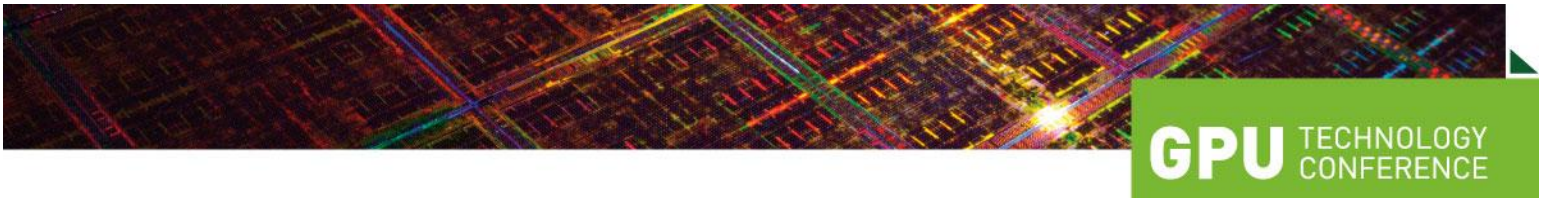

GPU Technology Conference, May 14-17, 2012
McEnery Convention Center, San Jose, California
www.gputechconf.com

Sessions on Computational Fluid Dynamics (subject to change)

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

S0102 - Flame On: Real-Time Fire Simulation for Video Games
Simon Green (NVIDIA), Christopher Horvath (Pixar)
Day: Tuesday, 05/15 | **Time:** 9:00 am - 9:25 am
Topic Areas: Computer Graphics; Computational Fluid Dynamics
Session Level: Intermediate

Fire and explosions are common elements in video games and other virtual environments. We present a real-time fire simulator inspired by the paper "Directable, High-Resolution Simulation of Fire on the GPU" [Horvath and Geiger 2009], but this time implemented entirely in CUDA and targeted at adding interactive fire to video games. This talk will describe both the tricks necessary to implement an efficient fluid simulator in CUDA, and techniques for rendering the results to achieve realistic looking fire.

S0296 - A GPU-Enabled SPH Method for Micro and Nanofluidic Simulations
Daniel Gaudlitz (FluiDyna GmbH)
Day: Tuesday, 05/15 | **Time:** 9:00 am - 9:25 am
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques
Session Level: Intermediate

With SPH methods multi-phase flows within complex geometries can be efficiently investigated. Also physical effects present in micro- and nanofluidic applications are described with little effort using the SPH methodology. In order to investigate microfluidic applications relevant to industry, large domains and high spatial resolutions are required. Therefore, a SPH method for accelerated computations on GPUs is currently developed. The code features dynamic casting of computational data into blocks of appropriate size to fit the GPU memory layout. Also tree-like data structures for efficient manipulation of particle distributions help to obtain significant performance gains on GPU hardware.

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

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

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

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

S0251 - RANS CFD Solver on Fermi
James Lin (Shanghai Jiao Tong University)
Day: Tuesday, 05/15 | **Time:** 10:00 am - 10:25 am
Topic Areas: Computational Fluid Dynamics
Session Level: Intermediate

SJTU-NS3D is an in-house CFD code co-developed by SJTU and COMAC for large civil airplane, solving 3D Reynolds Average Navier-Stokes (RANS) equations on structured grids by finite volume method, which could be used in designing wing model. In this talk, we will present the design and further optimization of CUDA version of SJTU-NS3D, and it achieves 20-fold speedup for standard M6 wing model and 37-fold speedup for wing model candidate from COMAC on single Fermi C2050.

S0300 - Jet: A Domain-Specific Approach to Parallelism for Film Fluid Simulation
Dan Bailey (Double Negative)
Day: Tuesday, 05/15 | **Time:** 10:00 am - 10:25 am
Topic Areas: Parallel Programming Languages & Compilers; Digital Content Creation & Film; Computational Fluid Dynamics
Session Level: Intermediate

Discover how a domain-specific language can not only provide fast parallel performance but a simpler user experience in an environment that highly values flexibility. This talk will present the Jet language and heterogeneous compiler built on the LLVM compiler framework that enables efficient generation of X86 machine code or NVIDIA PTX for stencil computation on structured grids. We show that moving target-specific optimizations upstream into the compiler can greatly improve the ability to manipulate the logic of the solver and thus lower the barrier-to-entry for artists and developers without compromising on performance.

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

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

S0364 - Interacting with Huge Particle Simulations in Maya with the GPU
Wil Braithwaite (NVIDIA)
Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm
Topic Areas: Digital Content Creation & Film; Computational Fluid Dynamics; Visualization
Session Level: Beginner

We present a plug-in for Maya which enables an artist to simulate huge particle counts in real-time by leveraging the NVIDIA GPU. Being able to interact with the simulation opens up new possibilities for modifying the workflow. We will demonstrate the plug-in, and provide insight into the algorithms used.

S0412 - A 2-Petaflops Stencil Application with Stereoscopic 3D Visualization - Gordon Bell Prize 2011
Takayuki Aoki (Tokyo Institute of Technology)
Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:25 pm
Topic Areas: Supercomputing; Computational Fluid Dynamics; Climate & Weather Modeling; Stereoscopic 3D
Session Level: Intermediate

Most stencil applications such as CFD and structure analysis are memory-bound problems. GPU has high performances in both computation and memory bandwidth suitable for them. The TSUBAME 2.0 supercomputer with 4224 GPUs has started since November 2010. We study a metal dendritic solidification by solving the phase-field model. The performance of 2.0 Petaflops was achieved for 4,096x6,500x1,0400 mesh on 4000 GPUs and we received the ACM Gordon Bell Prize in 2011. We also demonstrated several large-scale stencil applications (Lattice Boltzmann, weather prediction and so on) with stereoscopic 3D visualization.

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

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

S0247 - 3D ADI Method for Fluid Simulation on Multiple GPUs
Nikolai Sakharnykh (NVIDIA), Nikolay Markovskiy (NVIDIA)
Day: Tuesday, 05/15 | **Time:** 5:00 pm - 5:50 pm
Topic Areas: Algorithms & Numerical Techniques; Computational Fluid Dynamics
Session Level: Intermediate

Find out about a multiple GPU implementation of the Alternating Direction Implicit method for large 3D domains. The ADI technique is applied towards direct numerical fluid simulation. Modeling complex flows demands extremely large grids and a distributed computation is required for sharing the memory among multiple GPUs. In this session a novel distributed tridiagonal solver as well as parallelization and load balancing strategies will be covered in detail. Finally, a comprehensive performance analysis and scaling studies for different input geometries and possible future improvements will be discussed.

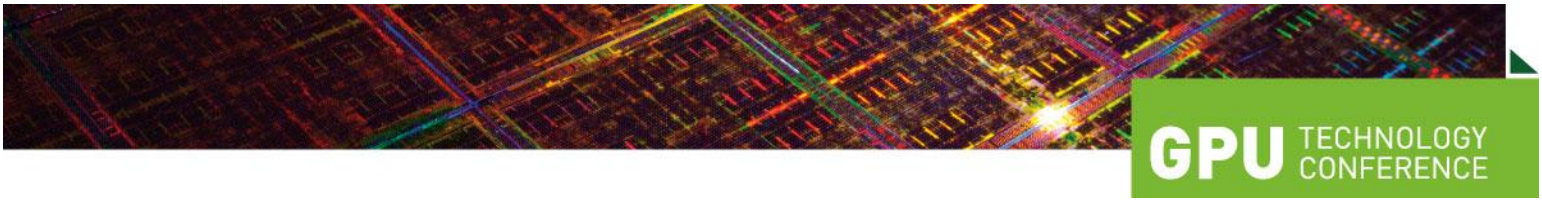

S0066 - Particleworks: Particle-based CAE Software Fully Ported on Multi-GPU
Yoshiaki Hanada (Prometech Software, Inc.), Issei Masaie (Prometech Software, Inc.)
Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am
Topic Areas: Computational Fluid Dynamics
Session Level: Intermediate

Get the latest information on Particle-based fluid simulation + multi-GPU computing as a commercial CAE software named "Particleworks" in Japan. In this session, we provide the information such as (1)Particle simulation trends in CAE, (2)Particle simulation development in Japanese industry, (3)Implementation and performance of full GPU porting and (4)Multi-GPUs scaling with the several clients' cases.

S0304 - Large Scale Computational Fluid Dynamics Simulations on Hybrid Supercomputers
John Humphrey (EM Photonics), Eric Kelmelis (EM Photonics)
Day: Wednesday, 05/16 | **Time:** 10:30 am - 10:55 am
Topic Areas: Computational Fluid Dynamics; Supercomputing
Session Level: Intermediate

Learn how to approach the all-too-common program of trying to retrofit a major application for speed in the modern era of the hybrid supercomputer. In this talk, we will focus on computational fluid dynamics (CFD) codes that are run on Top500 Supercomputers. Many of these applications have existed for 20 or more years, so the process of adding the GPU and getting wall-clock improvements in performance can be very challenging! Our talk will discuss how to properly target your effort, the impact of directives-based coding, and how to maintain efficiency across a hybrid cluster.

S0143 - Fluid-Structure-Interaction Using SPH and GPGPU Technology
Jean Luc Lacome (IMPETUS Afea SAS), Jerome Limido (IMPETUS Afea SAS)
Day: Wednesday, 05/16 | **Time:** 2:30 pm - 2:55 pm
Topic Areas: Computational Structural Mechanics; Algorithms & Numerical Techniques; Computational Fluid Dynamics
Session Level: Intermediate

There are two goals when developing engineering analysis software, one is accuracy and the other is speed. In the area of Fluid-Structure Interaction (FSI) computational time has always been the major impediment to solving large realistic engineering problems. In our implementation the fluid/structural dynamics solver uses a combination of GPU/CPU processing. The added benefit of using a powerful GPU workstation is that it is roughly 10 times less expensive than a regular CPU cluster. In this paper, we present the use of GPU Technology as implemented in the explicit dynamic finite element software IMPETUS Afea Solver®.

S0432 - New Ideas for Massively Parallel Preconditioners
John Appleyard (Polyhedron Software Ltd.), Jeremy Appleyard (Polyhedron Software Ltd.)
Day: Wednesday, 05/16 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Algorithms & Numerical Techniques; Computational Fluid Dynamics; Energy Exploration
Session Level: Advanced

Linear Solvers on serial machines tend to be highly recursive, but that's not an option on GPUs. In this paper we describe a new preconditoner for GMRES and similar Krylov subspace linear solvers that is highly parallel, but also provides effective mechanisms to reconcile remote driving forces in a spatially discretized system. We will present results, taken from some real-world studies using a commercial oil reservoir simulator, showing how it compares with a state of the art serial solver, and showing how performance scales in a domain decomposition formulation run on a multiple CPU+GPU cluster.

S0293 - Culises - A Library for Accelerated CFD on Hybrid GPU-CPU Systems
Daniel Gaudlitz (FluiDyna GmbH), Bjorn Landmann (FluiDyna GmbH)
Day: Wednesday, 05/16 | **Time:** 3:30 pm - 3:55 pm
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques
Session Level: Intermediate

The vast majority of CFD simulations relies on the solution of large-scale systems of linear equations (SLE), where the solution of a system can consume most of the total CPU time. We have developed a library (Culises) for state-of-the-art solution of SLE that is targeted on hybrid GPU-CPU platforms. Culises can be connected to MPI-parallelized CFD codes (e.g. OpenFOAM) via an application-specific interface. In this talk, we focus on efficient implementation of preconditioned Krylov subspace methods. Using the computing power of GPUs, Culises can significantly accelerate pure CPU computations for a multitude of industrial CFD applications.

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

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

S0286 - Scaling Applications to a Thousand GPUs and Beyond
Alan Gray (The University of Edinburgh), Roberto Ansaloni (Cray Italy)
Day: Wednesday, 05/16 | **Time:** 4:00 pm - 4:50 pm
Topic Areas: Supercomputing; Computational Fluid Dynamics; Parallel Programming Languages & Compilers; Application Design & Porting Techniques
Session Level: Intermediate

Discover how to scale scientific applications to thousands of GPUs in parallel. We will demonstrate our techniques using two codes representative of a wide spectrum of programming methods. The Ludwig lattice Boltzmann package, capable of simulating extremely complex fluid dynamics models, combines C, MPI and CUDA. The Himeno three-dimensional Poisson equation solver benchmark combines Fortran (using the new coarray feature for communication) with prototype OpenMP accelerator directives (a promising new high-productivity GPU programming method). We will present performance results using the cutting-edge massively-parallel Cray XK6 hybrid supercomputer featuring the latest NVIDIA Tesla 2090 GPUs.

S0367 - Physis: An Implicitly Parallel Framework for Stencil Computations
Naoya Maruyama (Tokyo Institute of Technology)
Day: Wednesday, 05/16 | **Time:** 4:30 pm - 4:55 pm
Topic Areas: Parallel Programming Languages & Compilers; Supercomputing; Development Tools & Libraries; Computational Fluid Dynamics
Session Level: Intermediate

This session presents how to implement finite difference methods in a concise, readable, and portable way, yet achieving good scalability over hundreds of GPUs, using the Physis high-level application framework. Physis extends the standard C language with a small set of custom declarative constructs for expressing stencil computations with multidimensional structured grids, which are automatically translated to CUDA for GPU acceleration and MPI for node-level parallelization with automatic domain-specific optimizations such as overlapped boundary exchanges. We demonstrate the programmability improvement and performance of Physis using hundreds of GPUs on TSUBAME2.0.

S0518 - GPU Computing: From Sand to Tank Dynamics
Dan Negrut (University of Wisconsin-Madison)
Day: Wednesday, 05/16 | **Time:** 5:00 pm - 5:25 pm
Topic Areas: Computational Structural Mechanics; Computational Fluid Dynamics
Session Level: Advanced

This talk explores the use of heterogeneous CPU/GPU computing, as enabled by an in-house developed Heterogeneous Computing Template (HCT), for physics-based simulations of mechanical systems. HCT draws on five components: advanced modeling techniques (formulating the governing equations); algorithmic support (solving these equations); proximity computation; domain decomposition/data exchange (for multi-node distributed CPU/GPU computing); and post-processing/visualization. These five components provide the foundation of a computational framework used to analyze mechanical systems with millions of interacting elements. Example applications will include granular terrain simulation, tracked and wheeled vehicle mobility studies (tanks, rovers), fluid-solid interaction and nonlinear finite element analysis.

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

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

S0264 - CU++: An Object-Oriented Framework for Computational Fluid Dynamics (CFD) Applications
Dominic Chandar (University of Wyoming)
Day: Thursday, 05/17 | **Time:** 9:30 am - 9:55 am
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques
Session Level: Intermediate

In this session, I will elucidate the power of blending C++ expression templates and CUDA which has resulted in a smart framework - CU++ for solving Computational Fluid Dynamics problems on structured and unstructured meshes. Briefly, CU++ allows a code developer with just C/C++ knowledge to write computer programs that will execute on the GPU with minimal knowledge of specific programming techniques in CUDA. It allows the user to reuse existing C/C++ CFD codes with minimal changes. Codes written in CU++ can also be compiled in serial mode to be executed on a CPU using the tool ugc.

S0305 - Classical Algebraic Multigrid for CFD with CUDA
Simon Layton (Boston University)
Day: Thursday, 05/17 | **Time:** 10:00 am - 10:25 am
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques
Session Level: Intermediate

Classical algebraic multigrid (AMG) is one of the most popular algorithms used in engineering, and the engine in many successful commercial packages. Among sparse linear solvers, it is known for being fast, parallel and scalable, yet it maps to GPU architecture with some considerable difficulty. We have tackled these difficulties and currently have a full CUDA implementation of classical AMG, which has been validated against the gold-standard, Hypre. Significant effort was dedicated to reducing thread divergence and optimizing memory access, and we continue to work on performance improvements. We are aiming for a competitive AMG code for fluid dynamics applications.

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

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

S0044 - A Massively Parallel Two-Phase Solver for Incompressible Fluids on Multi-GPU Clusters
Peter Zaspel (University of Bonn)
Day: Thursday, 05/17 | **Time:** 2:00 pm - 2:50 pm
Topic Areas: Computational Fluid Dynamics; Supercomputing; Algorithms & Numerical Techniques; Digital Content Creation & Film
Session Level: Intermediate

Join our presentation of a multi-GPU fluid solver for high performance GPU compute clusters. We use high-order scientific techniques to simulate the interaction of two fluids like air and water. Scientists, engineers and even the computer animation industry will profit from the enormous compute power of tens or hundreds of GPUs. A major focus in this talk will be on the applied GPU implementation techniques and the performance results including performance per Watt and performance per dollar results. We also highlight the lessons we learned from porting the complex CPU CFD code NaSt3DGPF to the GPU.

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

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

S0249 - GPU Accelerated Flow Solver Using Stabilized Finite Element Method
Hyung Taek Ahn (University of ULSAN), Hun Joo Myung (KISTI)
Day: Thursday, 05/17 | **Time:** 3:30 pm - 3:55 pm
Topic Areas: Computational Fluid Dynamics
Session Level: Intermediate

An efficient algorithm for solving Navier-Stokes equations for incompressible flows is presented. It is based on stabilized finite element method, namely SUPG(Streamline-Upwind Petrov-Galerkin) stabilization for momentum equation and PSPG(Pressure-stabilizing Petrov-Galerkin) stabilization for continuity equation. Unstructured meshes composed of tetrahedrons are utilized for simulating fluid flow around complex geometry. GPU speed up compared to its CPU counterpart will be presented for showing the efficiency of the current algorithm.

S0063 - Robust Preconditioned Conjugate Gradient for the GPU and Parallel Implementations
Rohit Gupta (Delft University of Technology)
Day: Thursday, 05/17 | **Time:** 4:00 pm - 4:50 pm
Topic Areas: Computational Fluid Dynamics; Algorithms & Numerical Techniques
Session Level: Intermediate

Get a closer look on how parallel conjugate gradient (CG) method can get an edge over it's optimized CPU implementation. We have developed preconditioning techniques for CG which are suited to the GPU and match Block-IC in terms of numerical performance. We present our results for two level preconditioned CG on the GPU and also compare it with multi-CPU implementations. Our results show that for large problem sizes(1 million unknowns and above) it is possible to achieve an order of magnitude and higher speedups for the two level preconditioned CG method.

S0218 - ASI Parallel Fortran: A General-Purpose Fortran to GPU Translator
Rainald Lohner (George Mason University)
Day: Thursday, 05/17 | **Time:** 4:30 pm - 4:55 pm

Topic Areas: Development Tools & Libraries; Computational Fluid Dynamics; Computational Physics; Parallel Programming Languages & Compilers
Session Level: Advanced
Over the last 3 years we have developed a general-purpose Fortran to GPU translator: ASI Parallel Fortran does. The talk will detail its purpose, design layout and capabilities, and show how it is used and implemented. The use of ASI Parallel Fortran will be shown for large-scale CFD/CEM codes as well as other general purpose Fortran codes.