Introduction to the Intel oneAPI Development Environment

Summary

16 - 17 February 2022, 9 am – 5 pm CET
Location: Virtual via Webex
Organiser: CASTIEL WP3 – Training, Twinning and Mentoring

Background

Intel[1] and EuroCC and CASTIEL projects[2] agreed to organise a two-day workshop about oneAPI to the members of EuroCC as a follow up of the introduction of Intel to EuroCC that took place in September 2021. It was agreed that an online workshop presentation with demos would take place on 16 - 17 February 2022, 9 am – 5 pm CET. It was open to anyone in the EuroCC and CASTIEL network interested in attending.

Intel launched the oneAPI initiative in 2019. The oneAPI initiative is a cross-industry, open, standards-based unified programming model that delivers a common developer experience across vendor independent accelerator architectures — for faster application performance, more productivity, and greater innovation.

The oneAPI industry initiative encourages collaboration on the oneAPI specification and compatible oneAPI implementations across the ecosystem.

End of 2020 and end of 2021, Intel released the first major two editions of various Intel® oneAPI Toolkits (Intel oneAPi 2021 and the most current Intel oneAPi 2022). These toolkits include the programming languages, programming models, libraries and tools that are built to the oneAPI specifications targeting Intel CPUs and Intel accelerators (integrated and discrete Intel GPUs and Intel FPGAs).

The Intel® oneAPI Base & HPC Toolkit solution is replacing the Intel® Parallel Studio XE tools suite. The Intel oneAPI Base&HPC Toolkit provides high performance compilers for Data Parallel C++/C/C++, Fortran, OpenMP and the performance libraries such as Intel® oneAPI Math Kernel Library and Intel® MPI Library and analysis tools all enhanced to support heterogeneous development. The Intel oneAPi Toolkit family though also has dedicated derivatives for photorealistic Rendering (oneAPI Rendering Toolkit) and for artificial intelligence purposes (oneAPI Ai Analytics Toolkit and the OpenVino development environment).

The workshop wanted to provide information about the oneAPI initiative and how to get on a standards-based path for heterogeneous programming.

Agenda

[1] https://www.intel.de/content/www/de/de/homepage.html
[2] https://www.eurocc-access.eu/about-us/the-projects/

Day 1, 16th February, 2022

			TOPIC	Presenter
9:00	9:10	0:10	Welcome and Introduction to Day 1	Edmund Preiss (Intel), Martina Blazkova (EuroCC-CASTIEL)
9:10	9:40	0:30	oneAPI – Introduction to a new Development Environment - Concept and oneAPI Standardization initiative - Intel's Tools Implementation – Intel oneAPI Toolkits and libs - Transition from Intel Parallel Studio XE to Intel oneAPI toolkits	Edmund Preiss (Intel)
9:40	10:00	0:20	Introduction to the DevCloud - Purpose: Demoing, testing and porting applications - Hardware and Software offerings - How to onboard & how to get an DevCloud account	Klaus-Dieter Oertel (Intel)
10:00	*10:05*	*0:05*	*Bio Break*	
10:05	12:05	2:00	Direct programming with oneAPI Compilers – with Demos - Intro to DPC++ heterogenous programming model and SYCL 2020 - Fortran and offloading with OpenMP5.0 - DPC++ features and examples o "Hello World" Example o Device Selection o Compilation and Execution Flow o Execution Model o Memory Model; Buffers, Unified Shared Memory (USM) o Kernel Model	Igor Vorobtsov (Intel)
12:05	*13:15*	*1:10*	*Lunch Break*	
13:15	14:15	1:00	Compatibility tool for porting purposes - with demo - Migration Cuda based GPU Applications to oneAPI (DPC++)	Igor Vorobtsov (Intel)
14:15	15:00	0:45	Intel oneAPI libraries (oneMKL, oneDAL, …) - with demos - Performance optimized libraries for AI, numerical simulations and other purposes	Gennady Fedorov (Intel)
15:00	*15:05*	*0:05*	*Bio Break*	
15:05	15:35	0:30	Debugging on GPU /Debugging on LLVM (TBD)	Alina Shadrina (Intel)
15:35	16:35	1:00	Intel's oneAPI AI Analytics Toolkit solutions for Artificial Intelligence (AI) - Intel performance optimized Python Distribution - Low level libs with basic oneDAL, py4Dal - Optimized ML libraries such as XGBoost and SciKit Learn, Modin - Optimized Tools Frameworks for Deep Learning	Shailen Sohbee (Intel)
16:35	16:40	0:05	Questions and Answers, Wrap up	Intel

Day 2, 17th February, 2022

			TOPIC	Presenter
9:00	9:05	0:05	Welcome and Introduction to Day 2	Edmund Preiss (Intel)
9:05	10:05	1:00	A 3rd Party oneAPI Case Study: GROMACS	Andrey Alekseenko (KTH, Sweden)
10:05	11:05	1:00	Intel OpenMP for Fortran Apps – with Demos - Parallelizing heterogenous applications with Intel OpenMP and OpenMP offloading	Alina Shadrina (Intel)
11:05	*11:10*	*0:05*	*Bio Break*	
11:10	11:50	0:40	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)
11:50	*13:00*	*1:10*	*Lunch Break*	
13:00	14:15	1:15	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)
14:15	*14:20*	*0:05*	*Bio Break*	
14:20	15:35	1:15	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)
15:35	*15:40*	*0:05*	*Bio Break*	
15:40	16:40	1:00	Programming for Distributed HPC Systems using Intel MPI	Dmitry Sivkov (Intel)
16:40	17:10	0:30	Dynamic Debugging with Intel Inspector - Demos - Identifying Memory and Threading Errors (Data Races and Deadlocks)	Heinrich Bockhorst (Intel)
17:10	17:15	0:05	Questions and Answers, Wrap up	Intel

Participants

The workshop was attended by 93 people (Day 1) and 81 people (Day 2) from EuroCC and CASTIEL. Representatives from different National Competences Centres participated in the workshop, e.g. Bulgaria, Germany, Finland, Turkey, Belgium, Austria, Lithuania, Romania, Czech Republic, Spain, Portugal, Croatia, North Macedonia, Sweden, Poland, Slovakia, Greece, Latvia, Luxembourg, Norway, Slovenia, Hungary, Montenegro.

Content of the Workshop

The workshop started with welcome from Martina Blazkova, CASTIEL WP3 and from Edmund Preiss, Manager Business Development SSG, EMEA at Intel Corporation. Edmund Preiss then briefly introduced the goals of the workshop. The workshop intended to teach the participants about the latest and expanded features of the Intel oneAPI toolkits. The presenters showed the participants how to use the tools for shared and distributed computing on heterogeneous hardware platforms including Intel CPUs, Intel hardware accelerators and Intel discrete graphics solutions.

The workshop covered the following topics:

- Intel oneAPI – Introduction and Overview of the New Development Environment
- Introduction to the Intel Dev Cloud for Testing and Porting Applications
- Direct programming with oneAPI Compilers DPC++ and Fortran/OpenMP 5 (with demos)

- oneAPI Compatibility Tool for porting CUDA applications to DPC++ (with demos)
- Intel oneAPI libraries oneMKL, oneDAL and others (with demos)
- oneAPI case study: easyWave, a library for tsunami simulations – from CUDA to DPC++ (with demos)
- Intel OpenMP 5 for Offloading to Accelerators in Fortran codes (with demos)
- Intel Hardware Roadmap/XPU and Architecture Specifics
- Application Profiling for Heterogeneous Hardware (with demos)
- Using performance optimized Intel oneAPI AI Analytics Toolkit for Classical Machine Learning and Deep Learning both for Training and Inference
- oneAPI case study: GROMACS

The attendees should understand:

- The transition from the Intel Parallel Studio Development tool to Intel oneAPI Toolkits (HPC, AI and Rendering Toolkits)
- How to use the Intel Toolskits to develop heterogeneous applications running on CPUs and HW accelerators
- Make use of the offered development environments including the Intel Dev Cloud 'sandbox'
- Being able to start developing and porting own applications with the oneAPI development environment

There was also enough time for answering questions from the participants on both days. The workshop finished with asking for feedback that was collected through a short online questionnaire.

Outcome

Intel considers the event a major success and time well spent. There was a high attendance of the target audience coming from many data centres across Europe where the oneAPI concept and the details and benefits of the Intel oneAPI Toolkits were presented.

Further, the feedback Intel has received from questionnaire was very encouraging. The average rating of 4.3 (vs 5 being the best rating) was a great testament for the perceived presentation quality of the workshop.

A feedback analysis session was carried out after the event that revealed aspects of what went well and what can be improved in the future. Notably to mention are:

1- What went well? /Highlights mentioned:
   - Demos
   - Presentations
   - (Technical) Expertise of presenters
2- Areas for improvements:
   - Splitting up workshops in half days
   - Longer bio breaks
   - Slower pace
   - Even more demos (and hands-on)
   - Better prep to avoid minor technical glitches

Intel would be more than happy engaging with EuroCC in another tools workshop. A viable target might be around Intel tools and expertise on Artificial Intelligence (both in terms of Deep Learning and Classical Machine Learning).

The Intel presentations were shared with all participants from the EuroCC and CASTIEL network. The recording and the presentations will be accessible on the EuroCC Access website[3].

In case of NCCs' interest to explore more about Intel or to start a collaboration with them, the NCCs will contact Edmund Preiss directly - edmund.preiss@intel.com. Some NCCs approached Intel directly after the event and individual follow up activities were defined.

Intel and CASTIEL WP3 will meet to discuss a potential further collaboration.

[3] https://www.eurocc-access.eu/services/resources/

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 951732. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, United Kingdom, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Switzerland, Turkey, Republic of North Macedonia, Iceland, Montenegro.