INTEL® CODE MODERNIZATION WORKSHOP 2017

Version 1.1

Inspiration for Developers - New Agenda!

2-DAYS WORKSHOP IN BERLIN

Hotel TRYP Berlin Mitte, Chausseestr. 33, 10115 Berlin

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

DAY 1 AGENDA WEDNESDAY 18.01.2017

08:30	09:30	**Registration and breakfast**
09:30	**09:45**	**Welcome – Day 1**
09:45	**10:15**	**THREE QUESTIONS EVEYONE KEEPS ASKING** 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
10:15	**11:00**	**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). Presenter: Stephen Blair-Chappell, Bayncore
11:00	11:30	**Coffee Break**
11:30	**12:15**	**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: Roger Philp, Bayncore
12:15	**13:00**	**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. Presenter: Roger Philp, Bayncore
13:00	14:00	**Lunch**
14:00	**14:45**	**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: Stephen Blair-Chappell, Bayncore
14:45	**15:30**	**TUTORIAL - 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: Presenter: Stephen Blair-Chappell, Bayncore
15:30	16:00	**Coffee Break**
16:00	**17:00**	**ENABLING MACHINE LEARNING ON INTEL ARCHITECTURE** 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
17:00	17:15	**Q&A**
17:30	**20:00**	**Networking with drinks & finger food**

INTEL® CODE MODERNIZATION WORKSHOP 2017

DAY 2 AGENDA THURSDAY 19.01.2017

08:30	09:30	**Breakfast**
09:30	**09:45**	**Welcome – Day 2**
09:45	**10:30**	**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
10:30	**11:15**	**CASE STUDY – Solving the N-Body Problem on INTEL® XEON PHI™ (KNL)** 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: Francois Fayard, Bayncore
11:15	11:45	**Coffee Break**
11:45	**12:30**	**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
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: Francois Fayard, Bayncore
14:15	**15:00**	**DELVING INTO THE MYSTERIES OF CPU DISPATCH AND HOW TO HARNESS IT TO YOUR ADVANTAGE** Most programmers are aware that the Intel compiler can generate multiple optimised code paths, with the 'best' path being chosen dynamically at runtime - a technique called CPU dispatch. In this session we show a number of examples where the more esoteric aspects of CPU Dispatch are used to/ bring both portability and improved performance. We also give examples of how manual CPU dispatch (as opposed to automatic CPU dispatch) can be used to great effect when developing low-level CPU-specific APIs or wrapper functions. Presenter: Stephen Blair-Chappell, Bayncore
15:00	15:30	**Coffee Break**
15:30	**16:15**	**CODE ACCELERATION THROUGH OFFLOADING FOR INTEL® ARCHITECTURE** In this session we look at the current trends in computational offloading and show how practitioners are using such techniques on first (KNC) and second generation (KNL) Intel® Xeon Phi™ Processors and Intel® GPUs. Presenter: Stephen Blair-Chappell, Bayncore
16:15	**17:00**	**CASE STUDY – OPTIMIZING BANDWIDTH BOUND STENCIL CODE ON INTEL® XEON PHI™ (KNL)** 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 when using the on-package high bandwidth memory (MCDRAM) the Intel® Xeon Phi™ processor will be shown and compared to system DDR memory. Results are analysed and using Intel® VTune Amplifier. Presenter: Francois Fayard, Bayncore
17:00	17:15	**Q&A**

In partnership with:

SOS SOFTWARE SERVICE
VALUE ADDED DISTRIBUTION

Consulting Partner: BAYNCORE

www.bayncore.com