GPU Technology Conference 2010 Sessions on Physics Simulation (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

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

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

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

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

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

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

4010 - Emerging Companies: CEO on Stage featuring Natural Motion, Optitex, and Useful Progress

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 Natural Motion, Optitex, and Useful Progress - covering the fields of computer graphics, physics simulation, and medical imaging.

Panelists will include Bill Tai (CRV), Paul Weiskopf (Adobe), and Tim Bajarin (Creative Strategies).

Speakers: Bill Tai, Charles River Ventures, Yoram Burg, OptiTex USA Inc., Sylvain Ordureau, UsefulProgress, Paul Weiskopf, Adobe, Tim Bajarin, Creative Strategies

Topics: General Interest, Medical Imaging & Visualization, Physics Simulation, Computer Graphics

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

2102 - Evacuate Now? Faster-than-real-time Shallow Water Simulation on GPUs

Learn how to simulate a half an hour dam break in 27 seconds! We present how shallow water simulation with interactive visualization is successfully mapped to modern graphics hardware. Featuring a live demo, we will present interactive shallow water simulations running on a standard laptop. The implementation has been verified against analytical and experimental data, supports multi-gpu simulation, and can run up-to 6300x6300 domain sizes at 320 million cells per second on the GTX 480.

Speaker: André Rigland Brodtkorb, SINTEF ICT

Topics: Physics Simulation, Computational Fluid Dynamics

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

2214 - Faster Simulations of the National Airspace System

Learn about twenty-four hour, fast-time simulations of traffic in the National Airspace System, which use GPU technology to help perform key steps in the trajectory prediction of flights. GPUs enabled us to improve the runtime by up to two orders of magnitude versus the previously required tens of minutes per execution. We will present a brief overview of the problem domain and a description of how the GPU has opened doors to uncharted research areas.

Speaker: Joseph Rios, NASA

Topic: Physics Simulation

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

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

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

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

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

2128 - Hybrid Quantum Mechanics/Electrodynamics (QM/ED) Modeling of Solar Cells on a CUDA Cluster

One of the greatest challenges of the twenty-first century is the utilization of renewable energy. In providing a theoretical explanation and guidelines for computer-aided design of dye-sensitized solar cell (DSSC), we recently developed a hybrid multi-scale quantum mechanics/classical electrodynamics (QM/ED) methodology.

Our numerical simulations were tested on a CUDA enabled Linux cluster using CP2K. We extended its CUDA implementation to MPI parallel environment. Our preliminary results demonstrated a superior performance advantage of hybrid MPI/GPGPU programming that could potentially shorten the total simulation wall time by an order of magnitude.

Speaker: Hanning Chen, Northwestern University

Topics: Quantum Chemistry, Energy Exploration, Molecular Dynamics, Physics Simulation

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

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

2037 - Numtech & GPGPU, a SME Point of View

Hear why and how Numtech, a french SME working in the field of atmospheric dispersion and expertise of meteorological events, is benchmarking GPGPU for its futures applications. A compressible and an incompressible interactive flow solvers are described.

Speaker: Emmanuel Buisson, Numtech

Topics: Computational Fluid Dynamics, Physics Simulation

Time: Thursday, September, 23rd, 09:30 - 9:50

2101 - Pricing American Options Using GPUs

This presentation focuses on the challenging problem of Pricing High-Dimensional American Options (PHAO) and how GPUs can be involved in this task. On the one hand, we present a method based on Malliavin calculus which is effective for parallel architecture. On the other hand, we
compare this method with Longstaff & Schwartz method which is more dedicated to sequential architecture. We will conclude with some ideas about the parallelization of the former method on a cluster of machines and finally we will discuss this method considering it as a reformulation of a non-linear parabolic problem using BSDEs.

Speaker: Lokman A. Abbas-Turki, Paris-Est University

Topics: Finance, Physics Simulation

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

2041 - PyCUDA: Even Simpler GPU Programming with Python

Explore PyCUDA, a robust, open-source toolkit that lets you control your GPU from the comfort of Python, a Matlab-like scripting language. Learn about Fermi tuning with PyCUDA, the new interfaces for CUBLAS and CUFFT, the ecosystem of third-party libraries built on PyCUDA, and examples illustrating PyCUDA's benefits to large-scale applications.

Speaker: Andreas Kloeckner, Courant Institute, NYU

Topics: Tools & Libraries, Computational Fluid Dynamics, Physics Simulation

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

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

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

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

2080 - Tackling Multi-Gigabit Design Challenges with a Practical Virtual EMI/ESD Lab

rn about efficient methodologies for performant and cost-effective EMI and ESD suppression techniques by means of massive GPU parallel processing for simulations. We will discuss solving ever more complicated EMI and ESD challenges very early in the design process using in a so called ‘Virtual EMI/ESD lab’.

Speakers: Davy Pissoort, KHBO-FMEC, Amolak Badesha, Agilent Technologies, Hany Fahmy, NVIDIA

Topics: Physics Simulation, Tools & Libraries

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

2112 - The Heisenberg Spin Glass Model on GPU: Myth versus Fact

Dive into implementations of the 3D Heisenberg spin glass model for GPUs. We will discuss results showing that fast shared memory gives better performance with respect to slow global memory only under certain conditions. Covers careful kernel tuning to achieve significant speedup with respect to a state-of-art high end multicore processor.

Speaker: Massimo Bernaschi, Istituto Applicazioni del Calcolo - C.N.R.

Topic: Physics Simulation

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

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

2246 - The challenges of integrating CUDA engines into an existing package, yet not sinking the boat

Based on a true story, come listen to a daring tale about the process of integrating a large CUDA component (physical engine) into an existing product (3D engine) replacing some of its functionality. The architectural difficulties and finer points that needed to be addressed. The tuning and testing of such a large system. While not effecting the stability of the original system.

Speaker: Eri Rubin, OptiTex

Topics: Physics Simulation, Tools & Libraries

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

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

2178 - Using GPUs to Track Changes in the Sun

Learn how GPU computing is enabling astrophysicists to study our closest star. NASA's recently launched Solar Dynamics Observatory is continuously streaming full-disk images of the Sun at visible, UV and EUV wavelengths. This presentation will discuss ways that GPU computing is helping scientists cope with the analysis of the immense data volumes as well as in numerical modeling of the Sun.

Speaker: Mark Cheung, Lockheed Martin Solar & Astrophysics Laboratory

Topics: Astronomy & Astrophysics, Computer Vision, Computational Fluid Dynamics, Physics Simulation

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