

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

Sessions on **Computer Graphics** (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.

SESSIONS

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.

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.

S0609 - Computational Graphics: An Overview of Graphics Research at NVIDIA

David Luebke (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:25 pm

Topic Areas: Computer Graphics

Session Level: Intermediate

The future of computer graphics presents many challenges. The worlds we render will be vastly more complex in geometry and artistic “texture”. Real-time rendering will use global illumination to achieve a far richer appearance, robustly. And content creation, which has grown to be the dominant cost of producing both games and film, must get simpler and less expensive. The NVIDIA Graphics Research group addresses these challenges with a focus on “Computational Graphics”: using general-purpose computation to enhance and extend the traditional pipelines and capabilities of real-time rendering. In this talk David Luebke, who leads graphics research, will give an overview of recent and ongoing work in computational graphics at NVIDIA Research.

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.

S0610 - Octree-Based Sparse Voxelization for Real-Time Global Illumination

Cyril Crassin (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 2:30 pm - 2:55 pm

Topic Areas: Computer Graphics

Session Level: Intermediate

Discrete voxel representations are generating growing interest in a wide range of applications in computational sciences and particularly in computer graphics. A new real-time usage of dynamic voxelization inside a sparse voxel octree is to compute voxel-based global illumination. When used in real-time contexts, it becomes critical to achieve fast 3D scan conversion (also called voxelization) of traditional triangle-based surface representations. This talk describes a new surface voxelization algorithm that produces a sparse voxel representation of a triangle mesh scene in the form of an octree structure using the GPU hardware rasterizer. In order to scale to very large scenes, our approach avoids relying on an intermediate full regular grid to build the structure and constructs the octree directly.

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.

S0611 - Edge-Aware Shaders for Real-Time Computer Graphics**Peter-Pike Sloan (NVIDIA)****Day:** Tuesday, 05/15 | **Time:** 3:00 pm - 3:25 pm**Topic Areas:** Computer Graphics**Session Level:** Intermediate

The most common approach in rendering is to define behavior at a point in terms of material properties and incident illumination. That approach works well when the geometry and material properties are well-known, and the light physics are simulated accurately. We present a technique to help situations where the model and/or physics is incomplete. This technique augments shaders with information about nearby edges, such as corners and boundaries between materials, and makes it natural to add richness procedurally near these visually critical regions.

S0403 - NURBS Tessellation with CUDA**Brent Oster (NVIDIA)****Day:** Tuesday, 05/15 | **Time:** 3:00 pm - 3:25 pm**Topic Areas:** Computer Graphics**Session Level:** Advanced

NURBS, or Non Uniform Rational B Splines, are a curved surface representation commonly used in computer aided design and digital content creation. This recursive representation gives a great deal of flexibility, allowing arbitrary surface order and knot vectors, enabling a single NURBS surface to contain many contiguous patches. However, this recursive representation is also expensive to compute, so a NURBS surface is often converted into multiple Bezier patches before being tessellated. In this implementation, we present an efficient method for directly tessellating NURBS surfaces using the NVIDIA CUDA computing API.

S0366 - OptiX Out-of-Core and CPU Rendering**David McAllister (NVIDIA, OptiX group)****Day:** Tuesday, 05/15 | **Time:** 3:30 pm - 4:20 pm**Topic Areas:** Ray Tracing; Computer Graphics**Session Level:** Intermediate

OptiX has broken some major barriers recently by enabling out-of-GPU-core memory rendering and by adding a CPU rendering back-end when an OptiX-capable GPU is not present in the system. OptiX users and CUDA developers will be interested in how we accomplished these feats within the existing GPU architecture. This talk will provide a brief introduction to OptiX and then dive into what the new features provide. We will then go under the covers and show how we pulled it off.

S0409 - Stochastic Rasterization**Eric Enderton (NVIDIA), Morgan McGuire (NVIDIA and Williams College)****Day:** Tuesday, 05/15 | **Time:** 3:30 pm - 4:20 pm**Topic Areas:** Computer Graphics; Digital Content Creation & Film**Session Level:** Intermediate

Learn how to render transparency, motion blur, and depth of field effects in real time using random sampling. These effects combine multiple objects in each pixel, making them expensive to compute directly. But recent research shows that, with stratified sampling and clever reconstruction, good image quality can be achieved with surprisingly small numbers of samples per pixel. We will explain how to do this on the GPU, and explore trade-offs of performance, quality, accuracy, and noise.

S0444 - Explore New Techniques in Volume Rendering/Segmentation with Open Inventor**Mike Heck (VSG)****Day:** Tuesday, 05/15 | **Time:** 3:30 pm - 3:55 pm**Topic Areas:** Computer Graphics**Session Level:** Advanced

The goal of this session is to show the improvements in quality, performance and flexibility of the volume rendering implementation of Open Inventor. The latest GPU techniques, such as virtual textures and ray casting, have been combined into a flexible shader API and applied on out of core data. The techniques of volume rendering, sugarcube rendering, basic and complex clipping, sculpting, editing and segmentation will be demonstrated using examples from a geobody extraction workflow. The great ease and flexibility of the shader pipeline API will be illustrated, and we will discuss the broad future perspectives of that technology.

S0261 - Scalable GPU Computing Service Architecture**Henrik Høj Madsen (LEGO), Michael Schøler (LEGO)****Day:** Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm**Topic Areas:** Cloud Computing; Computer Graphics; Ray Tracing**Session Level:** Intermediate

In this session we describe our GPU accelerated computing service which supports several internal business processes in a large scale company setup. The service supports diverse computational needs such as on-demand rendering, mesh optimization, a Massive Multiplayer Online Game (MMO), product visualizations and other demanding computational tasks. We present the architectural considerations for a service-oriented computational framework and the practical learning's and opportunities encountered during development a enterprise system using NVIDIA technologies such as CUDA, OptiX, OpenGL and OpenCL. Our aim is to share knowledge and present LEGO's vision for a GPU accelerated computational platform as a business-driven technology.

S0638 - Lenovo ThinkStation Accelerates Medical Research with Beckman Coulter (Presented by Lenovo)**Scott Ruppert (Lenovo), Tanmay Dharmadhikari (Beckman-Coulter)****Day:** Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm**Topic Areas:** Computer Graphics; Life Sciences**Session Level:** Intermediate

Lenovo ThinkStations utilize Nvidia Maximus technology to accelerate mission critical applications across multiple industries, including manufacturing, media & entertainment, and Life Sciences. Discover how GPUs are used to accelerate medical research from product experts with Lenovo and Beckman Coulter. Beckman Coulter has utilized Nvidia GPUs to reduce software development and test cycles by 50% with their Kaluza software. Kaluza is

a revolutionary flow cytometry analysis software solution that provides visualization tools, speed and an innovative simplicity to the flow community. See how Kaluza allows users to analyze 10 million cells in real time. Session attendees will receive a drawing entry to win a brand new ThinkPad Tablet.

S0021 - OptiX for DirectX Programmers - Eve Online's GPU-Raytraced Portraits
Bert Peers (CCP Games)

Day: Tuesday, 05/15 | **Time:** 4:30 pm - 4:55 pm

Topic Areas: Ray Tracing; Computer Graphics; Application Design & Porting Techniques

Session Level: Intermediate

By integrating NVIDIA's OptiX system for real-time GPU raytracing into a DirectX9 based engine, CCP Games enables high-quality raytraced player portraits for the single shard MMO Eve Online, reusing the game's assets and pipeline. We selectively add stochastic effects while closely maintaining the look of the DX9-based renderer that Art Direction aimed for. In this talk we approach OptiX from the point of view of a programmer familiar with DirectX, discuss integrating these two systems, and show how we reproduced some DirectX-based effects like transparency and subsurface scattering within OptiX.

S0612 - Parallel Quasi-Monte Carlo Sampling for Graphics

Alexander Keller (NVIDIA ARC)

Day: Tuesday, 05/15 | **Time:** 4:30 pm - 4:55 pm

Topic Areas: Computer Graphics

Session Level: Intermediate

Quasi-Monte Carlo methods are widely used in the valuation of financial derivatives and computer graphics. Such applications profoundly benefit from the performance advantages that quasi-Monte Carlo methods offer over Monte Carlo methods. At the core of such methods are simple algorithms generating quasi-Monte Carlo points, which are much more uniformly distributed than (pseudo-) random number ever can be. The structure and deterministic nature of quasi-Monte Carlo points allows for communication avoiding massively parallel deterministic algorithms that especially benefit from GPU architecture. In a self-contained tutorial we introduce the basics and principles of massively parallel quasi-Monte Carlo integration and provide insight and references such that the content directly can be put to practice.

S0267A - Mixing Graphics and Compute with Multiple GPUs

Alina Alt (NVIDIA)

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

Topic Areas: Computer Graphics; Application Design & Porting Techniques

Session Level: Beginner

In this session we will cover all the different aspects of interaction between graphics and compute. The first part of the session will focus on compute API interoperability with OpenGL (using CUDA and OpenCL APIs), while the second part of the session will delve into interoperability at a system level. In particular we will go through the challenges and benefits of dedicating one GPU for compute and another for graphics, how different system configurations affect data transfer between two GPUs, and how it translates into application design decisions helping to enable an efficient, cross-GPU interoperability between compute and graphics contexts. This talk is repeated on Thursday at 3:30 PM (session S0267B)

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.

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.

S0142 - VMD: High Performance Molecular Visualization and Analysis on GPUs

John Stone (University of Illinois at Urbana-Champaign)

Day: Wednesday, 05/16 | Time: 2:00 pm - 2:50 pm

Topic Areas: Molecular Dynamics; Algorithms & Numerical Techniques; Computer Graphics

Session Level: Intermediate

This talk will present recent successes in the use of GPUs to accelerate interactive molecular visualization and analysis tasks on desktop computers, and batch-mode simulation and analysis jobs on GPU-accelerated HPC clusters. We'll present Fermi-specific algorithms and optimizations and compare with those for other devices. We'll also present performance and performance/watt results for VMD analysis calculations on GPU clusters, and conclude with a discussion of ongoing work and future opportunities for GPU acceleration, particularly as applied to the analysis of petascale simulations of large biomolecular complexes and long simulation timescales.

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.

S0281 - Accelerate a Fully Functional Photo Editing Software with GPU**Kaiyong Zhao (Hong Kong Baptist University)****Day:** Wednesday, 05/16 | **Time:** 3:00 pm - 3:25 pm**Topic Areas:** Computational Photography; Computer Graphics**Session Level:** Beginner

Introduce how to design a fully functional GPU-based photo editing software, which provides features like layering and selecting, and integrates various adjusting tools and image filters. This design contains a fast layer rendering engine, an image filter framework which manages different filters supporting visual feedback for filter parameter adjustment. We will also introduce how to design undoing system for GPU-based image processing software. Specifically a CUDA-accelerated HDR tool will be presented in detailed.

S0410 - Computing Hausdorff Distances between Freeforms on the GPU**Sara McMains (UC Berkeley), Adarsh Krishnamurthy (UC San Diego)****Day:** Wednesday, 05/16 | **Time:** 5:00 pm - 5:25 pm**Topic Areas:** Algorithms & Numerical Techniques; Computer Graphics; Computer Vision**Session Level:** Intermediate

We present new GPU algorithms for computing the directed Hausdorff distance between freeform surfaces, with applications in shape matching, mesh simplification, and geometric approximation and optimization. Our algorithms run in real-time with very small error bounds for parametric models defined by complex NURBS surfaces and can be used to interactively compute the Hausdorff distance for models made of dynamic deformable surfaces. We discuss implementation decisions and tradeoffs between OpenGL, Cuda, and Thrust, and the advantages and disadvantages of parallel hierarchical culling methods for this application.

S0096 - Summed Area Ripmaps**Gernot Ziegler (NVIDIA)****Day:** Wednesday, 05/16**Time:** 5:30 pm - 5:55 pm**Topic Areas:** Algorithms & Numerical Techniques; Computer Vision; Computer Graphics**Session Level:** Intermediate

In this presentation, we show how ripmaps can replace Summed Area Tables (SATs) for the purpose of computing a large number of spatially varying box filter kernels throughout the input data, providing both higher accuracy and higher speed for typical use cases. For this purpose, we demonstrate an implementation of ripmap generation in CUDA C (accelerated by shared memory usage), and a texture-cache based box filter for spatially varying kernel sizes, which can be implemented in both CUDA C and graphics-based APIs (OpenGL and DirectX).

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.

S0619 - Hate to Wait? Flash Memory for Full-Throttle GPU Acceleration (Presented by Fusion-io)

Vincent Brisebois (Fusion-io), Robert Wipfel (Fusion-io)

Day: Thursday, 05/17 | Time: 9:00 am - 9:50 am

Topic Areas: Digital Content Creation & Film; Computer Graphics

Session Level: Intermediate

Are you guilty of ever not trying out an idea because of the time it would take to process the effect? With flash memory throttling your system like jet fuel for your GPU, you can finally make sluggish application performance a bad memory. This session will couple a technical overview of the latest in PCIe-attached flash memory technology for accelerating graphics processing with developer best practices and tuning for GPU applications using flash memory for image compositing, editing, video playback, 3D content creation, video capture and many other data-intensive tasks.

S0267B - Mixing Graphics and Compute with Multiple GPUs (Repeat Presentation)

Alina Alt (NVIDIA)

Day: Thursday, 05/17 | Time: 3:30 pm - 4:20 pm

Topic Areas: Computer Graphics; Application Design & Porting Techniques

Session Level: Beginner

In this session we will cover all the different aspects of interaction between graphics and compute. The first part of the session will focus on compute API interoperability with OpenGL (using CUDA and OpenCL APIs), while the second part of the session will delve into interoperability at a system level. In particular we will go through the challenges and benefits of dedicating one GPU for compute and another for graphics, how different system configurations affect data transfer between two GPUs, and how it translates into application design decisions helping to enable an efficient, cross-GPU interoperability between compute and graphics contexts. This talk is repeated on Tuesday at 5:00 PM (session S0267A)

S0138 - GPU Task-Parallelism: Primitives and Applications

Stanley Tzeng (University of California, Davis), Anjul Patney (University of California, Davis)

Day: Thursday, 05/17 | Time: 3:30 pm - 3:55 pm

Topic Areas: Application Design & Porting Techniques; Development Tools & Libraries; Computer Graphics

Session Level: Intermediate

We explore how a task-parallel model can be implemented on the GPU and address concerns and programming techniques for doing so. We discuss the primitives for building a task-parallel system on the GPU. This includes novel ideas for mapping tasking systems onto the GPU including task granularity, load balancing, memory management, and dependency resolution. We also present several applications which demonstrate how a task-parallel model is more suitable than the regular data parallel model. These applications include a Reyes renderer, tiled deferred lighting renderer, and a video encoding demo.

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.