

Computing Status

F. Bianchi
Torino

II SuperB Collaboration Meeting
Frascati, Dec 13, 2011

Outline

- Collaborative Tools.
- Software Management Tools.
- FastSim & Physics Tools.
- FullSim.
- Distributed Computing.
- Computing model & Pon ReCaS.
- R&D.

Web Portal: superb.infn.it

- Replace old sites.
 - Automated redirection to the portal is in place.
- Contents:
 - Public/private web pages for each division.
 - Access to Alfresco Doc Repository
 - Access to Wiki
 - Forum (not able to automatically send mails when a new message is added)
 - Should be fixed in a next release of Liferay
 - SuperB Calendar
- Not much into it yet.
 - There is a word like editor that allows to edit web pages
 - Alternative: send me pages in word and I'll take care of having them uploaded.
- The portal will soon be updated to the Liferay version 6.1

Administrative Database

- A «Master» database, updated with all SuperB collaborators, Institutions, Funding Agencies, affiliations, etc . is online
- The directory service is partially synchronized with the DB (to be improved)
- The web portal also provides an interface to the Administrative Database.
- A new procedure - linked with the administrative DB - to subscribe new SuperB users is online (with the old GUI, need to rewrite the interface).

Mailing Lists

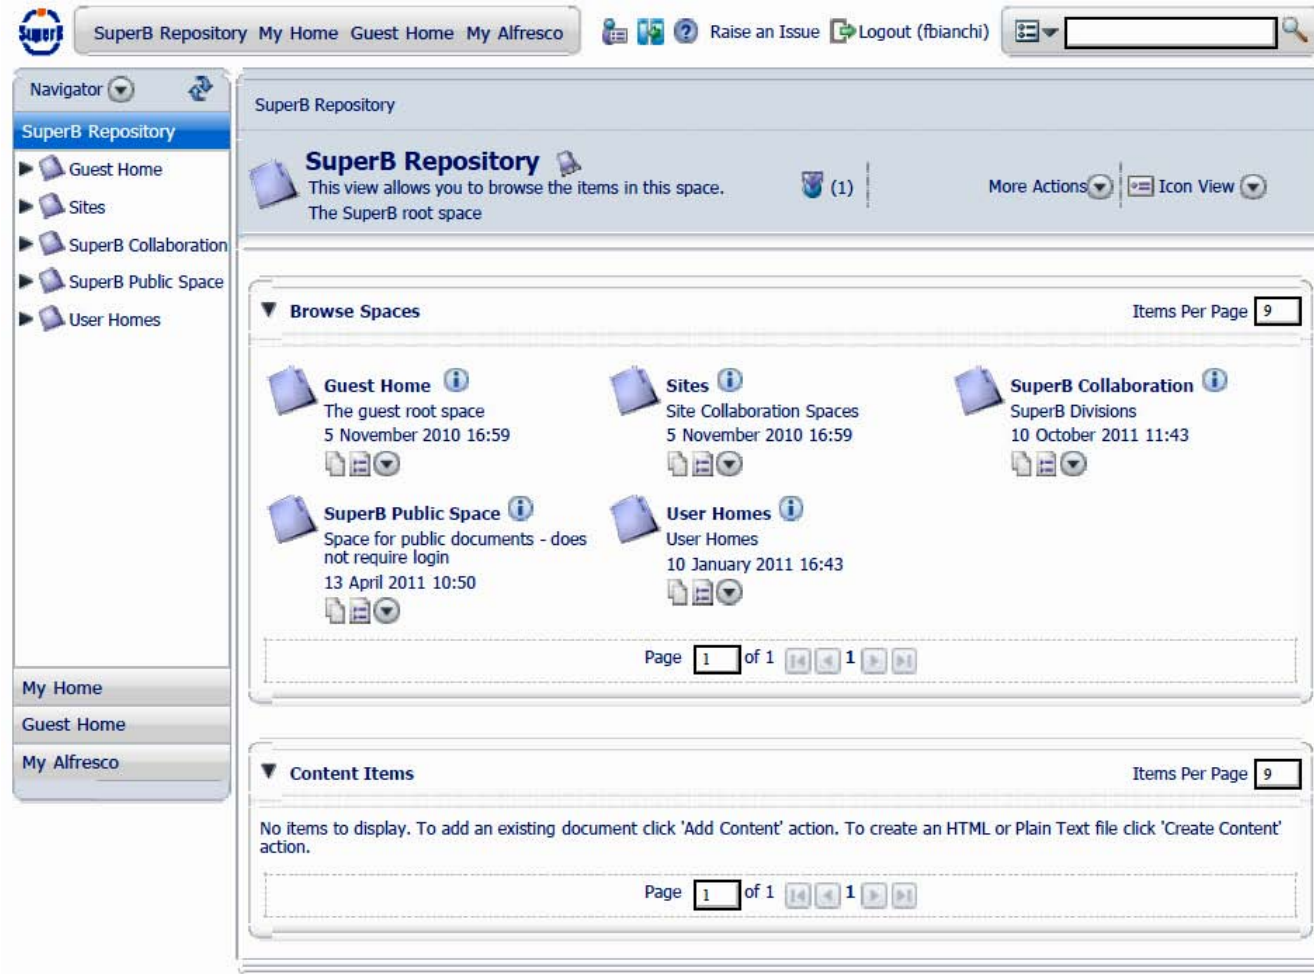
- Ready to integrate Sympa mailing lists with SuperB LDAP authentication, to automatically manage subscription to SuperB «general» mailing lists

Wiki

- Is now R/W protected.
 - Has been announced to the superb-members mailing list.
- LDAP authentication is required.
 - Two new LDAP group for Theorists and Computing Experts associated to SuperB have been created.
 - These are people who need to access internal SuperB documentation.
 - For the time being they have the access to the whole Wiki
 - This will be fixed soon.
- Mirroring of BaBar documentation into SuperB Wiki is now possible under the agreement between BaBar & SuperB.

Alfresco (1)

- Go to:
<https://sbdocserv.er.pd.infn.it:5210/alfresco/>
- Login with your LDAP account and you will land on this page.
- Click on "SuperB Collaboration".



Alfresco(2)

- And you will arrive on this page.
- You can also start from the web portal
 - Click on "Documents", then choose "Browse Documents" and you will land to this same page

The screenshot displays the Alfresco SuperB Collaboration web portal. The top navigation bar includes links for 'SuperB Repository', 'My Home', 'Guest Home', and 'My Alfresco', along with a search bar and a 'Logout (fbianchi)' button. A left-hand 'Navigator' pane shows a tree structure with 'SuperB Repository' selected, containing sub-items like 'Guest Home', 'Sites', 'SuperB Collaboration', 'SuperB Public Space', and 'User Homes'. Below this are buttons for 'My Home', 'Guest Home', and 'My Alfresco'. The main content area is titled 'SuperB Collaboration' and shows a list of 'Browse Spaces'. Each space entry includes a folder icon, a name, a division, and a timestamp. The spaces listed are: Accelerator (Accelerator Division, 29 March 2011 18:04), Computing (Computing Division, 15 March 2011 11:42), Detector (Detector Division, 4 April 2011 09:39), General (General Division, 26 October 2011 16:45), Physics (Physics Division, 5 April 2011 09:43), Project Board (Project Board, 20 October 2011 18:32), Proto-Council (Proto-Council, 20 October 2011 23:15), Site (Site Division, 5 April 2011 10:42), and Steering Committee (Steering Committee, 20 October 2011 18:33). Each entry has icons for document, folder, and download. At the bottom, there is a 'Content Items' section which is currently empty, displaying a message: 'No items to display. To add an existing document click 'Add Content' action. To create an HTML or Plain Text file click 'Create Content' action.'

SuperB Repository My Home Guest Home My Alfresco

SuperB Collaboration

SuperB Collaboration
This view allows you to browse the items in this space.
SuperB Divisions

Browse Spaces

Items Per Page 9

Space Name	Division	Last Modified
Accelerator	Accelerator Division	29 March 2011 18:04
Computing	Computing Division	15 March 2011 11:42
Detector	Detector Division	4 April 2011 09:39
General	General Division	26 October 2011 16:45
Physics	Physics Division	5 April 2011 09:43
Project Board	Project Board	20 October 2011 18:32
Proto-Council	Proto-Council	20 October 2011 23:15
Site	Site Division	5 April 2011 10:42
Steering Committee	Steering Committee	20 October 2011 18:33

Page 1 of 2

Content Items

Items Per Page 9

No items to display. To add an existing document click 'Add Content' action. To create an HTML or Plain Text file click 'Create Content' action.

Page 1 of 1

Alfresco (3)

- There is a space per division.
 - Plus a *General* space.
 - Plus restricted spaces that can be seen only by board members.
- Click on a space and you will see the documents into it
 - Or use the “Advanced Search” button to perform keyword based searches.
- Everyone with a LDAP account can now upload a document.
 - No authorization procedure is required.
 - An uploaded document is immediately visible.
- Documents are automatically numbered:
 - SB-<DivisionName>-<Year>-<number>
 - Es: SB-COM-2011-002
- Repository and web interfaces updated to version 4.0b

Sharing Documents

- A clear requirement.
- Alfresco Share 4.0 (team collaboration) deployed and integrated inside the SuperB portal, with LDAP authentication.
 - In evaluation phase: waiting for user feedback.

High Availability Cluster

- All the Collaborative Tools will be migrated to an High Availability cluster.
 - To have more stable hardware.
 - To have better response time (avoid timeout with some Java clients)
- Purchase of the cluster system is in progress.
- The design is done, we will proceed with the implementation as soon as all the hardware resources will be available.

Software Management Tools:

CMake (1)

- One of our goal was to get definitely rid of SRT as a base for building FastSim.
- Not completely accomplished but now we can avoid using old .mk files to declare dependencies among packages.
- dependencies are now handled completely by CMake via the usual CMakeLists.txt file
 - Configuration is now much faster

CMake (2)

- The user can specify, with the bash script `cmakeRun.sh`, the target he wants to build
 - E.g. only the binaries of a particular set of packages.
- Support for CMake is available also for the FullSim releases
 - Added a specific macro to manage Geant4 externals, plus some minor modifications to the main `CMakeLists.txt`
- Some work has started to try managing also packaging of releases

FastSim

- V0.2.7 → V0.3.x (CLHEP migration)
 - a few iterations following users' feedback
 - validation close to completion
- New round of development of detector response:
 - fDIRC
 - EMC
 - IFR
 - details in the Physics Tools session on Thu @11:30
- New background frames available
 - thanks to the bkg group
- A few bug fixes
 - e.g.: simulation of particle decays in flight

FastSim and Physics Tools

- Babar → SuperB tools documentation migration started
- Development of User Packages
 - hadronic recoil, tau, (charm) tagging, ...
- Ongoing activity to get ready for next Physics production (2012)

Joint Physics-Computing Physics session on physics tools

Thursday Dec 15 at 11:30

11:30->13:35 **Parallel VIII: Computing + Phys - Physics Tools** (Convener: Matteo Rama (LNF) / Aura Seminari)

Description:

Phone number: +39 06 6228 8548

or http://server10.infn.it/video/index.php?page=telephone_numbers

Meeting ID: 1556

11:30	Overview (10')	Matteo Rama (LNF)
11:40	fDIRC development status and plan (15')	NICOLAS ARNAUD (LAL ORSAY CNRS-IN2P3)
11:55	EMC development status and plans (15')	Chih-hsiang Cheng (Caltech)
12:10	IFR development status and plans (15')	Marcello Rotondo (PD)
12:25	FastSim tools for charm threshold running (15')	Gianluca Inguglia (Queen Mary University of London)
12:40	FastSim - BaBar FullSim comparison with HAD Breco (15')	Elisa Manoni (PG)
12:55	Physics needs (15')	John Walsh (PI)

FullSim: Optical Photon Simulation

- Thanks to work done by D. Roberts and A. Di Simone the optical photon simulation has reached a point where it makes sense to start using it in production
- On top of what already implemented in Bruno, it required a few minor tunings to BrnCore (Andrea), and quite some work from the detector side (Doug)
 - All collected into a new package BrnPID
- Minor tunings still needed from the BrnCore side, but the present implementation is *already* being used in production
 - Will provide extremely useful data to the PID experts

FullSim: Metadata Handling

- Complementary to bookkeeping database
- Ensures some level of redundancy for critical information
- "Metadata"= all information which does not change from one event to the other
 - Software release, OS, architecture
 - Geometrical details
 - Readout configuration
 - Simulation configuration
 - ...
- Implemented a general (though fairly simple) metadata handling
 - Can save and retrieve any information, both from macro files and on the C++ side
 - All metadata is automatically saved to the output file
 - If multiple outputs are produced (like in the bg frames), the information is copied to each of them
- *Already* in production
- Documentation available in the usual wiki
 - http://mailman.fe.infn.it/superbwiki/index.php/Geant4_SuperB_simulation_main_portal/Metadata

FullSim: EMC Deposit Timing

- Request was to be able to study evolution of energy deposits in time
- Implementation done by S. Tammaro (Tor Vergata)
 - User chooses the binning at runtime (in a macro file)
 - Bruno writes the binned information in hits as an array (TArrayD)
 - Existing integrated deposit left unchanged to smoothen the migration
 - Binning info written as metadata for offline analysis
 - Now waiting for first feedback from the detector experts
 - If the implementation is good enough, will push it forward and release it

Background Production: Software Validation

- The production for this meeting has been the first one for which we followed the validation procedure outlined at the last Elba meeting
 - Crucial to ensure good quality of the results
- This first iteration was rather informal. But all steps were followed
 - Feature-freeze, performance optimization, resource estimate
 - Pre-release validation, with sign-offs by experts
 - Post-release validation
 - 10% of total production, sign-offs from experts/analyzers required
 - Full production
- This resulted, as expected in some delays in the production
 - Time needed for the full cycle was ~10-15days
 - With time, we could cut this down to 7-10days, but it is unlikely that we will ever be able to be faster than that
- This resulted also, most importantly, in a “good” release, whose results we know will be free from any major bugs
 - i.e. no need to rerun all the production once again after this meeting because for some reason the present results are unusable...
 - indeed, we felt so confident in the good quality of the software we had, that we added one last-minute feature to the production (optical photons)
 - without a solid validation process, feature-freeze must be strictly enforced

Distr. Comp.: Production System

- The new system has been fully redesigned: database schema, job script, web pages and abstraction layers
- Improvement of base functionality:
 - Production definition
 - Submission interface for - on shift - people and for admin purpose
 - Monitor system: submission monitor plus job monitor integrated interfaces
- In progress:
 - Integration with Grid service availability monitor, nagios based
 - Integration with new Grid Logging service features
- Distributed Computing session on Wed Dec 14 11:00 - 13:00

User Tools for Grid Access

- The GANGA SuperB layer is under development.
- Use cases: analysis and personal production
 - Automatic job data preparation
 - Automatic running site selection
 - Job stage out setup configuration
 - On line job monitoring
 - Integration with bookkeeping DB
- Ganga hands on session on Wed Dec. 14 15:00 – 16:30
- Everyone is invited to participate !

Distributed Resources

- 25 sites are available to the SuperB VO.
 - From: Canada, France, Italy, Poland, UK and USA

Site	Min (cores)	Max (cores)	Disk (TB)	SRM layer	Grid Org.	Site contacts
RAL(T1)	200	1000	25	Castor	EGI	F. Wilson, C. Brew
Ralpp	50	500	5	dCache	EGI	F. Wilson, C. Brew
Queen Mary	300	3456	120	StoRM	EGI	A. Martin, C. Walker
Oxford Univ.	50	200	1	DPM	EGI	K. Mohammad, E. MacMahon
IN2P3-CC(T1)	500	1000	10	dCache	EGI	N. Arnaud, O. Dadoun
Grif	50	300	2	DPM	EGI	N. Arnaud, O. Dadoun
In2p3-lpsc	50	100	2	DPM	EGI	J.S. Real
In2p3-lres	50	100	2	DPM	EGI	Y. Patols
CNAF(T1)	500	1000	127	StoRM	EGI	A. Fella
Pisa	50	500	0.5	StoRM	EGI	A. Ciampa, E. Mazzeni, D. Fabiani
Legnaro	50	100	1	StoRM	EGI	G. Maron, A. Crescente, S. Fantinel
Napoli-infn	50	200	5	DPM	EGI	S. Pardi, A. Doria
Napoli-grisu	50	300	5	DPM	EGI	S. Pardi, A. Doria
Napoli-unina	50	300	5	DPM	EGI	S. Pardi, A. Doria
Bari	160	260	0.5	StoRM/Lustre	EGI	G. Donvito, V. Spinoso
Ferrara	10	50	0.5	StoRM	EGI	L. Tomassetti, A. Donati
Cagliari	10	50	1	StoRM	EGI	D. Mura
Perugia	10	50	1	StoRM	EGI	R. Cefala
Torino	50	100	2	DPM	EGI	S. Bagnasco, R. Brunetti
Frascati	50	100	2	DPM	EGI	E. Villucchi, G. Fortugno, A. Martini
Milano	50	100	2	StoRM	EGI	N. Neri, L. Vaccarossa, D. Rebatto
Catania	?	?	?	StoRM	EGI	G. Platania
Slac	400	400	10	NFS	OSG	S. Luiz, W. Yang
Caltech	200	400	4.5	NFS	OSG	F. Porter, P. Ongmongkolkul, S. Lo
OhioSC	?	?	?	dCache	OSG	R. Andreassen, D. Johnson
Victoria	50	100	5	dCache	EGI	A. Agarwal
Cyfronet	100	500	10	DPM	EGI	L. Flis, t.Szepienie, J. Chwastowski
Total	3090	11166	349			

Distributed Resources (2)

- Recent addition in Napoli:
 - New cluster with 200 cores, 50 TB storage, 10 Gb network available to R&D program

SuperB Computing Model

- Baseline is an extrapolation of BaBar computing model to a luminosity 100 times larger.
- “Raw data” from the detector will be permanently stored, and reconstructed in a two step process:
 - a “prompt calibration” pass on a subset of the events to determine calibration constants.
 - a full “event reconstruction” pass on all the events that uses the constants derived in the previous step.
- Monte Carlo data will be processed in the same way.
- Selected subset of Detector and MC data, the “skims”, will be made available for different areas of physics analysis.
 - Very convenient for analysis.
 - Increase the storage requirement because the same events can be present in more than one skim.
- Improvements in constants, reconstruction code, or simulation may require reprocessing of the data or generation of new simulated data.
 - Require the capability of reprocessing in a given year all the data collected in previous years.

An Attempt to Estimate the SuperB Detector Computing Needs

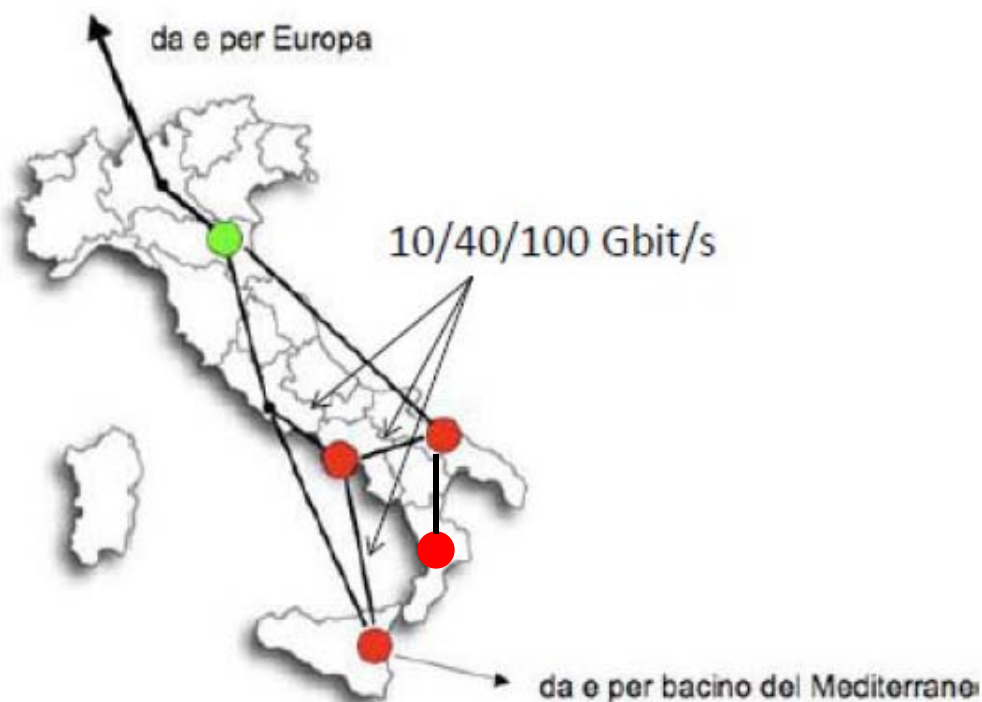
- Limited precision due to many assumptions:
 - Raw Event size $\sim 100\text{kByte}$ ($= 3 \times \text{BaBar}$)
 - Mini/Micro event size $= 2 \times \text{BaBar}$
 - CPU / unit lumi: $3 \times \text{BaBar}$
 - 2 copies of raw data
 - Skim expansion factor: 5
 - Some fraction of Mini on disk (100% \rightarrow 10%)
 - Equivalent amount of MC "lumi"
 - Raw data stored on tape
- Storage grows from $O(50)$ PB to $O(600)$ PB in 6 years.
- CPU grows from 500 to 12,000 KHepSpec in 6 years.

Cost and Challenges

- Crude estimate: O(5 Million Euros/year) assuming:
 - Validity of Moore's law.
 - Start of data taking in 2016.
 - Purchase the hardware one year in advance.
 - Infrastructure and power cost are not included.
- But:
 - Is Moore's law still valid ?
 - Code must be optimized for running on multi/many core architectures.
 - How to access efficiently and reliably hundreds of PB of data ?
 - Identify a strategy to avoid I/O bottleneck.
 - Understand how to share and replicate data among sites.
 - How use efficiently and reliably hardware resources widely dispersed ?
 - Adopt/develop a distributed resource management framework.
- R&D program is in place to address these issues.
 - Goal is to complete R&D and write a Computing TDR by mid 2013
 - Have a first version of the code by end 2014

Computing Infrastructure

- In Italy:
 - CNAF
 - 4 new centers in Bari, Catania, Cosenza, Napoli
- + centers in other participating countries



Pon ReCaS

- Approved (~45% of proposals have been approved).
- Funded with 13.7 ME
- Many constraints:
 - Money has to be spent by Dec 2014.
 - Money is allocated to each site (Bari, Catania, Cosenza, Napoli).
 - Within a site, money is allocated to various items (infrastructure, hardware, personnel, etc.)
- Moving money between site is not possible.
- Moving some money between items may be possible (to be negotiated with ministry)..

R&D: Parallelism in the framework

R&D group: V. Ciaschini, M. Corvo, P. Franchini, F. Giacomini, S. Longo, R. Stroili

- Our goal is to express the current framework's "Module Sequences" in terms of a graph of modules and dependencies among them
 - independent modules can potentially run in parallel
- We are investigating Intel® Threading Building Blocks (TBB)
 - TBB provides high-level, task-based parallelism primitives (in C++) that abstract platform details and threading mechanisms for scalability and performance
- Some tests are on-going on the current framework as proof-of-concept
- We are currently following also the R&D activities of other groups who are working on the same issues

R&D: Distributed Storage

- People: G. Donvito (Bari), D. DelPrete(Na), D. Diacono(Ba), A. Fella (Pisa), E. Luppi (Ferrara), M. Manzali (Ferrara), S. Pardi (Napoli), B. Santeramo (Bari), , G. Russo(Na), L. Tomassetti (Ferrara) + some Ferrara students
- R&D Data access:
 - HTTP and WebDAV in Wide Area Network scenarios
 - collaboration setup with CERN Data Management group and CNAF StoRM team
- Distributed Tier1 model study:
 - functional evaluation, data distribution policy, resource allocation
- Data access and storage system evaluation:
 - gLuster, NFS4, Hadoop, EOS
- Distributed Computing session on Wed 14 Dec. 11:00 - 13:00

Conclusions

- The computing group is supporting the collaboration by providing:
 - Collaborative Tools
 - Physics Tools: FastSim, etc.
 - FullSim
 - Production Tools
 - Bookkeeping Tools
- There is an active R&D program aimed at the design of the computing model.
- The approval of the Pon ReCaS is an important step forward into building the computing infrastructure.
- A severe lack of manpower is affecting us.
- Come and join the fun !