GPU Technology Conference 2010 Sessions on
Algorithms & Numerical Techniques (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.

2079 - A Fast, Scalable High-Order Unstructured Compressible Flow Solver

We will describe a scalable and efficient high-order unstructured compressible flow solver for GPUs. The solver allows the achievement of arbitrary order of accuracy for flows over complex geometries. High-order solvers require more operations per degree of freedom, thus making them highly suitable for massively parallel processors. Preliminary results indicate speed-ups up to 70x with the Tesla C1060 compared to the Intel i7 CPU. Memory access was optimized using shared and texture memory.

Speakers: David M. Williams, Stanford University, Patrice Castonguay, Stanford University

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

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

2058 - A Practical Introduction to Computational Fluid Dynamics on GPUs

Learn step-by-step procedures to write an explicit CFD solver based on final 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

2236 - A Work-Efficient GPU Algorithm for Level Set Segmentation

Explore a novel GPU level set segmentation algorithm that is both work-efficient and step-efficient. Our algorithm has O(logn) step-complexity, in contrast to previous GPU algorithms which have O(n) step-complexity. We apply our algorithm to 3D medical images and we show that in typical clinical scenarios, our algorithm reduces the total number of processed level set field elements by 16x and is 14x faster than previous GPU algorithms with no reduction in segmentation accuracy.

Speaker: Mike Roberts, Hotchkiss Brain Institute, University of Calgary, Canada

Topics: Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision, Computer Graphics

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

2061 - Accelerating Explicit FEM Shock & Blast Simulations

Explicit finite element codes are widely used to simulate the response of structures and mechanical equipment subjected to shock, blast and wave propagation phenomena. High resolution models require run times ranging from a few seconds to a few months are common and hence the payoff from GPU acceleration is tremendous. We describe the acceleration of our commercial finite element code NLFLEX using CUDA. We developed GPU kernels in CUDA based on our production code NLFLEX, for linear elasticity, explosives, elasto-plasticity and large deformation elasticity. We attained order of magnitude (10X) acceleration in single precision and approximately (5X) in double precision mode.

Speaker: Nachiket Gokhale , Weidlinger Associates Inc

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

Time: Thursday, September, 23rd, 10:30 - 10:50

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

2066 - Accelerating System Level Signal Integrity Simulation

Discuss how GPU acceleration for key parts of the ANSYS Nexxim Simulator resulted in significant speedup over multi-core processors. We will cover time consumption and data parallelism exposure considerations, and focus on key areas where GPU acceleration was applied including convolution and Eye rendering.

Speakers: Danil Kirsanov, ANSYS, Ekanathan Palamadai, ANSYS

Topics: Physics Simulation, Algorithms & Numerical Techniques, Signal processing

Time: Thursday, September, 23rd, 16:30 - 16:50

2110 - Acceleration of a Novel Rotorcraft Wake Simulation

Dive deep as we present the details of a new CUDA-based algorithm for accurate rotorcraft wake simulations. We use a vortex particle method, accelerated with a multipole tree algorithm, combined with a traditional grid-based CFD code. This CUDA algorithm can evaluate the velocity and velocity-gradient with an effective throughput approaching 300 billion interactions per second on a C1060. This gives 10x speed-up and 2.5x better accuracy compared to the parallel CPU version.

Speaker: Christopher Stone, Intelligent Light

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques

Time: Thursday, September, 23rd, 10:00 - 10:20

2235 - Advanced Medical Volume Rendering and Segmentation on the GPU

Learn how to speed up your interactive medical visualization pipeline by an order of magnitude and dramatically improve rendering quality at the same time. Leading researchers in medical imaging informatics describe recent advances in volume visualization and interactive segmentation. Emphasis is on the underlying parallel GPU algorithms and acceleration data structures.

Speakers: Mike Roberts, Hotchkiss Brain Institute, University of Calgary, Canada, Eric Penner, Hotchkiss Brain Institute, University of Calgary, Canada

Topics: Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision, Computer Graphics

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 realisitic 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

2108 - Binary Black Holes Simulations using CUDA

Get the latest information on how to evolve binary black holes simulations on GPUs.

Speaker: Abdul Mroue, CITA, Univ. Of Toronto

Topics: Astronomy & Astrophysics, Algorithms & Numerical Techniques, Physics Simulation

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

2082 - CU-LSP: GPU-based Spectral Analysis of Unevenly Sampled Data

Standard FFT algorithms cannot be applied to spectral analysis of unevenly sampled data. Alternative approaches scale as O(N^2), making them an ideal target for harnessing the raw computing power of GPUs. To this end, I have developed CU-LSP, a CUDA spectral analysis code based on the Lomb-Scargle periodogram. Preliminary benchmarking indicates impressive speed-ups, on the order of 400 relative to a single core of a modern CPU. An initial application of CU-LSP will be the analysis of time-series data from planet-search and asteroseismology satellites.

Speaker: Richard Townsend, University of Wisconsin-Madison

Topics: Astronomy & Astrophysics, Algorithms & Numerical Techniques, Signal processing

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

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

2105 - CUDA-FRESCO: An Efficient Algorithm for Mapping Short Reads

Learn about CUDA-FRESCO and how it addresses issues with MUMmerGPU. We will detail how CUDA-FRESCO overcomes MUMmerGPU's problems processing reads with errors or mismatches and delivers additional performance beyond MUMmerGPU's 5-12x speedup with less than 100bp query length.

Speaker: Chun-Yuan Lin, Department of CSIE, Chang Gung University

Topics: Life Sciences, Algorithms & Numerical Techniques, Tools & Libraries

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

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

2049 - Deflated Preconditioned Conjugate Gradient on the GPU

Explore how to use deflation as a second level preconditioning technique to speed up Block Incomplete Cholesky Preconditioned Conjugate Gradient Method. We use it to solve the Pressure correction equation involved in the solution of the Two-Phase Fluid Flow problem. Our implementation reaches speedup factors between 25-30, for more than 260,000 unknowns, when compared to the CPU.

Speakers: Rohit Gupta, Delft University Of Technology, Kees Vuik, Delft University Of Technology

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques

Time: Wednesday, September, 22nd, 14:30 - 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

2167 - Designing a Geoscience Accelerator Library Accessible from High Level Languages

Explore a library for geoscience applications on CUDA and OpenCL platforms. Target applications span atmosphere, ocean, geomorphology and porous media flows. These areas are linked by common numerical techniques encapsulated in our library. We will review the scope of the library, its meta-programming approaches, and its key design attributes. We will also demonstrate its support for multi-GPU parallelism within and across address spaces and provide examples of is use from high level languages including C, Fortran, and Python.

Speakers: Chris Hill, M.I.T, Alan Richardson, M.I.T

Topics: Programming Languages & Techniques, Algorithms & Numerical Techniques, Computational Fluid Dynamics, Tools & Libraries

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

2103 - Development of an Efficient GPU-Accelerated Model for Fully Nonlinear Water Waves

This work is concerned with the development of an efficient high-throughput scalable model for simulation of fully nonlinear water waves (OceanWave3D) applicable to solve and analyze large-scale problems in coastal engineering. The goal can be achieved through algorithm redesign and parallelization of an optimized sequential single-CPU algorithm based on a flexible-order Finite Difference Method. High performance is pursued by utilizing many-core processing in the model focusing on GPUs for acceleration of code execution. This involves combining analytical methods with an algorithm redesign of the current numerical model.

Speaker: Allan Peter Engsig-Karup, Technical University of Denmark

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

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

2046 - Efficient Automatic Speech Recognition on the GPU

Gain insights into implementation techniques for the speech recognition inference process based on the state-of-art weighted finite state transducer methods. We will discuss the capabilities of the GPU for handling large, irregular graph-based models with millions of states and arcs. We will also present solutions for four challenges in the implementation of speech recognition on the GPU and talk about our more than an order of magnitude faster performance on one core of a CPU.

Speaker: Jike Chong, Parasians, LLC

Topics: Machine Learning & Artificial Intelligence, Algorithms & Numerical Techniques, Audio Processing

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

2015 - Efficient Tridiagonal Solvers for ADI methods and Fluid Simulation

Learn about new techniques to efficiently implement the Alternating Direction Implicit method on GPU for large 2D and 3D domains with complex boundaries.

A novel tridiagonal solver for systems with variable sizes and a new hybrid approach will be covered in detail. Comprehensive performance analysis and key Fermi optimizations will be explored.

Various applications of tridiagonal solvers such as 3D direct numerical fluid simulation and a 2D depth-of-field effect for games will be briefly discussed.

Speaker: Nikolai Sakharnykh, NVIDIA

Topics: Algorithms & Numerical Techniques, Computational Fluid Dynamics

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

2021 - Efficient Volume Segmentation on the GPU

Explore a new technique in the detection of common regions in a 2D/3D data array. Connected components along the axes are linked before actual label propagation starts. The algorithm is completely gather-based, which allows for several optimizations in the CUDA C implementation. It enables real-time frame rates for the analysis of typical 2D images and interactive frame rates for the analysis of typical volume data.

Speakers: Allan Rasmusson, University of Aarhus (NVIDIA intern), Gernot Ziegler, NVIDIA

Topics: Algorithms & Numerical Techniques, Computer Vision, Imaging, Medical Imaging & Visualization

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

2098 - Enabling On Demand Value-At-Risk for Financial Markets

Learn how financial market risk managers can increase their ability to preempt exposure limit breaching and tighten risk control to increase investor confidence. Gain insight into the techniques for obtaining high performance Monte-Carlo based market value-at-risk (VaR) estimates over a hierarchy of risk aggregation levels. This session will focus on how the new Fermi platform can be used by financial institutions to enable on-demand estimates of the market VaR, and discuss important software architecture decisions, the benefits of the new GigaThread Engine and Parallel DataCache, as well as the guiding principles for constructing efficient algorithms on GPUs.

Speakers: Matthew Dixon, UC Davis, Jike Chong, Parasians, LLC

Topics: Finance, Algorithms & Numerical Techniques

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

2239 - Fast GPU Preconditioning for Fluid Simulations in Film Production

Explore how a less efficient, but highly parallel algorithm can still be a superior alternative to a sequential CPU method. This talk will present a simple CUDA-based Poisson solver to the

conjugate gradient method designed for solving well-conditioned matrices such as those that arise from the pressure projection stage of a Navier-Stokes fluid solver. In contrast to other active areas of research in this field, we show that a more brute force approach can still significantly out-perform the best CPU alternatives by sacrificing a high convergence rate in place of achieving much faster iterations.

Speaker: Dan Bailey, Double Negative

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques, Film

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

2087 - Fast High-Quality Panorama Stitching

We present a panorama stitching application implemented with CUDA C on the GPU. The image processing pipeline consist of SIFT feature detection and matching and Graphcut image stitching to achieve high-quality results. We demonstrate live panorama creation with a Webcam.

Speaker: Timo Stich, NVIDIA

Topics: Video Processing, Algorithms & Numerical Techniques, Computer Vision, Imaging

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

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

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, Dassualt 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

2020 - GPU-Accelerated Data Expansion for the Marching Cubes Algorithm

Learn how to accelerate marching cubes on the GPU by taking advantage of the GPU's high memory bandwidth and fast on-chip shared memory in a data expansion algorithm that can extract the complete iso-surface mesh from (dynamic) volume data without requiring any data transfers back to the CPU.

Speakers: Gernot Ziegler, NVIDIA, Chris Dyken, SINTEF

Topics: Algorithms & Numerical Techniques, Imaging, Medical Imaging & Visualization

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

2000 - Gravitational N-body Simulations: How Massive Black Holes Interact with Stellar Systems

Astrophysics is a field where super computing is a must to obtain new scientific results. in particular, the study of the interaction among massive black holes and surrounding stars is a hot topic, which requires heavy computations to have good representation of what happens in the inner regions of galaxies. We present the results obtained with our high precisioned N-body code, NBSymple, which exploits the joint power of a multi core CPU system together with the high performance NVIDIA Tesla C1060 GPUs.
The code is available at the website: astrowww.phys.uniroma1.it/dolcetta/nbsymple.html

Speakers: Roberto Capuzzo-Dolcetta, Sapienza Univ. of Roma, Alessandra Mastrobuono Battisti, Sapienza- University of Rome

Topics: Astronomy & Astrophysics, Algorithms & Numerical Techniques

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

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

2030 - High-Throughput Cell Signaling Network Learning with GPUs

Explore how GPUs are being used to enable high-throughput cell signaling network discovery and data-intensive computational systems biology more generally. Systems biology is transitioning from a largely reductive discipline to one focused on building predictive models of large-scale biological systems. New instrumentation will provide the necessary raw data for such an approach, the key challenge now is building the hardware and software tools to efficiently and interactively build these models. This session will describe how GPUs can and will play a key role in these efforts.

Speaker: Michael Linderman, Stanford University

Topics: Life Sciences, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence

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

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

2163 - Leveraging GPUs for Evolutionary Game Theory

Learn how GPUs are being used to accelerate the study of the emergence of cooperative behavior in biology, from the interactions of humans to viruses to bacteria. The work presented here achieves a speedup of 209x on a cluster of 4 Tesla GPUs.

Speaker: Amanda Peters, Harvard University

Topics: Algorithms & Numerical Techniques, Life Sciences

Time: Wednesday, September, 22nd, 10:00 - 10:20

2028 - Mathematica for GPU Programming

Mathematica is widely used in scientific, engineering, mathematical fields and education. In this session, new tools for general GPU programming in the next release of Mathematica are presented. These tools build on top of Mathematica's technology which provides a simple, yet powerful, interface to the large base of compiling tools. Applications of CUDA and OpenCL from within Mathematica will be presented. These examples will provide a general overview of the powerful development environment for GPU programming that Mathematica can offer not just for researchers but for anybody with basic knowledge of Mathematica and GPU programming.

Speaker: Ulises Cervantes-Pimentel, Wolfram Research

Topics: Programming Languages & Techniques, Algorithms & Numerical Techniques, Imaging, Tools & Libraries

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

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 Prevrhal, 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

2115 - Modified Smith-Waterman-Gotoh Algorithm for CUDA Implementation

It is axiomatic that computational throughput can be increased by exploiting the parallelism of GPU hardware –- but what if the computational algorithm is not easy to implement in parallel? We have modified one such algorithm -– the Smith-Waterman-Gotoh dynamic programming algorithm for local sequence alignment -– so as to make it more amenable to data-parallel computation. The result is a successful CUDA implementation that fully exploits GPU parallelism.

Speaker: Richard Wilton, The Johns Hopkins University

Topics: Life Sciences, Algorithms & Numerical Techniques

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

2088 - Nucleotide String Matching Using CUDA-Accelerated Agrep

Dive deep into the intelligent utilization of various CUDA memory spaces to remarkably speedup approximate DNA/RNA nucleotide sequence matching algorithm in bioinformatics by an amazing factor of 67 compared to multi-threaded quad core CPU counterpart. Our talk provides a very good example to demonstrate how to use indexable array to save frequently updated variables directly into GPU registers, how to organize shared memory into a 2D array to avoid bank conflict, and how to shuffle the data structure to satisfy the requirement for coalesced global memory access. Our CUDA implementation employs online approach and can be applied in real time.

Speaker: Hongjian Li, The Chinese University of Hong Kong

Topics: Life Sciences, Algorithms & Numerical Techniques

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

2171 - Parallel Algorithms for Interactive Mechanical CAD

The broad objective of our research is to develop mechanical Computer-Aided Design tools that provide interactive feedback to the designer. We have developed GPU algorithms for fundamental CAD operations (NURBS evaluation, surface-surface intersection, separation distance computation, moment computation, etc.) that are one to two orders of magnitude faster, and often more accurate, than current commercial CPU implementations. We will touch on strategies we have employed to meet GPU programming challenges, such as the separation

of CPU/GPU operations, imposing artificial structure on computations, and transforming problem definitions to suit GPU-computation models.

Speakers: Sara McMains, University of California Berkeley, Adarsh Krishnamurthy, University of California Berkeley

Topics: Algorithms & Numerical Techniques, Tools & Libraries, Computer Graphics

Time: Thursday, September, 23rd, 11:00 - 11:20

2068 - Parallelizing FPGA Technology Mapping using GPUs

FPGA technology mapping is an algorithm that is heavily data parallel, but contains many features that make it unattractive for GPU implementation. The algorithm uses data in irregular ways since it is a graph-based algorithm. It also makes heavy use of constructs like recursion which is not supported by GPU hardware. In this paper, we take a state-of-the-art FPGA technology mapping algorithm within Berkeley's ABC package and attempt to parallelize it on a GPU. We show that runtime gains of 3.1x are achievable while maintaining identical quality as demonstrated by running these netlists through Altera's Quartus II place-and-route tool.

Speaker: Doris Chen, University of Toronto

Topic: Algorithms & Numerical Techniques

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

2005 - Porting Large-Scale Legacy Fortran Codes

Explore a new automatic Fortran translator which has been developed and used to port the numerical subroutines of FEFLO , a general-purpose legacy Computational Fluid Dynamics code operating on unstructured grids, to run on the GPU. Data transfer to the CPU is minimized throughout the course of a CFD run. Benchmarks of large-scale production runs will be presented.

Speakers: Andrew Corrigan, Naval Research Laboratory & George Mason University, Rainald Löhner, George Mason University

Topics: Algorithms & Numerical Techniques, Computational Fluid Dynamics, Tools & Libraries

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

2032 - Practical Methods Beyond Monte Carlo in Finance

Murex will share its practical experience using GPUs to accelerate high-performance analytics based on GPU-enabled Monte Carlo and PDE methods. We will also briefly describe Murex's experience developing a high-level payoff scripting language that allows user-definable payoffs for single and cross-asset instruments.

Speaker: Pierre Spatz, Murex SAS

Topics: Finance, Algorithms & Numerical Techniques

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

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

2136 - Pseudo Random Number Generators for Massively Parallel Apps

Learn how to select the best and fastest pseudo random number generator for your massively parallel Monte Carlo simulation. Pseudo random numbers generators (PRNG) are a fundamental building block of these simulations and it is thus required to select suitable PRNGs with regard to the specific problem at hand while considering the parallel hardware architecture.

Recent developments in random number generations provide a wide variety of choices, each with different properties and trade-offs. We provide a comprehensive survey of the current state of the art for massively parallel PRNG and show a broad range of applications.

Speaker: Holger Dammertz, Ulm University

Topics: Algorithms & Numerical Techniques, Finance

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

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

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

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

2045 - Roe-Pike Scheme for 2D Euler Equations

Hear how we are improving our elsA and CEDRE computational fluid dynamics software by working on solving the Euler equations set on the GPU. We discuss how our implementation considers the associated Riemann problem and the Roe-Pike differencing scheme at several orders in space while also introducing immerse boundary conditions. Covers the significant speedup obtained through algorithmic and computational optimizations.

Speaker: Matthieu Lefebvre, ONERA

Topics: Computational Fluid Dynamics, Algorithms & Numerical Techniques

Time: Wednesday, September, 22nd, 14:00 - 14:20

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

2252 - Simulating Housefly Vision Elements Using OpenCL

An OpenCL GPU based computer simulation of a biologically motivated model, based on the anatomy of housefly's first optic ganglion, the lamina ganglionaris (the lamina layer) is presented. Specific to GPU technology, the computer model demonstrates: the implementation of a 2nd Order Runga-Kutta method to approximate coupled differential equations using GPU hardware; the mapping of a non-Cartesian coordinate system onto the Cartesian layout of the threads. Testing examined usage and access across device memory spaces to determine the optimal usage/access method for the ANN. This result was generalized for OpenCL GPU devices, using the capabilities of OpenCL.

Speaker: Karen Haines, WASP/The University of Western Australia

Topics: Neuroscience, Algorithms & Numerical Techniques, Signal processing

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

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

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

2140 - Superfast Nearest Neighbor Searches Using a Minimal kd-tree

Learn how to adapt a kd-tree spatial data structure for efficient nearest neighbor (NN) searches on a GPU. Although the kd-tree is not a natural fit for GPU implementation, it can still be effective with the right engineering decisions. By bounding the maximum height of the kd-tree, minimizing the memory footprint of data structures, and optimizing the GPU kernel code, multi-core GPU NN searches with tens of thousands to tens of millions of points run 10-40 times faster than the equivalent single-core CPU NN searches.

Speaker: Shawn Brown, UNC, Chapel Hill

Topics: Algorithms & Numerical Techniques, Databases & Data Mining, Machine Learning & Artificial Intelligence

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

2038 - The Best of Both Worlds: Flexible Data Structures for Heterogeneous Computing

Learn how to switch between array of structs (AoS) and struct of arrays (SoA) storage without having to change the data access syntax. A few changes to the struct and container definitions will enable you to evaluate the performance of AoS vs. SoA on your existing AoS code. We present a simple abstraction that retains the more intuitive AoS syntax array[index]component, yet allows you to switch between AoS and SoA storage with a single template parameter at class definition.

Speaker: Robert Strzodka, Max Planck Institut Informatik

Topics: Algorithms & Numerical Techniques, Tools & Libraries

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

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

2166 - The Triad of Extreme Computing-Fast Algorithms, Open Software and Heterogeneous Systems

The first wave of successful GPU accelerations has been crowded with highly-parallel methods that adapted well to the hardware. But the easy-pickings are now running out. The truly challenging applications require "going back to the algorithmic drawing board." To develop new versions of the most effective fast algorithms, such that our science can most benefit, an ideal

environment is created by the open software model, where efforts can be shared. We will describe one area of application --electrostatics of biomolecules in solution-- where we see at work the triad of extreme computing: fast algorithms, open software, and heterogeneous computing.

Speaker: Lorena Barba, Boston University

Topics: Algorithms & Numerical Techniques, Physics Simulation

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

2085 - Tridiagonal Solvers: Auto-Tuning and Optimizations

In this presentation, we will discuss and analyze the performance of three optimization techniques for tridiagonal solvers. We first present a hybrid Parallel Cyclic Reduction(PCR)-Gaussian Elimination(GE) tridiagonal solver, which combines work-efficient and step-efficient algorithms for high performance. We further discuss an auto-tuned variant of this technique which selects the optimal switching point between algorithms on a per-machine basis. Next, we present a technique to handle large systems, where shared memory constraints prohibit previous work to solve these systems directly. Finally, we will discuss optimizations on a cyclic reduction technique that avoid bank conflicts on current hardware.

Speakers: Andrew Davidson, University of California, Davis, Yao Zhang, University of California, Davis

Topics: Algorithms & Numerical Techniques, Computational Fluid Dynamics

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

2172 - Unveiling Cellular & Molecular Events of Cardiac Arrhythmias

George Mason University is using CUDA technology to get a 20x speed-up in simulations of intracellular calcium dynamics, thought to play a major role in the generation of cardiac arrhythmias. We will discuss the novel algorithms we have developed for Markov Chain Monte Carlo Simulation and their use in investigating elementary events of calcium release in the cardiac myocyte. The resulting extremely fast simulation time has generated new insights into how defects in the control of intracellular calcium may lead to cardiac arrhythmia.

Speaker: Tuan Hoang-Trong, George Mason University

Topics: Life Sciences, Algorithms & Numerical Techniques, Physics Simulation

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

2003 - Using CUDA to Accelerate Radar Image Processing

Come see how current GPU technology provides the means for the first portable real-time radar image processing algorithm. This session will outline how the GPU has afforded nearly three orders of magnitude improvement in performance for Synthetic Aperture Radar's (SAR) hallmark image processing algorithm. We will present algorithm details and further improvements.

Speakers: Richard Carande, Neva Ridge Technologies, Aaron Rogan, Neva Ridge Technologies

Topics: Signal processing, Algorithms & Numerical Techniques, Imaging, Video Processing

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

2122 - Using GPUs for Real-Time Brain-Computer Interfaces

Learn how GPU processing can provide researchers with an inexpensive and versatile alternative to dedicated signal processing hardware for real-time neural prosthetics. Topics will include an overview of algorithms, current state-of-the-art hardware, GPU processing in a real-time environment, multi-platform processing, and future directions in BCIs using GPU processing.

Speaker: Adam Wilson, University of Cincinnati

Topics: Neuroscience, Algorithms & Numerical Techniques, Signal processing

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