

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

Sessions on **Audio, Image & Video Processing** (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 - 10:30 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

S0395 - GPU Enablement in Adobe Photoshop

Jerry Harris (Adobe Systems), Jeff Chien (Adobe Systems)

Day: Tuesday, 05/15 | **Time:** 9:00 am - 9:50 am

Topic Areas: Digital Content Creation & Film; Audio, Image and Video Processing

Session Level: Beginner

Photoshop is one of the most popular products in history. It attempts to delight the customers with an immersive experience. Since CS4, Adobe has been tapping into the horsepower of the GPU to create a compelling playground for the imaginations of creative pros. Please join us to review the latest developments on how GPUs have been an enabling force.

S0404 - Computer Vision Libraries with GPUs

Eric Young (NVIDIA)

Day: Tuesday, 05/15 | **Time:** 9:30 am - 10:20 am

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

Session Level: Beginner

Learn how Computer Vision libraries can take advantage of GPUs. Computer Vision algorithms are extremely well suited for GPU architectures because they demand large computational power that GPUs offer over CPUs. This talk provides an overview of the different GPU libraries such as (OpenCV, GPUCV, PCL, and NPP Libraries) and

online resources (GPU4Vision and OpenNVIDIA) available for developers today. Examples and demonstrations of practical applications making use of these libraries will also be shown throughout the talk.

S0255 - Telecom Systems Simulations Acceleration via CPU/GPU Co-Processing: Turbo Codes Case Study

Paolo Spallaccini (Ericsson), Stefano Chinnici (Ericsson)

Day: Tuesday, 05/15 | **Time:** 10:00 am - 10:25 am

Topic Areas: Algorithms & Numerical Techniques; Audio, Image and Video Processing; Supercomputing

Session Level: Intermediate

Learn how the struggle for acceleration of simulations of a Serially Concatenated turbo code (SCCC) led to the knowledge of new techniques applicable to a broad range of non-natively parallel physical layer telecommunication systems simulations. The overall architectural features of CUDA became inspiring for newer parallelization techniques involving algorithm engineering; the simulation acceleration attained for iterative SCCC Decoder represents an example of efficiency of leveraging on heterogeneous GPU-CPU coprocessing concepts. The registrants will deep dive into data sets and tasks organization strategies as well as into results and insights, all widely presented and discussed.

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.

S0075 - Oculus Real-Time Modular Cognitive Vision System

Jeremie Papon (University of Gottingen), Alexey Abramov (University of Gottingen)

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

Topic Areas: Computer Vision; Audio, Image and Video Processing; Application Design & Porting Techniques; Machine Vision

Session Level: Intermediate

This session will explore ways to integrate GPU processing into a real-time computer vision architecture. While there has been a rapid push to move vision algorithms onto GPUs, integration into an efficient vision system architecture remains elusive. We will discuss our development of a modular vision system architecture that enables rapid prototyping of complex pipelines using multiple GPUs. The system incorporates modules for segmentation, disparity mapping, optical flow and particle filter tracking on the GPU. Our talk will explore the various difficulties associated with developing such a system and will give a hands-on demonstration of Oculus, our vision platform.

S0223 - Rapid Training of Acoustic Models Using GPUs**Jike Chong (Carnegie Mellon University), Ian Lane (Carnegie Mellon University Co)****Day:** Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm**Topic Areas:** Audio, Image and Video Processing; Machine Learning & AI**Session Level:** Intermediate

Learn how to realize robust and accurate speech recognition systems by training acoustic models on GPUs. For common languages, state-of-the-art systems are now trained on thousands of hours of speech data, which can take weeks even with a large cluster of machines. To overcome this development bottleneck, we propose a new framework for rapid training of acoustic models using highly parallel GPUs. With a single NVIDIA GTX580 GPU, our proposed approach is shown to be 51x faster than a sequential CPU implementation, enabling a moderately sized acoustic model to be trained on 1000-hour speech data in just over 9 hours.

S0049 - Using the GPU Direct for Video API**Thomas True (NVIDIA), Alina Alt (NVIDIA)****Day:** Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm**Topic Areas:** Audio, Image and Video Processing; Development Tools & Libraries; Digital Content Creation & Film; Machine Vision**Session Level:** Advanced

This tutorial will demonstrate how video I/O devices can take advantage of the GPU Direct for Video API to optimize the data transfer performance for digital video, film and broadcast applications and computer vision applications. The GPU Direct for Video API is a technology that permits the DMA transfer of data buffers between video I/O devices and the GPU through the use of a shared system memory buffer for immediate processing by OpenGL, DirectX, CUDA and OpenCL. This direct transfer can improve synchronization and eliminate latency between video capture, GPU processing and video output.

S0253 - Sensor Processing with Rugged Next-Generation NVIDIA GPUs (Presented by GE Intelligent Platforms)**Dustin Franklin (GE Intelligent Platforms)****Day:** Wednesday, 05/16 | **Time:** 9:00 am - 9:50 am**Topic Areas:** Audio, Image and Video Processing; General Interest; Machine Vision; Computer Vision**Session Level:** Intermediate

Swimming in sensors and drowning in data? Turn the tide on high-bandwidth sensors with rugged next-generation GPUs from NVIDIA. See how we deploy NVIDIA GPUs into the most extreme of environments, providing GPGPU capabilities onboard platforms where SWaP and GFLOPS/watt is key. Dig into four realtime CUDA sensor processing applications - Hyperspectral Imaging, Wide-Area Surveillance, 360° Situational Awareness, and GSM cellular SIGINT. Discuss the CUDA algorithms, interconnects, and rugged platforms behind each. Learn how we utilize GPUDirect and realtime Linux for improved latency and determinism.

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 speedup the analysis and to visualize the results.

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.

S0335 - Live 3D-Video with a Lightfield Camera**Christian Perwass (Raytrix GmbH)****Day:** Wednesday, 05/16 | **Time:** 2:00 pm - 2:50 pm**Topic Areas:** Computational Photography; Audio, Image and Video Processing; Stereoscopic 3D; Computer Vision**Session Level:** Beginner

In this session you will learn what a lightfield camera is, how it works and what you can do with it. Next to the theoretical presentation we give a live demo of the camera system developed by our company Raytrix that gives you 3D live video from a single camera through a single lens currently at up to 10fps with a maximum effective resolution of 3 megapixels synthesized from an 11 megapixel sensor using CUDA algorithms on a GTX580. Post-production features include pixel-wise focusing, depth zoom, variable stereo base-line and base-line rotation.

S0259 - A High Performance Platform for Real-Time X-Ray Imaging**Suren Chilingaryan (Karlsruhe Institute of Technology)****Day:** Wednesday, 05/16 | **Time:** 3:00 pm - 3:25 pm**Topic Areas:** General Interest; Supercomputing; Audio, Image and Video Processing; Algorithms & Numerical Techniques**Session Level:** Intermediate

We will share our experience on development of the GPU-based platform for synchrotron-based X-ray imaging aimed to analysis of dynamic processes. The complete data flow from the camera to the data storage will be discussed with a special focus on I/O issues, hardware platform, and ways to utilize the available system resources. An efficient GPU-implementation of filtered back projection will be presented highlighting differences of implementations for GT200, Fermi, and AMD Cypress architectures. We will introduce our software platform used to abstract current configuration of the imaging station and to simplify the development of parallel image processing algorithms.

S0073 - Cost-effective GPU Acceleration of a Video Restoration and Archiving Workflow**Klaus Gaedke (Technicolor)****Day:** Wednesday, 05/16 | **Time:** 3:30 pm - 3:55 pm**Topic Areas:** Audio, Image and Video Processing**Session Level:** Intermediate

The goal of this session is to present a complex GPU-accelerated video restoration and archiving workflow. The workflow consists of many different processing steps and a final review application. Fast and cost-effective processing and real-time display of the processed video material is a key requirement. It will be shown in detail how a GPU based acceleration can be achieved for many different processing steps and the review application based on the use of OpenCV, OpenCL, and OpenGL. Furthermore, an object oriented software architecture supporting the acceleration of several different processing tasks on the same graphics adapter will be presented.

S0131 - Multi-GPU Real-Time Ptychographic X-Ray Image Reconstruction**Filipe Maia (Lawrence Berkeley National Laboratory)****Day:** Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm**Topic Areas:** Audio, Image and Video Processing; Algorithms & Numerical Techniques**Session Level:** Intermediate

Learn how a new imaging technique, combined with the computational power of GPUs and the brightness of modern X-ray synchrotrons can quickly and easily produce images with nanometer level resolution. Ptychography is a recent X-ray imaging technique in which overlapping regions of a sample are exposed in quick succession and the resulting scattering is used to reconstruct a high resolution image of the sample. Discover why GPUs can substitute for the lack of X-ray lenses and how they enabled a dramatic reduction in the feedback time for users of the technique from days to seconds.

S0273 - Fast JPEG Coding on the GPU**Fyodor Serzhenko (Fastvideo)****Day:** Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm**Topic Areas:** Audio, Image and Video Processing; Algorithms & Numerical Techniques**Session Level:** Advanced

The goal of this session is to demonstrate how high speed JPEG compression and decompression can be efficiently implemented on the GPU using CUDA. In this session we will present: detailed analysis of Baseline JPEG compression and decompression processes and its constituent parts (such as Huffman Coding, RLE, Differential Coding, Quantization, Discrete Cosine Transform) and their suitability for the GPU architecture, analysis of achieved results and comparison with existing implementations, applications to high-speed imaging.

S0052 - Fast High Quality Image and Video Background Removal with CUDA**Timo Stich (NVIDIA)****Day:** Wednesday, 05/16 | **Time:** 4:30 pm - 4:55 pm**Topic Areas:** Audio, Image and Video Processing; Machine Learning & AI**Session Level:** Intermediate**S0133 - Improving Mars Rover Image Compression Via GPUs And Genetic Algorithms****Brendan Babb (University of Alaska Anchorage), Frank Moore (University of Alaska, Anchorage)****Day:** Thursday, 05/17 | **Time:** 9:00 am - 9:25 am**Topic Areas:** Machine Learning & AI; Audio, Image and Video Processing; Development Tools & Libraries**Session Level:** Beginner

Learn how to use Jacket to accelerate genetic algorithm (GA) image compression. Our research uses a GA to optimize lossy compression transforms that outperform state-of-the-art wavelet-based approaches for a variety of image classes, including fingerprints, satellite, medical, and images transmitted from the Mars Exploration Rovers. A typical training run evolves a population of transforms over many generations; since each transform must be applied to each image from the training set, each run entails thousands of independent, parallelizable fitness evaluations. By using MATLAB, and Jacket to perform 2D convolution on the GPU, we have greatly reduced the total computation time needed.

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.

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.

S0279 - Optimization Techniques for GPU and GPP**Lionel Lacassagne (Institute for Fundamental Electronics), Antoine Pedron (Institute for Fundamental Electronics)****Day:** Thursday, 05/17 | **Time:** 2:00 pm - 2:25 pm**Topic Areas:** Application Design & Porting Techniques; Computer Vision; Audio, Image and Video Processing**Session Level:** Intermediate

We present and evaluate optimizations techniques for GPU and GPP. We describe High level transforms targeting algorithm refactoring and low-level optimizations targeting hardware. The algorithm used to present these optimizations comes from image processing but is also representative of algorithms of different areas with same kind of local computations like stencils in computer science or finite difference methods in numerical analysis. We evaluate the impact of optimization on three generations of GPU and GPP, both mobile and desktop. We show that significant speedup can be achieved on GPP but also on GPU. Finally we compare together these two architectures.

S0157 - A Study of Persistent Threads Style Programming Model for GPU Computing**Kshitij Gupta (UC Davis), Jeff Stuart (UC Davis)****Day:** Thursday, 05/17 | **Time:** 3:00 pm - 3:50 pm**Topic Areas:** Parallel Programming Languages & Compilers; Audio, Image and Video Processing**Session Level:** Advanced

We present the usefulness of a new style of GPU programming called Persistent Threads, known to be useful on irregular workloads. First, we will begin by formally defining the PT model. We will then categorize use of PT into four use cases and present micro-benchmark analyses of when this model is useful over traditional kernel



formulations. Third, we will show a full speech recognition application that uses all four PT use cases. Finally, we will conclude our talk by suggesting appropriate modifications to GPU hardware, software, and APIs that make PT kernels both easier to implement and more efficient.