



2-day OpenACC Course at the Karlsruhe Institute of Technology

(27-28 January 2016)

CUDA and OpenCL are two programming languages that allow to develop GPU applications with great level of detail, but require advanced knowledge of GPU architecture and extensive code re-writing. Many cross-platform or prototype applications do not actually require deep GPU porting, if they perform data processing using standard algorithms, e.g. reduction. In such cases, fair performance numbers could be achieved with simplified *directive-based* language extensions. OpenACC is de facto an industry standard for C/Fortran directive extensions for porting code to GPUs. This training course walks through OpenACC programming technology, from basics to advanced practices used in real applications.

Hands-ons: All discussed topics will be accompanied with practical sessions, using PGI OpenACC compiler with NVIDIA backend. Exercises will be conducted either on attendee's preferred systems, or on the provided GPU server with NVIDIA Tesla K40 and K20 GPUs: <https://parallel-computing.pro/test-drive>.

All corresponding presentations and code samples will be available to attendees from the beginning of each training day.

Applied Parallel Computing LLC is delivering GPU training courses since 2009. Several dozens of courses have been organized all over Europe, both for commercial and academic customers. We work in close partnership with NVIDIA, CUDA Centers of Excellence and Tesla Preferred Partners. In addition to trainings, our company provides GPU porting/optimization services and [CUDA certification](#).

Day 1: Introduction to OpenACC

Morning (09:00-12:30)

09:00-10:30: lecture

- Advantages of OpenACC.
- Execution model: gangs, workers and vectors – three levels of coarse-grain and fine-grain parallelism. SIMD instructions.
- OpenACC memory model: host and accelerator address spaces.
- OpenACC directive syntax in C and Fortran. Main directives: *parallel* and *kernels* – offloading code regions to accelerator, *loop* – detailed parallelization parameters for each loop.
- Examples of *vector addition* and *reduction* explained.

10:30-12:30: [hands-on session](#)

- PGI compiler for NVIDIA and Radeon GPUs. GCC open-source OpenACC compiler.

- Compilation and deployment of OpenACC code examples.

Afternoon (13:30-17:00)

13:30-15:00: lecture

- Understanding PGI OpenACC compiler output. Compiler flags and environment variables for detailed analysis and performance reports. Compiler limitations in dependencies tracking, enforcing data independence.
- Data clauses: *deviceptr*, *copy*, *copyin*, *copyout*, *create*, *delete*, *present*, *present_or_(-copy, -copyin, -copyout, -create)*.
- Organizing data persistence regions using *data* directives.

15:00-17:00: [hands-on session](#)

- Profiling GPU kernels in OpenACC application. *Time* option, *PGI_ACC_TIME* environment variable. Profiling with *pgprof*.

- “Fill-in” exercise on implementing wave propagation stencil in OpenACC (*wave13pt*). Adding OpenACC directives step by step.

Day 2: Advanced OpenACC programming

Morning (09:00-12:30)

09:00-10:30: lecture

- Non-structured data lifetime with *enter data* and *exit data* directives.
- Additional data management directives: *cache*, *update*, *declare*.
- Organizing asynchronous execution using *async* clause and *wait* directive.
- *Atomic* directive. Allowed operations. Restrictions.

10:30-11:15: lecture

- Comparison of OpenACC and OpenMP 4 main directives: *parallel* and *kernels vs parallel* and *target*
- Comparison of OpenACC and OpenMP 4 data clauses: *copy*, *copyin*, *copyout* vs *map* and other directives
- Comparison of OpenACC and OpenMP 4 data persistence directives: *data* vs *target data* and *declare target*
- Getting LLVM-based OpenMP 4 compiler to work on NVIDIA GPUs
- Side by side comparison of *vector addition*, *reduction* and *wave13pt* examples in OpenACC and OpenMP 4.

11:15-12:30: hands-on session

- Advanced code optimization practices, by examples. Restrictions, common mistakes, workarounds.
- PGI OpenACC compiler pipeline. Accelerator-specific compiler flags, reviewing intermediate code.

Prerequisites

- Open (unfirewalled) connection to farm.parallel-computing.pro (188.44.42.33) over port 22
- Web browser capable for using server command line interface (Chrome or Firefox)
- (optional) PuTTY SSH client
- OneDrive or Dropbox to share files between client and server

Afternoon (13:30-17:00)

13:30-15:00: lecture

- OpenACC interoperation with CUDA C, CUDA Fortran and GPU-enabled libraries.
- External dependencies in OpenACC kernels, functions inlining. OpenACC *routine* directive, separate compilation and procedure calls.
- Accessing global variables.

15:00-16:30: hands-on session

- Jacobi iterative method with OpenACC. Step-by-step optimization.

16:30-17:00: Q & A