Longitudinal MPI decomposition with reduced grid in TM5-MP

-

Philippe Le Sager, KNMI

2015-10-06

-

23th ITM5 meeting

Reduced Grid

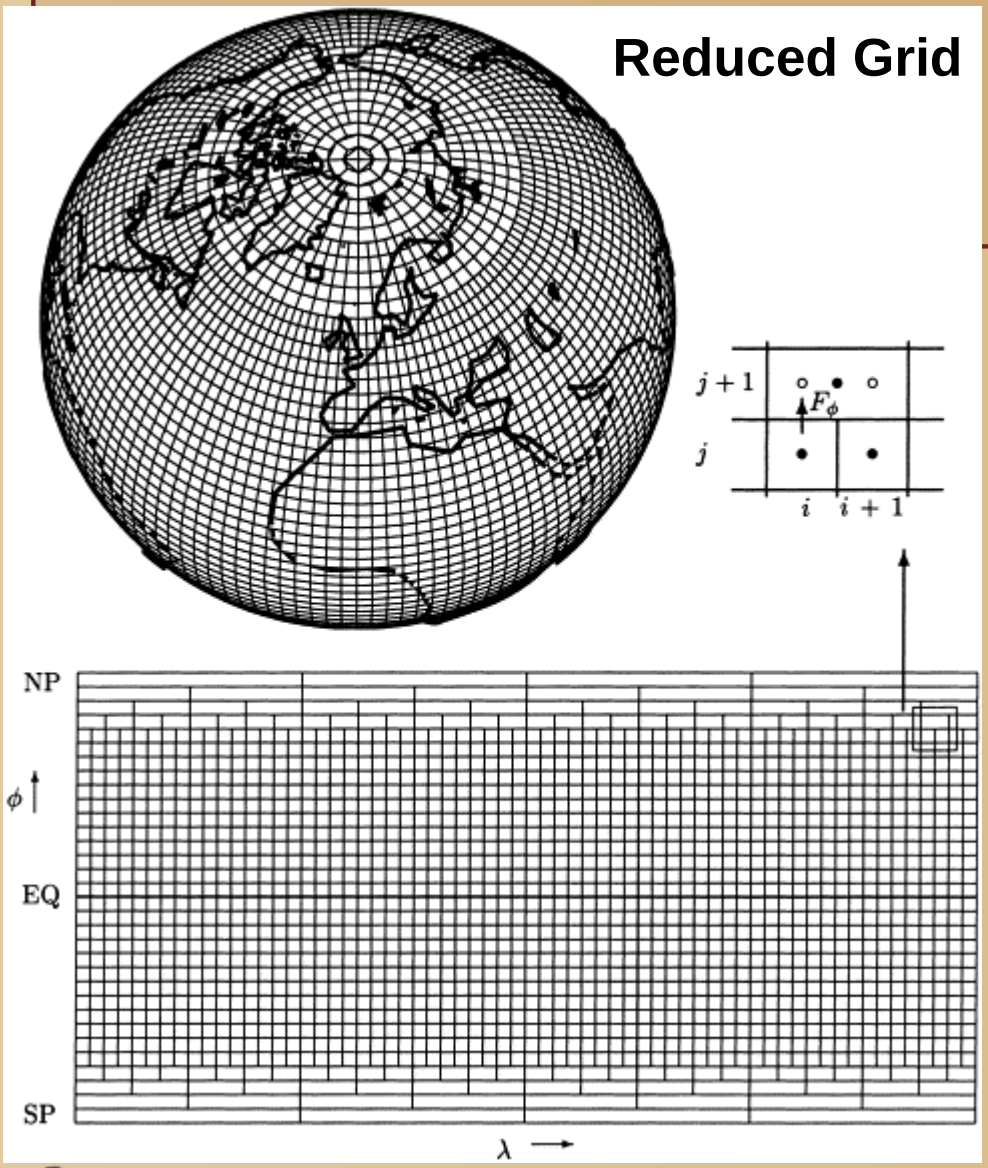

Current MPI decomposition

max N is 45 at 3x2

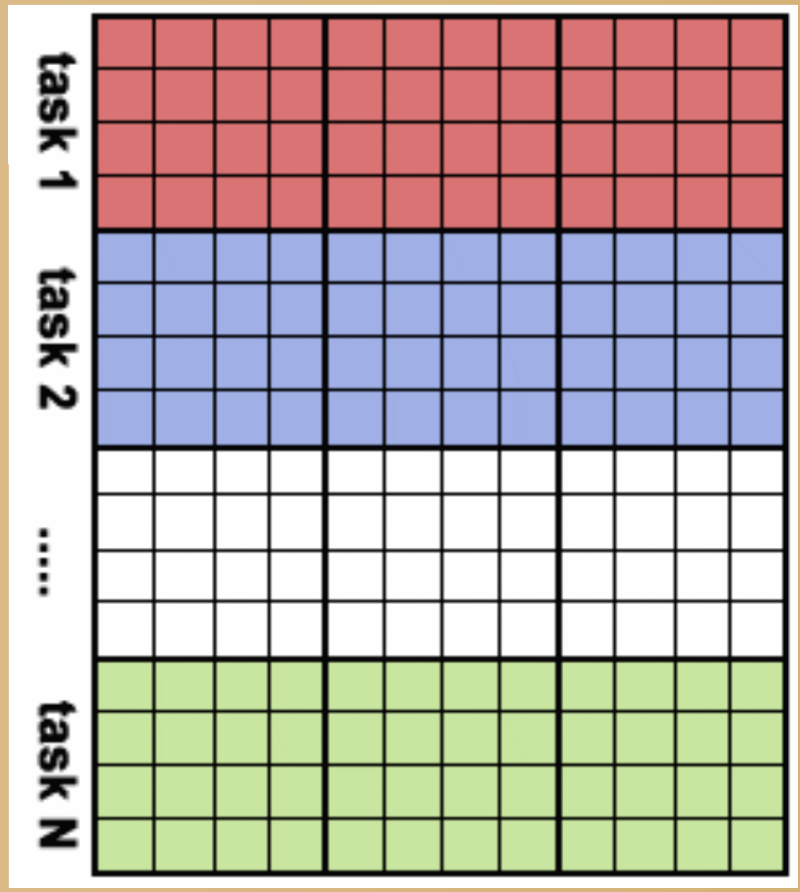

Implemented solution

(simple but works with ALL reduced grids)

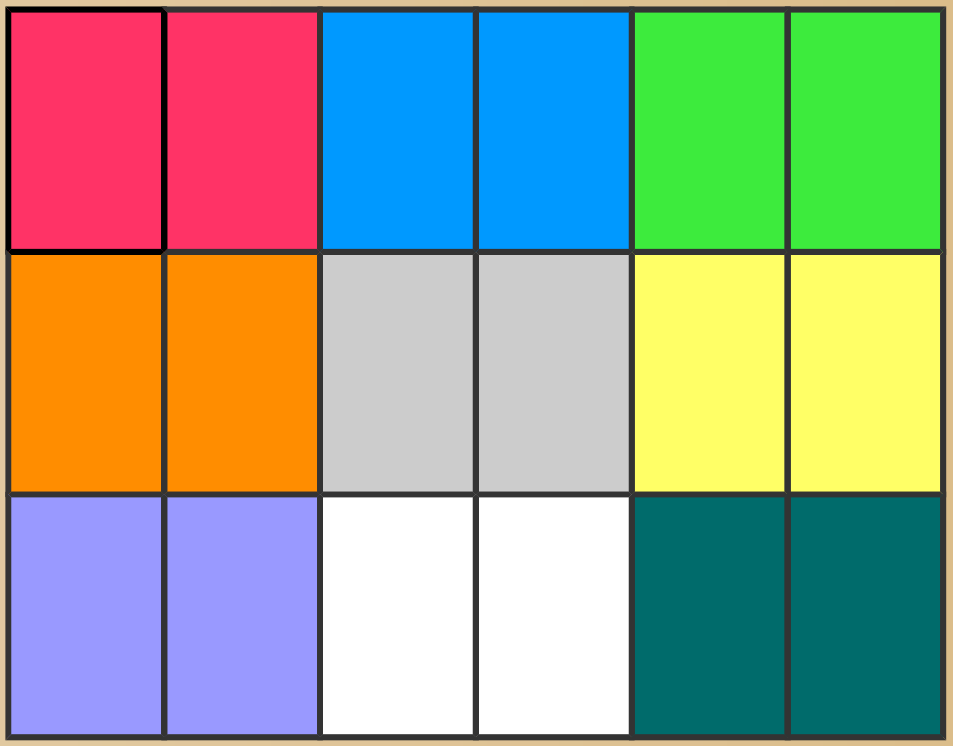

Regular grid

Implemented solution

(simple but works with ALL reduced grids)

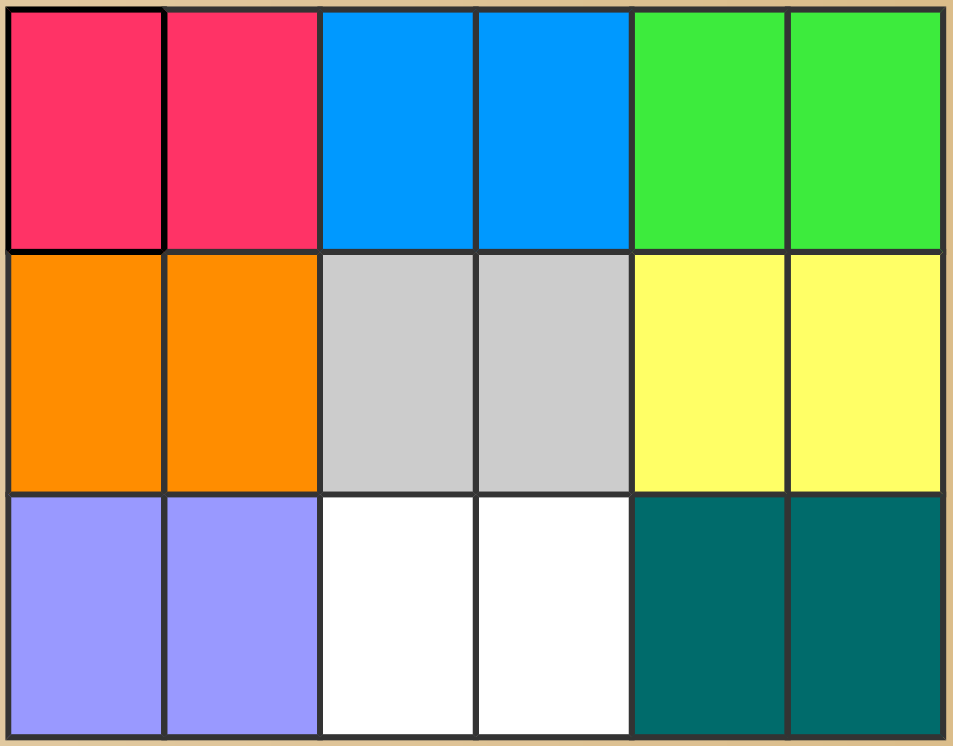

Regular grid

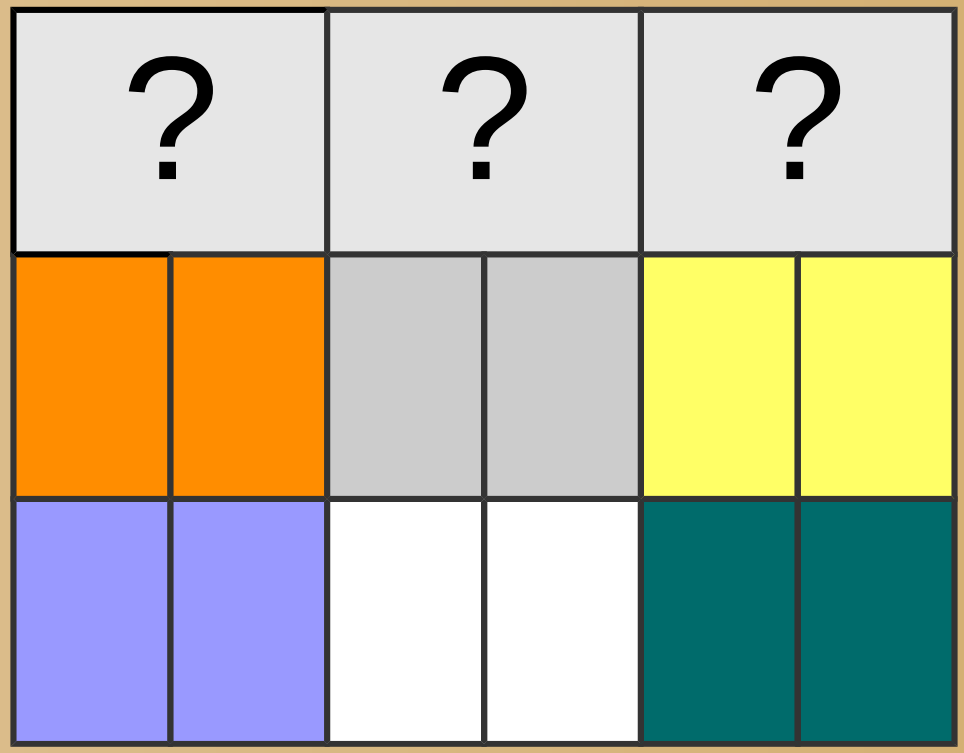

Reduced grid

Implemented solution

(simple but works with ALL reduced grids)

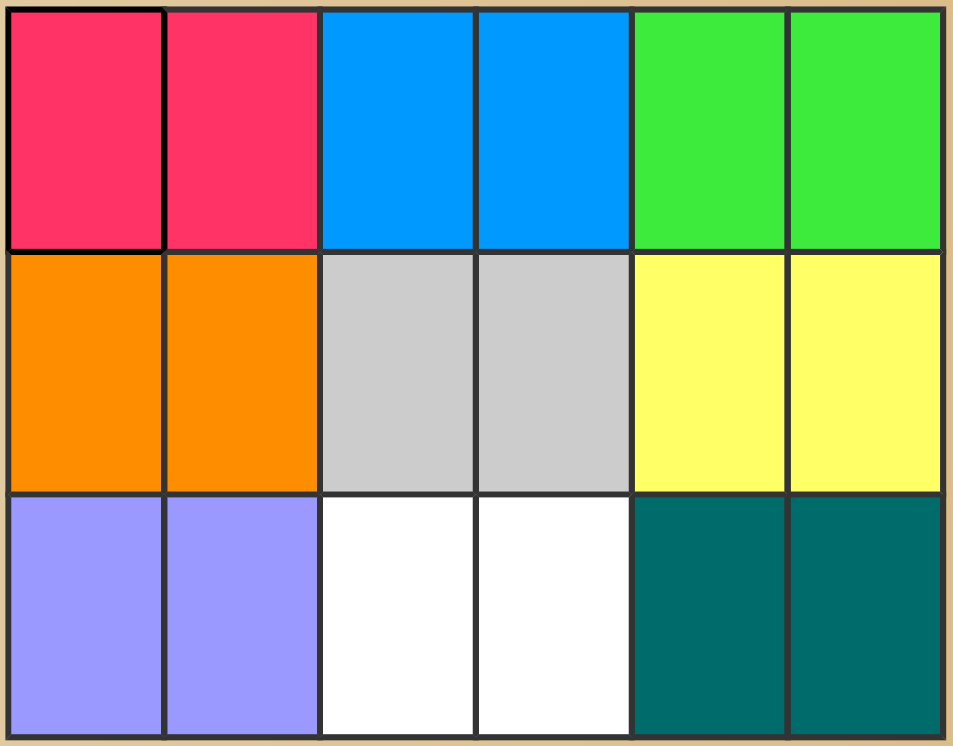

Regular grid

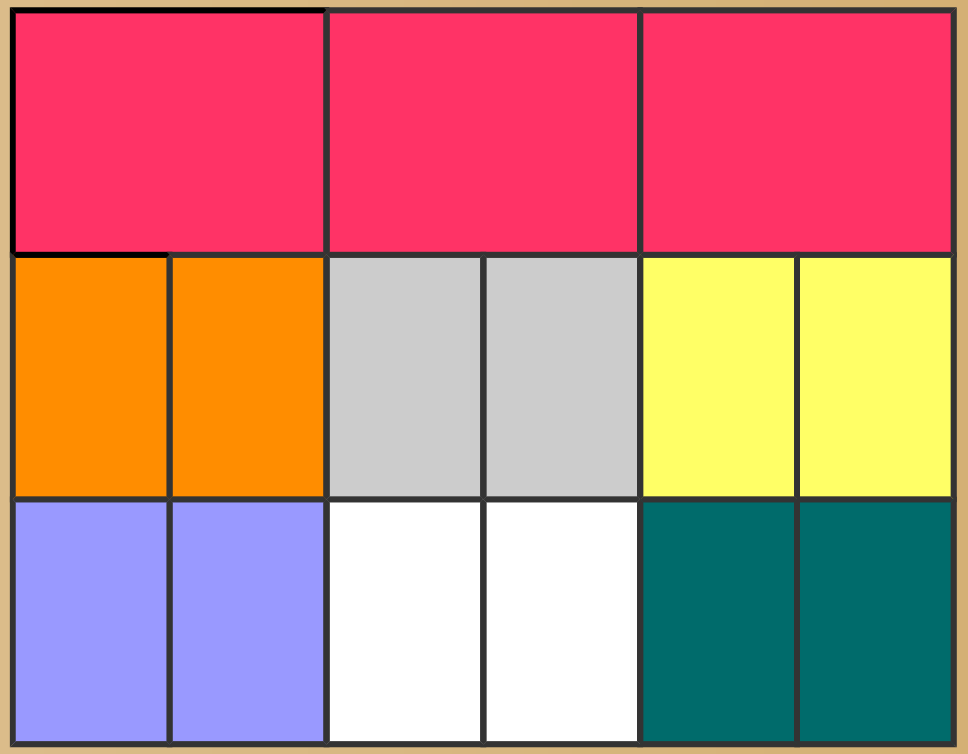

Reduced grid

(blue and green proc are idle)

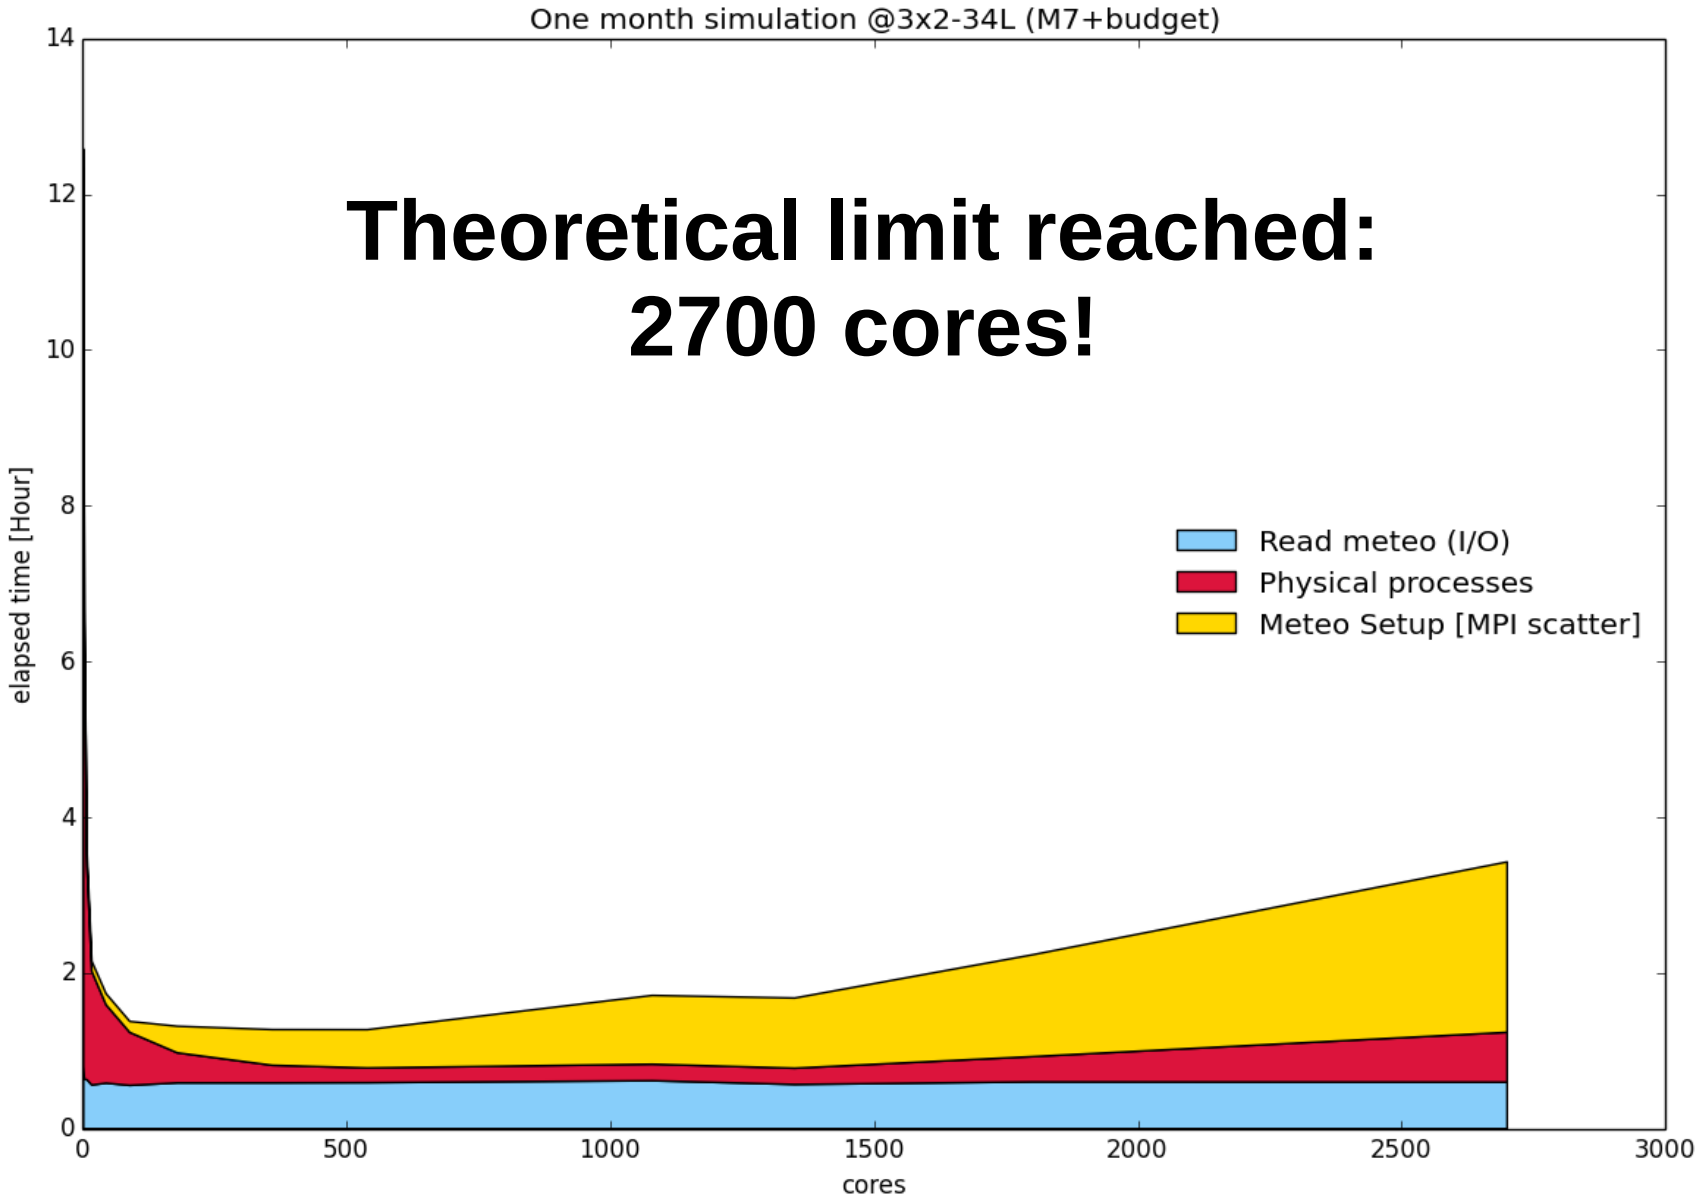
One month simulation @3x2-34L (M7+budget)
Theoretical limit reached:
2700 cores!
Read meteo (I/O)
Physical processes
Meteo Setup [MPI scatter]
elapsed time [Hour]
cores
0
2
4
6
8
10
12
14
0
500
1000
1500
2000
2500
3000

Good solution up to ~500 cores...

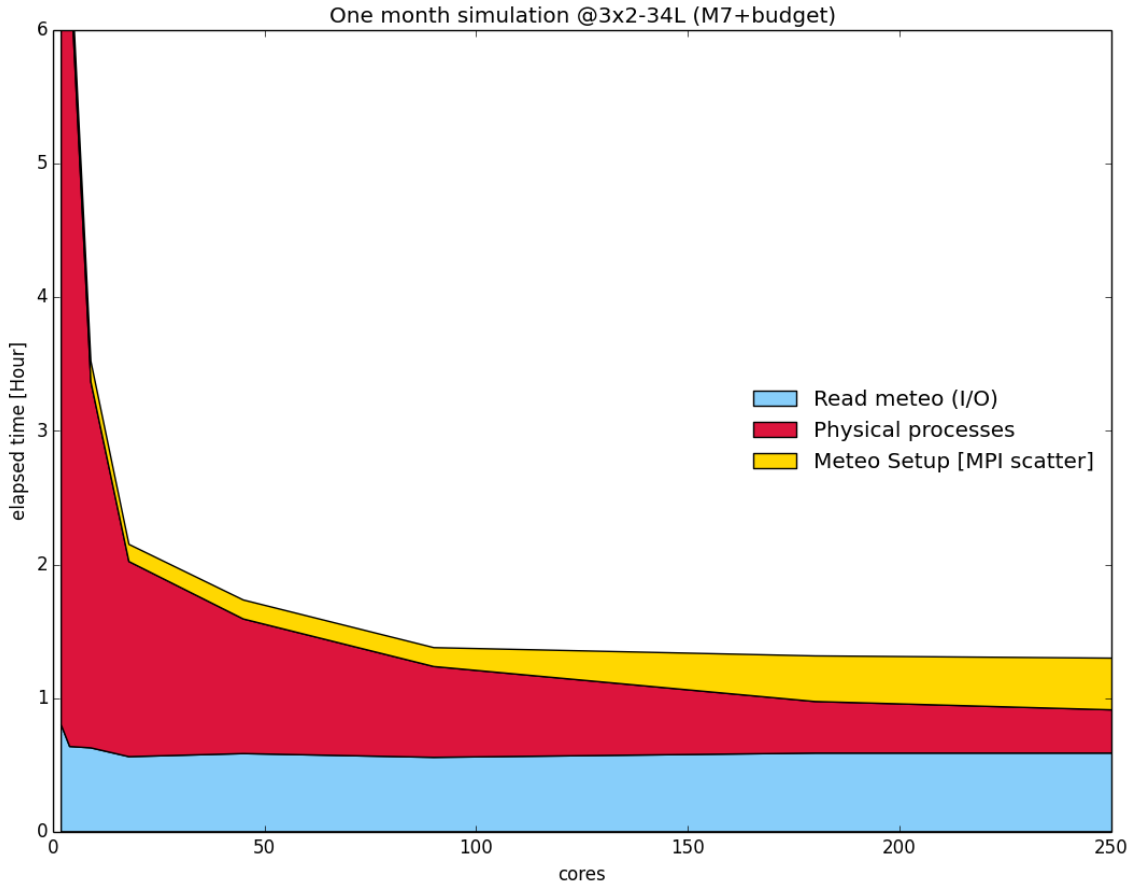

...but only 90 cores with the meteo:

reading and **scattering** meteo is the problem

(and machine dependent)

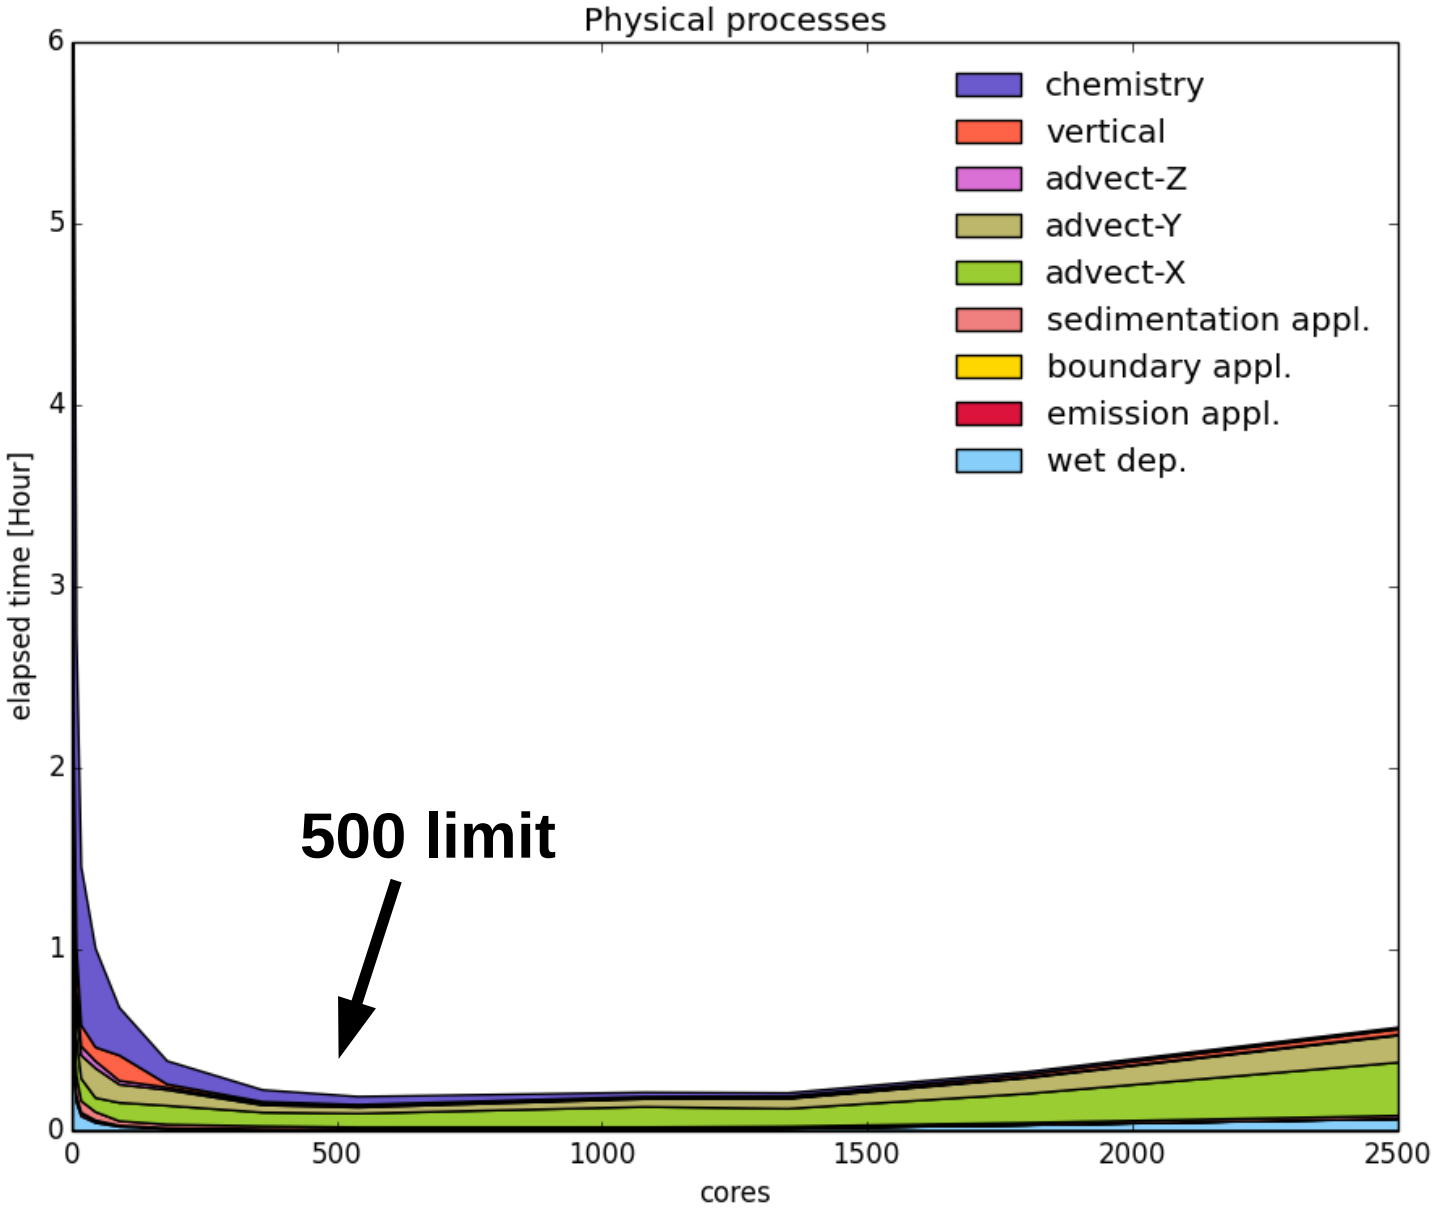
Physical processes
chemistry
vertical
advect-Z
advect-Y
advect-X
sedimentation appl.
boundary appl.
emission appl.
wet dep.
elapsed time [Hour]
cores
0
1
2
3
4
5
6
0
500
1000
1500
2000
2500
500 limit

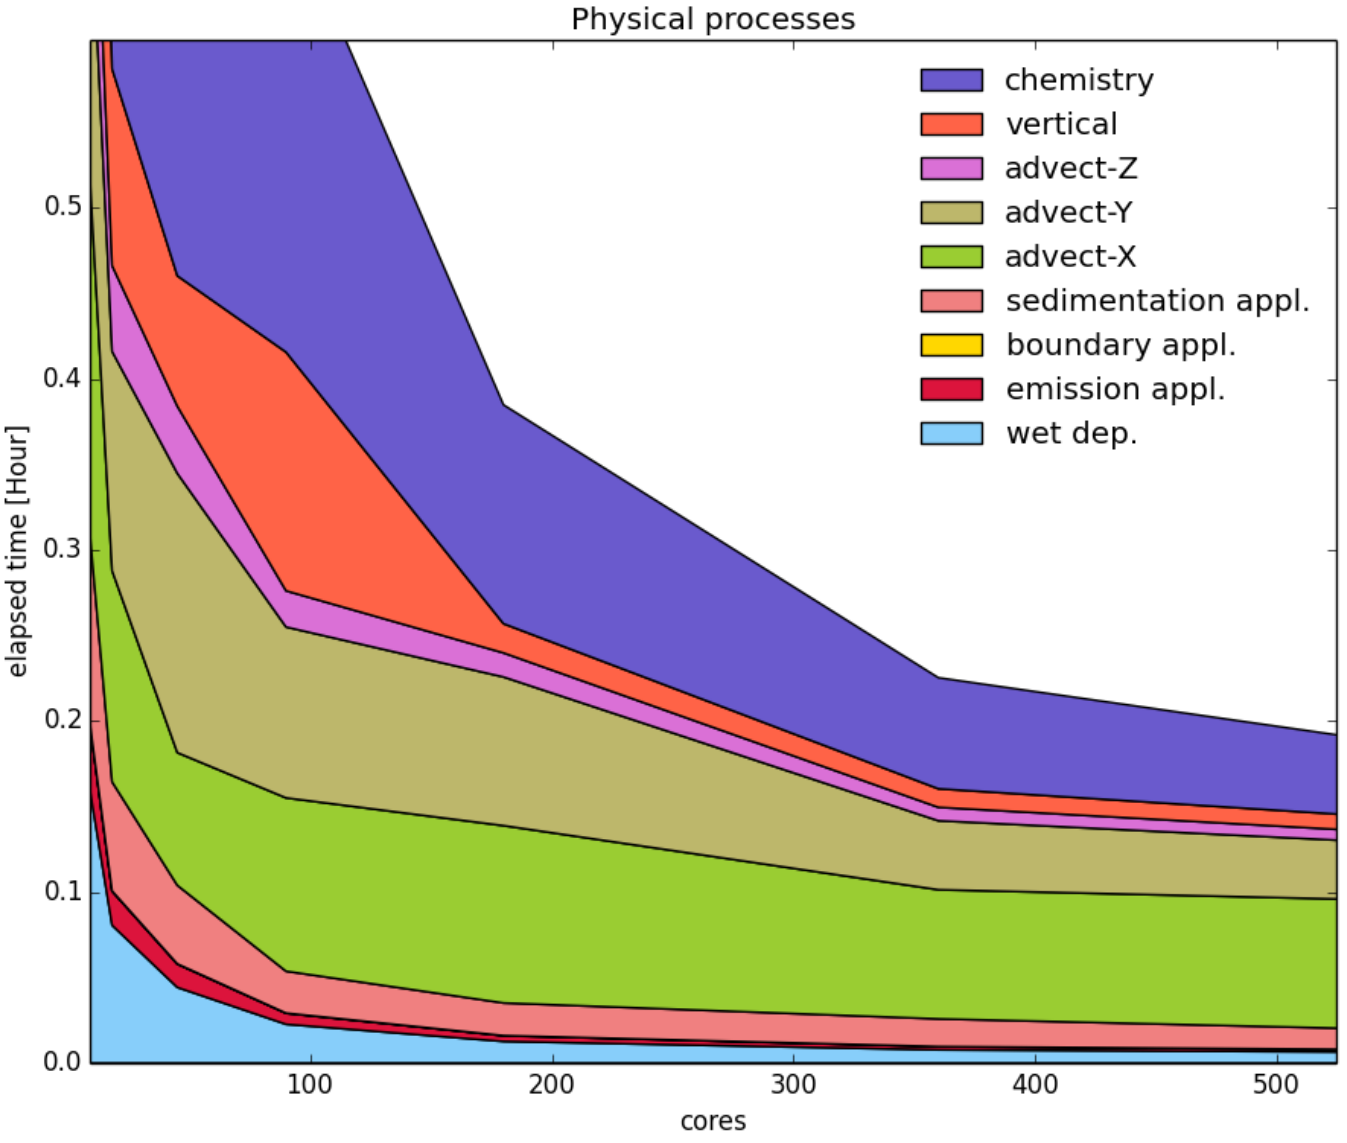

What's up with the convection?

Conclusion

- Good (although simple) solution that scales up to **~500 cores at 3x2**
- Bottleneck is meteo
- Next (1)
  - tests with netcdf w/ parallel IO to alleviate **meteo scattering**
  - open & read each met file **only once** (keep in memory all time steps available) to reduce IO.
- Next (2)
  - tests with 1x1