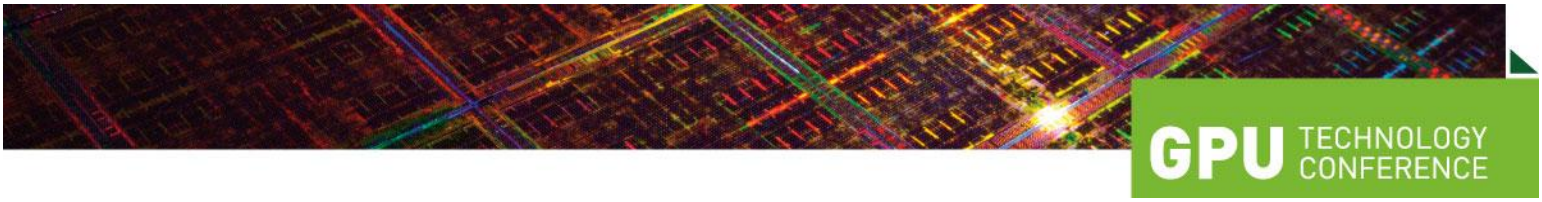

GPU Technology Conference, May 14-17, 2012
McEnery Convention Center, San Jose, California
www.gputechconf.com

Sessions on Computational Physics (subject to change)

IMPORTANT: Visit http://www.gputechconf.com/page/sessions.html for the most up-to-date schedule.

S0268 - Virtual Process Engineering - Realtime Simulation of Multiphase Systems
Wei Ge (Institute of Process Engineering, Chinese Academy of Sciences)
Day: Tuesday, 05/15 | **Time:** 9:00 am - 9:50 am
Topic Areas: Computational Fluid Dynamics; Molecular Dynamics; Computational Physics; Algorithms & Numerical Techniques
Session Level: Advanced

Realtime simulation and virtual reality with quantitatively correct physics for industrial processes with multi-scale and multiphase system is once a remote dream for process engineering, but is becoming true now with CPU-GPU hybrid supercomputing. Numerical and visualization methods for such simulations on thousands of GPUs will be reported with applications in chemical and energy industries.

S0258 - Sailfish: Lattice Boltzmann Fluid Simulations with GPUs and Python
Michal Januszewski (University of Silesia in Katowice; Google Switzerland)
Day: Tuesday, 05/15 | **Time:** 9:30 am - 9:55 am
Topic Areas: Computational Fluid Dynamics; Computational Physics; Development Tools & Libraries
Session Level: Intermediate

Learn how Run-Time Code Generation (RTCG) techniques allowed for fast development of a lattice Boltzmann (LB) fluid dynamics solver called Sailfish. Sailfish is completely open source, supports a wide variety of LB models (single and multiple relaxation times, the entropic model; single and binary fluids) and can take advantage of multiple GPUs. Even though the project is written predominantly in Python, no performance compromises are made. This talk will introduce the basic design principles of Sailfish and illustrate how RTCG allows to exploit the power of GPUs with minimal programmer effort.

S0031 - Unstructured Grid Numbering Schemes for GPU Coalescing Requirements
Andrew Corrigan (Naval Research Laboratory), Johann Dahm (University of Michigan)
Day: Tuesday, 05/15 | **Time:** 10:00 am - 10:25 am
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques; Computational Physics
Session Level: Advanced

Learn how to achieve high performance for computational fluid dynamics (CFD) solvers over unstructured grids using numbering schemes tailored for GPU coalescing requirements. Using these techniques, unstructured grid CFD solvers can make more effective use of memory bandwidth, which is an otherwise significant performance bottleneck that has so far led to relatively limited performance gains on GPUs in comparison to structured grid CFD solvers. Performance benchmarks will be shown using the Jet Engine Noise Reduction (JENRE) code.

S0321 - GPU-Based Monte Carlo Ray Tracing Simulation for Solar Power Plants
Claus Nilsson (Tietronix Software, Inc.), Michel Izygon (Tietronix Software, Inc.)
Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:25 pm
Topic Areas: Energy Exploration; Computational Physics; Ray Tracing
Session Level: Beginner

Learn about real time simulations of Concentrating Thermal Solar Power using GPU technology to enable performance optimization of these utility scale plants. By leveraging the power of GPUs and the parallel aspect of the field of thousands sun-tracking mirrors, we have been successful in cutting the computation time by orders of magnitude versus the previously required minutes and hours runtime. We will present an overview of the problem domain and describe how we used the GPU to derive a Monte Carlo physics ray tracing method to simulate the flux reflected by the mirrors onto the solar receiver.

S0046 - Application of the GPU to a Two-Part Computational Electromagnetic Algorithm
Eric Dunn (SAIC)
Day: Tuesday, 05/15 | **Time:** 2:30 pm - 2:55 pm
Topic Areas: Computational Physics; Algorithms & Numerical Techniques; Ray Tracing
Session Level: Beginner

The shooting and bouncing ray (SBR) method is one way to simulate electromagnetic field radiation. Like all methods, there are certain problems where it does not yield accurate results. In this presentation, we will explain one such case that consists of an antenna resonating between two metal plates. We will discuss how we used the graphics processing unit (GPU) to separate the problem into two parts. Each part is simulated individually with SBR producing an improved result. Such a GPU-accelerated, two-part approach can be applied to other more general hybrid simulations.

S0379 - GPU-based High-Performance Simulations for Spintronics
Jan Jacob (University of Hamburg - Institute of Applied Physics and Microstructure Research Center)
Day: Tuesday, 05/15 | **Time:** 2:30 pm - 2:55 pm
Topic Areas: General Interest; Computational Physics; Application Design & Porting Techniques
Session Level: Intermediate

The joint utilization of the electron's charge and spin in "spintronics" represents a promising technology for data processing and storage in nanostructures. The complex quantum effects like the spin-Hall effect in these devices require demanding numerical simulations providing a convenient link between idealized analytical models to often very complex results from measurements. The simulations involving multiplications and inversions of large matrices provide an ideal showcase for performance gain by employing GPGPUs in the execution of the algebraic routines on these matrices in computing environments with shared execution of algorithms on multiple nodes with multiple GPGPUs and CPU cores.

S0036 - Multiparticle Collision Dynamics on GPUs
Elmar Westphal (Forschungszentrum Juelich)
Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Computational Physics; Computational Fluid Dynamics; Molecular Dynamics
Session Level: Intermediate

See how we employ GPUs to simulate the interaction of millions of solvent and solute particles of a fluid system. Often the domain of large cluster system, the most time consuming part of our simulations can now be done on desktop PCs in reasonable time. This contribution shows how GPUs can effectively be used to accelerate existing programs and how techniques like streaming and increased data locality significantly enhance calculation throughput. It also shows how a GPU-optimized program structure yields usually expensive additional

functionality "almost free". Furthermore, a well-scaling single-node/multi-GPU implementation of the program is presented.

S0067 - PIConGPU - Bringing large-scale Laser Plasma Simulations to GPU Supercomputing
Michael Bussmann (Helmholtz-Zentrum Dresden-Rossendorf), Guido Juckeland (Center for Information Services and High Performance Computing, Technical University Dresden)
Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Computational Physics; Algorithms & Numerical Techniques; Application Design & Porting Techniques; Supercomputing
Session Level: Advanced

With powerful lasers breaking the Petawatt barrier, applications for laser-accelerated particle beams are gaining more interest than ever. Ion beams accelerated by intense laser pulses foster new ways of treating cancer and make them available to more people than ever before. Laser-generated electron beams can drive new compact x-ray sources to create snapshots of ultrafast processes in materials. With PIConGPU laser-driven particle acceleration can be computed in hours compared to weeks on standard CPU clusters. We present the techniques behind PIConGPU, detailed performance analysis and the benefits of PIConGPU for real-world physics cases.

S0221 - 1024 Bit Parallel Rational Arithmetic Operators for the GPU
Robert Zigon (Beckman Coulter)
Day: Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm
Topic Areas: Algorithms & Numerical Techniques; Computational Physics
Session Level: Intermediate

Learn how to create a set of rational arithmetic operators that manipulate 1024 bit operands on a Tesla C2050. These operators are used to create a numerically stable implementation for Bessel functions. Naive implementations of the Bessel functions produce unreliable results when they are used to solve Maxwell's equations by way of Mie theory. Maxwell's equations are used to model the scattering of light by small particles. Light scatter is used in Particle Characterization to measure the quality of materials like cocoa, cement and pharmaceuticals.

S0245 - Porting Legacy Plasma Codes to GPU
Peng Wang (NVIDIA)
Day: Tuesday, 05/15 | **Time:** 4:00 pm - 4:25 pm
Topic Areas: Computational Physics; Computational Physics
Session Level: Intermediate

Learn how to port legacy Fortran plasma codes to GPU. Many legacy plasma codes are written in Fortran and have many lines of codes. We will discuss techniques in porting such legacy codes easily and efficiently to CUDA C/C++. Performance analysis of major algorithmic patterns in plasma codes will be discussed. The discussion will use the GTC and GeFi plasma code as realistic examples.

S0058 - Advancing GPU Molecular Dynamics: Rigid Bodies in HOOMD-blue
Joshua Anderson (University of Michigan), Trung Dac Nguyen (University of Michigan)
Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:50 am
Topic Areas: Molecular Dynamics; Computational Physics
Session Level: Intermediate

Learn how rigid body dynamics are implemented in HOOMD-blue. Previous releases were capable of executing classical molecular dynamics -- where free particles interact via smooth potentials and their motion through time

is computed using Newton's laws. The latest version allows particles to be grouped into bodies that move as rigid units. Users can now simulate materials made of cubes, rods, bent rods, jacks, plates, patchy particles, bucky balls, or any other arbitrary shapes. This talk covers how these algorithms are implemented on the GPU, tuned to perform well for bodies of any size, and discusses several use-cases relevant to research.

S0125 - Memory Efficient Reverse Time Migration in 3D
Chris Leader (Stanford Exploration Project)
Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am
Topic Areas: Energy Exploration; Computational Physics
Session Level: Intermediate

Learn how we can image the interior of the Earth in three dimensions using Reverse Time Migration. We discuss how GPUs accelerate this method using parallel wave propagation kernels, texture memories and minimal device to host transfers. Further we discuss how the progression to 3D presents a multitude of new problems, particularly memory based - causing the system to be IO limited. By manipulating boundary positions and values to a pseudo-random form we show how many of these memory restrictions can be diminished and how detailed subsurface images can be fully constructed using GPUs.

S0236 - Advanced Optimization Techniques On a CUDA Implementation of Conjugate Gradient Solvers
Eri Rubin (OptiTex)
Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am
Topic Areas: Algorithms & Numerical Techniques; Algorithms & Numerical Techniques; Computational Physics; Application Design & Porting Techniques
Session Level: Intermediate

Linear systems are at the heart of allot of compute problems. In large sparse systems, there are 2 distinct approaches, the direct and iterative solvers. After many years of researching and testing both approaches, on CPU and GPU we have implemented a highly efficient CG solver on the GPU using a combination of unique techniques. In this talk we will go over these techniques and the improved performance they bring.

S0312 - GPU Implementation for Rapid Iterative Image Reconstruction in Nuclear Medicine
Jakub Pietrzak (University of Warsaw)
Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am
Topic Areas: Medical Imaging & Visualization; Computational Physics; Computer Graphics
Session Level: Intermediate

GPU implementation can greatly accelerate iterative techniques of 3D image reconstruction in nuclear medicine imaging. Single Photon Emission Computed Tomography (SPECT) is a functional imaging modality widely used in clinical diagnosis. To obtain high quality images within reduced scanning times high sensitivity collimators need to be used and their response function modeled in the reconstruction. This is in general very computationally intensive and unfeasible with CPU and algorithm implementations. Our software is able to perform the reconstruction of patient data within clinically acceptable times using relatively low cost and widely available hardware.

S0352 - GPU-Accelerated Parallel Computing for Simulation of Seismic Wave Propagation
Taro Okamoto (Department of Earth and Planetary Sciences, Tokyo Institute of Technology)
Day: Wednesday, 05/16 | **Time:** 10:30 am - 10:55 am
Topic Areas: Computational Physics; General Interest
Session Level: Advanced

We adopted GPU to accelerate large-scale, parallel finite-difference (FDTD) simulation of seismic wave propagation. Effective parallel implementation is needed because the size of the memory of a single GPU is too small for real applications. Thus we describe the memory optimization, the three-dimensional domain decomposition, and overlapping the communication and computation adopted in our program. We achieved so far a high performance (single-precision) of about 61 TFlops by using 1200 GPUs of TSUBAME-2.0, the GPU supercomputer in Tokyo Institute of Technology, Japan. As an important application, we show the results of the simulation of the 2011 Tohoku-Oki mega-quake.

S0269 - Accelerating 3D-RISM Calculations using GPUs
Yutaka Maruyama (Institute for Molecular Science), Fumio Hirata (Institute for Molecular Science)
Day: Wednesday, 05/16 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Life Sciences; Algorithms & Numerical Techniques; Computational Physics
Session Level: Intermediate

The three-dimensional reference interaction site model (3D-RISM) theory, is a powerful tool to investigate biomolecular processes in solution. Unfortunately, 3D-RISM calculations are often both memory intensive and time-consuming. We sought to accelerate these calculations using GPUs. To work around the problem of limited memory size in GPUs, we modified the less memory-intensive Anderson method for faster convergence of 3D-RISM calculations. Using this method on C2070, we reduced the computational time by a factor of eight compared to Intel Xeon (8 cores, 3.33GHz) with the conventional method.

S0055 - Particle Dynamics with MBD and FEA using CUDA
Graham Sanborn (FunctionBay)
Day: Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm
Topic Areas: Computational Structural Mechanics; Computational Physics; Computational Fluid Dynamics
Session Level: Intermediate

Many sphere particles are solved with DEM (Discrete Element Method) and simulated with GPU technology. Fast algorithm is applied to calculate hertzian contact forces between many sphere particles (from 100,000 to 1,000,000) and NVIDIA’s CUDA is used to accelerate the calculation. Many sphere particles and MBD and FEA entities are simulated within commercial software RecurDyn. Many models are built and simulated; fork lifter with sand model, oil in oil tank model, oil filled engine system and water filled washing machine model. All models are simulated with NVIDIA's GPU and the result is shown.

S0363 - Efficient Molecular Dynamics on Heterogeneous GPU Architectures in GROMACS
Szilárd Páll (KTH Royal Institute of Technology), Berk Hess (KTH Royal Institute of Technology)
Day: Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm
Topic Areas: Molecular Dynamics; Computational Physics; Life Sciences
Session Level: Intermediate

Molecular Dynamics is an important application for GPU acceleration, but many algorithmic optimizations and features still rely on code that prefers traditional CPUs. It is only with the latest hardware and software we have been able to realize a heterogeneous GPU/CPU implementation and reach performance significantly beyond the state-of-the-art of hand-tuned CPU code in our GROMACS program. The sub-millisecond iteration time poses challenges on all levels of parallelization. Come and learn about our new atom-cluster pair interaction approach

for non-bonded force evaluation that achieves 60% work-efficiency and other innovative solutions for heterogeneous GPU systems.

S0139 - GPU-Based Molecular Dynamics Simulations of Protein and RNA Assembly
Samuel Cho (Wake Forest University)
Day: Wednesday, 05/16 | **Time:** 5:00 pm - 5:25 pm
Topic Areas: Molecular Dynamics; Computational Physics
Session Level: Intermediate

Protein and RNA biomolecular folding and assembly problems have important applications because misfolding is associated with diseases like Alzheimer's and Parkinson's. However, simulating complex biomolecules on the same timescales as experiments is an extraordinary challenge due to a bottleneck in the force calculations. To overcome these hurdles, we perform coarse-grained molecular dynamics simulations where biomolecules are reduced into simpler components. Furthermore, our GPU-based simulations have a significant performance improvement over CPU-based simulations, which is limited to systems of 50-150 residues/nucleotides. The GPU-based code can simulate protein/RNA systems of 400-10,000+ residues/nucleotides, and we present ribosome assembly simulations.

S0129 - A Monte Carlo Thermal Radiation Solver in GPU/CPU Hybrid Architecture
Gaofeng Wang (Laboratoire E.M2.C, Ecole Centrale Paris), Oliver Gicquel (Laboratoire E.M2.C, Ecole Centrale Paris)
Day: Thursday, 05/17 | **Time:** 9:00 am - 9:25 am
Topic Areas: Computational Fluid Dynamics; Computational Fluid Dynamics; Computational Physics; Ray Tracing
Session Level: Intermediate

A Monte Carlo ray-tracing code is developed to predict radiative heat transfer behaviors in CFD simulation of combustion phenomena. Using emission-reciprocal method, each random ray casting of each node could be independently conducted for parallel computations. The code is efficiently implemented in hybrid GPU/CPU HPC resources using a dedicated dynamic load balancing strategy. A linear speedup scaling of hybrid HPC resources has been shown in demonstrating calculation of radiative heat transfer of a helicopter engine's combustion chamber, while adding one GPU in HPC resources pool is in sense of nine CPU cores supplements.

S0508 - Faster Finite Elements for Wave Propagation Codes
Max Rietmann (Institute for Computational Science / USI Lugano, Switzerland)
Day: Thursday, 05/17 | **Time:** 10:00 am - 10:25 am
Topic Areas: Algorithms & Numerical Techniques; Computational Physics
Session Level: Intermediate

Learn how to develop faster and better finite-element codes for wave propagation using GPUs and MPI combined with overlapping techniques to hide the cost of communications and of host/device memory copies. Different options based on mesh coloring or on atomic operations will be presented. The difficulty to define speedup will also be discussed (speedup versus what? using what definition of "cost"?). Examples will be given using SPECFEM3D, a highly optimized spectral finite-element code that has won the Gordon Bell Supercomputing award and the BULL Joseph Fourier award, and that can run on CPU or GPU clusters.

S0039 - Data-Driven GPGPU Ideology Extension
Alexandr Kosenkov (University of Geneva), Bela Bauer (Microsoft Research)
Day: Thursday, 05/17 | **Time:** 10:00 am - 10:25 am
Topic Areas: Application Design & Porting Techniques; Computational Physics; Parallel Programming Languages & Compilers; Development Tools & Libraries
Session Level: Advanced

In this session we will demonstrate how the GPGPU ideology can be extended so that it can be used on a scale of Infiniband hybrid system. The approach that we are presenting combines delayed execution, scheduling techniques and, most importantly, casts down the CPU multi-core ideology to the streaming multiprocessor's one enforcing full fledged "GPGPU as a co-processor" way of programming for large-scale MPI hybrid applications. Staying compatible with modern CPU/GPGPU libraries it provides more than a fine grained control over resources - more than you wanted that is.

S0217 - Efficient Implementation of CFD Algorithms on GPU Accelerated Supercomputers
Ali Khajeh-Saeed (University of Massachusetts, Amherst), Blair Perot (University of Massachusetts, Amherst)
Day: Thursday, 05/17 | **Time:** 10:30 am - 10:55 am
Topic Areas: Computational Fluid Dynamics; Computational Physics; Supercomputing; Application Design & Porting Techniques
Session Level: Intermediate

The goal of this session is to introduce the concepts necessary to perform large computational fluid dynamic (CFD) problems on collections of many GPUs. Communication and computation overlapping schemes become even more critical when using fast compute engines such as GPUs that are connected via a relatively slow interconnect (such as MPI on InfiniBand). The algorithms presented are validated on unsteady CFD simulations of turbulence using 192 graphics processors to update half-a-billion unknowns per computational timestep. The performance results from three different GPU accelerated supercomputers (Lincoln, Forge, and Keeneland) are compared with a large CPU based supercomputer (Ranger).

S0378 - VASP Accelerated with GPUs
Maxwell Hutchinson (University of Chicago)
Day: Thursday, 05/17 | **Time:** 2:00 pm - 2:50 pm
Topic Areas: Quantum Chemistry; Application Design & Porting Techniques; Computational Physics
Session Level: Intermediate

This session will detail the performance and capabilities of GPU-accelerated VASP, explain design decisions made in porting VASP to CUDA, and present a roadmap for GPU accelerated VASP development. We've achieved performance improvements up to around 20x on systems of around 100 ions and have implemented exact-exchange. We are working on ports of more conventional functionality.

S0071 - The High-Level Linear Algebra Library ViennaCL And Its Applications
Karl Rupp (TU Wien)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Development Tools & Libraries; Algorithms & Numerical Techniques; Computational Physics
Session Level: Intermediate

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.

S0087 - GPU Acceleration of Dense Stellar Clusters Simulation
Bharath Pattabiraman (Northwestern University), Stefan Umbreit (Northwestern University)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Astronomy & Astrophysics; Computational Physics; Algorithms & Numerical Techniques
Session Level: Intermediate

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.

S0368 - Unraveling the Mysteries of Quarks with Hundreds of GPUs
Ronald Babich (NVIDIA)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Computational Physics; Application Design & Porting Techniques; Algorithms & Numerical Techniques; Supercomputing
Session Level: Intermediate

Dive into the world of quarks and gluons, and hear how GPU computing is revolutionizing the way many calculations in lattice quantum chromodynamics (lattice QCD) are performed. The main computational challenge in such calculations is to repeatedly solve large systems of linear equations arising from a four-dimensional finite-difference problem. In this session, we'll discuss strategies for parallelizing such a solver across hundreds of GPUs. These include techniques and algorithms for reducing memory traffic and inter-GPU communication. The net result is an implementation that achieves better than 20 Tflops on 256 GPUs, realized in the open-source "QUDA" library.

S0091 - Sustainable Hybrid Parallelization of an Unstructured Hydrodynamic Code
Raphaël Poncet (Commissariat à l'Energie Atomique et aux Energies Alternatives)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Application Design & Porting Techniques; Algorithms & Numerical Techniques; Computational Fluid Dynamics; Computational Physics
Session Level: Advanced

The goal of this presentation is to share our methodology for porting a numerical code to hybrid supercomputing architectures using MPI coupled with directive-based languages (OpenMP for multicore CPUs, and HMPP for GPUs). Our code, VOLNA, is an unstructured partial differential equation hydrodynamic solver developed for the simulation of tsunamis. Our results demonstrate that using directive-based languages such as HMPP for GPU programming, one can retain good performance (e.g. speedup of 15 compared to 1 CPU core, 3 compared to 8 CPU cores) with minimal modifications of the original CPU source code (about 30 lines of directives in our case).

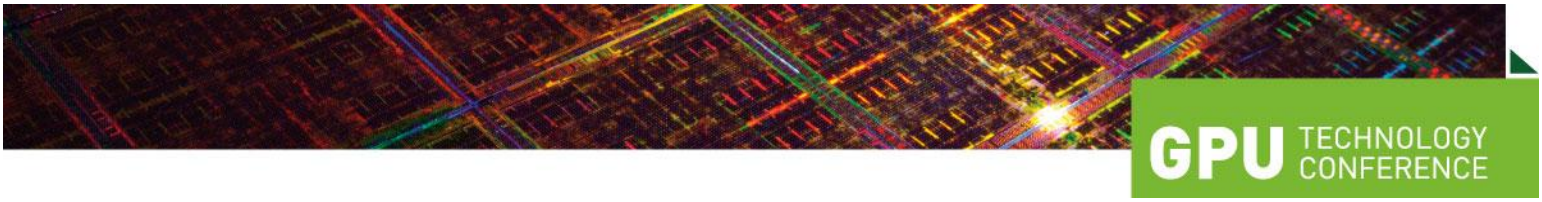

S0334 - The Fast Multipole Method on CPU and GPU Processors
Eric Darve (Stanford)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Computational Physics; Molecular Dynamics; Algorithms & Numerical Techniques
Session Level: Advanced

The fast multipole method (FMM) is a widely used numerical algorithm in computational engineering. Accelerating the FMM on CUDA-enabled GPUs is challenging because the FMM has a complicated data access pattern, mostly during the so-called multipole-to-local (M2L) operation. We have created several schemes to optimize the M2L and have attained a performance of over 350 (resp. 160) Gflop/s for single (double) precision arithmetic. The optimal algorithm was incorporated into a complete FMM code, which can accept any smooth kernel as specified by the user, making it very flexible. We have also developed a highly efficient CPU version.

S0282 - Leveraging NVIDIA GPUDirect on APEnet+ 3D Torus Cluster Interconnect
Davide Rossetti (Italian National Institue for Nuclear Physics)
Day: Thursday, 05/17 | **Time:** 4:30 pm - 4:55 pm
Topic Areas: Supercomputing; Computational Physics
Session Level: Intermediate

APEnet+ is a novel cluster interconnect, based on a custom PCI card which features a PCI Express Gen2 X8 link and a re-configurable HW component (FPGA). It supports a 3D Torus topology and has special acceleration features specifically developed for NVIDIA Fermi GPUs. An introduction to the basic features and the programming model of APEnet+ will be followed by a description of its performance on some numerical simulations, e.g. High Energy Physics simulations.

S0218 - ASI Parallel Fortran: A General-Purpose Fortran to GPU Translator
Rainald Lohner (George Mason University)
Day: Thursday, 05/17 | **Time:** 4:30 pm - 4:55 pm
Topic Areas: Development Tools & Libraries; Computational Fluid Dynamics; Computational Physics; Parallel Programming Languages & Compilers
Session Level: Advanced

Over the last 3 years we have developed a general-purpose Fortran to GPU translator: ASI Parallel Fortran does. The talk will detail its purpose, design layout and capabilities, and show how it is used and implemented. The use of ASI Parallel Fortran will be shown for large-scale CFD/CEM codes as well as other general purpose Fortran codes.