

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

---

---

## Sessions on **Digital Content Creation & Film** (subject to change)

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

---

---

### **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.

### **S0300 - Jet: A Domain-Specific Approach to Parallelism for Film Fluid Simulation**

**Dan Bailey (Double Negative)**

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

**Topic Areas:** Parallel Programming Languages & Compilers; Digital Content Creation & Film; Computational Fluid Dynamics

**Session Level:** Intermediate

Discover how a domain-specific language can not only provide fast parallel performance but a simpler user experience in an environment that highly values flexibility. This talk will present the Jet language and heterogeneous compiler built on the LLVM compiler framework that enables efficient generation of X86 machine code or NVIDIA PTX for stencil computation on structured grids. We show that moving target-specific optimizations upstream into the compiler can greatly improve the ability to manipulate the logic of the solver and thus lower the barrier-to-entry for artists and developers without compromising on performance.

### **S0024 - GPU-Accelerated Path Rendering**

**Mark Kilgard (NVIDIA)**

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

**Topic Areas:** Computer Graphics; GPU Accelerated Internet; Digital Content Creation & Film; Visualization

**Session Level:** Beginner

Standards such as Scalable Vector Graphics (SVG), PostScript, TrueType outline fonts, and immersive web content such as Flash depend on a resolution-independent 2D rendering paradigm that GPUs have not traditionally accelerated. This tutorial explains a new opportunity to greatly accelerate vector graphics, path rendering, and immersive web standards using the GPU. By attending, you will learn how to write OpenGL applications that accelerate the full range of path rendering functionality. Not only will you learn how to render sophisticated 2D graphics with OpenGL, you will learn to mix such resolution-independent 2D rendering with 3D rendering and do so at dynamic, real-time rates.

**S0364 - Interacting with Huge Particle Simulations in Maya with the GPU****Wil Braithwaite (NVIDIA)****Day:** Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm**Topic Areas:** Digital Content Creation & Film; Computational Fluid Dynamics; Visualization**Session Level:** Beginner

We present a plug-in for Maya which enables an artist to simulate huge particle counts in real-time by leveraging the NVIDIA GPU. Being able to interact with the simulation opens up new possibilities for modifying the workflow. We will demonstrate the plug-in, and provide insight into the algorithms used.

**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.

**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.

**S0409 - Stochastic Rasterization****Eric Enderton (NVIDIA), Morgan McGuire (NVIDIA and Williams College)****Day:** Tuesday, 05/15 | **Time:** 3:30 pm - 4:20 pm**Topic Areas:** Computer Graphics; Digital Content Creation & Film**Session Level:** Intermediate

Learn how to render transparency, motion blur, and depth of field effects in real time using random sampling. These effects combine multiple objects in each pixel, making them expensive to compute directly. But recent research shows that, with stratified sampling and clever reconstruction, good image quality can be achieved with surprisingly small numbers of samples per pixel. We will explain how to do this on the GPU, and explore trade-offs of performance, quality, accuracy, and noise.

**S0619 - Hate to Wait? Flash Memory for Full-Throttle GPU Acceleration (*Presented by Fusion-io*)****Vincent Brisebois (Fusion-io), Robert Wipfel (Fusion-io)****Day:** Thursday, 05/17 | **Time:** 9:00 am - 9:50 am**Topic Areas:** Digital Content Creation & Film; Computer Graphics**Session Level:** Intermediate

Are you guilty of ever not trying out an idea because of the time it would take to process the effect? With flash memory throttling your system like jet fuel for your GPU, you can finally make sluggish application performance a bad memory. This session will couple a technical overview of the latest in PCIe-attached flash memory technology for accelerating graphics processing with developer best practices and tuning for GPU applications using flash memory for image compositing, editing, video playback, 3D content creation, video capture and many other data-intensive tasks.

**S0044 - A Massively Parallel Two-Phase Solver for Incompressible Fluids on Multi-GPU Clusters****Peter Zaspel (University of Bonn)****Day:** Thursday, 05/17 | **Time:** 2:00 pm - 2:50 pm**Topic Areas:** Computational Fluid Dynamics; Supercomputing; Algorithms & Numerical Techniques; Digital Content Creation & Film**Session Level:** Intermediate

Join our presentation of a multi-GPU fluid solver for high performance GPU compute clusters. We use high-order scientific techniques to simulate the interaction of two fluids like air and water. Scientists, engineers and even the computer animation industry will profit from the enormous compute power of tens or hundreds of GPUs. A major focus in this talk will be on the applied GPU implementation techniques and the performance results including performance per Watt and performance per dollar results. We also highlight the lessons we learned from porting the complex CPU CFD code NaSt3DGPF to the GPU.