



INTEL® Code Modernization Workshop - 2018

Advanced Topics for Developers - New Agenda!

2-DAYS WORKSHOP IN OSLO

6-7 FEBRUARY 2017

Hotel Scandic Solli, Parkveien 68, 0260 Oslo, Norway

Day 1 A G E N D A Tuesday 6/2-18

Theme of the Day: CODE MODERNIZATION & Parallel Programming

Timing	Session name / description
--------	----------------------------

08:15	09:15	Registration & Light breakfast
-------	-------	--------------------------------

PARALLELISM, PERFORMANCE & OPTIMIZATION ON INTEL ARCHITECTURE – WHAT YOU SHOULD KNOW!

09:15	10:00	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.
-------	-------	---

HITTING THE BOUNDARY – SIX CASE STUDIES

10:00	11:00	In this session, we look at six case studies (KNL & SKL) showing the code modernization efforts needed to get best performance on Intel Architecture.
-------	-------	---

11:00	11:30	Coffee break
-------	-------	--------------

STRIDING TOWARDS PERFECTION- A STEP-BY-STEP NARRATIVE ON OPTIMIZING THE K-MEANS ALGORITHM

11:30	12:30	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™.
-------	-------	--

12:30	13:30	Lunch break
-------	-------	-------------

USING VTUNE TO ANSWER THE QUESTION ‘WHY IS MY PROGRAM RUNNING SO SLOW?’

13:30	14:15	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.
-------	-------	---

TUNING VECTORIZED CODE USING INTEL® VECTOR ADVISOR

14:15	15:00	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.
-------	-------	--

15:00	15:30	Coffee break
-------	-------	--------------



INTEL® Code Modernization Workshop - 2018

15:30	16:15	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 your Python codes using Intel®'s distribution.
16:15	17:00	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.
17:00	17:15	Open Forum with all presenters

Day 2 A G E N D A Wednesday 7/2-18

Theme of the Day: Artificial Intelligence & Deep Learning

08:15	09:15	Light breakfast
		AI CONCEPTS AND USE CASES
09:15	09:45	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.
		INTRODUCING THE NEW INTEL® CPU GENERATION FOR AI
09:45	10:30	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).
10:30	11:00	Coffee break
		INTEL® NERVANA™ SOFTWARE STACK – OVERVIEW & IMPLEMENTATION
11:00	11:30	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).
		PRACTICAL FRAMEWORKS SESSION 1: USING TENSORFLOW WITH KERAS
11:30	13:00	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
		PRACTICAL FRAMEWORKS SESSION 2: USING OPTIMIZED CAFFE FRAMEWORK
14:30	15:30	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
		PRACTICAL FRAMEWORKS SESSION 3: TIPS AND TRICKS FOR BEST PERFORMANCE
16:00	16:30	A series of 'Best Known Methods' and practical tips that will help get best performance when using Intel® Architecture