

GPU Technology Conference 2010 Sessions on High Performance Computing (subject to change)

IMPORTANT: Visit www.nvidia.com/gtc for the most up-to-date schedule and to enroll into sessions to ensure your spot in the most popular courses.

2133 - 3D Full Wave EM Simulations Accelerated by GPU Computing

3D Full Wave Electromagnetic simulations of RF components, antennas, printed circuit boards, can be quite time consuming. Computer Simulation Technology (CST) toolsuite includes the capability to activate GPU Computing. Examples will be shown of using Tesla C1060 and S1070 configurations to provide significant performance improvement of complex simulations.

Speaker: Fabrizio Zanella, CST of America

Topic: High Performance Computing

Time: Thursday, September, 23rd, 16:00 - 16:20

2283 - 500 Teraflops Heterogeneous Cluster

HPC Affiliated Resource Center (ARC) will be host of a very large interactive HPC. The large cluster (CONDOR) will integrate cell broadband engine processors, GPGPUs and powerful x86 server nodes, with a combined capability of 500 Teraflops. Applications will include neuromorphic computing, video synthetic aperture radar backprojection, matrix multiplications, and others. This presentation will discuss progress on performance optimization using the Heterogeneous Cluster and lessons learned from this research.

Speaker: Mark Barnell, Air Force Research Lab (AFRL)

Topic: High Performance Computing

Time: Thursday, September, 23rd, 16:00 - 16:50

2205 - A Highly Reliable RAID System Based on GPUs

While RAID is the prevailing method of creating reliable secondary storage infrastructure, many users desire more flexibility than offered by current implementations. To attain needed

performance, customers have often sought after hardware-based RAID solutions. This talk describes a RAID system that offloads erasure correction coding calculations to GPUs, allowing increased reliability by supporting new RAID levels while maintaining high performance.

Speaker: Matthew Curry, Sandia National Laboratories and the University of Alabama at Birmingham

Topic: High Performance Computing

Time: Tuesday, September, 21st, 17:00 - 17:50

2058 - A Practical Introduction to Computational Fluid Dynamics on GPUs

Learn step-by-step procedures to write an explicit CFD solver based on finite difference methods with staggered grid allocations and boundary fitted coordinates. We will discuss the derivation of the mathematical model, discretization of the model equations, development of the algorithms, and parallelization and visualization of the computed data using OpenCL and OpenGL. Compares case studies of natural convection, driven cavity, scaling analysis, and magneto-thermal convection computed using CSIRO's CPU/GPU supercomputer cluster to known analytical and experimental solutions.

Speaker: Tomasz Bednarz, CSIRO

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Physics Simulation

Time: Wednesday, September, 22nd, 10:00 - 10:50

2206 - Accelerated Computational Fluid Dynamics Employing GPUs

Coming soon.

Speaker: Daniel Gaudlitz, FluiDyna

Topics: Computational Fluid Dynamics, High Performance Computing

Time: Thursday, September, 23rd, 09:00 - 9:20

2240 - Accelerating LS-DYNA with MPI, OpenMP, and CUDA

When solving implicit problems, the computational bottleneck in LS-DYNA is the multifrontal linear solver. These operations are performed with double precision arithmetic, hence until the

arrival of the Tesla 2050, experiments with GPU acceleration were only a curiosity. This is no longer the case, and in this talk we will describe how LS-DYNA's hybrid (MPI and OpenMP) solver is further accelerated using GPUs to factor large dense frontal matrices.

Speaker: Bob Lucas, USC

Topics: High Performance Computing, Algorithms & Numerical Techniques

Time: Thursday, September, 23rd, 14:30 - 14:50

2208 - Acceleration of SIMULIA's Abaqus Solver on NVIDIA GPUs

Learn about Acceleware's and Dassault Systemes' integrated solution that performs an LDL^T factorization on GPUs within the Abaqus software package. We will discuss efficient GPU parallelization of the factorization algorithm and enabling the CPU and GPU to overlap their computations and data transfers. Includes an end user simulation case study and GPU performance measurements including 300 GFlops in single precision and 145 GFlops in double precision on NVIDIA Tesla C2050.

Speaker: Chris Mason, Acceleware

Topic: High Performance Computing

Time: Thursday, September, 23rd, 15:30 - 15:50

2089 - Analyzing CUDA Accelerated Application Performance at 20 PFLOP/s

Learn how applications can be executed over multiple GPUs located in multiple hosts, what the challenges are to scale one application to a 20 PFLOP/s machine and why tool support is a necessity. Receive an overview on the available performance analysis tools that support CUDA developers in generating applications with outstanding speedups.

Speaker: Guido Juckeland, TU Dresden - ZIH

Topics: High Performance Computing, Tools & Libraries

Time: Wednesday, September, 22nd, 17:00 - 17:50

2270 - Appro's GPU Computing Solutions

Learn how GPU's are changing the High Performance Computing landscape to deliver price/performance levels that were previously considered unachievable. Join Appro (<http://www.appro.com>), a leading provider of supercomputing solutions; to discuss the introduction of the Appro Tetra server, the most powerful GPU server available today in a 1U form factor and the availability of a new modular GPU expansion blade, both based on NVIDIA Tesla 20-series GPUs. The availability of these two products is a confirmation of Appro's commitment in providing the most innovative and powerful computing platforms at very attractive prices to the High Performance Computing markets.

Speaker: John Lee, Appro

Topic: High Performance Computing

Time: Tuesday, September, 21st, 15:00 - 15:50

2213 - BCSLIB-GPU: Significant Performance Gains for CAE

Hear product architects and developers describe the algorithmic depths and high level breath of the use of GPUs that have been employed to create BCSLIB-GPU, the GPU enablement of the industry standard sparse matrix software suite, BCSLIB-EXT. We provide a range of comparison data with Tesla and Fermi compared with multi-core CPU only systems and for a wide range of realistic demanding real world test problems.

Speaker: Danl Pierce, Access Analytics Int'l, LLC

Topics: Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing, Embedded & Automotive

Time: Thursday, September, 23rd, 15:00 - 15:20

2238 - Better Performance at Lower Occupancy

It is usually advised to optimize CUDA kernels for higher occupancy to hide memory and arithmetic latencies better. In this presentation, I show that increasing occupancy is not the only way and not always the best way to hide latency on GPU. Instead, it may be advantageous to rely on the parallelism within threads-instruction-level parallelism. This insight yields a simple optimization technique that is used in later versions of CUBLAS and CUFFT. I discuss the rationale behind the technique and illustrate it by speeding up matrix multiplication, starting with the basic implementation found in the NVIDIA GPU Computing SDK.

Speaker: Vasily Volkov, UC Berkeley

Topic: High Performance Computing

Time: Wednesday, September, 22nd, 15:00 - 15:50

2204 - Bridging GPU Computing and Neuroscience to Build Large-Scale Face Recognition on Facebook.

Abstract coming soon.

Speaker: Nicolas Pinto, MIT

Topics: Computer Vision, High Performance Computing, Machine Learning & Artificial Intelligence, Neuroscience

Time: Wednesday, September, 22nd, 14:00 - 14:50

2143 - CUDA Fortran Programming for NVIDIA GPUs

An introduction to programming NVIDIA GPUs using CUDA Fortran. Suitable for expert Fortran or CUDA C programmers who need to extract maximum performance from GPUs using an explicit GPU Fortran programming model. Introduces the CUDA Fortran language, and through examples, illustrates how to explicitly program GPUs in native Fortran 95/03 through creation of GPU kernel subroutines, management of host and device memory, definition of CUDA grids and thread blocks, launching kernels, and use of the CUDA Fortran runtime API. This talk includes a live component with a Windows laptop containing an NVIDIA GPU and the PGI CUDA Fortran compiler.

Speaker: Brent Leback, The Portland Group

Topics: Tools & Libraries, High Performance Computing, Programming Languages & Techniques

Time: Wednesday, September, 22nd, 15:30 - 15:50

2137 - CUDA for Real-Time Multigrid Finite Element Simulation of Soft Tissue Deformations

The take-away of this presentation is an efficient CUDA implementation of a finite hexahedra multigrid solver for simulating elastic deformable models in real time. Due to the regular shape of the numerical stencil induced by the hexahedral regime, computations and data layout can be restructured to avoid execution divergence and to support memory access patterns enabling the hardware to coalesce multiple memory accesses into single memory transactions. This enables to effectively exploit the GPU's parallel processing units and high memory bandwidth. Performance gains of up to a factor of 12 compared to a highly optimized CPU implementation are demonstrated.

Speakers: Christian Dick, Technische Universität München, Joachim Georgii, Technische Universität München

Topics: Physics Simulation, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 14:00 - 14:50

2057 - CUDA-Accelerated LINPACK on Clusters

This talk will illustrate the use of GPUs to accelerate the LINPACK benchmark on clusters with GPUs, where both the CPUs and the GPUs are used in synergy.

The acceleration is obtained executing DGEMM (matrix multiply) and DTRSM (for the solution of triangular systems) calls simultaneously on both GPU and CPU cores.

Details of the implementation will be presented together with results that shows how effective the solution is, both for performance and power efficiency.

Speakers: Everett Phillips, NVIDIA, Massimiliano Fatica, NVIDIA

Topics: High Performance Computing, Algorithms & Numerical Techniques

Time: Tuesday, September, 21st, 14:00 - 14:50

2153 - CULA - A Hybrid GPU Linear Algebra Package

Get the latest information on CULA, an implementation of hybrid GPU/CPU linear algebra solvers for NVIDIA GPUs. CULA launched at GTC2009 and has since received large speedups and many new features. We will cover all the features, old and new, along with performance, inner workings, and how users can integrate CULA into their applications. Learn how your existing linear algebra applications can benefit from a high quality library. Much more information is available at www.culatools.com and at our presentation and booth.

Speaker: John Humphrey, EM Photonics, Inc

Topics: High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

Time: Thursday, September, 23rd, 15:00 - 15:50

2070 - CUSPARSE Library: A Set of Basic Linear Algebra Subroutines for Sparse Matrices

The CUSPARSE library can impact and enable software solutions for computational science and engineering problems in the fields of energy exploration, physical simulations and life sciences among many others. It provides sparse linear algebra primitives that can be used to implement iterative linear system and eigenvalue solvers and can also serve as a building block for the state-of-the-art sparse direct solvers. CUSPARSE library is implemented using CUDA parallel programming model and provides sparse analogs to BLAS level-1,2,3 operations, such as matrix-vector multiplication, triangular solve and format conversion routines.

Speaker: Maxim Naumov, NVIDIA

Topics: Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing

Time: Thursday, September, 23rd, 11:00 - 11:50

2142 - Complex Geophysical Imaging Algorithms Enabled by GPU technology

Learn how computational expensive geophysical methods with 100s of TB of data become a commercial reality through the adoption of GPUs. The first part of the talk will give an overview of the computational challenges for imaging facing the oil and gas industry. The second part will show how the current most advanced methods are taking advantage of the GPU technology.

Speaker: David Nichols, Schlumberger

Topics: Energy Exploration, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 14:00 - 14:50

2040 - Derivatives & Bond Portfolio Valuation in a Hybrid CPU/GPU Environment

Learn how to compute traditional end of day computations in real time through the use of a hybrid GPU/CPU computing environment. We will detail how computing intensive tasks are delegated to the GPU while interface issues are dealt with by the CPU. We will discuss our methodology consisting of the following three components: (1) valuations; (2) by tenor risk measures; and (3) full distributions allowing for more complex analytics such as exotic options products valuation and counterparty value adjustments calculation.

Speaker: Peter Decrem, Quantifi

Topics: Finance, Algorithms & Numerical Techniques, High Performance Computing

Time: Thursday, September, 23rd, 14:00 - 14:50

2090 - Developing Highly Scalable Particle-Mesh Codes for GPUs: A Generic Approach

Dive deep into a multi-parallel Particle in Cell code that utilizes MPI, pthreads, and CUDA. Around this specific application a general C++ framework for transparent data transfers between GPUs has been developed and will be presented. Further techniques employed include interleaving of communication and computation, particle tiling and a study of how well CUDA performance can be transferred to OpenCL.

Speakers: Guido Juckeland, TU Dresden - ZIH, Michael Bussmann, Forschungszentrum Dresden-Rossendorf

Topics: Physics Simulation, Astronomy & Astrophysics, High Performance Computing

Time: Tuesday, September, 21st, 15:00 - 15:50

2231 - Driving on Mars, Redux: System Level Simulation of Dynamic Systems

Learn how GPU and HPC computing are used to predict through simulation the dynamics of large complex mechanical systems such as tracked vehicles including the Mars Rover. The presentation outlines the physics based approach and numerical solution methods that enabled the simulation of dynamic systems with millions of bodies on the GPU. The presentation will also explain how a HPC cluster is used to effectively render scenes with tens of thousands of bodies for generating animations that can be used by Engineers in the design process.

Speaker: Dan Negrut, University of Wisconsin

Topics: Physics Simulation, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 10:00 - 10:50

2176 - Easy GPU Meta-programming: A Case Study in Biologically-Inspired Computer Vision

Learn how to let the computer optimize your CUDA and OpenCL code for you with easy GPU Meta-programming and Scripting (e.g. PyCUDA). We will present a case study in which we consider the step-wise optimization of a 3D filter bank convolution, using a suite of open-source tools.

Speaker: Nicolas Pinto, MIT

Topics: Tools & Libraries, Computer Vision, High Performance Computing, Neuroscience

Time: Thursday, September, 23rd, 10:00 - 10:50

4005 - Emerging Companies: CEO on Stage featuring Empulse, Rocketick, and Softkinetic

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature Jedox Business Intelligence, Rocketick, and Softkinetic - covering the fields of computer vision, data mining and high performance computing.

Panelists will include Drew Lanza (Partner, Morgenthaler), Charles Carmel (Cisco), and Nathan Brookwood (Principal Analyst, Insight64).

Speakers: Drew Lanza, Morgenthaler, Uri Tal, Rocketick, Charles Carmel, Cisco, Nathan Brookwood, Insight 64, Kristian Raue, Jedox AG, Michel Tombroff, Softkinetic

Topics: General Interest, Computer Vision, Databases & Data Mining, High Performance Computing

Time: Wednesday, September, 22nd, 17:00 - 17:50

2138 - Faster, Cheaper, Better – Hybridization of Linear Algebra for GPUs

Learn how to develop faster, cheaper and better linear algebra software for GPUs through a hybridization methodology that is built on (1) Representing linear algebra algorithms as directed acyclic graphs where nodes correspond to tasks and edges to dependencies among them, and (2) Scheduling the execution of the tasks over hybrid architectures of GPUs and multicore. Examples will be given using MAGMA, a new generation of linear algebra libraries that extends the sequential LAPACK-style algorithms to the highly parallel GPU and multicore heterogeneous architectures.

Speakers: Stan Tomov, Hatem Ltaief, UNIVERSITY OF TENNESSEE

Topics: High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

Time: Thursday, September, 23rd, 09:00 - 9:50

2007 - Folding@home: Petaflops on the Cheap Today; Exaflops Soon?

Learn how Folding@home has used petascale computing with GPUs to make fundamental breakthroughs in computational biology and how this technology can make an impact in your work.

Speaker: Vijay Pande, Stanford University

Topics: Life Sciences, Cloud Computing, High Performance Computing, Molecular Dynamics

Time: Thursday, September, 23rd, 11:00 - 11:50

2086 - GPGPU DL POLY

Discover DL_POLY.

1. DL_POLY: an MD code ICHEC has ported to CUDA. The presentation especially focuses on the auto-tuning of the work distribution between CPU and GPU

Speaker: Gilles Civario, ICHEC

Topics: Molecular Dynamics, High Performance Computing

Time: Thursday, September, 23rd, 16:00 - 16:50

2147 - GPGPU Development for Windows HPC Server

Attend this demo-driven session to see how to schedule jobs to a Windows compute cluster that includes GPUs. We will also demonstrate GPU-enhanced versions of some commonly used HPC open-source codes, and show how NVIDIA Parallel Nsight™ can be used to debug GPU applications on a cluster. Provides a brief introduction to performance profiling tools that allow developers to analyze system, CPU and GPU events.

Speaker: Calvin Clark, Microsoft

Topics: High Performance Computing, Tools & Libraries

Time: Tuesday, September, 21st, 15:00 - 15:50

2155 - GPGPU in the real world. The ABAQUS experience

We describe the ABAQUS experience in integrating GPGPU acceleration into a complex, high performance commercial engineering software. In particular we discuss the trade-off we had to make and the benefits we obtained from this technology.

Speaker: Luis Crivelli, Dassault Systems Simulia Corporation

Topics: Physics Simulation, Algorithms & Numerical Techniques, Computational Fluid Dynamics, High Performance Computing

Time: Thursday, September, 23rd, 14:00 - 14:20

2179 - GPU - An R Library for Native GPU Objects

Come learn about the GPU R package. R is the widely popular open source statistical programming language. The GPU package extends R by providing GPU-based types, classes and methods implementing GPU versions of R vectors, matrices, lists and data frames. Subsequent operations with these are executed on the GPU. Users are not required to create special bindings or implement special syntax, nor do they need copy objects between CPU and GPU. The GPU packages allows programmers access to the computational power of GPUs with little modification to existing code.

Speaker: Christopher Brown, Open Data

Topics: Tools & Libraries, Algorithms & Numerical Techniques, High Performance Computing

Time: Tuesday, September, 21st, 16:00 - 16:20

2083 - GPU Accelerated Solver for the 3D Two-phase Incompressible Navier-Stokes Equations

This demonstrates the potential of GPUs for solving complex free surface flow problems using level set methods. These methods are capable of producing complex surface deformations, and therefore are used widely in computer graphics, as well as engineering applications. This work demonstrates that GPUs can be used to accelerate the most computationally expensive part of free surface flow calculations, and therefore allows much larger problems to be solved on workstation machines than was previously possible. These techniques will be exemplified by our current project to port our in-house fluid solver NaSt3DGPF to the GPU.

Speaker: Peter Zaspel, University of Bonn

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Physics Simulation

Time: Wednesday, September, 22nd, 16:00 - 16:50

2069 - GPU-Accelerated Business Intelligence Analytics

Join us and learn why GPU computing is a game changer for business intelligence (BI). We will discuss how GPUs can be used to accelerate BI analytics at much lower cost, higher performance, and better power efficiency than other alternatives.

Speaker: Ren Wu, HP Labs

Topics: Databases & Data Mining, Finance, High Performance Computing

Time: Wednesday, September, 22nd, 16:00 - 16:50

2217 - GPU-Based Conjugate Gradient Solvers for Lattice QCD

Learn how to perform state-of-the-art quantum chromodynamics (QCD) computation using NVIDIA GPUs at 1% of the cost of a conventional supercomputer and 10% of its power consumption. We will discuss how physicists around the world are using GPU clusters to solve QCD. We will focus upon how TWQCD have been using a large GPU cluster (200 GPUs) to simulate QCD, attaining 36 Teraflops (sustained).

Speaker: Ting-Wai Chiu, National Taiwan University

Topics: High Performance Computing, Physics Simulation

Time: Wednesday, September, 22nd, 16:00 - 16:50

2027 - GPU-Based Image Processing in Military Applications

There are more than 6000 Unmanned Aerial Vehicles (UAVs) in use in the US Military. The US Army alone has flown more than 1 million UAV flight hours. Every UAV captures at least one stream of video; some as many as 9. All this video needs to be processed and analyzed both during the mission, and post-mission. Traditionally, custom ASICs, and FPGAs were required for even the most rudimentary image processing tasks. Now, GPUs provide orders of magnitude more compute at a fraction of the cost. Hear how MotionDSP uses GPUs to provide previously impossible capabilities to military imaging.

Speaker: Sean Varah, MotionDSP Inc.

Topics: Video Processing, High Performance Computing

Time: Thursday, September, 23rd, 09:00 - 9:50

2282 - GPU-Enabled Biomedical Imaging

The purpose of this presentation is to describe several novel biomedical imaging applications which make extensive use of GPUs. In CT iterative reconstructions, for example, high performance computing is allowing us to see details and structures we previously were not able to discern.

Speaker: Homer Pien, MGH / HMS

Topics: Medical Imaging & Visualization, High Performance Computing, Imaging, Life Sciences

Time: Wednesday, September, 22nd, 17:00 - 17:50

2273 - GPUs In the Front Line of our Defenses (Sponsored by GE)

Find out how GPUs are accelerating defense and aerospace applications and providing superior information processing to drive the next generation of capabilities to protect both homelands and soldiers. Learn how rugged VPX hardware and software architectures are able to scale from small power- & weight-constrained vehicles through to large complex processing arrays, on platforms as diverse as unmanned aerial vehicles (UAV), through tracked ground vehicles, and to ship borne radar.

Speaker: Simon Collins, GE Intelligent Platforms

Topic: High Performance Computing

Time: Wednesday, September, 22nd, 15:00 - 15:50

2044 - GRASSY: Leveraging GPU Texture Units for Asteroseismic Data Analysis

Learn how to use the hidden computation capability of GPU texture units for general purpose computation. We describe GRASSY, a system for stellar spectral synthesis where the core problem is interpolation between pre-computed intensity value. We map these pre-computed tables to the GPU's texture memory. Interpolation then becomes a texture lookup where the hardware automatically performs the interpolation, albeit at very low precision. Our mathematical framework reasons about the impact of this precision and our performance results show 500X speedups. This work generalizes the GPU texture units as computation engines and opens up new problems for GPU acceleration.

Speaker: Matt Sinclair, UW-Madison

Topics: Astronomy & Astrophysics, High Performance Computing

Time: Wednesday, September, 22nd, 15:00 - 15:50

2272 - GStream: A General-Purpose Data Streaming Framework on GPUs

We present GStream, a general-purpose, scalable and C++ template run-time framework amenable to both the streaming problem and GPU architectures. GStream offers transparent streaming data transmissions and automatic memory synchronization over a rich collection of computing resources that are transparently allocated and reused.

Various problems other than streaming application, such as scientific computing, numerical codes and text processing, can be easily expressed using GStream and subsequently integrated with our GStream library. GStream's ease of use combined with efficient exploitation of GPU resources have the potential to lead to higher coding productivity and application performance through our data-centric specification paradigm.

Speakers: Xing Wu, North Carolina State University, Frank Mueller, North Carolina State University

Topics: Tools & Libraries, High Performance Computing

Time: Thursday, September, 23rd, 09:00 - 9:50

2062 - HOOMD-blue: Fast and Flexible Many-Particle Dynamics

See the newest capabilities and performance enhancements in HOOMD-blue, a general-purpose many-particle dynamics application written for GPUs. Speedups of 80-100x are attained for a wide range of simulation types. Topics for this presentation include an overview of HOOMD-blue, design and implementation details of the underlying algorithms, and a discussion on how generality is maintained without sacrificing performance.

Speaker: Joshua Anderson, University of Michigan

Topics: Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

Time: Thursday, September, 23rd, 15:00 - 15:50

2073 - High Performance Molecular Simulation, Visualization, and Analysis on GPUs

This talk will present recent successes in the use of GPUs to accelerate interactive visualization and analysis tasks on desktop computers, and batch-mode simulation and analysis jobs on GPU-accelerated HPC clusters. We'll present Fermi-specific algorithms and optimizations and compare with those for other devices. We'll also present performance and performance/watt results for NAMD molecular dynamics simulations and VMD analysis calculations on GPU clusters, and conclude with a discussion of ongoing work and future opportunities for GPU acceleration, particularly as applied to the analysis of petascale simulations of large biomolecular complexes and long simulation timescales.

Speaker: John Stone, University of Illinois at Urbana-Champaign

Topics: Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

Time: Wednesday, September, 22nd, 16:00 - 16:50

2100 - Hybrid GPU/Multicore Solutions for Large Linear Algebra Problems

Large linear algebra problems may be solved using recursive block decomposition in which GPUs efficiently compute the sub-blocks and multicore CPUs put the sub-blocks back together within a large shared memory space. This talk will present benchmark results for such a hybrid approach, implemented in Matlab® and using Jacket® to access the GPU compute power.

Speaker: Nolan Davis, SAIC

Topics: High Performance Computing, Algorithms & Numerical Techniques, Signal processing

Time: Thursday, September, 23rd, 16:00 - 16:50

2059 - Industrial Seismic Imaging on GPUs

At Hess Corporation, we have moved the most computationally intensive parts of our seismic imaging codes from CPUs to GPUs over the past few years. In this talk I will give an overview of seismic imaging, highlighting the physical and computational algorithms of these codes. I will discuss our software approach and the programming effort to port them to GPUs, concluding with a summary of our progress in adopting GPUs in production.

Speaker: Scott Morton, Hess Corporation

Topics: Energy Exploration, High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:50

2092 - Integrating CUDA into a Large-Scale Commercial Database Management System

In a large-scale database installation where data tables are distributed across multiple servers, computational throughput can be optimized by using GPUs on each server and integrating database management with GPU resources. In the Department of Physics and Astronomy at The Johns Hopkins University, we are experimenting with a set of software tools that closely couple SQL statements with GPU functionality. While still under development, the new framework is now routinely used in our research projects, e.g., to study the spatial clustering of galaxies as well as genomics.

Speakers: Richard Wilton, The Johns Hopkins University, Tamas Budavari, Johns Hopkins University, Alex Szalay, The Johns Hopkins University

Topics: Databases & Data Mining, Astronomy & Astrophysics, High Performance Computing, Tools & Libraries

Time: Wednesday, September, 22nd, 11:00 - 11:50

2033 - Integrating GPGPU Accelerated Pricing Models into an Existing Financial Services Infrastructure

Join Citadel Investment Group to explore our three year undertaking on the feasibility of GPGPU computing for option pricing. We will discuss our 140X performance boost and the hurdles we had to overcome to integrate GPUs into our existing infrastructure. Please note that our talk will not get into the details of the model (that's proprietary information), but we will share our innovative solution to drive a grid of GPUs.

Speaker: Scott Donovan, Citadel Investment Group

Topics: Finance, High Performance Computing

Time: Thursday, September, 23rd, 09:00 - 9:50

2287 - Internal GPUs on Dedicated x16 Slots - Are They Needed For HPC? (Sponsored by Dell)

We have benchmarked the real performance impact on a series of GPU accelerated applications to understand the benefits and drawbacks of different system level configurations. Come hear about the effects on performance of GPUs in shared slots and of GPUs that are externally connected.

Speaker: Mark Fernandez, Dell

Topic: High Performance Computing

Time: Wednesday, September, 22nd, 14:00 - 14:50

2261 - Introduction to GPU Ray Tracing with NVIDIA OptiX (Pre-Conference Tutorial)

Learn how to use NVIDIA OptiX to quickly develop high performance ray tracing applications for interactive rendering, offline rendering, or scientific visualization. This session will explore the latest available OptiX version.

Speakers: Dave McAllister, NVIDIA, Steve Parker, NVIDIA, Phillip Miller, NVIDIA

Topics: Ray Tracing, High Performance Computing, Computer Graphics

Time: Monday, September, 20th, 14:30 - 15:50

2144 - Large-Scale Visualization Using A GPU Cluster

Learn how to visualize extremely large-scale scientific data using GPGPU techniques on a GPU-accelerated visualization cluster. Recent advances in general-purpose GPU (GPGPU) computing provide a promising solution to compute-intensive scientific visualization. However, the largest scientific simulations produce datasets that are orders of magnitude larger than the memory available on current GPUs. Many distributed GPUs must be used in parallel. We present Longhorn, currently the world's largest GPU-enhanced cluster dedicated for visualization and data analysis, and describe the distributed memory architecture and GPGPU techniques to interactively visualize massive datasets using distributed GPUs on Longhorn.

Speakers: Byungil Jeong, TACC / UT-Austin, Paul Navratil, Texas Advanced Computing Center

Topics: Medical Imaging & Visualization, High Performance Computing

Time: Wednesday, September, 22nd, 16:00 - 16:50

2017 - Lessons Learned Deploying the World's First GPU-Based Petaflop System

Learn what to expect when deploying PetaFLOP or larger systems. The June 2010 list of the Top 500 computer systems featured the first GPU based cluster to exceed 1 PetaFLOP of floating point power -- a system that was built in a fraction of the time and the cost a CPU-only system

of that performance would have required. An overview of how system builders and administrators should prepare for large-scale HPC deployments.

Speaker: Dale Southard, NVIDIA

Topic: High Performance Computing

Time: Tuesday, September, 21st, 15:00 - 15:50

2117 - Migration of C and Fortran Apps to GPGPU using HMPP

GPGPU is a tremendous opportunity to many application fields. Migrating legacy software to GPGPU is a complex process that requires mastering the technological risks (e.g. loss of code portability, extensive code restructuration, debugging complexity) as well as costs. In this talk, we present a methodology based on HMPP (Heterogeneous Multicore Parallel Programming), allowing incremental processes that reduce the cost and risks of porting codes to GPGPU.

Speaker: Francois Bodin, CAPS entreprise

Topics: High Performance Computing, Tools & Libraries

Time: Wednesday, September, 22nd, 11:30 - 11:50

2211 - Modern Architecture for Massively Parallel Medical Tomographic Image Reconstruction on a GPU Cluster

Learn how to combine GPU and Cluster Programming with a real-world example. Many aspects of medical tomographic image reconstruction are embarrassingly parallel, but require massive compute power. We distribute the load onto a cluster of multi-GPU equipped nodes using Message Passing Interface (MPI) and CUDA. The Thrust library allows for a modern object-oriented approach.

Speakers: Sven Prevhal, Philips, Jingyu Cui, Stanford University

Topics: Medical Imaging & Visualization, Algorithms & Numerical Techniques, High Performance Computing, Tools & Libraries

Time: Wednesday, September, 22nd, 15:00 - 15:50

2081 - Morphing a GPU into a Network Processor

Modern Internet routers must meet two conflicting objectives, high performance and good programmability, to satisfy the ever-increasing bandwidth requirements under fast changing network protocols. A few recent works prove that GPUs have great potential to serve as the packet processing engine for software routers. However, current GPU's batched execution model cannot guarantee quality-of-service (QoS) requirement. In this work, we show how to convert a GPU into an effective packet processor through minimal changes in both hardware architecture and scheduling mechanism. Experimental results proved that the new GPU architecture could meet stringent QoS requirements, but maintain a high processing throughput.

Speaker: Yangdong Deng, Tsinghua University

Topics: General Interest, High Performance Computing

Time: Thursday, September, 23rd, 15:00 - 15:50

2141 - Moving the Frontier of Oil and Gas Exploration and Production with GPUs

Learn how the Oil and Gas Industry is embracing GPUs in order to tackle new and complex oil and gas plays around the world. The first part of this talk gives an overview of the business and geopolitical drivers of the industry, followed with the critical contribution of computation in the quest for secure supply of energy.

Speakers: Maurice Nessim, Schlumberger, Shashi Menon, Schlumberger

Topics: Energy Exploration, High Performance Computing

Time: Wednesday, September, 22nd, 10:00 - 10:50

2054 - NAMD, CUDA, and Clusters: Taking GPU Molecular Dynamics Beyond the Desktop

A supercomputer is only as fast as its weakest link. The highly parallel molecular dynamics code NAMD was one of the first codes to run on a GPU cluster when G80 and CUDA were introduced in 2007. Now, after three short years, the Fermi architecture opens the possibility of new algorithms, simpler code, and easier optimization. Come learn the opportunities and pitfalls of taking GPU computing to the petascale.

Speaker: James Phillips, University of Illinois

Topics: Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

Time: Thursday, September, 23rd, 14:00 - 14:50

2094 - Nearly Instantaneous Reconstruction for MRIs

GE's Autocalibrating Reconstruction for Cartesian Imaging (ARC) is a computationally intensive, widely used algorithm in MRI Reconstruction using Parallel Imaging. We demonstrate that an optimized CUDA implementation of ARC on a GPU can enable nearly instantaneous reconstruction and speedups of up to 10x over an optimized dual socket QuadCore CPU implementation. We will discuss challenges both with computational intensity and data read/write efficiency. We will also compare the Fermi C2050 with the C1060.

Speaker: Babu Narayanan, GE Global Research

Topics: Medical Imaging & Visualization, High Performance Computing

Time: Tuesday, September, 21st, 14:00 - 14:50

2008 - OpenCL Optimization

Learn how to optimize your OpenCL application to achieve maximum performance on NVIDIA GPUs. We will first briefly discuss how the OpenCL programming model maps onto NVIDIA GPU's architecture. We will then talk about memory, instruction, and NDRange optimization techniques, illustrating each with small code samples.

Speaker: Peng Wang, NVIDIA

Topics: Tools & Libraries, High Performance Computing

Time: Thursday, September, 23rd, 17:00 - 17:50

2248 - Parallel Processing on GPUs at the University of Utah

The University of Utah is a CUDA Center of Excellence. We have been doing both basic and applied research using CUDA. In this session, we plan to give 3-4 talks on ongoing research. Most of the work that we will be presenting has been peer reviewed at top conferences.

Speakers: Claudio Silva, University of Utah, Huy Vo, University of Utah

Topics: High Performance Computing, Life Sciences, Medical Imaging & Visualization, Tools & Libraries

Time: Wednesday, September, 22nd, 14:00 - 14:50

2106 - Particleworks: Particle-based CAE Software on Multi-GPU

Prometech Software, Inc. is an university launched technology venture in Japan and has been working in the field of particle-based computational fluid dynamics for several years. Through collaboratinos with major automotive and material companies in Japan, Prometech has implemented our Particle technology on Multi-GPU and delivered as a CAE software, "Particleworks".

In this session, we will discuss the theoretical background of our simulation (MPS; Moving Particle Simulation method), Multi GPU programming techniques of sparse matrix solver, performance results of Particleworks and the analysis examples of the Auto and Material.

Speaker: Issei Masaie, Prometech Software, Inc.

Topics: Computational Fluid Dynamics, High Performance Computing

Time: Thursday, September, 23rd, 11:30 - 11:50

2052 - Power Management Techniques for Heterogeneous Exascale Computing

Power consumption has become the leading design constraint for large scale computing systems. In order to achieve exascale computing, system energy efficiency must be improved significantly. Our approach will focus on investigating software methodologies to achieve energy efficient computing on heterogeneous systems accelerated with GPUs.

Speaker: Xiaohui Cui, Oak Ridge National Laboratory

Topic: High Performance Computing

Time: Tuesday, September, 21st, 16:00 - 16:50

2135 - Processing Petabytes per Second at the Large Hadron Collider at CERN

Learn how GPUs could be adopted by the ATLAS detector at the Large Hadron Collider (LHC) at CERN. The detector, located at one of the collision points, must trigger on unprecedented data acquisition rates (PB/s), to decide whether to record the event, or lose it forever. In the beginning, we introduce the ATLAS experiment and the computational challenges it faces. The second part will focus on how GPUs can be used for algorithm acceleration - using two critical

algorithms as exemplars. Finally, we will outline how GPGPU acceleration could be exploited and incorporated into the future ATLAS computing framework.

Speakers: Philip Clark, University of Edinburgh, Andy Washbrook, University of Edinburgh

Topics: High Performance Computing, Algorithms & Numerical Techniques, Physics Simulation

Time: Wednesday, September, 22nd, 16:00 - 16:50

2104 - Rapid Prototyping Using Thrust: Saving Lives with High Performance Dosimetry

Radiation poisoning is an everpresent danger for intervention teams that must visit nuclear sites. Virtual reality can help teams prepare for intervention, but efficient computation of radiation dosage is critical to study complex scenarios. Radiation protection research often uses codes based on the straight line attenuation method. As with other approaches, geometrical computations (finding all the interactions radiation rays/objects intersection) remain the simulation bottleneck. This talk will describe how we have used the Thrust high-level library for CUDA C/C++ to quickly prototype innovative algorithms and achieve a significant speed up.

Speaker: Lancelot Perrotte, CEA

Topics: High Performance Computing, Algorithms & Numerical Techniques, Physics Simulation, Ray Tracing

Time: Wednesday, September, 22nd, 11:00 - 11:50

2247 - Reconfiguring a Pool of GPUs on The Fly (Sponsored by NextIO)

The next generation of GPU appliances has arrived. See a live demonstration of the world's first modular GPU appliance capable of housing up to 8 Tesla GPU's supporting up to 4 server nodes. NextIO will demonstrate "drag and drop" hot plug reassignment of GPU's to a server without bringing down GPU applications on that server. The server does not have to be rebooted, the GPU applications do not stop and the server enclosure does not have to be opened in order to add or remove GPU resources to a server node. The appliance provides a reconfigurable pool of GPU resources.

Speaker: K.C. Murphy, NextIO

Topics: High Performance Computing, Cloud Computing

Time: Tuesday, September, 21st, 16:00 - 16:50

2218 - Redesigning Molecular Dynamics for GPUs and GPU Clusters

Generalized Born and Particle Mesh Ewald (PME) molecular dynamics are two computationally intensive algorithms for simulating biological molecules. While several adaptations of Generalized Born have attained excellent speedup on GPUs, high performance Particle Mesh Ewald has been more elusive. Here we describe in detail a recent port of PME implemented within AMBER 11 that has achieved performance on par with up to 128 nodes of a top ten supercomputer.

Speaker: Scott Le Grand, NVIDIA

Topics: Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

Time: Wednesday, September, 22nd, 15:00 - 15:50

2034 - Reformulating Algorithms for the GPU

Important applications in signal, data processing and bioinformatics that use dynamic programming are difficult to parallelize due to intrinsic data dependencies. We demonstrate a novel technique to extract parallelism out of data dependent algorithms and reformulate the same for GPUs.

This simple technique breaks the dependencies and resolves them at an optimal point later in time, thus obtaining remarkable speedup on GPUs. We present a case study from computational biology i.e., protein motif-finding. We also present how the same technique can be extended and applied to other relevant problems such as gene-prediction and phylogenetics.

Speakers: Narayan Ganesan, University of Delaware, Michela Taufer, University of Delaware

Topics: Life Sciences, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:20

2174 - Reverse Time Migration on GPUs

Learn how GPUs can be used to accelerate subsurface imaging for Oil & Gas exploration. We will discuss results and lessons learned while implementing a Reverse Time Migration algorithm on GPUs achieving significant performance improvements over a comparable CPU implementation.

Speaker: Alex Loddoch, Chevron

Topics: Energy Exploration, High Performance Computing

Time: Wednesday, September, 22nd, 15:00 - 15:50

2226 - Reverse Time Migration with GMAC

Get a close look at implementing Reverse Time Migration (RTM) applications across multiple GPUs. We will focus on how RTM applications can be scaled using the GMAC asymmetric distributed shared memory (ADSM) library to break the problem into manageable chunks. We will provide an introduction to GMAC and discuss handling boundary conditions and using separate kernels to improve efficiency.

Speakers: Javier Cabezas, Barcelona Supercomputing Center, Mauricio Araya, Barcelona Supercomputing Center

Topics: Energy Exploration, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 16:00 - 16:50

2293 - Scaling Up and Scaling Out GPUs with Supermicro's Twin™ Architecture (Sponsored by Supermicro)

Find out how Supermicro scales up and scales out GPU performance by using Twin™ architecture. In this session, we outline Supermicro's Twin™ architecture advantages across 1U/2U GPU servers and the design of personal supercomputer, and how we are able to scale and optimize GPU technology for datacenter environment and for professional workstation.

Speaker: Don Clegg

Topics: High Performance Computing, Computer Vision

Time: Wednesday, September, 22nd, 11:00 - 11:50

2078 - Shockingly fast and accurate CFD simulations

In the last three years we have demonstrated how GPU accelerated discontinuous Galerkin methods have enabled simulation of time-dependent, electromagnetic scattering from airplanes and helicopters.

In this talk we will discuss how we have extended these techniques to enable GPU accelerated simulation of supersonic airflow as well.

Speaker: Timothy Warburton, Rice University

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:50

2177 - Simplifying Parallel Programming with Domain Specific Languages

Explore a new approach in parallel programming which leverages Domain Specific Languages (DSLs) to simplify programming heterogeneous systems (multi-core processors and GPUs). This approach allows DSL users to take advantage of the power of GPUs without having working knowledge of lower level programming models such as CUDA. Topics will cover the advantages of the DSL approach in parallel programming, and the runtime implementation details with optimizations to have the performance benefits of using GPUs.

Speakers: HyoukJoong Lee, Stanford University, Hassan Chafi, Stanford University

Topics: Tools & Libraries, High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:50

2035 - Simulations of Large Membrane Regions

Learn how to study membrane-bound protein receptors by moving beyond the current state-of-the-art simulations that only consider small patches of physiological membranes. Towards this end, this session presents how to apply large-scale GPU-enabled computations of extended phospholipid bilayer membranes using a GPU code based on the CHARMM force field for MD simulations. Our code enables fast simulations of large membrane regions in NVT and NVE ensembles and includes different methods for the representation of the electrostatic interactions, i.e., reaction force field and Ewald summation (PME) methods. Performance and scientific results for dimyristoylphosphatidylcholine (PC) based lipid bilayers are presented.

Speakers: Michela Taufer, University of Delaware, Narayan Ganesan, University of Delaware, Sandeep Patel, University of Delaware

Topics: Molecular Dynamics, High Performance Computing, Physics Simulation

Time: Wednesday, September, 22nd, 11:30 - 11:50

2233 - Solving Your GPU Computing Needs (Sponsored by HP)

In this session we will go into detail and you will learn about HP's GPU enabled systems, from Workstations to our GPU enabled servers and clusters. You will get the latest information on configurations, options, GPU management and use cases.

Speakers: Dave Korf, HP, Will Wade, HP

Topic: High Performance Computing

Time: Tuesday, September, 21st, 14:00 - 14:50

2160 - StarPU: a Runtime System for Scheduling Tasks

See how StarPU provides task scheduling facilities for a hybrid platform and a powerful data management library that transparently takes care of data across the entire machine. We will discuss the significant performance improvements resulting from its flexible scheduler as well as its ability to mix parallel CPU kernels (eg. written in OpenMP or TBB) with CUDA/OpenCL and MPI.

Speaker: Cedric Augonnet, INRIA

Topics: Tools & Libraries, High Performance Computing

Time: Wednesday, September, 22nd, 10:00 - 10:50

2084 - State of the Art in GPU Data-Parallel Algorithm Primitives

Learn about the importance of optimized data-parallel algorithm primitives as building blocks for efficient real-world applications. Fundamental parallel algorithms like sorting, parallel reduction, and parallel scan are key components in a wide range of applications from video games to serious science. This session will cover the state of the art in data-parallel primitive algorithms for GPUs. Starting with an explanation of the purpose and applications of the algorithms, we will discuss key algorithm design principles, demonstrate current open source algorithm libraries for GPUs (CUDPP and Thrust), describe optimizations using new features in the Fermi architecture, and explore future directions.

Speaker: Mark Harris, NVIDIA

Topics: Algorithms & Numerical Techniques, High Performance Computing, Tools & Libraries

Time: Tuesday, September, 21st, 17:00 - 17:50

2278 - Strategies for Code Encapsulation in GPU Implementations

Code encapsulation is a common technique used to reduce code complexity that a given programmer has to understand. It allows the use of increasingly complex systems of hardware, software, and algorithms to tackle increasingly difficult scientific problems. Unfortunately, code encapsulation is not easily attainable in current GPU environments. We will share our OpenCL development experiences for achieving partial encapsulation in GPU implementations, and discuss best practices in this area.

Speaker: Brian Cole, OpenEye Scientific Software

Topics: Programming Languages & Techniques, High Performance Computing, Life Sciences

Time: Thursday, September, 23rd, 09:00 - 9:20

2119 - Supercomputing for the Masses: Killer-Apps, Parallel Mappings, Scalability and Application Lifespan

Hear the latest on how supercomputing for the masses is changing the world. We will look at some of the one- to three-orders of magnitude faster killer apps and see how they do it. We will discuss specific mapping to GPGPU hardware and techniques for high performance and near-linear scalability both within and across multiple GPGPUs. We will also consider software investment and the decades long longevity of some successful massively parallel investments in multithreaded software, scalability, balance metrics, lack of consensus on programming models, and lifecycle considerations.

Speaker: Robert Farber, PNNL

Topics: High Performance Computing, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence, Physics Simulation

Time: Tuesday, September, 21st, 11:00 - 11:50

2280 - TSUBAME2.0 Experience

Tsubame2.0 is the next-generation multi-petaflops supercomputer that been designed and built at Tokyo Tech, with more than 4000 NVIDIA Fermi GPUs. as a successor to the highly successful Tsubame1. Deep design considerations were made based on experiences on Tsubame1 retrofitted with the previous generation Tesla to maximize the versatility and the competitiveness of the system across considerable number of application domains, as well as accommodating as much strong scaling as possible. This resulted in a totally new custom

system design in collaboration with HP and NEC, rather than a machine with a retrofitted GPUs. The resulting supercomputer hopefully will become a design template of future large-scale GPU systems to come.

Speaker: Satoshi Matsuoka, Tokyo Institute of Technology

Topic: High Performance Computing

Time: Wednesday, September, 22nd, 10:00 - 10:50

2154 - The Impact of Data Movement on GPU Performance

GPU computing has taken the scientific computing landscape by storm, fueled by the massively parallel arithmetic hardware. When coding, researchers rely on best practices that have been developed in the short timespan of GPGPU. This session challenges a widely held belief that transfers to/from the GPU device must be minimized to achieve the best performance by presenting a case study on CULA, our library for dense linear algebra. The topics to be discussed include the relationship between computation and transfer time for synchronous/asynchronous transfers, and impact that data allocations have on memory performance and overall solution time.

Speakers: John Humphrey, EM Photonics, Inc, Daniel Price, EM Photonics, Inc.

Topics: High Performance Computing, Algorithms & Numerical Techniques, Tools & Libraries

Time: Wednesday, September, 22nd, 16:00 - 16:50

2286 - Towards Peta-Scale Green Computation - Applications of the GPU Supercomputers in the Chinese Academy of Sciences (CAS)

China now holds three spots in the June 2010 Top500 list of GPU-based supercomputers, and two of them, using NVIDIA GPUs, are related to CAS. Efficient use of these systems is more important than peak or Linpack performance. This session will cover some of the large-scale multi-GPU applications in CAS, ranging from molecular dynamics below nano-scale to complex flows on meter-scale and porous media on geological scales, as well as fundamental linear algebra and data/image analysis. The idea of keeping high-efficiency and generality of the computation platform by maintaining a consistency among the target physical system, the computational model and algorithm, and the computer hardware will be explained in detail and demonstrated through a number of super-computing applications in the chemical, oil, mining, metallurgical and biological industries.

Wei Ge, Institute of Process Engineering, Chinese Academy of Sciences, Xiaowei Wang,
Speakers: Institute of Process Engineering , Yunquan Zhang, Institute of Software, CAS, Long Wang,
Super Computing Center, Institute of Computer Network Information of CAS

Topic: High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:50

2134 - Ultra High Resolution Displays and Interactive Eyepoint Using CUDA

We'll go over the challenges we have overcome in building 100 million pixel seamless displays. One customer requirement involves interactive changes of the eyepoint as a person moves, relative to the screen, yet the distortions computed are quite non-linear. We discuss our use of a gpu to implement this procedure.

Speaker: Rajeev Surati, Scalable Display Technologies

Topics: Computer Graphics, High Performance Computing, Medical Imaging & Visualization

Time: Wednesday, September, 22nd, 10:00 - 10:50

2234 - Unstructured Finite Volume Code on a Cluster with Multiple GPUs per Node

Explore how a code written to run in parallel using OpenMP and on a single GPU was modified to run across multiple GPUs and nodes on a multi-CPU, multi-GPU cluster installed at the Naval Research Laboratory. We will discuss the performance of this code running in parallel using MPI/OpenMP and MPI/CUDA.

Speakers: Keith Obenschain, Naval Research Lab Code 6440, Andrew Corrigan, Naval Research Laboratory & George Mason University

Topics: Computational Fluid Dynamics, High Performance Computing

Time: Wednesday, September, 22nd, 15:00 - 15:50

2232 - What If You Had a Petabyte of Memory and/or a Petaflop of Compute? (Sponsored by SGI)

We will explore application spaces where GPU compute coupled with very large shared memory architectures and/or petaflops of compute are allowing new science or new business questions to be addressed.

Speaker: Bill Mannel, SGI

Topic: High Performance Computing

Time: To Be Determined