INTEL® CODE MODERNIZATION WORKSHOP 2016

Version 12 – Sep 12th

2-DAYS WORKSHOP IN AMSTERDAM

20 & 21 September 2016 at Rockstart Amsterdam
Herengracht 182 - 1016 BR Amsterdam, NL

rockstart.

More information and registration: http://www.inteldevconference.com

DAY 1 AGENDA TUESDAY 20.09.2016

08:30	10:00	Registration and Breakfast
10:00	10:15	**Welcome – Day 1**
10:15	10:45	**DELIVERING LEADING PERFORMANCE IN COMPUTING, ANALYTICS AND MACHINE LEARNING – ALL ABOUT INTEL® XEON™ & XEON PHI™** Learn about the parallel architecture, technical advances and features of the latest and future Intel processors, especially Intel® Xeon™ and Intel® Xeon Phi™ (aka Knights Landing, KNL). Presenters: TBD, Intel & Janko Strassburg, Bayncore
10:45	11:15	**TRENDS IN EFFICIENT PARALLEL COMPUTING AND PERFORMANCE** With all the advances in massively parallel and multi-core computing with CPUs and accelerators it is often overlooked whether the computational work is being done in an efficient manner. This presentation shows the well-known laws of parallel performance from the perspective of a system builder. It also covers through the use of real case studies, examples of how to program for energy efficient parallel application performance. Presenter: Martin Hilgeman, DELL
11:15	11:45	Coffee Break
11:45	12:30	**VECTORIZATION – NEW APPROACHES FOR PARALLELISM AT CORE LEVEL (SIMD)** Vectorization is one of the critical elements to maximize parallel performance. In this session we will show how to get started with vectorization and avoid common pitfalls. We will review cases where automatic vectorization fails, providing tips and best known methods for effective vectorization. This session is also illustrated with a couple of real-world case examples by using Intel® Parallel Studio XE suite. Presenter: Janko Strassburg, Bayncore
12:30	13:15	**CODING HIGH PERFORMANCE PYTHON* FOR DATA-INTENSIVE ALGORITHMS** This talk will introduce the recently released Intel® Distribution for Python which delivers high performance acceleration for scientific computing, data analytics, and machine learning. Learn how NumPy/SciPy can now leverage the full performance potential of parallel CPU architecture by linking performance libraries like Intel® MKL (Math Kernel Library), Intel® MPI (Message Passing Interface), Intel® TBB (Threading Building Blocks) and Intel® DAAL (Data Analytics Acceleration Library). Presenter: Andrey Nikolaev, Intel
13:15	14:15	Lunch
14:15	15:00	**PROFILING AND TUNING PYTHON* CODE FOR MAXIMUM PERFORMANCE** Performance bottlenecks and underutilization of compute resources result in higher cost per performance unit and watt. In this session you will learn how to identify and eliminate hotspots and realize performance potential in Python code using Intel® VTune™ Amplifier XE. Presenter: Stephen Blair-Chappell, Bayncore
15:00	15:30	Coffee Break
15:30	16:15	**ENABLING MACHINE LEARNING ON INTEL ARCHITECTURE** In the first part of this session we will discuss the overall ecosystem of Machine Learning algorithms, challenges and opportunities brought by the Big Data era, and how they are addressed by Intel software solutions. In the second part of the presentation we will review capabilities of Intel libraries that address Big Data problems. The rest of the session devotes to deep overview of Intel® Data Analytics Acceleration Library, software solution that provides the building blocks for all stages of the data processing, from data acquisition till modelling and scoring. Architectural aspects, content of the capabilities, and performance of the library on Intel's highly parallel CPUs will be discussed in this session also. Presenter: Andrey Nikolaev, Intel
16:15	17:00	**ENABLING DEEP LEARNING ON INTEL ARCHITECTURE** The first part of the session will discuss the capabilities available in Intel SW to enable high performance in Deep Learning applications. The rest of the session will review the features of two libraries, Intel® Math Kernel Library and Intel® Data Analytics Acceleration Library for support of Deep Learning, the performance of popular frameworks such as Caffe enabled with Intel SW, and explain similarities and differences of two solutions. Presenter: Andrey Nikolaev, Intel
17:00	17:30	**Open Q&A session**
17:30	20:00	**Networking with drinks & finger food**

INTEL® CODE MODERNIZATION WORKSHOP 2016

DAY 2 AGENDA WEDNESDAY 21.09.2016

08:30	**09:30**	Registration and Breakfast
09:30	**10:15**	**WHATS NEW IN IPS XE 2017?** Learn about some of the new features of Intel Parallel Studio XE 2017. We'll also take a look at some upcoming technologies, such as the roofline model support in Advisor Presenter: Stephen Blair-Chappell, Bayncore
10:15	**11:00**	**CASE STUDY – IMPORTANCE OF SOFTWARE TO ASML** Since 2007, John Koster (Senior Director Software Development at ASML) is department manager at ASML Development & Engineering. In this position he is leading the software research group. Special interest areas are software architecture, software metrology and software systems. Before, he did his studies and PhD in Mathematics and Computing Science at TU/e. He started his professional career at Philips Consumer Electronics, Eindhoven, where he held various engineering and project management positions in the area of digital video. Afterwards, he worked at Philips Semiconductors, Sunnyvale, responsible for software platform developments, followed by multiple product launches into European and American markets. John will be addressing the importance of software to ASML, which is often perceived as a hardware/machine company Presenter: John Koster, ASML
11:00	**11:30**	Coffee Break
11:30	**12:30**	**TUTORIAL - PROFILING AND IMPROVING ADVANCED VECTORIZED CODE** Using the latest enhancements to Intel Vectorization Advisor, this session shows how to detect vectorisation issues and strategies to overcome them. Included in this session is a practical hands-on example using DL_MESO Lattice Boltzmann code from Daresbury Labs. Presenter: Stephen Blair-Chappell, Bayncore
12:30	**13:30**	Lunch Break
13:30	**14:15**	**CASE STUDY – PARALLELIZING THE BLACK-SCHOLES MODEL** In this session, we will demonstrate how to go from serial code to optimized parallel code with a classic algorithm (the Black-Scholes option pricing model) using open standards, libraries and tools, to finally efficiently run on multi-core and many-core CPUs, including the latest Intel® Xeon Phi™ (KNL). Presenter: Stephen Blair-Chappell, Bayncore
14:15	**15:00**	**TUTORIAL: APPLYING LIBRARIES AND TOOLS TO REAL WORLD EXAMPLE** In this tutorial we will demonstrate capabilities of Intel SW for development of the recommendation system, the typical example of the application that requires distributed processing of the data. The tutorial will show how to use Python API of Intel® Data Analytics Acceleration Library, its capabilities together with Intel® MPI communication technology to train the respective model and produce recommendation for specific users. The tutorial will be completed with the analysis of the performance and scalability of the Intel SW based application. Presenter: Andrey Nikolaev, Intel
15:00	**15:30**	Coffee Break
15:30	**16:30**	**PREPARING FOR THE LATEST GENERATION OF INTEL ARCHITECTURE BEFORE HAVING ACCESS TO HARDWARE** In this session we'll show how to prepare for the latest generation of Intel architecture before you have access to the real hardware. We will give practical examples of how to prepare for the Intel® Xeon Phi™ Knights Landing to make sure that your code is 'KNL ready'. Presenter: Stephen Blair-Chappell, Bayncore
16:30	**17:00**	**CASE STUDY - IMPORTANCE OF CURRENT SOFTWARE TOOL CHAINS AND LIBRARIES** This session will show the impact on performance brought by up-to-date compilers and high performance libraries. Multiple examples of common software and custom code will be used to demonstrate how speed-ups can be achieved on various target platforms by using current software. Presenter: Janko Strassburg, Bayncore
17:00	**17:30**	**Q&A**

In partnership with:

LIFEBOAT DISTRIBUTION

hoCOMPUTER www.hocomputer.de

intel Software Elite Reseller

Consulting Partner:

BAYNCORE

www.bayncore.com