INTEL® SOFTWARE DEVELOPER Workshop
ROADSHOW WINTER 2018

OSLO
Scandic Solli Hotel, Parkveien 68, 0260 Oslo, norway

AGENDA – Day 1
- CODE MODERNIZATION & Parallel Programming -

Timing		Session name / description
08:15	**09:30**	**Registration with light breakfast**
09:30	**09:45**	**WELCOME & INTRODUCTION**
09:45	**10:30**	**PARALLELISM, PERFORMANCE & OPTIMIZATION ON INTEL® ARCHITECTURE – WHAT YOU SHOULD KNOW!** Starting with a brief overview of the latest Intel® silicon roadmap we look at how you can use Intel® Parallel Studio XE 2018 to get best performance on the new Intel® Xeon® Scalable Processors. We then discuss three key topics (Vectorization with AVX512, Threading, and Memory) that you need to address when modernizing code. *Stephen Blair-Chappell, Bayncore*.
10:30	**11:15**	**PRACTICAL SESSION 1: USING INTEL® PARALLEL STUDIO TO ANSWER THE QUESTION 'WHY IS MY PROGRAM RUNNING SO SLOW?'** In this session, we use three Intel® tools, Intel® Trace Analyzer and Collector, Intel® VTune Amplifier XE, and Intel® Vectorization Advisor to track down the reasons for slow running code in a Lattice Quantum Chromodynamics (LQCD) code. The example is based on a real problem reported by the HPC community. *Stephen Blair-Chappell, Bayncore*.
11:15	**11:45**	**Coffee break**
11:45	**12:45**	**STRIDING TOWARDS PERFECTION- A STEP-BY-STEP NARRATIVE ON OPTIMIZING THE K-MEANS ALGORITHM** A look at how code modernization techniques are being used in the scientific community to produce code that takes best advantage of the latest generation of CPU hardware. In this session we improve the performance of the k-mean clustering algorithm written in C++ by first working on the vectorization followed by improving the threading of the code. The final version is benchmarked on latest generation of Intel® Xeon®.. *Martin Golasowski, Bayncore*
12:45	**13:45**	**Lunch break**
13:45	**14:30**	**PRACTICAL SESSION 2: TUNING VECTORIZED CODE USING INTEL® VECTOR ADVISOR** In this session, we show how to use Intel® Vector Advisor to check how well your code is being vectorized and using the latest architecture available such as AVX512. Additionally, we look at various memory issues, such as non-contiguous memory accesses (unit stride vs. non-unit stride accesses), and how eliminating such issues can lead to significant speed up of vectorized code and improve the quality of code generated automatically by the compiler. *Roger Philp, Bayncore*
14:30	**15:15**	**USING THE INTEL® COMPILER TO CREATE FAST PORTABLE APPLICATIONS** In this session we take a close look at how you can use the Intel® compiler to bring performance and portability to your vectorized applications. We show how you can take full advantage of the latest instructions sets – such as AVX512 – and yet create programs that can still safely run on earlier generations of CPU. Additionally, we describe some of the recent compiler options supported by the latest version the Intel® compiler that improve the reproducibility of floating point results. *Martin Golasowski, Bayncore*
15:15	**15:45**	**Coffee break**
15:45	**16:30**	**OPTIMIZING PYTHON CODE USING THE INTEL® DISTRIBUTION OF PYTHON*** It used to be the case that you would never use the words 'performance' and 'python in the same sentence. The Intel® distribution of Python* changes all that. In this first of a two-part session we show how you can speed up you Python codes using by 'Cythonising' your Python code to achieve native performance. *Stephen Blair-Chappell, Bayncore*.
16:30	**17:15**	**PRACTICAL SESSION 3: ROOFLINE ANALYSIS USING INTEL® VECTOR ADVISOR** Learn how to run a Roofline Analysis using Intel® Vector Advisor. The Roofline model combines locality, bandwidth, and different parallelization paradigms into a single performance figure that shows the performance of the code under test. *Roger Philp, Bayncore*
17:15	**17:30**	**Q&A**
17:30	**19:00**	**NETWORKING COCKTAIL WITH DRINKS AND FINGER FOOD**

INTEL® AI ACADEMY
ROADSHOW WINTER 2018

OSLO
Scandic Solli Hotel, Parkveien 68, 0260 Oslo, norway

AGENDA – Day 2
- Artificial Intelligence & Deep Learning -

Timing		Session name / description
08:15	**09:30**	**Registration with light breakfast**
09:30	**09:45**	**WELCOME & INTRODUCTION**
09:45	**10:30**	**ACCELERATING AI FROM THE CLOUD TO THE EDGE** This session will cover Intel's vision for Artificial Intelligence and introduce the latest Intel® portfolio of Hardware, Software and Services from a software development and AI perspective. Besides the architectural details of the latest Intel® Xeon® Scalable processor family, we will also cover the whole spectrum of hardware solutions up to the recently announced Intel® Nervana™ Neural Network Processor (NNP). *Roger Philp, Bayncore*
10:30	**11:15**	**AI CONCEPTS AND USE CASES** In this session, we will explore the concepts and applications of Deep Learning, with a focus on real world applications using the Intel® CPUs for training and inference. *Mustafa Aldemir, Intel*
11:15	**11:45**	**Coffee break**
11:45	**12:15**	**INTEL® NERVANA™ SOFTWARE STACK – OVERVIEW & IMPLEMENTATION** This session will cover Intel® Nervana™'s software stack for AI, Machine Learning and Deep Learning: from low-level libraries like MKL / MKL-DNN, CPU-optimized frameworks (incl. neon, Caffe, TensorFlow, Theano), development tools like VTune, the Intel® Python* distribution, to the new Intel® Nervana™ Graph library (ngraph) *Mustafa Aldemir, Intel*
12:15	**13:00**	**OPTIMIZING PYTHON CODE USING THE INTEL® DISTRIBUTION OF PYTHON*** It used to be the case that you would never use the words 'performance' and 'python in the same sentence. The Intel® Distribution of Python* changes all that. In this second of a two-parts' session we show how you can speed up you Python codes 'out-of-the-box' by using the Intel® Distribution of Python*. In this session we use the Intel® optimized version of SciKit-Learn. *Martin Golasowski, Bayncore*
13:00	**14:00**	**Lunch break**
14:00	**14:30**	**MEET THE ENGINEERS**
14:30	**15:15**	**PRACTICAL FRAMEWORKS SESSION 1: USING TENSORFLOW*** In this tutorial we show how to use the Intel®-optimized version of TensorFlow* hosted on the high-level neural networks library Keras. As well as demonstrating of how to use these frameworks, the session will include an explanation of how the Intel® implemented optimizations were achieved. *Martin Golasowski, Bayncore*
15:15	**16:00**	**PRACTICAL FRAMEWORKS SESSION 2: DEEP LEARNING WITH INTEL BIGDL IN APACHE SPARK** Speed up your Spark-based machine learning apps by using Intel BigDL. In this session we provide an overview of the BigDL library and a short demonstration of its usage within the Apache Spark ecosystem. *Martin Golasowski, Bayncore*
16:00	**16:30**	**Coffee break**
16:30	**17:30**	**PRACTICAL FRAMEWORKS SESSION 3: END-TO-END DEVELOPMENT USING INTEL® OPTIMIZED CAFFE AND THE MOVIDIUS™ NEURAL COMPUTE STICK** In this session we discuss the advantages of using Caffe optimized for Intel® Architecture and show how to train deep network models using one or more compute nodes. We then show how these pre-trained models can be deployed on the Movidius™ Neural Compute Stick. *Roger Philp, Bayncore; Vishnu Madhu, Intel*
17:45	**18:00**	**Q&A and closing comments**