Convection permitting seasonal latitude-belt simulation using the Weather Research and Forecasting (WRF) model

T. Schwitalla, K. Warrach-Sagi, and V. Wulfmeyer
Institute of Physics and Meteorology
University of Hohenheim

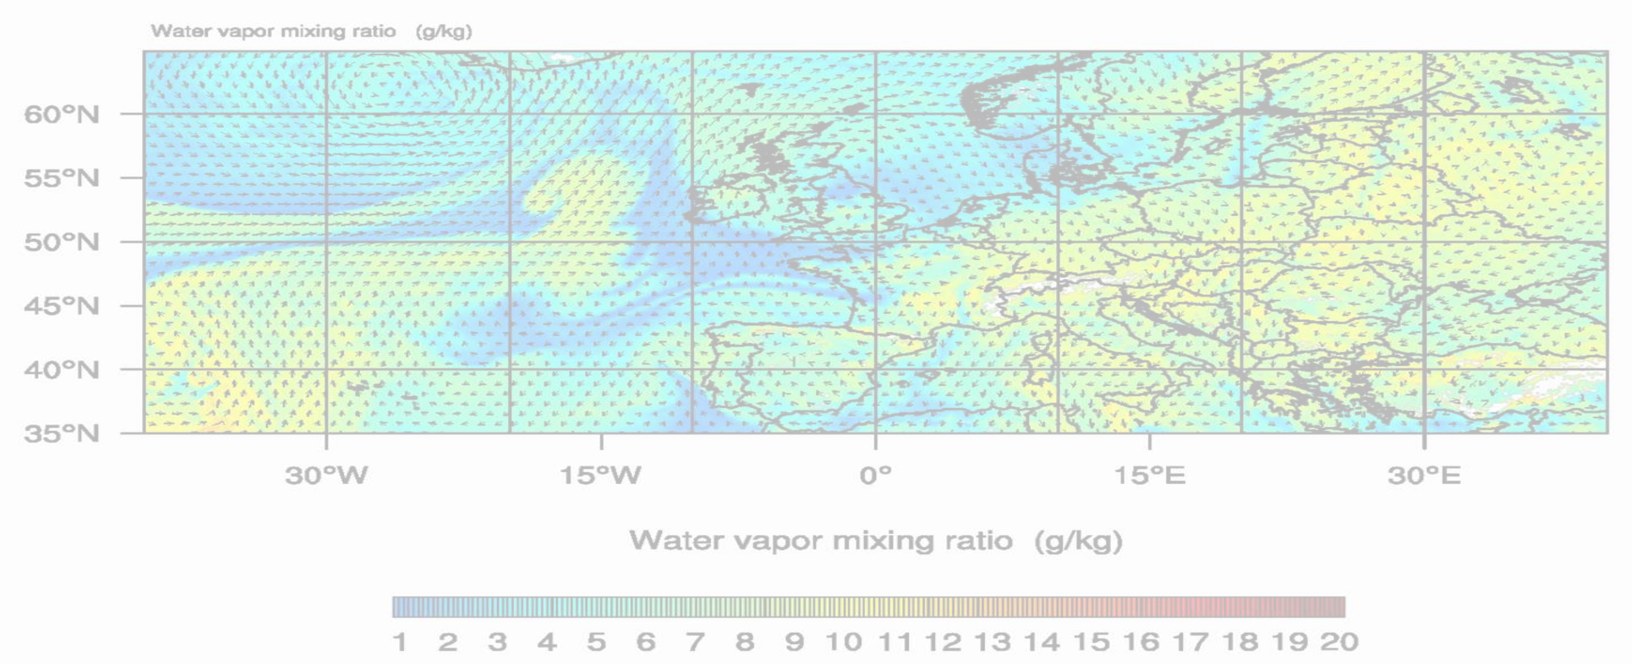

UNIVERSITY OF HOHENHEIM

Background

Usually, limited area models (LAM) are used for current numerical weather prediction

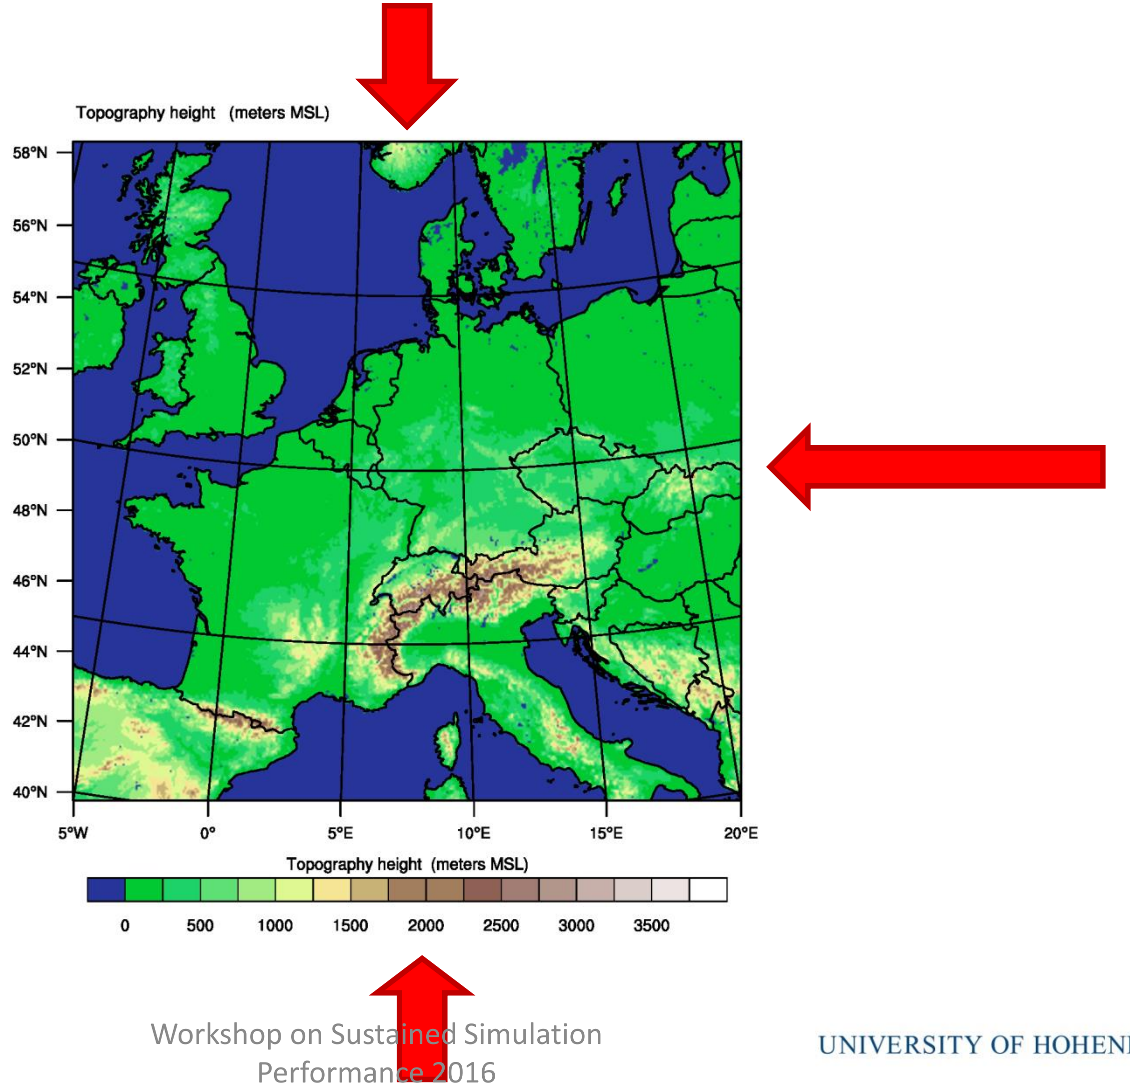

External boundary conditions are not always favorable – they may lead to a distortion due to model imbalance

Background

100km mesh size

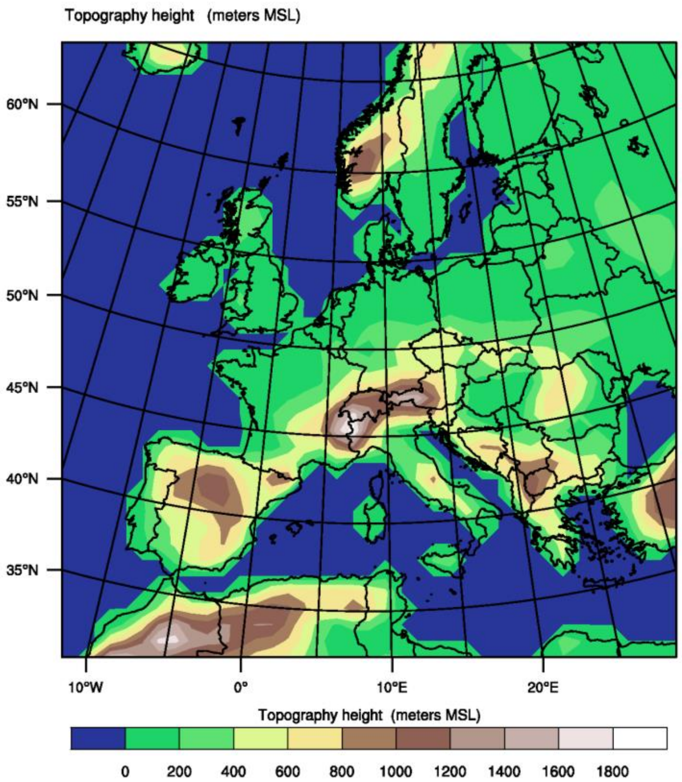

3km mesh size

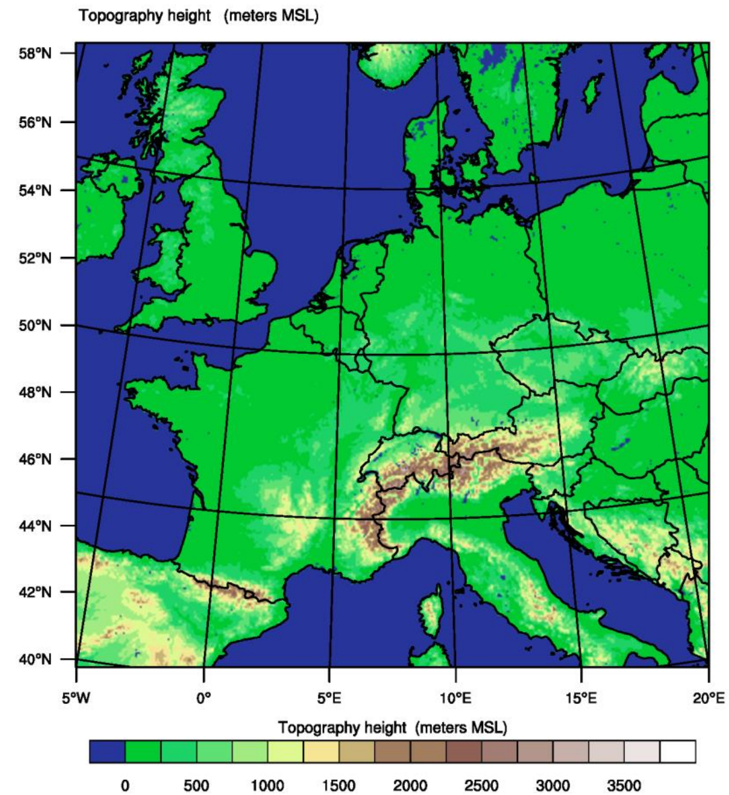

Background

First solution: Run a channel domain with 0.03° resolution

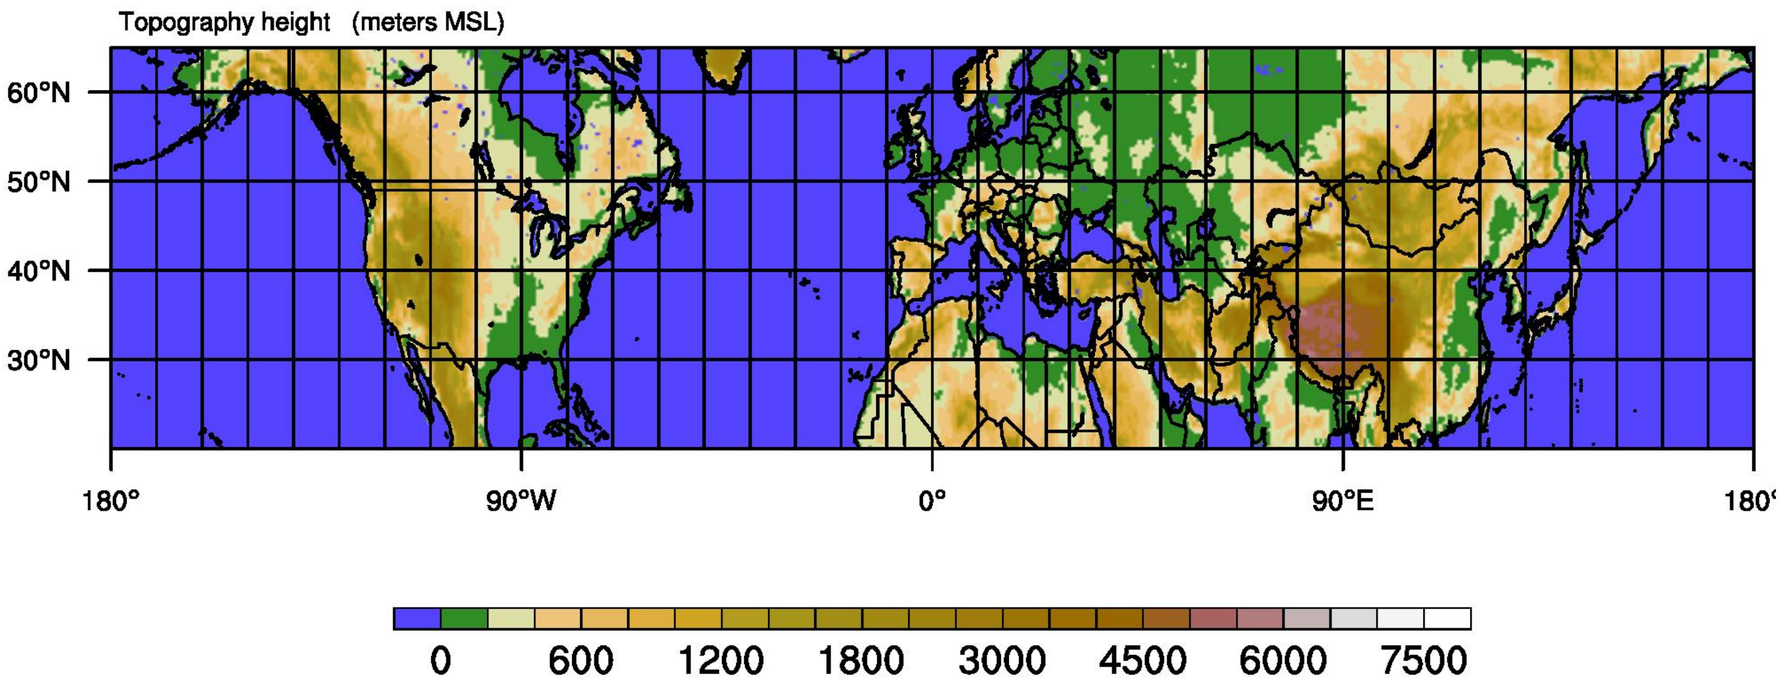

Experimental setup

- CP resolution of 0.03° (3.3 km) with 12000*1500*57 grid cells
- Model top 10 hPa with 14 levels up to 1500 m above ground
- Forcing data from ECMWF analysis every 6 h at the north/south
- 2-moment microphysics including ice, snow and graupel
- YSU Planetary boundary layer parameterization (non-local)
- NOAH Land surface model (4 soil layers, single layer snow model)
- Sea surface temperature data @6 km resolution (OSTIA project of UK Met Office)
- Simulation period July and August 2013
- No data assimilation

Technical aspects

- 3500 nodes of Cray XC40@HLRS (84000 cores in total)
- MPI/OpenMP hybrid mode
- Parallel NetCDF with LUSTRE file striping (set to 96)
- Output frequency was 30min for 3D data
- Output frequency of 15min for additional diagnostic files
- 128 restart files with 440GB each
- Including auxiliary files total data amount ~450TB
- Simulation without I/O takes about 1.5 days
- Total required time was 3.5 days

Results

Accumulated precipitation @12km

CMORPH analysis

WRF 0.12°

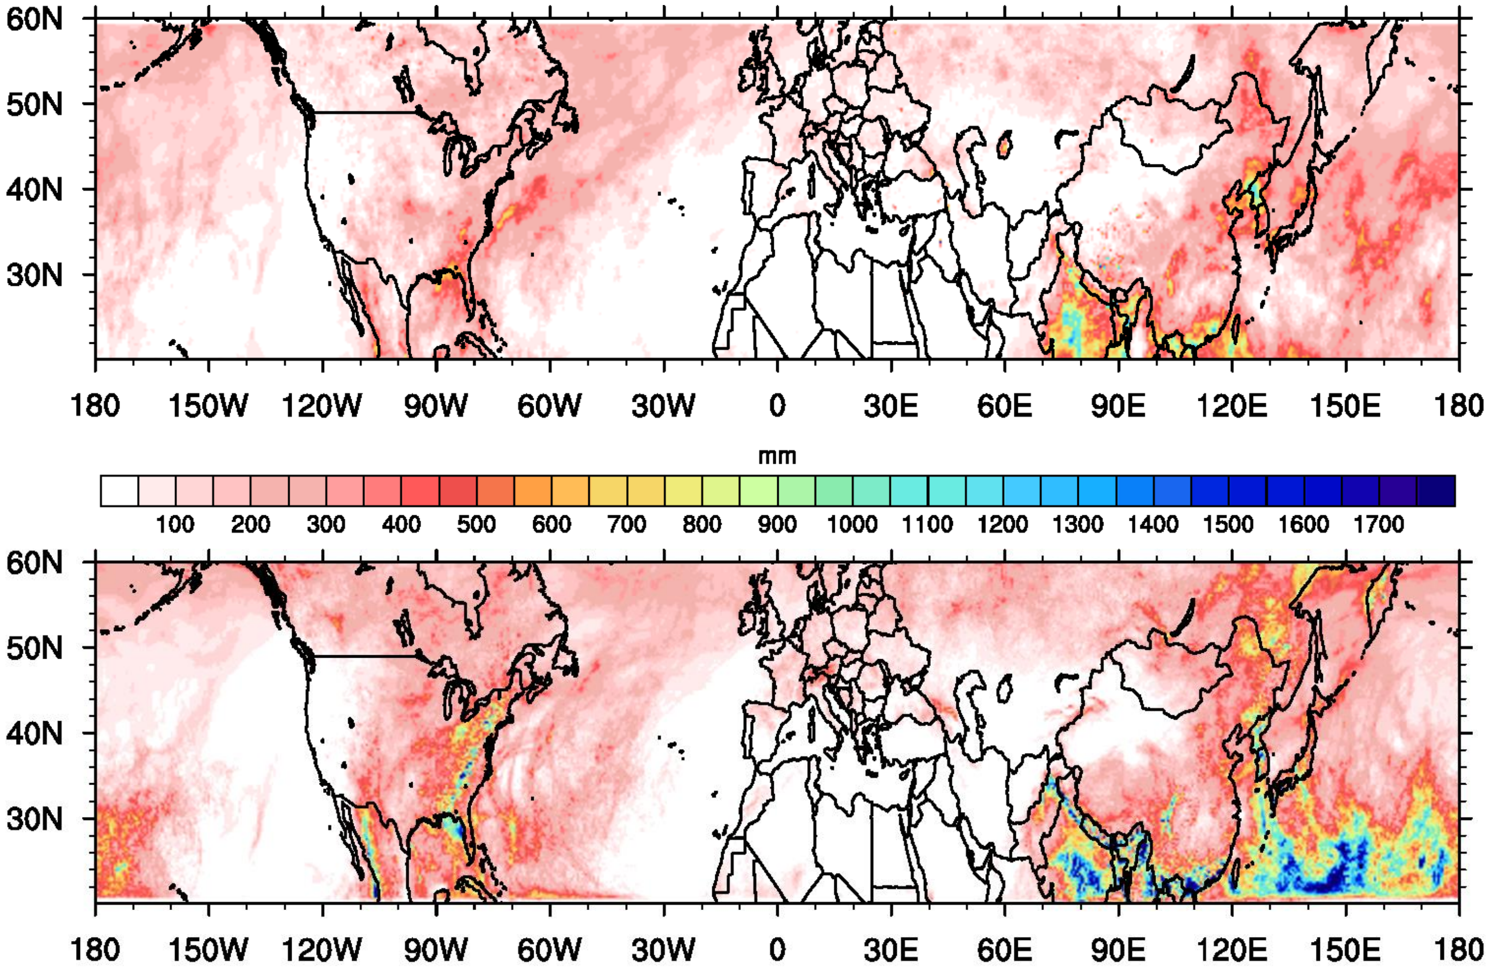

Results

Accumulated precipitation @ 3km

CMORPH analysis

WRF 0.03°

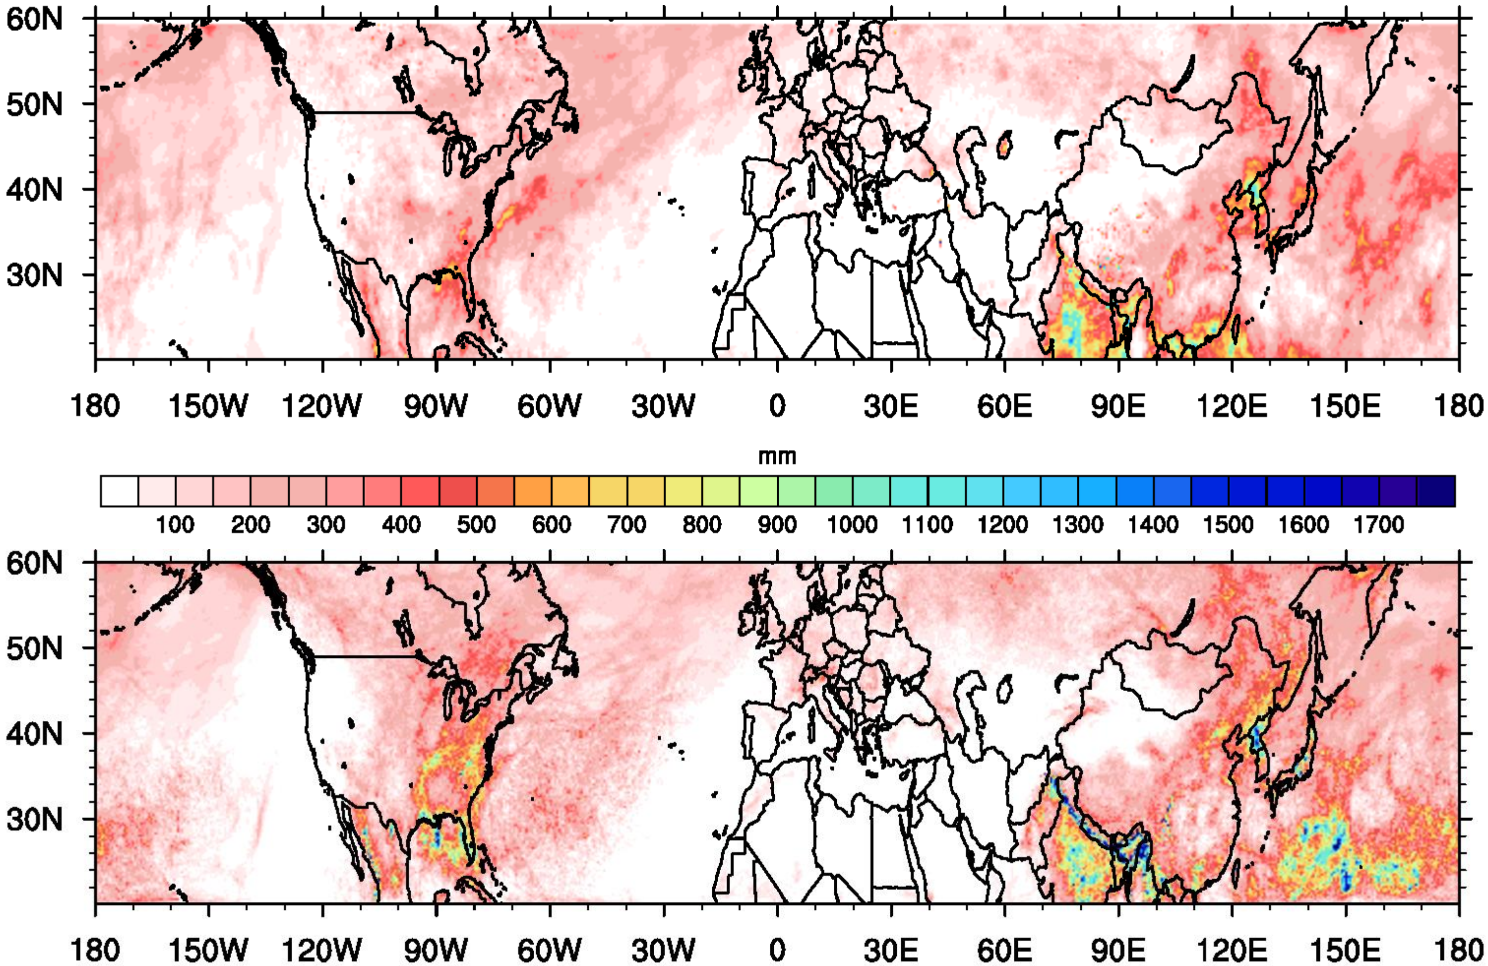

Accumulated precipitation

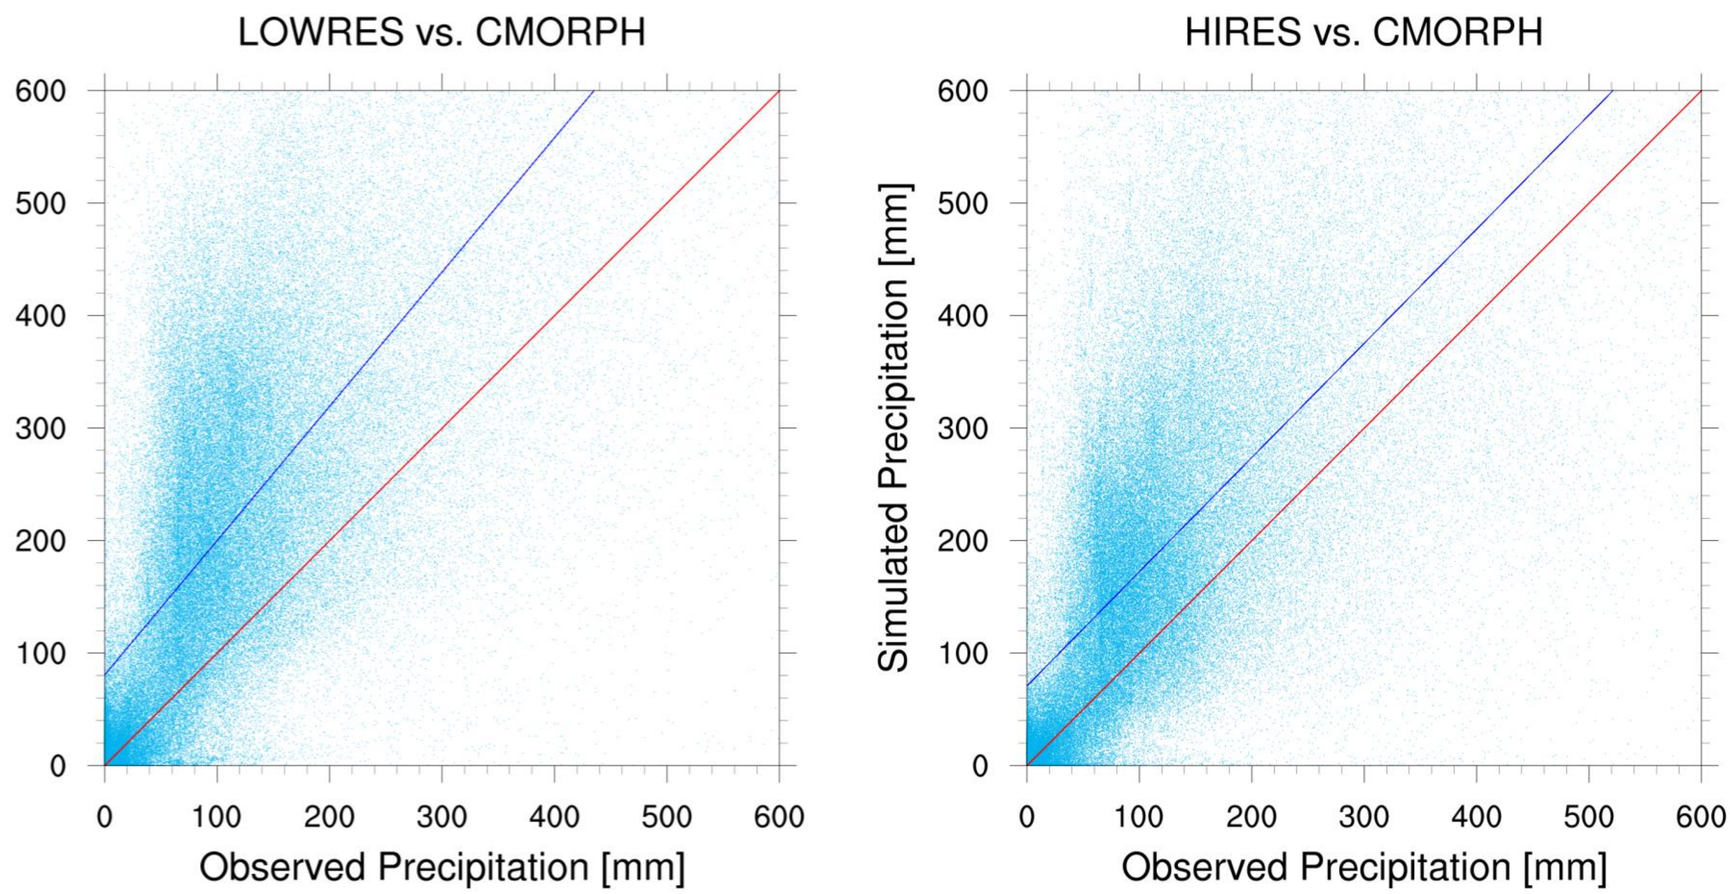

Precipitation PDF (Europe)

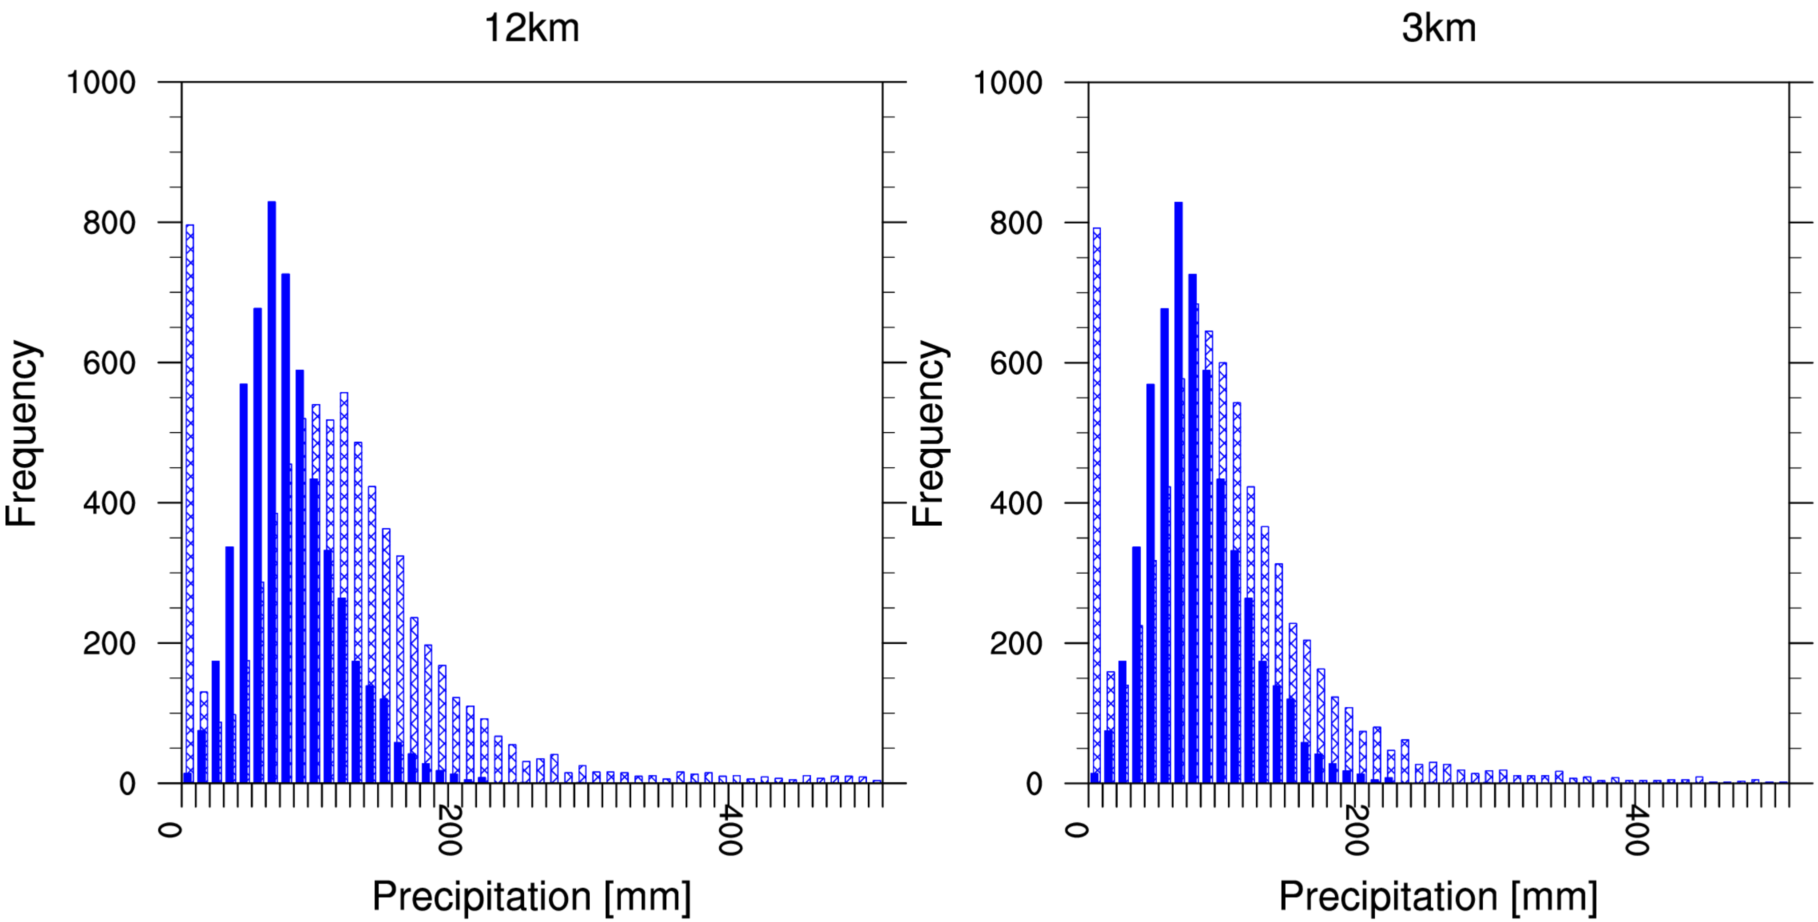

Summary

- First longer-term CP latitude belt simulation of the Northern hemisphere using the WRF model.
- Overestimation of storms in the Pacific Ocean
- Precipitation is overestimated at both resolutions

It‘s the way to go

Moving forward…

Larger system -> larger domain!

Before:

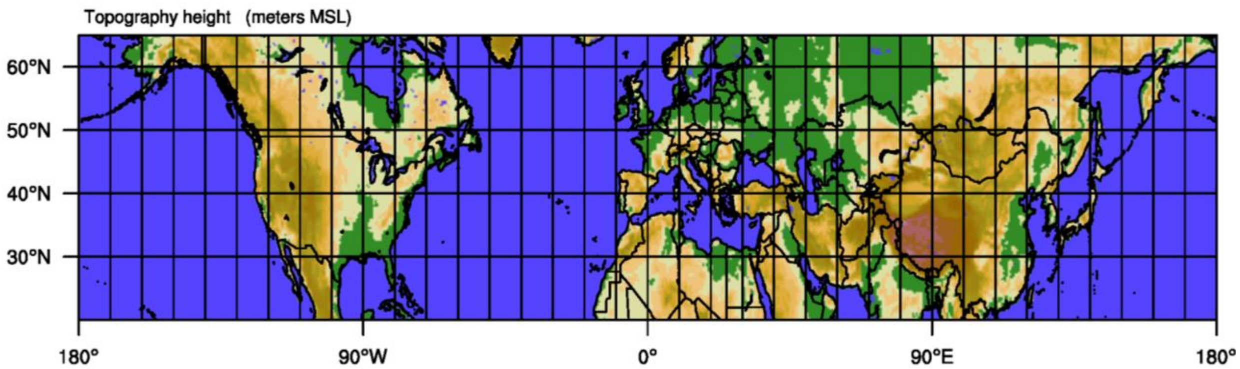

Now:

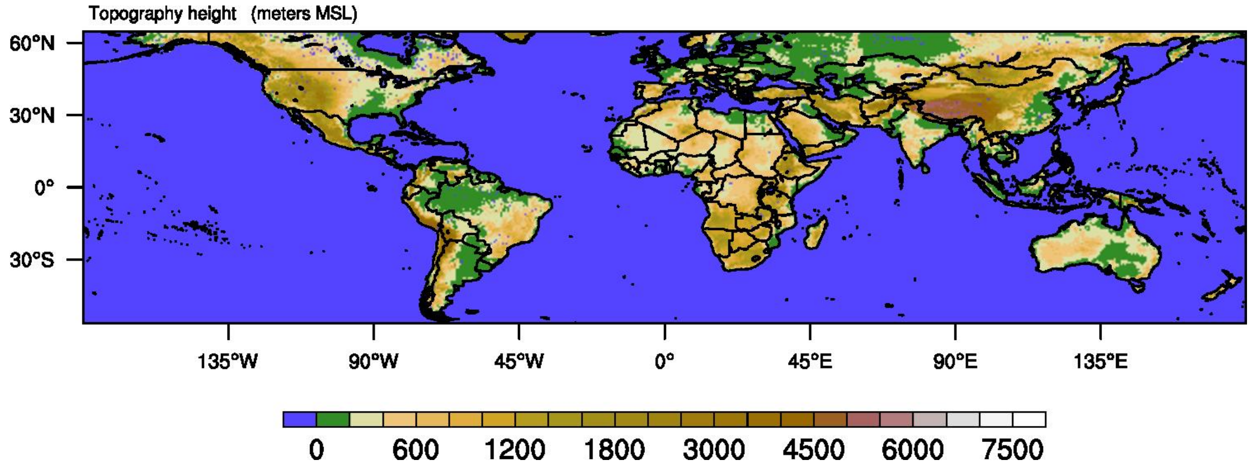

Experimental setup (2)

- CP resolution of 0.03° (3.3 km) with 12000*4060*57 grid cells
- Model top 10 hPa with 14 levels up to 1500 m above ground
- Forcing data from ECMWF analysis every 6 h at the north/south
- 2-moment microphysics including ice, snow and graupel
- YSU Planetary boundary layer parameterization (non-local)
- NOAH-MP Land surface model (4 soil layers, 3-layer snow model)
- SST data @6 km resolution from UK Met Office+ SST data from ECMWF
- Simulation period February to June 2015

Technical aspects (2)

- 4096 nodes of Cray XC40@HLRS (98304cores in total)
- MPI/OpenMP hybrid mode (6 OMP threads/node
- Simulation will run on ws9 (currently 26 large OSTs)
- Parallel NetCDF with LUSTRE file striping (set to 26)
- Output frequency was 6h for 3D data
- Output frequency of 30min for additional diagnostic files
- 20 restart files with 1.2TB each
- Simulation output is expected to be around 190TB
- Required input data is around 100TB
- Total required time is expected to be around 14 days

When you start to prepare your simulation

You wonder why you receive NetCDF error messages:

ERROR: Error in ext_pkg_write_field

This is because of a limitation of serial NetCDF

Solution: Each MPI tasks needs to write its own file

➢ This means 840*4*150 = 504000 files 100MB each.

Of course you wait for the next pitfall

You know that you have to use Parallel NetCDF.....

"One or more variable sizes violate format constraints"

→ One of your arrays is larger than 2^32-4 bytes (or 4GB).

Solution: Move from CDF2 standard to CDF5
standard when using PNetCDF

Let’s see what comes next…..

Maybe you want to check your data:

ncdump -h wrfout_d01_2015-02-01_12_00_00

ncdump: wrfout_d01_2015-02-01_12_00_00:
NetCDF: Unknown file format

And your plot program (NCL) tells you:

Variable: f
Type: file
(0) File Missing Value : -1

Solution: Ask the NCL and Cray developers to built these tools with CDF5 support.

But you never know….

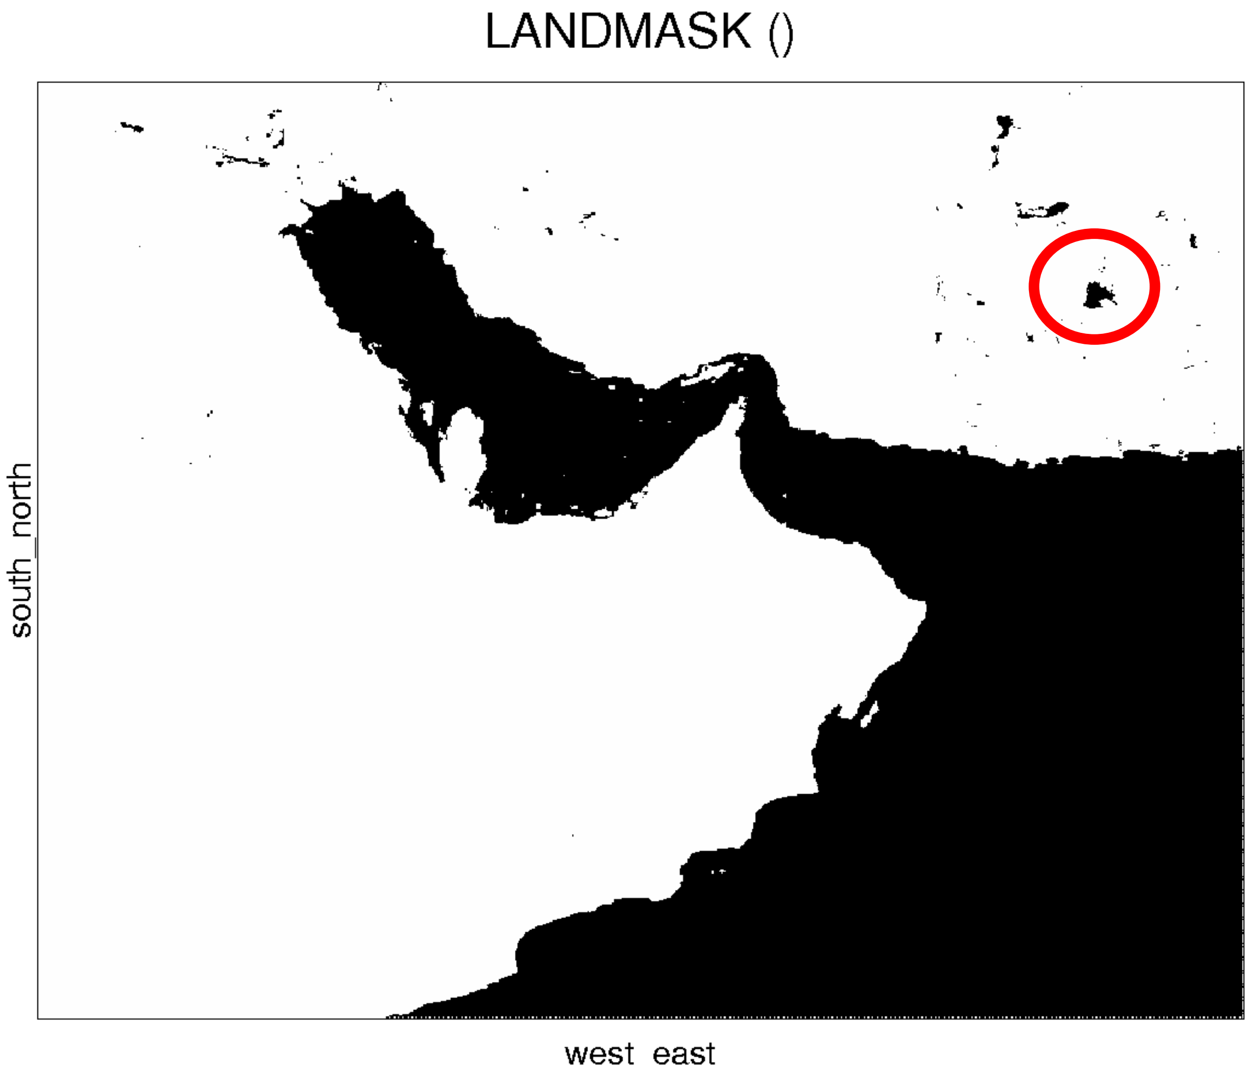

But you never know....

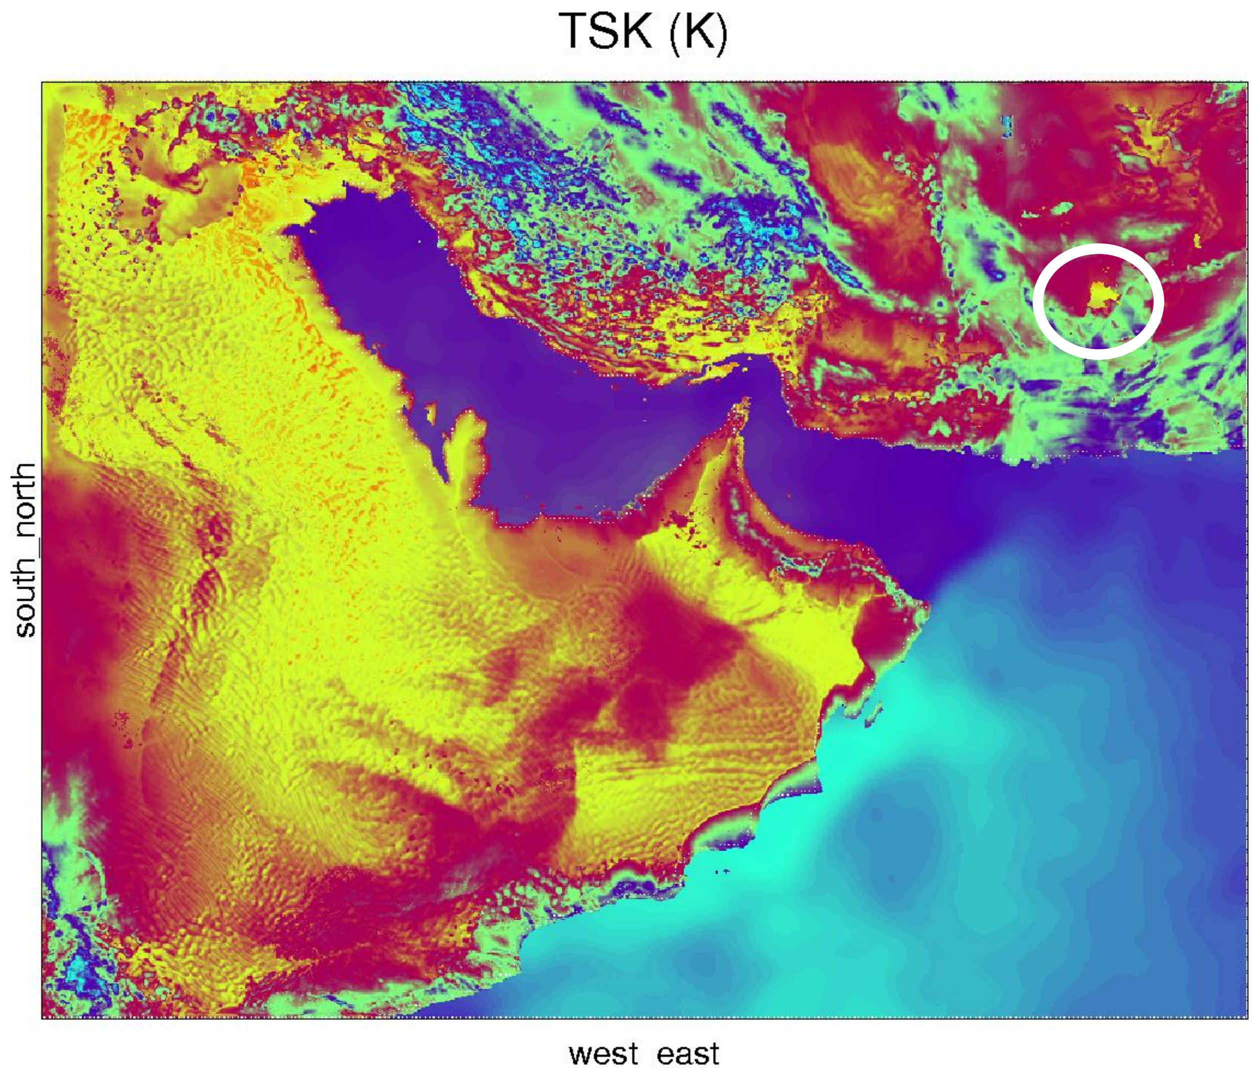

Hopefully the solution is..

- That you use the high-resolution SST data set from UK MetOffice (OSTIA)
- You also use SST data from the operational ECMWF analysis
- Combine both data sets (first check for OSTIA, then for ECMWF)
- Limit SST and skin temperatures over water surfaces to 34°C
- Limit the numerical time step to 10s due to convection and map scale factors

If all the above is the solution…

Then you eventually get your first result of the simulation:

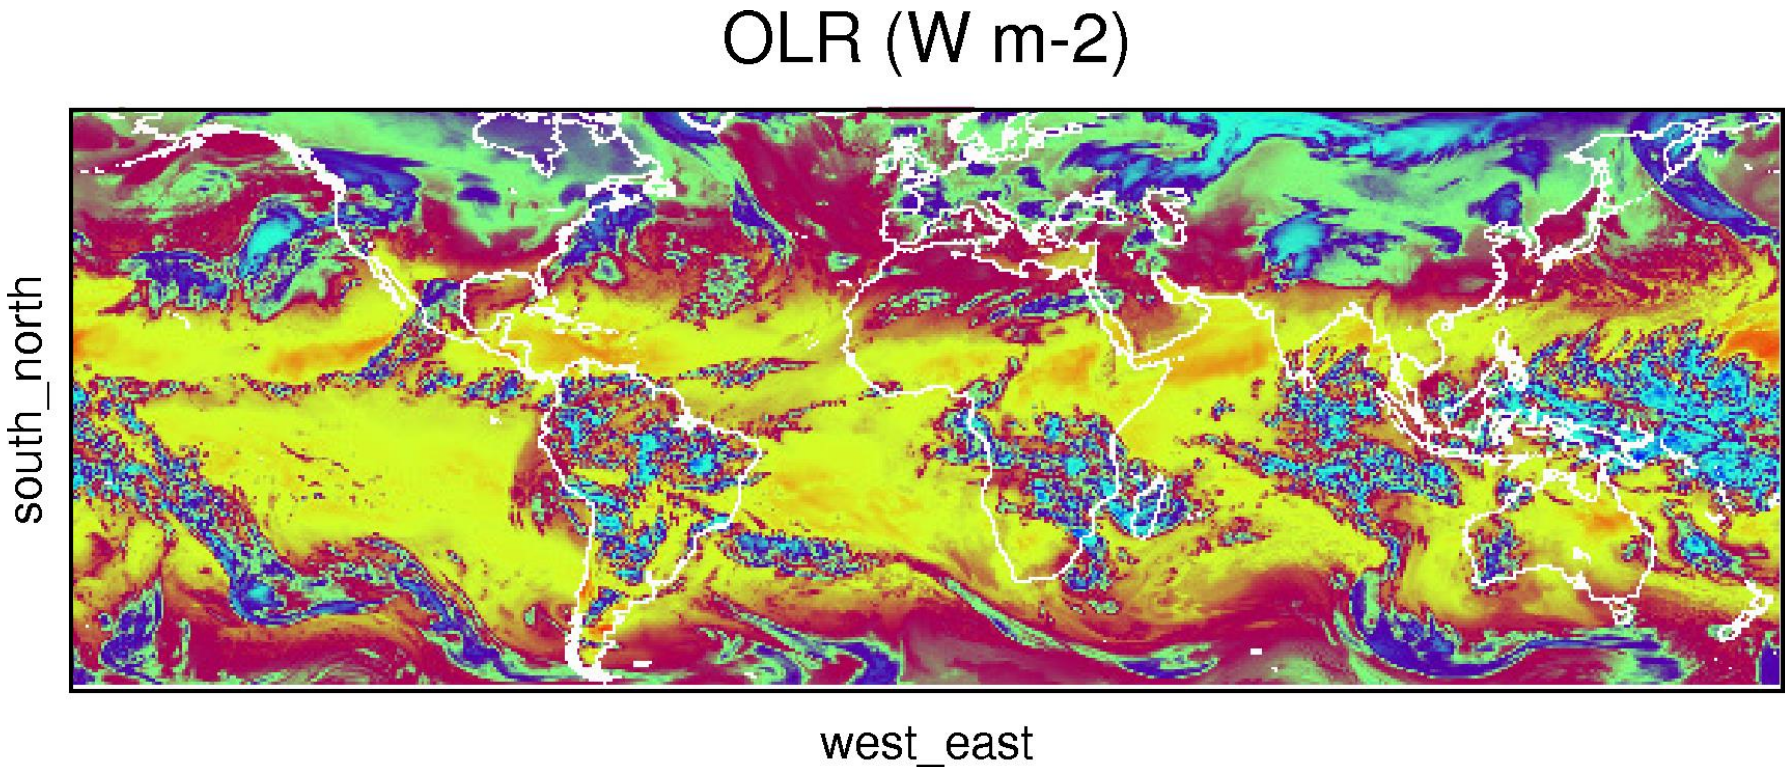

Scaling tests (pure computation,4096 nodes)

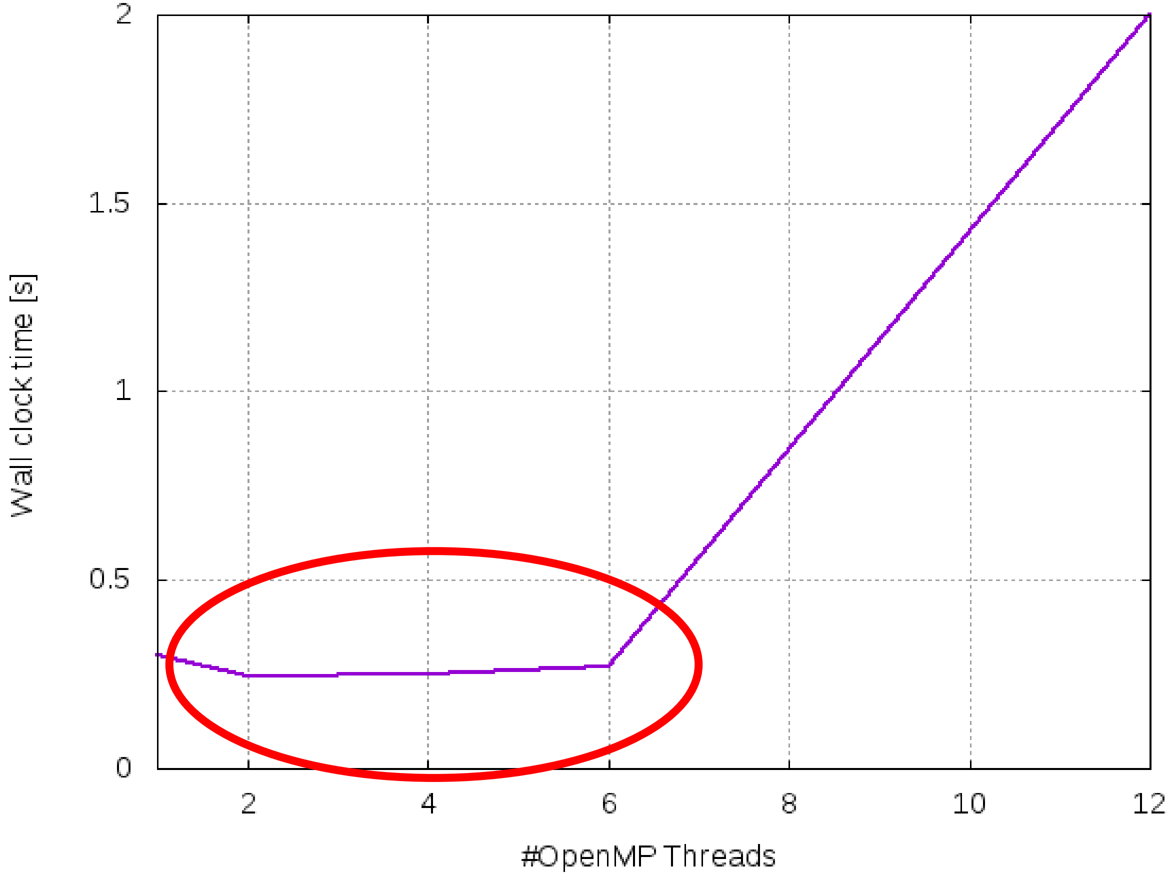

But.....

If you consider Parallel NetCDF I/O:

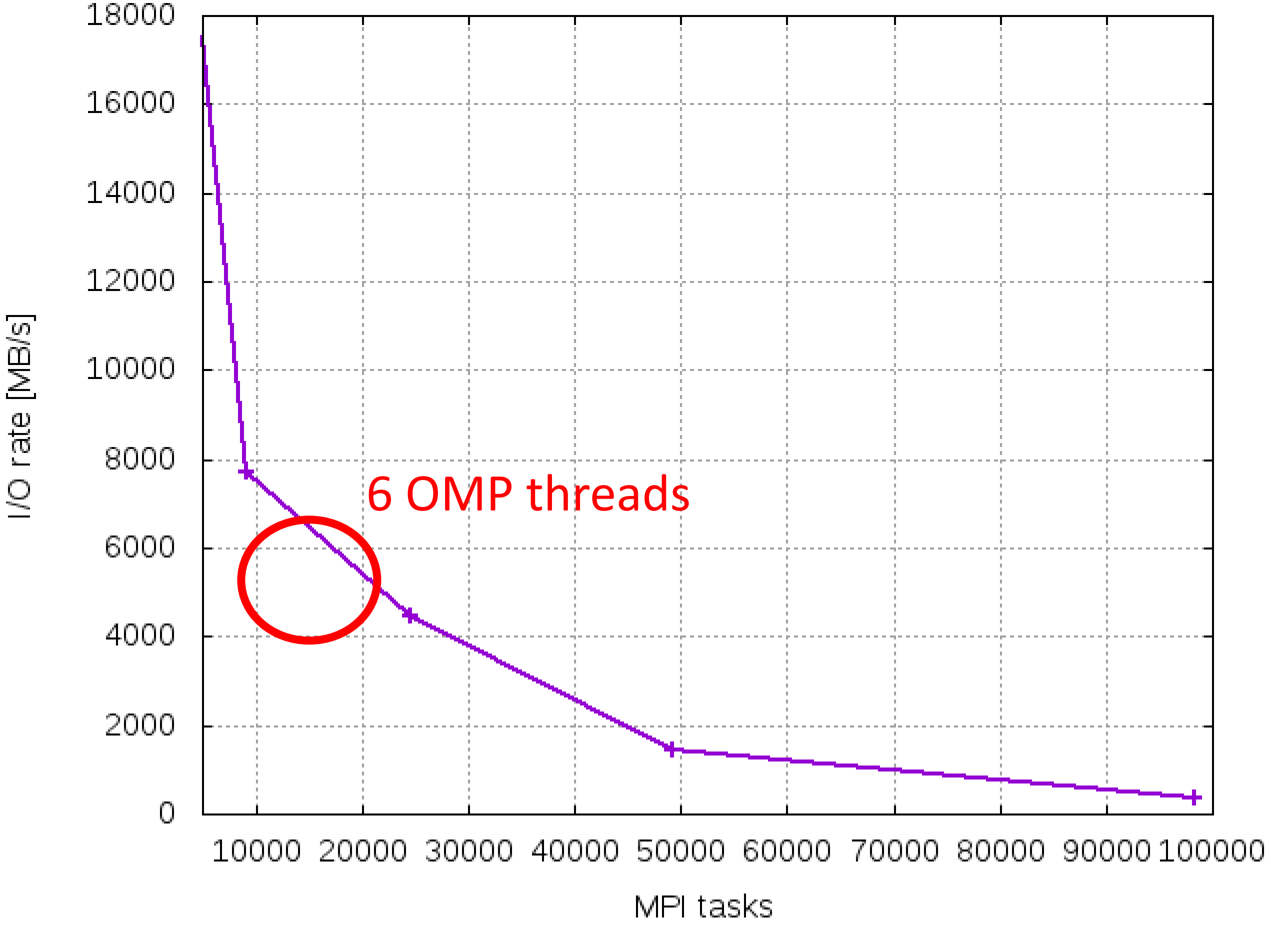

I/O solutions

- Reduce your domain size
- Further reduce number of MPI tasks
- Use NetCDF4 compression
- Increase number of OST's
- Try SIONLib from Jülich and LRZ
- Try Cray Data Warp

If you are too optimistic…

```
#0  0x00002aaaad8d7875 in raise () from /lib64/libc.so.6
#1  0x00002aaaad8d8e51 in abort () from /lib64/libc.so.6
#2  0x00000000224de9b5 in for__issue_diagnostic ()
#3  0x00000000224e55ea in for__signal_handler ()
#4  <signal handler called>
#5  0x0000000022135f3f in module_sf_sfclayrev_mp_psim_stable_ ()
#6  0x0000000022133195 in module_sf_sfclayrev_mp_sfclayrev1d_ ()
#7  0x000000002213103b in module_sf_sfclayrev_mp_sfclayrev_ ()
#8  0x0000000021a2a729 in module_surface_driver_mp_surface_driver_ ()
#9  0x00002aaaad604cb3 in __kmp_invoke_microtask ()
   from
/sw/hazelhen/hlrs/compiler/intel/Compiler/16.0.3.210/compilers_and_libraries_2016.3.210/linux/compiler/lib/intel64_lin/libiomp5.so
#10 0x00002aaaad5d3437 in __kmp_invoke_task_func (gtid=-77391588) at ../../src/kmp_runtime.c:7058
#11 0x00002aaaad5d460b in __kmp_fork_call (loc=0x7ffffb63191c, gtid=-77398260, call_context=(unknown: 1185205632),
argc=174930668, microtask=0x7ffffb62fd8c,
    invoker=0x7ffffb62fc0c, ap=0x7ffffb6429b0) at ../../src/kmp_runtime.c:2397
#12 0x00002aaaad5ad518 in __kmpc_fork_call (loc=0x7ffffb63191c, argc=-77398260, microtask=0x46a4cd80
<module_sf_sfclayrev_mp_psim_stab_>)
    at ../../src/kmp_csupport.c:339
#13 0x0000000021a32378 in module_surface_driver_mp_surface_driver_ ()
#14 0x00000000213a80bb in module_first_rk_step_part1_mp_first_rk_step_part1_ ()
#15 0x0000000020dc9ae1 in solve_em_ ()
#16 0x0000000020c7d91a in solve_interface_ ()
#17 0x000000002016d001 in module_integrate_mp_integrate_ ()
#18 0x0000000020085fe7 in module_wrf_top_mp_wrf_run_ ()
#19 0x0000000020085bff in MAIN__ ()
#20 0x0000000020085b7e in main ()
```