Applications Area Status

LCG Mini-Review

1st July 2008

Pere Mato/CERN

Outline

- Applications Area Overview
- Current project status
- Manpower
- PH R&D Projects
- Summary

Applications Area Organization

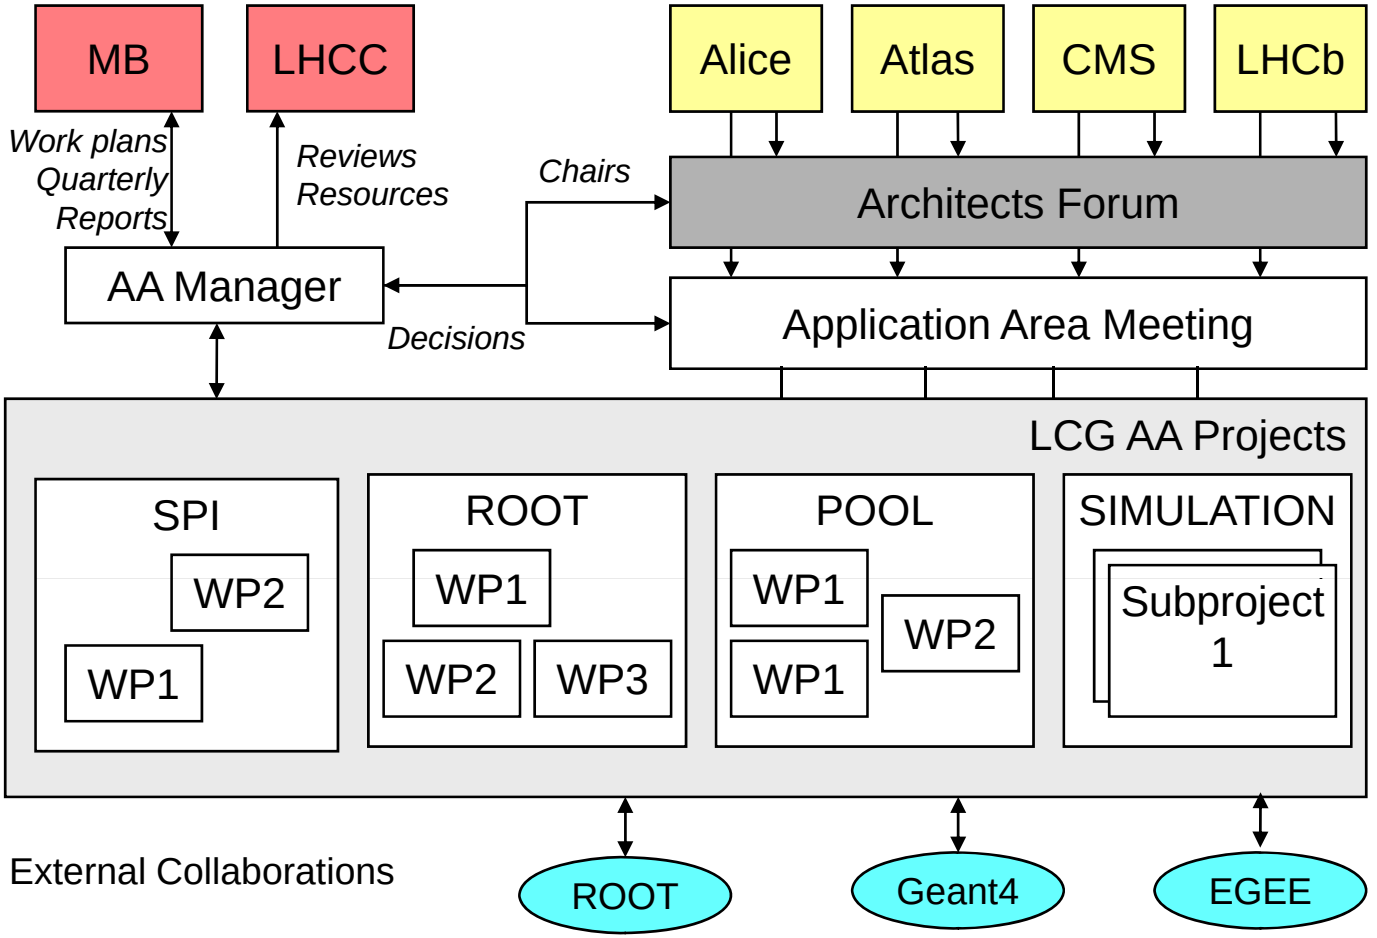

AA Projects

- SPI – Software process infrastructure (S. Roiser)
  - Software and development services: external libraries, savannah, software distribution, support for builds, tests, QA, etc.
- ROOT – Core Libraries and Services (R. Brun)
  - Foundation class libraries, math libraries, framework services, dictionaries, scripting, GUI, graphics, SEAL libraries, etc.
- POOL – Persistency Framework (D. Duellmann)
  - Storage manager, file catalogs, event collections, relational access layer, conditions database, etc.
- SIMU - Simulation project (G. Cosmo)
  - Simulation framework, physics validation studies, MC event generators, participation in Geant4, Fluka.

Unchanged Execution Approach

- Architects Forum (AF) meetings
  - Decision and action taking meetings
    - » Experiments participate directly in the planning, management, and architectural and technical direction of AA activities
  - Every two weeks, public minutes after internal circulation
    - » http://lcgapp.cern.ch/project/mgmt/af.html
- Application Area meetings
  - Informal forum of exchange of information between the AA projects and experiments, etc.
  - Lower frequency than before
- Workplans, quarterly reports and reviews
  - Discussed and agreed in the AF

AA General Highlights

- Configurations in use by experiments
  - Several new configurations (coherent releases of complete software stack including external packages) has been made available in last 6 months
    - » LCG_53, LCG_54, LCG_55
    - » Content details and schedule discussed in AF
  - Stability of basic functionality together with fast turn around for fixes is the common requirement from experiments
- SEAL migration (de-sealed)
  - Finally the CORAL/POOL project have removed dependencies to SEAL. The SEAL migration is now 100%.
  - LCG_55 configuration has been released and is being integrated in LHCb (ATLAS and CMS will do it later)

Software Process & Infrastructure

- Communication services (Hypernews, Savannah)
  - Heavily used by LHC experiments, LCG/AA, Grid, IT, …
  - Mainly user support
  - Adaptation in Savannah for Atlas (for LHC startup)
- Whole software stack successfully ported
  - New compilers (gcc 4.1) – as preview for slc5
  - New platforms (Mac OSX Leopard 32 bit)
- Operational tasks successfully handled
  - Manpower situation will improve again in August with new fellow joining

SPI Nightly Builds

- Used for early verification of code changes, new compilers/platforms and changes in "external" libraries
  - » by LCG Application Area development teams
  - » … and LHC experiments for their integration testing
- Allows faster release builds
  - Usually a complete release is done in ~ 1 day (before in the order of days/weeks)
- The nightly build system itself is also used by other parties for their software builds (LHCb, G4)
- Future developments
  - Speedup by exploiting the full potential of multi-core build servers
  - Distributed builds over several build nodes
  - Configuration via web service

Summaries of nightly builds for LCG AA projects

http://lcgapp.cern.ch/spi/cgi-bin/nightlies.py?day_cbox_today=&slot_cbox_dev2=&slot_cbox_dev1=&proj_cbox_LCGCMT=&proj_cbo

Summaries of nightly builds for LCG AA projects

News

Last update: Wed Jun 18 10:17:51 2008

[+] Display Criteria

[+] Colour coding

Wednesday (Today):

Builds are done every day
in different configurations
with info for builds and tests
on different platforms
including test evolution

Wednesday Slot : dev2 ("gaudipre") Gaudi release validation

Project	Version	slc4_ia32_gcc34 (Wed Jun 18 05:20 2008)		slc4_amd64_gcc34 (Wed Jun 18 02:34 2008)		slc4_amd64_gcc34_dbg (Wed Jun 18 02:57 2008)	
LCGCMT	LCGCMT_55	build	tests	build	tests	build	tests
GAUDI	GAUDI_v20r0-pre	build (8)	tests	build (8)	tests	build (8)	tests
LCGTEST	LCGTEST-preview	build	tests	build	tests	build	tests

Wednesday Slot : dev1 ("desealed") ROOT 5.18 patches + other projects de-sealed

Project	Version	slc4_ia32_gcc34 (Wed Jun 18 09:13 2008)		slc4_ia32_gcc34_dbg (Wed Jun 18 03:18 2008)		slc4_amd64_gcc34 (Wed Jun 18 04:36 2008)		slc4_amd64_gcc34_dbg (Wed Jun 18 02:43 2008)		osx105_ia32_gcc401_dbg (Wed Jun 18 05:46 2008)		win32_vc71_dbg (Wed Jun 18 07:25 2008)	
LCGCMT	LCGCMT_55-patches	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
ROOT	ROOT_5_18_00-patches	build	tests	build	tests	build	tests (2)	build	tests	build (2)	tests (2)	build	tests (22)
RELAX	RELAX_1_1_1x-patches	build	tests	build	tests	build	tests	build	tests	build	tests	build (336)	tests
CORAL	CORAL-preview	build (6)	tests	build (6)	tests	build (6)	tests	build (6)	tests	build (12)	tests (3)	build (194)	tests
COOL	COOL-preview	build	tests	build	tests	build	tests						
POOL	POOL-preview	build	tests	build	tests	build	tests						
GAUDI	GAUDI_HEAD	build (8)	tests	build (8)	tests	build (8)	tests						
GAUDIATLAS	GAUDI_ATLAS	build (21)	tests (3)	build (21)	tests (3)	build (17)	tests						
LCGTEST	LCGTEST-preview	build	tests	build	tests	build (2)	tests						

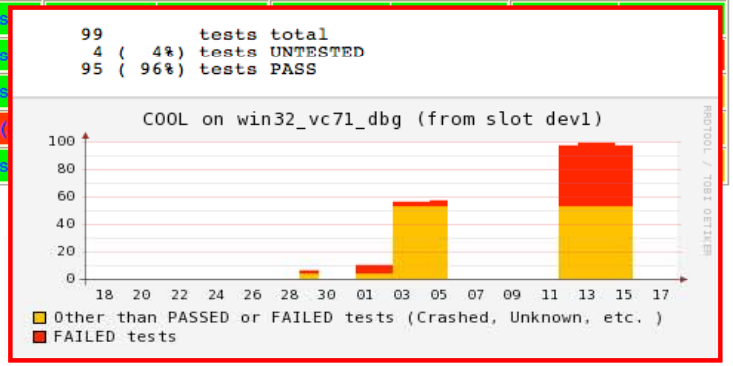

Stefan ROISER
Piotr Kolet

Done

0:36

ROOT Highlights

- The version of ROOT 5.18 was released on January 15
  - Several patches have been made available quickly after detection of problems by experiments
  - This is the baseline version used by the LHC experiments
- Latest production release 5.20 released last week
  - It includes a long list of new features and improvements
    - » Re-structuring of the source repository reflecting better the structure of the project in terms of work packages and to easy the maintenance of the release notes and other documentation
    - » first release of CINT based on Reflex made available for testing
    - » Improvements in speed in many areas, size of dictionaries, etc.
  - Probably most of the LHC experiments will skip this version and wait for the December release

ROOT Extended Schema Evolution

- ROOT I/O needed to support more complex changes in experiments "event models"
  - Be able to read event data in old format into new format also when stored 'column-wise'
  - The existing capabilities were largely insufficient
  - Obviously it requires user-provided transformation codes
- The specification of the extended schema evolution functionality has been defined and agreed by experiments
  - Specs document, presentation in AA meeting and AF discussion
  - An initial implementation is expected by end of August
  - Production quality release by the end of the year after having been validated by the experiments

PROOF

- New functionality driven mainly by the requirements of the ALICE CAF
  - Support for 'Datasets' to simplify data management
  - Package management
  - Better message logging of the sessions
- Currently working on
  - Integration with Condor for scheduling and resource management
  - PROOF-Lite for multi-core machines
  - Improving error handling and recovery

Persistency Framework

- Three deliverables:
  - POOL: Framework for the persistency of arbitrary C++ objects and relationships, with file-based (ROOT) or RDBMS back-ends
  - CORAL: General, technology independent interface to Relational Database
  - COOL: Framework for the handling of detector condition data associated with a time validity
- A "de-sealed" version of the persistency framework has been released
  - AA milestone achieved
  - Maintaining bug-fixes for the SEAL based version for some months
- Replacement manpower in POOL and CORAL fully effective

Status of POOL

- POOL code is production quality and ready for LHC data
  - relational and ROOT based persistency storage with file catalogs
  - no further development planned (and no man-power available)
- Performed review of POOL usage by the LHC experiments
  - none of the experiments exploiting the main motivation for POOL of having common interface to different storage technologies
    - » ATLAS and LHCb using RootStorageSvc
    - » CMS using RelationStorageSvc
  - none of the experiments using the full stack of POOL modules
    - » different parts of POOL functionality now provided by the experiments frameworks

What is used by whom

	Atlas		CMS		LHCb	
	event data	condition data	event data	condition data	event data	condition data
DataSvc	±			x		
PersistencySvc	x	x		x		
StorageSvc	x	x		x	x	
RootStorageSvc	x	x			x	
RelationalStorageSvc				x		
ImplicitCollation	x			x		
RootCollection	x					
RelationalCollection	x					
XMLCatalog	x	x				
LFCCatalog	~					
RelationalCatalog	~					
PyFileCatalog						

Conclusion for POOL

- POOL has provided solutions to a number of persistency-related issues
  - Some of these solutions has been incorporated to other projects (e.g. ROOT, Gaudi)
  - Original motivation of common interface to different persistency technologies turned out not to be relevant
- Plans for the future
  - short term (next months)
    - » minor clean up of the code (new CVS set up)
    - » updates for new versions of ROOT
  - medium term (over the next year)
    - » study of possible common approach to persistency (ROOT based) by ATLAS and LHCb (integrated in Gaudi/Athena)

CORAL

- Undertaken a review of outstanding bugs - many smaller items fixed/closed
- Started a CORAL server prototype
  - Requested by the ATLAS online community
  - A way to cope with some technical issues such as number of the database connections
- CORAL server prototype is build together with production components since CORAL 2.0.0
  - simple read-only tests pass, full read-only test suite (including COOL) are expected to pass by the end of the next week
  - First read-write production release of CORAL server planned for end of October

COOL

- Performance optimization for most IOV retrieval use cases is complete
  - For each use case, performance is now routinely tested in a variety of scenarios which have given problems in the past
  - Work is ongoing to simplify the code so that all use cases share a few well defined methods defining the optimized SQL strategies with good and stable performance
- A few functional enhancements have been provided as requested by the experiments (for the management of channel metadata and tags)
- Basic support for the CORAL server has been added and COOL tests defined

Simulation Project
Structure and tasks

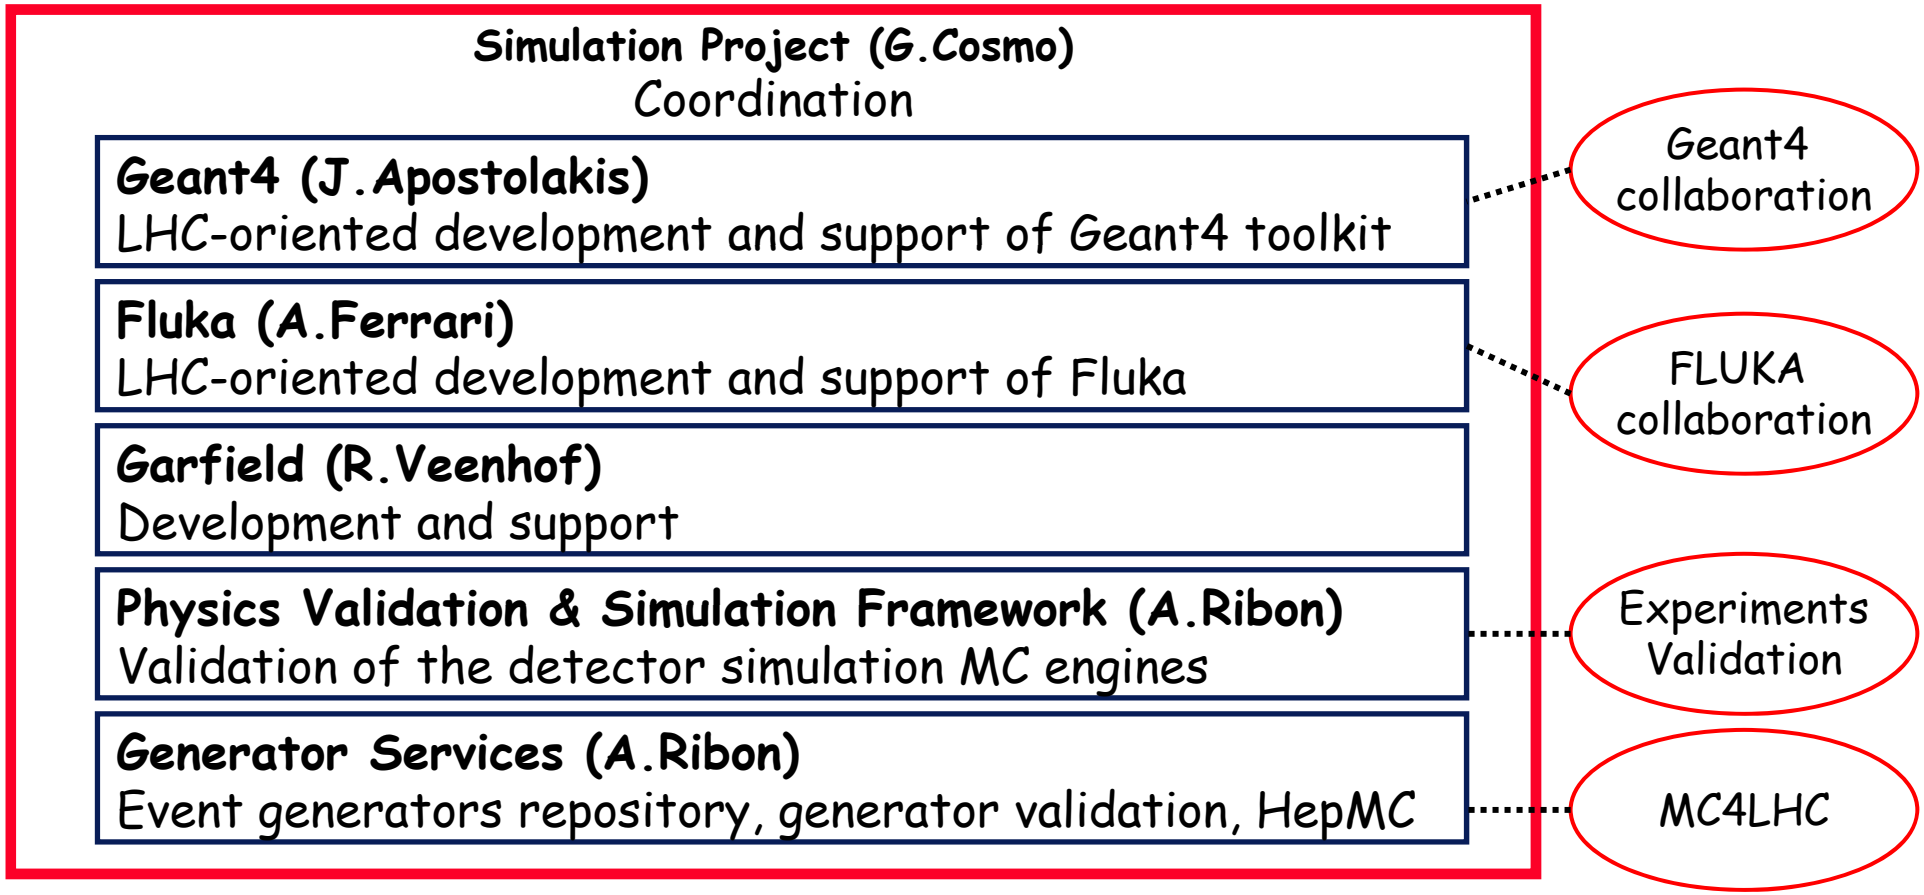

Geant4 - Highlights of 2008 (1) Release 9.1

- New physics models
  - First alpha-version of Liege intra-nuclear cascade
  - Native QMD low-energy nucleus-nucleus model
  - Re-scattering interface to Binary cascade
    - » Potential re-interaction inside nucleus of the product of a high-energy interaction
  - New physics-list options for analysis of test-beam data
    - » Inclusion of diffraction model for the description of shower shapes
- Better in CPU performance in hadronic physics (5-10 %)
- Improvements in electromagnetic physics
  - Improved straggling for ions at low energies
  - Revised high-energy gamma tail for muon Bremmstrahlung (NA49)
- New GDML plug-in for importing detector description setups
- New command-based scoring capabilities
- New optimised navigation technique for regular patterned geometries

Geant4 - Highlights of 2008 (2)
Interaction with experiments & feedback

- Patches to existing releases (8.3, 9.0, 9.1)
  - Improvements in event reproducibility and overall robustness
  - Fixes for problems reported by experiments in pre-production runs
    - » Fixes to quasi-elastic scattering following reports by HARP-CDP and NA61
    - » Inclusion of very short-live particles tracking (CMS)
- Assessment of experiments' needs for releases to be used in production
  - Convergence on adoption and support for the most recent releases
- New developments (9.2-β release in July 2008)
  - Improvements to Fritiof hadronic model for pion incident interactions
  - New alternative multiple-scattering model (optimal handling of large/small angles)
  - First implementation of GDML writer, completing support for GDML persistency
- Geant4 paper (NIM A 506 (2003) 250-303) designated as a "Current Classic" (Thomson Scientific's Essential Science Indicators)
  - http://sciencewatch.com/dr/cc/08-juncc/

LCG Generator Services

- GENSER
  - Structure stable and used by experiments
  - 24 generators installed (most with different versions)
  - More generators built also on Windows (for LHCb)
  - Extended set of tests
  - Evaluation of autotools for building all generators
- HepMC
  - New release process, more openly discussed and agreed
    - » HepMC 2.04.00 released on 12 June
      - Units are now mandatory: GEV or MEV ; MM or CM
- MCDB
  - Automatic uploading and downloading of samples
  - Integrated in the CMS software framework
  - To be tested in large, Grid-based productions

LCG Physics Validation

- Extension to Fluka of the ATLAS TileCal 2002 test-beam analysis has been concluded. Reasonably good agreement of Fluka and Geant4 QGSP_BERT with data.
- Including instrumental effects:
  - Readout time window
  - Birks quenching effects
- Show better agreement with data and consistency between different types of calorimeters.
- Main focus now on energy response and resolution
  - Non-smooth transitions between hadronic models
  - Try to reduce the use of parametrized models in theory-driven Physics Lists
- Most interesting Physics Lists: QGSP_BERT, FTF_BIC.

Hadronic Shower Shapes in Geant4

- Starting with version 8.3, the proper modeling of quasi-elastic and the inclusion of Bertini cascade model have improved significantly the description of hadronic shower shapes for pions.
- QGSP_BERT is now the default Physics List in ATLAS

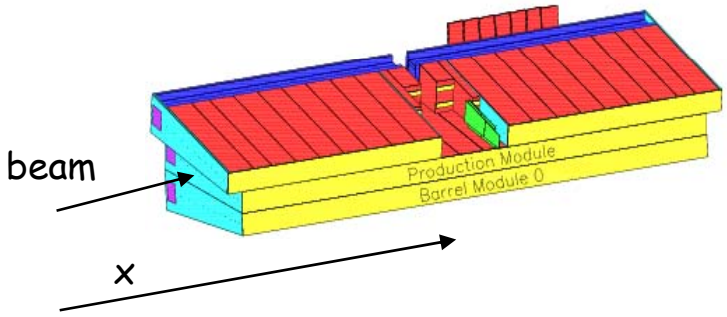

T. Carli, M.Simonyan, October 2007

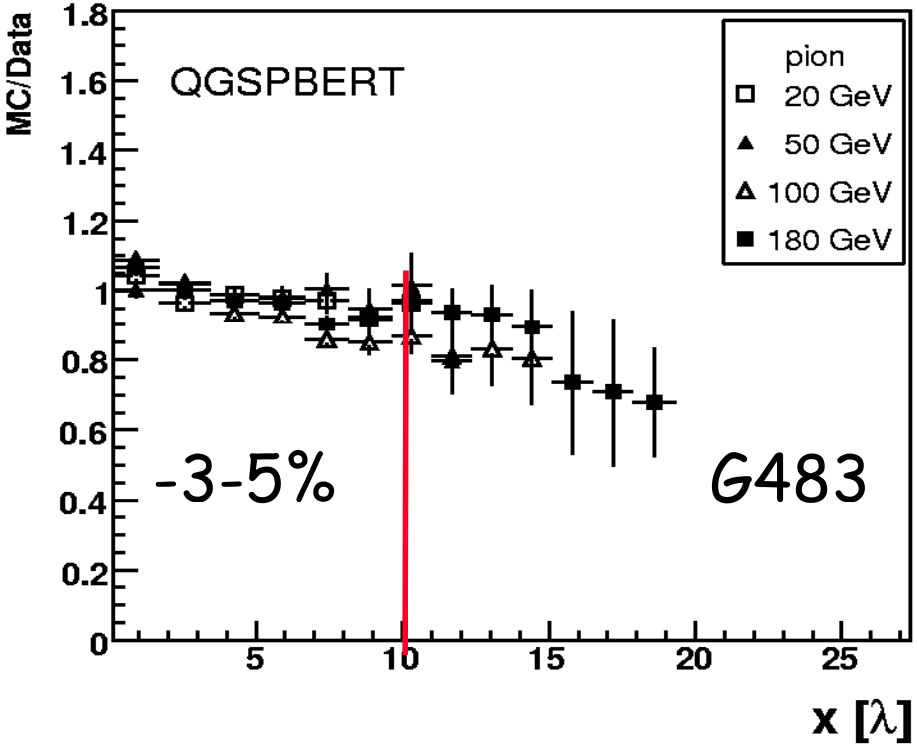

Manpower Table

Sum of FTE(2008)		Afiliation											
Project	Sub-Project	*ALICE*	*ATLAS*	*BNL*	*CERN*	*CMS*	*EGEE*	*FNAL*	*LCG-Russia*	*LCG-Spain*	*LHCb*	*Other*	*LCG-Italy*	***Grand Total***
MGR	(blank)				0.9									0.9
MGR Total					0.9									0.9
POOL	Catalog		0											0
	Collections		0.2											0.2
	Cool				0.8						0.2			1
	StorageMgr		0.2		0.9									1.1
	Coral				1.8	0.1								1.9
POOL Total			0.4		3.5	0.1					0.2			4.2
ROOT	Base				1.2			0.1						1.3
	Dictionary		0.1		1.6			0.2						1.9
	Geom	0.2												0.2
	Graf	0.6			1							0.2		1.8
	Gui			0.1	0.9									1
	I/O							0.7						0.7
	Math				1.9									1.9
	Mgr				1									1
	Proof				2.1									2.1
	Seal				0.1									0.1
ROOT Total		0.8	0.1	0.1	9.8			1				0.2		12
SIMU	Framework				0.5									0.5
	Garfield													
	Geant4				7.95									7.95
	Genser				0.5	0		0.1	1.9					2.5
	Mgr				0.25									0.25
	Validation				0.4									0.4
SIMU Total					9.6	0		0.1	1.9					11.6
SPI	(blank)				2.8									2.8
SPI Total					2.8									2.8
Grand Total		0.8	0.5	0.1	26.6	0.1		1.1	1.9		0.2	0.2		31.5

Estimated Resource Needs

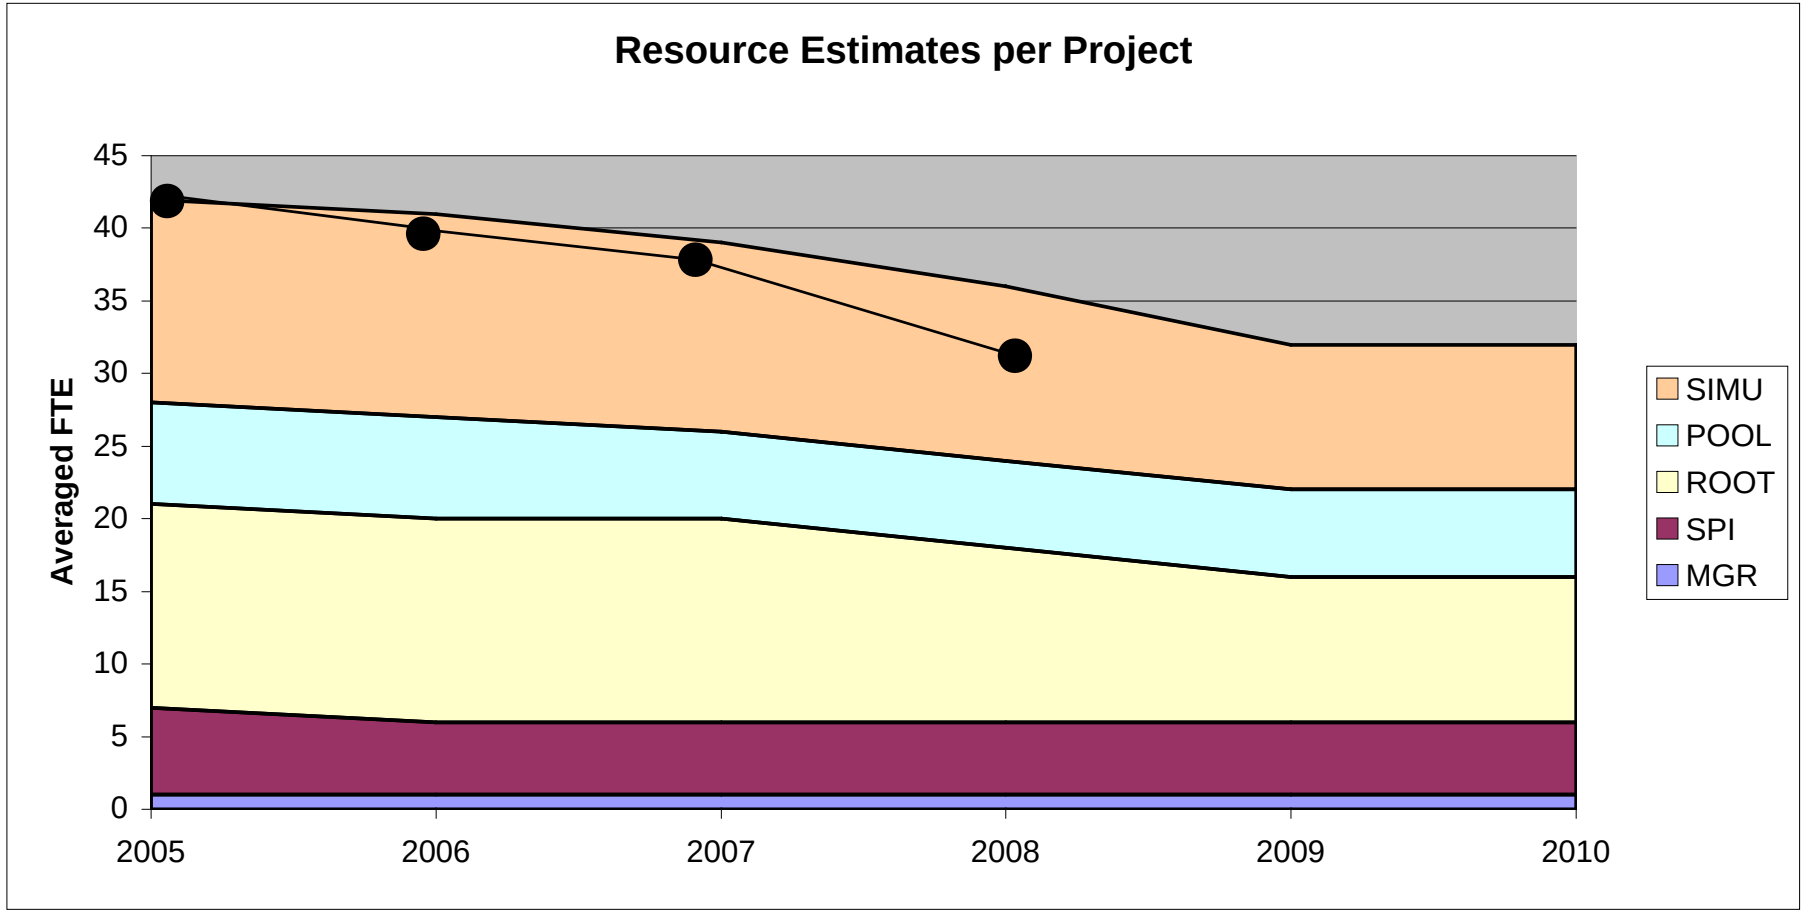

Staffing Remarks

- Big reduction as (more or less) expected in 2008 and 2009
  - No full replacement of LD contracts
    - » In PH: from 3 LD -> 0 IC + LD extensions
    - » In IT: from 4 LD -> 1 IC + 2 LD
  - Retirement
- Some activities are very much affected
  - Several milestones in Physics Validation on hold
  - Geant4 testing also affected
  - Rationalization of deliverables of Persistency Framework
  - PROOF development is also at big risk

PH - R&D Projects

- WP8 - Parallelization of Software Frameworks to exploit Multi-core Processors
  - Investigate current and future multi-core architectures
  - Measure and analyze performance of current LHC physics application software on multi-core architectures
  - Investigate solutions to parallelize current LHC physics software at application framework level and also investigate solutions to parallelize algorithms
- WP9 - Portable Analysis Environment using Virtualization Technology
  - Evaluation of the available virtualization technologies (virtual appliances)
  - Deployment of a read-only distributed file system with aggressive caching schema
  - Collect requirements from experiments and confront them with available technologies
  - Development of the "CernVM" prototype

PH - R&D Projects

- Both projects are hosted in PH-SFT with participation of LHC experiments
- Kickoff workshop took place in April with good participation from the LHC experiments and technology vendors
  - http://indico.cern.ch/conferenceDisplay.py?confId=28823
- Regular monitoring and control by the LHC experiments via the Architect's Forum (AF)
- Consultation and collaboration with IT department to ensure integration with IT services
  - Privileged relations with OpenLab

Summary

- Applications Area has continued to provide the software and services required by the experiments
  - The organization is mature and works reasonable well
  - Software functionality sufficient for first data-taking year
- Improved testing and integration of the complete software stack and introduced new release procedures
  - Software releases can be done very quickly on demand
  - Ready for LHC turn on
- Reduction of manpower is affecting seriously some of the projects
- Started the two new R&D work packages and put them under the existing control and monitor structure