



INTEL® CODE MODERNIZATION WORKSHOP - 2017

Advanced Topics for Developers - New Agenda!

2-DAYS WORKSHOP IN DUBLIN

12th & 13th June 2017

Spencer Hotel Dublin City, Excise Walk, IFSC, Dublin 1

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

DAY 1 AGENDA

MONDAY 12.6.2017

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

08:30 09:00 Registration and breakfast

09:00 09:10 **Welcome - Day 1**

09:10 09:40 **USING THE ADVANCED FEATURES OF INTEL PARALLEL STUDIO - WHAT'S ON OFFER**
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

09:40 10:30 **CASE STUDY**
An example of how advanced optimisation techniques along with the Intel Parallel Studio XE have been used in a real-world application.
Presenter: Francois Fayard, Bayncore

10:30 11:00 Coffee break

11:00 11:30 **SOMETHING OLD, SOMETHING NEW, SOMETHING BORROWED, SOMETHING BLUE**
A look at four important features that make Intel x86 architecture 'best in class'
Presenter: Stephen Blair-Chappell, Bayncore

11:30 12:15 **ARCHITECTURAL ANALYSIS USING VTUNE**
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

12:15 13:15 Lunch

13:15 14:15 **ARCHITECTURAL ANALYSIS USING VTUNE (Continued)**

14:15 14:45 **ROOFLINE PROFILING WITH INTEL VECTOR ADVISOR**
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

14:45 15:30 **USING THE INTEL VECTOR ADVISOR**
In this session, we show how to use Intel Vector Advisor to check how well your code is being vectorised. 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 vectorised code and improve the quality of code generated automatically by the compiler.
Presenter: Roger Philp, Bayncore

15:30 16:00 Coffee break

16:00 16:45 **ADVANCED CUSTOM ANALYSIS AND CONFIGURATION USING VTUNE**
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

16:40 17:00 **Q&A**

17:00 18:30 **Networking with drinks & finger food**



INTEL® CODE MODERNIZATION WORKSHOP FOR AI - 2017

DAY 2 AGENDA

TUESDAY 13.6.2017

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

08:30 09:00 Breakfast

CPU DISPATCH

09:00 09:45 In-depth exploration of the CPU Dispatching mechanism provided by the Intel Compilers.
Presenter: Stephen Blair-Chappell, Bayncore

HOW TO SUPER-CHARGE YOUR PYTHON CODES USING INTEL OPTIMISED PYTHON

09:45 10:30 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

10:30 11:00 Coffee break

BENCHMARKING AND ANALYSIS OF KEY MACHINE LEARNING PATTERNS AND KERNELS

11:00 12:15 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

12:15 13:15 Lunch break

PRACTICAL FRAMEWORKS SESSION 1: USING CAFFE ON INTEL ARCHITECTURE

13:15 14:15 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

PRACTICAL FRAMEWORKS SESSION 2: APPLICATION DEVELOPMENT USING KERAS, TENSORFLOW AND THEANO

14:15 15:00 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

15:00 15:30 Coffee break

ACCELERATING MACHINE LEARNING & CASE STUDY

15:30 16:45 A quick overview of the latest tools and technologies available from Intel in Artificial Intelligence and an associated Case Study.
Presenter: Roger Philp, Bayncore

16:45 17:00 Q&A

In partnership with:

grey matter®

BAYNCORE

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.