General Purpose GPU symposium

Introducing GPU architectures

Henk Corporaal

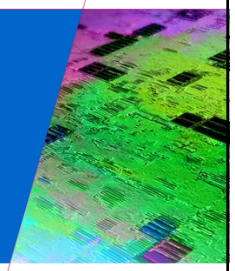

TU/e Technische Universiteit Eindhoven University of Technology

Where innovation starts

Why do we need GPUs?

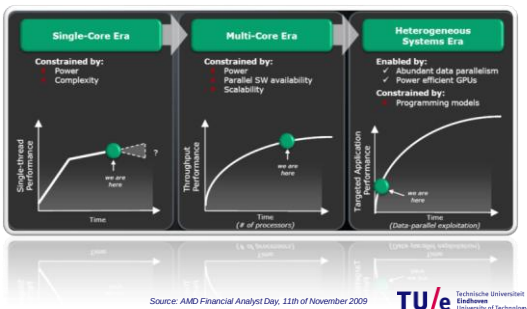

Source: AMD Financial Analyst Day, 11th of November 2009

What do GPUs bring us?

Acceleration for computer graphics

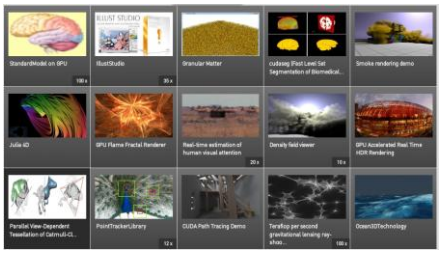

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

What do GPUs bring us?

Acceleration for image processing

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

What do GPUs bring us?

Acceleration for linear algebra

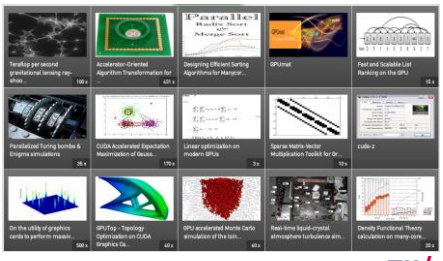

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

What do GPUs bring us?

Acceleration for molecular dynamics

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

What do GPUs bring us?

Acceleration for medical applications

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

What do GPUs bring us?

And many more other applications

Source: NVIDIA CUDA Zone, www.nvidia.com/cuda

Enormous computing capacity

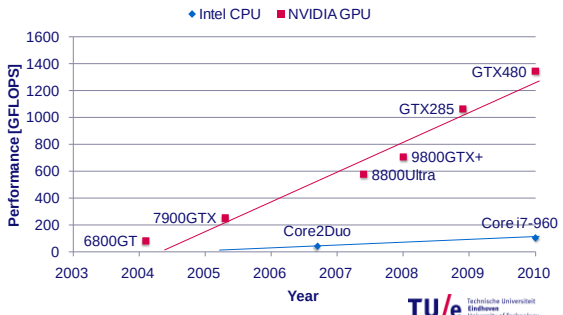

From system-on-chip to supercomputer

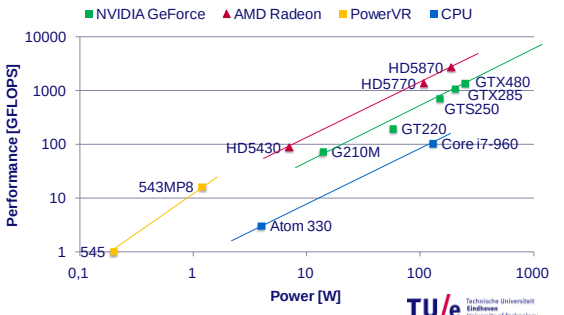

From system-on-chip to supercomputer

PowerVR low-power GPUs

From system-on-chip to supercomputer

Intel Larrabee x86-GPU

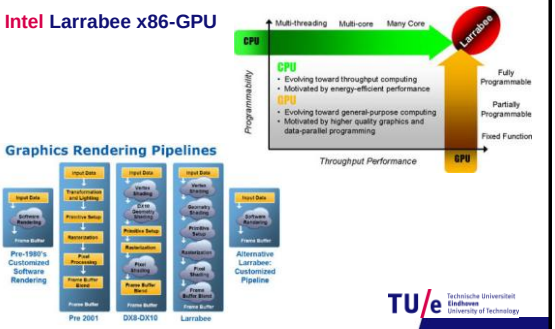

From system-on-chip to supercomputer

AMD ATI Radeon GPUs

1600 VLIW processing elements

2 TFLOPS per GPU

From system-on-chip to supercomputer

NVIDIA GeForce/Tesla GPUs

Tesla S2070: 4TFLOPS!

Programming languages

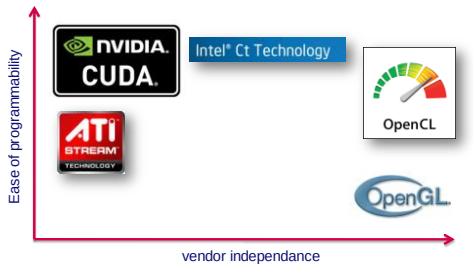

Interesting links

- **GPU research at the TU/e ES-group:** www.es.ele.tue.nl/gpu/
- **Mapping matrix multiplication on a GPU:** http://sites.google.com/site/5kk70gpu/matrixmul-example
- **Mapping reduction on a GPU:** http://developer.download.nvidia.com/compute/cuda/1_1/Website/projects/reduction/doc/reduction.pdf

Morning program

- **10:00 Introducing massively parallel processing using GPUs**
  Gert-Jan van den Braak (Electronic Systems, TU/e)
  Cedric Nugteren (Electronic Systems, TU/e)
- **10:30 GPU programming paradigms**
  Wouter Caarls (Biorobotics Lab, TUDelft)
- **11:00 Coffee Break**
- **11:30 Cross platform GP-GPU with OpenCL**
  George van Venrooij (Organic Vectory)
- **12:00 Automatic parallelization of C-code**
  Jos van Eijndhoven (VectorFabrics)
- **12:30 Lunch Break**