these features to become fully proficient at optimizing complex CUDA applications.

Speaker(s): NVIDIA Developer Tools Team

Topic(s): Development Tools & Libraries (Beginner)

THURSDAY, MAY 17, 14:30 (25 MINUTES)
ROOM M

S0022 Scalable Frameworks and Algorithms for Terascale Radio Astronomy Images

Learn how the oldest science is using the newest processors to solve a critical problem: how to accomplish traditional image analysis and visualization tasks when the images are terabytes in size? Simple, standard operations such as displaying 2-d slices, evaluating image statistics, and applying histogram equalization become manifestly challenging when images dramatically exceed single-node memory capacity. We will explain how our hybrid CPU-GPU cluster framework – which can volume render a 200GB image at >50fps! – will support traditional radio astronomy tasks for the colossal images that the Square Kilometre Array and its precursor, the Australian SKA Pathfinder, will generate.

Speaker(s): Christopher Fluke (Senior Lecturer, Swinburne University of Technology - Centre for Astrophysics and Supercomputing)

Topic(s): Astronomy & Astrophysics, Visualization (Intermediate)

THURSDAY, MAY 17, 14:30 (25 MINUTES)
ROOM C

S0032 Teraflop GPU Acceleration Of Large Matrix Algebra

Learn how Multipath's Fast Matrix Solver (FMS) is setting performance records using multiple GPU's solving large matrices in production applications. By (1) leveraging NVIDIA's CUBLAS library, (2) operating multiple GPU's in parallel and (3) overlapping data transfers with computation, FMS averages over 2 teraflops of performance, even on jobs lasting for days. The presentation also includes a description of what problems FMS solves and how it is incorporated into applications programs.

Speaker(s): Ronald Young (President, Multipath Corporation)

Topic(s): Development Tools & Libraries, General Interest (Beginner)

THURSDAY, MAY 17, 14:30 (50 MINUTES)
ROOM L

S0106 GPU Based Numerical Methods in Mathematica

A fast way of developing, prototyping and deploying numerical algorithms that can take advantage of CUDA capable systems is available in Mathematica 8. Over the past year, educators, scientists, and business users have taken advantage of the benefits that the support of GPU programming in Mathematica. By integrating and implementing CUDA/OpenCL in their programs, users make use of a hybrid approach, combining the speed-up that GPUs offer and a powerful numerical development system. In this presentation several examples describing numerical applications ranging from deconvolution of MRI imaging, linear solvers for FEM, systems of ODEs, line integral convolution visualization are presented.

Speaker(s): Ulises Cervantes-Pimentel (Senior Kernel Developer, Wolfram Research), Abdul Dakkak (Kernel Developer, Wolfram Research)

Topic(s): Algorithms & Numerical Techniques, Visualization, Application Design & Porting Techniques, Development Tools & Libraries (Intermediate)

THURSDAY, MAY 17, 14:30 (25 MINUTES)
MARRIOTT BALLROOM 3

S0231 Levenberg-Marquardt Using Block Sparse Matrices on CUDA

This session describes the experiences of constructing GPU based matrix-vector functions for block sparse matrices having multiple block sizes and a domain-specific numerical Jacobian generation function. The bundle adjustment algorithm is an optimization procedure which attempts to refine the relative camera pose, and 3D structure location variables, estimated from multiple sets of images. The Conjugate Gradient algorithm is used to solve the normal equations which appear in the inner loop to the non-linear least squares problem.

Speaker(s): Tetsuo Tawara (Software Engineer, Koozyt)

Topic(s): Application Design & Porting Techniques, Algorithms & Numerical Techniques (Intermediate)

THURSDAY, MAY 17, 15:00 (50 MINUTES)
ROOM C

S0071 The High-Level Linear Algebra Library ViennaCL and Its Applications

Get to know ViennaCL, an OpenCL high-level linear algebra software, which allows to get the speed of GPU computing at the convenience level of the C++ Boost libraries. Decrease the development and execution time of applications by utilizing our well-tested and widely used library, instead of spending days on learning details of GPU architectures and debugging. We provide examples that demonstrate not only how quickly existing applications are ported efficiently from single-threaded execution to fully utilizing multi-threaded environments, but also how to utilize the rich set of functionalities ranging from common BLAS routines to iterative solvers.

Speaker(s): Karl Rupp (Project Assistant, TU Wien)

Topic(s): Development Tools & Libraries, Algorithms & Numerical Techniques, Computational Physics (Intermediate)

THURSDAY, MAY 17, 15:00 (25 MINUTES)
ROOM M

S0087 GPU Acceleration of Dense Stellar Clusters Simulation

Computing the interactions between stars within dense stellar clusters is a problem of fundamental importance in theoretical astrophysics. This paper presents the parallelization of a Monte Carlo algorithm for simulating stellar cluster evolution using programmable Graphics Processing Units. The kernels of this algorithm exhibit high levels of data dependent decision making and unavoidable non-contiguous memory accesses. However, we adopt various parallelization strategies and utilize the high computing power of the GPU to obtain substantial near-linear speedups which cannot be easily achieved on a CPU-based system. This acceleration allows to explore physical regimes which were out of reach of current simulations.

Speaker(s): Bharath Pattabiraman (PhD Student, Northwestern University), Stefan Umbreit (Postdoctoral Associate, Northwestern University)

Topic(s): Astronomy & Astrophysics, Computational Physics, Algorithms & Numerical Techniques (Intermediate)

THURSDAY, MAY 17, 15:00 (25 MINUTES)
ROOM N

S0091 Sustainable Hybrid Parallelization of an Unstructured Hydrodynamic Code

The goal of this presentation is to share our methodology for