SuperB Computing

Status and goals for the meeting

Armando Fella – INFN CNAF
Representing the computing group

Annecy – 16 March 2010

Topics

- Distributed infrastructure setup
- Full Simulation production
- Fast Simulation production
- Computing R&D Workshop
- Goals of this meeting

February production

- Full sim and Fast sim largest productions ever
- Goals defined in Frascati achieved for full sim, for Fast sim: new functionality **OK**, production target partially achieved
  - O(1 Million) full sim events: **OK**
  - 1 ab^{-1} fast sim generics productions for det/geom. studies and physics studies + smaller samples: **0.25 ab^{-1} produced + 10^8 signal events produced**
    - no urgent case for background mixing: **O(0.01 ab^{-1}) actually produced**
- Large participation of sites to the production

Web UI production manager

- We have completed the production management GUI used for both Fast and Full Sim
- The interface has been fully exploited for all distributed job management
- Main components:
  - Bookkeeping DB interface to keep track of all produced data and conditions
  - Submission portal to generate the jobs execution scripts
  - Production monitor
- See L. Tomassetti talk in Parallel Comp. session Wed. 16-17:30

Full Grid integrated production model

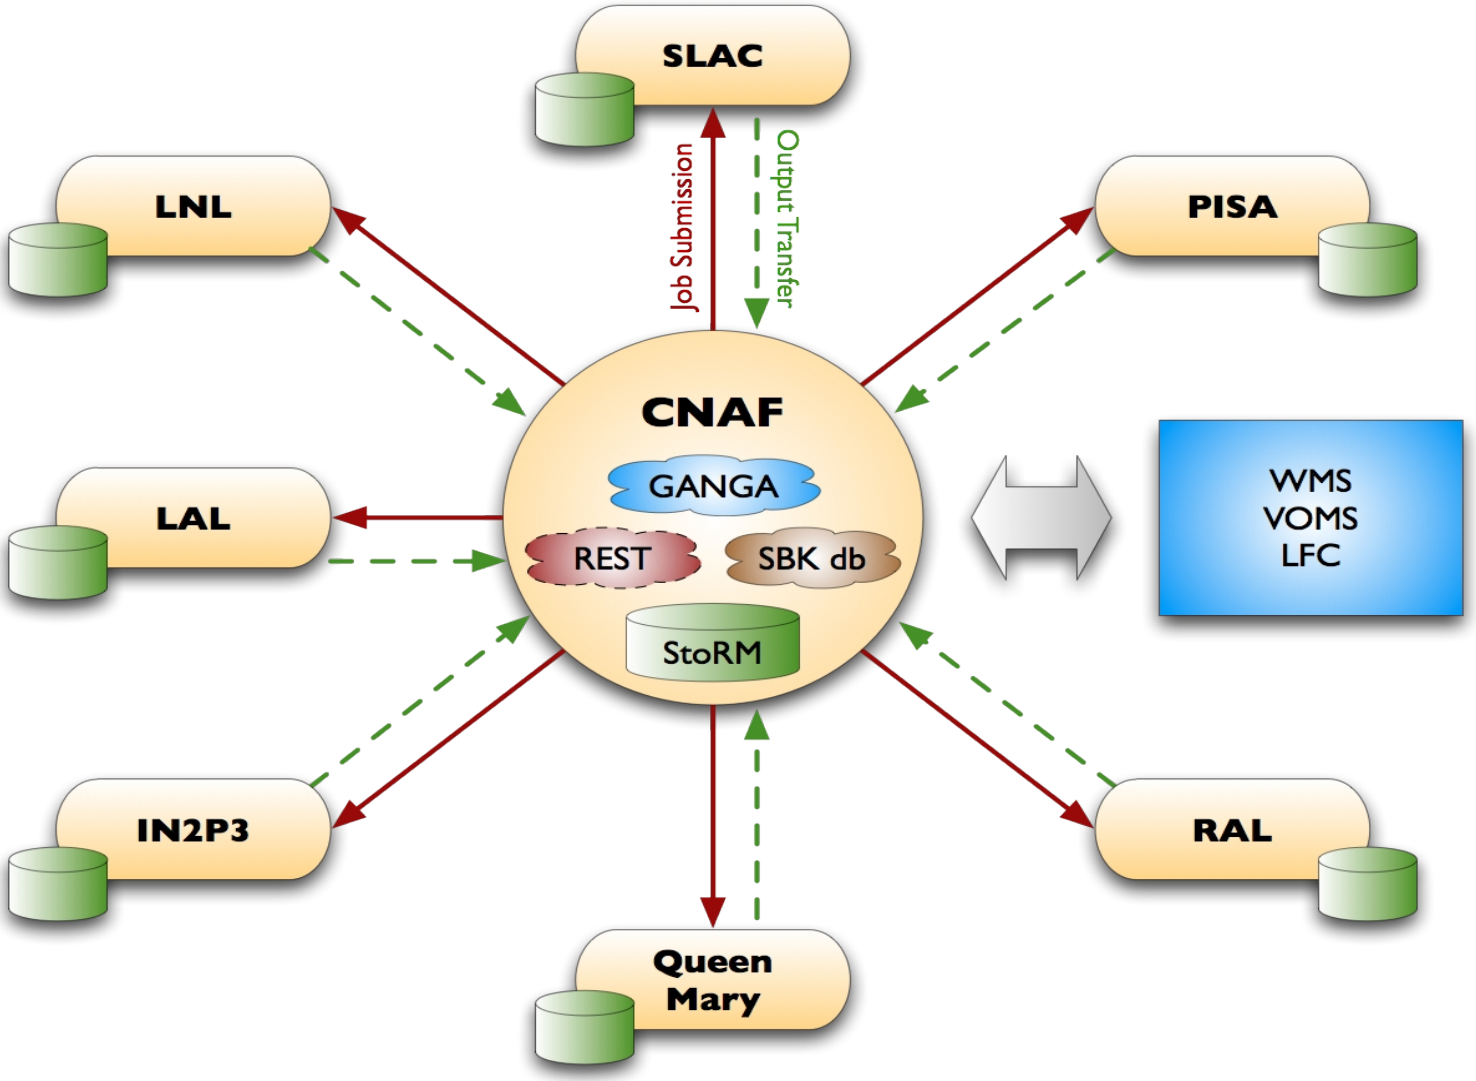

Full Grid integrated production work-flow

- The job input files, Bkg data and test release, are transferred via LCG-Utils to the Storage Elements of remote site
- The job submission is performed by GANGA on User Interface at CNAF
- The WMS routes the jobs to the matched site
- The job is scheduled by the site Computing Element to a Worker Node
- The job run time tasks are:
  - access the DB for initialization and status update via REST interface
  - retrieve/access input files by local site Storage Element
  - transfer the output to the CNAF Storage Element

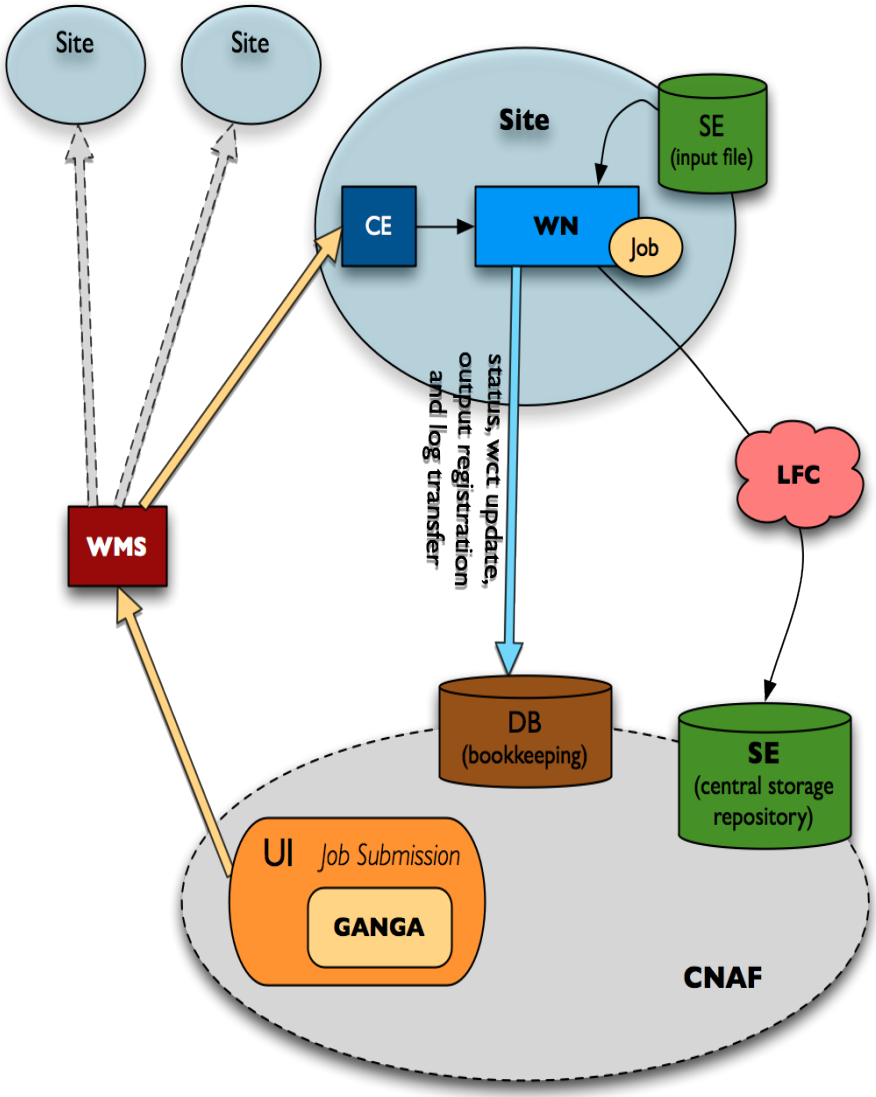

Sites setup and coordination

- Grid services setup at CNAF (detailed status follow)
- Distributed infrastructure enabling procedure :
  - SuperB Virtual Organization setup
  - Simulation software installation, WN compatibility packages installation
  - Transfer of Background data and test release to sites
  - Test session ad lib
- Coordination work with site contacts was crucial
  - Take advantage of BaBar distributed contacts network
  - Take advantage of Grid resources not in use by LHC tasks

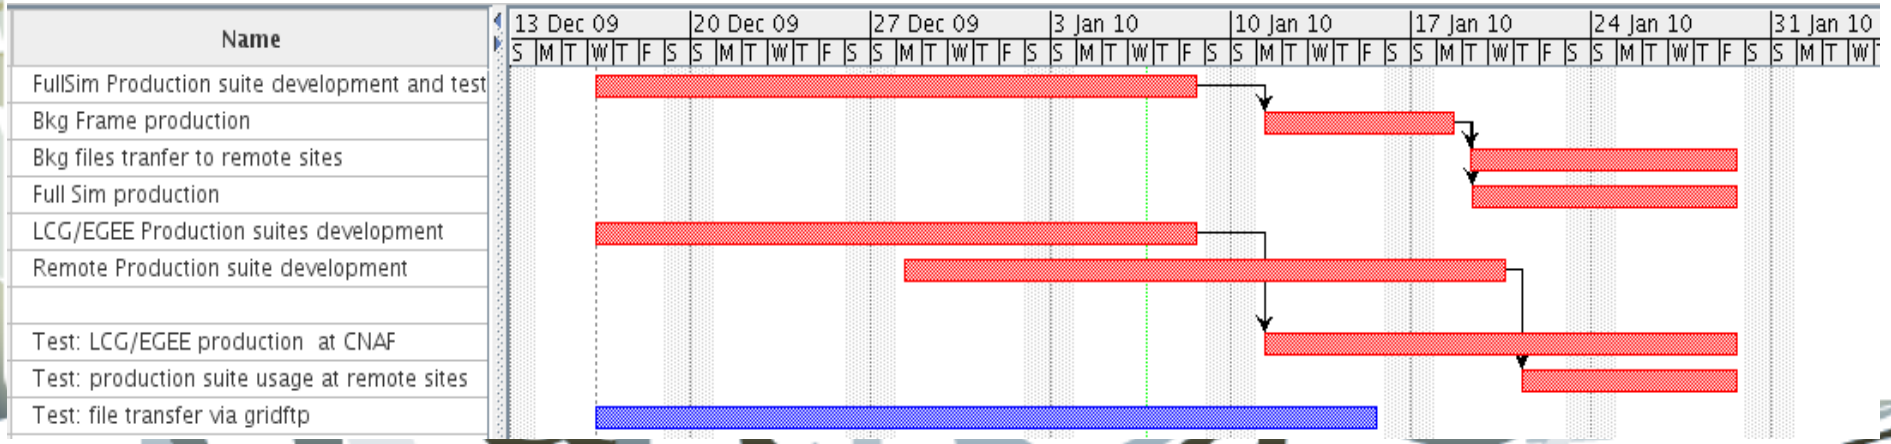

Full Simulation production

- The production has been performed at CNAF
- Job submission directly to LSF batch system
- Total amount of job submitted and fail rate: 9643 job, 0.02%
- Bruno release V00-01-12

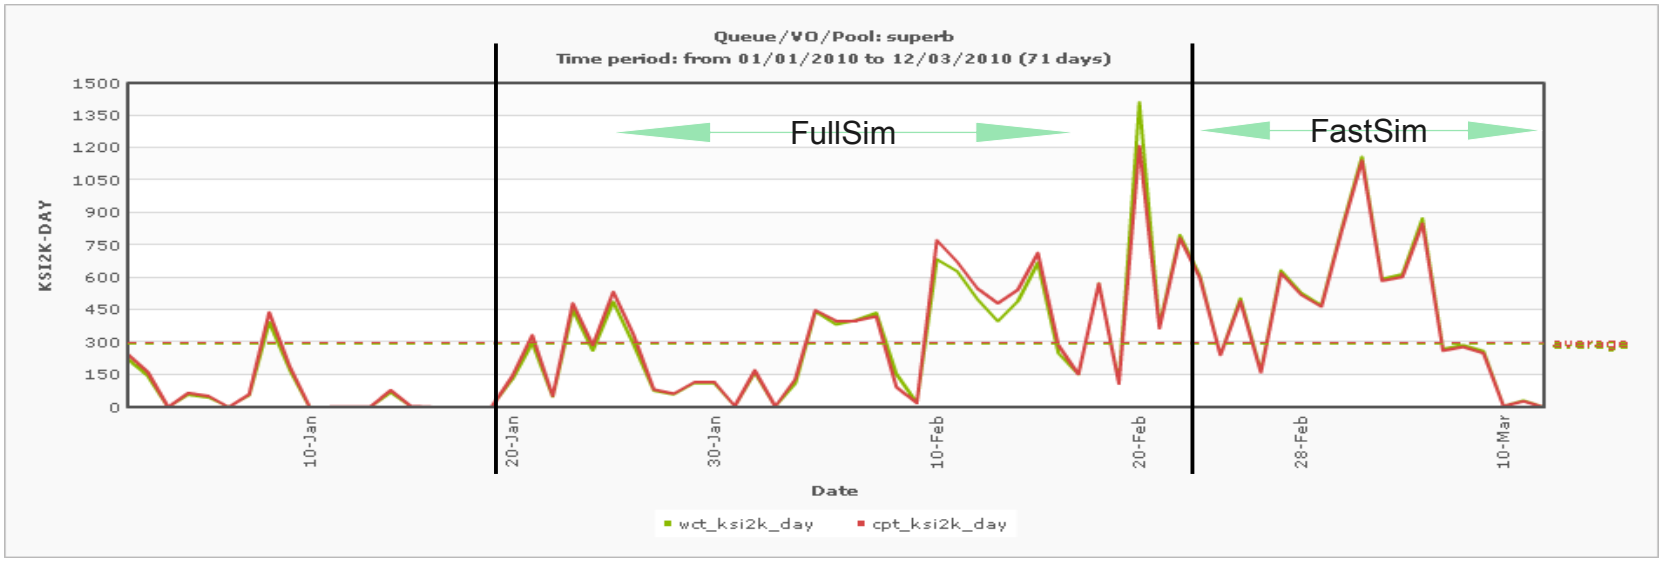

Full Simulation production: summary

- Background frames for Fast Simulation: **10^6** events
- Background studies: a total of **8x10^5** events:
  - **2x10^5** events of RadBhaBha (minDeltaE = 0.1) SuperB_Wolf_shielded, QGSP_BERT
  - **2x10^5** events of RadBhaBha (minDeltaE = 0.1) SuperB_unshielded, QGSP_BERT
  - **2x10^5** events of RadBhaBha (minDeltaE = 0.05) SuperB_Wolf_shielded, QGSP_BERT
  - **10^5** events of RadBhaBha (minDeltaE = 0.002) SuperB_Wolf_shielded, QGSP_BERT
  - **10^5** events of RadBhaBha (minDeltaE = 0.05) SuperB_Wolf_shielded, QGSP_BERT_HP

Full Sim details available here:

http://mailman.fe.infn.it/superbwiki/index.php/How_to_Grid/Site_setup

Prod series: 2010_02_Full_HP

geometry_type	generator_id	simtype	physics_list	output_type	status	nrun	nev	wct_avg	size_avg	size_tot
SuperB_Wolf_shielded	**RadBhaBha** *BRUNOBBBMINDE = 0.05*	fullsim	QGSP_BERT_HP	hits	**done**	**1 041**	**104 100** *(100)*	20 503.66 *(18 236.60 - 54 953.10)*	32.188 mb	32.722 gb
									Total File Size:	**32.722 gb**

Prod series: 2010_02_Full_DeltaE_0.05

geometry_type	generator_id	simtype	physics_list	output_type	status	nrun	nev	wct_avg	size_avg	size_tot
SuperB_Wolf_shielded	**RadBhaBha** *BRUNOBBBMINDE = 0.05*	fullsim	QGSP_BERT		**sys-failed**	**1**	**100** *(100)*	-		
SuperB_Wolf_shielded	**RadBhaBha** *BRUNOBBBMINDE = 0.05*	fullsim	QGSP_BERT	hits	**done**	**1 999**	**199 900** *(100)*	21 663.26 *(18 002.40 - 73 777.32)*	31.940 mb	62.351 gb
									Total File Size:	**62.351 gb**

Prod series: 2010_02_Full_DeltaE_0.02

geometry_type	generator_id	simtype	physics_list	output_type	status	nrun	nev	wct_avg	size_avg	size_tot
SuperB_Wolf_shielded	**RadBhaBha** *BRUNOBBBMINDE = 0.002*	fullsim	QGSP_BERT	hits	**done**	**1 000**	**100 000** *(100)*	43 529.36 *(37 564.47 - 155 846.62)*	66.784 mb	65.218 gb
									Total File Size:	**65.218 gb**

Prod series: 2010_02_Full

geometry_type	generator_id	simtype	physics_list	output_type	status	nrun	nev	wct_avg	size_avg	size_tot
SuperB_Wolf_shielded	**RadBhaBha** *BRUNOBBBMINDE = 0.1*	fullsim	QGSP_BERT	hits	**done**	**800**	**200 000** *(250)*	63 197.49 *(32 609.23 - 196 311.95)*	61.837 mb	48.310 gb
SuperB_unshielded	**RadBhaBha** *BRUNOBBBMINDE = 0.1*	fullsim	QGSP_BERT		**failed**	**2**	**500** *(250)*	-		
SuperB_unshielded	**RadBhaBha** *BRUNOBBBMINDE = 0.1*	fullsim	QGSP_BERT		**sys-failed**	**13**	**3 250** *(250)*	-		
SuperB_unshielded	**RadBhaBha** *BRUNOBBBMINDE = 0.1*	fullsim	QGSP_BERT	hits	**done**	**785**	**196 250** *(250)*	52 703.96 *(26 610.25 - 162 308.48)*	149.884 mb	114.901 gb
									Total File Size:	**163.211 gb**

Fast Simulation production: summary

- Generics production without background mixing: **10^9** events

- Generics production with background mixing: **10^8** events

  - only a fraction of the foreseen sample due to initially very high execution time issue

- BtoTauNu signal with background mixing: **3x10^6** events

- BtoKstarNuNu signal with background mixing: **6x10^6** events

- KplusNuNu signal with background mixing: **6x10^7** events

- BtoKNuNu signal with background mixing: **6x10^6** events

Fast Sim details available here:

http://mailman.fe.infn.it/superbwiki/index.php/February_production_status

done JOBS

Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_generic	MixBaBarBkg_NoPair.tcl	410	10 250 000	6 199 292
DG_BaBar	B0B0bar_generic	MixBaBarBkg_NoPair.tcl	404	10 100 000	6 299 508
DG_BaBar	ccbar	MixBaBarBkg_NoPair.tcl	215	10 750 000	2 357 493
DG_BaBar	uds	MixBaBarBkg_NoPair.tcl	203	19 490 000	2 703 634
DG_3	B+B-_generic	MixSuperbBkg_NoPair.tcl	526	5 260 000	9 798 367
DG_3	B0B0bar_generic	MixSuperbBkg_NoPair.tcl	516	5 160 000	9 817 195
DG_3	ccbar	MixSuperbBkg_NoPair.tcl	257	5 140 000	3 259 682
DG_3	uds	MixSuperbBkg_NoPair.tcl	252	10 080 000	3 782 590
DG_4	B+B-_generic	MixSuperbBkg_NoPair.tcl	520	5 200 000	9 292 856
DG_4	B0B0bar_generic	MixSuperbBkg_NoPair.tcl	623	6 230 000	11 393 686
DG_4	ccbar	MixSuperbBkg_NoPair.tcl	330	6 600 000	3 976 354
DG_4	uds	MixSuperbBkg_NoPair.tcl	252	10 080 000	3 679 825
DG_3	B+B-_generic	PacProduction.tcl	2 092	104 600 000	79 818 019
DG_3	B0B0bar_generic	PacProduction.tcl	2 060	103 000 000	79 992 260
DG_3	ccbar	PacProduction.tcl	1 129	112 900 000	24 731 284
DG_3	uds	PacProduction.tcl	1 556	311 200 000	51 541 216
DG_4	B+B-_generic	PacProduction.tcl	2 085	104 250 000	80 612 976
DG_4	B0B0bar_generic	PacProduction.tcl	2 025	101 250 000	84 969 268
DG_4	ccbar	PacProduction.tcl	1 101	110 100 000	27 916 869
DG_4	uds	PacProduction.tcl	2 624	524 800 000	89 399 621
Total			19 180	1 576 440 000	591 541 996

done JOBS

Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_SL_SL	MixBaBarBkg_NoPair.tcl	110	11 000 000	1 853 471
DG_BaBar	B+B-_taunu	MixBaBarBkg_NoPair.tcl	10	1 000 000	91 496
DG_BaBar	B0B0bar_SL_SL	MixBaBarBkg_NoPair.tcl	110	11 000 000	1 777 269
DG_3	B+B-_SL_SL	MixSuperbBkg_NoPair.tcl	110	11 000 000	5 031 753
DG_3	B+B-_taunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	189 493
DG_3	B0B0bar_SL_SL	MixSuperbBkg_NoPair.tcl	110	11 000 000	4 565 919
DG_4	B+B-_SL_SL	MixSuperbBkg_NoPair.tcl	108	10 800 000	6 070 554
DG_4	B+B-_taunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	190 041
DG_4	B0B0bar_SL_SL	MixSuperbBkg_NoPair.tcl	110	11 000 000	5 703 340
Total			688	68 800 000	25 473 335

done, running, submitted and prepared JOBS

Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_taunu_DX	MixBaBarBkg_NoPair.tcl	10	1 000 000	72 943
DG_3	B+B-_taunu_DX	MixSuperbBkg_NoPair.tcl	10	1 000 000	177 110
DG_4	B+B-_taunu_DX	MixSuperbBkg_NoPair.tcl	10	1 000 000	176 232
Total			30	3 000 000	426 286

done JOBS

Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_K+nunu	MixBaBarBkg_NoPair.tcl	10	1 000 000	90 304
DG_BaBar	B0B0bar_K0nunu	MixBaBarBkg_NoPair.tcl	10	1 000 000	101 452
DG_3	B+B-_K+nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	228 515
DG_3	B0B0bar_K0nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	241 725
DG_4	B+B-_K+nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	225 231
DG_4	B0B0bar_K0nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	240 036
Total			60	6 000 000	1 127 264

done JOBS

Geometry	Generator	tcl	Total Number of Jobs	Total Number of Events	Total CPU time, wct (s)
DG_BaBar	B+B-_Kstar+nunu	MixBaBarBkg_NoPair.tcl	10	1 000 000	105 231
DG_BaBar	B0B0bar_KstarOnunu_Kpi	MixBaBarBkg_NoPair.tcl	10	1 000 000	115 681
DG_3	B+B-_Kstar+nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	255 578
DG_3	B0B0bar_KstarOnunu_Kpi	MixSuperbBkg_NoPair.tcl	10	1 000 000	299 984
DG_4	B+B-_Kstar+nunu	MixSuperbBkg_NoPair.tcl	10	1 000 000	251 195
DG_4	B0B0bar_KstarOnunu_Kpi	MixSuperbBkg_NoPair.tcl	10	1 000 000	299 891
Total			60	6 000 000	1 327 559

Production system and operation credits

- INFN-Ferrara
  - Eleonora Luppi
  - Luca Tomassetti
  - Marco Ronzano
  - Giovanni Fontana
- INFN-CNAF
  - A. F.

Distributed production: Fast Sim results

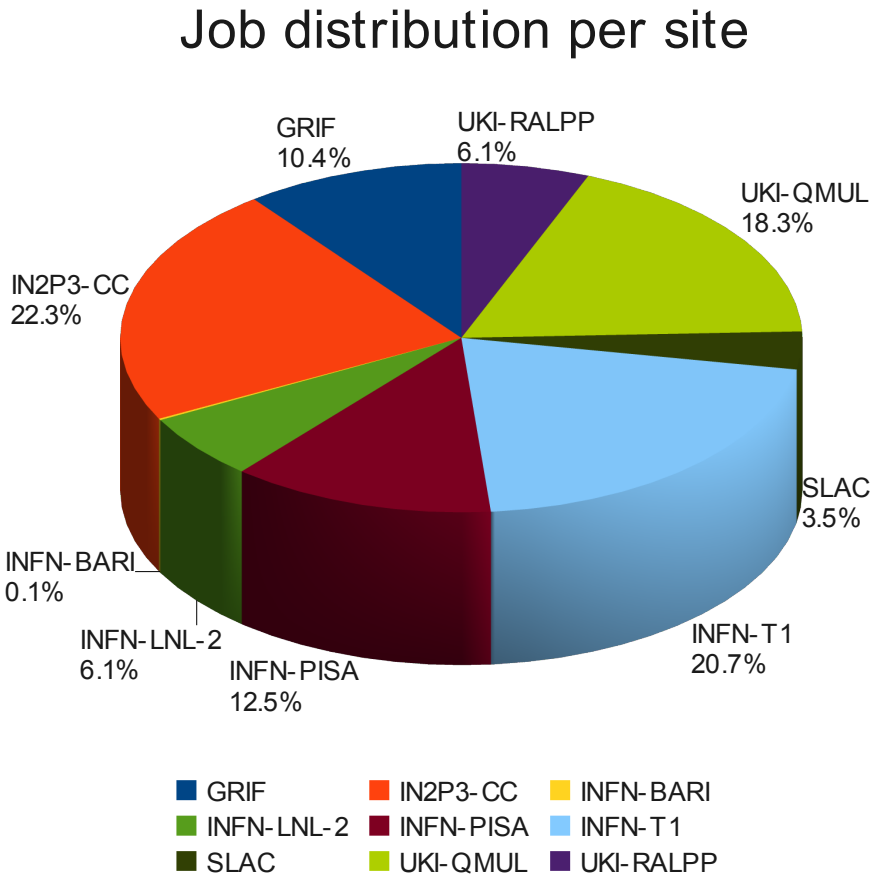

Site Name	Jobs done	Events generated
GRIF	2081	183650000
IN2P3-CC	4457	383930000
INFN-BARI	29	2850000
INFN-LNL-2	1214	120070000
INFN-PISA	2505	183310000
INFN-T1	4143	284660000
SLAC	699	69900000
UKI-QMUL	3672	337745000
UKI-RALPP	1217	94025000
TOTAL	**20017**	**~ 1.7x10^9**

Distributed Computing human network

The success of the distributed production efforts was due to the enthusiastic contributions of many people:

CNAF	Armando Fella
Caltech	Frank Porter, Piti Ongmongkolkul
SLAC	Steffen Luiz, Wei Yang
McGill	Steven Robertson
UVIC	Ashok Agarwal
Queen Mary	Adrian Bevan, Christopher Wilson
RAL T1, T2	Fergus Wilson, Chris Brew
CCIN2P3, GRIF	Nicolas Arnaud
INFN-Bari	Giacinto Donvito, Vincenzo Spinoso
INFN-LNL	Gaetano Maron, Alberto Crescente, Sergio Fantinelli
INFN-Napoli	Silvio Pardi, Alessandra Doria
INFN-Ferrara	Luca Tomassetti, Eleonora Luppi, Giovanni Fontana, Marco Ronzano
INFN-Pisa	Alberto Ciampa, Enrico Mazzoni, Dario Fabiani

Thanks/Congratulation to everybody

How the data can be accessed

- Users can access produced data into SuperB storage central repository at CNAF
- A simple Command Line Interface tool has been provided permitting the creation of data file list
  - Parametric data selection per Geom., Generator etc.
- Future plan includes the development of web based tool permitting data monitor and selection
- See L. Tomassetti talk in Parallel Comp. session Wed. 16-17:30

Data analysis and validation

FullSim and FastSim data analysis and validation is started

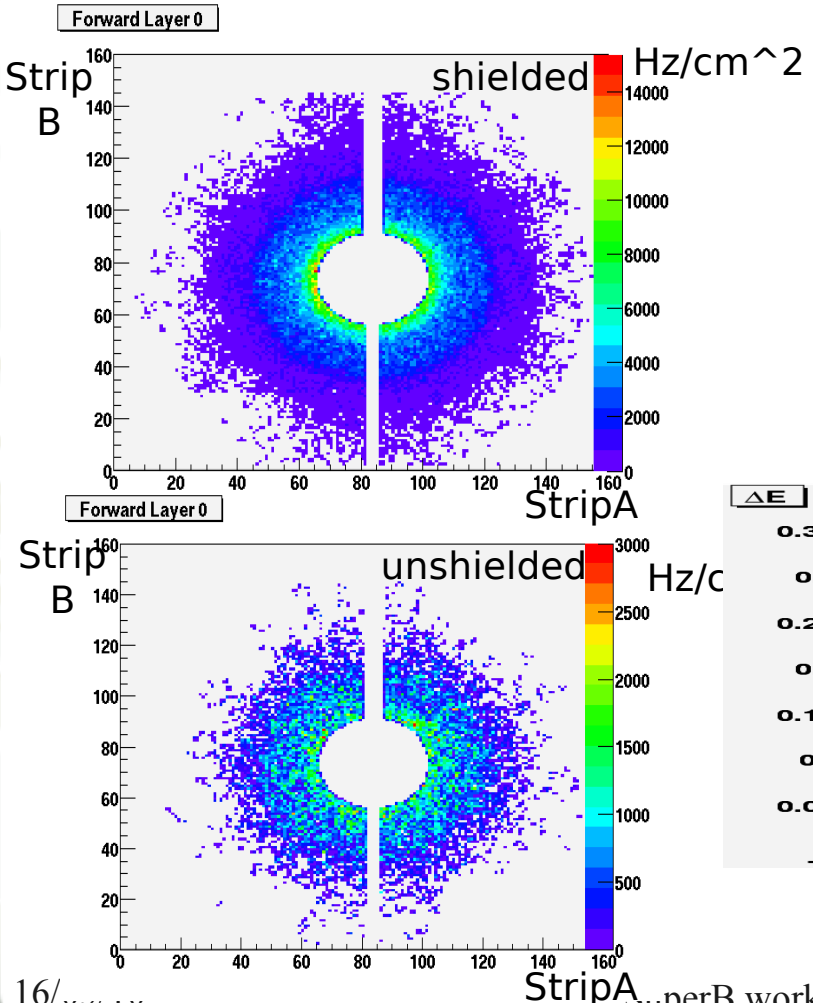

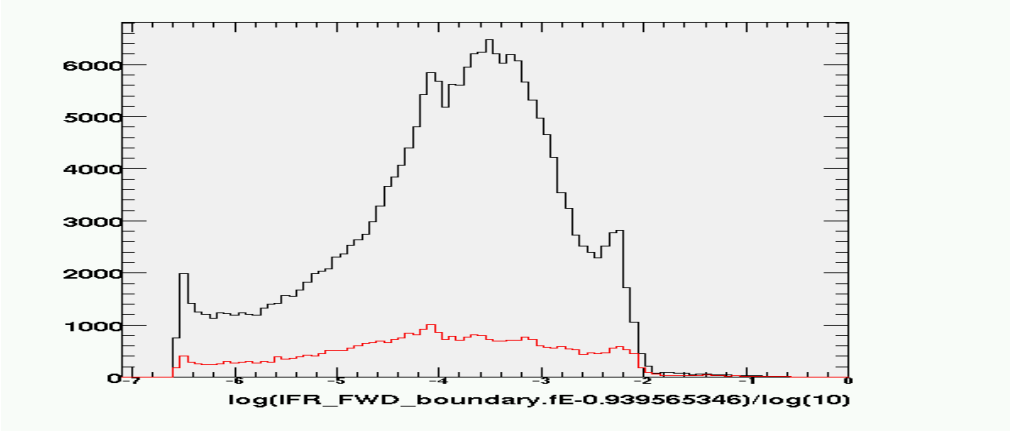

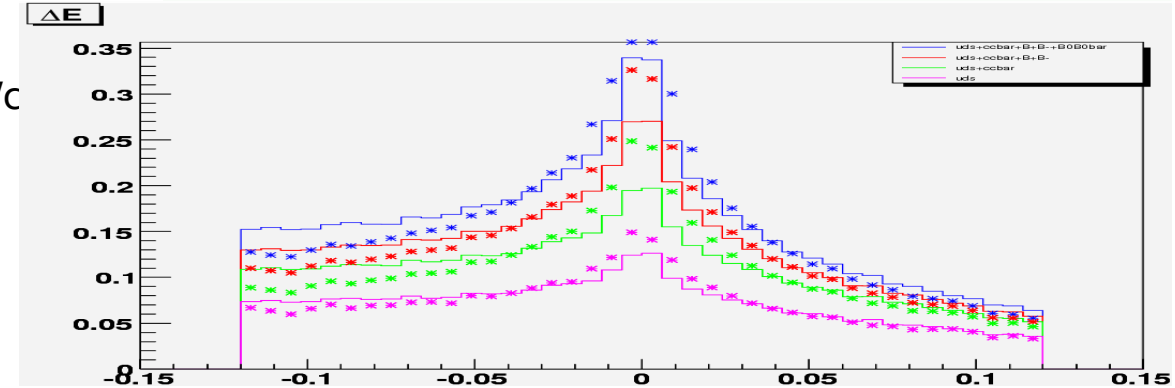

Contributions by E. Manoni and M. Munerato

SuperB R&D Computing WS

- Held in Ferrara, 9-12 March 2009, IUSS Univ. center
- Participation was larger than expected: nearly 50 registered participants from various countries:
  - several experts from CERN, SLAC, DESY, …
    - a few remote presentations
- The workshop succeeded in terms of interest shown by participations and achievement of the original goals
- All talks have been streamed online

Goals and outcomes

- Goals:
  - identify the key aspects we should work on in the R&D phase:
    - the "known known": how shall we attack the main problems we are already aware of (i.e.: BaBar code legacy) ?
    - the "known unknown": what are the key issues we should explore ? the ones that might have the largest impact on the design of our computing model ?
  - discuss what are the R&D activities that can be carried out in a time-scale of 9 to 15 months
- Outcome (to be finalized in a few weeks)
  - document describing the SuperB R&D plan, specifying: motivations, deliverables, timeline, resources needed

Workshop topics

- All the key topics for the SuperB computing model have been covered:
  - Impact of new CPU architectures, software architectures and frameworks
  - Code development: languages, tools, standards and QA
  - Persistence, data handling models and databases
  - Distributed Computing
  - User tools and interfaces
  - Performance and efficiency of large data storage

Where are WE?

- HEP code does not exploit the power of current processors
  - » One instruction per cycle at best
  - » Little or no use of vector units (SIMD)
  - » Poor code locality
  - » Abuse of the heap
- Running N jobs on N=8 cores still efficient but:
  - » Memory (and to less extent cpu cycles) wasted in non sharing
    - "static" condition and geometry data
    - I/O buffers
    - Network and disk resources
  - » Caches (memory on CPU chip) wasted and trashed
    - L1 cache local per core, L2 and L3 shared
    - Not locality of code and data (thread/core affinity)
- This situation is already bad today, will become only worse in future many-cores architectures

SuperB R&D Computing WS

- working groups have been formed, the outcome of the discussions, i.e. the R&D baseline proposals, were reported in detail in the closeout session
  - See: http://agenda.infn.it/conferenceDisplay.py?confId=2241
- What next ?
  - draft of working groups proposed R&D activities will be finalized in two weeks
    - with indication of expertise and man power requirements
  - definition of priorities and timescale will follow
  - a planning document will eventually be released to be discussed and agreed upon

Main goals for this meeting

- discuss the follow up of the R&D workshop and status of white papers
  - Comp. Tue. 2pm session
- discuss how to facilitate the distribution of simulation data for subsequent user analysis on remote sites
  - Comp. Wed. 16 - 17:30am session
- define the future production goals and timeline
  - Comp+Det+Phys. Thu. 2 pm session
- discuss the experience an the perspective of distributed production with the site contacts
  - See: parallel Comp. Session on Thu 16:00 - 17:30

BACKUP

Distributed resources

CPU				Disk space for data		
				(2 GB per day-core)		9
	minimum	maximum	"average" (2*min+max)/3		available	needed
	(cores)	*(cores)*	*(cores)*		*(TB)*	*(TB)*
CNAF	**250**	**750**	417	CNAF	**50**	**8**
SLAC	**200**	**1000**	467	SLAC	**few*10**	**8**
UVIC	**100**	**300**	167	UVIC	**2**	**1**
CALTEC	**100**	**300**	167	CALTEC	**20**	**3**
McGill	**60**	**60**	60	McGill	**-**	**1**
RAL T2	**50**	**100**		RAL T2	**10**	**7**
RAL T1	**300**	**600**	400	RAL T1	**10**	**7**
CCIN2P3	***300***	***600***	400	CCIN2P3	***0.5***	**7**
GRIF	**50**	**100**	400	GRIF	5	**5**
QMUL	**140**	**350**	400	QMUL	20	2
Pisa	**50**	**200**	100	Pisa	**1.5 / 2.5**	**2**
Legnaro-Pd	**50**	**100**	67	Legnaro-Pd	**2**	**1**
Napoli	**50**	**100**	67	Napoli	**-**	**1**
Bari	**50**	**100**	67	Bari	**1-2**	**1**
Ferrara	**10**	**50**	23	Ferrara	**-**	**0**
Total	1750	4660	3176.67	Total	119.5	
	may need to be updated				**may need to be updated**	

CNAF status

- CPU resources have been accessed via Virtual Machine worker node technology
- The production central repository have been accessed via GPFS and SRM (StoRM) systems
- SuperB Bookkeeping DB accessed by jobs via RESTfull interface
- Grid job submission manager via GANGA from User Interface
- The production relied on the following Grid services:
  - Logical File Catalog (LFC)
  - Authentication and Authorization service
    Virtual Organization Manager System (VOMS)
  - Job brokering and monitoring, Workload Manager System (WMS)