

GPU Technology Conference 2010 Sessions on Tools & Libraries (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.

2202 - A Programming Model and Tool for Automatic High-Performance C to CUDA Mapping

Discover our automatic C-to-CUDA mapper prototype, and how it optimizes execution and data movement for a broad class of loop codes. Coupled with our powerful mapper, C as an input language does not only offer portability but also performance and performance portability. Learn about our optimizations and some of the performance obtained through different uses of the mapper.

Speaker: Benoit Meister, Reservoir Labs

Topic: Tools & Libraries

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

2209 - Accelerating Computer Vision on the Fermi Architecture

GPUS have evolved from fixed function to general purpose, and continue to evolve with new features being added in every generation. This talk will discuss how to exploit the new features introduced by the Fermi architecture (such as concurrent kernel execution, writes to texture) to accelerate computer vision algorithms.

Speaker: James Fung, NVIDIA

Topics: Computer Vision, Tools & Libraries

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

2126 - Accelerating Signal Processing: Introduction to GPU VSIPL

Learn how to use the Vector Signal Image Processing Library to accelerate signal processing applications without needing to understand platform-specific programming and optimization techniques. We will discuss how GPU VSIPL implements the VSIPL API and uses CUDA-capable GPUs to maximize performance of several example applications.

Speaker: Dan Campbell, Georgia Tech Research Institute

Topics: Signal processing, Tools & Libraries

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

2001 - Acceleration of the Freesurfer Suite for Neuroimaging Analysis

See how GPU technology has dramatically accelerated the Freesurfer suite of tools used by thousands of researchers for the analysis of neuroimaging data.

Speaker: Richard Edgar, Mass. General Hospital

Topics: Medical Imaging & Visualization, Imaging, Tools & Libraries

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

2012 - Analysis-Driven Performance Optimization

The goal of this session is to demystify performance optimization by transforming it into an analysis-driven process. There are three fundamental limiters to kernel performance: instruction throughput, memory throughput, and latency. In this session we will describe:

- how to use profiling tools and source code instrumentation to assess the significance of each limiter;
- what optimizations to apply for each limiter;
- how to determine when hardware limits are reached.

Concepts will be illustrated with some examples and are equally applicable to both CUDA and OpenCL development. It is assumed that attendees are already familiar with the fundamental optimization techniques.

Speaker: Paulius Micikevicius, NVIDIA

Topic: Tools & Libraries

Time: Thursday, September, 23rd, 15:00 - 17:00

2164 - Analytical Performance Models to Improve the Efficiency of GPU Computing

Dive deep into a simple analytical model that provides insight into performance bottlenecks of parallel applications on GPU architectures. We will discuss how the model estimates the execution time of massively parallel programs. We will also cover how to optimize applications based on our developed performance analysis models.

Speaker: Hyesoon Kim, Georgia Tech

Topic: Tools & Libraries

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

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

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

Speaker: Guido Juckeland, TU Dresden - ZIH

Topics: High Performance Computing, Tools & Libraries

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

2213 - BCSLIB-GPU: Significant Performance Gains for CAE

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

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

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

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

2002 - CUDA Debugging on Linux

Boost your development speed by mastering the CUDA debugging tools NVIDIA provides on the Linux platforms. In this session you will learn the basics of cuda-gdb and cuda-memcheck, as well as their more advanced features with live demonstrations.

Speaker: Satish Salian, NVIDIA

Topic: Tools & Libraries

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

2143 - CUDA Fortran Programming for NVIDIA GPUs

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

Speaker: Brent Leback, The Portland Group

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

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

2216 - CUDA Libraries Open House

Learn about NVIDIA's CUDA libraries and meet the engineers that develop them. Lead developers will cover the capabilities, performance and future directions for NVIDIA's CUFFT, CUBLAS, CURAND, and NPP libraries (other libraries such as CUSPARSE and open source Thrust are covered in other talks). After the presentation, NVIDIA developers will remain in the room to chat and answer questions during the lunch break.

Speakers: Ujval Kapasi, NVIDIA, Philippe Vandermersch, Nvidia, Elif Albuz, Nathan Whitehead, NVIDIA, Frank Jargstorff

Topic: Tools & Libraries

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

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

2271 - Compose CUDA Masterpieces! Write better, Leverage More

Not all CUDA code is created equally. Learn how to step up your CUDA game. Also, learn how to build large, multi-person CUDA projects for your organization.

In very clear descriptions, learn the difference between naïve GPU code, intermediate GPU code, and advanced GPU mastery. We show how careful construction of CUDA kernels can affect application performance.

We also discuss how Jacket tools greatly facilitate the development of CUDA-based projects.

Finally, we will debut the Jacket runtime's new C/C++ library. With this library, the technical computing functions in Jacket's MATLAB engine are made available in C/C++.

Speaker: James Malcolm, AccelerEyes

Topic: Tools & Libraries

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

2050 - Copperhead: Data-Parallel Python for the GPU

Learn how to write Python programs that execute highly efficiently on GPUs using Copperhead, a data-parallel Python runtime. Using standard Python constructs like map and reduce, we will see how to construct data-parallel computations and embed them in Python programs that interoperate with numerical and visualization libraries such as NumPy, SciPy and Matplotlib. We will examine how to express computations using Copperhead, explore the performance of Copperhead programs running on GPUs, and discuss Copperhead's runtime model, which enables data-parallel execution from within Python.

Speaker: Bryan Catanzaro, University of California, Berkeley

Topic: Tools & Libraries

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

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 its 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

2297 - Developing CUDA Accelerated .NET Plugins for Microsoft Excel

Quantifi will demo its xLDevelopment environment, which provide developers with an easy to use development environment which allows cuda functionality to be in Microsoft Excel. With as little as four lines, one will also select the position of the function in the menu bar, xml markup language will display in the excel help functionality, and objects can be easily added to the object cache. These objects can then be inspected by the end user or developer. Performance information can also be displayed in the object cache. The environment provides the developer an environment where he can focus on developing high performance functionality, and all intermediate layers of interface are taking care of by the environment.

Speaker: Peter Decrem, Quantifi

Topics: Tools & Libraries, Finance

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

2125 - Developing GPU Enabled Visual Effects For Film And Video

The arrival of fully programable GPUs is now changing the visual effects industry, which traditionally relied on CPU computation to create their spectacular imagery. Implementing the complex image processing algorithms used by VFX is a challenge, but the payoffs in terms of interactivity and throughput can be enormous. Hear how The Foundry's novel image processing

architecture simplifies the implementation of GPU-enabled VFX software and eases the transition from a CPU based infrastructure to a GPU based one.

Speakers: Bruno Nicoletti, The Foundry, Jack Greasley, The Foundry

Topics: Film, Tools & Libraries, Video Processing

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

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

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

Speaker: Nicolas Pinto, MIT

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

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

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

2011 - Fundamental Performance Optimizations for GPUs

This presentation covers the major CUDA optimizations. Topics will include: maximizing memory throughput, kernel launch configuration, using shared memory, and improving GPU/CPU interaction. While C for CUDA is used for illustration, the concepts covered will apply equally to programs written with OpenCL and DirectCompute APIs.

Speaker: Paulius Micikevicius, NVIDIA

Topics: Programming Languages & Techniques, Tools & Libraries

Time: Wednesday, September, 22nd, 17:00 - 18:00

2156 - GMAC: Global Memory For Accelerators

Learn how to use GMAC, a novel run-time for CUDA GPUs. GMAC unifies the host and device memories into a unified virtual address space, enabling the host code to directly access the device memory, and removing the need for data transfers between host and device memories. Moreover, GMAC also allows pointers to be used by both, the host and device code indistinctly.

This session will present the GMAC run-time and show how to use it in current applications. This session will cover from the basics of GMAC to multi-threaded applications using POSIX threads, OpenMP and MPI.

Speaker: Isaac Gelado, Universitat Politecnica de Catalunya

Topic: Tools & Libraries

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

2147 - GPGPU Development for Windows HPC Server

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

Speaker: Calvin Clark, Microsoft

Topics: High Performance Computing, Tools & Libraries

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

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

2267 - GPU Computing with MATLAB®

MATLAB is a widely used tool for scientific, engineering and financial applications. As the popularity of GPUs has grown, there is strong interest from engineers and scientists who solve computationally intensive problems to be able to leverage GPUs within MATLAB and other products from MathWorks. This talk will discuss how MathWorks tools can help engineers and scientist to take advantage of GPU resources while continuing to work in the familiar MATLAB environment. A range of capabilities will be discussed and demonstrated.

Speaker: Loren Dean, MathWorks

Topic: Tools & Libraries

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

2039 - GPU Debugging with Allinea DDT

Discover how a debugger can help you fix those hard to find bugs in your GPU software, with this introduction to the special CUDA features in Allinea DDT.

Speaker: David Lecomber, Allinea Software

Topic: Tools & Libraries

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

2210 - GPU-Ocelot: An Open Source Debugging and Compilation Framework for CUDA

Learn how to debug and profile CUDA applications using GPU-Ocelot. Ocelot is a compilation and emulation framework for CUDA that includes debugging and profiling tools as well as backend compilers for NVIDIA GPUs and x86 CPUs. We will present examples of applications developed on x86 CPUs and deployed on NVIDIA GPUs. We will also discuss memory checking, race detection, and deadlock detection tools available within Ocelot.

Speakers: Gregory Damos, Georgia Institute of Technology, Andrew Kerr, Georgia Institute of Technology, Sudhakar Yalamanchili, Georgia Institute of Technology

Topic: Tools & Libraries

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

2072 - GPUs at the Computer Animation Studio

Learn five simple ways in which GPUs have been adopted in the production pipeline at Blue Sky Studios. Covers how we use GPUs to improve animation tools, add real-time anaglyph support, and accelerate noise functions including code samples from production tools.

Speaker: Hugo Ayala, Blue Sky Studios

Topics: Film, Stereoscopic 3D, Tools & Libraries

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

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

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

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

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

Topics: Tools & Libraries, High Performance Computing

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

2048 - H.264/AVC Video Encoding with CUDA and OpenCL

Join experts from MainConcept, a leading provider of video codecs to the professional market, as they demonstrate the latest version of their CUDA-based H.264/AVC Encoder.

Speaker: Thomas Kramer, MainConcept

Topics: Video Processing, Tools & Libraries

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

2300 - High-Performance Compressive Sensing using Jacket

This talk will present the ongoing work that I am doing in the L1-optimization group at Rice University. The purpose of the work is to merge both compressive sensing, for image/signal reconstructions and GPU computation, using NVIDIA's GPUs to enhance the technology of CS.

This talk will cover basic concepts in compressive sensing and the easy adaptation of operating on the GPU, in particular working with Jacket (by AccelerEyes). We will then cover some of our numerical experiments that encompass the use of different flavors of algorithms.

Speaker: Nabor Reyna

Topics: Imaging, Tools & Libraries

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

2219 - High-Productivity CUDA Development with the Thrust Template Library

Thrust is a parallel template library for developing CUDA applications. Modeled after the C++ Standard Template Library (STL), Thrust brings a familiar abstraction layer to the realm of GPU computing. Thrust provides host and device variants of the STL vector container to simplify memory management and facilitate data transfers. These containers are complemented with a large collection of generic data-parallel algorithms and a suite of useful iterator adaptors. Together, these features form a flexible high-level interface for GPU programming that greatly enhances developer productivity.

In this session we'll discuss Thrust's features and explain the basic design philosophy of the library.

Speaker: Nathan Bell, NVIDIA Research

Topic: Tools & Libraries

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

2299 - Integrating CUDA BLAS with IMSL Fortran

As GPU hardware becomes more prevalent in both research and commercial institutions, software that takes advantage of this specialized hardware is growing in demand. In many cases, it is infeasible or impossible to rewrite an existing program to run entirely on the GPU, so the goal is often to offload as much work as possible. As the IMSL Library team at Rogue Wave Software considers how best to tackle the GPU realm with a general mathematical library, the IMSL Fortran Library takes an initial step where the CUDA BLAS library is utilized to offload CPU work to GPU hardware. This presentation will discuss the approach and architecture of the solution. Benchmark results will show where success has been found. Plans for future products will also be covered.

Topic: Tools & Libraries

Time: Tuesday, September, 21st, 14:00 - 14:20

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

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

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

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

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

2004 - Languages, APIs and Development Tools for GPU Computing (Pre-Conference Tutorial)

Get a head start on the conference with this first-day introduction to key technologies for GPU Computing. This 90-minute tutorial session will cover the key features and differences between the major programming languages, APIs and development tools available today. Attendees will also learn several high level design patterns for consumer, professional and HPC applications, with practical programming considerations for each.

Speaker: Will Ramey, NVIDIA

Topics: Programming Languages & Techniques, Tools & Libraries

Time: Monday, September, 20th, 13:00 - 14:20

2026 - MatCloud: Accelerating Matrix Math GPU Operations with SaaS

We present MatCloud (www.mat-cloud.com), a cloud infrastructure and service for scientific computing using state-of-the-art GPU clusters. MatCloud is a service infrastructure exposed by a simple web terminal interface to run Matlab-like commands/scripts. Join us to see how GPU technology can not only be applied to cloud computing community, but also boost the adoption of cloud computing for its dramatic performance gains over traditional cloud infrastructures. MatCloud is an in-progress academic project and is under active development.

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

Topics: Cloud Computing, Tools & Libraries

Time: Tuesday, September, 21st, 17:00 - 17:20

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

2121 - Maximizing Throughput of Barco's GPU-Enabled Video Processing Server

Find out how Imec middleware realizes the full potential of GPU-enabled video processing servers to manage multiple video processing pipelines. We will discuss how the middleware monitors GPU and CPU execution to best balance the load. Covers how we achieved a 30% increase in throughput with only a minimal 0.05% overhead on Barco's GPU-enabled video processing server.

Speakers: Maja D'Hondt, imec, Roel Wuyts, imec

Topics: Video Processing, Tools & Libraries

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

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

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

Speaker: Francois Bodin, CAPS entreprise

Topics: High Performance Computing, Tools & Libraries

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

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

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

2249 - New Programming Tools GPU Computing

This session will focus on new parallel programming tools for GPU computing. The type of tools that fit into the session include (1) Planning tools for porting legacy applications to use GPU computing, (2) High-level programming and scripting tools for GPU computing, (3) Automation of common performance optimizations for GPU computing, (4) Performance analysis and diagnosis tools for GPU computing, (5) Tools that simplify heterogeneous parallel computing.

Speakers: Wen-mei Hwu, University of Illinois, Urbana-Champaign, Andrew Schuh

Topic: Tools & Libraries

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

2008 - OpenCL Optimization

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

Speaker: Peng Wang, NVIDIA

Topics: Tools & Libraries, High Performance Computing

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

2018 - OpenCL on the GPU (Pre-Conference Tutorial)

OpenCL is Khronos' new open standard for parallel programming of heterogeneous systems. This tutorial session will introduce the main concepts behind the standard and illustrate them with some simple code walkthrough. Attendees will also learn how to make efficient use of the API to achieve good performance on the GPU.

Speaker: Cliff Woolley, NVIDIA

Topic: Tools & Libraries

Time: Monday, September, 20th, 16:00 - 17:20

2227 - OpenGL 4.0 Tessellation for Professional Applications

The new generation of accelerated graphics is elevating visual computing to new heights. Tessellation, one of its most anticipated features, is already used in many scenarios to bring 3D graphics to an unprecedented level of realism. This talk will introduce tessellation using OpenGL 4.0. We will also describe how an existing application can be adapted to efficiently take advantage of this new feature and also how to overcome some of the challenges.

Speaker: Philippe Rollin, NVIDIA

Topics: Computer Graphics, Tools & Libraries

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

2124 - Operating System Abstractions for GPU Programming

GPGPU frameworks such as CUDA improve programmability, but GPU parallelism remains inaccessible in many application domains. This session argues that poor OS support causes this problem. OSES do not provide the kind of high-level abstractions for GPUs that applications expect for other resources like CPUs and file systems. We advocate reorganizing kernel abstractions to support GPUs as first-class computing resources, with traditional guarantees such as fairness and isolation. We demonstrate shortcomings in Windows 7 GPU support, and show that better OS abstractions can accelerate interactive workloads like gesture recognition by a factor of 10X over a CUDA implementation.

Speakers: Christopher Rossbach, Microsoft Research, Emmett Witchel, University of Texas at Austin

Topics: Programming Languages & Techniques, Tools & Libraries

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

2149 - Overview of Parallel Nsight for Visual Studio

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. This session is an entry level overview of the GPU computing and graphics development features of Parallel Nsight as well as a glimpse into the future of this powerful tool.

Speaker: Kumar Iyer, NVIDIA

Topic: Tools & Libraries

Tuesday, September, 21st, 11:30 - 11:50

Times:

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

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

2245 - Parallel Nsight for Microsoft Visual Studio (Pre-Conference Tutorial)

NVIDIA Parallel Nsight provides access to the power of the GPU from within the familiar environment of Microsoft Visual Studio. In this session, you will learn how to use Parallel Nsight to develop GPU computing and graphics applications.

Learn how to use the powerful Parallel Nsight debugger to identify errors in CUDA C/C++ kernels and HLSL shaders using GPU breakpoints and direct memory and variable inspection. See how Parallel Nsight displays system-wide performance characteristics, allowing you to create efficient GPU algorithms. Discover Parallel Nsight's Direct3D debugging features and appreciate how the familiar Visual Studio environment can accelerate productivity.

Speaker: Kumar Iyer, NVIDIA

Topic: Tools & Libraries

Time: Monday, September, 20th, 16:00 - 17:20

2151 - Parallel Nsight: Analyzing and Optimizing Massively Parallel Applications [Advanced]

Life altering products that provide early detection of breast cancer or simulate molecular behavior, accelerating drug discovery, are becoming reality thanks to the power of the GPU. As these technologies become mainstream, mainstream tools are required to support these development efforts. NVIDIA Parallel Nsight delivers the power of the GPU within the familiar environment of Microsoft Visual Studio. In this session, you will learn advanced techniques for visualizing your application's workloads and performance characteristics across the CPU, GPU, and operating system, and explore the depths of Parallel Nsight profilers, including GPU performance counters and how to use them.

Speaker: Sebastien Domine, NVIDIA

Topics: Tools & Libraries, Programming Languages & Techniques

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

2150 - Parallel Nsight: Debugging Massively Parallel Applications [Advanced]

Data parallel algorithms that provide real-time financial options pricing or identification of hidden oil reserves are utilizing the massively parallel nature of the GPU for industry changing performance gains. Developers require industry standard development tools to create the software that accomplishes these parallel tasks. NVIDIA Parallel Nsight delivers the power of the GPU within the familiar environment of Microsoft Visual Studio. In this session, you will learn advanced techniques for debugging CUDA C/C++ and DirectCompute code using Parallel Nsight, including conditional and data breakpoints as well as out of bound GPU memory access detection.

Speaker: Sebastien Domine, NVIDIA

Topics: Tools & Libraries, Programming Languages & Techniques

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

2248 - Parallel Processing on GPUs at the University of Utah

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

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

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

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

2053 - Pixel Bender: Building a Domain Specific Language on the GPU

Examine the challenges and advantages of building the Pixel Bender domain specific language for image processing for the GPU. We will examine how Pixel Bender was made to work within several Adobe applications across a wide range of hardware systems and platforms.

Speaker: Bob Archer, Adobe Systems Inc

Topic: Tools & Libraries

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

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

2023 - Processing Device Arrays with C++ Metaprogramming

I will describe tricks for building APIs using C++ metaprogramming that generate custom kernels for complex manipulation of device-side arrays in CUDA. Using a variation of Expression Templates, multiple operations can be fused into a single kernel that executes with reasonable efficiency.

Speaker: Jonathan Cohen, NVIDIA Research

Topics: Programming Languages & Techniques, Tools & Libraries

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

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

2148 - Rapid Prototyping and Visualization with OpenCL Studio

Learn about OpenCL Studio, an integrated OpenCL and OpenGL development environment for parallel programming and visualization. We will discuss building end user applications and using its integrated visualization capabilities to better understand the output and internal structure of parallel algorithms. We will also demonstrate its capabilities using several sample applications including particle systems, volumetric rendering, and image processing.

Speaker: Jochen Stier, Geist Software Labs

Topic: Tools & Libraries

Time: Tuesday, September, 21st, 15:00 - 15:20

2014 - Scalable Subsurface Data Visualization Framework

Mental Images' DiCE-based geospatial library is a CUDA and cluster-based visualization framework that enables scalable processing and rendering of huge amounts of subsurface data for interactive seismic interpretation.

Geospatial exploration in the oil and gas industries is concerned with scanning the earth's subsurface structure for detecting oil and for cost-effective drilling of detected oil reservoirs.

Efficient seismic interpretation requires the interpreters to be able to interactively explore huge amounts of volumetric seismic information with embedded stacked horizons to gain visual insight into the subsurface structure and to determine where oil recovery facilities and drilling infrastructure shall be built.

Speakers: Tom-Michael Thamm, mental images GmbH, Marc Nienhaus, mental images GmbH

Topics: Energy Exploration, Databases & Data Mining, Imaging, Tools & Libraries

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

2177 - Simplifying Parallel Programming with Domain Specific Languages

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

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

Topics: Tools & Libraries, High Performance Computing

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

2241 - Standing Out: Implementing a Great Stereo UI

Learn how to make S3D compatible user interfaces, HUDs, and in-game menus. The first part of this session will outline the common problems users encounter when displaying traditional 2D UI in stereoscopic 3D. The second part will focus on the different techniques, tips/tricks, and best practices developers can use to create high-quality S3D interfaces. The presentation will

highlight examples from several shipped titles, as well as showcase a complete 3D UI game demo running in S3D on multiple devices including PC and mobile.

Speaker: Brendan Iribe, Scaleform

Topics: Stereoscopic 3D, Tools & Libraries, Computer Graphics, Mobile & Tablet & Phone

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

2160 - StarPU: a Runtime System for Scheduling Tasks

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

Speaker: Cedric Augonnet, INRIA

Topics: Tools & Libraries, High Performance Computing

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

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

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

Speaker: Mark Harris, NVIDIA

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

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

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

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

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

2220 - Thrust by Example: Advanced Features and Techniques

Thrust is a parallel template library for developing CUDA applications which is modeled after the C++ Standard Template Library (STL). In this session we'll show how to implement decompose problems into the algorithms provided by Thrust. We'll also discuss the performance implications of "kernel fusion" and "array of structs" vs. "structure of arrays" memory layouts and how they relate to Thrust. Lastly, we'll present evidence that Thrust implementations are fast, while remaining concise and readable.

Speaker: Jared Hoberock, NVIDIA

Topic: Tools & Libraries

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

2225 - Tools for Managing Clusters of NVIDIA GPUs

Learn about the suite of tools NVIDIA provides to manage large installations of GPUs from the NVIDIA Tesla Series. The presentation will cover cluster management – tool and library –, as well as the GPUDirect technology that enables GPUs to communicate faster across the network.

Speaker: Peter Buckingham, NVIDIA

Topic: Tools & Libraries

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

2251 - TotalView Debugger for CUDA

Hear how the TotalView debugger is being extended to support GPU computation with CUDA. In addition to the basic challenges associated with debugging parallel programming, CUDA programming introduces a number of new concepts for which developers need visibility in debugging: a hierarchical memory, near-SIMD warps, streams, and kernels, among others. How do we create a tool that handles it all? We'll be discussing the status of our work and the challenges encountered in bringing this all together into a single package, TotalView for CUDA.

Speaker: Chris Gottbrath, TotalView Technologies, Inc., a Rogue Wave Software company

Topic: Tools & Libraries

Time: Wednesday, September, 22nd, 15:00 - 15:20

2111 - Using R for High-Performance Data Analysis

Data analysis is the art and the science of getting the correct quantitative models and their numerical parameters from the observed data. In this talk, we report on a project to integrate CUDA into the open source data analysis environment R. The combined use of the CPU and GPU resources can efficiently exploit the significant amount of data parallelism inherent in most data analysis problems and methods. This makes interactive analysis possible even for large, compute-intensive problems. The implementation and the achievable performance gains will be demonstrated on a concrete example from quantitative finance.

Speaker: Domokos Vermes, Worcester Polytechnic Insitute

Topics: Tools & Libraries, Databases & Data Mining, Finance, Life Sciences

Time: Tuesday, September, 21st, 16:30 - 16:50

2016 - VDPAU: PureVideo on Unix

Learn about VDPAU (Video Decode and Presentation API for Unix). VDPAU provides GPU-accelerated video decoding, post-processing, UI compositing, and display on Unix. VDPAU also supports sharing surfaces with OpenGL and CUDA ("interop"). This allows developers to implement their own post-processing algorithms or scene analysis, or to use decoded video surfaces as part of a scene rendered using OpenGL.

Speaker: Stephen Warren, NVIDIA

Topics: Video Processing, Tools & Libraries

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

2113 - WebGL: Bringing 3D to the Web

WebGL is a newly-emerging standard for 3D graphics and visual computing on the web. Supported and developed by major web browser vendors, WebGL enables rich interactive 3D graphics delivered through a web browser, on both desktop and mobile platforms. This session will contain an introduction to WebGL, and will focus application development issues unique to the web platform, optimization concerns, and how web technologies such as offline app support, HTML5 video and audio, File and WebSockets integrate with WebGL. Experienced OpenGL developers will learn how to transition their knowledge to WebGL development.

Speaker: Vladimir Vukicevic, Mozilla Corporation

Topics: GPU Accelerated Internet, Tools & Libraries, Computer Graphics

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