

The CMS Simulation Software

S. Abdouline, *FNAL, USA*, N. Amapane, *Universit of Torino, Italy*, F. Ambroglini, *INFN, Perugia, Italy*, P. Arce, *CIEMAT, Madrid, Spain*, S. Ashby, *CERN, Geneva, Switzerland*, S. Banerjee, *Tata Institute, Mumbai, India*, U. Berthon, *IN2P3-CRNS, Palaiseau, France*, T. Boccali, *SNS, Pisa, Italy*, M. Case, *UC Davis, USA*, C. Charlor, *IN2P3-CRNS, Palaiseau, France*, H. Cheung, *FNAL, USA*, T. Cox, *US DAVIS, USA*, N. Darmenov, *INRNE, Sofia, Bulgaria and CERN, Geneva, Switzerland*, A. De Roeck, *CERN, Geneva, Switzerland*, S. Dutta, *University of Pisa, Italy*, P. Elmer, *Princeton University, USA*, D. Elvira, *FNAL, USA*, F. Ferro, *INFN, Genova, Italy*, D. Evans, *FNAL, USA*, E. Garcia, *UI Chicago, USA*, P. Govoni, *INFN Milano-Bicocca, Italy*, O. Gutsche, *FNAL, USA*, X. Huang, *University of Puerto Rico (Mayagues), USA*, V. Innocente, *CERN, Geneva, Switzerland*, C. Jones, *Cornell University, USA*, A. Kharchilava, *SUNY Buffalo, USA*, P. Katsas, *University of Athens, Greece*, J. Kowalkowski, *FNAL, USA*, K. Lassila-Perini, *HIP, Helsinki, Finland*, S. Magni, *INFN Milano-Bicocca, Italy*, P. Meridiani, *INFN, Roma, Italy*, F. Moortgat, *CERN, Geneva, Switzerland*, A. Nikitenko, *Imperial College, London, UK*, D. Onoprienko, *KSU, USA*, I. Osborne, *Northeastern University, USA*, M. Paterno, *FNAL, USA*, M. Pioppi, *INFN, Perugia, Italy*, S. Piperov, *INRNE, Sofia, Bulgaria and FNAL, USA*, A. Rizzi, *SNS, Pisa, Italy*, E. Sexton-Kennedy, *FNAL, USA*, E. Shabalina, *UI Chicago, USA*, L. Silvestris, *INFN, Bari, Italy*, M. Spiropulu, *CERN, Geneva, Switzerland*, M. Stavrianakou, *FNAL, USA*, M. Strang, *SUNY Buffalo, USA*, W. Tanenbaum, *FNAL, USA*, J. Weng, *Karlsruhe University, Germany*, R. Wilkinson, *CALTECH, Pasadena, USA*, A. Yagil, *FNAL, USA*, J. Yarba, *FNAL, USA*, A. Zhokin, *ITEP, Moscow, Russia*

Abstract—In this paper we present the features and the expected performance of the re-designed CMS simulation software, as well as the experience from the migration process. Today, the CMS simulation simulation suite is based on the two principal components - Geant4 detector simulation toolkit and the new CMS offline Framework and Event Data Model. The simulation chain includes event generation, detector simulation, and digitization steps. With Geant4, we employ the full set of electromagnetic and hadronic physics processes and detailed particle tracking in the 4 Tesla magnetic field. The Framework provides "action on demand" mechanisms, to allow users to load dynamically the desired modules and to configure and tune the final application at the run time. The simulation suite is used to model the complete central CMS detector (over 1 million of geometrical volumes) and the forward systems, such as Castor calorimeter and Zero Degree Calorimeter, the Totem telescopes, Roman Pots, and the Luminosity Monitor. The designs also previews the use of the electromagnetic and hadronic showers parametrization, instead of full modelling of high energy particles passage through a complex hierarchy of volumes and materials, allowing significant gain in speed while tuning the simulation to test beam and collider data. Physics simulation has been extensively validated by comparison with test beam data and previous simulation results. The redesigned and upgraded simulation software was exercised for performance and robustness tests. It went into Production in July 2006, running in the US and EU grids, and has since delivered about 60 millions of events.

I. INTRODUCTION

OVER the past few years the CMS simulation software has undergone several critical changes, such as a replacement

of the modeling engine or, most recently, a major redesign of the core infrastructure.

The suite that manages both central and forward CMS detectors, as well as several test beam prototypes, is based on the Geant4 [1] toolkit and the new CMS offline Framework and Event Data Model. The implementation includes, on one end, the interface to the generator information, and on the other, sensitive detectors behaviour, hit collections, and modeling of the readout electronics response.

The software is always a subject of a very strict validation procedure, with the focus on the physics performance, as well as a number of technical aspects. The procedure includes comparison with test beam data and previous simulation results, and has recently been brought to a higher level, by commissioning of the automated physics and detector software validation suite.

II. THE CMS OFFLINE FRAMEWORK

The development of the new CMS Offline Framework and Event Data Model (EDM) has emerged from the need to better meet the requirements of the collaboration for production performance and features.

The Framework design ensures overall consistency, flexibility of use and extensibility, minimizes overheads from development and support of the common-use elements, and facilitates maintenance.

Overall, the job of the Framework is to allow combining in one application a variety of steps of the event processing chain,

starting from event generation, through detector simulation, and all the way up to reconstruction and analysis.

The Framework relies on the concept of an event processing module, which encapsulates well defined functionalities (ED-Producer) and is allowed to add to the event the persistent elements that compose the event data (EDProducts). Roughly speaking, each step in the event processing chain is represented by one or several such modules, adding corresponding data to the event. It should be noted that these modules must not interact directly with each other, but only through an event. The flow of the data through an event is defined almost entirely and almost “freely” by a user who composes a particular application; the user, however, must remember that the presence in the event of the data required by a specific module is a mandatory.

The event processing chain controlled by the Framework, may, of course, be cut between any two steps, and the state of the event is saved.

The natural persistent form of the event is a Root [2] tree, which allows straightforward access for monitoring and analysis.

One of the critical features of the Framework is that it ensures immutability of the already-existing event data and event reproducibility through recording the provenance of each component involved in a particular job.

The Framework machinery includes such important components of the EventProcessor, EventSetupProvider, ScheduleBuilder and ScheduleExecutor.

The EventProcessor handles the initialization and the event processing loop. An application execution begins with the EventProcessor to parse the job configuration information specified by the user. The action then moves to the ScheduleBuilder which creates user-scheduled modules. The EventSetupProvider allows access, as needed, to the common-use (non-event) data services; examples of such data services can be Geometry and/or Magnetic Field configuration. The ScheduleExecutor then runs the user-defined sequence of the event-based modules.

Speaking of the event processing, we must mention on-demand processing, referred to as unscheduled as opposed to the scheduled scheme described above. An unscheduled application is configured by specifying a selection of independent top-level EDProducts to be written out, or a selection of independent modules to be run, or some combination of the above, as well as the menu of the EDProducers that are known only to a specific module’s registry. A signaling mechanism based on the dispatcher-observer pattern, may be used to trigger an unscheduled application component, such as user monitoring action. For example, this machinery is largely used in the release validation procedure, when one needs to access quantities that are available only within detector simulation stage.

Regarding physics events generation, we must mention that CMS employs a variety of event generators, such as Pythia [3], Herwig [4], Alpgen [5], etc., all fully interfaced to the CMS Framework. The generator places information into the event in HepMC format [6], from where it may be accessed by other

modules for further processing. This may be produced and used on the fly, or it may be read from an external, pre-generated Root file and placed into the event.

Concerning the detector and magnetic field description, we utilize the Detector Description services from the CMS Geometry and Magnetic Field packages, that rely on the DDD/XML description files. These records may be used at the simulation step, as well as at other steps.

Visualization and event display are handled by separate packages, interfaced to the core and detector-specific simulation systems.

III. INTERFACE TO THE GEANT4 TOOLKIT

The interface application between Geant4 and CMS Framework consists of the two components - the Framework-based event processing module and a customized Geant4-based RunManager.

The Framework-based module defines what data to place into the event and communicates to the RunManager for initialization, event processing, and termination of the simulation application.

In its turn, the RunManager instantiates the Geant4 kernel and takes care of several Geant4-specific actions.

It accesses and handles Geometry and Magnetic Field records and assembles sensitive detectors.

One of the duties of the RunManager is to retrieve the generator information from the event and to apply user-specified selection criteria, such as η - and ϕ -range and/or pt -range of the generator’s particles. Those generator’s particles that survive the selection criteria are converted, by the RunManager, into Geant4-specific G4Event format.

The RunManager composes physics lists and handles storage and retrieval of the cross-section tables as built for a given detector geometry and physics model. From the variety of the physics models (lists) provided by the Geant4 toolkit several most widely-used are also interfaced to the CMS simulation machinery. The selection includes LHEP, QGSP, QGSC and FTFP, plus simplified physics subset configurations for debugging and material budget studies. The infrastructure is done in such a way that the choice of the physics list and of the production cuts is configurable at run time, as all other options. The design also features provisions for adding shower parametrizations, that would allow to tune the simulation to the real CMS data as well as improve the computational efficiency of the package.

The RunManager provides interfaces to such Geant4-specific actions as run, event, stacking, tracking and stepping actions.

It is also responsible for composing the final MonteCarlo truth record that contains the generators particles, with their track, vertices and decay records, plus selected tracks from the Geant4 simulation. These tracks have been selected because they have produced or at least initiated hits in the sensitive detector or because they have been identified as important for the interaction history and the eventual reconstruction of the full event tree.




Fig. 1. The cross-section view of the CMS detector




Fig. 2. Longitudinal view of the CMS detector

IV. THE CMS DETECTOR SIMULATION

As already stated in the Introduction, the CMS detectors include central (Tracker, Calorimeters and Muon subsystems), schematically presented in Fig. 1,2, and forward (CASTOR calorimeter, the Zero Degree Calorimeter (ZDC), Totem telescopes, the Roman Pots and the Luminosity Monitor).

The central detector alone is composed of more than one million of volumes, of many complex shapes.

The detector geometry configuration is interfaced to the Geant4 through the geometry builder machinery, associated with the Detector Description package, which is already mentioned earlier. The geometry builder converts Detector Description solids and materials into their Geant4 counterparts, as well as logical, physical and sensitive volumes. The later is the key concept in the collecting the records of particle interaction with the active components of the detector.

Magnetic field services also include interfacing to the Geant4, in such a way that the field type and configuration of stepping and propagation parameters can be selected at run time.

While implementing the simulation software, great attention has been paid to details that could affect the accuracy of modeling the detector performance in measuring electrons, photons and hadrons. For example, Tracker is the critical region not only due to its physics significance but also because interactions of particles in the material upstream the electromagnetic calorimeter directly affects the quality of reconstructing of the electrons and photons. This imposed the requirement of the very detailed description of all the active and passive volumes in the Tracker

region, for proper account of the material budget. Another example would be the complex geometry of the electromagnetic calorimeter that required careful approach, in order to allow simulation of the spacial characteristics of the electromagnetic showers as precise as possible. On the contrast, in the hadronic calorimeter the fine attention to the geometrical details may be somewhat less important as compared to the precision of modeling interactions of hadron with matter. However, this subject will be addressed in the following section.

A. Collections of the Simulated Hits

Common data formats for the simulated hit EDProducts are used for the two major types of detectors for tracking and calorimetry. The former includes not only Pixel and Silicon Strip Tracker devices, but also Muon subsystems.

Hit processing and collection are handled by the detector subsystems, with interfaces to the Geant4 methods invoked at the stepping and end of event actions.

B. Event Mixing

The event mixing concept, also referred to as pile-up, takes into account the fact that the LHC machine will produce a certain number of minimum bias interactions per bunch crossing on top of the trigger event. The average numbers of those minimum bias interactions are expected to be $\bar{3}$ in the “low luminosity” and $\bar{24}$ in the “high luminosity” phases of the LHC operation.

In addition to this so called in-time pile-up, it is necessary to account for pile-up coming from bunch crossings before and after the trigger event crossing. The number of crossings to include before and after the nominal one depends on the front end time response of the different detector components.

Since the pile-up depends on the machine luminosity, the pile-up events are simulated separately from the physics events. The simulation outputs can be merged together, at the level of simulated hits, according to the desired luminosity and conditions. This allows to reuse both types of samples and greatly reduces the amount of time required for the simulation of the effect.

The mixing is done in a separate step in the event processing chain, by a dedicated module whose task is to merge an event from a primary stream with a number of events from a secondary stream and to compose the event data element (EDProduct) called CrossingFrame. The CrossingFrame in its turn serves as the input to the digitization.

C. Modeling of the Electronics Response

The digitization is the process of translating the simulated hits into the electronic readouts response. It starts from simulated positions and energy losses in the sensitive detector elements and produces an output as close as possible to the CMS data as acquired in the DAQ systems.

It should be mentioned that digitization of each detector component is done by a dedicated software module in a

separated production step and is completely independent from other such modules.

In the tracing detector (Pixels and Silicon Strips) the electronic signal modeling takes into account Landau fluctuations, drift and diffusion of the energy loss distribution within the active element, as well as noise and coupling between channels.

Modeling digits in the electromagnetic calorimeters includes efficiency and non-uniformity of the light collection in the crystals. In the hadronic calorimeter digitization the number of photoelectrons is calculated based on the simulated energy deposition, with fluctuations, internal non-uniformities and noise taken into account. The energy depositions from up to 5 crossing prior and 3 subsequent crossings can be accounted for in this procedure.

In the Muon Drift Tube (DT) subsystem, modeling the electronic readout aims on obtaining the 4 ns resolution (about 220 microns intrinsic cell resolution), as measured in the test beam data. The algorithm relies on the particle direction and impact position with respect to the sense wire, and on the residual magnetic field effects.

The digitization of the Cathod Strip Chamber (CSC) subsystem models the response of the ADCs and discriminators connected to the strips and wires. Parametrizations of the amplifier and shaper response are convoluted with the ion drift collection time. Cross-talk, noise and readout dead time (200 ns) are taken into account.

The Resistive Plate Chamber (RPC) subsystem response to take place within 20 ns of the passage of a charged particle through the detector with a 3 ns Gaussian jitter, which also accounts for the contribution from the front-end electronics and the cables to the link board.

V. THE CMS DETECTOR SIMULATION VALIDATION

A. Comparison with the Test Beam Data

Simulation plays the key role in the detector design optimization, whose performance in the turn defines the systematic uncertainties associated with particles discovery, mass and cross section measurements.

Thus, it is essential to evaluate the accuracy of modeling physics processes in the simulation tool and to tune the software to better represent the real data.

For this reason, the Geant4 simulation results are compared to the test beam data, whenever available. At the earlier stage, benchmarking was also done against results from the package based on Geant3 [7].

Among subjects of great interest is the quality of the Geant4 simulation of the electromagnetic and hadronic interactions in the detector's material.

Comparison between Geant4-based simulation and electromagnetic calorimeter test beam data indicate that Geant4 models very well such important quantities as showers fluctuations and shower lateral profile. The later is directly connected with the fraction of incident energy in an ECAL cluster and plays important role in tuning shower reconstruction procedure and identification of photon and electron showers.

However, the effects that dominate the ECAL performance are not involved in the simulation of electromagnetic showers. These effects include electronics noise, photostatistics, non-uniformity of light collection within crystals, channels inter-calibration, the effect of showering in the material upstream the calorimeter, and large sensitivity of the crystal response to the radiation environment. These factors are taken into account in the digitization process, in a form of gross errors.

While the physics involved in modeling high energy electromagnetic interactions is reasonably well understood and implemented in the Geant4 toolkit, the precision of simulating hadronic interactions still presents discrepancies and remain the topic of highest attention.

CMS has taken benefit from its impressive collection of the HCAL test beam data and invested large efforts into studies of comparing energy resolution and linearity, e/π ratio and shower profile, as measured in response to the beams of pions, electrons and muons, against those quantities as simulated by the Geant4, with the use of such physics models as LHEP (parametric) or QGSP (microscopic).

The pion energy resolution and response linearity, as simulated by Geant4, appears to be in a good agreement with the test beam results, provided the large systematic uncertainties of the experimental data. However, results of modeling such intrinsic characteristics of a hadronic shower as longitudinal and transverse profiles depend on the choice of the physics model. Hopefully, further analysis of the test beam data and more detailed comparison versus Geant4 simulation results will allow to refine the quality of the physics models available with the toolkit.

Speaking of the simulation validation, we should also mention the efforts of studying processes that accompany passage of muons through matter, carried out by the Muon subsystem group. The studies included benchmarking of the simulated response of the Muon chambers to a single muon in the range from 10 GeV to 10 TeV versus the results of the Geant3-based package. The outcome of the studies indicated a number of improvements in Geant4 over Geant3, in such areas modeling muon bremsstrahlung, e^+e^- production, and muon-nuclear interactions.

The test beam studies of the Muon chambers also provided great basis for comparison of the simulation results to the real data. The test beam setup consisted of the chambers with or without an iron slab in-between and aimed on investigating the effect of muon showers in the passive material. The analysis involved response to muons in the p_T range from 50 GeV to 300 GeV and indicated that Geant4 slightly underestimates soft delta rays production in the cell volumes, while hard delta rays and electromagnetic showers are simulated correctly.

Future studies will involve analysis of the data from the ZDC test beam program and the corresponding simulation, which will be additional contribution into the efforts of validating physics model within Geant4.

B. Quality Assurance on the Release-to-Release Basis

The standard application of the central systems in the 4 Tesla magnetic field is also a subject of regular quality assurance procedure that follows the evolution of the CMS software and of the Geant4 versions, when applicable.

This is done by a fully automated software validation suite. In fact, the suite includes not only simulation software, but also addresses higher level quantities, such as reconstructed objects, etc. The task of the suite is to ensure, at the time of issuing a new release, that the CMS software behaves as expected, be it a specific component anywhere in the event processing chain, or the whole, integrated offline machinery.

The suite is still under extensive development but it has already proven highly useful, in particular in the recent evaluation of the CMS simulation software based on the latest release of Geant4.

Among features to be added to the CMS software validation suite in the nearest future is a complete interface to the CMS Production machinery, in order to take full advantage of the grid resources in running the tests. The results of the validation tests will be automatically published on a dedicated web server, in the same manner as in the Data Quality Monitoring outcome.

VI. CMS SIMULATION PERFORMANCE AND PRODUCTION

After undergoing a major redesign in the record-short time of less than 1 year, the upgraded CMS simulation software went into production in July 2006, as a part of Computing, Software and Analysis challenge (CSA2006).

Since then it has delivered over 60 millions of events for a variety of physics channels.

The software has proven to be stable and reliable, with the failure rate no more than 1 in 10^4 events.

The amount of CPU required to process an event depends on the event type. For example, on the 3.6GHz processor it has estimated at $\bar{37}$ seconds for a minimum bias or $\bar{200}$ seconds for a Higgs event with 4 leptons in the final state.

Overall, it provides basis to believe that, with a combination of full and fast simulations, CMS can reach the goal of having equal number of real and simulated events ($\bar{1.5} \times 10^9$ /year).

VII. CONCLUSION

In the CMS, the Object-Oriented Geant4-based simulation software has been successfully implemented, extensively validated for all CMS subdetectors and physics groups, tested for robustness and used for physics and detector studies.

It has been ported to the new Framework, on a very tight timescale and without loss of functionality.

It is efficient, modular, configurable and extendable.

One of its most important features is being able to adapt to ever-evolving requirements of such a live system as a physics experiment can be.

REFERENCES

- [1] S. Agostinelli et al., *Geant4: a simulation toolkit*, NIM A 506 (2003), 250-303
- [2] <http://root.cern.ch>
- [3] <http://www.thep.lu.se/~torbjorn/Pythia.html>
- [4] <http://hepwww.rl.ac.uk/theory/seymour/herwig>
<http://hepforge.cedar.ac.uk/herwig>
- [5] <http://mlm.home.cern.ch/mlm/alpgen>
- [6] <https://savannah.cern.ch/projects/hepmc>
- [7] R. Brun et al., *GEANT3 User Guide*, CERN Program Library W5013