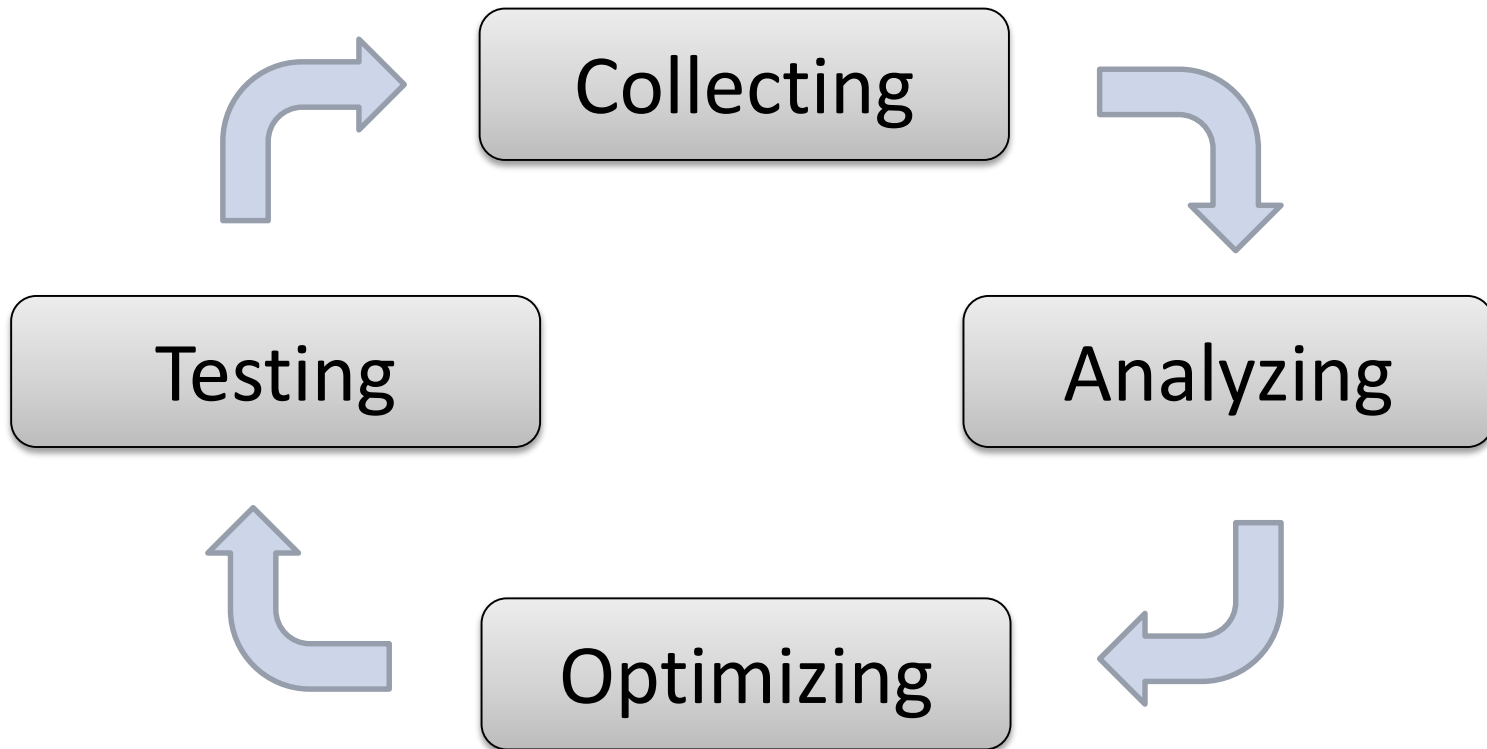


OpenMP and Performance

Christian Terboven, Dirk Schmidl
IT Center, RWTH Aachen University
Member of the HPC Group
{terboven,schmidl}@itc.rwth-aachen.de

- Performance Tuning aims to improve the runtime of an existing application.



- **A Hotspot is a source code region where a significant part of the runtime is spent.**

90/10 law

90% of the runtime in a program is spent in 10% of the code.

- **Hotspots can indicate where to start with serial optimization or shared memory parallelization.**
- **Use a tool to identify hotspots. In many cases the results are surprising.**

Performance Tools

■ Performance Analyses for

- Serial Applications
- Shared Memory Parallel Applications

■ Sampling Based measurements

■ Features:

- Hot Spot Analysis
- Concurrency Analysis
- Wait
- Hardware Performance Counter Support

- Standard Benchmark to measure memory performance.
- Version is parallelized with OpenMP.

Measures Memory bandwidth for:

y=x (copy)

y=s*x (scale)

y=x+z (add)

y=x+s*z (triad)

```
#pragma omp parallel for  
for (j=0; j<N; j++)  
    b[j] = scalar*c[j];
```

for double vectors x,y,z and scalar double value s

Function	Rate (MB/s)	Avg time	Min time	Max time
Copy:	33237.0185	0.0050	0.0048	0.0055
Scale:	33304.6471	0.0049	0.0048	0.0059
Add:	35456.0586	0.0070	0.0068	0.0073
Triad:	36030.9600	0.0069	0.0067	0.0072

- 1 Basic Analysis Types
- 2 Hardware Counter Analysis Types, choose Nehalem Architecture, on cluster-linux-tuning.
- 3 Analysis for Intel Xeon Phi coprocessors, choose this for OpenMP target programs.

The screenshot shows the 'Choose Analysis Type' dialog in Intel VTune Amplifier XE 2013. The 'Hotspots - Knights Corner Platform' analysis type is selected. The interface includes a tree view on the left, a central configuration panel, and a right sidebar with 'Start', 'Start Paused', and 'Project Properties' buttons.

Hotspots - Knights Corner Platform

Identify time-consuming code in your application. Advanced Hotspots analysis (formerly, Lightweight Hotspots) uses the kernel driver and extends the hotspots analysis by collecting call stacks, context switch and statistical call count data and analyzing the CPI (Cycles Per Instruction) metric. At the default level this analysis uses higher frequency sampling at lower overhead compared to Basic Hotspo...

List of Intel Xeon Phi coprocessor cards:

Analyze user tasks

Details

Events configured for CPU: Intel(R) Xeon(R) E5 processor

NOTE: For analysis purposes, Intel VTune Amplifier XE 2013 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.

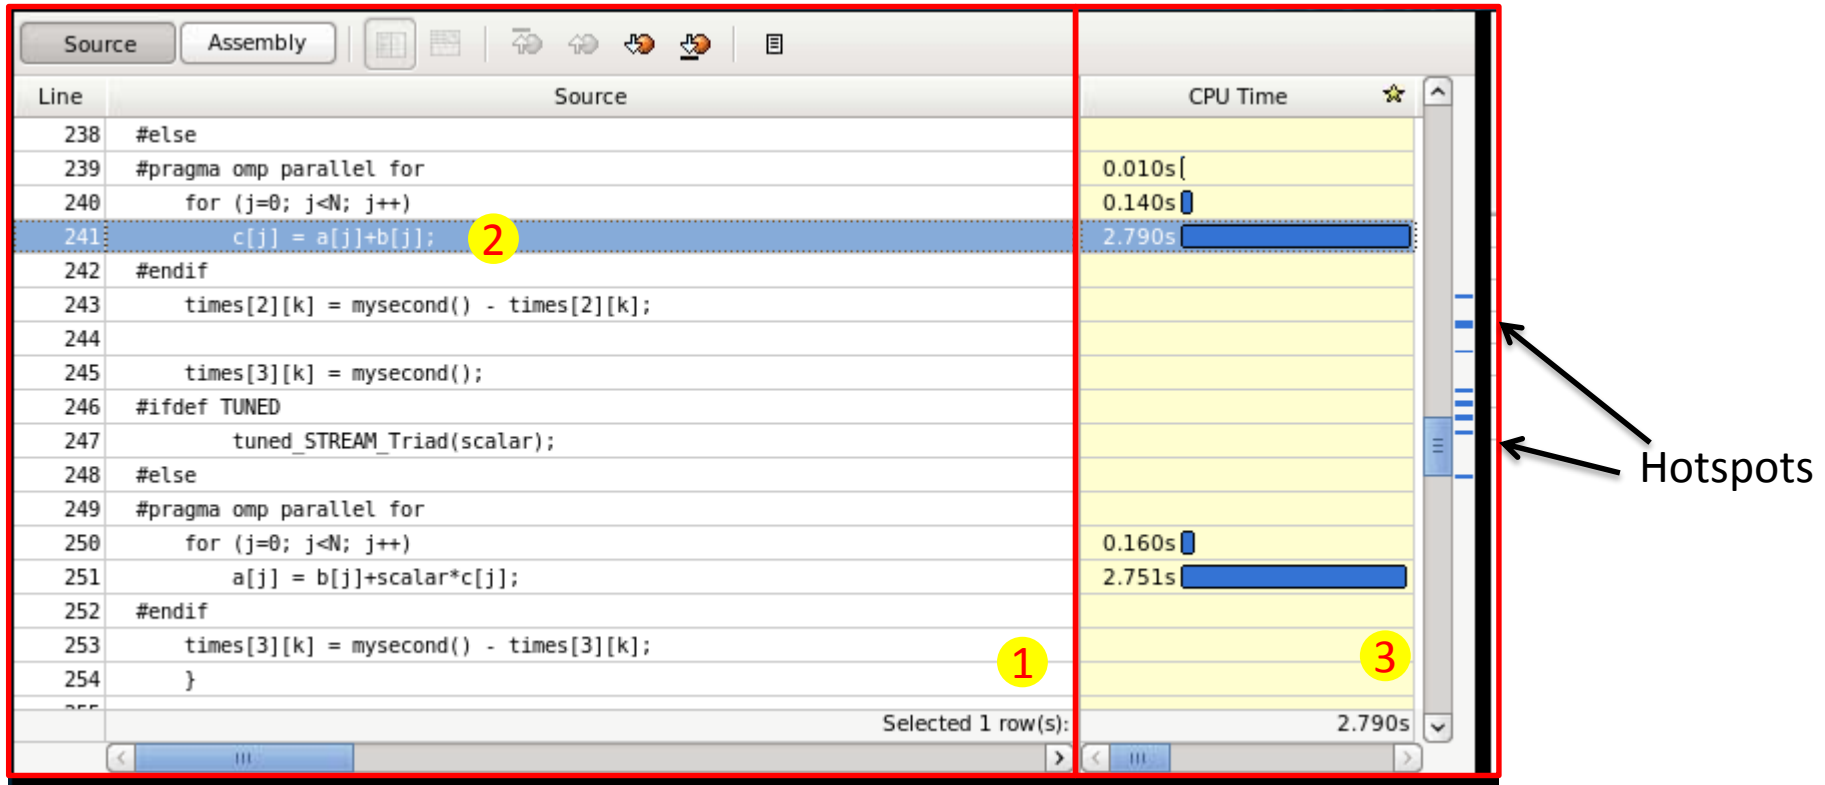
Event Name	Sample After	Event Description
CPU_CLK_UNHALTED	10000000	
INSTRUCTIONS_EXECUTED	10000000	

Amplifier XE – Hotspot Analysis



Double clicking on a function opens source code view.

- 1 Source Code View (only if compiled with -g)
- 2 Hotspot: Add Operation of Stream
- 3 Metrics View

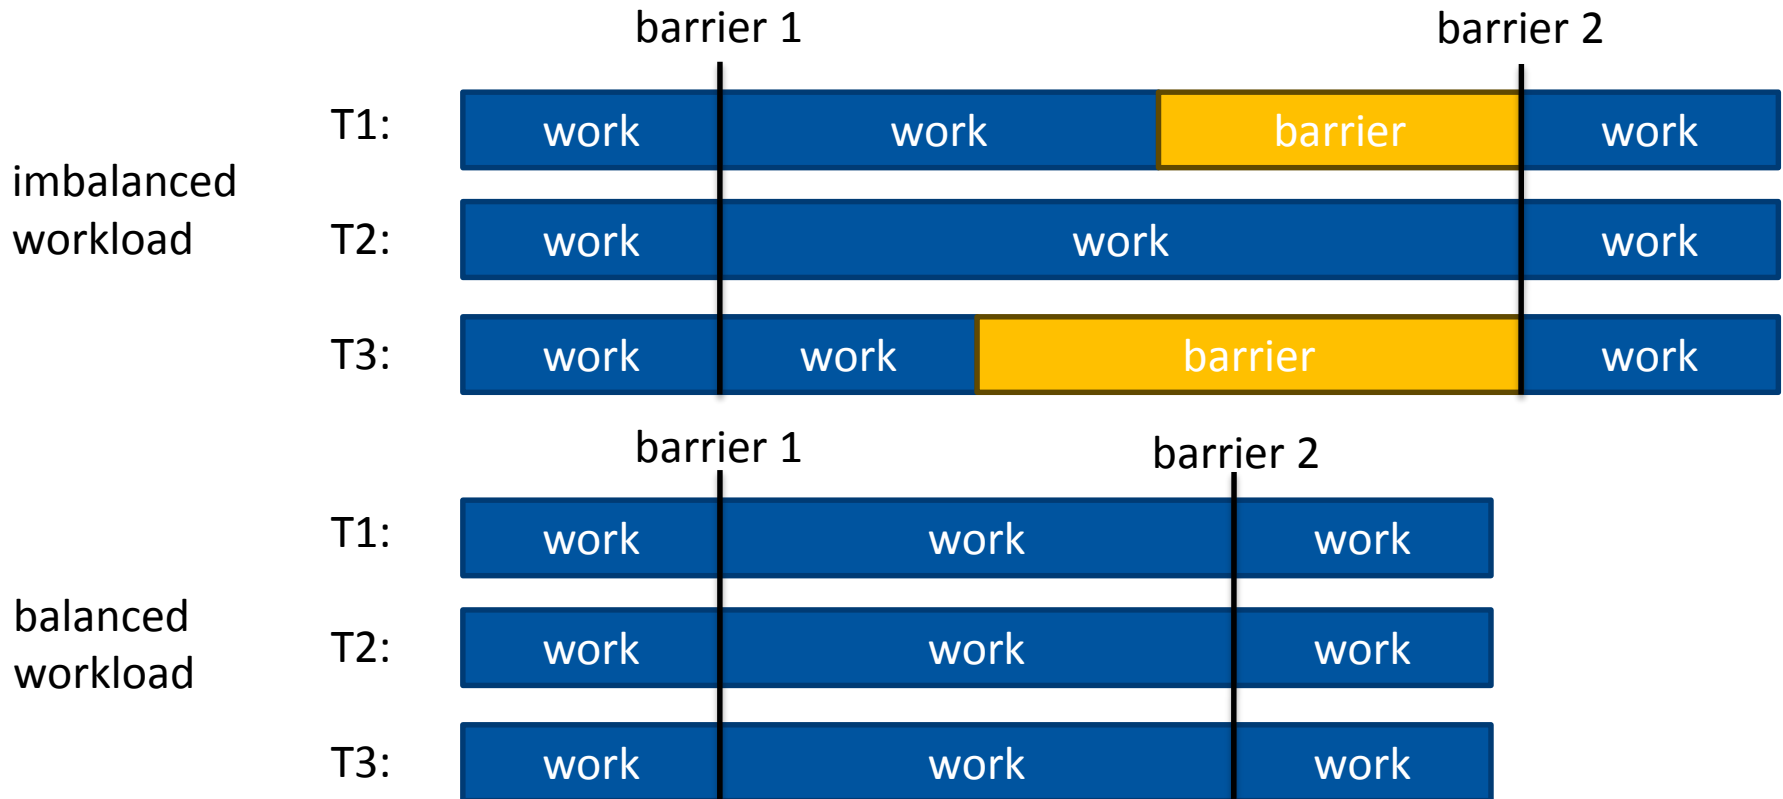


Load Balancing

■ Load imbalance occurs in a parallel program


→ when multiple threads synchronize at global synchronization points

→ and these threads need a different amount of time to finish the calculation.



■ Sparse Linear Algebra


- Sparse Linear Equation Systems occur in many scientific disciplines.
- Sparse matrix-vector multiplications (SpMxV) are the dominant part in many iterative solvers (like the CG) for such systems.
- number of non-zeros $\ll n*n$



Beijing Botanical Garden

Oben Rechts: Original Gebäude
Unten Rechts: Modell
Unten Links: Matrix

(Quelle: Beijing Botanical Garden and University of Florida, Sparse Matrix Collection)



Case Study: CG





■ $A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 2 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 4 & 0 & 4 & 4 \end{pmatrix}$

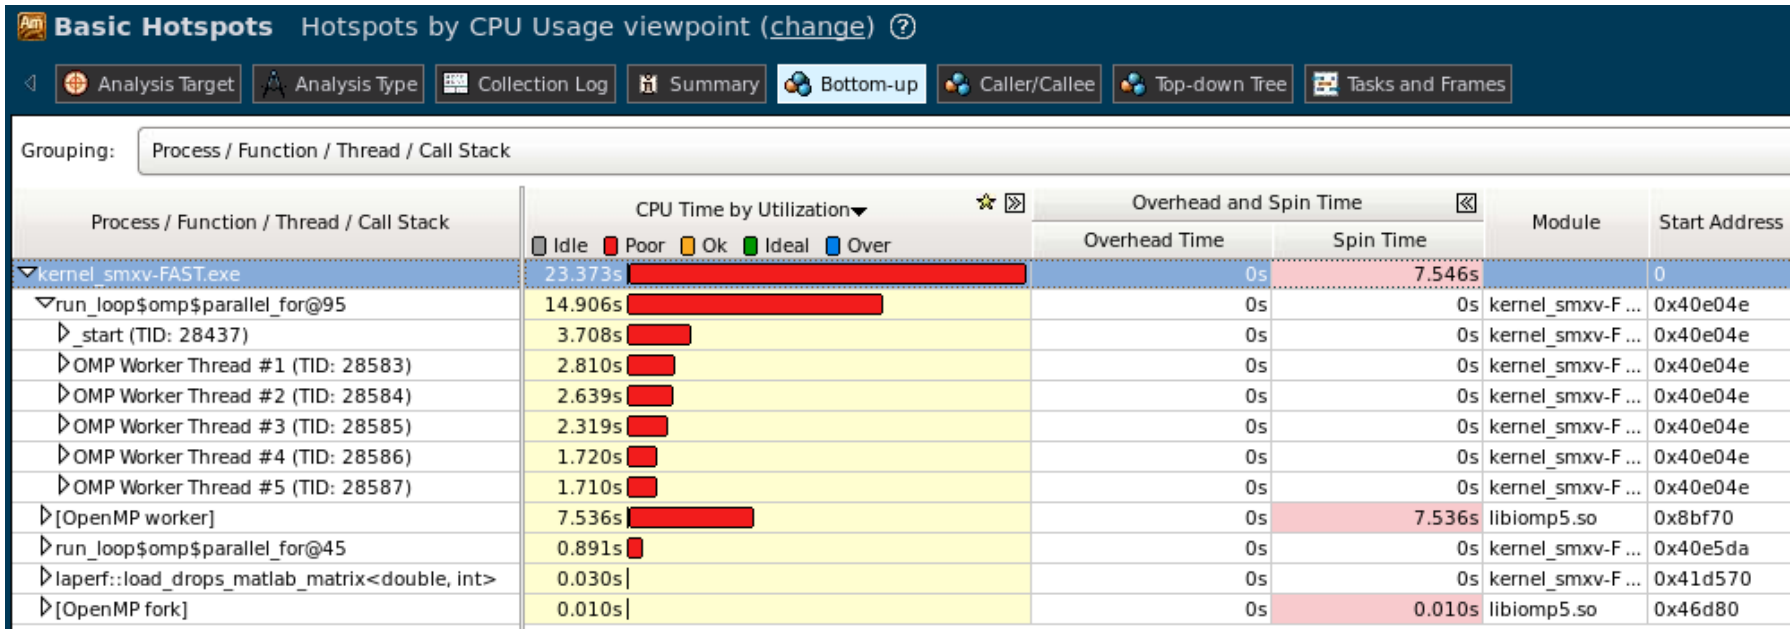
```
for (i = 0; i < A.num_rows; i++){
    sum = 0.0;
    for (nz=A.row[i]; nz<A.row[i+1]; ++nz){
        sum+= A.value[nz]*x[A.index[nz]];
    }
    y[i] = sum;
}
```

$\vec{y} = A * \vec{x}$

- Format: compressed row storage
- store all values and columns in arrays (length nnz)
- store beginning of a new row in a third array (length n+1)



- Grouping execution time of parallel regions by threads helps to detect load imbalance.
- Significant portions of Spin Time also indicate load balance problems.
- Different loop schedules might help to avoid these problems.




Basic Hotspots Hotspots by CPU Usage viewpoint (change) ?

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

Grouping: Process / Function / Thread / Call Stack

Process / Function / Thread / Call Stack	CPU Time by Utilization				Overhead and Spin Time		Module	Start Address
	Idle	Poor	Ok	Ideal	Over	Overhead Time		
kernel_smxv-FAST.exe	23.373s					0s	7.546s	0
run_loop\$omp\$parallel_for@95	14.906s					0s	0s	kernel_smxv-F ... 0x40e04e
_start (TID: 28437)	3.708s					0s	0s	kernel_smxv-F ... 0x40e04e
OMP Worker Thread #1 (TID: 28583)	2.810s					0s	0s	kernel_smxv-F ... 0x40e04e
OMP Worker Thread #2 (TID: 28584)	2.639s					0s	0s	kernel_smxv-F ... 0x40e04e
OMP Worker Thread #3 (TID: 28585)	2.319s					0s	0s	kernel_smxv-F ... 0x40e04e
OMP Worker Thread #4 (TID: 28586)	1.720s					0s	0s	kernel_smxv-F ... 0x40e04e
OMP Worker Thread #5 (TID: 28587)	1.710s					0s	0s	kernel_smxv-F ... 0x40e04e
[OpenMP worker]	7.536s					0s	7.536s	libiomp5.so 0x8bf70
run_loop\$omp\$parallel_for@45	0.891s					0s	0s	kernel_smxv-F ... 0x40e5da
laperf::load_drops_matlab_matrix<double, int>	0.030s					0s	0s	kernel_smxv-F ... 0x41d570
[OpenMP fork]	0.010s					0s	0.010s	libiomp5.so 0x46d80

- The Timeline can help to investigate the problem further.



- Zooming in, e.g. to one iteration is also possible.



Parallel Loop Scheduling

Load Balancing

- **for-construct: OpenMP allows to influence how the iterations are scheduled among the threads of the team, via the *schedule* clause:**
 - `schedule(static [, chunk])`: Iteration space divided into blocks of chunk size, blocks are assigned to threads in a round-robin fashion. If chunk is not specified: #threads blocks.
 - `schedule(dynamic [, chunk])`: Iteration space divided into blocks of chunk (not specified: 1) size, blocks are scheduled to threads in the order in which threads finish previous blocks.
 - `schedule(guided [, chunk])`: Similar to dynamic, but block size starts with implementation-defined value, then is decreased exponentially down to chunk.
- **Default on most implementations is `schedule(static)`.**



False Sharing

- **CPU is fast**

 - Order of 3.0 GHz

- **Caches:**

 - Fast, but expensive


 - Thus small, order of MB

- **Memory is slow**

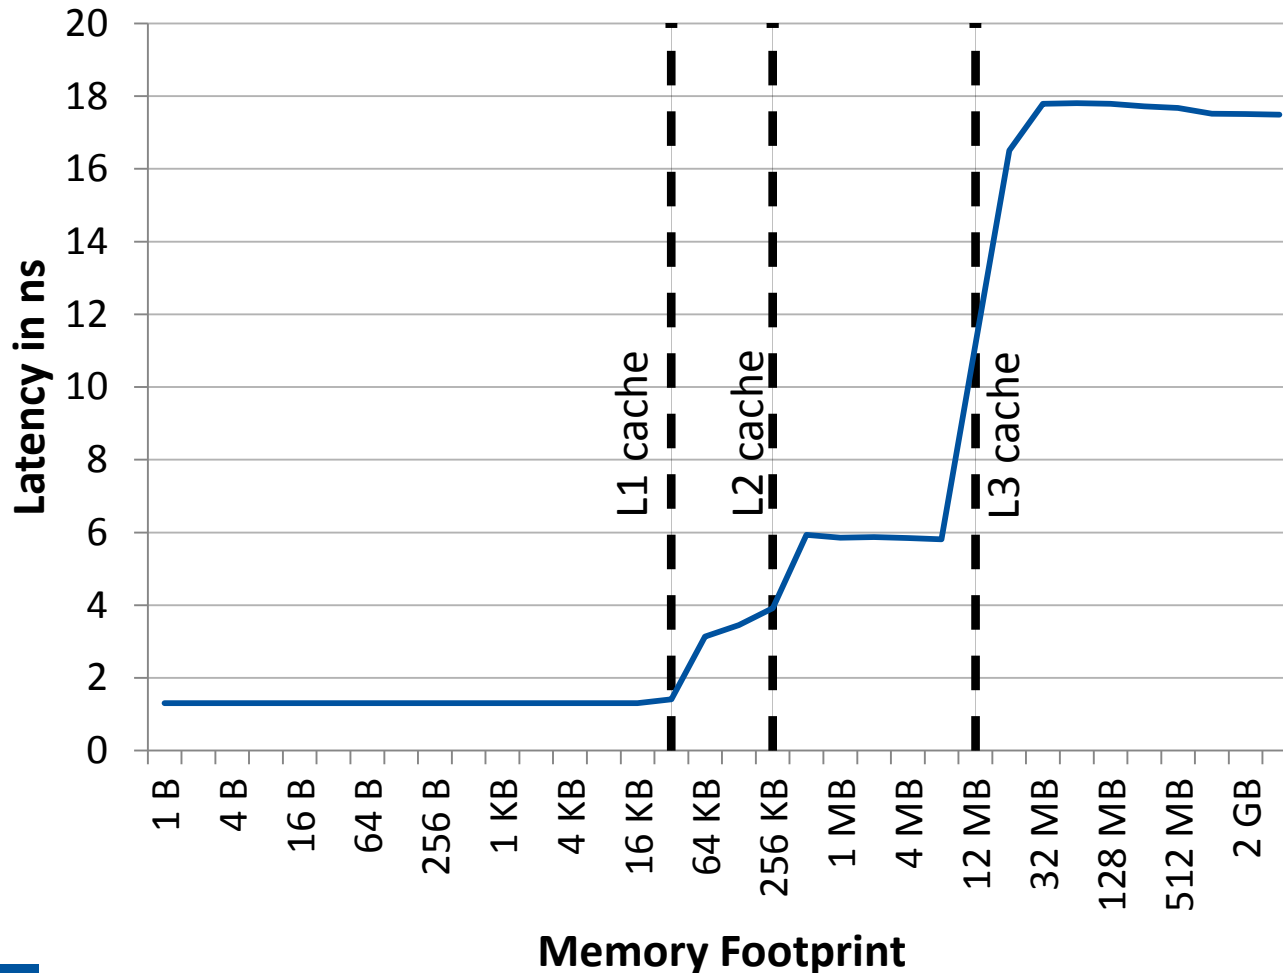
 - Order of 0.3 GHz

 - Large, order of GB


- **A good utilization of caches is crucial for good performance of HPC applications!**




Latency on the Intel Westmere-EP 3.06 GHz processor



- When data is used, it is copied into caches.
- The hardware always copies chunks into the cache, so called *cache-lines*.
- This is useful, when:
 - the data is used frequently (temporal locality)
 - consecutive data is used which is on the same cache-line (spatial locality)



- **False Sharing occurs when**
 - different threads use elements of the same cache-line
 - one of the threads writes to the cache-line
- **As a result the cache line is moved between the threads, also there is no real dependency**
- **Note: False Sharing is a performance problem, not a correctness issue**



Summing up vector elements again

It's your turn: Make It Scale!



```
#pragma omp parallel
{

#pragma omp for
  for (i = 0; i < 99; i++)
  {

      s = s + a[i];

  }

} // end parallel
```

```
do i = 0, 99
  s = s + a(i)
end do
```



```
do i = 0, 24
  s = s + a(i)
end do
```

```
do i = 25, 49
  s = s + a(i)
end do
```

```
do i = 50, 74
  s = s + a(i)
end do
```

```
do i = 75, 99
  s = s + a(i)
end do
```


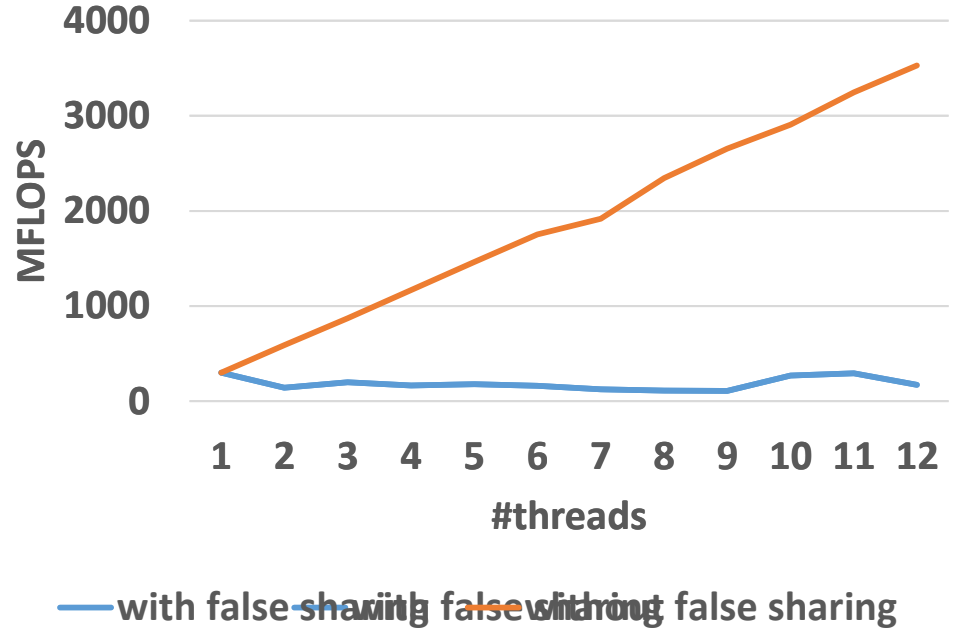

False Sharing



```
double s_priv[nthreads];  
  
#pragma omp parallel num_threads(nthreads)  
{  
    int t=omp_get_thread_num();  
  
    #pragma omp for  
    for (i = 0; i < 99; i++)  
    {  
        s_priv[t] += a[i];  
    }  
} // end parallel  
for (i = 0; i < nthreads; i++)  
{  
    s += s_priv[i];  
}
```

False Sharing

- no performance benefit for more threads
- Reason: false sharing of s_priv
- Solution: padding so that only one variable per cache line is used



False Sharing avoided



```
double s_priv[nthreads * 8];  
  
#pragma omp parallel num_threads(nthreads)  
{  
    int t=omp_get_thread_num();  
  
    #pragma omp for  
    for (i = 0; i < 99; i++)  
    {  
        s_priv[t * 8] += a[i];  
    }  
} // end parallel  
for (i = 0; i < nthreads; i++)  
{  
    s += s_priv[i * 8];  
}
```




NUMA Architectures

How To Distribute The Data ?

```
double* A;  
A = (double*)  
    malloc(N * sizeof(double));
```

```
for (int i = 0; i < N; i++) {  
    A[i] = 0.0;  
}
```



- **Important aspect on cc-NUMA systems**

- If not optimal, longer memory access times and hotspots

- **OpenMP does not provide support for cc-NUMA**

- **Placement comes from the Operating System**

- This is therefore Operating System dependent


- **Windows, Linux and Solaris all use the “First Touch” placement policy by default**

- May be possible to override default (check the docs)

- **Serial code: all array elements are allocated in the memory of the NUMA node containing the core executing this thread**


```
double* A;  
A = (double*)  
    malloc(N * sizeof(double));
```

```
for (int i = 0; i < N; i++) {  
    A[i] = 0.0;  
}
```



- **First Touch w/ parallel code: all array elements are allocated in the memory of the NUMA node containing the core executing the thread initializing the respective partition**

```
double* A;  
A = (double*)  
    malloc(N * sizeof(double));  
  
omp_set_num_threads(2);  
  
#pragma omp parallel for  
for (int i = 0; i < N; i++) {  
    A[i] = 0.0;  
}
```



- **Before you design a strategy for thread binding, you should have a basic understanding of the system topology. Please use one of the following options on a target machine:**

- Intel MPI's `cpuinfo` tool

- `module switch openmpi intelmpi`

- `cpuinfo`

- Delivers information about the number of sockets (= packages) and the mapping of processor ids used by the operating system to cpu cores.

- `hwloc`'tools

- `lstopo` (command line: `hwloc-ls`)

- Displays a graphical representation of the system topology, separated into NUMA nodes, along with the mapping of processor ids used by the operating system to cpu cores and additional info on caches.

- **Selecting the „right“ binding strategy depends not only on the topology, but also on the characteristics of your application.**
 - Putting threads far apart, i.e. on different sockets
 - May improve the aggregated memory bandwidth available to your application
 - May improve the combined cache size available to your application
 - May decrease performance of synchronization constructs
 - Putting threads close together, i.e. on two adjacent cores which possibly shared some caches
 - May improve performance of synchronization constructs
 - May decrease the available memory bandwidth and cache size
- **If you are unsure, just try a few options and then select the best one.**



■ Define OpenMP Places

- set of OpenMP threads running on one or more processors
- can be defined by the user, i.e. `OMP_PLACES=cores`


■ Define a set of OpenMP Thread Affinity Policies

- SPREAD: spread OpenMP threads evenly among the places
- CLOSE: pack OpenMP threads near master thread
- MASTER: collocate OpenMP thread with master thread

■ Goals

- user has a way to specify where to execute OpenMP threads for
- locality between OpenMP threads / less false sharing / memory bandwidth

■ Assume the following machine:



→ 2 sockets, 4 cores per socket, 4 hyper-threads per core

■ Abstract names for OMP_PLACES:

→ threads: Each place corresponds to a single hardware thread on the target machine.

→ cores: Each place corresponds to a single core (having one or more hardware threads) on the target machine.

→ sockets: Each place corresponds to a single socket (consisting of one or more cores) on the target machine.

■ Example's Objective:

→ separate cores for outer loop and near cores for inner loop

■ Outer Parallel Region: `proc_bind(spread)`, Inner: `proc_bind(close)`

→ spread creates partition, compact binds threads within respective partition


`OMP_PLACES=(0,1,2,3), (4,5,6,7), ... = (0-3):8:4 = cores`

```
#pragma omp parallel proc_bind(spread)
```


```
#pragma omp parallel proc_bind(close)
```

■ Example


→ initial



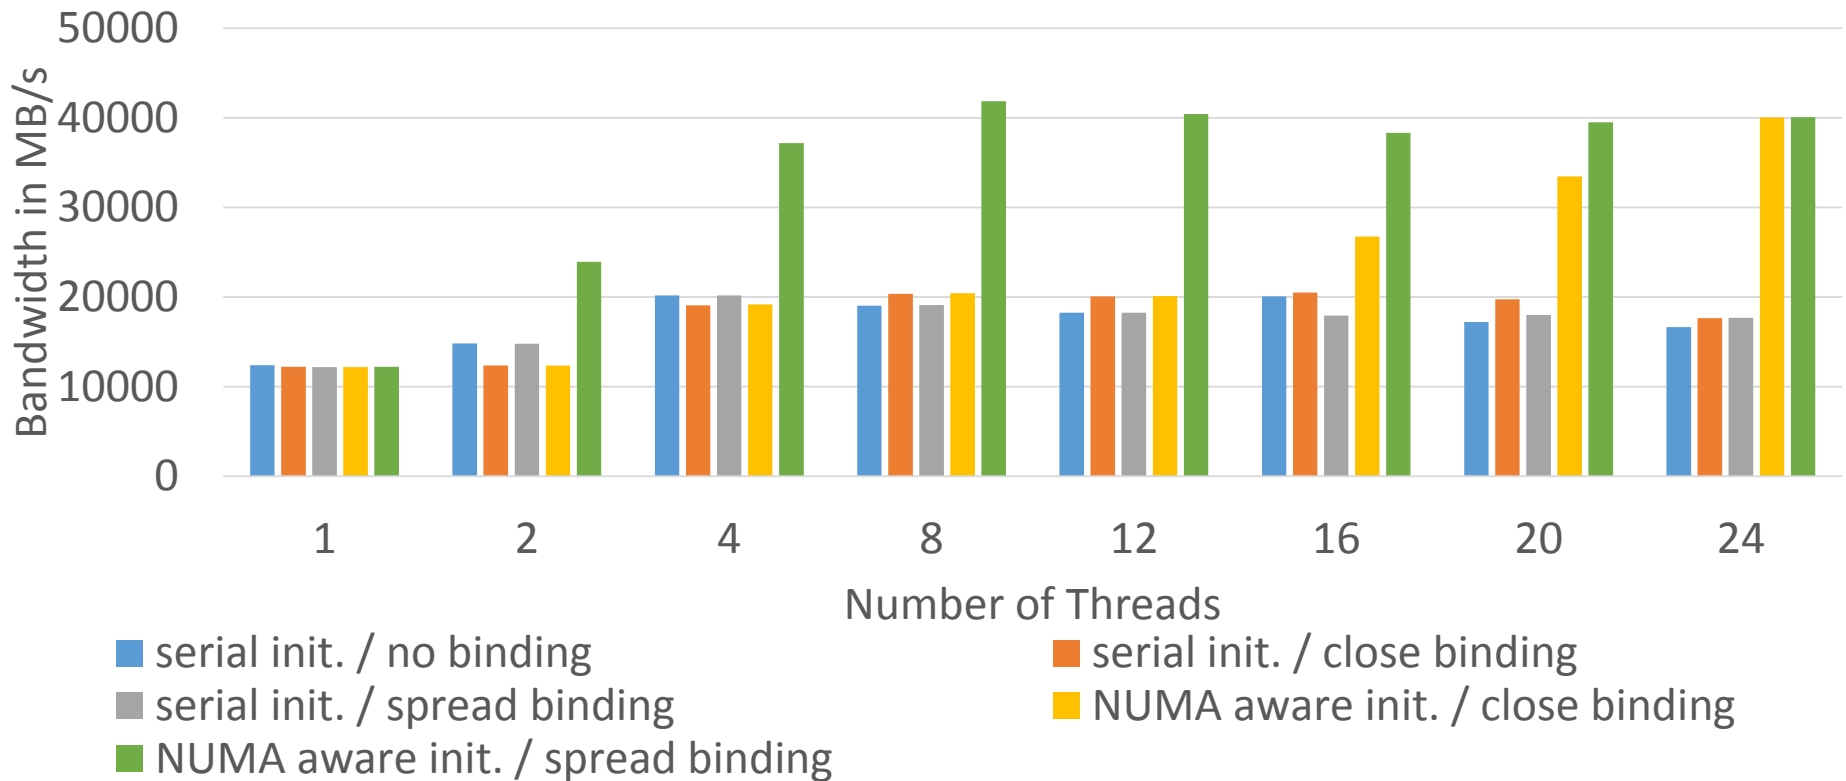
→ spread 4



→ close 4



- **Performance of OpenMP-parallel STREAM vector assignment measured on 2-socket Intel® Xeon® X5675 („Westmere“) using Intel® Composer XE 2013 compiler with different thread binding options:**



Detecting remote accesses

Hardware Counters

Definition: Hardware Performance Counters

In computers, hardware performance counters, or hardware counters are a set of **special-purpose registers** built into modern microprocessors to **store the counts of hardware-related activities** within computer systems. Advanced users often rely on those counters to conduct **low-level performance analysis** or tuning.

(from: <http://en.wikipedia.org>)

Derived Metrics

■ **Clock cycles per Instructions (CPI)**

→ CPI indicates if the application is utilizing the CPU or not

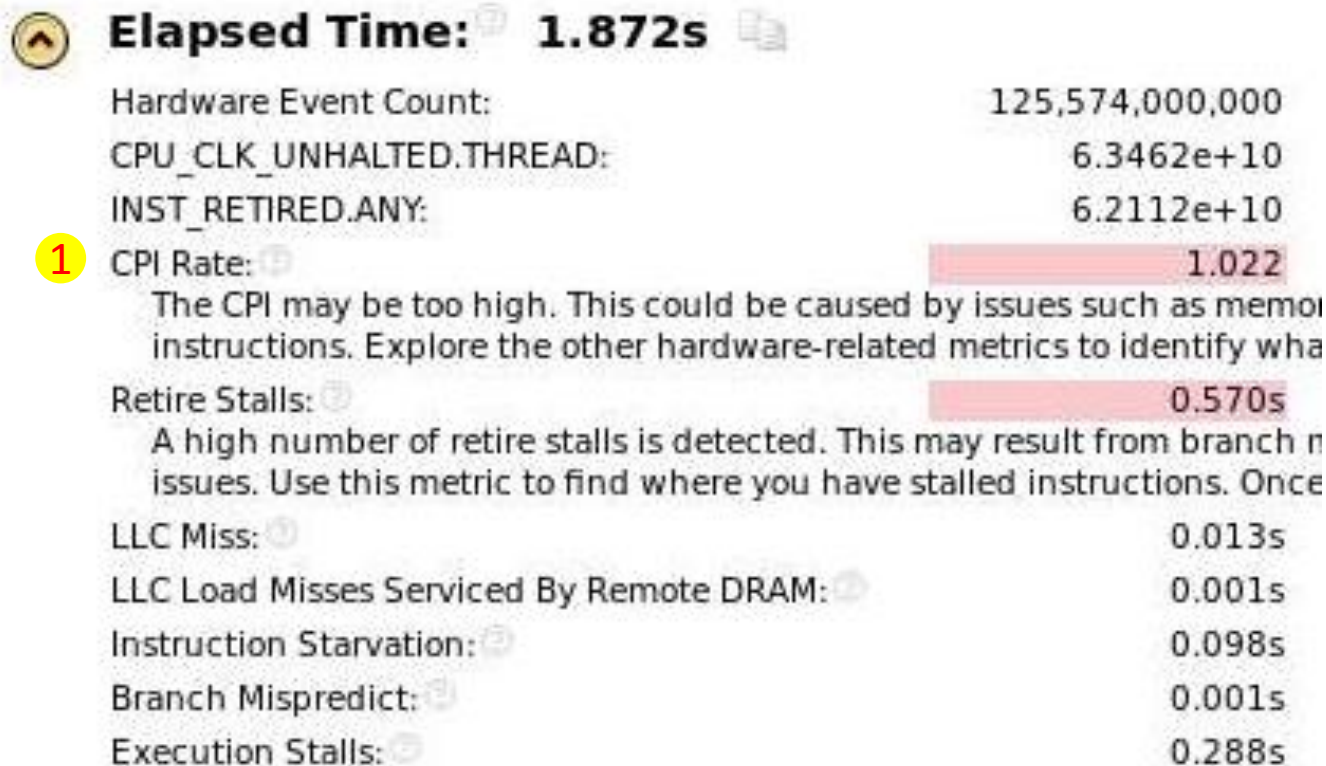
→ Take care: Doing “something” does not always mean doing “something useful”.

■ **Floating Point Operations per second (FLOPS)**

→ How many arithmetic operations are done per second?

→ Floating Point operations are normally really computing and for some algorithms the number of floating point operations needed can be determined.



1 **CPI rate (Clock cycles per instruction):** In theory modern processors can finish 4 instructions in 1 cycle, so a CPI rate of 0.25 is possible. A value between 0.25 and 1 is often considered as good for HPC applications.



- Stream example ($\vec{a} = \vec{b} + s * \vec{c}$) with and without parallel initialization.

→ 2 socket system with Xeon X5675 processors, 12 OpenMP threads

	copy	scale	add	triad
ser_init	18.8 GB/s	18.5 GB/s	18.1 GB/s	18.2 GB/s
par_init	41.3 GB/s	39.3 GB/s	40.3 GB/s	40.4 GB/s



- **Hardware counters can measure local and remote memory accesses.**

- MEM_UNCORE_RETIRED.LOCAL_DRAM_AND_REMOTE_CACHE_HIT
accesses to local memory

- MEM_UNCORE_RETIRED.REMOTE_DRAM
accesses to remote memory

- **Absolute values are hard to interpret, but the ratio between both is useful.**

■ Detecting bad memory accesses for the stream benchmark.

Source		Hardware Event Count: Total by Hardware Event Type	
So. Li. ▲	Source	MEM_UNCORE_RETIRED.LOCAL_DRAM_AND_REMOTE_CACHE_HIT	MEM_UNCORE_RETIRED.REMOTE_DRAM
229	#ifdef TUNED		
230	tuned_STREAM_Scale(scalar);		
231	#else		
232	#pragma omp parallel for		
233	for (j=0; j<N; j++)	20,000	20,000
234	b[j] = scalar*c[j];	3,820,000	3,940,000
235	#endif		


■ Ratio of remote memory accesses:

	copy	scale	add	triad
ser_init	52%	50%	50%	51%
par_init	0.5%	1.7%	0.6%	0.2%

Percentage of remote accesses for ser_init and par_init stream benchmark.

Case Study: CG Solver

Hotspot analysis of the serial code:




Hotspots are:

1. matrix-vector multiplication
2. scaled vector additions
3. dot product

Tuning:





- parallelize all hotspots with a parallel for construct
- use a reduction for the dot-product
- activate thread binding



Hotspot analysis of naive parallel version:


Event Name
MEM_UNCORE_RETIRED.LOCAL_DRAM_AND_REMOTE_CACHE_HIT
MEM_UNCORE_RETIRED.REMOTE_DRAM

A lot of remote accesses occur in nearly all places.

	MEM_UNCORE_RETIRED.LOCAL_...	MEM_UNCORE_RETIRED.REMOTE...
void matvec(const int n, const int int i,j;		
#pragma omp parallel for private(j)	20,000	0
for(i=0; i<n; i++){	0	0
y[i]=0;	0	0
for(j=ptr[i]; j<ptr[i+1]; j	6,740,000 	3,720,000 
y[i]+=value[j]*x[index[17,580,000 	6,680,000 
}		
}		




Tuning:

- Initialize the data in parallel
- Add parallel for constructs to all initialization loops



- Scalability improved a lot by this tuning on the large machine.

■ Analyzing load imbalance in the concurrency view:


So.. Line	Source	CPU Time: Total by... Idle Poor Ok Ic	Ove... and...
49	void matvec(const int n, const int nnz,		
50	int i,j;		
51	#pragma omp parallel for private(j)	22.462s 	10.612s
52	for(i=0; i<n; i++){	0.050s	0s
53	y[i]=0;	0.060s	0s
54	for(j=ptr[i]; j<ptr[i+1]; j++){	1.741s 	0s
55	y[i]+=value[j]*x[index[j]];	9.998s 	0s

■ 10 seconds out of ~35 seconds are overhead time

■ other parallel regions which are called the same amount of time only produce 1 second of overhead

■ Tuning:

→ pre-calculate a schedule for the matrix-vector multiplication, so that the non-zeros are distributed evenly instead of the rows



The Roofline Model


When to stop tuning?

- **Depends on many different factors:**


- How often is the code program used?
- What are the runtime requirements?
- Which performance can I expect?

- **Investigating kernels may help to understand larger applications.**


- Peak performance of a 4 socket Nehalem Server is 256 GFLOPS.




- **Memory bandwidth measured with Stream benchmark is about 75 GB/s.**



- The “Roofline” describes the peak performance the system can reach depending on the “operational intensity” of the algorithm.




Example: Sparse Matrix Vector Multiplication $y=Ax$



Given:

- x and y are in the cache
- A is too large for the cache
- measured performance was **12 GFLOPS**

- 1 ADD and 1 MULT per element
- load of value (double) and index (int) per element
- > 2 Flops / 12 Byte = 1/6 Flops/Byte



Questions?