

Multiarchitecture Programming for Accelerated Compute, Freedom of Choice for Hardware



# Welcome to the Intel/LRZ oneAPI Workshop

Software & Advanced Technologies Group (SATG)  
Software Products & Ecosystem  
June 2023

Edmund Preiß  
Software Dev Tools Business Development Mgr.



# Workshop Objectives

- Understand and apply:
  - The oneAPI programming model for mixed (heterogenous) Hardware Architectures
  - Intel's DevCloud – A Developer Sandbox
  - Building applications with ICX/DPC++/SYCL
  - Fundamentals of OpenMP offloading
  - How to use Intel's oneAPI libraries (oneMKL, ...) efficiently
  - Intel's profiling and performance analysis tools
  - Intel's DPC++ Compatibility tool - Migrate CUDA code to SYCL code

# Call to Action

- **Developers**

- Use your new Intel oneAPI Toolkits knowledge for application development and Code Modernization – i.e., for CPUs, dGPUs, ...
- Move CUDA code to SYCL – Let us know and work with us
  - [edmund.preiss@intel.com](mailto:edmund.preiss@intel.com)
- Create applications with Intel LLVM based C++ (ICX/DPC++) and Fortran (IFX) Compilers and use time saving and performant Libraries
- Practice with exercises available on Intel DevCloud

- **Data Centre Admins**

- Prepare and Update your data center with performance optimized Intel oneAPI Toolkits to serve your users and developers

# This event is sponsored by



intel.  
software elite  
reseller

<https://www.hocomputer.de>



Backup



**ho-COMPUTER thanks you for attending this session. Enjoy your new learning.**



**intel.**  
software elite  
reseller

# AGENDA

## Day 1: June 5<sup>th</sup>, 2023

TOPIC				Presenter
09:00	09:10	00:10	Welcome and Introduction to Day 1	Gerald Mathias (LRZ) Volker Weinberg Edmund Preiss (Intel)
09:10	09:30	00:20	oneAPI – Introduction to a mixed Architecture Development Environment - Motivation and oneAPI Standardization - Intel's oneAPI Toolkits Portfolio and Components - Intel oneAPI plug-ins for Nvidia and AMD hardware ( CPU and GPUs)	Edmund Preiss (Intel)
09:30	09:50	00:20	Introduction to the DevCloud - A sandbox for software development and benchmarking - Purpose: Demoing, testing and porting applications - Hardware and Software offerings - How to onboard & how to get a DevCloud account	Heinrich Bockhorst
09:50	10:40	00:50	Direct programming with oneAPI Compilers (Part 1) – with Demos - Intro to heterogenous programming model with SYCL 2020 - SYCL features and examples <ul style="list-style-type: none"> <li>o "Hello World" Example</li> <li>o Device Selection</li> <li>o Execution Model</li> <li>o Compilation and Execution Flow</li> <li>o Memory Model; Buffers, Unified Shared Memory (USM)</li> <li>o Performance optimizations with SYCL features</li> </ul>	Igor Vorobtsov (Intel)
10:40	11:10	00:30	<i>Break /Coffee</i>	
11:10	12:00	00:50	Direct programming with oneAPI Compilers (Part 2) – with Demos - Intro to heterogenous programming model with SYCL 2020 - SYCL features and examples <ul style="list-style-type: none"> <li>o "Hello World" Example</li> <li>o Device Selection</li> <li>o Execution Model</li> <li>o Compilation and Execution Flow</li> <li>o Memory Model; Buffers, Unified Shared Memory (USM)</li> <li>o Performance optimizations with SYCL features</li> </ul>	Igor Vorobtsov (Intel)
12:00	12:15	00:15	oneAPI Case Study - Seissol	Ravil Dorozhinski (TUM)
12:15	12:30	00:15	oneAPI Case Study - DPEcho	Salvatore Cielo (LRZ)
12:30	13:30	01:00	<i>Lunch (buffet Lunch in front of Auditorium)</i>	
13:30	14:45	01:15	Hands-on session applying SYCL and DPC++ knowledge (heterogenous programming) - practical exercises on dedicated target systems: <ul style="list-style-type: none"> <li>o install oneAPI tools</li> <li>o write your first SYCL programs</li> <li>o develop own kernels</li> <li>o test your own heterogenous programs</li> <li>o How to avoid errors</li> </ul>	Igor Vorobtsov + Matthias Kirchhart
14:45	15:15	00:30	<i>Break/Coffee</i>	Intel
15:15	16:30	01:15	Hands-on session - Programming with SYCL (continued) - practical exercises on dedicated target systems	Igor Vorobtsov + Matthias Kirchhart
16:30	16:45	00:15	Questions and Answers - Wrap up of the day	Intel
17:30	21:00	open	Social Networking Event	Intel+LRZ

# AGENDA

Day 2 : June 6<sup>th</sup>, 2023

TOPIC				
09:00	09:05	00:05	Welcome and Introduction to Day 2	Gerald Mathias (LRZ) Edmund Preiss (Intel)
09:05	10:05	01:00	Intel OpenMP for Offloading for Fortran – with Demos - Parallelizing heterogenous applications with OpenMP 5.2	Tobias Kloeffel (Intel)
10:05	10:40	00:35	Intel oneAPI libraries (oneMKL) for HPC - with demos - Performance optimized libraries for numerical simulations and other purposes	Gennady Fedorov (Intel)
10:40	11:10	00:30	<i>Break/Coffee</i>	
11:10	11:50	00:40	Target NVIDIA and AMD with oneAPI and SYCL Using SYCL based NVIDIA and AMD plugins with Demos	Rod Burns (Intel)
11:50	12:20	00:30	Open Source Compatibility tool for porting purposes(SYCLomatic) - with demo - Migration Cuda based GPU Applications to SYCL	Matthias Kirchhart (Intel)
12:20	12:45	00:25	Intel Debugging Tools for heterogenous programming ( CPU, GPU ) - with demos	Alina Shadrina / Pascal Baehr (Intel)
12:45	13:15	00:30	Programming for Distributed HPC Systems using Intel MPI	Rafael Lago (Intel)
13:15	14:15	01:00	<i>Lunch (buffet Lunch in front of Auditorium)+D60</i>	Intel
14:15	15:30	01:15	Hands-on session -OpenMP -Applying oneMKL lib with examples	Tobias Kloeffel + Gennady Fedorov (Intel)
15:30	16:00	00:30	<i>Break/Coffee</i>	Intel
16:00	16:45	00:45	Hands-on session - Using DPC+ Compatibility Tool	Matthias Kirchhart (Intel)
16:45	16:55	00:10	Questions and Answers - Wrap up of the day	Intel
17:00	17:45	00:45	Guided Tour of LRZ Cluster	LRZ



# AGENDA

Day 3 : June 7<sup>th</sup>, 2023

TOPIC			Presenter
09:00	09:05	00:05	Welcome and Introduction to Day 23 Gerald Mathias (LRZ) Edmund Preiss (Intel)
09:05	10:05	01:00	Application profiling for CPU and or mixed hardware with the Intel VTune - Demos - Vtune general / main functionality ( Hot spot analysis ,....) starting with CPU - Profiling Tools Interfaces for GPU - Profile heterogeneous SYCL/OpenMP Workloads with Intel VTune Profiler Rafael Lago
10:05	10:20	00:15	Break / Coffee
10:20	11:20	01:00	Application profiling for CPU and or mixed hardware with the Intel VTune - Demos - Vtune general / main functionality ( Hot spot analysis ,....) starting with CPU - Profiling Tools Interfaces for GPU - Profile heterogeneous SYCL/OpenMP Workloads with Intel VTune Profiler Rafael Lago
11:20	11:35	00:15	Break / Coffee
11:35	12:35	01:00	Application profiling for CPU and mixed hardware with the Intel Advisor - Demos - Advisor's main functionality ( Vectorization and Roofline ) starting with CPU - Estimate performance potential gains with Offload Advisor ( CPU -> HW Accelerator) - Analyse heterogeneous SYCL/OpenMP Workloads with Intel Advisor and Roofline analysis Stephen Blair Chappell (oneAPI external certified trainer)
12:35	14:05	01:30	Lunch (buffet Lunch in front of Auditorium)
14:05	15:05	01:00	Hands-on session VTune - CPU focus -Identifying hotspots and optimize apps Using the Profiler for System optimisation Rafael Lago + Stephen Blair Chappell
15:05	15:25	00:20	Break / Coffee
15:25	16:25	01:00	Hands-on session VTune (cont'd) -CPU focus Rafael Lago + Stephen Blair Chappell
16:25	16:35	00:10	Questions and Answers - Wrap up of the day Intel