General Purpose (computation) (on) GPUs

Why GPUs?

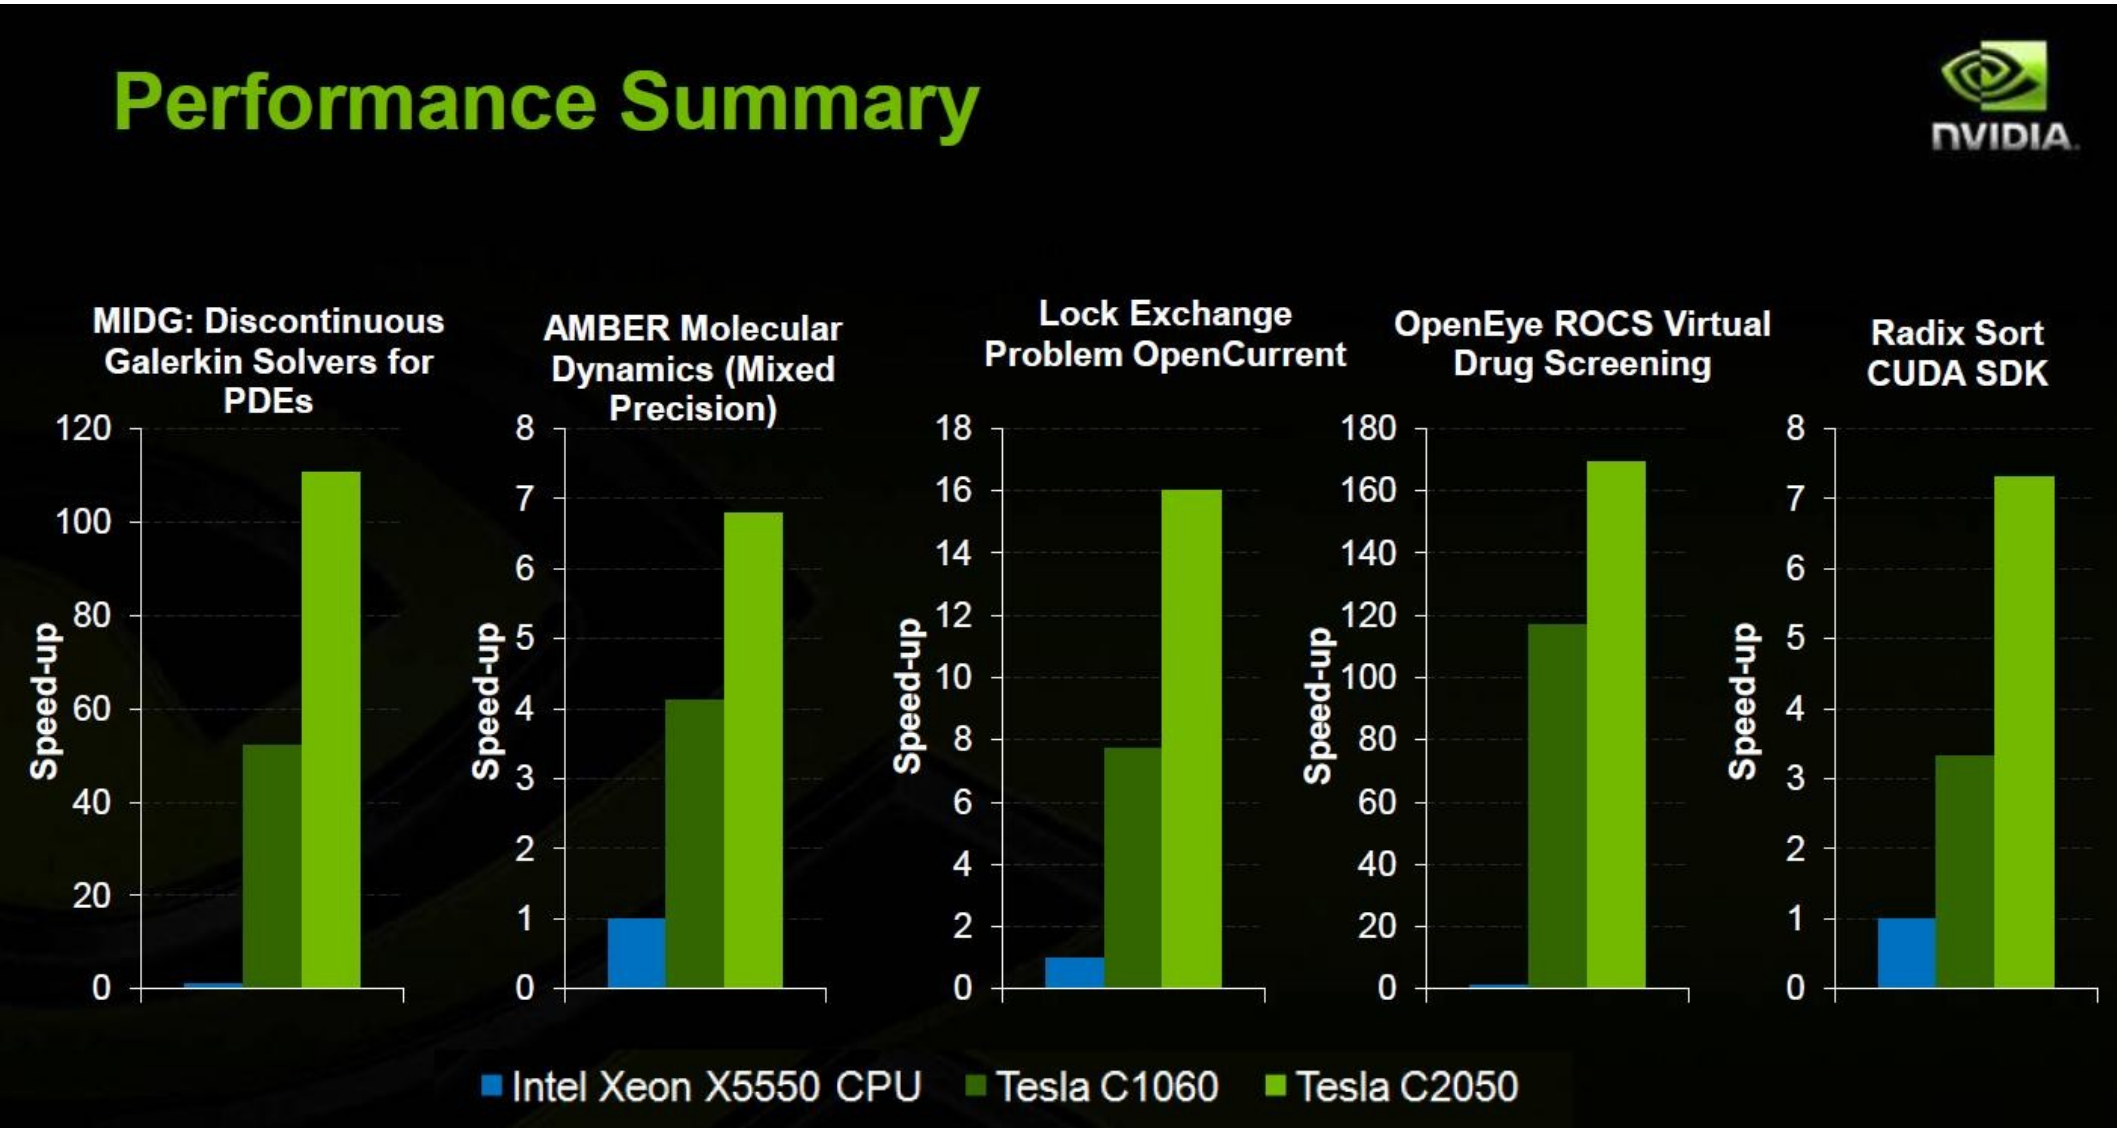
Performance Summary
NVIDIA
MIDG: Discontinuous Galerkin Solvers for PDEs
AMBER Molecular Dynamics (Mixed Precision)
Lock Exchange Problem OpenCurrent
OpenEye ROCS Virtual Drug Screening
Radix Sort CUDA SDK
Speed-up
0
20
40
60
80
100
120
1
2
3
4
5
6
7
8
10
12
14
16
18
100
120
140
160
180
Intel Xeon X5550 CPU
Tesla C1060
Tesla C2050

How?

4096 cores

2500+ cores

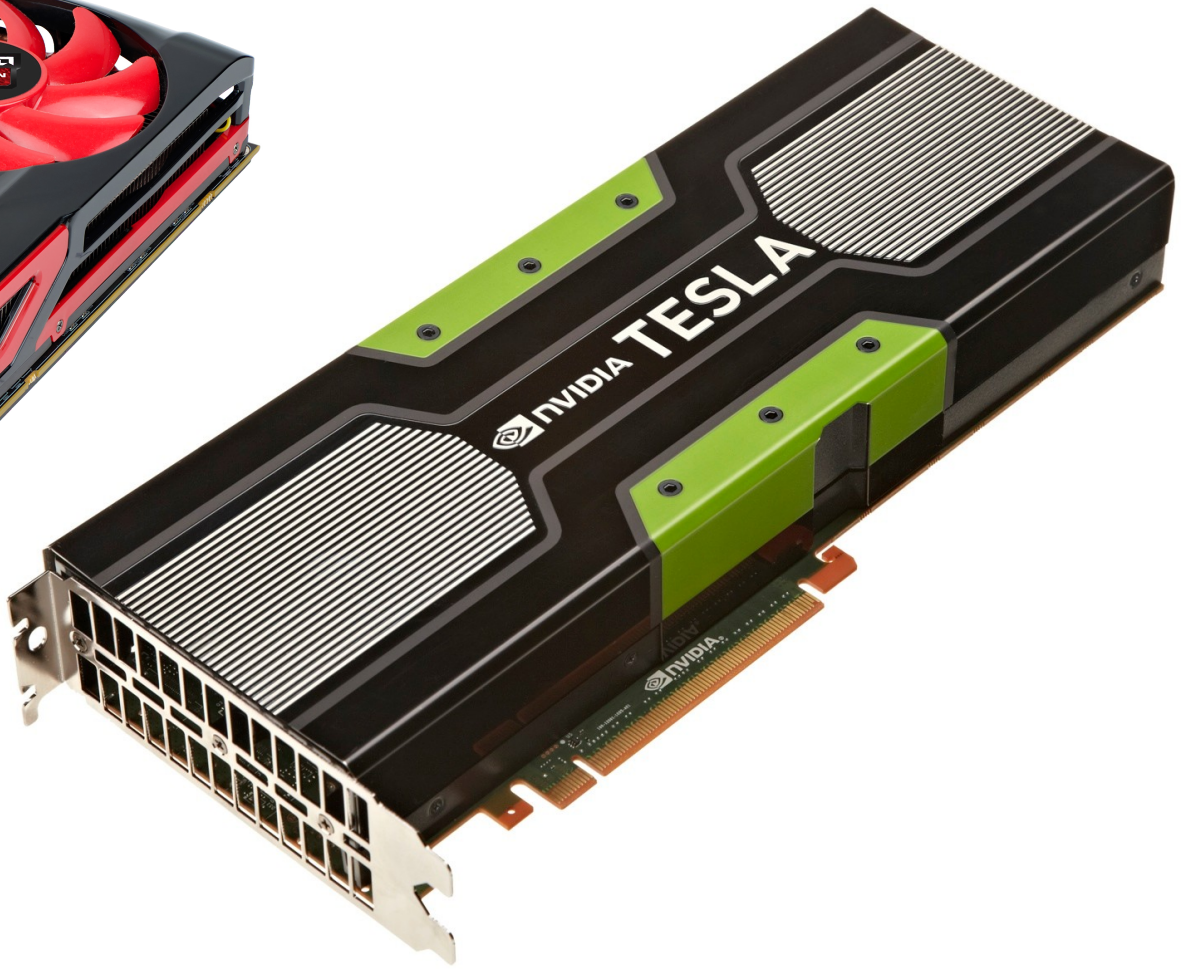

Massively Parallel

(Hundreds of?)Thousands of threads

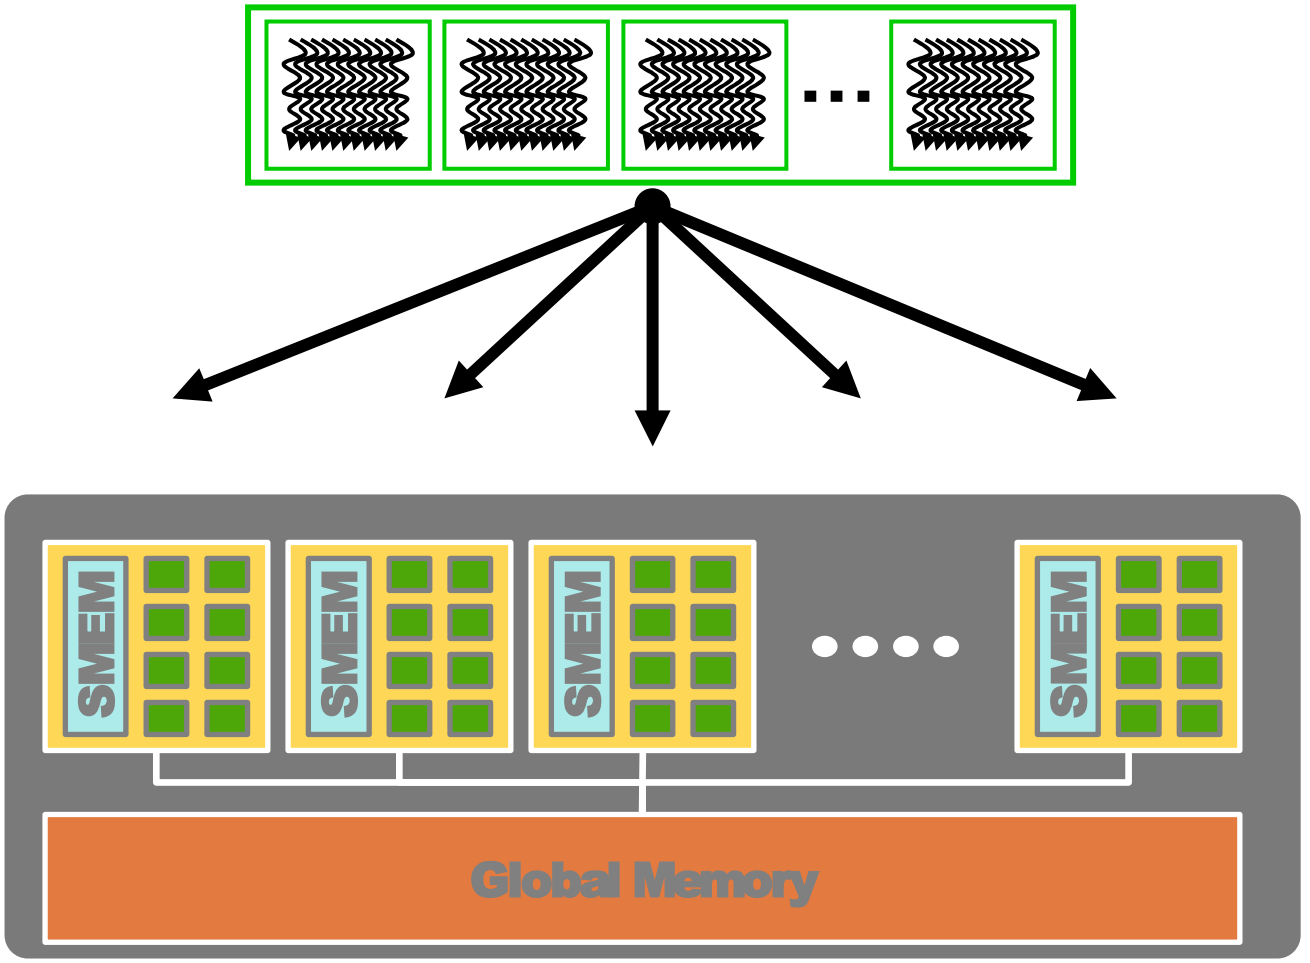

GPGPU

- Uni-DUE Forschungsprojekt
- Apply GPUs to your research
- Or just learn how they work

http://hpc.uni-due.de/teaching/st2014/gpgpu/

General Purpose (computation) (on) GPUs

Why GPUs?

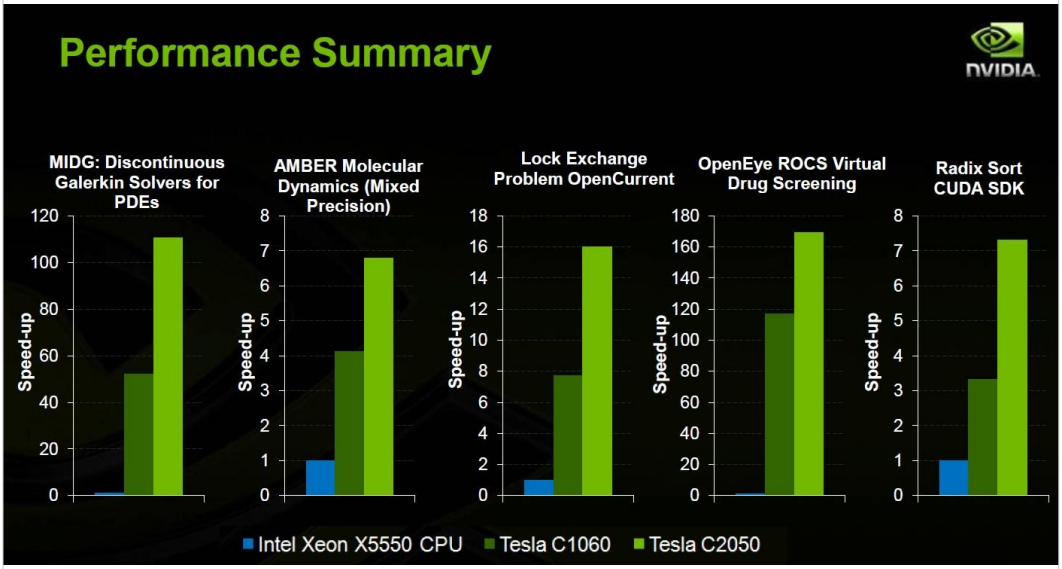

In recent years, the processing power of highly-SIMD RISC architectures, such as those employed by GPUs, has largely eclipsed conventional CPU performance.

How?

4096 cores

2500+ cores

The primary mechanism by which such performance is possible is due to the ability of manufacturers to pack so many cores in such a small area. This leads to accelerator cards which can quite literally run thousands of 'threads' at once.

Massively Parallel

(Hundreds of?)Thousands of threads

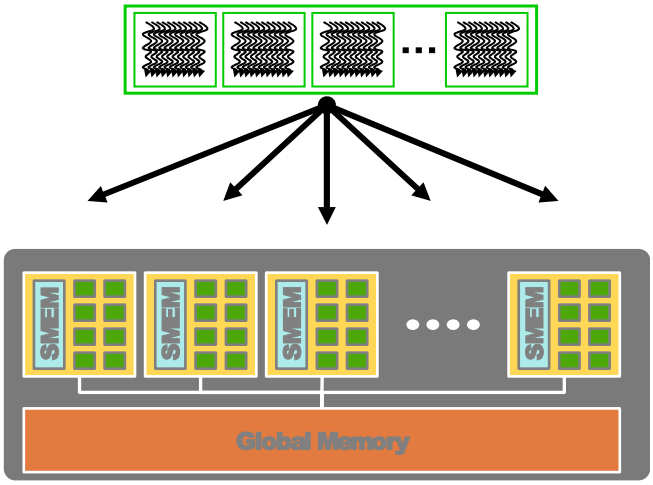

Of course, this processing power is only available if you know how to harness it. GPU processing requires developers to find parallelism at significantly finer scales than traditional techniques would require. Furthermore, programmers must be aware of architectural concerns, such as how threads may diverge and how to use memory efficiently, to get the peak performance from these processors.

GPGPU

- Uni-DUE Forschungsprojekt
- Apply GPUs to your research
- Or just learn how they work

http://hpc.uni-due.de/teaching/st2014/gpgpu/