

Reducing cache miss rate using thread overallocation to accelerate the MPI-OpenMP-based 2-D Hopmoc method

Frederico L. Cabral (fcabral@lncc.br)¹, Gabriel Costa (gcosta@lncc.br)¹, Carla Osthoff (osthoff@lncc.br)¹, Sanderson L. Gonzaga de Oliveira (sanderson@dcc.ufla.br)²
 1 - National Laboratory of Scientific Computing (LNCC), 2 - Federal University of Lavras (UFLA)

Abstract

This paper applies the MPI-OpenMP-based two-dimensional Hopmoc method using the explicit work-sharing technique with a recently proposed mechanism to reduce implicit barriers in OpenMP. Specifically, this paper applies the numerical algorithm to provide approximate solutions to the advection-diffusion equation. Additionally, this article splits the mesh used by the numerical method and distributes them to over-allocated threads. The mesh partitions became so small that the approach reduced the cache miss rate. Consequently, the strategy accelerated the numerical method in multicore systems. This paper then evaluates the results of implementing the strategy under different metrics. As a result, the use of the set of techniques improved the performance of the parallel numerical method.

HOPMOC versions

```

1 begin
2   Each MPI process computes its partition mesh
3   while time < FinalTime do
4     ...
5     #pragma omp for
6     for columns in the partition mesh do
7       for all lines in the input mesh do
8         MMOC or an explicit or implicit time semi-step
9       end
10    end
11  ...
12 end
13 end
    
```

(a) Pseudocode excerpt of parallel Hopmoc method using naive MPI-OMP

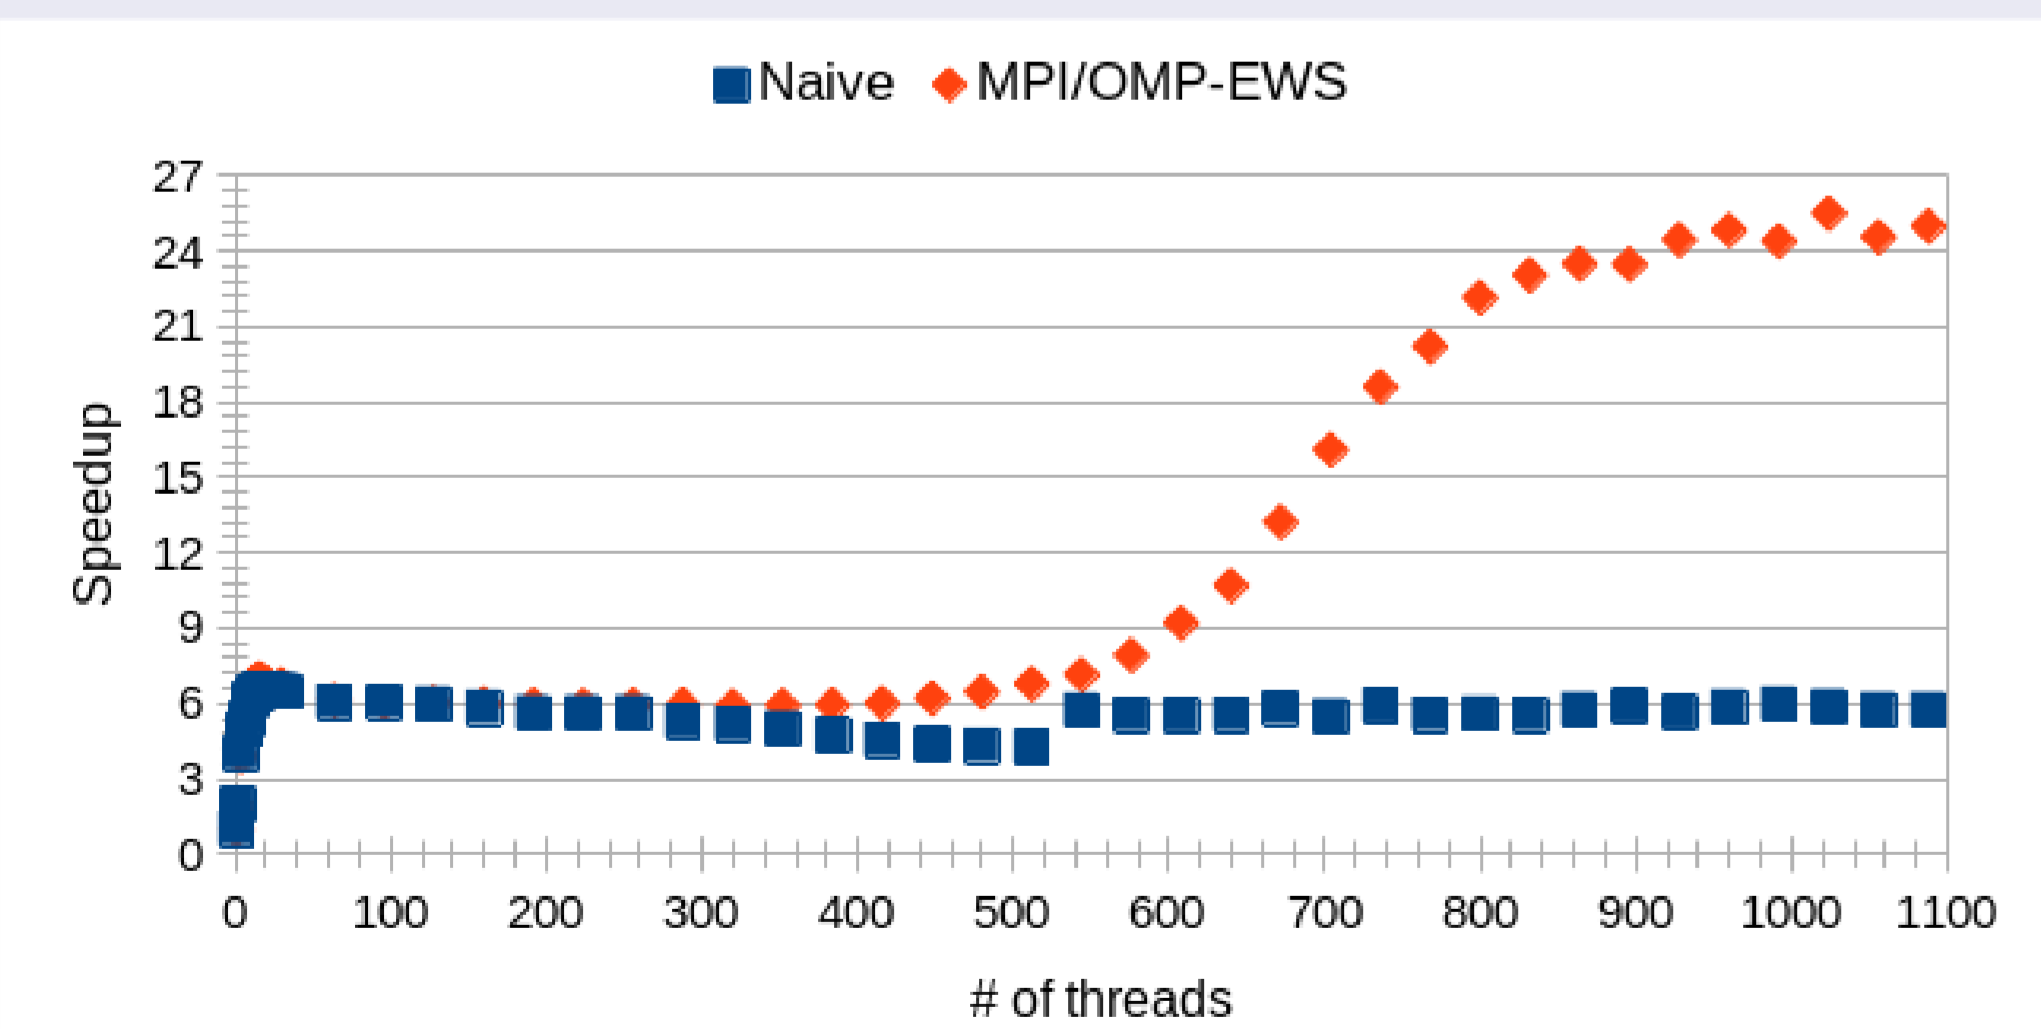
```

1 begin
2   Each MPI process computes its partition mesh
3   Each OMP thread computes its local
4   mesh inside the partition mesh allocated to a process
5   while time < FinalTime do
6     ...
7     lock
8     for columns in the thread local mesh do
9       for all lines in the input mesh do
10        MMOC or an explicit or implicit time semi-step
11      end
12    end
13    unlock
14    Each OMP thread waits for its neighbours
15  ...
16 end
17 end
    
```

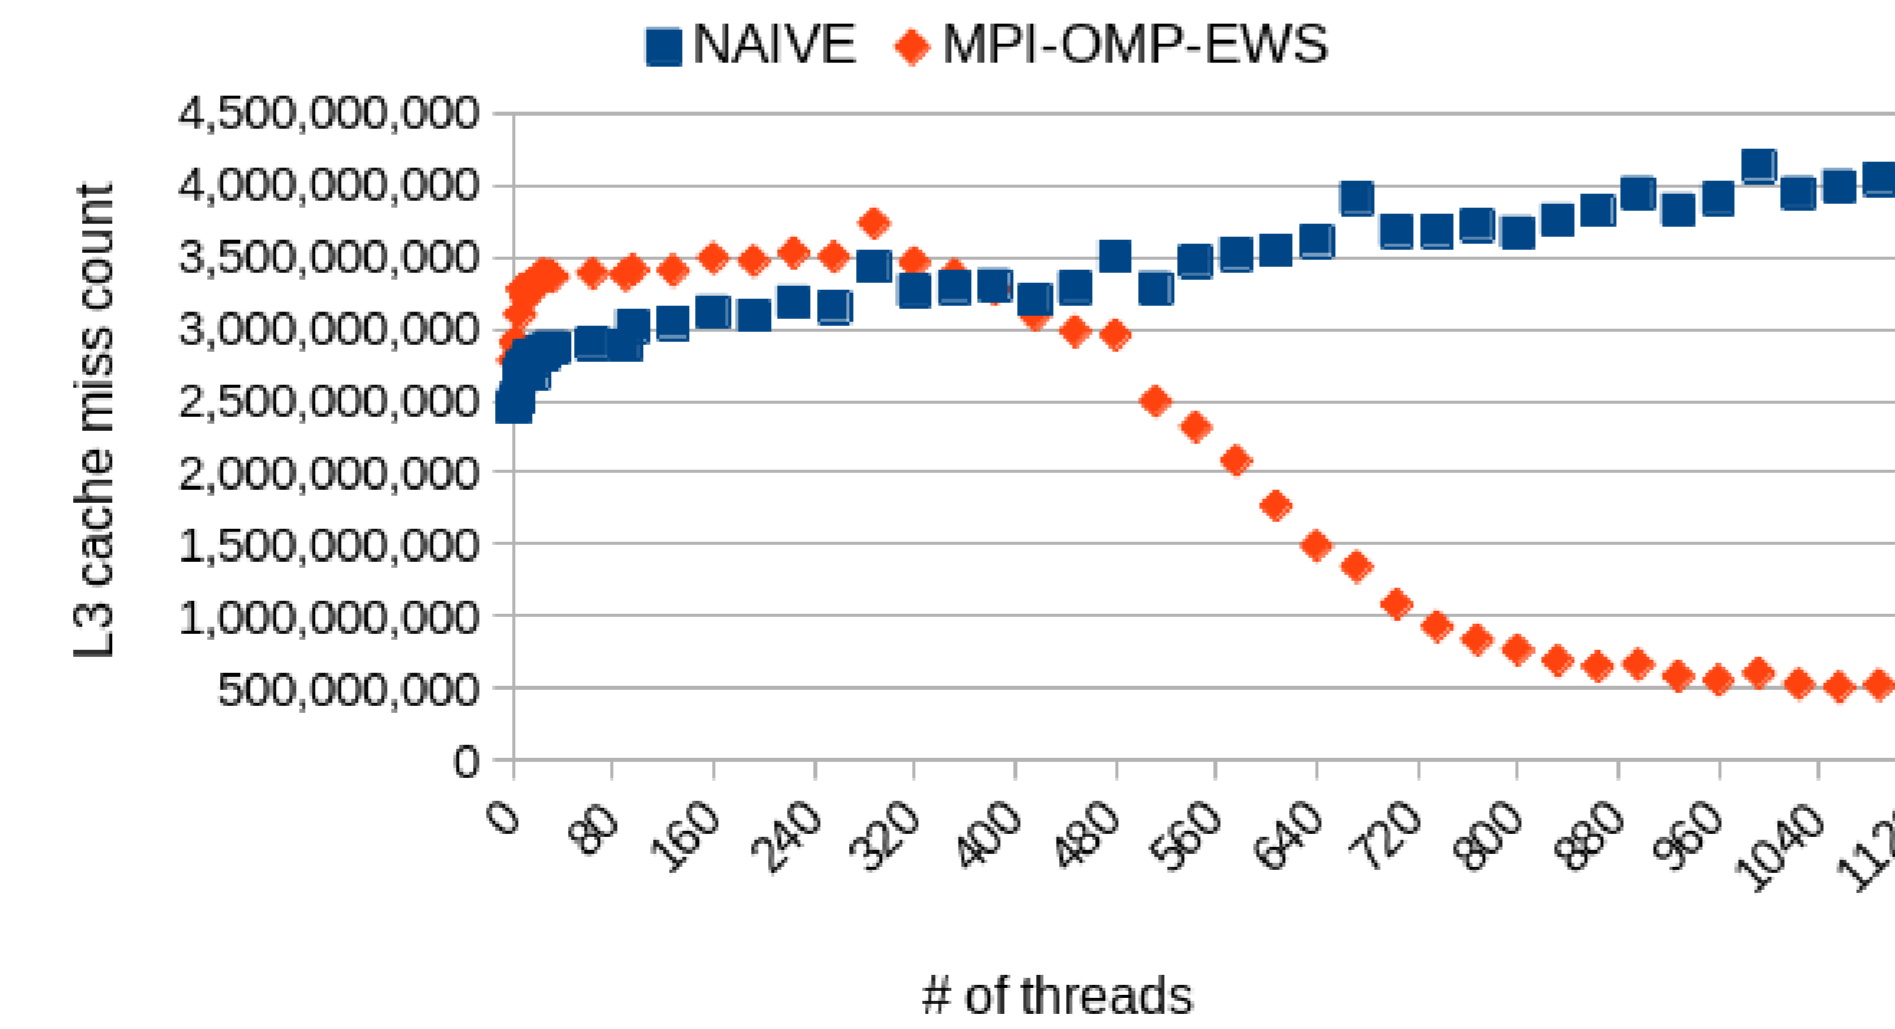
(b) Pseudocode excerpt of parallel Hopmoc method using MPI-OMP-EWS

Figure 1: Two distinct implementations of the Hopmoc method

Experiment performed on a Intel® Xeon® processor E5-2698 machine (32 physical cores)




(a) Speedups from executions of naive and MPI-OMP-EWS implementations of the two-dimensional Hopmoc method varying the number of threads from 1 to 1,100 threads.



(b) LLC miss count from executions of naive and MPI-OMP-EWS implementations of the two-dimensional Hopmoc method varying the number of threads from 1 to 1,200 threads

Figure 2: Speedup and LLC miss count



(a) L1 cache miss count from executions of naive and MPI-OMP-EWS implementations of the two-dimensional Hopmoc method varying the number of threads from 1 to 3,200 threads

(b) L2 cache miss count from executions of naive and MPI-OMP-EWS implementations of the two-dimensional Hopmoc method varying the number of threads from 1 to 3,200 threads

Figure 3: L1 and L2 cache miss counts

Experiment performed on a Intel® Xeon® Gold 6230R (104 physical cores and 208 threads)

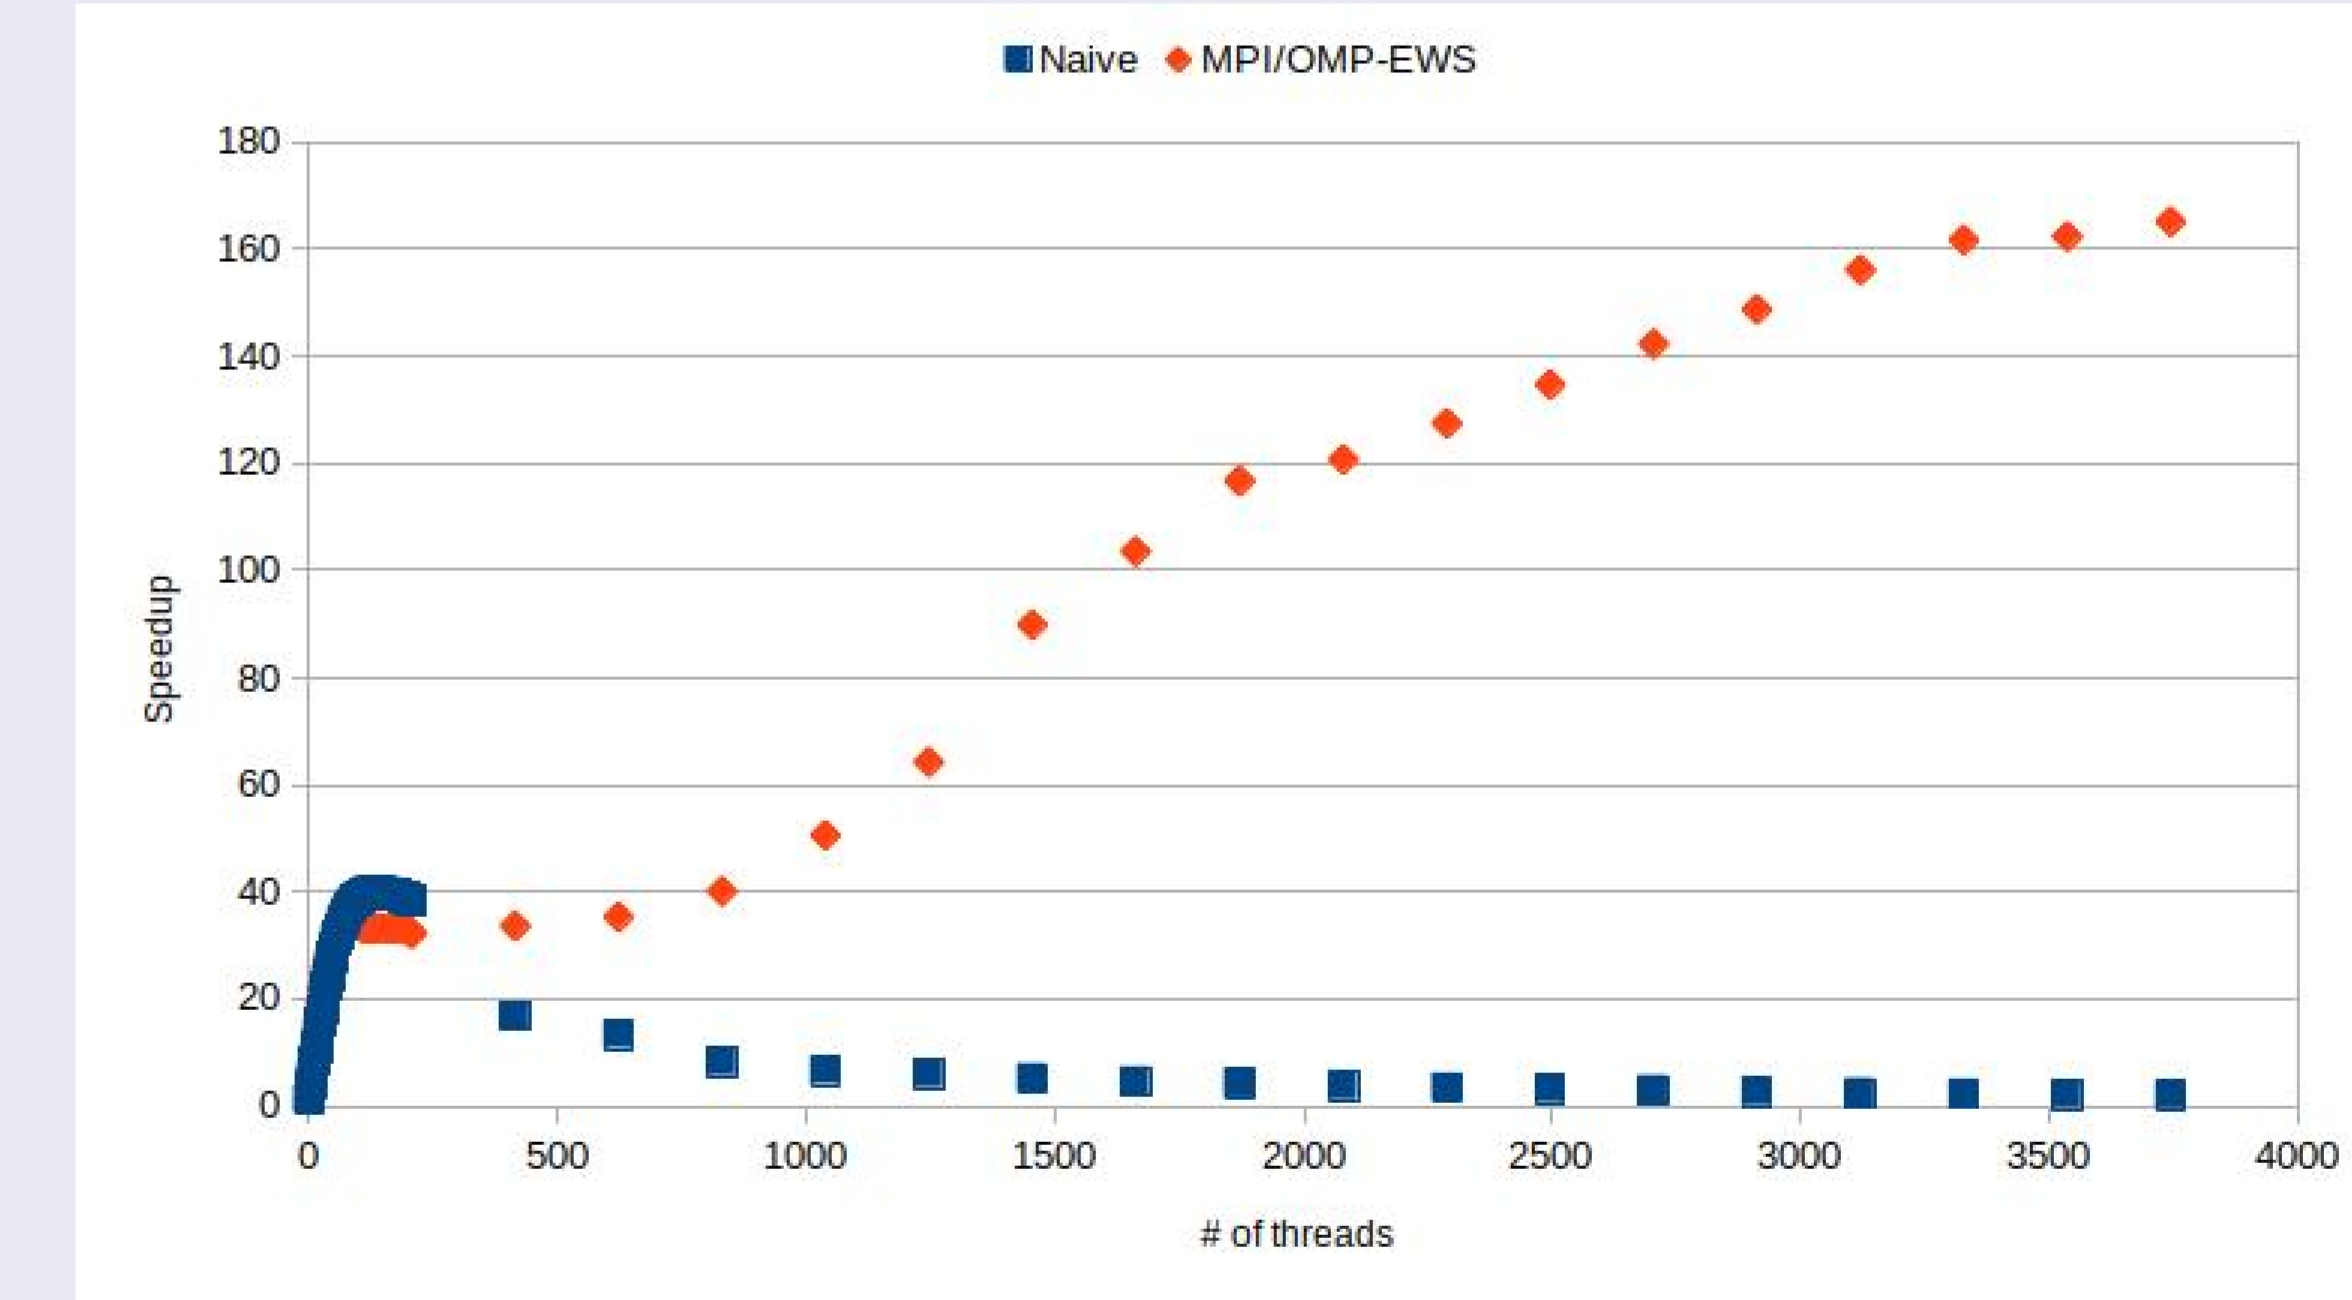


Figure 4: Speedups from executions of naive and MPI-OMP-EWS-based implementations of the two-dimensional Hopmoc method varying the number of threads from 1 to 3,744 threads.

Bibliography

- 1 Cabral, F.L., Gonzaga de Oliveira, S.L., Osthoff, C., Costa, G.P., Brandao, D.N., Kischinhevsky, M.; An evaluation of MPI and OpenMP paradigms in finite-difference explicit methods for PDEs on shared-memory multi- and many-core systems. *Concurrency and Computation: Practice and Experience*, 32(20), e5642 (2020). <https://doi.org/https://doi.org/10.1002/cpe.5642>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5642>, e5642 cpe.5642