

## Intel<sup>®</sup> Math Kernel Library. Intel<sup>®</sup> Integrated Performance Primitives, the latest Updates, upcoming features

May 2014 Gennady Fedorov

## Outline

- Intel<sup>®</sup> Cluster Studio XE
- Intel<sup>®</sup> MKL : Overview
- MKL 11.2 beta new Features
- Intel<sup>®</sup> IPP, Overview
- IPP <-> MKL
- IPP, MKL : Beyond current versions
- References

#### **Intel® Cluster Studio XE 2013 SP1** Tools to Scale Forward, Scale Faster – for HPC Clusters

Phase	Product	Feature	Benefit
	Intel® MPI Library	High Performance Message Passing (MPI) Library	<ul> <li>Enabling High Performance Scalability, Interconnect Independence, Runtime Fabric Selection, and Application Tuning Capability</li> </ul>
Build	Intel® Composer XE SP1	<ul> <li>C/C++ and Fortran compilers and performance libraries</li> <li>Intel® Threading Building Blocks</li> <li>Intel® Cilk™ Plus</li> <li>Intel® Integrated Performance Primitives</li> <li>Intel® Math Kernel Library</li> </ul>	<ul> <li>Enabling solution to achieve the application performance and scalability benefits of multicore and forward scale to many-core</li> </ul>
Verify	Intel® Inspector XE	Memory & threading dynamic analysis for code quality Static Security Analysis for code quality	<ul> <li>Increased productivity, code quality, and lowers cost, finds memory, threading, and security defects before they happen</li> <li>Now MPI enabled at every cluster node</li> </ul>
Verify & Tune	Intel® Trace Analyzer & Collector	MPI Performance Profiler for understanding application correctness & behavior	<ul> <li>Analyze performance of MPI programs and visualize parallel application behavior and communications patterns to identify hotspots</li> </ul>
Tune	Intel® VTune™ Amplifier XE	Performance Profiler for optimizing application performance and scalability	<ul> <li>Remove guesswork, saves time, makes it easier to find performance and scalability bottlenecks</li> <li>Now MPI enabled at every cluster node</li> </ul>



## Outline

- Intel<sup>®</sup> Cluster Studio XE
- Intel<sup>®</sup> MKL : Overview
- MKL 11.2 beta new Features
- Intel<sup>®</sup> IPP, Overview
- MKL : Beyond current versions
- References

## Intel<sup>®</sup> MKL Overview

#### Intel<sup>®</sup> Math Kernel Library

- Highly optimized threaded math routines
- Applications in science, engineering, finance
- Use Intel<sup>®</sup> MKL on Windows\*, Linux\*, Mac OS\*
- Use Intel<sup>®</sup> MKL with Intel compiler, gcc, MSFT\*, PGI
- Component of Intel<sup>®</sup> Parallel Studio XE and Intel<sup>®</sup> Composer XE



Editors' Ch Best HPC Pe T Paralle	View online noice Awards 2011 rformance Product or echnoogy Intel I Studio XE 2011		
Evans Data Corporation EDC	EDC North America Development Survey 2011, Volume II		
33% of math Intel's Ma	libraries users rely on ath Kernel Library		



## Intel<sup>®</sup> MKL Domains & Functions



**Optimization Notice** 

Copyright © 2014, Intel Corporation. All rights reserved. \*Other names and brands may be claimed as the property of others.

(intel)







#### Usage Models on Intel<sup>®</sup> Xeon Phi<sup>™</sup> Coprocessors

- Automatic Offload

- Compiler Assisted Offload

- Native Execution





5/30/2014

#### **Usage models Continued ..**

- Automatic Offload:
  - No code changes required
  - Automatically uses both host and target
  - Transparent data transfer and execution management
  - EX: Export MKL\_MIC\_ENABLE=1

Article for the List of AO Enabled Functions: <u>http://software.intel.com/en-us/articles/intel-mkl-automatic-offload-enabled-functions-for-intel-xeon-phi-coprocessors</u>





#### **How to Use Automatic Offload**

• Using Automatic Offload is easy



- What if there doesn't exist a MIC card in the system?
  - Runs on the host as usual without any penalty!
- Reporting of Automatic Offload profiling:
  - int mkl\_mic\_set\_offload\_report (int enabled);
    - off by the Default
    - **on** if "enabled" == 1 or 2
  - Env Variable: OFFLOAD\_REPORT



# **Work Division Control in Automatic Offload**

Using support functions

Examples	Notes
MKL_MIC_Set_Workdivision( MKL_TARGET_MIC, 0, 0.5)	Offload 50% of computation only to the $1^{st}$ card.

• Using environment variables

Examples	Notes
MKL_MIC0_WORKDIVISION=50	Offload 50% of computation only to the $1^{st}$ card.

 The support functions take precedence over environment variables





#### **Tips for Using Automatic Offload**

- AO works only when matrix sizes are right
  - **?**GEMM: Offloading only when M, N > 2048, K > 256
  - **?**TRSM/TRMM: Offloading only when M, N > 3072
  - **?**GETRF: M,N > 8192
  - **?**POTRF:N>=6144
  - **?**GEQRF:M=N>=8192
  - Square matrices may give better performance
- Work division settings are just hints to MKL runtime
- How to disable AO after it is enabled?
  - mkl\_mic\_set\_workdivision(MIC\_TARGET\_HOST, 0, 1.0)
  - MKL\_HOST\_WORKDIVISION=100



#### **Tips for Using Automatic Offload**

- Threading control tips:
  - Avoid using the OS core of the target, which is handling data transfer and housekeeping tasks. Example (on a 60core card):

MIC\_KMP\_AFFINITY=explicit,granularity=fine,proclist=[1-236:1]

- Also, prevent thread migration on the host: KMP\_AFFINITY=granularity=fine,compact,1,0
- Note: Different env-variables on host and MIC:

OMP\_NUM\_THREADSMIC\_OMP\_NUM\_THREADSKMP\_AFFINITYMIC\_KMP\_AFFINITY





# Intel MKL : Conditional Numerical Reproducibility (CNR) - motivation

Ever seen something like this?

C:\Users\me>test.exe 4.012345678901111

C:\Users\me>test.exe 4.012345678902222

C:\Users\me>test.exe 4.012345678902222

C:\Users\me>test.exe 4.012345678901111

C:\Users\me>test.exe 4.012345678902222

#### ...or this?





Intel® Xeon® Processor E3-1275

14

## Intel MKL : CNR - Why do results vary?

#### Root cause for variations in results

- double precision arithmetic example (a+b)+c ≠ a+(b+c)

2 <sup>-63</sup> +	1	+	-1 = 2 <sup>-63</sup>	(infinitely precise result)
(2 <sup>-63</sup> +	1)	+	-1 = 0	(correct IEEE single precision result)
2 <sup>-63</sup> + (	1	+	-1) = 2 <sup>-63</sup>	(correct IEEE single precision result)

#### Why does the order of operations change in Intel MKL?

- Memory alignment affects grouping of data in registers
- Number of threads most functions are threaded to use all available cores
- Dynamic task scheduling some algorithms use asynchronous task

#### Order matters when doing floating point arithmetic. Many optimizations require a change in order of operations.

## Conditional Numerical Reproducibility

#### Challenges Means of control

	Memory alignment	<ul> <li>Align memory — try Intel MKL memory allocation functions</li> <li>64-bytes alignment for processors in the next few years</li> </ul>
	Number of threads	<ul> <li>Set the number of threads to a constant number</li> <li>Use sequential libraries</li> </ul>
ĬM	Dynamic task scheduling	<ul> <li>New control to specify static scheduling</li> <li>Previous sequential operation was required</li> </ul>
S	Code branches	• New controls to maintain consistent code branches across different processors (mkl_cbwr_set (MKL_CBWR_AVX))

## Intel MKL 11.0 introduces functions to broaden the cases where reproducibility is possible.

16

## Why "Conditional"?

- Reproducibility is not guaranteed across...
  - operating systems
  - architectures



- Reproducibility across non-Intel processors is limited
- Reproducibility will not cover all future Intel processors
- Reproducibility will not be maintained from one version of Intel MKL to the next
- Reproducibility controls in Intel MKL only affect Intel MKL functions; other tools, and other parts of your code can still cause variation

#### New functionality expands the scope for reproducibility, but not to every case.

## Extended Eigensolver Routines(FEAST)

- Solution of standard generalized EV problems (Feast 2.0)
  - standard,  $Ax = \lambda x$ , where are
    - for A complex Hermitian and A real symmetric
  - generalized,  $Ax = \lambda Bx$ 
    - where A complex Hermitian, B Hermitian positive definite (hpd)
    - A real symmetric and B real symmetric positive definite (spd)
- Real/Complex and Single/Double precisions
- Compatible with FORTRAN 77, Fortran 90, and C
- matrix storage: Full, Band storages, Sparse format (SCR)
- Reverse communication interfaces (RCI)
- Predefined driver interfaces

## MKL 11.1 Highlights

- Intel<sup>®</sup> Xeon Phi<sup>™</sup> support from Windows\* host
- Better installation experience
  - A choice of components to install
  - Examples and tests packaged as archives
- HPL support for heterogeneous clusters
- Performance improvements
- CNR support for unaligned input data
- CNR support for arbitrary #threads delayed

## Outline

- Intel<sup>®</sup> Cluster Studio XE
- Intel<sup>®</sup> MKL : Overview
- MKL 11.2 beta new Features
- Intel<sup>®</sup> IPP, Overview
- MKL : Beyond current versions
- References



## MKL 11.2 beta - new Features

- Verbose mode
- Small Size Problem
- Cluster Pardiso
- CookBook
- Android
- Deprecations



### Verbose mode

Motivation:

MKL users : "It would be nice if MKL reports called functions names with parameters, that we won't spend additional time and resource to figure this out"

- BLAS, LAPACK, FFT
- The same standard MKL interface
- To enable verbose use:
  - environment variable MKL\_VERBOSE=1 or
  - function mkl\_verbose(1)
  - to disable mkl\_verbose(0)
- Disabled by the default

### Verbose Mode Output

**{MKL VERBOSE} thrID:**4158465744 **func:** DGEMM (T,T,5,2,4,0.5, 0xa1481c8,4,0xa148170,2,-1.2,0xa148290,5) **time:**0.2e-2s **freq:**3.2GHz **Nthr:8 dyn:**0 **CPU:**AVX **CNR:**off **AO**:0 **iface**:ilp64 **plat:**lnx **mem**:1 **mstat**:10MB

thrID: calling thread ID

func: called function name with input parameters

time: time (sec) spend in function

freq: current CPU frequency

Nthr: Number of MKL threads

**dyn**: {1|0} – MKL\_DYNAMIC

CPU: SSE2, SSE3, SSSE3, SSE4.1, SSE4.2, AVX, AVX2 for Intel & AMD for non Intel

**CNR**:OFF, AUTO, COMPATIBLE, SSE2, SSE3, SSSE3, SSE4\_1, SSE4\_2, AVX, AVX2

**AO:** {1|0},

plat: lnx, osx, win, mic,

mstat: MB used by function

iface: ilp64, lp64, cdecl, stdcall
mem: {1|0} - MKL\_FAST\_MM on/off

### Small Size Problem

- Significant overhead ( >100%) for small matrices
  - Function call overhead
  - Error checking
  - Architecture specific dispatching
  - Matrix size specific dispatching
  - Buffer allocation
- Unrolled C/Fortran codes may outperform MKL for small matrix multiplication

24

## Small Size Problem

#### To activate the feature:

• Compile your C or Fortran code with the macro and options:

Intel MKL Mode	Macro	<b>Compiler Option</b>
Threaded	MKL_INLINE	-DMKL_INLINE
Sequential	MKL_INLINE_SEQ	-DMKL_INLINE_SEQ

- Fortran application:
  - -fpp compiler option
  - Include files: mkl\_inline\_pp.fiand mkl\_inline.fi
- Link statically with MKL
- See examples into mkl\_user's guide

### Small Size Problem

#### Limitations of the Feature:

 if the MKL\_INLINE or MKL\_INLINE\_SEQ macro is used, MKL may skip error checking

#### IMPORTANT

With a limited error checking, you are responsible for checking the correctness of function parameters to avoid unsafe and incorrect code execution.

- Available for the Intel<sup>®</sup> 64 architecture only
- Only Intel<sup>®</sup> Compiler 14.0 or higher is supported
- You must link with Intel MKL statically
- Available for Intel<sup>®</sup> Xeon<sup>®</sup> processors and \*gemm functions
- Verbose mode, CNR and CBLAS interfaces are not supported

#### **Cluster Pardiso**



Cluster Pardiso. Interface/migration Supported since Intel<sup>®</sup> MKL 11.2 beta for : Linux\* & Microsoft Windows\*, 64-bit only



28

#### Cluster Pardiso. Features

- Zero one-based switch of Cluster Pardiso
- Support of distributed input data
- Correct treatment of internal errors
- MPI calls via MKL BLACS
- Single precision version both real & complex case
- LP64/ILP64 interfaces
- Matrix checker of CSR format
- Iterative refinement
- Reordering step is paralleled only by OpenMP
- Work with matrices with empty rows

#### Documentation - Cookbook

- Motivation: Requests from customers to show the "example from the life"
- Steps: Goal / Solution / Code Examples / Discussion / Resources
- MKL v. 11.2 beta 5 examples:

1) Finding an approximate solution to a nonlinear equation demonstrates a method of finding a solution to a nonlinear equation using Intel MKL PARDISO, BLAS, and Sparse BLAS routines

2) Factoring a block tridiagonal matrix uses Intel MKL implementations of BLAS and LAPACK routines

•3,4) Using Fast Fourier Transforms for computer tomography image reconstruction uses the Fast Fourier Transforms to reconstruct an image from computer tomography data

5) Noise filtering in financial market data streams uses Intel MKL summary statistics routines for computing a correlation matrix for streaming data

#### Note:

- +12 recipes would be added in the 11.2 Gold version
- Code examples in the cookbook are provided in Fortran for some recipes and in C for other recipes.

#### Cookbook - Recipe

#### Using LAPACK symmetric eigensolvers for Hermitian tridiagonal matrices

#### Soal

Compute eigenvalues and eigenvectors for a Hermitian tridiagonal matrix using LAPACK symmetric eigensolvers.

LAPACK provides symmetric eigensolvers for real-valued tridiagonal matrices, but no corresponding eigensolvers for complex Hermitian matrices.

#### Solution

A simple matrix transformation to a Hermitian tridiagonal matrix allows you to use one of the LAPACK eigensolvers.

- Multiply the Hermitian tridiagonal matrix by a matrix calculated to eliminate the imaginary parts of the offdiagonal elements, which transforms it to a real-valued matrix.
- Choose the LAPACK eigensolver routine according to the task you wish to perform and the algorithm you wish to use:
  - o sstev, dstev: Compute all eigenvalues and, optionally, eigenvectors.
  - sstevd, dstevd: Compute all eigenvalues and, optionally, eigenvectors using a divide-and-conquer algorithm.
  - o sstevx, dstevx: Compute selected eigenvalues and eigenvectors.
  - sstevr, dstevr: Compute selected eigenvalues and, optionally, eigenvectors using Relatively Robust Representations.
- Reverse the transformation to the eigenvalues and eigenvectors returned by the LAPACK routine in order to get the eigenvalues and eigenvectors of the original matrix.

#### Calculating eigenvalues and eigenvectors for a Hermitian tridiagonal matrix using DSTEV

```
PROGRAM complex_tridiagonal_hev_solver
IMPLICIT NONE
INTEGER N,INFO,I,J
PARAMETER (N=5)
C Matrix
COMPLEX*16 EC(N-1)
REAL*8 D(N)
C Eigenvectors
```

```
C Test the vectors are eigenvectors
      DO I=1,N
         DO J=1.N
              UZ(J) = -DCMPLX(D(I), 0D0) *VZ(J, I)
          END DO
          UZ(1) = UZ(1) + DC(1) * VZ(1,I) + EC(1) * VZ(2,I)
          DO J=2,N-1
            UZ(j)=UZ(j) + CONJG(EC(J-1))*VZ(J-1,I)+
     8
                            DC(J) *VZ(J,I) + EC(J) *VZ(j+1,I)
          END DO
          UZ (N) =UZ (N) + CONJG (EC (N-1)) *VZ (N-1, I) + DC (N) *VZ (N, I)
          PRINT *, "For ", I, "th eigenvalue ||S*U^j-lambda j*U^j||=",
                   DZNRM2 (N,UZ,1)
     8
      END DO
      STOP
      END
```

#### Discussion

In this example, S is a complex-valued tridiagonal Hermitian matrix:

$$S = \begin{bmatrix} d_1 & e_1 & & \\ c_1 & d_2 & e_2 & & \\ & c_2 & d_3 & e_3 & & \\ & \ddots & \ddots & \ddots & \\ & & & c_{N-2} & d_{N-1} & e_{N-1} \\ & & & & c_{N-1} & d_N \end{bmatrix}$$

Construct a diagonal matrix T, where  $|t_i| = 1$  for i = 1, 2, ... N:



A new matrix  $S_1 = T^H S T$  has the same eigenvalues as S.



## MKL for Android

- Android 32 bit
- Static linking
- Sequential version (1 thread)
- No message catalog
- Available since 11.1. update3 & 11.2 beta Update 1
- 64 bit will be available later TBD
- Dynamic Linking (Custom dll) not available TBD

## Outline

- Intel<sup>®</sup> Cluster Studio XE
- Intel<sup>®</sup> MKL : Overview
- MKL 11.2 beta new Features
- Intel<sup>®</sup> IPP, Overview
- MKL : Beyond current versions
- References

#### Intel® IPP Functions and Code Samples: Image Processing/Coding

Domain	Functions	Samples
Image Processing	<ul> <li>* Geometry transformations, such as resize/rotate</li> <li>* Linear and non-linear filtering operation on an image for edge detection, blurring, noise removal and etc for filter effect.</li> <li>* Linear transforms for 2D FFTs, DFTs, DCT.</li> <li>* image statistics and analysis</li> </ul>	* Tiled Image Processing / 2D Wavelet Transform /C++ Image Processing Classes/Image Processing functions Demo
<ul> <li>Computer</li> <li>Vision</li> <li>* Background differencing, Feature Detection (Corner Detection, Canny Edge detection), Distance Transforms, Image Gradients, Flood fill, Motion analysis and Object Tracking, Pyramids, Pattern recognition, Camera Calibration</li> </ul>		* Face Detection
Color Models	<ul> <li>* Convert image/video color space formats: RGB, HSV, YUV, YCbCr</li> <li>* Up/Down sampling</li> <li>* Brightness and contrast adjustments</li> </ul>	
JPEG Coding	<ul> <li>* High-level JPEG and JPEG2000 compression and decompression functions</li> <li>* JPEG/JPEG2000 support functions: DCT, Wavelet transforms, color conversion, downsampling</li> </ul>	* Integration with the Intel® JPEG Library (IJL) / Integration with the Independent JPEG Group (IJG) library /JPEG2000 encoder/decoder /JPEG viewer
Realistic Rendering	* Acceleration Structures, Ray-Scene Intersection and Ray Tracing * Surface properties, shader support, tone mapping	* Ray Tracing

#### Intel<sup>®</sup> IPP Functions and Code Samples: Video/Audio/Speech

Domain	Functions	Samples
Video Coding	* VC-1, H.264, MPEG-2, MPEG-4, H.261, H.263 and DV codec support functions	* Simple Media Player/ Video Encoder / h.264/DV decoding/video transcoder /Reverberation Demo/ Virtual Multi-Channel Audio Player/
Audio Coding	* Echo cancellation and audio transcoding, BlockFiltering, Spectral Data prequantization.	* Audio Codec Console application
Speech Coding	* Adaptive/Fixed Codebook functions, Autocorrelation, Convolution, Levinson-Durbin recursion, Linear Prediction Analysis & Quantization, Echo Cancellation, Companding	* G.168, G.167, G.711, G.722, G.722.1, G.722.2, AMRWB, Extended AMRWB (AMRWB+), G.723.1, G.726, G.728, G.729, RT-Audio, GSM AMR, GSM FR
Speech Recognition	* Feature Processing, Model Evaluation/Estimation/Adaptation, Vector Quantization, Polyphase Resampling, Advanced Aurora, Ephraim-Malah Noise Supression, AEC, Voice Detection	* Aurora, Advanced Aurora, Audio Processing, Gaussian Mixture, Speech Processing
Signal Processing	<ul> <li>* Transforms: DCT, DFT, MDCT, Wavelet (both Haar and user-defined filter banks), Hilbert</li> <li>* Convolution, Cross-Correlation, Auto-Correlation, Conjugate</li> <li>* Filtering: IIR/FIR/Median filtering, Single/Multi-Rate FIR LMS filters</li> <li>* Other: Windowing, Jaehne/Tone/Triangle signal generation, Thresholding</li> </ul>	* Signal Processing Function Demo

intel

# Intel® IPP Functions and Code Samples: Data Processing/Compression

Domain	Functions	Samples
Data Compression	<ul> <li>* Entropy-coding compression: Huffman, VLC</li> <li>* Dictionary-based compression: LZSS, LZ77</li> <li>* Burrows-Wheeler Transform (BWT), MoveToFront (MTF), Run-Length- Encoding (RLE), Generalized Interval Transformation (GIT)</li> <li>* Compatible feature support for zlib and bzip2</li> </ul>	* zlib, bzip2, gzip-compatible /General data compression examples
Cryptography	* Big-Number Arithmetic / Rijndael, DES, TDES, SHA1, MD5, RSA, DSA, Montgomery, prime number generation and pseudo-random number generation (PRNG) functions	* Intel IPP crypto usage in Open SSL*
String Processing	* Compare, Insert, change case, Trim, Find, Regexp, Hash	* "ippgrep" – regular expression matching
Vector Math	* Logical, Shift, Conversion, Power, Root, Exponential, Logarithmic, Trigonometric, Hyperbolic, Erf, Erfc	
Matrix Math	* Addition, Multiplication, Decomposition, Eigenvalues, Cross-product, transposition	
Common Functions	* CPUTypes, Thread number control, Memory Allocation	* Linkages/Different language support

## Outline

- Intel<sup>®</sup> Cluster Studio XE
- Intel<sup>®</sup> MKL : Overview
- MKL 11.2 beta new Features
- Intel<sup>®</sup> IPP, Overview
- MKL : Beyond current versions
- References

## MKL : Beyond 11.2 beta

- AO SVD
- AO Eigensolvers
- AO threshold mode for debugging
- AO int mkl\_mic\_set\_resource\_limit(double fraction);
- SpMV Ellpack format
- HPCG benchmarks

## MKL -- Deprecations since 11.2 beta

#### • The libraries are no longer provided:

- IA32: mkl\_p4p andmkl\_vml\_p4p
- Intel<sup>®</sup> 64 : mkl\_p4n and mkl\_vml\_p4n libraries
- Support for the following OS is Deprecated and will be removed
  - Windows XP\*
  - Windows Server 2003\*
  - Windows Vista\*
  - Red Hat\* Enterprise Linux\* 5
  - SUSE LINUX Enterprise Server\* 10
- Old MKL domain names used for threading settings will be deleted
  - MKL\_ALL, MKL\_BLAS, MKL\_FFT, MKL\_VML
  - MKL\_DOMAIN\_ALL, MKL\_DOMAIN\_BLAS, MKL\_DOMAIN\_FFT, MKL\_DOMAIN\_VML, MKL\_DOMAIN\_PARDISO

## IPP: beyond version 8.1

- IPP on Xeon Phi. Native, Assistant Offload
- Android Support
- Threading Strategy
- In-Place operations
- Deprecation Policy :

https://software.intel.com/sites/products/ipp-deprecated-features-feedback/

- UIC, UMC codec -- moving to legacy archive
- Audio and Video Codecs -- Media SDK
- IPP Preview GPU support (CV and IPPI)

### References

#### Intel<sup>®</sup> MKL page

- http://software.intel.com/en-us/articles/intel-mkl/
- http://software.intel.com/en-us/articles/intel-math-kernel-library-documentation/

#### **Intel MKL Forum:**

- <a href="http://software.intel.com/en-us/forums/intel-math-kernel-library/">http://software.intel.com/en-us/forums/intel-math-kernel-library/</a>
- Intel IPP Forum:
- <u>https://software.intel.com/en-us/forums/intel-integrated-performance-primitives/</u>
- Intel Experimental Software
  - <u>http://software.intel.com/en-us/whatif</u>



## Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Copyright © 2014, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

#### **Optimization Notice**

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

