



# Intel® VTune™ Amplifier: A Bridge to Performance, Parallelism, and Power

**Stanislav Bratanov**

Software and Services Group

*Intel Corporation*

November 21, 2013

**2nd CERN Advanced Performance Tuning workshop**

# Legal Disclaimer

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.

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, reference [www.intel.com/software/products](http://www.intel.com/software/products).

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Atom, Centrino Atom Inside, Centrino Inside, Centrino logo, Cilk, Core Inside, FlashFile, i960, InstantIP, Intel, the Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.

Copyright © 2010. Intel Corporation.

# Optimization Notice

## Optimization Notice

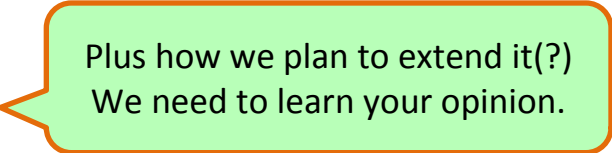
Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the “Intel® Compiler User and Reference Guides” under “Compiler Options.” Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools 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 Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® 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.

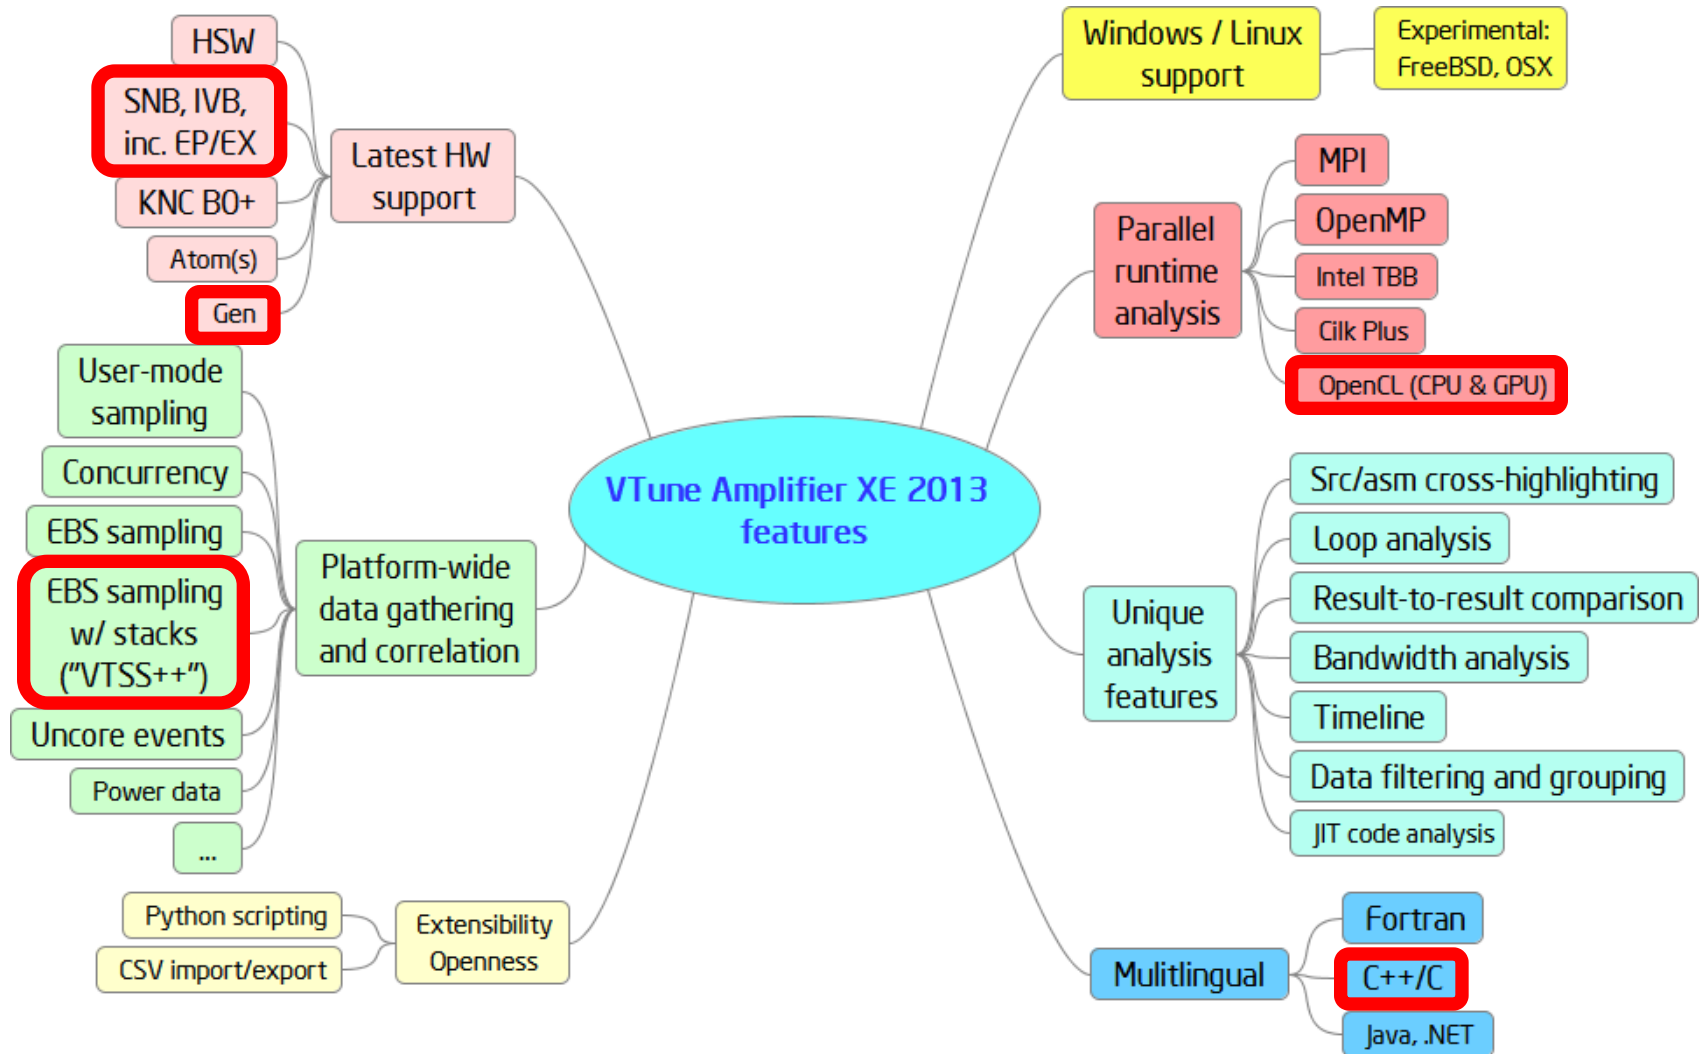
While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20101101

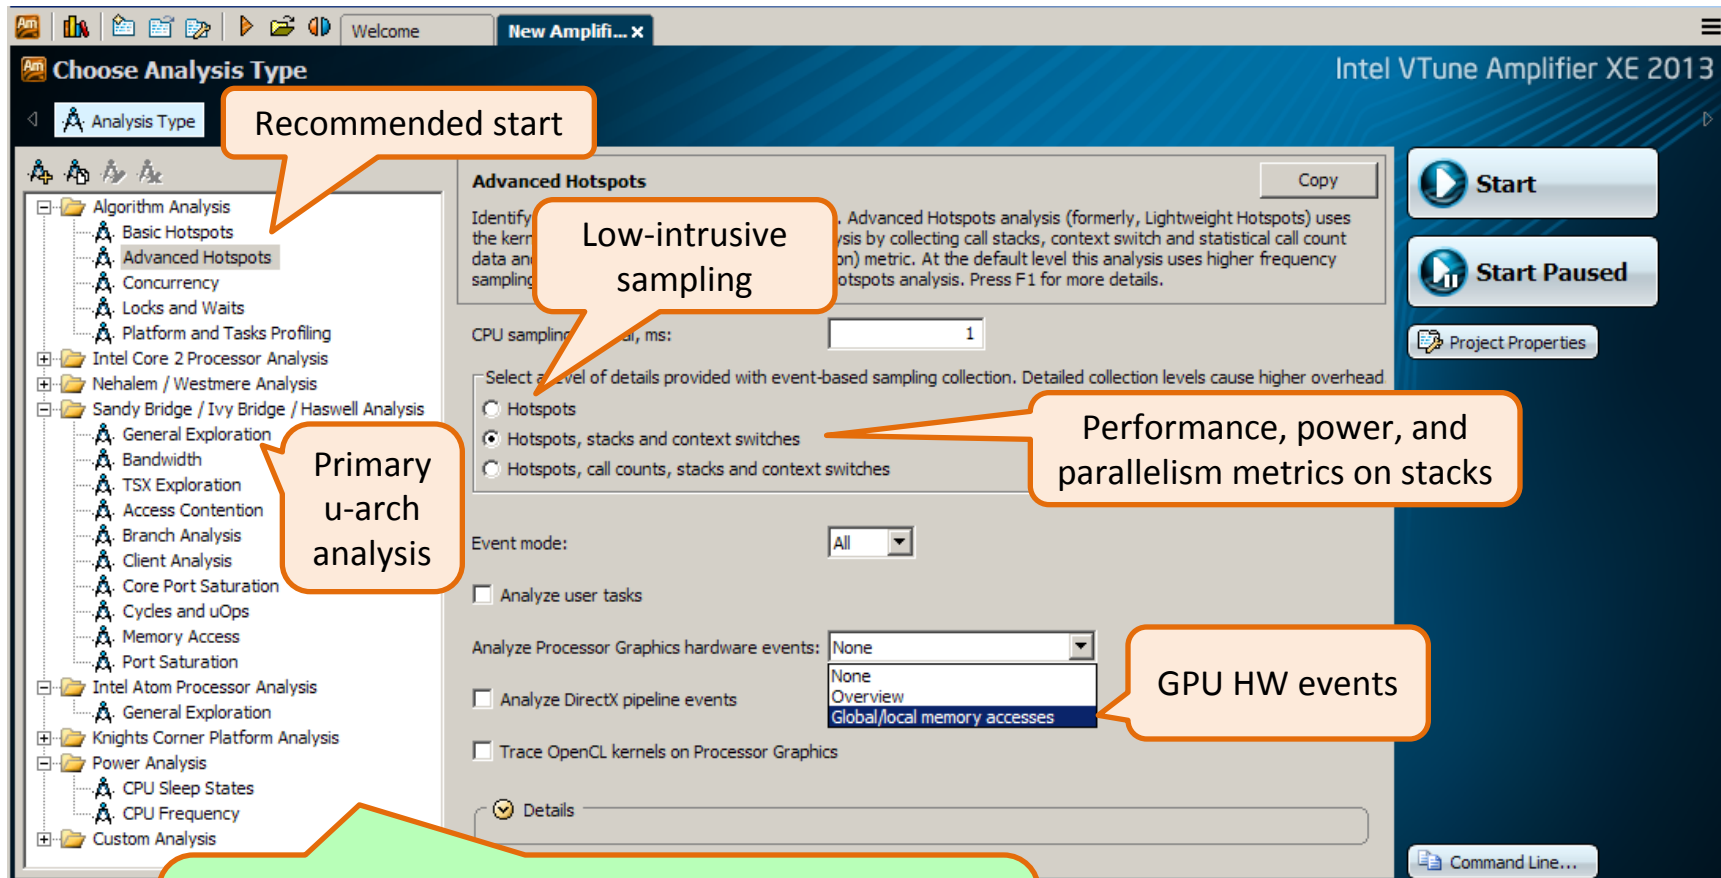
# Agenda

- How VTune works 
  - Optimizing for performance and power on CPU and GPU
- Case Study: NBody app
- Conclusions

# VTune is Big. Let's cover some of it



# First Things You See



TODO: Reorganize the hierarchy?  
uArch Analysis ->  
General Exploration ->  
CPU specific sets (that match the leaves of GE)



# Call Tree + Events + Threading Info

Primary hotspot

Synchronization hotspot (wait-spot)

HW events (e.g., clocks)

Thread contention

OS impact

Time lost on waits

Scheduled off CPU

/Function	Stack	CPU_CLK_UNH ... CORE by H/W ...	Synchronization Context Switches ...	Preemption Context Switches by H/W ...	Wait Time by H/W Context	Inactive Time by H/W Context
quantize_lines_xrpow		3,556,035,917	0	233	0	57,516,084
└─ quantize_lines_xrpow_ISO ← count_bits		3,556,035,917	0	233	0	57,516,084
└─ truncate_small_spectrums ← iteration_loop ← lam		3,180,425,513	0	209	0	51,279,756
└─ bin_search_StepSize ← truncate_small_spectrums		375,610,404	0	24	0	6,236,328
KiFastSystemCallRet		3,089,180,223	6,222	244	11,894,526,708	444,816,840
└─ NtWaitForSingleObject ← WaitForSingleObjectEx ←		2,993,182,443	5,730	218	11,817,527,328	435,746,772
└─ ZwRequestWaitReplyPort ← CsrClientCallServer		62,794,255	454	15	25,823,520	461,976
└─ ZwSetEvent ← SetEvent		20,723,639	32	6	50,760,216	8,319,132

We lost almost half of potential performance on contention: clocks wasted on contention are comparable with the time of useful work

Major contention on a WaitForSingleObject

TODO: Need to craft a special contention metric or a specialized viewpoint?



# Same + Active and Idle Power

Function / Call Stack	Hardware Event Count									Energy Core
	CPU_CLK_...	Idle Time	C3 Residency	C6 Residency	Wait Time	Inactive Time	Idle Wakeup	Synchronization...	Preemption...	
smvp	115,044,088,899	12,540,936,543	35,402,724	191,328,732	17,277,068,387	20,884,260	73,336	77,419	313	1,675,073,600
main	71,738,528,080	584,131,993	8,884,260	11,203,686	879,118,000	9,837,224	6,210	6,677	230	873,503,056
phi0	4,688,113,096	43,676	0	0	1,683,176	1,000,928	3	7	24	75,284,320
phi1	3,226,339,383	11,542	0	0	587,492	548,880	1	1	14	44,893,744
sin	2,935,350,290	0	0	0	883,050	544,230	0	2	14	48,974,016
phi2	2,739,433,958	14,192	0	0	1,021,170	510,048	1	1	14	33,486,224
[wow64cpu.dll]	2,186,456,142	18,877,721,482	100,790,496	1,441,949,274	89,381,484,762	398,825,650	317,483	324,945	68	567,768,464
WaitForSingleObject	1,838,001,558	18,872,158,940	100,790,496	1,441,949,274	89,379,240,911	398,806,258	317,313	324,773	67	567,005,904

Hotspots

HW events

Idle time

Cx state residency

Wait and inactive times

Wakeup from idle

Context switches

Consumed energy (uJoules)

Call stack

TODO: Need specialized metrics or viewpoints to automate idle power analysis?



System spent only ~10% of idleness in C6 state

Almost every wait brought the system to idle and then caused a wakeup

# Case Study

- NBody application:

- N bodies moving in the gravity field

- Source code attached:    
Not optimized    Optimized

- Runs on CPU and then on GPU

- 64k bodies for CPU, 256k bodies for GPU

- > to maintain comparable execution times (similar statistical errors)

- Intel® Core™ i7 3667U

- Intel® HD Graphics 4000

**CPU**

Name:	3rd generation Intel(R) Core(TM) Processor family
Frequency:	2.5 GHz
Logical CPU Count:	4

**GPU**

Name:	Intel(R) HD Graphics 4000
Vendor:	Intel Corporation
Driver:	9.18.10.3071 (3/18/2013)
Stepping:	8
EU Count:	16
Max EU Thread Count:	8
Max Core Frequency:	1.1 GHz

# Beginning with the Analysis

**General Exploration - Sandy Bridge / Ivy Bridge 0** General Exploration viewpoint (change) Intel VTune Amplifier XE 2013

Analysis Target Analysis Type Summary Bottom-up **Top-down Tree** Tasks and Frames nbody.cpp

Call Stack

Call Stack	Hardware Event Count: Total by Har ...	CPI Rate: Total	CPI Rate: Self
BaseThreadInitThunk	594,600,993,655	1.133	0.000
[Unknown stack frame(s)]	477,394,932,791	1.104	1.062
main	477,383,473,147	1.104	0.000
compute_bodies	446,662,587,670	1.109	0.000
[Loop@0x4014fe in compute_bodies]	395,063,701,545	1.047	0.000
[Loop@0x401518 in compute_bodies]	394,949,742,006	1.047	0.000
<b>[Loop at line 201 in compute_bodies]</b>	<b>346,418,487,695</b>	<b>1.047</b>	<b>1.047</b>
CIIsqrt	79,201,534,582	1.046	1.051
ExReleaseRunDownProt		1.050	1.050
[Loop at line 160 in comp		1.047	1.047
[Loop at line 160 in comput		1.020	1.020
[Loop at line 234 in compute_		2.022	2.022
[Loop at line 160 in compute_		1.024	1.024

Context Switch Call Stack

Viewing 1 of 1 selected stack(s)

100.0% (0.004s of 0.004s)

nbody-cpu-64-16.exe![\_Loop at line 160 in compute\_bodies]...

nbody-cpu-64-16.exe![\_Loop@0x401518 in compute\_bodies]...

nbody-cpu-64-16.exe![\_Loop@0x4014fe in compute\_bodies]...

nbody-cpu-64-16.exe!compute\_bodies+0x29e - nbody.cpp...

nbody-cpu-64-16.exe!main+0x177 - nbody.cpp:379

nbody-cpu-64-16.exe![\_Loop at line 160 in compute\_bodies]...

kernel32.dll!BaseThreadInitThunk+0xd - [Unknown]:[Unknown]

ntdll.dll!RtlInitializeExceptionChain+0x84 - [Unknown]:[Unk...

ntdll.dll!RtlInitializeExceptionChain+0x57 - [Unknown]:[Unk...

TODO: Need to estimate the number of iterations?

Select any region of inactivity and see sync call stack here

Thread

- cpx\_intrinsic\_thread (0x8)
- cpx\_intrinsic\_thread (0x1)
- cpx\_intrinsic\_thread (0xd)
- mainCRTStartup (0x1398)
- cpx\_intrinsic\_thread (0x4)

Hardware Events

Thread

- Running
- Context Sw...
- Hardware ...
- Hardware Events
- Hardware ...

No filters are applied. Any Process Thread: Any Thread

Timeline Hardware Event: CPU\_CLK\_UNHALTED.THREAD Call Stack Mode: User/system func Loop Mode: **Loops and functions**

40000ms 40100ms 40200ms 40300ms 40400ms 40500ms 40564.28ms 40700ms 40800ms

cpx\_intrinsic\_thread (0x8e0)

Hardware Event Count  
4,647,614

Context Switches  
Start: 40564.614ms Duration: 16.01us  
CPU: cpu\_0  
Reason: Synchronization  
Source File: nbody.cpp  
Source Line: 175

# Locating Threading Inefficiencies

Advanced Hotspots 0 Hardware Event Counts viewpoint (change) ?

Analysis Target Analysis Type Summary PMU Events Uncore Events Caller/Callee Top-

Grouping: Function / Call Stack

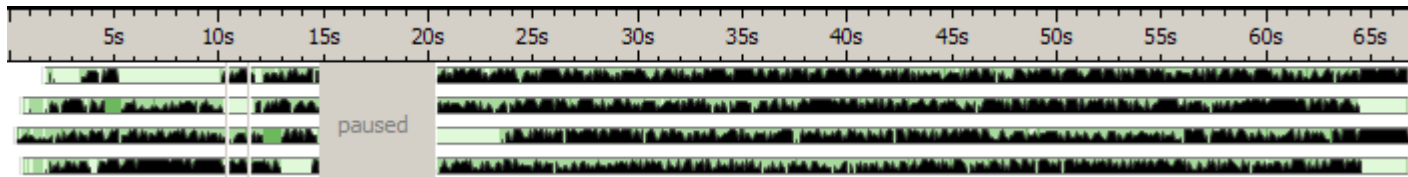
Function / Call Stack	CPU_CLK_UNHALTED.THREAD	Synchronization Context Switches
compute_bodies	290,757,552,857	5,339
CIIsqrt	35,533,930,925	618
math_exit	17,323,254,362	388
cp_x_exported_pick_nested_job	8,688,064,785	180
checkTOS_withFB	5,837,183,653	87
TlsGetValue	1,933,510,947	28
cp_x_intrinsic_get_tlsctx	1,296,428,368	
TlsGetValue	738,653,541	
[Unknown stack frame(s)]	635,280,540	
[NETwNe64.sys]	318,587,956	
KeAcquireSpinLockRaiseToDpc	47,787,540	
KeReleaseSpinLock	45,170,865	
[Outside any known module]	37,392,510	
RtlIsCriticalSectionLockedByThread	36,393,418	
KeSynchronizeExecution	25,851,512	0
hybDriverEntry	5,817,479	0
[storahci.sys]	5,778,814	0
[pinvm.dll]	4,704,889	0
[vtss.sys]	23,400,369	0
[wow64cpu.dll]	22,802,779	148
WaitForSingleObjectEx ← WaitForSingleObject	16,400,785	31
cp_x_exported_wait	11,150,121	21
compute_bodies ← main	6,540,565	8
_tmainCRTStartup ← BaseThreadInitThunk	2,256,224	1

Find thread synchronizations (that stem from ntdll/wow64)

The performance cost of thread contention is ~0% of the primary hotspot => no performance impact of thread contention

# No Problem, as Predicted

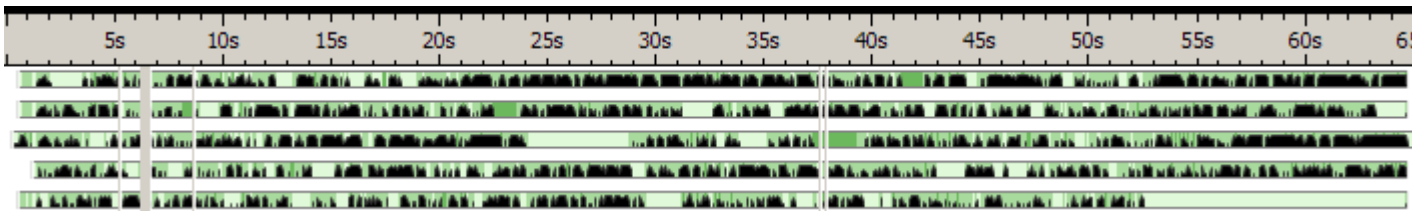
Elapsed Time: 67.238s



TODO: Need to emphasize this is an auto-pause we have to generate not to lose a single event count?

Slight CPU oversubscription is even better in this case: helps to hide various stalls

Elapsed Time: 65.066s



# Locating Performance Issues

General Exploration - Sandy Bridge / Ivy Bridge 0 General Exploration viewpoint (change)

Analysis Target Analysis Type Summary Bottom-up Top-down Tree Tasks and Frames

Grouping: Function / Call Stack

Function / Call Stack

Unfilled Pipeline Slots (Stalls)

Back-end Bound

Memory Bound

L1 Bound L2 Bound L3 Bound DRAM Bound Store Bound Core Bound Front-end Bound

	Loads Blocked by Store Forwarding	Split Loads	4K Aliasing	DTLB Overh ...	L2 Bound	L3 Bound	DRAM Bound	Store Bound	Core Bound	Front-end Bound
[Loop at line 201 in compute_bodies]					0.000	0.000	0.000	0.000	0.107	0.016
[Loop at line 234 in compute_bodies]					0.000	0.001	0.000	0.004	0.140	0.079
[Loop at line 160 in compute_bodies]				0	0.000	0.000	0.000	0.000	0.000	0.020
CIsqrt					0.000	0.000	0.000	0.000	0.065	0.033
math_exit					0.002	0.000	0.000	0.000	0.027	0.042
				0.002	0.001	0.000	0.000	0.000	0.107	0.016

Selected 1 row(s):

Expand each column marked pink until you come to the actual issue

Here is the problem, read tooltip to learn more

Loads are blocked during store forwarding for a significant proportion of cycles. Use source/assembly view to identify the blocked loads, then identify the problematically-forwarded stores, which will typically be within the ten dynamic instructions prior to the load. If the forwarding store is smaller than the load, change the store to be the same size as the load.

Threshold:  $(((((13 * LD\_BLOCKS.STORE\_FORWARD) / CPU\_CLK\_UNHALTED.THREAD) > 0.05) * (CPU\_CLK\_UNHALTED.THREAD / > 0.05)))$

Pick the actual HW event from the formula (**LD\_BLOCKS.STORE\_FORWARD** – typically counts bigger-size loads blocked by smaller stores to the same address) for further detailed analysis

# Locating Performance Issues

Source	LD_BLOCKS.STORE_FORWARD	Assembly	LD_BLOCKS.STORE_FORWARD
<pre>           // 1 / dist^(3/2)           float inv = rsqrtf(dist);           float cube = inv * inv * inv;            // compute force           float s = cache[jdx + 3] * cube;         </pre>	6,101,778,128	<pre>           fsub st0, dword ptr [ecx+edi+1+0x8]           fstp dword ptr [esp+0x48], st0           fld st0, dword ptr [esp+0x44]           fld st0, dword ptr [esp+0x40]           fld st0, dword ptr [esp+0x48]           fld st0, st1           fmulp st2, st0           fld st0, st2           fmulp st3, st0           fxch st0, st1           faddp st2, st0           fmul st0, st0           faddp st1, st0           fstp dword ptr [esp+0x14], st0           fld st0, dword ptr [esp+0x14]           fadd st0, dword ptr [0x42a008]           fstp dword ptr [esp+0x14], st0           fld st0, dword ptr [esp+0x14]           call 0x402600 &lt;_CIsqrt&gt;           Block 45:           fstp dword ptr [esp+0x14], st0           fld st0, dword ptr [esp+0x14]           add esi, 0x10           sub ebx, 0x1           fld1 st0           fdivrp st1, st0           fstp dword ptr [esp+0x14], st0         </pre>	248,540 0 0 461,507 0 0 131,738,979 220,391 10,723,596 123,269,940 0 8,909,793 5,396,063,183

Here is the culprit line...

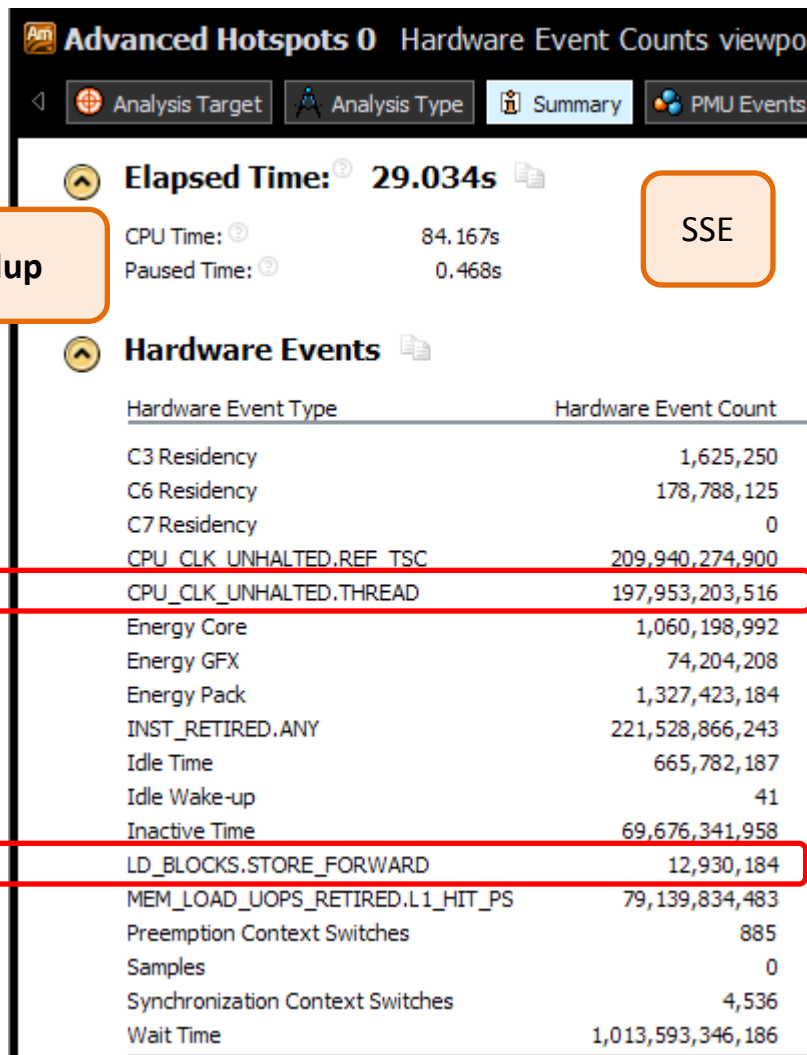
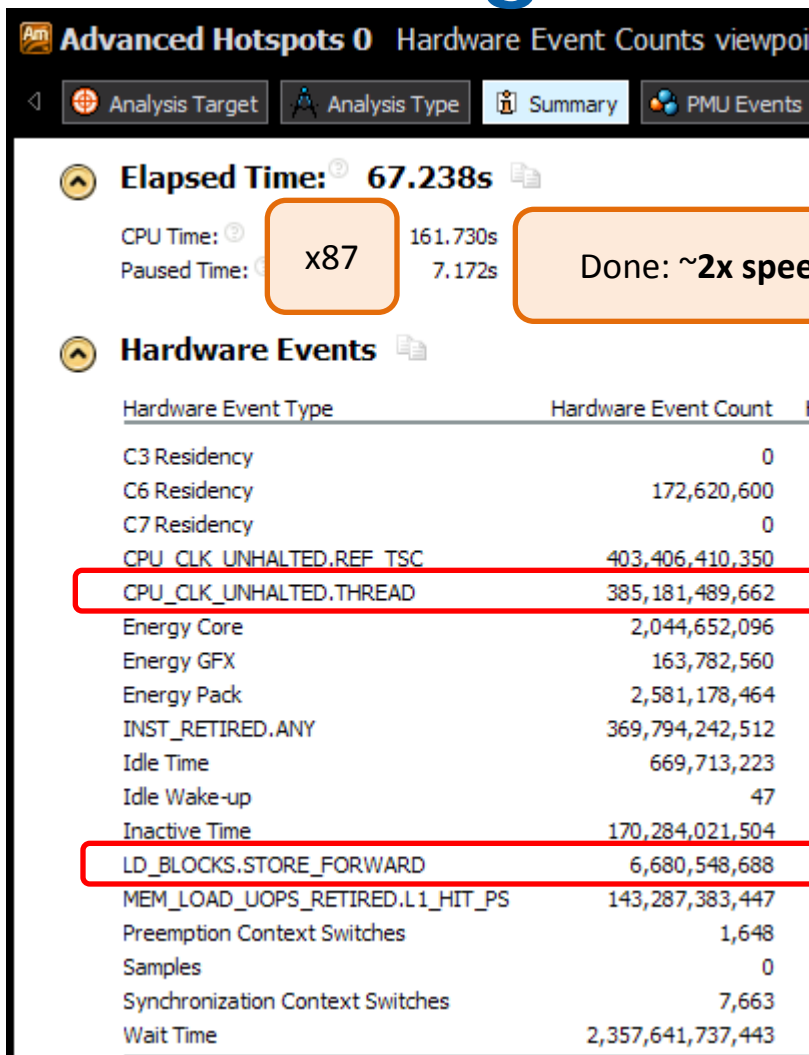
...and its corresponding disassembly highlighted

We'd better switch to SSE and eliminate both store-forwarding blocks, and DIV/SQRT latencies

TODO: Need a static-analysis best-case performance estimate to see potential gains?

TODO: Need an instruction stream view (unroll loops and calls) to see border effects?

# Eliminating Performance Issues





# Joules per Element: Better on GPU

- CPU: ~**20255** micro-Joules per element (64k elements)

Energy Core	1,060,198,992
Energy GFX	74,204,208
Energy Pack	1,327,423,184

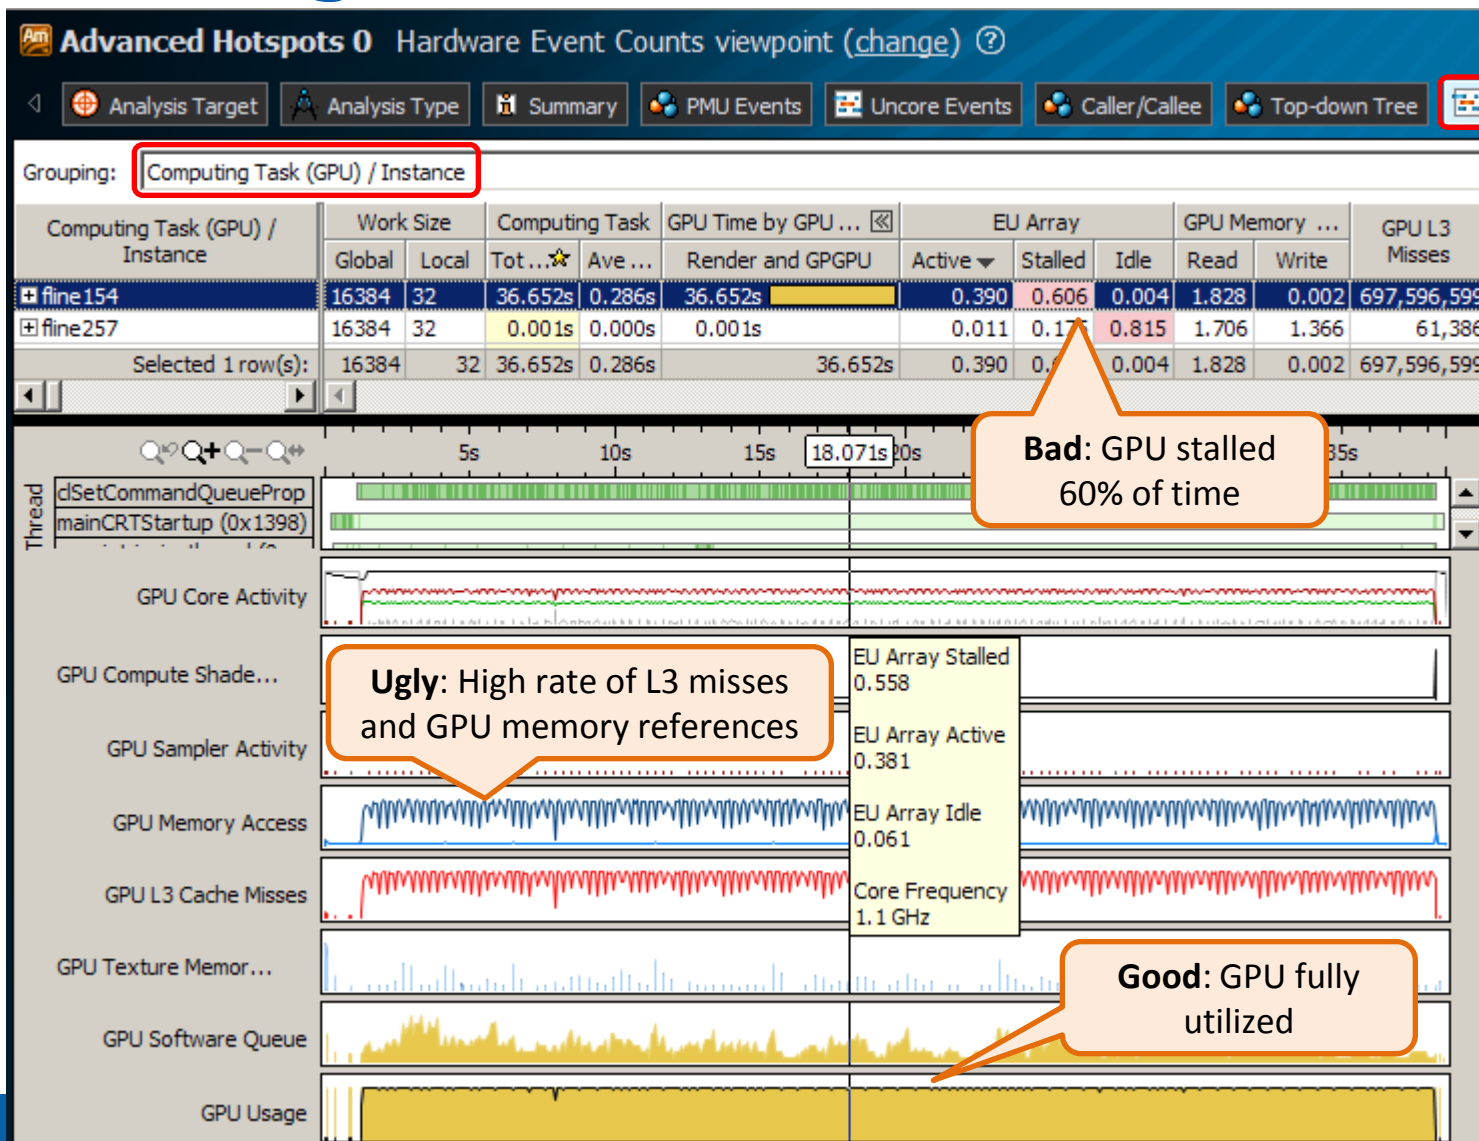
- GPU: ~**2332** micro-Joules per element (256k elements)

Energy Core	175,318,256
Energy GFX	346,957,008
Energy Pack	611,338,976

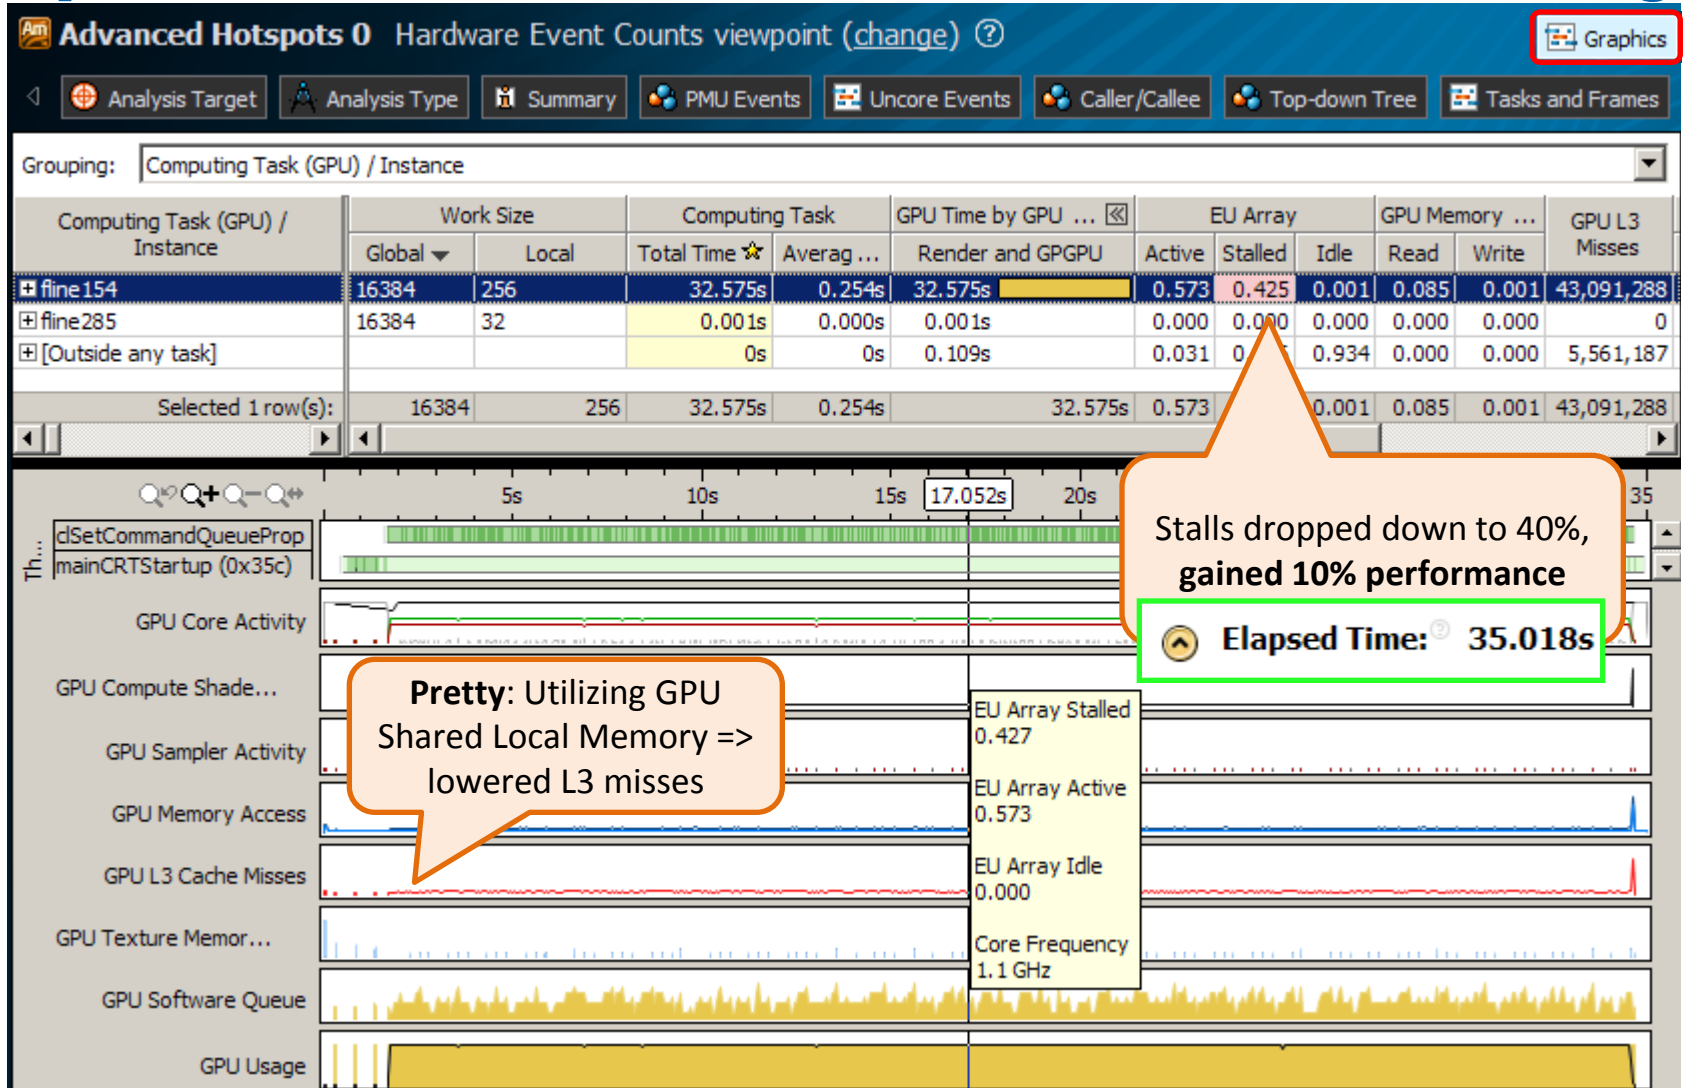
- GPU is ~**8X** more power efficient!

# Locating Issues on GPU

Elapsed Time: 38.643s



# Optimized for Shared Local Memory



# Avoid Long-Latency Functions

- Some math functions have long latencies in HW
  - Compare optimized and non-optimized versions w/o RSQRT:

Computing Task (GPU) / Instance	Work Size		Computing Task		GPU Time by GPU ...	EU Array		
	Global	Local	Tot ...	Ave ...		Active	Stalled	Idle
fine154	16384	32	35.181s	0.275s	35.181s	0.593	0.402	0.004
fine257	16384	32	0.001s	0.000s	0.001s	0.000	0.000	0.000

Elapsed Time: 37.315s

**Non-opt:** No speedup w/o RSQRT: memory stalls ruined the performance

Computing Task (GPU) / Instance	Work Size		Computing Task		GPU Time by GPU ...	EU Array		
	Global	Local	Tot ...	Ave ...		Active	Stalled	Idle
fine154	16384	256	11.894s	0.093s	11.894s	0.791	0.206	0.003
fine285	16384	32	0.001s	0.000s	0.001s	0.017	0.146	0.841

Elapsed Time: 14.017s

**Optimized:** ~2.5X speedup w/o RSQRT. ~2X lower stalls

# Locating Idle Power Leaks

Our code wakes the system up but lets the system stay in C6 for 15% of idleness  
(Idle\_Time / C6\_Residency)

Advanced Hotspots 0 Hardware Event Counts viewpoint (change) ?

Analysis Target Analysis Type Summary PMU Events Uncore Events Caller/Callee Tasks and Frames Graphics

Grouping: Function / Call Stack

Function / Call Stack	CPU_CLK_UNHALTED. ...	Idle Wake-up ▾	Idle Time	Wait Time	C6 Residency
[wow64cpu.dll]	18,256,902,833	126	1,143,310,881	837,608,717,747	185,591,275
[-] WaitForSingleObjectEx ← WaitForSingleObject	102,562,975	68	539,035,302	401,030,619,946	81,376,850
[-] [igdrcl32.dll] ← [igdrcl32.dll]	85,836,465	52	453,324,550	58,604,833,257	54,865,600
[-] dEnqueueReadBuffer ← D2H_od ← cpx_intrinsic_coload ← cpx_intrinsic_coload	80,982,781	51	453,312,330	58,603,311,271	54,865,600
[+] dEnqueueWriteBuffer ← H2D_od ← cpx_intrinsic_coload ← cpx_intrinsic_coload	4,853,684	1	12,220	1,521,986	0
[+] cpx_exported_wait	9,336,986	15	85,595,377	331,079,831,939	26,511,250
[+] cpx_intrinsic_thread ← BaseThreadInitThunk ← RtlInitializeExceptionChain ← RtlIn	4,347,915	1	115,375	11,269,466,766	0
[+] [pin.exe]	3,041,609	0	0	76,422,634	0
[+] crtExitProcess ← doexit ← exit ← tmainCRTStartup ← BaseThreadInitThunk ← R	0	0	0	65,350	0
[-] SleepEx ← Sleep	11,603,835,926	30	144,545,285	145,780,379	0
[-] [igdrcl32.dll] ← dSetCommandQueueProperty ← dSetCommandQueueProperty ←	11,603,835,926	30	144,545,285	7,422,044	0

Sleep() in **OpenCL** runtime both wastes active time and **doesn't** let the system go to C-states.  
Note that Sleep(0) is in many cases just an inefficient spin-wait

# Case Study Summary

- We scrutinized a parallel app, and:
  - Proved there are no threading issues
  - Found & eliminated a performance issue
  - Measured energy per element
  - Improved energy consumption 8x by moving from CPU to GPU
  - Found inefficiency in GPU memory usage
  - Optimized program and gained 10%
    - > Could gain 2.5x, but were impeded by RSQRT
  - Can lower idle power consumption by minimizing wake-ups in OpenCL runtime

# Conclusions

- Intel® VTune™ Amplifier XE:
  - Facilitates micro-architectural analysis
  - Uncovers software execution logic
  - Reveals threading inefficiencies and cost of parallelism
  - Correlates performance/power/parallelism metrics

– Is a

**bridge**

between

**you**



and  
detailed **SW**  
**analysis**

