


DLI Workshop on GPU Accelerated Computing with CUDA C/C++

25 July, 2018



DLI Workshop on GPU Accelerated Computing with CUDA C/C++

Dr Jeff Adie
Principal Solutions Architect, nVIDIA

Date:
July 25, 2018
Wednesday
2:30 pm – 4:30

Venue:
KB110,
Knowles Building,
Main Campus,
The University of Hong Kong

Abstract:

CUDA is the parallel computing platform and API built by NVIDIA for GPU programming. For the past 10 years CUDA has gone through tremendous development. Nowadays the most demanding HPC applications, including GROMACS, GAUSSIAN, AMBER, etc are accelerated by CUDA and NVIDIA GPUs because a single GPU can in some cases outperform 20 CPUs' computational power. From weather prediction and materials science to wind tunnel simulation and genomics, GPU-accelerated computing is at the heart of HPC's most promising areas of discovery. In this workshop, participant will gain hands-on experience in coding with CUDA and using it to accelerate applications.

Agenda:

- 1) Why GPU computing/GPU architecture (Presentation)
- 2) Introduction to Accelerated Computing (Hands-on workshop)
- 3) Accelerating Applications with CUDA C/C++ - (Hands-on workshop)

Prerequisites : Basic knowledge of CUDA programming model and basic C/C++ competency

For registration, please visit https://hkuems1.hku.hk/hkuems/ec_hdetail.aspx?ueid=58745

For enquiries, please contact Ms Lilian Chan by email at lilianyl@hku.hk.