

**3rd ASPERA Computing and  
Astroparticle Physics  
Workshop 3-4 MAY 2012,  
HANNOVER, GERMANY**



**Report of Contributions**

Contribution ID: 1

Type: **not specified**

## **AMD Accelerator technologies - CPU, GPU and APU architectures to enable scientific computing in the future. Making accelerators accessible through the Heterogenous System Architecture and Open standards.**

*Thursday, 3 May 2012 11:45 (45 minutes)*

AMD is developing processor technology in three formats - CPU, GPU, and more recently, APU. And across these areas, AMD is leading a drive toward Heterogeneous Systems Architecture (HSA), an open platform standard that takes advantage of CPU/GPU processor cores as a unified processing engine, which we are making into an open platform standard. This architecture enables many benefits for HPC, including greatly enhanced application performance and significantly low power consumption.

**Presenter:** NORDLUND, Leif

Contribution ID: 8

Type: **not specified**

## GPUs in gravitational wave data analysis

*Thursday, 3 May 2012 16:40 (40 minutes)*

Searches for gravitational-wave signals from inspiraling black hole or neutron star binaries push the limits of currently available computing resources with conventional CPU-based computer clusters. Previous efforts have used the advantages associated with GPU-hardware to accelerate computationally-intensive portions of the searches by porting those computations to run on the GPUs. Additionally, future computational savings could be obtained through further code optimization and novel analysis techniques, which will of course be impacted by the technologies available in the coming years. In this presentation, I will summarize the LIGO Scientific Collaboration and the Virgo Collaborations' efforts to accelerate inspiral searches using GPUs and will discuss how these efforts will be focused in the coming years.

**Presenter:** KEPPEL, Drew

Contribution ID: 9

Type: **not specified**

## GPUs in Fermi satellite data analysis

*Thursday, 3 May 2012 16:00 (40 minutes)*

The standard analysis of the Fermi LAT collaboration could be sped up by two orders of magnitude porting the most time-consuming Science Tools toward a GPU architecture. Using an NVIDIA S2050, with its Fermi architecture, we were able to accelerate the computation of the satellite “live-time cube”, reducing the execution time from 70 minutes (CPU) to 30 seconds (GPU). Other analysis tools could benefit from GPUs: in particular, the likelihood analysis and the upper limit computation. In this talk, we will present Uriel, the Ultrafast Robotic Interface for Extended Likelihood, and many different applications where GPUs can have an impact in gamma-ray astrophysics.

**Presenter:** Dr BASTIERI, Denis (Università di Padova)

Contribution ID: 12

Type: **not specified**

## Frontiers of Volunteer Computing

*Friday, 4 May 2012 09:00 (45 minutes)*

Ten years from now, as today, the majority of the world's computing and storage resources will reside not in machine rooms but in the hands of consumers. Through volunteer computing much of these resources will be available to science. The first PetaFLOPS computation was done using volunteered computers, and the same is likely to be true for the ExaFLOPS milestone. Volunteer computing has existed for a decade and is being used to do breakthrough science in areas ranging from molecular biology to radio astronomy; however, it is an emerging technology and has potential applications in many new areas, including those involving storage and processing of large data. The landscape of volunteer computing is shaped by many factors. Some of these involve hardware technology; mobile devices, Graphics Processing Units (GPUs), wired and wireless communication networks, memory, and storage. I will discuss trends in these areas. Other factors involve software: technologies like virtualization are making it easier for scientists to use volunteer computing, while the rise of proprietary software environments and vendor-controlled application markets is making it more difficult. Finally, I will discuss the organizational, economic, and marketing issues that must be addressed for volunteer computing to achieve its potential.

**Presenter:** Prof. ANDERSON, David (UC Berkeley)

Contribution ID: 13

Type: **not specified**

## Technology and Market Trend in Computing

*Friday, 4 May 2012 09:45 (45 minutes)*

For the past 15 years the CERN IT department has launched regular (every ~2-3 years) technology+market evaluations which are used as input for the computer center architecture and cost/budget planning activities.

The talk will try to give an overview of the various market and technology developments in the area of data processing and data storage. This will cover processors, memory, HDDs, SSDs and some future technologies.

Cost and technology trends for the next 3-5 years will be discussed.

**Presenter:** Dr PANZER-STEINDEL, Bernd

Contribution ID: 16

Type: **not specified**

## GPUs in real-time discovery of millisecond radio transients with LOFAR

*Thursday, 3 May 2012 17:20 (40 minutes)*

I will present a project that uses GPU technology with the next-generation LOFAR radio telescope, to search for bright millisecond bursts of radio emission from astrophysical sources. GPUs provide the computing power necessary to remove in real time the effects of propagation of the radio emission through the ionised interstellar medium. I will present details of the specific problem, our current approach for optimisation of the relevant GPU code, and why GPUs are currently the most appropriate solution compared to other multicore technologies. Finally, I will describe how our current work fits in the context of the Square Kilometer Array and its pathfinders for the science of astrophysical radio transients.

**Presenter:** Dr KARASTERGIOU, Aris

Contribution ID: 18

Type: **not specified**

## Energy Efficient Computing with GPUs

*Thursday, 3 May 2012 09:45 (45 minutes)*

The past five years have seen the use of graphical processing units (GPUs) for computation grow from being of interest to a handful of early adopters to a mainstream technology used in the world's largest supercomputers. One of the attractions of the GPU architecture is the efficiency with which it can perform computations. Energy efficiency is a key concern in the design of all modern computing systems, from the lowest power mobile devices to the largest supercomputers; it will be paramount in the push to exascale computing. We discuss the Echelon project, and how the NVIDIA GPU architecture will evolve over the coming 5-10 years. Echelon is a DARPA-funded project investigating efficient parallel computer architectures for the Exascale era.

**Presenter:** Dr LANFEAR, Timothy

Contribution ID: 19

Type: **not specified**

## Computing in the context of Cherenkov Telescope Array (CTA)

*Thursday, 3 May 2012 14:45 (45 minutes)*

The Cherenkov Telescope Array (CTA) –an array of tens of Cherenkov telescopes deployed on an unprecedented scale –will allow the European scientific community to remain at the forefront of research in the field of very high energy gamma-ray astronomy.

One of the challenges to design the CTA observatory is to handle the large amounts of data generated by the instrument and to provide simple and efficient user access at any level and according to astrophysical standards in order to serve the data and the software for data analysis to the physics community. The high data rate of CTA together with the large computing power requirements for Monte Carlo simulations, fundamental tool for data selection and calibration, demand dedicated computer resources which can be well handled through a DCI approach. Preliminary works and ideas about the organization of a coherent Data Management system for CTA will be presented.

**Presenter:** Dr LAMANNA, Giovanni (LAPP)

Contribution ID: 23

Type: **not specified**

## Welcome to the AEI, and introduction to ASPERA

*Thursday, 3 May 2012 09:30 (15 minutes)*

I'll first give you a lightning overview of the AEI. The Max-Planck Institute for Gravitational Physics, also known as the Albert Einstein Institute or AEI, is the world's largest institute devoted to the study of gravitation. The main focus in Hannover is the detection of gravitational waves. These were first predicted by Einstein about a century ago; we hope to make the first direct detections with the LIGO/VIRGO and GEO instruments in the coming five years. Searching for weak gravitational-wave signals in detector noise is a large-scale computing and data analysis problem, hence the relevance of this ASPERA workshop.

I'll then give you a lightning overview to ASPERA. ASPERA is a network of European national government agencies responsible for coordinating and funding national research efforts in Astroparticle Physics. Within the ERA-NET scheme of the European Commission, ASPERA started in July 2006. After a first successful three-year period, ASPERA, as ASPERA 2, is now continuing for another three-year programme towards the development of a sustainable body for astroparticle physics in Europe. ASPERA 2 is funded by the European Commission through the 7th Framework Programme.

**Presenter:** Prof. ALLEN, Bruce

Contribution ID: 24

Type: **not specified**

## Large Synoptic Survey Telescope (LSST): New Algorithms and Architectures

*Thursday, 3 May 2012 14:00 (45 minutes)*

When it enters operation, the Large Synoptic Survey Telescope will produce 15TB of image data each night, more than any other optical survey. In many respects, applying existing algorithms at this scale is a significant technical challenge on its own. However, the improved statistical errors and the fact that LSST is “deep, wide, and fast” will demand algorithms that are qualitatively different from those sufficient for surveys that are smaller in any one of these dimensions. In many cases, the computational demands for these more complex algorithms are considerably greater. In this talk, I will touch on several of these computational problems and the LSST collaboration’s plans to address them, with a focus on difference imaging, image coaddition, and galaxy shape measurement for weak lensing. A particular challenge for LSST is that the state-of-the-art in both algorithm development and hardware architecture may change significantly before the survey begins, and our approach must be flexible enough to take advantage of both.

**Presenter:** Dr JIM, Bosch

Contribution ID: 25

Type: **not specified**

## White Paper plan and writing assignments

*Friday, 4 May 2012 11:45 (1h 15m)*

**Presenter:** Prof. KATSANEVAS, Stavros (CNRS/IN2P3)

Contribution ID: 27

Type: **not specified**

## **General purpose, high-performance and energy-efficient x86 based computing with Many-Core Technologies (TFLOPS on a chip)**

*Thursday, 3 May 2012 11:00 (45 minutes)*

As we see Moore's Law alive and well, more and more parallelism is introduced to all computing platforms on all levels of integration and programming to achieve higher performance and energy efficiency. We will discuss the new Intel(r) Many Integrated Core (MIC) architecture for highly-parallel workloads with general purpose, energy efficient TFLOPS performance on a single chip. We will also discuss the journey to ExaScale including technology trends for high-performance and look at some of the R&D areas for HPC at Intel.

**Presenter:** CORNELIUS, Herbert

Contribution ID: 28

Type: **not specified**

## Computing Challenges at the Pierre Auger Observatory

*Friday, 4 May 2012 11:00 (45 minutes)*

Pierre Auger Observatory needs a lot of computing resources for simulation of cosmic ray showers with ultra-high energies up to  $10^{21}$  eV. We are able to use simultaneously several thousand cores and generate more than 1 TB of data daily in the current EGI grid environment. We are limited by available resources and a long duration of a single job for very high energies, which is already simplified by thinning parameter in the Corsika simulation program. Details of the time traces which would be useful for mass composition analyses and hadronic interaction physics get lost by thinning. Thousand times more computing power and correspondingly increased storage is needed for simulations without thinning. Significant speedup could be obtained by using many CPUs or even GPUs for generation of a single shower. We discuss the current trends in the middleware heading to a provision of a whole worker node with many cores to a single parallel job. The expected development in the Corsika and Geant4 towards parallelization and usage of GPU is needed for an efficient usage of the new infrastructure. Possibilities of computing in clouds are also discussed.

**Presenter:** CHUDOBA, Jiri (Acad. of Sciences of the Czech Rep. (CZ))

Contribution ID: 29

Type: **not specified**

## **the ASPERA process**

*Thursday, 3 May 2012 13:45 (15 minutes)*

**Presenter:** KATSANEVAS, Stavros (CNRS/IN2P3)