

How To Get OpenMP Tasks To Do Your Work

Leverage Tasking to Make Your Life Easier

Ruud van der Pas

Webinar, September 22, 2021

OpenMP[®]

My background is in mathematics and physics

*Previously, I worked at the University of Utrecht,
Convex Computer, SGI, and Sun Microsystems*




Currently I work in the Oracle Linux Engineering organization

I have been involved with OpenMP since the introduction

I am passionate about performance and OpenMP in particular

Your OneStop Place for OpenMP

<https://www.openmp.org>




A printed copy of the 5.1 OpenMP specs is also available on Amazon (700+ pages, 1.7 kg)


Food for the Eyes and Brains




*OpenMP 2.5 and intro
Parallel Computing*



*Covers the OpenMP Basics
to get started*



*Focus on the Advanced
Features*



*What goes on under the
hood*



Why Tasks are Needed

BT - The Time Before Tasking

The OpenMP Compiler Had To Know Everything

And in advance, (right) before execution

For example, the loop length, number of parallel sections, etc

*Gets hard with more dynamic problems like processing linked lists,
divide and conquer, recursion, etc*

A solution was ugly. At best

The Solution

Tasking came to the rescue!

In this talk we will show it all works

But no formal terminology, definitions, etc.

These are covered in the specs

OpenMP Application Program Interface

Version 3.0 May 2008

13
14
15
16
17
--

About 4+ pages on the tasking feature
(aside from definitions, scoping, etc)

2.6.3	parallel workshare Construct	58
2.7	task Construct	59
2.7.1	Task Scheduling	62
2.8	Master and Synchronization Constructs	63
2.8.1	master Construct	63
2.8.2	wait Construct	65

Talking about the Specifications - 5.1

OpenMP

OpenMP
Application Programming
Interface

Version 5.1 November 2020

About 22+ pages on the tasking feature
(aside from definitions, scoping, etc)

2.12	Tasking Constructs	161
2.12.1	task Construct	161
2.12.2	taskloop Construct	166
2.12.3	taskloop simd Construct	171
2.12.4	taskyield Construct	173
2.12.5	Initial Task	174
2.12.6	Task Scheduling	175
2.13	Memory Management Directives	177
2.19.5	taskwait Construct	261
2.19.6	taskgroup Construct	264



The Tasking Concept

What Is a Task?

A task is a chunk of independent work

You guarantee that different tasks can be executed simultaneously

```
#pragma omp task  
{"this is my task"}
```

When are Tasks Executed?

The OpenMP run time system decides on the scheduling of tasks

At certain points (implicit and explicit) tasks are guaranteed to be completed

Division of Labour

*This is what **you** need to do:*

- *Identify independent portions of work, the tasks*
- *Use **the #pragma omp task** construct to define the tasks*

The OpenMP runtime system handles everything else:

- *When a thread encounters a task construct, a new task is created*
- *It is up to the runtime system when this task is executed and by whom*
- *Execution of a task is either immediate, or may be delayed/deferred*
- ***Task synchronization** may be used to enforce task completion*

Task Creation and Execution in OpenMP

```
#pragma omp task
```

Defines a task

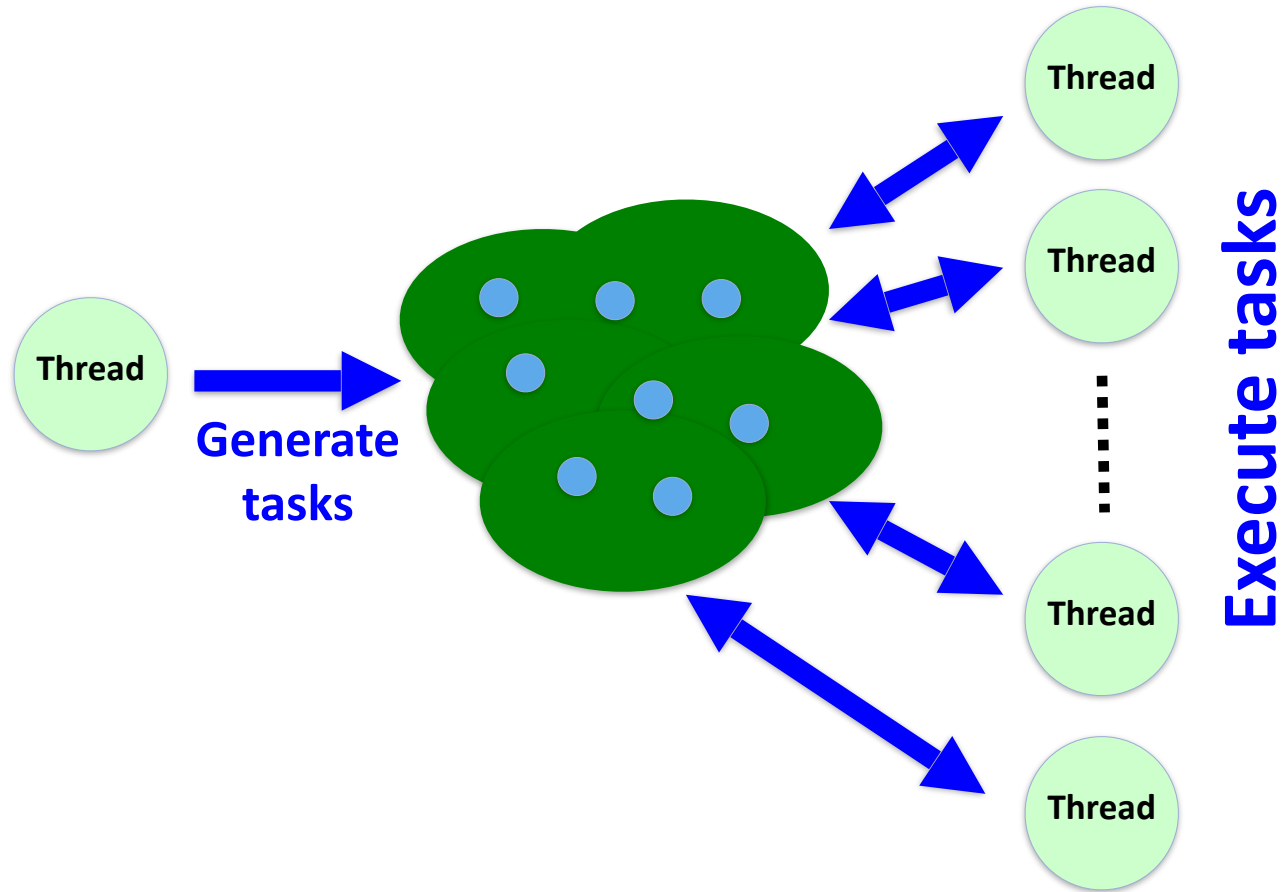
```
barrier (implied or explicit)
```

```
#pragma omp taskwait
```

Task Synchronization Points

```
#pragma omp taskgroup
```

A Common Execution Model with Tasks





A Day at the Races

Today's Task

Write a program that prints either “A race car” or “A car race” and maximize the parallelism

The Starting Point

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char *argv[]) {

    printf("A ");
    printf("race ");
    printf("car ");

    printf("\n");
    return(0);
}
```

```
$ gcc hello.c
$ ./a.out
A race car
$
```

What will this program print ?

The Parallel Version

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char *argv[]) {

    #pragma omp parallel
    {
        printf("A ");
        printf("race ");
        printf("car ");

    } // End of parallel region

    printf("\n");
    return(0);
}
```

***What will this program print
using 2 threads ?***

Let's Run It

```
$ gcc -fopenmp hello.c
$ export OMP_NUM_THREADS=2
$ ./a.out
A race car A race car
```

Note that this program could (for example) also print

“A A race race car car” or

“A race A car race car”, or

“A race A race car car”, or

.....

The Next Version

```
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char *argv[]) {
    #pragma omp parallel
    {
        #pragma omp single
        {
            printf("A ");
            printf("race ");
            printf("car ");
        }
    } // End of parallel region

    printf("\n");
    return(0);
}
```

*What will this program print
using 2 threads ?*

Let's Try It

```
$ gcc -fopenmp hello.c  
$ export OMP_NUM_THREADS=2  
$ ./a.out  
A race car
```

*But of course only 1 thread executes
now ...*

Time to Add the Tasks!

```
int main(int argc, char *argv[]) {  
    #pragma omp parallel  
    {  
        #pragma omp single  
        {  
            printf("A ");  
  
            #pragma omp task  
            {printf("race ");}  
  
            #pragma omp task  
            {printf("car ");}  
        }  
    } // End of parallel region  
  
    printf("\n");  
    return(0);  
}
```

*What will this program print
using 2 threads ?*

Mission Accomplished!

```
$ gcc -fopenmp hello.c
$ export OMP_NUM_THREADS=2
$ ./a.out
A race car
$ ./a.out
A race car
$ ./a.out
A car race
$
```

*Tasks may be executed in
arbitrary order*

That went well and quickly, so here is a final task to do

***Have the sentence end with “is fun to watch”
(hint: use a print statement)***

Add the Print Statement after the Tasks

```
int main(int argc, char *argv[]) {  
    #pragma omp parallel  
    {  
        #pragma omp single  
        {  
            printf("A ");  
            #pragma omp task  
            {printf("race ");}  
            #pragma omp task  
            {printf("car ");}  
            printf("is fun to watch ");  
        }  
    } // End of parallel region  
  
    printf("\n");  
    return(0);  
}
```



Oops!

```
$ gcc -fopenmp hello.c
$ export OMP_NUM_THREADS=2
$ ./a.out

A is fun to watch race car
$ ./a.out

A is fun to watch race car
$ ./a.out

A is fun to watch car race
$
```

*Tasks are executed at a Task Synchronization Point
(and not necessarily where you see them ...)*

Where is the Task Synchronization Point?

```
int main(int argc, char *argv[]) {  
    #pragma omp parallel  
    {  
        #pragma omp single  
        {  
            printf("A ");  
            #pragma omp task  
            {printf("race ");}  
            #pragma omp task  
            {printf("car ");}  
            printf("is fun to watch ");  
        } // Implied barrier  
    } // End of parallel region  
  
    printf("\n"); return(0);  
}
```

**Task synchronization
point**



Use the Taskwait Feature

```
int main(int argc, char *argv[]) {  
    #pragma omp parallel  
    {  
        #pragma omp single  
        {  
            printf("A ");  
            #pragma omp task  
            {printf("car ");}  
            #pragma omp task  
            {printf("race ");}  
  
            #pragma omp taskwait  
            printf("is fun to watch ");  
        }  
    } // End of parallel region  
    printf("\n"); return(0);  
}
```

*What will this program
print using 2 threads ?*

**Task synchronization
point**

*"The taskwait construct specifies a wait on
the completion of child tasks of the current task"*

Mission Accomplished!

```
$ gcc -fopenmp hello.c
$ export OMP_NUM_THREADS=2
$ ./a.out
$
A car race is fun to watch
$ ./a.out
A car race is fun to watch
$ ./a.out
A race car is fun to watch
$
```

The tasks are executed first now



Algorithm Examples

When to Use Tasks?

If you need them

Tasks do not replace everything

For example, if the loop construct works well, don't touch it

Consider tasks when things are much more dynamic

If you can't predict how much work will be done, for example

Or if the workload is unbalanced, or if ...

Three Scenarios

Scenario	An Example
<i>Irregular problems</i>	<i>Scan a linked list</i>
<i>Dependencies</i>	<i>Overlap I/O and computation</i>
<i>Recursion</i>	<i>A sorting algorithm</i>



Scan a Linked List

Scan Through a Linked List

```
1 my_pointer = head_of_list;
2 while (my_pointer != NULL)
3 {
4     (void) do_independent_work(my_pointer->value);
5     my_pointer = my_pointer->next;
6 } // End of while loop
```

This is cumbersome to do without tasking:

- *First count the number of passes through the while-loop*
- *Convert the while-loop to a for-loop*

The Problem and the Idea

The Problem

- *Scanning through the linked list is a serial process*
- *Check each record until the NULL pointer is encountered*


The Idea

- *Make the execution of function `do_independent_work` a task*
- *Use the `single` construct to have a single thread generate the tasks*
- *The other threads start executing the tasks as they become available*

The Relevant OpenMP Code Fragment

```
1 my_pointer = head_of_list;
2 #pragma omp parallel firstprivate(my_pointer)
3 {
4     #pragma omp single nowait ←
5     {
6         while (my_pointer != NULL)
7         {
8             #pragma omp task firstprivate(my_pointer)
9             {
10                (void) do_independent_work(my_pointer->value);
11            } // End of task
12            my_pointer = my_pointer->next;
13        } // End of while loop
14    } // End of single region ←
15 } // End of parallel region
```

Removes the implied
barrier here



Overlap I/O and Computation

A Common Problem

The processing part waits for data to arrive from disk

In case the data is read in chunks, tasking may help

*The main idea is to **overlap I/O and computations***

This keeps the processor (more) busy

*An elegant solution makes use of **task dependencies***

The Main Idea - Set Up a Pipeline



- *In reality, the two phases may not take the same time*
- *Regardless of that, this approach reduces the total time*
- *In the implementation there will be two tasks*
- *Through a dependence, we ensure processing does not start too early*
- *Although not shown here, a third stage may handle the output processing*

Task Dependencies

It is possible to set up dependencies between tasks

*This is implemented through the **depend** clause*

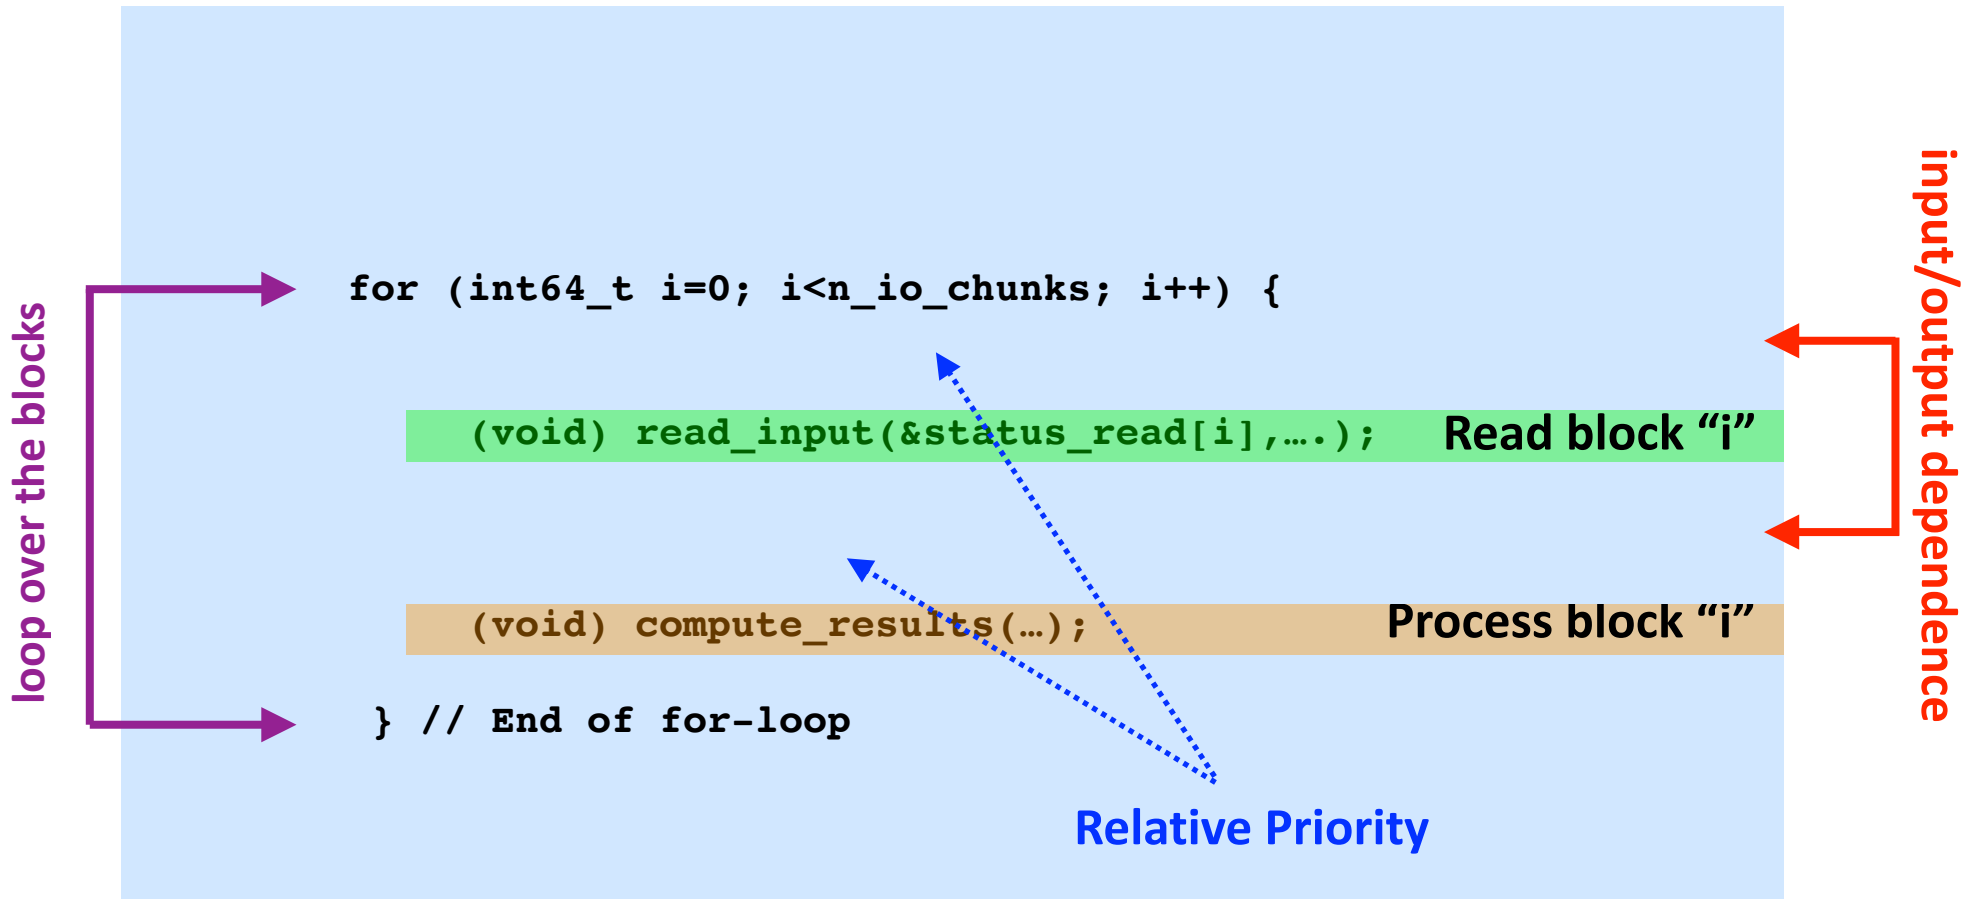
Various dependence types are supported

*In this case, we need an **out** and **in** dependence pair*

This is used to express that the computation depends on the I/O

Let's look at the code to see how this works

The OpenMP Code Structure





A Sorting Algorithm

The Quicksort Algorithm

A commonly used sorting algorithm

A divide-and-conquer strategy is used

The sequential algorithm:



1. Select a pivot
2. Re-order the elements in the array to be sorted such that:
 - All elements to the left of the pivot are smaller
3. Repeat for the parts to the left and right of the pivot

An Example of the Quicksort Algorithm in Action

0

1

2

3

4

8

5

7

3

9

select pivot

The Sequential Code

in the main program: `(void) quicksort(a, 0, n-1);`

```
1 int64_t quicksort(int64_t *a, int64_t lo, int64_t hi)
2 {
3     if ( lo < hi ) {
4         p = partition(a, lo, hi);
5
6         (void) quicksort(a, lo, p - 1); // Left
7
8         (void) quicksort(a, p + 1, hi); // Right
9     }
10    return(p);
11 }
```

The OpenMP Implementation

```
#pragma omp parallel ...
{
  #pragma omp single nowait
  { (void) ompQuicksort(a, 0, n-1);}
} // End of parallel region
```

In the main program

```
1 void ompQuicksort(int64_t *a, int64_t lo, int64_t hi)
2 {
3   if ( lo < hi )
4     {
5       int64_t p = partitionArray(a, lo, hi);
6       #pragma omp task ...
7         {(void) ompQuicksort(a, lo, p - 1);}   Left task
8       #pragma omp task ...
9         {(void) ompQuicksort(a, p + 1, hi);}   Right task
10      #pragma omp taskwait
11    }
12 }
```

Tuning The Parallel Sorting Algorithm

This algorithm requires some tuning

Tasks are not for free and carry some overhead

If an array section gets too small, it may be more efficient to switch to the sequential version

This is not so elegant, but it is efficient :-)

Tuning Algorithms that Use Tasks

The tasking system is very dynamic

In general, with tasking, there are some things to keep in mind

How many tasks should be created?

Should small tasks be merged at some point?

Should tasks be forced to execute as soon as possible?

How about relative priorities?

Is it necessary to keep a task tied to the same thread?

Tuning Controls for Tasking

Clause	Functionality
<i>final(expr)</i>	<i>If expr is true, stop generating more tasks</i>
<i>mergeable</i>	<i>Merge the data environment</i>
<i>untied</i>	<i>The task is not tied to a specific thread</i>
<i>if (expr)</i>	<i>If expr is false, generate an undeferred* task</i>
<i>priority(hint)</i>	<i>Give a priority hint to the runtime system</i>

**) In simple terms, an undeferred task is executed immediately*



The Parallel Quicksort in Practice

Preliminary Findings

We have run some tests with the parallel quicksort algorithm

Please keep in mind that this is very early work and incomplete

The quicksort algorithm is a good target for additional tuning

- *if-clause*
- *final-clause*
- *switch point for the sequential version*
- *how to handle idle threads*
- *NUMA controls*
- *...*

Testing Conditions

We used an 8 core VM with 16 hardware threads

Intel Xeon Platinum 8167M CPU @ 2.00GHz (“Skylake”)

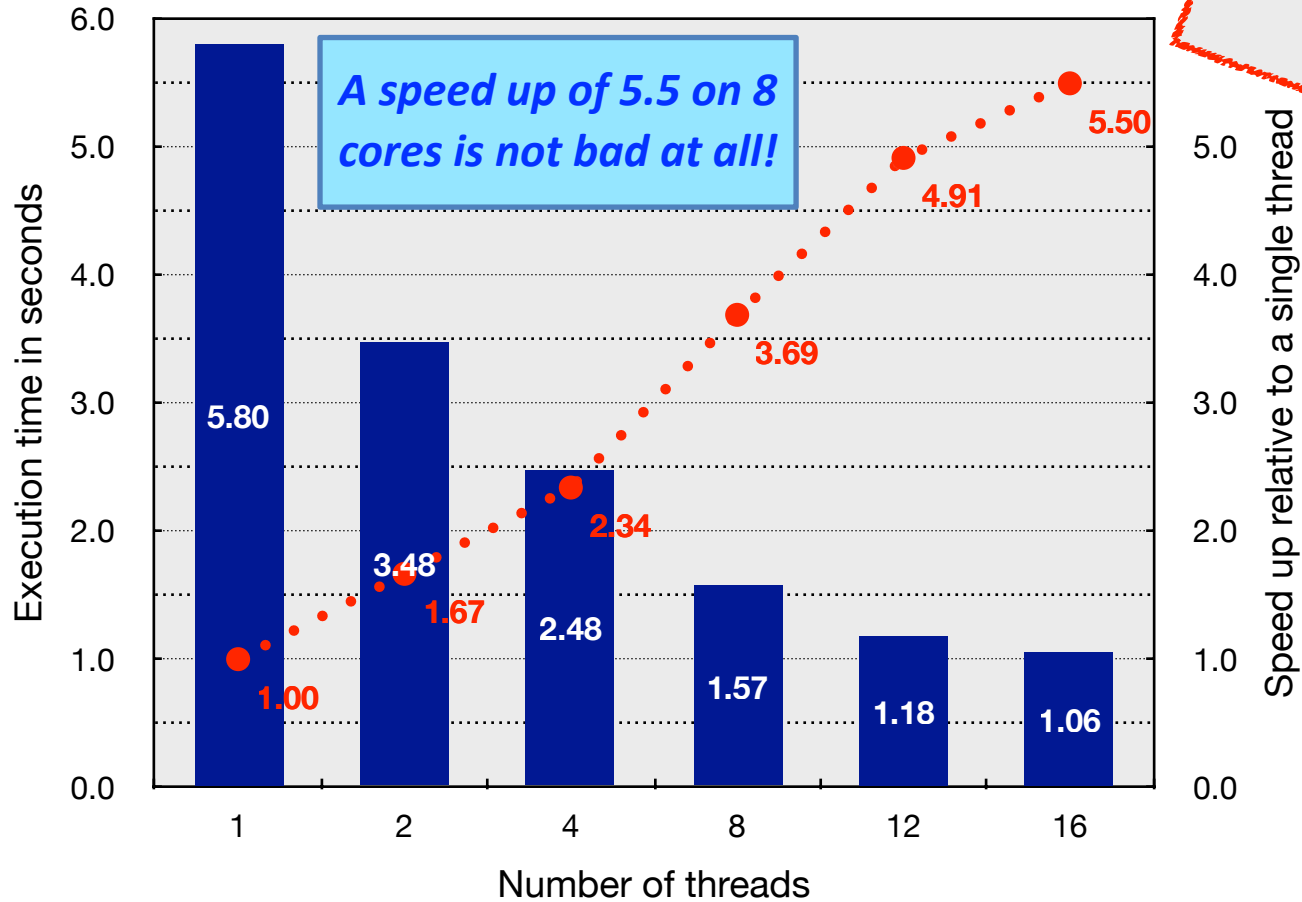
We sorted 40M 64 bit integers (320 MB of data)

The algorithmic parameters were:

- final-clause: cutoff at 4000 elements (0.01% of the length)*
- switch point for the sequential version: 400 (0.001% of the length)*


The Performance Using 8 Cores

Preliminary Results!



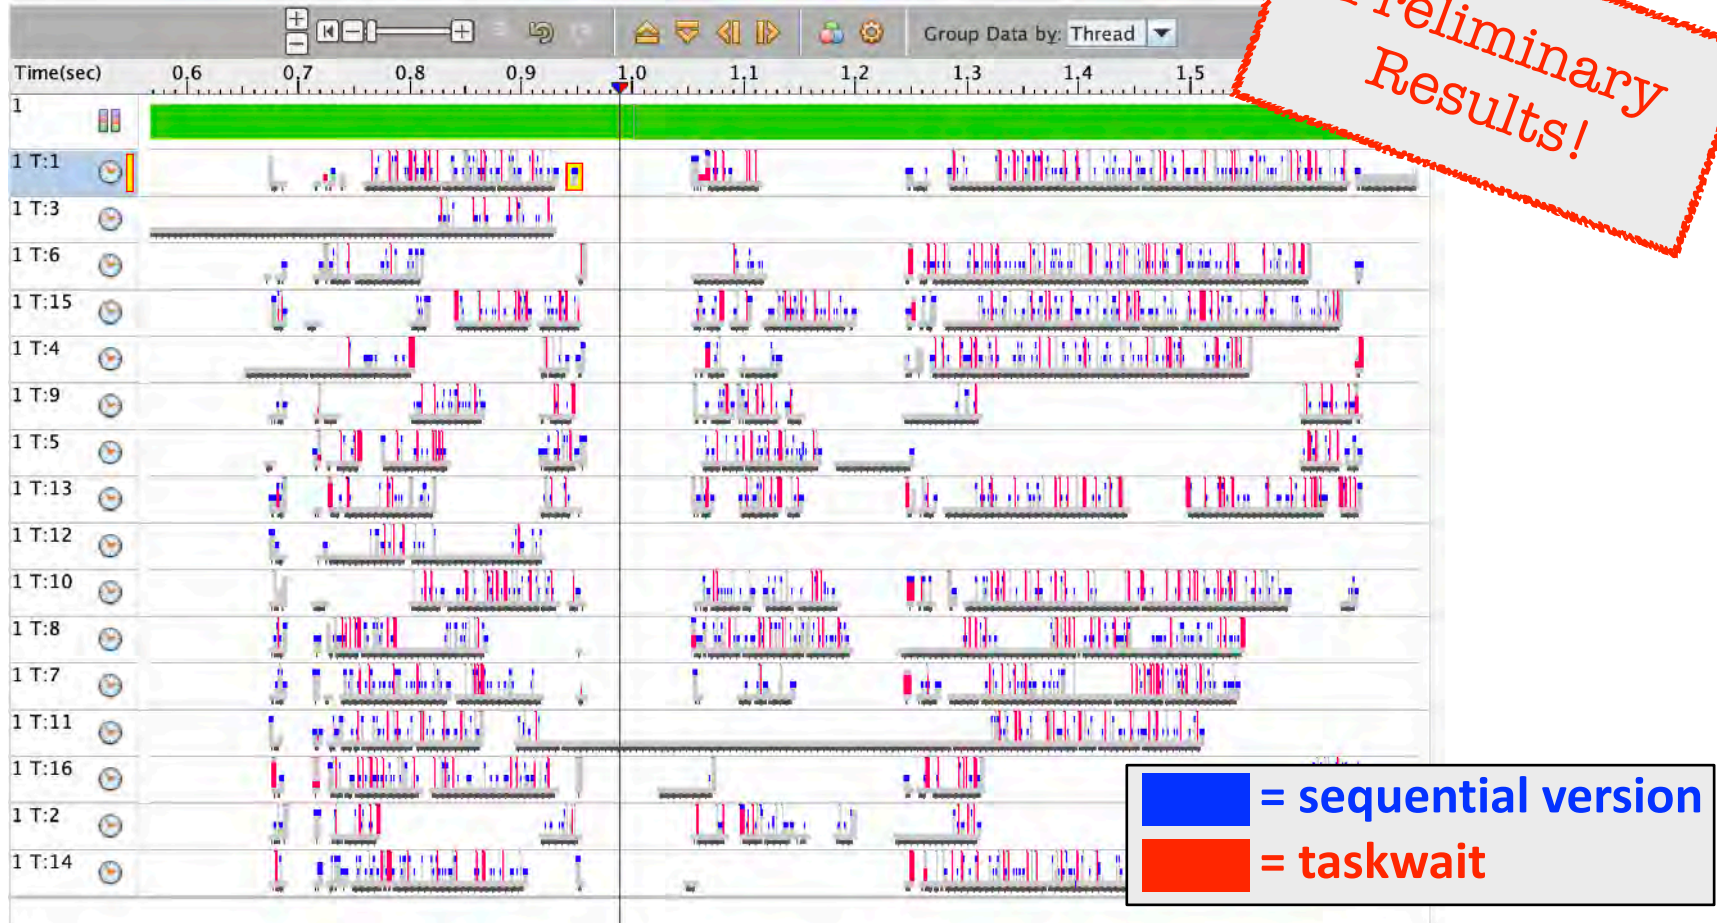
The Dynamic Behaviour Using 16 Threads

Preliminary Results!



Use Color Coding to Distinguish the Phases

Preliminary Results!





Many Stones Unturned

What Was not Talked About

The focus of this talk has been on the functionality

Many important topics have not even been touched upon

For example, the `taskloop`, `task_reduction`, and `taskyield` features

Data scoping has been skipped as well

These are useful/relevant topics though and there is more, so please check for more information about tasking

***At SC21 there is a full day tutorial on tasking on November 14
“Mastering Tasking with OpenMP”***

Additional OpenMP tutorials at SC21:

- ***The OpenMP Common Core - A Hands-on Introduction***
- ***Advanced OpenMP - Host Performance and 5.1 Features***
- ***Programming your GPU with OpenMP - A Hands-on Introduction***
- ***See also:***

<https://sc21.supercomputing.org/program/tutorials/#schedule>

The OpenMP Example Set

*You can download a set of example codes
They have been updated for 5.1*



The screenshot shows the OpenMP website's navigation bar with links for Home, Specifications, Community, Resources, News & Events, and About. Below the navigation bar, the main heading reads "OpenMP Examples - Updated with 5.1 Features". A sub-heading indicates "OPENMP RESOURCES UPDATE". The main content area features a large heading: "The OpenMP Examples document has been updated with new features found in the OpenMP 5.1 Specification". Below this, a paragraph explains that the OpenMP ARB works to maintain and promote the OpenMP API, and this effort includes publishing examples of OpenMP code. It states that the OpenMP API Examples document is a collection of programming examples intended to supplement the OpenMP API specification. The document is now available for the 5.1 version, which contains new features added in November of 2020. It includes new chapters on Loop Transformations, Directive Syntax, DMPT (roais) Interface, and Discretized Features. A list of new features, additions, and modifications is referenced in the Document Revision History at the end. The source codes for the OpenMP Examples can be downloaded from the "sources" directory for each chapter. A list of links is provided:

- <https://github.com/OpenMP/Examples>
- The codes for the OpenMP 5.1 Examples document have the tag v5.1

Thanks to our Examples Subcommittee Co-chairs:


- Henry Jiri (NASA Ames Research Center)
- Kent Milfeld (TACC, Texas Advanced Research Center)

At the bottom of the page, there is a button labeled "OPENMP EXAMPLES - VERSION 5.1" with a download icon.

Additional Reading Material

The book “Using OpenMP - The Next Step” covers many of the 4.5 features, including a full chapter on tasking

Although this covers 4.5, it provides a solid introduction into the more advanced topics and prepares you for 5.1



Thank You And ... Stay Tuned!

***Bad OpenMP
Does Not Scale***

Ruud van der Pas

Webinar, September 22, 2021