



**nVIDIA®**

## **OpenCL Update**

**Simon Green  
SIGGRAPH 2009 OpenCL BOF**

# NVIDIA and OpenCL

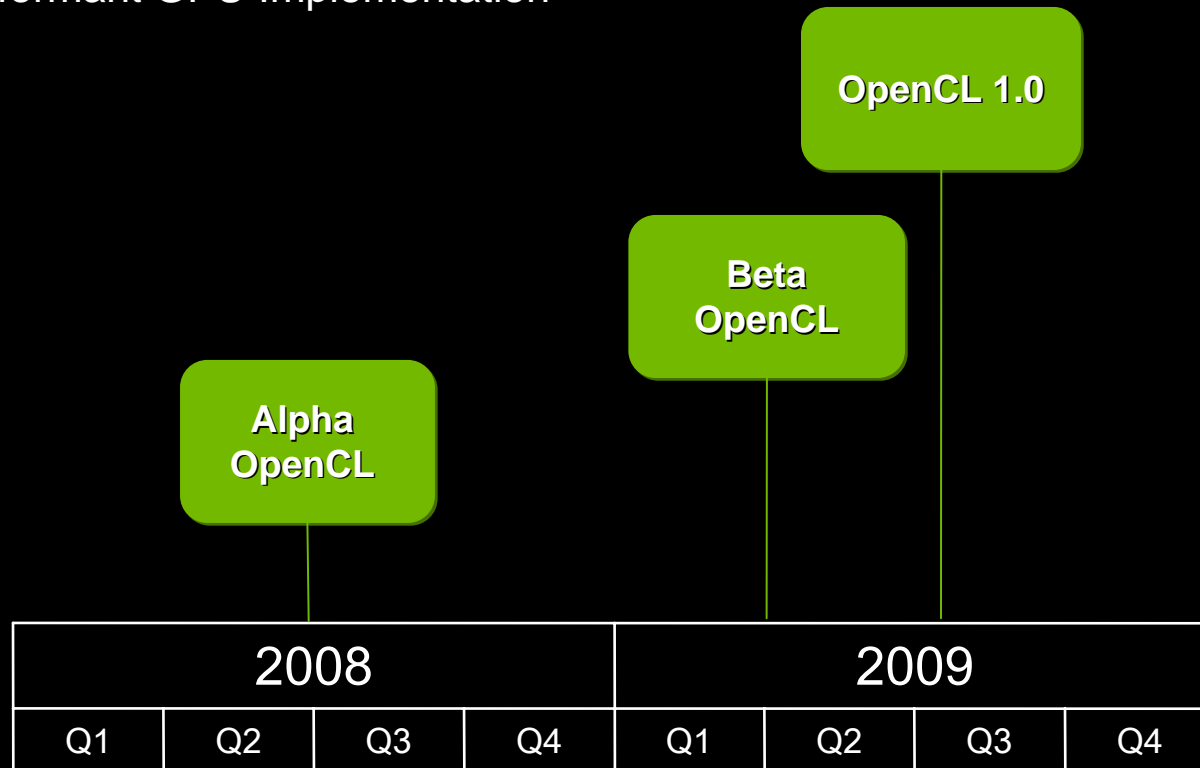
- NVIDIA has worked closely with Apple and the Khronos Group since the inception of OpenCL
- OpenCL was developed on NVIDIA GPUs
- OpenCL working group is chaired by NVIDIA VP Neil Trevett (also President of Khronos Group)
- NVIDIA is supplier of all GPUs for new Apple notebooks
- NVIDIA was first to show working OpenCL on GPU

hardware



# NVIDIA's OpenCL Timeline

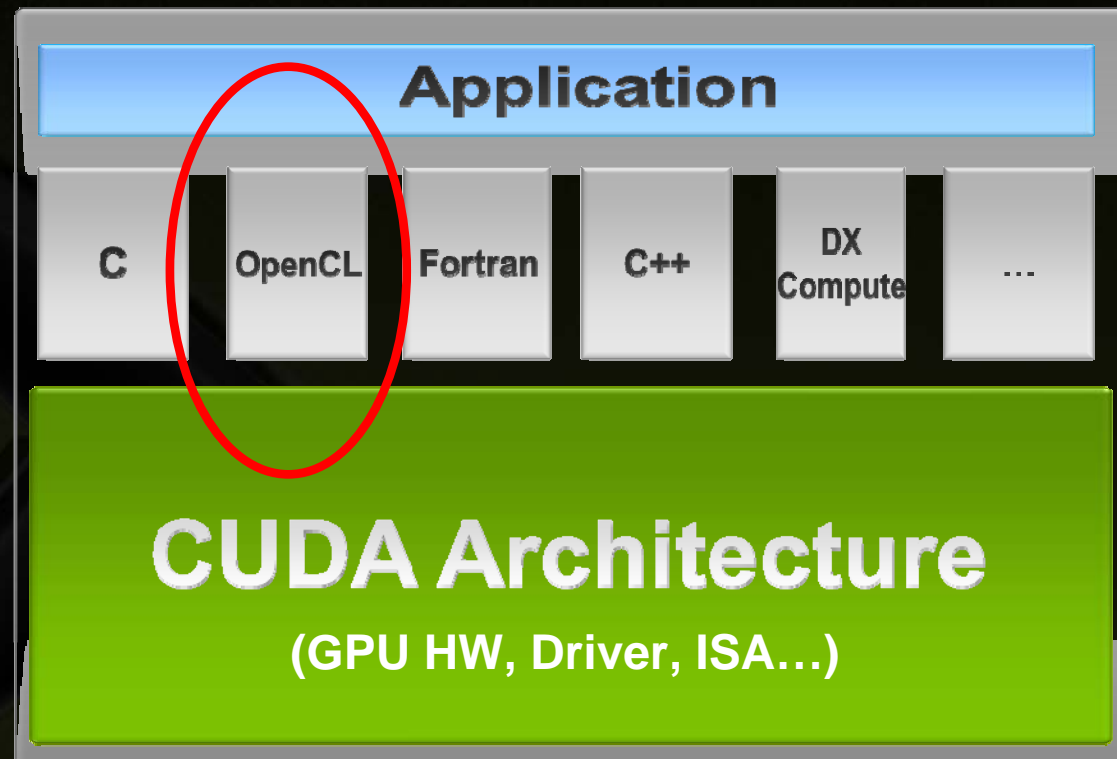
- 12 / 2008 1<sup>st</sup> operable OpenCL demo on GPU (Siggraph Asia)
- 4 / 2009 Drivers, compiler and SDK available to developers
- 5 / 2009 1<sup>st</sup> GPU implementation filed for conformance
- 6 / 2009 1<sup>st</sup> Conformant GPU Implementation



# OpenCL and the CUDA Architecture



- OpenCL
- C for CUDA
- DirectX Compute
- Fortran (PGI)
- C++



The CUDA Architecture supports all standard languages & APIs to tap the massive parallel computational power of the GPU

# Advantages of OpenCL

- **Cross-vendor support**
  - But be ready to optimize for differing hardware characteristics
- **Just-in-time compilation**
  - pass kernel source directly to the driver
- **Tight interoperability with OpenGL graphics**
- **API support for task parallelism**
- **Well-defined specification**
  - Guarantees behaviour across devices

# NVIDIA OpenCL

## Supported OSs

- 32 and 64 bit Windows XP and Vista (and soon Win 7)
- 32 and 64 bit Linux (Ubuntu, RHEL, etc.)
- Mac OS X Snow Leopard (indirectly via Apple)

## Development Environments

- Visual Studio 2005 (8) and Visual Studio 2008 (9) for Windows
- GCC for Linux

## Drivers and JIT compiler

- In SDK for Alpha & Beta
- Soon to be distributed with GPU drivers

## SDK

- Source code & white papers for sample applications (30 presently)
- Documentation: Getting Started Guide, Programming Manual, Best Practices (Optimization) Guide

# NVIDIA OpenCL SDK Samples

- Bandwidth test
- Black Scholes
- Box filter
- DCT
- Device Query
- Histogram
- Matrix Mul
- N-Body
- Post process GL
- Recursive Gaussian
- Reduction
- Scan
- Simple GL
- Simple multi GPU
- Sobel filter
- Transpose
- Vector Add
- Volume render
- Particles
- Ocean

# Tools – NVIDIA Visual Profiler


- Similar to existing CUDA Visual Profiler
- Allows visualization of profiling data from OpenCL applications
  - hardware performance counters
  - timing information
  - memory transfers
  - estimates fraction of peak performance
- Supported on Windows and Linux
- Will be released to all registered developers at the end of August
- See the live demo at the NVIDIA booth (Tools pod)



# NVIDIA OpenCL Implementation

- Based on LLVM compiler
  - Leverages existing NVIDIA optimization framework
- Supported Extensions
  - cl\_khr\_byte\_addressable\_store
  - cl\_khr\_global\_int32\_base\_atomics
  - cl\_khr\_global\_int32\_extended\_atomics
  - cl\_khr\_local\_int32\_base\_atomics
  - cl\_khr\_local\_int32\_extended\_atomics
- OpenCL / OpenGL interop currently not supported
- Available now for registered developers

# Demos!



# Conclusion

- OpenCL is great for developers, the industry and consumers
- The addition of OpenCL to DirectX Compute and C for CUDA offers a fantastic array of GPU Computing choices for developers
- Available now, get coding!

# OpenCL Information and Resources

- **NVIDIA OpenCL Web Page**

[http://www.nvidia.com/object/cuda\\_opencl.html](http://www.nvidia.com/object/cuda_opencl.html)

- **NVIDIA OpenCL Forum**

<http://forums.nvidia.com/index.php?showforum=134>

- **NVIDIA Registered Developer Extranet Site**

<https://nvdeveloper.nvidia.com/login.asp>

[http://developer.nvidia.com/page/registered\\_developer\\_program.html](http://developer.nvidia.com/page/registered_developer_program.html)

- **Khronos (current specification)**

<http://www.khronos.org/registry/cl/specs/opencl-1.0.43.pdf>

- **Khronos OpenCL Forum**

[http://www.khronos.org/message\\_boards/viewforum.php?f=28](http://www.khronos.org/message_boards/viewforum.php?f=28)

**Questions?**