

GPU COMPUTING WORKSHOP

7 MAY 2009

UWA UNIVERSITY CLUB



Over the past decade, advances in graphics processing units (GPUs) have seen them evolve into highly generalised parallel processors that can be applied to a wide range of computational problems. Impressive performance increases have been observed using GPUs in many computational applications including molecular dynamics, astrophysics simulations, pattern and sequence matching, three-dimensional imaging, and many more. This is your opportunity to meet and discuss the latest developments in general purpose computing on GPUs with other researchers and industry professionals, learn about GPU architectures and programming, and hear from others about their GPU applications.



This workshop and discussion forum aims to provide a detailed introduction to GPU computing with CUDA and NVIDIA Tesla computing solutions, and to present research in GPU and

Heterogeneous computing being undertaken in Western Australia. CUDA is a revolutionary parallel computing architecture from NVIDIA that is built upon a unified computing architecture and multiple software components. The architecture consists of an ISA and hardware compute engine. The available software tools include a C compiler and NVIDIA drivers for developers to build applications using C for CUDA, as well as useful libraries for high-performance computing.

Talks will range from demonstrating application capabilities of GPUs, all the way down to the low-level hardware and programming. Research experts Dr. Mark Harris (NVIDIA) and Professor Karen Haines (WASP@UWA) will present and provide the latest information on their work and other GPU and Heterogeneous computing topics. Topics to be covered will include an introduction to the CUDA architecture, programming model, and the programming environment of C for CUDA, as well as an overview of the Tesla GPU architecture, a live programming demo, and strategies for optimizing CUDA applications for the GPU. To better enable the uptake of this technology, Dragan Dimitrovici from XENON systems will provide an overview of CUDA enabled hardware options. The workshop will also include brief presentations of some of the projects using CUDA within Western Australia, including a background of Parallel Computing strategies required for optimizing your approach to GPU and heterogeneous programming as well as a review of the hardware requirements for getting started with CUDA.



Attendance is FREE for all participants (both iVEC and non-iVEC partners). The number of places for external (non-iVEC) attendees will be limited, so please express your interest as soon as possible! Contact Rosie Radica (rosie.radica@uwa.edu.au) to register with your name, organisation, occupation, iVEC Partner (if applicable), and any dietary requirements. Please use the header – GPU Workshop – in email sent to register. Sign-in on the day will be from 8:30am. Please only register if you are genuinely planning to attend.

Where: UWA University Club

Times: 9am - 5pm

