



# **GPGPU**

**General Purpose computation on Graphics Processing Units**

**Mark Harris, 2002**

**Perform demanding calculations on the GPU instead of the CPU!**

**At first, appeared to be a wild idea, but is now a very serious technology! Results were highly varied in the early years, but the GPU advantage has grown bigger and bigger.**



## **Key components of the GPGPU trend**

**High processing power in parallel**

**Programmability: Introduction of shader programs, much more flexible, programmable for any problem.**

**Floating-point buffers: Vital! Initially with poor precision. 32-bit floating-point decent... but not really impressive.**



## **GPGPU approaches**

- **Using fixed pipeline graphics**
  - **Shader programs**
    - **CUDA**
    - **OpenCL**
  - **Compute shaders**



## **Fixed pipeline GPGPU**

**Reformulate a problem to something that can be done by standard graphics operations.**

**Limited success 1999/2000. Not of any practical interest!**



## **Shader-based GPGPU**

**Portable! All GPUs can use shaders, no need for extra software, run using standard software/drivers.**

**All modern shader languages (GLSL, Cg, HLSL) are similar and easy to program in.**

**Requires a re-mapping of data to textures.**

**Very good results already in 2005: 8x speedups overall reported!**



## **CUDA-based GPGPU**

**Only works on NVidia hardware.**

**Requires extra software - which isn't very elegant.**

**Nice integration of CPU and GPU code in the same program.**

**Excellent results! 100x speedups are common - before optimizing! Even low-end GPUs give significant boosts.**



## **OpenCL-based GPGPU**

**Works on various hardware - not only GPUs.**

**Developed by Khronos Group, pushed by Apple.**

**Harder to get started, software looks pretty much like programming shaders.**



## **Compute shaders**

**Built into OpenGL**

**Similar to OpenCL**

**Good portability**





## **Use the source, Luke!**

**Three trivial examples:**

**Hello World! for CUDA**

**Hello World! for OpenCL**

**Hello World for GLSL**



## **In the Southfork lab**

**GTX660Ti**

**Kepler! Still respectable!**

**1300 cores!**

**Good mid-range board, about half the performance of a GTX 980. Great price/performance, lots of parallelism to play with, and pretty nice power consumption.**



## **In the Multicore lab**

### **GTX560Ti**

**Fermi! Good computing performance!**

**512 cores - but can still match a Kepler!**

**Alas, Fermi lacks some important feature of the Kepler. Thus, it is starting to feel a bit old.**



Information Coding / Computer Graphics, ISY, LiTH

**That's all, folks!**

**Next time: Introduction to CUDA**