



# Two examples\case studies using Intel® Xeon™ Phi

Tachyon Ray Tracing

Cloverleaf Hydrodynamics Mini-app

Stephen Blair-Chappell, Intel





# Tachyon ray tracer

Port to Intel® Xeon Phi™ with Intel® Cluster Studio XE



# Project goals

- Port to Intel® Xeon Phi™ and reach tangible performance gains vs initial Xeon-only baseline
- Test-drive Intel® Cluster Studio XE on Xeon Phi
- Create a case-study, with practical recommendations reusable in other cases

Not a goal: to create the best performing ray tracer. Refer to dedicated projects (e.g. Embree by Intel Labs)


# Tachyon ray tracer

Open source ray tracing demo

(<http://jedi.ks.uiuc.edu/~johns/raytracer/>)

Part of SpecMPI suite

Supports parallelism (MPI + OpenMP)



# Computational modes

Real-time rendering



Images (c) Audi, Dreamworks

Throughput computing



Production of *Puss in Boots*  
required **69 million** render hours



# Tachyon algorithm

3D model is a set of primitives (e.g. triangles)

3D space is pre-divided into grid, each voxel points to list of triangles contained/crossing it

Image pixel calculated using ray intersections (lights, reflections, shadows)

Hybrid parallelism: each frame is divided into chunks processed by MPI processes, a chunk is divided into lines processed by OpenMP threads



# Known issues of the algorithm

## ☹️ Communication profile:

- 1 master and n workers. Workers communicate to the master only.
- Master performs same computations + processing. A bottleneck and limited scalability.
- Each frame starts after a previous one. All workers have to wait for order from the master.

## ☹️ Work imbalance: lines and frames have different complexities

- 😊 Hybrid parallelism with dynamic OpenMP scheduling helps to relieve
- Static MPI scheduling still exhibits the issue across frames

Limited scalability across Xeon cluster. MPI+OpenMP hybrid better than MPI only

# Extra challenge - imbalance across Xeon and Xeon Phi

Xeon and Xeon Phi have different performances

How to split up the work ?

Which execution model to choose ?

Is ray tracing good for Xeon Phi ?



# Porting: Efficient apps for Xeon Phi

Tachyon's profile:

---

1. Allow massive parallelism  
(to load 60+cores x 4 threads)

☹ no slack: available parallel work  
(frame height) ~ # of threads


---

2. Run intensive computations  
(to efficiently use 512bit vectors)

☹ no vectorizable loops, only  
scalar computations

---

3. Provide memory efficiency  
(to meet current 4-8GB constraints)




Your application needs to meet certain requirements to use Xeon Phi best


# Target execution model – Symmetric MPI



**NATIVE model**



**OFFLOAD model**



**SYMMETRIC model**

Most flexible. Least number of code changes.

# Build for Xeon Phi

No code changes, only makefile:

-mmic	Target platform is Xeon Phi
-fp-model fast=2	Trade-off between accuracy and performance, OK for ray tracing

Very easy! Running code in a minute

# Why '-fp-model fast=2' ?

With default flag, a reciprocal ( $1/x$ ) computation unexpectedly became a hotspot on Phi (not on Xeon):

- Compiler generated heavy-weight code for higher precision

-fp-model fast=2 is a trade-off to favor performance (precision is still fine for ray tracing)

- Reciprocal calculation time reduced by >2x

# Run...

```
export I_MPI_MIC=enable
```

```
mpiexec.hydra \  
  -n 2 -host mynode1 <command-line> : \  
  -n 2 -host mynode2 <command-line> : \  
...  
  -n 2 -host mynoden <command-line> : \  
  -n 2 -host mynode1-mic0 <command-line> : \  
  -n 2 -host mynode1-mic1 <command-line> : \  
  -n 2 -host mynode2-mic0 <command-line> : \  
...  
  -n 2 -host mynoden-mic1 <command-line>
```

Same syntax. A Phi card is just like another node.

# First results



4 nodes x 2SNB –  
102 FPS



4 nodes x 1KNC – **32**  
**FPS ???**



4 nodes x (2SNB +  
1KNC) - **38 FPS !!!**

SNB – Sandy Bridge, 2<sup>nd</sup> generation Intel® Core™ processors

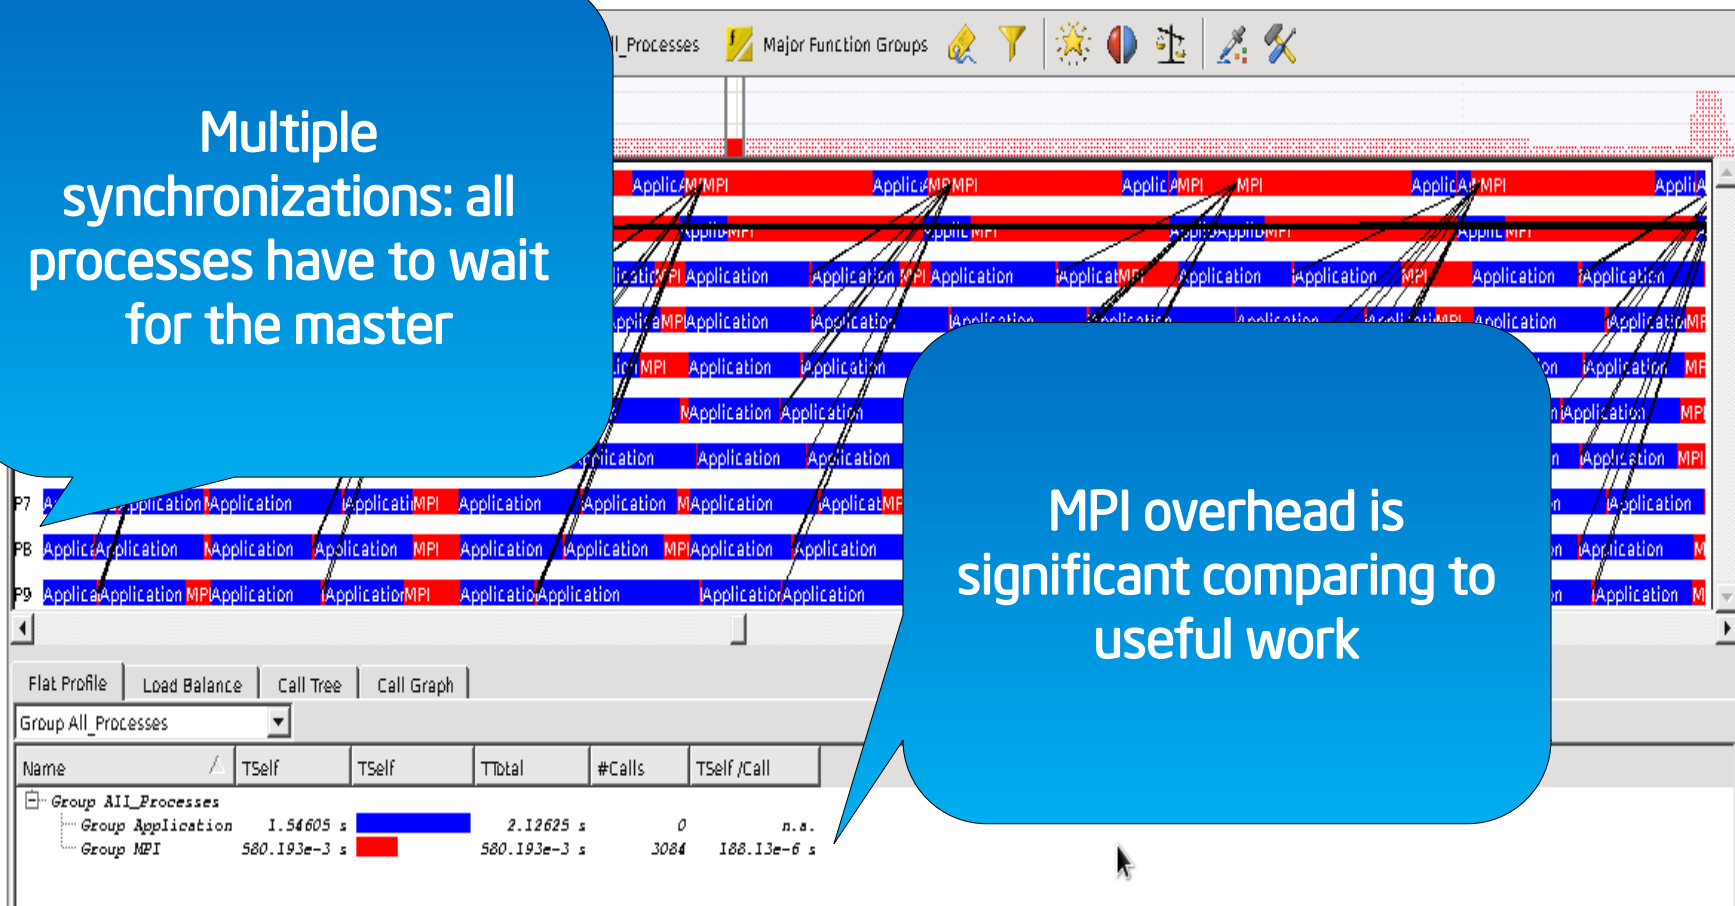
KNC – Knights Corner, Intel® Xeon Phi™ co-processors

Heterogeneous run slows down. Need to understand what happens

# Using Intel® Trace Analyzer and Collector

Multiple synchronizations: all processes have to wait for the master

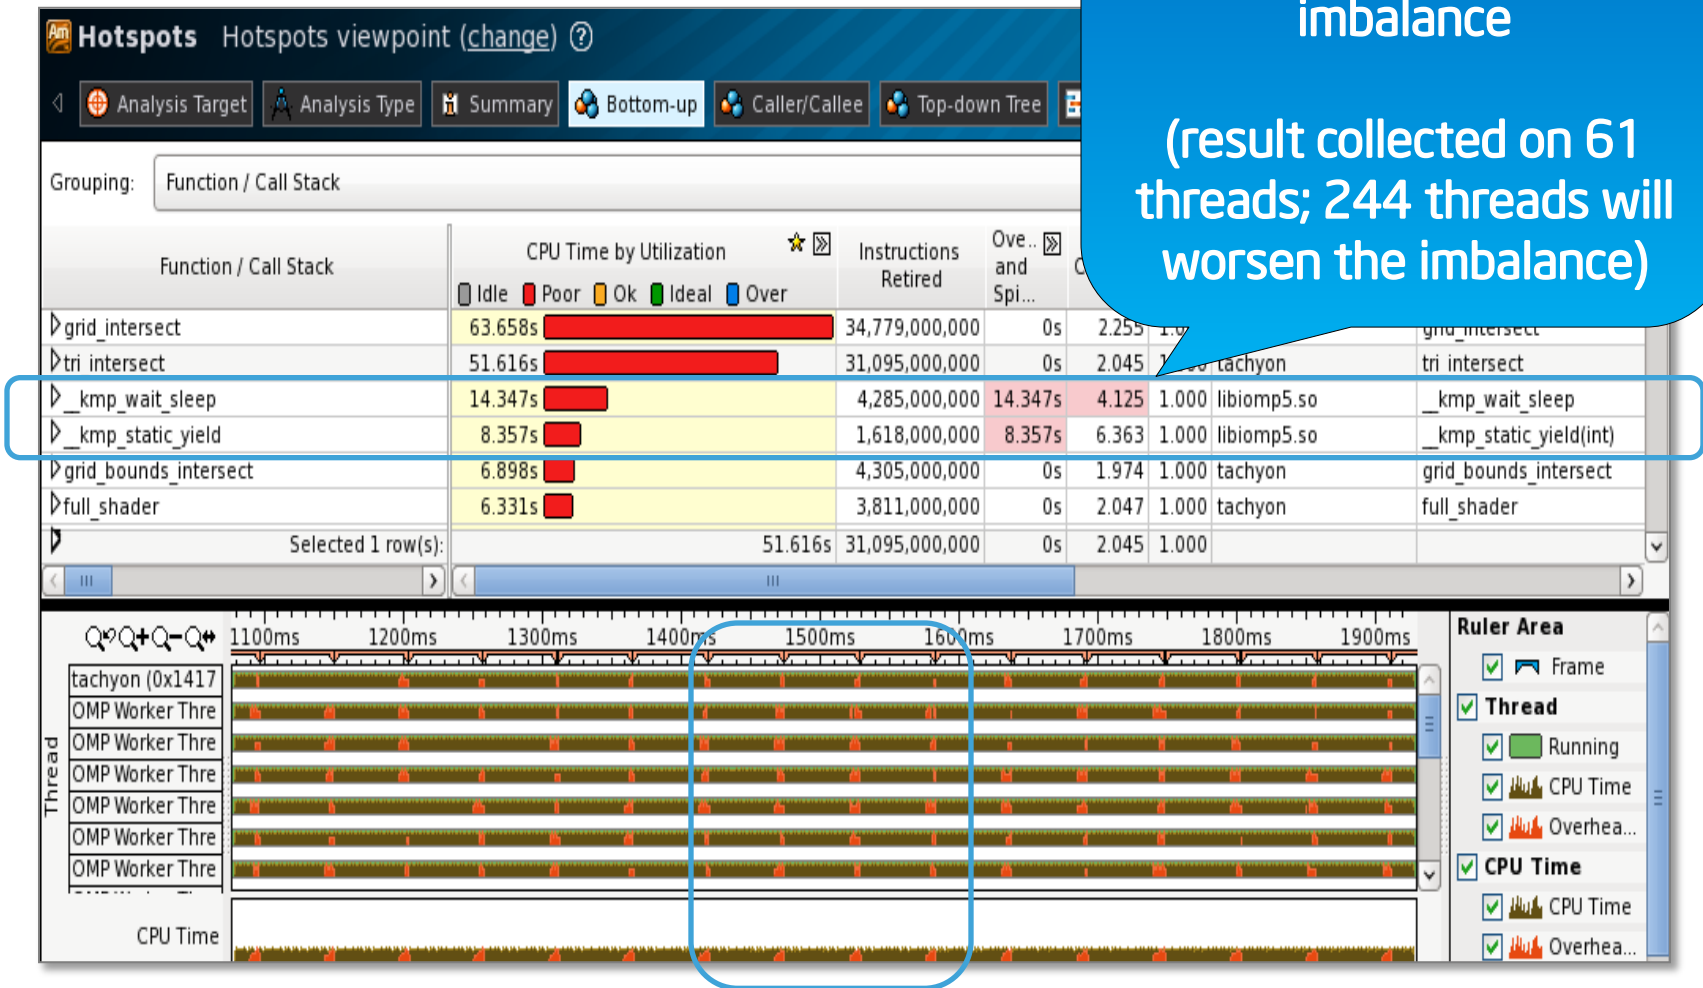
MPI overhead is significant comparing to useful work



# Using VTune™ Amplifier X

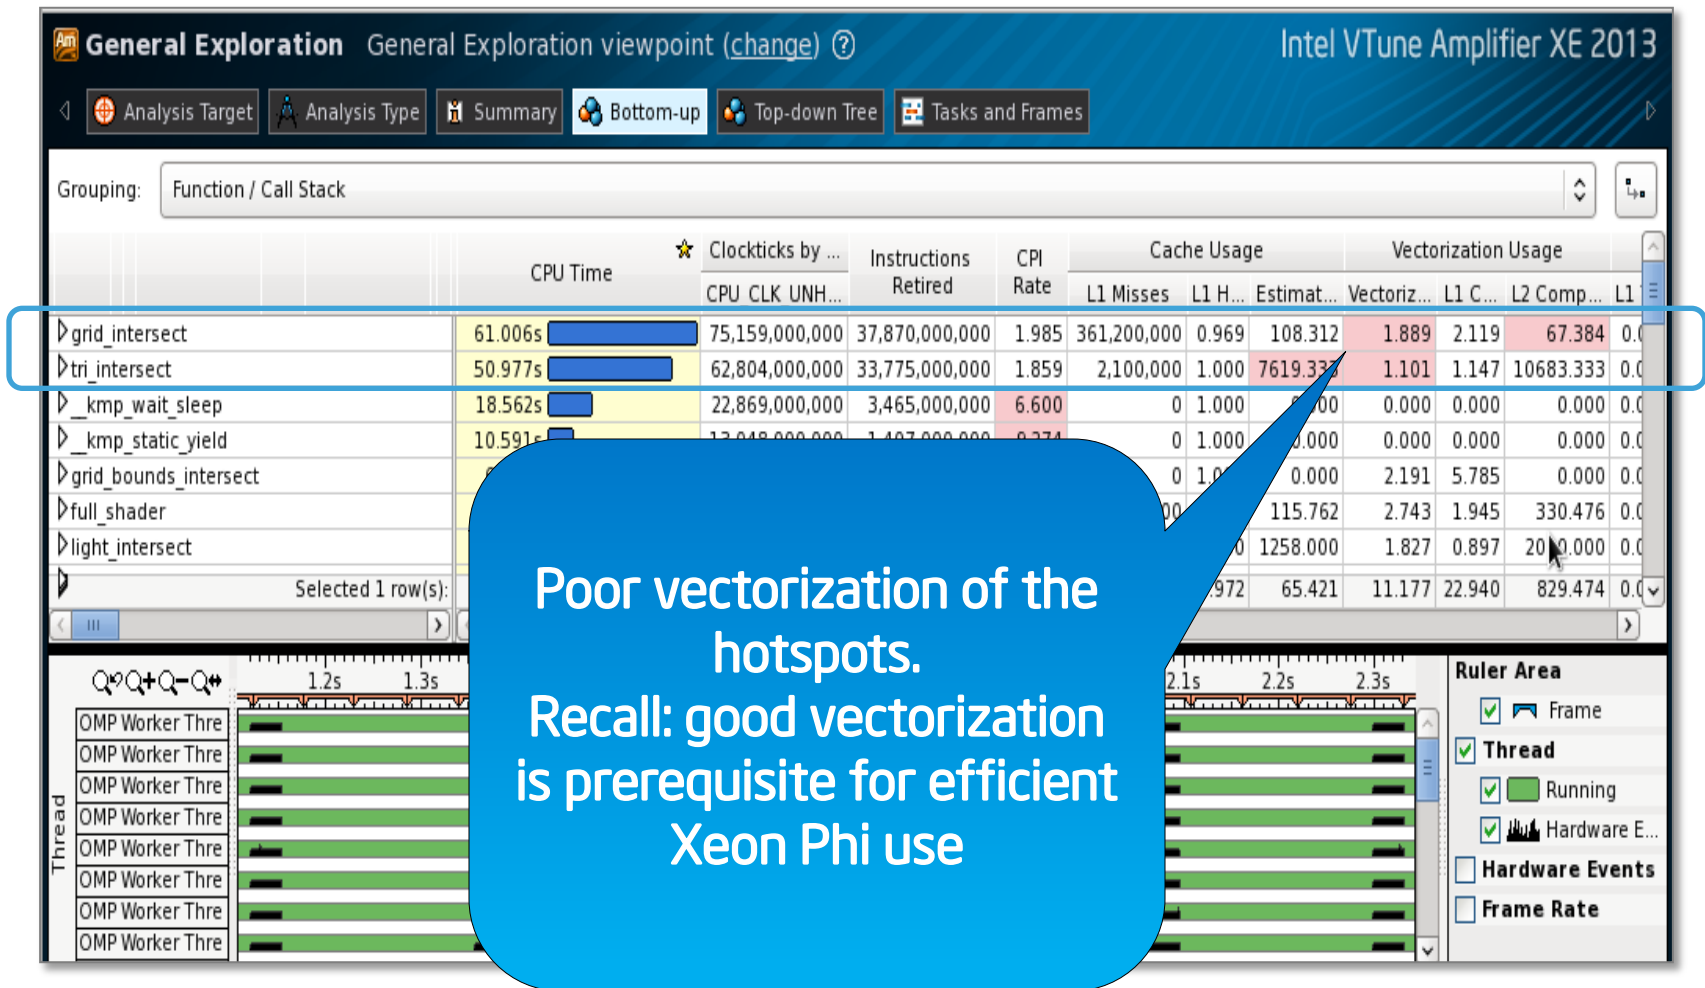
OpenMP overhead within  
each frame due to work  
imbalance

(result collected on 61  
threads; 244 threads will  
worsen the imbalance)





# Using VTune™ Amplifier XE



# Conclusions

- No vectorization – 512 bit registers (able to hold 16 floats) are wasted
- Insufficient parallelism – 240 hyper-threads are wasted




- Ranks on Xeon Phi run slower than on Xeon

Due to static MPI scheduling within each frame and frame-by-frame computation, Xeon's cannot start new frame until Xeon Phi's complete their lines.


Total performance suffers

# Improvement directions



This works for **both** Xeon Phi and Xeon

# #1 - Dynamic MPI scheduling



Each worker computes entire frame: asks a master for a frame #, computes and sends back entire frame

Master maintains a circular buffer, dispatches frame #, displays a frame. No computation by master

- Circular buffer to avoid memory growth

---

Significantly reduces # of communications

Reduced synchronizations: a worker does not wait for others anymore

---

Compensates Xeon vs Xeon Phi difference

Increases scalability

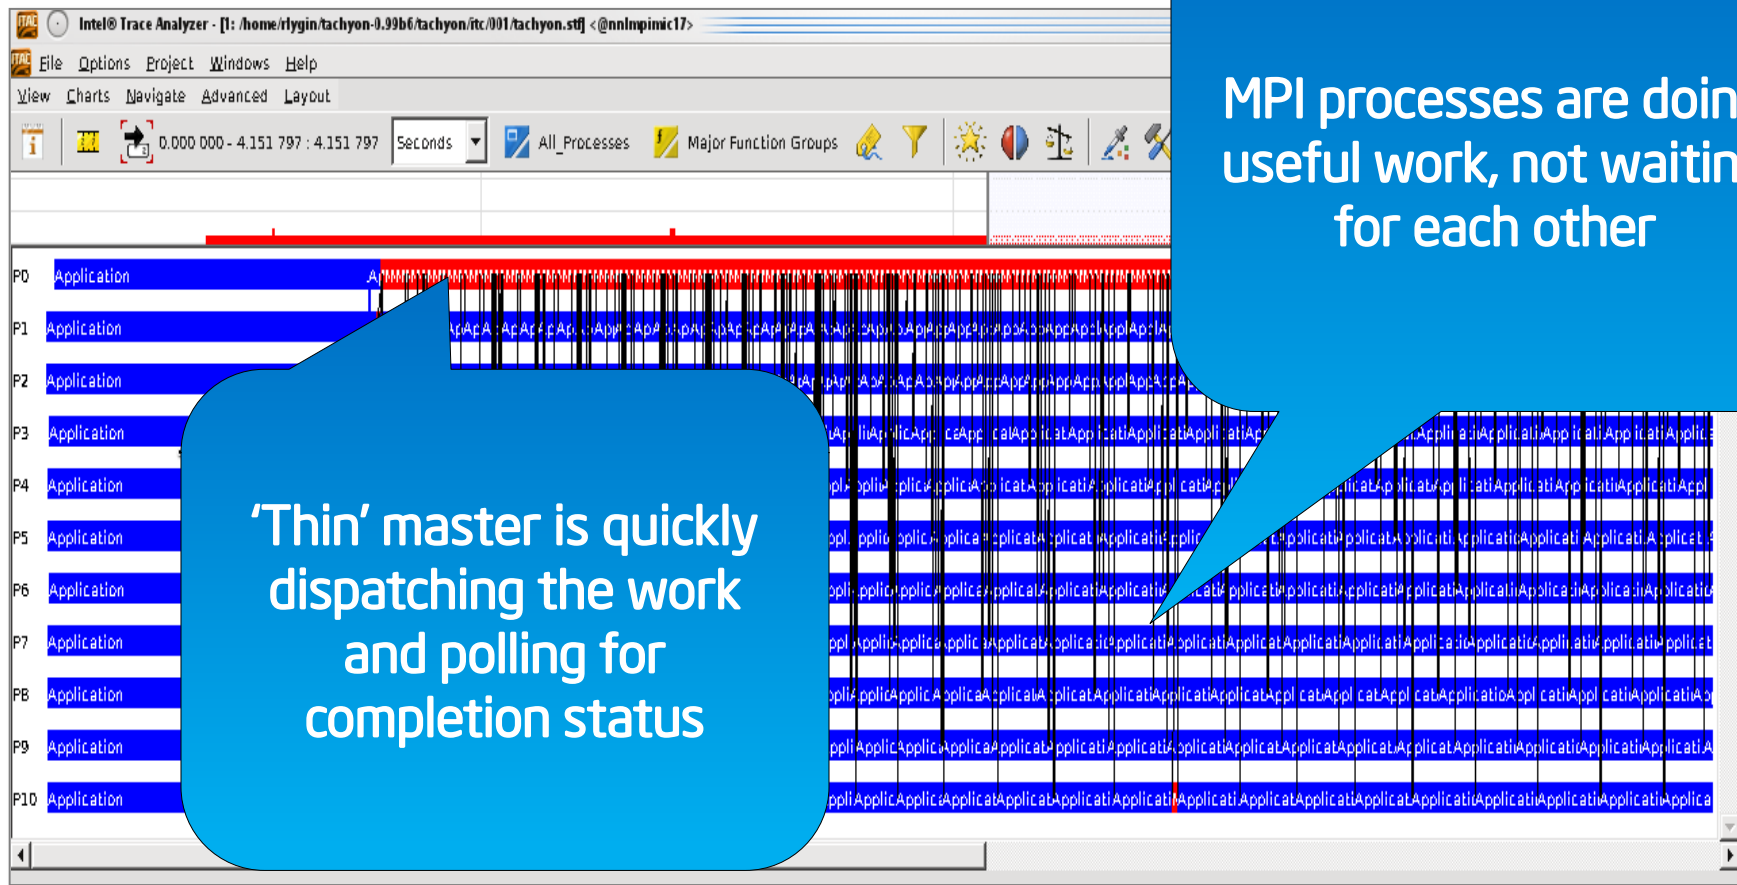
Improves both Xeon Phi and Xeon-only !

# Code change

- Producer-consumer like algorithm
- New algorithm – ~250 lines in main loop
- Not Xeon Phi specific: could be implemented to address limited Xeon scalability. Xeon Phi just triggered it.
- This is important: you optimize for Xeon, benefit everywhere!

Non-trivial but not a rocket science. Double ROI

# Re-running Intel Trace Analyzer and Collector



# Re-running Intel Trace Analyzer and Collector (cont'ed)

Each Xeon process (P1 and P2) processes 2x data of each Xeon Phi process (P3-P10).

Processes are no longer gated by each other

Total Data Volume [B] (Sender by Receiver)

	PD	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	Sum	Mean	
PD		240	232	100	104	104	104	100	104	104	100	1.31 k	131	55 M
P1	46.4 M											46.4 M	46.4 M	49.5 M
P2	48.8 M											48.8 M	48.8 M	44 M
P3	18.9 M											18.9 M	18.9 M	38.5 M
P4	19.7 M											19.7 M	19.7 M	33 M
P5	19.7 M											19.7 M	19.7 M	27.5 M
P6	19.7 M											19.7 M	19.7 M	22 M
P7	18.9 M											18.9 M	18.9 M	16.5 M
P8	19.7 M											19.7 M	19.7 M	11 M
P9	19.7 M											19.7 M	19.7 M	5.5 M
P10	18.9 M											18.9 M	18.9 M	
Sum	250 M	240	232	100	104	104	104	100	104	104	100	250 M		

## #2. Improve OpenMP parallelism



Create parallel slack by reducing chunk size: from a line to a few pixels.

- $\geq$  cache line (to avoid false sharing)

Keep dynamic scheduling  
(OMP\_SCHEDULE=dynamic)

---

Enables massive parallelism (# of chunks  $\gg$  HW threads)

Compensates different line complexities

- Also helps on Xeon



# Code change

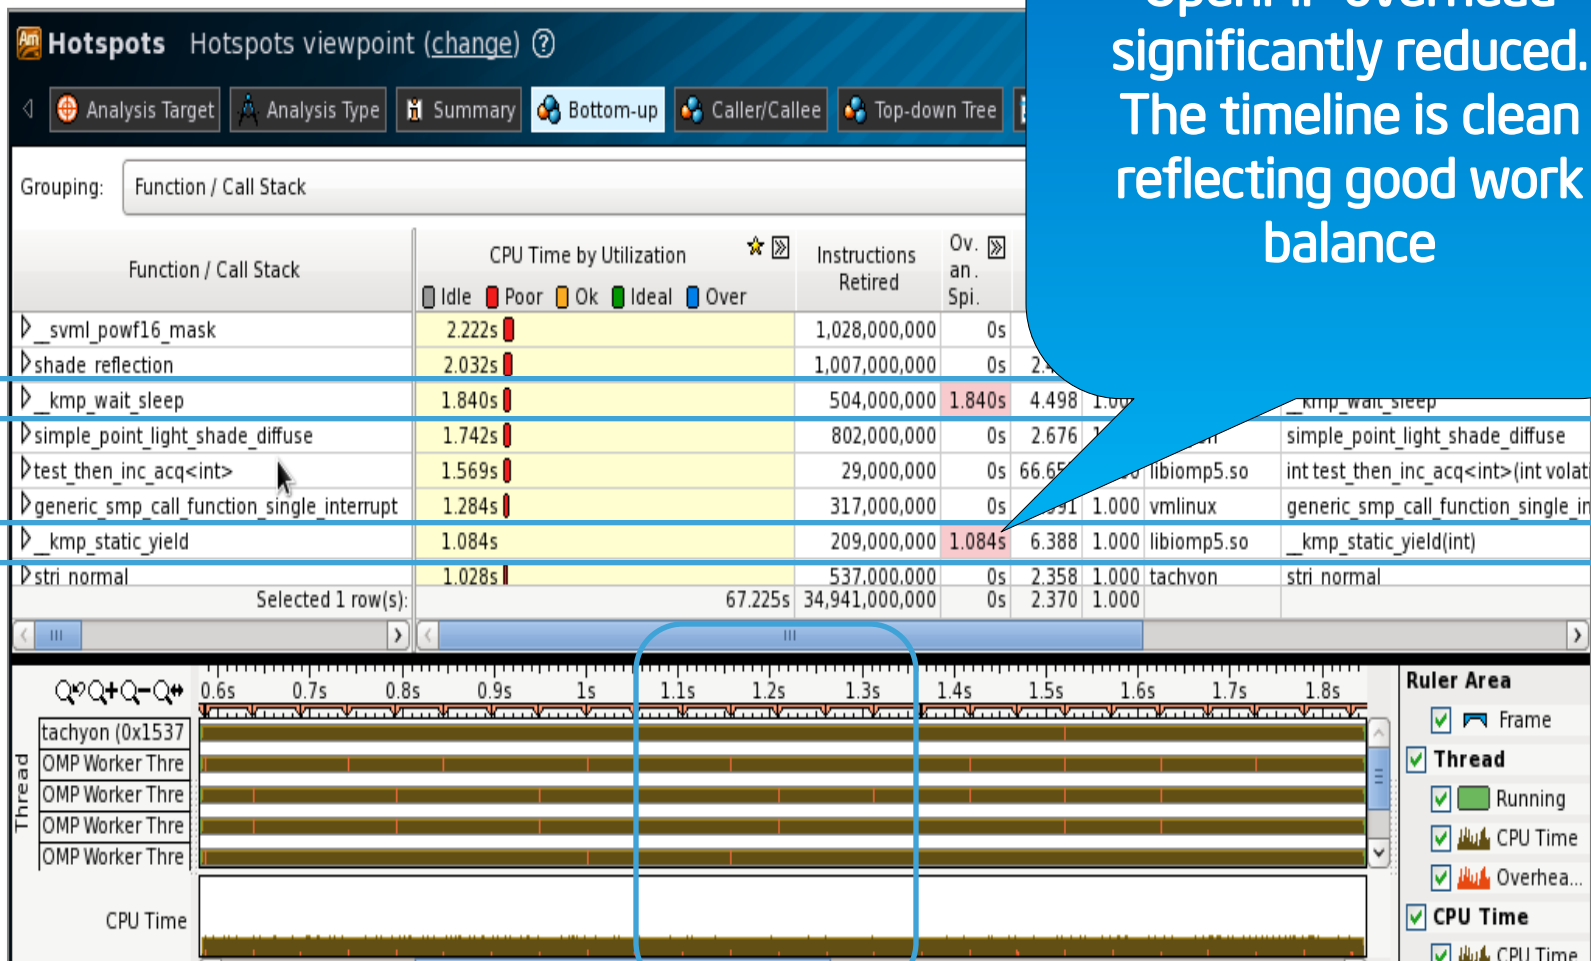
```
#pragma omp for schedule(runtime) nowait
#if defined(SINGLE_VAR_LOOP)
    for (p = 0; p < total_pixel; p += grain_size) {
        for (pp = 0; pp < grain_size; pp++) {
            int tp = p + pp;
            y = starty + (tp / xcount) * yinc;
            x = startx + (tp % xcount) * xinc;
            addr = hsize * (y - 1) + (3 * (x - 1));
        }
    }
#else /* SINGLE_VAR_LOOP */
    for (y=starty; y<=stopy; y+=yinc) {
        addr = hsize * (y - 1) + (3 * (startx - 1));
        for (x=startx; x<=stopx; x+=xinc, addr+=hskip) {
            primary.frng = cachefrng; /* each pixel uses
            col=scene->camera.cam_ray(&primary, x, y);
        }
    }
#endif
```

6 new lines  
– an  
OpenMP  
for-loop by  
pixel #,  
instead of  
by line #

Straightforward change. The same parallel model – OpenMP. Again, double ROI

# Re-running with Amplifier X

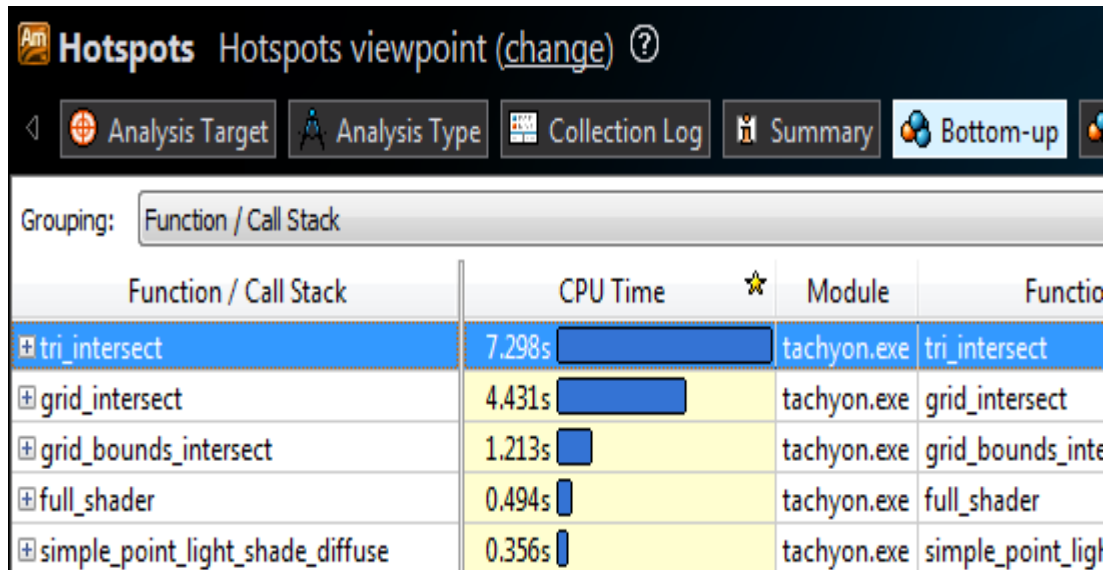
OpenMP overhead significantly reduced. The timeline is clean reflecting good work balance



# #3. Exploiting SIMD (Single Instruction Multiple Data)

How to utilize vectorization when:

- there are no loops in a hotspot function (*tri\_intersect*) ?
- the hotspot function is called on a linked list (*grid\_intersect*) ?



Hotspots Hotspots viewpoint (change) ?

Analysis Target Analysis Type Collection Log Summary Bottom-up

Grouping: Function / Call Stack

Function / Call Stack	CPU Time	Module	Function
tri_intersect	7.298s	tachyon.exe	tri_intersect
grid_intersect	4.431s	tachyon.exe	grid_intersect
grid_bounds_intersect	1.213s	tachyon.exe	grid_bounds_inte
full_shader	0.494s	tachyon.exe	full_shader
simple_point_light_shade_diffuse	0.356s	tachyon.exe	simple_point_lig

```
static void tri_intersect(const tri * trn, ray * ry) {
    vector tvec, pvec, qvec;
    flt det, inv_det, t, u, v;

    /* begin calculating determinant - also used to calc
    CROSS(pvec, ry->d, trn->edge2);

    /* if determinant is near zero, ray lies in plane of
    det = DOT(trn->edge1, pvec);

    if (det > -EPSILON && det < EPSILON)
        return;

    inv_det = 1.0 / det;

    /* calculate distance from vert0 to ray origin */
    SUB(tvec, ry->o, trn->v0);

    /* calculate U parameter and test bounds */
    u = DOT(tvec, pvec) * inv_det;
    if (u < 0.0 || u > 1.0)
        return;

    /* prepare to test V parameter */
    CROSS(qvec, tvec, trn->edge1);

    /* calculate V parameter and test bounds */
    v = DOT(ry->d, qvec) * inv_det;
    if (v < 0.0 || u + v > 1.0)
        return;


    /* calculate t, ray intersects triangle */
    t = DOT(trn->edge2, qvec) * inv_det;

    ry->add_intersection(t, (object *) trn, ry);
}
```

# Code change – new data structures

## Composite triangles:

### Composite triangles



V1: {x1, x2, ..., xn}  
{y1, y2, ..., yn}  
{z1, z2, ..., zn}

V2: ...

V3: ...

- SSE: 4 triangles, AVX: 8, Xeon Phi: 16
- Structure Of Arrays: register containing 4/8/16 float coordinates (x, y or z)
- Bit mask to describe 'real'/'void' triangles


## A small library of vector operations (+, -, dot-, cross-product, ...) using intrinsics

- Reused from Embree for SSE/AVX, extended for Phi

## Single C++ template intersection (et al) function

- No code duplication

Again, double ROI - improves both Xeon Phi and Xeon !



# Code change (cont'ed)

```
struct ssef
{
    enum { size = 4 }; // number of SIMD elements
    union { __m128 m128; float v[4]; int i[4]; }; // data

    __forceinline ssef ( ) {}
    __forceinline ssef ( const ssef& other ) { m128 = other.m128; }
    __forceinline ssef& operator=( const ssef& other ) { m128 = other.m128; return *this; }

    __forceinline ssef( const __m128& a ) : m128(a) {}
    __forceinline operator const __m128& ( void ) const { return m128; }
    __forceinline operator __m128& ( void ) { return m128; }

    __forceinline explicit ssef( const float* const a ) : m128(_mm_loadu_ps(a)) {}
    __forceinline ssef ( const float& a ) : m128(_mm_castps128_ps(_mm_shuffle_ps(_mm_set1_ps(a), _mm_set1_ps(0), 0))) {}

    struct avxf
    {
        enum { size = 8 }; // number of SIMD elements
        union { __m256 m256; float v[8]; }; // data

        __forceinline avxf ( ) {}
        __forceinline avxf ( const avxf& other ) : m256 (other.m256) {}
        __forceinline const avxf operator +( const avxf& a ) { return a; }
        __forceinline const avxf operator -( const avxf& a ) {
            const __m256 mask = _mm256_castsi256_ps(_mm256_set1_epi32(0x80000000));
            return _mm256_xor_ps(a.m256, mask);
        }
        __forceinline const avxf abs ( const avxf& a ) {
            const __m256 mask = _mm256_castsi256_ps(_mm256_set1_epi32(0x7fffffff));
            return _mm256_and_ps(a.m256, mask);
        }
        __forceinline const avxf sign ( const avxf& a ) { return _mm256_blendv_ps(a.m256, _mm256_set1_ps(0), a.m256); }
        __forceinline const avxf signmsk ( const avxf& a ) { return _mm256_and_ps(a.m256, _mm256_set1_ps(0x7fffffff)); }

        __forceinline const avxf rcp ( const avxf& a ) {
```

```
template<typename T>
static void tri_simd_intersect(const T* trn, ray * ry) {
    typedef typename T::value_type value;
    typedef typename T::scalar_type scalar;

    /* begin calculating determinant - also used to calculate U parameter */
    const value D (scalar (ry->d.x), scalar (ry->d.y), scalar (ry->d.z));
    const value pvec = cross(D, trn->edge2);

    /* if determinant is near zero, ray lies in plane of triangle */
    const scalar det = dot(trn->edge1, pvec);

    const scalar absDet = abs(det);
    typename T::boolean_type valid = trn->defined & (absDet >= helper<scalar>::epsilon);
    if (none(valid))
        return;

    const scalar inv_det = rcp(det);

    /* calculate distance from vert0 to ray origin */
    const value O (scalar (ry->o.x), scalar (ry->o.y), scalar (ry->o.z));
    const value tvec = O - trn->v0;

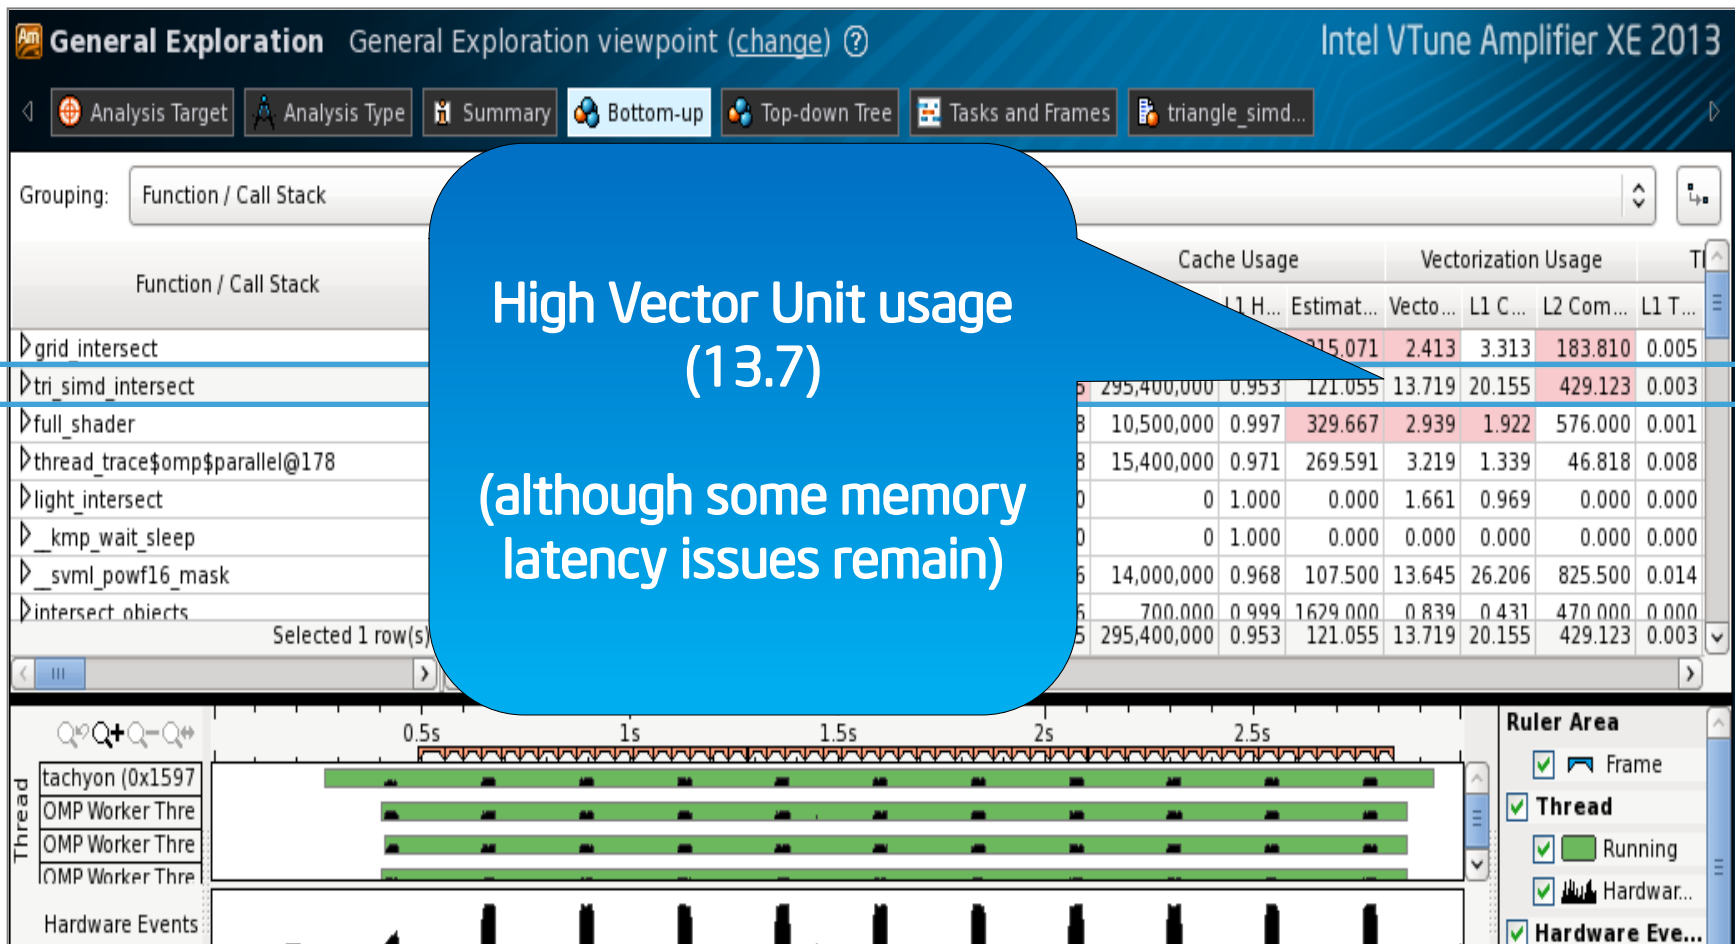
    /* calculate U parameter (without test bounds) */
    const scalar U = dot(tvec, pvec) * inv_det;
```

# SIMD benefits

- One intersection with multiple triangles at once
- Approach can be used for multi-rays intersections
  - used by Embree and Autodesk's ray tracer
- Small extra overhead during scene load (each grid cell rebuilds its list of simple triangles to composites) but benefit in heavy computations
- Intrinsics can be replaced with direct loops and compiler's auto-vectorization to improve portability

Again, double ROI - improves both Xeon Phi and Xeon !

# Re-running with Amplifier XE



# Updated results



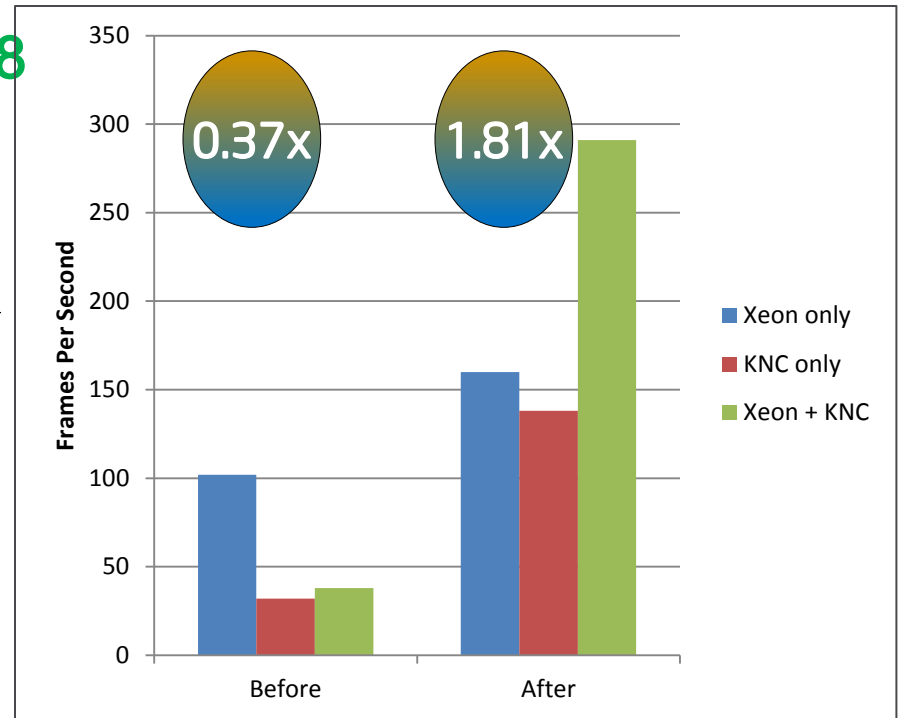
4 nodes x 2 SNB –  
102 → **160** FPS



4 nodes x 1 KNC – 32 → **138**  
FPS



4 nodes x (2SNB + 1KNC) – 38 → **291** FPS



Speed up on both Xeon  
and Xeon Phi

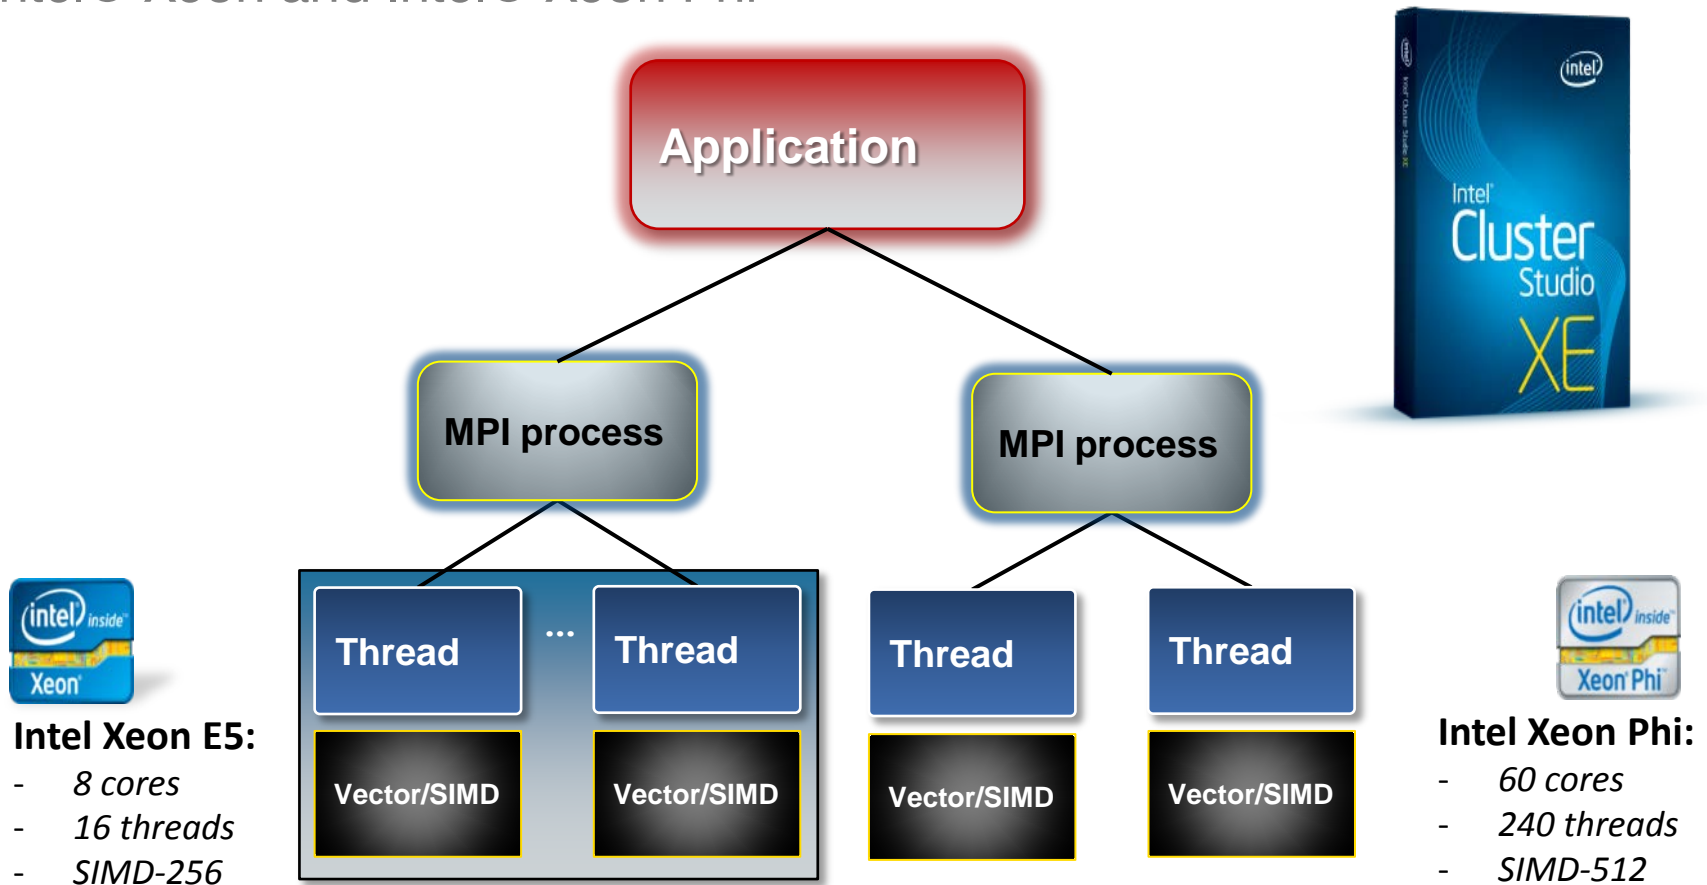
**7X speed up!**

Xeon and Xeon Phi add to each other



# Parallel programming for Intel architecture

Intel® Xeon and Intel® Xeon Phi™



Parallelism at all levels, with Intel software tools. Maximize your ROI!

# Next steps

Experiment with prefetching

Replace intrinsics with plain C and rely on vectorization by compiler

Experiment with replacing linked lists with arrays

Fine-tune with affinity settings (e.g. KMP\_AFFINITY=balanced)

# Summary

The application must meet certain criteria to benefit from Xeon Phi

You might need to apply reasonable efforts to achieve that

Good news:

- You can optimize for Xeon and benefit on Xeon Phi, and vice versa
- You use the same tools and programming models, same code




# Cloverleaf demo

Port to Intel® Xeon Phi™ with Intel® Cluster Studio XE



# What is Cloverleaf?




*Snapshots of a Cloverleaf simulation*

- Small open source benchmark
- Two dimensional compressible Euler equations across Cartesian grid
- Fortran framework
  - Fortran kernels
  - ANSI-C *kernels*
- Programmer's playground
  - CUDA
  - OpenMP
  - MPI
  - OpenCL
  - OpenACC

<http://warwick-pcav.github.com/CloverLeaf/>

# A Partitioned Data Set



- On Xeons,
  - simulation space split into 16 regions
  - Used MPI to run on multiple cores
- On Xeon Phi
  - 60 MPI tasks
  - 4 OpenMP thread

# What steps did you take to become Intel® Xeon Phi™ ready?

- At first we had no access to any MIC hardware
- So, development was carried out on a regular eight-core workstation PC
- We knew that to make best use of the Intel® Xeon Phi™ coprocessor, our code should be
  - well parallelized,
  - well vectorized.

# Anything went particularly well?

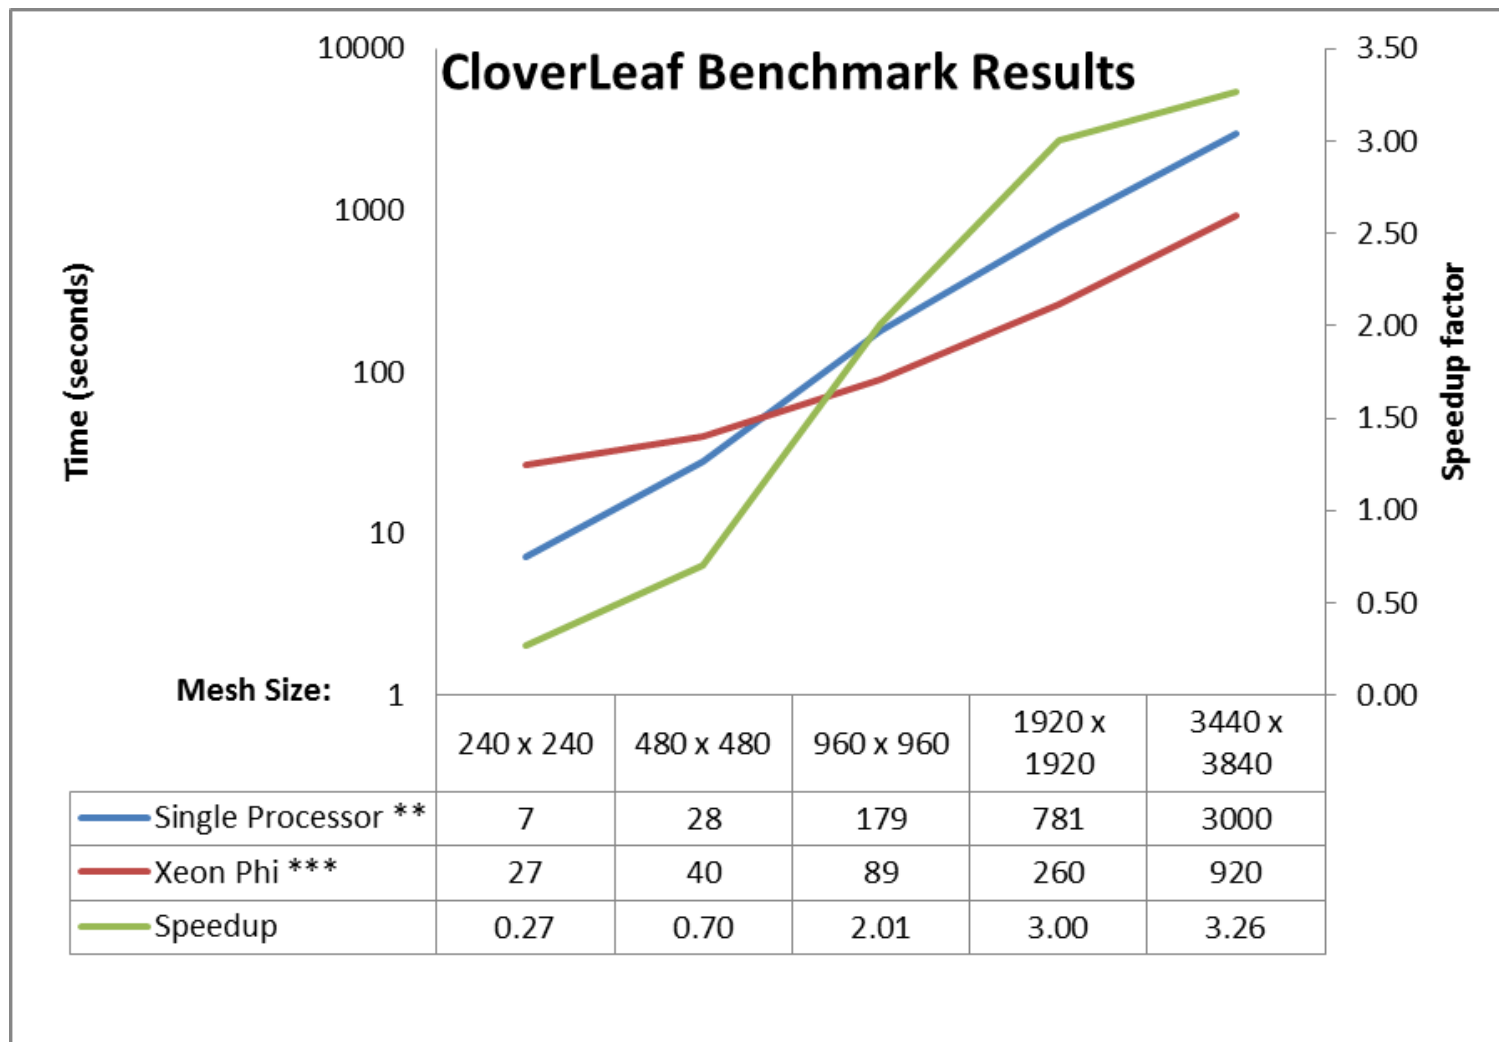
“We were surprised at how easy it was to get the first version of CloverLeaf running on the Intel® Xeon Phi™ coprocessor. We simply recompiled the existing code using the -mmic compiler option and ran executable natively on the coprocessor”



# What was the most difficult hurdle?

- Code originally not written with much consideration to how well it would vectorise
  - Compiler reports helped
  - Compiler reports were sometime confusing, with multiple message relating to one line
  - Has to resort to looking at assembler

# Speedup



\*\*Single socket eight core Intel® Xeon® E5-2687W processor, 3.1GHz, 32GB DDR3(1333Mhz) memory, with both Turbo boost and Hyperthreading enabled. \*\*\* Intel® Xeon Phi™ coprocessor had 61 cores, running at 1.09GHz, with 8GB of GDDR5 (5.5 GT/s) memory.

# Using Intel Trace Analyzer Collector (ITAC)








Name	△	TSelf	TSelf	TTotat	#Calls	TSelf /Call
Group All_Processes						
Group Application		934.354 s		1.02127e+3 s	60	15.5726 s
Group MPI		86.9119 s		86.9119 s	1089232	79.7919e-6 s

*Running Cloverleaf natively on Xeon Phi with a 3840 x 3840 mesh size.*

Group All_Processes						
Name	△	TSelf	TSelf	TTotat	#Calls	TSelf /Call
Group All_Processes						
Group Application		123.399 s		226.459 s	30	4.1133 s
Group MPI		103.06 s		103.06 s	4527860	22.7613e-6 s

*Running Cloverleaf natively on Xeon Phi with a 250 x 250 mesh size.*

# Overhead is spread over a number of MPI APIS.

[-] Group All Processes							
Group Application	123.399 s		226.459 s	30	4.1133 s		
MPI_Comm_size	568e-6 s		568e-6 s	30	18.9333e-6 s		
MPI_Comm_rank	1.108e-3 s		1.108e-3 s	30	36.9333e-6 s		
MPI_Finalize	11.7121 s		11.7121 s	30	390.402e-3 s		
MPI_Isend	17.7572 s		17.7572 s	1702750	10.4285e-6 s		
MPI_Irecv	9.64975 s		9.64975 s	1702750	5.66716e-6 s		
MPI_Waitall	42.3399 s		42.3399 s	1042500	40.6138e-6 s		
MPI_Barrier	2.5413 s		2.5413 s	270	9.41224e-3 s		
MPI_Reduce	466.145e-3 s		466.145e-3 s	11550	40.3589e-6 s		
MPI_Allreduce	18.5919 s		18.5919 s	67950	273.612e-6 s		

*The breakdown of MPI calls on the 250 x 250 mesh size.*

Not enough work is being done  
in each MPI task









# Two essential tools


Intel® VTune Amplifier XE

Intel® Trace analyzer and Collector (ITAC)



# Intel® Cluster Studio XE

Phase	Product	Feature
Build	 <b>Intel® Advisor XE</b>	Threading design assistant
	 <b>Intel® Composer XE</b>	<ul style="list-style-type: none"> <li>• C/C++ and Fortran compilers</li> <li>• Intel® Threading Building Blocks</li> <li>• Intel® Cilk™ Plus</li> <li>• Intel® Integrated Performance Primitives</li> <li>• Intel® Math Kernel Library</li> </ul>
	 <b>Intel® MPI Library</b>	High Performance Message Passing (MPI) Library
Verify & Tune	 <b>Intel® VTune™ Amplifier XE</b>	Performance Profiler for optimizing application performance and scalability
	 <b>Intel® Inspector XE</b>	Memory & threading dynamic analysis for code quality
	 <b>Intel® Trace Analyzer &amp; Collector</b>	<p>Static Analysis for code quality</p> <p>MPI Performance Profiler for understanding application correctness &amp; behavior</p>



Efficiently Produce Fast, Scalable and Reliable Applications. *Including on Xeon Phi*



# VTune



# Intel® VTune™ Amplifier XE

## Performance Profiler

Where is my application...

### Spending Time?

Function - Call Stack ▾	CPU Time ▾
algorithm_2	3.560s
do_xform ←	3.560s
algorithm_1	1.412s
BaseThreadInitTh	0.000s

- Focus tuning on functions taking time
- See call stacks
- See time on source

### Wasting Time?

Line		MEM_LOAD... LLC_MISS
475	float rx, ry, rz =	
476	float param1 = (AP	30,000
477	float param2 = (AP	
478	bool neg = (rz < 0	

- See cache misses on your source
- See functions sorted by # of cache misses

### Waiting Too Long?

Wait Time▼				Wait Count
Idle	Poor	Ok	Ideal	
176.504s	<div><div></div><div></div><div></div></div>			18,277
84.681s	<div><div></div><div></div><div></div><div></div></div>			5,499
84.612s	<div><div></div><div></div><div></div><div></div></div>			5,489

- See locks by wait time
- Red/Green for CPU utilization during wait

- Windows & Linux
- Low overhead
- No special recompiles

Advanced Profiling For Scalable Multicore Performance



# VTune Amplifier is a simple tool

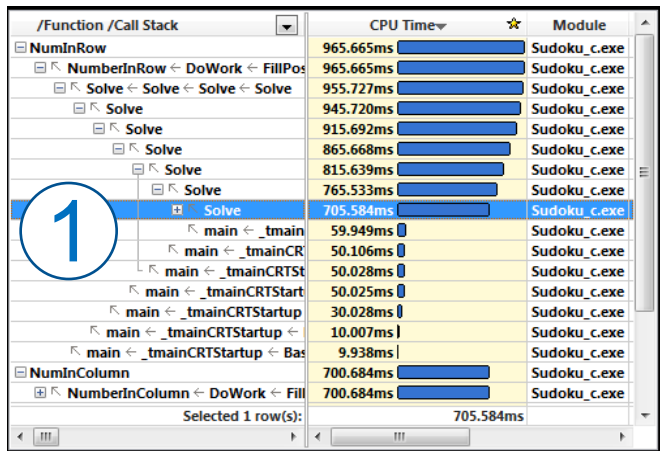
Imagine you have a cool car and you want to drive a little faster or fuel effective

All what you'd need you can find here.

VTune as other simple tools can provide basic information on performance of your engine.



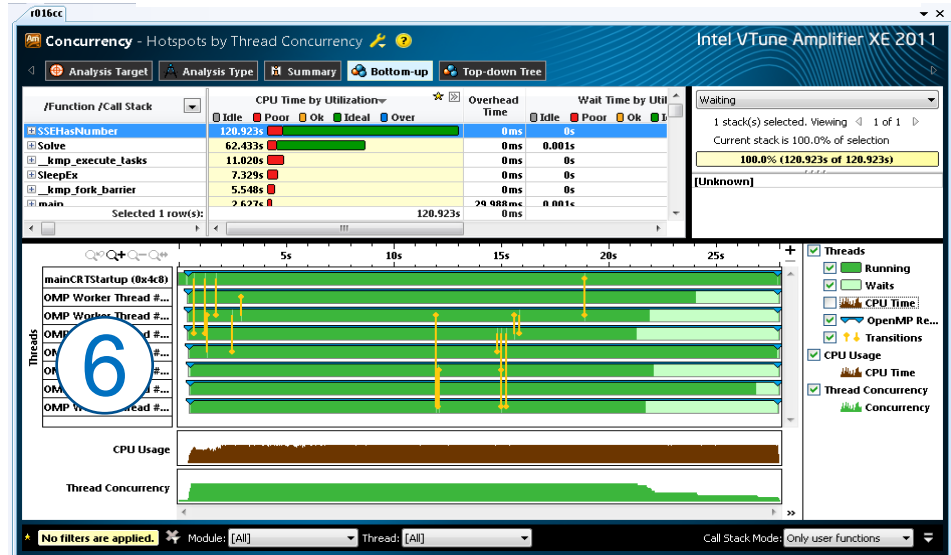
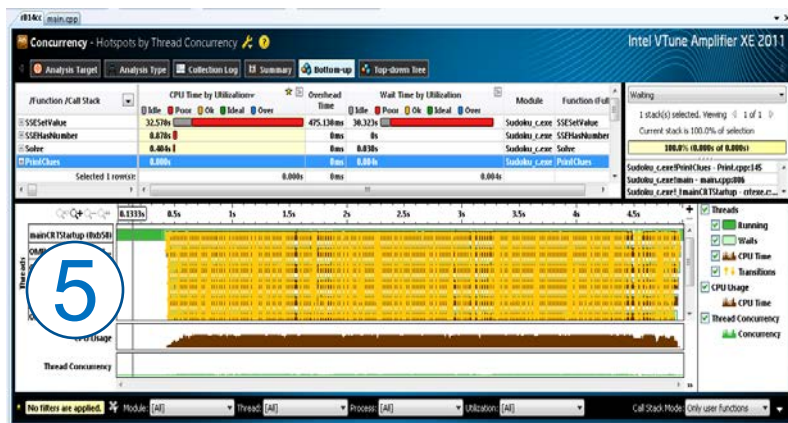

# An example of user mode analysis...



ID	Problem	Sources	Modules	State
P1	Unhandled application exception	functional	Sudoku_c.exe	New
P2	Cross-thread stack access	Generator.cpp; main.cpp	Sudoku_c.exe	New
P3	Cross-thread stack access	[Unknown]; xutility	Sudoku_c.exe; libomp5md.dll	Not fixed
P4	Data race	Generator.cpp	Sudoku_c.exe	Not fixed
P5	Data race	main.cpp	Sudoku_c.exe	Not fixed
P6	Data race	xtree	Sudoku_c.exe	Not fixed
P7	Data race	xmemory; xtree	Sudoku_c.exe	Not fixed
P8	Data race	xmemory; xtree	Sudoku_c.exe	Not fixed
P9	Data race	xmemory; xtree	Sudoku_c.exe	New
P10	Data race	xmemory; xtree	Sudoku_c.exe	New
P11	Data race	xmemory; xtree	Sudoku_c.exe	New

3

2



1. Hotspot Analysis
2. Implement
3. Find Threading Errors
- 4,5,6. Tune Parallelism


# VTune Amplifier is a complex tool

However, if you want your car to win a race...

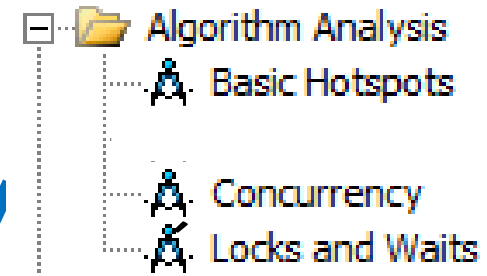
Your tools set has to be much more complex to analyze all aspects of engine functioning. You need to be more proficient in both: the tool's functionality and the engine internals!



## Profile the System



## Profile Applications

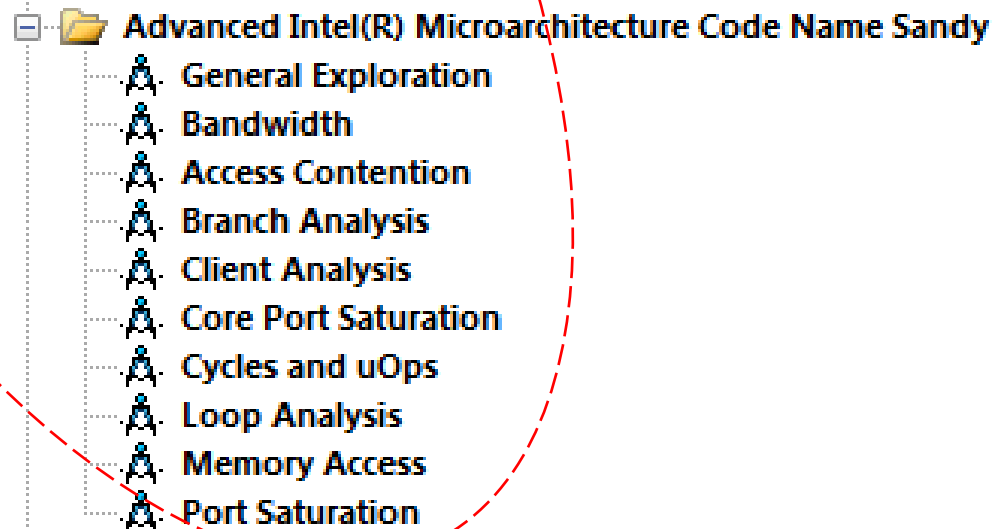


## User Mode

- Works on
- Intel
  - non-Intel
  - More overhead than lightweight hotspots

# VTune Amplifier XE

## Architectural Analysis



## Kernel Mode

- Works only on
- Intel



# An example of architectural analysis

Speedup by  
upgrading  
silicon

CPU	No Auto-Vectorisation	With Auto-Vectorisation	Speedup
<b>P4</b>	39.344	21.9	<b>1.80</b>
<b>Core 2</b>	5.546	0.515	<b>10.77</b>
<b>Speedup</b>	<b>7.09</b>	<b>45.52</b>	<b>76</b>


Speedup by swapping compiler

CPU EVENT	Without Vect	With Vect
CPU_CLK_UNHALTED.CORE	16,641,000,448	1,548,000,000
INST_RETIRED.ANY	3,308,999,936	1,395,000,064
X87_OPS_RETIRED.ANY	250,000,000	0
SIMD_INST_RETIRED	0	763,000,000

Full paper available here:

<http://edc.intel.com/Link.aspx?id=1045>

Verified  
using VTune





ITAC

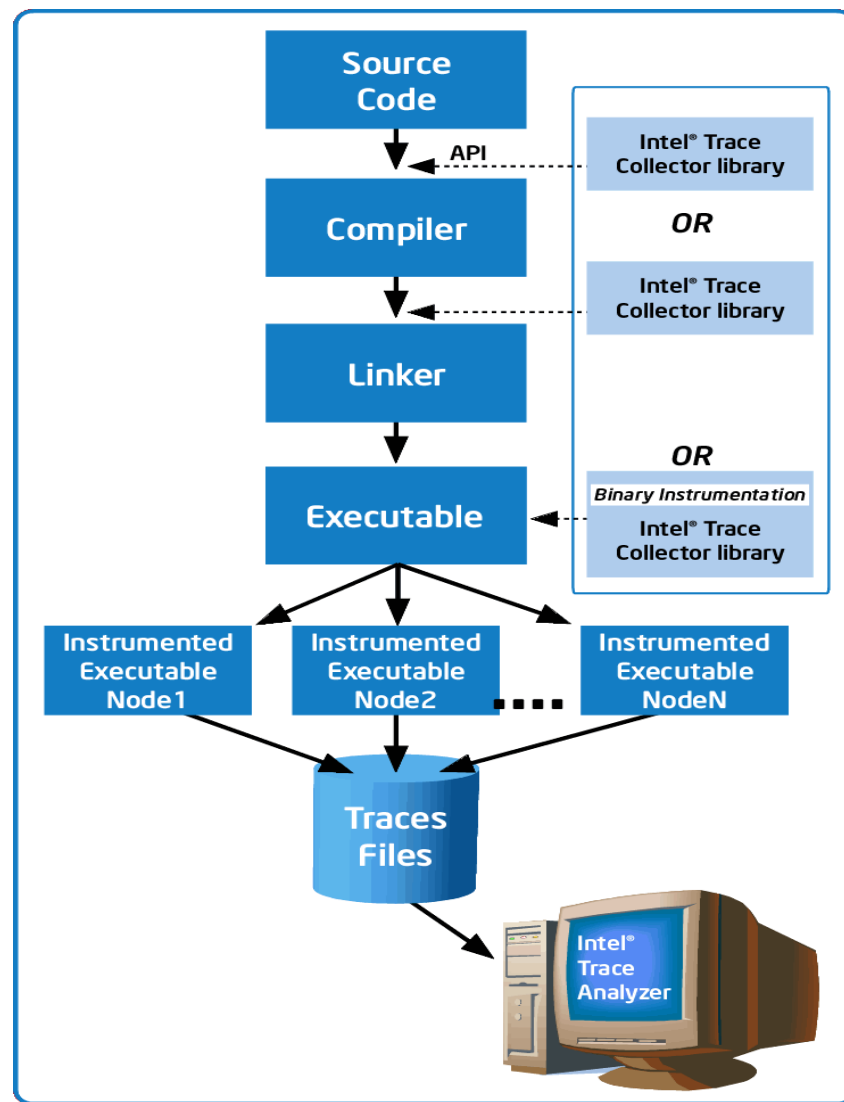




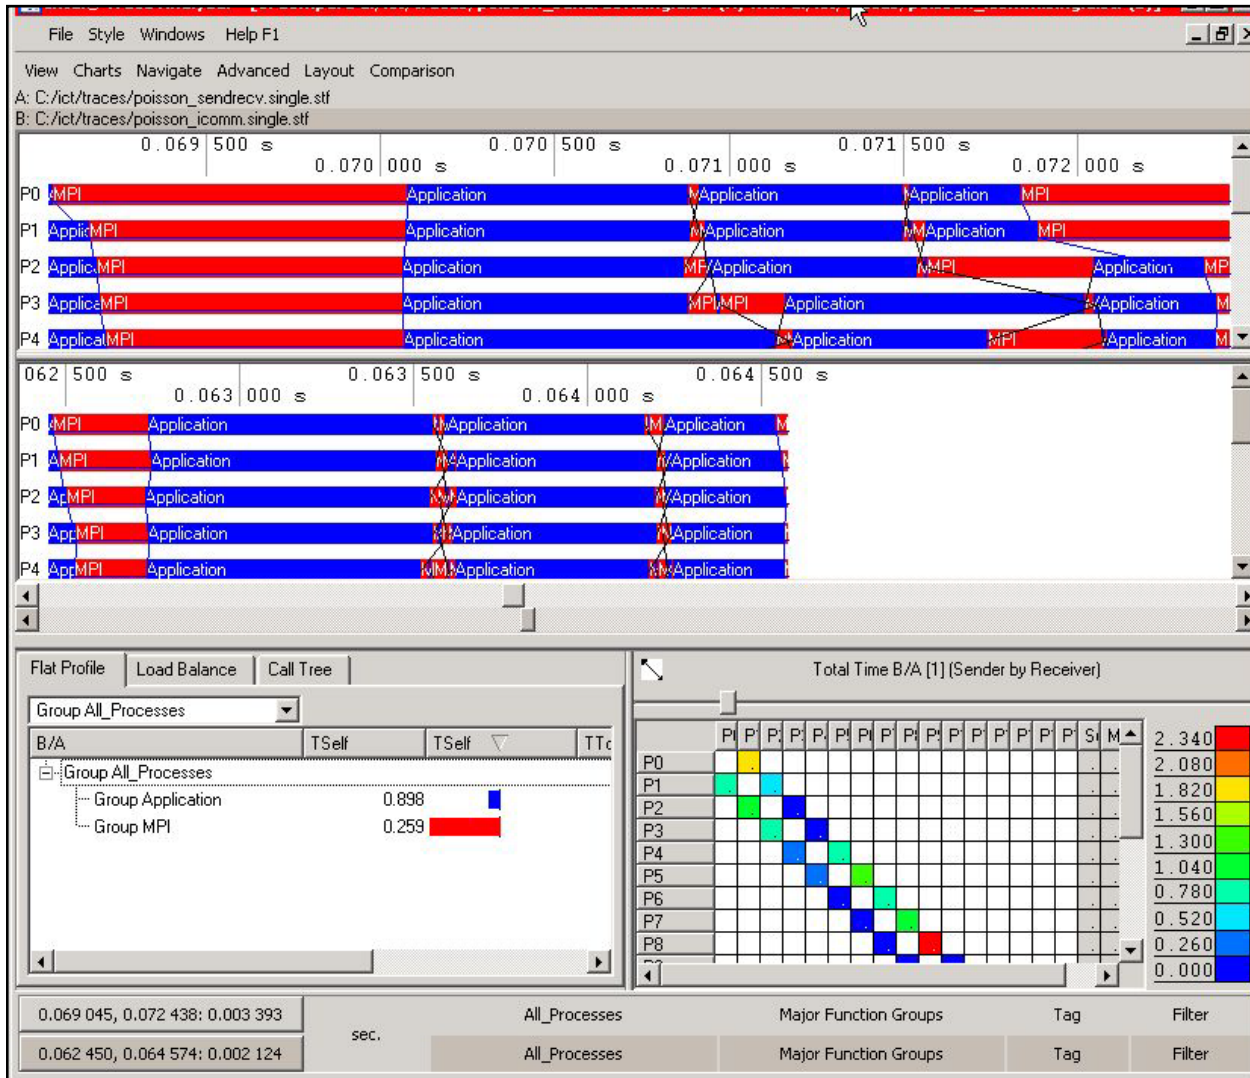
# Intel® Trace Analyzer and Collector Overview

Intel® Trace Analyzer and Collector helps the developer:

- Visualize and understand parallel application behavior
- Evaluate profiling statistics and load balancing
- Identify communication hotspots



# Intel® Trace Analyzer and Collector



Compare the event  
timelines of two  
communication profiles

Blue = computation  
Red = communication

Chart showing how the MPI processes interact




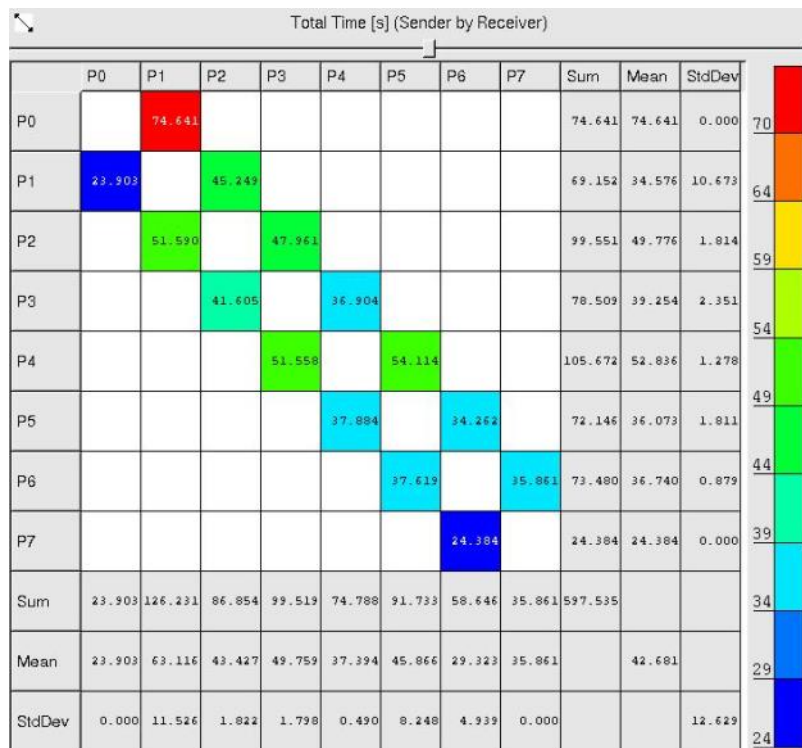
# Communication Profiles

Statistics about point-to-point or collective communication

Generic matrix supports grouping by several attributes in each dimension

Sender, Receiver, Data volume per msg, Tag, Communicator, Type

Available attributes: Count, Bytes transferred, Time, Transfer rate



# Typical Hands-on Xeon Phi training agenda

## Day 1 – Getting Ready

- 10.00 Welcome
- 10.30 Two Essential Requirements
- 11.00 Parallelism (L)
- 12.30 Lunch
- 1.30 Vectorisation (L)
- 4.00 Advance Profiling (Walkthrough)
- 5.00 End



25<sup>th</sup> & 26<sup>th</sup>  
June 2014  
Manchester

## Day 2-Xeon Phi Programming

- 09.00 Start
- 09.15 Native & Offload Programming for Xeon Phi (L)
- 11.30 A Case Study
- 12.00 Lunch
- 1.00 Vectorisation on Xeon Phi (L)
- 1.50 Parallelism on Xeon Phi (L)
- 3.40 Wrap-up
- 4.00 End