

X Seminar on Software for Nuclear, Subnuclear and Applied Physics



Report of Contributions

Contribution ID: 0

Type: **not specified**

Computing models for high-energy experiments (ie, how to survive an avalanche of data)

Thursday, 6 June 2013 09:00 (1h 30m)

Presenter: BOCCALI, Tommaso (PI)

Contribution ID: 1

Type: **not specified**

Quantitative methods for neuroimaging data analysis

Thursday, 6 June 2013 11:00 (1h 30m)

Presenter: RETICO, Alessandra (PI)

Contribution ID: 2

Type: **not specified**

Computing models for high-energy experiments (ie, how to survive an avalanche of data)

Friday, 7 June 2013 09:00 (1h 30m)

Presenter: BOCCALI, Tommaso (PI)

Contribution ID: 3

Type: **not specified**

Multi- and many-core computing for Physics applications

Monday, 3 June 2013 11:00 (1h 30m)

More and more often, processor manufacturers adopt the multi-core design approach as a way to further improve performances in spite of the fact that current micro-electronic technologies put a practical upper limit on clock frequency at approximately 3 GHz. A multi-core processor is a single chip integrating two or more independent CPUs. The number of cores within one chip is quickly growing: processors with 100 or more cores are expected in the near future. The many-core approach allows processors to scale according to Moore's law, but it bears a great impact on application design, further moving the challenge of sustaining performance from hardware to algorithms and software.

In this lecture we focus on architecture and programming aspects of recent developed many- and multi-core processors, analyzing the impact of their use on physics applications (mostly in theoretical physics but considering also experimental physics applications). We take into account multi-core processors based on "traditional" core-architecture such as the Sandybridge, as well as many-core systems based on GP-GPUs and on the most recent Intel Xeon-Phi. We also analyze programming strategies to exploit high performance computing using as test-bed real case physics applications.

Presenter: SCHIFANO, Sebastiano Fabio (FE)

Contribution ID: 4

Type: **not specified**

Data acquisition, slow control systems and image processing for Physics applications

Presenter: TOMASSETTI, Luca (FE)

Contribution ID: 5

Type: **not specified**

Tomographic image reconstruction: theory and applications to photon and proton tomography

Tuesday, 4 June 2013 11:00 (1h 30m)

Presenter: VANZI, Eleonora

Contribution ID: 6

Type: **not specified**

Knowledge and Technology Transfer @ INFN

Monday, 3 June 2013 09:30 (1 hour)

Presenter: FALCIANO, Speranza (ROMA1)

Contribution ID: 7

Type: **not specified**

From GPU-accelerated computing to GPU-accelerated data acquisition for physics experiments; the QUonG cluster, the APEnet+ network card and the APE project evolution

Friday, 7 June 2013 11:00 (1h 30m)

Graphical Processing Units have become established as reasonably cheap but very powerful numerical accelerators; they are employed more and more in modern clusters for scientific computing. On the other hand, the fat-tree topology that most of them employs for their high performance network infrastructure (like InfiniBand) has a number of shortcomings that become more and more severe when scaling up in node number and all the more so when nodes are equipped with GPUs. To mitigate the scaling issues, the APE group - within the framework of the European FP7 project EURETILE - is instead pushing an FPGA-based, PCI-Express Gen2 network card of its own design aimed at standard x86_64 servers, the APEnet+ board; APEnet+ not only leverages onto a 3-dimensional toroidal mesh topology - the same that APE parallel machines exploit since their inception in the 80's - but also on a novel, first-of-its-kind implementation of a Remote Direct Memory Access protocol towards the GPU memory.

With APEnet+, we built the QUonG (QCD-on-GPUs) cluster in Rome, a GPU-accelerated multi-core Xeon cluster dedicated to High Performance Computing.

Moreover, thanks to its low-jitter, high-throughput, direct-to-GPU-memory data injection capability, a version of the board called NaNet was developed, to use as a low latency interface between the readout boards and the GPUs where event detection is performed within the trigger system of the NA62 experiment at CERN.

We present a description of APEnet+, its design choices, its development history and a number of results obtained during investigations performed on QUonG and NaNet.

Presenter: SIMULA, Francesco (ROMA1)

Contribution ID: 8

Type: **not specified**

e-Infrastructures

Wednesday, 5 June 2013 11:00 (1h 30m)

Presenter: PAGANONI, Marco (MIB)

Contribution ID: 9

Type: **not specified**

Databases in experimental Physics: theory and practice

Wednesday, 5 June 2013 09:00 (1h 30m)

Primary author: TOMASSETTI, Luca (FE)

Presenter: TOMASSETTI, Luca (FE)

Contribution ID: **11**

Type: **not specified**

Geant4 simulation code: theory and practical session

Monday, 3 June 2013 16:00 (1h 30m)

Presenters: VARISANO, Annagrazia (LNS); ROMANO, Francesco (LNS)

Contribution ID: 12

Type: **not specified**

Data acquisition, slow control systems and image processing for physics applications

Tuesday, 4 June 2013 09:00 (1h 30m)

Presenter: TOMMASETTI, Luca (University & INFN Ferrara)

Contribution ID: 13

Type: **not specified**

Geant4 simulation code: theory and practical session

Wednesday, 5 June 2013 15:30 (1h 30m)

Presenters: VARISANO, Annagrazia (LNS); ROMANO, Francesco (LNS)

Session Classification: Geant4 simulation code: theory and practical session