

Task 2.2:

**Increasing the performance portability of
flagship codes for different accelerators**

30/05/2023



ChEESE-2P



Task	Task name	Start	End	Lead
Task 2.1	Code audits (Jose Gracia, HLRS)	3 (27)	15 (36)	HLRS
Task 2.2	Increasing the performance portability of flagship codes for different accelerators (Giorgio Amati, Luca Ferraro, CINECA)	1	48	CIN
Task 2.3	Optimisation for exascale (Piero Lanucara, CINECA)	6	48	CIN
Task 2.4	Continuous Performance Efficiency Monitoring (Marta Garcia, BSC)	18	48	BSC

Effort (PMs)													
CIN	CSIC	BSC	INGV	HLRS	TUM	LMU	UMA	CNRS	IPGP	BULL	CSC	UM	SU
67	12	18	18	18	22	14	12	28	1	14	8	8	7



Training/webinar

Planned Webinar in June (date to be confirmed)

- Remote
- Session will be recorded

13/06/2023:

- Porting to GPU a Fortran code using intrinsic parallelism

28/06/2023:

- OpenCL GPU Programming for HPC Applications

Forthcoming Webinar (July?)

- OpenMP offload portability performance



Porting to GPU a Fortran code using intrinsic parallelism

Time

- 13/6 , 10:30-12:30 (with Q&A session)

Description

- In this webinar we'll present the main features and issues of Fortran intrinsic parallelism for GPU and some figures of GPU performance for a 2D (mini-app) and a 3D cfd-code. Comparison with classical OpenAcc performance will be presented

Target audience and prerequisites

- Knowledge of Fortran

Instructor

- Giorgio Amati

Agenda

1. Porting to GPU a simple kernel
2. A more complex code: a Kinetic 2D CFD code
3. Porting to GPU a real 3D CFD code



OpenCL GPU Programming for HPC Applications/1

Time

- 28/6 , 10:00-12:00 (with Q&A session)

Description

- OpenCL for GPU computing has been flying under the radar ever since CUDA became popular and the myth spread that it would more efficient. With the three major vendors, Nvidia, AMD and Intel, now all providing powerful data-center GPUs in the high-performance computing sector, it becomes more important than ever to be able to write GPU software once and deploy it on all supercomputers without wasting any time on porting. In this workshop, I will introduce OpenCL and highlight the significant benefits that OpenCL offers over other proprietary languages, in particular it's excellent compatibility across all GPU hardware since 2009. On a brief case study, I will teach the main techniques of OpenCL performance optimization up to the physical roofline model limit. The workshop will conclude with an open discussion, where participants can seek advice for their own applications.



OpenCL GPU Programming for HPC Applications/2

Target audience and prerequisites

- Researchers using HPC; basic programming experience in C/C++ would be good

Instructor

- Moritz Lehmann

Agenda

1. Introduction to OpenCL,
2. GPU performance optimization techniques,
3. Multi-GPU parallelization,
4. Case-study,
5. Open discussion

Reference

- <https://github.com/ProjectPhysX/FluidX3D>
- [movie1](#)



Auxiliary Slide

