

# **Tasking in OpenMP 5.0**

Christian Terboven <terboven@itc.rwth-aachen.de>

June 22nd, 2021 NHR PerfLab Seminar Series



- De-facto standard Application Programming Interface (API) to write shared memory parallel applications in C,
  - C++, and Fortran

2

- Consists of Compiler Directives, Runtime routines and Environment variables
- Version 5.0 has been released at SC 2018
- Version 5.1 has been released at SC 2020





- De-facto standard Application Programming Interface (API) to write shared memory parallel applications in C,
  - C++, and Fortran

3

- Consists of Compiler Directives, Runtime routines and Environment variables
- Version 5.0 has been released at SC 2018
- Version 5.1 has been released at SC 2020





# **Motivation**



# **Sudoko for Lazy Computer Scientists**

Lets solve Sudoku puzzles with brute multi-core force

8 11 15 14 15 11 16 14 15 11 10 3 16 14 9 12 15 10 2 13 9 12 15 11 10 12 13 11 10 15 11 10 7 16 7 15 11 16 12 13 16 10 9 13 3 16 16 14 10 15 9 12 13 11 10 11 10 8 12 13 3 16 

(1) Search an empty field

- (2) Try all numbers:
  - (2 a) Check Sudoku
    - If invalid: skip
    - If valid: Go to next field

Wait for completion



• ]	۲h	is	pa	ara	all	el	alo	00	ritk	nm	<u>ו fi</u>	inc	ls	all	V	ali	id solutions	
		6						8	11			15	14			16	$\delta$	
1	5	11				16	14				12			6			(1) Search an empty fie first call contained in a	
1	3		9	12					З	16	14		15	11	10		#pragma omp paralle	ι
	2		16		11		15	10	1								#pragma omp single	
	,	15	11	10			16	2	13	8	9	12					1 (2) Try all numbers:	
1	2	13			4	1	5	6	2	3					11	10	(2 a) Check Sudoku	
	5		6	1	12		9		15	11	10	7	16			3		
		2				10		11	6		5			13		9	f It invalid: skip	
1	0	7	15	11	16				12	13						6	If valid: Go to ne: #pragma omp task	
	9						1			2		16	10			11	needs to work on a new	
	1		4	6	9	13			7		11		3	16			copy of the Sudoku board	
1	6	14			7		10	15	4	6	1				13	8	3	
1	1	10		15				16	9	12	13			1	5	4	Wait for completion	
			12		1	4	6		16				11	10			#pragma omp taskwai	t
			5		8	12	13		10			11	2			14	wait for all child tasks	
6	3	16			10			7			6				12			
5				5								~	-				T Center UNIVERS	<b>EN</b> ITY

#### **First Performance Evaluation**







7

# **Tasking Basics**



### What is a task in OpenMP?

- Tasks are work units whose execution
  - may be deferred or...
  - ... can be executed immediately
- Tasks are composed of
  - code to execute, a data environment (initialized at creation time), internal control variables (ICVs)
- Tasks are created...

9

- ... when reaching a parallel region  $\square$  implicit tasks are created (per thread)
- $\dots$  when encountering a task construct  $\square$  explicit task is created
- ... when encountering a taskloop construct  $\square$  explicit tasks per chunk are created
- $\dots$  when encountering a target construct  $\square$  target task is created



# **Tasking execution model**

- Supports unstructured parallelism
- unbounded loops

```
while ( <expr> ) {
    ...
}
```

```
void myfunc( <args> )
{
    ...; myfunc( <newargs> ); ...;
}
```

• Example (unstructured parallelism)





- Several scenarios are possible:
  - single creator, multiple creators, nested tasks (tasks & WS)
- All threads in the team are candidates to execute tasks



#### The task construct

11

<ul> <li>Deferring (or not) a unit of work (executable for any member of the team)</li> </ul>								
<pre>#pragma omp task [clause[[,] clause]] {structured-block}</pre>	<pre>!\$omp task [clause[[,] clause]]structured-block !\$omp end task</pre>							
Where palaus (as one of:	$\rightarrow$ if(scalar-expression)							

W	ne of:			$\rightarrow$ if(scalar-expression)	Cutoff Strategies	
	$\rightarrow$ firstprivate(list)			→ mergeable		
$\rightarrow$ shared(list)		Data Environment		$\rightarrow$ final(scalar-expression)		
$\rightarrow$ default(shared   none)				depend(dep-type: list)	Synchronization	
	$\rightarrow$ in_reduction(r-id: list)			$\rightarrow$ untied		
	$\rightarrow$ allocate([allocator:] list)	Miscollanoous		$\rightarrow$ priority(priority-value)	Task Scheduling	
	$\rightarrow$ detach(event-handler)	MISCENAREOUS		→ affinity(list)		



#### **The taskloop Construct**

• Task generating construct: decompose a loop into chunks, create a task for each loop chunk

#pragma omp taskloop [clause[[,] clause]...]
{structured-for-loops}

!\$omp taskloop [clause[[,] clause]...]
...structured-do-loops...
!\$omp end taskloop

W	herestotaatte tijs one of:	
	→ private(list)	
	→ firstprivate(list)	Data
	→ lastprivate(list)	Environment
	→ default(sh   <u>pr</u>   <u>fp</u>   none)	
	$\rightarrow$ reduction(r-id: list)	
	$\rightarrow$ in reduction(r-id: list)	
	→ grainsize(grain-size)	Chunks/Grain
12	-> numngtasks(num tasks) rboven   RWT	Chunks/Gram

<ul> <li>→ if(scalar-expression)</li> <li>→ final(scalar-expression)</li> <li>→ mergeable</li> </ul>	Cutoff Strategies
<ul> <li>→ untied</li> <li>→ priority(priority-value)</li> </ul>	Scheduler (R/H)
<ul> <li>→ collapse(n)</li> <li>→ nogroup</li> <li>→ allocate([allocator:] list)</li> </ul>	Miscellaneous



# **Task Synchronization**



# **Task synchronization: taskwait directive**

- The taskwait directive (shallow task synchronization)
  - It is a stand-alone directive

#pragma omp taskwait

- wait on the completion of child tasks of the current task; just direct children, not all descendant tasks; includes





# Task synchronization: taskgroup construct

- The taskgroup construct (deep task synchronization)
  - attached to a structured block; completion of all descendants of the current task; TSP at the end

```
#pragma omp taskgroup [clause[[,] clause]...]
{structured-block}
```





# **Task Scheduling**



#### **Task scheduling: taskyield directive**

- Task scheduling points (and the taskyield directive)
  - tasks can be suspended/resumed at TSPs 🖾 some additional constraints to avoid deadlock problems
  - implicit scheduling points (creation, synchronization, ...)
  - explicit scheduling point: the taskyield directive

```
#pragma omp taskyield
                                                    tied:
                                                                                             (default)
                                                                       foo()
  #pragma omp parallel
                                                                  single
  #pragma omp single
   {
      #pragma omp task untied
                                                    untied:
         foo();
                                                                        foo()
         #pragma omp taskyield
         bar()
17
```

# Task reductions (using taskgroup)

int res = 0; Reduction operation node t\* node = NULL; perform some forms of recurrence calculations -.... **#pragma omp parallel** associative and commutative operators The (taskgroup) scoping reduction clause **#pragma omp single** #pragma omp taskgroup task\_reduction(op: list) {structured-block} #pragma omp taskgroup \ task reduction(+: res) { // [1] while (node) { #pragma omp task in\_reduction(+: res) \ firstprivate(node) #pragma omp task in\_reduction(op: list) { // [2] {structured-block} res += node->value; THE (LASK) IN TEULUUT CLAUSE [PARTICIPATIN] node = node->next; } // [3]



# **Task reductions (+ modifiers)**

Reduction modifiers Former reductions clauses have been extended task modifier allows to express task reductions -Registering a new task reduction [1] -Implicit tasks participate in the reduction [2] -- Compute final result after [4] The (task) in\_reduction clause [participating] #pragma omp task in\_reduction(op: list) {structured-block} Task participates in a reduction operation [3] -

```
int res = 0;
node t* node = NULL;
....
#pragma omp parallel reduction(task,+: res)
{ || [1][2]
 #pragma omp single
   #pragma omp taskgroup
     while (node) {
       #pragma omp task in_reduction(+: res) \
                 firstprivate(node)
       { || [3]
         res += node->value;
       node = node->next;
} // [4]
```



# **Task Dependencies**



# OpenMP 4.0

• The depend clause was added to the task construct

# **OpenMP 4.5**

- The depend clause was added to the target constructs
- Support for doacross loops

# **OpenMP 5.0**

- lvalue expressions in the depend clause
- New dependency type: mutexinoutset
- Iterators were added to the depend clause
- The depend clause was added to the taskwait
- Dependable objects

# OpenMP 5.1

• New dependency type: inoutset



#### What's in the spec: sema depend clause (1)

- A task cannot be executed until all its predecessor tasks are completed
- If a task defines an in dependence over a variable
  - the task will depend on all previously generated sibling tasks that reference at least one of the list items in an out or inout dependence
- If a task defines an out/inout dependence over a variable
  - the task will depend on all previously generated sibling tasks that reference at least one of the list items in an in, out or inout dependence



# What's in the spec: sema depend clause (1)

• A task cannot be executed until all its predecessor tasks are completed

```
 If a task defin int x = 0;

 - the task will dep #pragma omp parallel #pragma omp single
    out or inout d
                        #pragma omp task depend(inout: x) //T1
                        \{\ldots\}

    If a task defin

   the task will dep
                      #pragma omp task depend(in: x)
                                                             //T2
  -
                      \{\ldots\}
    in, out or inc
                        #pragma omp task depend(in: x)
                                                             //T3
                        \{\ldots\}
                        #pragma omp task depend(inout: x) //T4
                        \{\ldots\}
```



of the list items in an

#### of the list items in an



## What's in the spec: sema depend clause (2)

• Set types: inoutset & mutexinoutset

24

```
int x = 0, y = 0, res = 0;
#pragma omp parallel
#pragma omp single
 #pragma omp task depend(out: res) //TO
  res = 0;
 #pragma omp task depend(out: x) //T1
  long_computation(x);
 #pragma omp task depend(out: y) //T2
  short_computation(y);
 #pragma omp task depend(in: x) depend(modexinoes)et/TBes) //T3
  res += x;
 #pragma omp task depend(in: y) depend(motostinoes)et/T4
  res += y;
 #pragma omp task depend(in: res) //T5
  std::cout << res << std::endl;</pre>
```



1. *inoutset property*: tasks with a mutexinoutset dependence create a cloud of tasks (an inout set) that synchronizes with previous & posterior tasks that dependent on the same list item

2. *mutex property*: Tasks inside the inout set can be executed in any order but with mutual exclusion



# Advanced features: deps on taskwait

- Adding dependences to the taskwait construct
  - Using a taskwait construct to explicitly wait for some predecessor tasks
    - Syntactic sugar!

```
int x = 0, y = 0;
#pragma omp parallel
#pragma omp single
{
    #pragma omp task depend(inout: x) //T1
    x++;
    #pragma omp task depend(in: y) //T2
    std::cout << y << std::endl;
    #pragma omp taskwait depend(in: x)
    std::cout << x << std::endl;
}
```



## **Advanced features: dependable objects (1)**



26

# **Clauses to optimize Task Scheduling**



## Task scheduling: programmer's hints

{structured-block}

Programmers may specify a priority value when creating a task
 #pragma omp task priority(pvalue)

- pvalue: the higher  $\square$  the best (will be scheduled earlier)

```
#pragma omp parallel
#pragma omp single
{
   for ( i = 0; i < SIZE; i++) {
     #pragma omp task priority(1)
        { code_A; }
   }
   #pragma omp task priority(100)
   { code_B; }
   ...
}</pre>
```



priority tasks



# affinity clause

- New clause: #pragma omp task affinity (list)
  - Hint to the runtime to execute task closely to physical data location
  - Clear separation between dependencies and affinity
- Expectations:
  - Improve data locality / reduce remote memory accesses
  - Decrease runtime variability
- Still expect task stealing
  - In particular, if a thread is under-utilized



#### **Selected LLVM implementation details**





# **Do you still remember the Motivation?**







## **Performance Analysis**

# Event-based profiling provides a good overview :





# Tracing provides more details:



## **Performance Evaluation (with cutoff)**





- **Cancellation:** Cancellation (since OpenMP 4.0) provides a best-effort approach to terminate OpenMP regions
  - Best-effort: not guaranteed to trigger termination immediately
  - Triggered "as soon as" possible
- Asynchronous API Interaction: This provides a mechanism to marry asynchronous APIs with the parallel task model of OpenMP
  - How to synchronize completions events with task execution?