


# Scientific Computing on GPUs

## Coding Demos

Dominik Götdeke, Jakub Kurzak, Jan-Philipp Weiß,  
André Heidekrüger and Tim Schröder



PPAM 2011 Tutorial  
Toruń, Poland, September 11  
<http://gpgpu.org/ppam11>



# Hands-on Concept

---

## Goals

- Getting a first impression how OpenCL code looks like
- Implementation of a simple vector addition for two different architectures

## Programming for performance

- Different architectures require different strategies
- OpenCL is *not* performance-portable!
  - Execution configuration
  - Data layout, memory access patterns
  - ...
- Covered by next two talks and the sample code

# Exercise 1

---

## Goals

- OpenCL concepts: Platforms and devices
- Introduction to our 'framework'
- Code will be available on the webpage

## Code

- 01-devicequery.c
- Live walkthrough
- Demo on a heterogeneous node in Dortmund

# Exercise 2

---

## Goals

- OpenCL concepts
  - Contexts, command queues
  - Kernel code, building and compiling OpenCL program objects, kernels, etc.
  - Device memory management
  - Execution configuration

## Code

- 02-axpy.c
- Live walkthrough and demo

## Tricky bits

- Different optimisations and execution configurations
- CPUs (and AMD GPUs): float4 to *explicitly* address SSE/VLIW
- NVIDIA GPUs: scalar code as CUDA cores are scalar