

Review of software projects at IN2P3

for the Conseil Scientifique of IN2P3

23rd - 24th June 2022

L. Arrabito - LUPM CNRS/IN2P3

with the contribution of

A. Beck, P. Delsart, J. Frankland, H. Grasland, B. Khélifi, S. Incerti, L. Maigne,
A. Matta, O. Stézowski, R. Terrier, A. Tsaregorodtsev, M. Verderi

1 Introduction

In this document we make a review of the software projects in which IN2P3 has a strong involvement. We do not intend to make a detailed review of each project, but we try to give a general overview of the software activities within IN2P3. Eleven projects, belonging to the different IN2P3 scientific domains (particle physics, astroparticle, nuclear physics and bio-medical physics) will be presented in this review. It has to be noticed that we will not cover those projects that are presented in dedicated thematic sessions of this Conseil Scientifique. We have grouped the 11 projects in the 3 following software categories :

- Simulation : Geant4, Geant4-DNA, GATE, nptool, Smilei
- Data processing and analysis : ACTS, AGATA, KaliVeda, Gammapy
- Workload, data and metadata management : DIRAC@IN2P3, AMI

Most of the presented projects are organized as international collaborations with several partners. For each project we try to give a brief description, the related scientific challenges, the involved human resources and funding, the development plans and the eventual R&D activities.

2 Simulation

IN2P3 is strongly involved since more than 25 years in the development of the well-known **Geant4** software toolkit for the simulation of the passage of particles in matter, widely used in HEP and nuclear physics communities. Moreover IN2P3 has a significant contribution in two other projects tightly connected to Geant4. The first is the **Geant4-DNA** project for the simulation of the physical interactions of ionizing radiation with the biological medium,

today fully integrated in Geant4, while the second is the **GATE** platform dedicated to simulations in medical imaging and radiotherapy, developed on the top of Geant4.

These 3 projects represent the main Monte Carlo simulation platforms in our disciplines and benefit of a great synergy. As an example, Geant4-DNA, originally developed as an independent software package, has been fully integrated in Geant4. The overall human resources at IN2P3 allocated to these projects is of about 19 FTE (5 for Geant4, 6 for Geant4-DNA and 8.4 for GATE) distributed in 9 laboratories.

Finally, it is worth to mention that the coordination of these 3 projects is ensured by IN2P3 permanent staff with Marc Verderi (LLR) spokesperson of the Geant4 collaboration, Sébastien Incerti (LP2i) for the Geant4-DNA collaboration and Lydia Maigne (LPC-Clermont) for the OpenGATE collaboration.

In the nuclear physics domain, the **nptool** software is specialized in the simulation of low energy transfer experiments and is also based on Geant4. Even if we present nptool in this section, it has to be mentioned that nptool is not only a simulation software, but it also built for the data processing and analysis of different types of nuclear physics experiments. The nptool project is led by LPC-Caen with the contribution of IJCLab for a total of 1 FTE.

The last simulation software that we present in this review, the **Smilei** project, is relatively different from the 4 Monte Carlo softwares mentioned above. Smilei is a software for the simulation of plasmas, built around the standard Particle-In-Cell (PIC) method. Contrary to the Monte Carlo softwares, which are well adapted to be executed on HTC resources (e.g. grids or clusters), Smilei is a massively parallel code designed to run on large super-computers. In terms of human resources, IN2P3 contributes with about 1 FTE at LLR. Moreover, the development of an HPC application like Smilei also requires specific equipment. IN2P3 is the major contributor to the parallel cluster ("3 Lab computing" at Ecole polytechnique) used for the Smilei development.

2.1 The Geant4 project

2.1.1 Scientific challenges

Geant4 is a software toolkit for the simulation of the passage of particles in matter. It is written in C++, is free and open source. Geant4 is also the name of the collaboration managing the software. Geant4 started in 1994 with the CERN RD44 project aiming at assessing the interest of Object Oriented (OO) technologies for simulation, in view of covering the various needs of the future LHC. The lack of flexibility of GEANT3 -the de facto HEP standard simulation package at that time- was a driving motivation : adding a new physics process or a new geometrical shape was requiring modifying about 60 routines, making such evolution next to impossible in practice. The OO approach allowed defining a generic kernel to which geometrical shapes or physics processes, in particular, are bound as "plug-in" components. This led to a large flexibility, which made Geant4 to be adopted in other domains by defining for example more detailed processes suited for precise tracking at low energies (with no impact for HEP). The space (human radioprotection, radiation effect on

electronic devices) and medical (e.g. : GATE) (radiotherapy, imaging) domains were quickly interested in Geant4, which is now, as for HEP and nuclear science, mission critical for these domains. Geant4 is also used in material science (e.g. : channeling), homeland security (simulation of scan systems in airports), "muography" (Gizeh pyramid, volcanos, imaging of remaining Fukushima reactor), etc.

IN2P3 was involved in Geant4 since the start of the project, with Michel Maire -now retired- but still active in Geant4-key developer of the electromagnetic physics module. IN2P3 people defined then the mechanism for fast simulation (LLR), have key contribution in visualization (IJCLab), contribute to geometry (IJCLab), drive the examples (IJCLab), are main authors of simulation at DNA scale, including post-irradiation chemistry (LP2i) –initially motivated by ESA for manned mission to Mars- and drive event biasing (LLR). All IN2P3 laboratories have management responsibilities in Geant4, as can be seen [here](#).

Beyond the development and maintenance activities, it has to be noted that IN2P3 members organize a yearly Geant4 tutorial in Orsay, managed by Ivana Hrivnacova (IJCLab), under the scope of PHENIICS doctoral school, but open to any students or researchers of public institutions, this tutorial being supported by ANF every two other years. This tutorial exists since almost 10 years. The same people organize tutorials for private companies under the scope of CNRS Entreprise. IN2P3 members participate to other tutorials worldwide. The Geant4 IN2P3 project has also contacts or common developments with other French institutions -CEA Saclay, INS2I/IRISA, IRSN- and has of course strong links with other institutions or laboratories through the collaboration : CERN, JLab, UK labs, KEK, ESA, KISTI, INFN, etc.

2.1.2 The software project

The Geant4 project consists in designing, evolving and maintaining the Geant4 software. Geant4 produces one public release per year (by beginning of December). Geant4 has a "Technical Forum", which meets typically 3 times a year, to discuss plans, requests, difficulties, etc. It has also an online user forum (under Discourse), and is using Bugzilla. Geant4 has regular publications and presentations in conferences by working group members and publishes general papers (3 at this time) when an important evolution is made.

2.1.3 Project timeline

As said above, Geant4 started in 1994 as the RD44 project, and then was turned into the Geant4 collaboration in 1999, after demonstration of the relevance of the OO technologies for simulation. Geant4 has a continuous evolution, but some important ones are the support for multi-threading processing in 2013 (with Geant4 version 10.0) now evolved to a "tasking" approach ; some extensions of the physics : low energy electromagnetic (started by 2000), Geant4-DNA (by 2007), or recently charm and bottom hadronic physics (i.e. : interactions of these particles in matter).

At IN2P3 level, the wish is to continue with current activities, continue investing on the DNA part –which finds a quite large community of users-, and possibly collaborate on

revision of tracking logic to exploit new hardware trend. Efforts in this direction are made at CERN and US (tentative to port on GPU), and we follow these. More details about the Geant4 global development plan can be found [here](#).

Geant4 at IN2P3 was reviewed by the Conseil Scientifique in March 2009, with some focus of low energy electromagnetic physics and DNA one, and the activity was well appreciated. Geant4 was reviewed in November 2021 as part of the "HL-LHC Computing Review Stage-2 : Common Software Projects"; the review panel made general positive comments, recommended even closer work with LHC experiments on Machine-Learning-based fast simulation and noted the risk of expertise disruption we pointed in the SWOT analysis below.

2.1.4 Human resources, equipment and funding

The number of IN2P3 members (including non permanent staff) involved in Geant4 and related activities is 10 for a total of 4.9 FTE. The human resources are essentially stable over the past year, but with arrival of an associate professor and hiring of research engineer at LP2i. The main activities of IN2P3 laboratories are :

- LP2i (2.06 FTE) : Geant4-DNA, low energy EM physics (inclusion of EPICS 2017 data)
- IJCLab (1.05 FTE) : Examples and analysis module
- LAPP (0.2 FTE) : Electromagnetic physics validation
- LLR (1.55 FTE) : Bethe-Heitler 5D module (electrons and muons), biasing options, fast simulation and spokespersonship

The Geant4 collaboration itself is a worldwide one, with about 130 members, for a total of about 30 FTE. CERN –with about 20-25% of the resources- is the biggest contributor today ; IN2P3 is about 12-15%, but with quite responsibilities, as mentioned above.

Geant4 has no "common funds", each institution provides the resources for its members. The IN2P3 budget is rather flat, with about 10 k euros/year, essentially for missions to pay the travels and stays for the yearly collaboration meeting, plus some management-like missions to CERN for Marc Verderi.

In terms of computing resources, the project members don't have special demands : each developer uses one laptop for developments and uses the CERN resources for code management (GitLab) and testing infrastructure. There is no plan to invest in new equipment (at least nothing of large scale) in near future. Some investment on new hardware might come if ideas mature enough.

2.1.5 SWOT self-analysis

The by far most important weakness of the Geant4 collaboration is the risk of expertise disruption. By the start of the HL-LHC most of the experienced developers will have left the collaboration. The rate of renewal is too low ; the short/fixed term contract policy is essentially incompatible with such long term project. This risk is not only the Geant4's one, given Geant4 is mission critical in many domains.

- Strengths : experienced collaboration and members –even retired people continue to contribute-, excellent physics and technical expertise, excellent complementarity between them, good communication, good climate in general, good collaborative work between individual up to large centers, strong CERN support for what validation and testing machinery is involved, attractiveness for new developers.
- Weaknesses : collaboration is aging, renewal rate is too low, excellent young members that would be natural persons to take over, but all are under fixed term contracts. High / very high risk of expertise disruption –with loss of development and loss of support- in the coming years, in particular in all physics domains, a risk that will be turned to a certainty by the start of the HL-LHC if keeping the current human resources trend.
- Opportunities : good reputation of Geant4 in many domains –public or private-, large and diverse community of users –sources of evolution-, opening and/or extension to new physics domains (e.g. : EIC/JLab, FCC/CERN, etc.), adaptation to new computing opportunities –when possible-. We can mention the CERN and FNAL tentative to vectorize the transport, stopped in 2021 after disappointing performances, and the AdePT (CERN) and Celeritas (US) projects to try to port Geant4 on GPU, with partial success at this point, but not yet realistic workflow, but other such R&D can be attractive to new developers.
- Threats : too low rate of human resources renewal, risk of too short overlap between “old” and “young” generation, short term contract policy incompatible with long term projects like Geant4. A situation that is incentive for the most talented developers –we saw that already several times- to go to private companies, because they’re having long term visions and stable contracts. Risk of seeing Geant4 evolving to an aggregate of functionalities with low consistency and low interoperability because of the loss of the global vision of the experienced developers. Risk of losing support of existing features because of loss of related expertise. The mitigation of these risks is in the hands of the funding agencies to define long terms support plans.

2.2 The Geant4-DNA project

2.2.1 Scientific challenges

A mechanistic understanding of the biological effects of ionizing radiation remains a major challenge in current radiobiology. The computational approach is now favored to respond to this challenge, in particular to meet the need for accurate tools for radiotherapy treatment planning, or to better estimate the risk to human health during long-term exposure to ionizing radiation in manned space missions. Numerous simulation tools have been developed over the last decades. They are able to simulate the damage induced to the DNA of the cell nucleus, which is still considered as the main site sensitive to ionizing radiation in cells, following a “bottom-up” approach (from DNA to macroscopic biological effects). Unfortunately, none of them are currently available to users, which prevents their widespread use and adaptability to different user needs.

2.2.2 Project organization

The Geant4-DNA project was initiated in 2001 by Petteri Nieminen of the European Space Agency (ESA), with the aim of providing the scientific community with the first open access platform for the simulation of biological damage caused by ionizing radiation at the DNA level. After the first developments initiated within the Geant4 collaboration on the simulation of the physical interactions of particles with liquid water in the context of studies focused on cellular irradiation and based on a collaboration between INFN and IN2P3, the ESA proposed to IN2P3 to coordinate this project from 2008. The Geant4-DNA collaboration was then created. It gathers today 60 international collaborators, including 13 at IN2P3 (4 laboratories involved : IPHC, LP2i, LPC-Clermont, SUBATECH).

A work program has been established, combining specific developments allowing the simulation of the physical interactions of ionizing radiation with the biological medium, of the associated radiolysis, of various geometries of biological targets, and of a number of applications for users, based on the disciplinary expertise available within the collaboration (not only radiation physics, but also radiation chemistry, radiation biology and computing).

2.2.3 State of the art

In an open science approach, Geant4-DNA is the first fully accessible platform that has been developed for the mechanistic modeling of biological effects of ionizing radiation. Many other "proprietary" codes exist and continue to be developed (see a detailed list [here](#) : in Table 1), highlighting the dynamics within this research field. These codes are unfortunately not freely available. On the contrary, Geant4-DNA is fully integrated into Geant4; thus, it becomes accessible to other simulation platforms that are based on Geant4, such as GATE or TOPAS (and its Geant4-DNA specific TOPAS-nBio interface) and GAMOS. This accessibility has allowed Geant4-DNA to become a reference simulation platform for the international community and the feedback provided by users to the Geant4-DNA developers not only ensures large-scale testing of the tool, but is also a source of improvements and development of new features.

2.2.4 Human resources and funding

The project is led by IN2P3 since 2008, re-elected every two years by the Geant4-DNA collaboration. The resources are made up of permanent researchers and engineers, involving many non-permanent junior researchers (in total, 14 PhD students, 17 postdoctoral fellows), in France and abroad. At IN2P3, four laboratories are members of the collaboration (IPHC, LP2i, LPC-Clermont, SUBATECH). Geant4-DNA has been a project within the Master Project "MOVI" until 2019 (annual funding : 5k euros), the support of the IN2P3 is now through the support of the research teams. In addition, Geant4-DNA has benefited from the strong and regular support of the institute through collaborative tools such as IRL (South Korea, Japan), IEA (Greece, Lebanon, Serbia), international exchanges (JINR, Czech R.). It has also regularly benefited from several external funding (Région Nouvelle Aquitaine, Aviesan/Inserm, ANR, ARC Australia, CampusFrance STAR South Korea, European Space

Agency - 3 contracts BioRad1, BioRad2, BioRad3 -, FNS Switzerland). Support for international exchanges is obviously essential. In 2022, a research engineer position was created at LP2i to reinforce the Geant4-DNA activity, which is a nice recognition of the efforts made. Several partners have also recruited young collaborators (e.g. ANSTO - Australia, Ioannina U. - Greece, IRSN - France, KEK - Japan, NPI - Czech R., Osaka U. - Japan, Unicamp - Brazil), allowing to perpetuate their contribution to Geant4-DNA.

In 2022 we count 10 permanent staff at IN2P3 for a total of about 6 FTE distributed among the following laboratories : IPHC, LP2i, LPC-Clermont, SUBATECH. In addition 2 PhD thesis are on-going (2 at LP2i including one in co-supervision with IRSN, 1 at IPHC).

2.2.5 Achievements

The developments of Geant4-DNA are numerous and all integrated in Geant4. They have been published in more than 120 publications and we regularly organize international conferences (we initiated the "Geant4 at the Physics-Medicine-Biology frontier" conference series in 2005, next iteration in 2022), workshops and tutorials, see all the information on the project [website](#).

Very briefly, the developments focused on : the development of models allowing the simulation of the physical interactions of elementary particles with biological matter ; modeling of water radiolysis ; modeling of realistic multi-scale geometries ; the possibility to combine the above-listed modeling methods to simulate early and late damage to irradiated DNA ; various application examples in physics, physical chemistry, radiolysis and radiobiology. More detailed information on these developments is available on the project [website](#).

Being integrated into Geant4, these developments benefit from the IT infrastructure for nightly testing codes on a hundred or so configurations (OS / compiler / etc.) deployed at CERN that we could not manage at our level (e.g. at the IN2P3 CC). These tests ensure a large software compatibility.

2.2.6 SWOT self-analysis

- Strengths : steering of the collaboration, disciplinary expertise available at IN2P3 (radiation physics, radiolysis, radiobiology) for the development of Geant4-DNA components and its applications to several domains (medical physics, space science, environmental science), link with our experimental platforms for validation of the toolkit (radiolysis, radiobiology).
- Weaknesses : mainly at the software level due to the rapid evolution of C++, the need to develop quality software and to optimize it (need for resources).
- Opportunities : open science, extensibility (new functionalities), international visibility (open-source code, organization of conferences and tutorials), strengthening the links between the Geant4, Geant4-DNA and GATE collaborations.
- Threats : Geant4 lifetime, growing competition (although the collaboration always welcomes any new development opportunity).

2.3 The GATE platform

GATE is an advanced open-source software developed by the international **OpenGATE** collaboration and dedicated to Monte Carlo simulations in medical imaging and radiotherapy. GATE is based on the Geant4 toolkit.

2.4 Project organization and human resources

The OpenGATE collaboration gathers 25 institutions (laboratories, companies and medical centres) worldwide involved in the development of the GATE platform. The collaboration is managed by a steering committee with representative members from all institutions involved. The collaboration receives support from different research agencies as well as from institutions involved in the collaboration. France is the major contributor with 11 laboratories (IN2P3, INSERM and CEA) for a total of 19 FTE, where the IN2P3 contribution is of 8.4 FTE distributed among the following laboratories : LPC-Clermont (2.5 FTE), IPHC (2 FTE), CPPM (1.3 FTE), IP2I (1.3 FTE), LPSC (1 FTE) and IJCLab (0.3 FTE). Moreover, IN2P3 has management responsibilities with Lydia Maigne (associate professor at LPC-Clermont) spokesperson of the collaboration since 2008.

2.5 State of the art

At the international level, GATE is the most popular Monte Carlo platform in TEP imaging and SPECT, with more than 2000 registered users (in particular in France and US). So far, US competing projects have been focusing on radiotherapy and radiobiology with the support of regular fundings from NIH. However all these codes, the most notable example being TOPAS, are not open source and can be used only as black boxes, like any commercial product. In order to compete with this offer, GATE and Geant4-DNA communities plan to develop python bindings to facilitate the installation procedure and to provide a more flexible user interface. Thanks to these developments GATE would be able to keep the leadership in this domain.

2.6 Development plans

In order to keep the GATE platform competitive, it's necessary to apply some structural changes in the code basis. The current code, written in C++, should be modernized and partially re-written. As said before, it is also foreseen to encapsulate some portions of the code in python to facilitate the installation procedure and to provide a more flexible user interface. The main development axes for the next 5 years are the following :

- Make GATE compatible with the multi-threading feature of Geant4
- Develop a new data analysis chain for medical imaging
- Develop python-based encapsulating techniques in common with Geant4 project to ease the installation of GATE and all its dependencies
- Support of complex geometries and analysis by developing dedicated python bindings

- Integrate variance reduction techniques and analytical methods for the following of accelerated particles in a dedicated library
- Improve the interface with third-party libraries as ITK for the handling of medical images
- Make the platform more compatible with third-party libraries for the image reconstruction (Castor, Stir, RTK, ...)
- Integrate machine learning techniques to speed-up the execution
- Develop new output formats to facilitate end-users analysis with python
- Improve adapted tests and benchmarks at each development stage
- Improve multi-OS compatibility
- Facilitate the development of multi-scale open source simulations (from molecule to patient)

In France, BioMaps, CREATIS and LaTIM laboratories are already strongly involved in these developments. Thanks to the recent recruitment of a 3-years contract research engineer (with advanced C++ skills) at LPC-Clermont, IN2P3 is now in the position to significantly contribute to this collective effort. At longer term, it will be important to keep these highly-skilled developers in the project by opening permanent positions. This strategy is also supported by the wider community within the GDR Mi2b (Outils et méthodes nucléaires pour la lutte contre le cancer) and will allow to enforce the already existing synergies among GATE, Geant4-DNA, Geant4 and Machine Learning community at IN2P3.

2.7 The nptool project

2.7.1 Scientific challenges

The main difficulty in planning, performing and analyzing low energy nuclear structure experiment is the fact that each setup is dedicated to specific physics case. The specific of the kinematic, the type of particle to detect and the beam characteristic make the use of a single static setup difficult. Modularity is therefore paramount at two levels. First assembling heterogeneous experimental setup comprise of several modular sub-system should be easily done. To these ends we use a set of human readable input files that describe both the physics at play and the detector setup. Secondly detector are used to produce various types of observable depending on the experiment and the combination of the detectors employed. Separating the experiment invariant part of the analysis from the experimentally specific one is therefore essential. By implementing a specific analysis project for each experiment the two get naturally separated while providing maximum flexibility to the end user.

2.7.2 Project description

The development of **nptool** started in late 2008, with the goal of streamlining the analysis and simulation process of low energy transfer experiment. Two axes of development were followed and formed the core philosophy of the framework. First a universal approach was

taken from the very start, making the framework both detector and physics agnostic. The second essential aspect of the design was to bring all the possibilities of both ROOT and Geant4 with minimum training for the end user.

Each detector is implemented as a group of classes that form a plugin library associated with an ASCII token. The token is used in the input file to trigger the loading of the library and the registration of the detector to a *Detector Factory*. The framework takes care of instantiating and initializing the required classes depending on the use case.

Over time nptool started integrating many useful libraries to help nuclear physicists perform complex analysis and share development time between collaborations. For instance, a full suite of classes is available for calibration of silicon detector, and is used commonly by the MUST2, SHARC, and TIARA collaborations. Tool for track reconstruction using RAN-SAC algorithm is shared between the MINOS-TPC and ACTAR-TPC analysis. Recently cross-talk rejection for neutron scintillator array was initiated and work to deliver a comprehensive library to the community is ongoing with collaboration between MSU, LPC-Caen, and CEA-DAM.

2.7.3 Development plans

With now more than 70 detectors system provided along the framework, maintenance of the package is proving more difficult. To this end a major revision of the framework architecture is in preparation where all the plugins will be externalized. The framework will come with utilities to install, uninstall, and update plugins required by the end user. In addition the possibility to duplicate and modify a given plugin within the scope of a project will be provided.

This architecture will hopefully help tackle the upcoming challenge of Open Science for short-lived experiments, helping the community to provide and publish self-contained, readable and archival, software packages associated with any data set.

2.7.4 State of the art

nptool framework is widely used worldwide to provide online and offline analysis, as well as Geant4 simulation, at majors facilities (GANIL, RIKEN, FRIB, ISOLDE, SPES, ISAAC-2, ELI-NP, Jyvaskyla) and universities or laboratories (IPHC, CEA-Saclay, UK National Laboratory, Surrey, York, Manchester, Padova, Santiago, TU-Darmstadt, Mumbai, Hope, Texas A&M).

Only one similar framework exists, FairRoot, developed at and for the FAIR facility in Germany. FairRoot is designed for needs of large-scale, fixed setup of multi-detector system such as CBM. Therefore FairRoot emphasizes more on the ability to run analysis and simulation on HPC clusters rather than offering the modularity needed by the low energy nuclear physics community.

2.7.5 Human resources

Today a core team of 6 developer exists, mostly based in IN2P3 laboratories. The collaboration is led by LPC-Caen, with support from IJCLab and CEA-DAM. Human resources are distributed as follows : 1 FTE at IN2P3 (0.8 FTE at LPC-Caen, 0.2 FTE at IJCLab), 0.5 FTE at CEA-DAM, 0.4 FTE in TU Darmstadt (Germany) and 0.1 FTE at STFC (UK).

2.8 The Smilei project

2.8.1 Scientific challenges

Plasma physics has been at the forefront of scientific research for more than half a century, covering a very wide range of scenarios and applications that range from the most extreme astrophysical events to laboratory experiments on laser-plasma interactions, from space weather and spacecraft design to fusion energy research. Recent technological developments have opened the field to an even wider range of applications of plasma technology, in particular through the development of compact particle accelerators, with applications in medicine, material science, and novel radiation sources. These are extremely challenging scientific problems, requiring state-of-the-art numerical codes and the most advanced computing resources to further our understanding of this field, and fuel the industrial exploitation and medical applications of the resulting technological developments.

2.8.2 Project description

Smilei is an open-source multi-physics massively parallel scientific application designed to simulate a broad range of plasma physics scenarios. The code is built around the standard Particle-In-Cell (PIC) method enriched by many additional physics modules : particle collision, field ionization, etc. Smilei also features state-of-the-art performances on most of the largest super-computers in Europe and some of the top ones in the world. Smilei is complemented by a toolkit suite including “Easi”, a python-based validation process used for continuous integration and available for all developers and “Happi”, a python-based visualization and post-processing library.

The project started in 2013 as a numerical tool for the Apollon laser community in a context where no truly open-source, high-performance solution existed. Today, Smilei is a state-of-the-art mature software solution used by many laboratories in Europe and worldwide and supported by several major institutions. Over 100 research articles (see [here](#)) have been published in peer-reviewed international journals. Efficient software engineering and high-performance computing (HPC) are ensured by the close collaboration between physicists (LULI¹ and LLR) and HPC experts (Maison de la Simulation, IDRIS, CINES).

1. <https://portail.polytechnique.edu/luli/>

2.8.3 The Smilei community

Smilei is also built around the important notion of a connected community offering collaborative tools (detailed documentation, GitHub, Element chatroom) and permanent/constant support to the users. Exchange and collaborations are pushed by the organisation of regular international workshops (every 2 years) gathering students, PhD and researchers. These workshops include a large amount of time of tutorials and practical sessions with direct interactions with the whole development team.

Smilei is also a learning tool for many universities that use the code to teach plasma physics and computational physics to Master students in Europe. A strong effort has been made on pedagogy. The **documentation** is designed for both new and advanced users as well as developers. A set of detailed open-source tutorials is also available online.

Every year, GENCI allocates approximately 15 million CPU hours for Smilei computation in France.

2.8.4 Smilei development plan for the laser wakefield acceleration applications

The context of PIC simulations for laser wakefield acceleration is dominated on the one hand by the emergence of new heterogeneous architectures for exascale and on the other hand by the necessity to produce quasi real time simulations. The Smilei team plans to match this context by developing code optimized for this new hardware but also with the introduction of more adapted numerical methods.

Today, it is clear that an efficient code must be able to benefit from GPU acceleration. This is the highest priority but alternative architectures of type ARM/RISC are also available in super-computing centers in Japan and soon in Europe. The related objectives for Smilei are :

- Porting the whole code on NVIDIA, AMD and Intel GPUs
- Optimize the code for ARM/RISC

Accuracy and performance of the simulations are also affected by the numerical methods at play. To match the exascale tendency, it is useful to devise numerical methods which reduce the overall size of the simulation and increase its arithmetic intensity. In this context, the main objectives are :

- "Perfectly Matched Layer" for Maxwell and envelope solvers
- Optimization of "particle merging"
- Multi-grid refinement
- Boosted frame

2.8.5 Human resources, equipment and funding

Smilei is developed by a core team of 7 permanent staff belonging to different French institutions (INP, INSIS, IN2P3, Ecole polytechnique and CEA).

IN2P3 contributes to this team with 1 research engineer and 1 post-doc at LLR for a total of 0.9 FTE. Additional contributions to Smilei from IN2P3 in the past include a PhD at LLR (half funded by IN2P3), a PhD at IJCLab and 4 years of post-doc at LLR 2 of which were funded by IN2P3.

National computing centers IDRIS and CINES are also involved in the code development (porting to Nvidia and AMD GPUs).

The development of a HPC application requires specific instruments and softwares. Smilei is developed mostly on a local parallel cluster at Ecole polytechnique called “3 Lab computing”. It is funded by Ecole polytechnique and the contributing laboratories, LULI, LSI and LLR. Over the last 9 years, 321k euros were invested in this cluster including 117k euros by IN2P3 which is the major contributor via the master project ALPe and the GALOP team at LLR.

3 Data processing and analysis

In this section we report about 4 software projects developed within HEP, astroparticle and nuclear physics communities. The first is the **ACTS** project, developed within the HEP community, aiming to provide a common software for track reconstruction. While traditionally each experiment has been developing its own suite of data processing softwares, ACTS project aims to mutualize the efforts and develop a common solution for the track reconstruction, which is one of the most complex and CPU-consuming task of the data processing for HEP experiments.

The second project presented here, **Gammapy**, has been developed within the astroparticle community and follows a similar approach to provide common tools to perform the analysis of high-level data produced by different experiments or observatories in the gamma-ray domain.

Similarly, in the nuclear physics domain, **nptool** (described in section 2.7 dedicated to simulation softwares), which also includes data processing and data analysis functionalities, has been designed as a generic tool usable for different experimental setups.

The other two main projects developed within the IN2P3 nuclear physics community, even if specific to a particular detector, provide common processing and analysis tools for their respective collaborations, thus allowing to mutualize the development efforts among different groups. These are the **AGATA software** for the on-line processing of data produced with the AGATA (Advanced GAMMA Tracking Array) detector and **KaliVeda** for the data processing and analysis for the INDRA and FAZIA detectors. Finally, it has to be noticed that some complementarity exists between KaliVeda and nptool. Indeed, KaliVeda is more specialized in treating reactions with a large number of particles in the final state (the reactions studied with INDRA and FAZIA detectors can produce more than 50 charged particles), while nptool is generally used in experiments producing one or two particles per reaction. On the other hand, as explained in section 2.7, nptool covers a wider range of experiments, including the detection of neutral particles (neutrons and gamma), which is not supported in KaliVeda.

3.1 ACTS project

3.1.1 Scientific challenges

The reconstruction of charged particle trajectories is one of the most complex and CPU consuming parts of event processing in HEP experiments. At future hadron colliders such as the High-Luminosity Large Hadron Collider (HL-LHC) or the Future Circular Collider (FCC), the significantly increased number of simultaneous collisions will result in a much more challenging tracking environment. Concurrent algorithms exploiting modern computing architectures with many cores and accelerators are necessary to maintain and improve the tracking performance. Finally, an important effort should be done to design softwares and algorithms that better benefit of the capacities of modern CPU, for instance through caching and vectorization techniques.

3.1.2 Project description

Based on the tracking experience at LHC, the **ACTS** project is an attempt to encapsulate the current ATLAS software into an experiment-independent and framework-independent software designed for modern computing architectures. It provides a set of high-level track reconstruction tools which are agnostic to the details of the detection technologies and magnetic field configuration. The software has been fully tested for thread-safety to support parallel execution of the code and its data structures are optimized for vector operations to speed up linear algebra operations.

The ACTS (A Common Tracking Software) project has been originally launched by the tracking group of the ATLAS experiment at CERN in the 2010s (see [link](#)), while today it gathers 11 core contributors from CERN, LBNL, CNRS/IN2P3, University of Geneva, Eotvos Lorand University in Budapest and Max-Planck Institute for Physics.

ACTS has been explored for a range of different detectors including Belle II, CEPC, sPHENIX, PANDA, FASER, and the future ATLAS Inner Tracking system (ITk) for the HL-LHC data-taking era. Today, there are no competing projects comparable to ACTS.

3.1.3 Status and R&D activities

The ACTS project is mature enough to allow its progressive adoption in production within the ATLAS experiment, which will allow an extensive validation of the software and will encourage its adoption by other experiments.

In parallel, several R&D activities are currently on-going along two main axes. The first axis concerns the development of tracking algorithms suitable to run on GPU, while the second is the study the application of machine learning techniques to the track reconstruction.

As it concerns the development of machine learning algorithms for tracking, the ATLAS group at IJCLab plays since long a major role in this domain, also organizing machine learning challenges ([TrackML](#)). More recently, the LAPP laboratory also started to actively contribute to this activity and will contribute even more in the context of the ATRAPP ANR (2022-2026). Finally, the L2IT laboratory in Toulouse, in collaboration with the Stanford

University, is working on the development of *Graph Neural Networks* tracking algorithms that should be integrated into ACTS by the end of 2022. In addition to these activities, the plan is to produce a first version of a complete track reconstruction chain on GPU by the end of 2022.

Beyond ACTS, several R&D projects in the field of track reconstruction are on-going within IN2P3. As it concerns tracking algorithms on GPU, some projects are much more advanced than the ACTS demonstrator, but they are specific to a given experiment. Among these, we mention the **Allen** project for LHCb and the **O2** project for Alice.

A collaboration between ACTS and Allen has already been attempted but it has not been pursued because of the major difficulty of porting the whole ACTS code to the specific memory management mechanism used in the Allen framework. The lesson learned from this experience is that it will be certainly easier to collaborate with projects that target only a specific component of the tracking chain, as for the CMS **Patatrack** demonstrator. This is also the approach chosen by several tracking machine learning projects, which probably explains the more advanced collaboration with ACTS in this domain.

Thanks to all these activities, IN2P3 contributions to the domain of the track reconstruction are well recognized at international level, but they could probably gain even more visibility with a more coordinated effort among the different groups.

3.1.4 Human resources

The core development team counts about 10 persons, the main contributors being CERN and US universities. IN2P3 contributes to the project less than 10% in terms of human resources with about 1.8 FTE at IJCLab and LAPP. However this contribution tends to grow-up thanks to the recent recruitment of 1 post-doc at IJCLab and 1 post-doc at LAPP. Moreover, as already mentioned, the work of the L2IT team on *Graph Neural Networks* algorithms will also be integrated in ACTS.

3.2 Gammapy

Gammapy is an open-source python research software allowing the astrophysical analysis of high-level data produced by gamma-ray telescopes from GeV to PeV, such as H.E.S.S., CTA, HAWC or Fermi-LAT. Even if most of these data are not yet in open access, the collaborations of these telescopes can analyse with Gammapy their high-level data, which consists of the list of gamma rays with their arrival direction and energy, as well as the associated instrumental response functions to reconstruct astrophysical quantities such as the source flux or its intrinsic size. This software has been selected as the main library to build the analysis tools of the CTA Observatory which will be delivered to the whole community (Press Review : [link](#)) and has been awarded by the jury prize of of the French Open Science awards for open source research software (see [link](#)).

Gammapy is a stand-alone library and application that can be installed on a personal computer or a cluster. Fully inserted into the open-source python ecosystem, it is affiliated with the Lead Developer **astropy** library and uses pre-existing python libraries (some

mandatory like `numpy` or `scipy`, and others optional like `emcee`). This architecture makes it possible to benefit from all the know-how of different communities (e.g. that of the `iminuit` statistical minimization library), while focusing on the central functionalities of our astrophysical applications. The result is a small library, with 51000 lines of code, 19400 lines of tests and 26500 lines for documentation (docstring and tutorial, excluding web pages).

Gammapy offers an API to implement most of the analysis methods published in the very high energy community. In addition, the library offers an innovative analysis method for the very high energy domain, a 3-dimensional analysis (spatial and spectral) of the data using a source model and a background noise model. Besides Gammapy provides reusable modules for other software projects. The most remarkable is `gammapy.maps` which implements multidimensional sky maps (i.e. which allow to represent quantities according to spherical coordinates and an arbitrary number of dimensions, e.g. time or energy). This module is reused by the `Fermipy` package which is used for the analysis of data from the LAT instrument of NASA’s Fermi observatory.

The Gammapy project is the result of an international collaboration, mainly between France and Germany, with strong contributions from Spain and Italy. Like the open-source `astropy` project to which we are affiliated, the Gammapy project is organized (see this [link](#)) with a Coordination Committee, two Lead Developers (APC/France, CfA-Harvard/U.S.A.) and two Project Managers (APC /France, ECAP/Germany).

The main laboratories actively involved in the development are :

- in France APC (IN2P3/Université de Paris Cité) with contributions from LTh (Observatoire de Paris/INSU), AIM (IRFU/INSU) LPNHE (IN2P3/Sorbonne), LUPM (IN2P3/Université Montpellier)
- in Germany MPI-K (Heidelberg) with contribution from ECAP (Erlangen) and TU – Dortmund
- in Spain, IAA-CSIC (Granada), UAM (Madrid)
- in Italy, INAF Palermo

From 2015 to 2022 included, the total number of FTE involved in the project is 12.3. In total, we estimate more than 30 FTE on the same period for the whole project. Finally, two permanent staff at IN2P3 also ensure the roles of Project Manager (Bruno Khélifi, researcher at APC) and Lead Developer (Régis Terrier, researcher at APC).

3.3 Data processing software for AGATA

3.3.1 Scientific challenges

AGATA is a European Collaboration which has the final objective to build a 4PI gamma tracker array. The project has started more than ten years ago, with two phases fulfilled so far. The first one, called demonstrator phase, had been a proof of principle demonstrating it is indeed possible to track gamma-rays, with energies ranging from few keV up to few MeV, using electrically segmented high purity Germanium detectors. During Phase1, the whole

system has been regularly updated, improved but at the same time used as a production tool (see scientific production [here](#)) to investigate various aspects of the nuclear structure.

The concept of gamma tracking is a rupture with the previous technologies used to detect gamma-ray at very high energy precision. The interaction of gamma-ray with matter is probabilistic and a photon could possibly then deposits only part of its energy in a given detection volume. Previous generation of gamma-ray multi-detectors relies on Germanium crystals, for the high energy resolution, surrounded by extremely high efficient, low resolution scintillators (Compton shields, ex : BGO) used to tag and then reject any event which may come from partially collected energy in Germanium.

Because of the high volume covered by the Compton shields, it is impossible to fully cover 4PI with Germanium. As well, the rejection system is not satisfactory. In gamma tracker, the idea is to build an almost full 4PI Germanium and a consequence is one should track any single gamma-ray to really determine its full initial energy. Such an approach should increase by several orders of magnitude the sensitivity to high multiplicity cascades. Since gamma-ray are tracked, the direction of emission is precisely known, the Doppler correction could be then determined much more precisely which is crucial for high recoil velocity experiments.

AGATA is the ultimate gamma-ray spectrometer for nuclear structure. With its tracking capabilities, it is also a Compton camera and any improvements in such direction is likely to be interesting for other fields like medical imaging, decommissioning of a nuclear power station, etc.

3.4 The software project

The initial R&D has been performed thanks to European funding. The concept of gamma tracking required two main key algorithms. The first one is called PSA for Pulse Shape Analysis. A Germanium crystal is electrically segmented (external face) into 36 segments, sensible to part of the active volume, but it has also a central contact sensible to the whole active volume. From the 36+1 digitized signals in a crystal, the PSA algorithm should deliver the list of all the interaction points (3D positions, energies deposited) produced by the interacting gamma-ray. The current algorithm used so far relies on comparisons of the collected signals with a database of signals (association of typical shapes with precise positions in 3D) and is computational-intensive. The second algorithm is the tracking one in charge of reconstructing all the trajectories in the whole array from the hits provided by the PSA applied to each crystal.

One of the objectives of the collaboration is to achieve online, in real time, these two algorithms requiring a complex data processing infrastructure which follows the array from one site to another one. Because of their complexity, the two key algorithms have not yet reached expected performances and a R&D program is active, in parallel of the data taking for scientific production, in order to improve them : machine learning technics are one of the main solution in that direction.

Inside the AGATA collaboration, IN2P3 is strongly involved in this R&D program. It has also strong responsibilities in the computing infrastructure in charge of running online complex data pipelines, most probably including GPU in the future. It is worth mentioning

also, the current AGATA online workflow manager system is an IN2P3 product (DCOD) developed by a team from the IJCLab, as well as the current tracking code (OFT).

3.4.1 Status and development plans

Because of the very high price of segmented Germanium detectors, AGATA is continuously growing which makes also the improvements of the whole system (infrastructure, detectors, electronic, computing) possible. AGATA is a nomad detector which has started to produce data in 2010 for two years at Legnaro (Italy) then for two years at GSI (Germany) and seven years at GANIL (France) : it has grown from 15 crystals up to 45 at the end of the GANIL campaign in July 2021. Then, AGATA has moved back to Legnaro and has been progressively re-installed with the first Legnaro beam on target realized end of April 2022.

AGATA Phase 2, with the goal to reach as much as possible almost 4PI (180 crystals) in 2030, has started by an international review (07/2020) of the previous phases and an IN2P3 review (KDP2 11/2020) allowing a new MOU has now being signed by almost of the European collaborators. With the first beam at Legnaro, the scientific production has re-started and some very last developments are about to be implemented. Amongst them a new electronics boards has been designed including an ethernet readout (PCI for the previous versions). This new electronic board will be added progressively to replace completely the current one around 2025.

From the point of view of the data processing, this step modified the current infrastructure which is adapted to the current electronic which requires one computer with an embedded PCI card per crystal. Such modification of the paradigm would allow to develop advanced solutions including dynamic load balancing and complex soft trigger into the workflow with the goal to better control the quality of the data produced (up to the final number of crystal in AGATA) but also to optimize the processing and the storage of only the most interesting events. As well new algorithms, likely being based on machine learning approaches, would modify the data pipeline requiring to have a system able to use GPU, possibly other hardware such as FPGA.

3.4.2 Human resources and equipment

Technical and human resources, for the IN2P3 contribution, have been evaluated in detail during the KDP2 review in 2020. The total number of FTE at IN2P3 contributing to software activities is about 5 FTE which represents about 30% of the total within the collaboration. The involved laboratories are : IJCLab, IP2I and GANIL.

While the collaboration is able to develop a R&D program on machine learning algorithms and their use in online data pipeline – promoting the use of containerized applications -, it will certainly benefit of much more strength and expertise in such domains to accomplish the foreseen work. It is one of the reason why the collaboration has tried to apply to a European call : "R&D for the next generation of scientific instrumentation, tools and methods" (HORIZON-INFRA-2022-TECH-01-01).

As it concerns the computing infrastructures for the data processing, within the current model, the raw data produced by AGATA have been saved for more than ten years in two grid data centers CC-IN2P3 and INFN Bologna. It represents about 500 TB for seven year of processing during the GANIL campaign.

Moreover, the R&D Phase 2 on algorithm already benefits from different IN2P3 facilities such as the IP2I GPU farm as well as the CC-IN2P3 one. Within the AGATA collaboration, financial resources are available also to set up small and specific benchmarks (few machines, some including GPU/FPGA to test integration in the current workflow). The software developed by the collaboration are most of the time on git repositories such as the one hosted at CC-IN2P3 (see for instance the [agapro](#) package containing most of the data processing bricks used so far in AGATA data pipelines). For this one, continuous integration capabilities are also used and are to be extended.

As machine learning algorithms are seen as promising, resources may be required to train supervised algorithms : it is foreseen then to use any available resources provided by IN2P3, other countries or Europe.

3.5 KaliVeda Heavy-Ion Analysis Toolkit

3.5.1 Scientific challenges

Heavy-ion collisions in the energy range of 10 to a few 100s of MeV/nucleon are used to study the thermodynamic properties of nuclear matter at sub-saturation densities and temperatures below the critical temperature of the nuclear liquid-gas phase transition. The final state of each collision is a multi-body event constituted of many different nuclear fragments, with widely different charge, Z , and mass, A . [INDRA](#) is an array of 336 charged particle identification modules covering 90% of the solid angle around the target in the laboratory frame, used for the study of such collisions since 1993. [FAZIA](#) is a new modular array specially developed to perform Z & A identification for nuclei up to $Z=25$, currently constituted of 12 16-telescope blocks placed at forward angles in the INDRA scattering chamber at GANIL since 2019. KaliVeda provides all software tools used for data reduction, particle identification, detector calibration and analysis of data produced by INDRA and FAZIA, as well as being used as the base framework for the slow control interfaces used to control the two arrays during experiments.

3.5.2 Project description

[KaliVeda](#) is a software framework written in C++ and based on ROOT. Beginning in 2002, it was initially developed in order to handle analysis of the large database of reactions studied with the INDRA charged particle multidetector array in several campaigns from 1993 onwards. The initial aim of the project was to provide a unified environment for analysis of existing data from different campaigns at CC-IN2P3 in Lyon. Another aim was to increase the efficiency of the data reduction for future campaigns by using a modular approach based on object-oriented programming in order to maximize the re-use of tried and tested code.

For each new experiment using INDRA and/or FAZIA, once the specific detector geometry used is implemented, reduced data for physics analysis can be produced by KaliVeda from raw data without further code development. Only the specific detector calibration parameters and particle identification grids need to be supplied. For the former, KaliVeda includes a library for fast analytic calculation of ranges and energy losses of charged particles in matter, adapted to the GANIL bombarding energy range. For the latter, KaliVeda provides a full set of graphical interfaces to facilitate the preparation of grids for Z- and A-identification of particles in either ΔE -E or pulse- shape discrimination (PSD) matrixes.

3.5.3 State of the art

Currently, all INDRA, FAZIA and INDRA-FAZIA experiments are included in the KaliVeda framework which is used for data analysis by the INDRA & FAZIA collaborations. Analysis of all data can be performed within the same framework either using the batch system at CC-IN2P3 (where data for both detectors are stored) or using the ROOT parallel-processing facility PROOF Lite on a multi-core laptop or desktop PC. Users implement their analysis by modifying a few methods in a template analysis class, using the many analysis tools provided by the framework and the histogramming (or TTree-filling) facilities provided by ROOT. The user then selects the data to analyse (by choosing the experimental campaign, and particular reaction to study) using a graphical interface tool, without dealing with details such as knowing where the data to analyse are stored, what the files containing the data are called or how to read them, etc.

3.5.4 Development plans

The development of KaliVeda is driven by the experimental program and analysis needs of the INDRA and FAZIA collaborations. A major evolution of the framework occurred in 2019 with the first coupling of the INDRA and FAZIA arrays, in order to handle the data of this experimental campaign. A similar evolution is currently underway as the DAQ system of INDRA has been completely replaced with digital electronics over the last year, and the E818 experiment which was run in April/May 2022 at GANIL was the first to use this new DAQ, again coupled with FAZIA.

3.5.5 Human resources and equipment

KaliVeda is an open-source project whose source code is currently hosted on github.com. In the future this will be migrated to the IN2P3 gitlab platform which is already used for continuous integration and to host the projects using KaliVeda as a base framework for the slow control interfaces used by INDRA and FAZIA. A core team of 3 maintainers/developers for a total of 0.7 FTE is spread over 3 IN2P3 laboratories (0.4 FTE at GANIL, 0.2 FTE at LPC-Caen, and 0.1 FTE at SUBATECH). The codebase also includes contributions from the many different students who have used KaliVeda for their PhD work.

4 Workload, data and metadata management

In this section we report about two software projects for the management of computing tasks (workload) and the handling data and scientific metadata in distributed and heterogeneous computing environments.

The first is the **DIRAC@IN2P3** project, i.e. the French component of the international DIRAC Interware project for the workload and data management over distributed infrastructures.

The second is **AMI**, a framework which provides advanced functionalities for the handling of scientific metadata coming from different sources and having different structures. Both of these frameworks have been originally developed within LHC collaborations and then evolved to become common solutions usable by several communities beyond the HEP domain.

4.1 The DIRAC@IN2P3 project

4.1.1 Genesis of project

The DIRAC@IN2P3 project was created in 2017 in order to consolidate efforts in CNRS laboratories related to software development and other activities within the DIRAC Consortium. The Consortium develops and promotes the **DIRAC** Interware – the software used for building distributed computing systems for scientific applications of arbitrary complexity. The Consortium unites 6 international partners including CNRS.

The DIRAC Interware project develops a framework and multiple components to provide a complete solution to manage large volumes of distributed computing and storage resources available to and shared by multiple scientific collaborations. It allows to manage complex workflows of modern scientific applications for both computational and data managements tasks. The DIRAC functionality is available to users via multiple interfaces including a versatile Web Portal. The DIRAC development framework offers programming interfaces that allow extending the already available components for the needs of particular applications. The software is constantly evolving in order to accommodate new technologies in the field of intensive computations and data management.

The DIRAC Interware project was initiated by CPPM in 2002 as a solution for a distributed computing system of the LHCb experiment at LHC, CERN. The success of the project inspired other large collaborations in the High Energy Physics domain to adopt it for their needs (Belle II at KEK, BES III at IHEP and others). In 2009 the software was generalized in order to make it suitable for any scientific community that needs management of intensive computation and large data volumes. In 2014 the DIRAC Consortium was created with the goal of further development and promotion of the DIRAC software. The Consortium includes currently 6 partners – CNRS, CERN, IHEP/China, KEK/Japan, Imperial College/UK and University of Montpellier.

Since 2012 the DIRAC software was updated in order to support multi-community installations. This allowed to provide services built upon this software as part of large computing

infrastructures. In particular, in 2012 DIRAC services were offered to the users of the France-Grilles National Grid Infrastructure (NGI). This facilitated access to computational grid to a large number of small scientific communities that can not afford complex developments because of lack of expertise and manpower. Similar services were later launched for other grid infrastructures - European Grid Infrastructure (EGI), GridPP NGI in UK, JINR in Russia and others. In particular, the EGI Workload Manager service based on DIRAC is now part of the European Open Science Cloud (EOSC) catalogue.

4.1.2 Project organization and achievements

The DIRAC@IN2P3 project unites 5 IN2P3 laboratories (CPPM, LUPM, CC-IN2P3, LPSC, IPHC) as well as the CREATIS laboratory of CNRS/INSERM. It gathers about 12 research engineers from IN2P3 laboratories for a total of 2.3 FTE. The project sets several goals : development of DIRAC components for the needs of new scientific communities and technology advances ; maintenance and operations of the EGI DIRAC service ; support of the large DIRAC user community. In the development area, the project main accomplishments are :

- Production system – the system that allow to manage very complex workflows of modern HEP and astrophysics collaborations involving hundreds of thousands of computational tasks and tens of millions of data files organized in datasets. This development generalizes the experience of large collaborations like LHCb, CTA, ILC/CLIC using DIRAC for their data production activities
- Integration of cloud computing resources into workflows managed by DIRAC
- Integration of computing resources available in stand-alone computing centers, in particular HPC centers
- Development of the Authentication and Authorization (AAI) framework for DIRAC users based on the OAuth2/OIDC technology which will replace the X509-based security infrastructure used by the grid systems

4.1.3 DIRAC@IN2P3 as service

The DIRAC@IN2P3 project takes also the responsibility for maintenance and operations of the [EGI Workload Manager service](#). The service is hosted at the CC-IN2P3 premises and is operated by a joint team of administrators from all the participating laboratories. The service is used by about 20 scientific communities from various scientific domains and has about 700 registered users. About 12 million computational tasks were executed by the service in 2021. Support for the service users is also an important task of the project to help new communities start using complex computing infrastructures.

In the future, the project will continue developments needed by the communities using DIRAC software and services for their specific workflows. It will also continue to maintain DIRAC services for the users of the EOSC as well as support other user communities.

4.2 The AMI project

AMI is a complete software ecosystem dedicated to scientific metadata. AMI provides ready-to-use solutions to all metadata related tasks, from end users Web applications down to low-level interactions with database systems. Although developed for the ATLAS experiment, AMI is a generic ecosystem easily adaptable and compatible with any existing database. Its main features are :

- Modern, modular Web framework to allow end-users to interact with metadata (query, display, update...)
- REST API to allow programatic interactions with the metadata from any programming language
- Task server to allow automatic aggregation and curation of data. The tasks workflow can be designed and steered with graphical interfaces based on Node-RED
- Full implementation of the Metadata Query Language (MQL), a high-level, user-oriented powerful SQL-like language which allows to write queries independent on the underlying data storage details
- Complete identification and authentication system allowing access to large ensembles of users with various roles
- Driver system allowing to connect to any database system (including existing databases)

The full suite, or part of it (server, task processing servers, Web framework, ...) can easily be installed from Docker Hub images. AMI is developed at LPSC laboratory in Grenoble. More detailed informations can be found at the project [website](#).

The project was born in 2000 in the context of ATLAS. A major release happened in 2015 (v2). Even if ATLAS remains the largest users community with 2300 active users, other experiments in different scientific domains have already adopted or they are evaluating AMI. Among these, we mention : [NIKA2](#) in radio astronomy, n2EDM (the new spectrometer for the neutron electric dipole moment search at PSI), SuperNemo (currently under evaluation). Finally in a totally different context, AMI has also been adopted for the LPSC administration forms.

4.2.1 Status and plans

The development team is entirely at LPSC laboratory in Grenoble and is composed of 3 research engineers, a student and a scientific coordinator for a total of 3 FTE. These resources are only just enough for the development and maintenance of the project, indeed in parallel to the development tasks the team is also in charge of operating the AMI instance for ATLAS.

Future plans consist in continuing the development and maintenance effort to cope with the needs of ATLAS for the Run3 and to prepare the HL-LHC. Beyond ATLAS needs, the challenge for the project is to further promote AMI to new users communities. As a result, on

one side the different scientific communities would benefit of a mature easy-to-deploy product for the metadata handling, and on the other side they could eventually contribute to the AMI development. The aim is to attract new developers from others IN2P3 laboratories to make AMI a national project at IN2P3 level or even at international level. Today, no concurrent projects exist, so that there is a great opportunity for AMI to be widely adopted by different scientific communities.

5 Conclusions

In this document, we have reviewed the main software projects in which IN2P3 has a strong participation. In two cases, i.e. KaliVeda and AMI projects, IN2P3 is the unique contributor, despite the users communities of these products are large and spread at international level. All IN2P3 scientific domains are covered by these projects, from particle physics to nuclear and astroparticle physics, including bio-medical physics.

As a first general remark, we notice that the approach of all the presented projects, is to develop common software solutions usable either by several groups within a single collaboration (e.g. AGATA and KaliVeda) either across multiple collaborations and scientific domains. As a consequence, the users communities of all these projects are very large and spread worldwide. Beyond the purely development task, some of the presented projects are also very active in dissemination activities, like the organization of tutorials for users and developers, which also contributes to the success of these projects.

From a technical point of view, some common features emerge from this review. All the Monte Carlo simulation softwares are based on Geant4, which favours a good synergy among these projects. The ROOT framework is still widely used by several simulation and data processing projects (Geant4, Geant4-DNA, GATE, nptool, AGATA software and KaliVeda).

Not surprisingly the most common language used by this category of softwares is C++, but the general trend is to provide python bindings for the high-level interfaces exposed to users. python is also the language used by Gammapy, which takes advantage of the rich python ecosystem for the analysis of astrophysical data (astropy, numpy, scipy,iminuit, ...). Finally, in the workload, data and metadata management category, AMI is mostly written in Javascript (e.g. for its Web framework), while DIRAC is written in python in a Service Oriented Architecture.

Concerning the R&D software activities, a number of projects (e.g. GATE, ACTS, AGATA) are exploring the application of machine learning techniques, which are seen as promising to improve overall performances. In parallel to the code optimization effort to better exploit modern CPU, an R&D activity on-going within several projects (e.g. Geant4, Smilei, ACTS, AGATA) concerns the development of adapted algorithms to exploit GPU accelerators. The use of alternative architectures of type ARM/RISC is also considered within the Smilei project, while AGATA investigates the integration of FPGA in its data processing workflow. Finally, in the context of the exascale, the Smilei project also studies the introduction of more adapted numerical methods to improve the accuracy and performance of the simulations.