



MADRID, SPAIN

NOVEMBER 7TH & 8TH AT WANDA METROPOLITANO, ATLÉTICO MADRID STADIUM

AGENDA - DAY 1 - NOV 7TH

- CODE MODERNIZATION & PARALLEL PROGRAMMING -

Timing	Session name / description
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. <i>Stephen Blair-Chappell, Bayncore</i>
10:15 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. <i>Stephen Blair-Chappell, Bayncore</i>
11:15 11:45	Coffee break
11:45 12:30	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™. <i>Georg Zitzlsberger, Bayncore</i>
12:30 13:30	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 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. <i>Roger Philp, Bayncore</i>
13:30 14:30	Lunch break
14:30 15:30	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. <i>Georg Zitzlsberger, Bayncore</i>
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 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. <i>Stephen Blair-Chappell, Bayncore</i>
16:30 17:30	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. <i>Roger Philp, Bayncore</i>
17:30 19:00	NETWORKING COCKTAIL WITH DRINKS AND FINGER FOOD

MADRID, SPAIN
NOVEMBER 7TH & 8TH AT WANDA METROPOLITANO, ATLÉTICO MADRID STADIUM
AGENDA - DAY 2 - NOV. 8TH
- ARTIFICIAL INTELLIGENCE & DEEP LEARNING -

Timing	Session name / description
08:15 09:15	Registration with light breakfast
09:15 09:30	WELCOME & INTRODUCTION
09:30 10:10	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). <i>Ralph de Wargny, Intel</i>
10:10 10:40	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. <i>Roger Philp, Bayncore</i>
10:40 11:15	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). <i>Georg Zitzlsberger, Bayncore</i>
11:15 11:45	Coffee break
11:45 12:45	PRACTICAL FRAMEWORKS SESSION 1: 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. <i>Stephen Blair-Chappell, Bayncore</i>
12:45 13:30	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 second of a two-parts' session we show how you can speed up your 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. <i>Georg Zitzlsberger, Bayncore</i>
13:30 14:00	CASE STUDY: 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. <i>Janko Strassburg, Bayncore</i>
14:00 15:00	Lunch break
15:00 15:30	MEET THE ENGINEERS <i>Stephen Blair-Chappell, Bayncore</i>
15:30 16:15	PRACTICAL FRAMEWORKS SESSION 2: 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 a 'live' VTune analysis of the frameworks and an explanation of how the Intel® implemented optimizations were achieved. <i>Georg Zitzlsberger, Bayncore</i>
16:15 16:30	Coffee break
16:30 17:15	DL INFERENCE USING THE MOVIDIUS™ NEURAL COMPUTE STICK Learn how trained models can be optimized for inference using innovative Movidius™ Neural Compute Stick. <i>Roger Philp, Bayncore</i>
17:15 17:30	Q&A and closing comments
17:30 18:30	Optional: guided tour of the Wanda Metropolitano, Atlético Madrid Stadium