



INTEL® CODE MODERNIZATION WORKSHOP 2016

Version 1.3 – 27th September

Inspiration for Developers - New Agenda!

2-DAYS WORKSHOP IN PADUA

18 & 19 October 2016 at Best Western Premier Hotel Galileo
Via Venezia 30, 35131 Padua



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

DAY 1

AGENDA

TUESDAY 18.10.2016

08:30 10:00 Registration and Breakfast

10:00 10:15 Welcome – Day 1

THREE QUESTIONS EVERYONE KEEPS ASKING

10:15 10:45 How can applications be prepared to run efficient on parallel target architectures. This explorative session will help with common issues and introduce optimization approaches.
Presenter: Stephen Blair-Chappell, Bayncore

DELIVERING LEADING PERFORMANCE IN COMPUTING, ANALYTICS AND MACHINE LEARNING – ALL ABOUT INTEL® XEON™ & XEON PHI™

10:45 11:15 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).
Presenter: Janko Strassburg, Bayncore

11:15 11:45 Coffee Break

CODING HIGH PERFORMANCE PYTHON* FOR DATA-INTENSIVE ALGORITHMS

11:45 12:30 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: Stephen Blair-Chappell, Bayncore



PROFILING AND TUNING PYTHON* CODE FOR MAXIMUM PERFORMANCE

12:30 13:15 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



13:15 14:15 Lunch

CASE STUDY – Solving the N-Body Problem on INTEL® XEON PHI™ (KNL)

14:15 15:00 Porting and optimizing of an n-body algorithm to the newest generation of Intel® Xeon Phi™ processors. Tuning techniques such as scalar optimizations, vectorization with structures of arrays and memory optimizations will be explained and their effect demonstrated.
Presenter: Roger Philp, Bayncore

CASE STUDY – STENCIL CODE OPTIMIZATION ON INTEL® XEON PHI™ (KNL)

15:00 15:45 We will demonstrate an optimization process of a 3D stencil code representing the diffusion of a solute in a solvent. The considerable impact on performance of the high bandwidth memory on-board the Intel® Xeon Phi™ processor will be shown. Results are analysed and validated using Intel® VTune Amplifier.
Presenter: Janko Strassburg, Bayncore

15:45 16:15 Coffee Break

ENABLING MACHINE LEARNING ON INTEL ARCHITECTURE

16:15 17:00 In this session we'll cover the overall technical ecosystem for machine learning (ML) / deep learning (DL) and how it is enabled by Intel software solutions. From popular open frameworks like Caffe or Theano, newly optimized for Intel, to advanced math libraries like Intel® Data Analytics Acceleration Library (Intel® DAAL) or Intel® Math Kernel Library (Intel® MKL) to cutting-edge SDKs and Intel TAP, we'll demonstrate how to enable the highest performance for ML/DL code on Intel's highly parallel CPUs.
Presenter: Roger Philp, Bayncore

SCALING DEEP LEARNING ON INTEL ARCHITECTURE AND SCALABLE SYSTEMS FRAMEWORK

17:00 17:45 In this talk, we will cover the elements of deep learning training and inference which map to Intel's multi-core and many-core platforms and present recent results which show the order of magnitude of improvement in performance and TCO.
Presenter: Janko Strassburg, Bayncore

17:45 18:00 Q&A

18:00 20:00 Networking with drinks & finger food



INTEL® CODE MODERNIZATION WORKSHOP 2016

DAY 2		AGENDA	WEDNESDAY 19.10.2016
08:30	09:30	Breakfast	
09:30	09:45	Welcome – Day 2	
09:45	10:45	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	
10:45	11:15	Coffee Break	
11:15	12:00	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	
12:00	12:30	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	
12:30	13:30	Lunch Break	
13:30	14:15	WHATS NEW IN INTEL PARALLEL STUDIO 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: Janko Strassburg, Bayncore	
14:15	15:00	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	
15:00	15:30	Coffee Break	
15:30	16:15	OPTIMIZE AND PERFORM WITH INTEL MPI Details of the new Intel® MPI library are provided, with special attention given to recent MPI 3.0 features. Profiling tools included in Intel® Parallel Studio XE 2017 for maximizing performance of HPC cluster applications are also described. Presenter: Roger Philp, Bayncore	
16:15	17:00	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	
17:00	17:30	Q&A	

In partnership with:



Consulting Partner:



Copyright © 2016 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.