# Introduction to OpenMP

#### **Rolf Rabenseifner**

rabenseifner@hlrs.de www.hlrs.de/people/rabenseifner/

#### University of Stuttgart High Performance Computing Center Stuttgart (HLRS) www.hlrs.de

Version 13, Aug 21, 2020 (for OpenMP 3.1 and older)





# **Outline**

•	Introduction into OpenMP	-3 (slide numbers)
•	Programming and Execution Model	- 20
	<ul> <li>Parallel regions: team of threads</li> </ul>	 _ 21
	– Syntax	- <u>25</u>
	<ul> <li>Data environment (part 1)</li> </ul>	- <u>28</u>
	<ul> <li>Environment variables</li> </ul>	- <u>29</u>
	<ul> <li>Runtime library routines</li> </ul>	- <u>30</u>
	<ul> <li>Exercise 1: Parallel region / library calls / privat &amp; shared variables</li> </ul>	- <u>33</u>
•	Worksharing directives	- <u>42</u>
	– Which thread executes which statement or operation?	- <u>43</u>
	– Exercise 2a: Pi	– <u>57</u>
	– Tasks	- <u>62</u>
	<ul> <li>Synchronization constructs, e.g., critical regions</li> </ul>	- <u>70</u>
	<ul> <li>Nesting and Binding</li> </ul>	- <u>77</u>
	– Exercise 2b: Pi	– <u>81</u>
٠	Data environment and combined constructs	- <u>86</u>
	<ul> <li>Private and shared variables, Reduction clause</li> </ul>	- <u>87</u>
	<ul> <li>Combined parallel worksharing directives</li> </ul>	- <u>92</u>
	<ul> <li>Exercise 3: Pi with reduction clause and combined constructs</li> </ul>	– <u>95</u>
	- Exercise 4: Heat	– <u>103</u>
٠	Summary of OpenMP API	- <u>123</u>
•	OpenMP Pitfalls & Optimization Problems	– <u>127</u> & <u>141</u>
•	Appendix (exercise solutions)	- <u>152</u>

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online Introduction to OpenMP  $\rightarrow$  Outline

Slide 2 / 151

### **OpenMP Overview: What is OpenMP?**

- OpenMP is a standard programming model for shared memory parallel programming
- Portable across all shared-memory architectures
- It allows incremental parallelization
- Compiler based extensions to existing programming languages
  - mainly by directives
  - a few library routines
- Fortran and C/C++ binding
- OpenMP is a standard



## **Parallel hardware architectures**

#### shared memory



#### Socket/CPU

#### → memory interface

#### UMA (uniform memory access) SMP (symmetric multi-processing)

All cores connected to all memory banks with same speed

#### **Parallel execution streams** on each core, e.g.,

x[ 0 ... 999] = ... on 1<sup>st</sup> core x[1000 ... 1999] = ... on 2<sup>nd</sup> core x[2000 ... 2999] = ... on 3<sup>rd</sup> core



#### Node → hyper-transport

#### ccNUMA (cache-coherent non-uniform memory access)

#CPUs x memory bandwidth Shared memory programming is possible **Performance problems:** 

- Each parallel execution stream should mainly access the memory of its CPU
   → First-touch strategy is needed to minimize remote memory access
- Threads should be pinned to the physical sockets



#### distributed memory



#### Cluster

#### $\rightarrow$ node-interconnect

#### NUMA (non-uniform memory access) !! fast access only on its own memory !! Many programming options:

- Shared memory / symmetric multiprocessing inside of each node
- distributed memory parallelization on the node interconnect
- Or simply one MPI process on each core

#### Shared memory programming with OpenMP

© 2000-2023 HLRS, Rolf Rabenseifner Introduction to OpenMP  $\rightarrow$  Intro

• REC  $\rightarrow$  <u>online</u>

#### MPI works everywhere

Slide, courtesy to Claudia Blaas-Schenner

Slide 4 / 151



**!\$OMP END PARALLEL DO** 

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \underline{\text{online}}]$ Introduction to OpenMP  $\rightarrow$  Intro



Slide 5 / 151





Slide 6 / 151



© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u> Introduction to OpenMP  $\rightarrow$  Intro



Slide 7 / 151 Call MPI Comm size(MPI COMM WORLD, size, ierror) Call MPI\_Comm\_rank(MPI\_COMM\_WORLD, myrank, ierror) m1 = (m+size-1)/size; ja=1+m1\*myrank; je=min(m1\*(myrank+1), m) jax=ja-1; jex=je+1 // extended boundary with halo

Call MPI\_Send(.....) ! - sending the boundary data to the neighbors Call MPI Recv(.....) ! - receiving from the neighbors, ! storing into the halo cells



end do

© 2000-2023 HLRS, Rolf Rabenseifner ● REC → online Introduction to OpenMP  $\rightarrow$  Intro



## Motivation: Why should I use OpenMP?



Time/Effort





## Further Motivation to use OpenMP

- OpenMP is the easiest approach to multi-threaded programming ٠
- Multi-threading is needed to exploit modern hardware platforms: ٠
  - Several CPUs together within one ccNUMA node
  - Several cores per CPU
  - Hyperthreading



#### Where should I use OpenMP?



Problem size



# Simple OpenMP Program

- Most OpenMP constructs are compiler directives or pragmas
- The focus of OpenMP is to parallelize loops with independent iterations
- · OpenMP offers an incremental approach to parallelism

```
Serial Program:Parallel Program:void main()void main(){double Res[1000];for(int i=0;i<1000;i++) {</td>double Res[1000];do_huge_comp(Res[i]);do_huge_comp(Res[i]);}}
```



#### Speedup, Efficiency, Scaleup, and Weak Scaling

•	Definition:	T(p,N) = time to solve problem of total size N on p processors	
•	Parallel speedup:	S(p,N) = T(1,N) / T(p,N)	Three different ways of reporting the success
		compute same problem with more processors	in <b>shorter time</b>
•	Parallel Efficiency:	E(p,N) = S(p,N) / p	
•	Scaleup:	Sc(p,N) = N / n with T(1,n) = T(p,N) compute <b>larger problem</b> with more processors	in <b>same time</b>
•	Weak scaling:	T(p, p•n) / T(1,n) is reported, i.e., problem size per process (n) is fixed	

• Problems:

Speedup & Amdahl's Law

- Absolute MFLOPS rate / hardware peak performance?
- Super-scalar speedup: S(p,N)>p, e.g., due to cache<sup>\*</sup>) usage for large p:
  - T(1,N) may be based on a huge number of N data elements in the memory in the one process, whereas
  - T(p,N) may be based on *cache based execution* due to only N/p data elements per process
- S(p,N) close to **p** or far less?  $\rightarrow$  see Amdahl's Law on next slide

\*) Faster memory for temporary copies of recently used data

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \overline{\text{REC}} \rightarrow \underline{\text{online}}]$ 

#### Amdahl's Law

 $T(p,N) = f \cdot T(1,N) + (1-f) \cdot T(1,N) / p$ f ... sequential part of code that can not be done in parallel S(p,N) = T(1,N) / T(p,N) = 1 / (f + (1-f) / p)For p  $\rightarrow$  infinity, speedup is limited by S(p,N) < 1 / f



-- S(p,N) = p (ideal speedup) → f=0.1% => S(p,N) < 1000 → f= 1% => S(p,N) < 100 -- f= 5% => S(p,N) < 20 → f= 10% => S(p,N) < 10



#### Amdahl's Law (double-logarithmic)

 $T(p,N) = f \cdot T(1,N) + (1-f) \cdot T(1,N) / p$ f ... sequential part of code that can not be done in parallel S(p,N) = T(1,N) / T(p,N) = 1 / (f + (1-f) / p)For p —> infinity, speedup is limited by S(p,N) < 1 / f





#### Who owns OpenMP? - OpenMP Architecture Review Board

**AMD**<sup>1)</sup> (Greg Rodgers) Argonne National Laboratory (Kalyan Kumaran) **ARM** (Graham Hunter) ASC/Lawrence Livermore NL<sup>1)</sup> (Bronis R. de Supinski) Barcelona Supercomputing Center (Xavier Martorell) Brookhaven National Laboratory (Vivek Kale) cOMPunity<sup>1)</sup> (Yonghong Yan) **CRAY**<sup>1)</sup>, a HPE company (Deepak Eachempati) Edinburgh Parallel Computing Centre<sup>1)</sup> (Mark Bull) Fujitsu<sup>1)</sup> (Naoki Sueyasu) **IBM**<sup>1)</sup> (Kelvin Li) **INRIA** (Olivier Aumage) **Intel**<sup>1)</sup> (Xinmin Tian) Lawrence Berkeley National Laboratory (Helen He) Leibniz Supercomputing Centre (Volker Weinberg) Los Alamos National Laboratory (Jamal Mohd-Yusof) Maui HPC Center (Alice Koniges) **Mentor Graphics**, a Siemens Business (Catherine Moore) **Micron** (Randy Meyer)

NASA<sup>1)</sup> (Henry Jin) **NEC<sup>1)</sup>** (Shin-ichi Okano) **NVIDIA** (Jeff Larkin) Oak Ridge National Laboratory (Oscar Hernandez) RWTH Aachen University<sup>1)</sup> (Dieter an Mey) Sandia National Laboratory (Stephen Olivier) Stony Brook University (Dr. Barbara Chapman) **SUSE** (Michael Matz) Texas Advanced Computing Center (Kent Milfeld) The University of Manchester (Antoniu Pop) University of Basel (Florina Ciorba) University of Bristol (Simon McIntosh-Smith) University of Delaware (Sunita Chandrasekaran) University of Tennessee (Piotr Luszczek)

From: <u>https://www.openmp.org/about/members/</u> Last update: Aug 21, 2020

2020: 11 (=**30%**) companies + 22 (**70%**) HPC centers 2008: 11 (=**70%**) companies + 5 (**30%**) HPC centers



<sup>&</sup>lt;sup>1)</sup> Already 2008 member of the ARB

#### **OpenMP Release History**



© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u> Introduction to OpenMP  $\rightarrow$  Intro



# **OpenMP** Availability

- OpenMP 1.0 (C/C++) and OpenMP 1.1 (Fortran 90) ٠ is available on all platforms in the commercial compilers
- Most features from OpenMP 2.0 are already implemented ٠
- OpenMP 2.5 no substantial new features/changes compared to 2.0
- OpenMP 3.0 task concept added
  - new features for loop worksharing
- **OpenMP 3.1** final and mergeable clauses for task construct
  - taskyield construct to allow user-defined task switching points
  - min and max reduction in C/C++
  - OMP NUM THREADS can handle a list for nested regions
  - OMP PROC BIND to bind threads to processors
- OpenMP 4.0 Thread affinity and OMP PLACES ٠
  - SIMD support
  - Support for accelerators
  - Tasking extensions
- **OpenMP 4.5** taskloop ٠
  - OpenMP 5.0 - reduction clause for taskloop and taskloop simd
    - default loop schedule **changed** from **monotonic** to **nonmonotonic**  $\rightarrow$  enables work stealing for dynamic & guided

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  online Introduction to OpenMP  $\rightarrow$  Intro

Slide 18 / 151

# **OpenMP Information**

- OpenMP Homepage: http://www.openmp.org/
- OpenMP user group: http://www.compunity.org/
- Barbara Chapman, Gabriele Jost, and Ruud van der Pas: Using OpenMP. MIT Press, 2008, ISBN-13: 978-0-262-53302-7.
- Georg Hager, Gerhard Wellein: Introduction to High Performance Computing for Scientists and Engineers. CRC Press, 2010, 256p, ISBN13: 978-1-439-81192-4
- R.Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald, and R. Menon: Parallel programming in OpenMP. Academic Press, San Diego, USA, 2000, ISBN 1-55860-671-8



# Outline — Programming and Execution Model

Introduction into OpenMP

#### Programming and Execution Model

- Parallel regions: team of threads
- Syntax
- Data environment (part 1)
- Environment variables
- Runtime library routines
- **Exercise 1:** Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

# **OpenMP Programming Model**

- OpenMP is a shared memory model.
- Workload is distributed between threads
  - Variables can be
    - shared among all threads
    - duplicated for each thread
  - Threads communicate through barriers ( $\rightarrow$  next slide).
- Unintended sharing of data can lead to race conditions:
  - race condition: when the program's outcome changes as the threads are scheduled differently.
- To control race conditions:
  - Use synchronization to protect data conflicts.

#### threads = parallel execution streams



Introduction to OpenMP → Programming and Execution Model

## **OpenMP Execution Model**



#### Fork-join model of parallel execution

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u> Introduction to OpenMP  $\rightarrow$  Programming and Execution Model

#### **OpenMP Parallel Region Construct**



# **OpenMP Parallel Region Construct Syntax**

- Block of code to be executed by multiple threads in parallel. Each thread executes the **same code in a replicated way**!
- Fortran:
   !\$OMP PARALLEL [ clause [[,] clause ] ... ]
   block
  - !\$OMP END PARALLEL
    - parallel/end parallel directive pair must appear in the same routine
  - C/C++:

ortran

C/C++

#pragma omp parallel[clause[[,]clause]...]new-line
 structured-block

- *clause* can be one of the following:
  - private(*list*)
  - shared(*list*)

[xxx] = xxx is optional

© 2000-2023 HLRS, Rolf Rabenseifner  $[ \bullet \text{REC} \rightarrow \text{online} ]$ Introduction to OpenMP  $\rightarrow$  Programming and Execution Model



- #pragma directives <u>case sensitive</u>
- Format: #pragma omp directive\_name [ clause [ [ , ] clause ] ... ] new-line
- Conditional compilation
  #ifdef \_OPENMP
   block,
   e.g., printf("%d avail.processors\n",omp\_get\_num\_procs());
  #endif
- Include file for library routines: #ifdef \_OPENMP #include <omp.h> #endif
- In the old OpenMP 1.0 syntax, the <u>comma [,]</u> between clauses was <u>not</u> allowed (some compilers in use still may have this restriction)



- Treated as Fortran comments <u>not</u> case sensitive
- Format: sentinel directive\_name [ clause [ [ , ] clause ] ... ]
- Directive sentinels:
  - Fixed source form: **!\$OMP** | C\$OMP | \*\$OMP [starting at column 1]
  - Free source form: **!\$OMP** [may be preceded by white space]
- Conditional compilation
  - Fixed source form: **!\$** | C\$ | \*\$
  - Free source form: **!\$**
  - #ifdef \_OPENMP [in my\_fixed\_form.F or .F90]
     block
     #endif
  - Example:
    - !\$ write(\*,\*) OMP\_GET\_NUM\_PROCS(),' avail. processors'
- Include file for library routines:
  - include 'omp\_lib.h' Or use omp\_lib [implementation dependent]

#### **OpenMP Data Scope Clauses**

- private (*list*)
   Declares the variables in *list* to be private to each thread in a team.
   They are uninitialized.
- shared (*list*)
   Makes variables that appear in *list* shared among all the threads in a team
- If not specified: default shared, but
  - stack (local) variables in called sub-programs are PRIVATE
  - Automatic variables within a block are PRIVATE
  - Loop control variable of parallel OMP
    - DO (Fortran)
    - for (C)
    - is PRIVATE

- ...



[see later: Data Model]

## **OpenMP Environment Variables**

- OMP\_NUM\_THREADS
  - sets the number of threads to use during execution
  - when dynamic adjustment of the number of threads is enabled, the value of this environment variable is the maximum number of threads to use
  - setenv OMP\_NUM\_THREADS 16 [csh, tcsh]
  - export OMP NUM THREADS=16 [sh, ksh, bash]
- OMP\_SCHEDULE
  - applies only to do/for and parallel do/for directives that have the schedule type RUNTIME
  - sets schedule type and chunk size for all such loops
  - setenv OMP SCHEDULE "GUIDED, 4" [csh, tcsh]
  - export OMP SCHEDULE="GUIDED, 4" [sh, ksh, bash]

# **OpenMP Runtime Library (1)**

- Query functions
- Runtime functions
  - Run mode
  - Nested parallelism
- Lock functions

C/C++

Fortran

- C/C++: add #include <omp.h>
- Fortran: add all necessary OMP routine declarations, e.g.,

!\$ INTEGER omp\_get\_thread\_num

or use include file

!\$ INCLUDE 'omp\_lib.h'

or module

!\$ USE omp\_lib

Existence of include file or module or both is implementation dependent.



# **OpenMP Runtime Library (2)**

- omp\_get\_num\_threads Function
   Returns the number of threads currently in the team executing the parallel region from which it is called
- Fortran
  - C/C++
    - integer function omp\_get\_num\_threads()
       C/C++:
      int omp\_get\_num\_threads(void);
    - omp\_get\_thread\_num Function Returns the thread number, within the team, that lies between 0 and omp\_get\_num\_threads()-1, inclusive. The master thread of the team is thread 0



C/C++

– Fortran:

- Fortran:

integer function omp\_get\_thread\_num()

- C/C++:

int omp\_get\_thread\_num(void);

## **OpenMP Runtime Library (3): Wall clock timers**

Portable wall clock timers similar to MPI\_WTIME

#### DOUBLE PRECISION FUNCTION OMP\_GET\_WTIME() double omp\_get\_wtime(void);

- provides elapsed time

```
START=OMP_GET_WTIME()
! Work to be measured
END = OMP_GET_WTIME()
PRINT *, `Work took `, END-START, ` seconds`
```

- provides "per-thread time", i.e. needs not be globally consistent



Fortran

C/C++

- DOUBLE PRECISION FUNCTION OMP\_GET\_WTICK() double omp\_get\_wtick(void);
  - returns the number of seconds between two successive clock ticks

# Outline — Exercise 1: Parallel region

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax

Library-calls

<del>. .</del>

Exercise

- Data environment (part 1)
- Environment variables
- Runtime library routines
- Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>



#### **OpenMP Exercise 1: Parallel region (1)**

- Goal: usage of
  - runtime library calls
  - conditional compilation
  - environment variables
  - parallel regions, private and shared clauses



Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1

## **OpenMP Exercise 1: Parallel region (2)**

- Compile serial program hello. [c|f90] and run ٠
- Compile as **OpenMP** program and run on 4 cores ٠
  - add OpenMP compile option, e.g.,
    - -fopenmp on gnu compilers gcc, and gfortran
    - -gopenmp on Intel compiler icc and ifort
  - export OMP NUM THREADS=4 ; ./a.out (with bash or similar)
  - setenv OMP NUM THREADS 4 ; ./a.out
- - (with tcsh or similar)
  - expected result: program is not parallelized, same output

```
bash$ gcc -fopenmp hello.c
     gfortran -fopenmp hello.f90
or
bash$ export OMP NUM THREADS=4; ./a.out
hello world -1
bash$
```

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \underline{\text{online}}]$ 

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1



# **OpenMP Exercise 1: Parallel region (3)**

- Please add the following additional code (i.e., don't change existing code!)
  - If compiled with OpenMP then
  - i := the rank of the current thread
- Compile as OpenMP and as non-OpenMP program and run on 4 cores
  - Expected result:
    - Without OpenMP: no difference

```
bash$ gcc hello.c
or gfortran hello.f90
bash$ export OMP_NUM_THREADS=4; ./a.out
hello world -1
bash$
```

• With OpenMP: Only the master thread is running, and therefore i = 0

```
bash$ gcc -fopenmp hello.c
or gfortran -fopenmp hello.f90
bash$ export OMP_NUM_THREADS=4; ./a.out
hello world 0
bash$
```

- After successful execution, you may compare with the provided solution:
  - ../../solution/hello/hello2.c or hello2.f90 (hello1 was your starting point)

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1
# **OpenMP Exercise 1: Parallel region (4)**

- Please add the following additional code (i.e., don't change existing code!)
  - Execute the i=omp\_get\_...() together with the print statement within a parallel region
- Compile as OpenMP and as non-OpenMP program and run on 4 cores
  - Expected result:
    - Without OpenMP: no difference to previous slide

```
hello world -1
```

 With OpenMP: Each thread executes its print statement after each thread stored a different value into i = 0, 1, 2, 3, but the variable i is shared! Therefore the output can be:

```
bash$ gcc -fopenmp hello.c
or gfortran -fopenmp hello.f90
bash$ export OMP_NUM_THREADS=4; ./a.out
hello world 3
hello world 1
hello world 1
hello world 3
```

- After successful execution, you may compare with the provided solution:
  - ../../solution/hello/hello3.c or hello3.f90

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1

# **OpenMP Exercise 1: Parallel region (5)**

- If you could not see the race condition, then you can apply a small trick:
  - Make the race-condition obvious:
  - Add a sleep(1); in C, or a CALL sleep(1) in Fortran
  - between the i= ... and the print statement
- Compile as OpenMP and as non-OpenMP program and run on 4 cores
  - Expected result:
    - Without OpenMP: no difference to previous slide

```
hello world -1
```

• With OpenMP: Each thread stores a different value into i = 0, 1, 2, 3,

but the variable *i* is still shared. And printing starts one second later <sup>(2)</sup> Therefore the output can be:



- After successful execution, you may compare with the provided solution:

../../solution/hello/hello4.c or hello4.f90

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  online

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1

# **OpenMP Exercise 1: Parallel region (6)**

- As last step, please add the private() clause to your parallel region directive:
  - Make the variable i private
- Compile as OpenMP and as non-OpenMP program and run on 4 cores
  - Expected result:
    - Without OpenMP: no difference to previous slide

```
hello world -1
```

 With OpenMP: Each thread stores a different value 0, 1, 2, 3 into its private i. And then it prints its private value. Therefore the output can be:

```
bash$ gcc -fopenmp hello.c
or gfortran -fopenmp hello.f90
bash$ export OMP_NUM_THREADS=4; ./a.out
hello world 3
hello world 1
hello world 0
hello world 2
```

- After successful execution, you may compare with the provided solution:
  - ../../solution/hello/hello5.c or hello5.f90

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1



#### **OpenMP Exercise 1: Summary**

- **Conditional compilation** allows to keep the serial version of the program in the same source files
- Testing an OpenMP library routine (always within conditional compilation!)
- compilers need to be used with special option for OpenMP directives to take any effect
- **Parallel regions** are executed by each thread in the same way unless worksharing directives are used (here, in this exercise, we used no worksharing)
- Decision about private or shared status of variables is important

© 2000-2023 HLRS, Rolf Rabenseifner  $\left[ \bullet \text{ REC} \rightarrow \text{online} \right]$ 

Introduction to OpenMP → Programming and Execution Model → Exercise 1

During the Exercise (25-30 min.) [

Please stay here in the main room while you do this exercise

And have fun with this short exercise 6

Please do not look at the solution before you finished this exercise, otherwise, 90% of your learning outcome may be lost



#### As soon as you finished the exercise, please go to your breakout room

and continue your discussions with your fellow learners:

Exercise + Discussion + Break – planned: 25-30 Minutes

Discussion: MPI, we compared it with "each process ~ a human being,

each message or collective communication ~ human communication",

Which analogy would fit to OpenMP – or is this a stupid Question?

• REC  $\rightarrow$  online © 2000-2023 HLRS, Rolf Rabenseifner

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1

# **Outline — Worksharing directives**

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables

#### Worksharing directives

- Which thread executes which statement or operation?
- Exercise 2a: Pi
- Tasks
- Synchronization constructs, e.g., critical regions
- Nesting and Binding
- Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ Introduction to OpenMP  $\rightarrow$  Worksharing directives

# **Worksharing and Synchronization**

- Which thread executes which statement or operation?
- and when?
  - Worksharing constructs
  - Master and synchronization constructs
- i.e., organization of the parallel <u>work</u>!!!

# **OpenMP Worksharing Constructs**

- Divide the execution of the enclosed code region among the members of the team
- Must be enclosed dynamically within a parallel region
- They do not launch new threads
- No implied barrier on entry
- sections directive
- for directive (C/C++)
- do directive (Fortran)
- workshare directive (Fortran)
- single directive

Task generating construct – another option for organizing the work in parallel

• task directive

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Worksharing directives

C/C++

Fortran



# **OpenMP** sections **Directives** - C/C++







# **OpenMP** sections **Directives** – Fortran

Fortran:

**!**\$OMP PARALLEL **!**\$OMP SECTIONS a=... b=... **!\$OMP SECTION** c=... d=... **!\$OMP SECTION** e=... f=.... **!\$OMP SECTION** g=... h=... **!**\$OMP END SECTIONS **!\$OMP END PARALLEL** 



#### **OpenMP** sections **Directives** – **Syntax**







f=7

10,14

f=7

i=

15,19

a(i)=





#### **OpenMP** do/for **Directives** – **Syntax**

• Immediately following loop executed in parallel



The corresponding for loop must have canonical shape
 → next slide

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online

Introduction to OpenMP  $\rightarrow$  Worksharing directives

#### **OpenMP** for **Directives** – **Canonical Shape**

- C/C++: #pragma omp for[clause[[,]clause]...]new-line for-loop
  - The corresponding for loop must have canonical shape:

var: must not be modified in the loop body; integer (signed or unsigned), or pointer type (C only), (OpenMP ≥ 3.0) or random access iterator type (C++ only)

C/C++

- *clause* can be one of the following: ٠
  - private(list)
  - reduction (*operator:list*)
  - collapse(n)

[see later: Data Model] [see later: Data Model]

 $(OpenMP \ge 3.0)$ 

- schedule(type[, chunk])
- nowait (C/C++: on #pragma omp for) (Fortran: on \$! OMP END DO)
- Implicit barrier at the end of do/for unless nowait is specified, i.e., if nowait is specified, threads do not synchronize at the end of the parallel loop
- collapse(*n*) with constant integer expression *n*: ٠ The iterations of the following *n* nested loops are collapsed into one larger iteration space.
- schedule clause specifies how the iterations (iteration space) ٠ of the (nested) loop(s) are divided among the threads of the team.

#### **OpenMP** schedule **Clause**

Within schedule(type[, chunk]) type can be one of the following:

- static: Iterations are divided into pieces of a size specified by *chunk*. The pieces are statically assigned to threads in the team in a round-robin fashion in the order of the thread number.
   Default chunk size: one contiguous piece for each thread.
- dynamic: Iterations are broken into pieces of a size specified by *chunk*.
   As each thread finishes a piece of the iteration space, it dynamically obtains the next set of iterations. Default chunk size: 1.
- guided: The chunk size is reduced in an exponentially decreasing manner with each dispatched piece of the iteration space. *chunk* specifies the smallest piece (except possibly the last).
   Default chunk size: 1. Initial chunk size is implementation dependent.
- auto: Scheduling is delegated to the compiler and/or runtime system (OpenMP ≥ 3.0)
- runtime: The decision regarding scheduling is deferred until run time. The schedule type and chunk size can be chosen at run time by setting the OMP\_SCHEDULE environment variable.

Default schedule: implementation dependent.

# Loop scheduling



Further information: OpenMP 3.0, Sect. 2.5.1, Description, pp 41-44

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Worksharing directives



- WORKSHARE directive allows parallelization of array expressions and FORALL statements
- Usage:

   !\$OMP WORKSHARE
   A=B+C
   ! Rest of block
   !\$OMP END WORKSHARE
  - Comenties:
- Semantics:
  - Work inside block is divided into separate units of work.
  - Each unit of work is executed only once.
  - The units of work are assigned to threads in any manner.
  - The compiler must ensure sequential semantics.
  - Similar to DO worksharing without explicit loops.

## **OpenMP** single **Directive** – **Syntax**

- The block is executed by only one thread in the team ٠ (not necessarily the master thread)
- Fortran: Fortran !\$OMP single [ clause [[,] clause ] ... ] block !\$OMP end single [nowait]
  - C/C++: #pragma omp single [ clause [ [ , ] clause ] ... ] new-line structured-block
    - Implicit barrier at the end of **single** construct ٠ (unless a nowait clause is specified)
    - To reduce the fork-join overhead, one can combine ٠
      - several parallel parts (for, do, workshare, sections)
      - and sequential parts (single)
      - in one parallel region (parallel ... end parallel)

C/C++

# Outline — Exercise 2a: Pi

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

#### **OpenMP Exercise 2a: pi Program (1)**

- Goal: usage of
  - worksharing constructs: do/for
  - critical directive → in Exercise 2b
- Working directory: ~/OpenMP/#NR/pi/

#NR = number of your PC, e.g., 01

- Serial programs:
  - Fortran 90: pi.f90
  - C: pi.c

- Calculating  $4 \int_0^1 \frac{1}{1+x^2} dx = 4 \arctan(x) \Big|_0^1 = 4 \Big( \frac{\pi}{4} - 0 \Big) = \pi$ 

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online

Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercise 2a





# **OpenMP Exercise 2a: pi Program (2)**

- compile serial program pi. [c|f90] and run
- Compile and run with OpenMP
  - Expected result: same pi-value and timing,
  - but the program reports the number of threads, and timing with omp\_get\_wtime
- add parallel region and do/for directive in pi. [c|f90] and compile
- set environment variable OMP\_NUM\_THREADS to 2 and run

– value of pi? (should be wrong!)

- run again
  - value of pi? (...wrong and unpredictable)
- set environment variable OMP\_NUM\_THREADS to 4 and run
  - value of pi? (...and stays wrong)
- run again
  - value of pi? (...but where is the race-condition?)
- After *this test with race-conditions*, you may compare your result with the provided solution:

- ../../solution/pi/pi1.c or pi1.f90

© 2000-2023 HLRS, Rolf Rabenseifner  $\left[ \bullet \overrightarrow{REC} \rightarrow \underline{online} \right]$ 

Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercise 2a

#### **OpenMP Exercise 2a: pi Program (3)**

- **add** private(x) **clause in** pi.[c|f90] **and compile**
- set environment variable OMP\_NUM\_THREADS to 2 and run
  - value of pi? (should be still incorrect ...)
- run again
  - value of pi?
- set environment variable OMP\_NUM\_THREADS to 4 and run
  - value of pi?
- run again
  - value of pi? (... and where is the second race-condition?)
- After *this test with race-conditions*, you may compare your result with the provided solution:

- ../../solution/pi/pi2.c or pi2.f90

- And where is the second race-condition?
  - → We'll resolve this in Exercise 2b after the talk "Synchronization constructs, e.g., critical regions"

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>



#### During the Exercise (15 min.)

Please stay here in the main room while you do this exercise And have fun with this short exercise (5)Please do not look at the solution before you finished this exercise, otherwise, 90% of your learning outcome may be lost As soon as you finished the exercise, please go to your breakout room and continue your discussions with your fellow learners: *Exercise* + *Discussion* + *Break* – *planned*: 15 *Minutes* Discussion: Do you have ideas how one can solve the problem based on the knowledge you learnt so far? Please do not program – only discuss!

# Outline — Tasks

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax

as ks

- Data environment (part 1)
- Environment variables
- Runtime library routines
- Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - **Tasks** (a task generating construct)
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Tasks

#### **OpenMP** task **Directive** – **Example**:

#### **Parallelized traversing of a tree**

# OpenMP 3.0



**OpenMP 3.0** 

- The task construct defines an explicit task.
- Fortran: !\$OMP task [clause [[,] clause]...] block !\$OMP end task
- **C/C++** C/C++:

#pragma omp task[clause[[,]clause]...]new-line
 structured-block

- Clauses:
  - untied
  - default(shared | none | private | firstprivate)
  - private(list)
  - firstprivate(*list*)
  - shared(*list*)
  - if (scalar expression)

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Tasks

- When a thread encounters a task construct, a task is generated from the code for the associated structured block.
- The encountering thread
  - may immediately execute the task,
  - or may defer its execution.

- The number of tasks can be limited, e.g., to the number of threads.
- Completion of a task can be guaranteed using task synchronization constructs → taskwait construct.
- When if (false) clause exists, then execution is "serial"
- Task scheduling points:
  - In the generating task: Immediately following the generation of an explicit task.
  - In the generated task: After the last instruction of the task region.
  - If task is "untied": Everywhere inside of the task.
  - In implicit and explicit barriers.
  - In taskwait.

At task scheduling points, tasks can be resumed or suspended. (Further constraints  $\rightarrow$  OpenMP 3.0, Sect. 2.7.1, page 62)

# **Outline** — Synchronization constructs

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online

# **OpenMP Synchronization**

- Implicit Barrier
  - end of parallel constructs
  - end of all other control constructs
  - implicit synchronization on control constructs can be removed with nowait clause
- Explicit
  - critical
  - ...

© 2000-2023 HLRS, Rolf Rabenseifner  $\left[ \bullet \text{REC} \rightarrow \text{online} \right]$ Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Synchronization

#### **OpenMP** critical **Directive**

- Enclosed code
  - executed by all threads, but
  - restricted to only one thread at a time

Fortran

C/C++

```
    Fortran:

            !$OMP CRITICAL [(name)]
            block
            !$OMP END CRITICAL [(name)]

    C/C++:
```

```
#pragma omp critical[(name)] new-line
    structured-block
```

• A thread waits at the beginning of a critical region until no other thread in the team is executing a critical region with the same name. All unnamed critical directives map to the same unspecified name.







Fortran: cnt = 0f=7 **!\$OMP PARALLEL !\$OMP DO** do i=1,20 if (b(i).eq.0) then **!**\$OMP CRITICAL cnt = cnt+1**!\$OMP END CRITICAL** endif a(i) = b(i) + f \* iend do **!\$OMP END DO !**\$OMP END PARALLEL



#### Fortran OpenMP critical — another example (Fortran)



# C/C++ OpenMP critical — another example (C/C++)

mx = 0;**mx=0** #pragma omp parallel private(pmax) pmax pmax pmax pmax pmax = 0; /\* or most negative number \*/ =0 **~=0** =0 =0 #pragma omp for private(r) nowait for (i=0; i<20; i++) 5,9 0.4r = work(i);r=... r=... pmax = (r>pmax ? r : pmax); pmax pmax pmax pmax } /\*end for\*/ endfor endfor endfor endfo /\*omp end for\*/ Only once per mx=.. #pragma omp critical thread mx=. mx= (pmax>mx ? pmax : mx); mx=. /\*omp end critical\*/ <mark>mx=..</mark> } /\*omp end parallel\*/
## Outline — Nesting and Binding

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

## **OpenMP Vocabulary**

- **Static extent** of the parallel construct: statements enclosed lexically within the construct
- **Dynamic extent** of the parallel construct: further includes the routines called from within the construct

#### Orphaned Directives:

Do not appear in the lexical extent of the parallel construct but lie in the dynamic extent

- Parallel constructs at the top level of the program call tree
- Directives in any of the called routines

[The terms *lexical extent* and *dynamic extent* are no longer used in OpenMP 2.5,

but still helpful to explain the complex impact of OpenMP directives. ]

#### **OpenMP Vocabulary**



Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Nesting & Binding

## **OpenMP Control Structures** — Summary

- Parallel region construct
  - parallel
- Worksharing constructs
  - sections
  - for (C/C++)
  - do (Fortran)
  - workshare (Fortran)
  - Single
- Task generating construct
  - task
- Combined parallel worksharing constructs [see later]
  - parallel for (C/C++)
  - parallel do (Fortran)
  - parallel workshare (Fortran)
- Synchronization constructs
  - critical

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>



C/C++ Fortran

## Outline — Exercise 2b: pi

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>



## **OpenMP Exercise 2b: pi Program (4)**

- Use your result from Exercise 2a (or solution pi2.[c|f90])
- Add critical directive in your pi.[c|f90] around the sumstatement and compile
- set environment variable OMP\_NUM\_THREADS to 2 and run
  - value of pi? (should be now correct!, but huge wtime!)
- run again
  - value of pi? (but not reproducible in the last bit!)
- set environment variable OMP\_NUM\_THREADS to 4 and run
  - value of pi? execution time? (Oh, does it take longer?)
- run again

- Look at wall-clock time, not at CPU time!
- value of pi? execution time?
- How can you optimize your code?
- After *successful execution*, you may compare your result with the provided solution:

- ../../solution/pi/pic.c or pic.f90

## **OpenMP Exercise 2b: pi Program (5)**

- move critical directive in pi. [c|f90] outside loop, and compile
- set environment variable OMP\_NUM\_THREADS to 2 and run
  - value of pi?
- set environment variable OMP NUM THREADS to 4 and run
  - value of pi? execution time? (correct pi, half execution time)
- run again
  - value of pi? execution time?
- After *successful execution*, you may compare your result with the provided solution:

- ../../solution/pi/pic2.c or pic2.f90

OpenMP-parallel with 4 threads			
computed pi = 3.1415926	5358967		
CPU time (clock)	=	0.01659	sec
wall clock time (omp get wtime)	=	0.01678	sec
wall clock time (gettimeofday)	=	0.01679	sec



### **OpenMP Exercise 2: pi Program - Solution**

#### Location: ~/OpenMP/solution/pi

- pi.[c|f90] original program
- pil.[c|f90] incorrect (no private, no synchronous global access) !!!
- pi2.[c|f90]

- incorrect (still no synchronous global access to sum) !!!
- pic.[c|f90] solution with critical directive, but extremely slow!
- pic2.[c]f90] solution with critical directive outside loop



Please stay here in the main room while you do this exercise

And have fun with this short exercise

Please do not look at the solution before you finished this exercise, otherwise, 90% of your learning outcome may be lost



As soon as you finished the exercise, please go to your breakout room

and continue your discussions with your fellow learners:

*Exercise* + *Discussion* + *Break* – *planned*: 25 *Minutes* 



## **Outline – Data environment and combined constructs**

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax

& Combined Constr.

En<.

Data

- Data environment (part 1)
- Environment variables
- Runtime library routines
- Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

#### **OpenMP Data Scope Clauses**

- private (*list*)
   Declares the variables in *list* to be private to each thread in a team.
- shared (*list*)
   Makes variables that appear in *list* shared among all the threads in a team
- If not specified: default shared
- **Exceptions**: private
  - stack (local) variables in called subroutines
  - Automatic variables within a block
  - Loop control variable of parallel DO (Fortran) and FOR (C) loops
  - Fortran only:
    - Loop control variable of a sequential loop enclosed in a parallel or task
       construct
    - Implied-do and forall indices
- Recommendation for C/C++:
  - Avoid private variables, use variables local to a block instead

Further information: OpenMP 3.0, Sect. 2.9.1, page 77

#### **Private Clause**

- private (variable) creates a local incarnation of the variable for each thread
  - value is uninitialized
  - private copy is not storage associated with the original

```
program wrong
      JLAST = -777
!$OMP PARALLEL
! $OMP
       DO PRIVATE (JLAST)
         DO J=1,1000
             JLAST = J
         END DO
! $OMP
       END DO
!$OMP END PARALLEL
      print *, JLAST
                          \rightarrow writes -777
                                                     (OpenMP \ge 3.0)
                                                     (OpenMP \leq 2.5)
                            or undefined value
```

- If initialization is necessary use firstprivate( var )
- If value is needed after loop use lastprivate ( var )
  - $\rightarrow$  var is updated by the thread that computes
    - the sequentially last iteration (on do or for loops)
    - the last section
- Nested private (var) with same var  $\rightarrow$  allocates again new private storage
- Sometimes shared/private is undefined → see OpenMP 3.0, Example A.30.2c/f pp. 234-235

#### **OpenMP** reduction **Clause**

- reduction (*operator:list*)
- Performs a reduction on the variables that appear in *list*, with the operator operator
- operator: one of
  - Fortran:

```
+, *, -, .and., .or., .eqv., .neqv.,
max, min, iand, ior, Or ieor
```

– C/C++:



Fortran

+, \*, -, &, ^, |, &&, or || With OpenMP 3.1 and later: max, min

- Variables must not be private in the enclosing context
- With OpenMP 2.0 and later, variables can be arrays (Fortran)
- At the end of the reduction, the shared variable is updated to reflect the result of combining the original value of the shared reduction variable with the final value of each of the private copies using the operator specified

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \text{online}]$ 



Fortran:

sm = 0	sm=0			
!\$OMP PARALLEL				
<pre>!\$OMP DO private(r), reduction(+:sm)</pre>	<b>Ni</b> ≢N			
do i=1,20	1,5	6,10	11,15	16,20
r = work(i)	<u>r=</u>	<b>r=</b>	<u>r=</u>	<b>r=</b>
sm = sm + r	sm= sm+r	sm= sm+r	sm= sm+r	sm= sm+r
end do	enddo	enddo	enddo	enddo
!\$OMP END DO				
<b>!\$OMP END PARALLEL</b>				

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  online



C / C++:

sm = 0; #pragma omp parallel #pragma omp for reduction(+:sm) for( i=0; i<20; i++) double r; r = work(i); sm = sm + r;} /\*end for\*/ /\*omp end for\*/ /\*omp end parallel\*/



## **OpenMP Combined** parallel do/for **Directive**

- Shortcut form for specifying a parallel region that contains a single do/for directive
- Fortran
   Fortran:
   !\$OMP PARALLEL DO [clause [[,] clause]...]
   do\_loop
   [!\$OMP END PARALLEL DO]
  - C/C++: #pragma omp parallel for [clause [clause]...] new-line for-loop
    - This directive admits all the clauses of the parallel directive and the do/for directive except the nowait clause, with identical meanings and restrictions

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ 

C/C++

# Fortran OpenMP Combined parallel do/for — an example

Fortran:



© 2000-2023 HLRS, Rolf Rabenseifner  $\left[ \bullet \text{REC} \rightarrow \underline{\text{online}} \right]$ 

# **C/C++** OpenMP Combined parallel do/for — an example

C / C++:



## Outline — Exercise 3: pi with reduction

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online



## **OpenMP Exercise 3: pi Program (1)**

- Goal: usage of
  - worksharing constructs: do/for
  - critical directive
  - reduction clause
  - combined parallel worksharing constructs:
     parallel do/parallel for
- Working directory: ~/OpenMP/#NR/pi/

#NR = number of your PC, e.g., 01

- Use your result pi.[c|f90] from the exercise 2
- or copy solution of exercise 2 to your directory:

- cp ../../solution/pi/pic2.\* .

• or copy slow "*critical inside loop*" solution of exercise 2 to your directory:

- cp ../../solution/pi/pic.\* .

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \underline{\text{online}}]$ 

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 3



- remove critical directive in pi.[c|f90],
   remove your additional partial sum variable, i.e., restore original numeric,
   and add reduction clause and compile
- set environment variable OMP\_NUM\_THREADS to 2 and run
  - value of pi?
- run again
  - value of pi?
- set environment variable OMP\_NUM\_THREADS to 4 and run
  - value of pi? execution time?
- run again
  - value of pi? execution time?
- After *successful execution*, you may compare your result with the provided solution:
  - ../../solution/pi/pir.c or pir.f90

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>



## **OpenMP Exercise 3: pi Program (3)**

- change parallel region + do/for to the combined parallel worksharing construct parallel do/parallel for and compile
- set environment variable OMP NUM THREADS to 2 and run ٠
  - value of pi?
- run again ٠
  - value of pi?
- set environment variable OMP NUM THREADS to 4 and run ٠
  - value of pi?
- run again ٠
  - value of pi?
- After *successful execution*, you may compare your result with the ٠ provided solution:
  - ../../solution/pi/pir2.c or pir2.f90
- At the end, compile again **without** OpenMP ٠
  - Does your code still compute a **correct** value of **pi**?

© 2000-2023 HLRS, Rolf Rabenseifner ● REC → online

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 3





### **OpenMP Exercise 3: pi Program - Solution**

#### Location: ~/OpenMP/solution/pi

- pi.[c|f90] O
- pi1.[c|f90]
- pi2.[c|f90]
- pic.[c|f90]
- pic2.[c|f90]
- pir.[c|f90]
- pir2.[c|f90]

- original program
- incorrect (no private, no synchronous global access) !!!
  - incorrect (still no synchronous global access to sum) !!!
  - solution with critical directive, but extremely slow!
  - solution with critical directive outside loop
  - solution with reduction clause
  - solution with combined parallel do/for
    and reduction clause

#### **OpenMP Exercise 3: pi Program - Execution Times F90**



© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 3

## **OpenMP Exercise 3: pi Program - Summary**

- Decision about private or shared status of variables is important
- Correct results with reduction clause and with critical directive
- Using the simple version of the critical directive is much more time consuming than using the reduction clause  $\Rightarrow$  no parallelism left
- More sophisticated use of critical directive leads to much better performance
- Convenient reduction clause
- Convenient shortcut form

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 3



Please stay here in the main room while you do this exercise

And have fun with this short exercise 3

Please do not look at the solution before you finished this exercise, otherwise, 90% of your learning outcome may be lost



As soon as you finished the exercise, please go to your breakout room

and continue your discussions with your fellow learners:

*Exercise* + *Discussion* + *Break* – *planned*: 20 *Minutes* 



Already in your breakout room, you may test different loop schedules (static, dynamic, guided, auto) and share your results with your colleagues.

Did you oversubscribe Vyour hardware?

## Outline — Exercise 4: Heat

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat Conduction Exercise
- Summary of OpenMP API
- OpenMP Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>



#### **OpenMP Exercise: Heat Conduction(1)**

- solves the PDE for unsteady heat conduction  $df/dt = \Delta f$
- uses an explicit scheme: forward-time, centered-space
- solves the equation over a unit square domain
- initial conditions: *f*=0 everywhere inside the square
- boundary conditions: *f*=*x* on all edges
- number of grid points: 20x20



## **OpenMP Exercise: Heat Conduction (2)**

- Goals:
  - parallelization of a real application
  - usage of different parallelization methods with respect to their effect on execution times
- Working directory: ~/OpenMP/#NR/heat/

#NR = number of your PC, e.g., 01

- Serial programs:
  - Fortran: heat.F90
  - C: heat.c
- Compiler calls:
  - See login slides
- Options:

Fortran

C/C++



© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 4: Heat



## **OpenMP Exercise: Heat Conduction (3)**

Tasks:



Parallelize heat.c or heat.F

- Use critical sections for global maximum: Use trick with partial maximum inside of the parallelized loop, and critical section outside of the loop to compute global maximum
- Or just with: reduction(max:dphimax)
- Hints:
  - Parallelize outer loop (index i in Fortran, k in C)
  - make inner loop index private!



Compile and run with 80x80 serial, and parallel with 1, 2, 3, 4 threads

• Result may look like

Serial: 0.4 sec, 1 thread: 0.5 sec, 2 threads: 2.8 sec, ...

• Why is the parallel version significantly slower than the serial one?

## **OpenMP Exercise: Heat Conduction (4)**

- Reason already in the **serial** program:
  - Bad sequence of the nested loops





- Inner loop should use contiguous index in the array, i.e.,
  - − First index in Fortran → "do i=…" must be inner loop
  - Second index in C/C++  $\rightarrow$  "for (k=...)" must be inner loop

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 4: Heat

## **OpenMP Exercise: Heat Conduction (5)**



Interchange sequence of nested loops for i and k forget to modify name

Compile and run parallel with 80x80 and with 1, 2, 3, 4 threads

- Result may look like
  - → 1 thread: 0.5 sec, 2 threads: 0.45 sec, 3 threads: 0.40 sec
- Reasons:
  - Problem is too small parallelization overhead too large



- $\rightarrow$  1 thread: 4.24 sec, 2 threads: <u>2.72</u> sec, 3 threads: <u>2.27 sec</u>
- $\rightarrow$  Don't worry that computation is prematurely finished by itmax=15000



**With 1000x1000 and** -Ditmax=500 and with 1, 2, 3, 4 threads

- $\rightarrow$  1 thread: 5.96 sec, 2 threads: 2.79 sec, 3 threads: 1.35 sec
- → Super-linear speed-up due to better cache reuse on smaller problem

#### OpenMP Exercise: Heat Conduction (6) Advanced exercise



#### OpenMP Exercise: Heat Conduction (7) Advanced exercise



Execute abort-statement (if (dphimax<eps) ... ) only each 20th it=... iteration

Move omp barrier directly after if (dphimax<eps) ... that this barrier is also executed only each  $20^{th}$  it=... iteration



Add schedule(runtime) and compare execution time

Fortran TODO Fortran only:

Substitute critical-section-trick

by reduction (max:dphimax) clause

© 2000-2023 HLRS, Rolf Rabenseifner  $\left( \bullet \text{ REC} \rightarrow \text{online} \right)$ 





## **OpenMP Exercise: Heat - Solution (1)**

#### Location: ~/OpenMP/solution/heat

- heat.[F|c] Original program
- heat\_x.[F|c] Better serial program with interchanged nested loops
- heatc.[F|c] Extremely slow solution with critical section inside iteration loop
- heatc2.[F|c] Slow solution with critical section outside inner loop, one parallel region inside time step iteration loop (it=...)
- heatc2\_x.[F|c] Fast solution with critical section outside inner loop, one parallel region inside iteration loop, interchanged nested loops
- heatc3\_x.[F|c] ... and parallel region outside of it=... loop
- heatc4\_x.[F|c] ... and abort criterion only each 20<sup>th</sup> iteration
- heats2\_x.[F|c] Solution with schedule(runtime) clause
- heatr2\_x.[F|c] Solution with reduction clause, one parallel region inside iteration loop

#### **OpenMP Exercise: Heat - Solution (2)**

- heatc2 → heatc2\_x
   Loss of optimization with OpenMP directives (and compilers)
- For controlling the parallelization:
  - Version 20x20: 1116 iterations
  - Version 80x80: 14320 iterations
  - Version 250x250: 15001 iterations [if itmax = 15000 (default)]

110996 iterations [if itmax is extended to 150000]

 heatc2\_x ←→ heatc3\_x Additional overhead for barriers and single sections (including implied barrier) must be compared with fork-join-overhead

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Data environment and combined constructs  $\rightarrow$  Exercise 4: Heat
### **OpenMP Exercise: Heat - Solution (3) – 20x20 Time**



Measurements on NEC TX-7 (asama.hww.de), 16 cores, May 4-5, 2006

#### **OpenMP Exercise: Heat - Solution (4) – 20x20 Efficiency**



### **OpenMP Exercise: Heat - Solution (5) – 80x80 Time**



#### **OpenMP Exercise: Heat - Solution (6) – 80x80 Efficiency**



© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \underline{\text{online}}]$ 

### **OpenMP Exercise: Heat - Solution (7) – 250x250 Time**



### **OpenMP Exercise: Heat - Solution (8) – 250x250 Efficiency**



### **OpenMP Exercise: Heat - Solution (9) – 1000x1000 Time**



#### **OpenMP Exercise: Heat - Solution (10)** – 1000x1000 Efficiency



© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ 

Introduction to OpenMP → Data environment and combined constructs → Exercise 4: Heat

#### **OpenMP Exercise: Heat - Execution Times F90 with 150x150**



## **OpenMP Exercise: Heat Conduction - Summary**

- Overhead for parallel versions using 1 thread.
- Be careful with compiler based optimizations.
- Datasets must be large enough to achieve good speed-up.
- Intel Inspector or another *race condition detection tool* should be used to guarantee zero race conditions.
- Be careful when using other than default scheduling strategies:
  - dynamic is generally expensive
  - static: overhead for small chunk sizes is clearly visible

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ 

# Outline — Summary of the OpenMP API

- Introduction into OpenMP
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables
- Worksharing directives
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks

Summary of OpenMP AP

- Synchronization constructs, e.g., critical regions
- Nesting and Binding
- Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs
  - Exercise 4: Heat
- Summary of OpenMP API
- OpenMP Pitfalls

## **OpenMP Components**

- Set of compiler directives
  - Control Constructs
    - Parallel Regions
    - Worksharing constructs
  - Data environment
  - Synchronization
- Runtime library functions
- Environment variables

### **OpenMP Architecture**



### **OpenMP Constructs**



## Outline — OpenMP Pitfalls

- Introduction into OpenMP ۲
- Programming and Execution Model
  - Parallel regions: team of threads
  - Syntax
  - Data environment (part 1)
  - Environment variables
  - Runtime library routines
  - Exercise 1: Parallel region / library calls / privat & shared variables \_
- Worksharing directives ۰
  - Which thread executes which statement or operation?
  - Exercise 2a: Pi
  - Tasks
  - Synchronization constructs, e.g., critical regions
  - Nesting and Binding \_
  - Exercise 2b: Pi
- Data environment and combined constructs
  - Private and shared variables, Reduction clause
  - Combined parallel worksharing directives
  - Exercise 3: Pi with reduction clause and combined constructs \_
  - Exercise 4: Heat
- Summary of OpenMP API ۲
- **OpenMP** Pitfalls

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet \text{REC} \rightarrow \underline{\text{online}}]$ Introduction to OpenMP  $\rightarrow$  Pitfalls

**OpenMP** Pitfalls





See Appendix D of the OpenMP 4.0 standard, e.g.:

- The size of the first chunk in SCHEDULE(GUIDED)
- default schedule for SCHEDULE(RUNTIME) ٠
- default schedule •
- default number of threads ٠
- default for dynamic thread adjustment ٠
- number of levels of nested parallelism supported ٠
- atomic directives might be replaced by critical regions ٠
- behavior in case of thread exhaustion ٠
- allocation status of allocatable arrays that are not affected by COPYIN ٠ clause are undefined if dynamic thread mechanism is enabled
- Fortran: Is include 'omp lib.h' or use omp\_lib ٠ or both available?



## Implied flush directive

- A FLUSH directive identifies a sequence point at which a consistent view of ٠ the shared memory is guaranteed
- It is implied at the following constructs: ٠
  - BARRIER
  - CRITICAL and END CRITICAL
  - END {DO, FOR, SECTIONS}
  - END {SINGLE, WORKSHARE}
  - ORDERED AND END ORDERED
  - PARALLEL and END PARALLEL with their combined variants
  - Immediately before and after every task scheduling point (OpenMP  $\geq$  3.0)
- It is NOT implied at the following constructs: ٠
  - Begin of DO, FOR, WORKSHARE, SECTIONS
  - Begin of MASTER and END MASTER
  - **Begin of SINGLE**



\*) Faster memory for temporary copies of recently used data

Recommendation: Do **not** loop with testing on data modifications caused by other threads (without correct and sufficient flushing) Use OpenMP synchronization methods More secure:

© 2000-2023 HLRS, Rolf Rabenseifner Introduction to OpenMP  $\rightarrow$  Pitfalls



Slide 129/151

### Two types of SMP errors

- Race Conditions
  - Data-Race: Two threads access the same shared variable and at least one thread modifies the variable and the accesses are concurrent, i.e. unsynchronized
  - The outcome of a program depends on the detailed timing of the threads in the team.
  - This is often caused by unintended share of data
- Deadlock
  - Threads lock up waiting on a locked resource that will never become free.
    - Avoid lock functions if possible
    - At least avoid nesting different locks





```
!$OMP PARALLEL SECTIONS
A = B + C
!$OMP SECTION
B = A + C
!$OMP SECTION
C = B + A
!$OMP END PARALLEL SECTIONS
```

- The result varies unpredictably based on specific order of execution for each section.
- Wrong answers produced without warning!



```
!$OMP PARALLEL SHARED (X), PRIVATE(ID, TMP)
       ID = OMP GET THREAD NUM()
!$OMP DO REDUCTION(+:X)
       DO 100 I=1,100
              TMP = WORK1(I)
              X = X + TMP
100
     CONTINUE
!$OMP END DO NOWAIT
       Y(ID) = WORK2(X, ID)
!$OMP END PARALLEL
```

- The result varies unpredictably because the value of X isn't ٠ dependable until the barrier at the end of the do loop.
- Solution: Be careful when you use **NOWAIT**. ٠



# **OpenMP programming recommendations**

- Write SMP code that is
  - portable and
  - equivalent to the sequential form.
- Use tools like "Intel<sup>®</sup> Inspector" (formerly "Intel<sup>®</sup> Thread Checker" and "Assure").



## **Sequential Equivalence**

- Two forms of sequential equivalence
  - Strong SE: bitwise identical results.
  - Weak SE: equivalent mathematically but due to quirks of floating point arithmetic, not bitwise identical.
- Using a limited subset of OpenMP e.g., no locks
- Advantages:
  - program can be tested, debugged and used in sequential mode
  - this style of programming is also less error prone





## **Rules for Strong Sequential Equivalence**

- Control data scope with the base language
  - Avoid the data scope clauses.
  - Only use private for scratch variables local to a block (eg. temporaries or loop control variables) whose global initialization don't matter.
- Locate all cases where a shared variable can be written by multiple threads.
  - The access to the variable must be protected.
  - If multiple threads combine results into a single value, enforce sequential order.
  - Do not use the reduction clause carelessly.
     (no floating point operations +,-,\*)
  - Use the ordered directive and the ordered clause.
- Concentrate on loop parallelism/data parallelism



### **Example for Ordered Clause: pio.c** / .f90



- "ordered" corresponds to "critical" + "order of execution"
- only efficient if workload outside ordered directive is large enough





## Rules for weak sequential equivalence

- For weak sequential equivalence only mathematically valid constraints are enforced.
  - Floating point arithmetic is not associative and not commutative.
  - In many cases, no particular grouping of floating point operations is mathematically preferred so why take a performance hit by forcing the sequential order?
  - In most cases, if you need a particular grouping of floating point operations, you have a bad algorithm.
- How do you write a program that is portable and satisfies weak sequential equivalence?
  - Follow the same rules as the strong case, but relax sequential ordering constraints.



## **Thread-safe library functions**

- Library functions (if called inside of parallel regions) ٠ must be thread-safe
- **Automatically switched** if OpenMP option is used: ٠
  - e.g., Intel compiler:
    - ifort -qopenmp -o my prog my prog.f or my prog.f90 (Fortran) •
    - (C, C++)• icc -qopenmp -o my prog my prog.c
- Manually by compiler option: ٠
  - e.g., IBM compiler:
    - xlf **r** -O -qsmp=omp -o my prog my prog.f (Fortran, fixed form)
    - xlf90 r -O -qsmp=omp -o my prog my prog.f90 (Fortran, free form)
    - xlc r -O -qsmp=omp -o my prog my prog.c
    - The "\_r" forces usage of reentrant library functions (not identical with thread-safe)
- **Manually by programmer:** Some library function are using an internal • buffer to store its state – one must use its reentrant counterpart:
  - e.g., reentrant erand48() instead of drand48()

gmtime r() qmtime()



(C)

## **Optimization Problems** – **Overview**

- Prevent unnecessary fork and join of parallel regions
  - if you can execute several loop / workshare / sections / single inside of one parallel region
- Prevent unnecessary synchronizations
  - e.g. with critical or ordered regions inside of loops
- Prevent false-sharing (of cache-lines)  $\rightarrow 3^{rd}$  next slide
- Prevent unnecessary cache-coherence or memory communication
  - E.g., same schedules for same memory access patterns  $\rightarrow$  6<sup>th</sup> next slide
  - First touch on (cc)NUMA architectures  $\rightarrow$  8<sup>th</sup> next slide
    - To locate arrays/objects already in a parallelized initialization to the threads where they are mainly used
  - Pin the threads to CPUs [not useful in time sharing on over-committed systems]
    - Otherwise, after each time slice, threads may run on other CPUs
    - numactl, or LIKWID (portable !!!)
    - dplace -x2 (SGI), O(1) scheduler (HP), SUNW\_OMP\_PROBIND envir. (Sun)
    - fplace -r -o 1,2 command (Intel), ...
       Do not pin the management threads

**Optimization Problems** 



## Get a feeling for the involved overheads

Operation	Minimum overhead (cycles)	Scalability
Hit L1 cache	1-10	Constant
Function call	10-20	Constant
Thread ID	10-50	Constant, log, linear
Integer divide	50-100	Constant
Static do/for, no barrier	100-200	Constant
Miss all caches	100-300	Constant
Lock acquisition	100-300	Depends on contention
Dynamic do/for, no barrier	1000-2000	Depends on contention
Barrier	200-500	Log, linear
Parallel	500-1000	Linear
Ordered	5000-10000	Depends on contention

All numbers are approximate!! They are very platform dependent !!

© 2000-2023 HLRS, Rolf Rabenseifner  $[\bullet]$ Introduction to OpenMP  $\rightarrow$  Pitfalls



## **Optimization Problems**

Prevent frequent synchronizations, e.g., with critical regions ٠

```
max = 0;
#pragma omp parallel private(partial max)
   partial max = 0;
#pragma omp for
   for (i=0; i<10000; i++)</pre>
   ł
     x[i] = ...;
     if (x[i] > partial max) partial max = x[i];
   }
#pragma omp critical
   if (partial max > max) max = partial max;
}
```

- Loop: partial max is updated locally up to 10000/#threads times
- Critical region: max is updated only up to **#threads** times ٠



## **False-sharing**



- Several threads are accessing data through the same cache-line.<sup>\*)</sup>
- This cache-line has to be moved between these threads.
- This is very time-consuming.

<sup>\*)</sup> The smallest accessible unit in memory and caches, usually 64 or 128 bytes, i.e., 8 or 16 doubles







#### False-sharing – results from the experiment



Although each thread accesses independent variables, the performance will be terrible, if these variables are located in the same cache-line

Measurements on NEC TX-7 with **128 bytes cache lines**,

timings with false-sharing (stride 1-8 with more than 1 thread) were varying from run to run.

Test program, see Appendix: piarr.c

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  online Introduction to OpenMP  $\rightarrow$  Pitfalls

Slide 145/151

### False-sharing – same with speed-up



Measurements on NEC TX-7 with **128 bytes cache lines**, timings with false-sharing (stride 1-8 with more than 1 thread) were varying from run to run.



### Communication overhead, e.g., due to different schedules



shared memory







### Communication overhead – results from the experiment

Calculation of pi with vector chunks Wall-clock time - the smaller the better! Timing on NEC TX-7, with n=10,000,000



Introduction to OpenMP  $\rightarrow$  Pitfalls

Slide 148/151

#### First touch



### Acknowledgements

Thanks to

- Rainer Keller (Hochschule für Technik, Stuttgart),
- Matthias Müller (RWTH Aachen),
- Isabel Loebich,

the original co-authors of this course.

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  REC  $\rightarrow$  online

Introduction to OpenMP  $\rightarrow$  Acknowledgements
# **OpenMP Summary**

- Standardized compiler directives for shared memory programming
- Fork-join model based on threads
- Support from all relevant hardware vendors
- OpenMP offers an incremental approach to parallelism
- OpenMP allows to keep one source code version for scalar and parallel execution
- Equivalence to sequential program is possible if necessary
  - strong equivalence
  - weak equivalence
  - no equivalence
- OpenMP programming includes race conditions and deadlocks, but a subset of OpenMP can be considered safe
- Tools like the Intel<sup>®</sup> Inspector help to write correct parallel programs



# Appendix

© 2000-2023 HLRS, Rolf Rabenseifner  $\bullet$  F Introduction to OpenMP  $\rightarrow$  Appendix



Slide 152(App.)

# Appendix – Content

- Example hello parallel region written in C and Fortran
- Example pi, written in C
- Example pi, written in Fortran (free form)
- Example heat
  - heatc2.c parallel version in C, with critical region (4 pages)
  - heatr.f parallel version in Fortran, with reduction clause (4 pages)
- <u>Cache-line false-sharing experiment</u>: piarr.c
- <u>Communication overhead / different schedules experiment</u>: pivec.c



# **Example hello** – **Parallel region**

- hello6.c final solution in C
- hello6.f90 final solution in Fortran

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1

# hello5.c – Parallel region



Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1



## hello5.f90 – Parallel region



**!\$OMP END PARALLEL** 

end program hello

© 2000-2023 HLRS, Rolf Rabenseifner  $\bigcirc$  REC  $\rightarrow$  <u>online</u>

Introduction to OpenMP  $\rightarrow$  Programming and Execution Model  $\rightarrow$  Exercise 1



# Example pi, written in C

- pi.c sequential code
- pi0.c sequential code with a parallel region, verifying a team of threads
- pic2.c parallel version with a critical region outside of the loop
- pir.c parallel version with a reduction clause
- pir2.c parallel version with combined parallel for
- pio.c parallel version with ordered region
- pio2.c parallel version with ordered execution if the number of threads is fixed

© 2000-2023 HLRS, Rolf Rabenseifner  $\left( \bullet \text{REC} \rightarrow \text{online} \right)$ 

Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercises 2a + 2b + 3: Pi

# pi.c – sequential code

The include and timing blocks are removed on the next slides

#include <stdio.h> /\* pi = integral [0..1] 4/(1+x\*\*2) dx \*/ w=1.0/n;#include <time.h> #include <sys/time.h> sum=0.0;#ifdef OPENMP for (i=1;i<=n;i++)</pre> include <omp.h> include block x=w\*((double)i-0.5);#endif #define f(A) (4.0/(1.0+A\*A)) sum=sum+f(x); const int n = 10000000;pi=w\*sum; the calculation int main(int argc, char\*\* argv) timing block C t2=clock();int i; ifdef OPENMP double w,x,sum,pi; wt2=omp get wtime(); clock t t1,t2; # endif struct timeval tv1,tv2; gettimeofday(&tv2, &tz); struct timezone tz; printf( "computed pi =  $\$24.16q\n$ ", pi); # ifdef OPENMP printf( "CPU time (clock) double wt1, wt2; = %12.4g sec\n", (t2-t1)/1000000.0 ); timing block A endif # ifdef OPENMP # ifdef OPENMP printf( "wall clock time pragma omp parallel (omp get wtime) = \$12.4g sec\n", wt2-wt1); pragma omp single printf("OpenMP-parallel with %1d # endif threads\n", omp\_get\_num\_threads()); printf( "wall clock time (gettimeofday) } /\* end omp parallel \*/ -= \$12.4q sec\n", # endif **OpenMP check** # pragma omp barrier (tv2.tv sec-tv1.tv sec) + gettimeofday(&tv1, &tz); (tv2.tvusec-tv1.tvusec)\*1e-6);# ifdef OPENMP return 0; wtl=omp get wtime(); # endif t1=clock(); timing block B © 2000-2023 HLRS, Rolf Rabenseifner | REC  $\rightarrow$  online



# pi0.c – only verification of team of threads – without parallelization

```
--- INCLUDE BLOCK ---
             #define f(A) (4.0/(1.0+A*A))
             const int n = 1000000;
             int main(int argc, char** argv)
                int i;
                double w,x,sum,pi;
                --- TIMING BLOCK A ---
             # ifdef OPENMP
                 int myrank, num threads;
             # pragma omp parallel private(myrank,num threads)
                  myrank = omp get thread num();
                  num threads = omp get num threads();
                  printf("I am thread \frac{2}{2} of \frac{2}{2} threads\n", myrank, num threads);
                 } /* end omp parallel */
             # else
                  printf("This program is not compiled with OpenMPn");
             # endif
                --- TIMING BLOCK B ---
             /* calculate pi = integral [0..1] 4/(1+x**2) dx */
                w=1.0/n;
                sum=0.0;
                for (i=1;i<=n;i++)</pre>
                  x=w*((double)i-0.5);
                  sum=sum+f(x);
                pi=w*sum;
                --- TIMING BLOCK C ---
                return 0;
© 2000-2023 HLRS, Rolf Rabenseifner | \bullet \text{REC} \rightarrow \text{online} |
                                                                                      f90
Introduction to OpenMP \rightarrow Worksharing directives \rightarrow (Exercise 2a)
```

### pi2.c

#### parallelization with race-condition, i.e., WRONG program

```
--- INCLUDE BLOCK ---
#define f(A) (4.0/(1.0+A*A))
const int n = 1000000;
int main(int argc, char** argv)
  int i;
  double w,x,sum,pi;
  --- TIMING BLOCK A ---
  --- PRINT NUM THREADS ---
  --- TIMING BLOCK B ---
/* pi = integral [0..1] 4/(1+x**2) dx */
  w=1.0/n;
  sum=0.0;
#pragma omp parallel private(x), shared(w,sum)
# pragma omp for
  for (i=1;i<=n;i++)</pre>
    x=w*((double)i-0.5);
                            Read/write race-conditions on shared variable sum.
    sum=sum+f(x);
                            Or wrong results if sum would be declared as private.
} /*end omp parallel*/
  pi=w*sum;
  --- TIMING BLOCK C ---
  return 0;
```

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercise 2a



#### pic2.c

#### parallelization with critical region outside of the loop

```
--- INCLUDE BLOCK ---
           #define f(A) (4.0/(1.0+A*A))
           const int n = 1000000;
           int main(int argc, char** argv)
             int i;
             double w,x,sum,sum0,pi;
             --- TIMING BLOCK A ---
             --- PRINT NUM THREADS ---
             --- TIMING BLOCK B ---
           /* pi = integral [0..1] 4/(1+x**2) dx */
             w=1.0/n;
             sum=0.0;
           #pragma omp parallel private(x,sum0), shared(w,sum)
             sum0=0.0;
           # pragma omp for nowait
             for (i=1;i<=n;i++)</pre>
                x=w*((double)i-0.5);
                sum0=sum0+f(x);
           # pragma omp critical
                sum=sum+sum0;
           } /*end omp parallel*/
             pi=w*sum;
             --- TIMING BLOCK C ---
             return 0;
© 2000-2023 HLRS, Rolf Rabenseifner | \bullet \text{REC} \rightarrow \text{online} |
```



### pir.c – parallelization with reduction clause

```
--- INCLUDE BLOCK ---
#define f(A) (4.0/(1.0+A*A))
const int n = 10000000;
int main(int argc, char** argv)
  int i;
  double w,x,sum,pi;
  --- TIMING BLOCK A ---
  --- PRINT NUM THREADS ---
  --- TIMING BLOCK B ---
/* calculate pi = integral [0..1] 4/(1+x*2) dx */
  w=1.0/n;
  sum=0.0;
#pragma omp parallel private(x), shared(w,sum)
# pragma omp for reduction(+:sum)
  for (i=1;i<=n;i++)</pre>
    x=w*((double)i-0.5);
    sum=sum+f(x);
} /*end omp parallel*/
 pi=w*sum;
  --- TIMING BLOCK C ---
  return 0;
```



### pir2.c - combined parallel for with reduction clause

```
--- INCLUDE BLOCK ---
#define f(A) (4.0/(1.0+A*A))
const int n = 10000000;
int main(int argc, char** argv)
  int i;
  double w,x,sum,pi;
  --- TIMING BLOCK A ---
  --- PRINT NUM THREADS ---
  --- TIMING BLOCK B ---
/* calculate pi = integral [0..1] 4/(1+x*2) dx */
  w=1.0/n;
  sum=0.0;
#pragma omp parallel for private(x), shared(w), reduction(+:sum)
for (i=1;i<=n;i++)</pre>
    x=w*((double)i-0.5);
    sum=sum+f(x);
/*end omp parallel for*/
 pi=w*sum;
  --- TIMING BLOCK C ---
 return 0;
}
```

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{REC} \rightarrow \text{online})$ Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercise 3



### pio.c – parallelization with ordered clause

```
--- INCLUDE BLOCK ---
           #define f(A) (4.0/(1.0+A*A))
           const int n = 10000000;
           int main(int argc, char** argv)
             int i;
             double w,x,sum,myf,pi;
             --- TIMING BLOCK A ---
             --- PRINT NUM THREADS
             --- TIMING BLOCK B ---
           /* calculate pi = integral [0..1] 4/(1+x*2) dx */
             w=1.0/n;
             sum=0.0;
           #pragma omp parallel private(x,myf), shared(w,sum)
           # pragma omp for ordered
             for (i=1;i<=n;i++)</pre>
                                               Most compute time should be
               x=w*((double)i-0.5);
                                               outside of the ordered region
               myf = f(x);
               pragma omp ordered
           #
                  sum=sum+myf;
                                                             CAUTION
           } /*end omp parallel*/
                                                   The sequentialization of the
             pi=w*sum;
                                                ordered region may cause heavy
             --- TIMING BLOCK C ---
                                                     synchronization overhead
             return 0;
© 2000-2023 HLRS, Rolf Rabenseifner \left[ \bullet \text{REC} \rightarrow \text{online} \right]
                                                                                 <del>1</del>90
Introduction to OpenMP \rightarrow Pitfalls
```

### pio2.c - ordered, but only if number of threads is fixed

```
double w,x,sum,sum0,pi;
         /* calculate pi = integral [0..1] 4/(1+x*2) dx */
            w=1.0/n;
            sum=0.0;
         #pragma omp parallel private(x,sum0,num threads), shared(w,sum)
            sum0=0.0;
         #ifdef OPENMP
            num threads=omp get num threads();
         #else
            num threads=1
         #endif
         # pragma omp for schedule(static, (n-1)/num threads+1)
            for (i=1;i<=n;i++)</pre>
                                                                               CAUTION
              x=w*((double)i-0.5);
                                                                 1. Only if the number of threads is fixed
              sum0=sum0+f(x);
                                                                    AND if the floating point rounding is the
                                                                    same.
         #pragma omp for ordered schedule(static,1)
            for (i=0;i<num threads;i++)</pre>
                                                                    then the result is the same on two
                                                                    different platforms and any repetition of
         #pragma omp ordered
                                                                    this program.
              sum=sum+sum0;
                                                                2.
                                                                    This program cannot be verified with
          } /*end omp parallel*/
                                                                    Assure because it has to call
            pi=w*sum;
                                                                    omp get num threads().
                                                                3. To use Assure, the second loop must be
                                                                    substituted by the critical region as
                                                                    shown in pic2.c
© 2000-2023 HLRS, Rolf Rabenseifner \left[ \bullet \text{REC} \rightarrow \text{online} \right]
                                                                                           190
Introduction to OpenMP \rightarrow Pitfalls
                                                                                                 Slide 165(App.)
```

### **Example pi, written in Fortran (free form)**

- pi.f90 sequential code
- pi0.f90 sequential code with a parallel region, verifying a team of threads
- pic2.f90 parallel version with a critical region outside of the loop
- pir.f90 parallel version with a reduction clause
- pir2.f90 parallel version with combined parallel for
- pio.f90 parallel version with ordered region
- pio2.f90 parallel version with ordered execution if the number of threads is fixed

© 2000-2023 HLRS, Rolf Rabenseifner  $\left[ \bullet \text{ REC} \rightarrow \underline{\text{online}} \right]$ 

Introduction to OpenMP  $\rightarrow$  Worksharing directives  $\rightarrow$  Exercise 2a + 2b + 3: Pi

# pi.f90 – sequential code



# pi0.f90 – only verification of team of threads – without parallelization

```
program compute pi
             implicit none
             --- TIMING BLOCK A ---
             integer i
             integer, parameter :: n=1000000
             real(kind=8) w,x,sum,pi,f,a
             !$ integer omp get thread num, omp get num threads
             !$ integer myrank, num threads
                 logical openmp is used
             ! function to integrate
             f(a) = 4.0 \ 8/(1.0 \ 8+a*a)
             !$omp parallel private(myrank, num threads)
             !$ myrank = omp_get_thread_num()
             !$ num threads = omp get num threads()
             !$ write (*,*) 'I am thread', myrank, 'of', num threads, 'threads'
             !$omp end parallel
                  openmp is used = .false.
             !$ openmp is used = .true.
                  if (.not. openmp is used) then
                    write (*,*) 'This program is not compiled with OpenMP'
                  endif
             --- TIMING BLOCK B ---
             ! calculate pi = integral [0..1] 4/(1+x*2) dx
             w=1.0 8/n
             sum=0.0 8
             do i=1,\overline{n}
                x=w*(i-0.5 8)
                 sum=sum+f(\overline{x})
             enddo
             pi=w*sum
             --- TIMING BLOCK C ---
             end program compute pi
© 2000-2023 HLRS, Rolf Rabenseifner \left[ \bullet \text{REC} \rightarrow \text{online} \right]
                                                                                     back
Introduction to OpenMP \rightarrow Worksharing directives \rightarrow (Exercise 2a)
```

# pi2.f90

#### parallelization with race-condition, i.e., WRONG program

```
program compute pi
implicit none
--- TIMING BLOCK A ---
integer i
integer, parameter :: n=1000000
real(kind=8) w,x,sum,pi,f,a
! function to integrate
f(a) = 4.0 \ 8/(1.0 \ 8+a*a)
--- PRINT NUM THREADS ---
--- TIMING BLOCK B ---
! calculate pi = integral [0..1] 4/(1+x**2) dx
w=1.0 8/n
sum=0.0 8
!$OMP PARALLEL PRIVATE(x), SHARED(w, sum)
!SOMP DO
do i=1,n
   x=w*(i-0.5 8)
                         Read/write race-conditions on shared variable sum.
   sum=sum+f(x)
                         Or wrong results if sum would be declared as private.
enddo
!$OMP END DO
!$OMP END PARALLEL
pi=w*sum
--- TIMING BLOCK C ---
end program compute pi
```



### pic2.f90

#### parallelization with critical region outside of the loop

```
program compute pi
             implicit none
             --- TIMING BLOCK A ---
             integer i
             integer, parameter :: n=1000000
             real(kind=8) w, x, sum, sum0, pi, f, a
             ! function to integrate
             f(a) = 4.0 \ 8/(1.0 \ 8+a*a)
             --- PRINT NUM THREADS ---
             --- TIMING BLOCK B ---
             ! calculate pi = integral [0..1] 4/(1+x**2) dx
             w=1.0 8/n
             sum=0.0 8
              !$OMP PARALLEL PRIVATE(x,sum0), SHARED(w,sum)
             sum0=0.0 8
              !$OMP DO
             do i=1, n
                 x=w*(i-0.5 8)
                 sum0=sum0+\overline{f}(x)
             enddo
              !$OMP END DO NOWAIT
              !SOMP CRITICAL
                 sum=sum+sum0
              !$OMP END CRITICAL
              !$OMP END PARALLEL
             pi=w*sum
             --- TIMING BLOCK C ---
             end program compute pi
© 2000-2023 HLRS, Rolf Rabenseifner \bigcirc REC \rightarrow online
Introduction to OpenMP \rightarrow Worksharing directives \rightarrow Exercise 2b
```



### pir.f90 – parallelization with reduction clause

```
program compute pi
implicit none
--- TIMING BLOCK A ---
integer i
integer, parameter :: n=1000000
real(kind=8) w,x,sum,pi,f,a
! function to integrate
f(a) = 4.0 \ 8/(1.0 \ 8+a*a)
--- PRINT NUM THREADS ---
--- TIMING BLOCK B ---
! calculate pi = integral [0..1] 4/(1+x*2) dx
w=1.0 8/n
sum=0.0 8
!$OMP PARALLEL PRIVATE(x), SHARED(w, sum)
!$OMP DO REDUCTION (+:sum)
do i=1,n
   x=w*(i-0.5 8)
   sum=sum+f(\overline{x})
enddo
!SOMP END DO
!SOMP END PARALLEL
pi=w*sum
--- TIMING BLOCK C ---
end program compute pi
```



### pir2.f90 - combined parallel do with reduction clause

```
program compute pi
implicit none
--- TIMING BLOCK A ---
integer i
integer, parameter :: n=1000000
real(kind=8) w,x,sum,pi,f,a
! function to integrate
f(a) = 4.0 \ 8/(1.0 \ 8+a*a)
--- PRINT NUM THREADS ---
--- TIMING BLOCK B ---
! calculate pi = integral [0..1] 4/(1+x*2) dx
w=1.0 8/n
sum=0.0 8
!$OMP PARALLEL DO PRIVATE(x), SHARED(w), REDUCTION(+:sum)
do i=1,n
   x=w^{*}(i-0.5 8)
   sum=sum+f(x)
enddo
!SOMP END PARALLEL DO
pi=w*sum
--- TIMING BLOCK C ---
end program compute pi
```

© 2000-2023 HLRS, Rolf Rabenseifner  $(\bullet \text{ REC} \rightarrow \text{online})$ 





### pio.f90 – parallelization with ordered clause







### pio2.f90 - ordered, but only if number of threads is fixed

```
!--unused-- use omp lib
!$ integer omp get num threads
real(kind=8) w, x, sum, sum0, pi, f, a
! calculate pi = integral [0..1] 4/(1+x**2) dx
w=1.0 8/n
sum=0.0 8
!$OMP PARALLEL PRIVATE(x,sum0,num threads), SHARED(w,sum)
sum0=0.0 8
num threads=1
!$ num threads=omp get num threads()
! $OMP \overline{DO} SCHEDULE (\overline{STATIC}, (\overline{n}-1)/num threads+1)
do i=1,n
   x=w*(i-0.5 8)
   sum0=sum0+\overline{f}(x)
enddo
                                                1.
!SOMP END DO
!$OMP DO ORDERED SCHEDULE(STATIC,1)
                                                    same.
do i=1, num threads
SOMP ORDERED
   sum=sum+sum0
!$OMP END ORDERED
                                                    this program.
enddo
!SOMP END DO
!SOMP END PARALLEL
pi=w*sum
. . .
```

#### CAUTION

 Only if the number of threads is fixed AND if the floating point rounding is the same,

then the result is the same on two different platforms and any repetition of this program.

- This program cannot be verified with Assure because it has to call omp\_get\_num\_threads().
- To use Assure, the second loop must be substituted by the critical region as shown in pic2.f90



Introduction to OpenMP  $\rightarrow$  Pitfalls



# heatc2\_x.c – Parallelization of main loop and critical region (page 1 of 4) – declarations

```
#include <stdio.h>
             #include <sys/time.h>
             #ifdef OPENMP
                 include <omp.h>
             #endif
             \#define min(A,B) ((A) < (B) ? (A) : (B))
             \#define max(A,B) ((A) > (B) ? (A) : (B))
             #define imax 20
             #define kmax 11
             #define itmax 20000
             void heatpr(double phi[imax+1][kmax+1]);
             int main()
               double eps = 1.0e-08;
               double phi[imax+1][kmax+1], phin[imax][kmax];
               double dx, dy, dx2, dy2, dx2i, dy2i, dt, dphi, dphimax, dphimax0;
               int i,k,it;
               struct timeval tv1, tv2; struct timezone tz;
             # ifdef OPENMP
                  double wt1,wt2;
             # endif
             # ifdef OPENMP
             #
                 pragma omp parallel
                  {
             #
                    pragma omp single
                    printf("OpenMP-parallel with %1d threads\n",
                 omp get num threads());
                  } /* end omp parallel */
             # endif
© 2000-2023 HLRS, Rolf Rabenseifner | \bullet \text{REC} \rightarrow \text{online} |
```

## heatc2\_x.c (page 2 of 4) - initialization

```
dx=1.0/kmax; dy=1.0/imax;
                dx^2=dx^*dx; dv^2=dv^*dv;
                dx2i=1.0/dx2; dy2i=1.0/dy2;
                dt=min(dx2,dy2)/4.0;
              /* start values 0.d0 */
              #pragma omp parallel private(i,k) shared(phi)
              #pragma omp for
                for (i=1;i<imax;i++)</pre>
                { for (k=0;k<kmax;k++)</pre>
                   { phi[i][k]=0.0;
              /* start values 1.d0 */
              #pragma omp for
                for (i=0;i<=imax;i++)</pre>
                { phi[i][kmax]=1.0;
              }/*end omp parallel*/
              /* start values dx */
                phi[0][0]=0.0;
                phi[imax][0]=0.0;
                for (k=1; k<kmax; k++)</pre>
                { phi[0][k]=phi[0][k-1]+dx;
                  phi[imax][k]=phi[imax][k-1]+dx;
                printf("\nHeat Conduction 2d\n");
                printf("\ndx = \$12.4q, dy = \$12.4q, dt = \$12.4q, eps = \$12.4q\n",
                        dx, dy, dt, eps);
                printf("\nstart values\n");
                heatpr(phi);
© 2000-2023 HLRS, Rolf Rabenseifner | \bullet \text{REC} \rightarrow \text{online} |
```

## heatc2\_x.c (page 3 of 4) - time step integration

```
gettimeofday(&tv1, &tz);
              # ifdef OPENMP
                   wt1=omp get wtime();
              # endif
              /* iteration */
                for (it=1;it<=itmax;it++)</pre>
                 { dphimax=0.;
              #pragma omp parallel private(i,k,dphi,dphimax0) \
                   shared(phi,phin,dx2i,dy2i,dt,dphimax)
                   dphimax0=dphimax;
              #pragma omp for
                   for (i=1;i<imax;i++)</pre>
                   { for (k=0; k<kmax; k++) \</pre>
                     { dphi=(phi[i+1][k]+phi[i-1][k]-2.*phi[i][k])*dy2i
                            +(phi[i][k+1]+phi[i][k-1]-2.*phi[i][k])*dx2i;
                        dphi=dphi*dt;
                       dphimax0=max(dphimax0,dphi);
                       phin[i][k]=phi[i][k]+dphi;
                                                                  In C, phi and phin are contiguous
                                                                 in the last index [k]. Therefore the
              #pragma omp critical
                                                                  k-loop should be the inner loop!
                   dphimax=max(dphimax,dphimax0);
                                                                   This optimization is done in all
              /* save values */
                                                                        heat... x.c versions
              #pragma omp for
                   for (i=1;i<imax;i++)</pre>
                   { for (k=0; k<kmax; k++)</pre>
                      { phi[i][k]=phin[i][k];
              }/*end omp parallel*/
© 2000-2023 HLRS
                   if(dphimax<eps) break;</pre>
                                                      • REC \rightarrow online
Introduction to Oper
                 }
                                                                                           Slide 177(App.)
```

### heatc2\_x.c (page 4 of 4) - wrap up

```
# ifdef OPENMP
    wt2=omp get wtime();
# endif
  gettimeofday(&tv2, &tz);
  printf("\nphi after %d iterations\n",it);
  heatpr(phi);
# ifdef OPENMP
    printf( "wall clock time (omp get wtime) = %12.4g sec\n", wt2-wt1 );
# endif
  printf( "wall clock time (gettimeofday) = %12.4g sec\n", (tv2.tv sec-
   tv1.tv sec) + (tv2.tv usec-tv1.tv usec)*1e-6 );
void heatpr(double phi[imax+1][kmax+1])
{ int i,k,kl,kk,kkk;
  kl=6; kkk=kl-1;
  for (k=0; k<=kmax; k=k+kl)</pre>
  { if(k+kkk>kmax) kkk=kmax-k;
    printf("\ncolumns %5d to %5d\n",k,k+kkk);
    for (i=0;i<=imax;i++)</pre>
    { printf("%5d ",i);
      for (kk=0; kk <= kkk; kk++)
      { printf("%#12.4g",phi[i][k+kk]);
      printf("\n");
return;
```



# heatr2\_x.f – Parallelization of main loop with reduction clause (page 1 of 4) – declarations

```
program heat
                    implicit none
                    integer i,k,it, imax,kmax,itmax
             c using reduction
                    parameter (imax=20, kmax=11)
                    parameter (itmax=20000)
                    double precision eps
                    parameter (eps=1.d-08)
                    double precision phi(0:imax,0:kmax), phin(1:imax-1,1:kmax-1)
                    double precision dx, dy, dx2, dy2, dx2i, dy2i, dt, dphi, dphimax
             ! times using cpu time
                    real t0
                    real t1
             !--unused-- include 'omp lib.h'
             !$
                    integer omp_get_num_threads
             !$
                    double precision omp get wtime
             !$
                    double precision wt0, wt1
             !$omp parallel
             !$omp single
                   write(*,*)'OpenMP-parallel with',omp get num threads(),'threads'
             !$
             !$omp end single
             !$omp end parallel
             С
                    dx=1.d0/kmax
                    dy=1.d0/imax
                    dx^2=dx^*dx
                    dy2=dy*dy
                    dx2i=1.d0/dx2
                    dy2i=1.d0/dy2
                    dt=min(dx2,dy2)/4.d0
© 2000-2023 HLRS, Rolf Rabenseifner \bigcirc REC \rightarrow online
```

### heatr2\_x.f (page 2 of 4) – initialization

```
!$OMP PARALLEL PRIVATE(i,k), SHARED(phi)
              c start values 0.d0
              !SOMP DO
                     do k=0, kmax-1
                       do i=1, imax-1
                          phi(i,k)=0.d0
                       enddo
                     enddo
              !$OMP END DO
              c start values 1.d0
              !$OMP DO
                     do i=0,imax
                       phi(i, kmax) = 1.d0
                     enddo
              !SOMP END DO
              !$OMP END PARALLEL
              c start values dx
                     phi(0,0) = 0.d0
                     phi(imax, 0) = 0.d0
                     do k=1, kmax-1
                     phi(0, k) = phi(0, k-1) + dx
                     phi(imax, k) = phi(imax, k-1) + dx
                     enddo
              c print array
                     write (*,'(/,a)')
                         'Heat Conduction 2d'
                    f
                     write (*, '(/, 4(a, 1pg12.4))')
                            'dx =',dx,', dy =',dy,', dt =',dt,', eps =',eps
                     write (*,'(/,a)')
                    f
                          'start values'
                     call heatpr(phi, imax, kmax)
© 2000-2023 HLRS, Rolf Rabenseifner | \bullet \text{REC} \rightarrow \text{online} |
```

### heatr2\_x.f (page 3 of 4) - time step integration



© 2000-2023 HLRS, Rolf Rabenseifner  $\left( \bullet \text{REC} \rightarrow \underline{\text{online}} \right)$ 

### heatr2\_x.f (page 4 of 4) – wrap up

```
call cpu time(t1)
!$
      wt1=omp get wtime()
c print array
      write (*,'(/,a,i6,a)')
          'phi after', it, ' iterations'
     f
      write (*,'(/,a,1pg12.4)') 'cpu time : ', t1-t0
      write (*,'(/,a,1pg12.4)') 'omp_get wtime:', wt1-wt0
!$
С
      stop
      end
С
С
      subroutine heatpr(phi, imax, kmax)
      double precision phi(0:imax,0:kmax)
С
      k1=6
      kkk=k1-1
      do k=0, kmax, kl
      if(k+kkk.gt.kmax) kkk=kmax-k
      write (*,'(/,a,i5,a,i5)') 'columns',k,' to',k+kkk
      do i=0, imax
      write (*, '(i5,6(1pq12.4))') i, (phi(i,k+kk),kk=0,kkk)
      enddo
      enddo
С
      return
      end
```

# False-sharing – an experiment (solution/pi/piarr.c)

```
n = 10000000; w=1.0/n; sum=0.0;
stride=1; if (argc>1) stride=atoi(argv[1]); /* unit is "double" */
#pragma omp parallel private (x, index) shared(w, sum, p sum)
# ifdef OPENMP
    index=stride*omp get thread num();
# else
                                       The thread-locality of psum[index]
    index=0:
# endif
                                          variables is forced manually,
  p sum[index]=0;
                                          and in the same cache-line
#pragma omp for
                                           if stride is small enough
  for (i=1;i<=n;i++)</pre>
    x=w*((double)i-0.5);
    p_sum[index+(x>1?1:0)] = p_sum[index] + 4.0/(1.0+x*x);
    /\overline{*} The term (x>1?1:0) is always zero. It is used to prohibit
       register caching of of p sum[index], i.e., to guarantee that
       each access to this variable is done via cache in the memory. */
#pragma omp critical
    sum=sum+p sum[index];
  pi=w*sum;
```





# **Communication overhead** – an experiment (solution/pi/pivec.c)





