



# INTEL<sup>®</sup> VTUNE<sup>™</sup> PROFILER

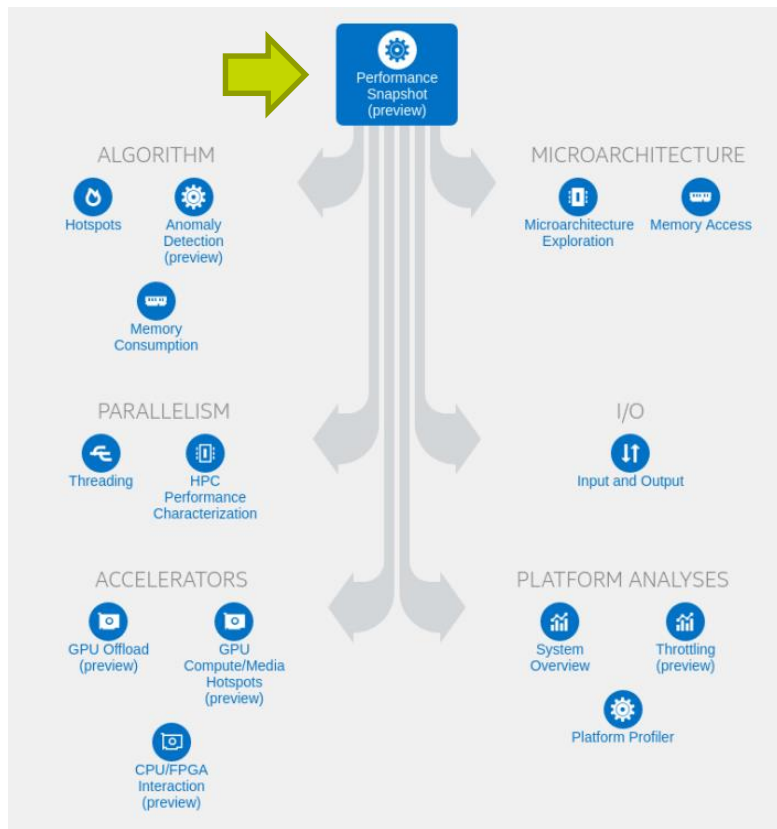
## A QUICK INTRO FOR HPC DEVELOPERS

Vladimir Tsymbol  
Rev. 09/2020

# Agenda

- Intel® VTune™ Profiler 2020 analysis types overview
- Basic analysis: Performance Snapshot and Hotspots
- HPC Performance Characterization analysis for hybrid MPI + OpenMP apps
- How to run VTune Profiler on MPI applications
- Memory Access analysis
- Accelerators (GPU and FPGA) analysis

# Intel® VTune™ Profiler 2020 Analysis Types



## Start with

- Command line
- Within a GUI

## Which type of analysis

- Performance Snapshot for initial assessment
- With little experience you know which one to use

## Simplest command line

- **vtune -collect performance-snapshot ./my\_app**
- Almost every GUI functionality is available in CL

## Launching analysis in a cluster environment

- The same CL, but wrapped into a mpi launcher commands

### Optimization Notice

# VTune Performance Snapshot

**Performance Snapshot (preview)** Performance Snapshot

Analysis Configuration Collection Log Summary

### Choose your next analysis type

Select a highlighted recommendation based on your performance snapshot.

- ALGORITHM**
  - Hotspots
  - Anomaly Detection (preview)
  - Memory Consumption
- MICROARCHITECTURE**
  - Microarchitecture Exploration 60.3%
  - Memory Access 20.7%
- PARALLELISM**
  - Threading 76.9%
  - HPC Performance Characterization
- ACCELERATORS**
  - GPU Offload (preview)
  - GPU Compute/Media Hotspots (preview)
  - CPU/FPGA Interaction
- PLATFORM ANALYSES**
  - System Overview
  - Throttling (preview)
  - Platform Profiler
- I/O**
  - Input and Output

**Elapsed Time<sup>®</sup>: 8.092s**

- CPU**
  - IPC<sup>®</sup>: 1.626
  - SP GFLOPS<sup>®</sup>: 0.000
  - DP GFLOPS<sup>®</sup>: 8.441
  - x87 GFLOPS<sup>®</sup>: 0.000
  - Average CPU Frequency<sup>®</sup>: 3.9 GHz
- GPU**

**Effective Logical Core Utilization<sup>®</sup>: 76.9%**  
(6.149 out of 8) ⬇

**Microarchitecture Usage<sup>®</sup>: 60.3%** ⬇  
of Pipeline Slots

**Memory Bound<sup>®</sup>: 20.7%** ⬇  
of Pipeline Slots

**Vectorization<sup>®</sup>: 48.7%** ⬇  
of Packed FP Operations

**GPU Active Time<sup>®</sup>: 0.0%** ⬇

# Performance Snapshot report in CL

vtune -collect performance-snapshot ./my\_app

Memory Bound: 77.5% of Pipeline Slots

| The metric value is high. This can indicate that the significant fraction of  
| execution pipeline slots could be stalled due to demand memory load and  
| stores. Use Memory Access analysis to have the metric breakdown by memory  
| hierarchy, memory bandwidth information, correlation by memory objects.

L1 Bound: 0.0% of Clockticks

L2 Bound: 0.0% of Clockticks

L3 Bound: 60.2% of Clockticks

| This metric shows how often CPU was stalled on L3 cache, or contended  
| with a sibling Core. Avoiding cache misses (L2 misses/L3 hits) improves  
| the latency and increases performance.

DRAM Bound: 13.1% of Clockticks

| This metric shows how often CPU was stalled on the main memory (DRAM).  
| Caching typically improves the latency and increases performance.

DRAM Bandwidth Bound: 0.0% of Elapsed Time

Store Bound: 0.0% of Clockticks

**NUMA: % of Remote Accesses: 20.9%**

| **A significant amount of DRAM loads were serviced from remote DRAM.**  
| **Wherever possible, try to consistently use data on the same core, or at**  
| **least the same package, as it was allocated on.**

# HPC Performance Characterization Analysis

Show important aspects of application performance in one analysis

- Entry point to assess application efficiency with definition of the next steps to investigate events with significant performance cost

## Threading: CPU Utilization

- Serial vs. Parallel time
- Top OpenMP regions by potential gain
- Tip: Use Hotspot OpenMP region analysis for more detail

## Memory Access Efficiency

- Stalls by memory hierarchy
- Bandwidth utilization
- Tip: Use Memory Access analysis

## Vectorization: FPU Utilization

- FLOPS † estimates from sampling
- Tip: Use Intel Advisor for precise metrics and vectorization optimization

>vtune –collect **hpc-performance** –data-limit=0 –r result\_dir ./my\_app

**HPC Performance Characterization**

Analysis Configuration | Collection Log | Summary | Bottom-up

Elapsed Time <sup>Ⓢ</sup>: 10.253s

SP GFLOPS <sup>Ⓢ</sup>: 129.325

Effective Physical Core Utilization <sup>Ⓢ</sup>: **52.7%** (23.181 out of 44) <sup>⚠</sup>

- Effective Logical Core Utilization <sup>Ⓢ</sup>: 52.3% (46.012 out of 88) <sup>⚠</sup>
- Serial Time (outside parallel regions) <sup>Ⓢ</sup>: 0.137s (1.3%)
- Parallel Region Time <sup>Ⓢ</sup>: 10.116s (98.7%)
  - Estimated Ideal Time <sup>Ⓢ</sup>: 5.623s (54.8%)
  - OpenMP Potential Gain <sup>Ⓢ</sup>: 4.493s (43.8%) <sup>⚠</sup>
  - Top OpenMP Regions by Potential Gain
  - Effective CPU Utilization Histogram

Memory Bound <sup>Ⓢ</sup>: **21.9%** <sup>⚠</sup> of Pipeline Slots

- Cache Bound <sup>Ⓢ</sup>: 25.1% <sup>⚠</sup> of Clockticks
- DRAM Bound <sup>Ⓢ</sup>: 6.7% <sup>⚠</sup> of Clockticks
- NUMA: % of Remote Accesses <sup>Ⓢ</sup>: 43.3%
- Bandwidth Utilization Histogram

FPU Utilization <sup>Ⓢ</sup>: **1.9%** <sup>⚠</sup>

- SP FLOPs per Cycle <sup>Ⓢ</sup>: 0.505 Out of 22
- Vector Capacity Usage <sup>Ⓢ</sup>
- FP Instruction Mix:
  - % of Packed FP Instr. <sup>Ⓢ</sup>:
    - % of 128-bit <sup>Ⓢ</sup>
    - % of 256-bit <sup>Ⓢ</sup>
  - % of Scalar FP Instr. <sup>Ⓢ</sup>: 88.9% <sup>⚠</sup>
- FP Arith/Mem Rd Instr. Ratio <sup>Ⓢ</sup>: 0.862
- FP Arith/Mem Wr Instr. Ratio <sup>Ⓢ</sup>: 2.459
- Top Loops/Functions with FPU Usage by CPU Time

**Optimization Notice:** A significant fraction of floating point arithmetic instructions are scalar. Use Intel Advisor to see possible reasons why the code was not vectorized.

### Optimization Notice



# Performance Aspects: CPU Utilization (1/4)

Collection Log Analysis Target Analysis Type Summary Bottom-up

Elapsed Time: 7.805s

SP GFLOPS: 14.041

CPU Utilization: 76.4%

Memory Bound: 63.2%

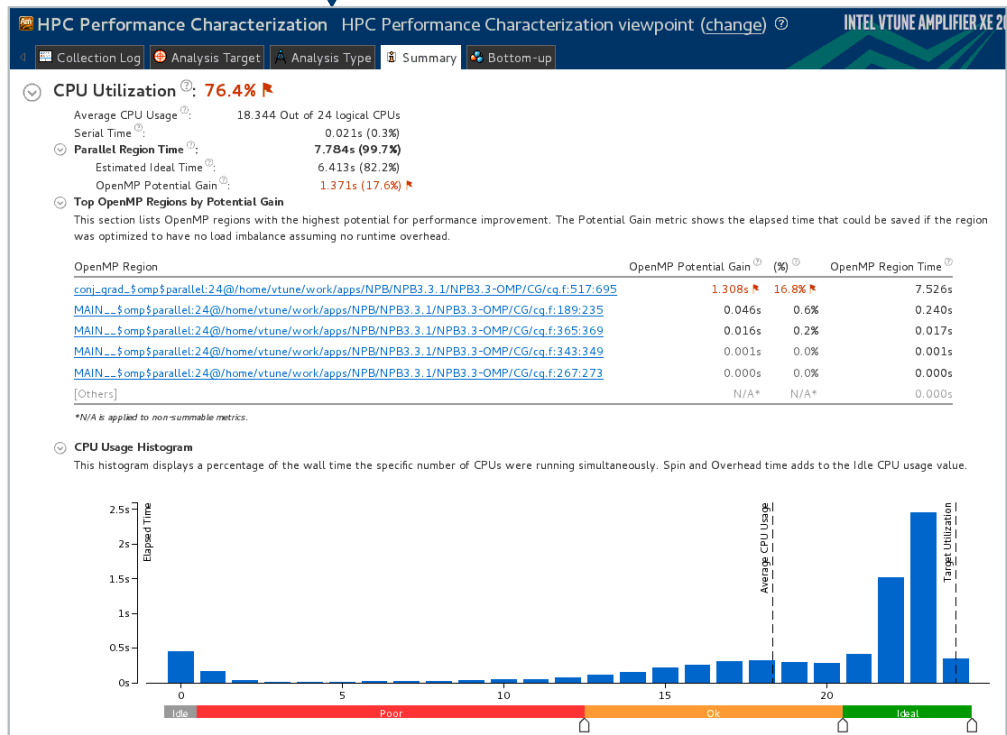
FPU Utilization: 1.3%

## CPU Utilization

- % of “Effective” CPU usage by the application under profiling (threshold 90%)
  - Under assumption that the app should use all available logical cores on a node
  - Subtracting spin/overhead time spent in MPI and threading runtimes

## Metrics in CPU utilization section

- Average CPU usage
- Intel OpenMP scalability metrics impacting effective CPU utilization
- CPU utilization histogram



### Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



# Performance Aspects: CPU Utilization (2/4)

VTune Profiler OpenMP\* Analysis: answering on customers' questions about performance in the same language a program was written in

Is serial time of my application significant to prevent scaling?

Serial Time (outside any parallel region): 4.020s (27.7%)

Serial time of your application is high. It directly impacts application Elapsed Time and scalability. Explore options for parallelization, algorithm or microarchitecture tuning of the serial part of the application.

How efficient is my parallelization towards ideal parallel execution?

Parallel Region Time: 10.469s (72.3%)

Estimated Ideal Time: 7.115s (49.1%)

How much theoretical gain I can get if invest in tuning?

Potential Gain: 3.354s (23.1%)

The time wasted on load imbalance or parallel work arrangement is significant and negatively impacts the application performance and scalability. Explore OpenMP regions with the highest metric values. Make sure the workload of the regions is enough and the loop schedule is..

## Top OpenMP Regions by Potential Gain

This section lists OpenMP regions with the highest potential for performance improvement. The Potential Gain metric shows the elapsed time that could be saved if the region was optimized to have no load imbalance assuming no runtime overhead.

OpenMP Region	Potential Gain (%)	Elapsed Time
<a href="#">conj_grad_\$omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:514:695</a>	3.294s 22.7%	10.208s
<a href="#">MAIN__\$omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:185:231</a>	0.059s 0.4%	0.260s
<a href="#">MAIN__\$omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:339:345</a>	0.001s 0.0%	0.001s
<a href="#">MAIN__\$omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:361:365</a>	0.001s 0.0%	0.001s
<a href="#">MAIN__\$omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:263:269</a>	0.000s 0.0%	0.000s
[Others]	0.000s 0.0%	0.000s

What regions are more perspective to invest?

Links to grid view for more details on inefficiency



# Performance Aspects: CPU Utilization (3/4)

Details in grid view: expansion by constructs with a barrier inside the region

Imbalance distribution by loop, single, reduction, user, join barriers

Advanced Hotspots Hotspots viewpoint (change) Intel VTune Amplifier XE 20

Collection Log Analysis Target Analysis Type Summary Bottom-up Caller/Callee Top-down Tree Platform

Imbalance on a loop barrier

Grouping: OpenMP Region / OpenMP Barrier-to-Barrier Segment / Function / Call Stack

OpenMP Region / OpenMP Barrier-to-Barrier Segment / Function / Call Stack	OpenMP Potential Gain							Elapsed Time	Number of OpenMP threads	Ins... Cou...	OpenMP Loop Schedule Type	OpenMP Loop Chunk	Avg OpenMP Loop Iteration Count
	Imbalance	Lock Con...	Creation	Schedu...	Redu...	Atomi...	Other						
conjugrad_omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:514:695	3.944s	0s	0.000s	0.002s	0.000s	0s	0.094s	11.095s	24	76			
conjugrad_omp\$loop_barrier_segment@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:500	3.725s	0s	0s	0.000s	0s	0s	0.000s	10.115s	24		Static	3125	75,000
conjugrad_omp\$loop_barrier_segment@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:683	0.149s	0s	0s	0s	0s	0s	0.004s	0.418s	24		Static	3125	75,000
conjugrad_omp\$loop_barrier_segment@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:625	0.033s	0s	0s	0.002s	0.000s	0s	0.002s	0.068s	24		Static	3125	75,000
conjugrad_omp\$loop_barrier_segment@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:650	0.015s	0s	0s	0.000s	0s	0s	0.001s	0.064s	24		Static	3125	75,000
conjugrad_omp\$loop_barrier_segment@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:664	0.014s	0s	0s	0.000s	0s	0s	0.001s	0.079s	24		Static	3125	75,000

Advanced Hotspots Hotspots viewpoint (change) Intel VTune Amplifier XE 2015

Collection Log Analysis Target Analysis Type Summary Bottom-up Caller/Callee Top-down Tree Tasks and Frames

Grouping: OpenMP Region / OpenMP Barrier-to-Barrier Segment / Function / Call Stack

OpenMP Region / OpenMP Barrier-to-Barrier Segment / Function / Call Stack	OpenMP Potential Gain							OpenMP Potential Gain (% of Collection T...					Elap... Time	Nu. of Ope... thr...	Ins. Co.	Ope... Loop... Chu.	Open... Loop Sched... Type
	Imba...	Lock Con...	Cre...	Scheduling	Red..	Oth.	Imba... (%)	Lock Cont...	Cre... (%)	Scheduli... (%)	Red.. (%)	Oth... (%)					
conjugrad_omp\$parallel:24@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:514:695	0.20...	0.0...	0.0...	3.127s	0.0...	0.0...	1.7%	0.0%	0.0%	25.9%	0.0%	0.0%	11.7...	24	76		
conjugrad_omp\$loop_barrier@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:572:580	0.00...	0.0...	0s	3.125s	0s	0s	0.1%	0.0%	0.0%	25.9%	0.0%	0.0%	11.1...	24	1	Dynamic	
conjugrad_omp\$loop_barrier@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:675:683	0.12...	0s	0s	0s	0s	0s	1.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.41...	24	312.	Static	
conjugrad_omp\$loop_barrier@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:621:625	0.02...	0s	0s	0.001s	0s	0s	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.07...	24	312.	Static	
conjugrad_omp\$loop_barrier@/home/vtune/work/apps/NPB/NPB3.3.1/NPB3.3-OMP/CG/cg.f:637:650	0.02...	0s	0s	0.000s	0s	0s	0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.07...	24	312.	Static	

Dynamic scheduling overhead on a parallel loop

## Optimization Notice

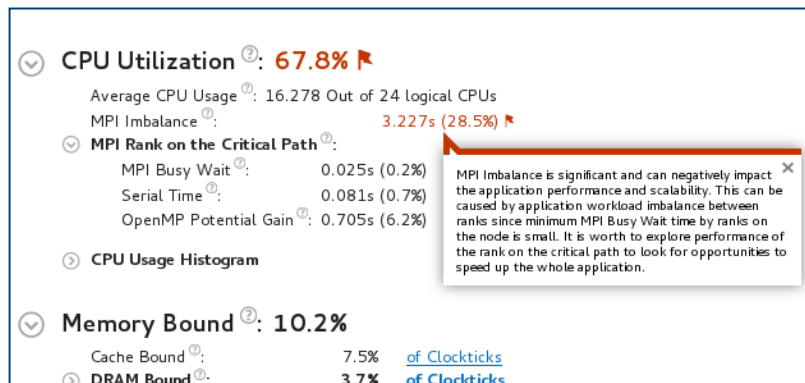
Copyright © 2020, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



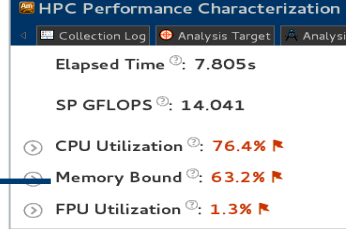
# Performance Aspects: CPU Utilization (4/4)

## Specifics for hybrid MPI + OpenMP apps

- MPI Imbalance metric for the node as average MPI busy wait time by rank (based on MPI progress sampling for MPICH-based MPIs (Intel MPI, CRAY MPI, ..\_))
- Serial (outside any parallel region) and OpenMP Potential Gain time for the MPI rank on critical path for the node (with minimal MPI busy wait per the node)
- Worth to explore details on serial hotspots and OpenMP inefficiencies for the MPI rank on critical path in grid view



# Performance Aspects: Memory Bound

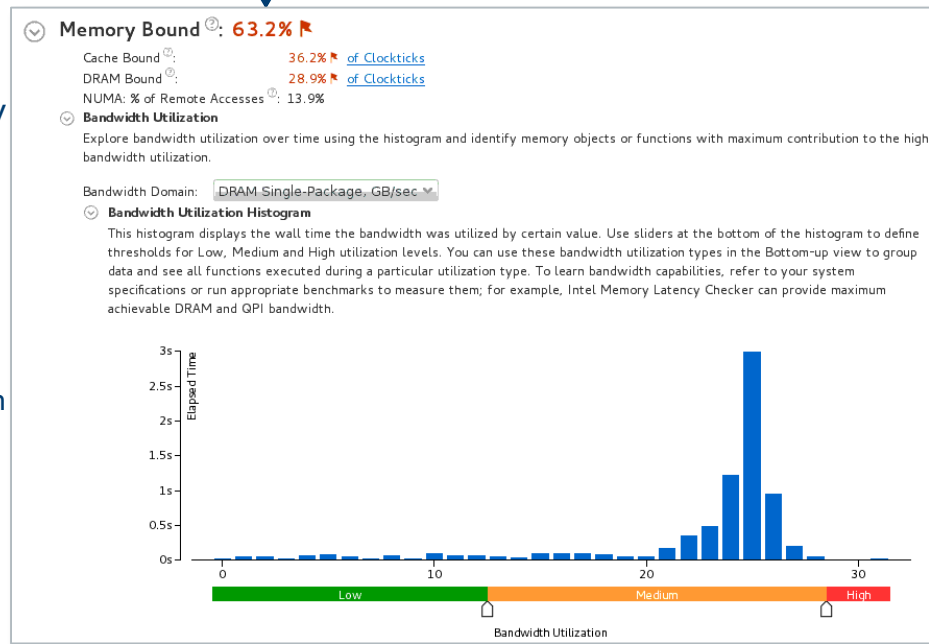


## Memory Bound

- % of potential execution pipeline slots lost because of fetching memory from different levels of hierarchy (threshold 20%)

## Metrics in Memory Bound section

- Cache bound
- DRAM bound
  - Issue description specifies if the code is bandwidth or latency bound with proper advice of how to fix
  - % of app elapsed time consuming high DRAM Bandwidth
  - NUMA: % of remote accesses
    - Important to explore if the code is bandwidth bound
  - Bandwidth utilization histogram



### Optimization Notice

# Performance Aspects: FPU Utilization

HPC Performance Characterization HPC Performance  
Collection Log Analysis Target Analysis Type Summary  
Elapsed Time: 101.194s  
GFLOPS Upper Bound: 24.612  
CPU Utilization: 12.7%  
Back-End Bound: 87.8%  
FPU Utilization Upper Bound: 1.4%

## FPU utilization

- % of FPU load (100% - FPU is fully loaded, threshold 50%)

## Metrics in FPU utilization section

- SP FLOPs per Cycle (vector code generation and execution efficiency)
- Vector Capacity Usage and FP Instruction Mix, FParith/Mem ratios (vector code generation efficiency)
- Top 5 loops/functions by FPU usage
  - Dynamically generated issue descriptions on low FPU usage help to define the reason and next steps:

Non-vectorized, vectorized with legacy instruction set, memory bound limited loops not benefiting from vectorization etc.

FPU Utilization: 1.3%  
SP FLOPs per Cycle: 0.211 Out of 16  
Vector Capacity Usage: 48.3%  
FP Instruction Mix:

- % of Packed FP Instr.: 93.1%
  - % of 128-bit: 93.1%
  - % of 256-bit: 0.0%
  - % of Scalar FP Instr.: 6.9%
- FP Arith/Mem Rd Instr. Ratio: 0.264
- FP Arith/Mem Wr Instr. Ratio: 6.298

**Top 5 hotspot loops (functions) by FPU usage**  
This section provides information for the most time consuming loops/functions with floating point operations.

Function	CPU Time	FPU Utilization	Vector Instruction Set	Loop Type
<a href="#">[Loop at line 575 in conj_grad_omp\$parallel@517]</a>	126.149s	1.6%	SSE2(128)	Body
<a href="#">[Loop at line 678 in conj_grad_omp\$parallel@517]</a>	5.004s	1.7%	SSE2(128)	Body
<a href="#">[Loop at line 575 in conj_grad_omp\$parallel@517]</a>	2.678s	2.1%	[Unknown]	Remainder
<a href="#">[Loop at line 573 in conj_grad_omp\$parallel@517]</a>	0.995s	4.0%	SSE2(128)	Body
<a href="#">[Loop at line 661 in conj_grad_omp\$parallel@517]</a>	0.952s	1.3%	SSE2(128); SSE2(128)	Body
[Others]	2.437s	N/A*	N/A*	N/A*

\*N/A is applied to non-summable metrics.

# Performance Aspects: Process/Thread Affinity

## Process/Thread Affinity Collection and Reporting

- Visualization of affinity masks collected at the end of application
- Command Line Report with threads by cores affinity
- Experimental GUI report with affinity shown along with thread execution on CPU and remote accesses

## How to

```
>vtune -c h-p -knob collect-affinity=true -r my_result ./my_app
```

## CL report:

```
>vtune -R affinity -r my_result
```

## GUI report (tech preview):

```
>vtune -R affinity -format=html -r my_result
```

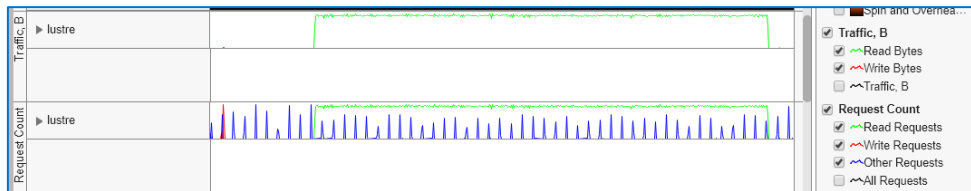


# Performance Aspects: Lustre Parallel File I/O

- Key Metrics:
  - I/O Wait Time, I/O Bandwidth, Packet Rate, Packet Size
- Availability:
  - Average on summary
  - Overtime on timeline with ability to filter by time for context summary
- Details by Parallel File System Shares
- To enable set “AMPLXE\_EXPERIMENTAL=lustre”, add “-knob parallel-fs-collection=true” to hpc-performance command line

📌 **Lustre File System IO Wait Time: 64.569s**

Bandwidth, B/s: 6,011,995.982  
Package Rate, pack/s: 25.750  
Average Package Size, KB: 233



📌 **Parallel File System detailed info:**

Parallel File System	Read Bytes	Read Requests	Read Wait Time	Write Bytes	Write Requests	Write Wait Time	All Requests	Request Wait Time
/ifs/ifs09	6,161,204,090	26,379	48.513s	381,814	12	0.011s	27,086	49.799s
/ifs/ifs12	0	0	0s	0	0	0s	6,058	11.961s
/ifs/ifs13	0	0	0s	0	0	0s	963	2.020s
/ifs/ifs11	0	0	0s	0	0	0s	635	0.789s

*\*N/A is applied to non-summable metrics.*

# HPC Performance Characterization – Command Line Reporting

- Generated after collection is done or with “-R summary” option of vtune
- Mimicking GUI metrics hierarchy

With issue descriptions that can be suppressed by “-report-knob show-issues=false” option

```
Elapsed Time: 7.805s
SP GFLOPS: 14.041
CPU Utilization: 76.4%
| The metric value is low, which may signal a poor logical CPU cores
| utilization caused by load imbalance, threading runtime overhead, contended
| synchronization, or thread/process underutilization. Explore CPU Utilization
| sub-metrics to estimate the efficiency of MPI and OpenMP parallelism or run
| the Loads and Waits analysis to identify parallel bottlenecks for other
| parallel runtimes.
|
| Average CPU Usage: 18.344 Out of 24 logical CPUs
Serial Time: 0.021s (0.3%)
Parallel Region Time: 7.784s (99.7%)
  Estimated Ideal Time: 6.413s (82.2%)
  OpenMP Potential Gain: 1.371s (17.6%)
  | The time wasted on load imbalance or parallel work arrangement is
  | significant and negatively impacts the application performance and
  | scalability. Explore OpenMP regions with the highest metric values.
  | Make sure the workload of the regions is enough and the loop schedule
  | is optimal.
|
Memory Bound: 63.2% of Pipeline Slots
| The metric value is high. This can indicate that the significant fraction of
| execution pipeline slots could be stalled due to demand memory load and
| stores. Use Memory Access analysis to have the metric breakdown by memory
| hierarchy, memory bandwidth information, correlation by memory objects.
|
Cache Bound: 36.2% of Clockticks
| A significant proportion of cycles are being spent on data fetches from
| caches. Check Memory Access analysis to see if accesses to L2 or L3
| caches are problematic and consider applying the same performance tuning
| as you would for a cache-missing workload. This may include reducing the
| data working set size, improving data access locality, blocking or
| partitioning the working set to fit in the lower cache levels, or
| exploiting hardware prefetchers. Consider using software prefetchers, but
| note that they can interfere with normal loads, increase latency, and
| increase pressure on the memory system. This metric includes coherence
| penalties for shared data. Check General Exploration analysis to see if
| contested accesses or data sharing are indicated as likely issues.
|
DRAM Bound: 20.9% of Clockticks
| The metric value is high. This indicates that a significant fraction of
| cycles could be stalled on the main memory (DRAM) because of demand loads
| or stores.
|
| The code is memory bandwidth bound, which means that there are a
| significant fraction of cycles during which the bandwidth limits of the
| main memory are being reached and the code could stall. Review the
| Bandwidth Utilization Histogram to estimate the roots of the issue.
| Consider improving data locality on NUMA multi-socket systems, which will
| reduce code memory bandwidth consumption.
|
NUMA: % of Remote Accesses: 13.9%
```

```
Elapsed Time: 7.805s
SP GFLOPS: 14.041
CPU Utilization: 76.4%
  Average CPU Usage: 18.344 Out of 24 logical CPUs
  Serial Time: 0.021s (0.3%)
  Parallel Region Time: 7.784s (99.7%)
    Estimated Ideal Time: 6.413s (82.2%)
    OpenMP Potential Gain: 1.371s (17.6%)
Memory Bound: 63.2% of Pipeline Slots
  Cache Bound: 36.2% of Clockticks
  DRAM Bound: 20.9% of Clockticks
  NUMA: % of Remote Accesses: 13.9%
FPU Utilization: 1.3%
SP FLOPS per Cycle: 0.211 Out of 16
Vector Capacity Usage: 48.3%
FP Instruction Mix
  % of Packed FP Instr.: 93.1%
  % of 128-bit: 93.1%
  % of 256-bit: 0.0%
  % of Scalar FP Instr.: 6.9%
FP Arith/Mem Rd Instr. Ratio: 0.264
FP Arith/Mem Wr Instr. Ratio: 6.298
Collection and Platform Info
Application Command Line: ./cg.B.x
User Name: vtune
```

# How to Run VTune on MPI Applications


Single or multiple node application launch:

```
<mpi_launcher> -n N <vtune_command_line> ./app_to_run
```

- >aprun -n 48 -N 16 vtune -collect hotspots **-trace-mpi** -r result\_dir ./my\_mpi\_app
- >mpirun -n 48 -ppn 16 vtune -collect hotspots -r result\_dir ./my\_mpi\_app

- Encapsulates ranks to per-node result directories suffixed with hostname

- result\_dir.hostname1 with 0-15, result\_dir.hostname2 with 16-31, result\_dir.hostname3 with 32-47

 • **-trace-mpi** option is mandatory to enable per-node result directories for MPIs not using PMI\_RANK (e.g. CRAY MPI)

- Works for software and driver-based/driverless collectors

- Highly recommended to set /proc/sys/kernel/perf\_event\_paranoid to 0 for perf-based collections to avoid collection overhead on modern multi-/many-core systems



# Selective Rank profiling of MPI applications with VTune

To reduce data collected by VTune on big scale runs use selective rank profiling using MPMD (Multiple Program Multiple Data) runs:

Example: profile rank 1 from 0-15:

```
>mpirun -n 1 ./my_app : -n 1 <vtune_command_line> ./my_app : -n 14 ./my_app
```

- In the case of Intel MPI launcher `-gtool` option can be used to simplify the launch configuration

Example: profile ranks 2, 3, 10-15 from 0-15:

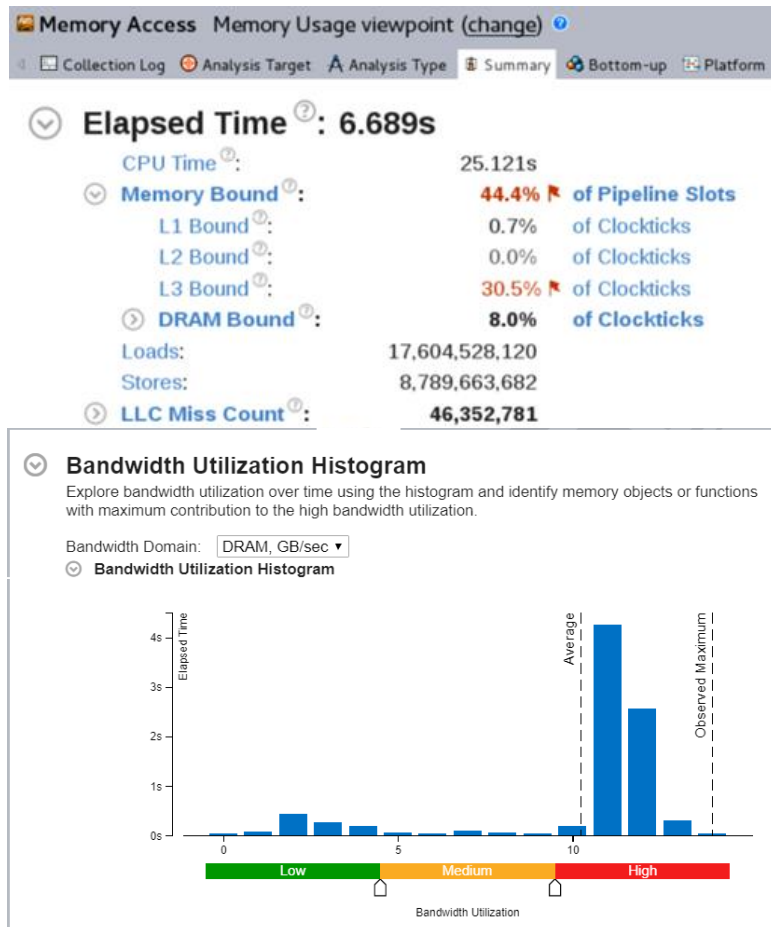
```
>mpirun -gtool "vtune -collect hotspots -r result_dir:2,3,10-15" ./my_app
```

# Memory Access Analysis

Specialized analysis type “Memory Access” to focus on memory-related issues:

- Performance problems by memory hierarchy (e.g. L1-, L2-, LLC-, DRAM-bound)
- Bandwidth-limited accesses
  - Covers DRAM, QPI, and PMEM bandwidth
- NUMA problems
- Performance metrics by memory objects (data structures)

```
>vtune -collect memory-access ./my_app
```



## Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



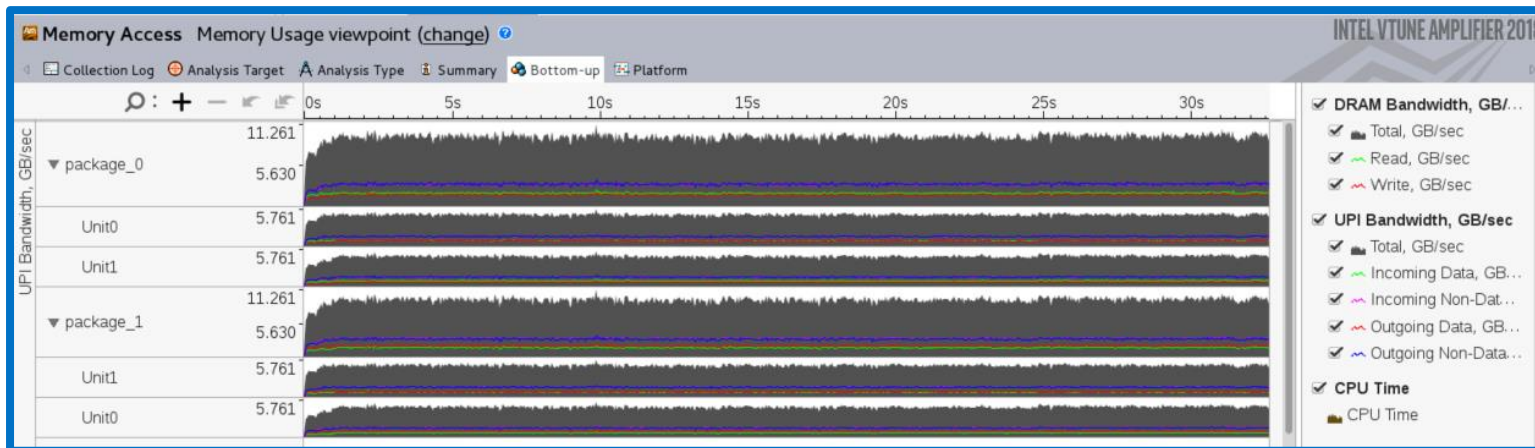
# Memory Bandwidth Analysis

## DRAM and PMEM bandwidth

- Can be expanded per package/channel
- Split to Read/Write
- Automatic DRAM pick measurement on Xeon

## QPI bandwidth per-package

- Can be expanded per-link
- Split to Incoming/Outgoing and Data/Non-Data



### Optimization Notice

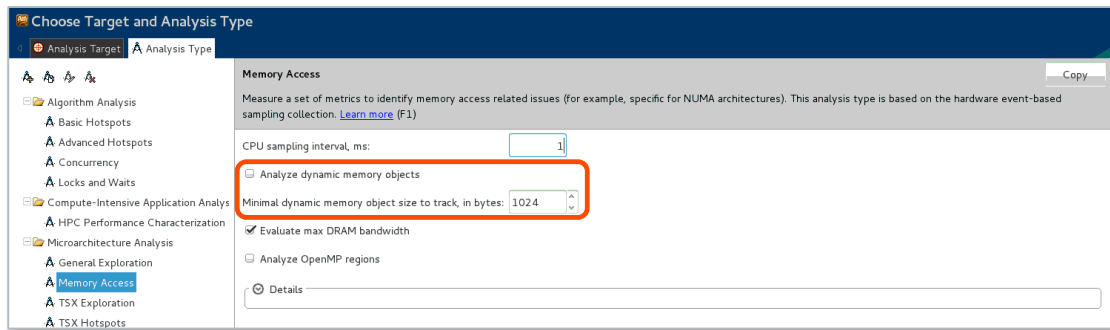
Copyright © 2020, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.

# View Performance Metrics by Memory Objects

## Memory Access analysis configuration options for profiling memory objects:

- **Analyze dynamic memory objects:** enables the instrumentation of memory allocation/de-allocation and map hardware events to memory objects
- **Minimal memory object size to track, in bytes:** specify a minimal size of memory allocations to analyze. This option helps reduce runtime overhead of the instrumentation. Default is 1K.

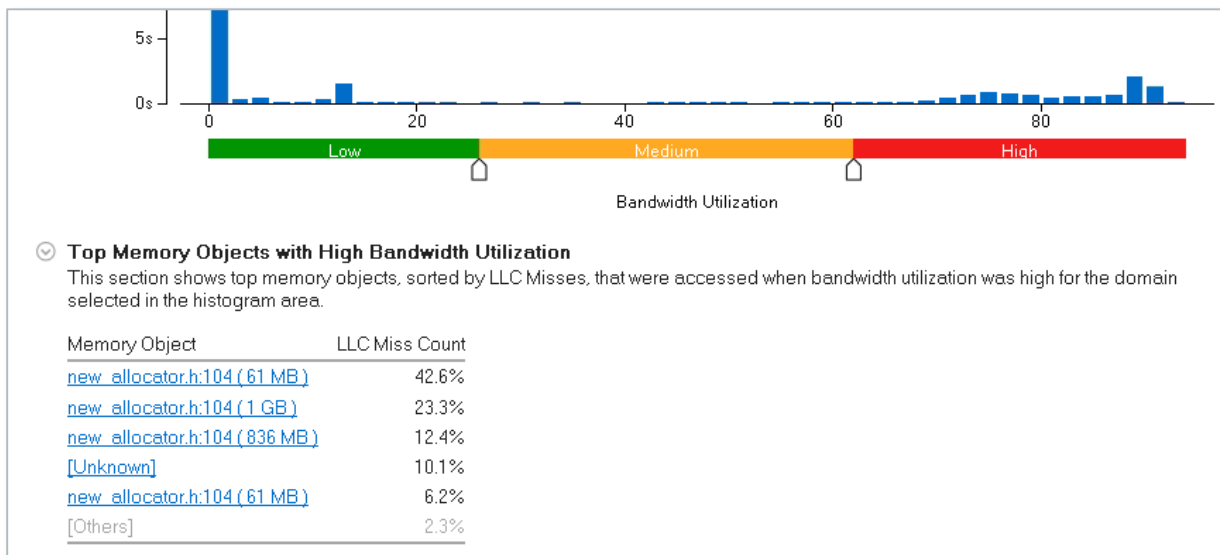
```
>vtune -c memory-access -knob analyze-mem-objects=true -knob mem-object-size-min-thres=2048 -- <app>
```



### Optimization Notice

# High Memory Bandwidth Analysis by Memory Objects

If memory object instrumentation is ON Bandwidth Utilization histogram is enriched with the table of Memory Objects by LLC Miss Count

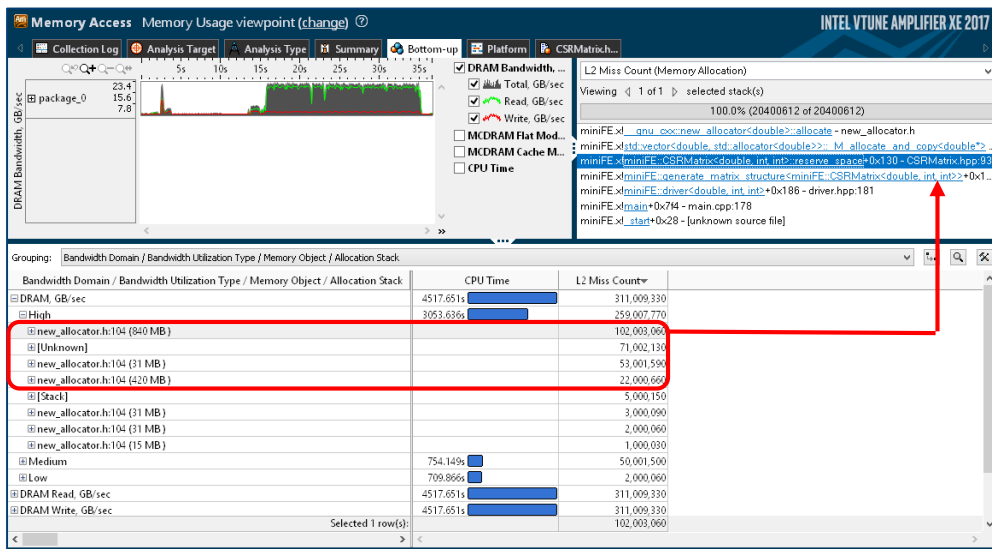


## Optimization Notice

# High Memory Bandwidth Analysis by Memory Objects

## Investigate the memory allocations inducing bandwidth

- “Bandwidth Domain/Bandwidth Utilization Type/Memory Object/Allocation Stack” grouping with expansion by “DRAM/High” and sorting by L2 Miss Count



The screenshot shows the source code view of the function 'void reserve\_space' in 'main.cpp'. The code is as follows:

```
88 void reserve_space(unsigned nrows, unsigned ncols_per_row)
89 {
90     rows.resize(nrows);
91     row_offsets.resize(nrows+1);
92     packed_cols.reserve(nrows * ncols_per_row);
93     packed_coefs.reserve(nrows * ncols_per_row);
94 }
95 #pragma omp parallel for
96 for(MINIFE_GLOBAL_ORDINAL i = 0; i < nrows; ++i) {
97     rows[i] = 0;
98     row_offsets[i] = 0;
99 }
```

A red arrow points from the top entry of the 'High' category in the table to line 93 of the code.

Focus on allocations inducing L2 misses

Allocation stack shows the allocation place in user's code

# GPU Application Analysis

## GPU Compute/Media Hotspots

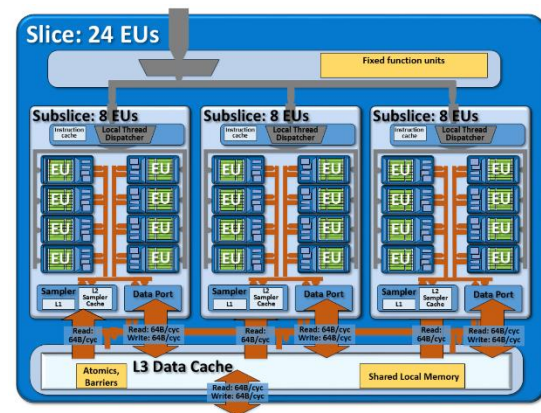
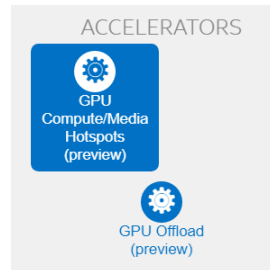
- Visibility into both host and GPU sides
- HW-events based performance tuning methodology
- Provides overtime and aggregated views

## GPU In-kernel Profiling

- GPU source/instruction level profiling
- SW instrumentation
- Two modes: basic Block latency and Memory access latency

Identify GPU occupancy and which kernel to profile.

Tune a kernel on a fine-grain level.



### Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.

# CPU/FPGA Interaction Analysis

Now process data sources collected with AOCL Profiler (new mode) in addition to OpenCL Profiling API (legacy mode).

Extended with FPGA device-side metrics, like Stalls, Global Bandwidth and Occupancy, and mapping FPGA kernel performance data to the source code

HOW

## CPU/FPGA Interaction (preview) ▾

Preview feature - should we keep it, change it, or drop it? [Send us your comments.](#)

Analyze the CPU/FPGA interaction issues via exploring OpenCL kernels running on the FPGA and identify the most time-consuming kernels. [Learn more](#)

**CPU sampling interval, ms**

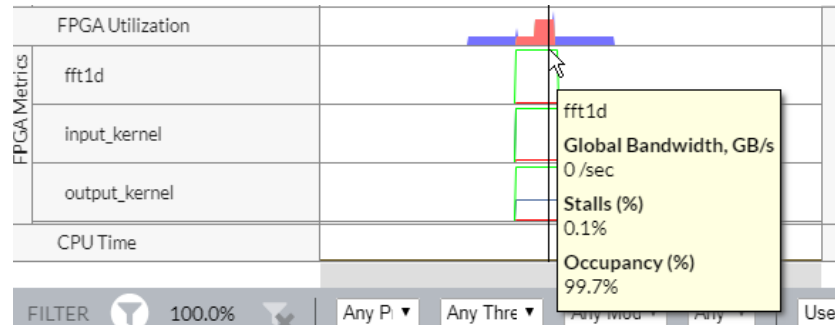
Collect stacks

**FPGA profiling data source**

AOCL Profiler ▾

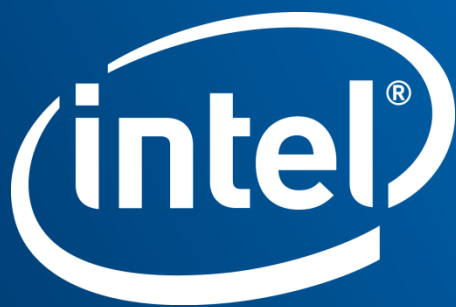
**AOCL Profiler**

OpenCL Profiling API



## Optimization Notice





# Legal Disclaimer & Optimization Notice

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. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

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.

Copyright © 2018, 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

## Optimization Notice