

Final SCAT meeting and satellite school

Universidad Técnica Federico Santa María

High Performance Heterogeneous Computing: hybrid CPU/GPU computing with NVIDIA CUDA

One of a series of mini-courses taking place 20-24 July 2009

Description

GPUs are attracting a great deal of attention in science, where typical applications are intensely computationally demanding. The reason is that GPUs are inexpensive and massively parallel.

In hybrid CPU/GPU systems, the GPU works as a coprocessor for the CPU. The CPU offloads computing intensive units of work to the GPU. Due to the massively parallel GPU architecture, the way in which units of work are processed by the GPU is greatly different from the processing of a general-purpose CPU. Therefore, an efficient implementation for the highly-specialized hardware of the GPU often requires not only to *recode* the serial application but to *rethink* the algorithm, and consequently, not all algorithms can be accelerated by the GPU. However, where it is possible to accelerate the computations using the GPU, the speed-ups can vary from one to two orders of magnitude in comparison with the most powerful CPU available today.

The focus of this workshop will be on hybrid CPU-GPU computing systems, where the GPU is used as a co-processor to accelerate some of the computations. In this workshop, we will discuss the problem of implementing algorithms for GPU-enabled computing systems from the point of view of an *heterogeneous computing system*.

We will provide an introduction to the CUDA architecture and programming model, and we will present real-case applications implemented for a hybrid CPU/GPU system. We will end with a *hands-on* session, where the attendees will apply the theory and gain real-application experience.

Lecturer:

Felipe A Cruz, University of Bristol, United Kingdom

This course will be in English.

For more information, email info@scat-alfa.eu or visit www.scat-alfa.eu