

GPU Technology Conference 2010 Sessions on Audio Processing (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.

2046 - Efficient Automatic Speech Recognition on the GPU

Gain insights into implementation techniques for the speech recognition inference process based on the state-of-art weighted finite state transducer methods. We will discuss the capabilities of the GPU for handling large, irregular graph-based models with millions of states and arcs. We will also present solutions for four challenges in the implementation of speech recognition on the GPU and talk about our more than an order of magnitude faster performance on one core of a CPU.

Speaker: Jake Chong, Parasians, LLC

Topics: Machine Learning & Artificial Intelligence, Algorithms & Numerical Techniques, Audio Processing

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

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

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

[Log in](#)[Log out](#)

