

SuperB Computing R&D Workshop

Report of Contributions

Contribution ID: 0

Type: **not specified**

Goals of the workshop

Tuesday, 9 March 2010 14:30 (20 minutes)

Presenter: MORANDIN, Mauro (INFN)

Session Classification: Plenary Session: Introduction

Contribution ID: 1

Type: **not specified**

The BaBar computing model and what we are likely to preserve in SuperB

Tuesday, 9 March 2010 14:50 (30 minutes)

Presenter: Dr BROWN, David (Lawrence Berkeley National Lab)

Session Classification: Plenary Session: Introduction

Contribution ID: 2

Type: **not specified**

SuperB and its computing requirements

Tuesday, 9 March 2010 15:20 (30 minutes)

Presenter: BIANCHI, Fabrizio (TO)

Session Classification: Plenary Session: Introduction

Contribution ID: 4

Type: **not specified**

Distributed Computing for HEP - present and future

Thursday, 11 March 2010 14:40 (30 minutes)

In this presentation, the main aspects related to distributed computing that a HEP experiment has to address are discussed. This is done analyzing what the current experiments, mainly at the CERN LHC, are using, either provided by Grid infrastructures or developed by themselves. After a brief introduction on the overall distributed computing architecture, the specific aspects that are treated are: security and access control, data and storage management, workload management (with particular attention to user analysis) and infrastructure management.

Presenter: GRANDI, Claudio (INFN-Bologna)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: 5

Type: **not specified**

Multithread and MPI usage in GRID

Thursday, 11 March 2010 15:30 (20 minutes)

The talk will be focused on the support for parallel programs, mpi and multithread, in EGEE, that is actually the Grid infrastructure of which InfnGrid is member of. I'll review the status of the support of MPI and the usage of the upcoming new syntax through the analysis of a Case Study.

Presenter: R. ALFIERI (Parma University & INFN)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: 6

Type: **not specified**

Introduction and goals

Tuesday, 9 March 2010 17:00 (15 minutes)

Presenter: INNOCENTE, Vincenzo (CERN)

Session Classification: Plenary Session: Impact of new CPU architectures (I)

Contribution ID: 7

Type: **not specified**

Modern CPU architectures: Challenges and opportunities for HEP

Tuesday, 9 March 2010 17:15 (1 hour)

In the last few years Computing has been characterized by the advent of “multicore cpus”. Effective exploitation of this new kind of computing architecture requires the adaptation of legacy software and eventually a shift of the programming paradigms to massive parallel. In this talk we will introduce the reasons that brought to the introduction of “multicore” hardware and the consequences on system and application software. The activities initiated in HEP to adapt current software will be reviewed before presenting the perspective for the future.

Presenter: INNOCENTE, Vincenzo (CERN)

Session Classification: Plenary Session: Impact of new CPU architectures (I)

Contribution ID: 8

Type: **not specified**

Experience and perspective in parallelization of data analysis (including experience from BaBar and future plan)

Tuesday, 9 March 2010 18:15 (30 minutes)

Increasing the dimension of data samples for data analyses and using advanced algorithms for background suppression (such as unbinned maximum likelihood fits) require high CPU performance. Recently, vendors like Intel and AMD have not incremented the performance of single CPU unit as in the past, but they are working on multi-core CPU. Currently we have up to 8 cores implemented on one single chip. Other possibilities are offered by the use of Graphical Process Units (GPUs), which offer hundreds of cores.

This fact represents a possible revolution in the development of new programs. Indeed we can parallelize the code obtaining great benefits from new multi-core architectures. So we have to reformulate some algorithms generally used for data analyses in High Energy Physics (e. g. Minuit for maximum likelihood fits). It is also possible to have faster execution of the code using the Message Passing Interface (MPI) paradigm, spreading the execution of the code over several CPUs in a cluster. These techniques of High Performance Computing (HPC) are well established in other fields, like computational chemistry and astrophysics. In High Energy Physics community there is not such a large use, but in the future it can be an elegant solution in all the cases where the data analyses will get more and more complicated. In particular the huge data samples expected to be collected in the next generations of experiments, such as the SuperB-factories, will represent a real challenge for the data analysis, requiring to rewrite the algorithms in parallel.

In this presentation I will give a description of the data analysis softwares used in the BaBar experiment and, taking experience from their complexity, I will give an overview on the possible scenario for the next generation of SuperB factories.

Presenter: LAZZARO, Alfio (MI)

Session Classification: Plenary Session: Impact of new CPU architectures (I)

Contribution ID: 9

Type: **not specified**

Discussion

Tuesday, 9 March 2010 18:45 (15 minutes)

Session Classification: Plenary Session: Impact of new CPU architectures (I)

Contribution ID: 10

Type: **not specified**

Coding for the GPU: highlights and application examples

Wednesday, 10 March 2010 08:30 (30 minutes)

Currently, technological solutions being adopted by more manufacturers are bringing to CPU architecture with an even more degree of parallelism, evolving from the multi-core era to the “many-core” era. In this scenario hundreds and, in short, thousands of processing cores are contained within the same processor. A so deep change in architectural paradigm compels an equally deep change in algorithms and programming paradigms. In this context, the GPU (Graphics Processing Unit) were created to accelerate typical 2D and 3D graphic processing, characterized by an extreme degree of parallelism, containing hundreds of processing cores. Due to their characteristics, these are now used to perform complex calculations also in more general fields. During the talk an introduction to GPU Computing will be presented, providing a small technological outlook on manycore scenarios and a brief introduction on architectures and programming languages such as CUDA and OpenCL. We report performance information obtained with algorithms developed in Perugia in the context of the European project Einstein Telescope (ET) and the INFN project macgate (Manycore Computing for Future Gravitational Observatories) and with others of more common interest, such as Monte Carlo methods, Random Number Generation and Simulations.

Presenter: BOSI, Leone Battista (PG)

Session Classification: Plenary Session: Impact of new CPU architectures (II)

Contribution ID: 11

Type: **not specified**

Exploiting concurrency to speed-up ROOT data analysis: caching, PROOF and multi-threading

Wednesday, 10 March 2010 09:00 (30 minutes)

Concurrency aims to improve computing performance by executing a set of computations simultaneously,

possibly in parallel. Since the advent of today's many-core machines the full exploitation of the available CPU power has been one of the main challenges for high-performance computing software

projects, including the HEP ones. However, in HEP data analysis the bottleneck is not (only) CPU but also the I/O hardware.

In this talk I will discuss what is done in ROOT improve the I/O performance exploiting the technological advancements. This starts by optimizing the way information is stored in files to improve the readout performance using cache techniques. Then a multi-process approach $\sqrt{}$ la PROOF

allows to exploit the full I/O bandwidth on multi-core machines; more generally, the PROOF technology can effectively increase the I/O capabilities of LAN (and perhaps WAN) setups.

Finally the use of multiple threads may help in reducing the time needed to get the data available in the application: I will present the current status of the parallel unzipping implementation, and discuss possible improvements and applications to the write case.

Presenter: GERRY GANIS

Session Classification: Plenary Session: Impact of new CPU architectures (II)

Contribution ID: 12

Type: **not specified**

Discussion

Wednesday, 10 March 2010 09:50 (40 minutes)

Session Classification: Plenary Session: Impact of new CPU architectures (II)

Contribution ID: 13

Type: **not specified**

Introduction and goals

Wednesday, 10 March 2010 11:00 (15 minutes)

Presenter: ELMER, Peter (Princeton University)

Session Classification: Plenary Session: Software architecture and framework

Contribution ID: 14

Type: **not specified**

Framework evolution: status and outlook

Wednesday, 10 March 2010 11:15 (30 minutes)

Object-Oriented event data processing software frameworks have been developed by all recent HEP experiments to structure and support their data processing programs such as the simulation, reconstruction and data analysis. These frameworks implement a given architectural vision and provide a set of features and functionalities that are very close among the various experiments. This presentation will try to identify which are these common architectural elements and the functionalities of an ideal framework based on the experience of the frameworks in use at the LHC experiments.

Presenter: Dr MATO, Pere (CERN)

Session Classification: Plenary Session: Software architecture and framework

Contribution ID: 15

Type: **not specified**

Discussion

Wednesday, 10 March 2010 11:45 (30 minutes)

Session Classification: Plenary Session: Software architecture and framework

Contribution ID: 16

Type: **not specified**

Persistency in HEP applications: current models and future outlook

Wednesday, 10 March 2010 14:45 (30 minutes)

Presenter: CALAFIURA, Paolo (Lawrence Berkeley National Lab)

Session Classification: Plenary Session: Persistence and data handling models

Contribution ID: 17

Type: **not specified**

The SciDB approach and its applicability to HEP computing

Wednesday, 10 March 2010 15:15 (40 minutes)

SciDB is a new open source database management system that emerged from the Extremely Large Database (XLDB) workshop series. Specifically designed to meet the needs of many scientific disciplines, it features an array-based data model that adds order and adjacency to the traditional relational set model. SciDB's array model and the SciDB software in its currently-projected form is *not* appropriate for storage of the largest HEP data type (event data), as I understand it.

So why discuss it at this workshop? Aside from its potential uses for more structured, though smaller, data sets, I believe there are several lessons to be learned from the design and implementation of SciDB that are generally applicable to large-scale scientific data management and analysis.

Presenter: K.-T. LIM

Session Classification: Plenary Session: Persistence and data handling models

Contribution ID: **18**

Type: **not specified**

Introduction, goals and an initial question

Thursday, 11 March 2010 08:30 (15 minutes)

Presenter: GAPONENKO, I (SLAC)

Session Classification: Plenary Session: Databases

Contribution ID: 20

Type: **not specified**

Critical aspects of user interaction with the analysis environment

Thursday, 11 March 2010 17:40 (30 minutes)

Using experience gained from the analysis environment in LHCb, BaBar and ATLAS, I will comment on the critical aspects that can make an analysis environment effective. I will do this by creating a set of requirements and then look at what lessons can be learned from past and current analysis environments with respects to these requirements. Some particular issues to consider are:

- The unwanted separation between users and developers of analysis tools
- Integration of user datasets into databases
- Duplication of data (Ntuples that are straight copies of DSTs)
- Software training
- Analysis in a distributed environment

Presenter: EGEDE, Ulrik (Imperial College (London))

Session Classification: Plenary Session: User tools and interfaces

Contribution ID: 23

Type: **not specified**

Databases in BABAR

Thursday, 11 March 2010 08:45 (30 minutes)

BABAR has been a unique experiment in a history of the High Energy Physics and not only due to the Physics itself, but also in a way this physics was extracted. It pioneered in involving hundreds of physicists, engineers and students into a practical C++ programming using OOAD methodologies. The experiment introduced a geographically distributed data processing and events simulation production system. Finally, it was the first (and so far - the only) experiment in a history of HEP which relied upon a commercial Object-Oriented Database to store detector events. At some point, almost every single bit of data, including detector events, were stored in databases. From the very beginning of the Experiment, databases were an integral part of the BABAR Computing Model. Over a course of the 9 years of data taking, our understanding on what are the right applications for databases and how they should be used had dramatically evolved. The experiment is now over, but its legacy still does matter. This talk is about our practical experience with various database technologies, a genesis of our understanding of how to design, implement and maintain a very complex system depending in its foundation on databases. It's about our successes and failures. Certain software artifacts of the experiment may also be found directly usable for the future Super-B.

Presenter: I. GAPONENKO**Session Classification:** Plenary Session: Databases

Contribution ID: 25

Type: **not specified**

Discussion

Thursday, 11 March 2010 10:00 (45 minutes)

Session Classification: Plenary Session: Databases

Contribution ID: 26

Type: **not specified**

Introduction and Goals

Friday, 12 March 2010 08:30 (10 minutes)

Presenters: Dr FURANO, Fabrizio (CERN); VAGNONI, Vincenzo Maria (BO)

Session Classification: Plenary Session: Performance and efficiency of large data storage

Contribution ID: 27

Type: **not specified**

Research, Development and Scientific Application of Gfarm File System

Friday, 12 March 2010 09:10 (25 minutes)

Gfarm File System is a wide-area distributed file system that federates local disks of compute nodes in a Grid or computer clusters. It is a high performance distributed parallel file system designed for I/O intensive scientific data analysis conducted in collaboration with multiple distant organizations. Gfarm is under vigorous development for better performance and usability. A parallel distributed workflow system is developed to make the most use of Gfarm features such as file affinity scheduling. A case study of large-scale data analysis in astronomy using Gfarm file system shows a good scalable performance.

In this talk we review recent researches and development of Gfarm.

Presenter: MASAHIRO TANAKA

Session Classification: Plenary Session: Performance and efficiency of large data storage

Contribution ID: 28

Type: **not specified**

Grid/Cloud computing evolution

Thursday, 11 March 2010 15:50 (20 minutes)

The presentation will first describe what the main current evolution trends for distributed computing seem to be, moving then on to explore how computing resources could be uniformly accessed via either grid or cloud interfaces using virtualization technologies. The integration of grids and clouds can on the one hand expand and optimize the use of available computing resources, while on the other hand allowing new scenarios for the solution of distributed computing tasks.

Presenter: SALOMONI, DAVIDE (INFN - CNAF)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: 29

Type: **not specified**

Middleware development in the EGI era

Thursday, 11 March 2010 15:10 (20 minutes)

After the end of the EGEE series of projects, the way improvements to the middleware are developed and deployed will change significantly, with a stronger focus on stability of the infrastructure and on quality assurance.

The middleware will evolve according to the requirements coming from users in terms of reliability, usability, functionality, interoperability, security, management, monitoring and accounting, also exploiting emerging developments like virtualization.

The presentation will briefly introduce the organizational framework for the middleware management in the EGI era, followed by an overview of the major expected developments.

Presenter: F. GIACOMINI (INFN - CNAF)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: **30**

Type: **not specified**

Discussion

Thursday, 11 March 2010 16:10 (35 minutes)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: 31

Type: **not specified**

Simulating storage system performance: a useful approach for SuperB ?

Friday, 12 March 2010 10:05 (25 minutes)

Capacity planning is a very useful tool to estimate future resource demand to carry on a given activity. Using well known modelling techniques it is possible to study the performance of a system before actually building it, and evaluate different design alternatives. However, capacity planning must be done properly in order to be effective. This presentation describes the benefits, requirements and pitfalls of performance modelling with a particular emphasis on storage systems.

Presenter: Dr MARZOLLA, Moreno (Bologna University)

Session Classification: Plenary Session: Performance and efficiency of large data storage

Contribution ID: 32

Type: **not specified**

Visualizations and interfaces for end users

Thursday, 11 March 2010 18:10 (30 minutes)

- A discussion of 4vector viewers as a visualization tool residing between the event display and final histograms. This has been useful for educating new students/analysts as well as an outreach tool.
- Viewpoints: a NASA developed multi-variate display package used to develop more of an intuition for multi-variate datasets. This has been used with great success with undergraduate students but should not be limited to that level of experience.
- A few comments about the revival of the 30-yo LASS data and what we might learn from that experience with respect to data format design.

Presenter: MAT BELLIS

Session Classification: Plenary Session: User tools and interfaces

Contribution ID: 33

Type: **not specified**

Introduction and goals

Wednesday, 10 March 2010 14:30 (15 minutes)

Presenter: Dr BROWN, David (Lawrence Berkeley National Lab)

Session Classification: Plenary Session: Persistence and data handling models

Contribution ID: 34

Type: **not specified**

Introduction and goals

Thursday, 11 March 2010 14:30 (10 minutes)

Presenters: FELLA, Armando (CNAF); LUPPI, Eleonora (Ferrara University & INFN)

Session Classification: Plenary Session - Distributed Computing

Contribution ID: 35

Type: **not specified**

R&D proposal from the sessions "Impact of new CPU architectures, software architectures and frameworks"

Friday, 12 March 2010 14:00 (30 minutes)

Presenters: ELMER, Peter (Princeton Univ.); INNOCENTE, Vincenzo (CERN)

Session Classification: Plenary Session

Contribution ID: **36**

Type: **not specified**

Discussion

Wednesday, 10 March 2010 15:55 (35 minutes)

Session Classification: Plenary Session: Persistence and data handling models

Contribution ID: 38

Type: **not specified**

R&D proposal from the session "Persistence, data handling models and databases"

Friday, 12 March 2010 14:45 (20 minutes)

Presenters: Dr BROWN, David (Lawrence Berkeley National Lab); VANYASHIN, Sasha (CERN)

Session Classification: Plenary Session

Contribution ID: 39

Type: **not specified**

R&D proposal from the session "Code development: languages, tools, standards and QA"

Friday, 12 March 2010 14:30 (15 minutes)

Presenters: DI SIMONE, Andrea (RM2); Prof. STROILLI, Roberto (Universita' di Padova and INFN); Dr LUITZ, Steffen (SLAC)

Session Classification: Plenary Session

Contribution ID: 41

Type: **not specified**

R&D proposal from the session "Distributed Computing"

Friday, 12 March 2010 15:05 (20 minutes)

Presenters: FELLA, Armando (CNAF); LUPPI, Eleonora (Ferrara University & INFN)

Session Classification: Plenary Session

Contribution ID: 42

Type: **not specified**

R&D proposal from the session "User tools and interfaces"

Friday, 12 March 2010 15:25 (20 minutes)

Presenter: BIANCHI, Fabrizio (TO)

Session Classification: Plenary Session

Contribution ID: 43

Type: **not specified**

R&D proposal from the session "Performance and efficiency of large data storage"

Friday, 12 March 2010 15:45 (20 minutes)

Presenters: Dr FURANO, Fabrizio (CERN); VAGNONI, Vincenzo Maria (BO)

Session Classification: Plenary Session

Contribution ID: 44

Type: **not specified**

Discussion

Friday, 12 March 2010 16:05 (25 minutes)

Session Classification: Plenary Session

Contribution ID: 45

Type: **not specified**

Discussion

Thursday, 11 March 2010 18:40 (20 minutes)

Session Classification: Plenary Session: User tools and interfaces

Contribution ID: 46

Type: **not specified**

Introduction and goals

Thursday, 11 March 2010 17:30 (10 minutes)

Presenter: BIANCHI, Fabrizio (TO)

Session Classification: Plenary Session: User tools and interfaces

Contribution ID: 47

Type: **not specified**

Discussion

Wednesday, 10 March 2010 18:35 (25 minutes)

Session Classification: Plenary Session: Code development: languages, tools, standards, QA, interplay with Online

Contribution ID: 49

Type: **not specified**

Databases in LHC experiments: usage and lessons learned

Thursday, 11 March 2010 09:15 (45 minutes)

Database applications and computing models of LHC experiments. Database integration with frameworks and into an overall data acquisition /production/processing and analysis chain of the experiment(s). Best practices in database development cycle: schema, contents, code deployment, etc. Choice of database technologies: RDBMS, hybrid, etc. Examples of database interfaces and tools, such as custom database interfaces and services on top of CORAL/COOL, etc. Scalability of database access in the distributed computing environment of LHC experiments. Distributed database applications: requirements, data complexity, update frequency, data volumes, usage, etc. Database areas in which substantial progresses was made and select remaining problems.

Presenter: VANIACHINE, A (Argonne National Laboratory)

Session Classification: Plenary Session: Databases

Contribution ID: 50

Type: **not specified**

Release building and validation at LHC

Wednesday, 10 March 2010 17:20 (30 minutes)

Presentation will be summary of release building and validation procedures done in LHC experiments at CERN. Brief description how this is done in LHC experiments with main focus on ATLAS.

The complex detectors determine the existing large collaboration with many users and developers providing

software code. Every experiment has its own developed and used tools but there are also many common

utilities and code used and shared between the experiments. ATLAS software comprises around 2000 software

packages, organized in 10 projects that are build on a variety of compiler and operating system combinations every night.

File-level parallelism, package-level parallelism and multi-core build servers are used to perform simultaneous builds of 6 platforms

that are merged into single installation. I will discuss the various tools that provide performance gains and the error detection and retry mechanisms that have been developed in order to counter-act network and other instabilities. Will be discussed test tools and the different levels of testing and validation framework which are implemented in order to give better quality view of the produced software.

Presenter: EMIL OBRESHKOV (CERN)

Session Classification: Plenary Session: Code development: languages, tools, standards, QA, interplay with Online

Contribution ID: 51

Type: **not specified**

Discussion

Tuesday, 9 March 2010 15:50 (40 minutes)

Session Classification: Plenary Session: Introduction

Contribution ID: 52

Type: **not specified**

CMS Online Software, Infrastructure Overview

Wednesday, 10 March 2010 18:05 (30 minutes)

The CMS online software encompasses the whole range of elements involved in the CMS data acquisition function. A central online software framework (XDAQ Cross-Platform DAQ Framework) has matured over a height years development and testing period and has shown to be able to cope well with the CMSrequirements.

The framework relies on industry standard networks and processing equipment. All subsystems of the DAQ have adopted the central online software framework with the philosophy of using a common and standardized technology in order to reduce the effort associated with the maintenance and evolution of the detector read-out system over a long lifetime of the experiment. Adhering to a single software infrastructure in all subsystems of the experiment imposes a number of different requirements.

High efficiency and configuration flexibility are among the most important ones. This presentation gives an overview of the software infrastructure, the architecture and design as well as brief outlook on the approaches and methods used to deal with the most technically relevant requirements.

Presenter: ORSINI, Luciano (CERN)

Session Classification: Plenary Session: Code development: languages, tools, standards, QA, interplay with Online

Contribution ID: 53

Type: **not specified**

Introduction and goals

Wednesday, 10 March 2010 17:00 (20 minutes)

Presenters: DISIMONE, Andrea (RM2); Prof. STROILL, Roberto (Universita' di Padova and INFN); Dr LUITZ, Steffen (SLAC)

Session Classification: Plenary Session: Code development: languages, tools, standards, QA, interplay with Online

Contribution ID: 54

Type: **not specified**

Overview and outlook of existing data access solutions

Friday, 12 March 2010 08:40 (30 minutes)

Presenters: Dr FURANO, Fabrizio (CERN); VAGNONI, Vincenzo Maria (BO)

Session Classification: Plenary Session: Performance and efficiency of large data storage

Contribution ID: 55

Type: **not specified**

Overview of the new technologies and evolution of storage systems for handling large volume of data

Friday, 12 March 2010 09:35 (30 minutes)

A report on some new emerging technologies suitable for managing, with high efficiency, huge amount of data will be given.

In particular details about Lustre, Hadoop and Ceph, as examples of different approaches to tackle the problem of providing input/output data to scientific applications, will be presented.

Presenter: DONVITO, Giacinto (INFN)

Session Classification: Plenary Session: Performance and efficiency of large data storage

Contribution ID: 56

Type: **not specified**

Discussion

Friday, 12 March 2010 11:00 (1 hour)

Session Classification: Plenary Session: Performance and efficiency of large data storage (II)

Contribution ID: 57

Type: **not specified**

Multi-level Parallel fit algorithms using MPI and CUDA

Wednesday, 10 March 2010 09:30 (20 minutes)

Linux clusters consisting of multi-core commodity-chip based nodes, augmented with GPGPU accelerators are becoming common at computing centers. In this talk we report on some early results of a study to investigate use of this computing paradigm to accelerate the fitting algorithms used in MINUIT. In particular we show that very good speedups are possible for the negative log likelihood (NLL) fit of a simple Gaussian. We also discuss a preliminary implementation of Breit-Wigner convoluted with a Gaussian and report on the challenges encountered implementing the complex error function of this algorithm on the GPU.

Presenter: TOMKO, Karen

Session Classification: Plenary Session: Impact of new CPU architectures (II)

Contribution ID: 58

Type: **not specified**

BaBar Online-Offline - Some Thoughts

Wednesday, 10 March 2010 17:50 (15 minutes)

A short overview of the BaBar Online/Offline interplay, lessons learned and some candidate topics for SuperB R&D

Presenter: Dr LUITZ, Steffen (SLAC)

Session Classification: Plenary Session: Code development: languages, tools, standards, QA, interplay with Online