



INTEL® CODE MODERNIZATION WORKSHOP - 2017

Advanced Topics for Developers - New Agenda!

2-DAYS WORKSHOP IN STOCKHOLM

8 & 9 MAY 2017

Coor Konferens, Tegelbacken 4A, 11152 Stockholm

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

DAY 1 AGENDA

MONDAY 8.5.2017

THEME OF THE DAY: OPTIMISATION TECHNIQUES USING INTEL PARALLEL STUDIO XE

08:00 09:00 Registration and Breakfast

09:00 09:10 Welcome – Day 1

USING THE ADVANCED FEATURES OF INTEL PARALLEL STUDIO - WHAT'S ON OFFER

09:10 09:40 A quick overview of some of the more advanced features of Intel Parallel Studio XE and what's new in the latest version.
Presenter: Stephen Blair-Chappell, Bayncore

ROOFLINE PROFILING WITH INTEL VECTOR ADVISOR

09:40 10:30 Learn how to use Intel Vector Advisor's roofline analysis to answer such questions as: Does my application work optimally on the current hardware? What is the most underutilized hardware resource? What limits performance? Is my application workload memory or compute bound? What is the right strategy to improve application performance?
Presenter: Francois Fayard, Bayncore

10:30 11:00 Coffee Break

ANALYSING MEMORY ACCESS PATTERNS (MAP) WITH INTEL VECTOR ADVISOR- AND THE IMPACT ON APPLICATION PERFORMANCE

11:00 12:15 In this session, we show how to use Intel Vector Advisor to check for 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 vectorised code and improve the quality of code generated automatically by the compiler.
Presenter: Roger Philp, Bayncore

12:15 13:15 Lunch

CASE STUDY

13:15 13:45 An example of how advanced optimisation techniques have been used in a real-world application.
Presenter: Francois Fayard, Bayncore

ARCHITECTURAL ANALYSIS USING VTUNE

13:45 15:30 In this session, we take a 'deep dive' into hardware *event based analysis* using Intel VTune Amplifier XE. We present a classic 'top down' architectural analysis methodology and show how such techniques can become your method of choice when carrying out both course-grained and fine-grained code optimization.
Presenter: Stephen Blair-Chappell, Bayncore

15:30 16:00 Coffee Break

ADVANCED CUSTOM ANALYSIS AND CONFIGURATION USING VTUNE

16:00 17:15 Using Intel VTune Amplifier XE 'out of the box' produces excellent results, however, there are occasions when by using your own custom configuration you can directly influence quality of the analysis. In this session, we show how to create and use your own custom analysis types; choose what hardware events to use; modify the collection parameters; and add API calls to your source code - leading to a better analysis and help track down those difficult-to-find bottlenecks in your code.
Presenter: Stephen Blair-Chappell, Bayncore

17:15 17:45 Q&A

18:00 20:00 Networking with drinks & finger food



INTEL® CODE MODERNIZATION WORKSHOP - 2017

DAY 2 AGENDA

TUESDAY 9.5.2017

THEME OF THE DAY: PRACTICAL MACHINE LEARNING AND DEEP LEARNING ON INTEL PLATFORMS

08:00 09:00 Breakfast

09:00 09:30 ACCELERATING MACHINE LEARNING - WHAT INTEL BRINGS TO THE TABLE

A quick overview of the latest tools and technologies available from Intel in Artificial Intelligence.
Presenter: Roger Philp, Bayncore

09:30 10:15 HOW TO SUPER-CHARGE YOUR PYTHON CODES USING INTEL OPTIMISED PYTHON

See how the latest Python distribution from Intel brings a significant performance boost in AI and Deep Learning/Machine Learning codes.
Presenter: Stephen Blair-Chappell, Bayncore

10:15 10:45 Coffee Break

10:45 12:00 BENCHMARKING AND ANALYSIS OF KEY MACHINE LEARNING PATTERNS AND KERNELS

In this session, we demonstrate the analysis of some classic ML patterns (e.g. SVM, K-Means and CNN) and computational kernels (e.g. GEMM, SVD PCA, etc.) , and show how they can be optimized to produce superior performance.
Presenter: Francois Fayard, Bayncore

12:00 13:00 Lunch Break

13:00 13:30 CASE STUDY

An example of an AI implementation using the Intel optimized frameworks and Intel tools.
Presenter: Roger Philp, Bayncore

13:30 14:30 PRACTICAL FRAMEWORKS SESSION 1: APPLICATION DEVELOPMENT USING KERAS, TENSORFLOW AND THEANO

In this tutorial we show how to use the Intel-optimized versions *TensorFlow* and *Theano* 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.
Presenter: Stephen Blair-Chappell, Bayncore

14:30 15:00 Coffee Break

15:00 16:15 PRACTICAL FRAMEWORKS SESSION 2: USING CAFFE ON INTEL ARCHITECTURE

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.
Presenter: Francois Fayard, Bayncore

16:15 16:30 Q&A

In partnership with:



Copyright © 2017 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.