

Joint INTEL / EuroCC-Castiel oneAPI Workshop

February 16th and 17th, 2022

Day 2

Edmund Preiss
SW Dev Tools BDM
edmund.preiss@intel.com



Goals and Objectives

- Provide insights into heterogenous (CPU and GPU) software programming with the Intel oneAPI Developments tools
 - Share knowledge from Intel specialists and from external speaker
 - Interactive Sessions with Q+A / Use Chat Room
 - Contact us if you have specific (i.e. feature) requests (edmund.preiss@intel.com)
 - In case you like to get
 - updated Intel SW Dev tools information and /or
 - Invitations to other local SW Dev tools training events
- then please share your email address with us (edmund.preiss@intel.com)

This is a ,public‘ event –
No Confidential Material to be shared

Agenda – Day 2

TOPIC			Presenter
09:05	00:05	Welcome and Introduction to Day 2	Edmund Preiss (Intel), EuroCC
10:05	01:00	A 3rd Party oneAPI Case Study: GROMACS	Andrey Alekseenko (KTH, Sweden)
11:05	01:00	Intel OpenMP for Fortran Apps – with Demos - Parallelizing heterogenous applications with Intel OpenMP and OpenMP offloading	Alina Shadrina (Intel)
11:10	00:05	<i>Bio Break</i>	
11:50	00:40	Intel HW (public) roadmap/XPU and architecture specifics - Server CPUs - Client CPUs (i.e., Alder Lake; iCore with integrated graphics) - Intel hardware accelerators GPUs (DG1, SG1, ATS) and Intel FPGAs	Jean-Laurent Philippe (Intel)
13:00	01:10	<i>Lunch Break</i>	
14:15	01:15	Application profiling for heterogenous hardware - Demos - Profile DPC++ and GPU Workload Intel VTune Profiler and Intel VTune Offload - Share experiences/key findings with Gromacs related porting and optimization efforts	Heinrich Bockhorst (Intel)
14:20	00:05	<i>Bio Break</i>	
15:35	01:15	Application profiling for heterogenous hardware - Demos - Profile DPC++ and GPU Workload with Intel Advisor include and Roofline analyser - Estimate performance potential gains with Offload Advisor (CPU -> HW Accelerator)	Klaus-Dieter Oertel (Intel)
15:40	00:05	<i>Bio Break</i>	
16:40	01:00	Programming for Distributed HPC Systems using Intel MPI	Dmitry Sivkov (Intel)
17:10	00:30	Dynamic Debugging with Intel Inspector - Demos - Identifying Memory and Threading Errors (Data Races and Deadlocks)	Heinrich Bockhorst (Intel)
17:15	00:05	- Questions and Answers - Wrap up	Intel

Additional 3rd Party Links of Recordings (including PDF)

■ Application: easyWave ; Presenter : Steffen Christgau (ZIB)

[A oneAPI Case Study: easyWave - A Tsunami Simulations Application](#)

[with demos ZIB's experiences with Intel oneAPI](#)

- Porting a tsunami application from CUDA to DPC++
- Running DPC++ code on GPU
- From CUDA to DPC++ back to Nvidia GPUs... and FPGAs
- A oneAPI case study with the tsunami simulation easyWave

Live Demos: [Christgau_playbook.tar.gz](#)

A oneAPI Case Study: easyWave

A Tsunami Simulation Application

Steffen Christgau

Supercomputing Department
Zuse Institute Berlin

■ Application: Ginkgo (Sparse Library), Presenter : Prof. Hartwig Anzt (KIT)

[A oneAPI Case Study: Ginkgo – a sparse linear algebra library for OneAPI Hardware](#)

Experience and potential of using oneAPI for the Ginkgo sparse linear algebra library

- Porting of numerical linear algebra kernels from Cuda to DPC++
- Pitfalls and solutions in the OneAPI Compatibility Tool
- Performance of Ginkgo's DPC++ backend on Intel GPU
- Live demo running Ginkgo in the Intel DevCloud



Ginkgo – a sparse linear algebra library for oneAPI hardware

Click on the Images to start the videos

