

GPU Technology Conference, May 14-17, 2012  
McEnergy Convention Center, San Jose, California  
[www.gputechconf.com](http://www.gputechconf.com)

---

## Sessions on **Computer Vision** (subject to change)

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

---

### **S0607 - High Performance 3D Perception**

Chris Slaughter (University of Texas Perception, Lynx Labs)

**Day:** Tuesday, 05/15 | **Time:** 9:00 am - 9:25 am

**Topic Areas:** Computer Vision

**Session Level:** Beginner

The path to general purpose graphics programming was driven by computer graphics: the process of rendering 3d models into 2d viewpoints. With the advent of flexible programming of GPGPU processing, this process can be reversed. 3D perception is the problem of inferring structure and motion of the physical world from 2d and 3d measurements. In this talk, we will demonstrate the role GPGPU plays in a diverse set of applications in high speed 3d perception and discuss optimization of these techniques for the GPGPU. We also demonstrate several capabilities of future systems which are enabled by GPGPU technologies.

### **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.

### **S0088 - Point Cloud Library (PCL) on CUDA**

Radu Rusu (Willow Garage, Inc) , Michael Dixon ( Willow Garage, Inc )

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

**Topic Areas:** Computer Vision; Algorithms & Numerical Techniques; Stereoscopic 3D; Machine Vision

**Session Level:** Intermediate

The Point Cloud Library (PCL - <http://pointclouds.org>) is a large scale, open project for 3D point cloud processing. The PCL framework contains numerous state-of-the art algorithms including filtering, feature estimation, surface reconstruction, registration, model fitting and segmentation. Due to the massively parallel nature of many of the above algorithms, GPGPU accelerations holds great potential for achieving real-time performance in numerous applications. In this work we demonstrate some of the recent advances in GPGPU programming for 3D point cloud processing, and outline plans for future development.

**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.

**S0062 - Histograms of Oriented Gradients with CUDA: Performance Analysis and Optimization Tips**

Anton Obukhov (Consultant)

**Day:** Tuesday, 05/15 | **Time:** 4:00 pm - 4:25 pm**Topic Areas:** Computer Vision; Machine Vision; Development Tools & Libraries**Session Level:** Advanced

Computer Vision is becoming increasingly popular and important nowadays. With the advent of powerful mobile devices and increasing power of desktop PCs, it is important to improve user experience by tackling the hardest problems of real-time interaction with the user. These include body parts tracking, face, and gesture recognition. This talk discusses a well-known Histogram of Oriented Gradients approach to object detection in images and its implementation with CUDA. A detailed performance analysis of different algorithm parts is conducted and optimizations for various usage cases are proposed. The role of OpenCV GPU module is highlighted and implementation details are provided.

**S0104 - GPU Implementation of Deep Learning for Intelligent Computer Vision**

Ben Goertzel (Novamente LLC)

**Day:** Tuesday, 05/15 | **Time:** 4:30 pm - 4:55 pm  
**Topic Areas:** Computer Vision; Algorithms & Numerical Techniques  
**Session Level:** Advanced

Learn how to use GPU supercomputing for intelligent computer vision, via deep learning algorithms. We will focus on a case study of visual object and event recognition in a humanoid robotics context, involving a port to CUDA of the DeSTIN "compositional spatiotemporal deep learning network" vision processing algorithm (originally implemented at the University of Tennessee in Knoxville for conventional serial computers). The audience will learn how to use the open-source DeSTIN CUDA code, and also how to port other deep learning algorithms to CUDA.

### **S0253 - Sensor Processing with Rugged Next-Generation NVIDIA GPUs (Presented by GE Intelligent Platforms)**

Dustin Franklin (GE Intelligent Platforms)

**Day:** Tuesday, 05/15 | **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.

### **S0355 - Seamless Scalable Displays - Using NVIDIA Warp + Intensity API**

Rajeev Surati (Scalable Display Technologies)

**Day:** Tuesday, 05/15 | **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:** Tuesday, 05/15 | **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.

**S0128 - VScreen: a Real-Time augmented video method**

Francisco J. Hernandez--Lopez (CIMAT A.C.) , Mariano Rivera (CIMAT A.C.)

**Day:** Tuesday, 05/15 | **Time:** 5:00 pm - 5:25 pm

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

**Session Level:** Beginner

We are presenting a tool for image editing that allows us to modify a region of any image or video by another image or video. This application is useful for advertisements, commercials, music videos, movies, etc. We named "Virtual Screen", or just VScreen, to our development. The main difference between editing (augmenting) videos and fixed images is that the occlusions need be managed. Moving objects in the foreground may occlude the augmented region in background. So that we use a procedure for foreground-background video segmentation, that is implemented in NVIDIA video cards to fulfill the real-time requirement.

**S0410 - Computing Hausdorff Distances between Freeforms on the GPU**

Sara McMains (UC Berkeley) , Adarsh Krishnamurthy (UC San Diego)

**Day:** Tuesday, 05/15 | **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:** Tuesday, 05/15 | **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 (e.g. OpenGL and DirectX).

**S0279 - Optimization Techniques for GPU and GPP**

Lionel Lacassagne (Institute for Fundamental Electronics) , Antoine Pedron ( Institute for Fundamental Electronics)

**Day:** Tuesday, 05/15 | **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 theses 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.

**S0038 - Designing Killer CUDA Applications for X86, multiGPU, and CPU+GPU**

Robert Farber

**Day:** Tuesday, 05/15 | **Time:** 4:00 pm - 4:25 pm

**Topic Areas:** Machine Learning & AI; Supercomputing; Databases, Data Mining, Business Intelligence; Computer Vision

**Session Level:** Intermediate

CUDA redefined software development with 10 to 1000-times faster GPU applications. Now a single CUDA source tree can support the x86 mass market (no GPU required) and 1/3 billion CUDA-enabled GPUs. MultiGPU and CPU+GPU apps utilize all system resources. GPUdirect, UVA, caches, prefetching, ILP (Instruction level Parallelism), automated analysis tools and more offer ease, capability, and performance. The overall impact on software investment, scalability, balance metrics, programming API, and lifecycle will be considered. Working real-time video and other examples from my book, "CUDA Application Design and Development" provide practical insight to enable augmented reality and your killer apps.