



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Scalability of OpenIFS48r1 + M7

Simone Vacondio

BSC-Earth Sciences Dpt. | Computational Earth Sciences

16/10/2025

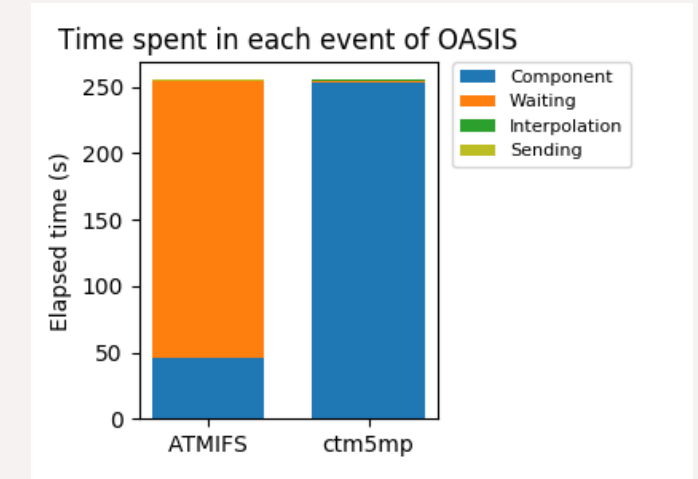
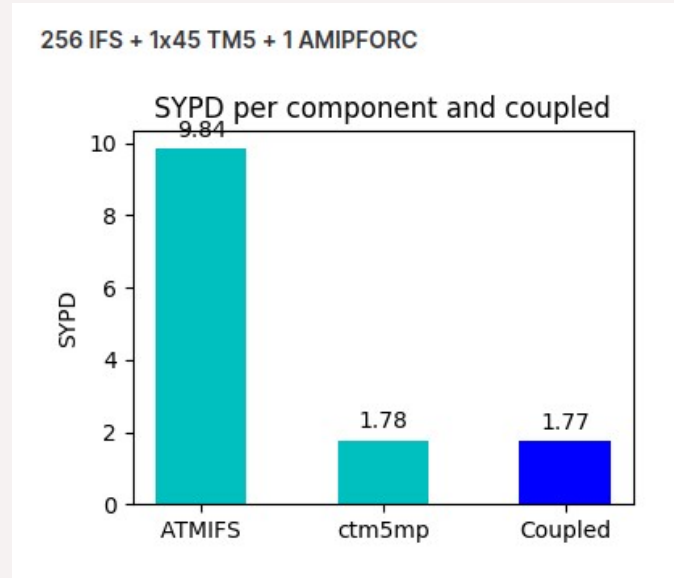
EC-Earth Consortium Meeting - October 2025

M7 in OpenIFS

From coupled component to hard-coded module

- M7/TM5: **bottleneck**
- **Concurrent coupling** (via OASIS)
 - pros: parallel efficiency (in principle)
 - cons: different grids, exchange of 3D fields, memory issues
- **Hard-coding**
 - pros: same grid, improved resolution
 - cons: no parallel efficiency (sequential coupling overhead)

EC-Earth 3 on MN4 | coupling the IFS to TM5



M. C. Acosta, S. Palomas, E. Tourigny, Earth and Space Science, 10, e2023EA002912 (2023)

EC-Earth 4 on MN5 | hard-coding M7 in OpenIFS

tasks	threads	no aerosols	aerosols + simpl. chem.	aerosols + full chem. (TM5)
1	112	27.068	7.893	5.761
2	56	37.711	10.746	8.102
4	28	48.133	12.972	9.558
8	14	58.125	13.845	9.885
14	8	59.915	14.472	9.699
28	4	55.946	12.687	9.709

(SYPD)

Details

- MN5 runs
- 112 CPUs/node
- 32 nodes

Scalability of OpenIFS-48r1+M7

Scalability of OIFS 48r1+M7

Running on MN5

The three main configurations for **aerosols** were tested:

- no aerosols
- aerosols (M7) @
simplified chemistry
- aerosols (M7) @ TM5
chemistry

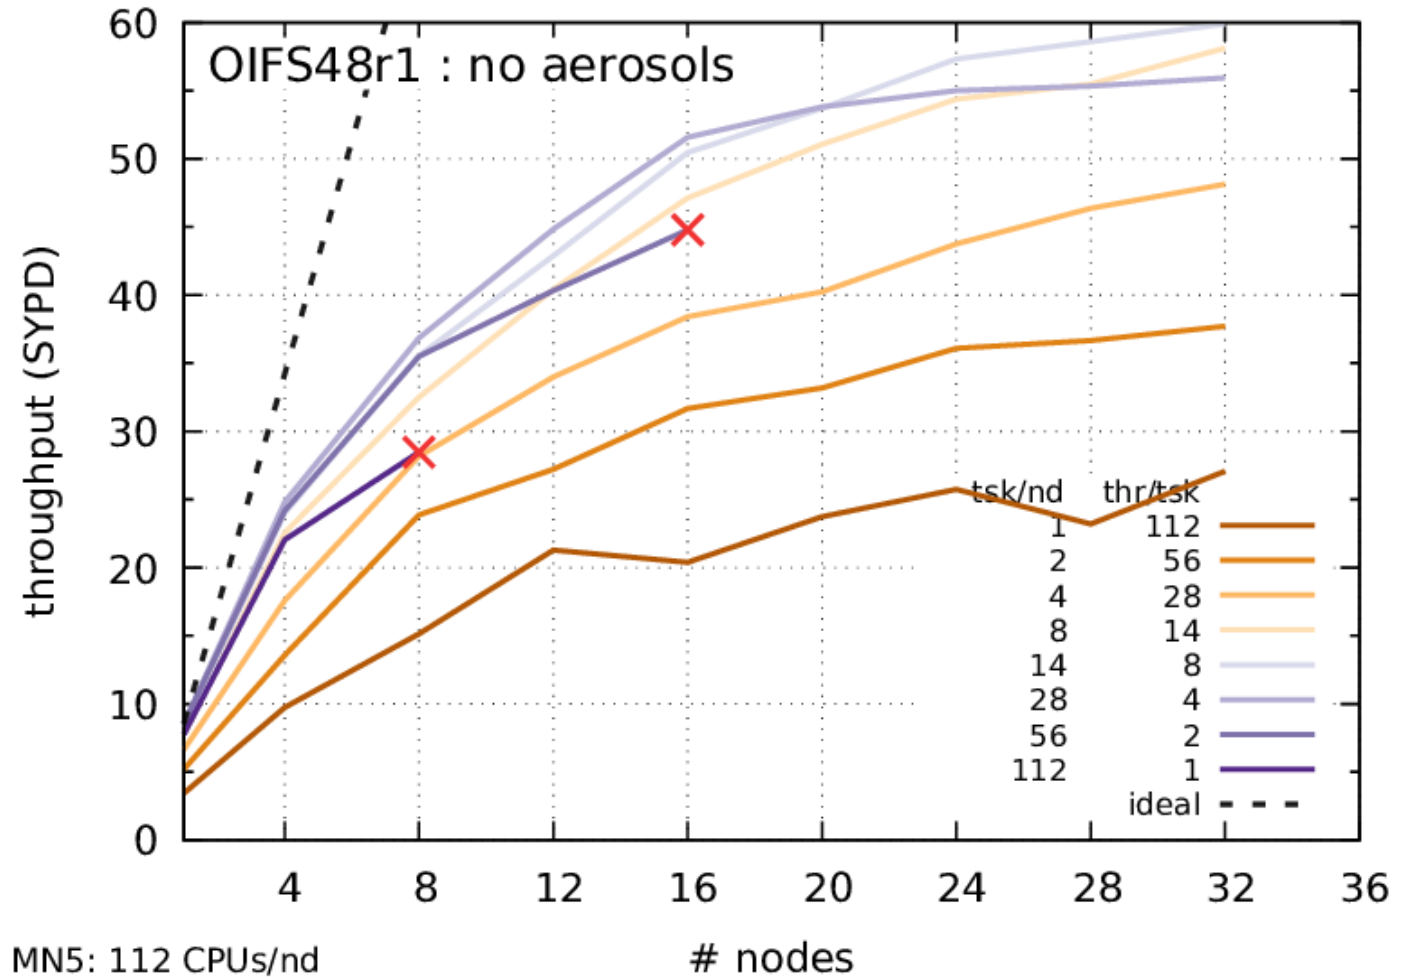
Also: **standalone OIFS48r1**

```
resolution           : TL255L91  
start date          : 1990-01-01  
emiss. scheme       : CMIP7  
simulated time      : 5 days  
model timestep      : 45 mins  
output freq         : 1 every 3.5 hrs.
```

Scalability of OIFS 48r1+M7

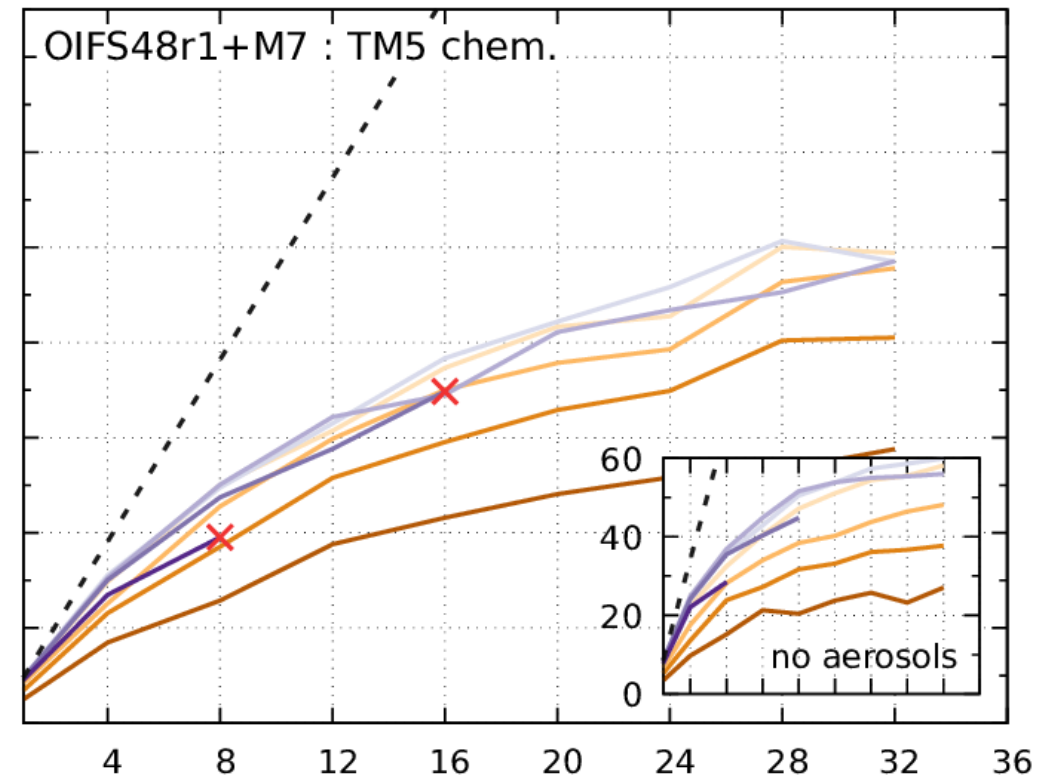
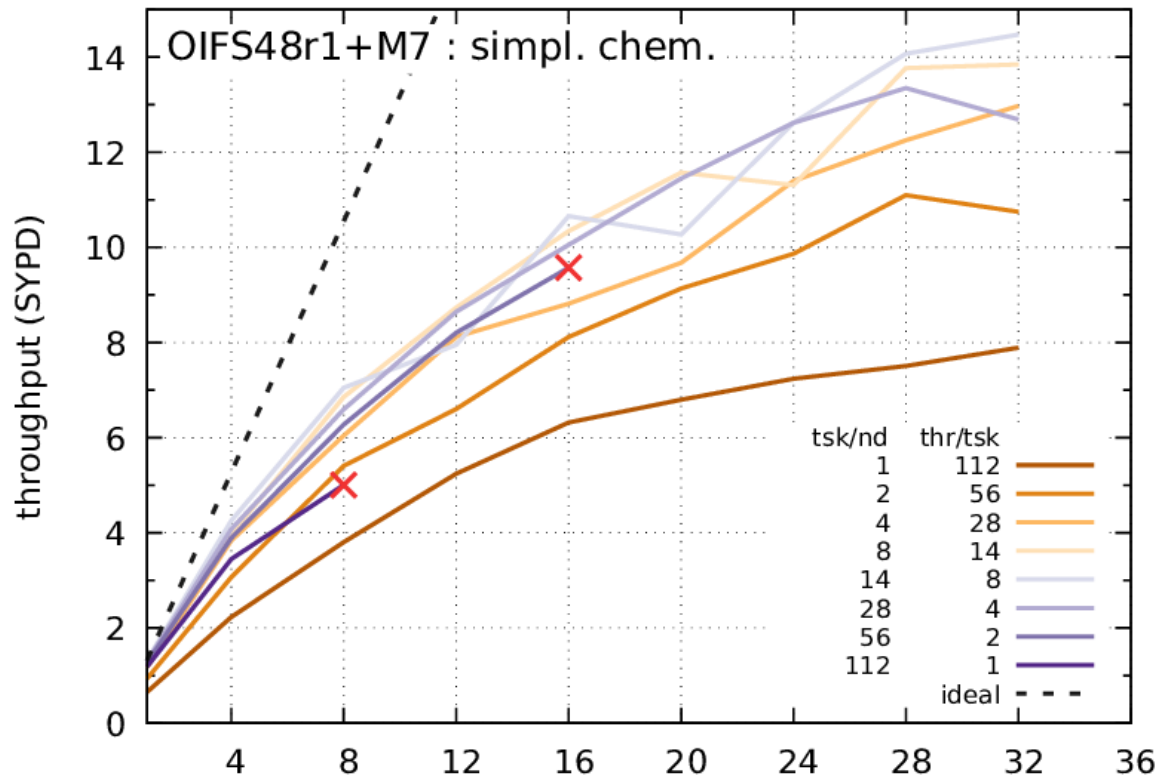
No aerosols

- # MPI tasks limited by resolution (~1000 max, **x**)
- **OpenMP** enables running on more nodes
- **MPI/OpenMP load**
 - no throughput gain for > 8 OMP thr/MPI tsk



configuration : no aerosols
grid resolution : 1_2 255
levels : 91
start date : 1990-01-01

emiss. scheme : CMIP7
simulated time : 5 days
model timestep : 45 mins
output period : 3.5 hrs.



MN5: 112 CPUs/nd

nodes

Scalability of OIFS 48r1+M7

With aerosols

- Performance hit
- ... however, similar scaling

- OpenMP scaling **still capped at 8 thr/tsk**
- perhaps **OpenMP in M7** could help with the bottleneck?

Tracing OpenIFS-48r1+M7

Tracing OIFS 48r1+M7

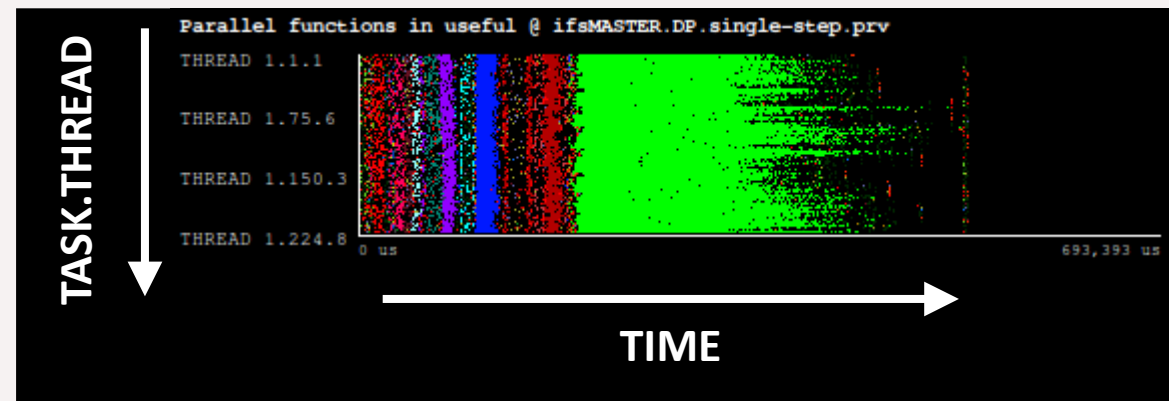
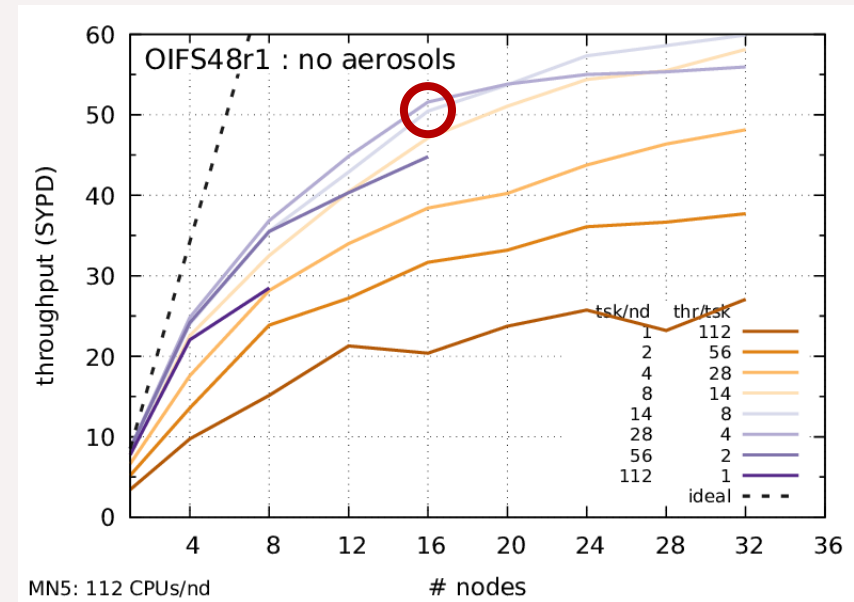
Traced configuration

- 16 nodes
- 14 tsk/nd
- 8 thr/tsk

Output

- Scatter plot
 - `task.thread` event

VS
time

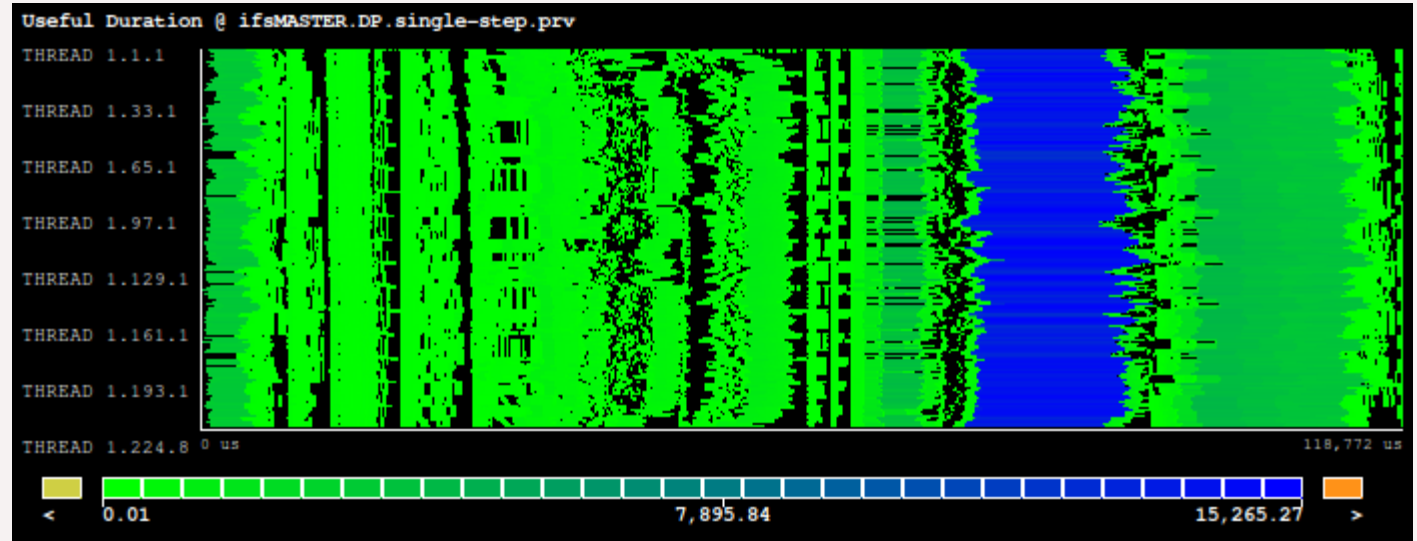


Tracing OIFS 48r1+M7


Useful duration vs MPI calls

- useful duration
 - = time not spent for MPI/OMP calls (comm., forking, etc.)
 - the **bluer** the stripe, the **longer** the chunk

Useful duration



MPI calls



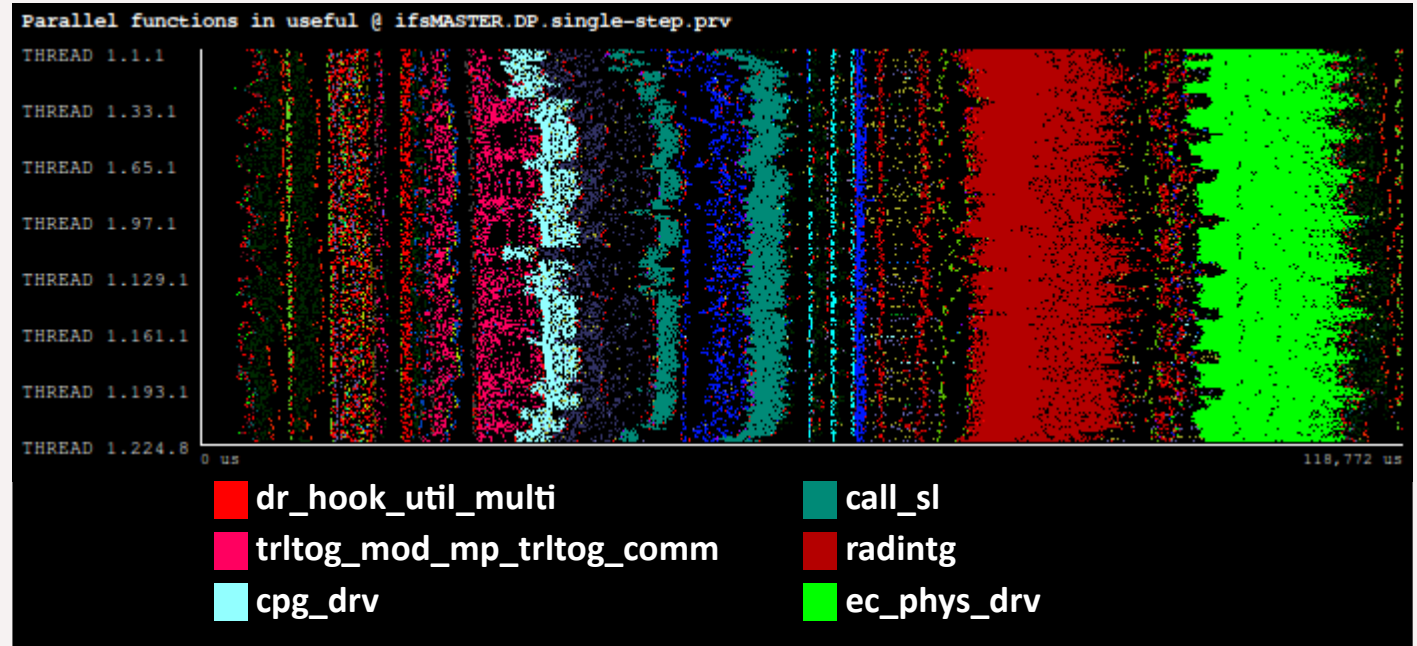
configuration : no aerosols
traced time : one timestep

Tracing OIFS 48r1+M7




OpenMP parallel functions

- Time spent in OpenMP
 - who : **colour**
 - how long : **length of chunk**
- Correlation with useful duration
 - most often,
in the present case

OpenMP parallel functions



configuration : **no aerosols**
traced time : **one timestep**



- call_sl
- raddrv
- radintg
- ec_phys_drv

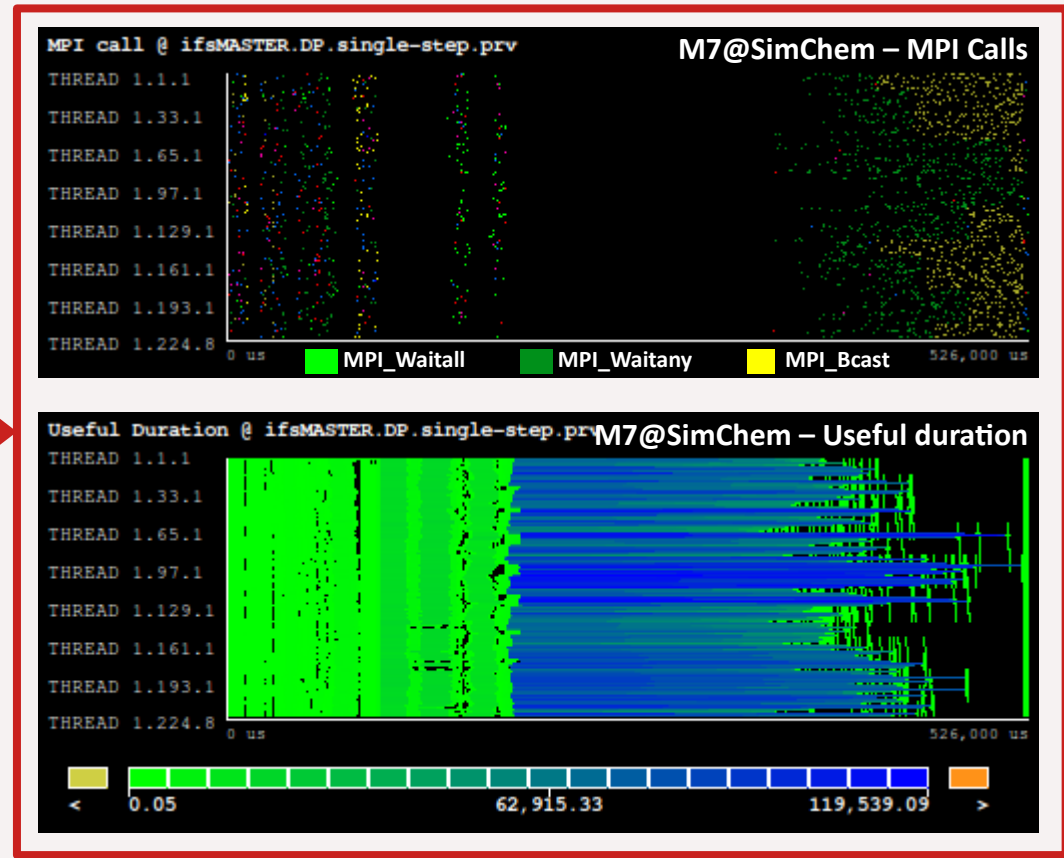


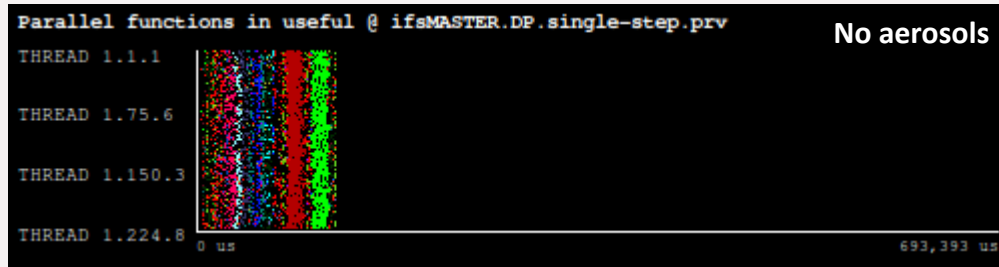
hamm7_interface
nested calls here

traced time : one timestep

Tracing OIFS48r1+M7

Comparing configurations

- Increased computation time for a number of routines
- ... but especially **ec_phys_drv**, which **calls M7 routines**
- **ec_phys_drv**'s trace suggests load imbalance

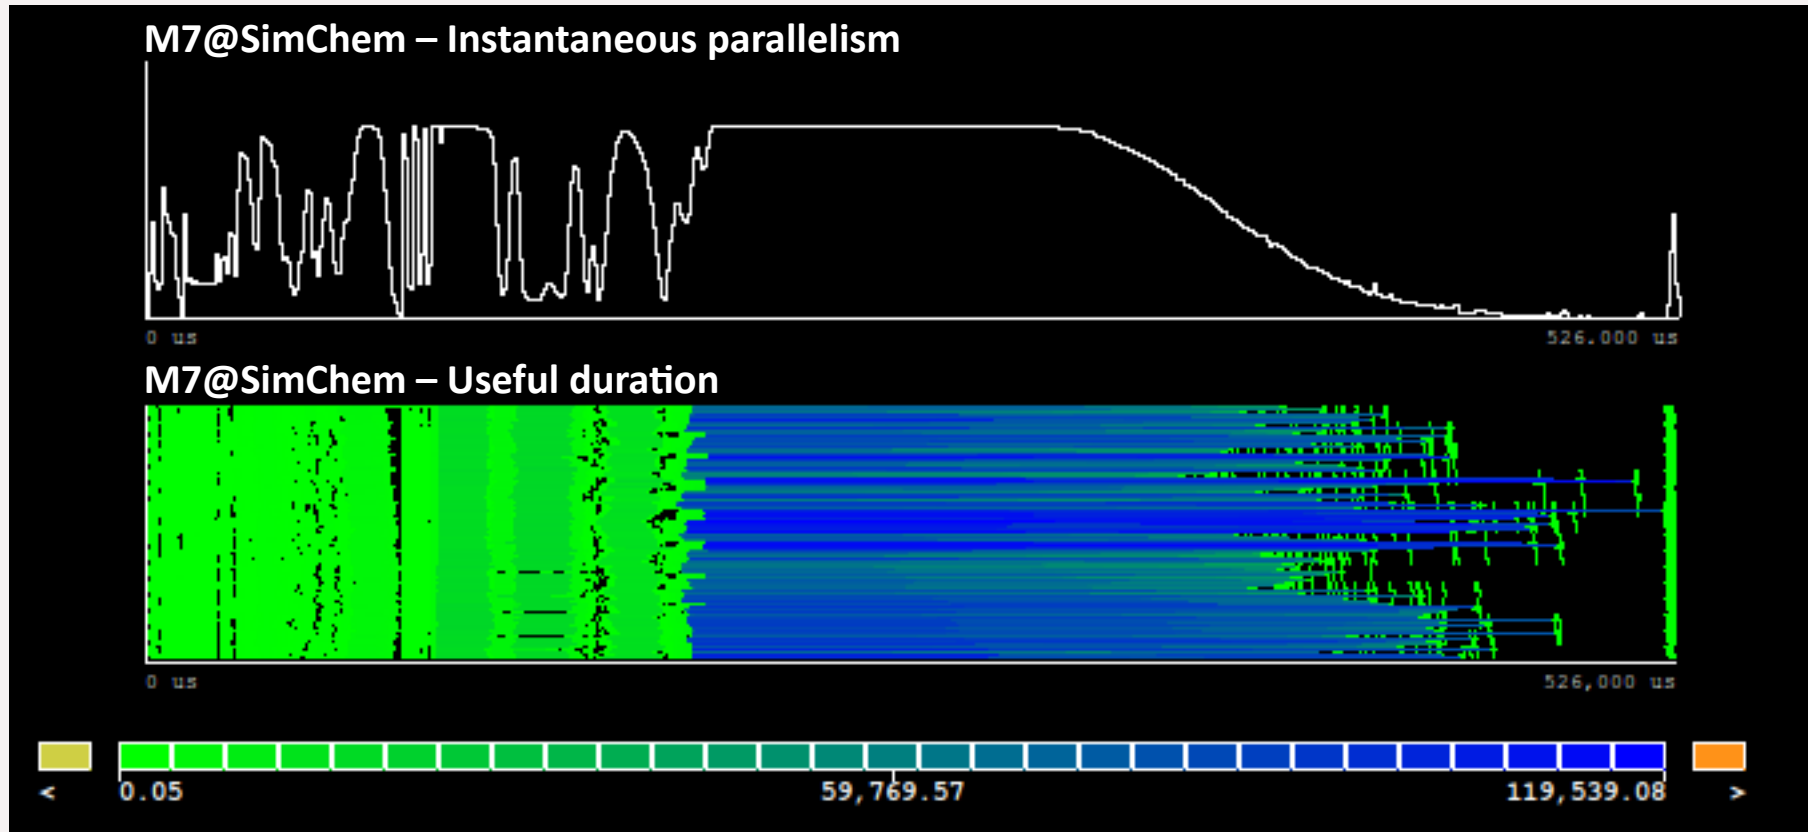


traced time : one timestep

Tracing OIFS48r1+M7

Load imbalance

- The **OMP** parallel duration of `ec_phys_drv` is indeed correlated with **useful duration**
- Time spent **waiting** and **broadcasting** at the end of the timestep



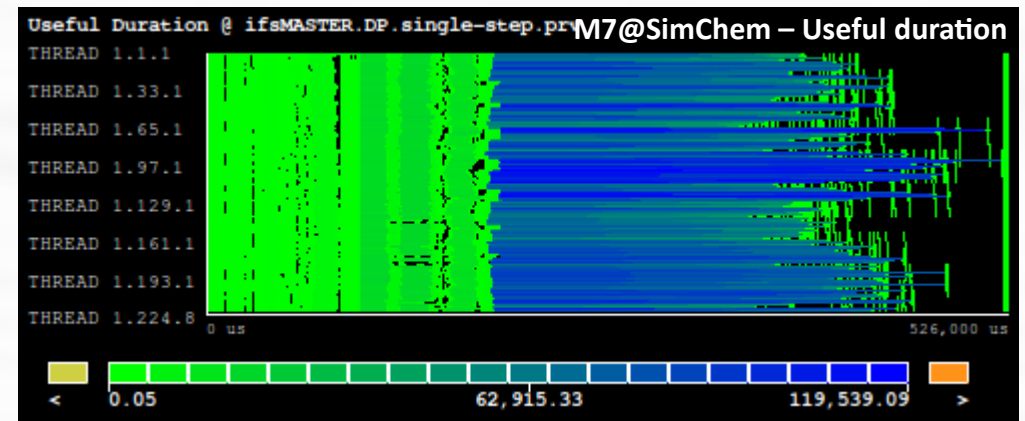

Tracing OIFS48r1+M7

One further metric

- Instantaneous parallelism
 - # MPI tasks doing **useful work** vs **time**
 - degradation in the black region of useful duration
 - further evidence of load imbalance

Conclusions

- The **M7 module** for aerosols introduces a **significant bottleneck** in the execution of OpenIFS48r1
- M7 does **not affect the scaling** of OIFS heavily
 - OpenMP performance cap at **8 OMP thr/MPI tsk**
 - an opening for addressing the bottleneck?
- M7 introduces **load imbalance** in the parallel execution of **ec_phys_drv**



Conclusions

- The **M7 module** for aerosols introduces a **significant bottleneck** in the execution of OpenIFS48r1
- M7 does **not affect the scaling** of OIFS heavily
 - OpenMP performance cap at **8 OMP thr/MPI tsk**
 - an opening for addressing the bottleneck?
- M7 introduces **load imbalance** in the parallel execution of `ec_phys_drv`

[Thank you!](#)

