



LUND
UNIVERSITY

LUND UNIVERSITY
FACULTY OF SCIENCES
DIVISION OF PARTICLE PHYSICS

Impact of different primary particles on Geant4 simulation execution time

Study on *protons* and *pions*

Author

Jessica Rebecca MADEIRA

Supervisors

Oxana SMIRNOVA
Caterina MARCON

January 25, 2021

Abstract

The High Luminosity Large Hadron Collider (HL-LHC) at CERN, an improvement to the LHC, scheduled to be operating in 2027, will further the potential for scientific breakthroughs in High Energy Physics. The project's demand in computational resources, however, is predicted to exceed what will be available then. Thus, efficiency optimization of Monte Carlo simulations of the collisions that occur in the LHC experiments, which today occupy 40% of those resources and are pivotal for the correct interpretation of the collected data, is the aim of recent research and development efforts.

This study focuses on the ATLAS experiment, specifically, and makes use of the Geant4 simulation software and its extensive libraries, used at LHC to simulate the passage of particles through matter. It aims at providing useful data for future full scale studies on software time response improvement. Parameters such as the type of build method - static or dynamic - and version of the GCC compiler, from recent studies, have been shown to have a considerable impact in reducing the execution time of the simulations. This research carries forward this analysis and studies the impact of different primary particles, which are created in the pp collisions, on the simulation time, specifically the pions, $\pi^{+/-}$, and protons, p . The software's virtual particles, geantino and charged geantino, were also studied.

A simulation benchmark was used, with a simplified version of the ATLAS detector, and was run through the Aurora cluster at Lund University. This was carried out for 10 and 20 GeV particles, using both static and dynamically compiled libraries. The statically compiled simulations were confirmed to decrease time by 10%, as was foreseeable. In addition, all considered particles exhibit simulation time distributions which agree with what would be expected from theory. The virtual particles confirm the large contribution that the simulation of interactions in the detector has on the execution time. Moreover, both the negative and positive pions registered a mean execution time about 4% smaller than the proton's, in agreement with the pions' smaller probability of interaction.

Contents

1	Introduction and Motivation	1
2	Background	3
2.1	Theoretical framework	3
2.1.1	The Standard Model	3
2.1.2	<i>Proton-proton</i> collisions	4
2.1.3	Particles interaction with matter	4
2.2	ATLAS detector	8
2.2.1	Inner detector	8
2.2.2	Calorimeters	9
2.2.3	Muon detector	10
2.3	Monte Carlo method - Geant4	10
3	Methodology	12
4	Results	14
5	Discussion	19
6	Conclusion	22
A	Lunarc Aurora system	25
B	Simulation time plots at 10 GeV	26
C	Additional reference plots	29

Chapter 1

Introduction and Motivation

Since the 1970's, with the development of the Standard Model which provided an accurate description of the fundamental structure of matter, Particle Physics has been successfully dedicated to studying the most elemental particles and the forces that act upon them. However, to date, all observed particles make up only 4% of the known Universe. Advancing into unprecedented higher energy regions could produce scientific breakthroughs in the discipline's most intricate unresolved phenomena, such as the existence of dark matter and energy. It was in search of answers to such occurrences and to test theory predictions that the Large Hadron Collider, LHC, and its various detectors were built at CERN, [Figure 1.1](#). In this two-ring-superconducting particle accelerator, bunches of protons are designed to collide with center-of-mass energy $\sqrt{s} = 14$ TeV at a luminosity¹ of 10^{34} cm⁻² s⁻¹. In addition to protons, the LHC is also operated with heavy ion beams at a lower peak luminosity. The collision products are tracked as they travel through the detectors' different materials [\[2\]](#).

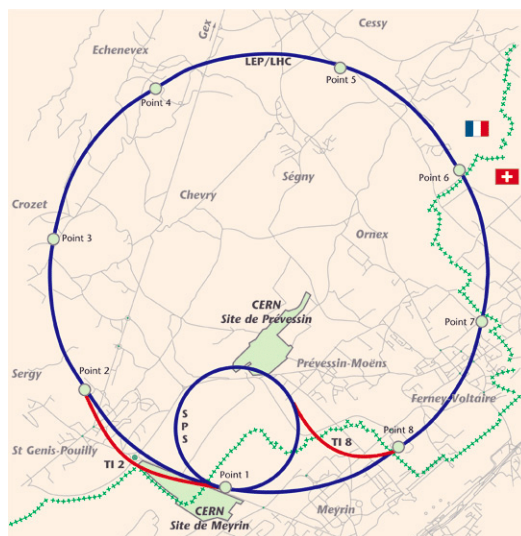


Figure 1.1: The geographical map of the LHC. It is situated 170 meters deep beneath the France–Switzerland border near Geneva, at CERN [\[3\]](#).

While the initial challenge was the build and design of all structures and experiments in

¹Luminosity measures the number of potential interactions per surface unit over a period of time the accelerator is able to produce [\[1\]](#).

the safest, most optimal and affordable way, to which simulation was indispensable, today, the biggest challenge is data handling. Simulation remains a powerful and crucial stage in all experiments at CERN: the degree of concordance between collected and simulated data confirms models and algorithms in output processing and event reconstruction, allowing for the correct interpretation of "real" data. However, considering the scarcity of the phenomena under investigation and in order to accurately study them, simulation models have to process very large amounts of computer-generated data. The ATLAS experiment [4], which built one of two high-granularity detectors designed to allow investigation over a wide variety of physics processes (Point 1 in Figure 1.1), consumes 40% of the computing resources available on detailed full detector simulations alone [5].

The particles produced in the collisions undergo multiple physical processes inside the detector: from hadronization, elastic collisions and decays of unstable particles to ionization of the detector's materials and inelastic nuclear collisions, which result in electromagnetic and hadronic showers. The electrical signal collected by the detector from the ionization of the material, the primary mechanism for energy loss at lower energies, is roughly proportional to the particle's energy deposition in the different components. At higher energies the cascade of secondary particles created when energetic particles collide with nuclei while travelling through the dense media in the detector is called a shower and is the dominant phenomenon. In the simulation process, after the system's structure is built and its materials defined, the particles are generated and transported through the different mediums. The resulting simulated energy distribution of a specific particle, which, ideally, should be close to the actual experiment's output, allows for its identification and analysis, as well as the discovery and study of new physics.

It was for the purpose of precisely simulating the collision products travelling in the LHC detectors that the Monte Carlo based simulation software Geant4, written in C++, was developed [6]. However, the LHC luminosity is increasing with upgrades. The High-Luminosity LHC, HL-LHC is projected to become operational in 2027. It will register collisions by an increased factor of 5 compared to today [7]. The projected computing requirements of the HL-LHC pose a challenge, not only financially but also regarding the extensive simulation time required for such high energies. Therefore, optimized and efficient Geant4 simulations are required to reduce the simulation's execution time and resolve the expected shortage in computing power, without forfeiting the quality of the data.

The quantitation of how different parameters of the simulation impact the execution time will, therefore, provide useful data for undergoing and future research and development efforts. Parameters such as the version of the C++ compiler and the type of build method used to link the extensive libraries provided by Geant4 to the executable of the simulation have been proven to impact considerably the execution time [8]. The study presented here extends on the scope of the parameters. Given that different particles will undergo specific interactions, can initiate showers of different structures and the detector will respond accordingly, studying simulated interactions of different primary particles is necessary to estimate their impact on the simulation time. Specifically, this investigation focuses on evaluating the differing impacts of the proton, p and the pion, $\pi^{+/-}$, as well as the instances where no particle interaction is detected and the particle is simply transported through the materials of the detector. The quantification was carried out for two particle energies: 10 and 20 GeV, using a benchmark simulation software suite, to allow for a controlled study.

Chapter 2

Background

2.1 Theoretical framework

2.1.1 The Standard Model

The Standard Model theory, SM, provides a description of the 17 building blocks of nature and their antiparticles, acted on by three of the four fundamental forces, omitting gravity: the strong, weak and electromagnetic forces. The Higgs boson, the last discovered particle of the SM, is the mechanism through which all particles acquire mass [9].

Particles are divided into two fundamental categories: *fermions* and *bosons*, with half-integer and integer spin, respectively. Whereas fermions are the fundamental constituents of matter, bound by the aforementioned forces, bosons are the forces' mediators - *gauge bosons* - and are exchanged between fermions during interactions. Experimentally confirmed, the massless *photon*, charged $W^{-/+}$ and the neutral Z^0 *bosons*, and the *gluon* are the force mediating particles of the electromagnetic, the weak and the strong forces, respectively. The gravitational force is thought to also be mediated by a boson, called the *graviton*.

Fermions are further subdivided into *quarks* and *leptons*, with 6 flavours each. To each of the massive and charged leptons, $e^{-/+}$, $\mu^{-/+}$ and $\tau^{-/+}$, corresponds an almost massless and neutral lepton, the *neutrino*. Quarks carry a fractional electrical charge and are colored particles (color charge property is the strong interaction analog to charge in the electromagnetic force). Therefore, they interact via the strong force through the exchange of gluons, and can not be found alone in nature. A composite particle, termed *hadron*, can be either made up of 3 quarks, denoted as *baryon*, or a quark-antiquark pair, called *meson*. Accordingly, they have half-integer and integer spins, respectively.

The theory of Quantum Electrodynamics, *QED*, describes how the electromagnetic force, *EM*, acts on charged particles, whereas the relativistic description of the strong force and colored particles is provided by the theory of Quantum Chromodynamics, *QCD*. The weak force affects and couples together different flavoured fermions, and so it is important in decay processes. With the EM force, they are described by the Electroweak theory. A coupling strength, g , and its associated dimensionless constant, α , characterize the strength of each one of the interactions.

2.1.2 Proton-proton collisions

Inside the LHC, bunches of protons collide 40 million times per second, at a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, to account for the low probability of interaction processes [6]. The underlying process in pp collisions is the interaction of colored partons, quarks and gluons. As mentioned, at high energies, these are described by the QCD theory and the process can be divided into three stages [10]:




Figure 2.1: Illustration of a hadron-hadron collision. Incoming particles are colored in blue [11].

1.Hard scattering: Interaction of two quarks, described by perturbative QCD effects. Hard scattering interactions are drawn in red in Figure 2.1. It produces high-energy final state partons that create additional colored partons to obey confinement: jet production.

2.Underlying event: The interactions of the remaining partons in the colliding beams can also create further particles, drawn in purple and dominated by non-perturbative effects.

3.Hadronization: Confinement requires that partons form hadrons. The process is drawn in light green and the dark green circles represent hadronic decays. The final colorless particles from the collision travel through the detector, and almost all will interact with the material.

2.1.3 Particles interaction with matter

Stable – electron, proton and photon – and long-lived particles, such as muons, neutrons and charged pions, can be directly detected in high energy collider experiments (neutrinos are also stable but undetectable by ATLAS). On the other hand, the detection of unstable particles, such as neutral pions, is performed from their decay products ($\pi^0 \rightarrow \gamma\gamma$).

The average distance between collisions in a material with N atomic density, denoted mean free path λ , of a particle with energy E , is defined as [12]:

$$\lambda = \frac{1}{N\sigma} \quad (2.1)$$

where $\sigma(E, Z, A)$ is the cross section of a specific process, defined as the effective area for a collision. It provides a measurement of the probability that a particle will interact in a material of atomic and mass numbers Z and A , respectively.

The total probability, $P_{int}(x)$, that a particle interacts up to a path length x is defined by the cumulative distribution function [13]:

$$\int_0^x P_{int}(x')dx' = 1 - \exp\left(-\frac{x}{\lambda}\right) \quad (2.2)$$

It is the nature of the particles' different interactions in matter that allows for their detection, identification and differentiation.

1. Interactions of charged particles




Figure 2.2: Plot of the muon's stopping powers with respect to the particle's momentum [14].

All charged particles travelling across a material experience the Coulomb force and interact electromagnetically with atomic electrons and nuclei. The dominant mechanism for energy loss is ionization of the medium. In the ionization process, the incident particle collides with the atomic electrons and has sufficient energy to remove them from the atom. These energetic electrons perform additional ionizations and are called δ -rays.

The ionization process is dominant up to around 10 MeV and 100 GeV for electrons (positrons) and muons, respectively [15]. For heavier charged hadrons, the energy at which dominant losses become radiative, can be even higher. The electronic stopping power is given by the Bethe-Bloch equation below, where material dependency is low since Z/A is roughly constant:

$$\frac{1}{\rho} \frac{dE}{dx} \approx - \frac{4\pi \hbar^2 c^2 \alpha^2 Z}{m_e v^2 m_u A} \left[\ln \frac{2m_e c^2 \beta^2 \gamma^2}{I_e} - \beta^2 \right], \quad (2.3)$$

where $m_u = 1.66 \times 10^{-27}$ kg is the atomic mass number, ρ is the material density, m_e is the electron's rest mass, c is the speed of light, $\beta = v/c$ is the relativistic velocity of the particle, α the dimensionless coupling strength and I_e the effective ionization potential of the material, approximated to $10Z$ eV. For highly relativistic particles, where $v \approx c$, energy loss depends logarithmically on $(\beta\gamma)^2$. This explains the rate's slow rise region in Bethe-Bloch plots, *e.g.* see Figure 2.2.

Particles with a velocity corresponding to the plot's minima, $\beta \approx 3$, are called minimum ionizing particles, MIPS. The higher the velocity of the particle the larger its transverse electric field, and therefore, ionization energy loss, will be. On the other hand, once the particle becomes too slow to ionize the material, it will lose its energy in atomic collisions. The consequent thermal excitation of the atoms heats the material.

In addition to the Coulomb interactions with electrons, particles also undergo Coulomb interactions with nuclei, resulting in the deflection of the particle's path in a material, through multiple elastic scatterings. This non-ionizing energy loss can lead to a displacement of the atoms in the lattice of a material, however, it is typically much smaller than the electronic energy loss [15]. At higher energies, depending on the particle type, other energy-loss mechanisms, such as radiative processes, are present. Hadrons also undergo inelastic nuclear collisions.

Energy deposition

Hadronic and EM showers are created when energetic hadrons and electrons and photons, respectively, travel through the high- Z materials of the detector. As mentioned, energetic δ -rays produced when the particles ionize the medium can transport energy away from the shower volume. Therefore, rather than characterizing showers through the particle's energy loss (Figure 2.3(b)), the shower development is more accurately described by the energy deposition in a medium, caused by Coulomb atomic and nuclear interactions of the charged particles, since total energy deposition is mostly constant (Figure 2.5) [15].

For the longitudinal component of the energy deposition, $\epsilon(z_i)$ it is assumed that a pencil beam impacts with an infinitely long block of material:

$$\epsilon(z_i) = \frac{1}{E_0} \frac{\Delta E(z_i)}{\Delta z} \quad (2.4)$$

where E_0 is the particle's energy, $\Delta E(z_i)$ is the energy deposited in the layer $[z_i, z_i + \Delta z]$. $\Delta z \rightarrow X_0$ and $\Delta z \rightarrow \lambda_{inel}$, in the high energy regime, are the characteristic lengths of EM and hadronic showers, called radiation length, X_0 and inelastic scattering length, λ_{inel} , respectively.

2. Interactions of electrons and photons




Figure 2.3: (a) Plot of photon cross sections with respect to the particle's energy, in lead. (b) Energy loss, per radiation length, as function of the electron's energy [14].

From Figure 2.3(b), at energies higher than the critical energy, E_c , electrons will mainly lose energy through bremsstrahlung. This phenomenon is related to the production of

electromagnetic radiation by the deflection of charged particles, but it is mostly exclusive to electrons, since these are abundant and easily deflected [15]:

$$-\left(\frac{dE}{dx}\right)_{\text{bremsstrahlung}} \propto \frac{E}{X_0} \quad (2.5)$$

where X_0 , radiation length, is defined as the distance traveled after which the electron loses energy through bremsstrahlung by a factor of $1/e$. It is approximated by the expression [9]:

$$X_0 \approx \frac{1}{4\alpha n Z^2 r_e^2 \ln(287/Z^{1/2})} \quad \text{where } r_e = \frac{e^2}{4\pi\epsilon_0 m_e c^2} = 2.8 \times 10^{-15} \text{ m}, \quad (2.6)$$

and r_e is defined as the classical radius of the electron and n is the number density of the nuclei.

For photons at energies higher than about 10 MeV (the value depends on material), the interactions are dominated by electron pair production, see Figure 2.3(a). At energies in the GeV range, a bremsstrahlung photon produces an electron pair, and vice-versa. This creates cascades of photons, electrons and positrons, called *electromagnetic showers*. The photon's constant absorption length at high energies is directly related to the radiation length, $\lambda_{abs} \approx \frac{9}{7}X_0$ [15].

Energy deposition plots from simulations provide an accurate description of the shower's development, considering the different nature of the initial interactions of electron and photon induced showers. The position of the longitudinal shower maximum, t_{max} , from a beam of energy E , is predicted through the formula (see Figure 2.5(top)):

$$t_{max} = \log \frac{E}{E_c} - 1.0 \quad (0.5 \text{ for } \gamma \text{ induced shower}) \quad (2.7)$$

The transverse component is considered to be independent of the beam energy and caused by multiple scattering of electrons, well described by the Moliere radius:

$$R_M \approx \frac{21 \text{ MeV}}{E_c} X_0 \quad (2.8)$$

For heavier particles, such as the muon, the energy at which the radiative stopping power becomes important will be higher, given that the rate of production of EM radiation is inversely proportional to the square of the particle's mass [9].

3. Interactions of hadrons

Charged hadrons are described by the Bethe-Bloch equation in the low energy range, given that the primary mechanism for the energy loss is ionization. Additionally, as stated earlier, both charged and neutral hadrons also undergo strong inelastic nuclear interactions at higher energies, GeV, see Figure 2.4.

The primary inelastic hadronic interaction produces forwards-directed collision products which interact further in the medium, originating cascades of particles that increase logarithmically with energy, *hadronic showers*. In analogy to the radiation length of EM




Figure 2.4: Proton-proton cross sections. Towards higher energies, inelastic cross sections are dominant over elastic ones [15].

showers, the longitudinal component of the shower is characterized by the inelastic scattering length, $\lambda_{inel} \propto \frac{A^{1/3}}{\rho}$. It corresponds to the average distance between interactions and is significantly larger than the radiation length [9]. Essentially, all neutral pions produced in these interactions immediately decay through the weak interaction into two photons. In other words, hadronic showers have an EM component. Given that the pion is the lightest hadron, the showers stop when the particle's energy falls below the pion threshold.

On average, 30% of incident energy is lost in nuclear excitation, evaporation and fission processes [9]. These processes produce lower energy particles – neutrons, protons, α and γ rays – which, together with neutrinos, carry energy out of the shower. Protons and the α particles from radioactive decays will range out through ionization, whereas neutrons, as relatively stable particles, will scatter elastically until they are thermally captured. Neutron capture in heavier materials induces fission processes which emit γ rays [15].

2.2 ATLAS detector

The ATLAS detector is divided into three main components: the inner detector, ID, the EM and hadronic calorimeters and the muon spectrometer, Figure 2.6.

2.2.1 Inner detector

In the ID, immersed in a $B = 2$ T solenoidal field, charged particles' positions and momenta are detected and measured. It is also here that electrons are identified. All its components rely on the ionization of an active medium to reconstruct the particle's trajectory, where the ionization will be proportional to the electrical signal output. The ID is comprised of [4]:

- **Silicon semiconductor pixel and strip detectors, SCT:** The pixel detectors are composed of thin layers of silicon connected to lower layers of electronics through solder spheres, while the latter has strips as the p-type material and the detector's bulk is n-type silicon. The reverse bias applied across the semiconductors results in a larger depletion layer and potential barrier. The collection of holes from ionization in the p-type materials provides a clear signal output [9].

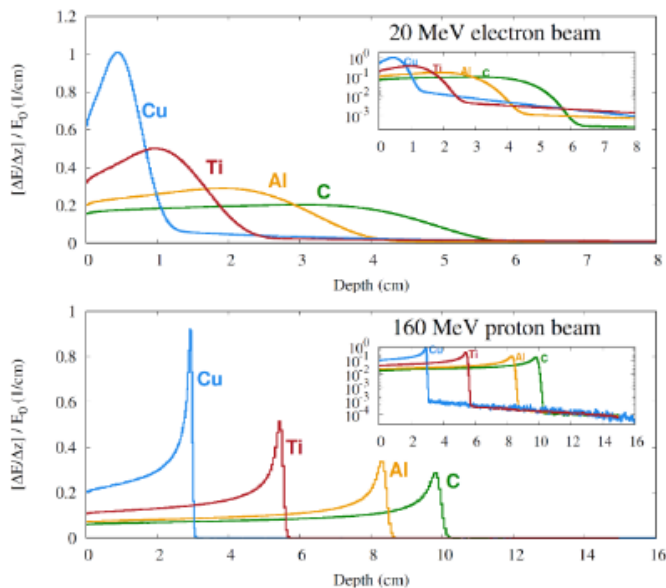


Figure 2.5: Longitudinal energy deposition profiles of electron beams (top) and proton beams (bottom) in different materials, defined by Eq.2.4. The broader peak of electrons is due to the Coulomb scattering of the electrons. For protons and α particles, the peak is more distinct - Bragg peak - since it also depends on nuclear elastic and inelastic scattering processes. The tails, more pronounced in heavier materials, are due to bremsstrahlung photons and neutrons, respectively [15].

- **Transition radiation tracker, TRT:** Straw tubes filled with a Xenon based gas mixture and electrode gold-plated tungsten wires in the center to collect the signal from the ionization. The characteristic emission of transition radiation in the X-Ray region caused by the passage of electrons between the straws and a radiator material, with different dielectric constants, allows for this particle's identification, since these X-rays perform additional ionizations [16].

The solenoidal magnet supplies an approximately uniform magnetic field in the direction of the beam-axis. The Lorentz force provided by the solenoid competes with the centripetal force, such that the transverse momentum, $p_T [GeV/c]$, of the particles can be calculated from [9]:

$$p_T = 0.3BR \quad (2.9)$$

where $R[m]$ is the radius of the track's curvature.

2.2.2 Calorimeters

Surrounding the inner detector are EM and hadronic calorimeters. They measure the energy of the interacting particles through their absorption by a high-Z material. It is in the calorimeters that EM and hadronic showers are detected and evaluated [6]:

- **Electromagnetic calorimeter:** Layers of lead plates as the high-Z absorber material and liquid argon (LAr) as the active medium. The ionization of the active layer by the electrons from the showers is measured by readout electrode circuits.




Figure 2.6: Illustration of the cross sectional view of the ATLAS detector, taken from Ref.[4].

- **Hadronic calorimeter** Composed of a tile calorimeter, and LAr hadronic end-cap and forward calorimeters. In the LAr calorimeters, copper and tungsten are the absorber materials. Layers of scintillator sheets and steel constitute the tile calorimeter. In scintillator materials, the ionization energy loss of the charged particles results in the emission of light (UV) [9].

2.2.3 Muon detector

The muon spectrometer makes use of the magnetic fields provided by the toroidal magnets to measure the momenta of muons, which mostly traverse the ID and the calorimeters undetected [17]. In muon drift tubes, the particle is tracked through the drift time of the electron from the ionization of a gas mixture towards high voltage electrode wires at the center of the tubes. Cathode strip chambers provide additional coordinate measurements, and thin gap chambers and resistive plate chambers are used as trigger chambers.

2.3 Monte Carlo method - Geant4

All the interactions mentioned above follow a random sequence and have probabilistic outcomes, thus it is imperative that the Geant4's simulation of particles through matter follows a statistical method. The Monte Carlo algorithms make use of statistical samplings with pseudo-randomised starting conditions or events: the sequence properties approximate properties of random sequences [18]. This is the technique adopted by the simulation software.

The Geant4 (Geant stands for *Geometry ANd Tracking*) has extensive libraries and requires users to follow organisational principles and provide a description of: the detector's

geometry and its materials, the physics processes assigned to each particle and the particle's generator and kinematics. The user can also implement additional class actions in order to have more control over the simulation's data structure [6]. Additionally, the software provides two virtual particles: the geantino and the charged geantino, massless and non-interacting particles. These are useful for geometry and tracking diagnostics.

The ATLAS collaboration makes use of predefined reference physics lists, which include different combinations of EM and hadronic physics. Other packages can be introduced within the guidelines of the framework; for example, a GDML¹ library for geometry description or the data analysis software ROOT [20][21].

During a simulation, the following objects are considered: *track* represents the state of the particle for a given instance, and is propagated and updated in a step-by-step way; *step* provides information regarding the change in the state of the particle and ends with an interaction or when the particle leaves the volume. Furthermore, an *event* is collection of tracks and *run* is a collection of events.

The Monte Carlo algorithm generates an interaction from the cumulative distribution function (CDF), Equation 2.2, by setting [13]:

$$\int P_{int}(x)dx = \eta \Leftrightarrow \frac{x}{\lambda} = -\ln(1 - \eta) \quad (2.10)$$

where $\eta \in [0, 1]$ is a pseudo random number. The material independent quantity $\frac{x}{\lambda}$ is defined as the Number of Mean Free Path, NMFP.

At the beginning of a step in the particle transportation process, the NMFP for each interaction associated to the particle is sampled and converted to a physical length from the cross section. The process with the smallest value will determine the step length. The particle is then transported, and, after an interaction, sampling of NMFPs continues until the particle disappears.

¹Geometry description language which provides an application-independent way of implementing new detector geometries, and also offers an exchange format for already existing ones. GDML uses a geometrical hierarchy to define the detector volumes, in a similar manner to the Geant4 software[19].

Chapter 3

Methodology

The study was carried out for individual particles at fixed energies and not complex collision events. To achieve this, randomization of primary species was turned off at runtime and each considered primary particle was monoenergetic and transported individually, with no background data to take into account. Given that pions are the lightest hadrons, these particles are created in abundance in the pp collisions at the LHC (see Figure 31 of Ref.[22]). On that basis, π^+ and π^- were chosen as primary particles to be simulated. In addition, since the proton is considerably larger in mass, and pion and proton induced showers differ in their development and composition, for comparison purposes, proton interactions were simulated as well. Lastly, geantinos and charged geantinos, which do not interact with the material and undergo transportation processes only, were employed to simulate instances of pure propagation, where no particle-detector interaction was involved.

The simulations were carried out for all particles at two chosen energies, 10 and 20 GeV. Tasks were executed within interactive sessions on the Aurora computer cluster at Lund University. The cluster uses the resource manager and job scheduler SLURM to efficiently schedule the tasks [23] (see Appendix A). A standalone Geant4 simulation suite by A. Dotti [24] was used as benchmark, and the execution times collected from the output of the simulations were used to plot time distribution histograms for further analysis. The simulation time of a run was considered as the time interval between the sampling of the first primary particle and the end of the last event.

In the compilation process – C++ is a compiled language – the library resources utilized by the application were linked both *statically* and *dynamically* for all primary particles. The compilation approach has a significant impact on the execution time of the Geant4 simulations, as shown before [8]. Dynamic libraries, or shared libraries, exist independently of the program executables. The libraries are loaded during compilation and shared within the program [25], thus dynamically linked libraries take up less computer resources and do not require a complete rebuild after modifications to the code are performed. For this reason, this build type is the most commonly used. However, despite the advantage of smaller executables and easy upkeep, this method results in a larger execution time, as compared to linking the libraries statically instead: static libraries are compiled into the application and are not shared. Each statically compiled executable therefore contains a copy of the library, which means access to libraries is faster. Despite taking up more computer memory, the execution time of statically linked simulations is, thus, shorter. For

consistency with previous studies, the time distribution differences between build-types were also quantified and analyzed.

Version 8.2.0 of the GCC compiler, proven to reduce, on average, the simulation time up to 30% [8] with respect to older compiler versions, and the latest Geant4 version, 10.6.2, were used to perform the study. The default optimization level of the compiler, which does not affect execution time, was set to O2. Instead of hard-coded detector geometry definitions of the ATLAS detector, the simulation makes use of a simplified GDML geometry implementation file which includes: the ID (pixel and SCT silicon detectors and the TRT), the LAr hadronic and tile calorimeter (hadronic calorimeter) and the muon calorimeter. The EM calorimeter is not considered here.

The reference physics list used is FTFP_BERT, recommended for high energy physics simulations [26]. The list contains all standard EM processes and employs hadronic models in the low and high energy range. The choice of energy was based on Figure 35 of Ref.[22] (see Figure C.3 in Appendix C): a considerable fraction of hadrons around this energy is created after hadronization in the detector, and so the choice was a representative sample. Additionally, Figure 30 (see Figure C.4) plots the ratio of the energy deposited in the hadronic calorimeter and the momentum measured in the ID, E/p , a measurement of the calorimeter's response, for simulations using the aforementioned physics list. The response becomes independent of the hadron type at beam energy of around 20 GeV, considered the high energy range (compensating calorimeter). One can therefore assume the difference in the time distributions of different particles should not be influenced heavily by the manner in which Geant4 invokes the used libraries in the hadronic calorimeter.

The solenoidal magnetic field of the benchmark is set to 4 T, and the particle source is an isotropic point placed at the origin. For each of the particles and build type, the benchmark simulation was run 100 times and 20 times at particle-energy 20 and 10 GeV, respectively. Given time constraints and in order not to exhaust the available computer resources in detriment of other Aurora users' jobs, the number of runs for 10 GeV particles was reduced. Each run was composed of 5000 events each, and one node¹ was exclusively allocated for every run, so that resources would not be shared with other processes. Each node used 4 processor cores.

The simulation time of the runs was read from the Geant4 output files using a C++ program, written so that only the integer value of the time was considered from the proton and pion output files. For geantinos and charged geantinos, the time was collected with two decimal places, taking into account the faster simulation times. Histograms of both library build-types for each primary particle were plotted via a ROOT macro file, written to collect the times from the script's output files. The resulting time plots were fitted with a Gaussian probability distribution, compared and analysed.

¹A node is a single computer in the cluster, composed of CPU cores [27]

Chapter 4

Results

The plot fit analysis was done for the time distributions at particle energy 20 GeV, presented in this section, as well as for the simulations run at 10 GeV (see plots and fits in Appendix B). The collection of the fit parameters and calculated percentage differences between particles and build types, for both considered energies, are presented in Table 4.1.




Figure 4.1: Time distributions for negative pions at 20 GeV. Static build (left) and dynamic build (right).

Concerning negative pions at 20 GeV, the histogram plots of the simulations are presented in Figure 4.1: measurements with statically linked libraries produce a simulation time mean of 990 ± 14 s. This corresponds to a decrease of 9.4% relative to measurements taken with dynamically linked libraries: the average simulation time from the Gaussian fit is 1083 ± 19 s.

The normal distributions for 20 GeV positive pions, Figure 4.2, suggest a small decrease

in the simulation time with respect to the negative pions, of about 0.7% and 0.4% for the static and dynamic compilation, respectively. This increase is within the uncertainty of the distributions. Regarding the static case, the execution time averages at 983 ± 12 s, while for the dynamic simulations, the distribution mean grows 9.7% to 1079 ± 18 s.




Figure 4.2: Time distributions for positive pions at 20 GeV. Static build (left) and dynamic build (right).

From the Gaussian fit of the proton histograms, [Figure 4.3](#), the results show that the time distribution of a statically compiled simulation peaks at 1023 ± 12 s. With respect to positive and negative pions, this corresponds to a 4.0% and 3.3% increase, respectively. For dynamically compiled simulations, the execution time increases by 10.5%: the fit reaches its maximum at 1130 ± 14 s. Relative to dynamically linked simulations of the negative and positive pions, the mean simulation time increases by 4.3% and 4.7%, respectively.

Regarding pure propagation cases – geantinos and charged geantinos – at 20 GeV, the time distribution plots and respective fits are presented in [Figure 4.4](#) and [Figure 4.5](#). From the static case histogram fits, geantino simulations take on average 3.0 ± 0.1 s. For charged geantinos, this mean increases by 152% with respect to geantinos (7.5 ± 0.1 s). Regarding the dynamic cases, the average execution time for the geantinos is 3.1 ± 0.1 s. This percentage is a 155% decrease with respect to the average time for simulating charged geantinos (8.0 ± 0.1 s). The percentage of increase between the static and the dynamic cases is smaller, as compared to the $\sim 10\%$ for the tested hadrons: for charged geantinos this corresponds to 6.0% and 4.7% for geantinos.

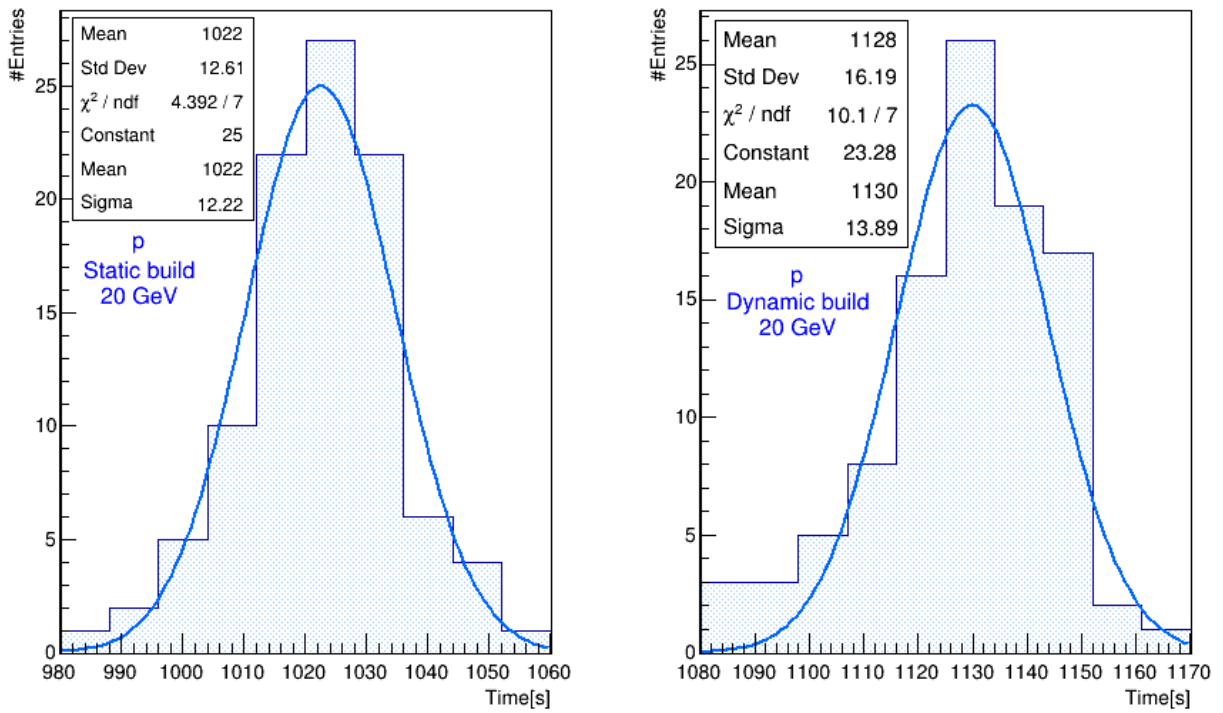


Figure 4.3: Time distributions for protons at 20 GeV. Static build (left) and dynamic build (right).




Figure 4.4: Time distributions for geantinos at 20 GeV: static build (left) and dynamic build (right).




Figure 4.5: Time distributions for charged geantinos at 20 GeV: static build (left) and dynamic build (right).

Particle type	Simulation time mean [s]	Percentage decrease w.r.t. protons [%]	Increase w.r.t static case [%]
Dynamic linking [20 GeV; 100 Runs]			
Proton, p	$1130 \pm 14(1.2\%)$		10.5
Positive pion, π^+	$1079 \pm 18(1.6\%)$	4.7	9.7
Negative pion, π^-	$1083 \pm 19(1.7\%)$	4.3	9.4
Geantino	$3.1 \pm 0.1(2.9\%)$	3.63×10^4	4.7
Charged geantino	$8.0 \pm 0.1(1.6\%)$	1.40×10^4	6.0
Static linking [20 GeV; 100 Runs]			
Proton, p	$1023 \pm 12(1.2\%)$		-
Positive pion, π^+	$983 \pm 12(1.2\%)$	4.0	-
Negative pion, π^-	$990 \pm 14(1.5\%)$	3.3	-
Geantino	$3.0 \pm 0.1(3.0\%)$	3.40×10^4	-
Charged geantino	$7.5 \pm 0.1(1.5\%)$	1.35×10^4	-
Dynamic linking [10 GeV; 20 Runs]			
Proton, p	$601 \pm 9(1.5\%)$		9.9
Positive pion, π^+	$577 \pm 5(0.8\%)$	4.2	9.6
Negative pion, π^-	$594 \pm 10(1.7\%)$	1.1	10.4
Geantino	$3.0 \pm 0.1(3.3\%)$	1.99×10^4	5.6
Charged geantino	$8.0 \pm 0.1(1.8\%)$	7.41×10^3	5.0
Static linking [10 GeV; 20 Runs]			
Proton, p	$546 \pm 6(1.1\%)$		-
Positive pion, π^+	$526 \pm 4(0.7\%)$	3.9	-
Negative pion, π^-	$538 \pm 8(1.5\%)$	1.5	-
Geantino	$3.2 \pm 0.1(2.0\%)$	1.70×10^4	-
Charged geantino	$8.4 \pm 0.1(0.7\%)$	6.40×10^3	-

Table 4.1: The fitted parameters from the measurements of all considered particles, for both static and dynamically linked simulations at 10 and 20 GeV. The yellow cells are the reference value for the analysis of the % decrease of the simulation time between particle types, taken to be the proton's. The last column lists the % increase of the simulation time, evaluated with respect to the static case.

Chapter 5

Discussion

As mentioned before, the time response of the software will differ with respect to the primary particle under simulation, given the diverse reconstructed pseudo-random interactions it undergoes within the simulated detector. Protons and pions are charged hadrons and thus they will experience the Lorentz force from the superconducting solenoidal magnets of the system, and their trajectory is deflected according to their different momentum at the considered energies. Given that the particles have the same energy, E , but different rest masses, m_0 , their momentum p , and therefore, relativistic velocity, β , will differ. Additionally, both particles interact with the ID and are stopped in the hadronic calorimeter, undergoing similar processes, despite the varying outcomes caused by the different masses. Consequently, from theory, one can expect the time distributions of the protons to deviate from the pions' ($m_{0,p} > m_{0,\pi}$), and the observed differences in the simulated data are consistent with expectations. Extrapolating exactly from time distributions a quantified relation to the physics processes and how they are employed by the methods of the reference list for different particles is not possible. Nonetheless, possible reasons that could explain these differences are addressed in what follows.

Evaluating the Gaussian fits of the plotted histograms, the normalized chi-square, χ^2/ndf , which measures "goodness of fit", shows that the level of agreement is reasonable and thus the Gaussian functions are considered as the time distributions in this study.

In the tested pions, as expected from previous studies, the static builds reduce the execution time by about 10%. This percentage decreases for the geantino and the charged geantino. Tracking geantinos makes use of less libraries, since these do not interact with the material, and the simulation simply transports the particle through the detector. Therefore, it requires a smaller simulation time and the difference between build types is less prominent, going down roughly to half the value of hadrons. A charged geantino will, however, be deflected by the magnetic field, despite not interacting either. Library requirements and the track length therefore increase, which explains the larger execution time and subsequently the larger percentage difference between the build types that was observed.

In view of the simulation of these particles and the observed data, one can say that the transport and magnetic deflection contribute significantly less to the execution time than the simulation of interactions. Furthermore, as expected, no difference in the time distribution plots was observed between the two energies considered for both the geantino and

the charged geantino (see [Figure B.4](#) and [Figure B.5](#)): the simulation of particle interactions is responsible for the time variation observed for different energy hadrons.

Pions are mesons, made up of two quarks ($\pi^+[u\bar{d}]$ and $\pi^-[d\bar{u}]$), with a rest mass of $m_0 = 139.57 \text{ MeV}/c^2$. Besides being deflected by the magnetic field, the particle loses energy through ionization following Bethe-Bloch's formula, [Equation 2.3](#), both in the ID and the hadronic calorimeter components present in the geometry. Additionally, at such energies, pions will also undergo inelastic nuclear collisions, which result in the creation of showers in the copper and tungsten layers of the LAr hadronic end-caps and steel sheets of the tile calorimeter. From [Figure C.2](#), the cross sections for both the positive and the negative pions are around 25 mb, therefore one could expect the mean simulation time of these particles to be the same within the parameter's error, given that both total cross sections are similar. The small fluctuations between π^+ and π^- observed in the measurements, however, are consistent with the slightly higher cross section value of the negative particle. The discrepancy between these two particles is more prominent towards the lower energy region, due to the different processes they are involved in. From the results of the simulations run at 10 GeV, despite having less entries for each particle and build type, this contrast is also visible, with an increase of 2.4% and 3.1% for the static and the dynamic cases, respectively. In this case, the results are not within the uncertainty range. However, due to time constraints, an in-depth proper investigation of the underlying differences in low energy interactions of the negative and positive particles could not be carried out.

Regarding the proton, this baryon is composed of three quarks (uud), has a mass of $m_0 = 938.27 \text{ MeV}/c^2$, and, from [Figure 2.4](#), at projectile momentum in the laboratory reference frame 20 GeV, the total cross section for pp collisions is about 40 mb. Similarly to pions, protons also interact with atomic electrons through the Coulomb force, and collide inelastically with the nucleons of the material. This is coherent with the percentage increase of $\sim 10\%$ between build types that is observed, given that the same set of hadronic models are employed on both hadrons.

Given the particles' different masses, the response to the magnetic field in the detector system by the software could be a source of differences in time distributions. The mass-energy equivalence $E^2 = (\mathbf{p} \cdot c)^2 + E_0^2$, where $E_0 = m_0 \cdot c^2$, gives that for particles at $E = 20 \text{ GeV}$, the larger the mass, the smaller the particle's momentum will be, and from [Equation 2.2.1](#), if the transverse component of the quantity decreases, the radius of curvature will be smaller. However, calculating the ratio of momentum between protons and pions gives $p_p/p_\pi \cong 0.98$, due to their high energy. If this small difference in the trajectory's radius indeed contributed to a difference between the measured times, it was not significant and could not be distinguished or quantified in the experiment.

Based on theory, the development of pion and proton induced showers from the nuclear collisions in the calorimeter differs. Due to conservation laws, specifically, baryon number conservation, B ($B = 0$ for mesons and $B = 1$ for baryons), the first leading particle from a proton interaction will most likely be another baryon. For charged pions, this likelihood decreases to conserve B and the first leading particle will mostly be either a charged or neutral pion [\[28\]](#). The larger probability for the creation of unstable neutral pions induces a larger EM component, from the particle's decay into two photons, in pion-induced showers.

Following that $X_0 < \lambda_{inel}$ and since hadronic showers are more diverse in their devel-

opment given that there are more possible final states for particles, it could be expected that simulating EM shower components would consume less computer resources and therefore lead to a smaller execution time. However, as mentioned in methodology, based on the E/p response plot of the calorimeter to charged hadrons, the implementation of the hadronic models in this region of the detector for different particles is not considered a substantial source of the observed differences in the measurements. Essentially, the ionization loss for particle tracking in the material of the ID is the dominant contributor to the observed variations in simulation times, and the parameter under study.

From [Equation 2.2](#), the total probability of interaction is proportional to $\propto 1 - e^{-1/\lambda}$, where λ depends on the inverse of the cross section. As a result, for smaller mean free paths, the probability of interaction will be larger. This means that a proton interacts more with the medium it traverses. Also, based on the stopping power plots in [Figure C.1](#), the energy loss of the pion is higher than the proton's at particle momentum 20 GeV. Since the pions will lose their energy more rapidly, it is the extra ionization processes simulated for protons, mostly due to its higher interaction probability but also longer detection before absorption in the calorimeter, that are the primary cause for the increase in time that is observed.

Regarding the simulations run at 10 GeV, the consistency between what can be expected from theory and observations persists. The particles can undergo the same set of processes within the detector: ionization of the active materials of the detector's components and shower development in the absorber material of calorimeter. This is confirmed by the same increase of $\sim 10\%$ between the two build types used. The increased mean simulation time for protons is still observed, however the difference is less disparate for negative pions. This execution time increase agrees with the steeper increase towards lower energies in the cross section plot of the negative pion, [Figure C.2](#). With respect to [Figure 2.4](#), at this energy, both protons and pions will be about as likely to interact with a material as they would if energy was doubled to 20 GeV. Therefore, since they make use of the same library resources, the execution time could be expected to be larger accordingly. By comparing the times of both considered energies for all three hadrons, in fact, the energy decrease in half of the 20 GeV particles led to an execution time decrease of 53 – 54%.

Chapter 6

Conclusion

Currently, parameterized Monte Carlo simulation models are the focus of several optimization efforts. These are expected to replace Geant4 in several regions of the detector, in particular the EM and hadronic calorimeter regions, in order to meet most of the computing requirements for the HL-LHC. In addition, machine learning options are being explored and physics models and detector geometries tuned [5]. However, full detailed simulations can not be excluded from experiment, and, more importantly, the Geant4 software is indispensable to the tracking of the particle's trajectory in the inner detector component. Computationally efficient fully detailed simulations are therefore required.

This study provides a quantitative discussion on the time distributions observed for protons, positive and negative pions, and neutral and charged geantinos, and on the impact the particle type has on the execution time of a simulation. This is carried out using a simplified geometry of the ATLAS detector and a standalone benchmark simulation. All observed data is in accordance with what was theoretically predicted and assumed.

The ionization of the material, mostly in the inner detector, is the primary cause for the observed differences in the time distributions of the hadrons. Additionally, the dominant consumer of library resources is the simulation of the particle's interactions. The simulation time of the proton exhibits on average a 3 – 4% increase with respect to the pions (mean time for protons is at 1130 s and for pions this average is roughly 1081 s), in agreement with the particle's higher cross section. The time distribution variation between both considered pions is within a reasonable uncertainty, given the slight variation in their probabilities of interaction. Lastly, by decreasing the particle's energy from 20 GeV to 10 GeV, the mean time decreased by roughly 50%, accordingly.

This is the first time the differences in simulation time due to the underlying mechanisms and structure of particular particles were quantified. As a conclusion, with this project, the goal is to provide useful data for future group projects and full scale, complete geometry studies, as well as for ongoing optimization and development efforts, in the hopes of meeting all the computer requirements expected to arise from the increased luminosity of the HL-LHC, by 2027.

Bibliography

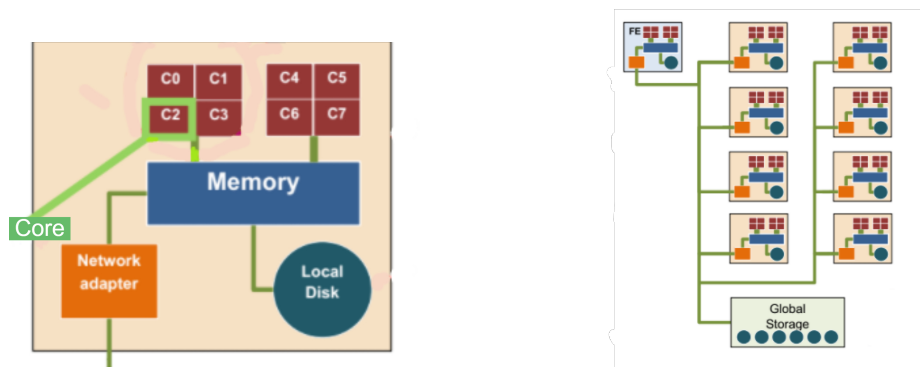
- [1] Herr W. and Muratori B. “Concept of luminosity”. In: *CERN Accelerator School and DESY Zeuthen: Accelerator Physics*. Sept. 2003, pp. 361–377 (page 1).
- [2] “LHC Machine”. In: *JINST* 3 (2008). Ed. by P. Evans L. Bryant, S08001. DOI: [10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001) (page 1).
- [3] Aad G. et al. “Search for a heavy particle decaying into an electron and a muon with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC.” In: *Physical review letters* 106 (June 2011), p. 251801 (page 1).
- [4] Aad G. et al. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008), S08003. DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003) (pages 2, 8, 10).
- [5] Bandieramonte M. “Requirements from HEP experiments”. In: HSF WLCG Virtual Workshop. Nov. 2020. URL: https://indico.cern.ch/event/941278/contributions/4084944/attachments/2147886/3621454/HSFWorkshop_HEPExperimentsRequirements.pdf (pages 2, 22).
- [6] Agostinelli S. et al. “GEANT4: A Simulation Toolkit”. In: *Nuclear Instruments and Methods in Physics Research* 506 (2003), pp. 250–303. DOI: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688> (pages 2, 4, 9, 11).
- [7] URL: <https://lhc-commissioning.web.cern.ch/schedule/HL-LHC-plots.htm> (page 2).
- [8] Marcon C., Smirnova O., and Muralidharan S. “Impact of different compilers and build types on Geant4 simulation execution time”. In: *EPJ Web Conference* 245 (2020), p. 05037. DOI: [10.1051/epjconf/202024505037](https://doi.org/10.1051/epjconf/202024505037) (pages 2, 12, 13).
- [9] Thomson M. *Modern Particle Physics*. New York, USA: Cambridge University Press, 2013, pp. 13–27 (pages 3, 7–10).
- [10] Altarelli G. Wells J. *QCD: The Theory of Strong Interactions*. Vol. 937. Wells J.(eds) Collider Physics within the Standard Model. Lecture Notes in Physics. Springer, 2017. Chap. 2. DOI: [10.1007/978-3-319-51920-3_2](https://doi.org/10.1007/978-3-319-51920-3_2) (page 4).
- [11] Höche S. “Introduction to parton-shower event generators”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*. 2015, pp. 235–295. DOI: [10.1142/9789814678766_0005](https://doi.org/10.1142/9789814678766_0005) (page 4).
- [12] Leo W. R. *Techniques for nuclear and particle physics experiments: a how-to approach; 2nd ed.* Berlin: Springer, 1994. DOI: [10.1007/978-3-642-57920-2](https://doi.org/10.1007/978-3-642-57920-2) (page 4).
- [13] Amako K. “Introduction”. In: IEEE Nuclear Science Symposium - Medical Imaging Conference. Geant4 Collaboration, 2003 (pages 4, 11).
- [14] Zyla P.A. et al (Particle Data Group). “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104) (pages 5, 6, 29, 30).

- [15] Lechner A. “Particle interactions with matter”. In: *CERN Yellow Rep. School Proc.* 5 (2018). Ed. by Bernhard Holzer. DOI: [10.23730/CYRSP-2018-005.47](https://doi.org/10.23730/CYRSP-2018-005.47) (pages 5–9).
- [16] Klinkby E. B. *Simulation of Transition Radiation and Electron Identification Ability of the ATLAS TRT*. Tech. rep. ATL-INDET-PROC-2011-043. Geneva: CERN, Dec. 2011. URL: <https://cds.cern.ch/record/1405697> (page 9).
- [17] *Muon Spectrometer*. Dec. 2017. URL: <https://atlas.cern/discover/detector/muon-spectrometer> (page 10).
- [18] Eckhardt R. “Stan Ulam, John Von Neumann, and the Monte Carlo Method”. In: *Los Alamos Science Special Issue* 15 (1987), pp. 131–136 (page 10).
- [19] Chytracsek R. et al. “Geometry Description Markup Language for Physics Simulation and Analysis Applications”. In: *IEEE Transactions on Nuclear Science* 53.5 (2006), pp. 2892–2896. DOI: [10.1109/TNS.2006.881062](https://doi.org/10.1109/TNS.2006.881062) (page 11).
- [20] URL: <https://root.cern> (page 11).
- [21] Walker B., Figgins J., and Comfort J.R. “A framework for Monte Carlo simulation calculations in Geant4”. In: *Nuclear Instruments and Methods in Physics Research A: Accelerators, Spectrometers, Detectors and Associated Equipment* 568.2 (2006). DOI: [10.1016/j.nima.2006.06.070](https://doi.org/10.1016/j.nima.2006.06.070) (page 11).
- [22] ATLAS Collaboration. “A measurement of the calorimeter response to single hadrons and determination of the jet energy scale uncertainty using LHC Run-1 pp collision data with the ATLAS detector”. In: *Eur.Phys.J.C* 77.26 (2017). DOI: [10.1140/epjc/s10052-016-4580-0](https://doi.org/10.1140/epjc/s10052-016-4580-0) (pages 12, 13, 31).
- [23] URL: <https://slurm.schedmd.com/overview.html> (page 12).
- [24] URL: <https://gitlab.cern.ch/adotti/Geant4HepExpMTBenchmark> (page 12).
- [25] *How to handle dynamic and static libraries in Linux*. URL: <https://opensource.com/article/20/6/linux-libraries> (page 12).
- [26] Dotti A. “Simulation of Showers with Geant4”. In: *International Conference on Calorimetry for the High Energy Frontier*. 2013, pp. 247–253 (page 13).
- [27] Vignesh T. et al. “Scheduling concurrent applications on a cluster of CPU–GPU nodes”. In: *Future Generation Computer Systems* 29.8 (2013), pp. 2262–2271. DOI: [10.1016/j.future.2013.06.002](https://doi.org/10.1016/j.future.2013.06.002) (page 13).
- [28] Akchurin N. et al. “On the differences between high-energy proton and pion showers and their signals in a non-compensating calorimeter”. In: *Nuclear Instruments and Methods in Physics Research A: Accelerators, Spectrometers, Detectors and Associated Equipment* 408.2-3 (1998), pp. 380–396. DOI: [10.1016/S0168-9002\(98\)00021-7](https://doi.org/10.1016/S0168-9002(98)00021-7) (page 20).
- [29] URL: <https://lunarc-documentation.readthedocs.io/en/latest/> (page 25).
- [30] Hein J., Lindemann J., and Ullner M. “Using the batch systems on Lunarc HPC platforms”. Lund University. Lecture. 2017 (page 25).

Appendix A

Lunarc Aurora system

This appendix intends to provide a brief explanation of the structure of a High Performance, HPC, cluster. More specifically, the one used at Lund University, called the Aurora cluster, managed by the Lunarc department. Aurora is composed of over 230 computer nodes, as schematized in Figure A.1b, ideal for intensive jobs. Each computer node, pictured in Figure A.1a, is a multiprocessor system and each computer processing unit, CPU, is made up of several cores. In the case of Aurora, each node contains 20 64 MB cores. These will manage and execute the computing processes [29].



(a) Structure of a computer node: composed of multiple cores and a local disk. (b) Computer cluster, made up of several nodes, which make use of the global storage.

Figure A.1: Schematics of a High Performance Computing Cluster, taken from [30].

Besides local storage disks which allow for fast memory access, the cluster also provides global storage to all nodes. This guarantees the best performance in job processing, including simulations in High Energy Physics. To manage and schedule the needs of multiple users, Aurora uses the batch system SLURM. In order to run a job, the user is required to provide a Bash job description containing information and instructions on the desired operation.

Appendix B

Simulation time plots at 10 GeV

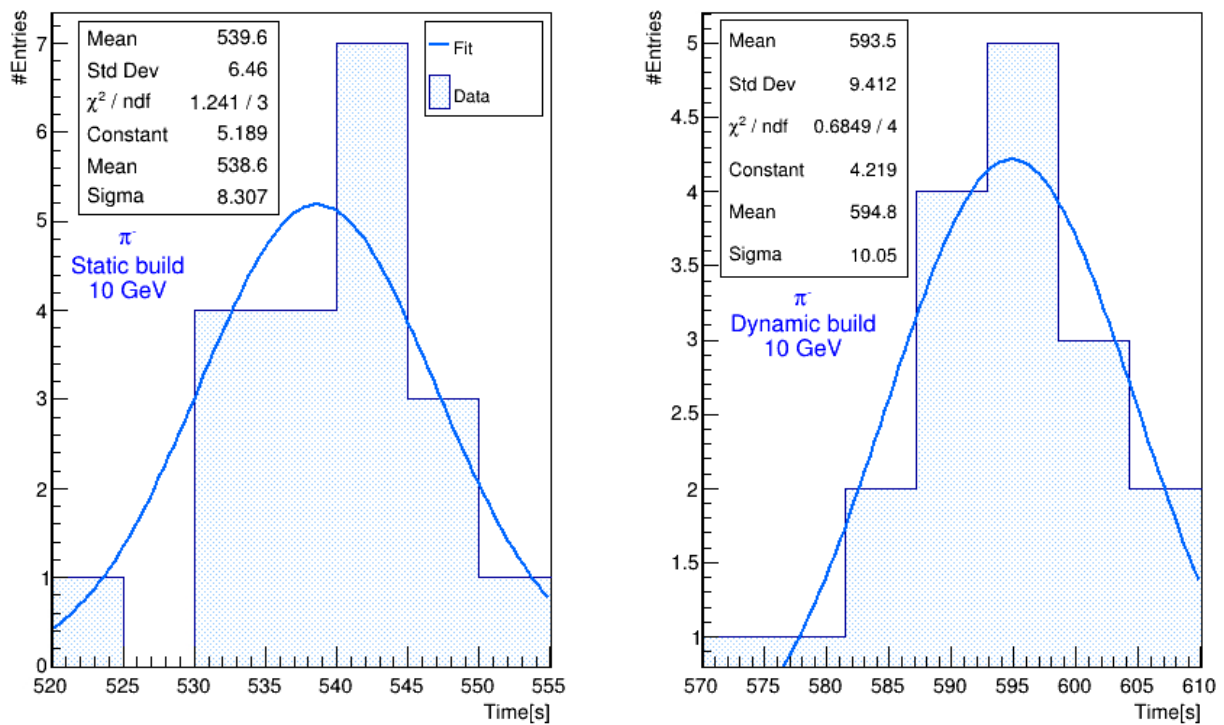


Figure B.1: Time distributions for negative pions at 10 GeV: static build (left) and dynamic build (right).

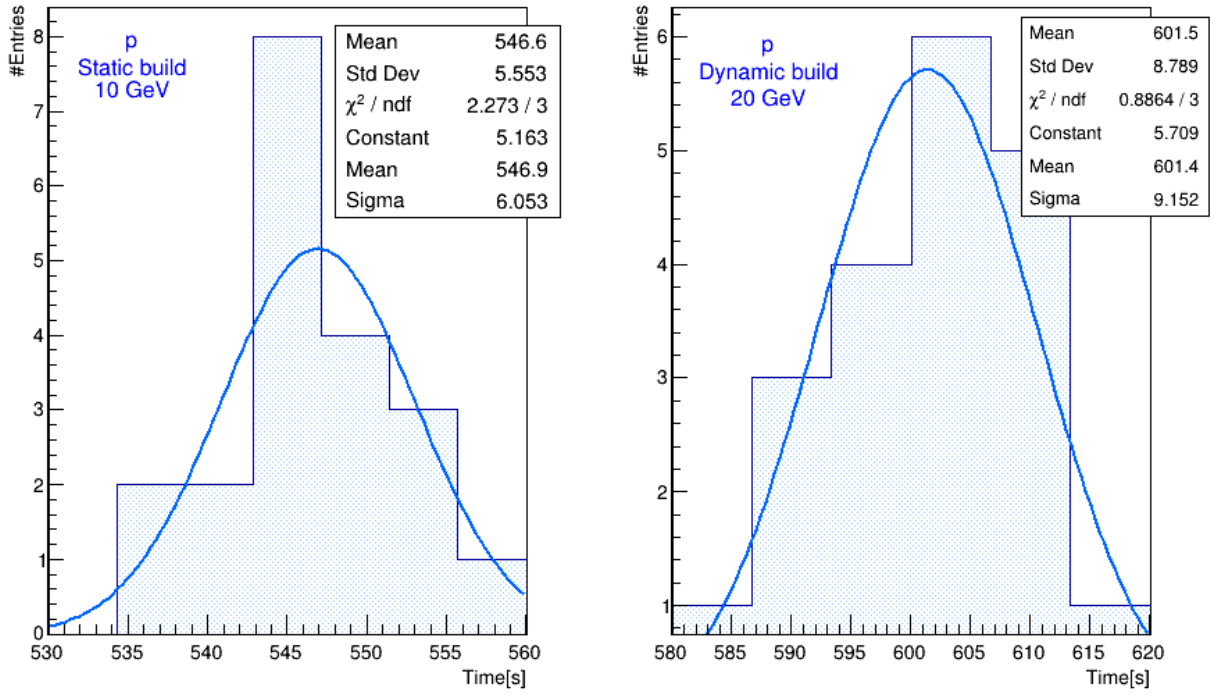


Figure B.2: Time distributions for protons at 10 GeV: static build (left) and dynamic build (right).

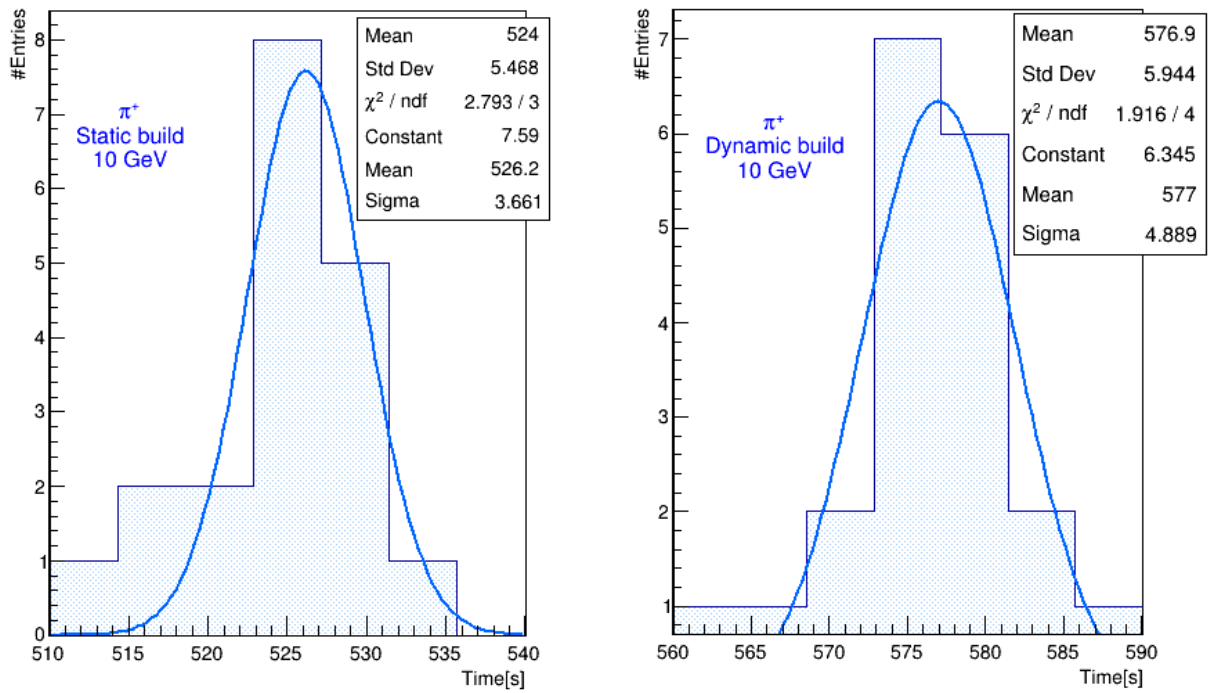


Figure B.3: Time distributions for positive pions at 10 GeV: static build (left) and dynamic build (right).




Figure B.4: Time distributions for geantinos at 10 GeV: static build (left) and dynamic build (right).




Figure B.5: Time distributions for charged geantinos at 10 GeV: static build (left) and dynamic build (right).

Appendix C

Additional reference plots




Figure C.1: Stopping power plots of muons, pions and protons in various materials[14].




Figure C.2: Stopping power plots of muons, pions and protons in various materials[14].




Figure C.3: Energy fraction carried by particles in several momentum ranges, as a function of jet transverse momentum, p_T [22].




Figure C.4: Calorimeter response to single particles of different types, with respect to to the calorimeter response to π^+ , using FTFP_BERT set of hadronic physics models[22].