

GPU Technology Conference 2010 Sessions on Imaging (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.

2298 - Accelerated Image Quality Assessment using Structural Similarity

Explores the GPU porting and performance analysis of the image quality assessment algorithm based on structural similarity index (SSIM). This index is a powerful tool for image quality assessment and the algorithm is highly suitable for GPU architecture, offering a rapid image quality assessment in many image restoration applications.

Speaker: Mahesh Khadtare, Computational Research Laboratories, Pune, INDIA.

Topics: Computer Vision, Imaging

Time: Thursday, September, 23rd, 11:30 - 11:50

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

2095 - Building High Density Real-Time Video Processing Systems

Learn how GPU Direct can be used to effectively build real time, high performance, cost effective video processing products. We will focus especially on how to optimize bus throughput while keeping CPU load and latency minimal.

Speaker: Ronny Dewaele, Barco

Topics: Video Processing, Imaging

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

2093 - Computational Photography: Real-Time Plenoptic Rendering

Get the latest information on GPU-based plenoptic rendering including a demonstration of refocusing, novel view generation, polarization, high dynamic range, and stereo 3D. Learn how GPU hardware enables plenoptic rendering tasks with high-resolution imagery to be performed interactively, opening up entirely new possibilities for modern photography.

Speakers: Andrew Lumsdaine, Indiana University, Georgi Chunev, Indiana University, Todor Georgiev, Adobe Systems

Topics: Imaging, Computer Vision, Stereoscopic 3D

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

2043 - Disparity Map Generation

Explore the algorithms and implementation of disparity maps on the GPU. We will discuss how a disparity map facilitates stereoscopic content creation, applications and approaches tried, and final results of real time calculations on GPUs.

Speaker: Henry Gu, GIC

Topics: Stereoscopic 3D, Computer Vision, Imaging

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

2021 - Efficient Volume Segmentation on the GPU

Explore a new technique in the detection of common regions in a 2D/3D data array. Connected components along the axes are linked before actual label propagation starts. The algorithm is completely gather-based, which allows for several optimizations in the CUDA C implementation. It enables real-time frame rates for the analysis of typical 2D images and interactive frame rates for the analysis of typical volume data.

Speakers: Allan Rasmusson, University of Aarhus (NVIDIA intern), Gernot Ziegler, NVIDIA

Topics: Algorithms & Numerical Techniques, Computer Vision, Imaging, Medical Imaging & Visualization

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

4007 - Emerging Companies: CEO on Stage featuring Aqumin, RTT, and Scalable Display

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature Aqumin, RTT, and Scalable Display - covering the fields of finance, computer graphics, and imaging.

Panelists will include Norman Winarsky (SRI), Savitha Srinivasan (IBM), and Rob Enderle (Enderle Group).

Speakers: Norman Winarsky, SRI, Andrew Jamison, Scalable Display Technologies, Savitha Srinivasan, IBM, Rob Enderle, Enderle Group, Jeroen Snepvangers, Michael Zeitlin, Aqumin

Topics: General Interest, Finance, Imaging, Computer Graphics

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

4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel, and Total Immersion

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature Cinnafilm, Perceptive Pixel, and Total Immersion - covering film, imaging, and computer vision.

Panelists will include Bill Tai (CRV), Paul Weiskopf (Adobe), and Tim Bjarin (Creative Strategies).

Speakers: Bill Tai, Charles River Ventures, Lance Maurer, Cinnafilm, Inc., Bruno Uzzan, Total Immersion, Paul Weiskopf, Adobe, Tim Bjarin, Creative Strategies

Topics: General Interest, Computer Vision, Film, Imaging

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

2215 - Extending OpenCV with GPU Acceleration

OpenCV is a widely popular computer vision library, with millions of downloads and hundreds of thousands of users. Applications span many industries including robotics, industrial machine vision, automotive, film & broadcast, medical, and consumer applications. NVIDIA and the OpenCV development team are collaborating to provide CUDA implementations of the most demanding algorithms, thus enabling a new level of real-time capability and higher quality results.

This talk will introduce OpenCV, and summarize the new CUDA enabled capabilities, and provide an overview of future plans.

Speaker: Joe Stam, NVIDIA

Topics: Computer Vision, Imaging, Stereoscopic 3D, Video Processing

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

2087 - Fast High-Quality Panorama Stitching

We present a panorama stitching application implemented with CUDA C on the GPU. The image processing pipeline consists of SIFT feature detection and matching and Graphcut image stitching to achieve high-quality results. We demonstrate live panorama creation with a Webcam.

Speaker: Timo Stich, NVIDIA

Topics: Video Processing, Algorithms & Numerical Techniques, Computer Vision, Imaging

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

2051 - GPGPU in Commercial Software: Lessons From Three Cycles of the Adobe Creative Suite

Learn about leveraging GPUs for commercial software. We will discuss lessons learned creating and using the Adobe Image Foundation libraries to accelerate image and video processing using

GPUs and multi-core. These libraries are used by most of Adobe's applications as well as integrated by hobbyist and professional applications with different levels of experience with GPUs and diverse user bases.

Speaker: Kevin Goldsmith, Adobe Systems, Incorporated

Topics: Imaging, Video Processing

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

2224 - GPU Acceleration in Adobe Creative Tools

Hear experts explain how Adobe Creative Suite 5 harnesses the power of CUDA technology in several of its core software applications. We will focus on the complete redesign of the core video playback and rendering engine in Adobe Premiere Pro CS5 and how it uses the power of GPUs to deliver superior performance and change the game for Adobe in professional video production.

Topics: Video Processing, Imaging

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

2020 - GPU-Accelerated Data Expansion for the Marching Cubes Algorithm

Learn how to accelerate marching cubes on the GPU by taking advantage of the GPU's high memory bandwidth and fast on-chip shared memory in a data expansion algorithm that can extract the complete iso-surface mesh from (dynamic) volume data without requiring any data transfers back to the CPU.

Speakers: Gernot Ziegler, NVIDIA, Chris Dyken, SINTEF

Topics: Algorithms & Numerical Techniques, Imaging, Medical Imaging & Visualization

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

2282 - GPU-Enabled Biomedical Imaging

The purpose of this presentation is to describe several novel biomedical imaging applications which make extensive use of GPUs. In CT iterative reconstructions, for example, high performance computing is allowing us to see details and structures we previously were not able to discern.

Speaker: Homer Pien, MGH / HMS

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

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

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

2096 - High-Speed CT Reconstruction in Medical Diagnosis & Industrial NDT Applications

We present the software platform CERA developed by Siemens, which utilizes (multiple) graphics processing units (GPUs) in order to deliver high-speed CT reconstructions, and describe its implementation challenges using CUDA and OpenCL. We further show how GPU acceleration enables the utilization of reconstruction approaches which provide highly improved reconstruction quality in NDT applications.

Speaker: Holger Scherl, Siemens AG

Topics: Medical Imaging & Visualization, Imaging

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

2139 - Interactive Histology of Large-Scale Biomedical Image Stacks

Get the latest information on leveraging GPU computing to process and visualize large-scale biomedical image stacks. We will discuss both display-aware processing and GPU-accelerated texture compression for histology applications on the GPU.

Speakers: Won-Ki Jeong, Harvard University, Jens Schneider, King Abdullah University of Science and Technology

Topics: Medical Imaging & Visualization, Imaging, Life Sciences

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

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

2145 - Photo Editing on the GPU with MuseMage

See how MuseMage greatly accelerates image processing and editing while providing real-time feedback by harnessing the power of GPUs. We will discuss the majority of MuseMage tools which are fully implemented on GPUs.

Speakers: Kaiyong Zhao, HKBU, Yubo Zhang, UC Davis

Topic: Imaging

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

2169 - Real-time Volumetric Medical Ultrasound Applications for GPU Computing

Real-time volumetric medical ultrasound requires computationally intensive rapid processing of data for visualization of acquired acoustic data. Clinical applications of GPU-based technologies in obstetrics and cardiology will be discussed.

Speaker: Roez Lazebnik, Siemens Healthcare

Topics: Medical Imaging & Visualization, Imaging, Stereoscopic 3D, Computer Graphics

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

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

2242 - Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion

Ever since the 1827 discovery of Brownian motion by observing pollen grains, quantifying motion under the microscope has led to breakthroughs in physics, biology and engineering. Here, I present methods we have developed using confocal microscopy to deduce 3D structure and dynamics from 2D image sequences. We analyze the motion of diffusing colloidal particles and swarms of bacteria free to swim in 3D, which we observe at the single-organism level. We rely heavily on GPU computing to process our large data sets, making extensive use of NPP, CuFFT and optical-flow CUDA algorithms originally developed for machine vision in automobiles.

Speaker: Peter Lu, Harvard University

Topics: Computer Vision, Imaging, Life Sciences

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

2003 - Using CUDA to Accelerate Radar Image Processing

Come see how current GPU technology provides the means for the first portable real-time radar image processing algorithm. This session will outline how the GPU has afforded nearly three orders of magnitude improvement in performance for Synthetic Aperture Radar's (SAR) hallmark image processing algorithm. We will present algorithm details and further improvements.

Speakers: Richard Carande, Neva Ridge Technologies, Aaron Rogan, Neva Ridge Technologies

Topics: Signal processing, Algorithms & Numerical Techniques, Imaging, Video Processing

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