High Performance GPU Computing with NVIDIA, CUDA, and Fermi.

Thursday 19 August 2010, 8:30-17:00

ARRC auditorium, 26 Dick Perry Avenue, Technology Park, Kensington.

Presented by

Dragan Dimitrovici (Founder, XENON Technology Group),
Derek Gerstmann (Research Fellow, WASP/CMCA, UWA)
Mark Harris (Senior Developer Technology Engineer, NVIDIA)
Christopher Harris (Research Associate, ICRAR, UWA)
Valerie Maxville (Education Program Leader, iVEC)
Munish Mehta (Senior Bioinformatics Officer, UWA)
Lawrence Murray (Research Scientist, CSIRO Mathematics, Informatics and Statistics)
Andrew Rohl (CEO, iVEC),
Randall Wayth (Curtin Institute of Radio Astronomy, ICRAR/Curtin University)

Abstract

CUDA is a parallel computing architecture and programming environment from NVIDIA that enables dramatic increases in computing performance by harnessing the power of the GPU (graphics processing unit). Computing is evolving from "central processing" on the CPU to "co-processing" on the CPU and GPU. To enable this new computing paradigm, NVIDIA has developed the CUDA parallel computing architecture. With over 100 million CUDA-enabled GPUs sold to date, software developers, scientists and researchers are finding broad-ranging uses for CUDA, including image and video processing, computational biology and chemistry, fluid dynamics simulation, CT image reconstruction, seismic analysis, financial computing, ray tracing, and much more.

The latest CUDA-enabled GPU architecture from NVIDIA, code-named "Fermi", is now available in the form of the Tesla 20 series GPU computing solutions, which support many “must have” features for technical and enterprise computing. These include ECC memory for uncompromised accuracy and scalability, support for C++ and 8x the double precision performance compared to Tesla 10-series GPU computing products. NVIDIA Tesla GPUs are being used in 100s of clusters and data centers around the world, including the Nebulae cluster, currently the 2nd fastest supercomputer in the world.

In this workshop you will learn about CUDA, the Fermi architecture, and Tesla GPU Computing products. You will learn about the basics of programming GPUs using CUDA C and C++, the variety of available computational libraries for CUDA, tools for profiling and debugging CUDA applications, and approaches for optimizing CUDA parallel applications. You will also learn about CUDA-enabled desktop, workstation, and cluster computing solutions provided by Xenon Systems. The workshop will also include presentations on some of the ways these technologies are being used by researchers in Western Australia.

Agenda

8:30-9:00	Registration and refreshments.
9:00-9:15	Introduction and welcome. Valerie Maxville.
9:15-9:45	iVEC, an introduction and GPU directions. Andrew Rohl.
9:45-10:30	Introduction to NVIDIA CUDA, Tesla and the Fermi Architecture. Mark Harris.
10:30-10:50	Refreshments.
10:50-11:10	CUDA enabled hardware from XENON Systems. Dragan Dimitrovici.
11:10-12:00	CUDA parallel programming model and live CUDA programming example. Mark Harris.
12:00-12:40	Lunch
12:40-1:00	Genocentric epistasis analysis using CUDA. Munish Mehta.
1:00-1:20	GPU accelerated Monte Carlo inference on large-scale environmental models? Lawrence Murray.
1:20-1:40	The Murchison Widefield Array real-time system. Randall Wayth.
1:40-2:00	Detecting fast transient radio sources with GPUs? Christopher Harris.
2:00-2:20	Incremental volume visualisation for large streaming datasets. Derek Gerstmann.
2:20-2:40	Refreshments.
2:40-4:00	CUDA debugging and profiling tools. Optimizing CUDA application performance and Fermi update. Mark Harris.
4:00-4:30	Panel from the presenters for Q&A. Valerie Maxville convening.

Please RSVP early!

The workshop is free but due to seating limitations and for catering purposes bookings are essential. Please RSVP to admin@wasp.uwa.edu.au or 64888740 (mornings) including name and contact details.