

Introduction to OpenMP

Dr. Christian Terboven



Introduction to OpenMP

Dr. Christian Terboven

False Sharing




Hardware-specific Concept

General OpenMP Concept



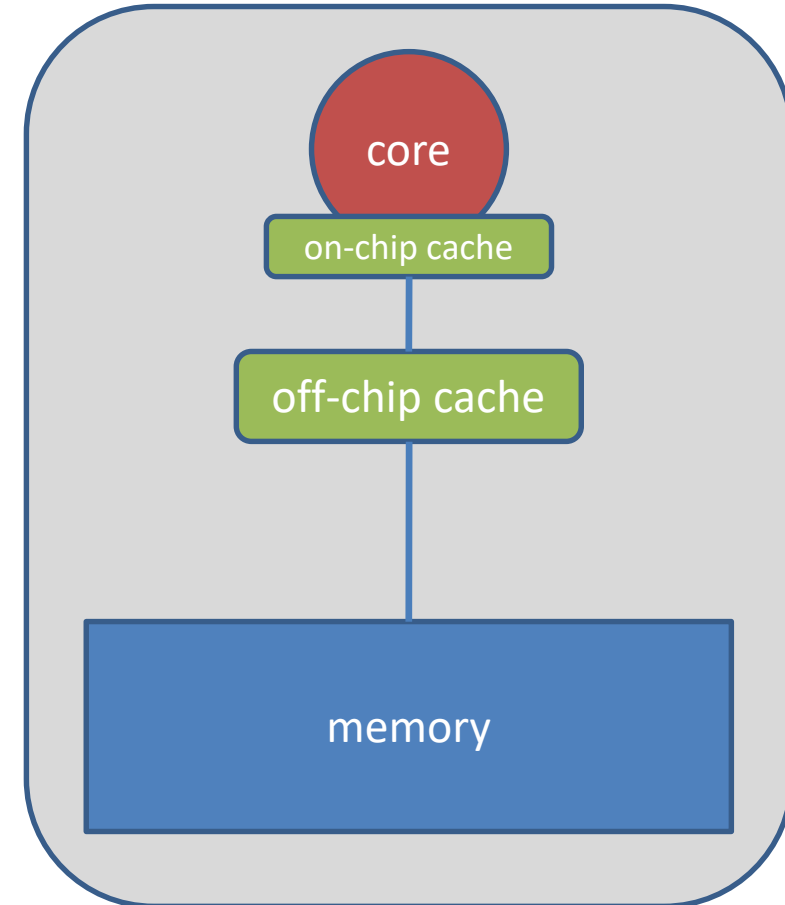
- There is a growing gap between core and memory performance:
 - memory, since 1980: 1.07x per year improvement in latency
 - single core: since 1980: 1.25x per year until 1986, 1.52x p. y. until 2000, 1.20x per year until 2005, then no change on a *per-core* basis



– Source: John L. Hennessy, Stanford University, and David A. Patterson, University of California, September 25, 2012




- 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!




Visualization of the Memory Hierarchy


- 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, although there is no real data dependency
- Note: False Sharing is a performance problem, not a correctness issue



Summing up vector elements again

```
#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
```


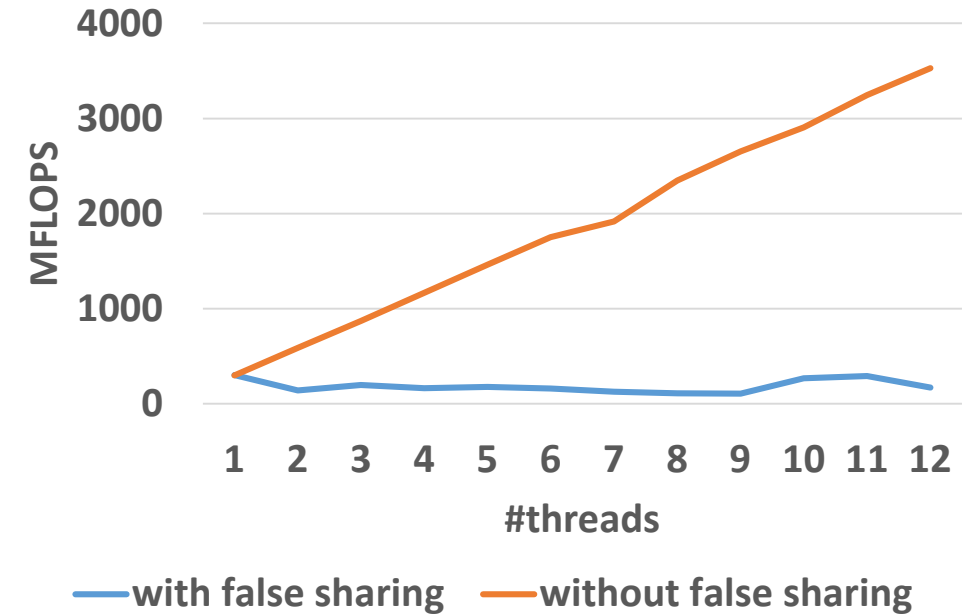


```
1  double s_priv[nthreads];
2  #pragma omp parallel num_threads(nthreads)
3  {
4      int t=omp_get_thread_num();
5      #pragma omp for
6      for (i = 0; i < 99; i++)
7      {
8          s_priv[t] += a[i];
9      }
10 } // end parallel
11 for (i = 0; i < nthreads; i++)
12 {
13     s += s_priv[i];
14 }
```




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



```
1  double s_priv[nthreads * 8];
2  #pragma omp parallel num_threads(nthreads)
3  {
4      int t=omp_get_thread_num();
5      #pragma omp for
6      for (i = 0; i < 99; i++)
7      {
8          s_priv[t * 8] += a[i];
9      }
10 } // end parallel
11 for (i = 0; i < nthreads; i++)
12 {
13     s += s_priv[i * 8];
14 }
```





Introduction to OpenMP

Dr. Christian Terboven


Example: PI



Example: Pi

```
1 double f(double x)
2 {
3     return (4.0 / (1.0 + x*x));
4 }
5
6 double CalcPi (int n)
7 {
8     const double fH = 1.0 / (double) n;
9     double fSum = 0.0;
10    double fX;
11    int i;
12
13    #pragma omp parallel for
14    for (i = 0; i < n; i++)
15    {
16        fX = fH * ((double)i + 0.5);
17        fSum += f(fX);
18    }
19    return fH * fSum;
20 }
```


$$\pi = \int_0^1 \frac{4}{1+x^2}$$



What is wrong with this code?

```
1 double f(double x)
2 {
3     return (4.0 / (1.0 + x*x));
4 }
5
6 double CalcPi (int n)
7 {
8     const double fH = 1.0 / (double) n;
9     double fSum = 0.0;
10    double fX;
11    int i;
12
13    #pragma omp parallel for
14    for (i = 0; i < n; i++)
15    {
16        fX = fH * ((double)i + 0.5);
17        fSum += f(fX);
18    }
19    return fH * fSum;
20 }
```


$$\pi = \int_0^1 \frac{4}{1+x^2}$$



Example: Pi

```
1 double f(double x)
2 {
3     return (4.0 / (1.0 + x*x));
4 }
5
6 double CalcPi (int n)
7 {
8     const double fH = 1.0 / (double) n;
9     double fSum = 0.0;
10    double fX;
11    int i;
12
13    #pragma omp parallel for private(fX,i) reduction(+:fSum)
14    for (i = 0; i < n; i++)
15    {
16        fX = fH * ((double)i + 0.5);
17        fSum += f(fX);
18    }
19    return fH * fSum;
20 }
```


$$\pi = \int_0^1 \frac{4}{1+x^2}$$



```
1 double f(double x)
2 {
3     return (4.0 / (1.0 + x*x));
4 }
5
6 double CalcPi (int n)
7 {
8     const double fH = 1.0 / (double) n;
9     double fSum = 0.0;
10    double fX;
11    int i;
12
13    #pragma omp parallel for private(fX,i) reduction(+:fSum)
14    for (i = 0; i < n; i++)
15    {
16        fX = fH * ((double)i + 0.5);
17        fSum += f(fX);
18    }
19    return fH * fSum;
20 }
```

$$\pi = \int_0^1 \frac{4}{1+x^2}$$

What if we had forgotten this?

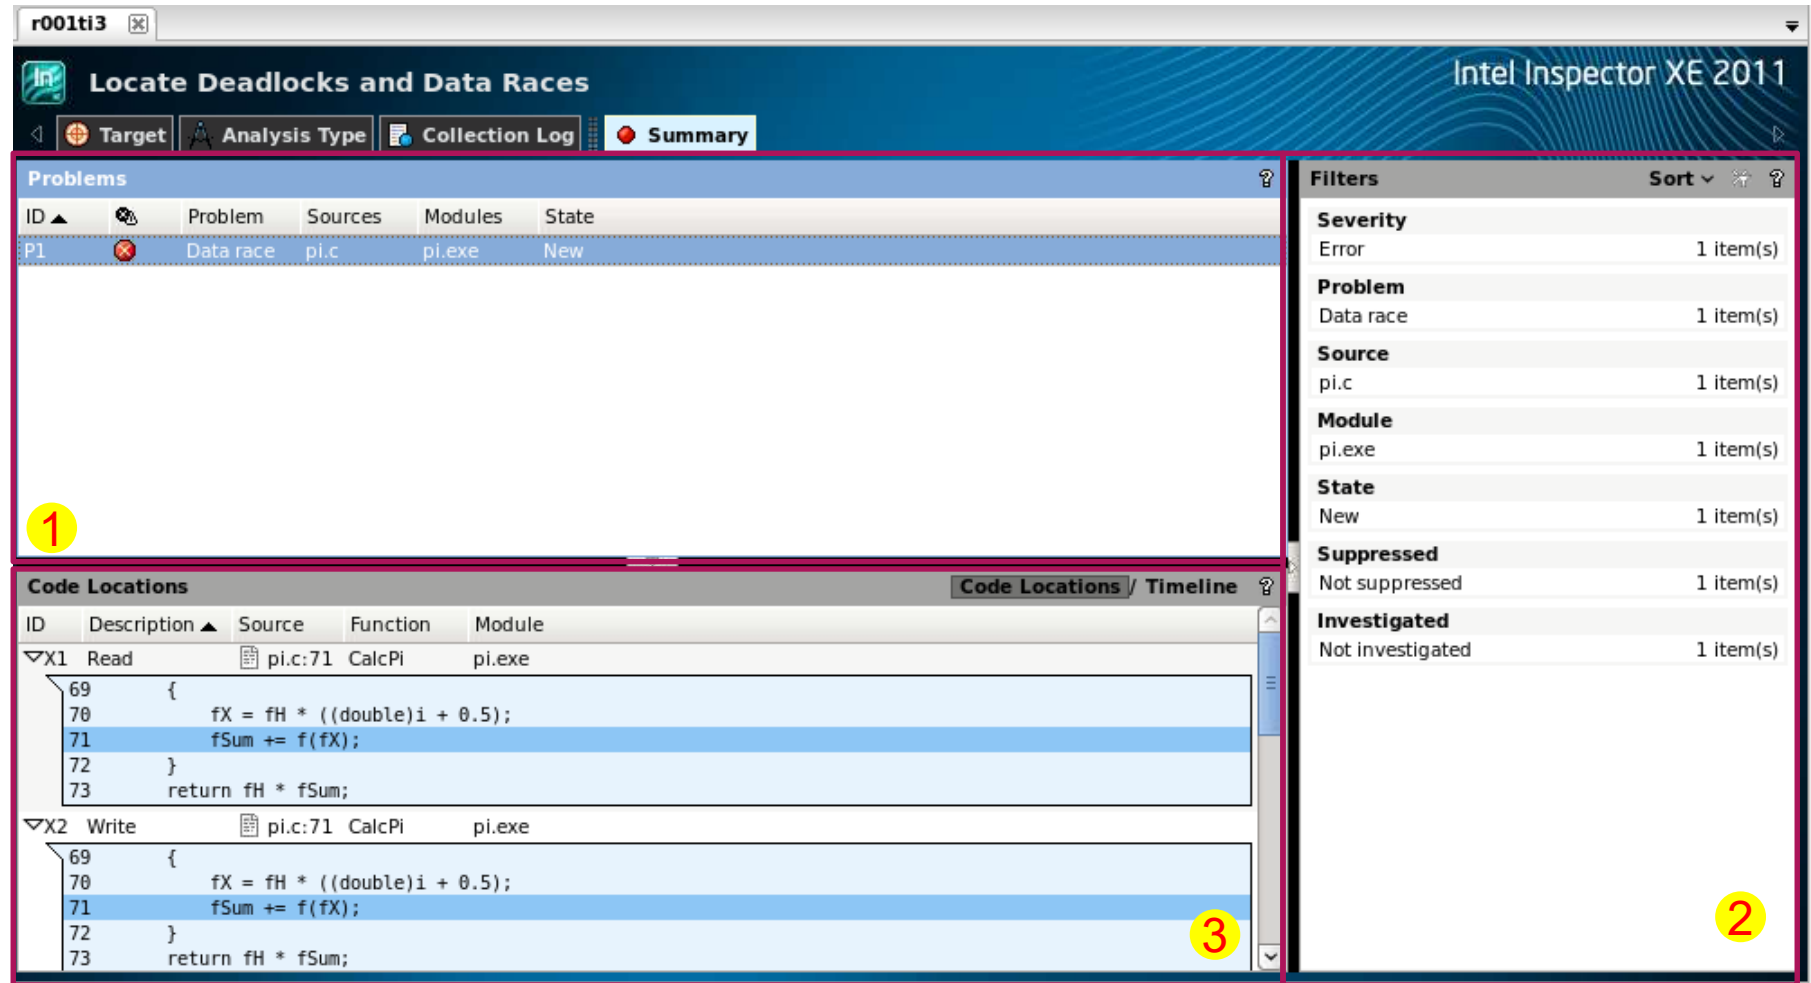


- Data Race: the typical OpenMP programming error, when:
 - two or more threads access the same memory location, and
 - at least one of these accesses is a write, and
 - the accesses are not protected by locks or critical regions, and
 - the accesses are not synchronized, e.g. by a barrier.
- Non-deterministic occurrence: e.g. the sequence of the execution of parallel loop iterations is non-deterministic and may change from run to run
- In many cases *private* clauses, *barriers* or *critical regions* are missing
- Data races are hard to find using a traditional debugger
 - Use tools like *Intel Inspector XE*, *ThreadSanitizer*, *Archer*



Inspector XE – Results

- 1 detected problems
- 2 filters
- 3 code location



The screenshot displays the Intel Inspector XE 2011 interface. The main window is titled "Locate Deadlocks and Data Races" and shows a summary of detected problems. A table in the "Problems" section lists a single problem (P1) of type "Data race" in source file "pi.c" within module "pi.exe", with a state of "New". A yellow circle with the number "1" is placed over this table. To the right, a "Filters" panel shows a list of filters: Severity (Error, 1 item(s)), Problem (Data race, 1 item(s)), Source (pi.c, 1 item(s)), Module (pi.exe, 1 item(s)), State (New, 1 item(s)), Suppressed (Not suppressed, 1 item(s)), and Investigated (Not investigated, 1 item(s)). A yellow circle with the number "2" is placed over the "Investigated" filter. Below the "Problems" section, the "Code Locations" panel shows two locations (X1 and X2) for the data race. Both locations are in source file "pi.c" at line 71, within the "CalcPi" function of module "pi.exe". The code snippets for both locations are identical, showing a loop where a variable "fX" is calculated and then added to "fSum". A yellow circle with the number "3" is placed over the code snippet for location X2.

The missing reduction is detected.



```
1 double f(double x)
2 {
3     return (4.0 / (1.0 + x*x));
4 }
5
6 double CalcPi (int n)
7 {
8     const double fH = 1.0 / (double) n;
9     double fSum = 0.0;
10    double fX;
11    int i;
12
13    #pragma omp parallel for private(fX,i,fSum)
14    for (i = 0; i < n; i++)
15    {
16        fX = fH * ((double)i + 0.5);
17        fSum += f(fX);
18    }
19    return fH * fSum;
20 }
```

What if we just
made the fSum
variable private?

fSum == 0
(no update to
global variable)



- Results for $n = 2 \cdot 10^9$:

# Threads	Runtime [sec.]	Speedup
1	1.141	1.00
2	0.575	1.96
4	0.298	3.93
8	0.161	7.08

System: CLAI2018 Node (Intel Xeon 8160)

Compiler: Intel Compiler 19.0, Flags: `-fopenmp -O3`

- Scalability is good (for up to 8 threads):
 - About 100% of the runtime has been parallelized.
 - As there is just one parallel region, there is virtually no overhead introduced by the parallelization.
 - Problem is parallelizable in a trivial fashion ...



Questions?

