

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

Sessions on **Visualization** (subject to change)

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

TUTORIALS

S0023 - NVIDIA OpenGL for 2012

Mark Kilgard (NVIDIA)

Day: Monday, 05/14 | **Time:** 9:00 am - 10:20 am

Topic Areas: Computer Graphics; Development Tools & Libraries; Visualization; Audio, Image and Video Processing

Session Level: Intermediate

Attend this session to get the most out of OpenGL on NVIDIA Quadro and GeForce GPUs. Topics covered include the latest advances available for Cg 3.1, the OpenGL Shading Language (GLSL); programmable tessellation; improved support for Direct3D conventions; integration with Direct3D and CUDA resources; bindless graphics; and more. When you utilize the latest OpenGL innovations from NVIDIA in your graphics applications, you benefit from NVIDIA's leadership driving OpenGL as a cross-platform, open industry standard.

S0341 - See the Big Picture Scalable Visualization Solutions for System Integrators

Doug Traill (NVIDIA)

Day: Monday, 05/14 | **Time:** 10:30 am - 11:50 am

Topic Areas: Visualization

Session Level: Beginner

NVIDIA Quadro Scalable Visualizations Solutions provide many feature for System Integrators who are building large scale displays. Come join us in this tutorial session on how to configure multi-projector systems, stereoscopic and immersive displays.

S0530 - Multi-Display Roundtable

Andrew Page (NVIDIA)

Day: Monday, 05/14 | **Time:** 1:00 pm - 2:20 pm

Topic Areas: Visualization

Session Level: Beginner

Join NVIDIA product manager and application engineers for multi-display systems for an interactive discussion on the current trends in video walls, blended multi-projector systems and its deployment.

SESSIONS

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.

S0024 - GPU-Accelerated Path Rendering

Mark Kilgard (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm

Topic Areas: Computer Graphics; GPU Accelerated Internet; Digital Content Creation & Film; Visualization

Session Level: Beginner

Standards such as Scalable Vector Graphics (SVG), PostScript, TrueType outline fonts, and immersive web content such as Flash depend on a resolution-independent 2D rendering paradigm that GPUs have not traditionally accelerated. This tutorial explains a new opportunity to greatly accelerate vector graphics, path rendering, and immersive web standards using the GPU. By attending, you will learn how to write OpenGL applications that accelerate the full range of path rendering functionality. Not only will you learn how to render sophisticated 2D graphics with OpenGL, you will learn to mix such resolution-independent 2D rendering with 3D rendering and do so at dynamic, real-time rates.

S0254 - Graphics in the Cloud - How NVIDIA is Enabling Cloud Visualization

Will Wade (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm

Topic Areas: Cloud Computing; Visualization; Computer Graphics

Session Level: Intermediate

Engineers, artists, scientists, and gamers are the most demanding visual thinkers on the planet, and as such have not been willing to move their computing environments to the infamous "cloud". These remotely accessed systems are seen as slow and not up to the visual experience that users expect when dealing with these types of applications. NVIDIA aims to change that perception with the NVIDIA Virtual Graphics Platform. In this session you will hear about the technologies behind accelerating graphics in the cloud, and some of the industry partnerships that are enabling it.

S0328 - Best Practices in GPU-Based Video Processing

Thomas True (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm

Topic Areas: Audio, Image and Video Processing; Digital Content Creation & Film; Computer Vision; Medical Imaging & Visualization

Session Level: Intermediate

The combination of the GPU's massively parallel compute engine with extremely high memory bandwidth and new programming paradigms such as CUDA and OpenCL have made the GPU well suited for image and video processing applications. This session will explore best practices and techniques for the development of efficient GPU-based video and image processing applications. Topics to be discussed include image segmentation and threading models for efficient parallelism, optimal memory usage strategies to reduce expensive data movement as well as

multi-GPU considerations. Case studies and examples specific to video and image processing will be presented.

S0413 - Delivering 3D Professional Graphics from the Cloud with Citrix XenDesktop

Derek Thorslund (Citrix Systems, Inc.)

Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm

Topic Areas: Cloud Computing; Computer Graphics; Visualization

Session Level: Beginner

Recent technological advances have made it practical to deliver 3D professional graphics applications from the Cloud (private or public) with a high quality user experience and at an attractive cost. Organizations can keep their intellectual property safe in the data center since only fully-rendered screen images are sent over the network. Users in remote locations no longer have to wait for large file transfers. And they can access 3D models from a wide variety of devices, including iPads and Android tablets. Learn how Citrix XenDesktop, XenServer and Receiver technologies have made all of this a reality for many organizations today.

S0436 - Integrated GPU Acceleration with Real Time Visualization Of Terabyte Data

Kelly Walker (Hue)

Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:25 pm

Topic Areas: Visualization; Energy Exploration

Session Level: Beginner

Computation and visualization doesn't necessarily have to act as two separate entities. This talk explains the integration of real-time compute with real-time visualization. Industry and academia have provided attractive solutions for compiler-directive optimized code for computations. To support cases that involves massive yet ad-hoc data I/O and computation with interactive visualization, Hue developed a different model which bridges the gap between "complete system rewrite" and "compiler directive optimized code". The talk explains how highly optimized data I/O mechanisms coupled with predefined input and output definitions for kernels provide excellent scalability and interactivity during runtime.

S0356 - Optimized Texture Transfers

Shalini Venkataraman

Day: Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm

Topic Areas: Visualization

Session Level: Advanced

Many real world graphics applications need to transfer textures efficiently in and out of GPU memory in the form of 2D images, 2.5D terrains or 3D volumes as well as their time-varying counterparts. This talk covers a new feature of recent GPUs called copy engines, which are available in NVIDIA's Fermi generation and above GPU's. Copy engines allow true GPU asynchronous transfers where data transfer and rendering OpenGL commands can occur simultaneously. This talk provides programming pointers on how to take advantage of this feature.

S0359 - VMware and NVIDIA: Delivering 3D Workstations from the Cloud

Aaron Blasius (VMware), Warren Ponder (VMware)

Day: Tuesday, 05/15 | **Time:** 5:00 pm - 5:50 pm

Topic Areas: Visualization; Cloud Computing

Session Level: Advanced

This session will detail the delivery of the most demanding Workstation class workloads from the private cloud using technologies from NVIDIA and VMware. We will cover the configuration and performance metrics of the combined VMware, NVIDIA direct pass through hardware accelerated graphics solution. Using sample workloads,

we will demonstrate how customers can realize the operational and security benefits of cloud based personal computing without sacrificing performance.

S0017 - 4D Medical Image Processing with CUDA

Anders Eklund (Linköping University)

Day: Wednesday, 05/16 | **Time:** 9:00 am - 9:50 am

Topic Areas: Medical Imaging & Visualization; Audio, Image and Video Processing; Neuroscience; Visualization

Session Level: Advanced

Learn how to do 4D image processing with CUDA, especially for medical imaging applications. In this session we will give a couple of examples of how 4D image processing can take advantage of the computational power of the GPU. We will present how to use the GPU for functional magnetic resonance imaging (fMRI) analysis and true 4D image denoising. Most of our examples use the GPU both to speed-up the analysis and to visualize the results.

S0353 - Programming Multi-GPUs for Scalable Rendering

Shalini Venkataraman (NVIDIA)

Day: Wednesday, 05/16 | **Time:** 9:00 am - 9:50 am

Topic Areas: Visualization

Session Level: Advanced

Multi-GPU configurations are becoming common affordable options for OpenGL applications to scale performance, data size, display size and image quality. We show how to structure your application for multi-gpu rendering by using multiple threads and OpenGL contexts and handle the synchronization and data transfer. We conclude with a discussion of common scaling approaches: - Image Scaling - Each GPU is responsible for rendering and displaying its view frustum. - Performance Scaling - Maintaining a constant frame rate by dedicating GPUs for certain rendering tasks. - Data Scaling - Decomposition of large datasets that across multiple GPUs.

S0322 - Scaling Up Your Displays

Shalini Venkataraman (NVIDIA)

Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am

Topic Areas: Visualization; Computer Graphics

Session Level: Beginner

This talk will describe how to scale up from one to many displays for high end visualization. You will learn about NVIDIA's new Warp and Blend capability that allows you to create a truly seamless logical display comprised of many individual display outputs. With this new capability you can project your graphics onto curved surfaces and implement the correct transformation entirely on the GPU without any external hardware to get the correct display transformations.

S0312 - GPU Implementation for Rapid Iterative Image Reconstruction in Nuclear Medicine

Jakub Pietrzak (University of Warsaw)

Day: Wednesday, 05/16 | **Time:** 10:00 am - 10:25 am

Topic Areas: Medical Imaging & Visualization; Computational Physics; Computer Graphics

Session Level: Intermediate

GPU implementation can greatly accelerate iterative techniques of 3D image reconstruction in nuclear medicine imaging. Single Photon Emission Computed Tomography (SPECT) is a functional imaging modality widely used in clinical diagnosis. To obtain high quality images within reduced scanning times high sensitivity collimators need to be used and their response function modeled in the reconstruction. This is in general very computationally

intensive and unfeasible with CPU and algorithm implementations. Our software is able to perform the reconstruction of patient data within clinically acceptable times using relatively low cost and widely available hardware.

S0348 - GPUs Open New Avenues in Medical MRI

Chris A. Cocosco (University Medical Center Freiburg, Dept. of Radiology, Medical Physics)

Day: Wednesday, 05/16 | **Time:** 10:30 am - 10:55 am

Topic Areas: Medical Imaging & Visualization

Session Level: Beginner

See how GPUs enable exciting new developments in medical Magnetic Resonance Imaging (MRI). Their computational power makes now practical new MRI techniques that can bring shorter imaging sessions, better images, and more insight into human physiology. Learn about the characteristics of the general computational approach for obtaining the final image, and how it can be implemented using an iterative conjugate gradient algorithm. The algorithm exhibits massive parallelism and fits well the GPU architecture. Learn about its CUDA implementation details and Matlab integration. See throughput measurements of Tesla GPUs compared to top of the line many-core and large RAM CPU systems.

S0355 - Seamless Scalable Displays- using NVIDIA Warp + Intensity API

Rajeev Surati (Scalable Display Technologies)

Day: Wednesday, 05/16 | **Time:** 10:30 am - 10:55 am

Topic Areas: Visualization; Audio, Image and Video Processing; Computer Vision; Computer Graphics

Session Level: Beginner

In this talk we will discuss how we use the NVIDIA Warp and Intensity API to create seamless displays made up of multiprojectors based on our camera feedback systems. We will show and discuss case studies in production including a 25 megapixel touch wall, military dome simulation systems, VR Walls, VR Caves, and immersive conference rooms that are made affordable and enabled by this technology.

S0342 - Volumetric Processing and Visualization on Heterogeneous Architecture

Wei Li (Siemens Corporation)

Day: Wednesday, 05/16 | **Time:** 2:00 pm - 2:25 pm

Topic Areas: Visualization; Supercomputing

Session Level: Advanced

Volumetric data is typically very large and involves intensive computation for processing and visualization. We have developed an OpenCL-based framework that can utilize all available resources in a system or a cluster of systems. The framework manages one or more OpenCL devices. A large volume is partitioned into bricks. Each OpenCL device is associated with a set of brick producers that generates the contents of bricks while optionally utilizing other bricks as input. The framework is also composed of a scheduler that distributes brick workloads to different devices and chooses an optimized processing order aiming at certain criteria.

S0105 - Hardware Acceleration for Vessel Visualization Tasks

Christoph Kubisch (NVIDIA)

Day: Wednesday, 05/16 | **Time:** 2:30 pm - 2:55 pm

Topic Areas: Medical Imaging & Visualization; Computer Graphics

Session Level: Beginner

To analyze datasets visually, systems with fast feedback loops on user interaction are beneficial. In this session rendering and preprocessing techniques for medical volume data will be presented using OpenGL and CUDA. In

the context of the coronary artery disease the analysis of individual vessel branches is important. We show how local transfer function application and generation by means of histogram analysis can help navigating and finding details in the datasets. Furthermore, domain-specific acceleration and illustration techniques for volume rendering are also applied to datasets from brain aneurysms.

S0141 - GPU-Accelerated Optical Coherence Tomography Imaging

Kang Zhang (GE Global Research)

Day: Wednesday, 05/16 | **Time:** 3:30 pm - 3:55 pm

Topic Areas: Medical Imaging & Visualization

Session Level: Beginner

We developed a series of GPU-based technologies to accelerate the imaging reconstruction and visualization for optical coherence tomography (OCT). Several GPU-based algorithms such as non-uniform fast Fourier transform, numerical dispersion compensation, simultaneous phase modulation and multi-GPU implementation were developed to achieve improved impulse response, better SNR, doubled imaging range and higher system stability. The GPU-accelerated 4D-OCT system was validated by imaging both in vivo and ex vivo biological tissues. This technology overcomes the imaging reconstruction and visualization bottlenecks that widely exist in current ultrahigh speed OCT systems and opens the way to interventional OCT imaging for applications in guided microsurgery.

S0507 - Interactive and Scalable Subsurface Data Visualization Framework

Tom-Michael Thamm (NVIDIA ARC), Marc Nienhaus (NVIDIA ARC)

Day: Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm

Topic Areas: Visualization; Cloud Computing

Session Level: Intermediate

The goal is to present an interactive visualization framework for large geo-spatial data. This framework has been developed by NVIDIA Advanced Rendering Center for the oil and gas (Hydrocarbhone) industry. The Cuda based application is running on the cloud at interactive frame-rates. The visualization is remote on clients in a browser, including tablets. The scalable visualization framework can handle terra bytes of.

S0202 - Terascale Volume Visualization in Neuroscience

Johanna Beyer (King Abdullah University of Science and Technology), Markus Hadwiger (KAUST)

Day: Wednesday, 05/16 | **Time:** 4:30 pm - 4:55 pm

Topic Areas: Visualization; Neuroscience

Session Level: Intermediate

Learn how to create a scalable volume visualization system for interactive rendering of terascale EM data. We will describe the major design principles, how we can avoid the standard approach of pre-computing a 3D multi-resolution hierarchy such as an octree, and how to handle continuous streaming of newly acquired data. For rendering we build upon a visibility-driven approach and 3D virtual texturing, and perform interactive volume rendering of a "virtual" volume, where the corresponding physical storage is only represented and populated in a sparse manner with 2D instead of 3D image data on the fly during rendering.

S0303 - GPU Acceleration for Threshold Based Region Growth Algorithms.**Supratik Moulik (University of Pennsylvania), Jason Walsh (University of Pennsylvania 3D lab)****Day:** Thursday, 05/17 | **Time:** 9:00 am - 9:50 am**Topic Areas:** Medical Imaging & Visualization; Bioinformatics**Session Level:** Beginner

Come learn how the massively parallel computing power of modern GPUs help to create faster and more accurate volume rendered images for the medical imaging community. Attendees of this session will gain insight into how GPUs can accelerate region growth algorithms and how these algorithms can be optimized for the latest generation of NVIDIA hardware. Topics covered will include fundamental of region growth, GPU implementations, and practical examples of vessel tracking algorithms based on GPU accelerated algorithms.

S0326 - Next Generation InfoWall**Alina Alt (NVIDIA)****Day:** Thursday, 05/17 | **Time:** 9:00 am - 9:50 am**Topic Areas:** Visualization; Computer Graphics**Session Level:** Intermediate

Learn how you can use a multiple display configuration to render video content captured from multiple sources, utilizing the power of GPUs to achieve unprecedented performance.

S0324 - Content Generation and Real-Time Hologram Computation for Holographic 3D-Displays**Enrico Zschau (SeeReal Technologies GmbH)****Day:** Thursday, 05/17 | **Time:** 10:00 am - 10:25 am**Topic Areas:** Visualization; Stereoscopic 3D; Algorithms & Numerical Techniques; Audio, Image and Video Processing**Session Level:** Beginner

This session will introduce SeeReal's sub-hologram technology to massively reduce hologram computation effort in comparison to classic holography and how SeeReal implemented those still compute intensive algorithms to execute on the GPU to enable viewing of interactive, rich 3D-content on holographic 3D-displays using off-the-shelf graphics hardware. In contrast, you will explore why classic holography does not suit well for interactive applications. Furthermore guidelines to create appropriate 3D-content are presented, including aspects regarding transparency in holograms. Finally the specification and some impressions of SeeReal's 20 holographic prototype will be presented, which allows viewing of live computed holograms showing 3D-content and 3D-video.

S0053 - Real Time GPU-Based Marine Scenes Simulation**Jérôme Graindorge (ALYOTECH), Julien Houssay (ALYOTECH)****Day:** Thursday, 05/17 | **Time:** 10:00 am - 10:25 am**Topic Areas:** Climate & Weather Modeling; Visualization**Session Level:** Intermediate

Marine survey, carried out by sea or by air, is of major concern for current defense and security applications. Essential surveillance/ observation/ identification systems involve electro-optics (visible and infra-red) and radar. Optimizing their performance requires amounts of expensive observational data spanning the wide variability of the marine environment. Computer simulation provides a valuable flexible and inexpensive alternative. Since 2007, ALYOTECH, in partnership with the IFREMER (French Research Institute for Exploration of the Sea), has been developing a GPU-based real-time ocean scene simulator for visible, infrared and radar sensors, in order to meet the challenging requirements arising from marine survey issues.

S0382 - Hybrid System Architectures for High-Speed Processing in Optical Coherence Tomography**Brian Applegate (Texas A&M University Department of Biomedical Engineering)****Day:** Thursday, 05/17 | **Time:** 10:00 am - 10:25 am**Topic Areas:** Medical Imaging & Visualization; Life Sciences; Application Design & Porting Techniques; Development Tools & Libraries**Session Level:** Intermediate

Several factors are spurring the development of hardware and software to accomplish high-speed processing for Optical Coherence Tomography (OCT), e.g. ultrahigh speed (>1 MHz) volumetric imaging and clinical applications (e.g. intravascular imaging). The computation power of GPUs ensures that it will be an essential part of the solution. We are exploring the development of a hybrid system in which the computational burden is shared between GPUs and other processors. This will make it possible to extract crucial diagnostic information in real or near real time. Technical challenges and recent progress will be discussed.

S0252 - Building Real-Time Professional Visualization Solutions with OpenCL**Kristof Denolf (Barco), Ronny Dewaele (Barco)****Day:** Thursday, 05/17 | **Time:** 10:30 am - 10:55 am**Topic Areas:** Audio, Image and Video Processing; Visualization**Session Level:** Intermediate

Professional visualization solutions, like high-quality high-resolution medical displays or very large screens for surveillance or entertainment, benefit from GPU's image and graphics compute capabilities to achieve real-time performance, but add specific constraints, like low-latency, multiple HD streams and strict synchronization. This talk first motivates the industrial relevance of development in OpenCL on heterogeneous devices. It then explains the techniques currently explored to meet the specific design constraints, with a main focus on parallel data transfer and compute. The lessons learned are illustrated with a real-life example.

S0022 - Scalable Frameworks and Algorithms for Terascale Radio Astronomy Images**Christopher Fluke (Swinburne University of Technology - Centre for Astrophysics and Supercomputing)****Day:** Thursday, 05/17 | **Time:** 2:30 pm - 2:55 pm**Topic Areas:** Astronomy & Astrophysics; Visualization**Session Level:** Intermediate

Learn how the oldest science is using the newest processors to solve a critical problem: how to accomplish traditional image analysis and visualization tasks when the images are terabytes in size? Simple, standard operations such as displaying 2-d slices, evaluating image statistics, and applying histogram equalization become manifestly challenging when images dramatically exceed single-node memory capacity. We will explain how our hybrid CPU-GPU cluster framework - which can volume render a 200GB image at >50fps! - will support traditional radio astronomy tasks for the colossal images that the Square Kilometre Array and its precursor, the Australian SKA Pathfinder, will generate.

S0106 - GPU Based Numerical Methods in Mathematica**Ulises Cervantes-Pimentel (Wolfram Research), Abdul Dakkak (Wolfram Research)****Day:** Thursday, 05/17 | **Time:** 2:30 pm - 3:20 pm**Topic Areas:** Algorithms & Numerical Techniques; Visualization; Application Design & Porting Techniques; Development Tools & Libraries**Session Level:** Intermediate

A fast way of developing, prototyping and deploying numerical algorithms that can take advantage of CUDA capable systems is available in Mathematica 8. Over the past year, educators, scientists, and business users have taken advantage of the benefits that the support of GPU programming in Mathematica. By integrating and

implementing CUDA/OpenCL in their programs, users make use of a hybrid approach, combining the speed-up that GPUs offer and a powerful numerical development system. In this presentation several examples describing numerical applications ranging from deconvolution of MRI imaging, linear solvers for FEM, systems of ODEs, line integral convolution visualization are presented.

S0411 - Artifact-Free Cloud-Based CAD Rendering

Sara McMains (UC Berkeley), Sushrut Pavanaskar (UC Berkeley)

Day: Thursday, 05/17 | **Time:** 4:30 pm - 4:55 pm

Topic Areas: Algorithms & Numerical Techniques; Computer Graphics; Cloud Computing; Visualization

Session Level: Beginner

Cloud computing for mechanical CAD provides centrally stored and synchronized models for concurrent engineering. For compactness, trimmed parametric NURBS surface representations are optimal for data transfer to client devices, which must evaluate and render models locally. Direct GPU rendering without pre-tessellation is an attractive solution in this context, both for speed and to preserve fidelity to the original geometry. However, existing data-parallel direct rendering approaches for NURBS suffer from rendering artifacts at trim boundaries. This talk proposes a solution to address these rendering artifacts that are still preventing wide-scale adoption of all such direct rendering algorithms for trimmed parametric models.