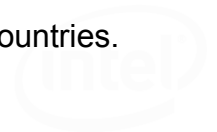


Performance Monitoring on Intel® Core™ i7 Processors*

Ramesh Peri

Principal Engineer & Engineering Manager
Profiling Collectors Group (PAT/DPD/SSG)
Intel® Corporation
Austin, TX 78746

* Intel, the Intel logo, Intel Core and Core Inside are trademarks of Intel Corporation in the U.S. and other countries.



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.

- All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Customers, licensees, and other third parties are not authorized by Intel to use Intel code names in advertising, promotion or marketing of any product or service.
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel [Performance Benchmark Limitations](#)
- Copyright © 2009, Intel Corporation. All rights reserved.



Risk Factors

The above statements and any others in this document that refer to plans and expectations for the first quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Many factors could affect Intel's actual results, and variances from Intel's current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the corporation's expectations. Current uncertainty in global economic conditions pose a risk to the overall economy as consumers and businesses may defer purchases in response to tighter credit and negative financial news, which could negatively affect product demand and other related matters. Consequently, demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including conditions in the credit market that could affect consumer confidence; customer acceptance of Intel's and competitors' products; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of new Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient supply of components from suppliers to meet demand. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; capacity utilization; excess or obsolete inventory; product mix and pricing; variations in inventory valuation, including variations related to the timing of qualifying products for sale; manufacturing yields; changes in unit costs; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. The recent financial crisis affecting the banking system and financial markets and the going concern threats to investment banks and other financial institutions have resulted in a tightening in the credit markets, a reduced level of liquidity in many financial markets, and extreme volatility in fixed income, credit and equity markets. There could be a number of follow-on effects from the credit crisis on Intel's business, including insolvency of key suppliers resulting in product delays; inability of customers to obtain credit to finance purchases of our products and/or customer insolvencies; counterparty failures negatively impacting our treasury operations; increased expense or inability to obtain short-term financing of Intel's operations from the issuance of commercial paper; and increased impairments from the inability of investee companies to obtain financing. Intel's results could be impacted by adverse economic, social, political and physical/infrastructure conditions in the countries in which Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports.



Agenda

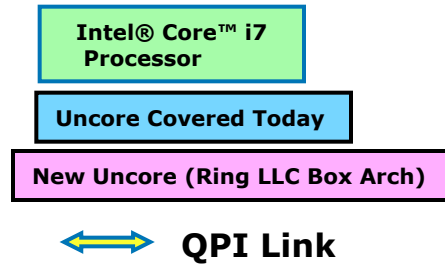
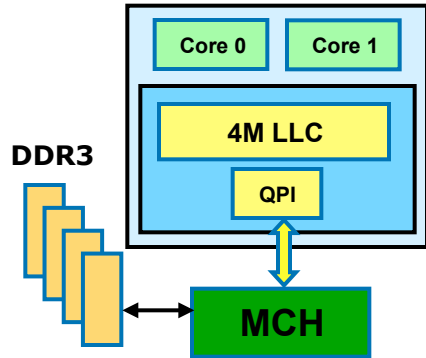
Goal: Provide a sample of the PMU features of Core™ i7

- **Core™ i7 Processor Architecture Overview**
- **Performance features of Core™ i7**
 - **Loop Stream Detector**
 - **Macro-Fusion**
 - **Memory Access**
 - **False Sharing**
 - **Load Latency Threshold**

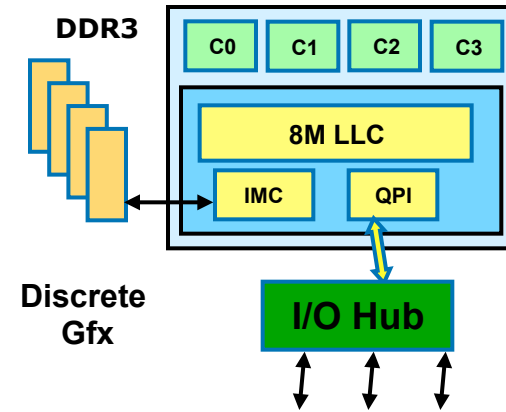


Platforms

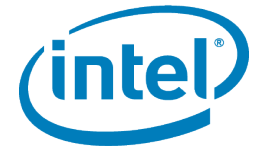
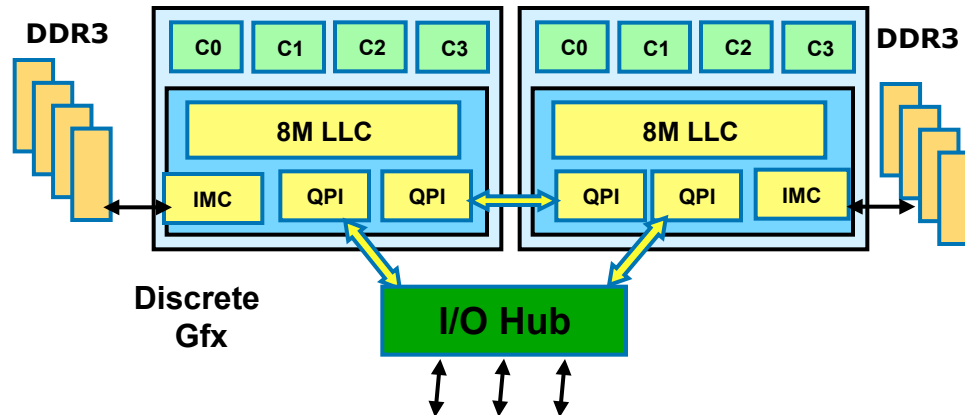
Mobile/Desktop



HE Desktop /mobile UP server & WS



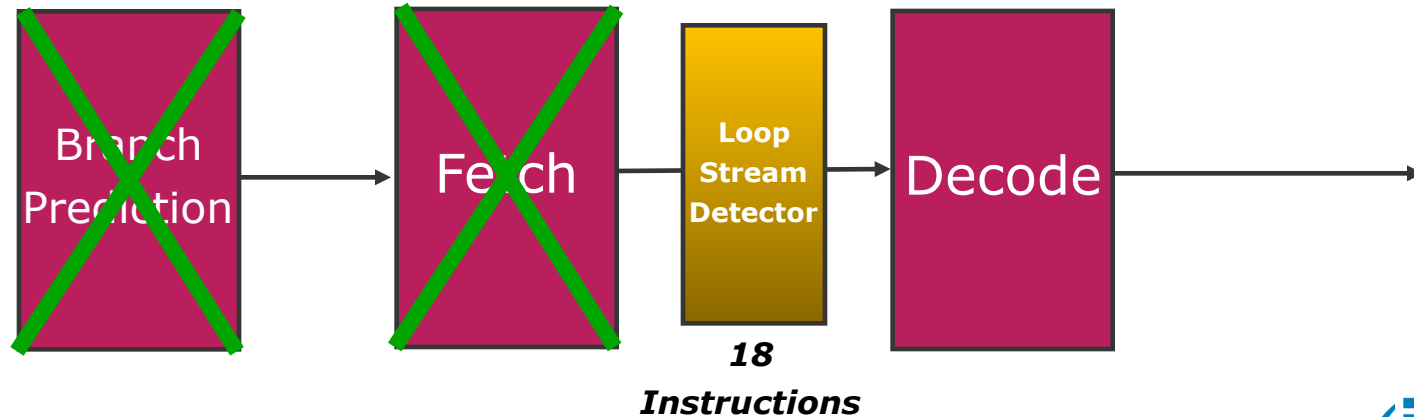
DP Servers & WS



Loop Stream Detector - Recap

- Loops are very common in most software
- Take advantage of knowledge of loops in HW
 - Decoding the same instructions over and over
 - Making the same branch predictions over and over
- Loop Stream Detector identifies software loops
 - Stream from Loop Stream Detector instead of normal path
 - Disable unneeded blocks of logic for **power savings**
 - **Higher performance** by removing instruction fetch limitations

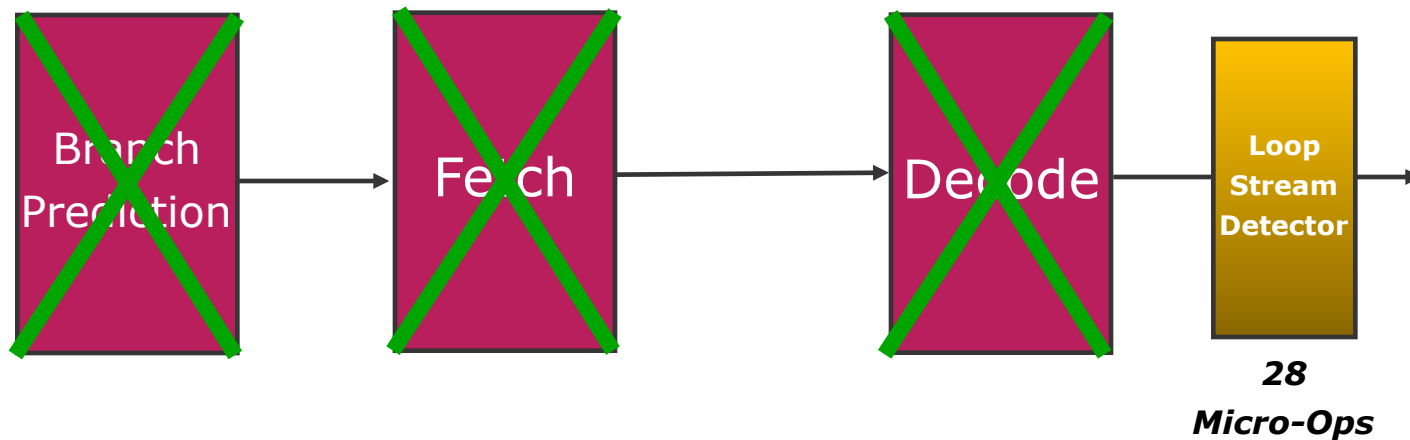
Core 2 Loop Stream Detector



Core™ i7 Loop Stream Detector

- Same concept as in prior implementations
- **Higher performance:** Expand the size of the loops detected
- **Improved power efficiency:** Disable even more logic

Nehalem Loop Stream Detector

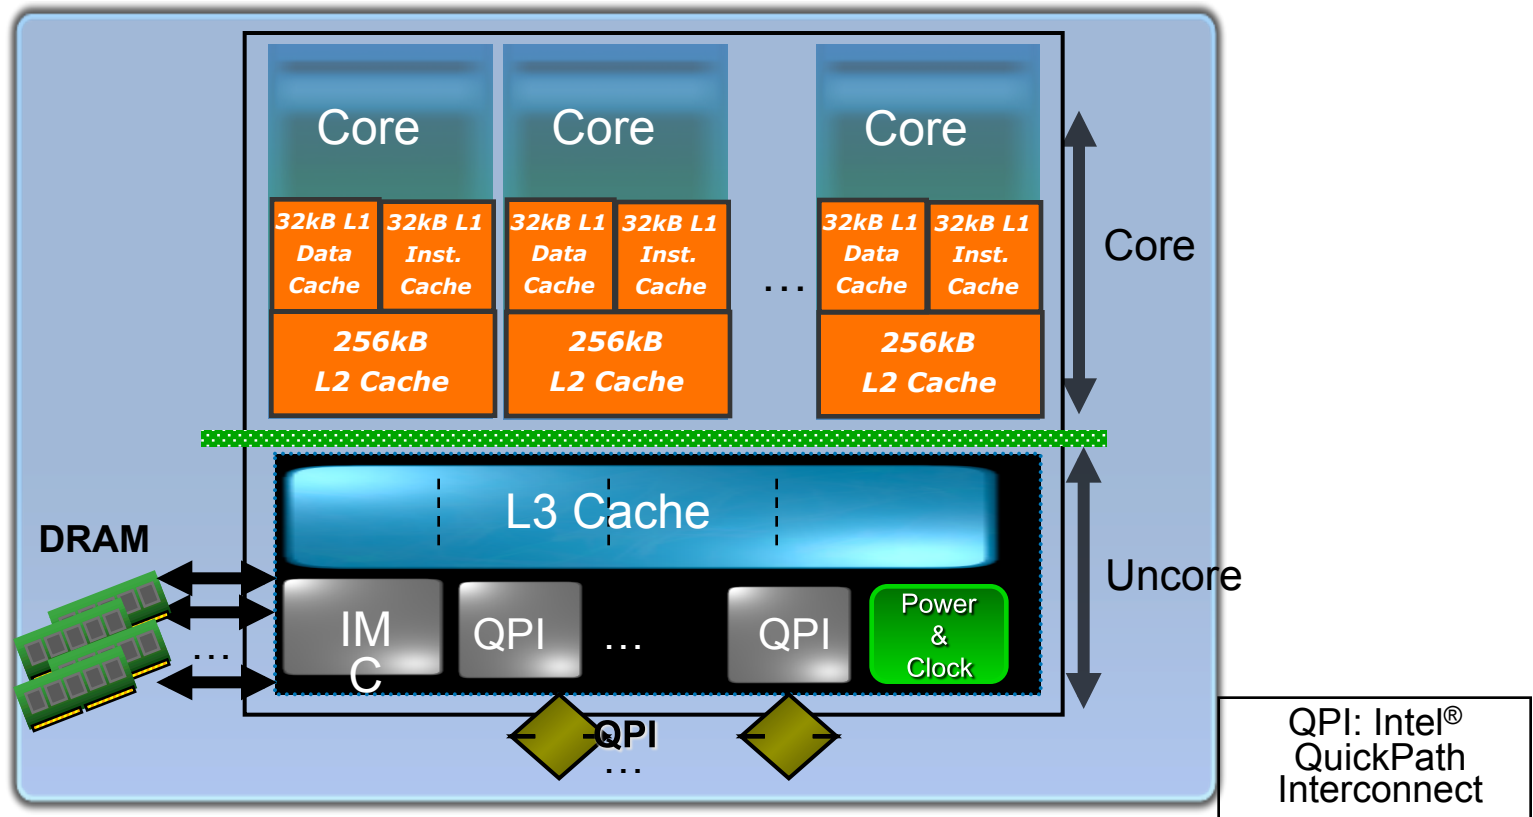


Loop Stream Detector – is it working ?

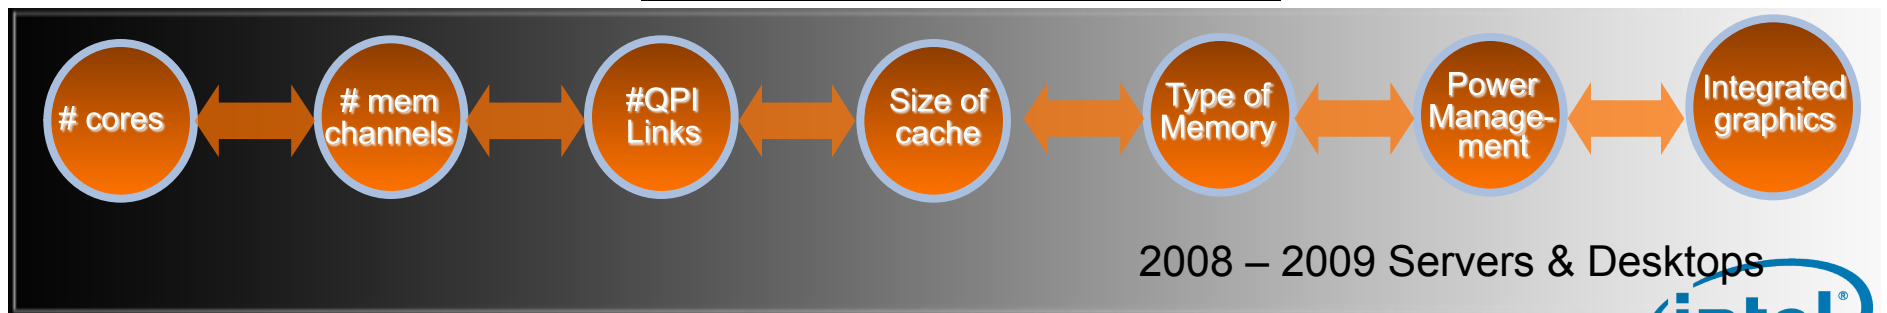
- LSD.UOPS is an event which provides the number of uops delivered by loop stream detector
- A tool which can tell
 - For every hot loop in the program whether the loop stream detector is active for that loop or not



Core™ i7



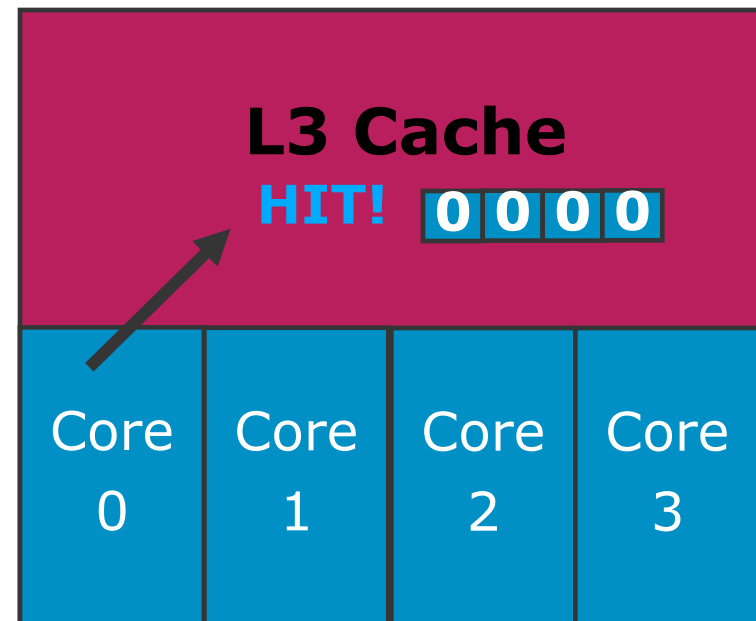
Differentiation in the "Uncore":



L3 Cache

- **L3 is inclusive with respect to L1 & L2**
 - **Provides good scalability**
 - Is implemented by maintaining a set of “core valid” bits per cache line in the L3 cache

Inclusive



Core valid bits limit unnecessary snoops

Some PMU events using this feature

- MEM_UNCORE_RETIRED.OTHER_CORE_L2_HITM
 - Load instructions retired that HIT “modified” data in sibling core
- MEM_LOAD_RETIRED.OTHER_CORE_L2_HIT_HITM
 - Load instructions retired that HIT “modified” or “unmodified” data in sibling core

These events can enable many capabilities like true/false sharing, race detection tools



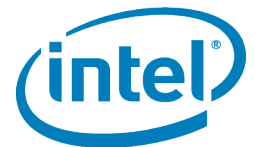
Some limitations of these events

- No Store HITM
 - Currently only Load HITM
- No ability to identify the source core id of the access
- No ability to raise PMI based on address range for data addresses
- PEBS EIP puts current EIP
 - which is the next instruction address
 - Requires some clever techniques and effort to figure out actual address
- HITM in the presence of SMT is not useful



Macro Fusion - Recap

- Introduced in Core™ 2 Microarchitecture
- TEST/CMP instruction followed by a conditional branch treated as a single instruction
 - Decode as one instruction
 - Execute as one instruction
 - Retire as one instruction
- Higher **performance**
 - Improves throughput
 - Reduces execution latency
- Improved **power efficiency**
 - Less processing required to accomplish the same work



Core™ i7 Macro Fusion

- Goal: Identify more macrofusion opportunities for increased *performance* and *power efficiency*
- Support all the cases in Core™ 2 Microarchitecture **PLUS**
 - CMP+Jcc macrofusion added for the following branch conditions
 - JL/JNGE
 - JGE/JNL
 - JLE/JNG
 - JG/JNLE
- Core™ only supports macrofusion in 32-bit mode
 - Core™ i7 supports macrofusion in both 32-bit and 64-bit modes



How effective is Macro Fusion ?

- MACHINE_CLEAR.ASSIST_FUSION
 - Counts the number of fusion assists
- Ensure that “all” appropriate CMP/Jxx sequences are all fused



False Sharing

What is it and why is it a Problem

- Cache coherency protocols require that all cores use the most current version of every cacheline
- Shared lines can be modified by any thread
 - Causing lines to be renewed regularly, if any thread writes to any byte in the line
 - (replace an invalid state copy with new valid copy)
 - Line renewal can cause a cache miss by other threads
 - and a 40-300 cycle execution stall
 - Depending on cacheline location
- False sharing is when different threads access non-overlapping regions of a cacheline

False Sharing Causes Avoidable 40-300 Cycle Stalls For Every Read Following a Write by Another Thread



Data Address Profiling and False Sharing Detection

Sampling during app execution

Symbolization & Data Address reconstruction

Aggregation

Precise Event Sampling:

events associated with memory operations, e.g.
MEM_INST_RETIRED.LOADS,
MEM_INST_RETIRED.STORES...

Pin threads affinity

Iterate over Samples and PEBS records in ebs.tb5

Using the binary, identify the instruction that overflowed event counter -> IP-1

Sample: IP, data address, threadID..
To aggregate addresses into cachelines:

Binary

...	...
0x7F5D	mulpd xmm1, xmm2
0x7F61	movaps xmm2, XMMWORD PTR [rsp+0230h]
0x7F69	mulpd xmm7, xmm2
0x7F6D	subpd xmm3, xmm1
0x7F71	movsd xmm1, MMWORD PTR [rcx+rbx+rho_i.0+018h]
0x7F7A	movhpd xmm1, MMWORD PTR [rcx+rbx+rho_i.0+020h]
0x7F83	mulpd xmm1, xmm8
...	...

IP-1

```
00000000055CC9900
00000000055CC9908
00000000055CC9910
00000000055CC9916
00000000055CC9928
00000000055CC9938
&&...FFFFC0
```

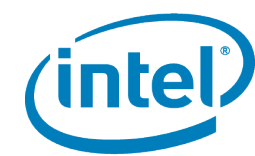
Sample record:
IP, process, module, threadID..

PEBS record:
IP, rax, rbx, rcx

ebs.tb5

Same cacheline accessed by different threads at different offsets
True and False Sharing
Next foils Illustrate GUI Navigation

Cacheline Address / Offset / Thread ...	Contributors	MEM...L1D_MISS
▼ 0x00000000055cc900	Offsets: 5 Threads: 3	31 (0.0%)
▼ Offset:0x00(0)	Threads: 1	21 (0.0%)
▶ Thread:00003fbb(0014)	Functions: 3	21 (0.0%)
▼ Offset:0x38(56)	Threads: 1	5 (0.0%)
▶ Thread:00003fbd(0009)	Functions: 3	5 (0.0%)
▼ Offset:0x08(8)	Threads: 2	1 (0.0%)
▶ Thread:00003fbb(0014)	Functions: 3	0 (0.0%)
▶ Thread:00003fbc(0015)	Functions: 1	1 (0.0%)
▼ Offset:0x10(16)	Threads: 1	3 (0.0%)
▶ Thread:00003fbc(0015)	Functions: 2	3 (0.0%)
▼ Offset:0x28(40)	Threads: 1	1 (0.0%)
▶ Thread:00003fbc(0015)	Functions: 1	1 (0.0%)



Synthetic Example: Heavy Contention on this Line -- Multiple Threads Accessing Different Offsets Indicate False Sharing (Identified by Rose Highlighting)

Intel(R) Performance Tuning Utility - 2007-12-15-08-33-27 - Eclipse Platform

File Edit Navigate Project Run Window Help

2007-12-15-08-22-51 2007-12-15-08-33-27

Function	Module	Collected Data Refs (%Total)	LLC Misses (%Total)	Avg. Latency	Total Latency (%Total)	Cachelines #	Pages # (%Total)	MEM_LOAD_RETIRED.L2_MISS (%Total)
sort	main_share.exe	8,594,000,000 (100.0%)	400,000 (100.0%)	3	26,186,000,000 (100.0%)	1,029	24 (85.7%)	400,000 (100.0)

Total Selected:

Granularity: Function Process: main_share.exe Thread: All Module: All Filter by selection

Experiment Summary Console Cachelines View Top by: Collected Data Refs

Cacheline Address / Offset / Thread / Function	Collected Data ...	LLC Misses (%T...	Avg. Latency	Total Latency (...	Contention (%...	MEM_LOAD_RE...	MEM_LOAD_RE...	INST_RETIRED...	Contributors
0x0042a3c0	1,959,600,000 ...	400,000 (100...	3	6,252,000,000 ...	909,100,000 (...	400,000 (100...	39,200,000 (8...	1,920,000,000 ...	Offsets: 2 Threads: 2
0x0064ff40	836,000,000 (...	0 (0.0%)	3	2,508,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	836,000,000 (...	Offsets: 1 Threads: 1
0x0054ff40	764,000,000 (...	0 (0.0%)	3	2,292,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	764,000,000 (...	Offsets: 1 Threads: 1
0x0054ff80	366,000,000 (...	0 (0.0%)	3	1,098,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	366,000,000 (...	Offsets: 2 Threads: 1
0x0064ff80	276,000,000 (...	0 (0.0%)	3	828,000,000 (...	0 (N/A)	0 (0.0%)	0 (0.0%)	276,000,000 (...	Offsets: 2 Threads: 1
0x004369c0	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 7 Threads: 1
0x0042e580	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
0x0042f380	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
0x004327c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 4 Threads: 1
0x00440900	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
0x0042e9c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
0x004396c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
0x004399c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
0x00440dc0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
0x00430280	10,000,000 (0...	0 (0.0%)	3	30,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	10,000,000 (0...	Offsets: 5 Threads: 1
0x0042f9c0	10,000,000 (0...	0 (0.0%)	3	30,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	10,000,000 (0...	Offsets: 5 Threads: 1

Total Selected:

2007-12-15-08-33-27 (Basic Data Access Profiling)



Expanding the "arrow" we see the 2 threads access the line at Different Offsets...This is False Sharing

The screenshot displays the Intel(R) Performance Tuning Utility interface. The top table shows the 'sort' function in 'main_share.exe' with 8,594,000,000 data references and 400,000 LLC misses. Below this, the 'Cachelines View' shows a detailed list of cache lines. The first cache line, at address 0x0042a3c0, is highlighted in red and circled. Its 'Contributors' column shows 'Offsets: 2 Threads: 2', indicating that two threads access this cache line at different offsets. This is a classic sign of false sharing.

Function	Module	Collected Data Refs (%Total)	LLC Misses (%Total)	Avg. Latency	Total Latency (%Total)	Cachelines #	Pages # (%Total)	MEM_LOAD_RETIRED.L2_MISS (%Total)
sort	main_share.exe	8,594,000,000 (100.0%)	400,000 (100.0%)	3	26,186,000,000 (100.0%)	1,029	24 (85.7%)	400,000 (100.0)

Cacheline Address / Offset / Thread / Function	Collected Data ...	LLC Misses (%T...	Avg. Latency	Total Latency (...	Contention (%...	MEM_LOAD_RE...	MEM_LOAD_RE...	INST_RETIRED...	Contributors
0x0042a3c0	1,959,600,000 ...	400,000 (100...	3	6,252,000,000 ...	909,100,000 (...	400,000 (100...	39,200,000 (8...	1,920,000,000 ...	Offsets: 2 Threads: 2
▶ Offset:0x04(4)	1,050,500,000 (...	100,000 (25.0%)	3	3,319,000,000 (...	0 (N/A)	100,000 (25.0%)	20,400,000 (46...	1,030,000,000 (...	Threads: 1
▶ Offset:0x00(0)	909,100,000 (1...	300,000 (75.0%)	3	2,933,000,000 (...	0 (N/A)	300,000 (75.0%)	18,800,000 (42...	890,000,000 (...	Threads: 1
▶ 0x0064ff40	836,000,000 (...	0 (0.0%)	3	2,508,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	836,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff40	764,000,000 (...	0 (0.0%)	3	2,292,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	764,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff80	366,000,000 (...	0 (0.0%)	3	1,098,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	366,000,000 (...	Offsets: 2 Threads: 1
▶ 0x0064ff80	276,000,000 (...	0 (0.0%)	3	828,000,000 (...	0 (N/A)	0 (0.0%)	0 (0.0%)	276,000,000 (...	Offsets: 2 Threads: 1
▶ 0x004369c0	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 7 Threads: 1
▶ 0x0042e580	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
▶ 0x0042f380	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
▶ 0x004327c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 4 Threads: 1
▶ 0x00440900	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x0042e9c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x004396c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x004399c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x00440dc0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1

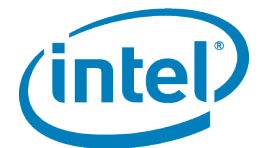


Select the falsely shared cacheline (now blue) and Filter the Hotspot view to only Display Accesses to that Line (multiple lines also work)

The screenshot displays the Intel(R) Performance Tuning Utility interface. At the top, a table lists functions and their performance metrics. Below this, a 'Total Selected:' section shows filter settings for Granularity (Function), Process (main_share.exe), Thread (All), and Module (All). A red circle highlights the 'Filter by selection' button, which is used to filter the 'Cachelines View' below. The 'Cachelines View' table shows various cache lines, with the first line (0x0042a3c0) highlighted in blue, indicating it is the selected line. A red arrow points from the 'Filter by selection' button to this blue-highlighted line.

Function	Module	Collected Data Refs (%Total)	LLC Misses (%Total)	Avg. Latency	Total Latency (%Total)	Cachelines #	Pages # (%Total)	MEM_LOAD_RETIRED.L2_MISS (%Total)
sort	main_share.exe	8,594,000,000 (100.0%)	400,000 (100.0%)	3	26,186,000,000 (100.0%)	1,029	24 (85.7%)	400,000 (100.0)

Cacheline Address / Offset / Thread / Function	Collected Data ...	LLC Misses (%T...	Avg. Latency	Total Latency (...	Contention (%...	MEM_LOAD_RE...	MEM_LOAD_RE...	INST_RETIRED...	Contributors
0x0042a3c0	1,959,600,000 ...	400,000 (100...	3	6,252,000,000 ...	909,100,000 (...	400,000 (100...	39,200,000 (8...	1,920,000,000 ...	Offsets: 2 Threads: 2
▶ Offset:0x04(4)	1,050,500,000 (...	100,000 (25.0%)	3	3,319,000,000 (...	0 (N/A)	100,000 (25.0%)	20,400,000 (46...	1,030,000,000 (...	Threads: 1
▶ Offset:0x00(0)	909,100,000 (1...	300,000 (75.0%)	3	2,933,000,000 (...	0 (N/A)	300,000 (75.0%)	18,800,000 (42...	890,000,000 (1...	Threads: 1
▶ 0x0064ff40	836,000,000 (...	0 (0.0%)	3	2,508,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	836,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff40	764,000,000 (...	0 (0.0%)	3	2,292,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	764,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff80	366,000,000 (...	0 (0.0%)	3	1,098,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	366,000,000 (...	Offsets: 2 Threads: 1
▶ 0x0064ff80	276,000,000 (...	0 (0.0%)	3	828,000,000 (...	0 (N/A)	0 (0.0%)	0 (0.0%)	276,000,000 (...	Offsets: 2 Threads: 1
▶ 0x004369c0	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 7 Threads: 1
▶ 0x0042e580	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
▶ 0x0042f380	14,000,000 (0...	0 (0.0%)	3	42,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0...	Offsets: 6 Threads: 1
▶ 0x004327c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 4 Threads: 1
▶ 0x00440900	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x0042e9c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x004396c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x004399c0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1
▶ 0x00440dc0	12,000,000 (0...	0 (0.0%)	3	36,000,000 (0...	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0...	Offsets: 5 Threads: 1



Only Events Referencing the Selected Line(s) are now in the Hotspot View Double Click to reach source/ASM view

The screenshot displays the Intel(R) Performance Tuning Utility interface. The main table shows performance metrics for various functions. The 'sort' function in the 'main_share.exe' module is selected and circled. Below the main table, the 'Cachelines View' provides a detailed breakdown of cache line events, including collected data references, LLC misses, and latency for each cache line address.

Function	Module	Collected Data Refs (%Total)	LLC Misses (%Total)	Avg. Latency	Total Latency (%Total)	Cachelines #	Pages # (%Total)	MEM_LOAD_RETIRED.L2_MISS (%Total)
sort	main_share.exe	1,959,600,000 (22.8%)	400,000 (100.0%)	3	6,252,000,000 (23.9%)	1	1 (3.6%)	400,000 (100.0)

Cacheline Address / Offset / Thread / Function	Collected Data ...	LLC Misses (%T...	Avg. Latency	Total Latency (...)	Contention (%...	MEM_LOAD_RE...	MEM_LOAD_RE...	INST_RETIRED...	Contributors
0x0042a3c0	1,959,600,000 ...	400,000 (100....	3	6,252,000,000 ...	909,100,000 (...	400,000 (100....	39,200,000 (8...	1,920,000,000 ...	Offsets: 2 Threads: 2
▶ Offset:0x04(4)	1,050,500,000 (...	100,000 (25.0%)	3	3,319,000,000 (...	0 (N/A)	100,000 (25.0%)	20,400,000 (46....	1,030,000,000 (...	Threads: 1
▶ Offset:0x00(0)	909,100,000 (1...)	300,000 (75.0%)	3	2,933,000,000 (...	0 (N/A)	300,000 (75.0%)	18,800,000 (42...	890,000,000 (1...	Threads: 1
▶ 0x0064ff40	836,000,000 (...	0 (0.0%)	3	2,508,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	836,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff40	764,000,000 (...	0 (0.0%)	3	2,292,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	764,000,000 (...	Offsets: 1 Threads: 1
▶ 0x0054ff80	366,000,000 (...	0 (0.0%)	3	1,098,000,000 ...	0 (N/A)	0 (0.0%)	0 (0.0%)	366,000,000 (...	Offsets: 2 Threads: 1
▶ 0x0064ff80	276,000,000 (...	0 (0.0%)	3	828,000,000 (...	0 (N/A)	0 (0.0%)	0 (0.0%)	276,000,000 (...	Offsets: 2 Threads: 1
▶ 0x004369c0	14,000,000 (0....	0 (0.0%)	3	42,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0....	Offsets: 7 Threads: 1
▶ 0x0042e580	14,000,000 (0....	0 (0.0%)	3	42,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0....	Offsets: 6 Threads: 1
▶ 0x0042f380	14,000,000 (0....	0 (0.0%)	3	42,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	14,000,000 (0....	Offsets: 6 Threads: 1
▶ 0x004327c0	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 4 Threads: 1
▶ 0x00440900	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 5 Threads: 1
▶ 0x0042e9c0	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 5 Threads: 1
▶ 0x004396c0	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 5 Threads: 1
▶ 0x004399c0	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 5 Threads: 1
▶ 0x00440dc0	12,000,000 (0....	0 (0.0%)	3	36,000,000 (0....	0 (N/A)	0 (0.0%)	0 (0.0%)	12,000,000 (0....	Offsets: 5 Threads: 1



The Pointer "sum" is Causing the False Sharing

The screenshot displays the Intel(R) Performance Tuning Utility interface. The left pane shows the source code for a sorting function:

```
1 int sort(int* data, volatile int* sum, int size...  
2 {  
3  
4     int i;  
5     for(i=0; i<size; i++)*sum += data[i]*data[i];  
6     return *sum;  
7 }
```

The right pane shows the assembly code for the same function. The instructions are grouped into blocks. Block 3, starting at address 0x156F, is highlighted in blue and contains the following instructions:

- 0x156F 5 mov edx, DWORD PTR [ebp-4]
- 0x1572 5 mov eax, DWORD PTR [ebp+08h]
- 0x1575 5 mov ecx, DWORD PTR [ebp-4]
- 0x1578 5 mov esi, DWORD PTR [ebp+08h]
- 0x157B 5 mov edx, DWORD PTR [eax+edx*4]
- 0x157E 5 imul edx, DWORD PTR [esi+ecx*4]
- 0x1582 5 mov eax, DWORD PTR [ebp+0ch]
- 0x1585 5 mov ecx, DWORD PTR [eax]
- 0x1588 5 add ecx, edx
- 0x1589 5 mov edx, DWORD PTR [ebp+0ch]
- 0x158C 5 mov DWORD PTR [edx], ecx
- 0x158E 5 jmp sort+00eh

The performance metrics for Block 3 are: 1,959,600,000 instructions collected, 400,000 LLC misses, 6,252,000 instructions total, and 400 MEM_L. The total selected for this block is 400 instructions.

The bottom status bar shows "Total Selected (4 instructions):" for the current view.



Load Latency Threshold Event

- **Ability to trigger count on minimum latency**
 - Core cycles from load execute->data availability
- **Linear address in PEBS buffer**
 - Allows driver to collect physical address
 - Only total measurement of local/remote home access
- **Data source captured in bit pattern**
 - Actual NUMA source revealed
- **Only ONE latency event/min thresh can be taken per run**
 - Minimum latency programmed with MSR
 - Global per core
 - 0x3F6 MS_PEBS_LD_LAT_THRESHOLD bits 15:0
- **Can use to detect a variety of properties about memory accesses**
 - Local vs remote etc.
 - Can be filtered at with hot spots to detect causes for them



Call to Action

- **Download PTU from <http://software.intel.com/en-us/articles/intel-performance-tuning-utility/> and have fun**
 - **Makes all the events of Core™ i7 available**
- **If you have any questions or comments you can reach me at ramesh.v.peri@intel.com or ask them in the discussion forums at <http://software.intel.com/en-us/forums/>**



Legal Acknowledgements

- Intel, the Intel logo, Intel Core and Core Inside are trademarks of Intel Corporation in the U.S. and other countries.
- Vtune is a trademark of Intel Corporation in the U.S. and other countries.
- Itanium is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.
- Other names and brands may be claimed as the property of others.

