencies in some of the detectors are to be varied.

1.3.3 Reconstruction

The reconstruction of particles and other physics objects in the ATLAS detector has been developed over many years, and was initially implemented in a program called ATRECON [24], mostly written in Fortran77 and using ZEBRA as memory manager. At present a new ATLAS software framework called Athena [25] is in use, written in C++ and with a fully object-oriented design. ATRECON runs on simulated GEANT-3 data and Athena uses GEANT-4 as an external package to the ATLAS software. The reconstruction proceeds in two stages. First, data from each detector is reconstructed in a stand-alone mode. Second, the information from all detectors is combined to get the most accurate measurements and identification of the final objects used in the analysis: photons, electrons, muons, leptons, kaons, jets, b-jets, $E_{\mathrm{T}}^{\mathrm{miss}}$, primary vertex, *etc.* The output of the various algorithms is stored in standardized combined n-tuples, which allow analysis in

a PAW [26] framework using GEANT-3 and now allow analysis in a ROOT [27] framework using GEANT-4. Event Summary Data (ESD) and Analysis Object Data (AOD) [28] can not be read directly with ROOT: their storage format is POOL (Pool Of persistent Objects for LHC (Data Management) [29])/ROOT, which has to be accessed via an Athena job (see chapter 5). Figure 1.9 shows a view of a high-p_{T} reconstructed $\mathrm{H} \rightarrow \mathrm{ZZ}^{\star} \rightarrow \mathrm{ee}\mu\mu$ decay, with reconstructed tracks in the inner detector and the muon system and with the reconstructed energy clusters of the two electrons and a high-p_{T} jet in the calorimeters.

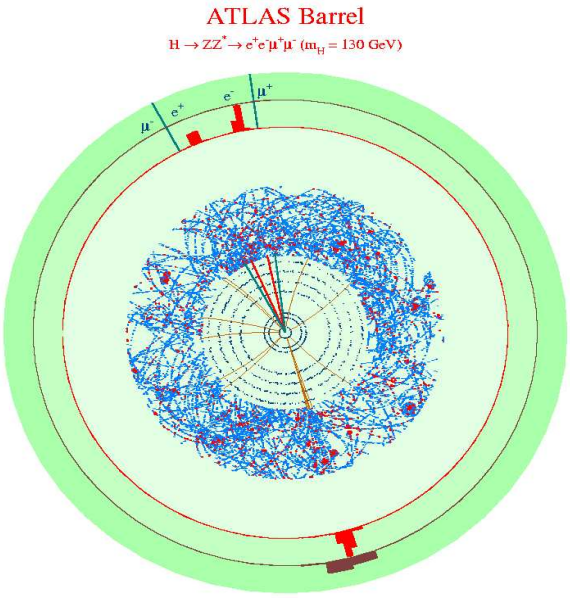

Figure 1.9: *Display in the transverse plane of a simulated high-p_{T}* $\mathrm{H} \rightarrow \mathrm{ZZ}^{\star} \rightarrow \mathrm{ee}\mu\mu$ *decay at a luminosity* $5 \times 10^{34}\ cm^{-2}\ s^{-1}$.

Tracks of charged particles are reconstructed from hits in the precision tracker and in the TRT. Two rather different algorithms with comparable performances have been developed:

- iPatRec starts from a combinatorial search in the precision tracker and extrapolates the track from the silicon to the TRT.
- xKalman finds tracks in the TRT with a histograming method and follows them inward using Kalman-filtering techniques.
- newTracking pattern strategy starts from Space Points in Pixels and SCT with an extension into TRT (inside-out tracking). Similar to xKalman and iPatRec.

Track reconstruction can be performed over the full inner detector, or over a limited range around *seeds* found by the other detectors (electromagnetic clusters, jets,

muons) and around Monte-Carlo truth information for checks. The use of seeds in the reconstruction algorithm can save a factor of up to 100 in CPU time for events including the pile-up expected at high luminosity. It is possible to run LVL1 and LVL2 trigger algorithms prior to reconstruction.

In a second stage, information from several detectors is combined. Muons reconstructed in the muon system are refined by matching the track to an inner detector track. This improves the momentum resolution, especially at moderate p_{T}, and yields accurate track parameters at the vertex. Lower-p_{T} (down to 2 GeV) muons are found by matching an inner detector track to the tile calorimeter cells. Photon conversions and $\mathrm{K_s}$ decays are searched for by pairing inner detector tracks. The primary vertex is reconstructed using all the tracks in the event. High-p_{T} (above 10 GeV) photon identification requires electromagnetic cluster shower-shape variables and the absence of reconstructed tracks in the inner detector, except for identified conversions. High-p_{T} electron identification requires a track reconstructed in the inner detector with transition-radiation hits and a measured momentum matching a calorimeter energy deposition compatible with an electromagnetic shower. Softer non-isolated electrons (down to p_{T} of 1 GeV) are identified by extrapolating inner detector tracks to the EM calorimeter. Finally, τ-leptons are identified from narrow jets associated with a small number of charged tracks. Candidate b-jets are tagged by algorithms combining the impact parameter of high-quality tracks or by vertexing algorithms.

1.3.4 Fast Simulation and Reconstruction

Fast particle-level simulation and reconstruction is an intermediate step between simple parton-level analysis of the event topology, which in general yields much too optimistic results for physics processes at hadron colliders, and very sophisticated and CPU-consuming full detector simulation and reconstruction. This kind of approach is needed for quick and approximate estimates of signal and background rates for specific channels. In addition, fast simulation and reconstruction is the only practical tool for high-statistics studies of complex background processes. A complete package for fast detector simulation and physics analysis has been implemented over the past few years, called ATLFAST [30].

ATLFAST can be used for fast simulation of signal and background processes, including the most crucial detector aspects: jet reconstruction in the calorimeters, momentum/energy smearing for leptons and photons, magnetic field effects and missing transverse energy. It provides, starting from the list of particles in the event, a list of reconstructed jets, isolated leptons and photons, the expected missing transverse energy, and reconstructed charged tracks. Values for the rejections against non b-jets are also provided as a function of the efficiencies for identifying b-jets. In most cases, the detector-dependent parameters are tuned to what is expected for the performance of the ATLAS detector from full simulation and reconstruction.

The ATLFAST package aims to reproduce as well as possible the expected detector

performance in terms of resolution and particle-identification for important physics signals. It does not attempt, at present, to reproduce accurately the expected efficiencies for lepton and photon isolation. In the case of hadronic jets, the jet reconstruction (and veto) efficiency is often dominated by physics effects, which are straightforward to model in the fast simulation. For any specific channel, the predictions from ATLFAST in terms of resolution and reconstruction efficiency, should always be confirmed with full-simulation results. The acceptances, jet reconstruction efficiencies, jet-veto efficiencies, and mass resolutions have shown good agreement between fast and full simulations. Not all the detector effects can be readily parametrized in fast simulation and only the basic information of the detector geometry is used in the package. This basic information is for example: the coverage for precision physics and for the calorimetry, the size of the barrel/end-cap transition region for the EM calorimeter, and the granularity of the hadronic calorimeters. No effects related to the detailed shapes of particle showers in the calorimeters, the charged track multiplicity in jets, *etc.*, are taken into account. In particular, energy isolation of leptons is only simulated in a crude way.

To illustrate the consistency between fast and full simulation and reconstruction, the mass reconstruction of the WH with H $\rightarrow$ bb process with $m_{\mathrm{H}} = 100$ GeV is presented in figure 1.10 using ATLFAST. Jets with $p_{\mathrm{T}} > 15$ GeV before energy recalibration are accepted and the efficiency for reconstruction is about 80% per b-jet from fast and full simulation. The mass peak can be reconstructed with an expected resolution of 15 GeV with full simulation (and 12.5 GeV with fast simulation, see figure 1.10 to be compared with 15 GeV with full simulation) and a correct position of the peak in the distribution (after jet energy recalibration), nevertheless with some fraction of the signal appearing as non-Gaussian tails. Most of these effects can be attributed to the final-state radiation and hadronisation.

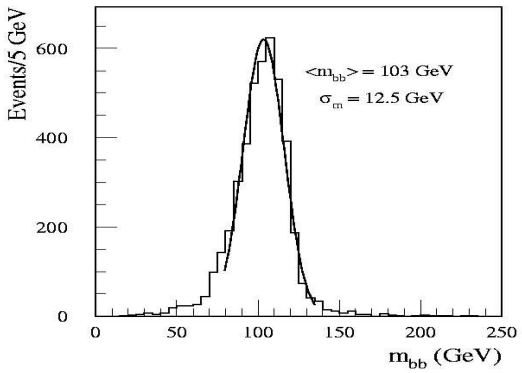

Figure 1.10: *Distribution of reconstructed invariant mass, $m_{\mathrm{b\bar{b}}}$ as obtained using fast simulation of* WH *production with $m_{\mathrm{H}} = 100$ GeV and* H $\rightarrow$ bb *decay at low luminosity operation.*

1.3.5 Worldwide LHC Computing Grid (WLCG)

The LHC Computing Grid Project (LCG) [31] was approved by the CERN Council on 20 September 2001 to develop, build and maintain a distributed computing infrastructure for the storage and analysis of data from the four LHC experiments. The project was defined with two distinct phases. In phase 1, from 2002 to 2006, the necessary software and services would be developed and prototyped, and in phase 2, covering the years 2006-2008, the initial services would be constructed and brought into operation in time for the first beams from the LHC machine. Towards the end of phase 1 it was foreseen that a Technical Design Report (TDR) [32] would be produced presenting the planning for the second phase in light of the experience gained during phase 1.

Over 2004 a Memorandum of Understanding [33] has been developed defining the Worldwide LHC Computing Grid Collaboration (WLCG). The members of the collaboration are CERN, as host laboratory, and the major computing centres that commit to provide resources for LHC data storage and analysis. It also defines the organizational structure for the management, deployment and operation of these resources as phase 2 of the LCG Project.

Requirements of the LHC experiments

The Large Hadron Collider (LHC) starts to operate in 2008 and will produce roughly 15 Peta-bytes (15 million Giga-bytes) of data annually, that thousands of scientists around the world will access and analyse. The mission of the LHC Computing Project (LCG) is to build and maintain a data storage and analysis infrastructure for the entire high-energy physics community that will use the LHC.

When the planned operation of the LHC accelerator has been achieved, access to experimental data needs to be provided for 5,000 scientists in some 500 research institutes and universities worldwide who are participating in the LHC experiments. In addition, all the data needs to be available over the estimated 12-year lifetime of the LHC. The analysis of the data, including comparison with theoretical simulations, requires several thousands of kSI2k[1] at 2004 measures of processing power. A traditional approach would be to centralize all of this capacity at one location near the experiments. In the case of the LHC, however, a novel globally distributed model for data storage and analysis - a computing Grid - was chosen because it provides several key benefits. In particular:

- The significant costs of maintaining and upgrading the necessary resources for such a computing challenge are more easily handled in a distributed environment, since individual institutes and participating national organizations can fund local computing resources and retain responsibility for them, while still contributing to the global goal.

[1] kiloSpecInt2000, kSI2k, is a measure of processor performance in terms of speed and throughput relative to an established benchmark. The benchmarking and testing are performed once every five years by spechbench.org. For comparison purposes, a 3.2 GHz Pentium 4 has a 1.267 kSI2k rating.

- It's a challenge to build a distributed system that has no single points of failure. Multiple copies of data and automatic reassigning of computational tasks to available resources ensure the load balancing of resources and facilitates access to the data for all the scientists involved, independent of geographical location. Spanning all time zones also facilitates round-the-clock monitoring and support.

Of course, a distributed system also presents a number of significant challenges. These include ensuring scalability and reliability of services as main problems, adequate levels of network bandwidth between the contributing resources, maintaining coherence of software versions installed in various locations, equipped with heterogeneous hardware, managing and protecting the data so that it is not lost or corrupted over the lifetime of the LHC, and providing accounting mechanisms so that different groups have fair access, based on their needs and contributions to the infrastructure. These are some of the challenges that the LCG Project is addressing.

LCG system architecture

The LCG Project will implement a Grid to support the computing models of the experiments using a distributed four-tiered model.

- The original raw data emerging from the data acquisition systems will be recorded at the Tier-0 center at CERN. The maximum aggregated bandwidth for raw data recoding for a single experiment (ALICE) is 1.25 GB/s. The first-pass reconstruction will take place at the Tier-0, where a copy of the reconstructed data will be stored. The Tier-0 will distribute a second copy of the raw data across the Tier-1 centers associated with the experiment. Additional copies of the reconstructed data will also be distributed across the Tier-1 centers according to the policies of each experiment.
- The role of the Tier-1 centers varies according to the experiment, but in general they 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 access to large amounts of data. At present 11 Tier-1 centers have been defined, most of them serving several experiments.
- The role of the Tier-2 centers is to provide computational capacity and appropriate storage services for Monte Carlo event simulation and for end-user analysis. The Tier-2 centers will obtain data as required from Tier-1 centers, and the data generated at Tier-2 centers will be sent to Tier-1 centers for permanent storage. More than 100 Tier-2 centers have been identified.
- Other computing facilities in universities and laboratories will take part in the processing and analysis of LHC data as Tier-3 facilities. These lie outside the scope of the LCG Project, although they must be provided with access to the data and analysis facilities as decided by the experiments.

Prototype implementations of system components

The LCG system has to be ready for use with full functionality and reliability at the start of LHC data taking. In order to ensure this readiness, the system is planned to evolve through a set of steps involving the hardware infrastructure as well as the services to be delivered. At each step it is planned to use a prototype installation of the Grid middleware[2] for extensive testing. The experiments use it to perform their Data Challenges (DCs), progressively increasing in scope and scale, stressing the LCG system with activities that are more and more similar to those performed when the experiments are running. The service providers at CERN and in the Tier-1 and Tier-2 centers also use the prototype Grid for a series of Service Challenges, aimed at identifying reliability, performance, and management issues at an early stage, to allow robust solutions to be developed and deployed.

Organization and management of the project

The LCG Project will collaborate and inter-operate with other major Grid development projects, network providers and production environments around the world.

The EGEE (Enabling Grids for E-science in Europe) [34] project, partially funded by the European Union, has the goal of deploying a robust Grid infrastructure for science. It interconnects a large number of sites in 34 countries around the world, integrating several national and regional Grid initiatives in Europe, such as INFN Grid in Italy, DutchGrid in the Netherlands and GridPP in the UK. EGEE was designed as a project that would extend the Grid that had been set up by phase 1 of the LCG Project during 2003, with the aim of starting with a working Grid and demanding applications, and evolving to support other sciences. The operation of the LCG Grid therefore effectively merged with that of EGEE when the latter began in April 2004. In return EGEE provides very substantial funding for Grid operations and support at its partner sites, the large majority of which provide computing resources for the LHC experiments. The majority of the sites in the Worldwide LCG Collaboration will be connected to EGEE. The current phase of the EGEE project is funded to March 2006. It is expected that a second phase will be funded for two more years, until March 2008. It is clearly important that the LCG Project continues to be closely involved in the operation of EGEE until the longer term funding model is understood.

The LCG Project also has relationships with other Grid operations environments, such as the Open Science Grid (OSG)/Grid3 [35] in the US and the Nordic Data Grid Facility (NDGF)/NorduGrid [36]. The LCG Project has a formal relationship with the US experiment projects whose resources are connected to OSG, but not with OSG itself.

[2]Middleware is computer software that connects software components or applications. The software consists of a set of enabling services that allow multiple processes running on one or more machines to interact across a network. It is used most often to support complex, distributed applications.

1.3.6 ATLAS Computing Model

The ATLAS Computing Model [37] defines the requirements on the architecture and performance of the software and the overall computing system.

The ATLAS Computing Model embraces the Grid paradigm and a high degree of decentralization and sharing of computing resources. The required level of computing resources relays even stronger than previous CERN-based experiments on off-site computing facilities.

The primary event processing occurs at CERN in a Tier-0 facility. The raw data is archived at CERN and copied (along with the primary processed data) to the Tier-1 facilities around the world. These facilities archive the raw data, provide the reprocessing capacity, provide access to the various processed versions, and allow scheduled analysis of the processed data by physics analysis groups. Selected datasets produced by the physics groups are copied to the Tier-2 facilities for further analysis. The Tier-2 facilities also provide the simulation capacity for the experiment, with the simulated data housed at Tier-1s. In addition, Tier-2 centers will provide analysis facilities, and some will provide the capacity to produce calibrations based on processing raw data. An additional CERN Analysis Facility should play an important role in the calibration and algorithmic development work.

One of the principal challenges for ATLAS computing is to develop and operate a data storage and management infrastructure able to meet the demands of a yearly data volume of O(10 PB) utilized by data processing and analysis activities spread around the world. The ATLAS Computing Model establishes the environment and operational requirements that ATLAS data-handling systems must support, and, together with the operational experience gained to date in test beams and Data Challenges, provides the primary guidance for the development of the data management systems.

The ATLAS Computing Model foresees the distribution of raw and processed data to Tier-1 and Tier-2 centers, so as to be able to exploit fully the computing resources that are made available to the Collaboration. Additional computing resources will be available for data processing and analysis at Tier-3 centers and other computing facilities to which ATLAS may have access. A complex set of tools and distributed services, enabling the automatic distribution and processing of the large amounts of data, has been developed and deployed by ATLAS in cooperation with the LHC Computing Grid (LCG) Project and with the middleware providers of three large Grid infrastructures: LCG/EGEE, OSG/Grid3 and ARC/NorduGrid. The tools are designed in a flexible way, in order to have the possibility to extend them to use other types of Grid middleware in the future. These tools, and the service infrastructure on which they depend, were initially developed in the context of centrally managed, distributed Monte Carlo production exercises. They will be re-used wherever possible to create systems and tools for individual users to access data and compute resources, providing a distributed analysis environment for general usage by the ATLAS Collaboration. The first

version of the Production System (ProdSys) was deployed in the summer 2004 and has been used since then. It was used for Data Challenge 2 (DC2), for the production of simulated data for the 5th ATLAS Physics Workshop (Rome, June 2005), for the reconstruction and analysis of the 2004 Combined Test-Beam data and a newer version is being used for the ongoing Data Challenge 3 (DC3, also called Computing System Commissioning, CSC).

ATLAS has two different subjects as main computing operations. The first one is the preparation, distribution and validation of the ATLAS software among the institutes that are involved in the ATLAS collaboration. The second one is to supervise the computing and data management operations run centrally on Tier-0, Tier-1s and Tier-2s. The ATLAS Virtual Organization[3] (VO) [38] will allow production and analysis users to run jobs and access data at remote sites using the ATLAS-developed Grid tools.

In the past few years the Computing Model has been tested and developed by running Data Challenges of increasing scope and magnitude, as was proposed by the LHC Computing Review in 2001. We have run two major Data Challenges since 2002, performed other massive productions and currently running the third Data Challenge in order to provide simulated data to the physicists and to reconstruct and analyse real data coming from test-beam activities; this experience is now useful in setting up the operations model for the start of LHC data-taking in 2008.

[3]In Grid computing, a Virtual Organization is a group of individuals or institutions who share the computing resources of a 'Grid' for a common goal.

Chapter 2

b-tagging at high and very high p_T

Many physics processes within and beyond the Standard Model (SM) involve b-quark decays. A light Standard Model Higgs boson has a large branching ratio into b-quarks and many supersymmetric signatures include b-quarks in their final states. Furthermore, new heavy particles, expected in models beyond SM are decaying into b-quarks and produce high p_T b-jets. The p_T of these b-jets depends on the scale of the new physics and jets could have transverse momenta of a few hundred GeV. The potential of ATLAS to identify high and very high p_T jets is therefore relevant for both expected and unexpected physics.

The ATLAS inner tracker is designed to provide an efficient tagging of b-jets. In this chapter the capability of the ATLAS detector to identify high p_T b-jets is studied. Two different event samples have been analyzed, a heavy Higgs boson ($\mathrm{m}_H = 400$ GeV) in associated production with a W in section 2.1, and the Little Higgs s-channel production of a heavy $\mathrm{Z}_H(\mathrm{m}_{Z_H} = 2$ TeV) in section 2.2. In both cases, the heavy particle decays into b-quark pairs.

2.1 b-tagging at high p_T

In this section, a heavy Higgs boson with $m_H = 400$ GeV, in associated production with a W, decaying into b-quark pairs is analyzed to perform b-tagging studies at high p_T. The simulation of the events is discussed in section 2.1.1. An introduction to the b-tagging algorithm used by ATLAS is given in section 2.1.2. Finally, the results for high p_T jets are presented in section 2.1.3.

2.1.1 Event simulation

The ATLAS full simulation is a detailed simulation of the expected detector response (see section 1.3). The event generation involves a chain of Monte Carlo programs. First, events are generated with PYTHIA 6.203. Then the response of the ATLAS detector is simulated using the ATLSIM program. Particles from the generator are tracked through the detector using GEANT 3.21. All relevant energy loss mechanisms are taken into account. The GEANT energy losses are converted to detector signal digits using a detailed description of the front-end electronics.

For the detector description, the Data Challenge 1 (DC1) layout was used. This layout is more realistic than the layout used in 1999 for the analysis presented in the ATLAS Technical Design Report. The b-tagging is mostly affected by modifications in the inner detector, in particular in the pixel system. But also SCT, TRT and the magnetic field have been revised as well. The present design of the pixel detector allows to insert and remove the detector in an easy way. The new pixel detector is composed of 3 barrel layers and 3 end-cap wheels. The radius of the first pixel layer (or B-layer) has been increased from 4.3 to 5.05 cm to accommodate the larger beam pipe. All the changes have the overall result of increasing the dead material by a factor of $\sim$ 1.5, allowing a more realistic detector description.

In addition, changes have been made to other systems (in material and design). The most important are:

- SCT tilt sign has been reversed, the detectors design has changed from n-on-n to p-on-n.
- TRT geometry has been changed, a new modular design has been introduced for better η coverage.
- A more realistic non-uniform magnetic field has been included in the simulation of the inner tracker.

At the LHC, multiple collisions occur in each bunch crossing. In the simulation, minimum bias events are therefore added to each hard process. These events account for the vast majority of interactions resulting from beam collisions, and are also called minimum bias or pile-up events. During the initial low luminosity phase, $\mathcal{L} = 10^{33} cm^{-2} s^{-1}$, on average 2.3 minimum bias events are expected in each bunch crossing, increasing to 23 events for the high luminosity phase, $\mathcal{L} = 10^{34} cm^{-2} s^{-1}$. The deterioration of the b-tagging efficiency is however not very significant, as discussed later.

Data samples

Several data samples with large number of events were produced. The process $W + H(400)$ with the Higgs boson ($m_H = 400$ GeV) decaying into $b\overline{b}$, $u\overline{u}$ and

$c\overline{c}$. High statistics (50,000 events) were used for the decays involving $u\overline{u}$ and $c\overline{c}$ and 20,000 events for the decay into b-quarks. These samples were part of the standard ATLAS samples generated for b-tagging purposes. In addition, a dedicated simulation was performed for the Little Higgs decay of Z_H ($m_{Z_H} = 2$ TeV) into b, u and c quarks. In this case 20,000 events were simulated for each sample.

Reconstruction

The output of the simulation is used by the pattern recognition algorithm to perform the track reconstruction. Two main reconstruction algorithms exist in ATLAS: iPatRec [39] and xKalman [40]. Since both algorithms yield very similar results, only results using xKalman are presented in the following. The xKalman version used corresponds to ATRECON release 6.6.0.

The pattern recognition algorithms reconstruct track parameters using the inner detector hit information. Using reconstructed tracks and their parameters, jets can be identified and reconstructed. The ATRECON framework allows the output of ATLFAST and full simulation, so ATLFAST jets can be used in the analysis. This procedure avoids the lengthy procedure of reconstructing showers in the calorimeters.

The tuning of ATLFAST with full simulation data suggests that the jets are realistic. Furthermore, b-tagging results using different types of jets (including ATLFAST jets) shows no significant discrepancies [41] [42]. ATLFAST jets are based on a cone algorithm, where calorimetric clusters are combined to form jets when the angular distance is smaller than $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} < 0.4$. The jet direction is calculated as the weighted average of the cluster directions and the energy as the sum of the cluster energies. A jet is labeled as b-jet in ATLFAST if a b-quark of $p_T > 5$ GeV (after final state radiation) is found in a cone of $\Delta R = 0.2$ around the reconstructed jet. Jets are reconstructed only for $|\eta| < 2.5$. The same criteria are applied to c-jets.

Figure 2.1 shows the jet p_T distribution for the two samples studied in the following. The dashed line is the p_T distribution of jets from the decay of the heavy Higgs boson and the continuous line corresponds to jets from the decay of Z_H. The cuts applied to reconstructed jets are $|p_T| > 15GeV$ for the first sample and $|p_T|$ larger than 500 GeV for the second sample. Only $|\eta| < 2.5$ values are considered. For the decay into u-jets the increment of p_T of the jets implies an increase on the number of tracks inside the jets. For u-jets from the first sample, the mean charged particle multiplicity is 7.8, increasing up to 13.5 for the second sample (see again figure 2.1).

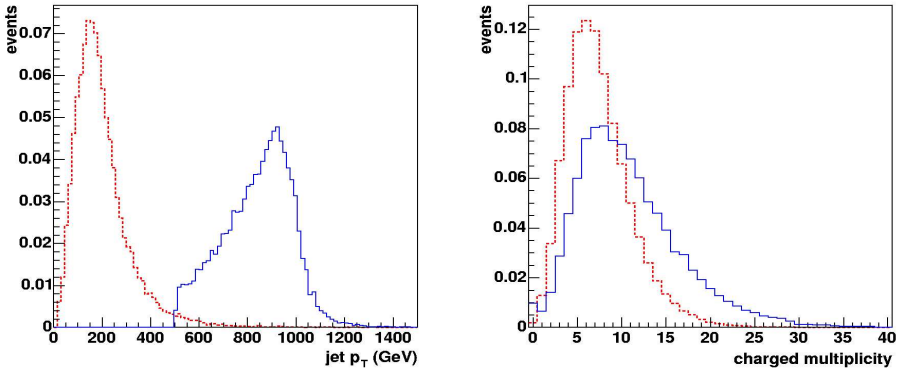

Figure 2.1: *Momentum of b-jets for jets from the decay of the heavy Higgs boson (dashed line) and the heavy* Z_H *(continuous line). On the right side the corresponding multiplicities of charged particles in u-jets are shown.*

2.1.2 2D and 3D algorithms

The b-tagging algorithms take advantage of the long lifetime of B-hadrons, on average 1.56 ps [43] or $c\tau \sim 460\mu m$, that form displaced vertexes. Particles not coming from a B-hadron decay point generally to the primary vertex. This property allows the tagging of b-jets by either reconstructing the secondary vertex or examining the impact parameter of tracks associated to the jet.

A source of misidentification for b-tagging arises from charmed hadrons that have a long lifetime and can be confused with B-decays. Other sources are interactions of the primary particles with the dead material of the detector, like conversions of photons or hadronic interactions. Therefore the b-tagging performance depends strongly on the amount of dead material inside the tracking volume.

There are several b-tagging algorithms that provide a good b-tagging performance (simple track counting, ALEPH method,...). In the inner detector Technical Design Report (TDR) and in previous studies, the best performance was achieved with an algorithm based on the likelihood ratio method, that is used in the following. This method is just based on impact parameter information and is described in more detail in reference [41].

First, some cuts are applied in order to select a good quality sample of tracks. A cut on the p_T of the track ($|p_T| < 1GeV$) helps to decrease the number of tracks severely affected by multiple scattering. An angular cut of $\Delta R = 0.4$ around the jet axis is performed. Some additional quality cuts are applied to the selected track sample:

- number of precision hits (pixels + SCT), $N_{pr} \geq 9$,
- number of pixel hits, $N_{pix} \geq 2$,
- a hit in the first pixel layer (B-layer),

- transverse impact parameter, $d_0 < 1mm$.

The first two cuts guarantee enough precision hits for the reconstruction. The B-layer hit is necessary to obtain a good resolution in the impact parameter. These cuts eliminate tracks coming from long-lived particles and tracks created by the interaction of primary tracks with dead material. Finally, the cut on d_0 excludes photon conversions. Another set of cuts is applied to further eliminate secondary tracks.

- the longitudinal impact parameter $z_0 < 1.5mm$,
- the χ^2 of the track $\chi^2 < 3$,
- no shared hits in the B-layer,
- a maximum of two shared hits in the precision tracker.

The first cut is important only in the the presence of pile-up, where z_0 is the primary interaction vertex. The last two cuts have a more significant effect. Those cuts are useful to eliminate tracks from conversions or secondary tracks originated in the detector material.

With tracks passing these cuts, the resolution of the inner detector can be calculated. The reviewed design modifications result in a loss in performance with respect to the performance presented in the TDR. The most significant changes can be observed in the single track transverse momentum and impact parameter d_0 resolution. Figure 2.2 shows the difference in resolution between the TDR layout and the DC1 layout for d_0 (left) and $1/p_T$ (right). The increase of material is responsible for most of the degradation of the resolution in the low p_T region. The impact parameter resolution is also affected by the increase of the B-layer radius. The effect of these changes in the b-tagging performance for both layouts, TDR and DC1, is illustrated in table 2.1 and 2.2. When the p_T increases the difference is reduced and for the heavy Higgs boson sample the values are similar. The difference between TDR and DC1 is not only caused by the change in layout, also the software is important. The improvement of the tracking software is the reason for a good performance at high p_T using the DC1 layout. These changes have been studied in detail in references [41] [42] [44].

$R(\epsilon_b)$	TDR	DC1	ratio
$\epsilon_b = 50\%$	326 ± 37	180 ± 8	0.55
$\epsilon_b = 60\%$	124 ± 9	57 ± 2	0.46

Table 2.1: *Rejection of u-jets for b-tagging efficiencies 50% and 60%. Results for the ATLAS benchmark samples with a light Higgs* (m_H *= 120 GeV for TDR and DC1, respectively).*

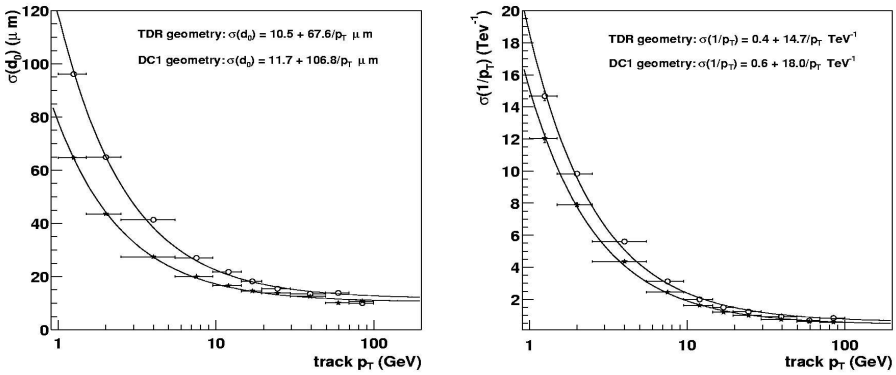

Figure 2.2: *Dependence on $|\eta|$ of resolutions in transverse impact parameter,* d_0 *(left) and 1/*p_T *for the TDR geometry (open circles) and the DC1 layout (stars).*

$R(\epsilon_b)$	TDR	DC1	ratio
$\epsilon_b = 50\%$	126 $\pm$ 9	139 $\pm$ 6	1.10
$\epsilon_b = 60\%$	65 $\pm$ 3	55 $\pm$ 1	0.85

Table 2.2: *Rejection of u-jets for b-tagging efficiencies 50% and 60%. Results for ATLAS benchmark samples with a heavy Higgs (m_H = 400 GeV). TDR and DC1 results are compared.*

It has been pointed out that the impact parameter is an important ingredient of the b-tagging algorithm. In figure 2.3 the definition of the impact parameter of a track (d_0) is shown graphically. The impact parameter is defined as the closest distance between the primary vertex and the track helix in the transverse plane (d_0) or in the longitudinal direction (z_0). Here the d_0 transverse impact parameter is normally referred as impact parameter, since the precision in the transverse plane is much better than in the longitudinal direction. The longitudinal impact parameter can be used for b-tagging as well, but is less discriminating as discussed later.

The impact parameter is signed depending on the point where the track crosses the jet axis. The sign is negative if the track appears to originate before the primary vertex and positive otherwise. The sign is also explained in figure 2.3. For long-lived particles the value is large and positive. Particles arising from the primary vertex have zero impact parameter. This fact is used to create likelihood functions. For that purpose, the track significance (S_{d_0}) is defined. S_{d_0} is the ratio of the signed impact parameter and its error:

$$S_{d_0} = \frac{d_0}{\sigma_{d_0}}$$

where the error σ_{d_0}, includes both the measurement error and the beam spread of 15 μm. This value is calculated for each track inside the jet and yields a different distribution for each jet flavour. Figure 2.4 shows the u-jet (dashed line) and

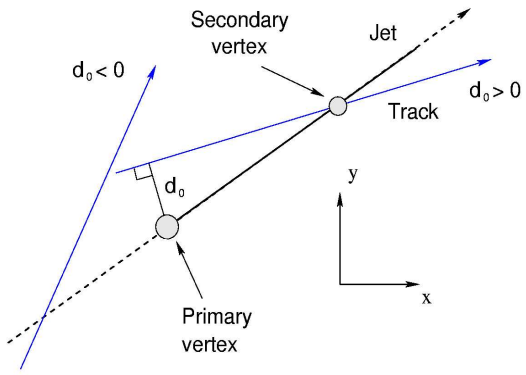

Figure 2.3: *Definition of the transverse impact parameter* (d_0)*. Note that this quantity may be negative.*

b-jet (continuous line) significance distributions for the sample $H(400)$ (left) and $Z_H(2000)$ (right). The shape of the distribution is nearly gaussian for u-jets. The small tail present on the positive side is due to material interactions and long lived particles: K^0, Λ,... The b-jet curve shows a pronounced asymmetry towards the positive side. The excess of positive significance tracks results from B-hadron decays. The excess in the negative side is due to an incorrect determination of the impact parameter sign. A clear difference is also observed between the two samples. The excess of positive side tracks is less pronounced in the second sample, implying a p_T dependence of the b-tagging discrimination.

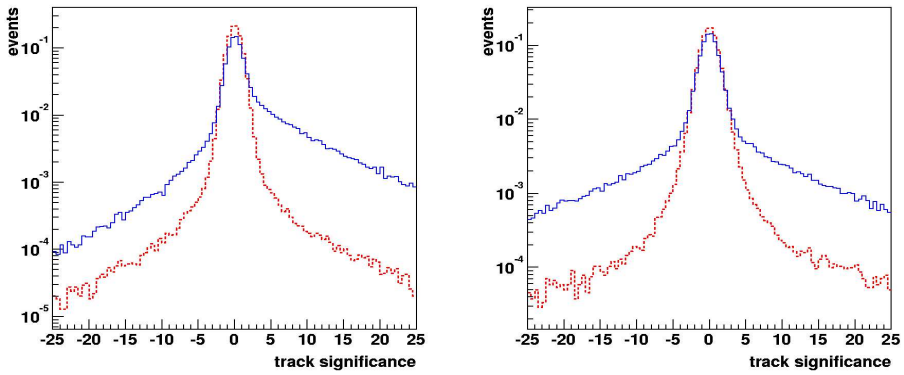

Figure 2.4: *Signed track significance distributions for u (dashed line) and b jets (continuous line) from the decay of the heavy Higgs boson* (m_H *= 400 GeV) on the left and the heavy gauge boson* (m_{Z_H} *= 2000 GeV) on the right.*

The ratio between the two distributions for u and b-quarks is a measure of the capability to discriminate both type of quarks. The weight of a track (w_t) is calculated using the ratio:

$$w_t(S_{d_0}) = log\Big[\frac{f_b(S_{d_0})}{f_u(S_{d_0})}\Big] \qquad\qquad 2D\ -\ weight$$

where $f_u(S_{d_0})$ is the fraction of tracks in the u-jet sample with significance S_{d_0}. For the 2 dimensional algorithm only the transverse impact parameter is used. An improvement is expected by adding the longitudinal impact parameter in the 2D+Z or 3D algorithm. The weight is then defined as follows:

$$w_t(S_{d_0}, S_{z_0}) = log\Big[\frac{f_b(S_{d_0}, S_{z_0})}{f_u(S_{d_0}, S_{z_0})}\Big] \qquad\qquad 3D\ -\ weight$$

where S_{z_0} is the signed significance of the longitudinal impact parameter (an error of 30 μm is assumed for z_0 following simulation studies detailed in reference [12]). The weight of each track depends only on S_{d_0} for the 2D, but on a 2-variable function (S_{d_0},S_{z_0}) for the 3D algorithm. In this way, the proper correlation between the transverse and longitudinal impact parameters is fully taken into account.

The jet weight is defined as the sum of individual track weights:

$$w_{jet} = \sum_i w_t^i$$

where w_t^i is the weight of track i and the sum runs over all tracks associated to the jet. The jet weight distribution for both types of jets (u and b) and both samples are shown in figure 2.5. The expanded shape of b-jet distribution allows the separation of u and b-jets. The narrower distribution for b-jets in the very high p_T sample is indicative of a reduced discrimination power for both types of jets.

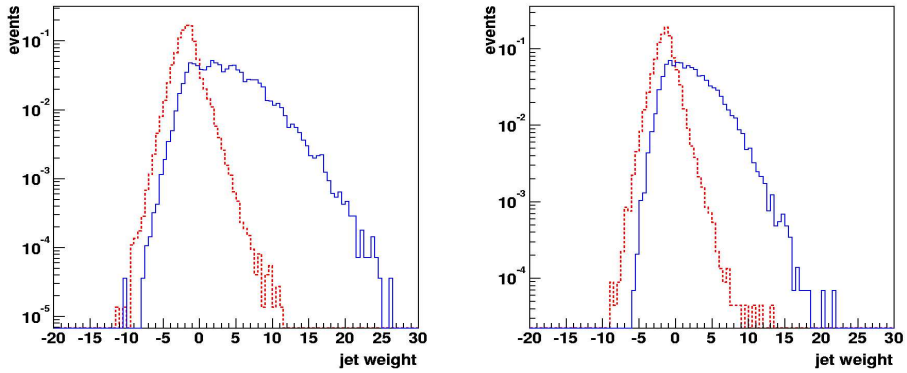

Figure 2.5: *Jet weight distributions distributions for u (dashed line) and b jets (continuous line) from the decay of the heavy Higgs boson* (m_H *= 400 GeV) on the left and the heavy gauge bosons* (m_{Z_H} *= 2000 GeV) (right) for the 2D algorithm.*

Track weights have been extracted from the same set of MC data that are used later to perform the various analysis. In the running experiment, these weights are expected to be calculated from both MC and real data.

2.1.3 High p_T tagging

Once the jet weights have been obtained, the u-jet and b-jet tagging efficiencies can be calculated as the integrals of normalized weights distributions. Thus a correspondence between b-tagging efficiencies and u-tag efficiency is established. In figure 2.6 (left), rejection for u-jets (R_u), defined as the inverse of u-jet efficiency, versus the b-tagging efficiency (ϵ_b) is shown for the H(400) sample. On the right-hand side, the rejection of u-jets versus the efficiency for tagging b-jets from the $Z_H(2000)$sample is displayed.

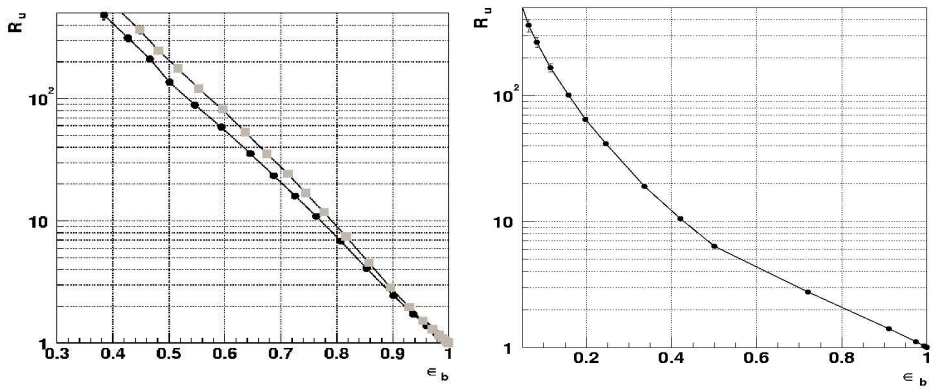

Figure 2.6: *u-jet rejection* (R_u) *versus the efficiency of tagging b-jets from the heavy Higgs boson decay (left) and the heavy gauge boson sample,* $Z_H(2000)$ *(right). Values using the 2D (dark markers) and 3D (light markers) algorithms are presented. For the heavy gauge boson sample,* $Z_H(2000)$*, there is not an improvement using 3D.*

As expected, the rejection power for u-jets is larger for the H(400) sample. With increased p_T the jet complexity increases leading to worse discrimination between u and b-jets. In figure 2.6, 2D and 3D algorithms are compared as well. For very high p_T jets, no significant improvement is observed using the 3D algorithm. The rejection for u-jets is a steep function of the b-tagging efficiency. The choice of ϵ_b depends on specific efficiency and purity requirements in a given analysis. In many ATLAS searches, a typical assumption for light jet rejection is $R_u = 100$, for a b-tagging efficiency of 60%. The rejection (at an efficiency of 50%) obtained using just the 2D algorithm and the two samples under study in the present analysis is: efficiency is:

$$\begin{array}{rcl} H(400): & R_u = 139 \pm 6 & 50\% \text{ efficiency} \\ Z_H(2000): & R_u < 10 & 50\% \text{ efficiency} \end{array}$$

The rejection for very high p_T sample is worse, as expected. An increased rejection may be obtained by reducing the b-tagging efficiency, as discussed in the next section. The rejection for c-jets is determined in the same way. Table 2.3 summarizes the rejection values for the H(400) sample, for both 2D and 3D algorithms.

	u-rejection	c-rejection
2D ($\epsilon_b = 50\%$)	139 $\pm$ 6	11.3 $\pm$ 0.3
3D ($\epsilon_b = 50\%$)	207 $\pm$ 9	14.0 $\pm$ 0.2

Table 2.3: *Rejection for u and c jets for a 50% b-tagging efficiency for the 2D and 3D algorithms evaluated for the sample WH(400).*

The rejection values quoted up to now are averages over all jet momenta and pseudo-rapidity values. Figure 2.7 shows the p_T dependence (left) and η dependence (right) of the u-jet rejection for the H(400) sample. As expected, the rejection has a significance dependence on p_T. A peak is observed for momenta around 100 GeV. For momenta larger that some 500 GeV the average flying distance of B-hadrons exceeds the 5 cm radius of the B-layer, therefore tracks from B-decays miss the B-layer hit and their impact parameter resolution is considerably degraded. This effect is discussed in more detail in the next section. Figure 2.7 (right) shows the dependence on pseudo-rapidity. The rejection is a factor 4 worse for large values of η. The increase in material in the forward region is the cause of this deterioration.

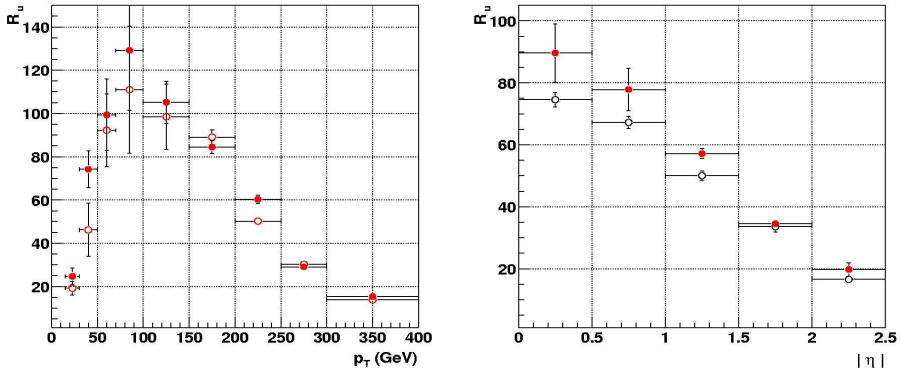

Figure 2.7: R_u *corresponding to b-tagging efficiencies of 60% versus* p_T *(left) and* $|\eta|$ *(right). Jets from the H(400) sample without pile-up (filled circles) and with pile-up (open circles) as for high luminosity.*

2.2 *b*-tagging at very high p_T

In this section, the Little Higgs s-channel production of a heavy Z_H (m_{Z_H} = 2 TeV) decaying into *b*-quark pairs is analyzed to perform *b*-tagging studies at very high p_T. The very high p_T performance is discussed in section 2.2.1. An introduction to the secondary vertex algorithm used by ATLAS is given in section 2.2.2. Finally, the results for very high p_T jets are presented in section 2.2.3.

2.2.1 Very high p_T performance

In all decays analyzed in next chapter, the tagging of very high p_T *b*-jets is required. The average p_T of these jets is reported in table 2.4:

$< p_T >$	$Z_H \to b\bar{b}$	$Z_H \to t\bar{t}$	$W_H \to tb$
M = 1 TeV	400	200	250
M = 2 TeV	800	400	500

Table 2.4: *Average* p_T *(in GeV) of final state b-jets for the various channels studied in next chapter.*

The tagging of *b*-jets has been studied in detail within ATLAS (see previous section), but in all these studies the average p_T of *b*-jets was typically below 200 GeV. In order to investigate the *b*-tagging performance for larger p_T values, the following event samples have been generated and analyzed, as commented before:

$$Z_H(2\ TeV) \to b\,\overline{b},\ c\,\overline{c},\ u\,\overline{u}$$

Each sample contained 20,000 events and was processed using the full reconstruction provided by DC1 software. The *b*-tagging results relevant for the analysis discussed in next chapter are summarized below.

As explained in the previous sections, the *b*-tagging performance is degraded as the p_T of the jet increases. This degradation can however be compensated by a reduction of the *b*-tagging efficiency as shown in figure 2.8. This figure shows the overall *b*-tagging performance obtained using the *b*-jet sample from $Z_H(2000)$ decays. The average p_T of these jets is 800 GeV. For $\epsilon_b = 50\%$, the rejection obtained with the 2D algorithm is $R_u < 10$, a very poor value. However, rejections close to $R_u \sim 100$ can be achieved if the efficiency is reduced to $\epsilon_b \sim 20\%$.

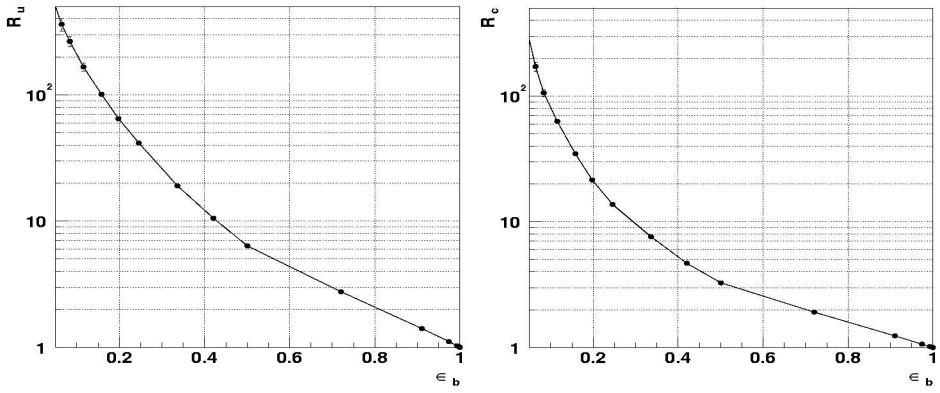

Figure 2.8: *Rejection of u-jets (left) and c-jets (right) versus b-tagging efficiency (ϵ_b). These plots have been obtained with the* Z_H *(2000) sample and the 2D algorithm.*

The standard 2D b-tagging algorithm was used. Since this algorithm is not optimized for high p_T jets, the results should be considered conservative.

The efficiencies (ϵ) and rejections ($R = 1/\epsilon$) reported in table 2.5 have been applied in the analysis presented in the next chapter.

mass	decay	ϵ_b	R_c	R_u
1 TeV	$b\bar{b}$	0.1	140	1000
	$t\bar{t}$	0.5	10	100
	tb	0.5	7	40
2 TeV	$b\bar{b}$	0.1	45	90
	$t\bar{t}$	0.2	28	130
	tb	0.2	26	75

Table 2.5: *Efficiencies and rejections applied in the b-tagging analysis presented in the next chapter.*

All these numbers are extracted from figures 2.9 and 2.10.

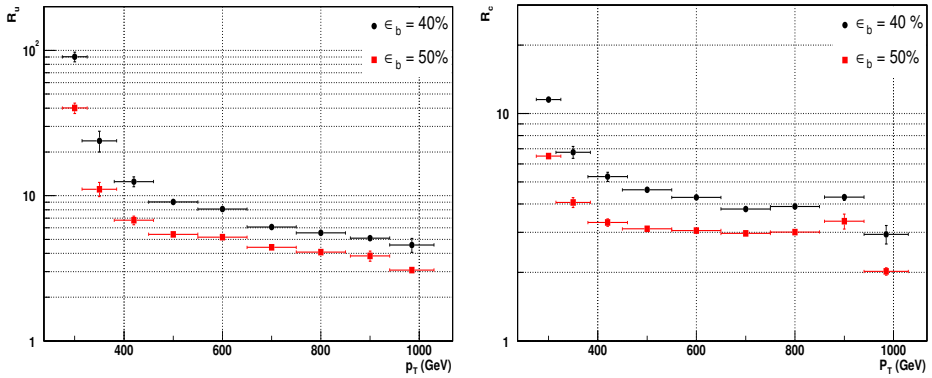

Figure 2.9: *Rejection of u-jets (left) and c-jets (right) versus average* p_T *of b-jets for* $\epsilon_b = 40\%$ *and 50%. These plots has been obtained with the* $Z_H(2000)$ *sample and the 2D algorithm.*

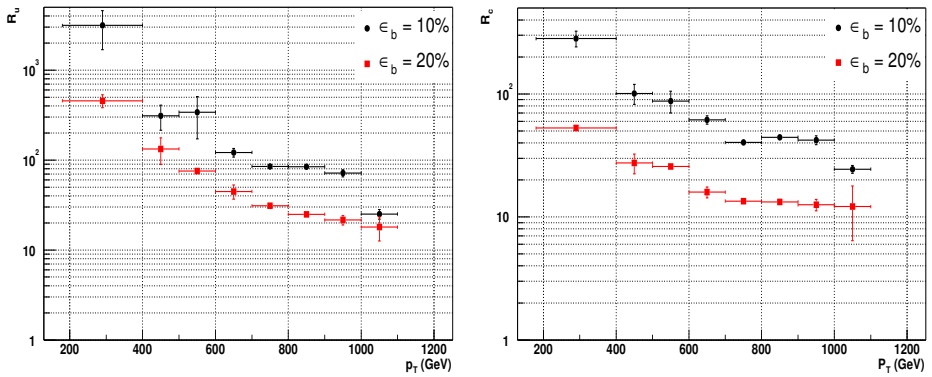

Figure 2.10: *Rejection of u-jets (left) and c-jets (right) versus average* p_T *of b-jets for* $\epsilon_b = 10\%$ *and 20%. These plots have been obtained with the* Z_H *sample and the 2D algorithm.*

2.2.2 The SV1 algorithm

The 2D and 3D algorithms presented before are optimized to achieve large efficiencies and still provide large rejection factors. A different class of algorithms, based on secondary vertexes rather than track impact parameters, has also been considered by ATLAS and used to improve the *b*-tagging performance. Furthermore, secondary vertex algorithms (SV1 in the following) yield a better performance at very high p_T, as discussed in the next section.

The main purpose of a secondary vertex algorithm is to detect the presence of secondary vertexes due to B(D) hadron decays inside a jet, with the highest efficiency. The procedure used by the SV1 algorithm is used in the VKalVrt vertexing package. More details about the algorithm and its performance can be found in reference [45]. The search starts with all track pairs inside a jet and impact parameter significance larger than 2.0. If the track pair is compatible with a secondary

vertex, the invariant mass is calculated. As shown in figure 2.11.a) and 2.11.b), most of these tracks are K_s^0 and Λ decays. The position of the secondary vertex is calculated as well, as shown in figure 2.11.c), where secondary interactions (mainly photon conversions) in the pixel layers are clearly seen. All these secondary vertexes are eliminated before applying the b-tagging algorithm. All remaining tracks pairs producing vertexes are combined into a single vertex. If the χ^2 is unacceptably large, the track with the largest χ^2 contribution is removed and the procedure is iterated until an acceptable χ^2 is obtained. Using the H(400) $b\,\overline{b}$ and $u\,\overline{u}$ samples, the probability to find a secondary vertex in a b-quark jet is 77%, and the same probability for u-jets is only 12%. Figure 2.12 shows the (normalized) radius for secondary vertexes found in u and b-jets. The b-jet sample shows clearly the exponential tail expected from long lived B-hadrons.

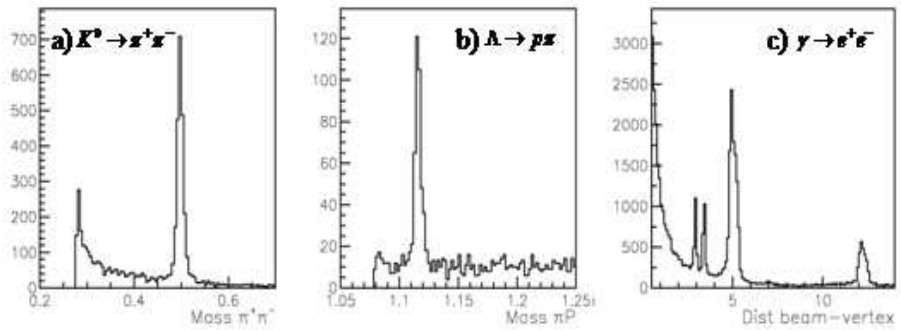

Figure 2.11: *Invariant mass of track pairs used by secondary vertex algorithm (a) and (b). The third plot (c) shows the radius of the secondary vertex. The various peaks appearing in this plot correspond to the position of the beam-pipe-walls and pixel layers. These plots have been obtained with the H(400) sample.*

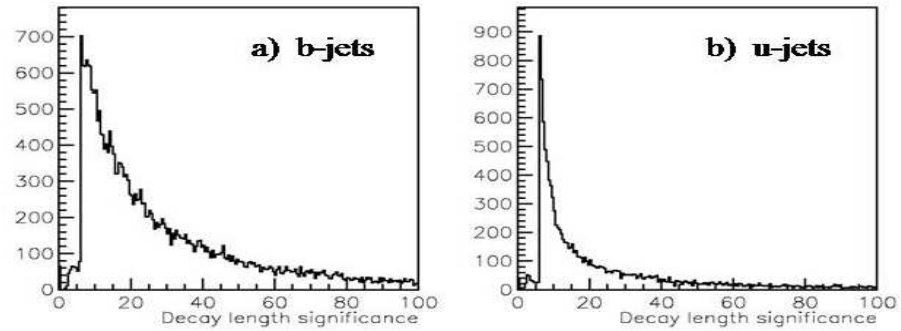

Figure 2.12: *Decay length significance for secondary vertexes found by the SV1 algorithm in b-jets (a) and u-jets (b). These plots have been obtained with the H(400) sample.*

The SV1 algorithm is intended to provide a discriminating variable that could easily be combined with the 2D and 3D algorithms. Therefore, a weight should be calculated using variables uncorrelated with track impact parameters, excluding for example the decay length of the vertex. The following 3 variables have been selected:

- the invariant mass M of all tracks associated with the secondary vertex,
- the fraction of energy F between these tracks and the total jet, and
- the number N of 2-track secondary vertexes combined in order to obtain the final vertex.

The discriminating power of all these variables is clearly seen in figure 2.13.a) to 2.13.f). The SV1 algorithm yields a weight defined as:

$$W_{SV1} = log\Big[\frac{b(M,F,N)}{u(M,F,N)}\Big]$$

where b(M,F,N) and u(M,F,N) are the distributions obtained for b and u-jets, respectively. Note that 3-dimensional distributions are used, rather than the simple weight products in order to properly take into account the correlations between the individual distributions. Figure 2.14, for example, shows the correlation between the variables F' and M', and this correlation is far from negligible. For practical purposes, the variables actually used by the algorithm are M', F' and N' defined as follows:

$$M' = \frac{M}{M+1} \qquad F' = F^{0.7} \qquad N' = logN$$

The vertex tag can easily be combined with the 3D tag by a simple sum of weights:

$$W_{b-tag} = W_{3D} + W_{SV1}$$

and the resulting weight has an improved performance as demonstrated in the following table 2.6:

ϵ_b \ R_u	2D	3D	3D+SV1
$\epsilon_b = 50\%$	187 ± 13	317 ± 29	671 ± 90
$\epsilon_b = 60\%$	65 ± 3	113 ± 6	261 ± 22

Table 2.6: *Performance of the various b-tagging algorithms on the H(400) samples, for efficiencies of 50% and 60%.*

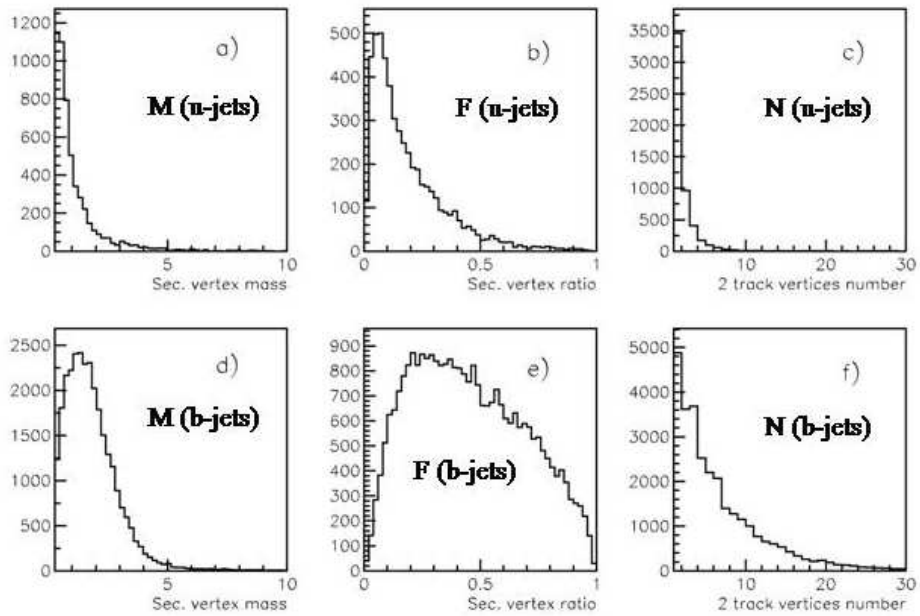

Figure 2.13: *Discriminating functions M, F and N used in the SV1 algorithm for u and b-jets. The definitions of M, F and N are given in the text. These plots have been obtained using the H(400) sample.*

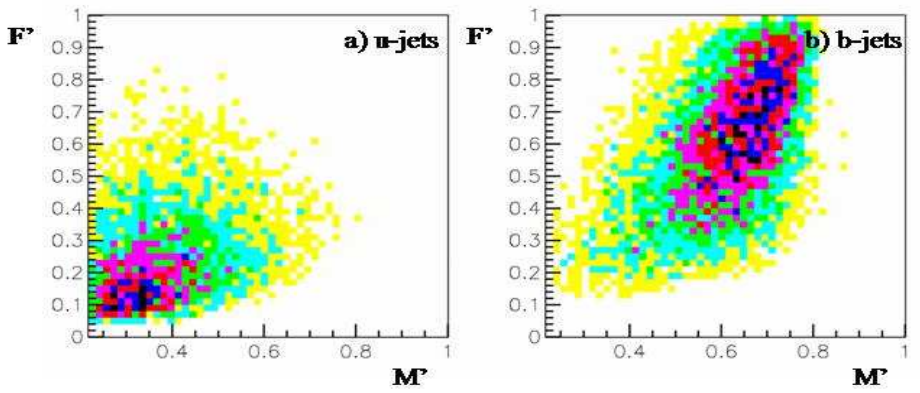

Figure 2.14: *Correlation between the variables F' and M' used by the SV1 algorithm. These plots have been obtained using the H(400) sample.*

2.2.3 SV1 versus 2D at very high p_T

The SV1 algorithm may be used at high p_T to reinforce the 2D/3D algorithms, as commented in the previous sections. The situation is somewhat different at very high p_T, as explained below (see also figure 2.15). The decay length of a typical B-hadron inside a very high p_T jet with $p_T \sim$ 1 TeV is:

$$L_B = c\tau_B\gamma_B = c\tau_B[p(B)/m_B]$$

where τ_B is the lifetime of B-hadrons and γ_B the relativistic boost. Since the typical fragmentation energy of B-hadrons in b-jets is of the order of 70% and $c\tau_B \sim 500\mu$, L_B is in the range of 5 to 10 cm. This means that the B-hadron decay tracks miss in general the B-layer hit. As a result the impact parameter suffers from large errors and the 2D/3D algorithms have a poor discriminating power. This poor performance affects also initially the SV1 algorithm, but may be recovered by a constrained fit of the secondaries to a common secondary vertex. The SV1 performance is therefore expected to be less degraded than the 2D/3D performance. In fact, at very high p_T, the 3D algorithm yields no improvement over 2D, so 2D is used in the following.

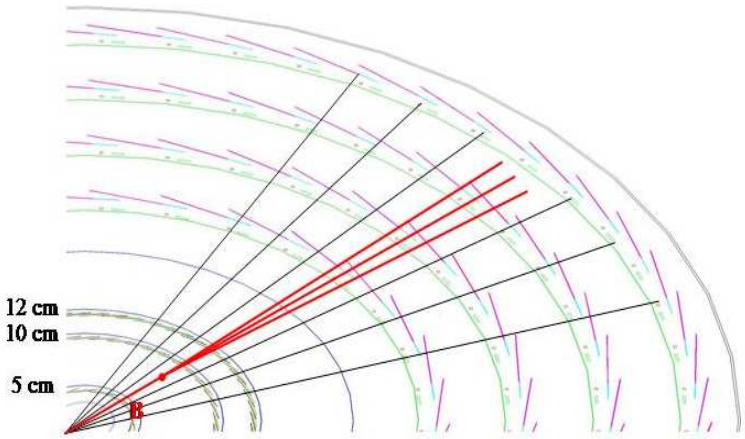

Figure 2.15: *A typical very high* p_T *b-jet with* $p_T \sim$ *1 TeV decaying inside the pixel detector: B-layer (5 cm), pixels (10 and 12 cm).*

In order to verify these hypotheses, the $Z_H(2000)$ event samples were reprocessed using the Athena version 10.0.1 (Rome production, see chapter 4). This version incorporates a fully operative SV1 vertex tagger. The AOD output in this simulation is otherwise similar of the CBNT output of the DC1 production. This analysis of these data confirms that the SV1 tagger provides a better performance at very high p_T (see figure 2.16 and 2.17, showing the integrated rejection R_u as a function of ϵ_b and the rejection R_u as a function of p_T for a fixed efficiency $\epsilon_b = 40\%$, respectively).

More recently, the samples have been reprocessed using Athena version 12.0.6 (CSC samples) that incorporates a more realistic simulation of the ATLAS detector, and the results are under study (see section 5.3.1).

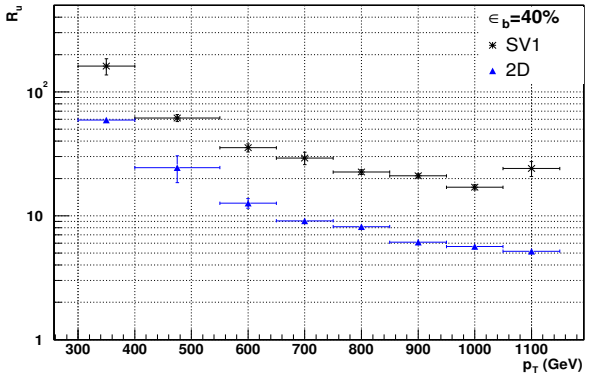

Figure 2.16: *Rejection of u-jets versus b-tagging efficiency (ϵ_b) using the 2D (DC1, Rome) and the SV1 (Rome) algorithms. The Z_H (2000) sample has been used to obtain this plot.*

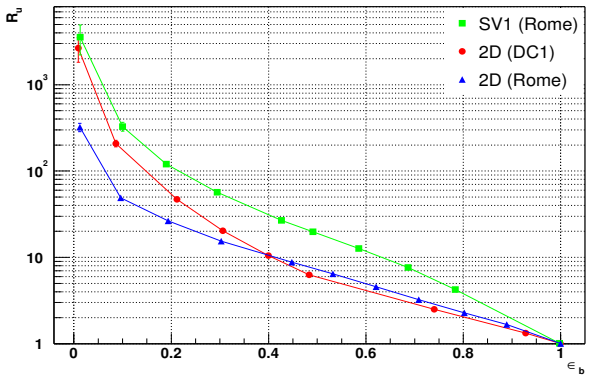

Figure 2.17: *Rejection of u-jets versus average* p_T *of b-jets for ϵ_b = 40% using the 2D and SV1 algorithms. The Z_H (2000) sample has been used to obtain this plot.*

As a final comment, the *b*-tagging performance at very high p_T used in the next chapter, based uniquely on the 2D algorithm, has been proved to be too optimistic in view of the previous results. This performance can however be improved by using the SV1 algorithm, ensuring the full validity of the results.

Chapter 3

Little Higgs searches

The Little Higgs model is a novel theory introduced to solve the hierarchy problem of the Standard Model (SM) without Supersymmetry (SUSY). In this model, new heavy particles cancel the radiative corrections of the Standard Model to the Higgs mass, allowing a low mass Higgs boson. These new particles are within the reach of the LHC. The simplest model based on the Little Higgs ideas is the so-called Littlest Higgs model, which has been used as reference for the simulation studies presented in this chapter.

The principles and phenomenology of this model are discussed in section 3.1. Finally, the possibility to discover the leptonic and hadronic decay channels of heavy gauge bosons using the ATLAS detector are discussed in sections 3.2 and 3.3, respectively.

3.1 The Little Higgs model

A low mass Higgs boson is problematic within the SM, as discussed before. Various new theories have been proposed in order to solve the problem. In supersymmetric theories, the divergent fermionic loops contributing to the Higgs mass cancel by additional bosonic loops. The supersymmetry breaking scale should be close to 1 TeV in order to avoid fine tuning [46]. Recently, a new class of theories with the desired cancellation of quantum corrections to the Higgs mass have been proposed, with fermionic loops canceled by additional fermionic loops. These are the so-called Little Higgs theories [47]. In this new model, the Higgs appears as a pseudo-Goldstone boson (i.e. a massive Goldstone boson) resulting from a spontaneously broken new global symmetry at large scale Λ_S. As in the case of supersymmetry, this scale should not be much larger than a few TeV in order to avoid fine tuning.

Several implementations of the Little Higgs scenario exist, depending on the global initial symmetry group, and the group resulting after breaking this symmetry. Some of the proposed models are $SU(5)/SO(5)$ [48], $SU(6)/SP(6)$ [49], $SU(3)^2/SU(3)$ [50] and more generally $SU(3)^n/SU(3)^k$ [51]. In all these theories the braking scale is Λ_S of the order of 10 TeV. For energies below the TeV scale, all these theories

approach the SM. The free parameters of the various models are however constrained in order to avoid deviations of the SM excluded by the very accurate LEP measurements.

All Little Higgs models predict new heavy particles. A heavy top-like particle T with 2/3 electric charge is required in order to cancel the contribution of the top quark to the Higgs mass. A heavy gauge boson sector is needed to cancel the loops of SM electroweak gauge bosons. A scalar sector appears as a result of the spontaneous symmetry breaking. All the new scalar particles are heavy due to radiative corrections, except the Higgs boson where quantum loops partially cancel, es explained before. All these new particles could have masses of the order of 1 TeV and would be therefore within the reach of LHC experiments.

The minimal model containing the Little Higgs mechanism is the so-called Littlest Higgs model [48]. In this model, the global symmetry $SU(5)$ is broken to $SO(5)$, yielding 14 massless Goldstone bosons. The $SU(5)$ group contains $[SU(2)\otimes U(1)]^2$, i.e. two copies of the SM electroweak symmetry $SU(2)\otimes U(1)$. This symmetry is broken into the SM symmetry, yielding heavy partners of the electroweak gauge bosons. These new bosons are called A_H, $\mathrm{W}_H^{\pm}$ and Z_H. The SM electroweak symmetry is broken at the Fermi scale to the U(1) symmetry of QED. The various breaking processes yield 7 massive gauge bosons (including Z and $\mathrm{W}^{\pm}$) and therefore only 7 out of the initial 14 Goldstone bosons survive. These scalars become massive particles called ϕ^0, ϕ^+, ϕ^{++} (with corresponding antiparticles), and the light Higgs boson of the SM, called h in the following.

Figure 3.1 shows the particle mass spectrum expected in the Littlest Higgs model, as described in reference [52]. The various parameters of the model, production cross sections and decay rates for the new particles are discussed below.

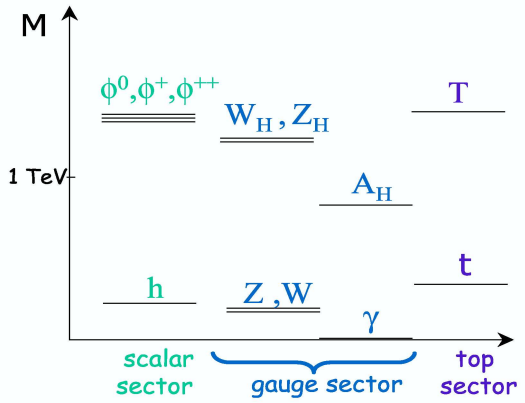

Figure 3.1: *Mass spectrum predicted by the Littlest Higgs model.*

3.1.1 Heavy top

In the SM, fermions become massive via Yukawa interactions with the Higgs boson. In the Littlest Higgs model, two Yukawa couplings are needed as well, in order to

give mass to the SM top and to the heavy top T. These couplings are called λ_1 and λ_2. They are related to the top mass and the T mass by the following relations

$$\mathrm{M}_T \approx f\sqrt{\lambda_1^2+\lambda_2^2} \qquad\qquad \frac{1}{\lambda_1^2}+\frac{1}{\lambda_2^2} \approx \frac{v^2}{\mathrm{m}_t^2}$$

where $\mathrm{f} = \Lambda_S/4\pi$ is proportional to the $SU(5)$ symmetry breaking scale and is of the order of 1 TeV, v is the Fermi scale (244 GeV), m_t is the top mass and M_T the T mass. These relations yield a lower limit on the T mass. An upper limit is also required to avoid fine tuning, so M_T should be inside the mass window

$$\frac{2\mathrm{m}_t}{v} f \leq \mathrm{M}_T \leq 2\ TeV(\frac{\mathrm{m}_h}{200\ GeV})^2$$

Figure 3.2 (left), extracted from reference [52], displays the T production cross-sections at the LHC ($\sqrt{s} = 14\ TeV$), as a function of M_T. The expected number of events is also indicated in the scale on the right side of the figure, for 3 years of data taking at high luminosity. The two production processes included in the figure are $T\bar{T}$ pair production (dashed line) and single T production (dotted line). It has been assumed that $\lambda_1 = \lambda_2 = 1$ in both cases. The dotted lines correspond to different choices of these parameters for the Wb single T process ($\lambda_1/\lambda_2 = 2$ and 1/2). The single T process is dominant for masses larger than 700 GeV.

The heavy top T has 3 dominant decay modes, with the following partial widths [52]:

$$\Gamma(T \to th) = \Gamma(T \to tZ) = \frac{1}{2}\Gamma(T \to bW) = \frac{\kappa_T^2}{32\pi}\mathrm{M}_T$$

where $\kappa_T = \lambda_1^2/\sqrt{\lambda_1^2+\lambda_2^2}$. Other decay modes are suppressed by a factor v^2/f^2.

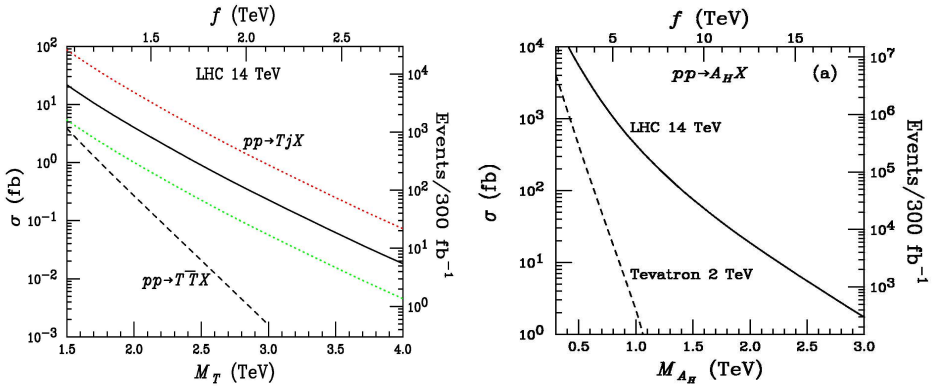

Figure 3.2: *Production cross-sections for the heavy top quark T versus mass (left) and cross-section for the gauge boson* A_H *versus mass (right).*

3.1.2 Heavy gauge bosons

The new vector boson sector includes A_H, Z_H and $W_H^{\pm}$, the massive partners of the photon γ and the gauge bosons Z and $W^{\pm}$, respectively. The neutral partner of the photon, A_H, is generally assumed to be lighter than the other heavy bosons. These heavy bosons, Z_H and $W_H^{\pm}$, have the same mass within this model, whereas the neutral partner of the photon, A_H, is generally assumed to be lighter and could therefore be produced at LHC, but also at Tevatron. The dominant production mechanism for A_H is $pp \rightarrow A_H X$. The cross-section as as function of mass is displayed in figure 3.2 (right side) for LHC energy.

The mass of A_H is dependent on various parameters of the model:

$$M_{A_H}^2 = m_Z^2 s_W^2 \frac{f^2}{5 s'^2 c'^2 v^2} + O(\frac{v^2}{f^2})$$

where $c_W = \cos\theta_W$ and $s_W = \sin\theta_W$, θ_W being the Weinberg angle. In the same way $s' = \sin\theta'$ and $c' = \cos\theta'$, where θ' is another mixing angle, a new free parameter of this model.

The expected dominant decays of A_H are $A_H \rightarrow f\overline{f}$ and $A_H \rightarrow Zh$. The couplings of A_H to fermions are however not fixed by the model and as a result decays and production rates cannot be calculated without further assumptions, like anomaly cancellations. The mass of the other heavy gauge bosons is:

$$M_{Z_H}^2 \approx M_{W_H}^2 \approx m_W^2 \frac{f^2}{s^2 c^2 v^2} + O(\frac{v^2}{f^2})$$

where $s = \sin\theta$ and $c = \cos\theta$, θ being another mixing angle, a new free parameter of this model. These relations provide again lower limits on the mass of the heavy gauge bosons. Upper limits can also be obtained to avoid fine tuning. As a result the following mass window is obtained:

$$m_W \frac{2f}{v} \leq M_{W_H} \approx M_{Z_H} \leq 6\ TeV(\frac{m_h}{200\ GeV})^2$$

In the Littlest Higgs model, the mass of the SM boson Z is also affected by corrections of the order of v^2/f^2. As a result, the tree level SM relation $M_W/M_Z = c_W$ is no longer valid. Precision measurements on electroweak observables can constrain the model and provide a lower bound to the scale f of a few TeV. The Littlest Higgs model may be modified in order to avoid this difficulty [53] [54], but these modified models are not considered here.

The cross-section for Z_H production at the LHC as a function of mass is presented in figure 3.3 (left side). The W_H cross section is larger by a factor of about 2. There is only one free parameter in the computation of these cross-sections, for a given mass of Z_H. This parameter is $\cot\theta$, where θ is the mixing angle defined before. The cross-section is proportional to $\cot^2\theta$. A value of $\cot\theta = 1$ has been

selected to compute the cross-section presented in figure 3.3. A large number of events containing these heavy gauge bosons are therefore expected at the LHC, about $5 \cdot 10^6$ Z_H events per year at high luminosity, for a mass of 1 TeV.

The branching ratios of Z_H versus cotθ are presented in figure 3.3 (right side). The dominant decay modes of Z_H are into fermion pairs and into the SM gauge boson Z in association with the Higgs boson. A general formula for the partial width of a heavy vector resonance into fermion pairs, neglecting all fermion masses, is

$$\Gamma(V \to f\overline{f'}) = \frac{C}{12\pi}(\mathrm{g}_V^2 + \mathrm{g}_A^2)\mathrm{M}_V$$

where $V = Z$ or W, C is a color factor (1 for leptons and 3 for quarks), and finally g_V and g_A are the vector and axial couplings of V_H to fermions. To leading order, the values of g_V and g_A are:

$$for\ \mathrm{Z}_H : \mathrm{g}_V = -\mathrm{g}_A = \frac{g}{2} cot\theta \mathrm{T}_3; \quad for\ \mathrm{W}_H : \mathrm{g}_V = -\mathrm{g}_A = \frac{g}{2\sqrt{2}} cot\theta$$

where g is the SM electroweak coupling constant and T_3 the third component of weak isospin for fermion f. The partial width of the heavy gauge boson decay to the corresponding SM gauge boson and a Higgs boson is:

$$\Gamma(\mathrm{V}_H \to Vh) = \frac{g^2 cot^2 2\theta}{192\pi}\mathrm{M}_{\mathrm{V}_H}$$

where all constants are already defined. In figure 3.3 (right side) the branching ratios of all these decays are displayed as a function of cotθ. For values of cotθ smaller than 0.5, the decay into $V\,h$ is dominant. This decay mode allows discrimination with other theories predicting heavy gauge bosons, and is therefore particularly relevant in order to test the model. For large values of cotθ, the decay rate into any quark pair ($b\overline{b}$, $t\overline{t}$, ...), or the sum of charged lepton pairs, is equal to 12 %. This result is however common to many other theories.

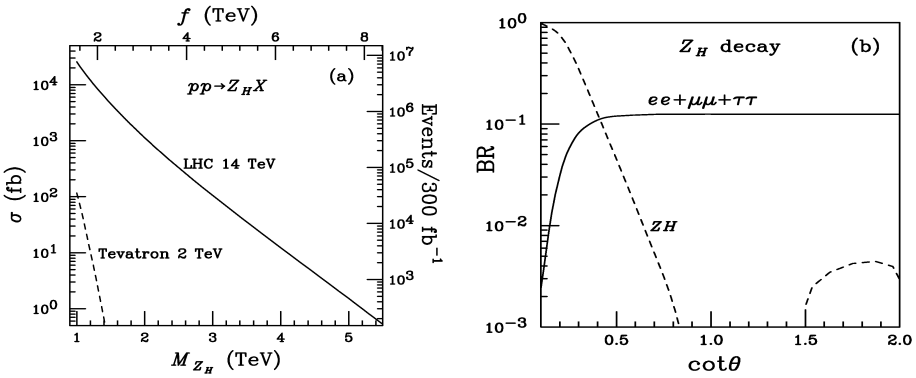

Figure 3.3: *Production cross-section for* Z_H *versus mass for cotθ = 1 (left) and branching ratios for the different decay modes versus cotθ (right).*

3.1.3 Scalar sector

The new scalar particles ϕ^0, ϕ^+, and ϕ^{++}, are expected in the model with a large mass. At leading order, these particles are almost degenerate in mass. The mass formula is

$$M_\phi^2 \approx \frac{2m_h^2 f^2}{v^2} \frac{1}{[1-(4v'f/v^2)^2]}$$

where v' is a new free scale parameter. The various constraints within the model yield the following mass limits:

$$\frac{\sqrt{2}m_h}{v} f \leq M_\phi \leq 10\ TeV$$

In this scalar sector, the doubly charged particle ϕ^{++} is particularly relevant to test the model. If the coupling to W pairs is not too small and the mass is not too heavy, it may be accessible at LHC. The ϕ^{++} has a peculiar signature since it is a doubly charged particle. The decay $\phi^{++} \rightarrow W^+W^+$ is dominant ($\sim$ 100%) and should provide the best signature at LHC.

3.2 Search for leptonic decays of Z_H and W_H using the ATLAS detector

The Littlest Higgs model predicts the existence of one charged (W_H) and two neutral (Z_H and A_H) heavy gauge bosons. W_H and Z_H are almost degenerate in mass and are expected to be heavier than A_H. All these bosons are likely to be discovered via their decays into leptons. However, in order to distinguish these gauge bosons from those predicted by other models, the characteristic decays $Z_H \to Zh$ and $W_H \to Wh$ must be observed, where h is the standard Higgs boson. Once the masses of the new bosons are specified, the angle θ determines the couplings of Z_H and θ' those of A_H. In the case of Z_H, the branching ratios into e^+e^- and $\mu^+\mu^-$ increase with $\cot\theta$ and reach an asymptotic value of 4%. A search for the decays $Z_H \to l^+l^-$ and $W_H \to l\nu$, l being an electron or a muon, is described below using the ATLAS experiment. Further details of these searches can be found in [55].

3.2.1 Measurement of $Z_H \to l^+l^-$ and $W_H \to l\nu$

A search for a peak in the invariant mass distribution of either e^+e^- or $\mu^+\mu^-$ is sensitive to the presence of A_H or Z_H. As an example, figure 3.4 shows the e^+e^- mass distribution arising from a Z_H of mass of 2 TeV for $\cot\theta = 1$ and $\cot\theta = 0.2$. The production cross-section for the former (latter) case is 1.2 (0.05) pb [52]. Events were required to have an isolated e^+ and e^- of $p_T > 20$ GeV and $|\eta| < 2.5$ which provides a trigger. The Standard Model background shown in the plot arises from the Drell-Yan process. In order to establish a signal, a minimum of 10 events in the peak of at least 5σ significance is required. Figure 3.5 shows the accessible region as a function of $\cot\theta$ and M_{Z_H}. A similar search for A_H can be carried out, but as discussed before, the couplings of A_H to fermions are not fixed by the model.

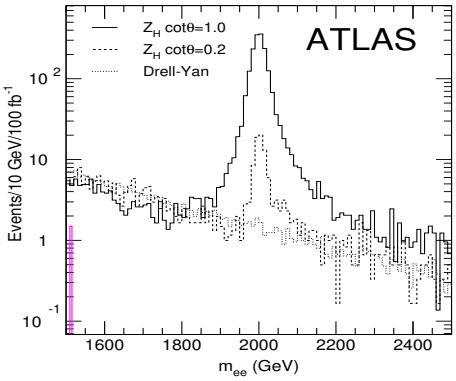

Figure 3.4: *Plot showing the e^+e^- mass distribution arising from a Z_H of mass of 2 TeV for $\cot\theta = 1$ (upper, solid, histogram) and $\cot\theta = 0.2$ (middle, dashed, histogram). The lowest, dotted histogram shows the distribution from background only.*

The decay $W_H^{\pm} \to \ell\nu$ manifests itself via events that contain an isolated charged lepton and missing transverse energy. Events were selected by requiring an isolated electron with of $p_T > 200$ GeV, $|\eta| < 2.5$ and $E_{\mathrm{T}}^{\mathrm{miss}} > 200$ GeV. The transverse mass from $E_{\mathrm{T}}^{\mathrm{miss}}$ and the observed lepton is calculated and the signal appears as a peak in this distribution. The main background arises from $\ell\nu$ production via a virtual W. The same requirements as before are imposed in order to establish a signal. Figure 3.5 shows the accessible region as a function of $\cot\theta$ and M_{W_H}.

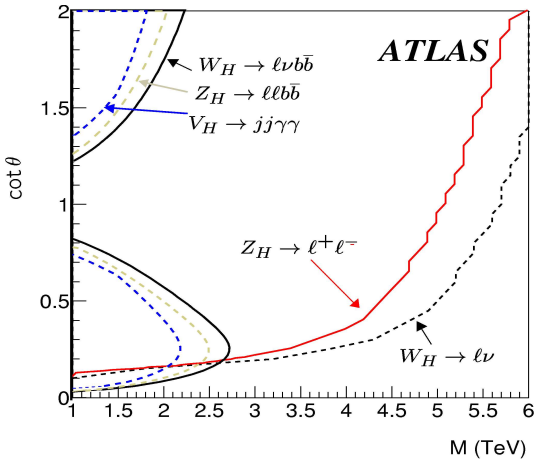

Figure 3.5: *Plot showing the accessible regions for* 5σ *discovery of the gauge bosons* W_H *and* Z_H *as a function of mass and* $\cot\theta$ *for various final states. The regions to the left of the lines are accessible with 300* fb^{-1}.

3.2.2 Measurement of $Z_H \to Zh$ and $W_H \to Wh$

Observation of the cascade decays $Z_H \to Zh$, $A_H \to Zh$, and $W_H \to Wh$ provides crucial evidence that an observed new gauge boson is of the type predicted in the Little Higgs Models. The decay $Z_H \to Zh$ is discussed first. For a choice of 120 GeV for the Higgs mass, the decay $Zh \to \ell^+\ell^- b\bar{b}$ results in a final state with two $b-$jets that reconstruct to the Higgs mass and a $\ell^+\ell^-$ pair that reconstructs to the Z mass. The coupling $Z_H Zh$ is proportional to $\cot 2\theta$. When combined with the coupling of Z_H to quarks that controls the production cross-section, the $\cot\theta$ dependence of the rate is shown in figure 3.6, which illustrates the difficulty of observation when $\cot\theta \approx 1$.

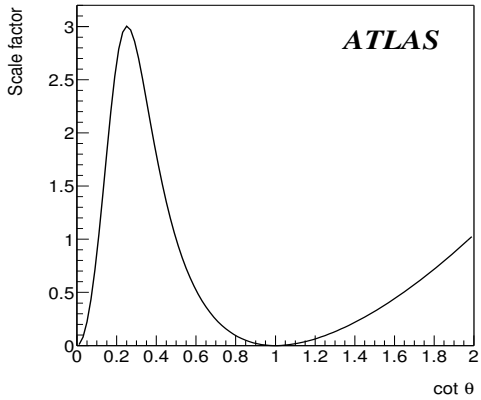

Figure 3.6: *The* $\cot\theta$ *dependence of the production rate times branching ratio for process* $Z_H \to Zh$.

A signal for the decay $Z_H \to Zh$ is extracted using the event selection described below.

- Two leptons of opposite charge and same flavor with $p_T > 6(5)$ GeV for muons (electrons) and $|\eta| < 2.5$. One of them is required to satisfy $p_T > 25$ GeV in order to provide a trigger.
- The lepton pair has a mass between 76 and 106 GeV
- Two reconstructed $b-$jets with $p_T > 25$ GeV and $|\eta| < 2.5$, which are within $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} < 1.5$.
- The $b-$jet pair should have a mass between 60 and 180 GeV.

The mass of the reconstructed Zh system is shown in figure 3.7 for a Z_H mass of 1 TeV and $\cot\theta = 0.5$. The presence of a leptonic Z decay in the signal ensures that the background arises primarily from $Z + jet$ final states.

A similar method can be used to reconstruct the $W_H \to Wh \to \ell\nu b\overline{b}$ decay. The $b-$jet selections were the same as above while the lepton selection is now as follows.

- One isolated e or μ with $p_T > 25$ GeV and $|\eta| < 2.5$.
- $E_{\mathrm{T}}^{\mathrm{miss}} > 25$ GeV.

The missing transverse energy is assumed to arise only from the neutrino in the leptonic W decay, and the W momentum is then reconstructed. The background, dominated by $W + jets$ and $t\bar{t}$ events, is larger than in the previous case, nevertheless a signal can be extracted.

The decay $h \to \gamma\gamma$ has a much smaller rate but provides a very characteristic signal. Figure 3.5 includes the accessible regions for decays involving a Higgs boson.

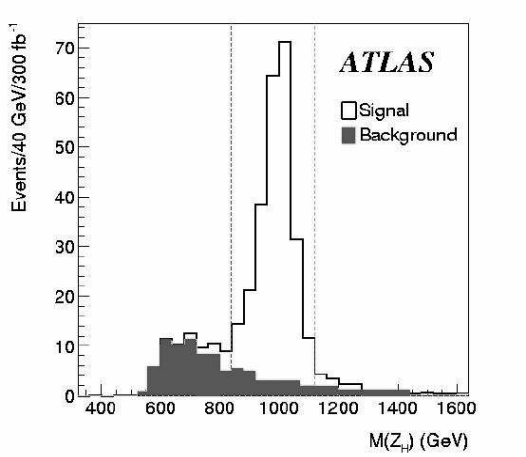

Figure 3.7: *Invariant mass of the Zh system reconstructed from the $\ell^+\ell^- b\bar{b}$ final state showing the signal from a Z_H of mass 1000 GeV with* $\cot\theta = 0.5$.

3.2.3 Discovery potential of ATLAS

The decays $Z_H \rightarrow l^+ l^-$ and $W_H \rightarrow l\,\nu$ can be observed by ATLAS up to 5-6 TeV heavy gauge boson masses, if $cot\theta = 1$ assuming $\mathcal{L} = 300 fb^{-1}$. For other values of $cot\theta$, the result is shown in figure 3.5. The decay channels $Z_H \rightarrow Z\,h$ and $W_H \rightarrow W\,h$ can also be observed, but in this case the accessible region is much more reduced and depends strongly on the value of $cot\theta$.

3.3 Search for hadronic decays of Z_H and W_H using the ATLAS detector

The heavy gauge bosons Z_H and W_H are predicted by the Littlest Higgs model with equal mass. The potential discovery channels for Z_H and W_H at the LHC are provided by the leptonic decay modes as discussed in the previous section. The conclusion was that the heavy gauge bosons can be observed using the leptonic decays $Z_H \to l\,\overline{l}$ and $W_H \to l\,\overline{\nu}_l$ up to masses of 5 to 6 TeV, for cotθ values around 1. The lepton l may be either an electron or a muon, but the electronic channel is preferred due to a better invariant mass resolution.

In the following, the possibility to identify the much more challenging hadronic channels is investigated: $Z_H \to b\,\overline{b}$, $Z_H \to t\,\overline{t}$ and $W_H \to t\,b$. These channels are more difficult to detect but may provide very useful information on Z_H and W_H couplings to quarks, once a resonance is found in any of the leptonic modes discussed before.

3.3.1 Cross-sections and Branching Ratios

According to theory [52], the couplings of Z_H and W_H to fermions depend on a single parameter, the θ mixing angle. In practice $\cot\theta$ is used to compute production and decay rates. This value is not really constrained by the model, but for practical reasons, the value $\cot\theta = 1$ is assumed as an example in the following. The cross-sections at LHC have been calculated using the PYTHIA [56] Monte Carlo generator for hadronic collisions, with proper fermionic couplings for Z_H and W_H to quarks and leptons. The structure functions of the proton called CTEQ5L have been used selected to compute all cross-sections. These production cross-sections for heavy boson masses of 1 and 2 TeV are listed in the table 3.1.

M (TeV)	σ (Z_H) pb	σ (W_H) pb
1	19	37
2	0.9	1.9

Table 3.1: *Z_H and W_H production cross-sections at LHC assuming* $\cot\theta = 1$.

As cross-sections depend on $\cot^2\theta$, the result for any value of $\cot\theta$ is obtained by scaling the values presented in the table.

The Branching Ratios and total widths of Z_H and W_H can also be computed from the information provided in [52]. Neglecting QCD corrections and fermion masses (also for top quark), the total widths of Z_H and W_H are equal. The result is:

$$\Gamma/M = [\ 3.4\ \cot^2\theta + 0.071\ \cot^2 2\theta\ ]\ \% \tag{3.1}$$

where the first term includes all possible fermion decays and the second term only the higgs decays $Z_H \to Z\ h$ and $W_H \to W\ h$, where h is the Standard Model Higgs boson. The Branching Ratio for a specific channel is easily calculated dividing the

corresponding partial width by the total width. The result is displayed in the following table 3.2 as a function of cotθ:

$\cot\theta$	BR($Z_H \to b\,\overline{b}$)	BR($W_H \to t\,b$)	Γ(M = 1 TeV) (GeV)
0.1	0.24%	0.48%	17.7
0.5	11.9%	23.8%	8.90
1.0	12.5%	25%	34.0
1.5	12.5%	25%	76.6
2.0	12.5%	25%	136.4

Table 3.2: *Branching ratios for the decays* $Z_H \to b\,\overline{b}, t\,\overline{t}$ *and* $W_H \to t\,b$ *and total width of* Z_H *and* W_H*, as function of* $\cot\theta$.

For $\cot\theta = 1$, the branching ratio of the $Z_H \to b\,\overline{b}, t\,\overline{t}$ decays is equal to 12.5% and the branching ratio of $W_H \to t\,b$ decay is equal to 25%. In fact, these values stay constant as function of $\cot\theta$, except at lower values, where the decay into the Higgs boson becomes dominant.

As a result of gauge symmetry, all branching ratios of Z_H and W_H into fermions are equal, except for the color factor of 3 (in case of quarks) or 1 (in case of leptons). It follows that

$$\Gamma(Z_H \to b\overline{b}) = \Gamma(Z_H \to t\overline{t}) = \frac{3}{24} \cdot \sum \Gamma(Z_H \to f\overline{f})$$

$$\Gamma(W_H \to tb) = \frac{3}{12} \cdot \sum \Gamma(W_H \to f\overline{f}')$$

and therefore the branching ratios shown in table 3.2 are calculated using equation 3.1 in the following way:

$$BR(Z_H \to b\overline{b}) = BR(Z_H \to t\overline{t}) = \frac{(3/24) \cdot 3.4 \cdot \cot^2\theta}{3.4 \cdot \cot^2\theta + 0.071 \cdot \cot^2 2\theta}$$

$$BR(W_H \to tb) = \frac{(3/12) \cdot 3.4 \cdot \cot^2\theta}{3.4 \cdot \cot^2\theta + 0.071 \cdot \cot^2 2\theta}$$

For $\cot\theta = 1$ (therefore $\cot 2\theta = 0$), the Branching Ratios into heavy quarks are:

$$BR(Z_H \to b\,\overline{b}) = BR(Z_H \to t\,\overline{t}) = 1/8 = 12.5\% \tag{3.2}$$

$$BR(W_H \to t\,b) = 1/4 = 25\% \tag{3.3}$$

3.3.2 Event simulation

Most of the events have been generated using the Monte Carlo program PYTHIA coupled to ATLFAST [57] to simulate the detector response. ATLFAST is the ATLAS fast simulation, i.e. a simulation where the detector response is parametrized

without any previous reconstruction of charged tracks in the inner detector or showers in the various calorimeters. This fast simulation is sufficient in the analysis presented here, except for b-jet identification, where a detailed information on all charged tracks inside jets is required , as discussed before [chapter 2]. However, the b-tagging information is supplied by the rejection curves presented in chapter 2. These curves are used in the following whenever b-jet identification is needed.

The Monte Carlo generator ALPGEN [58] has also been used. This MC includes the exact matrix elements for a large number of multi-parton processes of interest in the studies made at of the Tevatron and the LHC. ALPGEN provides therefore a more realistic computation of $W + jets$ events than PYTHIA. However, the fragmentation and hadronisation steps of PYTHIA for $W + jets$ events, are also used in ALPGEN.

The heavy gauge bosons Z' and W' of PYTHIA were used for Z_H and W_H. The corresponding couplings were taken from reference [52]. The production mechanisms 141 ($f_i\overline{f}_i \rightarrow \gamma/Z/Z'$) and 142 ($f_i\overline{f}_i \rightarrow W'$) were selected to generate Z_H and W_H events, respectively.

The following sets of events were used in the analysis:

- **Signal**: $Z_H \rightarrow b\,\overline{b}, t\,\overline{t}$ and $W_H \rightarrow t\,b$ with $M(Z_H) = M(W_H) = 1$ and 2 TeV. Samples of $10,000$ events were generated for each decay and mass value except for $Z_H \rightarrow t\,\overline{t}$ ($20,000$ events) and $W_H \rightarrow t\,b$ ($100,000$ events) in both cases with $M = 1$ TeV.

- **Background**: Standard Model final states containing $b\overline{b}$, $t\overline{t}$, $2\,jets$ and $W + jets$. Since only events with very high p_T jets and large center-of-mass energy are relevant in the analysis, it was imposed at the generation level that $\sqrt{\hat{s}} > 0.5$ (1.0) TeV and $\hat{p}_T > 0.1$ (0.25) TeV, where $\sqrt{\hat{s}}$ and $\hat{p}_T$ are the center-of-mass energy and transverse momentum of the hard process. The low cuts were used in the background estimation for M = 1 TeV, and the high cuts for M = 2 TeV. Samples of $100,000$ events were generated for each final state, except for the $t\,\overline{t}$ sample with low cuts (2M events).
  Other backgrounds like single top production and $Z + jets$ events were not simulated since the expected yields for the present analysis are negligible. For single top production, the program ZTOP [59] predicts, for $p_T(top) >$ 100 GeV, a cross-section of $\sigma = 30$ pb, to be compared with an expected cross-section for double-top of $\sigma = 164$ pb. The cross-section for $Z + jets$ events is expected to be a factor 10 smaller than the cross-section for $W + jets$, once the leptonic branching ratios of the gauge bosons are taken into account. Furthermore, $Z + jets$ events are suppressed in the present analysis by requiring missing energy.

The MC event samples are such that the statistical fluctuations correspond approximately to the following integrated luminosities for real data:
$\mathcal{L} = 3 \cdot 10^4\ pb^{-1}$ for M = 1 TeV and $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$ for M = 2 TeV.

Backgrounds can be 'reducible' or 'irreducible'. Reducible backgrounds are confused with the signal because they have a different final state but a similar detector

response. Irreducible backgrounds have the same final state but a different kinematics. The cross-section for reducible and irreducible events are reported in tables 3.3 and 3.4. The selection cuts are tuned to eliminate those backgrounds as much as possible They are described in the next section.

M (TeV)	σ $(gg \to jj)$ pb	σ $(W + jets)$ pb
1	303000	42.4
2	5700	1.8

Table 3.3: *$gg \to jj$ and $W + jets$ reducible background cross-sections at LHC.*

M (TeV)	σ $(gg \to b\overline{b})$ pb	σ $(gg \to t\overline{t})$ pb
1	391	164
2	7.7	6.4

Table 3.4: *$gg \to b\overline{b}$ and $gg \to t\overline{t}$ irreducible background cross-sections at LHC.*

3.3.3 Event selection

The cuts used in the present analysis are summarized below. Many of the cuts used in this analysis have been used in previous ATLAS work (see for example reference [55]). A search for resonances decaying into top quarks is also found in the Physics TDR, section 18.1.4.2 [60], where some of the cuts used here have been applied as well.

For the analysis of the decay $Z_H \to b\,\overline{b}$, the following cuts are applied for M = 1 (2) TeV, respectively:

- 2 b-jets with $p_T > 250$ (500) GeV and $|\eta| < 2.5$.

The b-tagging of very high p_T jets for this mode and also for the following modes is discussed in the next section. Two background sources have been analyzed: $b\,\overline{b}$ production (irreducible background) and 2 jet production (reducible background).

For the analysis of the decay $Z_H \to t\,\overline{t}$, the following cuts are applied for M = 1 (2) TeV, respectively, assuming that one of the t-quarks decays leptonically:

- A lepton (electron or muon) with $p_T > 25$ GeV and $|\eta| < 2.5$, and in addition $(E_T)_{miss} > 25$ GeV (this cut is also useful to fulfill the trigger requirements). The lepton p_T distribution before cuts is shown in figure 3.8.

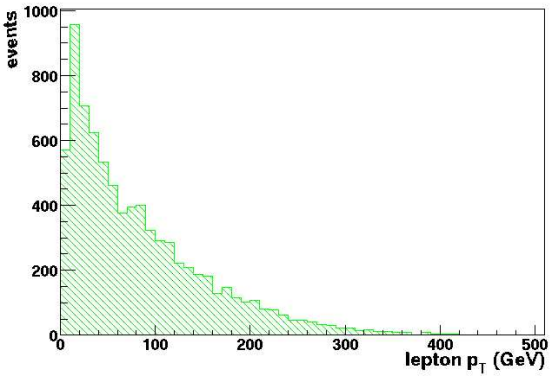

Figure 3.8: *Distribution of the lepton* p_T *before cuts for* $M = 1$ *TeV in the process* $Z_H \rightarrow t\,\overline{t}$*. The event scale is arbitrary.*

- 2 b-jets with $p_T > 25$ GeV and $|\eta| > 2.5$, the first jet such that ΔR(b-lepton) < 2 and the second such that ΔR(b-lepton) > 2. The distance ΔR is defined in the usual way as

$$\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$$

  where ϕ is the azimuthal angle and η is the pseudo-rapidity.

  It is noted that the ΔR cut is useful to identify the b-jet from the leptonic decay of the top-quark and the b-jet from the hadronic decay of the top-quark (see figure 3.9), but does not provide any significant discrimination against the background, since this background is mainly irreducible.

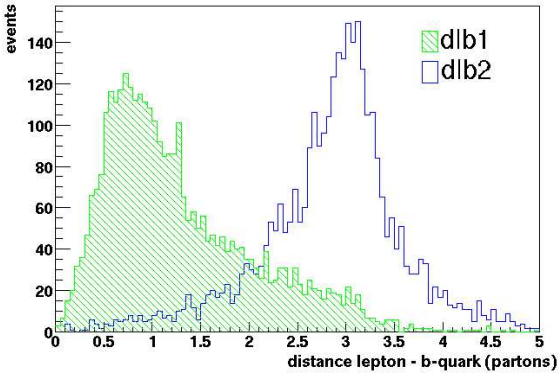

Figure 3.9: *Distribution of* ΔR*(b-lepton) in the process* $Z_H \rightarrow t\,\overline{t}$ *for* $M = 1$ *TeV. The b-jet from the top leptonic decay is called b1 and the b-jet from the top hadronic decay is called b2. The event scale is arbitrary. These distributions are obtained at the parton level.*

- The first t-quark is reconstructed using the first b-quark, the lepton, and the missing energy, assuming that the missing neutrino is parallel to the lepton. The second t-quark is reconstructed using the second b-quark and all other jets, within a distance of 2. The p_T of the reconstructed t-quarks (see figure 3.10) should satisfy the cut $p_T > 250$ (500) GeV.

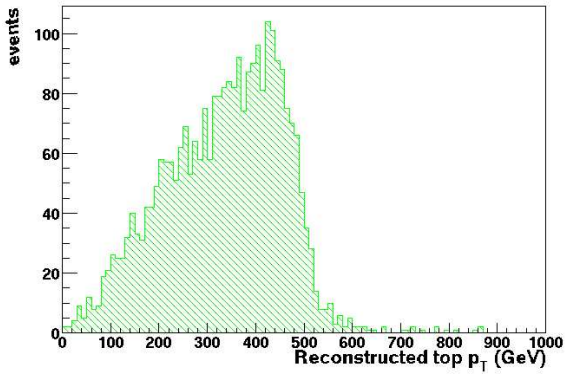

Figure 3.10: *Distribution of the reconstructed top p_T before cuts for $M = 1$ TeV in the process $Z_H \rightarrow t\,\overline{t}$. The event scale is arbitrary.*

The approximation of assuming the neutrino and the lepton parallel (inspired by τ searches in ATLAS) yields good results for M $>$ 1 TeV. The longitudinal momentum of the neutrino can also be calculated using the equation $m_W^2 = (p_l + p_\nu)^2$, where p_l and p_ν are the 4-momenta of the lepton and the neutrino. This second approximation yields also goods results for M $<$ 1 TeV, but has not been used in the present analysis.

For the analysis of the decay $W_H \rightarrow t\,b$, where the t-quark is assumed to decay leptonically, the selection proceeds in the same way as in the $Z_H \rightarrow t\,\overline{t}$ case, except that no energy is collected around the second b-quark, since in this case this b-quark is the only component of the second jet. For both decays, $Z_H \rightarrow t\,\overline{t}$ and $W_H \rightarrow t\,b$, the dominant background is $t\,\overline{t}$ production (irreducible). Another background source, $W + jets$ production (reducible), has been considered as well.

3.3.4 Signal and background

The number of signal and background events is calculated as follows:

$$N = \mathcal{L} \cdot \sigma \cdot \epsilon$$

where $\mathcal{L}$ is the luminosity ($3 \cdot 10^5 pb^{-1}$) and the other parameters are described below.

- **Signal**:
  σ is the Z_H and W_H production cross-section (see table 3.1), BR is the Branching Ratio of the Z_H and W_H decays (eq. 3.2 and 3.3) and ϵ is the total efficiency calculated as follows:

  $$\begin{array}{ll} Z_H \to b\,\overline{b}: & \epsilon = \epsilon_{kin} \cdot \epsilon_b^2 \\ Z_H \to t\,\overline{t}: & \epsilon = \epsilon_{kin} \cdot \epsilon_b^2 \cdot \epsilon_l \cdot 2 \cdot BR(W \to l\,\nu_l) \cdot BR(W \to j\,j) \\ W_H \to t\,b: & \epsilon = \epsilon_{kin} \cdot \epsilon_b^2 \cdot \epsilon_l \cdot BR(W \to l\,\nu_l) \end{array}$$

  where ϵ_b is the b-tagging efficiency, ϵ_l is the lepton efficiency ($\epsilon_l = 0.9$ is used in this analysis) and ϵ_{kin} is the kinematic efficiency (see table 3.5). Note that $BR(t \to W\,b) = 100\%$ is omitted in the previous formula.

- **Irreducible background**:
  σ is the irreducible background production cross-section (table 3.4) and ϵ is calculated as before using background events.

- **Reducible background**:
  σ is the reducible background production cross-section (table 3.3) and ϵ is a mixture of efficiencies for different types of jets. These efficiencies include kinematic efficiencies, branching ratios and b-tagging efficiencies. In this analysis, 5 different kinds of events, bc, bu, cc, cu and uu, are used. Therefore

  $$\epsilon = \frac{N_{bc} \cdot \epsilon_b \cdot \epsilon_c + N_{bu} \cdot \epsilon_b \cdot \epsilon_u + N_{cc} \cdot \epsilon_c^2 + N_{cu} \cdot \epsilon_c \cdot \epsilon_u + N_{uu} \cdot \epsilon_u^2}{N_{bc} + N_{bu} + N_{cc} + N_{cu} + N_{uu}}$$

  The c-tagging and u-tagging components of these efficiencies, resulting from the b-tagging analysis, are reported below.

The kinematic efficiencies achieved for signal events by the selection cuts described before are summarized in table 3.5. This efficiencies decrease at large mass values because leptons from top decays are less isolated.

M(GeV)	$b\overline{b}$	$t\overline{t}$	tb
1000	49%	27%	25%
2000	51%	21%	18%

Table 3.5: *Kinematic efficiencies (all BR and tagging efficiencies excluded).*

The branching ratios of W decays used in this analysis are the following:

$$BR(W \to j\,j) \sim 67\%$$

$$BR(W \to l\,\nu_l) \sim 22\%$$

being l an electron or a muon (τ lepton is excluded).

For all decays, the tagging of very high p_T b-jets is required. The average p_T of these jets is reported in table 3.6:

$< p_T >$	$Z_H \to b\bar{b}$	$Z_H \to t\bar{t}$	$W_H \to tb$
M = 1 TeV	400	200	250
M = 2 TeV	800	400	500

Table 3.6: *Average p_T (in GeV) of final state b-jets for the various channels studied.*

The tagging of b-jets has been studied in detail within ATLAS (see for example [61]), but in all these studies the average p_T of b-jets was typically below 200 GeV. In order to investigate the b-tagging performance for larger p_T values, the following event samples have been generated and analyzed:

$$Z_H(2\ TeV) \to b\,\overline{b},\ c\,\overline{c},\ u\,\overline{u}$$

Each sample contained $20,000$ events and was processed using the full reconstruction provided by ATLAS software. The b-tagging results are already discussed in the previous chapter.

The standard b-tagging algorithm was used. Since this algorithm is not optimized for high p_T jets, the results should be considered conservative, and further improvements might be expected in the future.

The efficiencies (ϵ) and rejections ($R = 1/\epsilon$) reported in table 3.7 have been applied in the analysis presented in this note.

mass	decay	ϵ_b	R_c	R_u
1 TeV	$b\bar{b}$	0.1	140	1000
	$t\bar{t}$	0.5	10	100
	tb	0.5	7	40
2 TeV	$b\bar{b}$	0.1	45	90
	$t\bar{t}$	0.2	28	130
	tb	0.2	26	75

Table 3.7: *Efficiencies and rejections applied in the b-tagging analysis.*

The reconstructed mass peaks are shown in figure 3.11. The corresponding mean values and widths resulting from gaussian fits are reported in table 3.8. Since the natural width of Z_H and W_H is 34 GeV for M = 1 TeV (assuming $\cot\theta = 1$), the reconstructed mass width is dominated by fragmentation and experimental effects (detector resolution). The invariant mass distributions showing the signal on top of the expected background are shown in figure 3.12. By applying mass cuts of $\pm$ 200 (400) GeV around the expected mass values of M = 1 (2) TeV, it is possible to obtain the number of signal and background events reported in tables 3.9 and 3.10 for an integrated luminosity of $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$, corresponding to 3 years of running at high luminosity. The significance (significance = $signal/\sqrt{background}$) of the expected signal is indicated as well for all decays. The significance has in general a large value around M = 1 TeV, well above the discovery threshold of 5,

but only the decay $W_H \to t\ b$ yields a signal which can be clearly separated from the background as shown in figure 3.12.

M (GeV)	$Z_H \to b\bar{b}$	$Z_H \to t\bar{t}$	$W_H \to tb$
1000	962 ± 120	994 ± 119	953 ± 111
2000	1949 ± 251	1938 ± 246	1855 ± 232

Table 3.8: *Central value of the reconstructed mass (M) and width* ($\pm\Delta M$) *for the various decay channels analyzed.*

M = 1 TeV	$b\bar{b}$	$t\bar{t}$	tb
signal	3461	13170	34178
irreducible bkg	63764	116130	23231
red. bkg (PYTHIA)	24388	126	257
red. bkg (ALPGEN)	—	1130	2191
significance (ALPGEN)	11.7	38.5	214.4

Table 3.9: *Signal and background inside a mass window of* ± *200 GeV around M = 1 TeV for* $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$ *and* $\cot\theta = 1$.

M = 2 TeV	$b\bar{b}$	$t\bar{t}$	tb
signal	174	74	203
irreducible bkg	1605	520	100
red. bkg (PYTHIA)	18623	2	3
red. bkg (ALPGEN)	—	30	48
significance (ALPGEN)	1.2	3.2	16.7

Table 3.10: *Signal and background inside a mass window of* ± *400 GeV around M = 2 TeV for* $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$ *and* $\cot\theta = 1$.

For the $t\,\bar{t}$ and $t\,b$ channels, the irreducible background consists of $t\,\bar{t}$ events and the reducible background of $W + jets$ events. The reducible background can be computed using both PYTHIA and ALPGEN. The values obtained with ALPGEN are larger and have been used, in a conservative approach, to calculate the various significances.

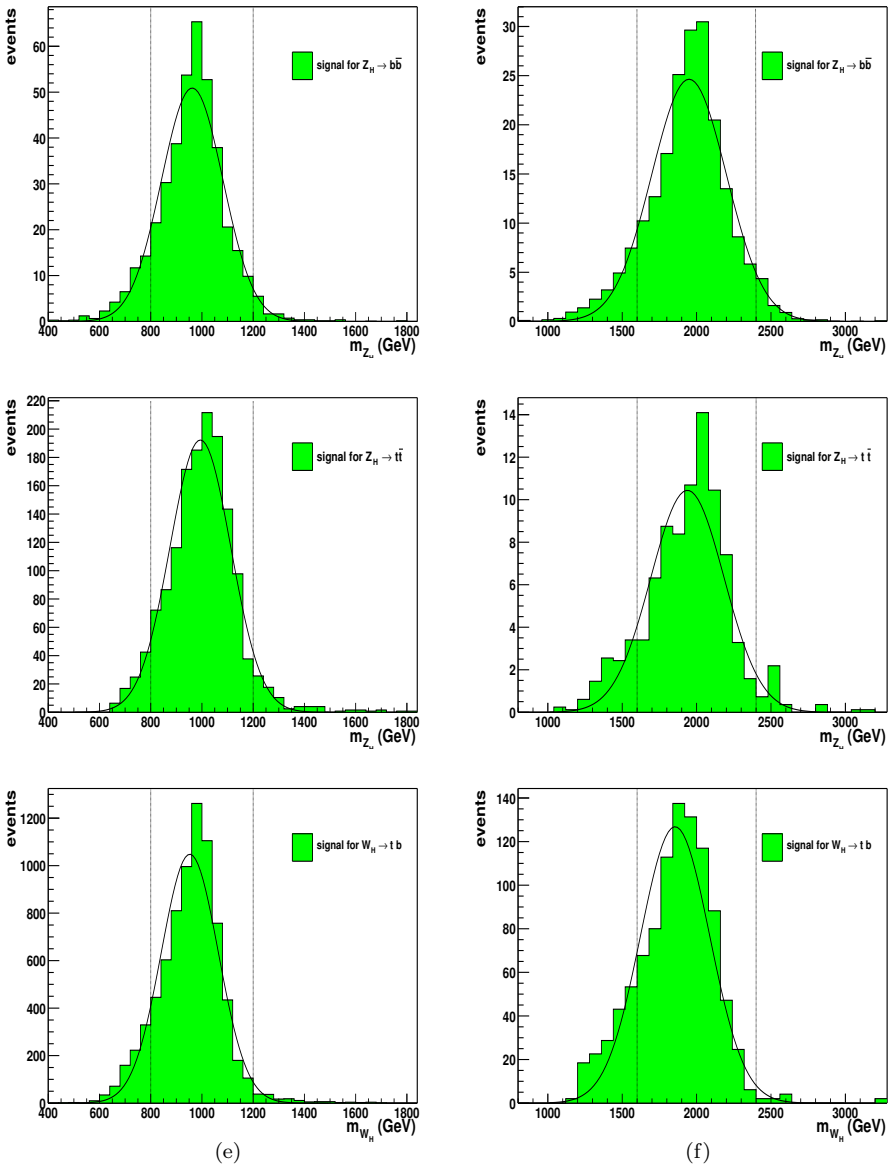

(e) (f)

Figure 3.11: *Reconstructed mass peaks for the 3 decay modes $Z_H \to b\,\overline{b}$, $t\,\overline{t}$ and $W_H \to t\,b$ and the 2 mass values, $M = 1$ and 2 TeV. The curves on top of the data are gaussian fits. The event scale is arbitrary.*

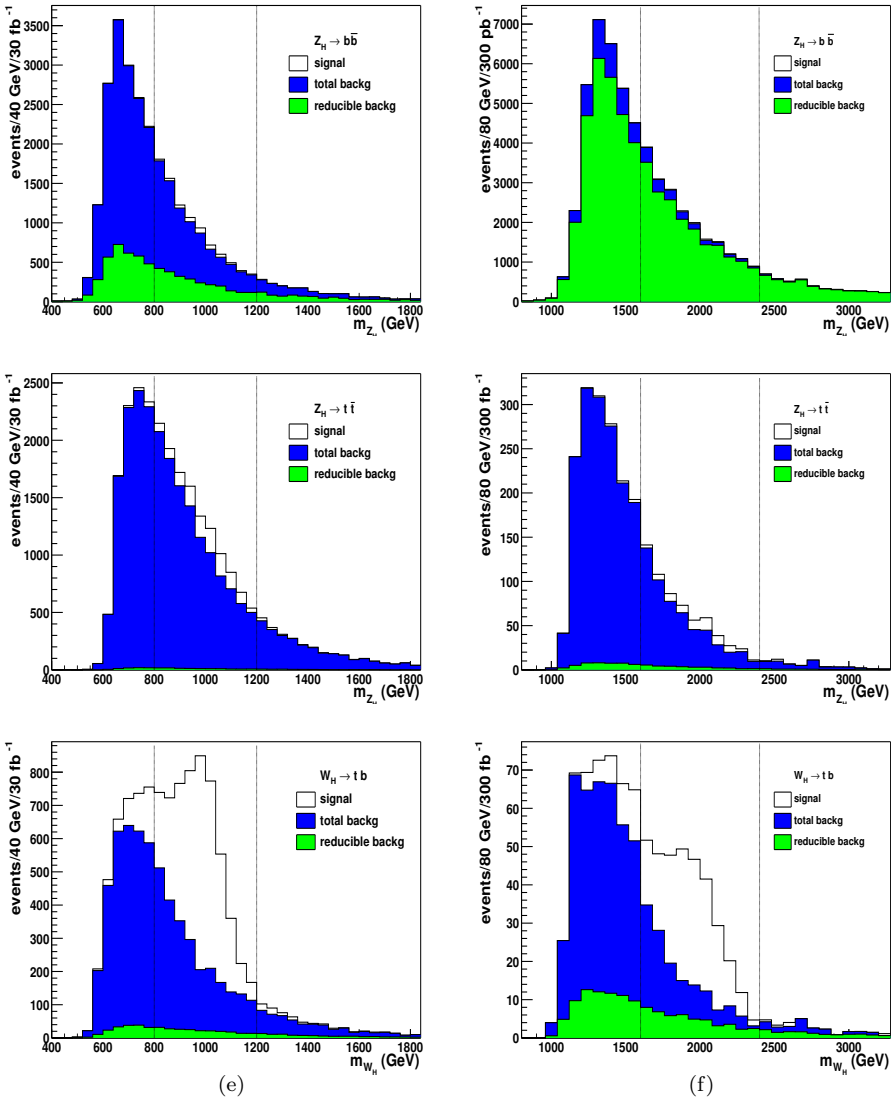

(e) (f)

Figure 3.12: *The same as in figure 3.11, but including both signal and background. The mass windows used to calculate the significance are indicated in the figures. A luminosity of* $\mathcal{L} = 3 \cdot 10^4\ pb^{-1}$ *is assumed for* $M = 1\ TeV$ *and* $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$ *for* $M = 2\ TeV$*. Note: reducible background obtained with ALPGEN.*

3.3.5 Discovery potential of ATLAS

According to the analysis presented before, it is possible to deduce the discovery regions for the various decays indicated in figure 3.13. The discovery region is defined as the region with significance exceeding the value of 5 and a number of expected events above 10. This region has been obtained assuming that the cross-section is proportional to $(\cot\theta)^2$ and extrapolating the results (for both signal and background) obtained for masses of 1 and 2 TeV to other mass values, assuming that the significance decreases exponentially with M (this hypothesis is found valid in most of heavy particle searches).

As a summary, the analysis of expected signals and backgrounds shows that the decays $Z_H \rightarrow b\,\overline{b}$ and $t\,\overline{t}$ are difficult to detect, but the decay $W_H \rightarrow t\,b$, on the contrary, might yield a signal clearly separable from the background, up to W_H masses of 2.5 TeV, if cotθ is assumed to be equal to 1. The detection of the $W_H \rightarrow t\,b$ allows a measurement of BR($W_H \rightarrow t\,b$) and, by comparison with BR($W_H \rightarrow l\,\nu$), provides an interesting test for the couplings of the W_H boson.

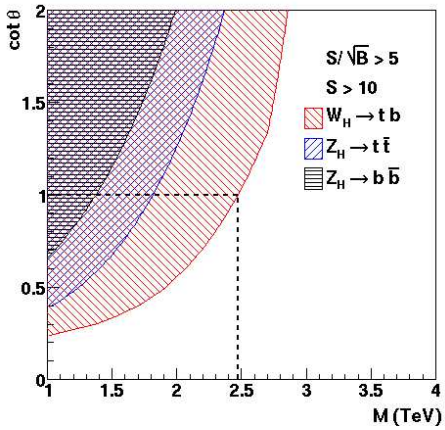

Figure 3.13: *Discovery region for the decay $Z_H \rightarrow b\,\overline{b}, t\,\overline{t}$, $W_H \rightarrow t\,b$ in the Mass-Cotθ plane, assuming an integrated luminosity of $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$.*

Chapter 4

ATLAS Production System and Data Challenges

The ATLAS collaboration is preparing for data taking, scheduled to start in 2008. After data collection by the on-line trigger, the expected data volume recorded for offline reconstruction and analysis is of the order of few Peta-bytes per year, to be processed by institutions distributed worldwide. Therefore the required computing resources are two orders of magnitude larger than for previous experiments in particle physics. The LHC experiments need a worldwide distributed data management and computing system to handle this huge amount of data. The LHC Computing Review in 2001 recommended that LHC experiments should carry out Data Challenges (DCs) of increasing size and complexity to be prepared for LHC data taking. The goal of the ATLAS Data Challenges is the validation of the ATLAS computing model.

A description of the ATLAS Production System (ProdSys), that allows to manage massive productions of Monte Carlo data using Grid resources, is shown in section 4.1. The main results obtained for ATLAS DCs are discussed in section 4.2, where a description of the IFIC contribution to the ATLAS Computing System Commissioning (CSC) simulation production running an ATLAS ProdSys instance is given also.

4.1 ATLAS Production System

The ATLAS experiment requires a large amount of simulated data. The ATLAS Production System (ProdSys) [62] allows to manage the production of a large amount of Monte Carlo data using Grid resources in an automatic way, with minimal human intervention. This system consists of a set of four components (see figure 4.1): a central production database (ProdDB), a Distributed Data Management (DDM) [63] system, a supervisor component and an executor component. In the following sections these components are described.

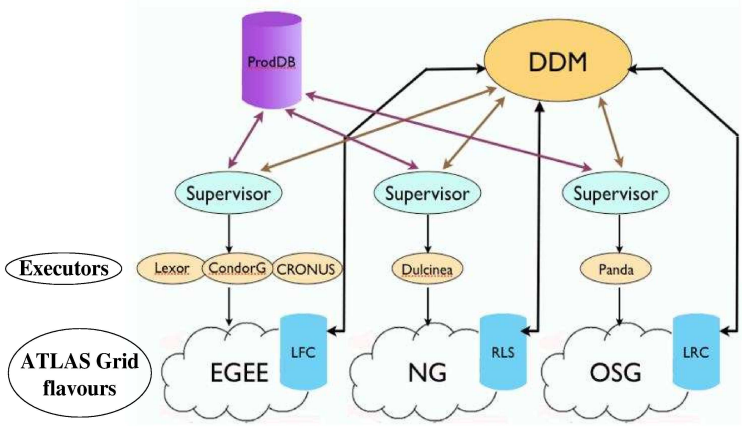

Figure 4.1: *ATLAS Production System architecture.*

The architecture of the Production System was designed in 2003 with the following goals in mind:

- The system should be as flexible as possible while retaining simplicity.
- The system should target automatic (i.e. with minimal human intervention) production on all three Grids in use by ATLAS in 2004, i.e. LHC Computing Grid (LCG) [28]/Enabling Grids for E-science in Europe (EGEE) [30], Open Science Grid (OSG)/Grid3 [31] and Advanced Resource Connector (ARC)/NorduGrid [32]. Legacy batch systems had also to be supported as back-up solution.
- The system should use Grid middleware services if possible.

The ATLAS ProdSys has been used since the ATLAS Data Challenge 2 (DC2) [64] and is currently being used for the ATLAS Data Challenge 3 (DC3, also called Computing System Commissioning, CSC, see section 4.2.5). Some changes and improvements have been applied in the meanwhile to the ATLAS ProdSys, for example, the supervisor component.

The ATLAS ProdSys distinguishes between two levels of abstraction: task and job (see figure 4.2). A task transforms input datasets into output datasets by applying a task transformation.

Datasets are usually quite large and consist of many logical files. In this case, a job transforms input logical files into output logical files by applying a job transformation. In this way, one could say that a task is split into several jobs and these jobs are managed by ATLAS ProdSys on the different Grid flavour resources, so the overall Production System relies on the performance of the individual Grid systems.

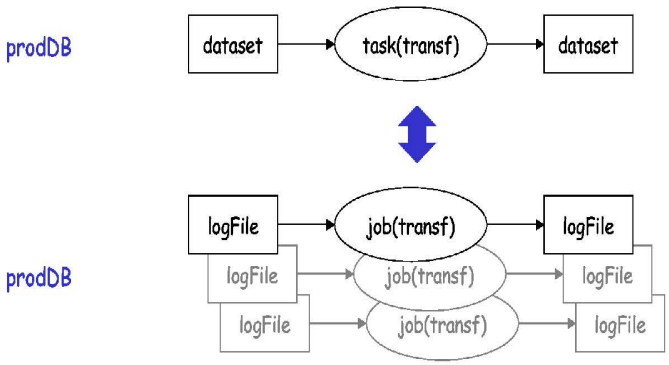

Figure 4.2: *Conceptual view of datasets and transformations.*

According to the ATLAS full simulation chain, one can classify these jobs as: Event Generation (evgen), Simulation (simul), Digitalization (digit) and Reconstruction (recon) jobs.

4.1.1 Distributed Data Management

The scope of the ATLAS Distributed Data Management system, called Don Quijote 2 (DQ2) [65], encompasses the management of file-based data of all types: event data, conditions data, user-defined file sets containing files of any type.

ATLAS, as mentioned previously, uses three Grids: LCG/EGEE, OSG/Grid3 and NDGF/ARC providing each a set of tools to manage distributed data. Therefore, there needs to be an ATLAS specific layer on top of the Grid middleware to bookkeep and present data in a form physicist expect, manage data flow and provide a single entry point to all distributed ATLAS data.

This ATLAS specific layer on top of the Grid middleware is called DQ2, which tries to hide as much as possible the Grid middleware for users.

DQ2 is based on the concept of versioned datasets, defined as a collection of files or other datasets, for example raw data files from a particular detector run.

ATLAS has DQ2 central catalogues which define datasets and their locations. In this case, a dataset is also the unit of data movement. There is no DQ2 global physical file replica catalogue, since it's expected to obtain more than 100,000 files and replicas created every day at the LHC start-up. The physical file resolution is done by Grid specific catalogues (see figure 4.1: LFC [66] for LCG/EGEE, RLS [67] for NorduGrid and LRC [68] for OSG/Grid3) at each Tier-1 cloud holding only data on that cloud.

In order to enable data movement, there is a set of distributed 'site services' which use a subscription mechanism to pull data to a site. As content is added to a dataset, then the site services copy it to subscribed sites.

On the other hand, there is a set of tools also for users to access this data, called DQ2 end-user tools [69].

4.1.2 Job Transformations

The job transformation is the script that sets the run-time environment, allows possible compilation of patches to software in a release, runs the Athena [25] executable, parses the log file for the known warning and error messages, and supervises everything up to the end. Currently it is implemented as a shell script. It can include any data file that is missing from the release or that could be needed for a particular job.

There are two kinds of transformations named production (system independent) and KitValidation (KVT, test oriented). They differ in the call method, with positional arguments for production and switched parameters (name-value pairs) for KVT. Building the transformation for production jobs is at present a manual operation; the transformation is then inserted into a Pacman [70] cache and loaded by Grid jobs at start-up.

Work is in progress to provide a generic transformation for all usages, including non-production jobs. The intention is to move to Python [71] (instead of shell) scripts, as in this way the transformation can be seen as an extension of the jobOptions [72]. The transformation will also check the integrity of the expected output file at the end of the job (this needs the knowledge of the number of output events as counted inside the Athena job).

4.1.3 Production Database

There is only a single central production database. The database holds different tables for task/job definitions and executions (see figure 4.3). Some of these records: job transformations, job definitions, job executions and logical files are explained in this section.

A job-transformation record describes a particular executable script for a given ATLAS software (Athena) release. The description includes the signature of the transformation, describing each formal parameter together with its type and its meta-type[1], indicating how the values should be passed to the executable.

Each job-definition record points to its associated job transformation record. Other fields allow to keep track of operational parameters as the current attempt number to execute a given job (lastAttempt), of the supervisor component handling the job supervisor, the job priority, etc. The bulk of the job definition is, however, stored as an XML [73] tree in the field jobXML. It includes the actual values to be assigned to the formal parameters of the transformation and additional information about logical input and output files.

[1] An example of type is dataset (DS) and an example of meta-type is: inputDS, outputDS. A meta-type is a type that describes the properties of another type.

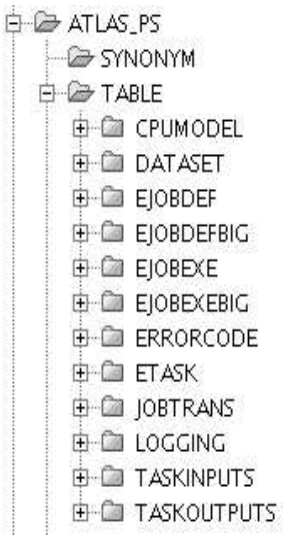

Figure 4.3: *ATLAS Production DataBase (ProdDB) tables.*

Each job definition can be associated with zero, one, or more job-execution records, corresponding to each attempt to execute the job. Each attempt has a unique number appended to the names (both logical and physical) of all files produced, ensuring interference-free operation even in the case of lost and/or zombie [2] jobs. The execution record includes information like start- and end-time of the job, resource consumption, the location where the outputs were stored, etc.

In the last table, logicalFile, the Production System stores all meta-data related to logical files. Most of the information is redundant with respect to the information stored in the meta-data catalogues of the Grids (size, guid, md5sum, logicalCollection).

The production database used in 2004-2005 for Data Challenge 2 and subsequent productions was implemented as an Oracle [74] database hosted at CERN. A MySQL version of the production database was also available for small-scale productions.

4.1.4 Supervisor

The next component, called the supervisor, is a common component for all ATLAS Grid flavours that picks free jobs from the production database and hands them on to one of the executors connected to it. The information about jobs is exchanged using XML, usually wrapped in XMPP [75] (using the Jabber protocol) or wrapped in SOAP [76] using web services. As its name suggests, the supervisor follows up the job, asking at regular intervals the job status until the job is declared 'done'

[2] A zombie job is a job for which the supervisor has lost its reference, due to a software problem with the executor or the underlying middleware, but is still running.

by the executor. At that point, the supervisor verifies for successful jobs the existence of all expected outputs and, if all is correct, renames them to their initial logical names by dropping the attempt number from the temporary logical name. Additionally, the files are added to the logicalFile table together with any meta-data produced by the job and returned by the executor. In case of failed jobs, the supervisor simply schedules the job for re-execution in the production database, so that it can be picked up again if the maximum number of attempts is below the maximum allowed.

The supervisor does not perform any resource brokering by themselves. A simple protocol is used to assign jobs on the basis of the executor capacity. The supervisor asks the executor how many jobs can be handled (possibly qualified with resource requirements) and the executor replies with a number (possibly qualified with characteristics). The supervisor may then send a number of jobs to the executor, that in turn may choose to process or refuse them. The non-binding nature of the protocol allows both very sophisticated and very simple implementations to co-exist on both the executor and supervisor code.

For efficiency reasons an implementation of the supervisor can have an internal state, but the design does not require this. All required state information is saved in the database. Having a stateless component obviously makes it more resilient against crashes, as they can be restarted at any time.

The supervisor implementation used in 2004-2005 for DC2 and Rome production is called Windmill [77]. Each Windmill instance is connected to a specific executor and manages all jobs processed by this executor.

Eowyn [78] is the second generation implementation of the supervisor component being at present used for Data Challenge 3 production (or also called Computing System Commissioning, see section 4.2.5). The most important difference to the first generation is the possibility of running the supervisor and the executor in a common instance reducing dramatically the complexity. This interface exchanges information using Python objects and requiring executors to provide immediate response. This again is a significantly simplification to the original asynchronous protocol. If required, an asynchronous message protocol can still be implemented in an additional layer.

To test Eowyn simulators have been written for the production database, the facility/executor and the Distributed Data Management system. This not only allows for rapid standalone development, testing all possible error conditions, but also allows simulation of the effect of certain conditions like DDM downtime or slowness, high job failure rates, ProdDB downtime or slowness, ... faster than real time. The simulators expose the same interface as the real components. The central Eowyn module does not see any difference. Most simulators allow the user to set the average time that a certain call takes and the average rate at which a certain call will fail.

Eowyn is designed to be multi-threaded. The activities of Eowyn can be split into three categories :

- Activities depending only on the production database,
- Activities depending on the production database and the executor and
- Activities depending on the production database and the DDM.

It can be expected that in case of problems with the executor, the other activities can proceed. A supervisor process spends most of its time waiting for either the database, the executor, or the DDM system. Multi-threading allows to minimize the cycles lost in waiting and, more importantly, avoids the problem of balancing the different activities.

4.1.5 Executors

The task of the executors is to interface the supervisor to the different Grid flavours or legacy systems. They translate the system-neutral job definition into the system-specific language. The executors implement a system-neutral interface with the usual methods: submit, getStatus, kill, etc. Again, the design does not require the executor to keep an internal state.

Five executors were developed and deployed in 2004-2005 and used for Data Challenge 2 and subsequent productions:

- Lexor [79] for the LCG/EGEE Grid flavour;
- Lexor-CondorG (Lexor-CG) [80] for the LCG/EGEE Grid flavour as a test of direct job submission to the Computing Elements;
- Dulcinea [81] for the NorduGrid;
- Capone [82] for Grid3;

Lexor

Lexor is the executor that submits jobs to the LCG/EGEE Grid flavour. Lexor has been fully coded in Python. First, a module implementing a generic LCG/EGEE job class was developed. The code is based on the SWIG [83] API to the workload-management system, developed by the European DataGrid (EDG) [84] project for Grid User Interface (UI) [85]. Via this API, this class offers methods for defining, submitting, and managing the job on the Grid. For job definition, several methods allow the manipulation of the underlying Job Description Language (JDL) [86] description, either by direct access to the Classified Advertisements (classAd) [87] attributes or via higher-level functions. Other methods are used to tune and perform the job submission, to monitor the job status (and extract single pieces of information from the status), and finally to retrieve the output sandbox of the job. A job-cancellation method is also provided, but is not yet used in the Production System.

When in submission mode, the supervisor periodically asks for the number of jobs that the executor is willing to manage. To answer this question, a query to the LCG information System is issued in order to retrieve the number of free CPUs in the sites matching the requirements included in the supervisor request. Computing this number turned out to be a complicated task, because of the aggregated nature of the information published by the LCG Information System. Often the real amount of available CPUs to the ATLAS VO [37] was overestimated, because some of the advertised CPUs were not accessible by ATLAS jobs. This problem was temporarily addressed by defining one queue for each VO, but a refining of the information scheme had to be considered. In the current version of the information system resource availability is published separately for each Virtual Organization.

Once submitted, a job is dispatched to the Computing Element (CE) selected by the broker and queued until the local batch system starts executing it on a Worker Node (WN). Before the required transformation can be started, several actions have to be performed on the WN, so the job has to be wrapped by an appropriate script. The wrapper has mechanisms to cope with temporary problems in the underlying Grid as well. An example of such problems is the occasional unavailability of ATLAS software, due to NFS [88] problems between the CE exporting the software directory and the WN. Jobs being executed on such hosts fail at the very beginning, thus freeing a WN that attracts an increasing number of jobs. To avoid this 'black-hole' effect, the wrapper was modified so the job sends an e-mail message to the submitter and then sleeps for a few hours. When it wakes up, the process is repeated until the software is found or the job is killed. This keeps the WN busy, preventing further attempts to run additional jobs.

The stage-in (copy of input data) and the stage-out (copy of output data to local SE) phases of the wrapper are also enclosed inside a sleep-and-retry cycle. The LCG commands used to stage files are in fact very sensitive to failures or long delays in the services they rely on.

Lexor-CG

Lexor-CondorG (Lexor-CG) was developed from the original Lexor executor to address the problem of slow job-submission rates to LCG resources. It was noted that the interaction with the LCG Workload Management System (WMS) [89] took up to 45 seconds per job, and this delay prevented to keep the available CPUs filled. A simple calculation, assuming $4,000$ available CPUs and jobs lasting 8 hours, illustrates that a submission rate of at least 1 job per 7 seconds is required. Indeed, during the Rome production (early 2005) there were eight instances of Lexor required in order to approach this submission rate. The resource requirement, in terms of hardware, services (Resource Brokers, RBs), and particularly manpower, was deemed not to scale, and a different submission solution was sought.

CondorG [90] is standard Grid middleware for remote job submission to CEs, and indeed it forms part of the LCG RB. However, it needs an additional server, since the LCG RB is not used in this case and only CondorG is used. In the case

of the RB, it chooses the destination CE, and CondorG submits to the named site. However, when given information about resources, CondorG can also do the resource brokerage. This information is taken from the Berkeley Database Information Index (BDII) [91] and converted into the Condor [92] ClassAd format.

The fundamental difference, compared with the original Lexor executor, is that the resource brokerage and the submission are done by separate components (see figure 4.4). The Negotiator and one or more Schedulers run on different machines, an scalability is achieved by increasing the number of Schedulers only. Furthermore, the scheduler is sufficiently light to run much closer to the UI, perhaps on the same machine.

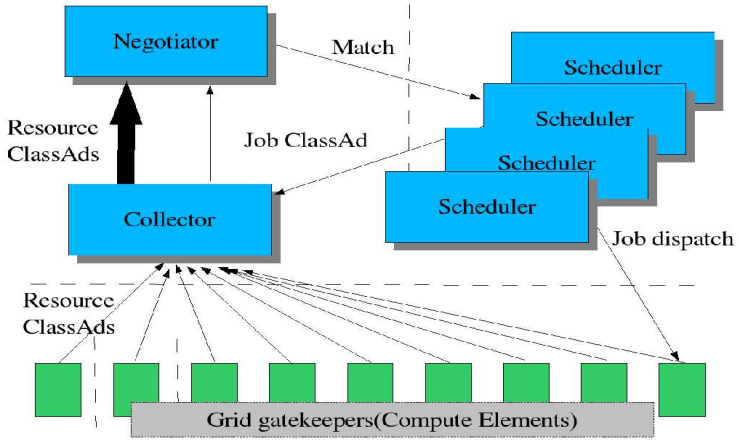

Figure 4.4: *ATLAS CondorG executor architecture.*

The interaction with the local Scheduler is therefore much faster, and job submission takes 0.5 s. Similarly, the status and getOutput requests are instantaneous, as the response is similar to that of a local batch system.

There are, however, two perceived deficiencies with this approach. First, if the UI machine hosts the Scheduler, then it cannot be turned off, which is inconvenient if the UI machine is, for example, a laptop. For Rome production this issue was not relevant because the Windmill services were running on the UI. In general the UI could be a high-availability machine where a user can connect from remote to submit jobs. In terms of fault-tolerance, the CondorG scheduler regains its state after a reboot. Lastly, CondorG does provide a separation of the UI and Scheduler functionality that is included in the recent gLite/EGEE [93] release. This option is not tested at present.

A second concern is the lack of central logging and book-keeping (L&B) [94] when using CondorG. We should stress 'central' because there is in fact a local record of the stages in the job life, and a mechanism exists to extract this to a MySQL database. Despite the LCG central L&B has been identified as a potential cause for poor performance, so not having this architecture is an advantage of CondorG because ATLAS has made no explicit use of the Grid-specific logging and book-

keeping services so far. This does not prevent the L&B information being migrated to a central place, asynchronous to job submission, in the case the information is going to be used in the future.

Lexor was used as the basis for the Lexor-CG executor because its modular design allowed the easy exchange of the LCG submission with the CondorG submission. Everything else, except the job submission mechanism, remained the same and was re-used. During production operation, improvements to Lexor were also applied to Lexor-CG.

Allocating jobs to available resources has been solved in a different way in case of Lexor executor. The number of free CPUs published on the Grid has often been found to be not accurate and then was avoided. Instead jobs were submitted as quickly as possible until a configurable number of jobs were queued at a site and then submission was suspended.

This procedure exhibits one feature of CondorG unavailable in the LCG RB. CondorG executor knows in general how many jobs have been submitted to a site. This knowledge is used to avoid sending jobs to a site due to stale information in the information system. In comparison, the RB model does not allow for such knowledge due to principle considerations. In case of difficulties with the feedback due to instabilities of the information service that can lead to sites being flooded with jobs.

The input and output sandbox mechanisms are not part of the CondorG executor and were implemented within the executor. The GRAM [95] protocol ensures that a single executable reaches the worker node, and that the standard output and error gets back to the submission machine. So the executable was built to be a self-extracting archive containing the input sandbox. Similarly the output files are transmitted in a common archive into the *stdout*. This provides the same functionality as the LCG RB sandbox mechanism, but minimizes the number of file transfers, that require time-consuming authentication and authorization.

Dulcinea

The Advance Resource Connector (ARC) is a Grid middleware developed by the NorduGrid collaboration. NorduGrid's ARC has been deployed at a number of computing resources around the world. These resources are running various Linux distributions and use several different Local Resource-Management Systems (LRMS). Despite using different LRMS and specific information providers, the different sites can present the information about available resources in a uniform way in ARC's information system. This information system is based on the Globus [96] Monitoring and Discovery Service (MDS). The information provided is used by the brokering algorithm in the ARC user interface to find suitable resources for the tasks to be performed, as well as by various monitoring tools like the NorduGrid Grid Monitor. NorduGrid's ARC is fully integrated with the Globus Replica Location Service (RLS) [67]. This allows jobs sent to an ARC-enabled site to specify input-file locations as an RLS catalogue entry instead of a physical-file location.

A job can also, if desired, automatically register created output files in an RLS catalogue. The input and output file handling is done by the gatekeeper on a site. This offloads the task of transferring input and output files from the user jobs.

The ATLAS Production System executor for NorduGrid's ARC, Dulcinea, is implemented as a C++ shared library. This shared library is imported into the Production System Python framework. The executor calls the ARC user-interface API and the Globus RLS API to perform the tasks. The job description received from the supervisor in the form of an XML message is translated by the Dulcinea executor into an Extended Resource Specification Language (XRSL) [97] job description. This job description is then sent to one of the ARC-enabled sites, selecting a suitable place using the resource-brokering capabilities of the ARC user-interface API. In the brokering, among other things, the availability of free CPUs and the amount of data needed to be stage-in on each site to perform a specific task is taken into account.

The look-up of input-data in the RLS catalogue and the stage-in of these files to the site is done automatically by the ARC Grid Manager. The same is true for stage-out of output data to a Storage Element (SE) and the registration of these files in the RLS catalogue. The Dulcinea executor has just to add the additional RLS attributes needed for the DonQuijote data management system to the existing file registrations.

The Dulcinea executor also takes advantage of the further capabilities of the ARC middleware. The executor does not have to keep any local information about the handled jobs, but relies on job information provided by the Grid information system.

Because of the diversity of the ARC-enabled resources used for Linux installation on the various sites, the ATLAS software release is recompiled for various distributions. The ATLAS software release is then installed on the sites desiring to participate in ATLAS productions. After testing the installation on each site, the ATLAS run-time environment is published in the site's Grid information system. By letting the Dulcinea executor request the run-time environment in its job description, only sites with the proper software installation are considered in the brokering.

Capone

The Grid3 executor system, Capone, communicates with the supervisor and handles all the interactions with Grid3 resources and services. Job requests from the supervisor are taken by Capone, which interfaces to a number of middleware clients on the Virtual Data Toolkit (VDT) [98]. These are used with information in Grid servers to submit jobs to the Grid.

ATLAS job parameters are delivered to Capone in an XML-formatted message from the supervisor. These messages include job-specific information such as specific parameters expected by the transformation (the executable), identification of

input and output files, and resource-specific requirements for the job. Capone receives these messages as either Jabber or Web service requests and translates them into its own internal format before starting to process the job. The transformation to be executed is determined from the input data contained in the supervisor message. The Globus RLS is consulted to verify the existence of any required input files and serves meta-data information associated with these files.

The submission and subsequent monitoring of the jobs on the Grid is performed using CondorG. On the CE, if necessary, all input files are first staged-in before the execution proceeds. The ATLAS software itself, an Athena executable, is invoked from a sandbox using a VDT-supplied executable. A wrapper script, specific to the transformation to be executed, is called first to ensure that the environment is set up correctly before starting any ATLAS-specific executables. The wrapper script also checks for errors during execution and reports results back to the submitter. In addition, the wrapper script performs additional functions such as evaluating an MD5 checksum on all output files. The Condor status of each job is checked periodically by Capone and, when the remote job is completed, Capone resumes control of the job.

Capone checks the results of the remote execution: the program exit code and the presence of all expected output files. This is a delicate step since a large variety of errors may be discovered here, ranging from Input/Output problems encountered during the stage-in process, to errors in the execution of Athena, to site characteristics that prevent Condor to exiting correctly. The next area in the CE is the output SE. The transfer also involves evaluation of the MD5 checksum and file size of the destination files and checking the integrity of the transfers. Furthermore, some important meta-data, like the GUID (a Globally Unique file IDentifier) [99], is recovered from service files in the remote-execution directory. The process ends with the registration of the output files to RLS inserting logical filenames, physical filenames and meta-data attributes required by DonQuijote, the ATLAS data-movement utility that allows data transfers between Grids.

There is a new system called PanDA (Production and Distributed Analysis) which has been developed by US-ATLAS and is used for CSC/DC3 production as a regional (OSG) 'executor', interacting with an ATLAS production 'supervisor'. A more detailed description about PanDA is found in section 5.2.1.

4.2 ATLAS Data Challenges

The ATLAS Collaboration is preparing for data taking and analysis at LHC, scheduled to start operation during 2008. According to the Computing Model, after reduction of data by the on-line trigger, the expected volume of data recorded for offline reconstruction and analysis could be of the order of few Peta-bytes (10^{15} bytes) per year, to be analyst's by institutes distributed worldwide. The computing resources are estimated to be two orders of magnitude larger than for previous experiments in particle physics. So the experiment needs a worldwide distributed data management and computing system, to handle the huge amount of data and to ensure that resources are available.

The LC Computing Review in 2001 recommended that the LHC experiments should carry out Data Challenges (DCs) of increasing size and complexity to prepare for the LHC start-up. The goals of the ATLAS Data Challenges are the validation of the ATLAS Computing Model [100], the complete software chain, the data model, and the correctness of the technical computing choices.

The collaboration decided to perform these DCs in the context of the LCG, to which ATLAS is committed, but also to use both the middleware and the resources of two other Grid projects, OSG/Grid3 and ARC/NorduGrid. The task of the LCG project is to prepare the computing infrastructure for the simulation, processing and analysis of the LHC data for all LHC collaborations. The main emphasis of the LCG project is the deployment of Grid technologies for LHC computing. Both OSG/Grid3 and ARC/NorduGrid have similar approaches using the same foundations (Globus) as LCG, but with somewhat different middleware.

A series of Data Challenges was started in 2002:

- ATLAS Data Challenge 1 (DC1) [21], in 2002-2003, 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. Some early testing of the Grid infrastructures was done on LCG and Grid3.
- Data Challenge 2 (DC2) [64], in the second half of 2004, and the Rome production [22], in the first half of 2005, were completely run on Grid. After DC2, a new Monte Carlo production was necessary to produce the required event samples for an ATLAS Physics Workshop (called Rome Workshop) scheduled in June 2005.
- Data Challenge 3 (DC3), also called Computing System Commissioning (CSC), started at the end of 2005 and currently ongoing, is fully running on the Grid.

4.2.1 Data Challenges scope and scale

The ATLAS experiment requires a large amount of simulated data in order to optimize the design of the sub-detectors (detector performance), to evaluate the

physics performance that is being achieved, and to test the software and computing infrastructure. Samples of simulated data consist of a large number ($\sim 10^7$) of simulated events. Each event is the result of a collision between two protons. The process of the full ATLAS simulation chain (see section 1.3) using the Athena framework requires the following steps:

- Event Generation: the final state configurations of proton-proton collisions are generated using programs relying on theoretical calculations, phenomenological models and experimental inputs (see section 1.3.1).
- Detector Simulation: the interaction of the generated particles inside the detector is simulated taking into account the geometry and distribution of the detector material (see section 1.3.2).
- Digitization: the detector response is derived from the particle interactions and it is converted into digital form, to resemble the real output of the detector . In addition, because of the high rate of collisions at LHC, digitized signals from several simulated events can be piled-up to create samples with a realistic experimental background. The digitized events (with or without pile-up) can be used to test the software chain to be applied to real LHC data (see section 1.3.2).
- Reconstruction: particle trajectories and energies from the detector signals are reconstructed (see section 1.3.3).

The ATLAS full simulation requires a chain of different programs with different characteristics in terms of memory usage and CPU time consumption. Typically, a simulation (long) job runs for 24 hours, whereas a digitization or reconstruction (short) job runs typically for 3 to 4 hours.

4.2.2 Data Challenge 1 (DC1) experience

All Data Challenges included physics content in order to engage the physicist community with the exercise and leading to a meaningful validation of the software. For DC1 [21], in 2002-2003, a major goal was to provide simulated data required for the preparation of the ATLAS High Level Trigger TDR [101].

The DC1 production was split into 3 different phases. The first phase (event generation and simulation) was achieved during the summer of 2002, involving 40 institutes from 19 countries (see figure 4.5). The second phase (October 2002-March 2003), where 'pile-up' was included, was performed with the participation of 56 institutes from 21 countries. In the third phase, distributed reconstruction of the most demanding high-statistics samples was carried out using the 9 largest sites (April-June 2003). The output data of this phase was stored in 8 sites (Alberta, BNL, CCIN2P3 Lyon, CERN, CNAF Bologna, FZK Karlsruhe, Oslo, RAL).

Phase 1 (Generation and Simulation)
Due to the huge amount of computing time needed, it was essential to use all computing resources available in ATLAS institutes around the world. At peak time,

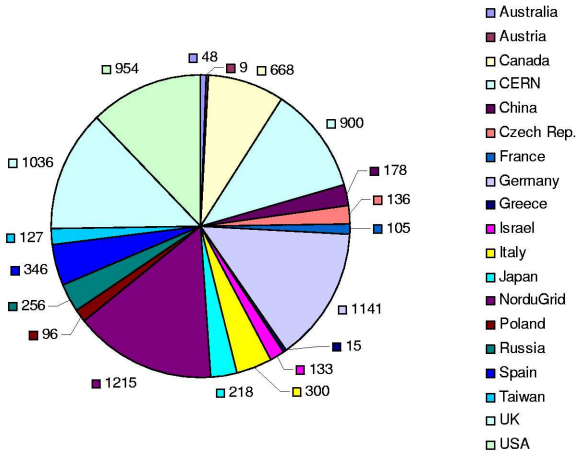

Figure 4.5: *Number of normalized processors per country accessible to DC1 (Countries are displayed clockwise, starting with Australia (48)). The numbers of processors per site varied between 9 and 900 resulting in a total of* $\sim 5,000$ *CPUs at peak time corresponding to a processing power of* $\sim$ *2.2 MSI2k.*

$\sim$3200 processors located in 40 institutes from 19 countries were used. This corresponds to $\sim$1400 kSI2k, i.e. $\sim$50% of the CPU estimated for just one Regional Facility at LHC start-up (2007).

During phase 1:

- About 50 million events were generated using PYTHIA.
- About 51 million events were passed through detailed detector simulation using ATLSIM (ATLAS simulation using GEANT 3): about 40 million were single-particle events (muons, photons, electrons, pions) and the remaining $\sim$ 11 million were complete events.

The total data volume produced during phase 1 was about 24 Tera-bytes and about 8 Tera-bytes for generated events; the total CPU time necessary to generate all the events was about 200 kSI2k-days, and the CPU time required to simulate all events, about 1.4 MSI2k-days.

Phase 2 (Pile-up production)
In DC1 phase 2 pile-up was added to a sub-set of data samples. New countries, China, Greece and new institutes from Canada, Italy, NorduGrid, UK and USA joined the effort in the course of phase 2, so 56 institutes from 21 countries participated in DC1 phase 2, adding to a total of 1.5 MSI2k-s. About 3.9 million events were produced at low luminosity and 2.650 million events at high luminosity. This part of DC1 consumed about 3.2 MSI2k-days and produced a data volume of about 34 TB in 32,000 partitions.

Phase 3 (Reconstruction)
To facilitate the access to the large distributed datasets, and since not all production sites were accessible via Grid tools, the data were replicated to 8 different sites. Therefore, the processing of the data was mostly done in countries where these sites are located. About 6.4 million events were processed during the reconstruction phase. This part of DC1 consumed about 4.4 MSI2k-days and produced a data volume of about 200 GB in 25,000 partitions.

4.2.3 Data Challenge 2 (DC2) experience

ATLAS has been the first experiment at LHC where a whole Data Challenge exercise has been completely performed using the Grid infrastructure. ATLAS ProdSys was used in DC2 [64] to manage the production of simulation jobs into the three different Grid flavours: LCG/EGEE, OSG/Grid3 and ARC/NorduGrid.

The whole ATLAS DC2 exercise took from the beginning of July 2004 to the end of December 2004, where in total 10 million events were processed in $\sim$ 260 thousand jobs, consuming $\sim$ 200 kSI2k years of CPU and producing $\sim$ 60 TB of data. During this exercise, 69 sites spread over 20 countries participated.

The phase 1 of the ATLAS DC2 started in July 2004 and was divided into several parts listed as follows: Event generation, detector simulation, pile-up and digitization. The output, called 'byte stream data', was similar to detector 'Raw Data'; Data transfer to CERN ($\sim$ 35 TB in 4 weeks);
Event mixing, where events from different physics channels were mixed in ad-hoc proportions. For pile-up, the event size was 3.3 MB and the computer power required to process these events, 144 kSI2k-s. For event mixing, the corresponding values were 3 MB and 5400 SI2K-s.

The ATLAS DC2 finished the simulation task at the end of September 2004; 10 million events were generated and simulated using the three Grid flavours. The contribution from each flavour is displayed in figure 4.6.

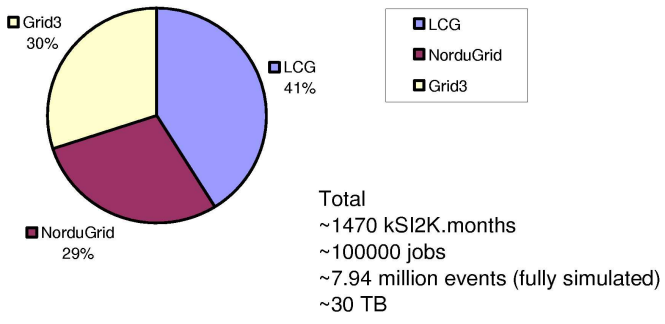

Figure 4.6: *The chart plots the contribution of each Grid flavour in the generation and simulation parts for the ATLAS Data Challenge 2.*

Between the start of DC2 in July 2004 and the end of September 2004, the automatic Production System submitted about 235,000 jobs, reaching approximately 2,500-3,500 jobs per day, distributed over the three Grid flavours. Overall, they consumed $\sim$ 1.5 million SI2k-months of CPU ($\sim$5,000 CPU months on that average present day CPU) and produced more than 30 TB of physics data. About 6 TB of these data were moved using Don Quijote servers.

Additionally, this exercise showed in LCG that all the involved elements (middleware, Production System, deployment and monitoring tools) needed improvement:

- The Information System in a Grid infrastructure must be very stable and robust, since several services and client tools rely on it. For example, in LCG, it is crucial for the Workload Management System to discover resources and for the ATLAS DDM service obtain storage status information. Unfortunately, this was not the case at the time of DC2.
- Site configuration problems were the main source of failures during DC2. In addition, a problematic site appears 'free', so it continuously attracts jobs hat fail after a few seconds of execution time, creating a 'sink-hole' effect.
- The WMS is a key component for production in a distributed environment and some problems were found during production (see section 4.2.4).

The generation and simulation contributions to the ATLAS DC2 from each Grid are explained in more detailed in the following subsections: DC2 production using Grid3, NorduGrid and LCG.

DC2 production using Grid3

The Grid3 collaboration deployed an international Data Grid. The facility was operated jointly by the U.S. Grid projects and the U.S. participants in LHC experiments.

The deployed infrastructure (see figure 4.7) has been in operation since November 2003 involving 27 sites, a peak of 2,800 processors, work loads from 10 different applications exceeding 1,300 simultaneous jobs, and data transfers among sites larger than 2 TB/day.

Figure 4.8 shows the jobs contribution to the ATLAS DC2 in the simulation part. A specific Production System has submitted jobs to Grid3 sites at full scale using the shared facilities. Around 30,000 jobs finished successfully, 2.4 million events and 8 TB were produced, and more than 0.5 million CPU-hours were consumed.

During three months, July-September 2004, about one third of ATLAS DC2 production was executed on Grid3, keeping pace with peer projects using NorduGrid and the LCG. Efficiencies steadily improved as new experience was incorporated into the system. It was found that two areas of Grid infrastructure are critical to distributed frameworks: job-state management, control and persistency, and trouble-shooting failures in end-to-end integrated applications. The next steps in the project should make progress in those areas.

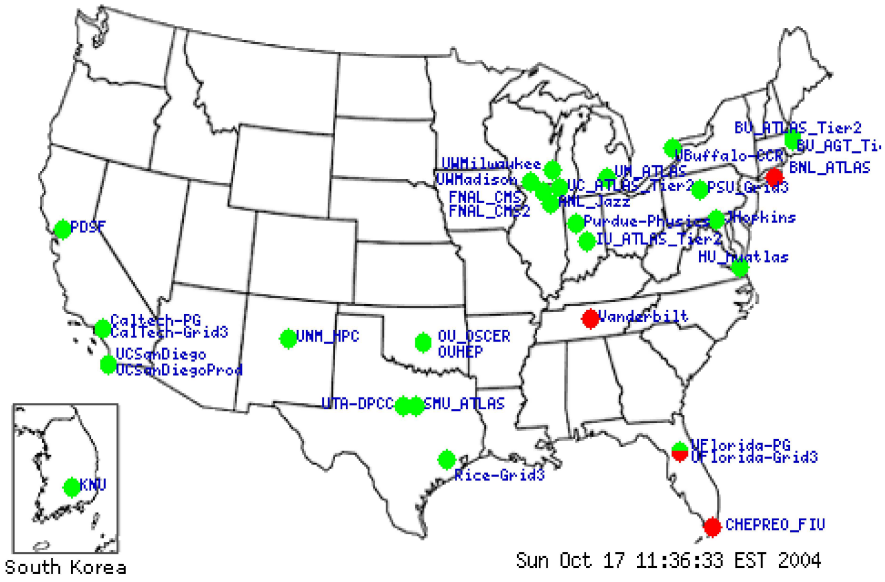

Figure 4.7: *Geographical distribution of Grid3.*

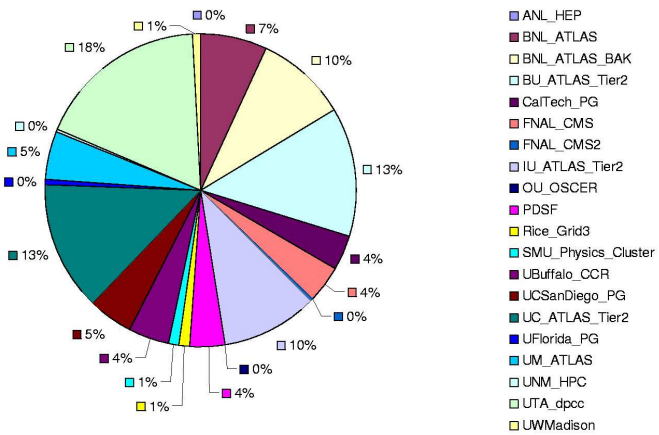

Figure 4.8: *Job contribution by site in Grid3 for the simulation part in DC2.*

DC2 production using NorduGrid

The NorduGrid project is established mainly across Nordic countries but includes sites from other countries as well. It was the first project to reach production quality level and contributed to a significant part of the DC1 production. It provided resources for DC2 and support for production on non-RedHat 7.3 platforms (e.g. Mandrake 8.0, Debian 3.0).

The NorduGrid resources include (see figure 4.9) the original small test-clusters at the different physics-institutions, but also some of the biggest supercomputer clusters in Scandinavia. It was one of the largest operational Grids in the world with approximately 4,000 CPU's, storage capacity of 14 TB, involving 40 sites from 11 countries.

Figure 4.10 shows the jobs contribution to the ATLAS DC2 in the simulation part. Around 30,000 jobs were processed and 2.4 million events produced.

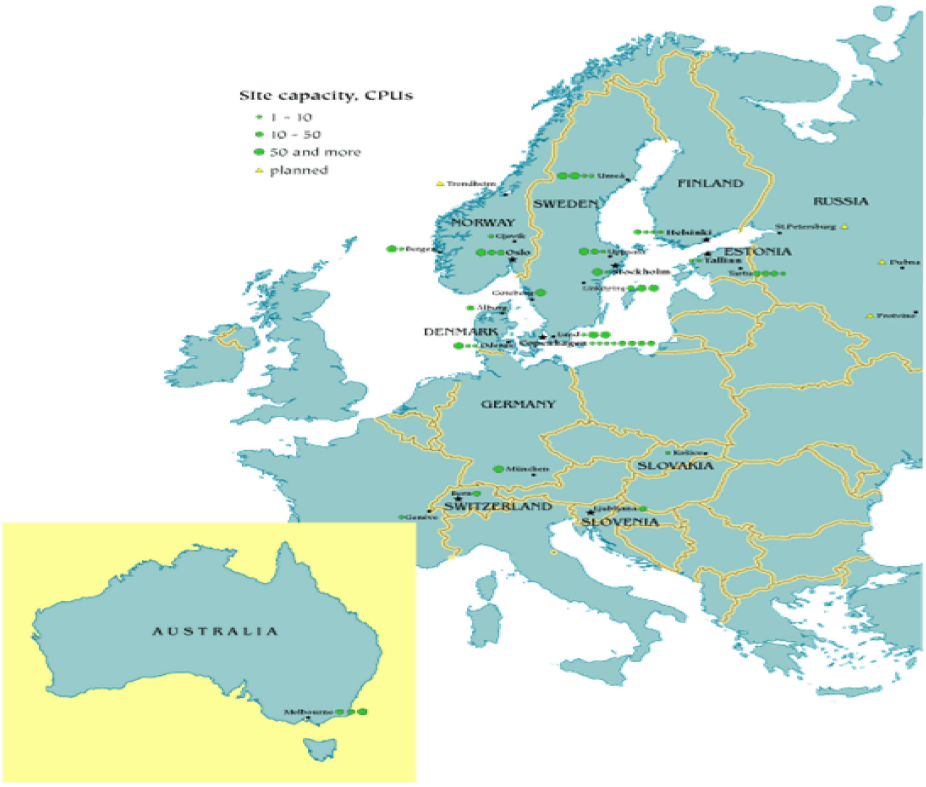

Figure 4.9: *Geographical distribution of NorduGrid.*

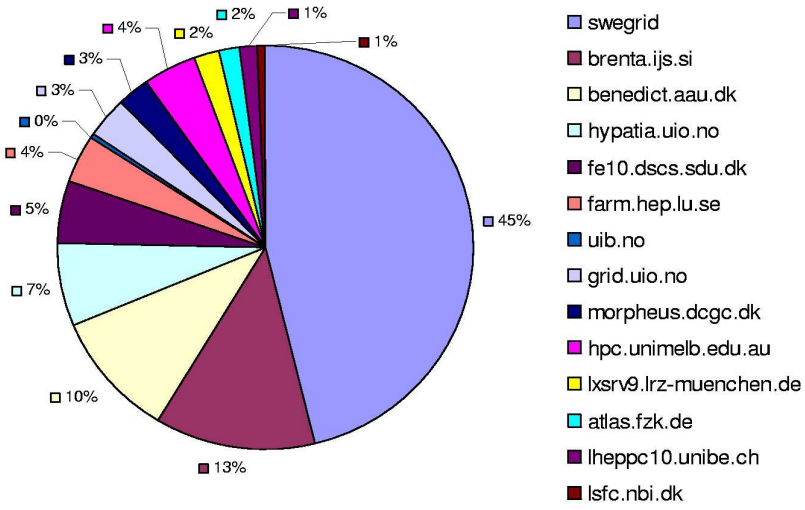

Figure 4.10: *Job contribution by site in NorduGrid for the simulation part in DC2.*

The ARC middleware and the Dulcinea executor provided a stable service for ATLAS DC2 in 2004. Twenty-two sites in 7 countries operated as a single resource and contributed approximately 30% of the total ATLAS DC2 production. Most of the resources were not 100% dedicated to DC2 production. The number of middleware-related problems was negligible, except for the initial instability of the RLS server. Most job failures were due to site-specific problems.

DC2 production using LCG

The LCG project is built upon the most stable developments of the European Data Grid middleware (EDG) [75], the US Virtual Data Toolkit (VDT) project (Globus) and European Data Trans Atlantic Grid (DataTag) [102] monitoring tools.

The deployed infrastructure (see figure 4.11) has been operating since 2003 with 82 sites from 22 countries, providing a peak of 7,269 processors and a total storage capacity of 6,558 TB.

Figure 4.12 shows the LCG jobs contribution in the generation and simulation part. LCG sites processed the production during this period using shared facilities. Around 40,000 jobs finished successfully and 3.1 million events were produced.

The heavy work-load, imposed on these services by the high number of concurrent requests from the running jobs, sometimes caused them to become unresponsive for a while or even to crash. This affected all the running jobs that, after successfully executing the transformation, failed to stage the produced files to the output Storage Element. A retry mechanism did not completely solve the problem, but

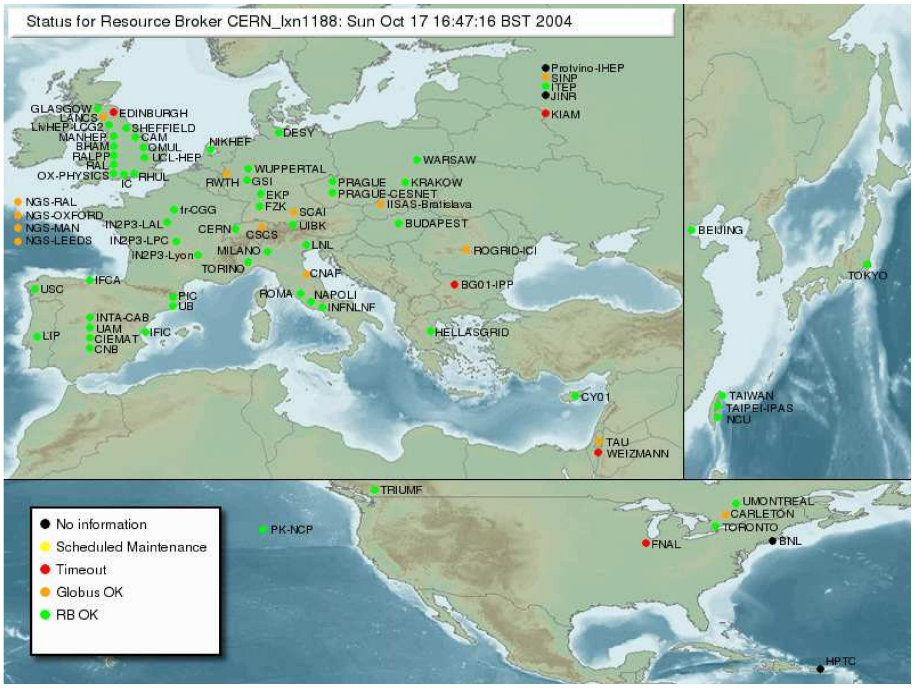

Figure 4.11: *Geographical distribution of LCG.*

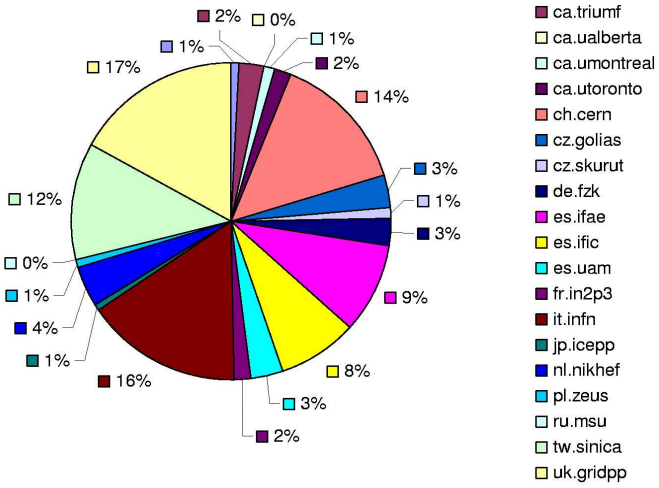

Figure 4.12: *Job contribution by site in LCG for the simulation part in DC2.*

allows the recovery of many jobs with LCG commands hanging. In fact, killing by hand the command process triggers a subsequent attempt, which succeeded.

During DC2 production in 2004, Lexor managed more than 100,000 jobs, including generation, simulation, digitization and pile-up. More than 30 LCG sites were involved, providing all about 3,000 CPUs. The job management was shared, on average, by 6 Lexor instances owned by different managers, each instance taking care of a maximum of about 800 jobs at a time. This threshold was chosen in order to limit the size of queries issued by supervisors to the central production database and to circumvent limitations of the Resource Broker.

4.2.4 Rome production experience

ATLAS has completed two Data Challenges, followed by a large production intended to provide data for physics studies, to be discussed during the Rome Workshop in June 2005. The Rome production [22] activity, due to methodology, large number of events and goals, can be considered totally equivalent to a real Data Challenge: beside the mere event production, it offered a unique opportunity to test improvements in the production framework, Grid middleware and reconstruction software, and spot remaining difficulties. Like DC2, the Rome production was fully carried out in a Grid environment. ATLAS provided resources managed by the three different Grid infrastructures.

Some improvements with respect to the DC2 exercise were applied to the Rome production exercise, since the Production System, the LCG operation and support teams gathered a lot of experience during DC2 and benefit of this experience at the time of Rome production:

- The Information System was improved in several aspects. This reduced the single point of failure effect for both job submission and data management during the job execution.
- The LCG operation team developed a series of automatic tools for site sanity controls, called Site Functional Tests (SFT) [103], running every day at every site testing the correct configuration of the nodes and the interaction with Grid services. This helped in reducing the 'sink-hole' effect. During the Rome production the ATLAS production managers were excluding problematic resources, while for large part of DC2 every LCG site was used. This was done mostly on the basis of the LCG Site Functional Tests.
- The WMS became more robust before the Rome production, thanks to several bug fixes and optimizations in the WMS work-flow and its interaction with external services.

The number of successful jobs processed on LCG, Grid3 and NorduGrid system in DC2 (before January 2005) and Rome production (after January 2005) is shown in figure 4.13.

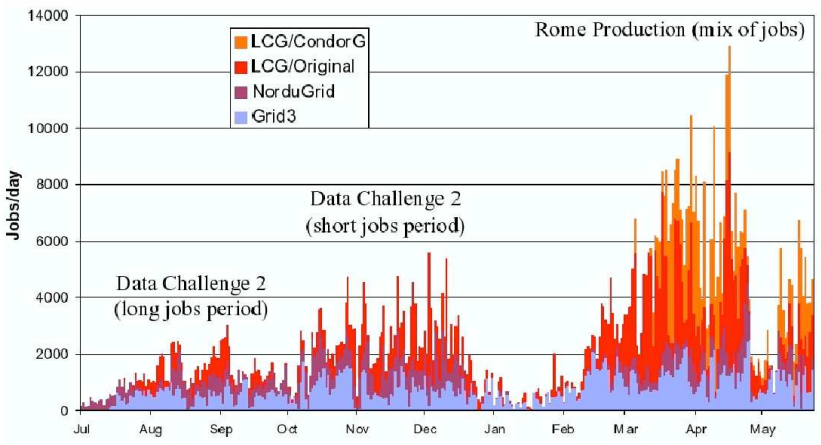

Figure 4.13: *Number of jobs run every day on the native LCG, Grid and NorduGrid infrastructure during DC2 (from July 2004 to January 2005) and Rome production (from January 2005 to June 2005).*

The total number of processed jobs and the total CPU usage by each Grid flavour during the Rome production are shown in figures 4.14 and 4.15.

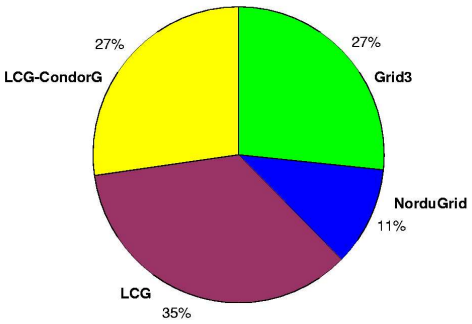

Figure 4.14: *Job contribution by Grid flavour for the Rome production.*

The production for the Rome workshop using the LCG flavour amounted to a total of 380 thousand jobs, out of them 109 thousand being simulation jobs, 106 thousand digitization jobs, 125 thousand reconstruction jobs and 40 thousand pile-up jobs (the reconstruction was performed on digitized events and partially on piled-up events).

In total, 1.4 million files were stored in LCG, adding to 45 TB of data. These figures show a clear improvement with respect to DC2, where 91.5 thousand jobs were processed on LCG and no reconstruction was performed.

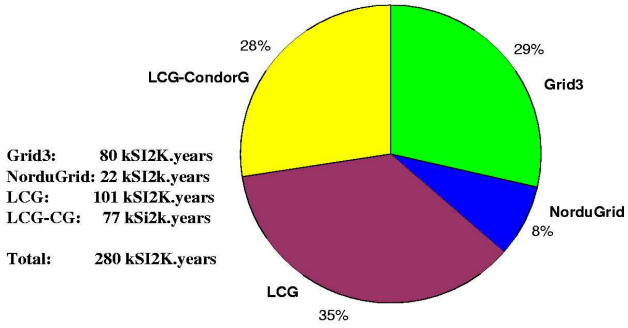

Figure 4.15: *CPU usage by Grid flavour for the Rome production.*

Rome production on LCG proved to be a big step forward compared to DC2 simulation. However, several components of Grid middleware and Production System required improvement, regarding in particular reliability and performance. Table 4.1 summarizes the main causes of jobs failures in the Rome production, together with the failure rates. As shown, the most problematic issue comes from DDM ($\sim$ 26 %).

System	Causes	Rate (%)
WMS	Total	1.6
DDM	Download input files	26.4
DDM	Upload output files	0.8
ATLAS/LCG	Athena-crash	9.1
ATLAS	Proxy expired	0.3
LCG	Site misconfiguration	0.9
Unclassified		9.0

Table 4.1: *Causes of jobs failures during Rome production. The percentage represents the number of failed jobs for the given reason, in respect of the number of submitted jobs.*

At the time of Rome production, a reliable file transfer service was not in place. Replication of files between Storage Elements and data movement from/to the Worker Node (WN) was performed through LCG data management client tools (lcg-utils) [104]. Such tools did not provide timeout and retry capabilities and, as a consequence, hanging commands were killed only by the batch system, when the job reached the maximum time limit of the queue.

A file is a 'Grid file' if it is both stored in a SE and registered in the file catalogue. Data management high level tools (like lcg-utils) should ensure consistency between entries in the catalogue and files in the SE. The LCG data management service does not ensure consistency between files in the SE and the corresponding entries in the file catalogue. At the end of Rome production, a tedious process of catalogue and SEs manual clean-up was needed. In the future, the new version of

lcg-utils should guarantee automatic transactions and partially solve the problem.

One of the main problems in data management in Grid infrastructures is data access on Massive Mass Storage (MMS) systems (hybrid systems with a tape library and a disk front-end). Data need to be moved (staged) from tape to disk before they can be accessed. The middleware should be able to ensure the presence of data on disk when an application or a service needs to access them. Moreover, it should be possible to 'pin' data, i.e. to ensure they remain staged on disk until they are needed or in use. Unfortunately, this is not the case in LCG and a good strategy for data access on MMS systems (MMS) was really missing during the Rome production. A large fraction of the stage-in problems mentioned in table 4.1 arose during data access on mass storages. The situation will hopefully improve with the introduction of a Storage Resource Manager (SRM) [105] as a front-end to every SE. The SRM is an interface, agreed on between experiment and middleware developers, to standardize access to storage and management.

It must not be forgotten that the aim of the production and Data Challenges is not only the validation of the Computing Model, the software and middleware infrastructure, but also the generation of event samples which can be analyzed by physicists. In this aspect, the ATLAS strategy for file distribution must be rethought. At the end of the Rome production, the output files were chaotically spread around 143 different SEs. Such disorganized distribution of files complicated the analysis of the reconstructed event samples and, since in the ATLAS production chain the output of one phase represents the input of the subsequent, it affected the production itself.

Failures in the LCG WMS RB were limited, however there are some standing issues. The performance of the WMS, generally acceptable in normal conditions, was observed to degrade under stress. A single Resource Broker could not serve efficiently more than $2,000$ jobs and the submission time of an ATLAS job could vary from 10 to more than 40 seconds depending on the load. As a temporary solution, several RBs dedicated to ATLAS and with different hardware solutions have been deployed at CERN, CNAF (Bologna, Italy) and IFIC (Valencia, Spain). To cope with the lack of performance in job submission, ATLAS has developed in parallel an executor, Lexor-CG, using the condor Condor-G service.
The gLite WMS RB is supposed to overcome several of those limitations. It will be possible in fact to submit several jobs in bulk mode, i.e. in a single interaction.

The Lexor-CG executor submitted part of the jobs using the Rome production in early 2005. The production rate on LCG was immediately doubled, making full usage of LCG resources. This rate was maintained over the course of Rome production and two Lexor-CG instances were able to match the production rate of eight plain Lexor instances. This could be directly attributed to the higher submission rate.

During the reconstruction phase it was noted that more executors were an advantage. The shorter jobs, with four output files requiring renaming and validation, placed Windmill as the bottleneck. Removing the job-submission-rate bottleneck was an important step towards identifying this problem, to be addressed in future

production. In this way, Eowyn replaced Windmill for future production. In this way, Eowyn replaced Windmill for future productions.

The WMS RB relies on a user defined 'ranking' expression to select the best available resource matching the job requirements. During the Rome production was very difficult to find a good ranking expression for job distribution. At that time, there was not provided an adequate VO specific view of computing resources. The situation improved dramatically when ATLAS dedicated queues were configured at almost every site.

The information at the 'site level' for several attributes like the frequency of CPU cycles, the memory, etc. refer to a representative WN which is not always the worst performing of the farm. For this reason, the requirements specified at job submission time might be not fulfilled in reality. In the ATLAS production for instance, almost 10 % of the jobs failed with a crash of the experiment application (Athena).

The job monitoring and error diagnostic is essential in order to efficiently run a production in a distributed environment. Job failures have to be categorized and prioritized, so that people can invest time solving the most serious and urgent issues. The production manager need a complete overview of the available resources and the distribution of running and waiting jobs to be able to identify problematic behaviors of some services. However the monitoring system still needs to be improved.

Capone was the Grid3 executor system for Rome production and DC2. Then, PanDA was introduced instead (see section 5.2.1).

4.2.5 IFIC contribution to Data Challenge 3 (DC3) exercise

In preparation for ATLAS physics data in 2007, Data Challenge 3 (DC3) (also called 'Computing System Commissioning', CSC) in 2006-2007 is testing all components of the Computing Model that have not so far been integrated with mainstream software developments, such as the simulation of the trigger chain and the calibration and alignment procedures.

Detailed planning for CSC is under way, with the following major tests planned:

- Tier-0 processing.
- Data distribution.
- Distributed production and reconstruction of simulated data.
- Distributed Analysis.
- Reprocessing.
- Calibration Data Challenge.
- Data streaming.

Finally, after testing all these components, a component integration towards the full system test (called 'Final Dress Rehearsal') is scheduled for Summer or early Autumn 2007.

After the Rome production experience, several components have been improved in the ATLAS Production System:

- Eowyn replaced Windmill as supervisor component in ProdSys.
- DQ2 is fully integrated for production in DDM.
- All components in executors have been improved.

Statistics for number of successful jobs using Athena 11.0.X during the CSC exercise is shown in figure 4.16.

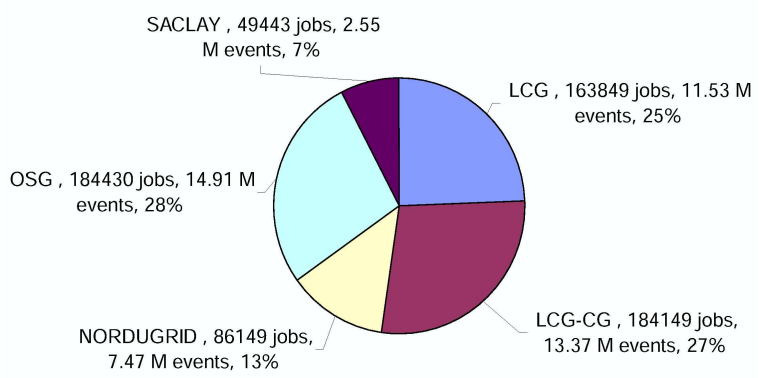

Figure 4.16: *Number of successful production jobs using Athena release 11.0.X during CSC exercise.*

Currently, the CSC/DC3 exercise is an ongoing production, so it has not finished yet. The problems that are found in LCG/EGEE during production are discussed in the ATLAS LCG/EGEE weekly production meetings and in the Computing Operations weekly meetings. Then, these problems are fixed during the continuous production.

On the other hand, IFIC has been involved in ATLAS DCs since they started in 2002. IFIC was the only Spanish institute that contributed to the DC1 effort using a infrastructure called GoG (Grupo de Ordenadores para el Grid) [106]. As commented before, these resources were not part of a Grid environment (see section 4.2.3).

For DC2, IFIC was involved into the ATLAS MC production, running a Windmill/Lexor ProdSys instance. This instance picked up jobs defined in the ProdDB by the Supervisor (Windmill), then the Supervisor parsed the *xml* files to the Executor (Lexor) that submitted jobs using the Workload Management System RB.

This instance was operational at IFIC from July 2004 to September 2004 and was installed on one dedicated UI, just for production. Some other Grid services were also installed for production:
One RB, one BDII and one Jabber.

After DC2, a new simulation production was started called, the 'Rome production'. IFIC participated with the following Grid services for this production:
One RB, one BDII and one Proxy Server, in addition to the usual GoG resources.

Finally, IFIC has been involved in the ATLAS CSC/DC3 simulation production for LCG/EGEE running an Eowyn/Lexor ProdSys instance. This instance, installed at one dedicated UI just for production, was running at IFIC from 25th January 2006 to 9th August 2006, and managed 62,910 successful jobs. As for the previous DC2, other Grid services were deployed only for production at IFIC:
One RB, one BDII and one Jabber.

My main role in the ATLAS DCs has been focused on this last exercise, the CSC/DC3 MC simulation production. I have been one of the main responsibles to install and maintain the Eowyn/Lexor ProdSys instance at IFIC.
On the other hand, I have been one of the ATLAS MC production managers in the CSC/DC3 exercise also, while the instance was installed and running at IFIC.

The statistics of this instance are shown in figures 4.17 and 4.18, where the total number of successful jobs for LCG/EGEE is shown as well.

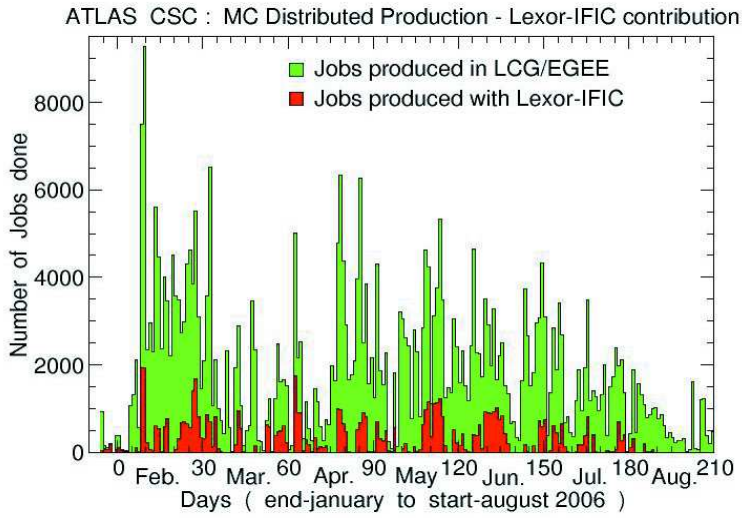

Figure 4.17: *CSC production jobs managed by the Lexor instance at IFIC.*

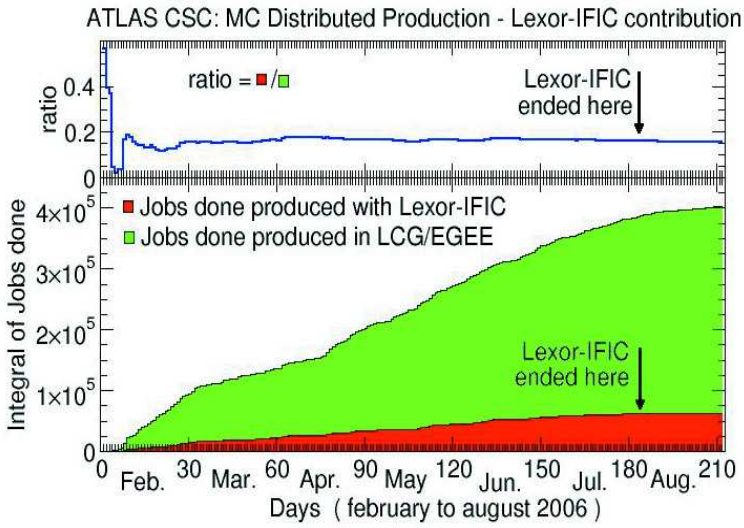

Figure 4.18: *IFIC contribution to the total CSC production jobs for LCG/EGEE Grid flavour.*

A more detailed study of this contribution is displayed in table 4.2.

month	ATLAS LCG/EGEE	Lexor IFIC	contrib. (%)
January	2136	548	25.7
February	90358	13800	15.3
March	45846	8529	18.6
April	64875	10984	16.9
May	71179	11392	16.0
June	66059	11324	17.1
July	46181	6052	13.1
August	7080	281	4.0

Table 4.2: *Total number of successful jobs per month managed by the Lexor instance at IFIC and for the whole LCG/EGEE.*

where the total number of successful jobs for LCG/EGEE was $393,714$ and approximately 16% of them were managed by the instance running at IFIC.

On the other hand, the statistics of finished successful jobs run at IFIC's CE (lcgce.ific.uv.es) during this time are shown in table 4.3.

where the total number of finished successful jobs processed at IFIC were $5,855$, i.e. 1.5% of the total LCG/EGEE.

month	ATLAS LCG/EGEE	IFIC's CE	contrib. (%)
January	2136	6	0.28
February	90358	2802	3.10
March	45846	784	1.70
April	64875	681	1.00
May	71179	671	0.90
June	66059	343	0.52
July	46181	432	0.94
August	7080	136	1.90

Table 4.3: *Total number of finished successful jobs per month processed at IFIC's CE and for the whole LCG/EGEE.*

In August 2006, the ATLAS ProdSys replaced the use of the LCG RB by the gLite EGEE RB, which offers a better job management performance. At the beginning, the ATLAS ProdSys found some problems using the gLite EGEE RB and some patches were required to apply on it. For these technical reasons, IFIC stopped running its ProdSys instance.

In September 2006, once these problems with the gLite EGEE RB were solved, there was a modification in the strategy of the ATLAS CSC/DC3 MC production for LCG/EGEE. Instead of running several instances, the MC production management was focused only on four instances: two instances running the CondorG executor and other two instances running the Lexor executor. In this case, the two Lexor instances were selected to run at CERN and Milan.
Since then, my contribution and from the rest of people involved in the ATLAS LCG/EGEE MC production group have joined to this effort by doing weekly shifts of the MC production, being these instances centralized only on a few sites.

Chapter 5

ATLAS Distributed Analysis and Tier-3 prototype

The ATLAS collaboration is preparing for the LHC start-up, schedule to start in 2008. The expected data volume recorded is of the order of few Peta-bytes per year, to be processed by institutions distributed worldwide. The LHC experiments need a worldwide Distributed Data Management (DDM) and computing system to handle this huge amount of data.

As part of the preparation for data taking and physics analysis, the ATLAS Distributed Analysis (ADA) project has as a goal to enable physicists to analyze the data produced by the detector and reconstructed by the Production System (ProdSys), and to carry out production of moderately large samples.

A description of the ATLAS Analysis Computing Model and Distributed Analysis strategy is in section 5.1. The current prototypes for ATLAS Distributed Analysis, including an example of application using ADA with the Production System, are discussed in section 5.2. Finally, a brief description of a prototype for a local user analysis facility, called Tier-3, at IFIC is given in section 5.3.

5.1 ATLAS Distributed Analysis

The primary goal of the ATLAS Distributed Analysis (ADA) project is to enable to ATLAS physicists to perform their analysis using resources shared on the Grid. This is achieved by providing easy access to the vast computing resources of the various Grids, in a way that hides most of the complexities of Grid environments.

The Distributed Analysis model is described in the ATLAS Computing Model [100] and assumes that data will be distributed by the ATLAS Distributed Data Management (DDM) system to various computing facilities. User jobs are in turn assigned to sites in function of the availability of data. A typical analysis job consists of a script that configures and executes a user defined algorithm in Athena (the ATLAS software framework). The script specifies the input data

from a collection of potentially interesting events and produces one or more files containing plots and histograms.

As in many other large collaborations, different ATLAS physics groups have different work models and the Distributed Analysis system needs to be flexible enough to support all current and newly emerging work models and still remain robust enough.

With data distributed at computing facilities located around the world, an efficient system to manage access to these data is critically important to achieve an effective data analysis.

Users performing data analysis need a random access mechanism to allow rapid pre-filtering of data based on certain selection criteria. These data have to be readily accessible by the processing system as well. Analysis jobs produce also large amounts of data. Users need to be able to store and gain access to data in the Grid environment. Since in the Grid environment data are not stored centrally, an automated management system incorporating the concept of file ownership and user quota is required.

To meet these requirements, a Distributed Data Management system [107] has been developed (see section 4.1). It provides a set of services to move data between Grid-enabled computing facilities whilst maintaining a series of databases to track data movements. The vast amount of data is also grouped into datasets based on various criteria (e.g. physics characteristics, production batch run, etc.) for a more efficient query and retrieval. DDM consists of three components: a central dataset catalogue, a subscription service and a set of client tools for dataset lookup and replication. The central dataset catalogue is a collection of independent internal services and catalogues that collectively acts as a single dataset book-keeping system. Subscription services enable data to be automatically pulled onto a site. A user can ensure that he/she is working on the latest version of a dataset by simply subscribing to it. Any subsequent changes to this dataset (i.e. additional files, version changes, etc.) trigger a fresh download of the updated version automatically. Client tools provide users with the means to interact with the central dataset catalogue. Typical actions include listing, retrieving and inserting datasets.

5.1.1 ATLAS Event Data Model (EDM)

The ATLAS detector will produce one Peta-byte of data per year, a vast amount of information that prohibits the wide distribution of raw data to worldwide collaborators. To enable physicists to analyze the data at remote sites three additional stages of datasets are introduced:

- The Event Summary Data (ESD) [28] containing detailed output of the detector reconstruction which is to be produced from the raw data. It contains sufficient information to allow particle identification, track refitting, jet calibration, etc. Thus allowing for the rapid tuning of reconstruction algorithms

and calibrations. The target size for the ESD is 500 kB per event.

- The Analysis Object Data (AOD) [28] is a summary of the reconstructed event, and contains sufficient information for general analysis. Several tailor–made streams of AOD's are foreseen for the different needs of the physics community. The AOD can be produced from the ESD and thus it is unnecessary in general to navigate back and process the raw data, adding significant cost and time benefits. The target size for the AOD is 100 kB per event.
- Derived Physics Data (DPD) is an n-tuple-style representation of event data for end-user analysis and histogramming. The inclusion of DPD in the Computing Model is an acknowledgment of the common practice by physicists of building subsamples in a format suitable for direct analysis and display by means of standard analysis tools (ROOT, for example), though software providers certainly expect that analysis, histogramming, and display via standard tools will be possible with AOD as input.

Inevitably, there will be some overlap between the different reconstruction realms: for example, some objects will exist in both AOD and ESD. There will also be 'TAGs' [108] on each event, indicating some general features of the event, and thus allowing a quick access. The target size for the TAGs is 1 kB per event.

5.1.2 ATLAS Analysis Computing Model

The ATLAS analysis model is divided into two components (see figure 5.1):

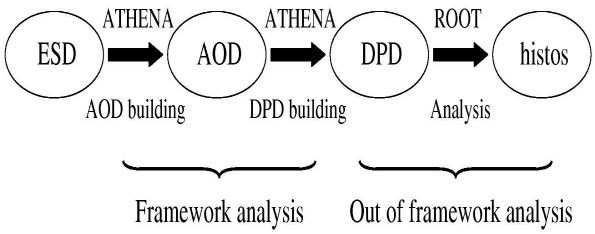

Figure 5.1: *ATLAS Analysis Model.*

- Scheduled central production of augmented AODs, n-tuples and TAG collections from ESD. Derived files moved to other Tier-1s and Tier-2s.
- User-augmented AOD streams, n-tuples, new selections, etc. and individual user simulation and CPU-bound tasks matching the official MC production. This could imply modest to large job traffic between Tier-2s (and Tier-1s, Tier-3s).

Distributed Analysis emphasis is on a batch model to implement the ATLAS Computing Model. In fact, interactive solutions are difficult to implement on top of

the current middleware layer. Distributed Analysis effectiveness depends strongly upon the hardware and software infrastructure. Analysis occurs in parallel to production. For this reason, Job Priorities [109] are being developed and deployed at some sites for testing purposes.

Analysis is divided into 'group' and 'on-demand' types:

- Group analysis (Tier-2s) is characterized by access to full AOD and some ESD data:
  - This is resource intensive and must be a scheduled activity
  - can back-navigate from AOD to ESD at same site
  - can harvest small samples of ESD (and some RAW) to be sent to Tier-2s
  - must be agreed by physics and detector groups
  - will produce deep copies of subsets, dataset definitions and TAG selections
- On-demand analysis (Tier-1/Tier-3)
  - could be restricted Tier-2s, specializing some Tier-2s for some groups (note that all Tier-2s are for ATLAS-wide usage), and CERN Analysis Facility (CAF).
  - Role and group based quotas (to be determined per group not per user) are essential.
  - Data selection over small samples with Tier-2 file-based TAG dataset selector and TAG queries over large samples by batch job to database TAG at Tier-1s/large Tier-2s.
  - Analysis data: group-derived DPDs, ROOT Trees, subsets of ESD and RAW.

5.1.3 Distributed Analysis Strategy

The ATLAS Distributed Analysis system is in development and several approaches are being studied and evaluated. In this way, ATLAS has adopted a multi-pronged approach to Distributed Analysis by exploiting its existing Grid infrastructure directly via the various supported Grid flavours and indirectly via the ATLAS Production System, ProdSys, (see figure 5.2).

Figure 5.2 shows various front-end clients enabling distributed analysis on the existing Grid infrastructures. These front-end clients (PanDA [110]/Pathena [111], ATCOM [112] and GANGA [113]) are intended to perform Distributed Analysis according to figure 5.2.

PanDA (Production and Distributed Analysis) is a job management system associated with OSG designed specifically for both distributed production and Distributed Analysis. PanDA has native support for the ATLAS DDM system allowing to accept DDM datasets as input (pre-staging it whenever required) and

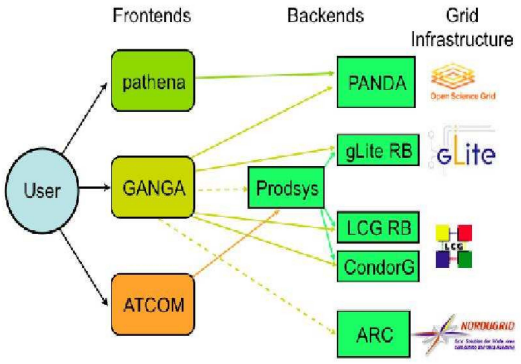

Figure 5.2: *Distributed Analysis strategy.*

producing DDM datasets as output (retrievable using DDM tools). PanDA offers users a comprehensive system view presenting heterogeneous distributed resources as a single uniform resource, accessible via a standard interface. It also has extensive web-based job monitoring and browsing capabilities. PanDA does not have a Graphical User Interface (GUI) but looks to GANGA in order to provide a graphical job definition and submission interface. PanDA has achieved this by exposing a useful set of client API.

Pathena is a Python [71] script designed to enable access to OSG resources via the PanDA job management system. It is in the process of becoming a drop-in replacement for the executable used in the ATLAS software framework. Users are able to exploit distributed resources for their analysis activities with minimal inconvenience. Pathena makes the submission of analysis jobs to the PanDA system a painless two-stage process involving an optional build step (where user code can be compiled) followed by an execution step (with built-in job splitting capabilities). A further merge step is in development, allowing the resulting output datasets from split jobs to be consolidated.

ATCOM (the ATLAS Commander) was the dedicated graphical user interface front-end to the ATLAS Production System, designed to be used by a few expert users involved in large-scale organized production of ATLAS data. It had the potential to be used for Distributed Analysis purposes as well. ATCOM was discontinued due to lack of support and priority to GANGA and Pathena user interfaces.

GANGA (Gaudi [114]/Athena and Grid Alliance) is a powerful user friendly front-end tool for job definition and management, jointly developed by the ATLAS and LHCb [10] experiments. GANGA provides distributed analysis users easy access to the whole Grid infrastructure supported by ATLAS. This is achieved by interfacing to an array of submission back-end mechanisms. Submission to the Production System and ARC are planned for the near future.

GANGA currently provides two user interface clients: a Command Line Interface

(CLI) and a Graphical User Interface (GUI). In addition, it can also be embedded in scripts for non-interactive/repetitive use. GANGA, intended to satisfy both ATLAS and LHCb requirements (unlike Pathena and ATCOM that are designed for specific ATLAS tasks), has been designed from the onset to be a highly extensible generic tool with a component plug-in architecture. This pluggable framework makes the addition of new applications and back-ends an easy task. A synergy of GANGA and DIANE [115] (a job-distribution framework) has been adopted in several instances.

GANGA has been used to submit various types of jobs to the Grid including a few High Energy Physics (HEP) experiments, like Babar [116] and NA48 [117]. There are also other HEP experiments, like PhenoGrid [118] and Compass [119], in the preliminary stages of looking to exploit GANGA for their own applications. In addition, other kind of applications have used GANGA also for the search for drugs to combat Avian flu, regression testing of GEANT 4 to detect simulation result deviations and the optimization of the evolving plan for radio frequency sharing by 120 countries.

GANGA is currently in active development with frequent software releases and has an increasing pool of active developers.

5.2 Prototypes for ATLAS Distributed Analysis

Jobs for simulation and reconstruction use in general GAUDI/Athena software that results from a coordinated, experiment-wide effort, and is installed in many sites. The user submitting jobs, possibly a production manager, performs the job configuration, involving selection of the algorithms to be run, definition of the algorithm and specification of inputs and outputs. The situation is similar for an analysis job, except that the users running a given analysis usually want to load one or more self-written algorithms and therefore use code that may be available only in individual work areas. Another major difference between analysis and production consists in the amount of input data. Typically an analysis job requires Giga-bytes or even Tera-bytes of data collected in datasets distributed among many storage elements. The dataset location is performed by accessing various meta-data and file catalogues and hence a mechanism for data location has to be included as an integral part of Distributed Analysis.

In the following, different types of Distributed Analysis tools (ProdSys, Pathena and GANGA) are described. ProdSys, the first prototype, started in 2005. Finally, the ProdSys effort moved to GANGA, the Distributed Analysis tool for LCG/EGEE and has been adapted for all ATLAS Grid flavours. Pathena is another Distributed Analysis tool developed by ATLAS in the US.

5.2.1 PanDA

The PanDA production and Distributed Analysis system has been developed by US-ATLAS to meet ATLAS requirements for full scale production and Distributed Analysis processing. PanDA development began in August 2005 and took over US-ATLAS production in December 2005.

Within the ATLAS Production System, PanDA is used as a regional (OSG) 'executor', interacting with an ATLAS production 'supervisor', in order to receive and report production work. PanDA also operates as a system in its own right, for both production and analysis work-loads. In September 2006, additional effort started in collaboration with the OSG/Grid3 and Condor to generalize PanDA as a generic high level work-load manager, usable by anyone in the OSG or even a wider Grid community.

Jobs are submitted to PanDA (see figure 5.3) via a simple python client interface, so users can define job sets, associated datasets and input/output files. Job specifications are transmitted to the PanDA server via http, with submission information returned to the client. This client interface has been used to implement PanDA front ends for ATLAS production (Panda Executor Interface), Distributed Analysis (Pathena), and US regional production. The PanDA server receives work from these front-ends into a global job queue, operated by a brokerage module to prioritize and assign work on the basis of job type, priority, input data and location, and available CPU resources. Allocation of job blocks to sites is followed by the dispatch of input data to those sites, handled by a data service interacting

with the ATLAS Distributed Data Management system. Data pre-placement is a strict precondition for job execution: jobs are not released for processing until the data arrives at the processing site. When data dispatch completes, jobs are made available to a job dispatcher.

An independent subsystem manages the delivery of pilot jobs to worker nodes via various scheduling systems. A pilot launched on a worker node contacts the dispatcher and receives an available job suitable to the given site. If no appropriate job is available, the pilot may immediately exit or may pause and ask again later, depending on the configuration. An important attribute of this scheme for interactive analysis, where minimal latency from job submission to launch is important, is that the pilot dispatch mechanism bypasses any latencies in the scheduling system for submitting and launching the pilot itself.

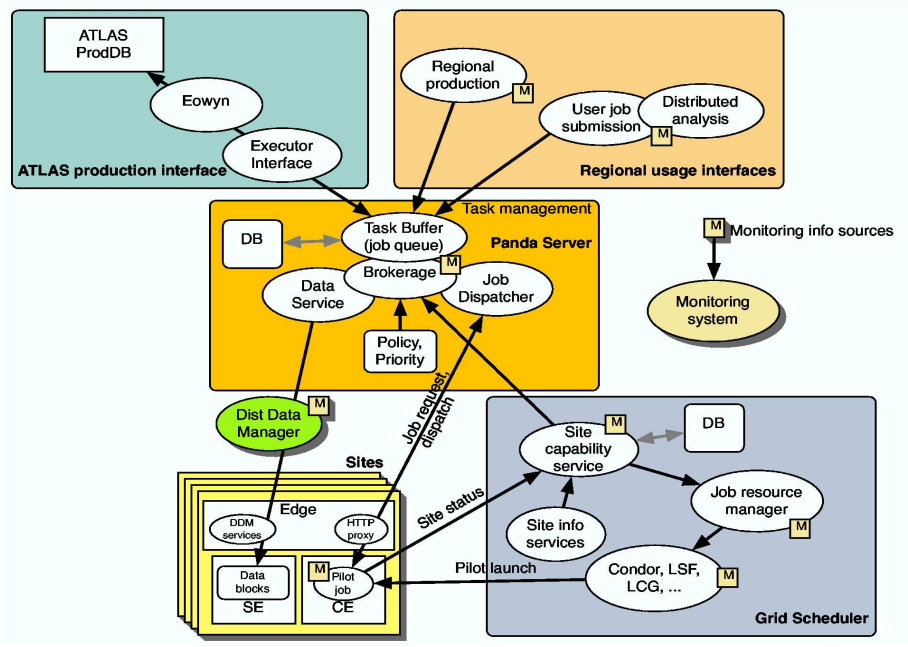

Figure 5.3: *PanDA architecture.*

Job flow in PanDA is shown in figure 5.4:

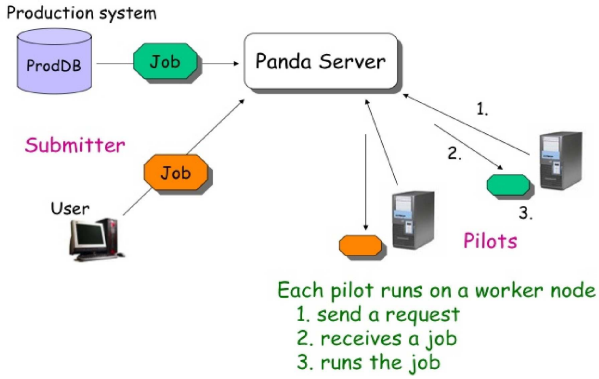

Figure 5.4: *PanDA job flow.*

A user job (see figure 5.5) is composed of sub-jobs, i.e. one buildJob and many runAthena's. The buildJob receives source files from the user, compiles and produces libraries. The completion of buidJob triggers runAthena's. Each runAthena retrieves the libraries and runs Athena. Output files are added to an output dataset. Then DDM moves the dataset to the user local area.

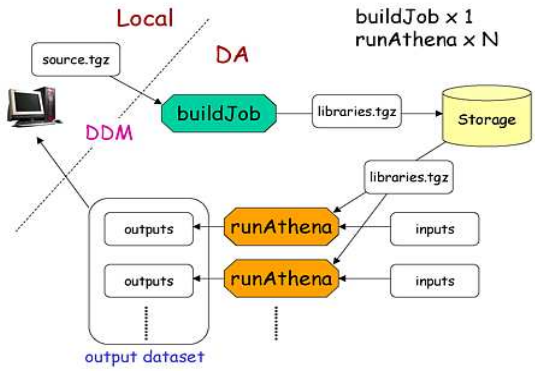

Figure 5.5: *Job submission using Pathena.*

The input dataset contains input files, such as POOL/ROOT files, byte-stream files or TAG collections. Input datasets must be registered in DDM before submission. Official datasets are in any case registered. When a user intends to process a private dataset, dq2_put, a DQ2 user script, is available.

Pathena is a glue-script intended to submit user-defined jobs to distributed analysis systems. It provides a consistent user-interface to Athena users. The following tasks are performed:
1. archiving of user's work directory
2. sending the archive to PanDA

3. extracting job configuration from jobOptions
4. defining jobs automatically
5. submitting jobs

5.2.2 GANGA

GANGA is an easy-to-use front-end for job definition and management, implemented in Python. It is being developed to meet the needs of a Grid User Interface (UI) by the ATLAS and LHCb experiments in High Energy Physics, and is a key tool for distributed-analysis systems [120, 121].

The experiment data-processing applications, including simulation, reconstruction and physics analysis, are based on the GAUDI/Athena C++ framework. They provide core services, such as message logging, data access, histogramming, and allow run-time configuration.

GANGA deals with configuration of ATLAS and LHCb applications, allows switching between testing on a local batch system and large-scale processing on the Grid, and helps to keep track of results. All this information is contained inside an object called GANGA job.

A job in GANGA is constructed from a set of building blocks (figure 5.6). All jobs have to specify the software to be run (application) and the processing system (back-end) to be used. Many jobs 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) for dividing a job into sub-jobs that can be processed in parallel, and for combining the various outputs. In this case, after splitting, the job becomes a master job and provides a single point of access for all sub-jobs.

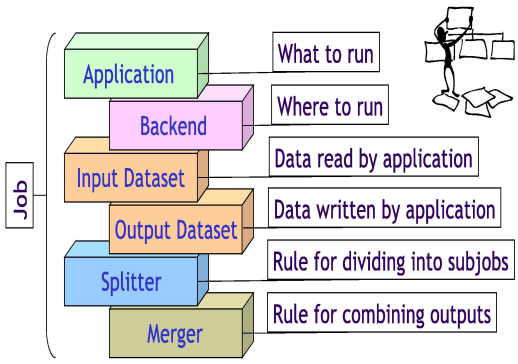

Figure 5.6: *Building blocks for constructing a GANGA job.*

The functionality of GANGA is divided between components (figure 5.7). GANGA Core links them together and performs most common tasks. It is represented by Application and Job Managers, Job Repository, and File Workspace. All components communicate via the GANGA Public Interface (GPI), designed to help

GANGA users and external developers to write scripts and plug-in modules more easily. The client allows access to GPI commands in any of these three ways: through a shell – the Command Line Interface in Python (CLIP); using GPI scripts, or through a Graphical User Interface.

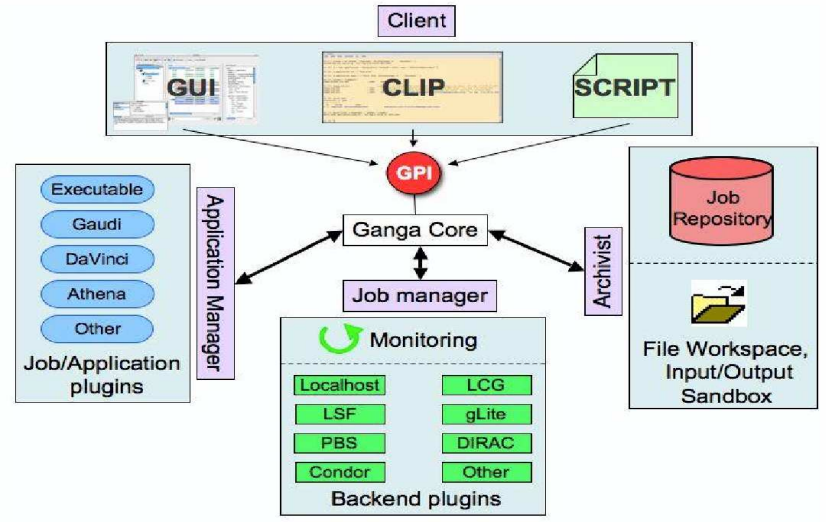

Figure 5.7: *Schematic representation of the GANGA architecture. The main functionality is divided between Application Manager, Job Manager and Archivist, and is accessed by the client through the GANGA Public Interface (GPI). The client can run the Graphical User Interface (GUI), the Command-Line Interface in Python (CLIP) or GPI scripts.*

The GANGA components are the following:

- The Application Manager deals with defining the task to be performed within a job, including the application to be run, the user code to be executed, the values to be assigned to any configurable parameters, and the data to be processed. GANGA calls the Application Manager for every job before submission. The task of the Application Manager is to return a configuration object which contains backend-independent information. In case of split jobs the operation of the Application Manager is optimized, so as to perform only configuration actions for sub-jobs that are not performed by the master job. For jobs with a complicated configuration, this procedure can reduce the required CPU by a large scale factor.
- The Job Manager takes care of any job splitting requested, packs the required user code, performs submission to the back-end, monitors the job progress, and retrieves output files when the job completes. It takes as input the configuration object prepared by the Application Manager and calls the application runtime handler. It also allows a dynamic connection to a new back-end, and running the same application on different Grid or batch systems. The Job Manager retrieves the information about a job status with the

help of the Monitoring component that polls plug-in modules representing back-ends where jobs are submitted.

- Acting as a simple book-keeping system, the Job Repository stores all GANGA jobs. The Job Repository and File Workspace provide persistency for jobs and their input/output sand-boxes so that GANGA can be restarted having access to the same information. If necessary, this information can be searched for particular jobs according to job meta-data, also stored in the Job Repository. The Job Repository can be either local or remote. Both types of repository have the same interface, so there is no difference from the program point of view. However they differ substantially in implementation and usage.

- The local Job Repository is a light-weight database written entirely in Python, that can be easily run on any computer platform. It is mainly designed for individual users and provides the possibility to work with GANGA off-line. In contrast, the remote Job Repository is a service centrally maintained for usage by many users. It has the advantage of job access from any location with internet connection. It supports secure connection and user authentication and authorization based on Grid proxy certificates. The performance tests of both the local and remote repositories show good scalability for up to 10 thousand jobs in one user session, with the average time of job creation being 0.2 and 0.4 second for the local and remote repository, respectively.

GANGA has default parameters that can be redefined at start-up using one or more configuration files, using a syntax readable by the standard Python ConfigParser module. The configuration files allow selection of the Python packages that should be initialized at GANGA start.

Users can use a GANGA client in different ways:

- GANGA's Command Line Interface with Python provides interactive job definition and submission from an enhanced Python shell, IPython [122], with many functionalities. Users need to enter just a few commands to set application properties and submit jobs to run the application on selected back-ends. Switching from one back-end to another is trivial as well.

- GPI scripts allow sequences of commands to be executed in the GANGA environment, and are ideal for automating repetitive tasks. GANGA includes commands that can be used outside the Python/IPython environment to create GPI scripts containing job definitions; to perform job submission based on these scripts, or on scripts exported from CLIP; to query job progress; and finally to kill jobs.

- The GUI (figure 5.8) aims to further simplify the user interaction with GANGA. It is based on the PyQt [123] graphics toolkit and GPI, and consists of a set of dockable windows and a central job monitoring panel. The main window includes:

– a toolbar, which simplifies access to all the GUI functionality;
– a logical folder organizer implemented as a job tree;
– a job monitoring table showing the status of jobs selected from the job tree;
– a job-detail panel, which allows inspection of job definitions.

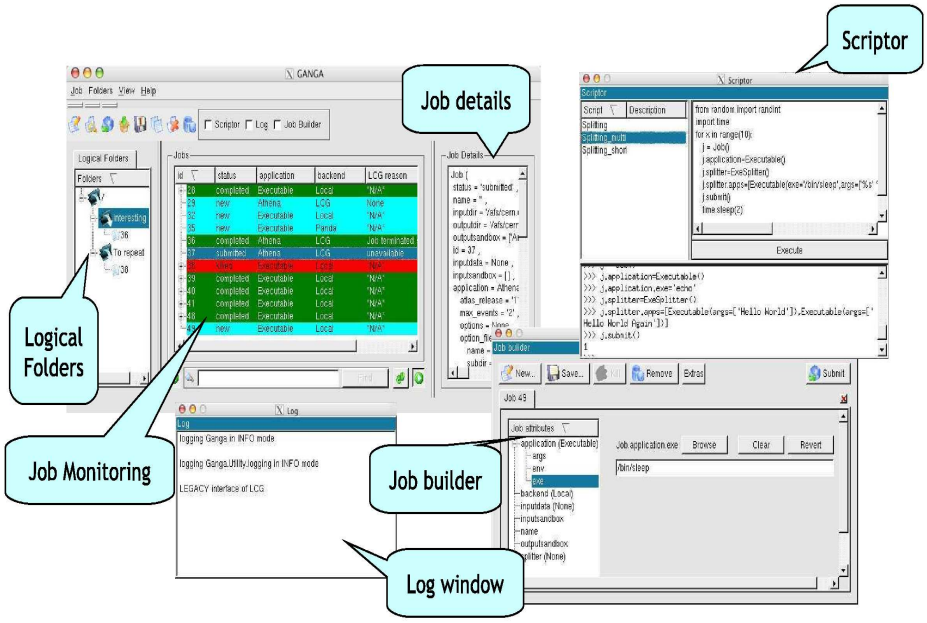

Figure 5.8: *Screen-shot of the GANGA GUI, showing: a main window, displaying job tree, monitoring panel and job details; an undocked scriptor window; a job builder window; and finally a log window.*

5.2.3 Distributed Analysis using the Production System

The ATLAS Production System has been used to submit physics analysis jobs to the LCG Grid flavor. Using the production framework for analysis has the advantage to take profit from the experience of the large scale production. The executors have evolved and provide nowadays a robust and scalable framework. Only limited additional resources are necessary to perform the required modifications to support analysis. This experience allowed to compare the execution of analysis tasks using such a system versus using direct submission to the Grid infrastructure. This activity was part of the ATLAS Data Challenges (DCs), intended to validate the Computing Model and to ensure the correctness of technical choices.

Analysis jobs are different than production jobs, since their key issue is that they access rapidly the data. Therefore, the experience gained operating the system is essential also to perform analysis using Grid resources. Table 5.1 (the same table

as table 4.1, in chapter 4) shows a summary of most frequent errors found during the ATLAS DC2 and Rome production.

System	Causes	Rate
Workload management system (WMS)	Total	1.6%
Data Management (DM)	Download input files	26.4%
DM	Upload output files	0.8%
ATLAS/LCG	Athena-crash	9.1%
ATLAS	Proxy expired	0.3%
LCG	Site misconfig.	0.9%
unclassified		9%

Table 5.1: *Summary of the main causes for job failures and their rates for the ATLAS DC2 and Rome production. As shown, the most problematic issue comes from Data Management* $\sim$ *26%.*

One of the main problems in Grid is data management, concretely data access on MMS systems (hybrid systems with a tape library and a disk front-end). Data need to be moved (staged) from tape to disk before they can be accessed. The middleware should be able to ensure the presence of data on disk when an application or a service needs to access them. Moreover, it should be possible to pin data, i.e. to ensure they remain staged on disk until they are needed or in use (in mass storage systems, automatic clean-up agents can free-up disk space if data is not accessed since long time). Unfortunately, this is not the case in LCG and a good strategy for data access on MMS systems was really missing. In particular, if data was not staged on disk, commands would fail during file access.

The lesson for analysis is the need to carefully prepare the data access. Distribution of data on several sites and local access to the data is a very important issue to minimize failures. Despite the possibility to run analysis jobs via the Production System, not all functionalities to support Distributed Analysis are currently available. An initial service was operated, with the aim to support ATLAS users in the future and a temporary solution based on the LCG tools and the LFC (LHC File Catalogue) was used. The ATLAS DDM system, DQ version 2, is currently integrated.

In the following, the technical issues that had to be addressed are discussed.

- A dedicated database was setup for analysis jobs to separate private work from the ongoing production.
- A generic analysis transformation was created, that compiled user code or the user package on the worker node, processed AOD input files and produced histograms, and n-tuple files as output.
- A Graphical User Interface (ATCOM) was used to define jobs (see figure 5.9) and tasks. The tool was also used to monitor (see figure 5.10) their current status.

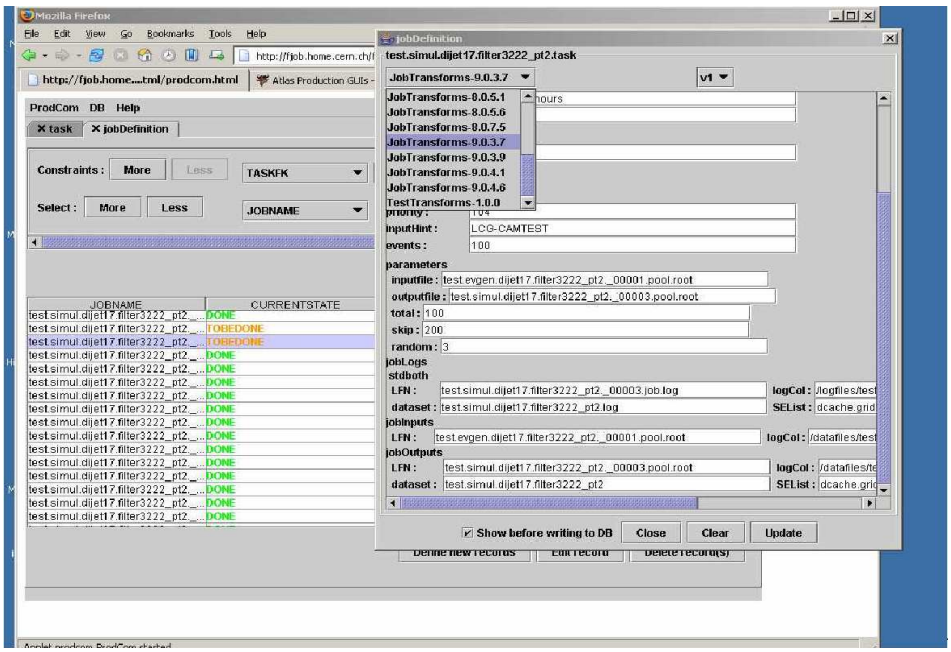

Figure 5.9: *Screen-shot of ATCOM used for job definition.*

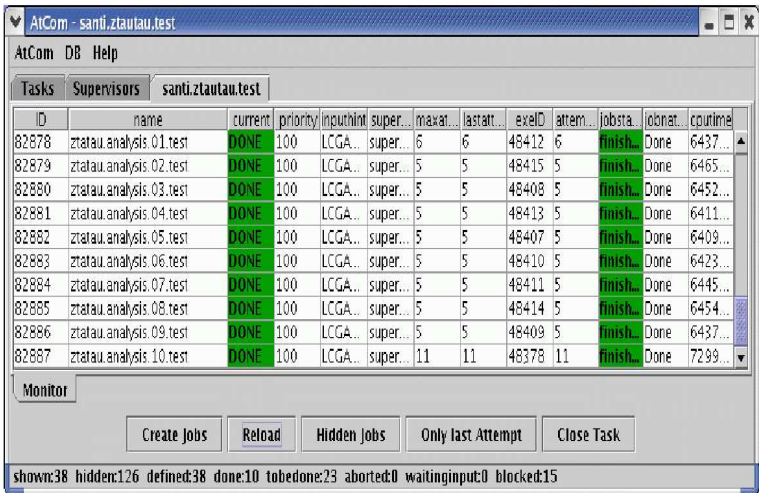

Figure 5.10: *Screen-shot of ATCOM used for job monitoring.*

The security was addressed by controlling the access to the database using standard Grid certificates. The jobs were submitted to the Grid using the certificate of the user operating the executor. This was considered a temporary solution until a full security model was introduced in the Production System.

To perform an analysis, the user has to define a task and associated jobs according to the conventions of the ATLAS Production System. The task contains summary information about the jobs to be processed (input/output datasets, transformation parameters, resource and environmental requirements), and individual jobs are defined by specific parameters needed for execution. The task and the jobs are defined in Grid-independent way.

This analysis setup was evaluated by several early adopters, in fact, I am one of them (see section 5.2.4). These users defined their analysis tasks via interfaces, and an operation team followed the processing of the jobs submitted via the executor. Analysis was done running a supervisor and a Lexor/CondorG instance. The analysis setup was tried only on one LCG Grid flavour. AOD input data was available at several sites, where jobs were submitted. The middleware did not provide means to separate production and analysis jobs. Queues that were filled with simulation jobs from the ongoing production caused long waiting times for the analysis jobs. In total, 155 GB of merged datasets were used for Distributed Analysis tests. The datasets were registered using the CERN LFC catalogue and replicated in the sites shown in table 5.2.

Site	Storage Element	Computing Element
LNF	atlases.lnf.infn.it	atlasce.lnf.infn.it
IFIC	castorgrid.ific.uv.es	lcg2ce.ific.uv.es
SINICA	lcg00123.grid.sinica.edu.tw	lcg00126.grid.sinica.edu.tw
CNAF	grid007g.cnaf.infn.it	gridit-ce-001.cnaf.infn.it
NA	griditse.na.infn.it	griditce01.na.infn.it
PIC	castorgrid.pic.es	ce01.pic.es
MI	t2-se-01.mi.infn.it	t2-ce-01.mi.infn.it
CERN	castorgrid.cern.ch	lxgate13.cern.ch
ROMA	t2-se-01.roma1.infn.it	t2-ce-01.roma1.infn.it

Table 5.2: *Sites where datasets were replicated for ATLAS Distributed Analysis tests.*

The two selected algorithms were a $t\,\overline{t}$ and a SUSY reconstruction program provided by the physics groups. In case of the $t\,\overline{t}$ reconstruction, the algorithm was already included in the Athena release and the required software was already installed in the Grid sites. In case of the SUSY algorithm, the source code was transferred with the job and the shared library compiled on the worker node. Two datasets were used. Fifty events per file, a total of $7,000$ files for the first dataset and $1,000$ for the second one. Eighty jobs (70 for the first one and 10 for the second one) with 100 input files each were defined with ATCOM. These 80 jobs were submitted with Eowyn-Lexor/CondorG and were processed in several LCG sites (ific.uv.es, gridpp.rl.ac.uk, farmnet.nikhef.nl, etc.). Each job produced 3 output

files (n-tuple, histogram, and log) stored on Castor-based [124] Storage Elements (SE) (CERN, IFIC, Taiwan). The jobs were running in several sites and were instructed to store the output in just one SE close to the user location. The ROOT package was used to merge the histogram output files and to analyze the results.

It is clear that all these tests were only a first step towards the validation of the Computing Model and Distributed Analysis. More realistic tests have to involve many users working in concurrent mode. This requires not only progress in the application, but also progress in the Grid middlewares and the site configurations. An example is the ongoing discussion on Job Priorities in LCG, that should allow the coexistence of production and user analysis activities in a Grid site with sufficient resources.

For the main test, few resources were used in 9 sites. The system was able to process jobs of 10 thousand events in about 10 min. This performance was difficult to achieve due to the instability of the major software components, still in development phase. Nevertheless, this first attempt is encouraging and promising.

The limited size of this test was implied by the limited size of data currently available at that time. With the LHC start-up, much more data and resources for analysis are expected. Continuation of this activity requires evolution of the setup in parallel with the ongoing evolution of the Production System. Improvements in the underlying system, as the data management, would facilitate th analysis tasks in the future.

Compared to direct submission, one could observe a more robust execution and by taking profit of the advanced monitoring capabilities of the system. A drawback was that such a system represents an additional infrastructure element, to be operated by the experiment. For ATLAS and also other experiments, a relevant question is to evaluate the advantages of using such a system for analysis, requiring additional effort to operate it as a service for users. Distributed Analysis in ATLAS is still a work in progress. With the start of LHC, one could expect a dramatic increase of data volumes. This implies also that general ATLAS users employ resources on the Grid to perform their analysis.

5.2.4 Example of application using the ProdSys prototype

For the ATLAS Distributed Analysis group, I have been a physicist trying to perform an analysis study using one of the prototypes for the ATLAS Distributed Analysis mentioned before. Concretely, I have used the ATLAS Production System prototype taking profit from my previous experiences involved in the ATLAS Data Challenges and the ATLAS Production System.

In this section, an example of Distributed Analysis application is shown for $Z_H \rightarrow t\bar{t}$ using the ProdSys prototype, a heavy Z gauge boson decaying into tops in the Little Higgs model (see section 3.3). This dataset was made in the official 'Rome Production' for the Exotics working group using the Athena full chain simulation version 10.0.1. A total of 400 AOD's were produced, each AOD containing 50

events (20,000 events in total). To perform the analysis, according to the Distributed Analysis using the Production System, I defined a task and associated jobs (see 5.2.3 and 4.1). Task and jobs were created in the ATLAS Development DataBase (DevelDB) using ATCOM, a pilot tool to provide a GUI for the Production System used for DC2. The application supported only creating, editing and submitting jobs. A task and four jobs were defined. Each job had as input 100 AOD's (5,000 events in total).

IFIC participated in DC2 running a Lexor instance and using some dedicated resources for production (see section 4.2.6). In this way, I used this Lexor instance to pick jobs defined in the DevelDB and submit them to LCG resources. PhpOradmin tool is an ATLAS database web interface. I used this web interface to select the development database and check records. ATCOM is a java GUI that gives interactive access to the ATLAS DataBase and can satisfy the monitoring requirements. In fact, I used this tool to monitor the progress status of tasks and jobs.

Some performance tests were made running these analysis jobs at IFIC and CERN, because jobs should be submitted where data are located. Tables 5.3 and 5.4 show the time consumed by the following steps running an analysis job: copy of input data (stage in), algorithm CPU processing time (Athena) and copy of output data to local SE (stage out).

IFIC	stage in	Athena	stage out
Job 1	10 min 55.9 sec.	4 min. 7.8 sec.	1 min. 6.5 sec.
Job 2	12 min. 37.7 sec.	4 min. 5.6 sec.	1 min. 5.7 sec.
Job 3	10 min. 42.7 sec.	3 min. 58.9 sec.	1 min 5.2 sec.
Job 4	10 min 45.1 sec.	4 min. 5.0 sec.	1 min 6.8 sec.

Table 5.3: *Time for each job run at IFIC to copy the required input data to the WN (stage in), CPU processing time (Athena) of the analysis job and time to copy the output data to the local SE (stage out).*

CERN	stage in	Athena	stage out
Job 1	10 min 42.7 sec.	5 min. 16.4 sec.	1 min. 7.5 sec.
Job 2	18 min. 6.4 sec.	5 min. 6.4 sec.	1 min. 8.2 sec.
Job 3	18 min. 55.3 sec.	5 min. 14.6 sec.	1 min. 6.7 sec.
Job 4	18 min. 11.5 sec.	4 min. 38.3 sec.	1 min. 10.3 sec.

Table 5.4: *Time for each job run at CERN to copy the required input data to the WN (stage in), CPU processing time (Athena) of the analysis job and time to copy the output data to the local SE (stage out).*

Output data were stored using the IFIC/CERN CASTOR storage system. These output ROOT files containing histograms were merged using ROOT. Finally, after merging, figure 5.11 shows the analysis result.

These analysis study results were taking into account as feedback by the ADA ProdSys group, since input from users was needed to improve the system.

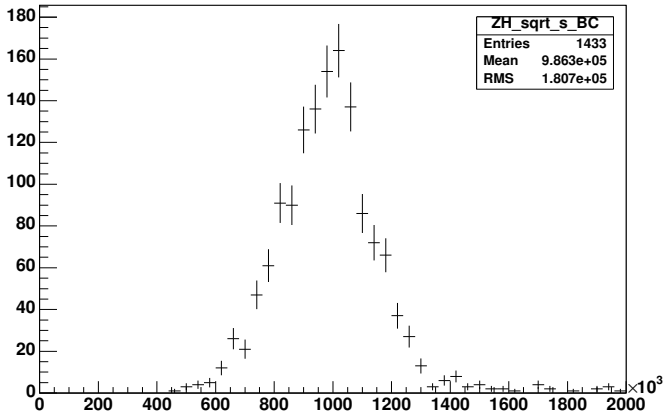

Figure 5.11: *$Z_H \rightarrow t\,\bar{t}$ invariant mass distribution after merging the output ROOT file histograms from the analysis jobs. Invariant mass in* GeV *units is showed in the* y *axis.*

My experience in the ATLAS Distributed Analysis system is only limited to the LCG/EGEE Grid flavour, so it means that I have not used Pathena yet. But, I have also used GANGA to perform some analysis tests, but no in such an exhaustive way as for the ADA ProdSys prototype. However, comparing the GANGA and the ProdSys prototypes:

- GANGA provides two user interface clients: a Command Line Interface and a Graphical User Interface. ProdSys provides a GUI (ATCOM) also to define and monitor tasks and jobs, but it doesn't hide all the Grid environment complexities for users as GANGA. In this sense, GANGA is more user-friendly than ProdSys. Users only need to fulfill some fields using the GANGA GUI to submit jobs.
- ProdSys has been used by 'expert users' who have some background about the system. However, GANGA is used by general users who don't need to have any background about the back-end systems.
- GANGA supports different bank-ends to run the jobs (see figure 5.2), including ProdSys as one of these back-ends. However, ProdSys only supports itself.
- ProdSys is easy to use, but it needs a ProdSys instance running, which makes the system more difficult to run. In addition, it's unclear who should maintain this ProdSys instance for analysis purposes: expert users, ProdSys team or ADA group.

Finally, the ATLAS Management Computing Board decided to integrate the Distributed Analysis collaboration efforts in just one project, GANGA, that can use the ProdSys and other systems also as back-ends to perform analysis. So, the ADA ProdSys effort has been moved to GANGA.

5.3 Towards a local user analysis facility: the Tier-3 prototype

According to the ATLAS Analysis Computing Model (see section 5.1.2), analysis is divided into 'group' and 'on-demand' types. This analysis will be performed by physics groups, not by individual users, on Tier-2s. Therefore, the Tier-2 resources will be preferentially used by physics analysis groups. This means that users from universities and institutes need some extra computing resources, in addition to Tier-2 resources, to perform their own work and then contribute with their studies and algorithms to the group effort. There is a hierarchical model in the ATLAS Computing Model (see section 1.3.7) and despite the official infrastructure, a Tier-3 center is necessary to be introduced for this purpose.

The ATLAS Computing Model has a hierarchical model (Event Filter, Tier-0/Tier-1/Tier-2) with specific roles and responsibilities:

- Event Filter Farm at CERN
  - Assembles data (at CERN) into a stream to the Tier-0 Center
- Tier-0 Center at CERN
  - Data archiving: Raw data to mass storage at CERN and to Tier-1 Centers
  - Reconstruction, calibration and alignment processing: Fast production of Event Summary Data (ESD) and Analysis Object Data (AOD)
  - Distribution: ESD, AOD to Tier-1 Centers and storage at CERN
- Tier-1 Centers distributed worldwide (10 centers)
  - Data steward: Re-reconstruction of raw data they archive, producing new ESD, AOD
  - Coordinated access to full ESD and AOD (all AOD, 20-100% of ESD depending upon site)
- Tier-2 Centers distributed worldwide (approximately 30 Centers)
  - Monte Carlo simulation, producing ESD, AOD sent to Tier-1 Centers
  - On demand user physics analysis of shared datasets
  - Scheduled working group activity
  - AOD processing, TAG extraction

Some local computer resources, beyond Tier-1 and Tier-2, for users from universities and institutes are required to do physics analysis, called Tier-3. A Tier-3 could be as small as a modern desktop computer on each physicist's desk, or as large as a Linux farm, perhaps operated as part of a shared facility from an institution's own resources.

The main goal of the Tier-3 is physics analysis on site with seamless access to all ATLAS Grid resources. For this purpose, analysis tools and access to data are very important, so these tools have to be installed in a Tier-3 infrastructure to facilitate the operations for users. Note that Tier-3 centers are outside the official infrastructure, but following the previous scheme, the specific role and responsibilities for a Tier-3 would be as follows:

- Tier-3 Centers
  - Interactive analysis
    - Protype:
      - reconstruction
      - calibration
      - selection algorithms
    - Final selection:
      - plots, studies, ...
  - Statistical analysis on AOD/DPD

where physicists from an institute or a university could perform physics analysis on site and have access to the different ATLAS simulation and analysis facilities: tools, data, etc. Actually, according to the ATLAS Computing Model, users should send analysis jobs to sites where data are available and extract relevant output as n-tuples or similar.

IFIC Tier-3 resources are going to be split into two parts:

- Some resources are going to be coupled to IFIC Tier-2 resources in a Grid environment. These extra Tier-2 resources will be used preferentially by Tier-3 users. While resources are idle, then they can be used by the ATLAS community.

- A PC farm to perform interactive analysis, outside Grid.

Physicists need to produce small MC simulation productions in order to develop analysis algorithms, such as selection algorithms and background studies, for example.

As a starting point, in order to perform a Tier-3 prototype at IFIC, the user requirements (IFIC users) to perform physics analysis are taking into account. One of these requirements for users is to produce small MC simulations and store the output data for further analysis. In the next section, section 5.3.1, the requirements for such a production using the Athena full chain simulation are described, based on my previous experience.

5.3.1 Requirements for a private Monte Carlo production

The Monte Carlo simulation production is going to be run at Tier-2 centers distributed worldwide (approximately 30 centers) according to the ATLAS Computing Model. In the following, a real user example is given: I have produced a CSC/DC3 Monte Carlo simulation private production inside the ATLAS Exotics working group to perform the study of b-tagging at very high p_T.

This CSC/DC3 private production consists of a set of three datasets. The physics processes are hadronic decays of Z_H in the Little Higgs model, where this Z_H decays into $b\,\overline{b}$, $u\,\overline{u}$, $c\,\overline{c}$. For each one of these hadronic decays a dataset has been made. Each dataset consists of $20,000$ events, so the total number of events is $60,000$. Each event has been processed using the Athena full chain simulation (FullSim events):

- Generation
- Simulation (GEANT 4) and Digitization
- Reconstruction

Each step of this CSC/DC3 private production is detailed as follows:

Generation step

Generator (Pythia) Athena jobOptions (one for each physics process) were validated by the Exotics MC production managers, then they submitted these jobOptions files to the ATLAS ProdSys: $50,000$ events per dataset were generated into 10 evgen (event generation) output files ($5,000$ events/evgen file) using the Athena release 12.0.5.1. Each evgen file size is approximately 193 MB. Information about these evgen jobs is recorded by the ATLAS ProdDB.

Simulation and Digitization step

In this case, two different steps from the Athena full chain simulation (GEANT 4 simulation and digitization) are run in the same job (Grid job). First, GEANT 4 simulation (simul) is run and, once the simulation output file is produced, then the digitization (digit). So, this job can be split into GEANT 4 simulation and the digitization step. Simulation is the step after generation. In fact, the generation output (evgen file) is the input data for GEANT 4 simulation. A typical simulation job is set to process only 50 events, since this is the most CPU-consuming step in the full chain. A simulation event usually takes 1400 kSI2k.sec = 700 sec. One evgen file is the input for several simulation jobs:

$$\frac{5000\ events/evgen\ input\ file}{50\ events/simul\ job} = 100\ simul\ jobs/evgen\ input\ file$$

this dataset contains 20, 000 events, so 4 evgen files and 400 simul jobs are produced to collect 20, 000 events inside this dataset. In fact, there are three datasets, so the total number of evgen files are 12 and the total number of simul jobs are 1200. For each job a simulation output file is produced, called simul.HITS. Each simul output data file size is approximately 88 MB.

As commented before, simulation and digitization are handled by a common job. Once simulation is finished, the output is used as input for digitization, that produces as output a digit.RDO file. Then additional 400 digitization output data files per dataset are produced, which means 1200 digitization output data files in total for the three datasets. Each digitization data file contains 50 events (50 events/ digit output file) and its size is approximately 138 MB. A digit event takes 21 kSI2k.sec = 10.5 sec. The job submission scheme is displayed in figure 5.12, where the user is logged into a UI (via ssh) and submits jobs using a Resource Broker (RB). Then the RB manages the jobs over all available resources according to the user job requirements.

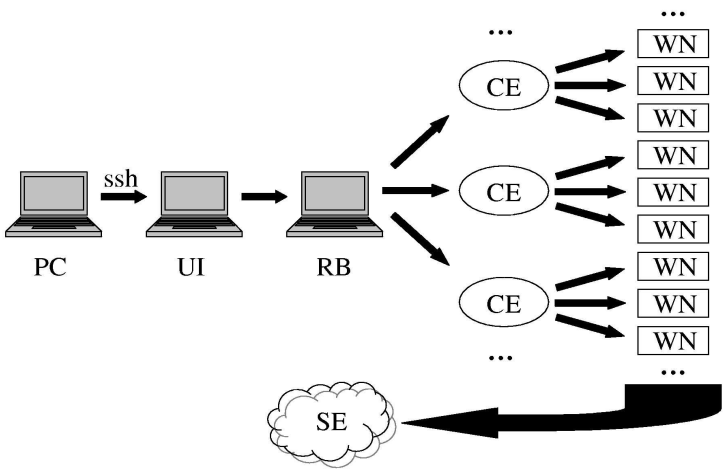

Figure 5.12: *Job submission scheme of simulation private production.*

Reconstruction step

Reconstruction is the last step in the Athena full chain simulation. The output data files from digitization (digit.RDO files) are the input for reconstruction. As all official CSC/DC3 productions, a reconstruction (recon) job contains 250 events (250 events/recon job), which means that a reconstruction job has as input 5 digit data files (50 events/digit file). Therefore a total of 80 reconstruction jobs per dataset and a total of 240 reconstruction jobs for the three datasets are required. Each reconstruction job produces 3 output data files: ESD, AOD and NTUP with 250 events per file. Each of these files has the following size:

ESD $\sim$ 760 MB
AOD $\sim$ 134 MB
NTUP $\sim$ 37 MB

ESD and CBNT NTUP are the first files (they require $\sim$105 seconds/event on an Intel(R) Pentium(R) 4 CPU 3.20 GHz and cache size of 512 KB) to be produced using as input the digit data files, then AOD (takes $\sim$ 4.1 seconds/event on an Intel(R) Pentium(R) 4 CPU 3.20 GHz and cache size of 512 KB) is produced using ESD as input, and finally NTUP (SAN [125] or HighPtView [126]) is produced using AOD as input. Therefore a reconstruction job has three steps to produce ESD, AOD and NTUP. Note that all these data files from the CSC private Athena full chain production have been stored at IFIC CASTOR SE (castorgrid.ific.uv.es). This means that all output data from jobs processed at any WN has been copied to castorgrid.ific.uv.es.

This private MC simulation production needs a storage capacity of 0.5 TB to store all output files from each step (see table 5.5). A total of 60 thousand Athena full simulation events has been produced. For future productions, at the LHC start-up, larger productions and some local resources will be needed to perform these private productions.

	events/file	size/file (MB/file)	# files per dataset	# datasets	total size (GB)
EvGen	5000	193	4	3	2.3
Simul	50	88	400	3	105.6
Digit	50	138	400	3	165.6
Recon	250	931	80	3	223.4
					496.9

Table 5.5: *Number of events and size per output data file at each step of the Athena full chain simulation. Number of files per dataset, total number of datasets and total size for each MC production step. Finally, the total size of the production is shown also.*

5.3.2 The Tier-3 prototype at IFIC

LHC data taking is going to start in 2008, less than one year left. For this reason, Tier-3 analysis facilities should be ready for physics analysis.

IFIC (as many other centers, institutes and universities) has a Tier-3 prototype in which different goals and steps should be done. As user at IFIC, my experience in Monte Carlo production and physics analysis is taken into account for several aspects. In a near future, after finishing my PhD, I'm going to be involved to help in the development and deployment of the Tier-3 setup at IFIC. Here, the first steps towards a Tier-3 infrastructure at IFIC are shown and described.

Users from IFIC could get access to the ATLAS Collaboration Tier-2 resources and Tier-3 resources (extra Tier-2 resources) from their desktops or laptops (see figure 5.13). These extra Tier-2 resources would be used preferentially by IFIC

users. Otherwise, if these resources are free, they could be used, in addition to the IFIC Tier-2 resources, by the ATLAS collaboration, ensuring an efficient use. On the other hand, other special requirements setup could be deployed for IFIC users also, like a PC farm to perform interactive analysis.

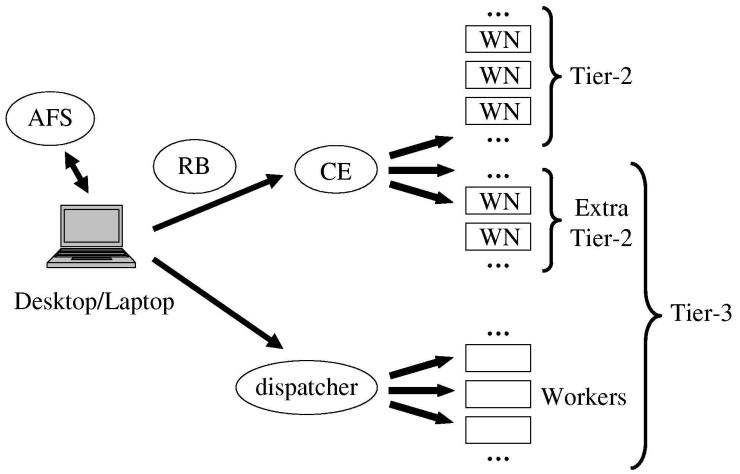

Figure 5.13: *User access from desktops or laptops to the ATLAS resources to perform physics analysis at IFIC. The IFIC Tier-3 resources will be coupled to the IFIC Tier-2 resources in a Grid environment, considered as extra Tier-2 resources. In addition, a PC farm will be deployed to perform interactive analysis tests in a non-Grid environment.*

An individual physicist can get access to the ATLAS software and analysis tools using its desktop or laptop. These software and analysis tools are already installed at IFIC as a first test. These tools are installed in a share file system area, called AFS [128], where IFIC users have access to them, in the same way as LXPLUS [129] at CERN, in order to avoid installing these tools at each desktop or laptop.

These ATLAS software (Athena, Atlantis [127], etc) and Distributed Analysis tools (GANGA) are installed under this path: /afs/ific.uv.es/project/atlas/software. Other useful tools, like ROOT or the Grid UI environment are also installed at IFIC AFS, but under another different path. The following tools can be setup by a user at IFIC:

- Athena releases:
  /afs/ific.uv.es/project/atlas/software/releases
- DDM/DQ2 client tools:
  /afs/ific.uv.es/project/atlas/software/ddm/pro03/dq2/dq2.sh
- DDM/DQ2 end-user tools:
  /afs/ific.uv.es/project/atlas/software/ddm/dq2_user_client/setup_IFIC.sh
- GANGA:
  /afs/ific.uv.es/project/atlas/software/ganga

- Grid UI environment:
  /afs/ific.uv.es/sw/LCG-share/sl4/etc/profile.d/grid_env.sh
- ROOT:
  /afs/ific.uv.es/acs/root/

Users can use their desktop as a private exclusive-use UI and developer station. In this way, some first steps shown in figure 5.13 has already started, like access to ATLAS Grid resources and tools as commented previously, but it needs more effort to install several ATLAS software tools (such as different Athena releases and other versioned tools).

IFIC Tier-3 Grid resources are going to be coupled to IFIC Tier-2 resources, but separating clearly them using different queues, shares or priorities for the ATLAS Collaboration Tier-2 and IFIC Tier-3 users. IFIC Tier-2 resources are used by the whole ATLAS Collaboration, reaching the ATLAS resources requirements (disk storage and CPU). Extra IFIC Tier-2 resources (Tier-3) are going to be preferentially used by IFIC Tier-3 users. These resources, WNs and SEs, will be used to run local simulation and analysis productions and to store interesting data for analysis:

- Production of MC samples of special interest for the local institution
- AOD private productions for further analysis
- AOD analysis

Having access to the ATLAS software and Grid resources, users from IFIC can perform local checks, running small Grid or non-Grid test jobs and develop their analysis code before submitting larger simulation or analysis productions to the Tier-2 or Tier-1. In case they need more resources.

The development of analysis code could motivate also a local copy of a small number (perhaps a few thousand) of ESD, AOD or RAW data events. Having the RAW, ESD, AOD and TAG for the same events.

In addition to the IFIC Tier-3 Grid resources, a PROOF (Parallel ROOT Facility) [130] farm for evaluation is going to be deployed for interactive analysis of n-tuples. This requires no direct connection to the ESD or AOD, but it does require access to the data when these n-tuples are generated. At this point, we should note that this infrastructure, Tier-3 Grid and non-Grid resources, is going to use the same SE, having access to the data in both cases. So, the n-tuples will be generated using the Tier-3 Grid or non-Grid resources if AODs are copied at the local Tier-3 SE. Careful deployment of the SE will be made to ensure the correct operation of the Tier-2 and avoid overloading from the Tier-3.

Interactive analysis is a very important issue for physics analysis. For this reason, interactive analysis with Derived Physics Data (DPD) n-tuples (HighPtView, SAN, ...) are analysed using ROOT. A PROOF farm with few powerful PCs

(outside the Grid environment) for evaluation is going to be installed to do this interactive analysis. One important issue is that this farm must be well connected to the SE in order to get a fast access to data.

Some data access performance tests, in which I have been involved, have been performed using the IFIC local SE, Lustre [131], and a b-tagging physics sample, $Z_H \rightarrow b\,\overline{b}$ with $M(Z_H) = 2$ TeV. AODs and DPDs for the same events have been used to analyze them with Athena and ROOT, respectively. This data has been accessed in different ways: local disk, RFIO [132] (CASTOR) and Lustre SE.

Figure 5.14 shows the results of the data access performance test using different ways of local storage, the Athena framework to analyze the data and 38 AOD's as input data files. Each AOD has 250 events (9500 events in total) and the average size per AOD is around 161 MB. As figure 5.14 shows, the slowest way to access this data is using the local disk, analyzing around 6.43 events per second. RFIO takes 6.96 events per second and the fastest way, Lustre, takes around 7.06 events per second. In general, the band-with used is around 4 MB/s.

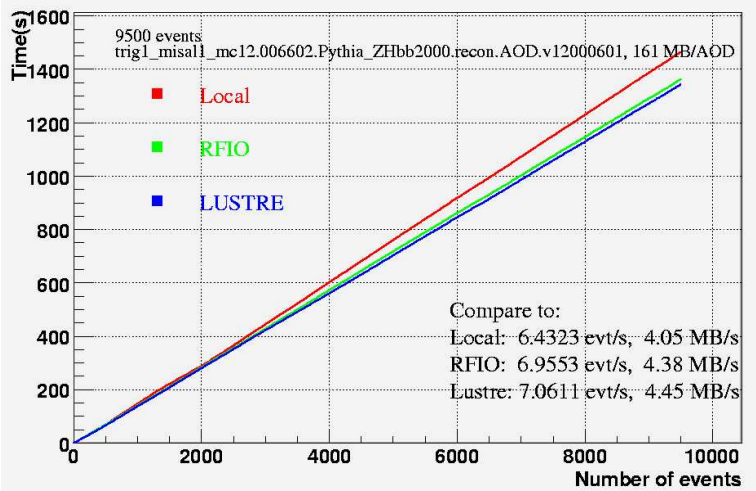

Figure 5.14: *Data access performance test using the IFIC local SE, the Athena framework and AOD's as input data files.*

Figure 5.15 shows the results of the data access performance test using different ways of local storage, the ROOT analysis tool to analyze the data and just one DPD with 9500 events as input data file. The size of this DPD is around 178 MB and contains all data from the previous AOD's (9500 events), extracted from these AOD's (see figure 5.1). As figure 5.15 shows, the slowest way to access this data is RFIO, in this case, analyzing around 97.75 events per second. Lustre takes around 101.24 events per second and the fastest way, local disk, takes around 102.46 events per second. The band-with used is around 2 MB/s, lower than in figure 5.14. This is due to the data structure and data size, since the event size in a DPD is much smaller than in an AOD.

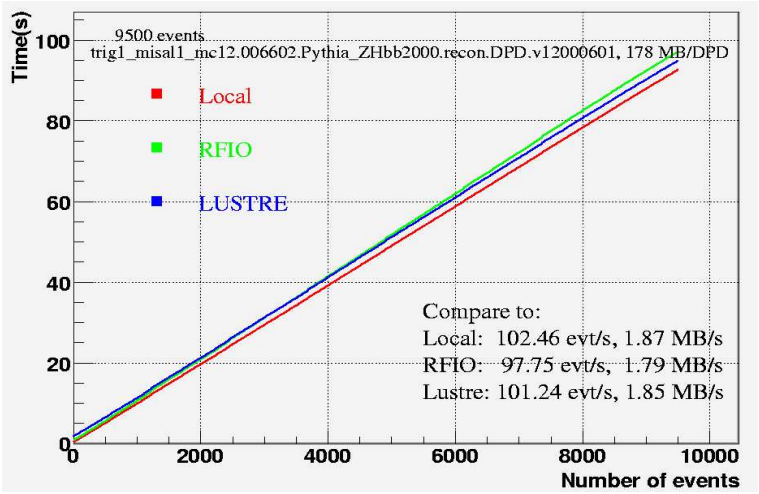

Figure 5.15: *Data access performance test using the IFIC local SE, ROOT analysis tool and one DPD as input data file.*

As shown in figure 5.14 and 5.15, the results obtained for data access performance tests using AODs and DPDs are different and show that DPDs are analyzed faster than AODs. A fast data access for analysis is very important.

The aim of the work being done at the Tier-3 at IFIC is to provide a flexible computing resource, so the physicist could make an analysis work in an easy and convenient way. The Tier-3 will provide interactive and non-interactive CPU as well as enough storage capacity and bandwidth to the data sources, the closest Tier2 (Federated Spanish T2) and other resources abroad.

Conclusions

In this thesis, the search for the Little Higgs gauge bosons Z_H and W_H using the ATLAS detector is discussed. The analysis of the new gauge bosons includes a study of the capability of ATLAS to perform an efficient b-tagging for high and very high p_T jets. The results running a ProdSys instance, called Lexor, installed at IFIC to manage the ATLAS CSC Monte Carlo production in LCG/EGEE and the test of the ATLAS Distributed Analysis system using the ProdSys are shown as well.

The discovery of the Little Higgs gauge bosons, Z_H and W_H, decaying hadronically, has been investigated. The best channels to search for these particles are leptonic decays. The modes involving hadronic decays are however useful to check the model. The analysis has been performed for masses of 1 and 2 TeV for both heavy gauge bosons. The discovery region is defined as the region with significance exceeding the value of 5 and a number of expected events above 10. This region has been obtained assuming an integrated luminosity of $\mathcal{L} = 3 \cdot 10^5\, pb^{-1}$ and extrapolating the results obtained for masses of 1 and 2 TeV. As a summary, the analysis of expected signal and background shows that the decays $Z_H \rightarrow b\,\overline{b}$ and $t\,\overline{t}$ are difficult to detect, but the decay $W_H \rightarrow t\,b$, on the contrary, might yield a signal clearly separable from the background, up to W_H masses of 2.5 TeV, if the parameter of the model $\cot\theta$ is assumed to be equal to 1. The detection of the $W_H \rightarrow t\,b$ allows a measurement of BR($W_H \rightarrow t\,b$) and, by comparison with BR($W_H \rightarrow l\,\nu$), provides an interesting test for the couplings of the W_H boson.

In order to perform this analysis and reduce the background, the b-tagging capability of ATLAS for b-jets coming from Little Higgs heavy particle decays has been studied. Two different samples containing b-jets have been analysed: a heavy Higgs boson ($m_H = 400$ GeV) in associated production with a W, and the Little Higgs s-channel production of a heavy Z_H(2 TeV). In both cases, the heavy particle is assumed to decay into b-quark pairs. The rejection obtained using the 2D algorithm is:

$H(400):$	$R_u = 139 \pm 6$	50% b-tagging efficiency
$Z_H(2000):$	$R_u < 10$	50% b-tagging efficiency

In many ATLAS searches, a typical assumption for light jet rejection is $R_u = 100$, for a b-tagging efficiency of 60%. The rejection for the very high p_T sample is worse, as expected. An increased rejection may be obtained by reducing the b-tagging efficiency:

$$Z_H(2000): \quad R_u \sim 1000 \quad 10\% \text{ b-tagging efficiency}$$

At very high p_T, the impact parameter suffers from large errors and the 2D algorithm has a poor discriminating power. The SV1 performance is expected to be less degraded than the 2D performance. The b-tagging performance at very high p_T, based on the 2D algorithm, can however be improved by using the SV1 algorithm.

A Lexor ProdSys instance for the ATLAS CSC/DC3 production has run at IFIC from 25th January 2006 until 9th August 2006, managing 62,910 successful jobs (a successful job is a job ending properly). The total number of successful jobs in ATLAS LCG/EGEE was 393,714 and around 16% of them were managed by the instance run at IFIC. The total number of succesful jobs processed at IFIC were 5,855, i.e. 1.5% of the total ATLAS LCG/EGEE. Some resources at IFIC were devoted to running this instance and contributing to the ATLAS CSC/DC3 Monte Carlo production:

- One UI, where the Supervisor and Executor were installed.
- One RB to manage the jobs picked up by the Executor.
- A BDII to provide information about the ATLAS LCG/EGEE grid resources.
- And a Jabber to convert XML into JDL.

Distributed Analysis using the ProdSys has been tested taking profit of the experience gained using the large scale production. Several problems were found at the beginning, in particular data management problems. Presently the system has been improved and is more robust. All the effort devoted to the distributed analysis using the ProdSys has been used to improve GANGA, a successful analysis tool to perform Distributed Analysis in ATLAS.

Conclusiones

En esta tesis, se discute la búsqueda de los bosones *gauge* del modelo *Little Higgs* Z_H y W_H utilizando el detector ATLAS. El análisis de los nuevos bosones *gauge* incluye un estudio de la capacidad del detector ATLAS para realizar una eficiente detección para *jets* de tipo b de alto y muy alto p_T.

El descubrimiento de los bosones *gauge* del modelo *Little Higgs*, Z_H y W_H, se ha investigado considerando sus desintegraciones hadrónicas. Los mejores canales para la búsqueda de estas partículas son sus desintegraciones leptónicas. Los modos que incluyen desintegraciones hadrónicas son sin embargo útiles para comprobar el modelo. Los análisis se han realizado para masas de 1 y 2 TeV para ambos bosones *gauge* pesados. La región de descubrimiento se define como la región con significancia que excede el valor de 5 y que tiene un número esperado de sucesos mayor que 10. Esta región se ha obtenido considerando un luminosidad integrada de $\mathcal{L} = 3 \cdot 10^5\ pb^{-1}$ y extrapolando los resultados obtenidos para las masas de 1 y 2 TeV. Como resumen, el análisis de la señal y fondo esperado muestra que las desintegraciones $Z_H \rightarrow b\,\overline{b}$ y $t\,\overline{t}$ son difíciles de detectar, pero que la desintegración $W_H \rightarrow t\ b$, por el contrario, podría proporcionar una señal claramente separable del fondo, hasta una masa del W_H de 2.5 TeV, si el parámetro del modelo ($\cot\theta$) se considera igual a 1. La detección de $W_H \rightarrow t\ b$ permite una medida de BR($W_H \rightarrow t\ b$) y, por comparación con BR($W_H \rightarrow l\ \nu$), proporciona un interesante test para los acoplos del bosón W_H.

Para realizar este análisis y reducir el fondo, se ha estudiado la capacidad de ATLAS para identificar *b-jets* procedentes de las desintegraciones de partículas pesadas del modelo *Little Higgs*. Dos muestras diferentes que contienen *b-jets* se han analizado: un bosón de *Higgs* pesado ($m_H = 400GeV$) en asociación con un W, y la producción de un Z_H (2 TeV). En ambos casos, la partícula pesada se desintegra en pares de *quarks* de tipo b. El factor de rechazo (R) obtenido utilizando el algoritmo 2D es:

$H(400):$	$R_u = 139 \pm 6$	50% eficiencia para identificación de *quarks* de tipo b
$Z_H(2000):$	$R_u < 10$	50% eficiencia para identificación de *quarks de tipo* b

En muchas búsquedas realizadas con ATLAS, una hipótesis típica para el factor de rechazo de *jets* ligeros es $R_u = 100$, suponiendo una eficiencia de identificación de *quarks* de tipo b del 60%. El factor de rechazo para una muestra de muy alto p_T es peor, como se esperaba, pero se puede obtener un incremento del factor de rechazo reduciendo la eficiencia de identificación:

$Z_H(2000)$: $R_u \sim 1000$ 10% de eficiencia de identificación de *quarks* de tipo b

A muy alto p_T, el parámetro de impacto es más difícil de obtener y el algoritmo 2D tiene un poder discriminativo reducido. Sin embargo, el algoritmo SV1 se degrada en menor medida que el 2D. La identificación de *quarks* de tipo b a muy alto p_T, basado en el algoritmo 2D, puede por tanto ser mejorado utilizando el algoritmo SV1.

Una instancia del sistema de producción, Lexor, para la produccin de ATLAS CSC/DC3, estuvo operativa desde el 25 de enero hasta el 9 de agosto del año 2006, administrando 62, 910 trabajos de forma satisfactoria (un trabajo satisfactorio es un trabajo que ha acabado correctamente). El número total de trabajos satisfactorios en ATLAS LCG/EGEE fue de 393, 714 y de éstos, aproximadamente un 16% fueron gestionados por la instancia ubicada en el IFIC. El número total de trabajos satisfactorios procesados en el IFIC fue de 5, 855, es decir 1.5% del total de trabajos procesados por ATLAS LCG/EGEE. Se utilizaron los siguientes recursos del IFIC dedicados a esta instancia y para contribuir a la producción de Monte Carlo de ATLAS CSC/DC3:

- Un UI, donde el Supervisor y Ejecutor estaban instalados.
- Un RB para gestionar los trabajos gestionados por el ejecutor.
- Un BDII para proporcionar información sobre los recursos Grid de ATLAS LCG/EGEE.
- Y un *Jabber* para convertir mensajes XML en JDL.

El sistema de Análisis Distribuido utilizando ProdSys se ha testeado aprovechando la experiencia obtenida utilizando las producciones Monte Carlo a gran escala. Varios problemas fueron encontrados al principio, en particular relacionados con las gestión de datos. Posteriormente el sistema ha mejorado y ha mostrado mayor robustez. Todo el esfuerzo depositado en el sistema de análisis distribuido utilizando ProdSys se ha utilizado para mejorar GANGA, una herramienta de análisis para realizar Análisis Distribuido en ATLAS.

Agradecimientos

Mi estancia en el IFIC durante los cuatro años y medio desde que empecé el doctorado ha sido de lo más satisfactoria y provechosa, tanto a nivel laboral como personal. Durante este tiempo, he tenido la oportunidad y el placer de conocer gran cantidad de gente de la que guardo gratos recuerdos.

Me gustaría agradecer la ayuda mostrada por varias personas, tanto en el ámbito personal como académico.

En primer lugar, a José Salt por darme la oportunidad de llevar a cabo esta tesis doctoral en el IFIC.

A mis directores, Eduardo Ros y Santiago González, por su eficiente orientación en la realización de esta tesis y por el trabajo que he compartido con ellos, enriqueciéndome acedémica y profesionalmente.

A D. Antonio Ferrer, E. Higón y M^{a} V. Castillo (mi tutora) por haberme mostrado los primeros pasos en la Física de Altas Energías.

A los compañeros del CERN, con los que he pasado tan buenos momentos y que han hecho posible que las estancias fuera de casa hayan sido mucho más amenas. Entre ellos: Josemi (CIEMAT), Raúl (Cerdenyola), Marcos Fernández, Nacho (UAM), Carlos (UAM), Nacho Sevilla, Jesús (CIEMAT), Blanca Perea, Juan López (CERN) y Juan Alcaraz (CIEMAT).
También me gustaría agradecer a Armin Nairz, Luc Goossens y, especialmente, Gregor Mair las mañanas y tardes en las que hemos tomado café juntos, haciendo un pequeño descanso en el CERN.

Me gustaría nombrar a algunos compañeros de la facultad, de los que guardo gratos recuerdos y que aunque cada uno ha seguido un camino diferente, aún seguimos en contacto: Albert Ruiz, Mireia Romaguera, Jordi Gomis, Lluís Ros y Gabriel Jover.

Hay compañeros/as del IFIC a los que me gustaría nombrar, con los que también he compartido parte de tiempo en el CERN: David y Diego Calderón por cuidarme cuando tan sólo era un novato, Gustavo Conesa (compañero de piso en mi primera estancia en el CERN), Maria José Costa (aventurera donde las haya).

Me gustaría agradecer la compañía y labor de compañeros/as cercanos en las horas de trabajo en el IFIC: Azzedine 'M. Kaci', Farida Fassi y Gabriel Amorós.

A los compañeros/as del IFIC con los que he compartido tan buenos momentos. Entre ellos: Álvaro, Álex, Paqui, Juanan, Simona, Carlos, Carmen, Esteban, Julio, Ciro y Marcel. Especialmente nombrar a Javi, Belén y Santi:

- A Santi, además de su labor como director, por las estancias tan agradables en el CERN siendo compañeros de piso, de trabajo y de ocio durante varios meses.
- A Belén por su incondicional apoyo, siempre dispuesta a ayudarme.
- A Javi por sus charlas e incondicional ayuda en el trabajo, ayudándome a solucionar una gran cantidad de problemas encontrados.

Para finalizar, especialmente a mis padres por su confianza y apoyo en mi aventura científica.

Bibliography

[1] K. Hagiwara et al. (the Particle Data Group). *Review of particles physics*. Phys. Rev. **D 66** (2002) (http://pdg.lbl.gov)

[2] F.J. Harset et al. Phys. Rev. Lett. **B 46** (1973) 138

[3] UA1 Collaboration, G. Arnison et al. Phys. Lett. **B 122** (1983) 103.

[4] The LEP Collaborations. *A Combination of Preliminary Electroweak Measurements and Constraints on the Standard Model*, hep-ex/0511027 (2005).

[5] D. Costanzo. *Higgs physics at the Large Hadron Collider*, hep-ex/0105033.

[6] L. Ibaez. *Physics Beyond the Standard Model*. CERN Academic Training Lectures, 29 March - 2 April 2004, http://cdsweb.cern.ch/record/690000

[7] G. Altarelli. *Status of the Standard Model and Beyond*, hep-ph/0306055 (2003)

[8] G.F. Giudice, R. Ratazzi and J.D. Wells, Nucl. Phys. 595:250, 2001.

[9] ATLAS Collaboration. *The Large Hadron Collider Technical Design Report*. CERN/AC/95-05, 1995.

[10] LHCb Collaboration. *LHCb Technical Proposal*. CERN/LHCC/98-4, 1998.

[11] ALICE Collaboration. *ALICE Technical Proposal*. CERN/LHCC/95-71, 1995.

[12] ATLAS Collaboration. *ATLAS Detector and Physics Performance Technical Design Report*. CERN/LHCC/99-14, 99-15, 1999.

[13] ATLAS Collaboration. *Liquid Argon Calorimeter Technical Design Report*. CERN/LHCC/96-41, 1996.

[14] A. Marchesini et al. *HERWIG: A monte carlo event generator for simulating hadron emission reactions with interfering gluons*. Version 5.1 - April 1991. Comput. Phys. Commun., 67:465-508, 1992.

[15] Frank E. Paige and Serban D. Protopopescu. *ISAJET 5.30: A monte carlo event generator for p p and anti-p p interactions*. To appear in Proc. of 1986 Summer Study on the Physics of the Superconducting Super Collider, Snowmass, Colo., Jun 23 - Jlu 11, 1986.

[16] T. Sjostrand. *High-energy physics event generation with PYTHIA 5.7 and JETSET 7.4*. Comput. Phys. Commun., 82:74-90, 1994.

[17] Parton Distribution Functions:
http://durpdg.dur.ac.uk/hepdata/pdf.htm

[18] Johannes Blumlein, Helmut Bottcher and Alberto Guffanti.
*Non-singlet QCD analysis of deep inelastic world data at O(alpha(s)**3)*, 2006.

[19] R. Brun, F. Bruyant, M. Marie, A.C. McPherson and P. Zanarini.
Geant3. CERN-DD/EE/84-1.

[20] S.Agostinelli et al. *GEANT4: A simulation toolkit*.
Nucl. Instrum. Meth., A506:250-303, 2003.

[21] The ATLAS DC1 Task Force. *A step towards a computing Grid for the LHC experiments: ATLAS Data Challenge 1*. CERN-PH-EP/2004-028

[22] S. Campana et al. *Analysis of the ATLAS Rome Production experience on the LCG Computing Grid*. IEEE Computer Society Press, USA, e-Science05 p. 82-89, December 2005 [10]

[23] DICE (ATLAS Geometry Description & the Detector Digitization):
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/DOCUMENTS/simulation.html

[24] ATRECON (ATLAS Reconstruction program):
http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/DOCUMENTS/reconstruction.html

[25] The Athena Framework:
https://twiki.cern.ch/twiki/bin/view/Atlas/WorkBookAthenaFramework

[26] PAW Home Page:
http://paw.web.cern.ch/paw/

[27] R. Brun et al. *ROOT - An Object-Oriented Data Analysis Framework*:
http://root.cern.ch

[28] ATLAS Event Data Model:
https://twiki.cern.ch/twiki/bin/view/Atlas/EventDataModel

[29] D. Duellmann et al. *The LCG POOL Project Overview*:
http://pool.cern.ch/papers/pool-overview-chep2003.pdf

[30] ATLFAST (ATLAS Fast Simulation package):
http://www.hep.ucl.ac.uk/atlas/atlfast/

[31] LHC Computing Grid (LCG) project:
http://lcg.web.cern.ch/LCG/

[32] LCG Technical Design Report (TDR), LCG-TDR-001, CERN-LHCC-2005-024:
http://lcg.web.cern.ch/LCG/tdr/

[33] Memorandum of understanding for Collaboration in the Deployment and Exploitation of the WLCG:
http://lcg.web.cern.ch/lcg/C-RRB/MoU/WLCGMoU.pdf

[34] EGEE project:
http://www.cern.ch/egee

[35] Open Science Grid (OSG) project:
http://www.opensciencegrid.org/
The Grid 2003 Project:
http://www.ivdgl.org/grid2003/index.php

[36] Nordic Data Grid Facility (NDGF):
http://www.ndgf.org/
NorduGrid project:
http://www.nordugrid.org

[37] ATLAS Computing Model:
http://www.gridpp.ac.uk/eb/ComputingModels/atlas_computing_model.pdf

[38] Virtual Organization (VO):
http://en.wikipedia.org/wiki/Virtual_organization

[39] R. Clift, A. Poppleton. *IpatRec: inner detector pattern-recognition and track-fitting*. ATLAS note, ATL-SOFT-94-009, 1994.

[40] I. Gavrilenko. *Description of Global Pattern Recognition Program (XKalman)*. ATLAS note, ATL-INDET-97-165, 1997.

[41] J.E. García, S. González, E. Ros, M. Vos.
The b-tagging performance of the complete ATLAS DC1 layout using WH events. ATLAS internal documentation, ATL-COM-INDET-2003-017, 2003.

[42] S. Corread, V. Kostyukhin, J. Lêvêque, A. Rozanov, J.B. De Vivie de Régie. *b-tagging with DC1 data*. ATLAS note, ATL-PHYS-2004-006, 2004.

[43] K. Hagiwara et al. Phys. Rev. **D** 66 (2002) 010001

[44] M. Vos. *The ATLAS inner tracker and the detection of light super-symmetric Higgs bosons*. PhD thesis, CERN-THESIS-2003-028.

[45] V. Kostyokhin. *Seconday vertex b-tagging*. ATL-PHYS-2003-032.

[46] G. Kane. *Supersymmetry*. Perseus Publishing, 1997.

[47] N. Arkani-Hamed, A.G. Cohen, H. Georgi.
JHEP 0207, 020 (2002) and Phys. Lett. **B 513**, 232 (2001).

[48] N. Arkani-Hamed, A.G. Cohen, E. Katz and A.E. Nelson. JHEP **0207** (2002) 034, hep-ph/0206021.

[49] I. Low, W. Skiba and D. Smith. Phys. Rev. D **66** (2002) 072001, hep-ph/0207243.

[50] N. Arkani-Hamed, A.G. Cohen, E. Katz, A.E. Nelson, T. Gregoire and J.G. Wacker. JHEP **0208** (2002) 021, hep-ph/0206020.

[51] T. Gregoire and J.G. Wacker. JHEP **0208** (2002) 019, hep-ph/0206023.

[52] T. Han, H.E. Logan, B. McElrath and L.T. Wang.
Phenomenology of the little Higgs model. Phys. Rev. D **67** (2003) 095004, hep-ph/0301040.

[53] S. Chang and J.G. Wacker. Phys. Rev. D **69** (2004) 035002, hep-ph/0303001.

[54] R. Casalbuoni, A. Deandrea and M. Oertel. JHEP **0402** (2004) 032, hep-ph/0311038.

[55] G. Azuelos et al. *Exploring Little Higgs models with ATLAS at the LHC*. Eur.Phys.J., C39S2:13, 2005, hep-ph/0402037.

[56] T. Sjöstrand, S. Mrenna, P. Skands. *PYTHIA: 6.4 Physics and Manual*, hep-ph/0603175 (http://www.thep.lu.se/ torbjorn/Pythia.html)

[57] E. Richter-Was, D. Froidevaux, L. Poggioli.
ATLFAST: 2.0 a fast simulation package for ATLAS. ATL-PHYS-98-131.

[58] M. L. Mangano et al. *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, hep-ph/0206293.

[59] Z. Sullivan. *Understanding single-top-quark production and jets at hadron colliders*. Phys. Rev. **D 70** (2004) 114012.

[60] ATLAS TDR. *Detector Physics Performance Technical design Report, Vol. II*. CERN/LHCC/99-15 (1999).

[61] M. Vos et al. *The b-tagging performance of the complete ATLAS DC1 layout using WH events*. ATL-COM-INDET-2003-017.

[62] ATLAS Production System (ProdSys):
https://twiki.cern.ch/twiki/bin/view/Atlas/ProdSys

[63] ATLAS Distributed Data Management:
https://twiki.cern.ch/twiki/bin/view/Atlas/DistributedDataManagement

[64] S. González de la Hoz et al.
ATLAS Data Challenge 2: A massive Monte Carlo production on the Grid.
ATL-SOFT-PUB-2005-001 and Springer Lecture Notes in Computer Science, LNCS 3470, pp. 30-39, February 2005 [45].

[65] M. Branco, D. Cameron, T. Weanus.
A Scalable Distributed Data Management System for ATLAS.
Computing in High Energy and Nuclear Physics 2006 (CHEP'06).
Mumbai (India), February 2006.

[66] LCG File Catalog (LFC):
http://www.gridpp.ac.uk/wiki/LCG_File_Catalog

[67] A. Chervenak, E. Deelman, I. Foster et al.
Giggle: A Framework for constructing Scalable Replica Location Services.
Proceedings of SuperComputing 2002, Baltimore (USA), 2002

[68] Local Replica Catalog (LRC):
http://www.gridpp.ac.uk/wiki/Local_Replica_Catalog

[69] DQ2 client and end-users:
https://twiki.cern.ch/twiki/bin/view/Atlas/UsingDQ2

[70] PACMAN:
http://physics.bu.edu/ youssef/pacman/index.html

[71] G. Van Rossum and F.L. Drake, Jr (eds.).
Python Reference Manual, Release 2.4.3 (Python Software Foundation 2006).
http://www.python.org

[72] Athena job options (jobOptions):
https://twiki.cern.ch/twiki/bin/view/Atlas/JobOptions

[73] XML:
http://en.wikipedia.org/wiki/XML

[74] Oracle DataBase:
http://en.wikipedia.org/wiki/Oracle_database

[75] XMMP,
http://en.wikipedia.org/wiki/Extensible_Messaging_and_Presence_Protocol

[76] SOAP:
http://en.wikipedia.org/wiki/SOAP

[77] Windmill project:
http://heppc12.uta.edu/windmill/

[78] Eowyn - Supervisor component for ATLAS ProdSys:
https://twiki.cern.ch/twiki/bin/view/Atlas/EowynSupervisor

[79] D. Rebatto.
The LCG-2 executor for the ATLAS DC2 Production System.
Proceedings of CHEP04, Interlaken (Switzerland), 2004.

[80] CondorG executor for ATLAS ProdSys:
https://twiki.cern.ch/twiki/bin/view/Atlas/CondorG

[81] R. Sturrock et al. *Performance of the NorduGrid ARC and the Dulcinea executor in ATLAS DC2.* Proceedings of CHEP04, Interlaken (Switzerland), 2004.

[82] M. Mambelli et al.
ATLAS Data Challenge Production on Grid3.
Proceedings of CHEP04, Interlaken (Switzerland), 2004.

[83] Simplified Wrapper and Interface Generator (SWIG):
http://www.swig.org/

[84] The European DataGrid project:
http://eu-datagrid.web.cern.ch/

[85] Grid User Interface (UI):
http://www.gridpp.ac.uk/deployment/users/ui.html

[86] Job Description Language (JDL):
http://trinity.datamat.it/projects/EGEE/wiki/wiki.php?n=JDL.AboutJDL

[87] ClassAD language:
http://www.cs.wisc.edu/condor/classad/

[88] Network file system (NFS):
http://en.wikipedia.org/wiki/Network_File_System

[89] Workload Management System (WMS):
http://glite.web.cern.ch/glite/wms/

[90] J. Frey, T. Tannenbaum, I. Foster, M. Livny and S. Tuecke.
Condor-G: A Computation Management Agen for Multi-Institutional Grids.
HPDC10, San Francisco (USA), 2001, http://www.cs.wisc.edu/condor/condorg

[91] Berkeley Database Information Index (BDII):
https://twiki.cern.ch/twiki//bin/view/EGEE/BDII

[92] Condor High Throughput Computing project:
http://www.cs.wisc.edu/condor/

[93] gLite EGEE middleware:
http://glite.web.cern.ch/glite/

[94] Logging and book-keeping (L&B):
http://glite.web.cern.ch/glite/lb/

[95] Grid Resource Allocation Management (GRAM):
http://www.birds-eye.net/definition/acronym.cgi?what+is+GRAM=
Grid+Resource+Allocation+Management&id=1172439013

[96] Globus Toolkit:
http://www-unix.globus.org/toolkit/

[97] Extended Resource Specification Language (XRSL):
http://progress.psnc.pl/English/architektura/xrsl.html

[98] Virtual Data Toolkit (VDT):
http://vdt.cs.wisc.edu/

[99] Globally Unique Identifier (GUID):
http://en.wikipedia.org/wiki/Globally_Unique_Identifier

[100] ATLAS Computing Technical Technical Design Report, CERN-LHCC-2005-022, http://atlas-proj-computing-tdr.web.cern.ch/atlas-proj-computing-tdr/PDF/Computing-TDR-final-July04.pdf

[101] ATLAS Collaboration, ATLAS High-Level Trigger, *Data Acquisition and Controls Technical Design Report*. CERN-LHCC-2003-022 (2003)

[102] Data TransAtlantic Grid:
http://datatag.web.cern.ch/datatag

[103] Site Functional Tests (SFTs):
http://sft-fzk.gridka.de:8080/sft/php/sftestcases.html

[104] LCG commands (lcg-utils):
http://www.gridpp.ac.uk/wiki/LCG_Utils

[105] Storage Resource Management Middleware Project:
http://sdm.lbl.gov/srm/

[106] Grupo de Ordenadores para el Grid (GoG):
http://ific.uv.es/grid/gog

[107] ATLAS Database and Data Management Project:
http://atlas.web.cern.ch/Atlas/GROUPS/DATABASE/project/ddm

[108] Tag For Event Selection:
https://twiki.cern.ch/twiki/bin/view/Atlas/TagForEventSelection

[109] Job Priorities:
https://twiki.cern.ch/twiki/bin/view/Atlas/JobPriorities

[110] T. Weanus, K. De et al. *PanDA - Production and Distributed Analysis*:
http://twiki.cern.ch/twiki/bin/view/Atlas/Panda

[111] T. Maeno. *Distributed Analysis on Panda*:
http://uimon.cern.ch/twiki/bin/view/Atlas/DAonPanda

[112] The ATLAS Commander (AtCom):
http://uimon.cern.ch/twiki.cern.ch/twiki/bin/view/Atlas/AtCom

[113] GANGA:
http://ganga.web.cern.ch/ganga/

[114] P. Mato. *GAUDI - Architecture design document, LHCb-98-064 (1998)*:
http://proj-gaudi.web.cern.ch/proj-gaudi/welcome.html
http://atlas-computing.web.cern.ch/atlas-computing/packages/athenaCore.php

[115] DIANE:
http://it-proj-diane.web.cern.ch/it.proj-diane/

[116] BaBar:
http://www-public.slac.stanford.edu/babar/

[117] NA48:
http://na48.web.cern.ch/NA48/

[118] PhenoGrid:
http://www.phenogrid.dur.ac.uk/

[119] COMPASS:
http://wwwcompass.cern.ch/

[120] D. Liko et al. *The ATLAS strategy for Distributed Analysis in several Grid infraestructures*. Proc. 2006, Conference for Computing in High Energy and Nuclear Physics (Mumbai, India, 2006).

[121] U. Egede et al. *Experience with distributed anlysis in LHCb*. Conference for Computing in High Energy and Nuclear Physics (Mumbai, India, 2006).

[122] iPython:
http://ipython.scipy.org

[123] PyQT:
http://www.riverbankcomputing.co.uk/pyqt

[124] CASTOR - CERN Advanced STORage manager:
http://castor.web.cern.ch/castor/

[125] Structured Athena Aware Ntuple:
https://twiki.cern.ch/twiki/bin/view/Atlas/StructuredAthenaAwareNtuple

[126] HighPtView:
https://twiki.cern.ch/twiki/bin/view/Atlas/HighPtView

[127] Atlantis event display for ATLAS:
http://www.hep.ucl.ac.uk/atlas/atlantis/

[128] Andrew File System (AFS):
http://en.wikipedia.org/wiki/Andrew_file_system

[129] PLUS (Public Login User Service) - LXPLUS:
http://plus.web.cern.ch/plus/

[130] PROOF: the Parallel ROOT Facility:
http://root.cern.ch/twiki/bin/view/ROOT/PROOF

[131] Lustre Cluster File System:
http://wiki.lustre.org

[132] Remote File Input/Output access:
http://hikwww2.fzk.de/hik/orga/ges/infiniband/rfioib.html