Enabling the Next Generation of Computational Graphics with NVIDIA® Nsight™ Visual Studio Edition

Jeff Kiel

Director, Graphics Developer Tools

NVIDIA®

Computational Graphics Enabled

- Problem: Complexity of Computation Graphics
- Solution: Kepler with Maximus™ & CUDA® 5.0
- How To: NVIDIA® Nsight™ Visual Studio Edition
- Questions?

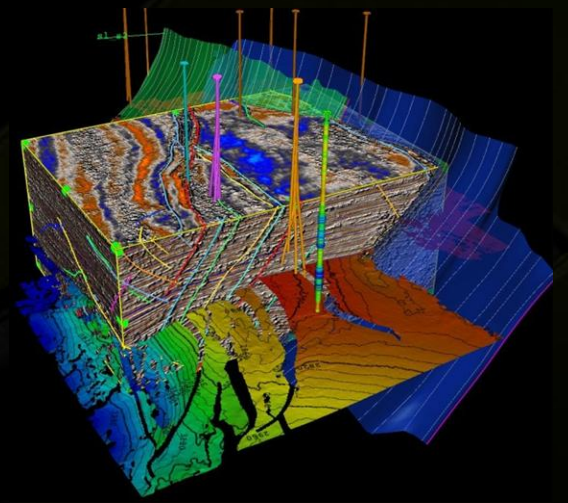

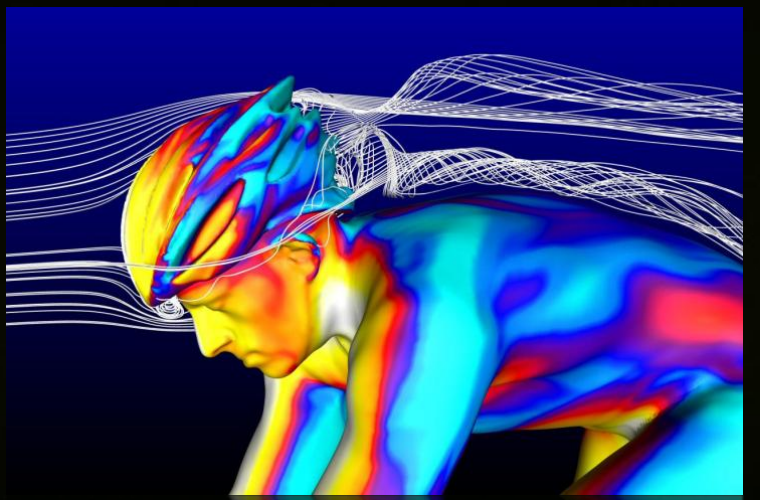

Complexity of Computational Graphics

Real-World Physics During Interactive Design

- ✓ Mechanical Dynamics
- ✓ Fluid Dynamics
- ✓ Light Simulation

Real-World Materials During Interactive Design

- ✓ Physical Properties
- ✓ Photorealism
- ✓ Global Lighting

Complexity of Computational Graphics

Real-World Physics in Scene Design

- ✓ Smoke
- ✓ Fluid Dynamics
- ✓ Light Simulation

High-Performance Parallel Processing

- ✓ Color Correction
- ✓ Effects
- ✓ GPU Rendering

Smoke effects created on NVIDIA Maximus for MIB3 courtesy Sony Pictures

Real-time rendering on NVIDIA Maximus for GDC courtesy Dawnrunner Prod.

Complexity of Computational Graphics

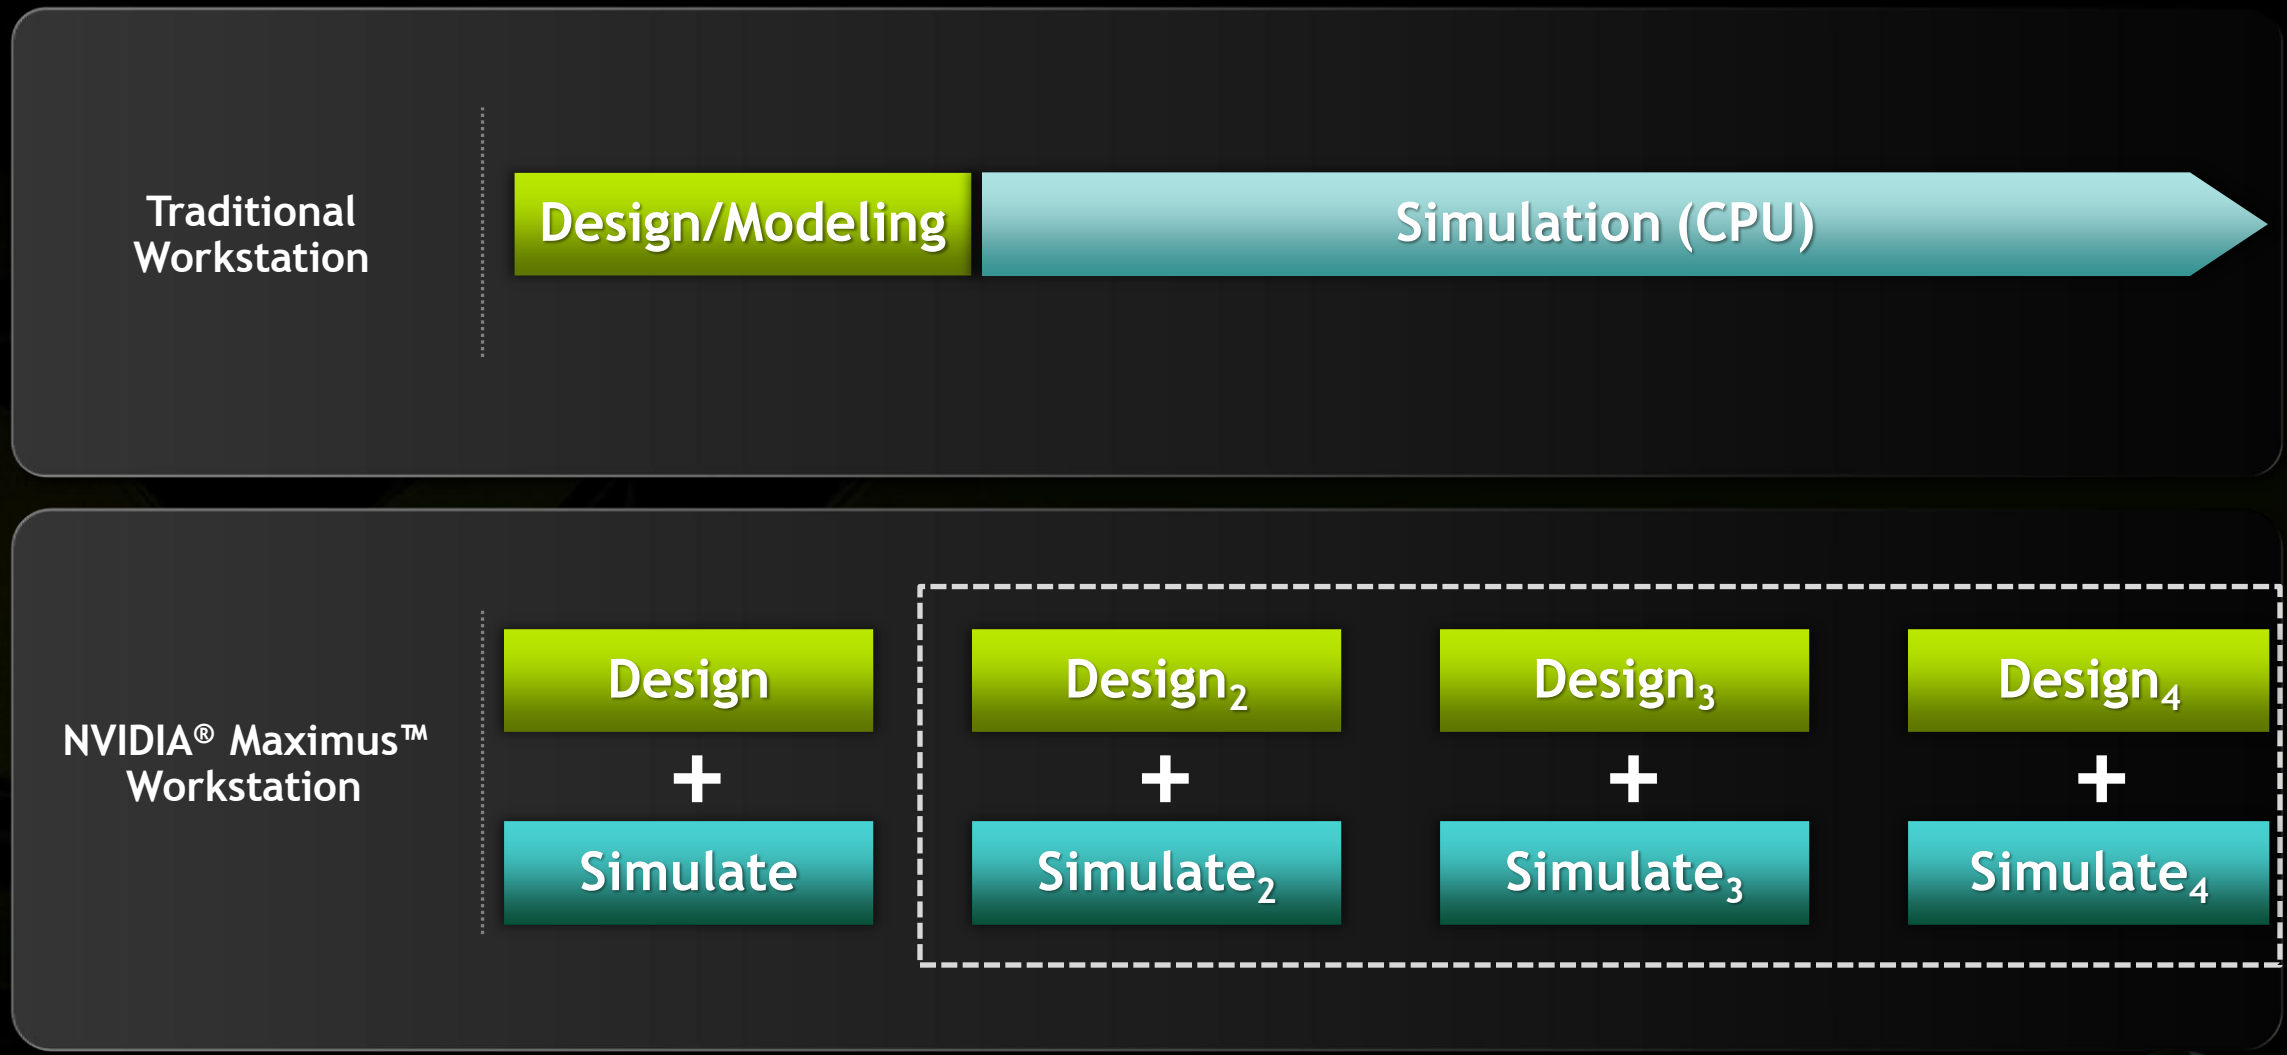

Solution: Kepler, Maximus & CUDA 5.0 …tamed by Nsight

- Kepler Architecture
  - SMX, Dynamic Parallelism, Bindless Textures, FXAA/TXAA, PCIe Gen3, HW H.264 encode, 4 displays
- Maximus
  - Quadro runs graphics (OpenGL/Direct3D)
  - Tesla handles computation
  - Seamlessly supported by over 22 professional applications
- CUDA 5.0
  - Dynamic Parallelism
  - RDMA for GPUDirect
  - GPU Library Object Linking

Quadro K5000
Visualization

+

Tesla K20
Computation

DESIGNED FOR NVIDIA CUDA™

Complexity of Computational Graphics

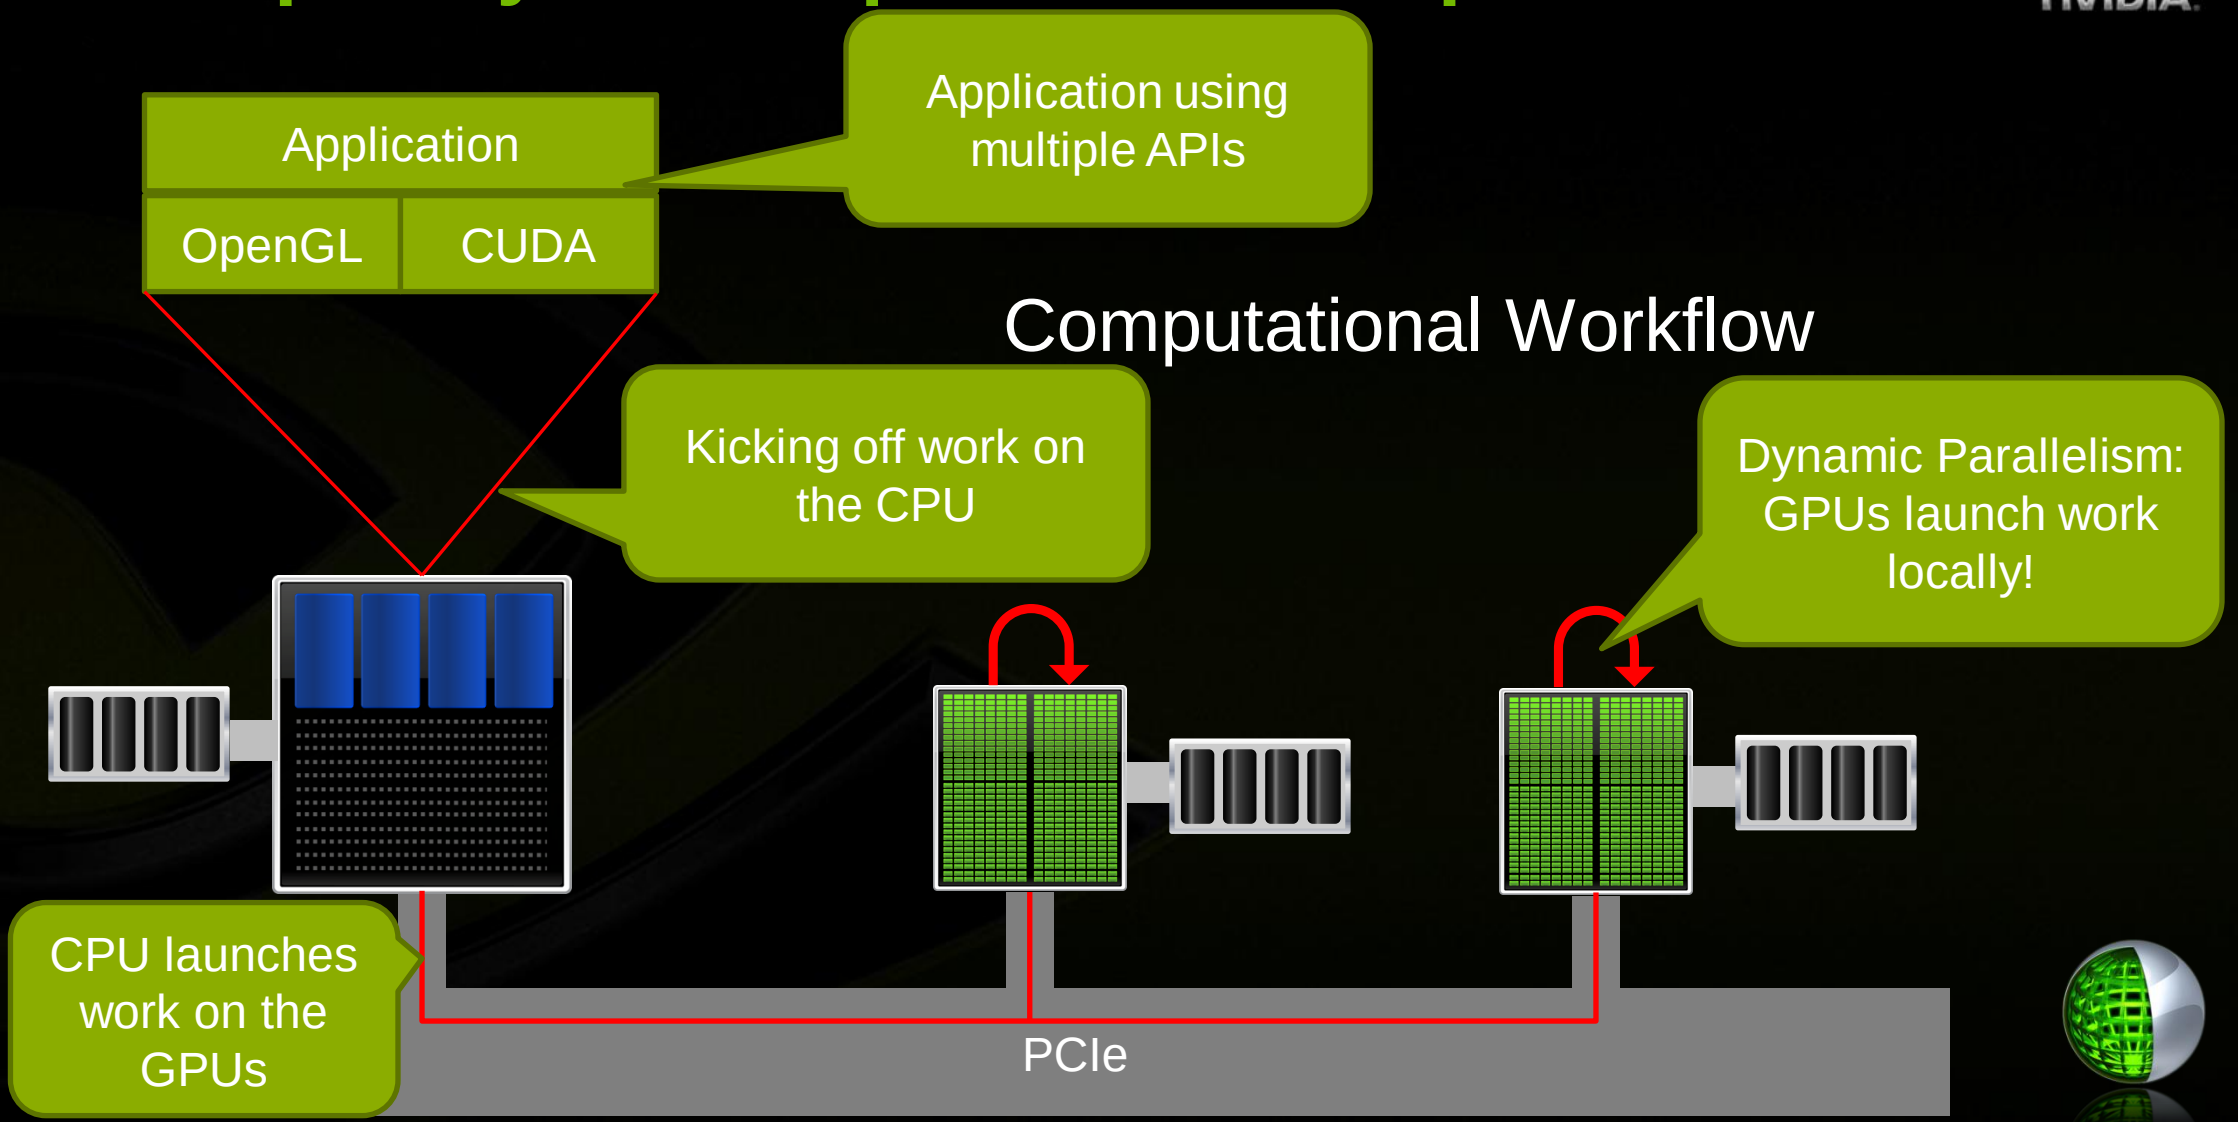

Complexity of Computational Graphics

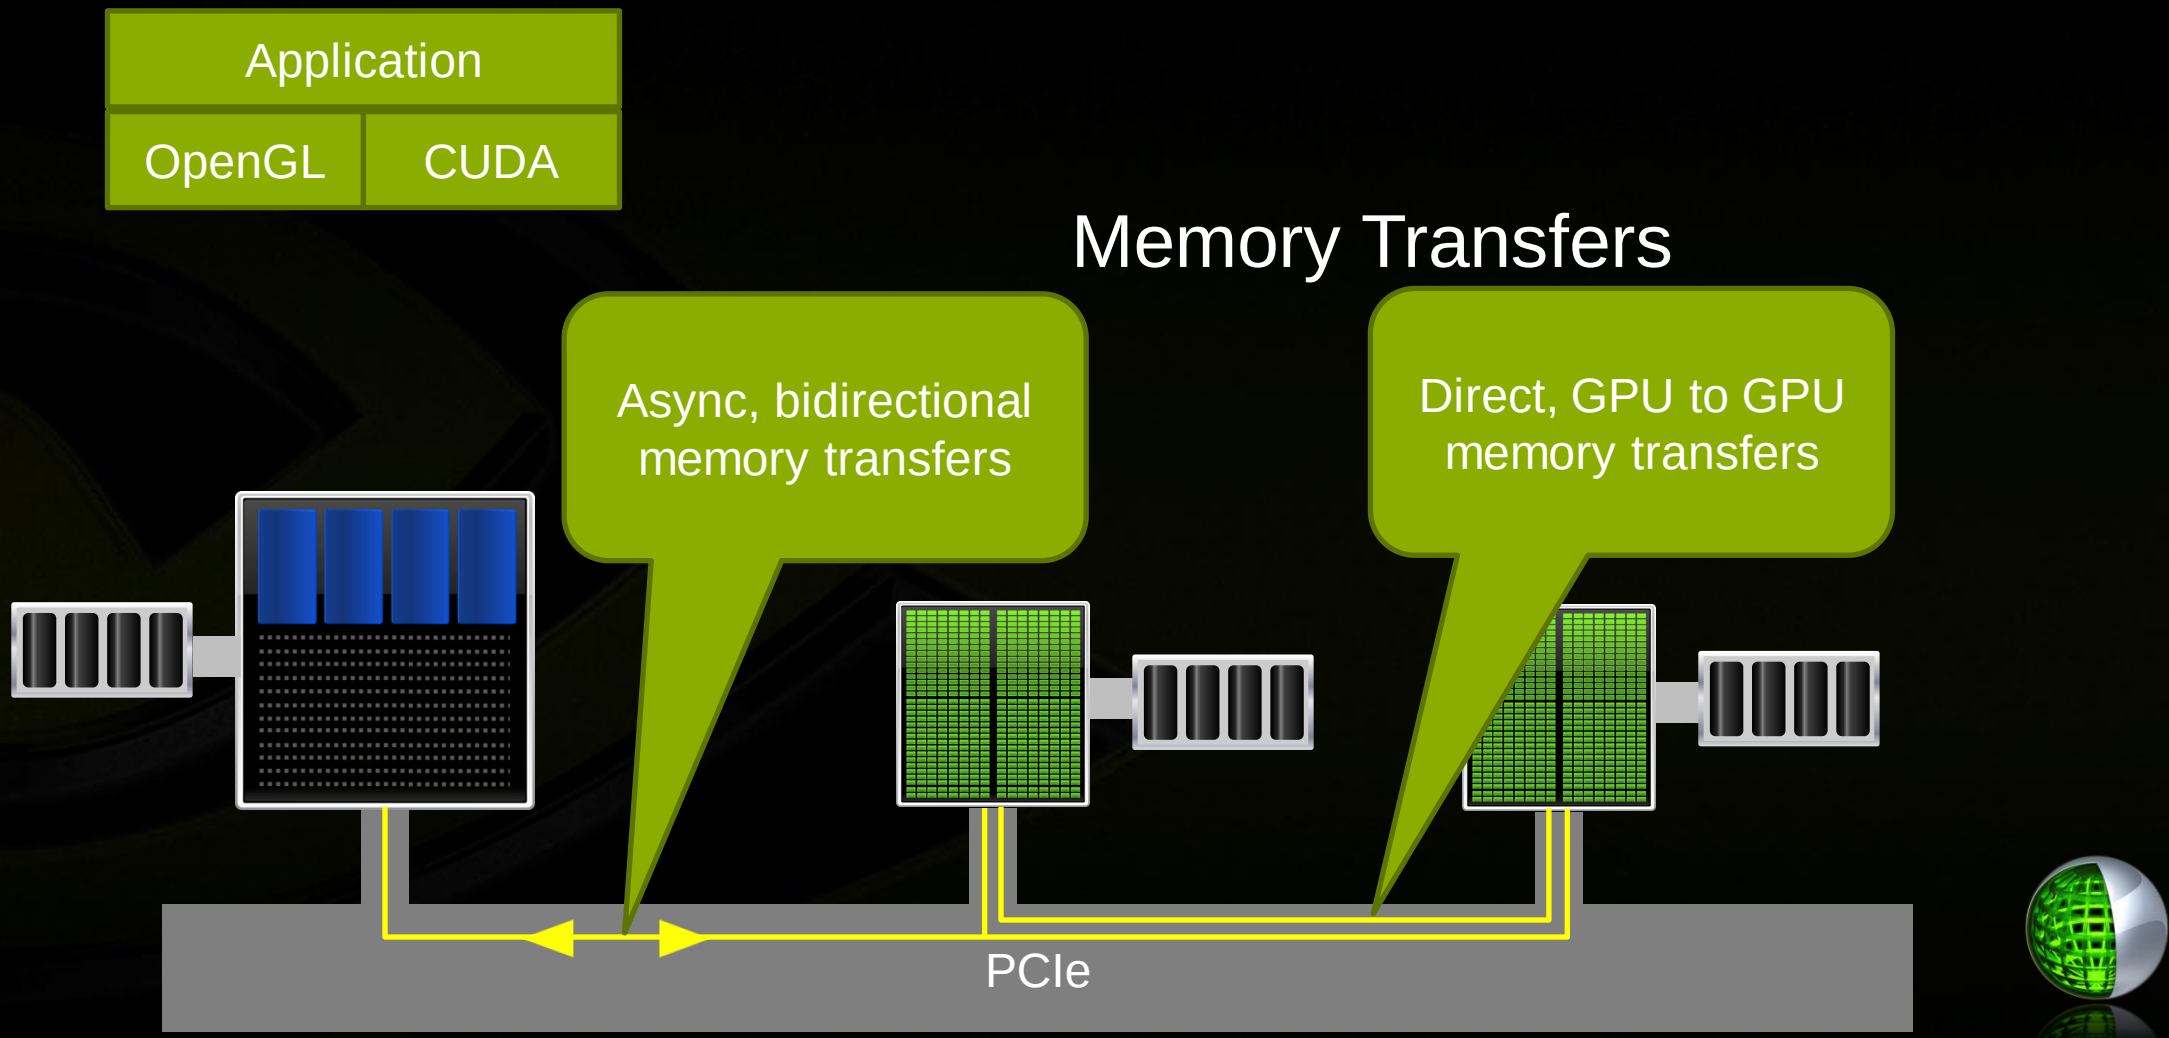

NVIDIA® Nsight™ Visual Studio Edition

Visual Studio integrated development for GPU and CPU

Build **Debug** **Profile**

Microsoft® Visual Studio® PARTNER

NVIDIA Nsight® Microsoft® Visual Studio® Edition

NVIDIA

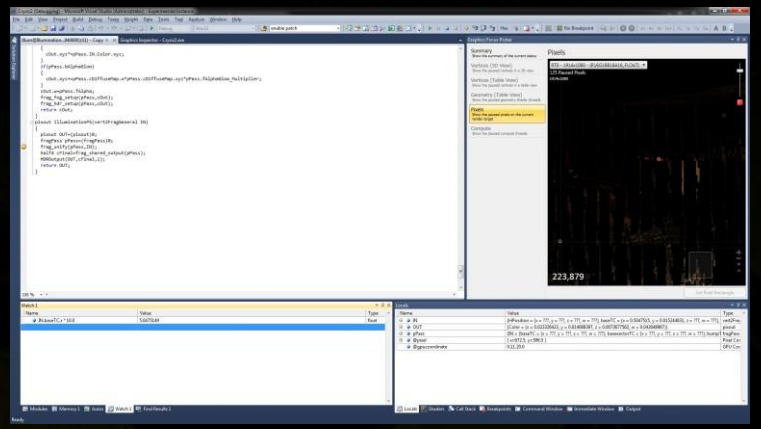

GPU Debugger

GPU native Compute and Graphics debugging

GPU breakpoints including complex conditionals

GPU memory views and exception reporting

Dynamic Shader Editing

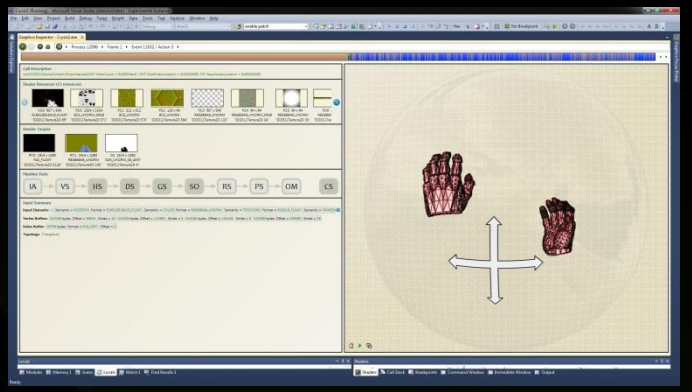

Frame Debugger

Real-time inspection of Graphics APIs

Investigate GPU pipeline state

See contributing fragments with Pixel History

Profile frames to find GPU bottlenecks

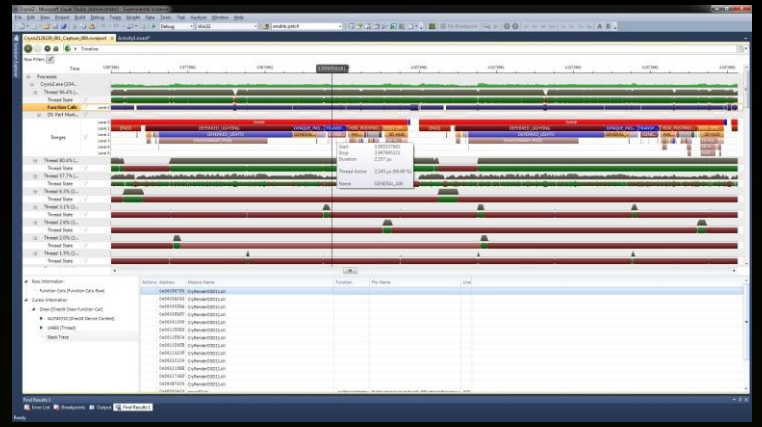

System Analysis

View CPU & GPU events on a single timeline

Examine workload dependencies/memory transfers

CPU/OS, File I/O, Compute, Graphics API Trace

Capture call stack and jump to source

Free and Integrated into Visual Studio 2008 and 2010

NVIDIA Nsight 3.0 – CUDA Features

- Maximus support - debug CUDA and Graphics
- CUDA 5.0
- Advanced CUDA Profiling
  - Source correlated profiling
  - Dynamic kernels and memory transfer tracing
- View interop surfaces while debugging CUDA
- Debug and trace dynamic kernels
- CUDA library usage

NVIDIA Nsight 3.0 – Graphics Features

- Single GPU Debugging for Graphics
- Support for OpenGL
  - All new GUI
  - GLSL Shader Debugging
  - Frame Debugger
  - Frame Profiler
  - Pixel History
  - Starting with 4.2 Core functions, more to come…

Demo Application: FluidsGL

- Developed by Cliff Woolley & Nolan Goodnight
- Compute the fluid simulation on one GPU
  - Simple fluids solver for the Navier-Stokes equations for incompressible flow
  - Jos Stam's FFT-based stable fluids system
- Render the results in OpenGL on second GPU
  - Diffuse texture & normal map
  - Sample in UV space & screen space then blend
  - Gives look that there is a 2nd draw call in the background

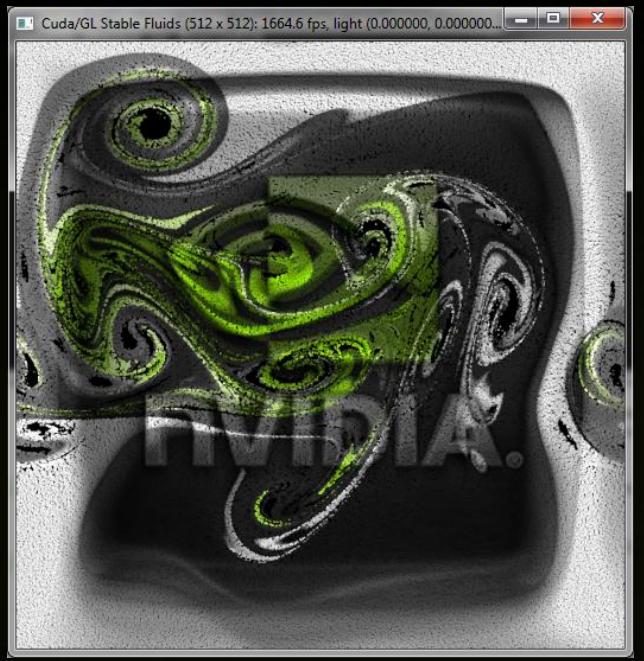

Demo: Setup your CUDA project

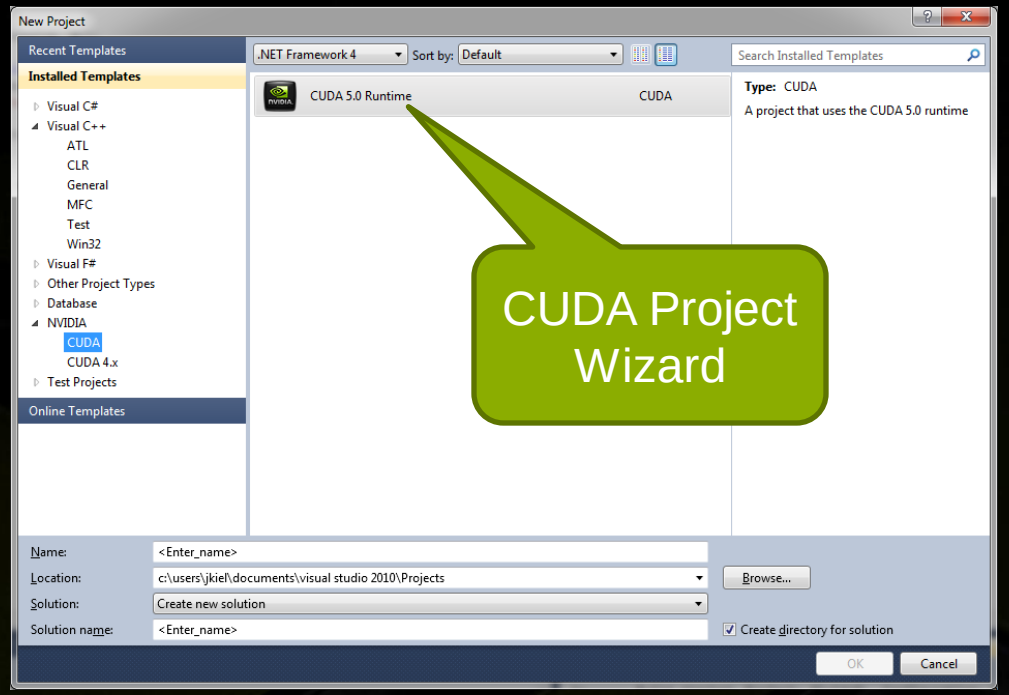

CUDA Project Wizard

CUDA Compile Settings

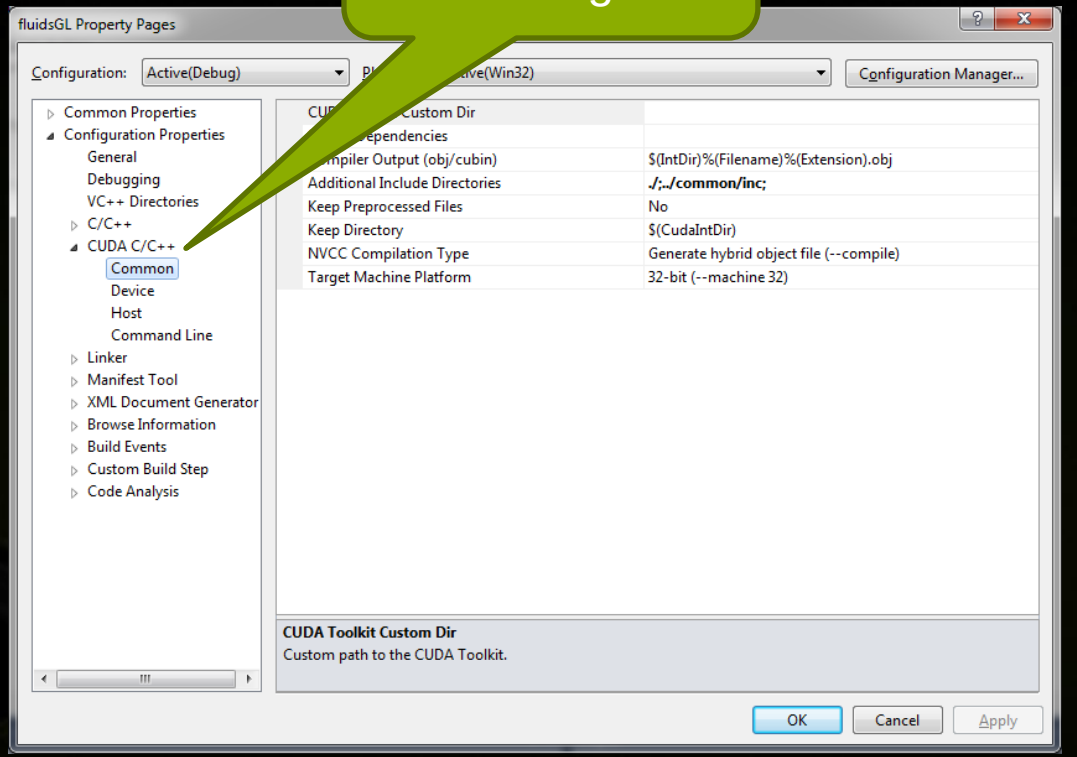

Demo: Launching…

Start Nsight Monitor

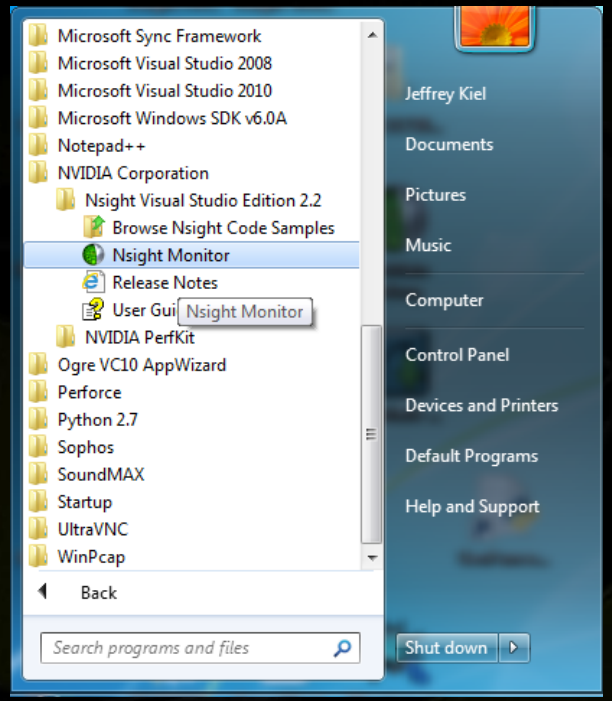

Configure Parallel Nsight Project Settings

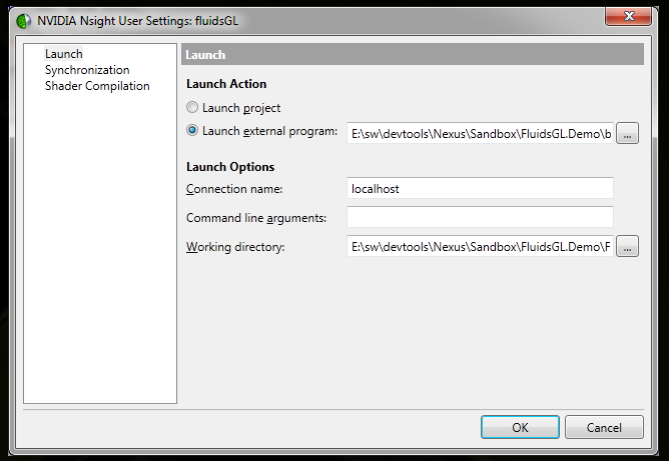

Launch Your Application

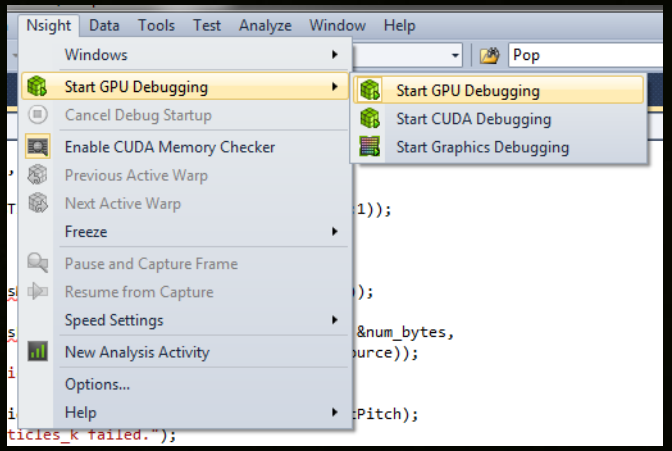

Demo: Single GPU & Memory Checker

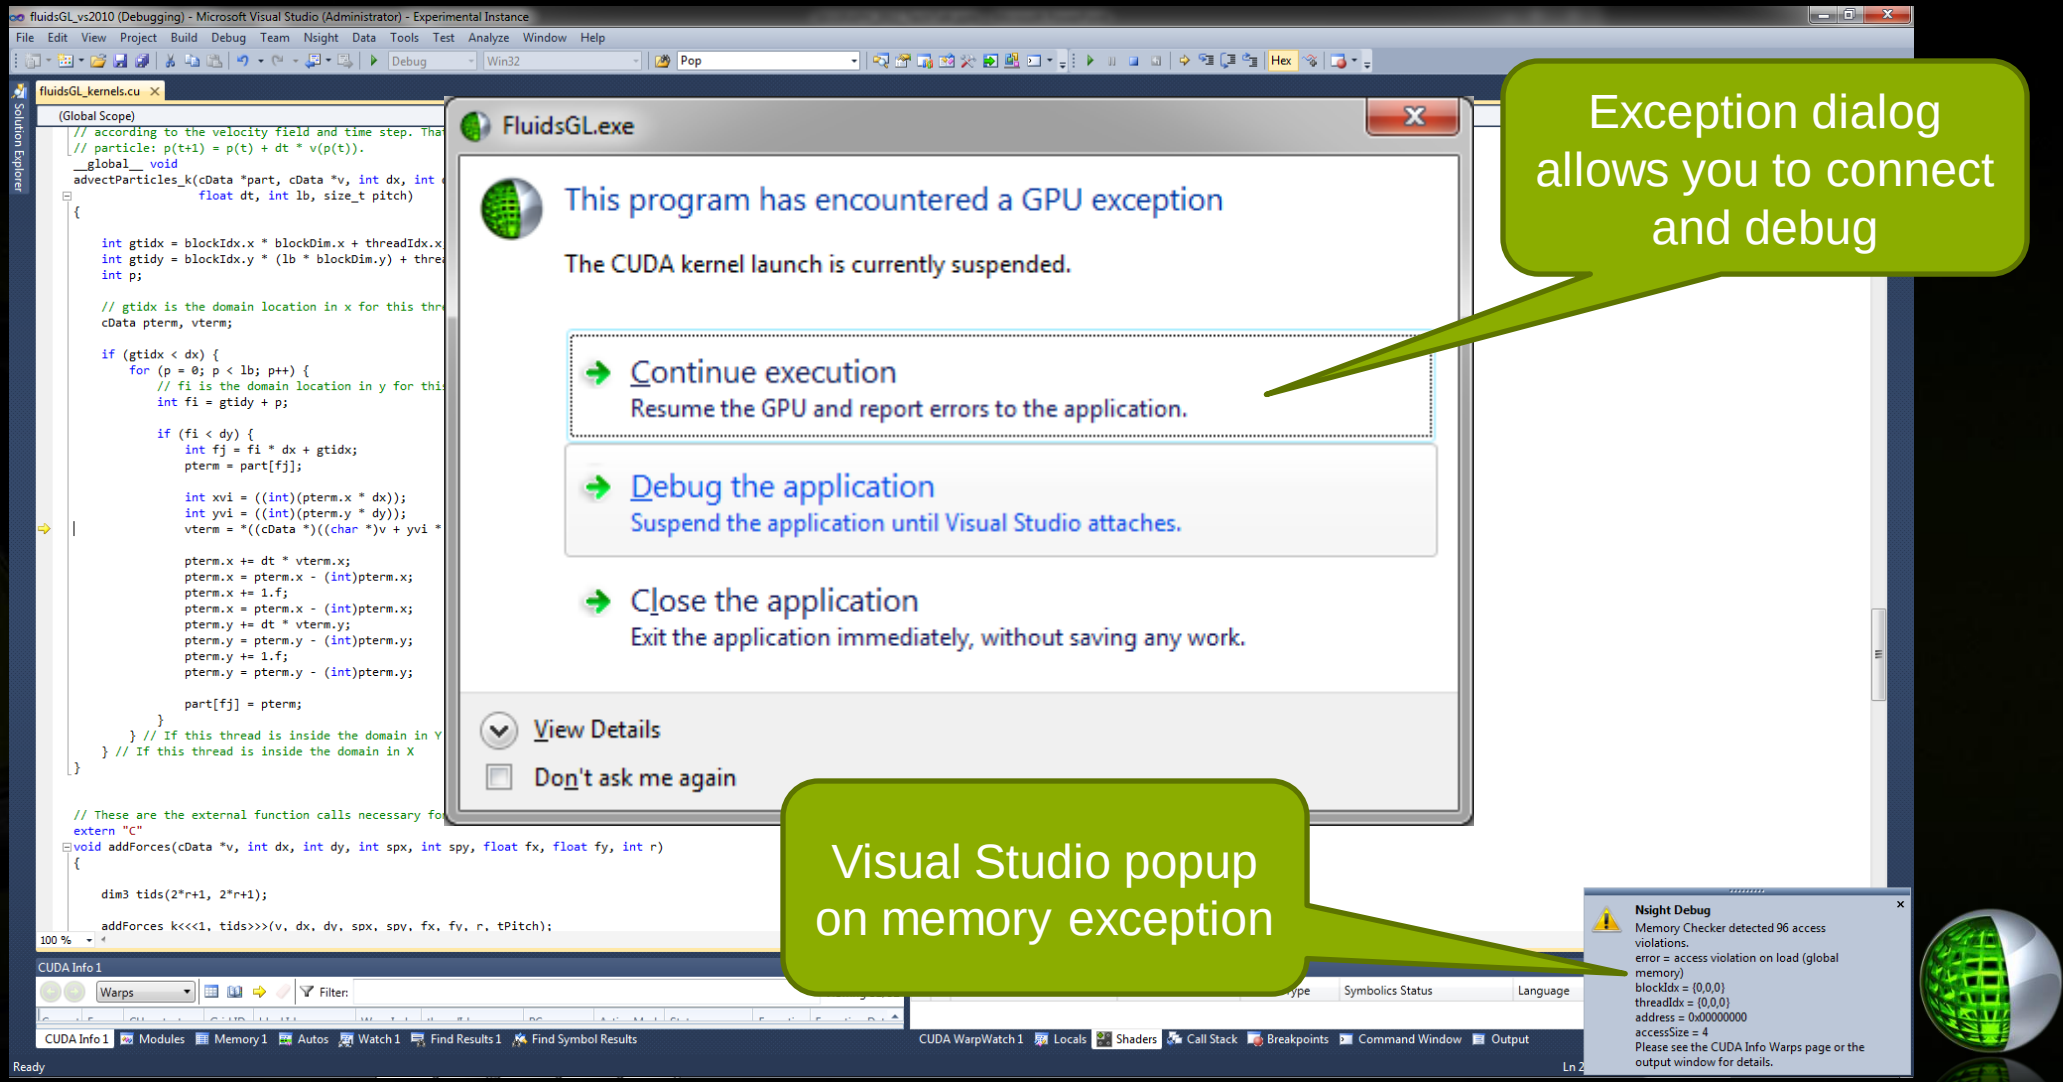

Demo: CUDA Kernel Debugging

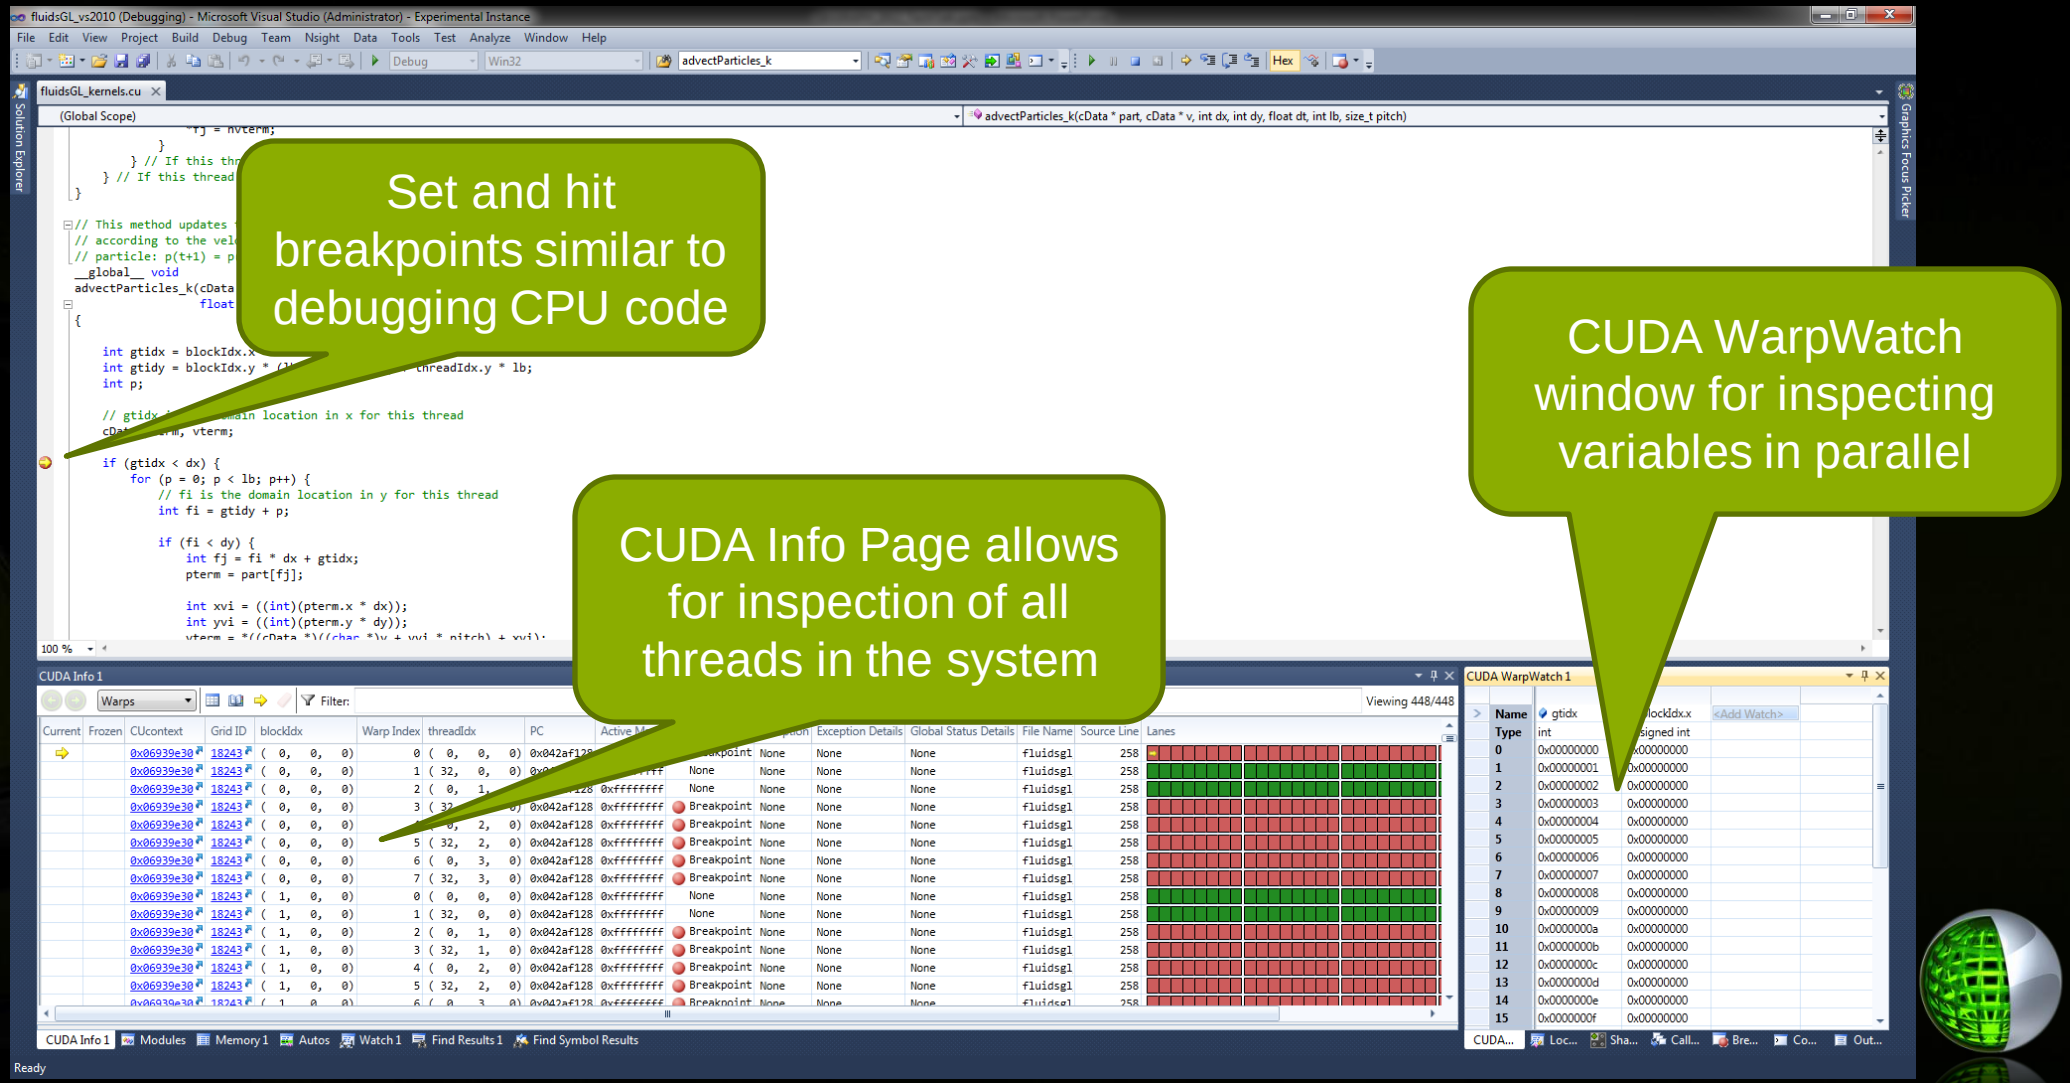

Demo OpenGL Shader Debugging

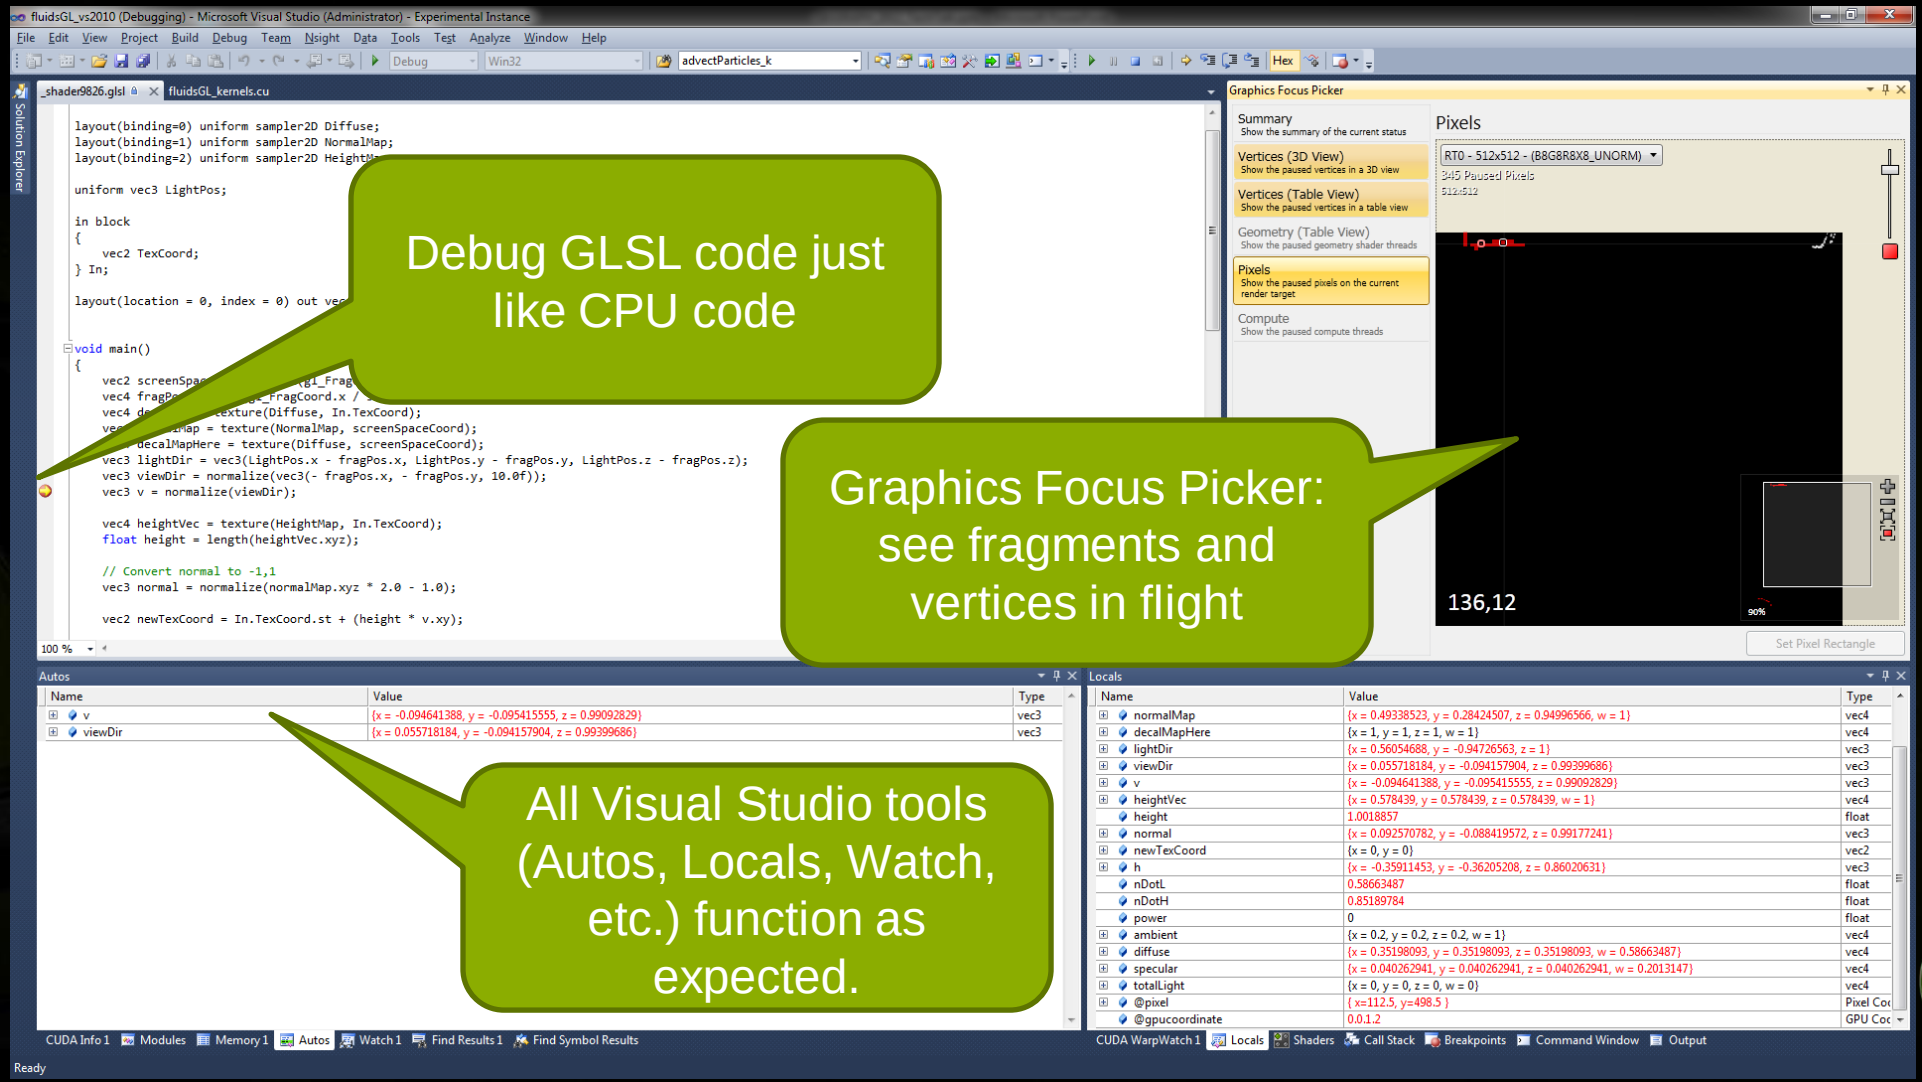

Demo: OpenGL Frame Debugger

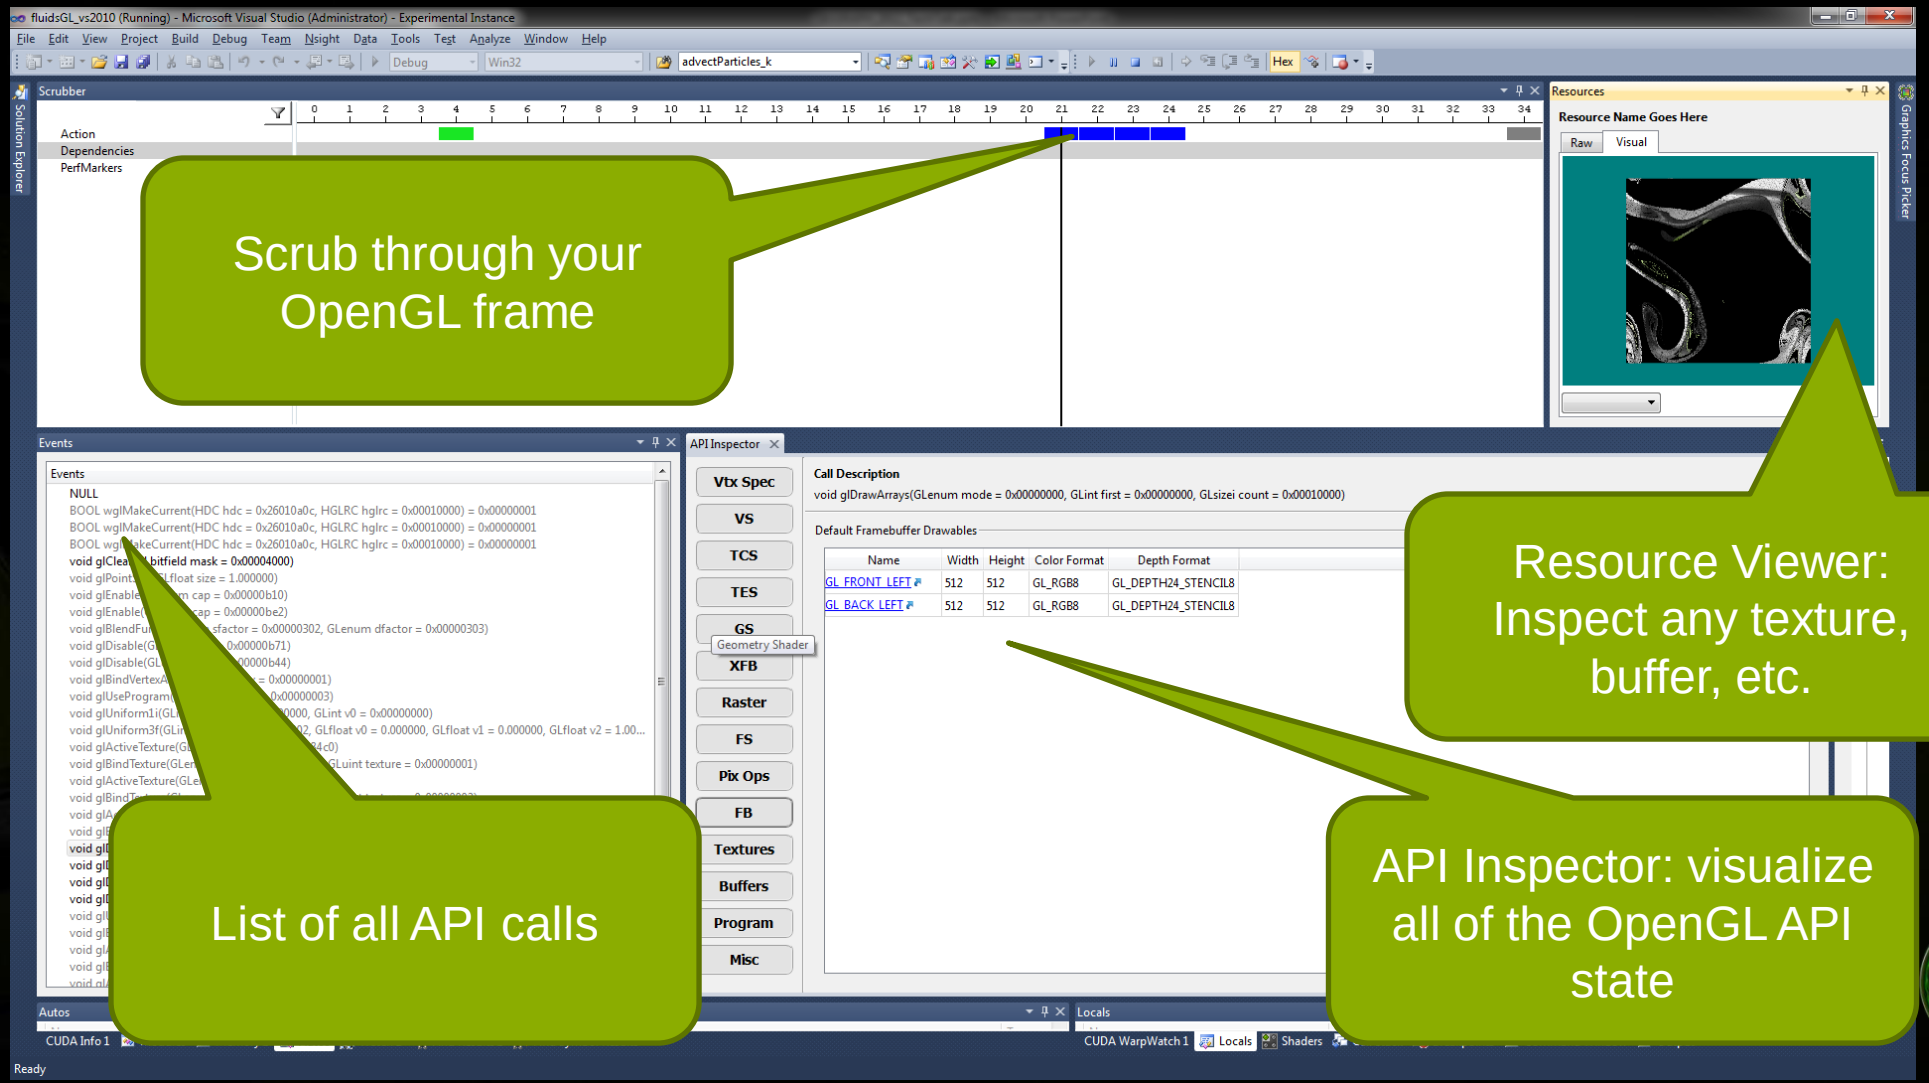

Demo: FluidsGL Analysis

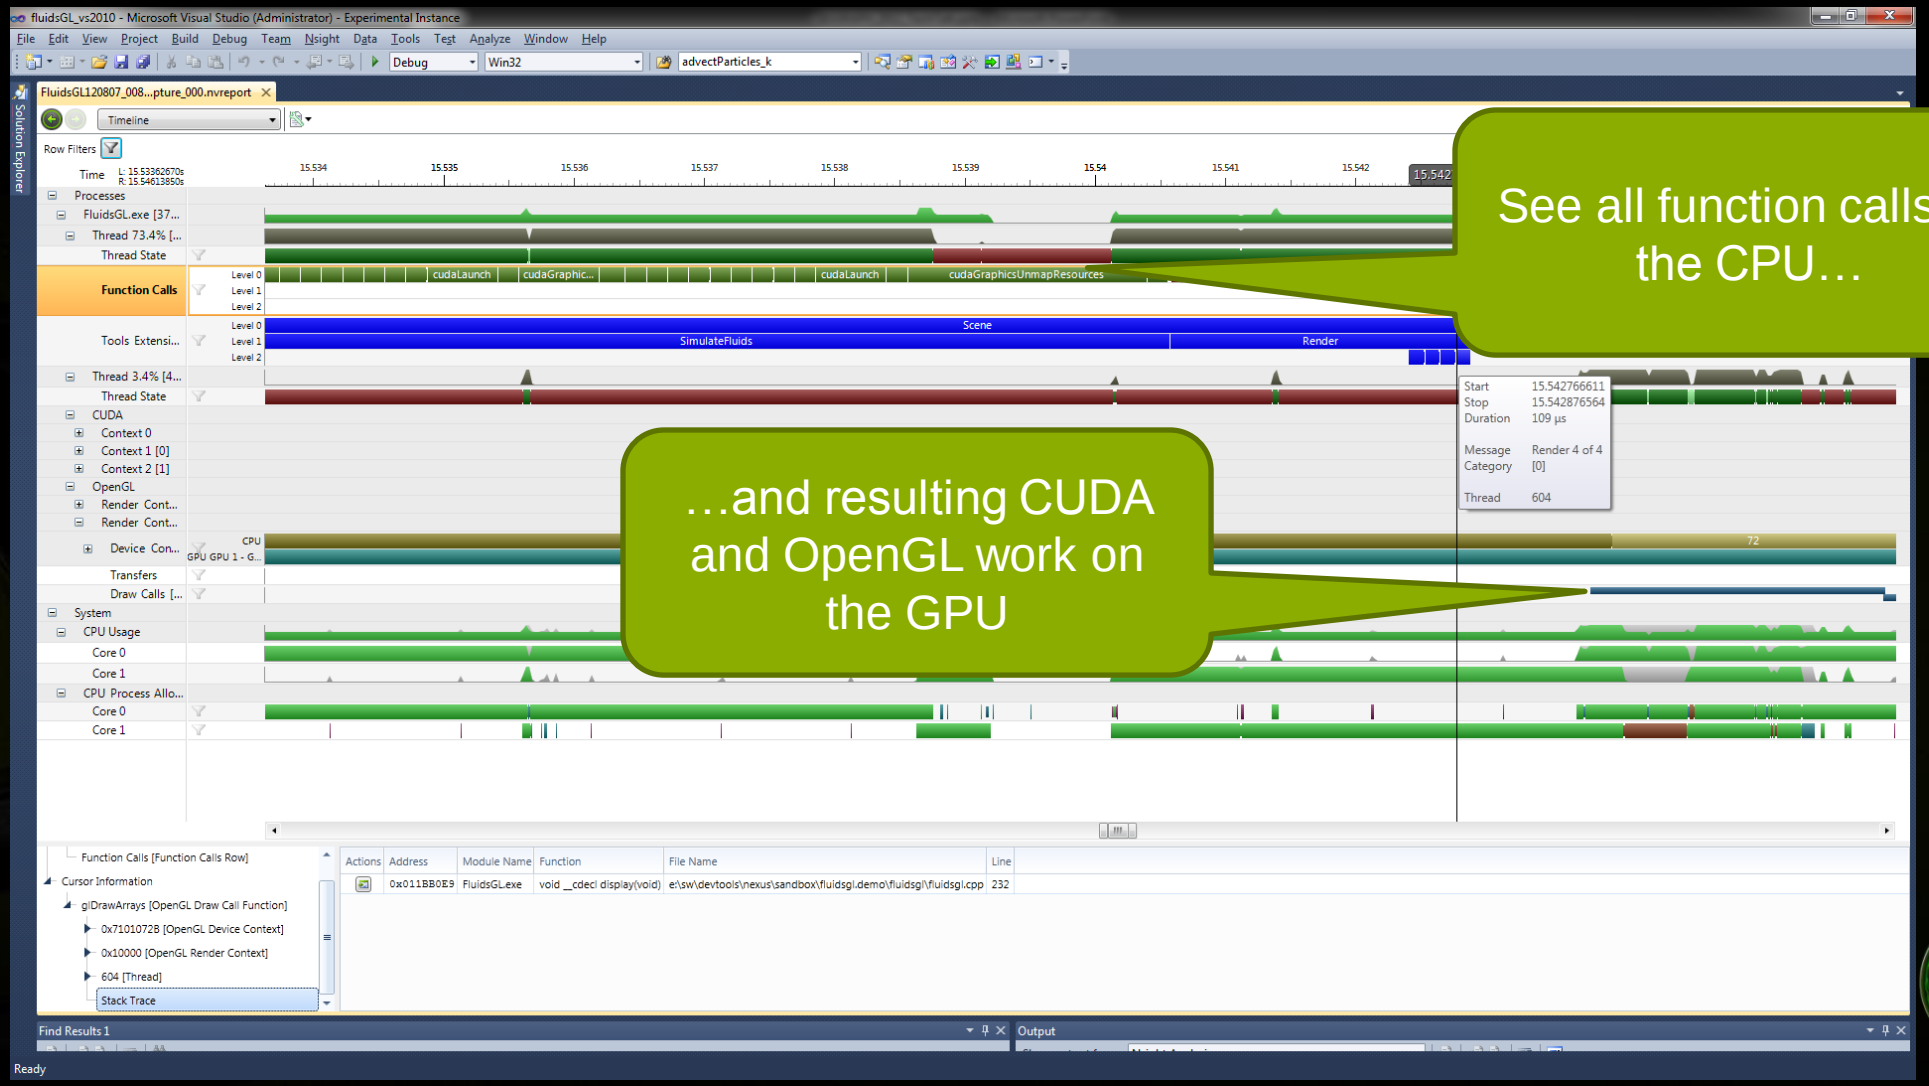

Demo: OptiX Fracture Demo Analysis

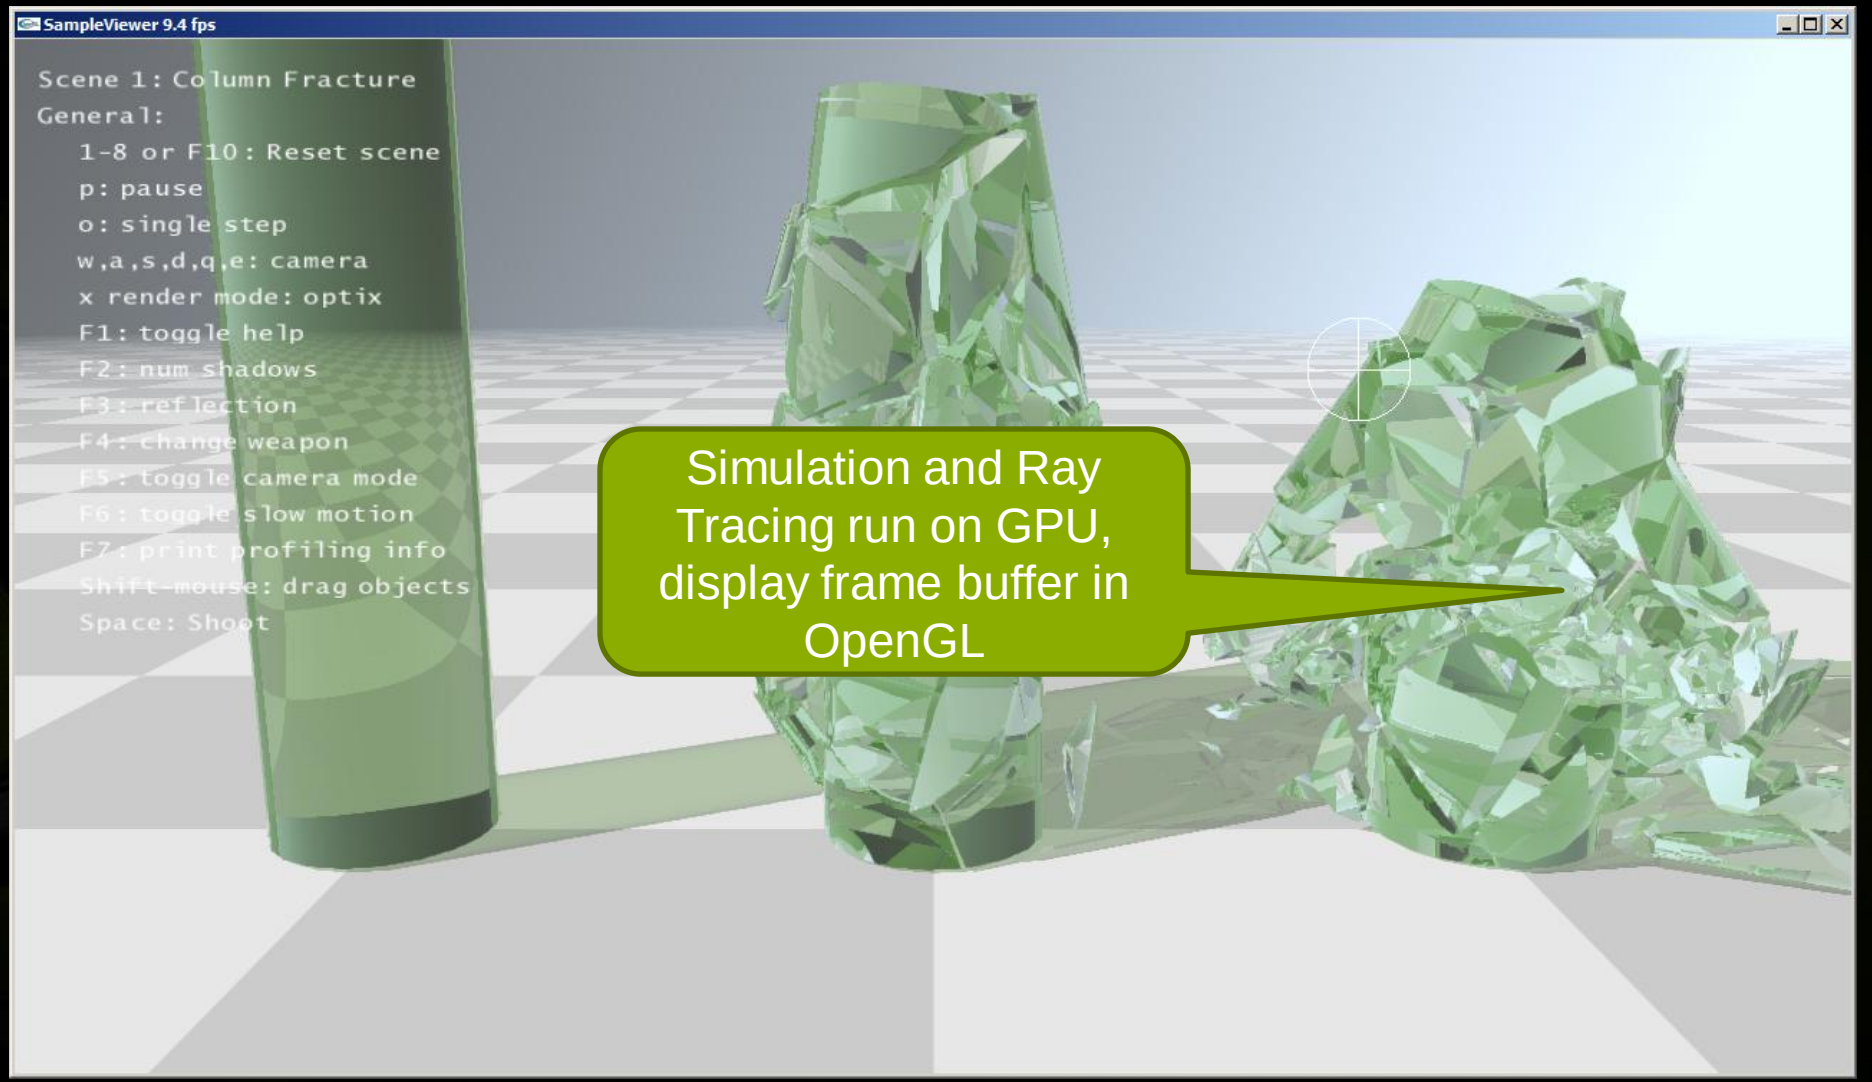

Demo: OptiX Fracture Demo Analysis

Track GPU usage across all processes

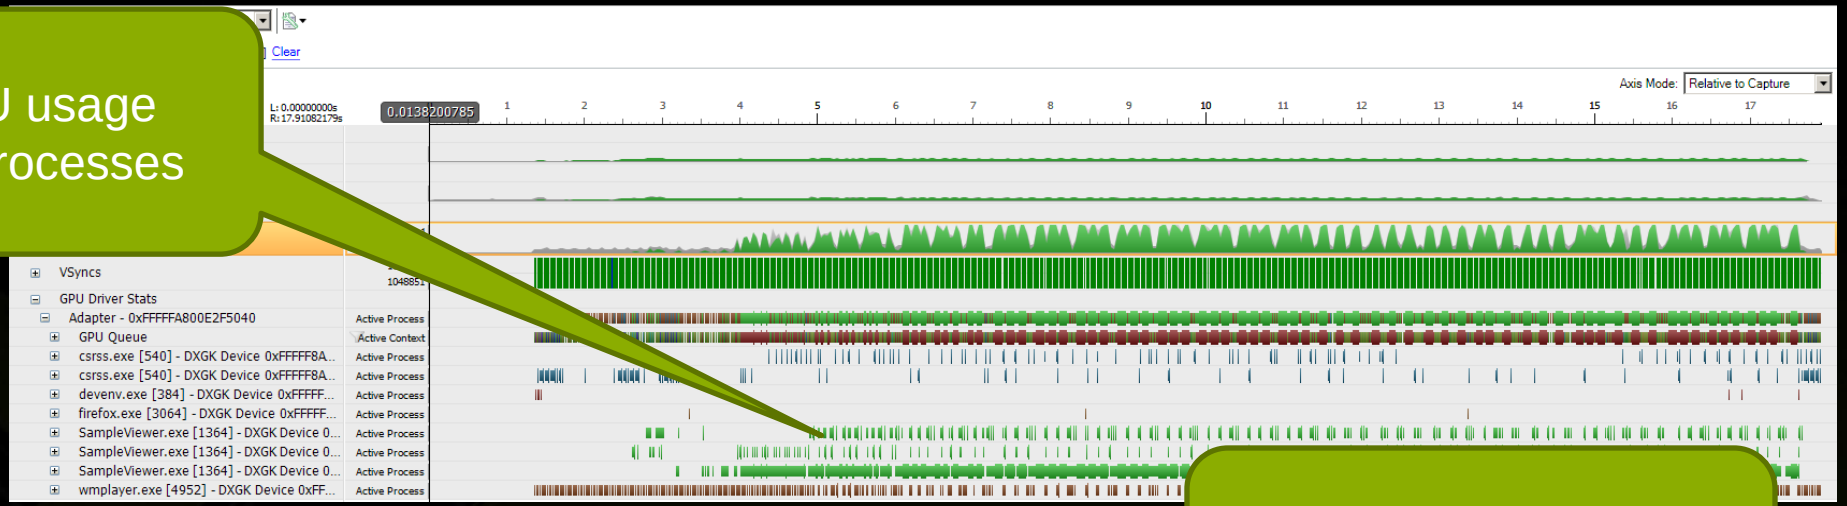

Highlight the process of interest

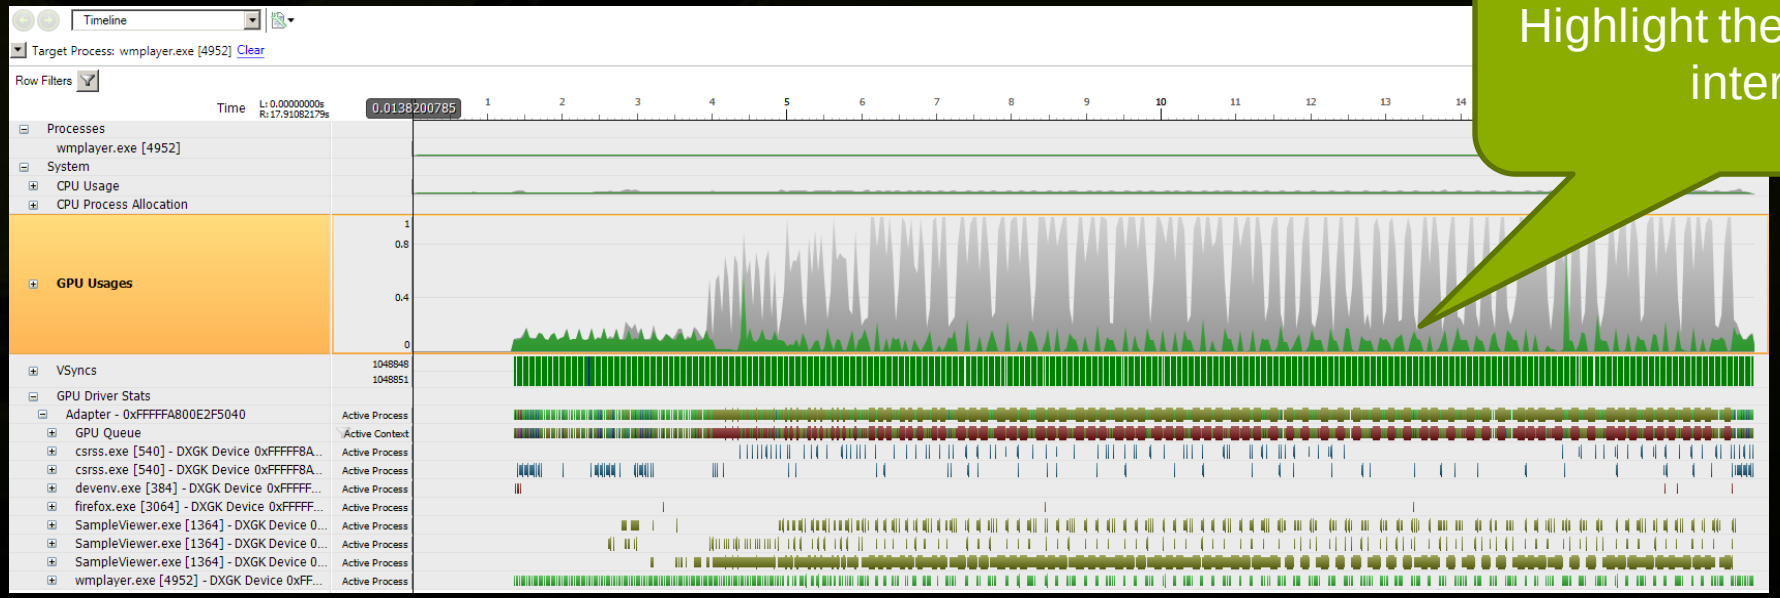

Demo: OptiX Fracture Demo Analysis

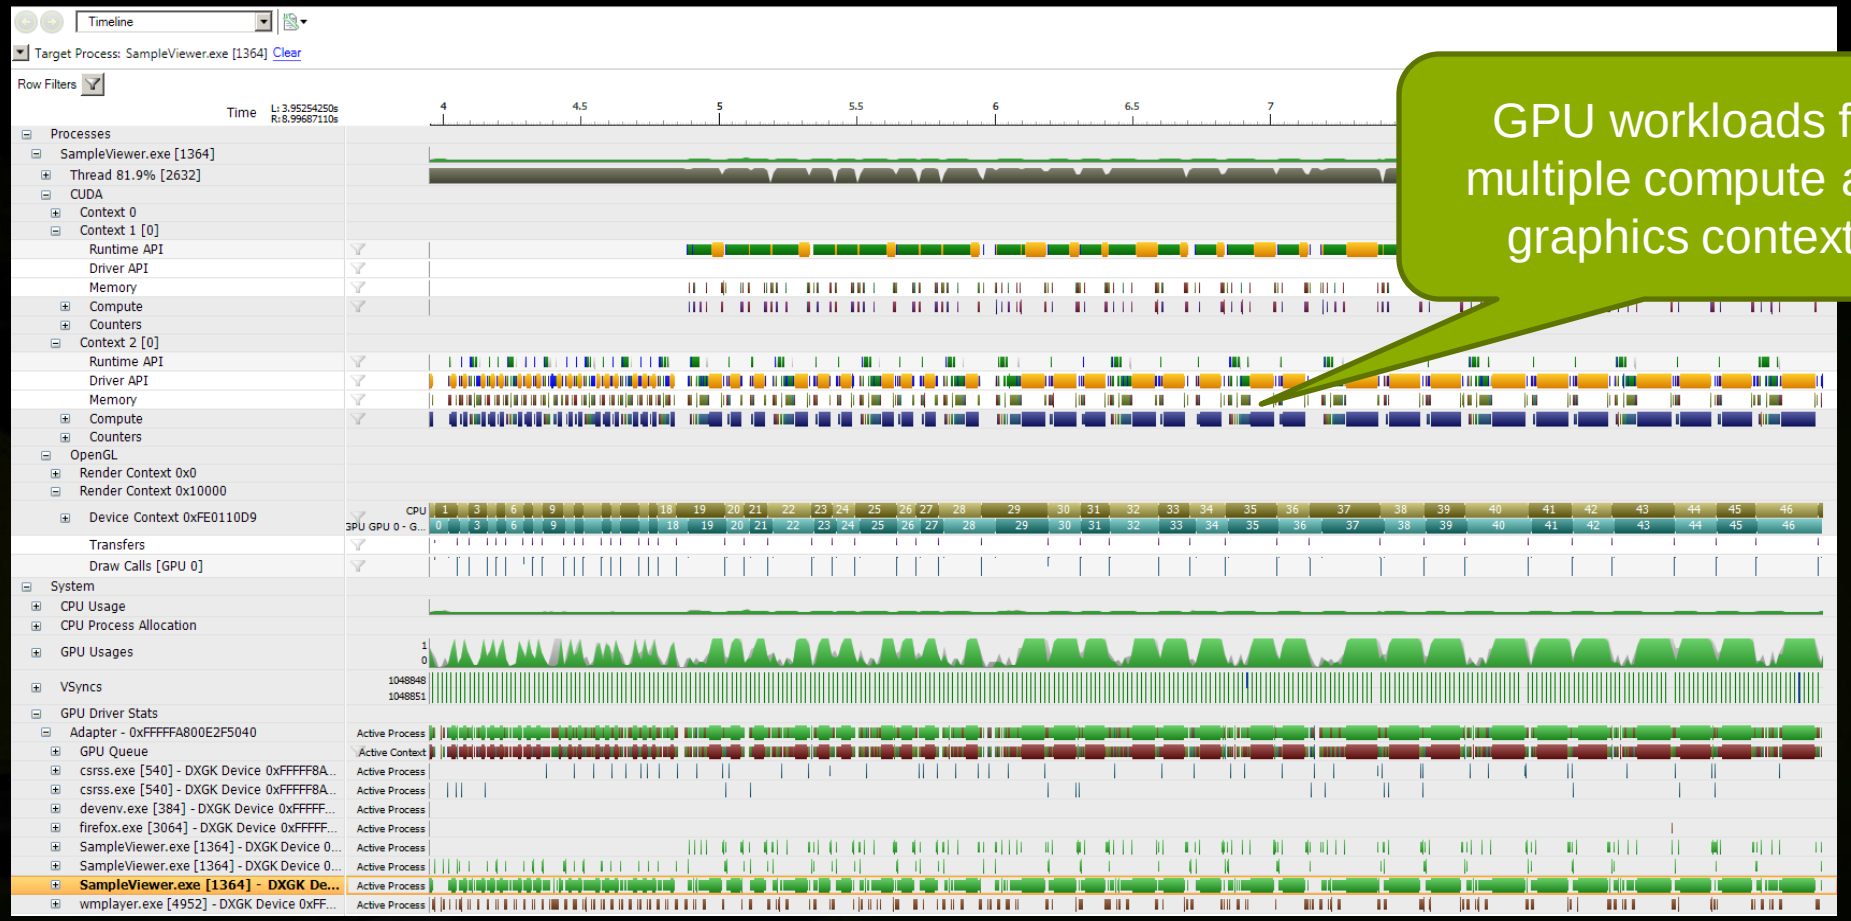

Demo: OptiX Fracture Demo Analysis

Inspect OS GPU work queue to look for contention

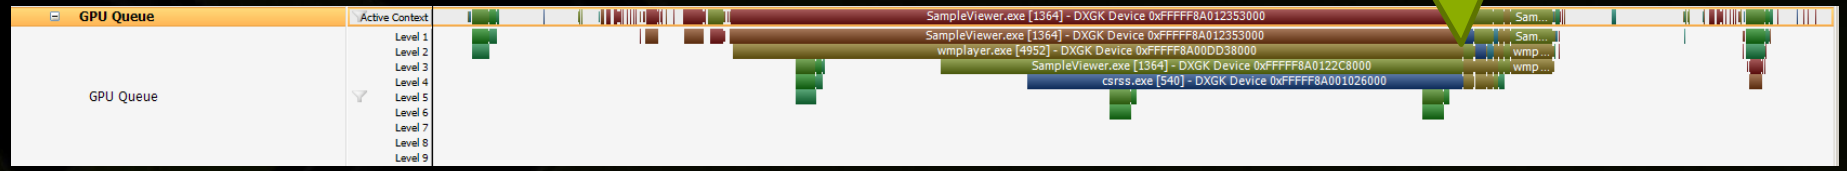

Source Correlated Profiling

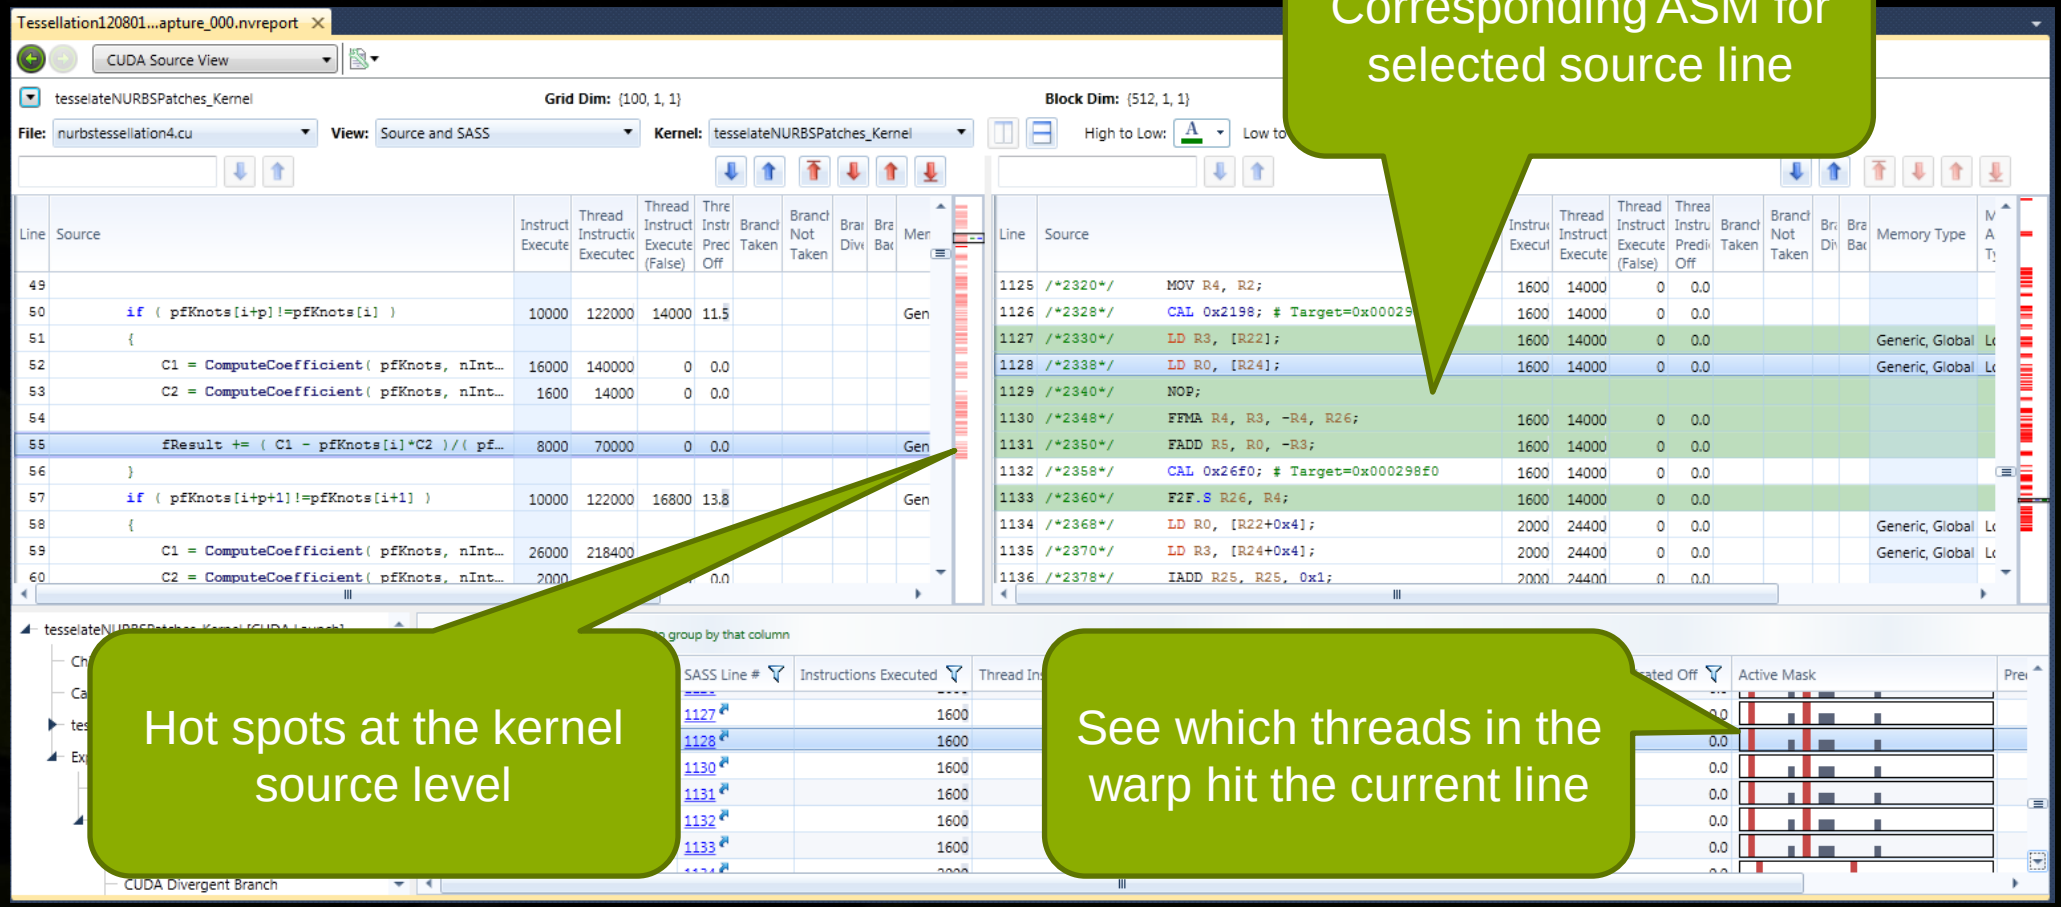

Supported Configurations: Desktop

- ✓ CUDA & Graphics Analysis
- ✓ CUDA & Graphics Profiling
- ✓ CUDA & Graphics Debugging
- ✓ CUDA Memory Checker
- ✓ Graphics Frame Debugger
- ✓ Graphics Pixel History

Single GPU

2 GPUs
Tesla®+GPU

Multi-OS

Remote PC

Supported Configuration: Laptop

- ✓ CUDA Analysis
- ✓ CUDA Profiling
- ✓ CUDA Debugging
- ✓ CUDA Memory Checker

Optimus™ Laptop

PerfKit returns to Linux!

Windows download posted, Linux to come Sept 2012

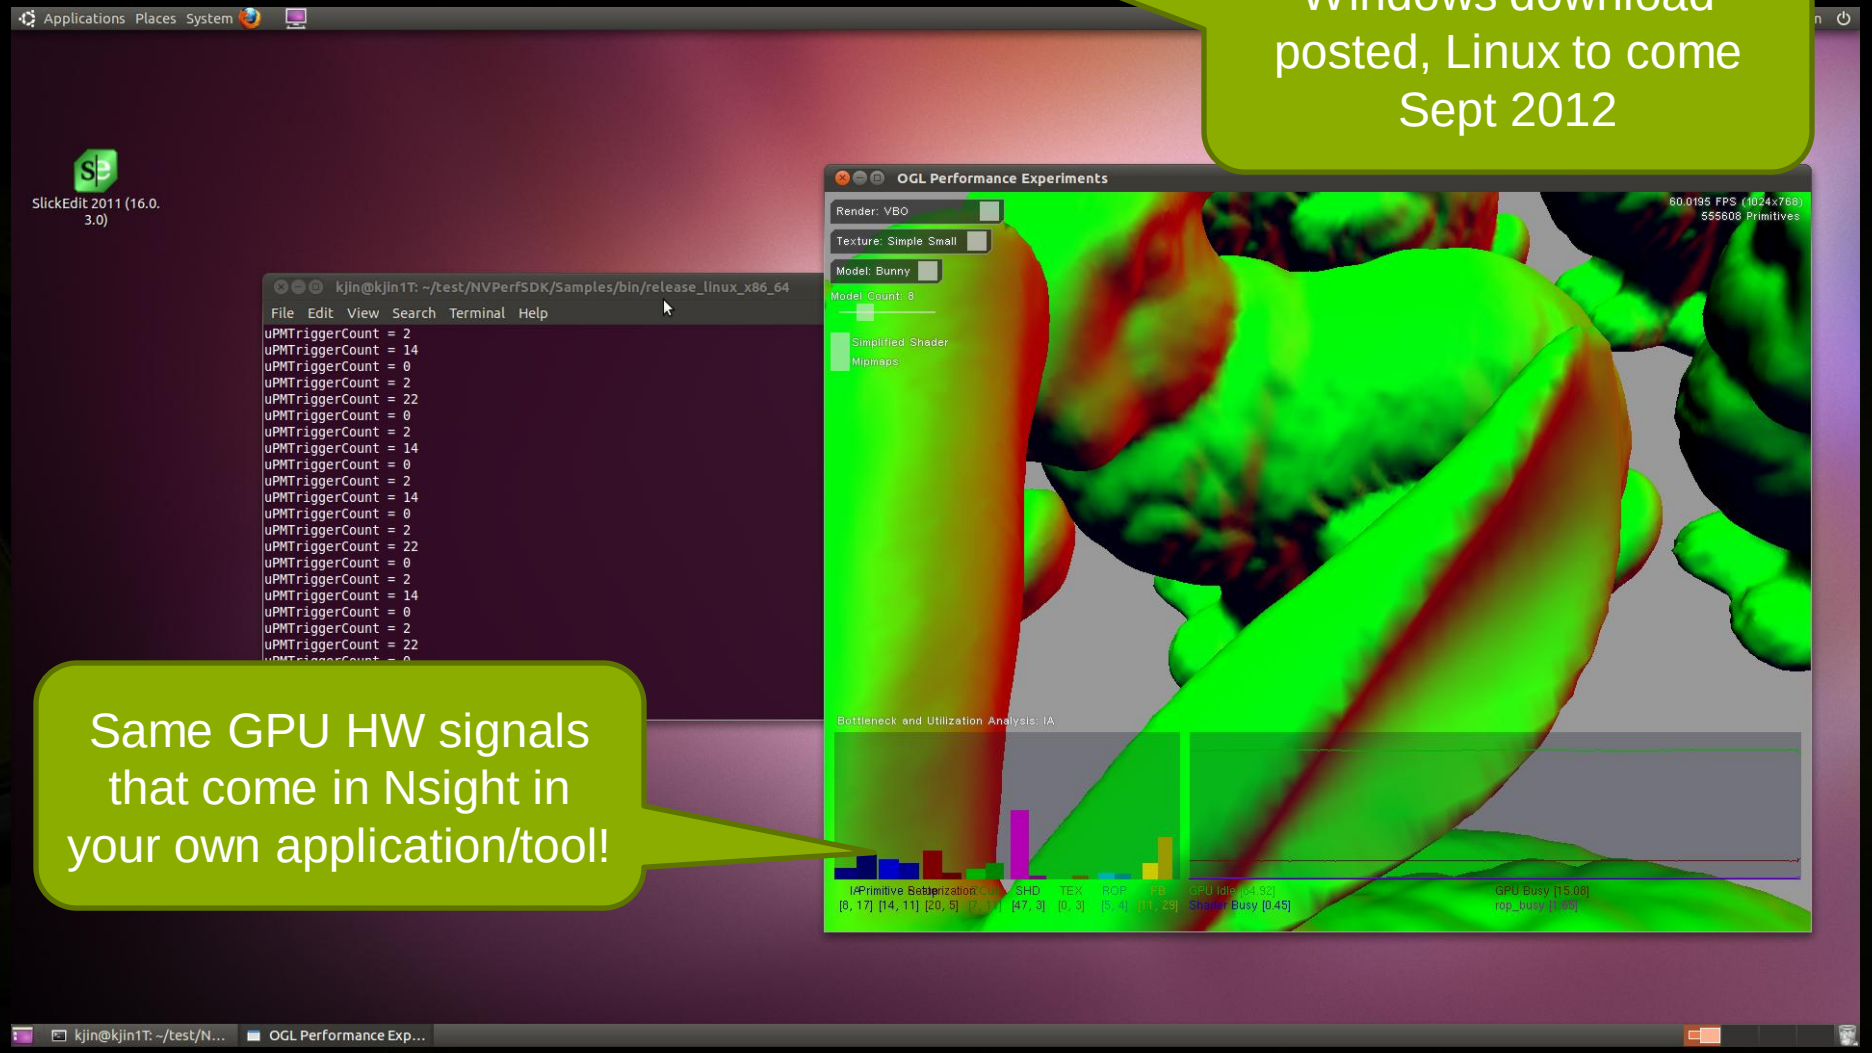

Same GPU HW signals that come in Nsight in your own application/tool!

Recap: Nsight Tames Computational Graphics

- Maximus & CUDA 5.0 support
- Simultaneous CUDA & Graphics Debugging
- Single GPU Debugging for Graphics and CUDA
- OpenGL Frame Debugger and Profiler
- Tracing CPU & GPU workloads and memory transfers on all APIs
- Source correlated performance information

Thank you…questions?

- Join our Registered Developer Program
- Grab latest version from
  http://developer.nvidia.com/nsight

GPU TECHNOLOGY CONFERENCE

GTC 2013 | March 18-21 | San Jose, CA

The Smartest People. The Best Ideas. The Biggest Opportunities.

Showcase Your Work Among the Elite of GPU Computing!

Call for Submissions
- Call for Sessions: August 2012
- Call for Posters: October 2012

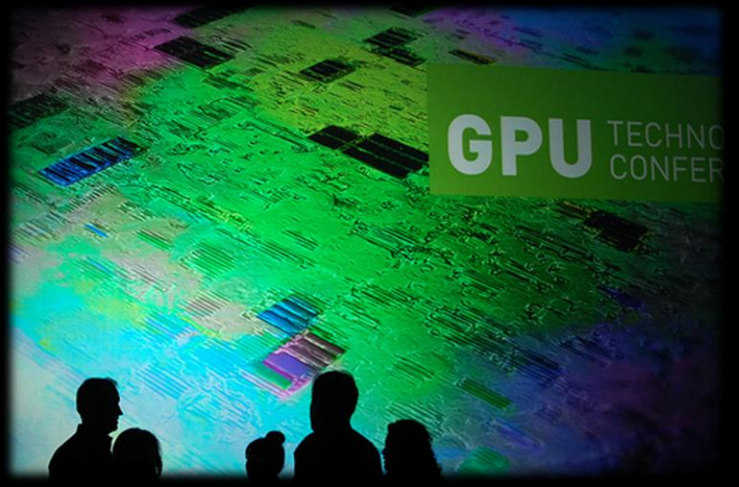

As a speaker you'll receive many benefits that include:

- A complimentary Full Conference pass
- Discounted passes for your colleagues
- Access to exclusive speaker-only on-site amenities
- Audio capture and hosting of your session
- A unique opportunity to promote your name, your expertise, and your company

Learn more at www.gputechconf.com

GTC TALKS SERIES