PHYSICS VALIDATION OF THE LHC SOFTWARE

F. Gianotti*, CERN, Geneva, Switzerland

Abstract

The LHC software will be confronted to unprecedented challenges as soon as data taking will start. This paper summarizes the main software requirements coming from the LHC environment, detectors and physics, and presents some examples of software validation studies and performances achieved to-date [1]. The present status of the main software components, in terms of the path accomplished so far, is also discussed. This paper reflects more the point of view of an LHC physicist and end-user than of a software expert.

REQUIREMENTS AND CHALLENGES FOR THE LHC SOFTWARE

The LHC is an unprecedented project in terms of physics goals, required performance of the detectors and of the trigger systems, and complexity of the experimental environment. The main consequences for the software are:

- The experiments are going to explore a very broad particle energy range, from a fraction of a GeV up to several TeV. Accurate simulation and efficient reconstruction over this range are needed.
- The environment will be much harsher than at previous machines. Central heavy-ion collisions will produce about 10000 charged particles per event in the ALICE detector. When running at the LHC design luminosity of $L = 10^{34}$ cm^{-2} s^{-1}, about 20 pp collisions are expected to occur simultaneously in the ATLAS and CMS detectors at each crossing of the two beams (every 25 ns), giving rise to the so-called "event pile-up". This environment must be simulated in all its complexity, and the reconstruction must be able to extract the interesting physics objects (e.g. a Higgs boson decaying into two photons) from the potential confusion generated by the accompanying pile-up particles.
- The interaction rates will be huge (see Table 1), e.g. 10^9 pp interactions per second in ATLAS and CMS at design luminosity, whereas the affordable rate-to-storage is about 100 Hz. As a consequence, the experiment high-level triggers (HLT), which will be based mostly on software selections, will have to provide rejection factors of order 10^3. The online software must therefore be fast (event processing times in the range 1 ms to 1 s are envisaged, depending on the trigger level), and robust and reliable, since an event rejected at the trigger level is lost forever.
- Unprecedented detectors are being built, in terms of variety of technologies, complexity, and performance. The measurement precisions of these devices are expected to be between the permil and the percent level in most cases. The software should not limit this performance, which implies for instance the implementation of a very accurate detector description (a visual example is shown in Fig. 1), including all the active elements, but also services, cables, support structures, cracks, etc.
- Finally, and most importantly, the LHC has very ambitious physics goals. The experiments will explore an unknown territory up to the multi-TeV scale, and must be prepared to cope with a large number of topologies and to extract rare New Physics processes from huge backgrounds (as an example, the production cross-section for a Higgs boson of mass 150 GeV is ten orders of magnitude smaller than the total pp cross-section). To meet these goals, a very precise understanding of the detector performance and of the backgrounds to New Physics is needed, and good control of all possible systematic effects. This can only be achieved if many Monte Carlo generators are available, different levels of detector simulations (fast, full, parametrized), and several complementary recon-

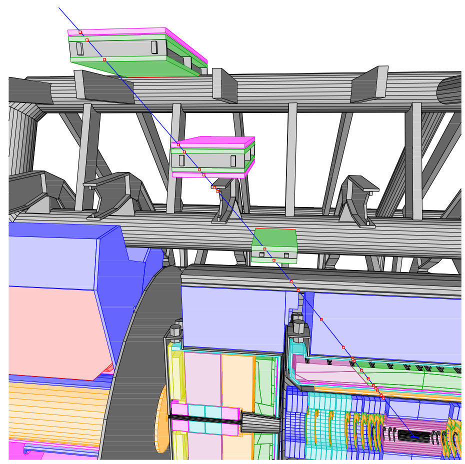

Figure 1: Three-dimensional display of a muon track traversing the ATLAS detector, showing the details of the geometrical description [1].

* fabiola.gianotti@cern.ch

[1] Only physics-related aspects of software and computing, and no technical issues, are addressed in this paper.

Table 1: Expected interaction rates, input rates to the high-level triggers, rates to storage and event sizes for the four LHC experiments. Luminosities are given in units of cm^{-2} s^{-1}.

	ATLAS, CMS pp, $L = 10^{34}$	**LHCb** pp, $L = 2 \times 10^{32}$	**ALICE** central Pb-Pb, $L = 10^{27}$
Interaction rate	10^9 Hz	10^7 Hz	8 kHz
Input rate to HLT	$\sim$100 kHz	10^6 Hz	$<$ 1 kHz
Rate to storage	100-200 Hz	$\sim$200 Hz	$\sim$50 Hz
Event size	1-2 MB	$\sim$100 kB	$\sim$25 MB

struction algorithms. The software, in particular the framework, must be modular and flexible enough to allow the seamless use of this variety of tools.

In summary, two main requirements for the LHC software can be outlined. First, it must cope with the above challenges, i.e. it should not become the limiting factor to the trigger capabilities of the experiments, to the detector performance nor to the LHC physics reach. Second, in spite of these challenges and of the unavoidable underlying complexity, the software must be easy-to-use, stable and highly interactive. Only in this way will each one of the about 4000 LHC physicists be able, irrespectively of his/her geographical location, age, and scientific profile, to run the experiment software, modify the part of it closest to physics (e.g. the reconstruction algorithms), analyze the data, and extract the physics results.

EXAMPLES OF SOFTWARE VALIDATION AND PERFORMANCE

Some examples of validation of software components against the main physics requirements, and of the performance achieved to-date, are presented below.

Simulation

In general, particle interactions in the LHC detectors must be described with precisions between a few permil and a few percent in most cases. These and other simulation requirements are discussed in detail in Ref. [2], and only a few examples are reported here.

The simulation packages must be able to model the individual microscopic collisions produced by the incident particles in the thin layers (e.g. $\sim$ 300 μm Si sensors) of the LHC trackers, including all relevant physics processes (ionization, δ-ray emission, multiple scattering, bremsstrahlung, hadronic interactions, synchrotron radiation, etc.) down to energies as low as $\sim$ 1 keV ($\sim$ 10 eV) in Silicon (gas). This is needed for a reliable estimate of the detector performance and ageing, and for the determination of the reconstruction efficiency.

Concerning the muon spectrometers, one of the issues is that high-energy muons (in the few hundred GeV range) sometimes undergo catastrophic energy losses in the upstream calorimeters, and produce showers of hits in the muon chambers. These processes should also be simulated at the percent level, since they have an impact on the detector performance and reconstruction efficiency.

Concerning the calorimeters, their response can be studied and calibrated with data (test-beam data and physics samples to be collected at the LHC) up to a few hundred GeV. Beyond this range, i.e. in the TeV region (where New Physics might be observed ...), the experiments will rely on the prediction of the simulation, which will have been validated with data at lower energies. An illustrative example of the impact of possible simulation inadequacies on the LHC physics potential is quark Compositeness. One of the main questions to be addressed at the LHC is: are quarks elementary or are they composite particles ? If quarks are composite, new interactions (in addition to QCD) are expected to contribute to the production of high-p_T jets, leading to an excess of di-jet events in the TeV region compared to the Standard Model expectation. The size of this effect depends on the scale Λ where the New Physics lies: the higher Λ, the smaller the excess. Thanks to its energy and luminosity, the LHC has the intrinsic power of exploring scales of New Physics Λ up to about 40 TeV, through the above-mentioned processes. A potential limitation to this reach comes from the fact that a fake excess of high-p_T jets can be produced if the calorimeter response to jets is non-linear at the level of 1-2% in the $\sim$3 TeV energy range, and if this non-linearity is not predicted by the simulation. Such a fake excess could mimic a Compositeness scale $\Lambda \simeq 30$ TeV. This shows that the LHC intrinsic reach of $\sim$40 TeV could be limited to less than 30 TeV by the inadequacy of the simulation tools. To eliminate this problem, the simulation must reproduce the ratio of the calorimeter response to electrons and pions (the so-called "e/π ratio"), which governs the calorimeter linearity, to a few percent.

Since several years, a lot of effort has been invested in comparing the main simulation packages (Geant4 [3] and Fluka [4]) with test-beam data, and on improving the physics content of the simulation to reproduce the data with increasing accuracy. As an example of achieved performance, Fig. 2 shows the energy spectrum of 300 GeV test-beam muons measured with a prototype of the ATLAS Tile hadron calorimeter. Fluka describes well the data, in particular the high-energy tails due to catastrophic losses. As another example, Fig. 3 demonstrates that the Geant4 simulation is able to reproduce the e/π ratio of three different hadron calorimeters (the ATLAS Tile Fe-scintillator calorimeter, the ATLAS HEC Cu-liquid argon calorimeter,

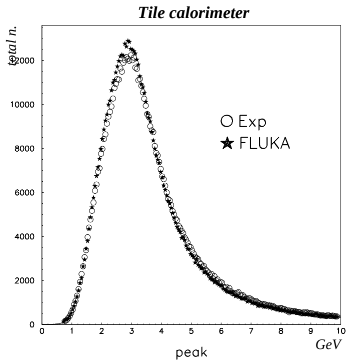

Figure 2: Absolute comparison of the energy spectra of 300 GeV muons incident on a prototype of the ATLAS Tile calorimeter, as measured with test-beam data (open symbols) and as predicted by a Fluka simulation of the test-beam set-up (stars). From Ref. [5].

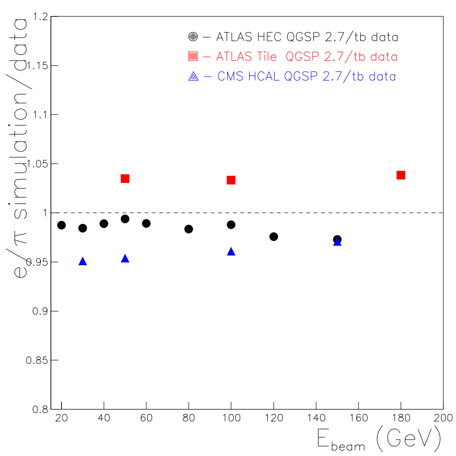

Figure 3: The ratio between the Geant4 prediction (QGSP 2.7 physics list) and the test-beam measurements of the e/π ratio, as a function of energy, for three LHC calorimeters (see text). From Ref. [6].

and the CMS HCAL brass-scintillator calorimeter) to better than 5%, which is close to the goal dictated by physics, as discussed above.

Robustness and CPU performance are also crucial issues for the simulation engines, given the huge number of physics processes that need to be simulated with accuracy down to very low energies, and the millions of volumes used to describe the complex detectors of the four LHC experiments. These aspects have been addressed over the last couple of years by the experiment Data Challenges [7].

Reconstruction

Three cases are discussed below, which are particularly relevant for the LHC operation and physics goals, and represent new challenges compared to previous machines.

The first example is the CMS trigger, which is divided into two levels: level one, hardware-based, will reduce the initial pp interaction rate of 1 GHz to about 100 kHz; the high-level triggers [8], software-based, will use the offline framework and reconstruction code and will have to provide an additional event rejection of about 1000 (see Table 1). Such a huge rejection factor to be achieved with a purely software trigger is due to the CMS choice of operating without a dedicated second-level trigger, unlike in previous hadron collider experiments and in ATLAS. This trigger scheme entails very demanding software requirements in terms of CPU (the average time available at the HLT to process one event and take a decision is $\sim$40 ms) and of code robustness and reliability. In order to meet these requirements, in particular the CPU constraints, a strategy based on data loading on demand, regional reconstruction and partial reconstruction has been adopted.

This can be illustrated with one example, namely the possibility of triggering on b-quark jets, which is a difficult case and also one very relevant to physics since the third fermion family could play a special rôle in New Physics. The distinctive feature of jets originating from the fragmentation of b-quarks is that these jets are produced at a distance of typically a few millimetres from the primary pp interaction vertex, due to the long lifetimes of B-mesons. Therefore, the HLT reconstruction algorithm first determines the position of the primary vertex using track segments from the Pixel detector only; then jet tracks are reconstructed only inside a cone around the jet direction provided by the level-one trigger; for each track, reconstruction is stopped as soon as seven hits are found on the trajectory (instead of ten hits as required offline), since this is sufficient to keep the rate of fake tracks low enough; finally, if there are at least two tracks inside the jet cone pointing to a secondary vertex, the jet is classified as a b-jet. This reconstruction procedure takes today a total time of 300 ms per event at low luminosity ($L = 10^{33}$ cm^{-2} s^{-1}), which, extrapolated to the year 2007 CPU power, should meet the requirement of about 40 ms per event. The achieved physics performance of the online algorithm is very similar to the offline performance, as shown in Fig. 4. These results are very encouraging, although they have been obtained with a perfect detector. The next step is therefore to test the robustness of the fast and partial HLT reconstruction against alignment problems, readout inefficiencies, dead channels, etc.

Reconstruction of very high-energy muons (in the TeV region) is another difficult case which has not been addressed by previous experiments. One of the LHC physics goals is to look for new particles in the TeV mass range decaying into muons, such as a possible additional gauge boson $Z' \rightarrow \mu\mu$. These searches require a muon reconstruction efficiency of larger than 90%, since the expected signal rates are tiny in most theories, and a muon momentum resolution of better than 10%, in order to observe a relatively narrow signal peak on top of the Standard Model background. In the ATLAS muon spectrometer (discussed here as an example) the sagitta, i.e. the deviation of the

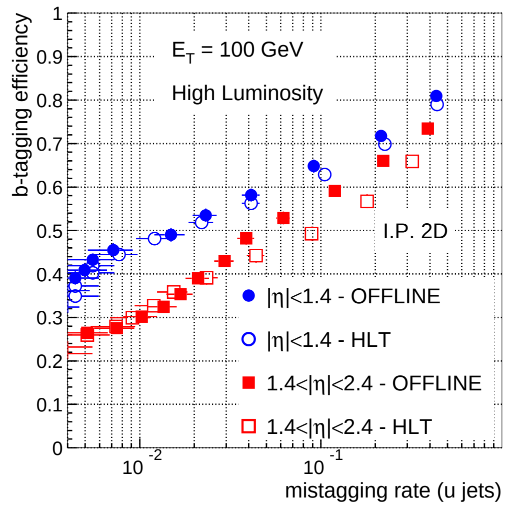

Figure 4: Expected efficiency for tagging b-quark jets of E_T=100 GeV, as a function of the fraction of fakes from u-quark jets, as obtained from a full simulation of the CMS tracker at design luminosity in the barrel (circles) and end-cap (squares) regions. The closed (open) symbols show the performance of the offline (online) b-tagging algorithm [8].

muon trajectory from a straight track, is only about 500 μm at 1 TeV. A momentum resolution of about 10% demands therefore a measurement of the sagitta to the challenging precision of $\sim$50 μm.

Four main ingredients, which all have important implications on the software required functionalities, contribute to achieve the needed performance. First, an accurate description (to better than 10%) of all the materials traversed by muons from the production vertex to the external chambers (see Fig. 1) and of the energy losses along the muon path is necessary, because these losses affect the muon trajectory and momentum. Second, since the muon spectrometer is designed to provide the required 50 μm sagitta resolution, alignment errors should be kept below $\sim$30 μm to be negligible compared to the intrinsic detector performance. Therefore about 10^5 probes will be used to monitor the parameters of the 2000 chambers (temperature, gas pressure, etc.), whose measurements will have to be stored in the database, regularly updated, and used in the alignment procedure every few hours. Furthermore, since the magnetic field is highly non-uniform in the open air-core geometry of the ATLAS spectrometer, the field map must be very detailed and must be accessed typically 500000 times during the reconstruction of the muon trajectory. Finally, catastrophic energy losses in the calorimeters, which become more and more important at high energy, produce confusion of hits in the spectrometer, which requires robust and sophisticated pattern recognition techniques.

The performance achieved to-date with the detector simulation and reconstruction, shown in Fig. 5, is close to the specifications. In particular, a momentum resolution of better than 10% has been obtained at $\sim$1 TeV, and and efficiency of almost 90% (smaller than at lower energies because of the impact of catastrophic energy losses).

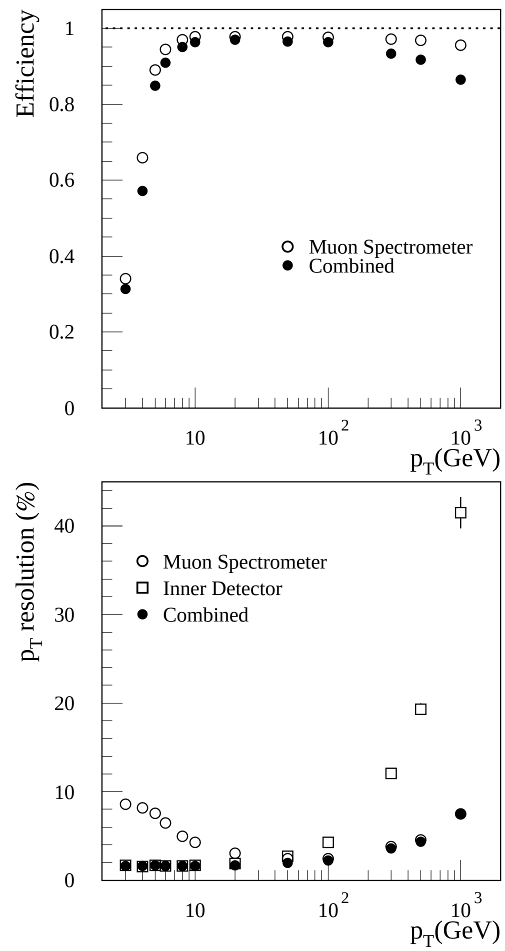

Figure 5: Expected muon reconstruction efficiency (top) and momentum resolution (bottom) in ATLAS as a function of p_T, from full simulation and reconstruction. Shown are the results obtained using the inner detector alone (squares), the muon spectrometer alone (open circles), and both detectors together (closed circles).

In contrast to the previous case, the last example discussed here, i.e. tracking in the ALICE experiment using the Silicon layers and the Time Projection Chamber, addresses the very low-energy region below $\sim$1 GeV and the very high-multiplicity environment expected in central heavy-ion collisions at the LHC. The robustness of the ALICE reconstruction algorithm [9] can be appreciated from Fig. 6, which shows that the performance, a track reconstruction efficiency of more than 90% for a rate of fake tracks of a few percent, is not significantly deteriorated if the particle multiplicity is increased from the expected 6000 charged particles per rapidity unit to 8000 charged particles per rapidity unit. One of the key features for a pattern recognition operating in such a busy environment is to be able to deal with extended cluster shapes in the various detector layers, which are expected from track overlaps.

but necessary since the LHC project will have a lifetime of about 15 years.

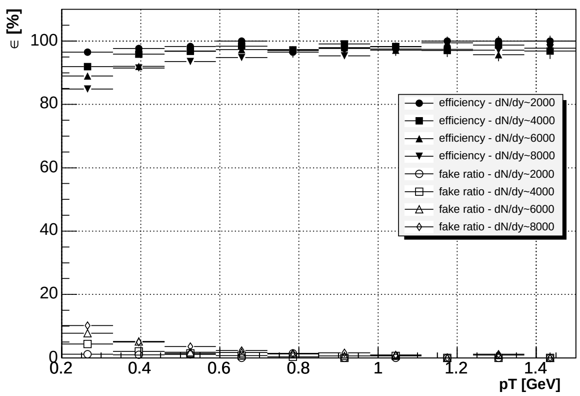

Figure 6: Track reconstruction efficiency (closed symbols) and fraction of fake tracks (open symbols), as a function of the track transverse momentum, as obtained from a full simulation and reconstruction of the ALICE Silicon and TPC detectors for events with various particle multiplicities [9].

Other areas

Examples of additional issues relevant for LHC physics are:

- Analysis environment. There is wide consensus that modern analysis frameworks are much more than the tools to plot an histogram. Rather, they must provide the full chain from accessing the data at the various levels of the hierarchy (raw data, reconstructed data, Analysis Object Data, calibration data, simulated data, etc.), to running any reconstruction, trigger, calibration algorithm, up to the extraction of the final physics results. Interactivity in all these operations and simple interfaces are the key requirements for efficient physics analysis at the LHC, also because this is the part of the software to which the end-users will be mostly exposed.
  The two LHC experiments most advanced in this area, and indeed very close to what the users require, are LHCb and ALICE. In the LHCb PYTHON-based interactive framework [10], called Bender, the user can navigate the data structure, display an event, run the reconstruction and analysis algorithms, and plot the final results using several histogramming tools (see Fig. 7). Similar functionalities are provided by the ALICE ROOT-based environment [11], where the users can browse not only the data but also the status of the various tasks, and therefore can check what their application (e.g. a sequence of reconstruction algorithms running interactively) is doing.
- Schema evolution. With the software released in e.g. April 2020, it should be possible to access and analyze in a transparent way the data recorded in e.g. June 2009. Implementation of this requirement is highly non-straightforward from the technical point of view, but necessary since the LHC project will have a lifetime of about 15 years.
- GRID. The distributed computing infrastructure of the LHC era can be a spectacular technical and social success if it will enable everybody to do analysis anywhere anytime, thereby involving in the exploitation of LHC physics also people from countries historically at the periphery of the scientific community. However, it could also be a major failure if it will not work as expected, thus becoming an obstacle to efficient analysis rather than an aid, and slowing down the delivery of physics results.
  From the end-user point of view three requirements are particularly relevant: transparent access, i.e. physicists must be screened from the GRID complexity; technical support from experts at the various Tiers at any time over the whole LHC lifetime; availability of simple monitoring and diagnostic tools, enabling the user to check the status of a job, understand the reasons for an execution failure, etc.

PRESENT STATUS OF THE LHC SOFTWARE

At about three years from the LHC start-up, it is legitimate to address the question of where do we stand today with the LHC software. Figure 8 shows, for the various software components, the fraction of accomplished path, from a rough estimate based on inputs from the four LHC experiments. The general picture is fairly positive:

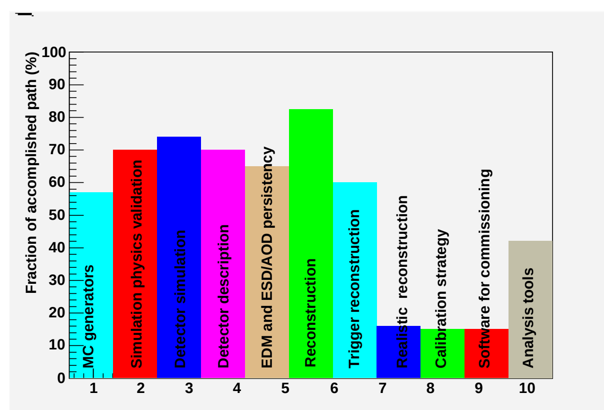

Figure 8: For several components of the LHC software, an approximate estimate of the fraction of work already completed, averaged over the four LHC experiments.

the basic infrastructure is in place, and more than 50% of the work has been done in many areas. Also, fluctuations from experiment to experiment are quite small in most cases. However, three aspects are still at an infancy stage in all experiments. Realistic detectors, i.e. detectors with dead channels, missing high-voltage connections, mis-calibrations, mis-alignments, etc., have not been implemented in the software yet; in particular, the reconstruc-

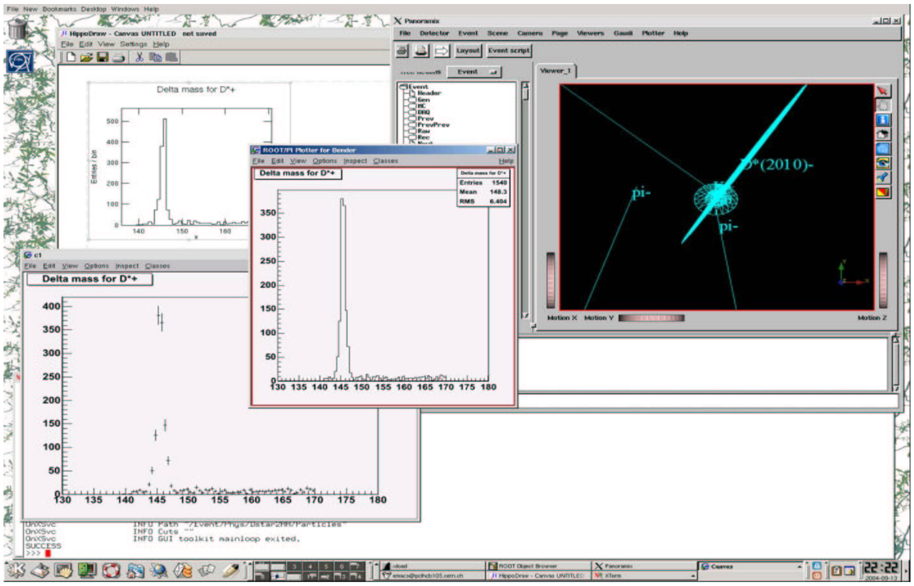

Figure 7: An example of analysis and event visualization with Bender, the LHCb PYTHON-based analysis framework [10].

tion algorithms assume perfect detectors in most cases. Furthermore, a clear calibration strategy is still missing: will the detector calibrations be performed at the HLT or in the Tier0 or in both, which data streams will be used for this purpose, how many times these streams will need to be reprocessed at the raw-data level (which may have implications on the Computing Models) ? Finally, the software for commissioning the detector, with e.g. cosmic muons and beam-halo muons, has not been developed yet, in spite of the fact that these are the first data samples that the experiments are going to record.

CONCLUSIONS

In about three years from now, with the advent of the LHC operation, particle physics will enter a new epoch, hopefully the most glorious and rich in discoveries of its history. It would be a major failure if these prospects were limited by inadequate software tools and computing infrastructure. It would also be a major failure if, because of the software complexity, only a small fraction of the LHC community were able to perform data analysis.

Enormous progresses have been made over the last years in the development of the experiment and LCG software, and physics validation studies carried out in several areas have given very encouraging results, indicating close-to-adequate performance in many cases. However, a lot remains to be done to meet the two main requirements mentioned above. As an example, the LHC software and Computing Models have been developed so far in view of the steady-state LHC operation, which may be reached only in the years 2009-2010. In contrast, at the beginning of data taking the software will be confronted with the most atypical and difficult situations, and a lot of flexibility will be needed to cope with them. These difficulties potentially include non-calibrated and non-aligned detectors with all sorts of problems, machine backgrounds, higher-than-expected trigger rates, the need for fast and frequent reprocessing of part of the data for calibration purposes, not to mention a few thousand physicists accessing simultaneously the computing infrastructure (databases, GRID) and using the software. It is therefore time for the software and computing to address the initial phases of the LHC operation, not to hinder the fast delivery of physics results and postpone a possible early discovery.

REFERENCES

[1] ATLAS Collaboration, “Detector and physics performance Technical Design Report”, CERN/LHCC/99-14.

[2] A. De Roeck, F. Gianotti, A. Morsch and W. Pokorski, “Simulation physics requirements from the LHC experiments”, CERN-LCGAPP-2004-02 http://lcgapp.cern.ch/project/mgmt/doc.html.

[3] S. Agostinelli *et al.* [Geant4 Collaboration], Nucl. Instrum. Meth. A **506**, 250 (2003).

[4] A. Fasso’, A. Ferrari, J. Ranft and P. Sala, in Proceedings of the Monte Carlo 2000 Conference, Lisbon, Springer-Verlag (2001), 159-164. See also the Fluka WEB page: www.fluka.org.

[5] M. Antonelli, G. Battistoni, A. Ferrari and P.R. Sala, “Study of Radiative Muon Interactions at 300 GeV”, in Proceedings of the VI International Conference on Calorimetry in High Energy Physics, Frascati, June 8–14, 1996, A. Antonelli, S. Bianco, A. Calcaterra and F.L. Fabbri eds., Frascati Physics Series, 561-570 (1997).

[6] C. Alexa *et al.*, “Geant4 hadronic physics validation with LHC test-beam data: first conclusions”, CERN-LCGAPP-2004-10, http://lcgapp.cern.ch/project/mgmt/doc.html.

[7] D. Stickland, these Proceedings.

[8] CMS Collaboration, "Data Acquisition and High-Level Trigger Technical Design Report", CERN/LHCC/2002-26.

[9] M. Ivanov, these Proceedings.

[10] P. Mato, these Proceedings.

[11] F. Carminati, these Proceedings.