

'General purpose GPU computing: Transforming your desktop into a personal super-computer'

Date(s) Tuesday 28th February 2012 (16:00-17:00)

Download [Add to Calendar \(/research/activity/cncr/news/previous/28Feb12-cncr-seminar.aspx?ical=true\)](/research/activity/cncr/news/previous/28Feb12-cncr-seminar.aspx?ical=true)

Speaker: Dr Thomas Nowotny, Department of Informatics, University of Sussex

Driven by the ever-increasing computational demands of the games industry graphical processing units (GPUs) have developed from simple co-processors to powerful compute platforms. Unlike modern CPUs that are still essentially using a sequential compute model, GPUs are massively parallel with half a thousand cores on a single chip in modern graphics cards. With the introduction of the CUDA (common unified device architecture) application programming interface, NVIDIA, one of the leading GPU manufacturers, has made their GPUs accessible to general purpose computing applications. Given a suitable computational problem and an optimized parallel implementation, modern GPUs can achieve speed-ups of up to 50 to 100 times over a classical CPU core of a recent CPU. In this talk I will introduce the opportunities and challenges of general purpose GPU computing and illustrate them with the GPU enhanced neuronal network (GeNN) framework we are developing at Sussex. I will argue that building meta-compilers that generate code from a simpler domain-specific problem description is one of the better approaches to GPU programming. GeNN is based on such an approach which offers many decisive advantages: The system can provide for a large number of potential model elements, it can optimise for specific model structures and GPU hardware properties at compile time and it can be extended more easily than pre-compiled solutions. The present beta version shows encouraging competitive computing performance and is available under the GPL v2 licenses at <http://genn.sourceforge.net>.