

# More on OpenMP Tasking

Christian Terboven, Dirk Schmidl

IT Center, RWTH Aachen University  
{terboven,schmidl}@itc.rwth-aachen.de

# Sudoku

- This parallel algorithm finds all valid solutions

	6					8	11			15	14			16	
15	11				16	14				12			6		
13		9	12					3	16	14		15	11	10	
2		16		11		15	10	1							
	15	11	10			16	2	13	8	9	12				
12	13			4	1	5	6	2	3				11	10	
5		6	1	12		9		15	11	10	7	16		3	
	2				10		11	6		5			13	9	
10	7	15	11	16				12	13					6	
9						1			2	16	10			11	
1		4	6	9	13			7		11		3	16		
16	14			7		10	15	4	6	1				13	8
11	10		15				16	9	12	13			1	5	4
		12		1	4	6		16				11	10		
		5		8	12	13		10			11	2			14
3	16			10			7			6				12	

**(1) Search an empty field**

**(2) Insert a number and check Sudoku**

**(3 a) If invalid:  
Delete number,  
Insert next number**

**(3 b) If valid:  
Go to next field**

**Wait for completion**

# Parallel Brute-force Sudoku



- This parallel algorithm finds all valid solutions

6								11		15	11			16
15	11													
13		9	12											
2		16		11										
	15	11	10											
12	13			4	1	5	6	2	3				11	10
5		6	1	12		9		15	11	10	7	16		3
	2				10		11	6		5		13		9
10	7	15	11	16				12	13					6
9						1		2	16	10				11
1		4	6	9										
16	14			7										3
11	10		15											4
		12		1	4	6		16				11	10	
		5		8	12	13		10			11	2		14
3	16			10										

first call contained in a  
`#pragma omp parallel`  
`#pragma omp single`  
 such that one tasks starts the  
 execution of the algorithm

`#pragma omp task`  
 needs to work on a new copy  
 of the Sudoku board

`#pragma omp taskwait`  
 wait for all child tasks

(1) Search an empty field

(2) Insert a number  
 and check Sudoku

(3 a) If invalid:  
 Delete number,  
 Insert next number

(3 b) If valid:  
 Go to next field

Wait for completion

## ■ OpenMP parallel region creates a team of threads

```
#pragma omp parallel
{
  #pragma omp single
    solve_parallel(0, 0, sudoku2, false);
} // end omp parallel
```

→ Single construct: One thread enters the execution of `solve_parallel`

→ the other threads wait at the end of the `single` ...

→ ... and are ready to pick up threads „from the work queue“

## ■ Syntactic sugar (either you like it or you don't)

```
#pragma omp parallel sections
{
  solve_parallel(0, 0, sudoku2, false);
} // end omp parallel
```

## ■ The actual implementation

```
for (int i = 1; i <= sudoku->getFieldSize(); i++) {
    if (!sudoku->check(x, y, i)) {
#pragma omp task firstprivate(i,x,y,sudoku)
    {
        // create from copy constructor
        CSudokuBoard new_sudoku(*sudoku)
        new_sudoku.set(y, x, i);
        if (solve_parallel(x+1, y, &new_sudoku)) {
            new_sudoku.printBoard();
        }
    } // end omp task
    }
}
```

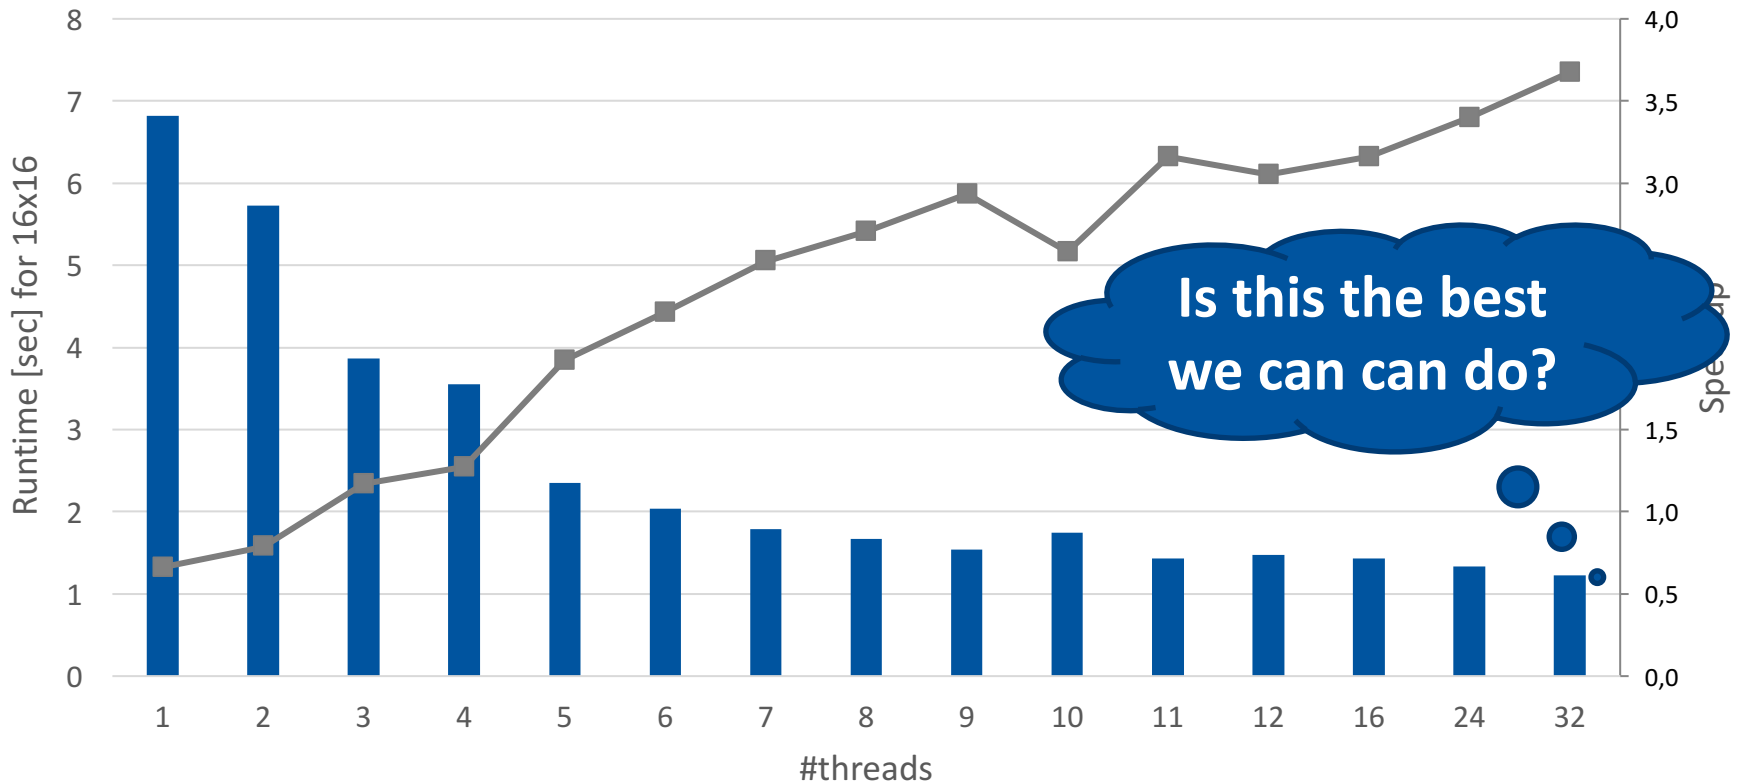
**#pragma omp task**  
need to work on a new copy of  
the Sudoku board

```
#pragma omp taskwait
```

**#pragma omp taskwait**  
wait for all child tasks

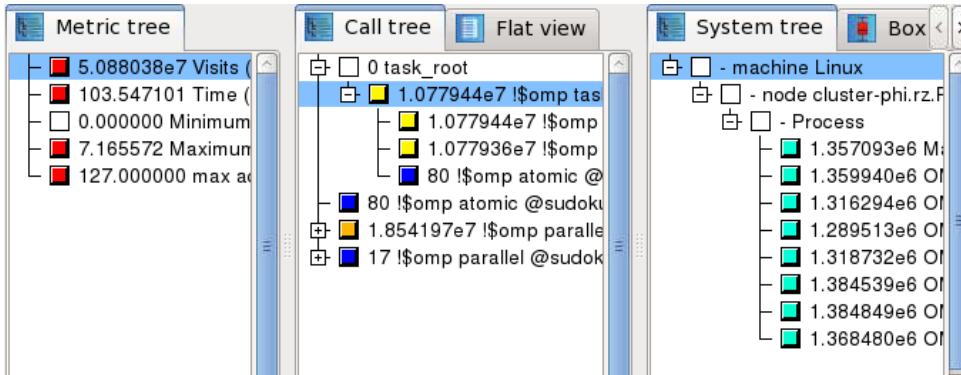
Sudoku on 2x Intel Xeon E5-2650 @2.0 GHz

■ Intel C++ 13.1, scatter binding    ■ speedup: Intel C++ 13.1, scatter binding

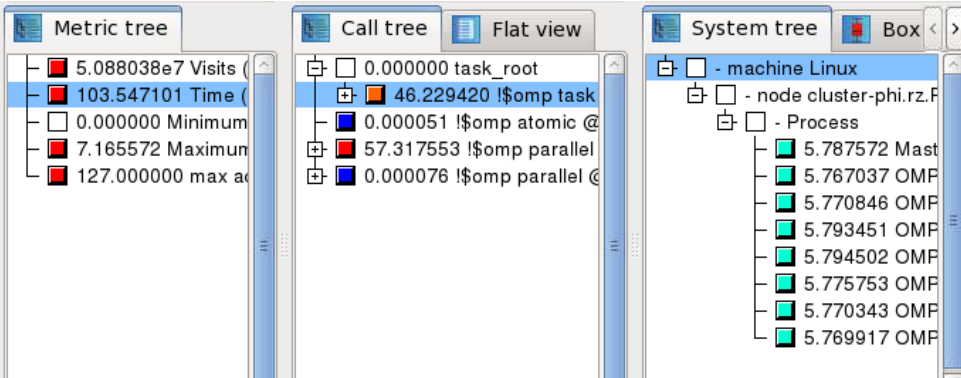


# Performance Analysis

Event-based profiling gives a good overview :






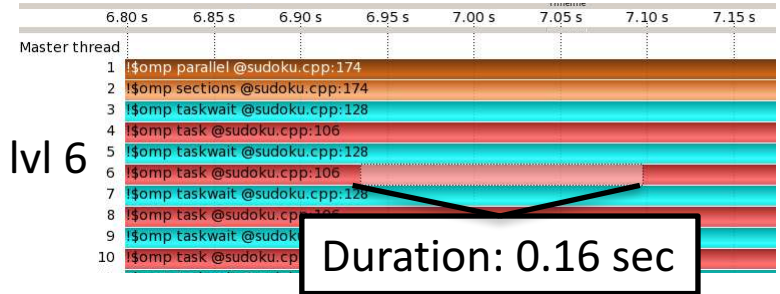
Every thread is executing ~1.3m tasks...



... in ~5.7 seconds.

=> average duration of a task is ~4.4 μs

Tracing gives more details:




Tasks get much smaller down the call-stack.





Sudoku on 2x Intel Xeon E5-2650 @2.0 GHz



# Scheduling

- **Default: Tasks are *typed* to the thread that first executes them → not necessarily the creator. Scheduling constraints:**
  - Only the thread a task is tied to can execute it
  - A task can only be suspended at task scheduling points
    - Task creation, task finish, `taskwait`, `barrier`, `taskyield`
  - If task is not suspended in a barrier, executing thread can only switch to a direct descendant of all tasks tied to the thread
- **Tasks created with the `untied` clause are never tied**
  - Resume at task scheduling points possibly by different thread
  - ~~No scheduling restrictions, e.g., can be suspended at any point~~
  - But: More freedom to the implementation, e.g., load balancing

- **Problem: Because untied tasks may migrate between threads at any point, thread-centric constructs can yield unexpected results**
  
- **Remember when using untied tasks:**
  - Avoid `threadprivate` variable
  - Avoid any use of thread-ids (i.e., `omp_get_thread_num()`)
  - Be careful with `critical region` and *locks*
  
- **Simple Solution:**
  - Create a tied task region with

```
#pragma omp task if(0)
```

- The `taskyield` directive specifies that the current task can be suspended in favor of execution of a different task.

→ Hint to the runtime for optimization and/or deadlock prevention

C/C++

```
#pragma omp taskyield
```

Fortran

```
!$omp taskyield
```

# taskyield Example (1/2)



```
#include <omp.h>

void something_useful();
void something_critical();

void foo(omp_lock_t * lock, int n)
{
    for(int i = 0; i < n; i++)
        #pragma omp task
        {
            something_useful();
            while( !omp_test_lock(lock) ) {
                #pragma omp taskyield
            }
            something_critical();
            omp_unset_lock(lock);
        }
}
```

## taskyield Example (2/2)



```
#include <omp.h>

void something_useful();
void something_critical();

void foo(omp_lock_t * lock, int n)
{
    for(int i = 0; i < n; i++)
        #pragma omp task
        {
            something_useful();
            while( !omp_test_lock(lock) ) {
                #pragma omp taskyield
            }
            something_critical();
            omp_unset_lock(lock);
        }
}
```

The waiting task may be suspended here and allow the executing thread to perform other work; may also avoid deadlock situations.





# Questions?