GPU Technology Conference 2010 Sessions on Computational Fluid Dynamics (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.

2009 - 4D Visualization and Analysis of Flow

4D flow or vector data is now common in CFD simulations as well as acquisition techniques like 4D flow MRI to study abnormal blood flow patterns. We show how by mixing compute and graphics combined with stereo we are now able to interactively analyze and visualize the resulting data to understand abnormal flow patterns. Topics include flow field rendering, computing derived quantities, merging volumetric rendering with computed geometry such as particles and surfaces, and integration 3d vision stereo.

Speaker: Shalini Venkataraman, NVIDIA

Topics: Medical Imaging & Visualization, Computational Fluid Dynamics, Stereoscopic 3D

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

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

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

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

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

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

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

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

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

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

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

2292 - Implementation of High-Order Adaptive CFD Methods on GPUs

A discontinuous high-order formulation named the Correction Procedure via Reconstruction (CPR) is recently implemented on Nvidia GPUs. The CPR formulation is related to the discontinuous Galerkin (DG) method, and unifies several methods such as the DG, spectral volume and spectral difference into a single framework efficient for hybrid meshes. In preliminary 2D inviscid flow computations, a single GPU has been able to deliver a speedup of 44 over a CPU of the same generation. Extension is being made for viscous flow computation, and results will be presented at the final presentation.

Speakers: Z.J. Wang, Iowa State University, Lizandro Solano, Arun Somani

Topic: Computational Fluid Dynamics

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

2295 - Large-scale CFD Applications and a Full GPU Implementation of a Weather Prediction Code on the TSUBAME Supercomputer

Many CFD applications have been successfully accelerated on GPUs, but for large-scale simulations that require memory beyond a single GPU, communication is required between GPUs over cluster nodes through PCI-Express and interconnects. To overcome performance bottlenecks and preserve parallel scalability, an overlapping technique between computation and communication is essential. This work presents results of an LBM for incompressible flow, and a Tsunami simulation solving the shallow water equation for simulations on the NVIDIA Tesla-based TSUBAME supercomputer of Tokyo Tech. In addition results will be presented on a complete GPU implementation of a production-level weather prediction code developed by the JMA that achieves 15 TFLOPS for an 80-fold speedup.

Speaker: Takayuki Aoki, Tokyo Institute of Technology

Topic: Computational Fluid Dynamics

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

2118 - Large-scale Gas Turbine Simulations on GPU Clusters

This talk describes a strategy for implementing structured grid PDE solvers on GPUs. Techniques covered include the use of source-to-source compilation and the use of sparse matrix vector multiplications for complicated boundary conditions. A new production-quality solver for flows in turbomachines called Turbostream that uses these techniques is presented. The impact of the use of GPUs on the turbomachinery design process is demonstrated by two 64-GPU simulations that have recently been performed on the University of Cambridge's GPU cluster.

Speaker: Tobias Brandvik, University of Cambridge

Topic: Computational Fluid Dynamics

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

2170 - Lattice Boltzmann Multi-Phase Simulations in Porous Media using GPUs

Learn how a very efficient implementation of multiphase lattice Boltzmann methods (LBM) based on CUDA delivers significant benefits for predictions of properties in rocks. This simulator on NVIDIA hardware enables us to perform pore scale multi-phase (oil-water-matrix) simulations in natural porous media and to predict important rock properties like absolute

permeability, relative permeabilites, and capillary pressure. We will show videos of these simulations in complex real world porous media and rocks.

Speaker: Jonas Toelke, Ingrain

Topics: Computational Fluid Dynamics, Energy Exploration

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

2109 - Migration of a Complete 3D Poisson Solver from Legacy Fortran to CUDA

We describe our journey of migrating a legacy direct solver library for Poisson equations written in Fortran77 to CUDA in order to harness the computational power provided by the Tesla device ("Fermi"). This legacy library is still widely used today as it is the most complete library that can deal with three different boundary conditions (Dirchlet, Neumann and Cyclic) and two grid configurations (staggered and centered) independently in any of the three dimensions (x, y, z); giving a total of over 200 configurations.

Speaker: Huynh Phung, A*STAR Institute of High Performance Computing

Topics: Tools & Libraries, Computational Fluid Dynamics

Time: Wednesday, September, 22nd, 10:30 - 10:50

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

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

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

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

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

2022 - Solving PDEs on Regular Grids with OpenCurrent

OpenCurrent is an open source library with support for structured 3D grids and various PDE solvers that operate on them, including a multigrid Poisson solver and an incompressible Navier-Stokes solver. It also includes extensions for splitting grids across multiple GPUs. This talk will provide a basic introduction to the code base and its design principles.

Speaker: Jonathan Cohen, NVIDIA Research

Topic: Computational Fluid Dynamics

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

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

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

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

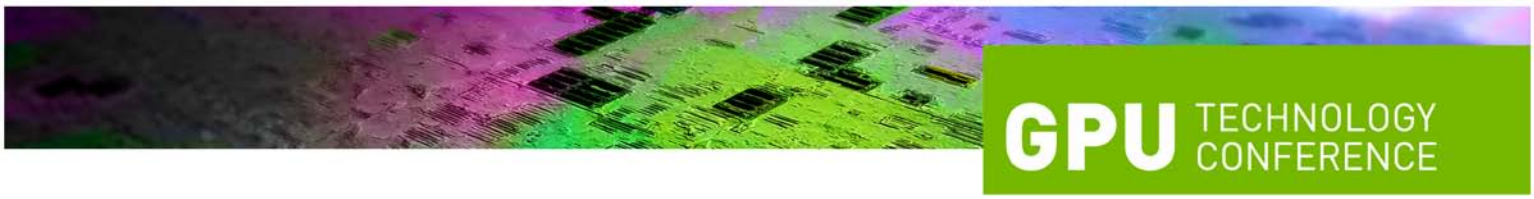
GPU TECHNOLOGY
CONFERENCE