



1. Applications Area – July 1 – September 30 2005

Pere Mato – Applications Area Manager

The finalization of Planning the AA for the Phase 2 of LCG has been achieved during the quarter. Important discussions took place during summer in the context of the architect's forum (AF) between the projects and the experiment representatives to polish some of the concerns with regard to the merge of the SEAL and ROOT projects, which is the main change on AA with respect to LCG phase 1. In practice the merge consists of two parts: firstly the merge of the development teams into a single team, and secondly the evolution of the software product into a single set of core software libraries. The merger of the SEAL+ROOT development teams happened immediately with Rene Brun as the leader of the new merged team. The software merge into a single set of libraries was originally foreseen to be completed within a very aggressive timescale but this has now been adapted to allow more time for the experiments to migrate their software. During summer discussions it was agreed that the new ROOT/SEAL project would continue to support the list of SEAL packages currently used in the LHC experiments and a new ROOT work package was introduced in the plan to reflect this agreement and a convenor for the work package was identified. The net result is that the adiabatic migration to a single set of libraries will take longer than initially foreseen. The plan document was finally endorsed by the PEB committee on September 27th.

Alberto Aimar will be leaving the Applications Area at the end of the year and Andreas Pfeiffer has already started to replace him as SPI project leader. This change together with the foreseen reduction in manpower in SPI will imply some changes in the scope and organization of the SPI activities. This is going to be finalized during next quarter.

The ROOT workshop has been an important milestone. The workshop represents an opportunity to present to the world the latest developments in ROOT and get the feedback from the experiments other than LHC. The version of ROOT released before the workshop included a number of packages originally from the SEAL project among many other new and improved functionalities. The feedback from the workshop participants was positive concerning the merge of the SEAL functionality (mathlibs, reflex, etc.)

The adaptation of the POOL frameworks to the new C++ reflection library (Reflex) has been a major activity during last quarter and has basically finished. The experiments are currently validating their software with the new versions. This represents an important milestone towards the single dictionary system for all AA software that we are pursuing.

WBS 1.1.1 - Software process and infrastructure (SPI)

Project Leader: Alberto Aimar, CERN IT/LCG

In the 3rd quarter 2005 SPI continued to support and improve the existing services; such as, among others, the savannah portal, the external software service, and software testing frameworks. All activities of SPI were steered by decisions of the Architect Forum and by the recommendation of the LCG AA Internal Review of March 2005.

Starting on October 1st 2005 Andreas Pfeiffer becomes the leader of the SPI project.

Savannah service

The service is stable and widely established among the users' community. The web-based portal is now in use by more than 1400 registered users in the LHC experiments and in several CERN departments. User support and bug fixing were the standard activities of this quarter.

The major achievement of this quarter for Savannah was the move of the service from RH73 to a new server running SLC3. SLC3 will be the only Linux platform supported at CERN by the end of 2005. The move was very successful and it only stopped the service for one hour. But the preparation had required very thorough testing and implied the port of Savannah to more recent versions of the tools on which the system is based (apache, mysql, php, etc.).

The project hosted the main developer of the open source Savannah for one month. This



allowed the alignment with the latest open source development sources and the merge of new features, improvements and bug fixes.

External software service

SPI continued to maintain and improve the external software service. In this quarter several new versions of external packages were installed. And also several grid related packages needed by the LCG projects (POOL in particular) were upgraded in this quarter, in the versions available in the quarter (glite, VDT, Globus, RLC clients, LFC, etc).

The main improvement in this quarter was the automation of the installation and deployment of external software packages and the completion of the generation of the web of the External Software web site. Another important achievement is the maintenance of the build information in a single XML file to maintain: SPI now automatically generates the SCRAM files from a single and clear XML format. Using the same system to also create the files for CMT is already implemented; this functionality will be tested and released in the next quarter.

Software download service

SPI has completed and documented in the current quarter a system based on the pacman distribution system. We are currently waiting for the experiments to verify the system in order to improve it. The system allows the download of both binaries and source of all LCG software, including all external packages needed. The binary cache contains the recent versions of the LCG software (SEAL, POOL, etc.) and all needed external packages (boost, python, root, etc.) are available for download from SPI via pacman.

An evaluation of the porting of pacman to Win32 produced a report highlighting the difficulties of such a task (even if pacman is written in python) and it was approved by the Architects Forum that this porting is postponed.

Testing frameworks

In this quarter SPI applied the automatic generation of test coverage reports, based on GCOV and LCOV packages. SPI generates the reports for the LCG projects, and worked to generate those of some LHC frameworks and HEP applications (ROOT, CLHEP, etc.) widely, used in the community. The service is focusing on user support and porting to new platform. In this quarter as part of the common activities with EGEE the frameworks provided by SPI (CppUnit, PyUnit, QMTest) are used by all LCG projects and by the EGEE JRA1 middleware. SPI provided support and customisation of the tools as needed. All the testing frameworks and tools (CppUnit, Oval, QMTest, etc.) were installed for the SLC3 platform.

Build and librarian service

SPI continues to maintain solutions that can both integrate with specific build systems (SCRAM and CMT). SPI is now using a single generic description for all dependencies and configurations (in XML) and developed a small tool to automate the generation of the configuration files for several specific tools (CMT, SCRAM, pacman, etc). As planned this tool (MoCoM) is now in production. The several “postbuild” functions (doxygen, lxr, generation of tar files, QA reports, etc) have been automated and are now executed at every release using a new small tool developed in the SPI project.

Code documentation service

In this quarter tasks of this service have been automated and are part of a “postbuild” tool that generates doxygen and lxr documentation at every release and a snapshot is produced every night for each project (SEAL, POOL, COOL, GDML, etc).

Policies and QA

In this quarter there was further improvement of the automated QA report generation and the development of reporting tools to automatically extract metrics for the projects hosted on the Savannah service, not only of the LCG projects. The report generation is a tool configurable for any project hosted on the Savannah service. Other kind of automatic reports are those



related to “test coverage”. This is done to verify how much of the tests of each project really test the code before it is delivered. These reports are now produced for several LCG and HEP packages (POOL, COOL, GDML, ROOT, etc.).

Coming activities

No major changes are foreseen in most of the services: all resources are taken by user support and development of the current automated solutions. During this quarter a major reduction of resources was announced for the next year; therefore some further planning of the future milestones is being discussed by the new project leader with the other projects and users.

WBS 1.1.2 - Persistency framework (POOL and conditions DB)

Project Leader: Dirk Duellmann, CERN IT/ADC

Object Storage and References

During the last quarter there have been six releases leading to POOL 2.1.3 (Reflection based) and 2.2.2 (Reflex based). The main work item has been the move to the new LCG dictionary interface as part of the LCG Application Area strategy. This required maintaining two development branches in order to support the older Reflection based production (2.1.x series) and in parallel prepare and evolve the Reflex based 2.2.x series. The move to Reflex has required adaptations of all POOL storage components but is well advanced now and awaiting final validation from the experiments. Once this is obtained the Reflection based development branch can be frozen and the CORAL and ROOT5 integration and validation can start. The POOL code that used to cross populate the ROOT/CINT dictionary was moved into the *cintex* package and is now maintained by ROOT/SEAL.

The POOL storage components have integrated several bug fixes from porting activities in the context of openLab to complete a native 64bit version of POOL.

The LFC file catalog adaptor (introduced in POOL V2.1) has been picked up by several experiments for validation and the LFC developers have provided many enhancements and bug fixes in response to experiment feedback. The planning of a consistent configuration and version evolution between grid middleware components and application area software is still an issue and needs to be addressed across the areas boundaries.

Database Access and Distribution

The POOL relational abstraction layer (RAL) continued to evolve and several bug fixes and performance improvements have been made available. A significant development in the RAL area is the proof-of-concept FroNtier plugin, which allows reading relational data via http from FroNtier servers taking advantage of distributed proxy caches. The FroNtier plugin can be used transparently by existing RAL applications and will not require other database software on the client side. A production version is planned after first validation tests by the CMS experiment.

The outcome of the review of the RAL component interfaces and functionality (http://pool.cern.ch/ral_review.html) has been included in the upcoming major release of RAL, which is named CORAL (COmmon Relational Access Layer).

CORAL has no direct dependencies on other POOL components and will be released as separate package, which can be used as relation database foundation library with or without POOL. To simplify the migration from RAL to CORAL, both implementations can coexist in one application.

First CORAL releases have been made available to current RAL users for validation. CORAL include significant functional improvements (new attribute list, bulk operations also for update and delete, BLOB support, etc) but will also improve the deployment in a grid environment by providing connection retry and database service lookup facilities in a consistent way. These developments build on experience with database connection handling obtained with a similar package in ATLAS. To support upcoming validation and production activities an extensible monitoring framework has been integrated in CORAL. This framework allows extracting



information about frequency and latencies of database operations and provides this either as diagnostic printout locally or to accumulate database performance information in network based monitoring back-ends. The CMS experiment is currently implementing a plugin to integrate CORAL monitoring with Mona Lisa.

Collections and Metadata

The POOL collection components have received several improvements and optimisation driven by comprehensive ATLAS performance tests of the RAL and ROOT based collection implementations. The set of POOL collection utilities to administer (merge, copy, etc) larger collections has been completed.

Conditions Database

Five production releases have been produced leading up to COOL 1.2.5. COOL has added important functionality like multi-channel bulk retrieval, support for very large character data (CLOB) and full support for SQLight as backend. The project has spent significant effort on documentation and a doxygen based user guide and end user examples are now available. COOL is preparing the move to CORAL, which will make service lookup, connection retry, and monitoring available for COOL users.

On experiment request COOL added an API to retrieve the table names for integration of the COOL schema into larger relational data models on the experiment side. This avoids duplication of data that would arise in case of two independent relational models. The deployment issues arising from a database coupling between COOL maintained data and experiment data models need further study and a prototype application is being setup in collaboration with ATLAS.

Several major performance improvements have been achieved as result of the ongoing COOL validation test at CERN Tier 0. These tests are performed on the basis of ATLAS and LHCb requirement estimations and are extended also to LCG Tier 1 sites in the context of the LCG 3D project.

An important new development, which has been released, is the PyCOOL python binding which form the basis for user level GUI tools. PyCOOL has also used to provide COOL data distribution tools, which allow to extract subsets of condition data (either tag or time based) e.g. to create read-only subsets of conditions data for LCG Tier 2 sites running COOL services against MySQL database back-ends.

The manpower situation in the project has gradually improved and both client experiments ATLAS and LHCb are now contributing significantly to the software development and integration tests.

WBS 1.1.5 – Simulation

Project Leader: G. Cosmo, CERN PH/SFT

The Simulation Framework subproject (W. Pokorski)

The simulation framework project activity over the last quarter has been devoted mainly to areas of GDML, geometry objects persistency for Geant4 and Monte Carlo truth handling.

- In the area of GDML, the GDML writer for ROOT has been implemented in Python. It uses the PyRoot Python binding. Starting from the in-memory ROOT geometry, it allows exporting the geometry in GDML. A new release 2.3.0 of GDML is out.
A request for a technical student will be reiterated in order to progress with some of the activities planned for the next months on GDML.
- The feasibility study of the Geant4 geometry objects persistency using ROOT has been completed. Using the practical example of the full LHCb geometry setup, it has been demonstrated that object persistency for the detector description can be achieved using ROOT. LCGDict and Reflex packages have been used for the generation of the dictionary of the Geant4 geometry classes needed by ROOT I/O. A list of necessary changes to the Geant4 code has been prepared.



- The Monte Carlo truth handling subproject has been started. The first meeting took place where experiments presented their way of dealing with the Monte Carlo truth. Further meetings are scheduled where it will be evaluated how much of the Monte Carlo truth handling mechanism can be implemented in common. Potential extensions to HepMC as the Object Oriented event record will be also discussed.
- The Flugg-based application which is under development for the ATLAS calorimeter test-beam has been ported to the latest (2005.6) release of Fluka. Further work in this area is planned for the last quarter of 2005 with results available by the end of the year.

The Geant4 subproject (J. Apostolakis)

One of the milestones in the 3rd quarter included the set of developments contributed to the Geant4 development release made in September 2005. Effort concentrated on identifying the cause for the significant variation in visible energy for some sampling calorimeters and on creating prototype revisions to improve the stability of physical observables when changing production thresholds in these and similar setups.

- A key effort undertaken was the study on the stability of EM quantities from sampling calorimeters against changes in cut. The issue affected the visible energy in sampling calorimeters in Atlas, LHCb and the ILC, but was first observed by users in sensitive medical and space applications, before it was seen in HEP studies. After extensive investigation the underlying cause was identified in the treatment of angular distributions and the stepping mechanics associated with multiple scattering. Refinements and improvements in the multiple scattering physics process, created by Laszlo Urban (MFPI Budapest) in close collaboration with Michel Maire (LAPP), were tested on simplified LHC sampling calorimeter setups and other setups in the EM verification/validation test suite. A first version of the relevant revisions was provided in the September development release as an available option.
- It was also implemented a revised version of the transition radiation (XTR) process, enabling gamma emission within the radiator (and changing it to a discrete process), and introducing classical XTR parameterization models as an alternative. A presentation was made at the "TRDs for the 3rd Millennium", the 3rd Advanced TRD workshop for accelerator and space applications at Ostuni (Italy) in September 2005.
- Effort in hadronic physics concentrated on verification and validation testing. It was created a prototype of revised physics lists, in order to incorporate additional or improved physics options. These will include refined processes for stopping particles and new options for hadronic elastic interactions. Additional verification tests for hadronic processes have been implemented: a first version for "in flight" hadronic processes, and a refined beta version for "stopping" processes. A validation test at the detector level in the form of an advanced example rare-decay calorimetry was created. A talk on deep inelastic scattering (DIS) for the high energy simulation of neutrino-nuclear interaction was presented at the Conference on "Frontiers of Science" at Milan (Italy) in September 2005. A prototype implementation of the propagator interface in Binary Cascade to strings models was also implemented.
- New precision tests for CSG solids were created and utilized to verify the surface intersection performance. A first result of these tests was to identify issues in the sphere and torus shapes. Fixes were created that resolved long-standing problems due to additional solutions and to precision respectively, affecting particular configurations of these shapes.

Additional shapes (specific solids) were introduced including an ellipsoid with cuts in Z (starting from a user contribution, and revised together with an INFN G4 collaborator), and a full elliptical cone with cut in Z. In addition a tetrahedron, contributed by users at Vanderbilt University, was added.

A new ability was added to solids, to generate random points on their surface. This will enable the development of better volume overlap/coincidence detection in the future.

- A prototype port to CLHEP 2.0.2.1 was created, to enable the checking of portability and study the potential migration to this version of CLHEP for release 8.0 of



December 2005. This work confirmed the ability to maintain compatibility of the Geant4 source code simultaneously with CLHEP versions 1.9.2.1 and 2.0.2.1 and to utilize the definitions for units and physical constants in the global namespace; adoption of units and physical constants in the 'CLHEP' namespace will not be enforced as a first step of the migration.

Fixes and other developments:

- A problem was reported recently by experiments using the Geant4 fast-simulation framework, where electrons were deviated before it could apply itself. After consultation with the local Geant4 team, a user in ATLAS undertook an extensive investigation and found the reason of the problem and proposed a valid fix for it. The fix has been tested and is now included in the September development release.
- The September development release also includes a fix in the geometry modeler for the handling of theta-angle limited spherical shells.
- Fixes were also introduced for angular distribution of muon-Bremsstrahlung and pair-production from muons at high energies.
- There is ongoing support to enable experiments to simulate exotic particles simulation (I-hadrons, monopoles ...) covering several aspects including creating processes, interfacing with tracking, and enabling the simulation of magnetic monopoles.
- Design iterations are proceeding to allow coupled navigation in two or more geometries in parallel. This will support the use cases of LHCb and others, in which tracking in field, biasing and tallying (and/or fast-simulation ...) each have their own geometry, and charged tracks proceed in lock-step through these.
- Physics lists: technical issues and alternative physics options are under investigation.
- Data: a cataloguing and first assessment of the external data libraries used is ongoing, included a research for the conditions of use of the latest and historical versions of some data libraries.

The deliverable on improvement to the regression suite for release validation and testing infrastructure is behind schedule. Effort has been concentrated in implementing the recent change in coordination of the primary testing activity and in improving the day to day testing scheme which this revealed. A large part of the integration testing is now undertaken by Ian McLaren.

The new Collaboration Agreement has been sent to the Collaboration Board (CB) members for the final endorsement.

The annual Geant4 Collaboration Workshop will take place in Bordeaux, November 3rd-10th; it will also include the public Users' Workshop. Preparation of the program and topics of particular relevance are underway. Progress on the issue of the license for the Geant4 code is anticipated.

Personnel:

- Support for Sergey Sadilov's ended end of July 2005.
- Oliver Link fellowship has been extended until end of 2005.

The Fluka subproject (A. Ferrari)

The activity on Fluka during this quarter has been particularly important and allowed to achieve a few critical milestones, as well as a major goal of the joint INFN-CERN "Collaboration Agreement on the Development and Maintenance of the Fluka code" of December 2003.

The planned full release of Fluka, the first one under the INFN-CERN Collaboration Agreement, occurred on July 26th. Originally supposed to be a beta-release, it was actually released as an official version and the experience since then showed it has full production quality, with minimal bugs found up to now. This release fulfills the 7th milestone (beta-release, source included limited to CERN and INFN) in the Fluka LCG planning. The binary release is open to everybody, while the source release is at present limited to CERN and INFN with plans to open it to everybody in the next future. Both the binary and source releases are covered by the Fluka User license (<http://www.fluka.org/License.shtml>). Written acknowledgment of the License is required for the source release.



The Fluka Coordination Committee¹ issued the following announcement for the July release:

"The Fluka Coordinating Committee (FCC) is very pleased to announce that as of July 26th 2005 the main tasks falling under its mandate have now been completed. The current status may be summarised as follows:

- *a new version of Fluka (Fluka 2005.6) containing significant scientific and technical enhancements is now ready for public release*
- *the release includes source code and binaries together with release notes*
- *the technical Users' Manual has been finished and will be published immediately as a CERN Yellow Report*
- *the Fluka Software License describing the conditions of use of the Fluka code has been established to coincide with the release*
- *according to the Fluka Collaboration Agreement, CERN and INFN engage themselves and their personnel to recognize as Fluka only what is obtained from the official web site (www.fluka.org). In particular, the use of non-authorized versions or parts of them and improper referencing in official publications will be avoided*
- *users of Fluka needing assistance to migrate to this new version are encouraged to contact a member of the recently established Fluka Scientific Committee(*) for assistance."*

The announcement was endorsed by the INFN and CERN top management, in particular for CERN by Prof. Engelen. In the same meeting of July, the FCC approved special derogations to the License for the ALICE experiment.

From the technical and scientific point of view, the following activities and/or developments were successfully completed (most of them were prerequisites for the release):

- Completion of PEMF elimination and thorough testing of the new solution (1st milestone of the Fluka LCG planning)
- Completion of the (online) radioactive decay module (2nd milestone of the Fluka LCG planning). This module is already in heavy production use in problems associated with CNGS and LHC. It has been validated against the "old" offline procedure thanks also to an effort by the RP Group
- Cleanup and proper documentation/acknowledgments of the source code (6th milestone of the Fluka LCG planning)
- Release of the first report on Fluka. The joint CERN/INFN/SLAC report has been published as a CERN Yellow Report CERN-2005-10 (8th milestone of the Fluka LCG planning)
- Thorough testing and cross-checks of the new version (9th milestone of the Fluka LCG planning)
- Extension of the PEANUT model to antinucleons and increase of the PEANUT model maximum energy to 5 GeV for most particles (partial fulfillment of the 4th milestone of the Fluka LCG planning)
- Photon cross sections now taken from EPDL97
- Photon coherent scattering cross sections now include the effect of the real and imaginary part of the anomalous scattering form factor
- Atomic elastic form factors are now from EPDL97 as well
- Photon coherent scattering: completely new sampling
- Photoelectric: edge-by-edge fit/tabulations are now detailing every edge (down to few eV's)
- Photon pair production: new sampling scheme accounting for e+/e- differences at low energies
- Implementation of preprocessor like directives in the input file
- Substitution of all "external" mathematical routines with algorithms adapted from the public domain SLATEC library

¹ The current membership of the FSC is as follows: Giuseppe Battistoni (INFN/Milano), Federico Carminati (CERN/PH), Alberto Fasso (SLAC), Alfredo Ferrari (CERN/AB), Andrea Ottolenghi (INFN and University of Pavia), Johannes Ranft (Siegen University), Stefan Roesler (CERN/SC), Paola Sala (INFN/Milano), Vasilis Vlachoudis (CERN/AB)



- New random number generator, with 53 random bits, based on the latest Marsaglia et al. paper. It is claimed to fulfill all known randomness tests

The pace of development will surely slow down a bit in the last quarter, with the priority of the main developers partially switched back to applications following the momentous release period.

The Physics Validation subproject (A. Ribon)

The following achievements were obtained during the 3rd quarter:

- Verification of longitudinal hadronic shower shape: the study has not yet been completed, but progress has been achieved in the ATLAS TileCal 2002 test-beam analysis (see also point below), and results are expected by the end of October. For CMS, results are not yet ready.
- 3rd simple benchmark: A.Ribon has started the third simple benchmark, using the latest Fluka and Geant4 releases. Very preliminary comparisons between simulation and data, for laboratory rapidity and transverse momentum square distributions, are already available. Now a more in-depth study of the experimental analysis, and a better understanding of possible contaminations in the real data not yet included in the simulation, must follow. The study will have to proceed in collaboration with Fluka and Geant4 experts. The results are expected by the end of the year.
- Calorimeter test-beam validation: M.Gallas has completed the migration from the original ATLAS TileCal 2002 test-beam setup to the new ATLAS framework, with the latest Geant4 version, and is now testing it with the experts involved in the original test-beam analysis. Once this phase is completed, the new production will start and by the end of October it is expected to reproduce the old comparison between Geant4 and test-beam data with the latest Geant4 release (including the verification of the hadronic shower shapes). W.Pokorski updated the setup to the recent Fluka release (Fluka2005.6), and everything is now ready to tackle the final step of creating Geant4-hit objects from the Fluka physics output. Regular meetings between M.Gallas, W.Pokorski, and A.Ribon will start in October; results are expected by the end of the year.
- Lot of activities are undergoing in ATLAS to analyze the 2004 combined test-beam data, and a status report is expected in the Physics Validation meeting of October 26th. First results are expected for the end of year.

The Garfield subproject (R. Veenhof)

The main activities over the past quarter have been:

- Participation to the TRD-2005 conference in Ostuni with the opening talk.
- Completion of the interface with the new version of Magboltz. The new release includes anisotropic cross-sections for all noble gases as well as CS2. The documentation of the cross-sections has been brought up to date.

Note: S.Biagi, the author of Magboltz has retired. The problem for future support of Magboltz needs to be sorted out.

The Generator Services subproject (P. Bartalini)

The most relevant events which took place during the 3rd quarter are the release of GENSER 1.1.0 on July 25th and the open meeting on Herwig++ on August 29th.

Current manpower situation for the project:

- Coordinator: Paolo Bartalini (University of Florida)
- Librarian : Andreas Pfeiffer (CERN)
- Liaisons : Alberto Ribon (CERN)
- WP1 : Mikhail Kirsanov, Alexander Toropin (LCG Russia, GENSER integration)
- WP3 : Sergey Belov and Anton Gusev (LCG Russia, MCDB development)
- WP4: Marianne Bargiotti (INFN, Monte Carlo validation)



The librarian is currently taking other responsibilities in LCG and we will need to find a replacement. LCG Russia confirms 2005 resources for the year 2006. New INFN participation (at 0.3 FTE) on Monte Carlo validation (WP4) will be available for the next three years.

WP1 (GENSER): Release 1.1.0 has been made available end of July. The documentation has been improved. New tools (python scripts) are being put in place for package installation/compilation and post-installation checks. There is an ongoing activity on testing the Fortran packages with the new gcc-4 compiler.

WP3 (MCDB, HEPML): MC database is now usable, but until now few users have tested it. In order to move forward and collect the necessary feedback it becomes important to advertise it more widely. This has started and will be continued in Q4 2005. During the quarter the activity concentrated on the improvement of the user interface, on the tests with large files and on the authentication systems (in particular the one supporting GRID certificates).

Simulation Framework: the prototype was reviewed and works but it had not really followed the design. In parallel with the reorganisation of the software in CMS, the project was restarted to follow a more coherent vision; there will be a task force mid October working on it. It is however not clear at what level the project can be driven in LCG, given the limited contact with the people involved and the different objectives more CMS-oriented.

WP4: Basic sanity checks: GENSER tests have been moved in separate sub-package.

Monte Carlo Validation: a new activity started on the validation of NRQCD models recently inserted in Pythia 6.3.

WBS 1.1.6 – Core libraries and services (ROOT)

Project Leader: Rene Brun, CERN PH/SFT

During this quarter, the organization in work-packages created at the time of the SEAL+ROOT merge has been consolidated. A new work-package SEAL has been added to continue support for the existing SEAL libraries until the migration from the experiments is complete. The ROOT Users Workshop took place at CERN (28-30 September). This has been a major milestone for the project. The workshop had 115 registered users, 41 talks and 9 posters. All work-packages had at least one report at the workshop. For more information see <http://root.cern.ch/root/R2005/Welcome.html>.

The next major milestone is the release of version 5.oX before Christmas.

BASE work package (F. Rademakers)

A number of new classes and new features were added to the ROOT Core libraries in the past quarter. Here is the list of them:

- New TMacro class to allow for storing a C++ macro in a ROOT file.
- Added support for multi-level TAB completion in the ROOT shell.
- Added TFileMerger class to allow for easy copying of two or more files using the many TFile plugins
- Introduced a XML DOM (Document Object Model) parser internally using libxml2
- Added CASTOR 2 support in TCastorFile (both CASTOR v1 and v2 protocols are supported)
- Finalized the port to MacOS X Tiger (10.4). On Tiger the default Fortran compiler is gfortran (optionally g95 is possible)
- New port to AIX5 with gcc and at the same time improved the port to xIC
- Work has been done to make more classes thread safe

Dictionary work package (Ph. Canal)

This work package is responsible for the development of the reflexion system and its integration with the CINT and Python interpreters. In early May an important workshop took place at CERN to discuss the use of the Reflex dictionary package within CINT. This will be a



very important change for CINT. The Reflex package from SEAL has been incorporated into the ROOT CVS structure. Reflex has been ported to more platforms and released end of June with ROOT 5.02. Like Reflex, the Cintex package from SEAL has also been incorporated into ROOT CVS and ported to more platforms.

The CINT package was originally developed and maintained by Masa Goto with his own code management system. CINT has now been moved to the ROOT CVS structure. The old components of CINT written in K&R C have been changed to compile with C++. An important restructure of the code has been made to take advantage of CVS. This simplifies substantially the CINT code and improves its maintenance.

A new byte-code generator is currently under development since one year. The new generator is aimed at more C++ compliance (in particular regarding the scoping rules and support for namespaces).

Investigations of the changes necessary to interface CINT with Reflex are going on. Development of a Reflex dictionary using the current CINT Api is nearing completion.

Update to the existing code is still going on including improvement of the handling of function pointer, port to gcc v4.00 and improvement of the support for 64 bits quantities.

I/O and Trees work package (Ph. Canal)

This work package is responsible for the development of the ROOT object persistency system (basic I/O, Trees and Trees queries). Many performance improvements have been applied to the TTree processing. Support for `std::pair` was improved in particular in the case with the original library. We introduced the first version of the TTreeSQL class which allows the processing and creation of SQL tables.

Many improvements (in particular to support STL collections) have been made in versions 4.04 and the development version 5. Several new cases of schema evolution have been implemented. The Tree query mechanism (TTreeFormula) has been extended to support Tree friends, event lists and indices in case of chains of Trees, in particular when used in the PROOF system. A new test suite "roottest" has been made available and ported to all ROOT supported systems. It consists of several hundred tests (mainly I/O and Trees) that supersedes the previous test suite (still used to run the most typical tests and for benchmarking).

PROOF work package (F. Rademakers)

The PROOF work package passed successfully an important milestone with a demo of a large set of new features during the recent ROOT workshop. These new features are:

- Support for multi-sessions. The limit to only one PROOF session per ROOT client session has been removed. Each session has a unique identifier, a GUI progress dialog, which is started when the first processing action is started, and a list of results
- Asynchronous Running Mode. This feature makes possible to process selectors queries asynchronously, i.e. without blocking the client ROOT session. Queued requests are processed sequentially and the end of processing a copy of the results is sent back to the client who can proceed to finalize the query. By default processing is synchronous (blocking).
- Finalize, Retrieve, Archive and Remove Functionality. A set of new methods have been added to the TVirtualProof interface to administrate the available queries.

Other Improvements include:

- Optimization of slave server startup reducing considerably the PROOF startup time on large clusters.
- Support for Friend Trees.
- Support for Indexes. Indexes can be used with PROOF to connect events from different trees.

Future developments include a powerful GUI that allows easy access to the many new PROOF features. The capability to disconnect and reconnect the client ROOT session from



running PROOF sessions is being implemented. And finally, we are studying ways to fairly and optimally schedule queries of many simultaneous users on large PROOF clusters.

MATH work package (L. Moneta)

The main activities of this work package have been focused on the Mathcore and MathMore libraries, where new versions have been produced for the ROOT development release of mid-September (v5.04.02). These new developments have also been presented at the ROOT Workshop. MathCore contains improvements in the physics and geometry vector package. Classes to describe the various types of 3D rotation classes are now present, including the conversion between the various types. The MathMore library has been release for the first time. It contains mathematical functions and numerical algorithms like derivation, integration and root finders implemented using functions of the GNU Scientific Library (GSL). A C++ wrapper has been developed to hide the GSL implementation. In addition the GSL functions are compiled inside the MathMore library starting from a GSL tar file to avoid having an external dependency. For more detailed information see the ROOT release notes

(<http://root.cern.ch/root/Version50400.news.html>). Both Mathcore and MathMore can be downloaded and built as standalone packages, thanks to an autoconf based front-end.

This quarter saw some progress also in the minimization. Thanks to the work performed by a summer student (A. McLennan), the new C++ version of MINUIT has been extensively tested and few bugs have been discovered and fixed. (see the summer student report at

http://seal.web.cern.ch/seal/documents/minuit/FunctionMinimization_report.pdf). In addition, a new interface of the new C++ MINUIT to ROOT has been developed and it is being ready to be integrated in the next release of ROOT.

Other activities include improvements done in the TF1 class for the algorithms to find the minimum and maximum of a one dimensional function and the addition in ROOT of a new statistical tool, sPlot, to unfold data distributions.

In addition, the ROOT Statistics tools have been presented in September by A. Kreshuk at PHYSTAT05 conference (see <http://www.physics.ox.ac.uk/phystat05/Talks/KRESHUK.ppt>).

GUI work package (I. Antcheva)

There is a substantial progress in the user interface developments in ROOT, as well as continued support for users using the ROOT GUI classes for the application design and many tiny improvements.

The ROOT object browser was upgraded to support the control of tree nodes hierarchy. This development is important when browsing detector geometries.

The new Style Manager graphical user interface was recently introduced in v5.04. Its main purpose is to manage different styles in a ROOT session. It allows users to import a style from a canvas or a macro, to select a style for editing, to export it in a C++ macro, to apply a currently selected style on a selected object in a canvas or on all canvases, to set it as the global gStyle for the current session. This interface is composed of two parts:

- the top level interface manages a list of all available styles for the current ROOT session and shows the currently selected one;
- the style editor deals with the settings of the currently selected style.

A preview of the selected canvas helps for precision work. It can be updated dynamically at run-time or by request to show how the edited style looks. All changes made in the style editor can be cancelled and the edited style can be restored to the last saved state in a macro.

The Qt layer validation and tests have been continued and the current status can be seen at: <http://antcheva.home.cern.ch/antcheva/QtToDo1.html>

Graphics work package (O. Couet)

Many new major developments have been done in the various graphics domains.

OpenGL developments have greatly improved the rendering quality and the user interaction with pictures:



- Shapes outline style. This provides a 'line-drawing' style, with outlined filled polygons.
- Extended camera interactions (Rotate (Orbit), Truck (Pan), Dolly (interactive move in/out), Zoom (field of view)).
- Improve transparent objects picking and drawing.

In addition a major internal re-structuring of the GL viewer has been done to support repeated additions of duplicated placed shapes, dynamic scene builds, level of detail scheme, common display list cache and multi-pass rendering. See all details in the 5.04/00 ROOT release notes (<http://root.cern.ch/root/Version50400.news.html>).

The GL viewer can now be displayed directly embedded into a TPad. This allows a much better integration of 3D scenes with all other graphics produced in ROOT. It required some reorganization in the GL viewer internals. These developments are still in progress. In particular the polylines, polymarkers and composite shapes drawing are missing as well as the PostScript output. Also there are still some performances issues to improve

The new GL viewer was successfully demonstrated at the ROOT workshop. A new printing facility in batch is available thanks to TASImage. TASImage is an interface to image processing library using libAfterImage. A new class TImageDump using TASImage allows to save canvas in GIF, JPEG, etc., image formats in batch mode. Before this class was introduced it was not possible to generate bitmap files in batch.

GEOM work package (A. Geatha)

The most important changes in the geometry package since version 5 are related with consolidation of navigation algorithms and few new features:

Navigation

- Tolerance on boundaries introduced for distance computation algorithms of all basic shape classes. The behavior and stability of these algorithms when tracking very close to boundaries become particularly important when interfacing TGeo with FLUKA and GEANT4. Currently all these algorithms provide consistent answers if the rounding error for the position of a particle "on boundary" is less than 1E-10. This value was chosen by testing various FLUKA native geometries.
- New global navigation query available within the geometry manager class. The method provides linear propagation of a track and crossing of next boundary, computing the normal vector to the crossed surface. As option, the method can also compute the safety distance. The method is currently used by FLUKA-TGeo interface as well as by the ray-tracing algorithm of TGeo.
- Several minor fixes and optimizations for navigation algorithms

New features

- New type of geometry constructs: assemblies of volumes. These represent unions of different volumes without needing a container. An assembly derives from a volume object and uses the same user interface for defining the components. The main advantage of assemblies compared to normal volumes is that they cannot produce overlaps due to their virtual container. Navigation within these structures preserve all optimizations existing at the level of volumes (bounding box + voxelization), therefore their usage do not imply overheads in performance.
- Scaled shapes available. One can use now any basic shape supported by TGeo (or user defined, deriving from TGeoShape) and scale it arbitrary on the 3 axis. To create a scaled shape one has to define a normal shape and a scale transformation. Scaled shapes may find use cases when trying to represent deformations.
- Possibility to work with geometry segmentation. One may want to use TGeo to describe the detailed segmentation of their geometry (e.g. structure of pads/wires of the detection element) without being active in simulation. This can be currently easily done by setting a flag at the level of TGeoVolume class.
- A lot of improvements and new features 3D visualization of geometry based on GL.



SEAL work package (L. Moneta)

During this quarter 5 new releases of SEAL have been produced according to requests from experiments or POOL. These releases include new versions of ROOT and Boost packages and fixes and improvements mainly in the Dictionary subsystem (Reflection, Reflex and DictionaryGenerator).

Milestone performance during the quarter

Level 1 Milestone WBS 1.1

ID	Due	Done	Summary	Notes
AA0	31/03/2005	27/09/2005	Development, support and resource plan through 2008	Postponed until Q3. PEB endorsed plan document on 27th Sept.
AA1	30/09/2005		Phase 1 AA software complete and deployed	Not clear when this should be considered done.
AA2	30/09/2005	28/09/2005	ROOT workshop at CERN. Various prototypes addressing different topics of the SEAL+ROOT merge will be completed and detailed plans will be made available to the wider ROOT community.	Math and Reflex libraries released as part of ROOT 5.04.00 ready for the ROOT workshop.

WBS 1.1.1 – SPI

ID	Due	Done	Summary	Notes
SPI03	31/07/2005	30/09/2005	Pacman caches for binaries and sources	Done. The binaries caches are available and can be tested by the experiments.
SPI10	30/09/2005	30/09/2005	Testing tools updated to the latest versions on the AFS area and available to LCG projects and outside.	Done. The latest versions have been upgraded, and are available. Will go in release configuration once it is agreed with the LCG projects.
SPI11	01/08/2005	30/09/2005	Doxygen and LXR automated and documented for the LCG software and for the packages (ROOT, AIDA, CLHEP)	Done. There are automatic scripts that are run after every new build and provide the documentation with Doxygen and LXR. In addition there is a snapshot on the LCG code that nightly runs Doxygen and LXR.

WBS 1.1.2 - Persistency

ID	Due	Done	Summary	Notes
RAL03	31/08/2005		RAL validation release (implementing the RAL review outcome and separate packaging)	On going work. Not yet finished.



LHC Computing Grid Project – Progress Report June 2005

CAT02	31/08/2005	19/07/2005	Production release of LCG catalog adaptors (level 2)	Done. All catalog adaptors released in version 2.1.1 of POOL.
CDB01	30/04/2005	31/05/2005	First COOL production release	Done. COOL 1.1.0 released.

WBS 1.1.5 – Simulation

ID	Due	Done	Summary	Notes
SF515	30/06/2005		GDML-Schema extension to support divisions and reflections	Postponed due to lack of manpower. Request for a technical student will be reiterated.
SF516	30/06/2005	15/07/2005	Feasibility study of POOL/ROOT based Geant4 geometry persistency	Achieved. A prototype using ROOT was implemented; requirements have been analysed and necessary modifications will be implemented in Geant4.
SF517	30/06/2005	03/08/2005	Python implementation of GDML writers for Geant4 and ROOT	Achieved. Implemented in GDML 2.3.0 release early August. For Geant4 a "light" Python layer over the existing C++ implementation will be provided.
SF530	30/09/2005		Flugg application to one of the LHC calorimeter test-beam simulation	In progress, ATLAS TileCal 2002 setup updated and geometry exported to GDML; interface with Flugg in progress.
VD512	30/06/2005		New validation results on longitudinal shower shapes with Geant4 7.0 completed	In progress, still waiting for feedback from ATLAS and CMS for the verification with the most recent version of Geant4. Results from ATLAS are expected by end October 2005.
VD522	30/09/2005		3rd simple benchmark for physics validation completed	In progress, first preliminary results achieved. Final results to be expected by end of 2005.
VD524	30/09/2005		Validation of Fluka against one LHC calorimeter test-beam	In progress, ATLAS TileCal 2002 setup with Flugg being completed. See also SF530
G4518	30/06/2005	30/06/2005	Release 7.1: refinements to ionization processes, additional string fragmentation and verification of proton reactions at high energies.	Achieved. On time as scheduled.



G4523	30/09/2005	30/09/2005	Development release: cascade interface to strings, stability study for EM observables and review of LPM effect; prototype port to CLHEP 2.x.	Achieved. Planned developments included. First prototype study for migration to CLHEP 2.0.2.1 performed.
-------	------------	------------	--	--

WBS 1.1.6 Core libraries and Services (ROOT)

ID	Due	Done	Summary	Notes
BAS1	01/08/2005		New updated plugin manager	Postponed. The experiments continue to use the SEAL plugin manager. Scheduled for the december release.
BAS2	15/07/2005	12/07/2005	A new edition of the ROOT Users Guide	Achieved. Version 4.04 of Users Guide available.
BAS3	15/07/2005	15/07/2005	The ROOT bug tracking system moved to Savannah	Achieved.
DIC1	30/06/2005	30/06/2005	Reflex and Cintex in the ROOT v5 development release	Achieved. Introduced in the ROOT 5 development release of June
PRF7	30/09/2005	30/09/2005	Demonstration of new PROOF system at ROOT 2005	Achieved. New features of PROOF demonstrated in the ROOT workshop.
MAT1	30/06/2005	30/06/2005	First version of new MathCore library	Achieved. Introduced in the ROOT 5 development release of June
MAT2	31/08/2005	20/09/2005	First version of the MathMore and random number generator library	Achieved. Library released as part of ROOT 5.04/00.
MAT3	30/09/2005		New C++ Minit in be released in ROOT	In progress. Delayed.
GRA1	30/09/2005	30/09/2005	Demonstrate the new GL viewer at the ROOT workshop.	Achieved. Demonstrated in the ROOT workshop.