

Intended Audience: Software Developers

- Interested in performance optimizing your application
 - >Don't need to be a performance expert
 - >But should be an expert in the application!
- Working on a platform with a 2nd generation Intel® Core™ processor
- Using Intel® VTune™ Amplifier XE performance analyzer
 - >The performance information here applies to other tools (PTU, etc) but is focused on VTune Amplifier XE

Software and Services Group



1

How to Use this Presentation

- Read through the slides once, then again while collecting data
- Remember performance analysis is a process that may take several iterations
- Software Optimization should begin *after you have*:
 - >Utilized any compiler optimization options (/O2, /QxSSE4.2, etc)
 - >Chosen an appropriate workload
 - >Measured baseline performance



Using **Intel® VTune™ Amplifier XE** to Tune Software on the **2nd generation Intel® Core™** **processor family**

Software and Services Group

Ver. 1.01

<http://software.intel.com/en-us/articles/using-intel-vtune-amplifier-xe-to-tune-software-on-the-2nd-generation-intel-core-processor-family/>

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

- Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.
- The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.
- Copies of documents which have an order number and are referenced in this document or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website.
- Intel® Hyper-Threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology-enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. For more information including details on which processors support HT Technology, see [here](#).
- Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain computer system software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.
- 64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information.
- Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see <http://www.intel.com/technology/turboboost>.
- Intel, the Intel logo, VTune, inTru, and Core are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- *Other names and brands are the property of their respective owners.
- Copyright © 2011, Intel Corporation



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

Software and Services Group



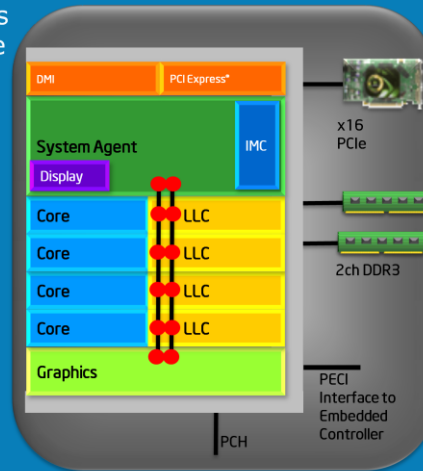
5

Agenda

- Intel® Microarchitecture Codename Sandy Bridge
- The new Intel® VTune™ Amplifier XE
- The Software Optimization Cycle
 - > Hotspots
 - > Methods for Determining Efficiency
 - > Common Architectural Causes of Inefficiency:
 - **Cache Misses**
 - Other Data Access Issues
 - Execution Stalls
 - Branch Mispredicts
 - Front End Stalls

Intel® Microarchitecture Codename Sandy Bridge

- Manufactured on Intel® 32nm process technology – delivering a performance and energy boost
- Full integration of processor cores, memory controller, last-level cache, and graphics and media processing
- Intel® Turbo Boost 2.0 Technology
- Intel® Hyper-Threading Technology
- Hardware-based media and graphics accelerators:
 - Intel® Quick Sync Video
 - Intel® HD Graphics
 - Intel® Clear Video HD Technology
 - Intel® InTru™ 3D Technology
- Intel® Advanced Vector Extensions (AVX) Instructions



Software and Services Group



7

For how to tell which product numbers are Sandy Bridge, check here:

http://www.intel.com/products/processor_number/about/core.htm

For more on these features see

<http://www.intel.com/consumer/products/processors/core-family.htm> or <http://www.intel.com/technology/architecture-silicon/2ndgen/index.htm>

The New Intel® VTune™ Amplifier XE

- The latest version of VTune Analyzer was released in November 2010
- VTune Amplifier XE features:
 - Multiple Collection Types
 - > Hotspots (statistical call tree)
 - > Thread Concurrency
 - > Locks and Waits Analysis
 - > Event-based Sampling
 - Timeline View Integrated into all Analysis Types
 - Source/Assembly Viewing
 - Compatible with C/C++, Fortran, Assembly, .NET
 - Visual Studio Integration, Command-line, or Standalone interface for Windows* or Linux*



Software and Services Group



New in Amplifier XE: Pre-Configured Profiles!

The Intel® Microarchitecture Codename Sandy Bridge: General Exploration profile should be used for a top-level analysis of potential issues. It is the subject of this guide.

Intel(R) Microarchitecture Code Name Sandy Bridge - General Exploration
Analyze general issues affecting the performance of your application. This analysis type is based on the hardware event-based sampling collection. Press F1 for more details.

Details
To modify collector options for a predefined analysis type, right-click the analysis type in the tree, select Copy from Current entry in the pop-up menu, and edit the copy of the selected analysis type configuration.
Events configured for CPU: Intel(R) Core(TM) Processor 2xxx Series

NOTE: For analysis purposes, Amplifier XE may adjust the Sample After values in the table below by a multiplier. The multiplier depends on the value of the Duration time estimate option specified in the Project Properties dialog.

Event Name	Sample After	LBR Filter	Description
ARITH_FPU_DIV_ACTIVE	2000000	None	Cycles that the divider is busy with
BR_MISP_RETIRED_ALL_BRANCHES_PS	400000	None	all macro branches (Precise Event)
CPU_CLK_UNHALTED_REF_TSC	2000000	None	Reference cycles when thread is in
CPU_CLK_UNHALTED_THREAD	2000000	None	Cycles when thread is not halted (f
D8B2MTE_SWITCHES.PENALTY_CYCLES	2000000	None	D8B-to-MTE switch true penalty c
DTLB_LOAD_MISSES.STLB_HIT	1000000	None	First level miss but second level hit
DTLB_LOAD_MISSES.WALK_DURATION	2000000	None	Cycles PPH1 is busy with this walk.
INSTR_CACHE_MISSES	2000000	None	*Number of Instruction Cache, Stre
INSTR_CACHE_MISSES_CYCLES	2000000	None	*Number of cycles that Uops were

**Video online!
See notes.**

All the events required are pre-configured – no research needed!
Simply click Start to run the analysis.

Software and Services Group 9

The logic for identifying issues on Microarchitecture Codename Sandy Bridge is embedded into the interface. All the formulas and metrics used are the same as the ones given in this guide. You no longer have to apply formulas and rules to the data yourself to figure out what it means – using this guide and the interface tuning features, you can easily pinpoint problems and possible solutions.

The formulas and metrics are only applied to the General Exploration profile, and the General Exploration viewpoint (which is automatic). For the other profiles, it will just show the raw data.

Also view our video demo of this interface at:

<http://software.intel.com/en-us/videos/channel/parallel-programming/the-intel-vtune-amplifier-xe-analysis-and-results-interface-for-intel-microarchitecture-codename-sandy-bridge/1265162566001>

Enhanced General Exploration View for Intel® Microarchitecture Codename Sandy Bridge

The enhanced view is present when running the General Exploration profile with the General Exploration viewpoint selected (the default).

Function	Hardware Event Count by Hardware Thread		CPI Rate	Filled Pipeline Slots		Unfilled Pipeline Slots (Stalls)		Module	Function (Full Name)
	CPU_CLK_THREAD	INST_RETIRE_ANY		Retired Pipeline Slots by Assist	Cancelled Pipeline Slots	Back-end Bound Pipeline Slots	Front-end Bound Pipeline Slots		
sphere_intersect	14,284,000,000	17,624,000,000	0.910	0.322	0.003	0.573	0.033	analyze_locks.exe	sphere_intersect
grid_intersect	5,214,000,000	4,068,000,000	1.282	0.265	0.132	0.498	0.104	analyze_locks.exe	grid_intersect
grid_bounds_intersect	1,124,000,000	976,000,000	1.152	0.330	0.046	0.548	0.095	analyze_locks.exe	grid_bounds_intersect
tbodytask_scheduler_initialize	830,000,000	226,000,000	3.673	0.246	-0.019	0.325	0.440	tbody.dll	tbodytask_scheduler_initialize(mt)
GdipCreateSolidFill	386,000,000	526,000,000	0.720	0.611	-0.249	0.503	0.135	gdipplus.dll	GdipCreateSolidFill
[rdpdd.dll]	256,000,000	564,000,000	0.454	0.469	0.078	0.203	0.250	rdpdd.dll	[rdpdd.dll]
pos2grid	232,000,000	182,000,000	1.275	0.397	-0.086	0.655	0.034	analyze_locks.exe	pos2grid
tri_intersect	160,000,000	170,000,000	0.941	0.625	-0.125	0.475	0.025	analyze_locks.exe	tri_intersect
shader	154,000,000	151,000,000	1.000	0.390	-0.208	0.792	0.026	analyze_locks.exe	shader(struct ray *)
Raypnt	146,000,000	184,000,000	0.514	0.603	-0.247	0.644	0.000	analyze_locks.exe	Raypnt(struct ray *, double)
rdowd.svs1	116,000,000	224,000,000	0.496	0.345	0.034	0.517	0.103	rdowd.svs	rdowd.svs1

Selected 1 row(s): 14,284,000,000 CPU_CLK_THREAD, 17,624,000,000 INST_RETIRE_ANY

All collected data is presented in hierarchical format (see next slide), with helpful metrics already calculated (see issue slides).



Enhanced General Exploration View for Intel® Microarchitecture Codename Sandy Bridge

General Exploration - General Exploration

Analysis Target | Analysis Type | Collection Log | Summary | Bottom-Up

Grouping: Function

Function	Hardware Event ...			Filled Pipeline Slots		Unfilled Pipeline Slots (Stalls)			
	CPU_... THRE...	INS. AN.	CPI Rate	Retired Pipeline Slots by Assists	Cancelled Pipeline Slots	Back-end Bound Pipeline Slots	Front-end Bound Pipeline Slots		
sphere_intersect	14,284,000, ...	17, ...	0.810	0.322	0.068	0.571	0.059		
grid_intersect	5,214,000,000	4.0 ...	1.282	0.265	0.132	0.498	0.104		
grid_bounds_intersect	1,124,000,000	976 ..	1.152	0.320	0.046	0.548	0.085		
tbb:task_scheduler_init:initialize	830,000,000	226 ..	3.673	0.246	-0.019	0.325	0.448		
GdipCreateSolidFill	386,000,000	536 ..	0.720	0.611	-0.200	0.502	0.125		
[rdpdd.dll]	256,000,000	564 ..	0.454	0.469	0.0	0.0	0.0		
pos2grid	232,000,000	182 ..	1.275	0.397	-0.0	0.0	0.0		
tri_intersect	160,000,000	170 ..	0.941	0.625	-0.1	0.0	0.0		
shader	154,000,000	154 ..	1.000	0.390	-0.2	0.0	0.0		
Raypnt	146,000,000	284 ..	0.514	0.603	-0.2	0.0	0.0		
lrdpdd.svs	116,000,000	234 ..	0.496	0.345	0.0	0.571	0.000		
Selected 1 row(s):								0.000	
						0.498	0.000	0.000	0.000
						0.548	0.000	0.000	0.000
						0.325	0.000	0.000	0.000
						0.503	0.003	0.041	0.000
						0.203	0.000	0.000	0.000

Unfilled Pipeline Slots (Stalls) breakdown:

Back-end Bound Pipeline Slots	ICache Misses	ITLB Overhead	DSB to MITE Switch Cost
0.571	0.000	0.000	0.000
0.498	0.000	0.000	0.000
0.548	0.000	0.000	0.000
0.325	0.000	0.000	0.000
0.503	0.003	0.041	0.000
0.203	0.000	0.000	0.000

Hierarchical data display corresponds to how available execution slots in each core's pipeline are utilized.

Expand a column to see a breakdown of issues pertaining to its category of pipeline utilization: Retired, Cancelled, Back-end Bound, or Front-end Bound Pipeline Slots

Enhanced General Exploration View for Intel® Microarchitecture Codename Sandy Bridge

General Exploration - General Exploration

Analysis Target Analysis Type Collection Log Summary Bottom-up

Grouping: Function

Function	Hardware Event ...			Filled Pipeline Slots		
	CPU_... THRE...	INS. AN.	CPI Rate	Retired Pipeline Slots by Assists	Cancelled Pipeline Slots	Machine Clears
sphere_intersect	14,284,000, ...	17, ..	0.810	0.322	0.000	0.006
grid_intersect	5,214,000,000	4,0 ..	1.282	0.265	0.331	0.015
grid_bounds_intersect	1,124,000,000	076 ..	1.152	0.320	0.228	0.000
tbb:task_scheduler_init:initialize	830,000,000	226 ..	3.673	0.246	0.000	0.000
GdipCreateSolidFill	386,000,000	536 ..	0.720	0.611	0.000	0.000
[rdpdd.dll]	256,000,000	564 ..	0.454	0.469	0.250	0.000
pos2grid	232,000,000	182 ..	1.275			
tri_intersect	160,000,000	170 ..	0.941			
shade	154,000,000	154 ..	1.000			

Pre-computed metrics for each category of pipeline utilization saves users analysis time.

For a given hotspot, if a cell is highlighted pink, it means the value for that metric is over VTune's pre-determined threshold and should be investigated.

Software and Services Group



12

Note that issue highlighting occurs under 2 conditions:

1. The value for the metric is over VTune's pre-determined threshold
2. The associated function uses 5% or greater of the CPU clockticks sampled

The Old Way vs. The New Way

General Exploration - Hardware Event Counts

Analysis Target Analysis Type Summary Bottom-up

Grouping: Function

Function	INST_RETIR... ANY by Package	CPU_CLK_UN... THREAD by Package	CPU_CLK_UN... REF_TSC by Package	BR_MISP_RET... ALL_BRANC... by Package	OFFCORE_R... ANY_REQUE... LLC_MISS_L... DRAM_1 by
sphere_intersect	17,624,000,000	14,284,000,000	12,210,000,000	0	
grid_intersect	4,068,000,000	5,214,000,000	4,404,000,000	86,400,000	
grid_bounds_intersect	976,000,000	1,124,000,000	990,000,000	12,800,000	
[rdpdd.dll]	564,000,000	256,000,000	220,000,000	3,200,000	
GdipCreateSolidFill	536,000,000	386,000,000			
Raypnt	284,000,000	146,000,000			

The Old Way: To see if there is an issue with branch misprediction, multiply event value (86,400,000) by 20 cycles, then divide by CPU_CLK_UNHALTED.THREAD (5,214,000,000). Then compare the resulting value to a threshold. If it is too high, investigate.

General Exploration - General Exploration

Analysis Target Analysis Type Summary Bottom-up

Grouping: Function

Function	Filled Pipeline Slots		
	Retired Pipeline Slots by Assists	Cancelled Pipeline Slots Branch Mispredict	Machine Clears
sphere_intersect	0.322	0.000	
grid_intersect	0.265	0.331	0.0
grid_bounds_intersect	0.320	0.228	0.0
tbb::task_scheduler_init::initialize	0.246	0.000	0.0
GdipCreateSolidFill	0.611	0.000	0.0

The New Way: Look at the Branch Mispredict metric, and see if any cells are pink. If so, investigate.

Software and Services Group



Complexities of Performance Measurement

- Two features of the 2nd generation Intel® Core™ processor family have a significant effect on performance measurement:
 - Intel® Hyper-Threading Technology
 - Intel® Turbo Boost 2.0 Technology
- With these features enabled, it is more complex to measure and interpret performance data
 - Most events are counted per thread, some events per core
 - See VTune Amplifier XE Help for specific events
- Some experts prefer to analyze performance with these features disabled, then re-enable them once optimizations are complete

Software and Services Group



14

Both Intel® Hyper-Threading Technology and Intel® Turbo Boost 2.0 Technology can be enabled or disabled through BIOS on most platforms.

Contact with the system vendor or manufacturer for the specifics of any platform before attempting this. Incorrectly modifying bios settings from those supplied by the manufacturer can result in rendering the system completely unusable and may void the warranty.

Don't forget to re-enable these features once you are through with the software optimization process!

The "Software on Hardware" Tuning Process

1. Identify Hotspots

- Determine efficiency of hotspots
 - > If inefficient, identify architectural reason for inefficiency

2. Optimize the issue

3. Repeat from step 1!

Software and Services Group



15

Note: While VTune Amplifier XE's Concurrency, Timeline and Locks and Waits features can also be helpful in threading an application, this slideset is not aimed at the process of introducing threads.

The process described here could be used either before or after threading.

However, we *do* recommend that you follow a top-down process when optimizing: beginning with system tuning (if appropriate), then algorithmic tuning, then microarchitectural tuning. The name of Software on Hardware tuning just means we are tuning software for specific hardware.

Remember for all upcoming slides – that you should only focus on hotspots! Only try to determine efficiency, identify causes, and optimize in hotspots!

Step 1) Identify the Hotspots

- **What:** Hotspots are where your application spends the most time
- **Why:** You should aim your optimization efforts there!
 - >Why improve a function that only takes 2% of your application's runtime?
- **How:** VTune Amplifier XE *Hotspots* or *Lightweight Hotspots* analysis type
 - >Usually hotspots are defined in terms of the CPU_CLK_UNHALTED.THREAD event (aka "clockticks")

For the 2nd generation Intel® Core™ processor family, the CPU_CLK_UNHALTED.THREAD counter measures unhalting clockticks on a per thread basis. So for each tick of the CPU's clock, the counter will count 2 ticks if Hyper-Threading is enabled, 1 tick if Hyper-Threading is disabled. There is no per-core clocktick counter.

There is also a CPU_CLK_UNHALTED.REF counter, which counts unhalting clockticks per thread, at the reference frequency for the CPU. In other words, the CPU_CLK_UNHALTED.REF counter should not increase or decrease as a result of frequency changes due to Turbo Mode 2.0 or Speedstep Technology.

Step 1) Determine Efficiency

- Determine efficiency of the hotspot using one of three methods:
 - % Pipeline Slots Retired / Cycle
 - Changes in CPI
 - Code Examination
- Note: Some of these methods are more appropriate for certain codes than others... see notes on the following slides

% Pipeline Slots Retired and Changes in CPI methods rely on VTune Amplifier XE's event-based sampling. The Code Examination method relies on using VTune Amplifier XE's capability as a source/disassembly viewer.

Efficiency Method 1: % Retired Pipeline Slots / Cycle

- **Why:** Helps you understand how efficiently your app is using the processors
- **How:** General Exploration profile, Metric: *Retired Pipeline Slots*
- **What Now:**
 - For a given hotspot:
 - In general, > 90% retiring (.9 or higher) is good. Go to efficiency method 3.
 - 50 - 90% for client apps - Consider investigating stall reduction using the following issue slides.
 - < 60% for server apps – consider investigating stall reduction.

Function	CPU_CLK_THREAD	INST_RETIRE_ANY	CPI Rate	Retired Pipeline Slots	Cancelled Pipeline Slots
glibcrt_initialize	14,284,000,000	17,624,000,000	1.23	0.22	
grid_intersect	5,214,000,000	4,068,000,000	0.78	0.265	
grid_bound_i_intersect	1,124,000,000	976,000,000	0.87	0.320	
fbbrtask_scheduler_init_initialize	830,000,000	226,000,000	0.27	0.246	
GdipCreateSolidFill	386,000,000	536,000,000	1.39	0.611	
[rdpdd.dll]	256,000,000	564,000,000	2.20	0.463	
pos2grid	232,000,000	182,000,000	0.79	0.39	
tri_intersect	160,000,000	170,000,000	1.06	0.55	
shader	154,000,000	154,000,000	1.00	0.390	
Rayprint	146,000,000	284,000,000	1.94	0.603	
[indowd.vv1]	116,000,000	234,000,000	2.01	0.345	
Selected 1 row(s):		14,284,000,000	17,624,000,000		

Software and Services Group



18

Formula:

$$\left(\frac{\text{UOPS_RETIRED.RETIRE_SLOTS}}{4 * \text{CPU_CLK_UNHALTED.THREAD}} \right)$$

Thresholds: Investigate if -
% Retiring < .9

This metric is based on the fact that when operating at peak performance, the pipeline on a 2nd generation Core CPU should be able to retire 4 micro-operations per clock cycle (or “clocktick”). The formula looks at “slots” in the pipeline for each core, and sees if the slots are filled, and if so, whether they contained a micro-op that retired.

The thresholds are general guidelines. Depending on the domain, some applications can run with less slots allocated retiring than the thresholds above and still be very efficient. For example, it is common for database workloads to be running with only 20-25% of allocated slots retiring per clocktick (due to heavy I/O).

Efficiency Method 2: Changes in Cycles per Instruction (CPI)

- Why:** Another measure of efficiency that can be useful when comparing 2 runs
 - >Shows average time it takes one of your workload's instructions to execute
- How:** General Exploration profile, Metric: *CPI*
- What Now:**
 - >CPI can vary widely depending on the application and platform!
 - >If code size *stays constant*, optimizations should focus on reducing CPI

Function	CPU_CLK_UNHALTED.THREAD	INST_RETIRED.ANY	CPI Rate	Retired Pipeline Slots	Cancelled Pipeline Slots
isphere_intersect	14,284,000,000	17,624,000,000	0.810	0.322	
gnd_intersect	5,214,000,000	4,068,000,000	1.282	0.265	
gnd_bounds_intersect	1,124,000,000	976,000,000	1.152	0.320	
tbb:task_scheduler_init:initialize	830,000,000	226,000,000	3.673	0.246	
GdipCreateSolidFill	386,000,000	536,000,000	0.720	0.611	
[rdpdd.dll]	256,000,000	564,000,000	0.454	0.469	
pos2gnd	232,000,000	182,000,000	1.275	0.397	
tri_intersect	160,000,000	170,000,000	0.941	0.625	
shader	154,000,000	154,000,000	1.000	0.390	
Raypnt	146,000,000	284,000,000	0.514	0.603	
[rdowd.vsx]	116,000,000	224,000,000	0.496	0.345	
Selected 1 row(s):					

Formula:

$CPU_CLK_UNHALTED.THREAD / INST_RETIRED.ANY$

Threshold:

In the interface, CPI will be highlighted if > 1 . This is a very general rule based on the fact that some very well tuned apps achieve CPIs of 1 or below. However, many apps will naturally have a CPI of over 1 – it is very dependent on workload and platform. It is best used as a comparison factor – know your app's CPI and see if over time it is moving upward (that is bad) or reducing (good!).

Note that CPI is a ratio! Cycles per instruction. So if the code size changes for a binary, CPI will change. In general, if CPI reduces as a result of optimizations, that is good, and if it increases, that is bad. However there are exceptions! Some code can have a very low CPI but still be inefficient because more instructions are executed than are needed. This problem is discussed using the Code Examination method for determining efficiency.

Another Note: CPI will be doubled if using Intel® Hyper-threading. With Intel® Hyper-Threading enabled, there are actually 2 different definitions of CPI. We call them "Per Thread CPI" and "Per Core CPI". The Per Thread CPI will be twice the Per Core CPI. Only convert between per Thread and per Core CPI when viewing aggregate CPIs (summed for all logical threads).

Note: Optimized code (i.e: SSE instructions) may actually lower the CPI, and increase stall % – but it will increase the performance. CPI is just a general efficiency metric – the real measure of efficiency is work taking less time.

Efficiency Method 3: Code Examination

- **Why:** Methods 1 and 2 measure how long it takes instructions to execute. The other type of inefficiency is executing too many instructions.
- **How:** Use VTune Amplifier XE's capability as a source and disassembly viewer
- **What Now:**
 - Failure to utilize modern instructions results in larger code size
 - See next 2 slides for potential issues

The screenshot displays the Intel VTune Amplifier XE 2011 interface. The main window shows a disassembly view of assembly code. The left pane shows the source code with line numbers 31 through 51. The right pane shows the corresponding assembly instructions with their addresses and performance metrics. The performance metrics include CPU rate by CPU_CLK_UNHALTED, CPU_CLK_THREAD, and INSTRUCTION. The assembly instructions are grouped into blocks (Block 7, Block 8, Block 9, Block 10, Block 11). The instructions are: `mov dword ptr [ebp+0x14], 4`, `mov eax, dword ptr [ebp+0x4]`, `cmp eax, dword ptr [ebp+0x4]`, `jne [ebp+0x4]-Block 8>`, `mov dword ptr [ebp+0x20], 0`, `mov eax, dword ptr [ebp+0x4]`, `add eax, 0x1`, `mov dword ptr [ebp+0x20], 4`, `mov eax, dword ptr [ebp+0x4]`, `cmp eax, dword ptr [ebp+0x4]`, `jne [ebp+0x4]-Block 10>`, `mov eax, dword ptr [ebp+0x4]`, `add eax, 0x1`, `mov eax, dword ptr [ebp+0x4]`, `add eax, 0x1`, `mov eax, dword ptr [ebp+0x4]`.

Software and Services Group



20

This method involves looking at the disassembly to make sure the most efficient instruction streams are generated. This can be complex and can require an expert knowledge of the Intel instruction set and compiler technology. What we have done is describe how to find 2 easy-to-detect issues and suggest how they may be fixed using new features of Intel® Microarchitecture Codename Sandy Bridge.

Code Study 1: Convert Legacy Floating Point or Integer SSE Code to Intel® Advanced Vector Extensions (AVX)

- **Why:** Using SIMD instructions can greatly increase floating point performance. For existing FP or Integer SSE code, converting to AVX instructions has several advantages, including support for wider vector data (up to 256-bit), 3- and 4-operand syntax that allows NDS operations, and power savings.
- **How:** Examine your assembly code for existing SSE instructions (using xmm registers), MMX instructions (using mmx registers), or for floating point instructions that are not packed (such as faddp, fmul, or scalar SSE instructions like addss)
- **What Now:**
 - Convert Intel SSE code to 128-bit AVX instructions automatically using the Intel Compiler /QxAVX (Windows*) or -xAVX (Linux*) switches
 - Optimize to AVX – See the [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#), chapter 11

Software and Services Group



21

For more on AVX, see: <http://software.intel.com/en-us/articles/intel-avx-new-frontiers-in-performance-improvements-and-energy-efficiency/>

SSE instructions will look like: addps xmm4, xmm5.

+ addss is a **s**(calar) Intel® SSE instruction – which is better than x87 instructions– but is not as good as packed SSE instructions.

Code Study 2: Take Advantage of Improvements in Intel® Advanced Encryption Standard (AES) Instructions

- **Why:** Existing AES instruction throughput has been improved on Sandy Bridge microarchitecture, which can result in significant performance increases in parallel encryption/decryption.
- **How:** If the application's functionality is in the domain of encryption/decryption, check to see if AES instructions are being used. Blocks of aes instructions with xmm registers as operands may be using parallel modes:

```
aesenc  %xmm5, %xmm7
aesenc  %xmm5, %xmm8
aesenc  %xmm5, %xmm9
aesenc  %xmm5, %xmm10
```

- **What Now:**
 - If AES is being used in parallel modes (such as ECB, CTR, and CBC-Decrypt), increase performance by redefining the number of blocks to be processed in parallel. (8 on Sandy Bridge compared to 4 on Westmere).
 - If AES is not being used and the application does any encryption or decryption, try it! See the [Intel® Advanced Encryption Standard \(AES\) New Instructions Set](#)

Step 1) Identify architectural reason for inefficiency

- If Methods 1 or 2 are used to determine code is inefficient, investigate potential issues *in rough order of likelihood*

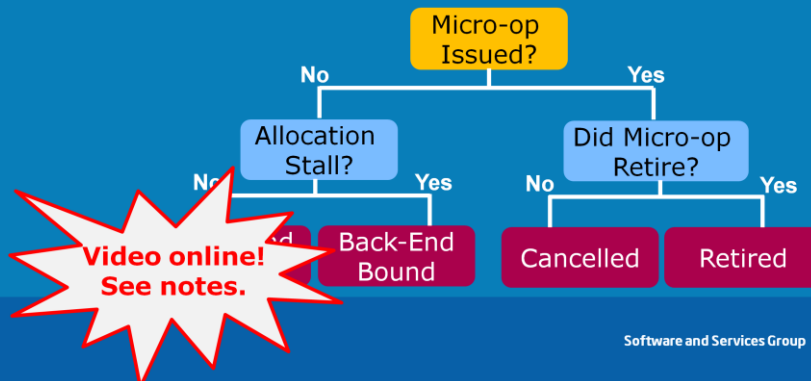
- **Cache Misses** Back-End Bound
 - Contested Accesses
 - Other Data Access Issues
 - Blocked Loads, Cache Line Splits, 4K Aliasing Conflicts, DTLB Misses
 - Allocation Stalls
- **Microcode Assists** Retired
- **Branch Mispredicts, Machine Clears** Cancelled
- **Front End Stalls** Front-End Bound

These are issues that result in inefficient pipeline use and high CPI. In addition to being in rough order of likelihood, these issues have been classified into the 4 categories of pipeline slot usage identified in the [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#).

The General Exploration profile also groups metrics according to these 4 categories. If desired, you can see how your application used the available pipeline slots per cycle using the first 4 metrics to the right of CPI: Retired Pipeline Slots, Cancelled Pipeline Slots, Back-End Bound Pipeline Slots, and Front-End Pipeline Slots.

Issue Classification

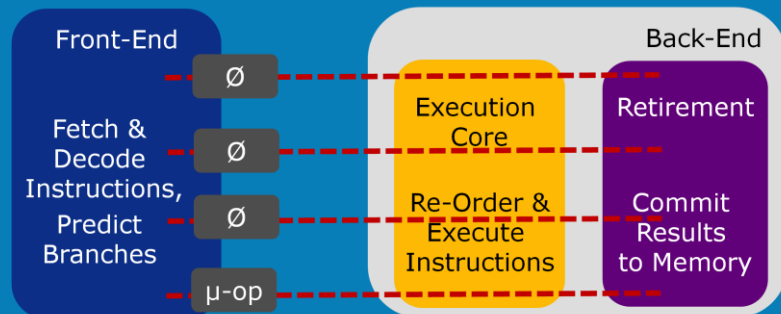
- The General Exploration View for Intel® Microarchitecture Codename Sandy Bridge uses the methodology outlined in the Optimization Reference Manual to analyze performance by studying “Slots” in a core’s execution pipeline
- Performance issues are classified according to what happened for each possible slot in the pipeline, per cycle:



Note that the way this methodology allows us to classify what percentage of all pipeline slots end up in each category, for each cycle and for each core. It is possible that for a given dataset, there may be a significant percentage of pipeline slots in multiple categories that merit investigation. Ideally a large percentage of slots will fall into the “Retired” category, but even then, it may be possible to make your code more efficient.

For a complete description of this methodology, see the [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#), Appendix B.3. You can also view our 10-minute video, which describes the methodology in more detail, here: <http://software.intel.com/en-us/videos/channel/parallel-programming/performance-analysis-methodology-for-intel-microarchitecture-codename-sandy-bridge/1265132823001>.

The Pipeline Slot Methodology, Illustrated



Case 1: Front-End does not provide micro-operations for all 4 pipeline slots

Front-End Bound

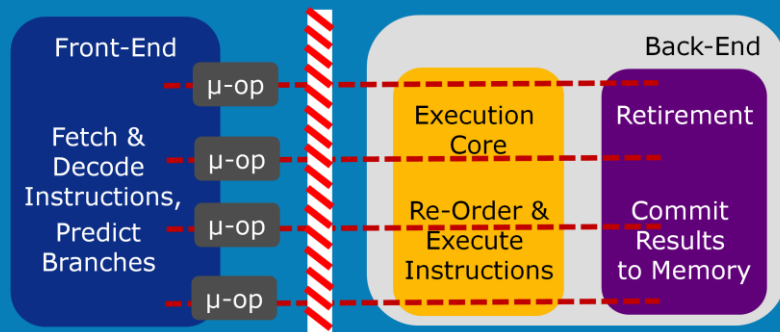
Software and Services Group



25

The Front-End consists of several different structures. It is responsible for fetching instructions, decoding them into micro-operations, and then delivering those micro-operations to the Back-End of the pipeline. For Intel® Microarchitecture Codename Sandy Bridge, a maximum of 4 micro-operations can be delivered to the Back-End portion of the pipeline per cycle (per core). Front-End issues are generally caused by delays in fetching code (due to caching or ITLB issues) or in decoding instructions (due to specific instruction types or queuing issues). Front-End issues are generally resolved by compiler techniques.

The Pipeline Slot Methodology, Illustrated



Case 2: Back-End cannot accept micro-operations for all 4 pipeline slots

Back-End Bound

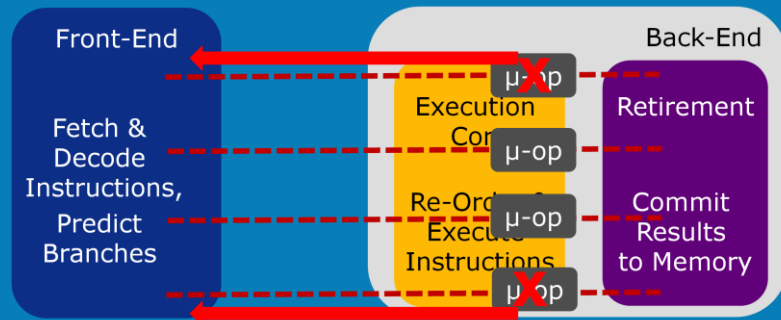
Software and Services Group



26

The Back-End of the pipeline is responsible for accepting micro-operations from the Front-End, then re-ordering them as necessary to schedule their execution in the various execution units, retrieving needed operands from memory, executing the operations, then committing the results to memory. If the Back-End is not able to accept micro-operations from the Front-End, it is generally because internal queues are full. Most of the time this is due to data access issues – the Back-End's structures are being taken up by micro-operations that are waiting on data from the caches.

The Pipeline Slot Methodology, Illustrated



Case 3: Micro-operations make it to the Back-End, but then get removed from the pipeline

Cancelled

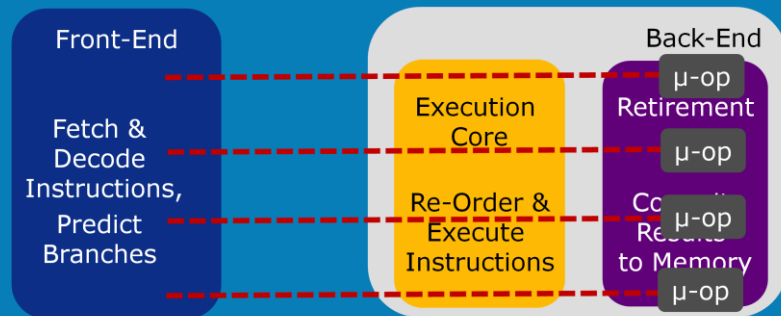
Software and Services Group



27

Cancelled micro-operations most likely happen because the Front-End mis-predicted a branch. This is discovered in the Back-End when the branch operation is executed. At this point, if the target of the branch was incorrectly predicted, the micro-operation and all subsequent incorrectly predicted operations are cancelled and the Front-End is re-directed to begin fetching instructions from the correct target.

The Pipeline Slot Methodology, Illustrated



Case 4: Micro-operations make it to the Back-End, Execute, and then Retire

Retired

Software and Services Group



28

In general, having as many pipeline slots retiring per cycle as possible is the goal. Only one issue is identified for this category – which deals with how to get micro-operations to this stage faster.

Step 2) Optimize the Issue

- For each potential issue, there are several important pieces of information:
 - Why? – Why you should be concerned about this potential problem.
 - How? – Which profile and metric to use in the Amplifier XE interface. If the data is highlighted, then it should be investigated.
 - What Now? – Helps you move to Step 2 of the Tuning Process (Optimize the Issue). Gives suggestions for follow-up investigations or optimizations to try.
 - Event Names and Metric Formulas are given in the Notes. These are not included on the slide because they are already embedded in the Amplifier XE logic. You only need to use the pre-configured profiles and metrics pointed out in order to know if you may have a problem.

Cache Misses

- **Why:** Cache misses raise the CPI of an application
 - Focus on long-latency data accesses coming from 2nd and 3rd level misses
- **How:** General Exploration profile, Metrics: *LLC Hit*, *LLC Miss*
- **What Now:**
 - If either metric is highlighted for your hotspot, consider reducing misses:
 - Use the cacheline replacement analysis outlined in the [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#), section **B.3.4.2**
 - Use software prefetch instructions
 - Block data accesses to fit into cache
 - Use local variables for threads
 - Pad data structures to cacheline boundaries
 - Change your algorithm to reduce data storage

Software and Services Group



30

Formulas:

% of cycles spent on memory access (LLC misses):

$$(\text{MEM_LOAD_UOPS_MISC_RETIRED.LLC_MISS_PS} * 180) / \text{CPU_CLK_UNHALTED.THREAD}$$

% of cycles spent on last level cache access (2nd level misses that hit in LLC):

$$((\text{MEM_LOAD_RETIRED.L3_HIT_PS} * 26) + (\text{MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HIT_PS} * 43) + (\text{MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM_PS} * 60)) / \text{CPU_CLK_UNHALTED.THREAD}$$

Thresholds: Investigate if –

% cycles for LLC miss \geq .2,

% cycles for LLC Hit \geq .2

Doing the cacheline replacement analysis study (the first suggestion) can be very helpful in resolving these issues. The approach is to figure out at what level most of the application's loads are being satisfied, then look one level above. This is because if loads are coming from a particular level of cache or memory, it is because the data has been replaced at the level above (for example, if loads are coming from memory, the data has been replaced in the LLC). The study shows how to identify the areas of your code that are causing the replacements to occur. Once these are identified, you can try changing the algorithm, doing non-temporal stores, or one of the other suggestions above.

Contested Accesses

- **Why:** Sharing modified data among cores can raise the latency of data access
- **How:** General Exploration profile, Metrics: *Contested Accesses*
- **What Now:**
 - If either metric is highlighted for your hotspot, locate the source code line(s) that is generating HITMs by viewing the source. Look for the MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM_PS event which will tag to the next instruction after the one that generated the HITM.
 - Then use knowledge of the code to determine if real or false sharing is taking place. Make appropriate fixes:
 - For real sharing, reduce sharing requirements
 - For false sharing, pad variables to cacheline boundaries

Formula:

% of cycles spent accessing data modified by another core:

$(\text{MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM_PS} * 60) / \text{CPU_CLK_UNHALTED.THREAD}$

Thresholds: Investigate if –

% cycles accessing modified data > .05

This metric is also called write sharing. It occurs when one core needs data that is found in a modified state in another core's cache. This causes the line to be invalidated in the holding core's cache and moved to the requesting core's cache. If it is written again and another core requests it, the process starts again. The cacheline ping pong-ing between caches causes longer access time than if it could be simply shared amongst cores (as with read-sharing).

Write sharing can be caused by true sharing, as with a lock or hot shared data structure, or by false sharing, meaning that the cores are modifying 2 separate pieces of data stored on the same cacheline.

Note that in the case of real write sharing that is caused by a lock, Amplifier XE's Locks and Waits analysis should also indicate a problem. This hardware-level analysis will detect other cases as well though (such as false sharing or write sharing a hot data structure).

Other Data Access Issues: Blocked Loads Due to No Store Forwarding

- **Why:** If it is not possible to forward the result of a store through the pipeline, dependent loads may be blocked
- **How:** General Exploration profile, Metric: *Loads Blocked by Store Forwarding*
- **What Now:**
 - If the metric is highlighted for your hotspot, investigate:
 - View source and look at the LD_BLOCKS_STORE_FORWARD event. Usually this event tags to next instruction after the attempted load that was blocked. Locate the load, then try to find the store that cannot forward, which is usually within the prior 10-15 instructions. The most common case is that the store is to a smaller memory space than the load. Fix the store by storing to the same size or larger space as the ensuing load.

Formula:

Blocked Store Forwarding Cost = $(LD_BLOCKS_STORE_FORWARD * 13) / CPU_CLK_UNHALTED.THREAD$

Threshold: Investigate if –
Cost \geq .05

Store forwarding occurs when there are two memory instructions in the pipeline, a store followed by a load from the same address. Instead of waiting for the data to be stored to the cache, it is "forwarded" back along the pipeline to the load instruction, saving a load from the cache. Store forwarding is the desired behavior, however, in certain cases, the store may not be able to be forwarded, so the load instruction becomes blocked waiting for the store to write to the cache and then to load it.

Other Data Access Issues: Cache Line Splits Back-End Bound

- **Why:** Multiple cache line splits can result in load penalties.
- **How:** General Exploration profile, Metric: *Split Loads, Split Stores*
- **What Now:**
 - If the metric is highlighted for your hotspot, investigate by viewing the metrics at the sourcecode level. The split load event, MEM_UOP_RETIRED.SPLIT_LOADS_PS, should tag to the next executed instruction after the one causing the split. If the split store ratio is greater than .01 at any source address, it is worth investigating.
 - To fix these issues, ensure your data is aligned. Especially watch out for mis-aligned 256-bit AVX store operations.

Software and Services Group



33

A cacheline split is any load or store that traverses a 64-byte boundary.

Formulas:

Split Load Cost = (MEM_UOP_RETIRED.SPLIT_LOADS_PS * 5) / CPU_CLK_UNHALTED.THREAD

Split Store Ratio = MEM_UOP_RETIRED.SPLIT_STORES_PS / MEM_UOP_RETIRED.ANY_STORES_PS

Thresholds: Investigate if –

Split load cost \geq .1 or

Split store ratio is $>$ 0.01

Beginning with the Intel® Core™ architecture, the penalty for cacheline splits has been reduced to only 5 cycles. However, if there are repeated splits occurring, the penalty can grow, and even just a 5-cycle increase in latency can make a difference in application performance.

Other Data Access Issues: 4K Aliasing

Back-End Bound

- **Why:** Aliasing conflicts result in having to re-issue loads.
- **How:** General Exploration profile, Metric: *4K Aliasing*
- **What Now:**
 - If this metric is highlighted for your hotspot, investigate at the sourcecode level.
 - Fix these issues by changing the alignment of the load. Try aligning data to 32 bytes, changing offsets between input and output buffers (if possible), or using 16-Byte memory accesses on memory that is not 32-Byte aligned.

Software and Services Group



34

Formula:

$$\text{Aliasing Conflicts Cost} = (\text{LD_BLOCKS_PARTIAL.ADDRESS_ALIAS} * 5) / \text{CPU_CLK_UNHALTED.THREAD}$$

Threshold: Investigate if

Aliasing conflicts cost \geq .1

This occurs when a load is issued after a store and their memory addresses are offset by (4K). When this is processed in the pipeline, the issue of the load will match the previous store (the full address is not used at this point), so pipeline will try to forward the results of the store and avoid doing the load (this is store forwarding). Later on when the address of the load is fully resolved, it will not match the store, and so the load will have to be re-issued from a later point in the pipe. This has a 5-cycle penalty in the normal case, but could be worse in certain situations, like with un-aligned loads that span 2 cache lines.

Other Data Access Issues: DTLB Misses

- **Why:** First-level DTLB Load misses (Hits in the STLB) incur a latency penalty. Second-level misses require a page walk that can affect your application's performance.
- **How:** General Exploration profile, Metric: *DTLB Overhead*
- **What Now:**
 - If this metric is highlighted for your hotspot, investigate at the sourcecode level.
 - To fix these issues, target data locality to TLB size, use the Extended Page Tables (EPT) on virtualized systems, try large pages (database/server apps only), increase data locality by using better memory allocation or Profile-Guided Optimization

Formula:

$$\text{DTLB Overhead} = ((\text{DTLB_LOAD_MISSES.STLB_HIT} * 7) + \text{DTLB_LOAD_MISSES.WALK_DURATION}) / \text{CPU_CLK_UNHALTED.THREAD}$$

Threshold: Investigate if-

DTLB Overhead \geq .1

On target data locality to TLB size: this is accomplished via data blocking and trying to minimize random access patterns.

Note: this is more likely to occur with server applications or applications with a large random dataset

Allocation Stalls

- **Why:** Certain types of instructions can cause allocation stalls because they take longer to retire. These increase latencies overall.
- **How:** General Exploration Profile, Metric: *LEA Stalls, Flags Merge Stalls*
- **What Now:**
 - If this metric is highlighted for your hotspot, investigate at the sourcecode level.
 - Try to eliminate uses of 3-operand LEA instructions, Look for certain uses of an LEA instruction (see section 3.5.1.3 of [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#)) or partial register use (see section 3.5.2.4 of [Intel® 64 and IA-32 Architectures Optimization Reference Manual](#)) and fix.

Formulas:

Flags Merge Stalls =

PARTIAL_RAT_STALLS.FLAGS_MERGE_UOP_CYCLES/CPU_CLK_UNHALTED.THREAD

LEA Stalls =

PARTIAL_RAT_STALLS.SLOW_LEA_WINDOW/CPU_CLK_UNHALTED.THREAD

Threshold: Investigate if-

Flags Merge Stalls > .05

LEA Stalls > .05

Long-latency instructions cause the Back-end to refuse instructions from the front-end (allocation stalls).

Microcode Assists

- **Why:** Assists from the microcode sequencer can have long latency penalties.
- **How:** General Exploration Profile, Metric: *Assists*
- **What Now:**
 - If this metric is highlighted for your hotspot, re-sample using the additional assist events to determine the cause.
 - If FP_ASSISTS.ANY / INST_RETIRED.ANY is significant, check for denormals. To fix enable FTZ and/or DAZ if using SSE/AVX instructions or scale your results up or down depending on the problem
 - If $((\text{OTHER_ASSISTS.AVX_TO_SSE_NP} \times 75) / \text{CPU_CLK_UNHALTED.THREAD})$ or $((\text{OTHER_ASSISTS.SSE_TO_AVX_NP} \times 75) / \text{CPU_CLK_UNHALTED.THREAD})$ is greater than .1, reduce transitions between SSE and AVX code

Formula:

Assist % = $\text{IDQ.MS_CYCLES} / \text{CPU_CLK_UNHALTED.THREAD}$

Threshold: Investigate if –

Assist Cost $\geq .05$

There are many instructions that can cause assists when there is no performance problem. If you see MS_CYCLES it doesn't necessarily mean there is an issue, but whenever you do see a significant amount of MS_CYCLES, check the other metrics to see if it's one of the problems we mention.

Branch Mispredicts

Cancelled

- **Why:** Mispredicted branches cause pipeline inefficiencies due to wasted work or instruction starvation (while waiting for new instructions to be fetched)
- **How:** General Exploration Profile, Metric: *Branch Mispredict*
- **What Now:**
 - If this metric is highlighted for your hotspot try to reduce misprediction impact:
 - Use compiler options or profile-guided optimization (PGO) to improve code generation
 - Apply hand-tuning by doing things like hoisting the most popular targets in branch statements.

Software and Services Group



38

Formula:

Mispredicted branch cost: $(20 * BR_MISP_RETIRED.ALL_BRANCHES_PS) / CPU_CLK_UNHALTED.THREAD$

Threshold: Investigate if -

Cost is $\geq .2$

Note that all applications will have some branch mispredicts - it is not the number of mispredicts that is the problem but the impact.

To do hand-tuning, you need to locate the branch causing the mispredicts. This can be difficult to track down due to the fact that this event will normally tag to the first instruction in the correct path that the branch takes.

Machine Clears

- **Why:** Machine clears cause the pipeline to be flushed and the store buffers emptied, resulting in a significant latency penalty.
- **How:** General Exploration Profile, Metric: *Machine Clears*
- **Now What:**
 - If this metric is highlighted for your hotspot try to determine the cause using the specific events:
 - If MACHINE_CLEAR.SMC is significant, investigate at the sourcecode level. This could be caused by 4K aliasing conflicts or contention on a lock (both previous issues).
 - If MACHINE_CLEAR.MEMORY_ORDERING is significant, the clears are being caused by self-modifying code, which should be avoided.

Formula:

Machine Clear cost: $((\text{MACHINE_CLEAR.SMC} + \text{MACHINE_CLEAR.MEMORY_ORDERING} + \text{MACHINE_CLEAR.MASKMOV}) * 100) / \text{CPU_CLK_UNHALTED.THREAD}$

Threshold: Investigate if -

Cost is $\geq .02$

Machine clears are generally caused by either contention on a lock, or failed memory disambiguation from 4k aliasing (both earlier issues). The other potential cause is self-modifying code (SMC).

Front-end Stalls

- **Why:** Front-end stalls (at the Issue stage of the pipeline) can cause instruction starvation, which may lead to stalls at the execute stage in the pipeline.
- **How:** General Exploration profile, Metric: *Front-end Bound Pipeline Slots*
- **What Now:**
 - If this metric is highlighted for your hotspot, try using better code layout and generation techniques:
 - Try using profile-guided optimizations (PGO) with your compiler
 - Use linker ordering techniques (/ORDER on Microsoft's linker or a linker script on gcc)
 - Use switches that reduce code size, such as /O1 or /Os

Formula:

$$\text{IDQ_UOPS_NOT_DELIVERED.CORE} / (\text{CPU_CLK_UNHALTED.THREAD} * 4)$$

Threshold: Investigate if –

Front-End Bound uOps \geq .15

Assists or excessive Branch Mispredicts, all on previous slides, could be the reason for front-end issues, so check for and resolve those problems first. This issue may also be caused by instruction cache misses (on server apps), which are generally fixed by better code layout.

Good Luck!

For more information:

VTune User Forums:

<http://software.intel.com/en-us/forums/intel-vtune-performance-analyzer/>

VTune Amplifier XE Videos:

<http://software.intel.com/en-us/articles/intel-vtune-amplifier-xe/>

Intel® 64 and IA-32 Architecture Software Developer's Manuals:

<http://www.intel.com/products/processor/manuals/index.htm>

Optimization Guide for Intel® Microarchitecture Codename Nehalem:

<http://software.intel.com/file/15529>

For optimization of the integrated graphics controller on
Intel® Microarchitecture Codename Sandy Bridge:

www.intel.com/software/gpa

Software and Services Group



41

