Applications Area Report
October 2013 – March 2014

Apr 2014

Pere Mato

ROOT

The ROOT team is working hard towards delivering the ROOT 6 for end of May 2014. There has been several beta releases (beta2 on February 7th, and beta3 in April 3rd) to facilitate the integration of this major release in the software stack of the LHC experiments. The progress in the development of this version and issues in the integration are being reported in weekly planning meetings attended by the stakeholders. Minutes of these meetings are available attached to the agendas.

A total of 8 new patch releases 5.34/11-18 have been made available to the users. These releases fix many problems and issues, as well as they add new functionality that has been developed for the master (version 6) and back-ported to the current production release.

Over the period of reference a total of 588 new issues (bugs, feature requests) has been created and resolved a total of 568 issues.

Persistency Framework

New major releases of CORAL and COOL have been prepared for ATLAS in the LCGCMT_67 configuration to provide new features in both packages, such as support for a new relational schema with "vector payload" in COOL and protocol changes and threading improvements in the CORAL server. These and other enhancements required backward incompatible changes in both APIs, which were not possible during the LHC data taking and have been postponed to the current LS1 phase. This is the first release that is only supported with C++11 build options (using gcc47 and gcc48 on SLC6), although C++11 extensions are still disabled in CORAL and COOL as they would not be compatible with ROOT5. Two patch versions with urgent bug fixes in COOL query performance (LCGCMT_67a) and in the frontier client (LCGCMT_67b) have also been prepared.

The port of PyCool to the latest ROOT6 beta3 version has been completed. Several issues in successive ROOT6 candidates have been identified and solved during this process, thanks to the good coverage of the COOL test suite. A release candidate LCG_68_root6 based on ROOT6 beta3 and a reference release LCG_68 based on ROOT5 have been produced for LHCb, using almost identical code bases for CORAL and COOL, except for the replacement of Boost by C++11 classes in the COOL API used with ROOT6. One major issue with C++ exception handling in PyCool is still pending, but this will only be fixed when ROOT6 moves to a more recent JIT version.

Progress is being made in the port of the CORAL and COOL build system from CMT to cmake, in the port of the ticketing system from savannah to JIRA, and in the port of the CORAL server test infrastructure from quattor to Puppet.

Simulation

The new major release Geant4-10.0 was announced on December 4th as planned. It introduces the ability to efficiently perform event-level parallelism on multi-threaded (MT) capable hardware. Builds can either be done in sequential or MT mode (choice to be made at installation time), compatible with the API defined in previous releases for the sequential mode. This allows for easier transition/porting of the user's code to the new release. A sequential application can still build and run without changes in API with an MT-enabled Geant4 build. Full event reproducibility is guaranteed in release 10.0, with exact numerical reproduction of events in MT-mode vs. sequential mode. Excellent linear scalability of performance up to a very large number of threads O(100) has been measured on different hardware systems, with memory footprint reduction (footprint varying from 30 to 50 MB/thread depending on the application), and additional gain in event throughput when running in hyper-threading mode on capable hardware.

On the physics side, release 10.0 includes several improvements and extensions in Bertini Cascade (improved two-body final-state angular distributions for nucleon-nucleon and gamma-nucleon collisions; improved the phase-space generation for multi-body final states; added the capability to handle muon capture) and Fritiof/FTF (improvements to hadron-nucleon diffraction dissociation; extended to handle nucleus-nucleus collisions and model re-tuning with enlarged set of thin-target data) models.

Production of isomers (long-lived meta-stable nuclides) is enabled by default. Obsolete CHIPS and LEP/HEP models and codes have been completely removed. Improvements in neutron capture, low-energy neutron and inelastic cross-sections. Improvements are in particular expected in the reproduction of hadronic showers when using the FTFP_BERT physics lists. High-precision neutron data are now provided in compressed form, reducing considerably the size of the G4NDL library.

For EM physics, PAI ionisation models have been reviewed and improved, in particular for the tail delta-electron production spectra for incident electrons. An interface to angular generator is added for the majority of ionisation models and the magnetic monopole ionisation model has been upgraded for very non-relativistic cases. New micro-dosimetry models for ionisation in silicon based on dielectric functions computation are added. The Urban multiple-scattering model has unified and upgraded to the latest version.

Use of fast math functions for log/exp/pow in critical areas of EM processes brings overall CPU speedup.

The AIDA Unified Solids classes are made available in release 10.0 as an optional component, allowing the possibility to replace the standard geometrical primitives as choice at installation time; a limited number of shapes are available, in particular an optimised version of the polycone shape with improved CPU performance. The obsolete BREPs module has been removed. Also, new UI commands for overlaps checks are provided, now based on the built-in generation of random points on solids' surfaces.

Release 10.0 also comes with a new framework for physics-based biasing, based on wrapper and helper classes, and a new module for low-temperature "phonon" physics.

A new patch release 10.0.p01 was also made available last February. A set of fixes back-ported from release 10.0 have been collected in a new patch to the release 9.6 series, release 9.6.p03, announced on March 20th.

On March 18-19th 2014, the second LPCC Detector Simulation Workshop was held at CERN. The agenda is available here:

https://indico.cern.ch/conferenceDisplay.py?ovw=True&confId=279530

The goal of the Workshop was to update the assessment of the status of LHC detector simulations, including the validation against data, the definition of the future needs for analysis work and for upgrade R&D studies, and the progress in the development of the tools. Two important problems, related to electromagnetic physics, have been solved directly by the experiments, without need of changes in Geant4: the Geant4 simulation of the ALICE TPC; the simulation of the impact parameter distribution in LHCb.

Two other important problems, related to Geant4 electromagnetic physics, have seen some progress, although not yet a complete solution, which means that further work is needed by both Geant4 and by the experiments: the tails of the Z->ee lineshape; and the lateral electron shower shape.

Related to hadronic physics, there is a general satisfaction, although there is a strong need of more validation studies, in particular with test-beam data, and with recent versions of Geant4 (>= 9.6). The changes in lateral hadronic showers that have been reported by the developers, based on simplified calorimeters, still need to be confirmed by data. For the visible energy, an important issue to be investigated by ATLAS is the determination of the coefficients of the Birks quenching in scintillators.

For the Geant4 developers, the main focus will be on energy resolution, which seems to be underestimated at high energies.

Good agreement between Fluka and Geant4 (with high-precision neutron physics) has been reported by ATLAS on cavern radiation background; moreover, CMS has reported preliminary good agreement between Geant4 HP and the much faster XS approach.

Regarding the Geant4 versions of interest for Run 2, it seems that G4 9.6 will be used in at least two cases (ATLAS and LHCb), whereas for G4 10, there is general interest to study further its multi-threaded capabilities, and could be also used for production in Run 2 (by CMS and maybe ALICE).

The date for the next, third LPCC Detector Workshop has not yet been decided; however, it has been agreed to have half-day topical meetings, roughly once per year, on dedicated topics of interest related to Geant4 simulations.

CernVM file system

For the file system the main activity of the past months was still the consolidation of the 2.1 version. Version 2.1.15 has been released on Oct 17, 2013 after full validation and recommendation by WLCG. A consolidation and bug fix release 2.1.17 has been published on March 10, 2014 and is currently under deployment. The most relevant improvement is a TBB-driven file processing engine, with drastic performance improvement on multi-core nodes.

In December a schedule for end-of-life of version 2.0 has been proposed and accepted. The schedule foresees three major phases: the migration of all clients by end March 2014 (achieved); the migration of all Stratum 1's by June 2014 (in progress); the migration of Stratum 0 during Q3/2014. Consequently, the end-of-life for all version 2.0 components has been fixed for end of Q3/2014.

Discussion are going on with the main Stratum 1 / Stratum 0 administrators to provide a bootstrap CernVM-FS repository containing public keys and configuration files,

disentangling the core CernVM-FS software from CERN-related configurations. This will allow to better coordinate

operations of the several Stratum 0 and Stratum 1 servers not hosted at CERN or Fermilab but still using grid resources. The target is to have the configuration repository sometime during Q2 or Q3/2014.

CernVM appliance

Version 2.7 of CernVM, the last one based on SL5 and with Conary as package manager, was released on Oct 13, 2013 with an end-of-life foreseen for Sept 30, 2014.

The first production release of CernVM 3, based on SL6 and using the new uCernVM bootloader technology, was also released end of January 2014. Work has continued to consolidate the cluster definition features of the online contextualization interface and to include at tool to dynamically resize clusters of virtual machines on standard clouds. Together with CernVM 3, this work is at the base of the Virtual Analysis Facility, a system integrating all the CernVM ecosystem to instantiate a service - e.g. PROOF - on a cluster on virtualized resources. The Virtual Analysis Facility has been tested successfully on the Openstack cloud systems at CERN, both for PROOF and Q&A services, and is at the base of an italian activity of cloud evaluation involving ALICE, ATLAS and CMS.

Data Preservation

Effort to promote the CernVM solution for software environment preservation has continued in the context of the DPHEP collaboration. A set of demonstrator exercises has been defined involving ALEPH and CMS. The first step consisted in using the uCernVM bootloader technology to recreate SLC4 and SLC5 virtual machines and it was successfully accomplished. The ongoing activities aim to have all the relevant components fully functional.

SPI

In the last period, three major (67, 68, 68root6) and one minor (67a) LCG releases were created. While LCG 68 serves as production release, LCG68_root6 is a preview-release for validation of the upcoming ROOT 6 release. It only contains changes required by the changes between ROOT 5 and ROOT 6. These two releases are as well the first releases with LCG and GENSER releases being delivered together, using the new CMake based infrastructure. On platforms for upcoming releases, the gcc47 builds have been stopped. As replacement builds with the clang compiler will be provided. There exist compatibility problems with the LCG software stack on the latest MacOS X version (Mavericks), which could not be resolved yet.

On Savannah-to-JIRA migration activities, the efforts have been re-strengthened and migrations are ongoing. As first experiment, ALICE finished their migration to JIRA. For CMS the migration is expected to finish well before the next report. Timescales for ATLAS and LHCb are still under discussion.