



Sandia  
National  
Laboratories

Exceptional service in the national interest

# Characterizing the Performance of Task Reductions in OpenMP 5.X Implementations

Jan Ciesko, Stephen Olivier

Thursday, 29<sup>th</sup> September, 2022

UT, Chattanooga, Tennessee,



## Agenda

- Reductions in OpenMP
- Implementations in LLVM/Clang and GCC
- Benchmarks
- Evaluation
- Conclusion



# Reductions in OpenMP

- OpenMP supports reductions since V1.0
- Tasking since V3.0 in 2008
- Demand for task-parallel reductions
- I started working on this in 2013
- Presented to OMP LC in 2015.
- Proposal made it into the spec in 2018 (OpenMP 5.0)
- Task reductions are conceptually concurrent and are orthogonal to the depend clause
- Compiler support evaluation is important for performance portability
  - How well did we do in the spec?
  - Did implementers implement the spec?

## OpenMP C and C++ Application Program Interface

Version 1.0 – October 1998

### 2.3 parallel Construct

The following directive defines a *parallel region*, which is a region of the program that is to be executed by multiple threads in parallel. This is the fundamental construct that starts parallel execution.

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

The *clause* is one of the following:

```
if(scalar-expression)
private(list)
firstprivate(list)
```

8

004-2229-001

OpenMP C and C++ Application Program Interface

Directives [2]

```
default(shared | none)
shared(list)
copyin(list)
reduction(operator: list)
```



# Reductions in OpenMP

- Task reductions allow reductions computed by arbitrary task graphs
- Challenge: definition of participating tasks and of the scope of the reduction computations
- Clauses for task reductions:

Clause	Semantic	Description
<b>reduction(task, op: var)</b>	Scoping, Participation	Scopes a task reduction for a <b>parallel or work-sharing region</b> *
<b>reduction(op: var)</b>	Scoping, Participation	Scopes a task reduction for a <b>taskloop</b> region and makes created tasks participants.
<b>task_reduction(op: var)</b>	Scoping	Scopes a task reduction for a <b>taskgroup</b> region.
<b>in_reduction(op: var)</b>	Participation	Denotes participation of a <b>task</b> , <b>target task</b> , or <b>taskloop</b> in a <b>task reduction</b> .

- 'task' reduction modifier on construct only for 'parallel', 'for', 'sections' or 'scope'

## Rule of Thumb:

1. Participating tasks must have a enclosing scope that defines a task reduction
2. The reduction computation completes by the time the scope ends



# Implementations in GCC and LLVM/Clang

- Privatization and data reuse

```
1 void func(int &sum) {
2 #pragma omp taskgroup task_reduction(+: sum)
3 #pragma omp task in_reduction(+: sum)
4         sum++;
5 }
```



```
1 void func (int & sum) {
2     struct .omp_data_s.0 .omp_data_o.1;
3     ...
4     .omp_data_o.1.sum = sum_2(D);
5     __builtin_GOMP_task (_Z4funcRi._omp_fn.0, &.omp_data_o.1, 0B, 8, 8, 1,
6     0, 0B, 0, 0B);
7     return;
8 }
9 void _Z4funcRi._omp_fn.0 (struct .omp_data_s.0 & restrict .omp_data_i) {
10 ...
11     void * D.2516[1]; //new double pointer
12     _3 = .omp_data_i_2(D)->sum; //reference to original reduction storage location
13     D.2516[0] = _3;
14     __builtin_GOMP_task_reduction_remap (1, 0, &D.2516); ● //redirect
15     sum_6 = D.2516[0]; //dereference
16     _10 = *sum_6;
17     _11 = _10 + 1; //use
18     *sum_6 = _11;
19     return;
20 }
```

```
gcc -fopenmp -c task.c
-fdump-tree-optimized -
-o task.o.gcc
```

```
clang -Xclang -S -
emit-llvm -
Xpreprocessor -
fopenmp -c task.c -o
task.o.s.llvm (not
shown)
```

\*GCC Version 13,  
20220518



# Benchmarks

## Applications\*

- Fibonacci
- Powerset
- Powerset-UDR
- Dot-product

6 implementations each:

- Parallel task-reduction
- Taskloop reduction
- Taskgroup reduction
- Manual per-thread data privatization
- Atomics
- Stack

☰ README.md

## OMP Task Bench (OMP-TB)

OMP-TB is a collection of benchmarks to measure tasking performance and tasking-related features in OpenMP. Currently it includes benchmark as listed below. Benchmarks in the `reductions` sub-directory target task-parallel reduction support. In general, such benchmarks are useful to evaluate compiler language support as well as its efficient implementation.

### OMP-TB Benchmarks

- reductions/dot (Dot Product)
- reductions/fib (Fibonacci)
- reductions/powerset (Powerset Permutations)
- reductions/powerset-final (Powerset Permutations using the final OpenMP clause)
- reductions/powerset-UDR (Powerset Permutations using user-defined reductions)
- reductions/others/array\_sum (Array Sum)
- reductions/others/knapsack (Knapsack)
- reductions/others/knightstour (Knights Tour)
- reductions/others/max\_height\_tree (Max Height)
- reductions/others/nbinarywords (n-Permutations)
- reductions/others/nqueens (N-Queens)
- reductions/others/TSP (Travelling Salesman Problem)

[1] [https://github.com/sandialabs/openmp\\_task\\_bench](https://github.com/sandialabs/openmp_task_bench)

\* OMP-TB [1] contains further examples



# Implementation: Fibonacci

- Parallel task reduction using manual cut-off

In results: "parallel-red"

```
1 void fib(int n, int &sum) {  
2     if (n < 2)  
3         sum += n;  
4     else {  
5         if (n < cut_off) {  
6             fib(n - 1, sum);  
7             fib(n - 2, sum);  
8         } else {  
9             #pragma omp task firstprivate(n) in_reduction(+ : sum)  
10            fib(n - 1, sum);  
11  
12             #pragma omp task firstprivate(n) in_reduction(+ : sum)  
13             fib(n - 2, sum);  
14         }  
15     }  
16 }
```

```
17 ...  
18 #pragma omp parallel reduction(task, + : sum) \  
19 num_threads(conf.num_threads)  
20 #pragma omp single  
21 #pragma omp task firstprivate(n) in_reduction(+ : sum)  
22     fib(n, sum);
```

task

Note: we avoided  
the use of final



# Implementation: Fibonacci

- Taskgroup reductions using manual cut-off

```
1 void fib(int n, int &sum) {  
2     if (n < 2)  
3         sum += n;  
4     else {  
5         if (n < cut_off) {  
6             fib(n - 1, sum);  
7             fib(n - 2, sum);  
8         } else {  
9             #pragma omp task firstprivate(n) in_reduction(+ : sum)  
10            fib(n - 1, sum);  
11  
12             #pragma omp task firstprivate(n) in_reduction(+ : sum)  
13            fib(n - 2, sum);  
14         }  
15     }  
16 }
```

```
17 ...  
18 #pragma omp parallel shared(n, sum) num_threads(conf.num_threads)  
19 #pragma omp single  
20 #pragma omp taskgroup task_reduction(+ : sum) ←  
21 #pragma omp task firstprivate(n) in_reduction(+ : sum)  
22     fib(n, sum);
```

main

task

In results: “taskgroup-red”



# Implementation: Fibonacci

- Using the stack and manual cut-off

In results: "stack"

```
1 int fib(int n) {  
2     int x, y; ←  
3     if (n < 2)  
4         return n;  
5     else {  
6         if (n < cut_off) {  
7             x = fib(n - 1);  
8             y = fib(n - 2);  
9         } else {  
10             #pragma omp task shared(x) firstprivate(n)  
11             x = fib(n - 1);  
12  
13             #pragma omp task shared(y) firstprivate(n)  
14             y = fib(n - 2);  
15  
16             #pragma omp taskwait ←  
17         }  
18     }  
19     return x + y;  
20 }  
21 }
```

22 ...  
23 #pragma omp parallel shared(sum) num\_threads(conf.num\_threads)  
24 #pragma omp single  
25 #pragma omp task shared(sum) firstprivate(n)  
26 sum = fib(n);

main

task



# Implementation: Fibonacci

- Using explicit threadprivate and manual cut-off

```
1 #pragma omp threadprivate(mysum)
2 void fib(int n) {
3     if (n < 2)
4         mysum += n;
5     else {
6         if (n < cut_off) {
7             fib(n - 1);
8             fib(n - 2);
9         } else {
10            #pragma omp task firstprivate(n)
11            fib(n - 1);
12
13            #pragma omp task firstprivate(n)
14            fib(n - 2);
15        }
16    }
17 }
```

task

In results: "threadpriv"

```
18 ...
19 #pragma omp parallel num_threads(conf.num_threads)
20 {
21     mysum = 0;
22     #pragma omp single
23     #pragma omp task
24         fib(n);
25 }
26
27 #pragma omp parallel num_threads(conf.num_threads)
28 {
29     #pragma omp single
30         nthreads = omp_get_num_threads();
31     #pragma omp for reduction(+ : sum)
32         for (int i = 0; i < nthreads; i++)
33             sum += mysum;
34 }
```

main



## Evaluation

System configuration:

- **LLVM/Clang 14.0 (release)**
- **GCC 13 (code version dated 20220518)**
- Intel® Xeon® Skylake Platinum 8160 Processor, dual-socket, 48 cores, (blake.sandia.gov)
- 192GB RAM
- Flags: -fopenmp, -Wall, -Wextra, -pedantic, -Werror, -O3.
- Env: OMP\_PROC\_BIND = close, OMP\_PLACES = cores (one thread per core with incremental core IDs)


Methodology

- Each benchmark ran 5x, average time and standard deviation recorded
- Generate very large task counts to accumulate overheads
- Downside: tasking vs reduction overheads blend together


# Evaluation: Fibonacci

- 1) Stack has lowest OH, OH with # tasks, has TW
- 2) TP has low OH but requires manual reduction
- 3) Comparable performance

- Time over number of tasks, N=33, 48 threads, variable cut-off, **1.2k – 11405k tasks**



(a) LLVM/Clang




(b) GCC/g++


# Evaluation: Powerset

- 1) Similar behavior of taskwait to Fib
- 2) Taskgroup and taskloop diverge (3), due to differences in reduction list item lookup and explicit task creation (taskloop)

- Time over number of tasks, N=18, 48 threads, variable cut-off, **2 - 265k tasks**



(a) LLVM/Clang




(b) GCC/g++


# Evaluation: Powerset UDR

- 1) Stack accesses causing cache faults (LLVM)
- 2) Differences in reduction list item look-up and explicit task creation (taskloop)

- Time over reduction size (UDR), **4B – 131kB**, 48 threads, 262k tasks



(a) LLVM/Clang




(b) GCC/g++

## Evaluation: Dot-product (not recursive)

- 1) Execution with work for 48 threads
- 2) Parallel for is invariant as it does not use tasks

- Time over number of tasks,  $N=2^{24}$ , 128MB, 48 threads, variable cut-off, **2 - 131k tasks**





## Conclusion

- Both compilers support task-parallel reductions as in the spec
- Performance is comparable to manually implemented and optimized reductions
- For large task counts, tasking overheads dominate (incl. taskwait)
- We recommend the use of these constructs

### Future work

- Evaluation of tasking implementations in LLVM/Clang and GCC

### Links:

[1] [https://github.com/sandialabs/openmp\\_task\\_bench](https://github.com/sandialabs/openmp_task_bench)