Séminaires de l'IDRIS

PGI Accelerator Compilers and CUDA Fortran Overview

Jeudi 28 janvier 2010 (10h30-12h)

Doug Miles

The Portland Group

This presentation gives an overview of programming NVIDIA GPUs using the PGI Accelerator programming model in C and Fortran, and using PGI CUDA Fortran. It introduces the compute-specific details of NVIDIA GPUs, gives a basic example of programming in CUDA C, and provides a detailed overview of the explicit CUDA Fortran programming model implemented in the PGI 10.0 compilers. This provides the basis for an overview of PGI Accelerator directive-based programming for x64+NVIDIA systems, including motivations, goals, current status and capabilities, and upcoming features. The PGI Accelerator programming model is a high-level implicit model for offloading compute-intensive code regions from a host CPU to an Accelerator via Fortran directives or C pragmas, simplifying the coding process and ensuring that the resulting programs are portable to other platforms and compilers.

Doug Miles has worked in the HPC industry since 1985 at Floating Point Systems, Cray, The Portland Group and STMicroelectronics, primarily in applications engineering and management positions.
He has been the director of technical and business operations at PGI since 2003.

L'accès à ce séminaire est libre mais l'enregistrement est obligatoire à l'adresse `http://www.idris.fr:data/seminaires`