

## **GPU Technology Conference 2010 Sessions on Signal Processing** (subject to change)

*IMPORTANT: Visit [www.nvidia.com/gtc](http://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.*

---

### **2126 - Accelerating Signal Processing: Introduction to GPU VSIPL**

Learn how to use the Vector Signal Image Processing Library to accelerate signal processing applications without needing to understand platform-specific programming and optimization techniques. We will discuss how GPU VSIPL implements the VSIPL API and uses CUDA-capable GPUs to maximize performance of several example applications.

Speaker: Dan Campbell, Georgia Tech Research Institute

Topics: Signal processing, Tools & Libraries

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

### **2066 - Accelerating System Level Signal Integrity Simulation**

Discuss how GPU acceleration for key parts of the ANSYS Nexxim Simulator resulted in significant speedup over multi-core processors. We will cover time consumption and data parallelism exposure considerations, and focus on key areas where GPU acceleration was applied including convolution and Eye rendering.

Speakers: Danil Kirsanov, ANSYS, Ekanathan Palamadai, ANSYS

Topics: Physics Simulation, Algorithms & Numerical Techniques, Signal processing

Time: Thursday, September, 23rd, 16:30 - 16:50

### **2082 - CU-LSP: GPU-based Spectral Analysis of Unevenly Sampled Data**

Standard FFT algorithms cannot be applied to spectral analysis of unevenly sampled data. Alternative approaches scale as  $O(N^2)$ , making them an ideal target for harnessing the raw computing power of GPUs. To this end, I have developed CU-LSP, a CUDA spectral analysis code

based on the Lomb-Scargle periodogram. Preliminary benchmarking indicates impressive speed-ups, on the order of 400 relative to a single core of a modern CPU. An initial application of CU-LSP will be the analysis of time-series data from planet-search and asteroseismology satellites.

Speaker: Richard Townsend, University of Wisconsin-Madison

Topics: Astronomy & Astrophysics, Algorithms & Numerical Techniques, Signal processing

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

## **2175 - Hello GPU: High-Quality, Real-Time Speech Recognition on Embedded GPUs**

In this presentation, we will talk about our experiences of implementing an end-to-end automatic speech recognition system that runs in faster than real-time on embedded GPUs, targeted towards small form-factor consumer devices. Focusing specifically on some of the challenges encountered during the design process, a major portion of our talk will focus on giving insights into modifications we made to well-established speech algorithms to fit well within the GPU programming model. We will show how these changes helped us in realizing a highly optimized system on platforms with limited memory bandwidth and compute resources.

Speaker: Kshitij Gupta, UC Davis

Topics: Embedded & Automotive, Audio Processing, Signal processing, Mobile & Tablet & Phone

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

## **2100 - Hybrid GPU/Multicore Solutions for Large Linear Algebra Problems**

Large linear algebra problems may be solved using recursive block decomposition in which GPUs efficiently compute the sub-blocks and multicore CPUs put the sub-blocks back together within a large shared memory space. This talk will present benchmark results for such a hybrid approach, implemented in Matlab® and using Jacket® to access the GPU compute power.

Speaker: Nolan Davis, SAIC

Topics: High Performance Computing, Algorithms & Numerical Techniques, Signal processing

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

## **2076 - Implementing CUDA Audio Networks**

Learn how to implement a commercial software library that exploits CUDA for audio applications. We focus on the overall threading architecture and the underlying math for implementing general purpose audio processing in CUDA devices. Covers the use of inter-process communication to make a plug-in implementation loadable in 32 bit hosts installed in 64 bit systems, distributing the GPU load on remote servers, and creating a CUDA network for high-end purposes such as a big recording facility.

Speaker: Giancarlo Del Sordo, Acustica Audio

Topics: Audio Processing, Signal processing

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

## **2042 - Interactive 3D Audio Rendering Systems**

Learn how to leverage GPUs for interactive audio rendering. This session will give a short overview of the architecture of current GPUs, emphasizing some key differences between GPU and CPUs programming models for audio processing. We will illustrate the benefits of GPU-accelerated audio rendering with results from 3D audio processing and sound scattering simulations. Finally, we will discuss best practices for GPU implementations as well as future opportunities for audio rendering on massively parallel architectures.

Speaker: Nicolas Tsingos, Dolby Laboratories

Topics: Audio Processing, Ray Tracing, Signal processing

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

## **2116 - Real-time Multichannel Audio Convolution**

Learn how a synthesis of 3D sound scenes can be achieved using a peer-to-peer music streaming environment and GPU. We will discuss the technical and cost benefits to this approach, while noting that it frees the CPU for other tasks.

Speakers: Jose Antonio Belloch, Institute of Telecommunications and Multimedia Applications, Universidad Politecnica de Valencia, Alberto Gonzalez, Universidad Politecnica de Valencia, Antonio M. Vidal, Universidad Politecnica de Valencia

Topics: Audio Processing, Signal processing

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

## **2252 - Simulating Housefly Vision Elements Using OpenCL**

An OpenCL GPU based computer simulation of a biologically motivated model, based on the anatomy of housefly's first optic ganglion, the lamina ganglionaris (the lamina layer) is presented. Specific to GPU technology, the computer model demonstrates: the implementation of a 2nd Order Runge-Kutta method to approximate coupled differential equations using GPU hardware; the mapping of a non-Cartesian coordinate system onto the Cartesian layout of the threads. Testing examined usage and access across device memory spaces to determine the optimal usage/access method for the ANN. This result was generalized for OpenCL GPU devices, using the capabilities of OpenCL.

Speaker: Karen Haines, WASP/The University of Western Australia

Topics: Neuroscience, Algorithms & Numerical Techniques, Signal processing

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

## **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

## **2122 - Using GPUs for Real-Time Brain-Computer Interfaces**

Learn how GPU processing can provide researchers with an inexpensive and versatile alternative to dedicated signal processing hardware for real-time neural prosthetics. Topics will include an overview of algorithms, current state-of-the-art hardware, GPU processing in a real-time environment, multi-platform processing, and future directions in BCIs using GPU processing.

Speaker: Adam Wilson, University of Cincinnati

Topics: Neuroscience, Algorithms & Numerical Techniques, Signal processing

Time: Wednesday, September, 22nd, 15:00 - 15:50