INTEL® SOFTWARE DEVELOPER WORKSHOP
ROADSHOW FALL / WINTER 2017

ZURICH
November 15 & 16 at FIFA World, Seestrasse 27, CH-8002 Zürich

AGENDA DAY 1 – WEDNESDAY NOV. 15
- CODE MODERNIZATION & PARALLEL PROGRAMMING -

Timing		Session name / description / presenter
08:15	**09:15**	Registration with light breakfast
09:15	**09:30**	**WELCOME & INTRODUCTION**
09:30	**10:15**	**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 both the new Intel® Xeon® Scalable Processors (Purley / Skylake-SP) as well as the Intel® Xeon Phi™ processor family (Knights Landing and Knights Mill). We then discuss three key topics (Vectorization with AVX512, Threading, and Memory) that you need to address when modernizing code.
10:15	**11:15**	**HITTING THE BOUNDARY – SIX CASE STUDIES** In this session, we look at six case studies showing the code modernization efforts needed to get best performance on Intel® architecture – Intel® Xeon™ and Intel® Xeon Phi™.
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® and Intel® Xeon Phi™.
12:45	**13:45**	Lunch break
13:45	**14:30**	**USING VTUNE TO ANSWER THE QUESTION 'WHY IS MY PROGRAM RUNNING SO SLOW?'** In this session, we use Intel® VTune™ Amplifier XE 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.
14:30	**15:15**	**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 and 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.
15:30	**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 session we show how you can speed up you Python codes using Intel®'s distribution.
16:30	**17:15**	**INTEL® CPU DISPATCH – HOW 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.
17:15	**17:30**	**Q&A**
17:30	**19:00**	**NETWORKING COCKTAIL WITH DRINKS AND FINGER FOOD**

INTEL® NERVANA™ AI ACADEMY
ROADSHOW FALL / WINTER 2017

ZURICH
November 15 & 16 at FIFA World, Seestrasse 27, CH-8002 Zürich

AGENDA DAY 2 – THURSDAY NOV. 16
- ARTIFICIAL INTELLIGENCE & DEEP LEARNING -

Timing		Session name / description / presenter
08:15	**09:15**	Registration with light breakfast
09:15	**09:30**	**WELCOME & INTRODUCTION**
09:30	**10:00**	**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.
10:00	**10:45**	**INTRODUCING THE NEW INTEL® CPU GENERATION FOR AI** This session will introduce the architectural details and the key features of the latest Intel® server CPUs from a software development and AI perspective. We will cover both the new Intel® Xeon® Scalable Processors (Purley / Skylake-SP) as well as the Intel® Xeon Phi™ processor family (code name Knights Landing and Knights Mill).
10:45	**11:15**	Coffee break
11:15	**11:45**	**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).
11:45	**13:00**	**PRACTICAL FRAMEWORKS SESSION 1: USING OPTIMIZED 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 a 'live' VTune analysis of the frameworks and an explanation of how the Intel® implemented optimizations were achieved.
13:00	**14:00**	Lunch break
14:00	**14:30**	**CASE STUDY 1: USING AI ON INTEL® ARCHITECTURE** A look at how Deep Learning on Intel® architecture is being used in image analysis in the medical and retail fields
14:30	**15:30**	**PRACTICAL FRAMEWORKS SESSION 2: USING OPTIMIZED CAFFE FRAMEWORK** In this session we show how to build Caffe optimized for Intel® architecture, train deep network models using one or more compute nodes, and deploy networks. In addition, various functionalities of Caffe are explored in detail including how to fine-tune, extract and view features of different models, and use the Caffe Python API.
15:30	**16:00**	Coffee break
16:00	**16:30**	**CASE STUDY 2: MANUFACTURING PACKAGE FAULT DETECTION USING DEEP LEARNING** A proof of concept focused on adopting deep-learning technology based on Caffe* for manufacturing package fault detection.
16:30	**17:00**	**PRACTICAL FRAMEWORKS SESSION 3: TIPS AND TRICKS FOR BEST PERFORMANCE** A series of 'Best Known Methods' and practical tips that will help get best performance when using Intel® architecture.
17:00	**17:45**	**DL INFERENCE USING FPGA & THE MOVIDIUS™ NEURAL COMPUTE STICK** Learn how trained models can be optimized for inference using Intel®'s FPGA and also using the innovative Movidius™ Neural Compute Stick.
17:45	**18:00**	**Q&A**