www.bsc.es

EXCELENCIA SEVERO OCHOA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

2nd BSC Severo Ochoa Retreat

Earth Sciences Applications and Collaborations with CS

George Markomanolis

ES: Alba Badia, Jose M. Baldasano, Sara Basart, Enza Di Tomaso, Santiago Gassó, Oriol Jorba, Francesc Martinez, Vin cenzo Obiso, Kim Serradell,

CS: Judit Gimenez, Julian Morillo, Enric Tejedor

Barcelona, 23 January 2014

Outline

- Introduction
  - SO Applications
  - NMMB/BSC-CTM
  - Data Assimilation
- Collaboration projects with CS
- COMPSs on Fix-Vrbl tasks
- Performance NMMB/BSC-CTM
- Future work

Severo-Ochoa Earth Sciences Application

- Development of a Unified Meteorology/Air Quality/Climate model
  - Towards a global high-resolution system for global to local assessments

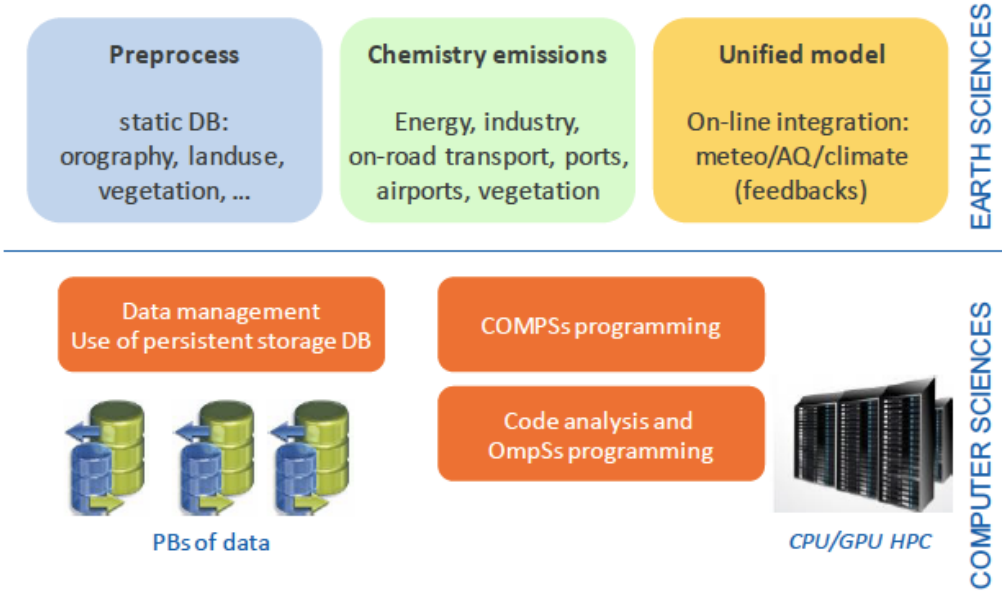

- Extending NMMB/BSC-CTM from coarse regional scales to global high-resolution configurations
- Coupling with a Data Assimilation System for Aerosols

- International collaborations:

Meteorology

National Centers for Environmental Predictions

Climate
Global aerosols

Goddard Institute Space Studies

Air Quality

Uni. of California Irvine

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

NMMB/BSC-Chemical Transport Model (Overview)

NMMB/ BSC-CTM

Nonhydrostatic Multiscale Model on the B-grid (NMMB)
meteo variables/parameters

→ *Janjic and Gall (NCAR/TN 2012)*
→ *Janjic and Vasic (EGU2012)*
→ *Janjic et al. (MWR 2011)*
→ *(...)*

BSC Chemical Transport Model
(gas/aerosol variables: mass mixing ratios)

- *GAS-PHASE CHEM (52 species)* → *Jorba et al. (JGR 2012)*
- *DUST (8 bins)* → *Pérez et al. (ACP 2011)* → *Haustein et al. (ACP 2012)*
- *SEA-SALT* → *Spada et al. (ACP 2013)*

+

- **Global aerosols**
- **Secondary aerosols**
- **Ocean model**

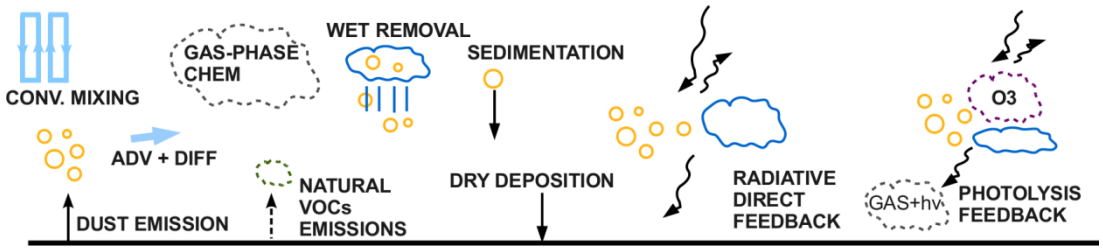

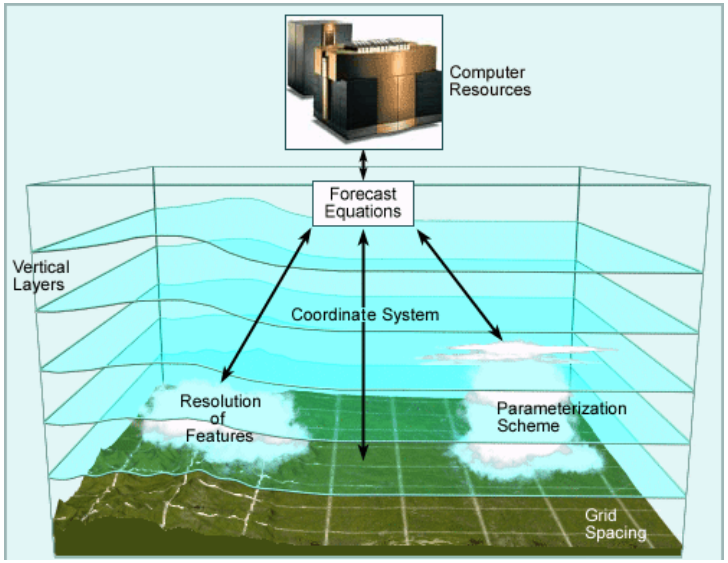

Data Assimilation – Motivations

Atmospheric models are far from being perfect

Dust optical depth: 2014 01 06 h+00

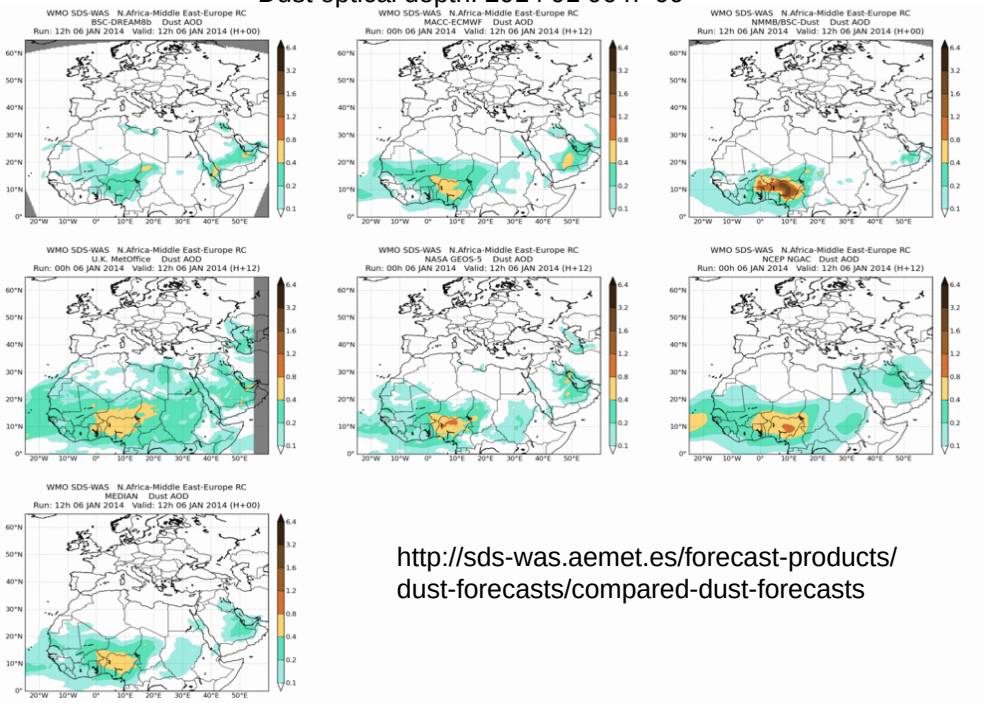

http://sds-was.aemet.es/forecast-products/dust-forecasts/compared-dust-forecasts

A considerable amount of accurate earth observations is available

http://www.wmo.int/pages/prog/gcos/

Data assimilation 'optimally' combines **models** and **observations**

Data Assimilation – Workflow

t = 1

ensemble background

(short-term forecast)

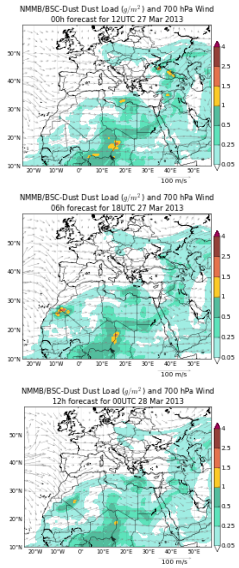

Kalman filter*

ensemble analysis

(initial conditions)

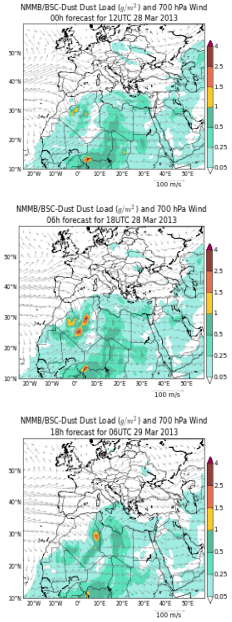

model

t = 2

ensemble background

(short-term forecast)

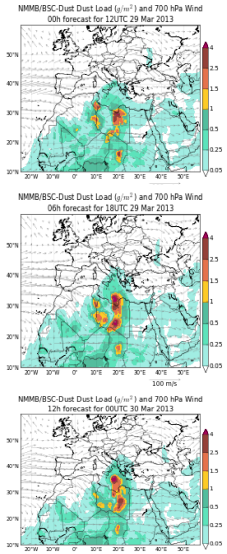

Kalman filter*

...

Observations

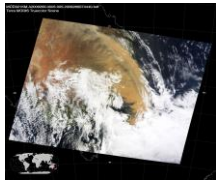

http://modis-atmos.gsfc.nasa.gov/

mean analysis

(initial conditions)

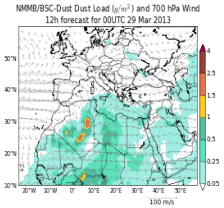

model

long-term forecast

Observations

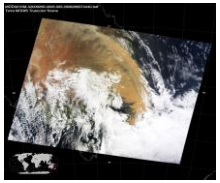

http://modis-atmos.gsfc.nasa.gov/

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

* In collaboration with N. Schutgens (Uni. Oxford, UK)

Data Assimilation – Workflow

- BASH script starts the submission of the assimilation job
  - We want all the ensembles to be executed in parallel
  - We have 40 ensembles, we provide 20 cores for each execution and one ensemble for long-forecast. We should need totally 82 nodes (1,312 exclusive cores)

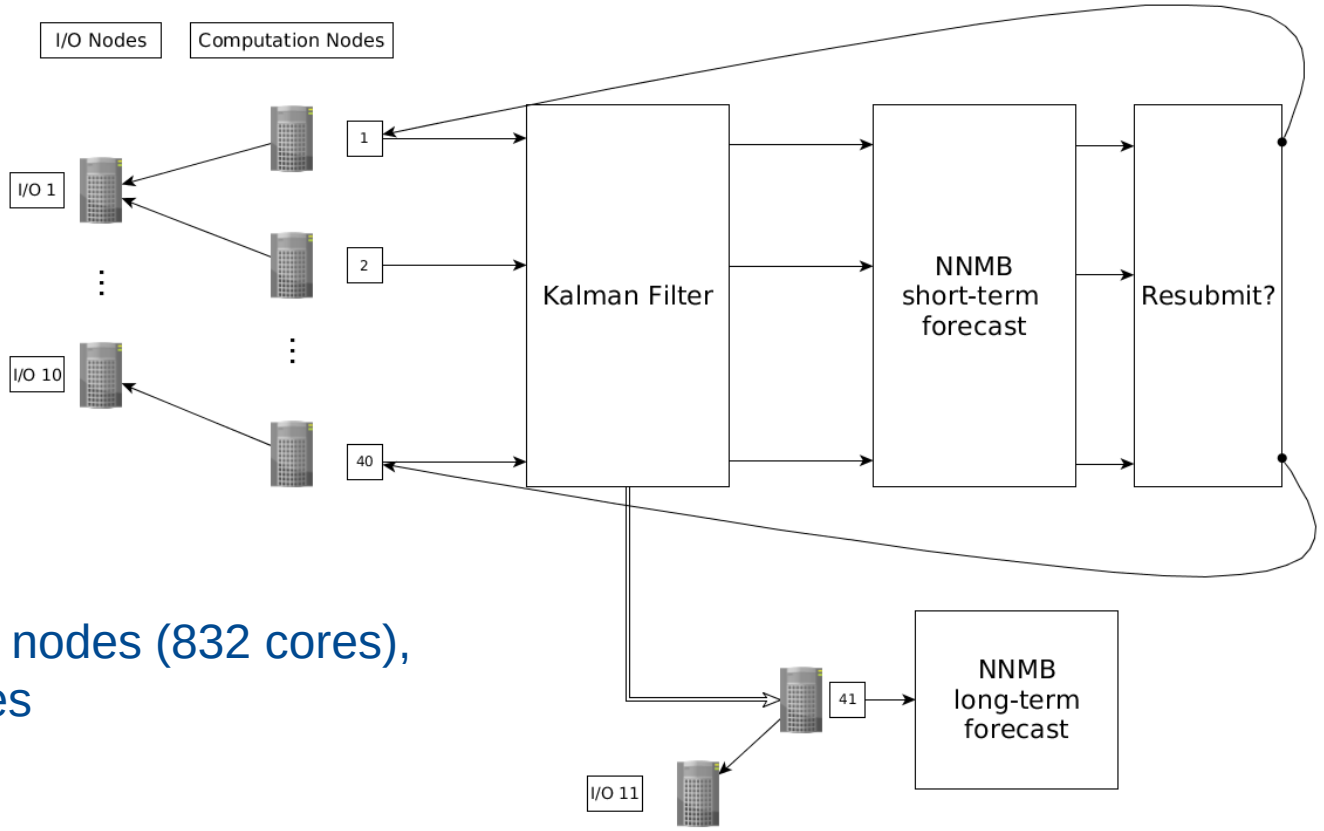

- Now, we need 52 nodes (832 cores), ~36% less resources

http://modis-atmos.gsfc.nasa.gov/

Collaborations with Computer Sciences

- Use of CompSS:
  - Porting Model pre-processing to CompSs
  - Test of the application at low and high resolution
  - Limitations on the approach when dealing with high-resolution configurations

- Exploring Self-contained Objects technology:
  - Porting of HERMESv2.0 emission model from C++ to Java
  - Exploring I/O management in ES application

- Performance analysis:
  - Definition of model configurations
  - Initial extraction of traces – problems when using instrumented functions

- Use of OmpSs:
  - Compilation of NMMB/BSC-CTM with mercurium
  - Preliminary runs with multi-threads
  - Expected significant impact in tracers transport and chemistry

1. Use of COMPSs

Execution diagram: Focus on the Preprocessor

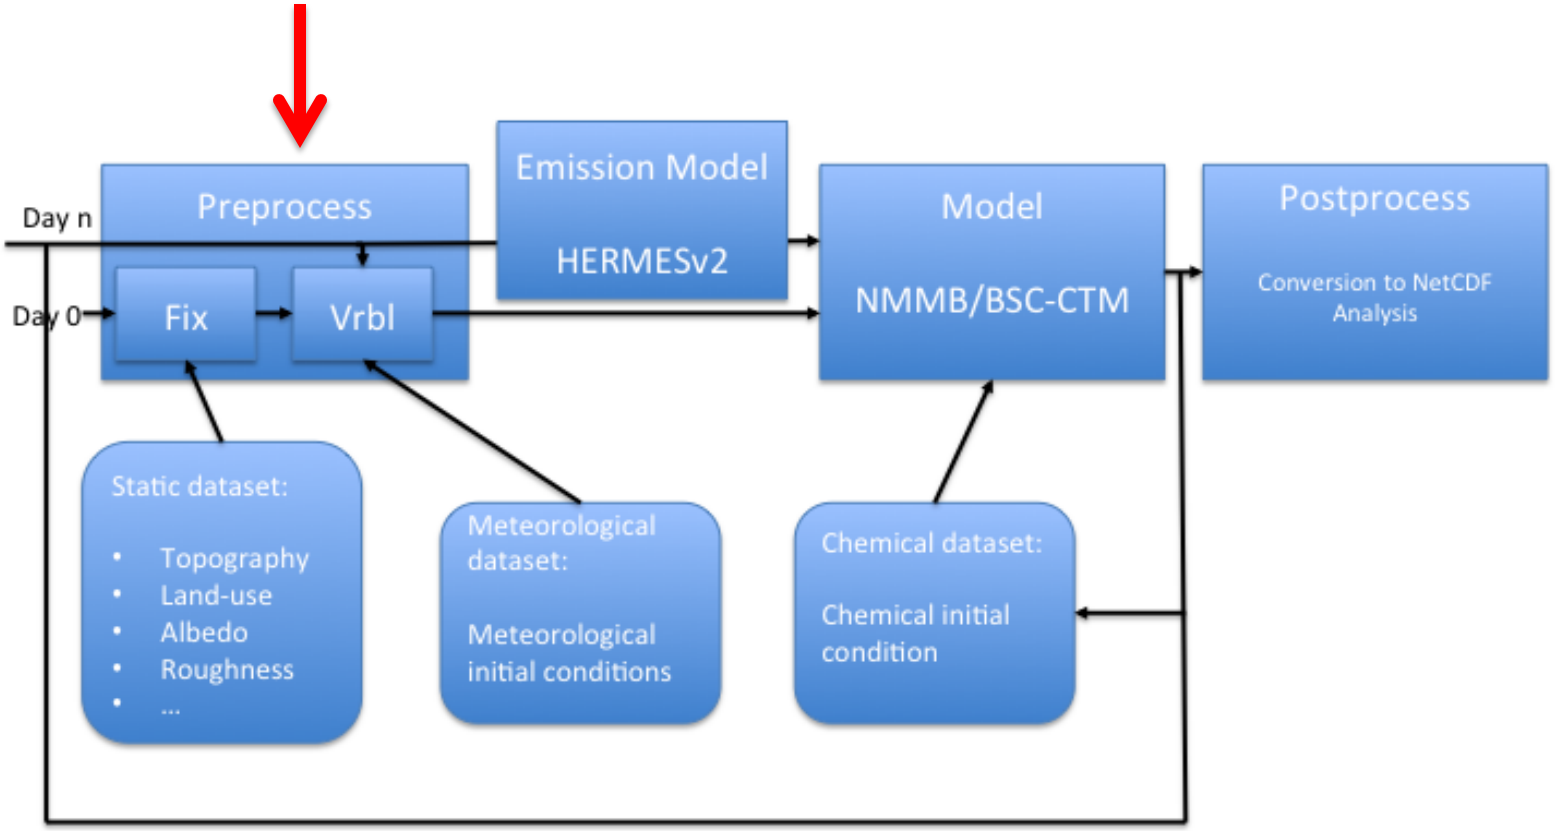

Tested in two cases: Global domain 1ºx1.4º resolution
Global domain 12km x 12km

Preprocessor Step 1 – COMPSs

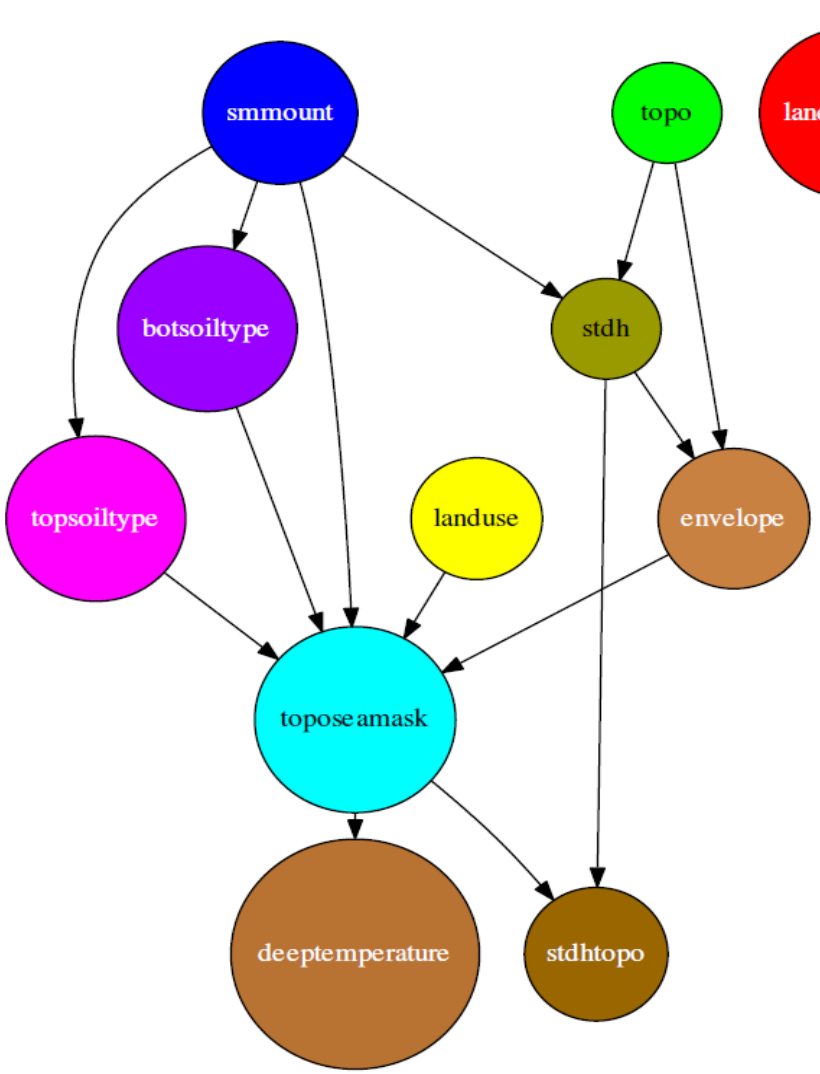

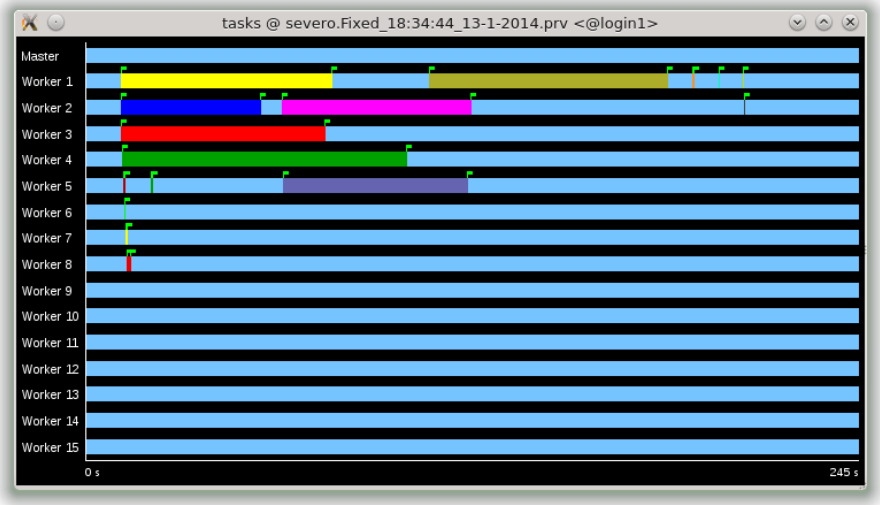

We used 5 cores and achieved a speedup of 2.7 times!

Preprocessor Step 2 – COMPSs

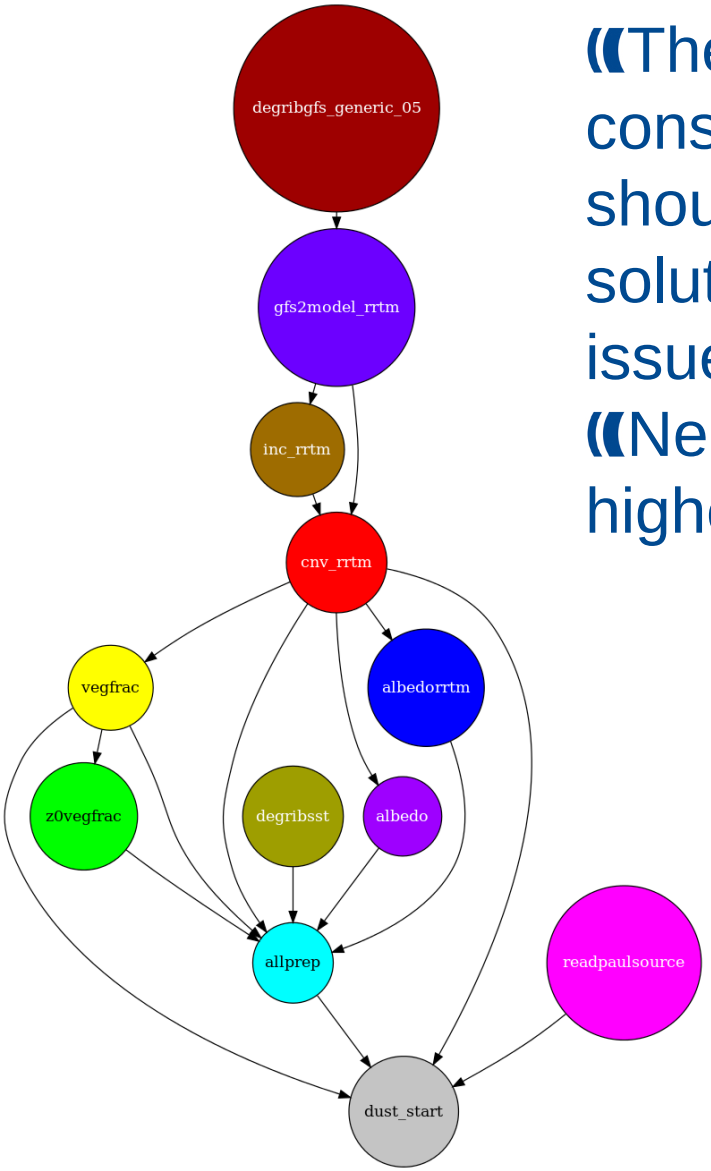

- The serial part *allprep* consumes a lot of time, we should investigate a hybrid solution because of memory issues
- Need to be improved for higher resolution forecasts

Test on a bigger case

- We applied this method to generate 12km global resolution input files (more than 6GB output files)

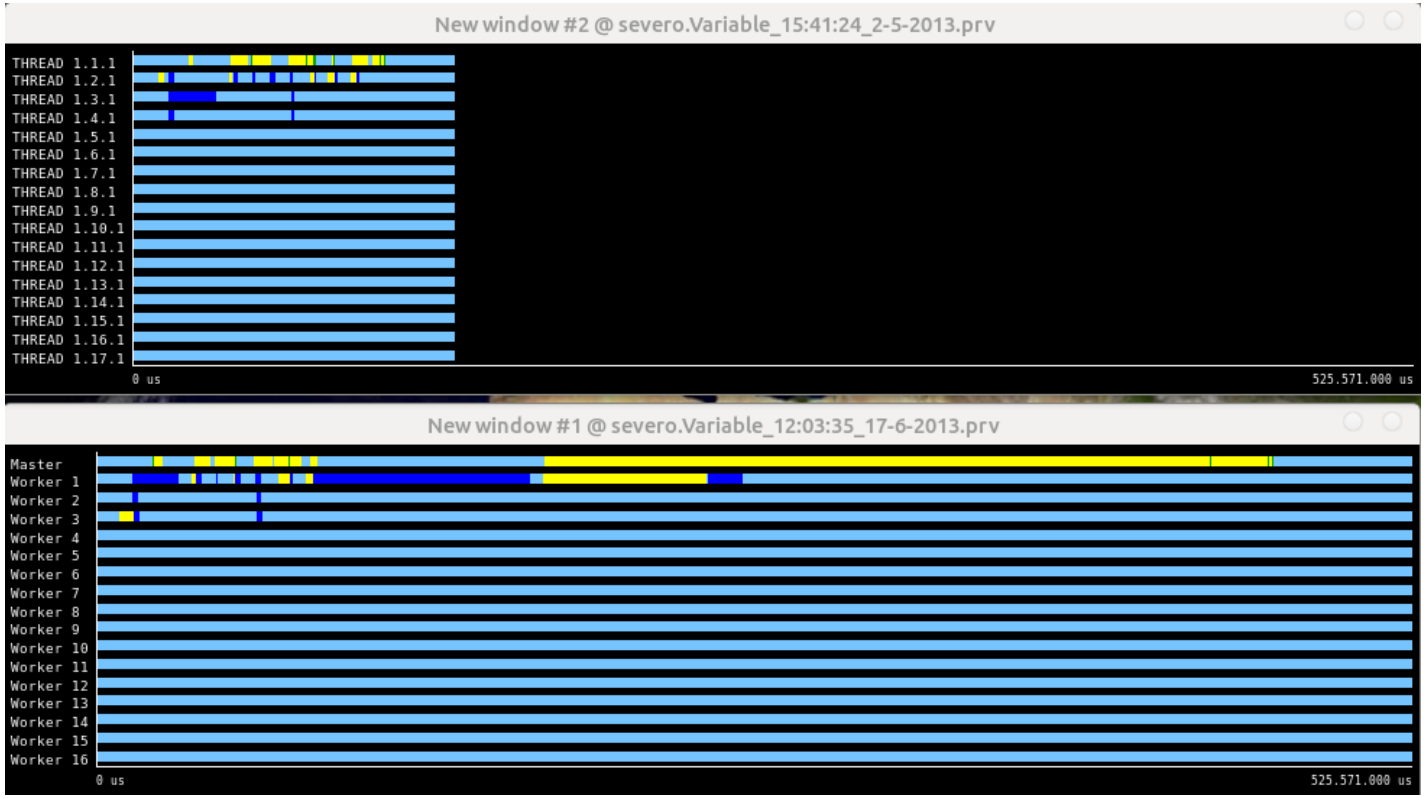

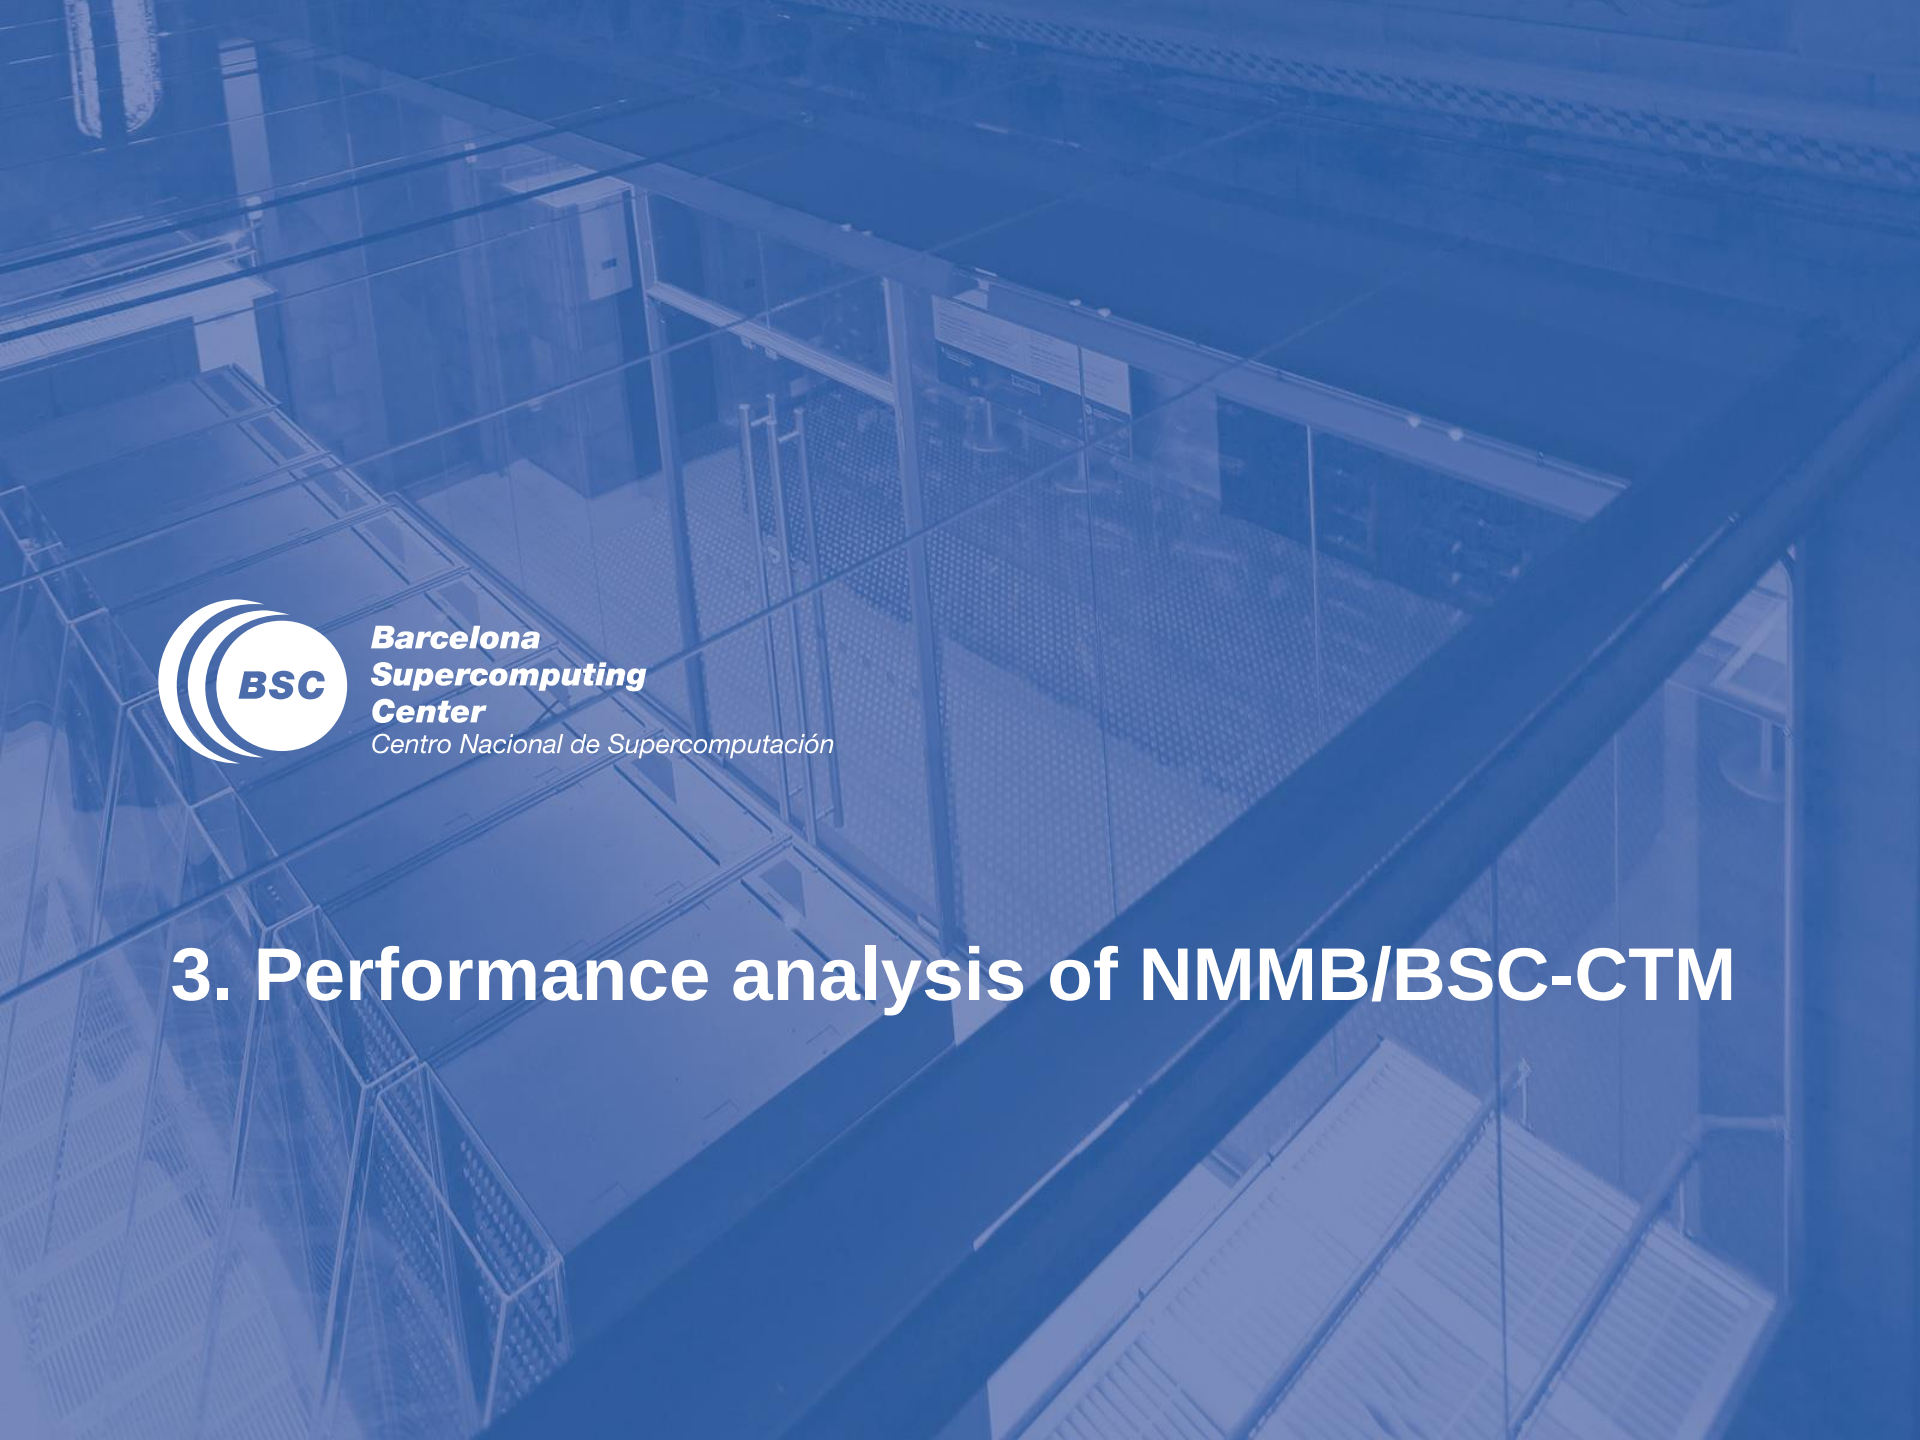

3. Performance analysis of NMMB/BSC-CTM

Execution diagram – Focus on the Model

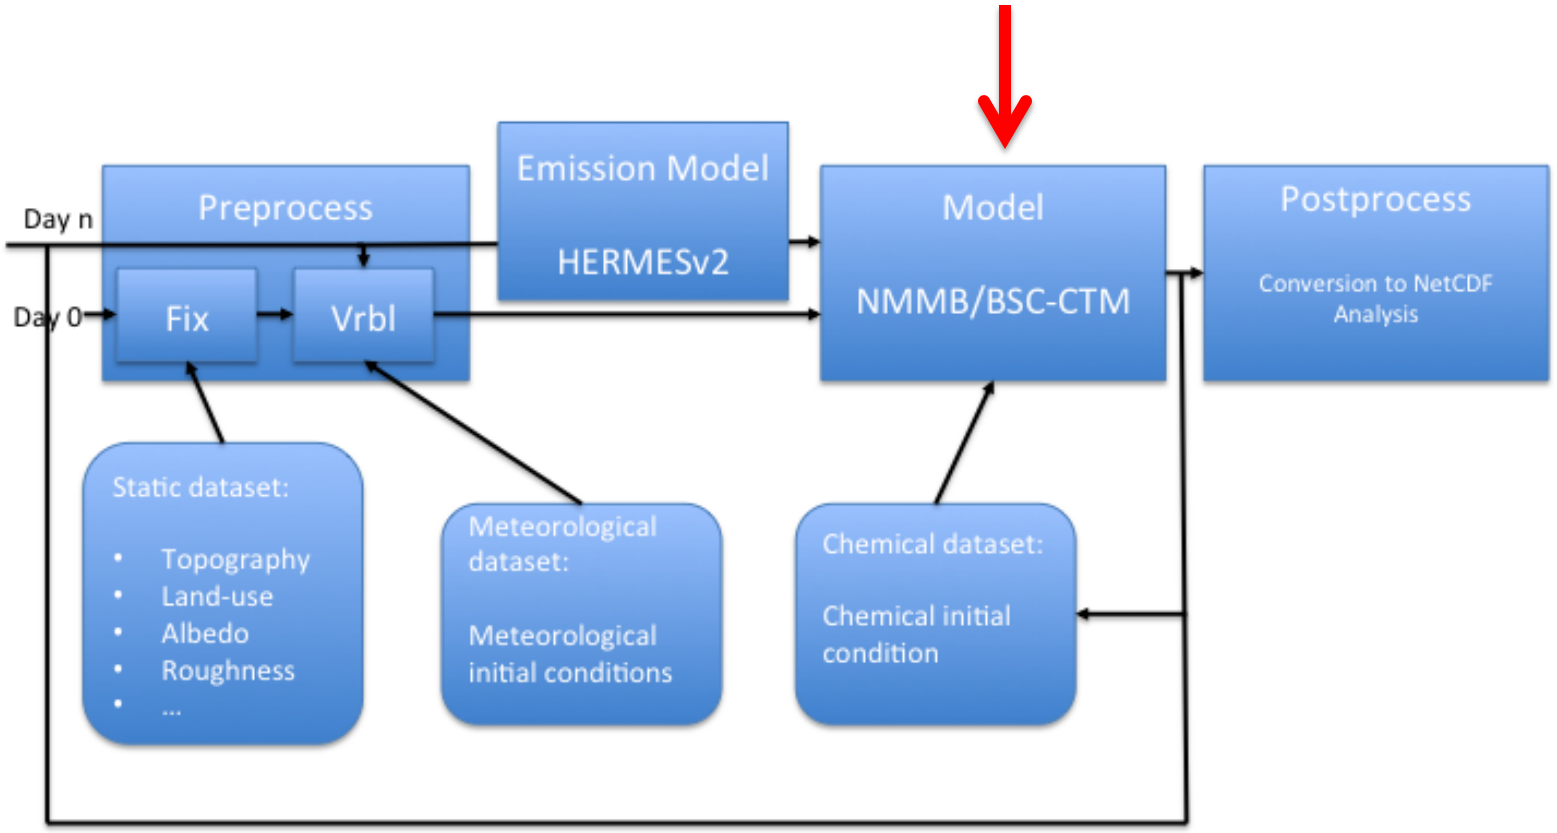

Study domain: Global domain 24km x 24km resolution

Paraver

- One hour simulation of NMMB

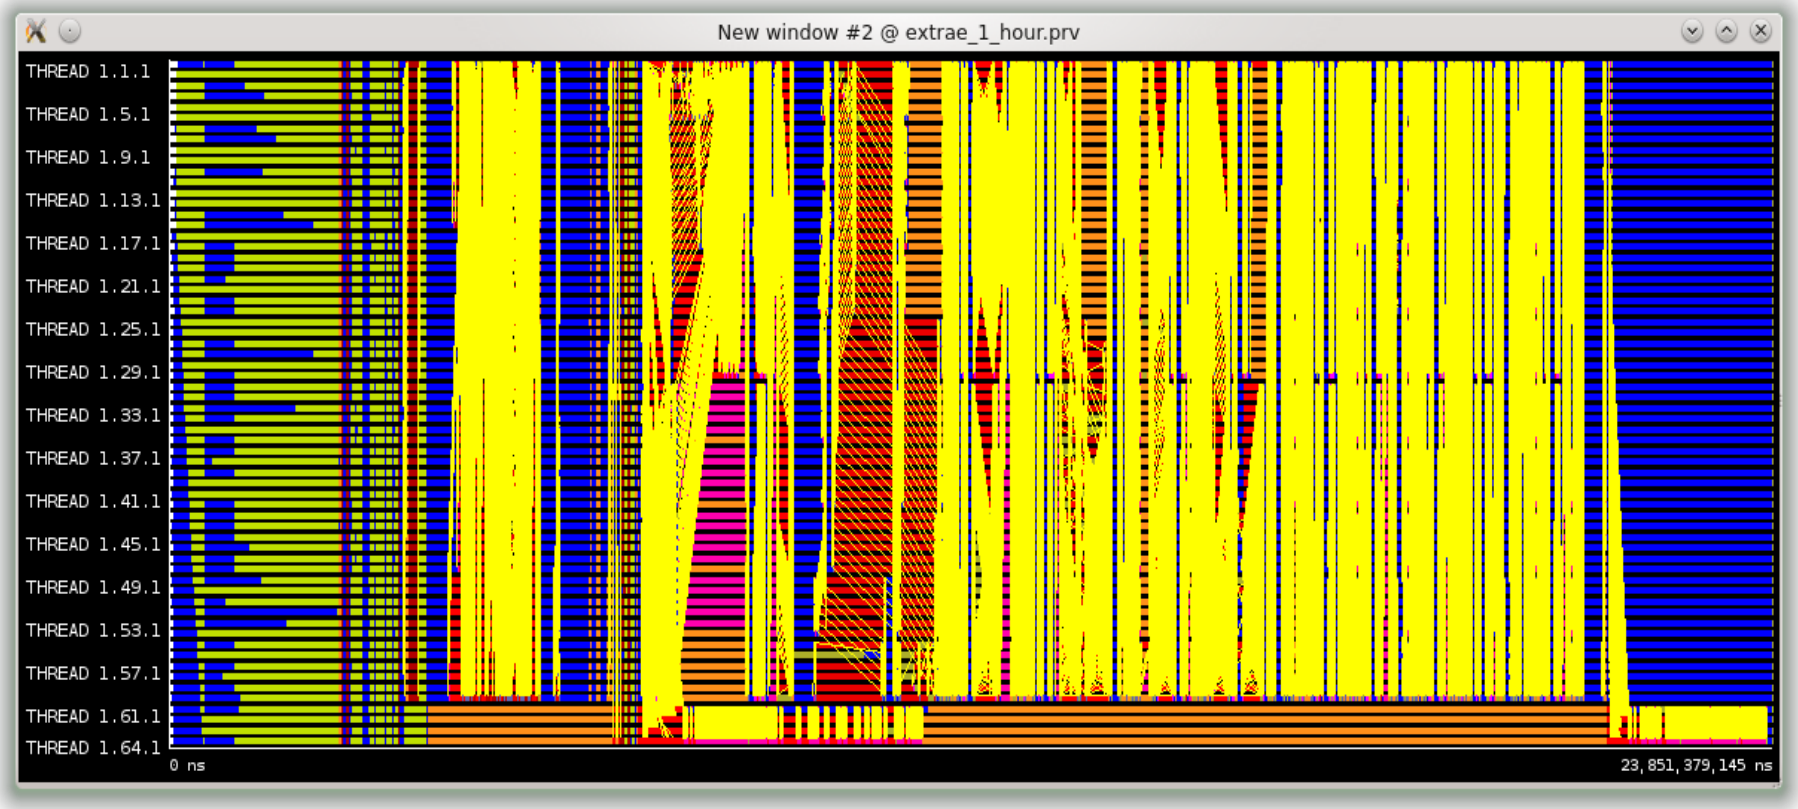

- Last four processes are used for I/O

Issue with I/O

(There is no parallel I/O implemented!

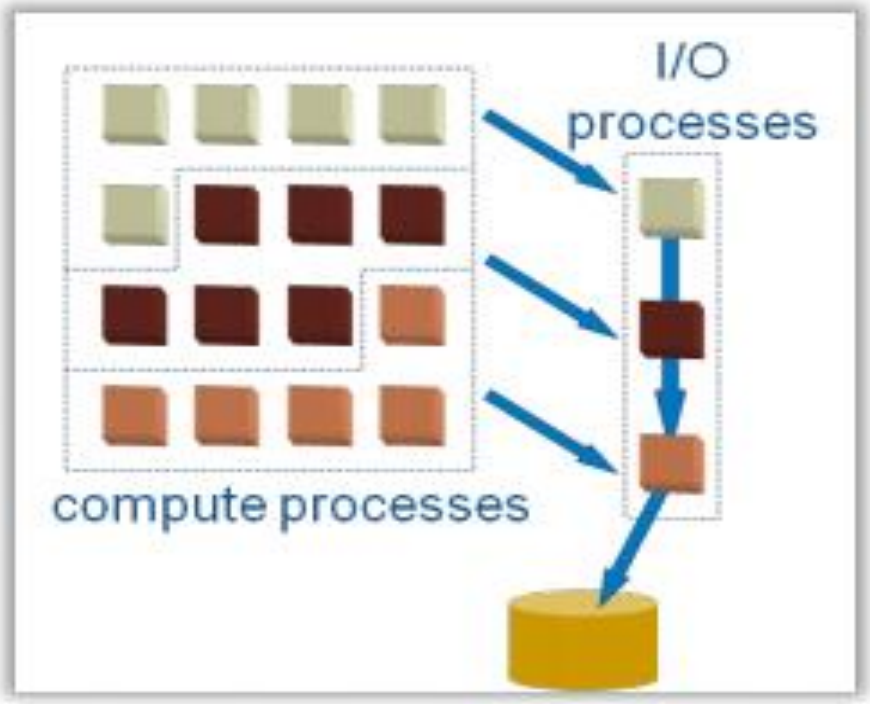

Last hour

With I/O

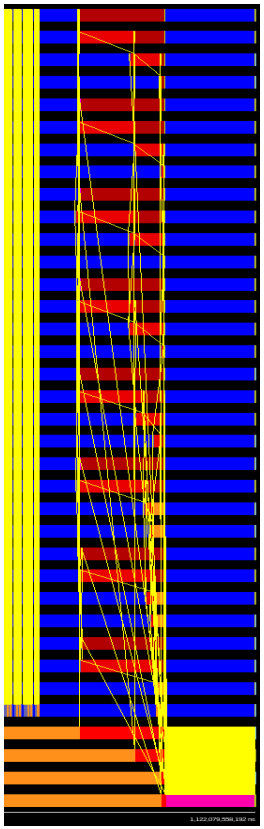

Without I/O

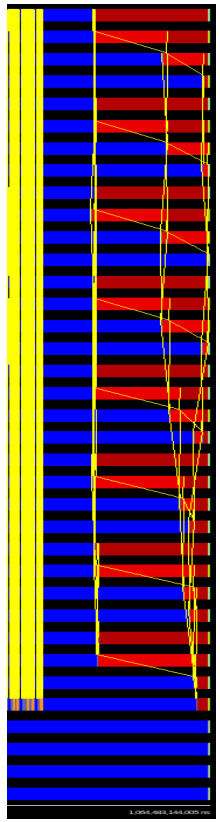

Issue with I/O – Mapping

- Initial mapping for an experiment with 64 cores where the last 4 ranks are the write tasks

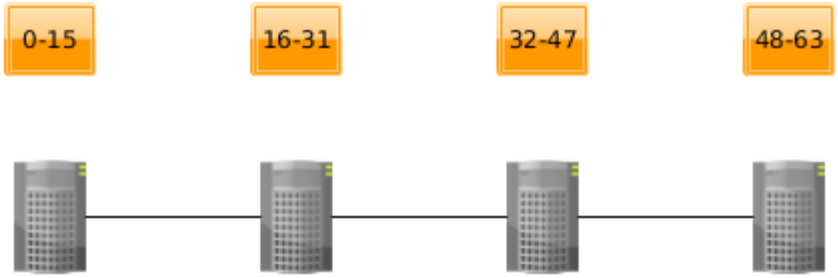

- Final mapping

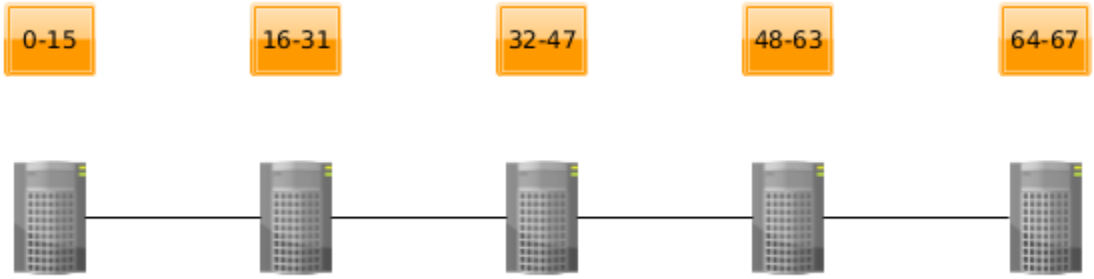

Performance of different mapping and more I/O servers

- The new mapping improved the execution time between 2.73 and 3.85 times

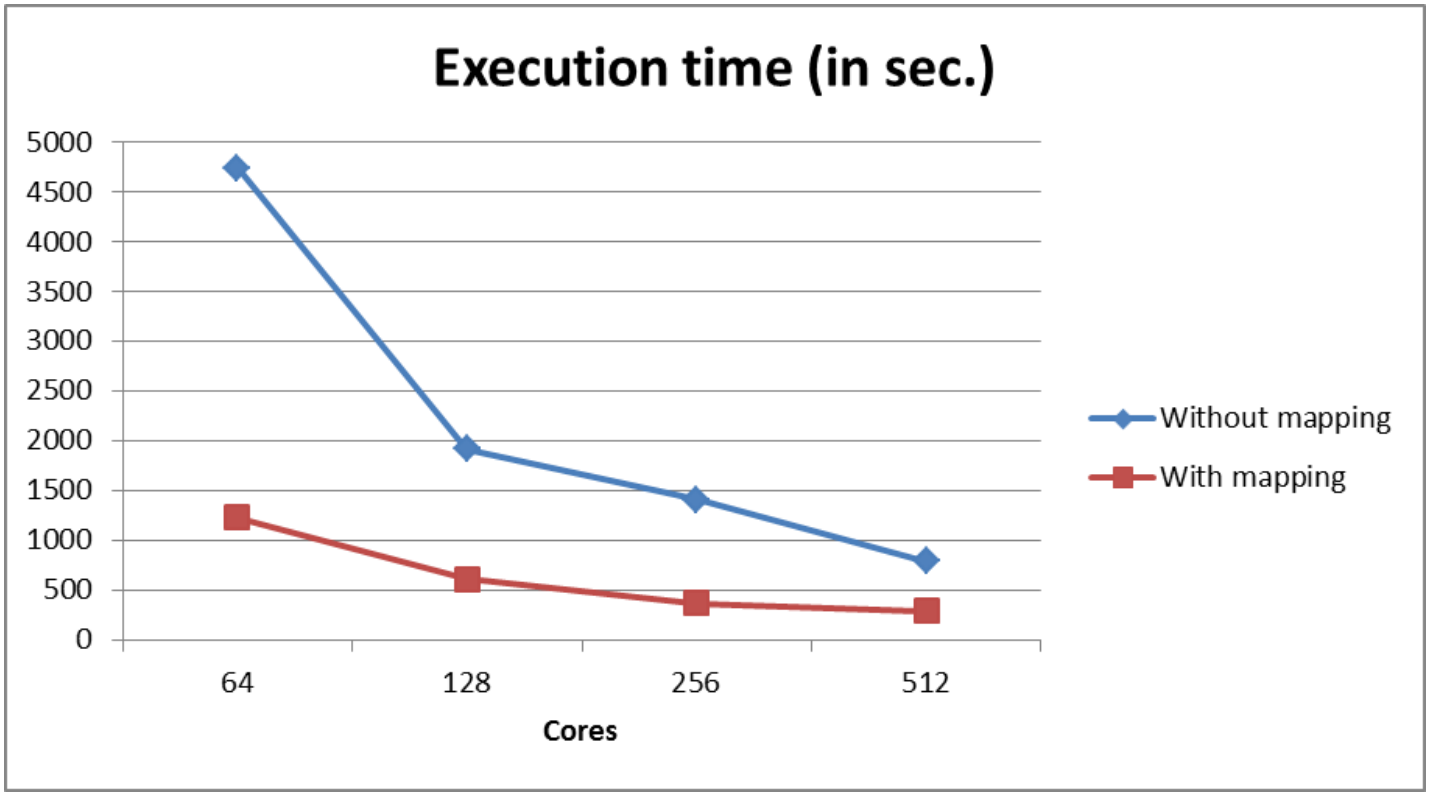

Paraver

- Three hours simulation of NMMB, global, 24km, 64 layers

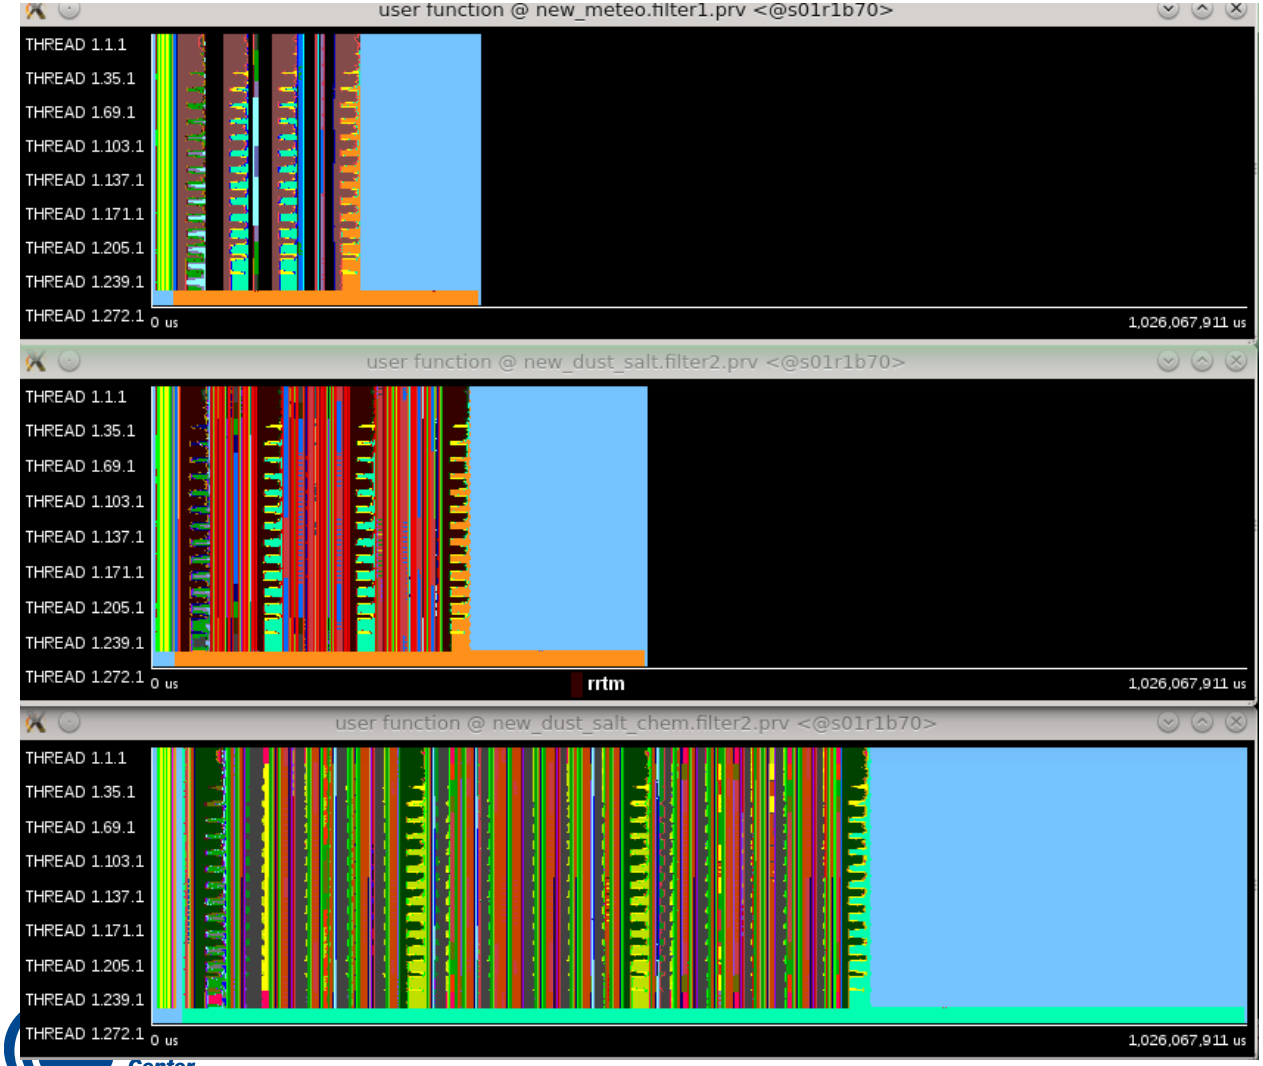

meteo: 9 tracers

meteo + aerosols:
9 + 16 tracers

meteo + aerosols +
gases: 9 + 16 + 53

Paraver – Useful computation - Meteo

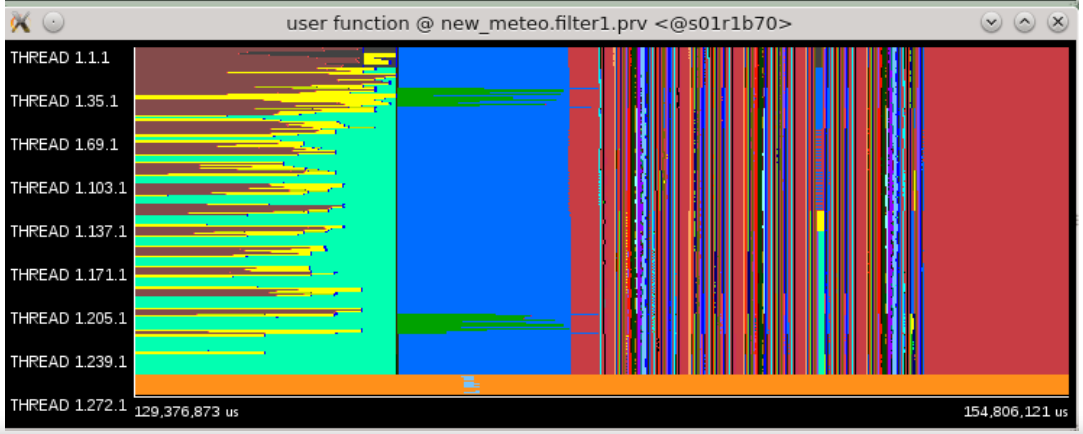

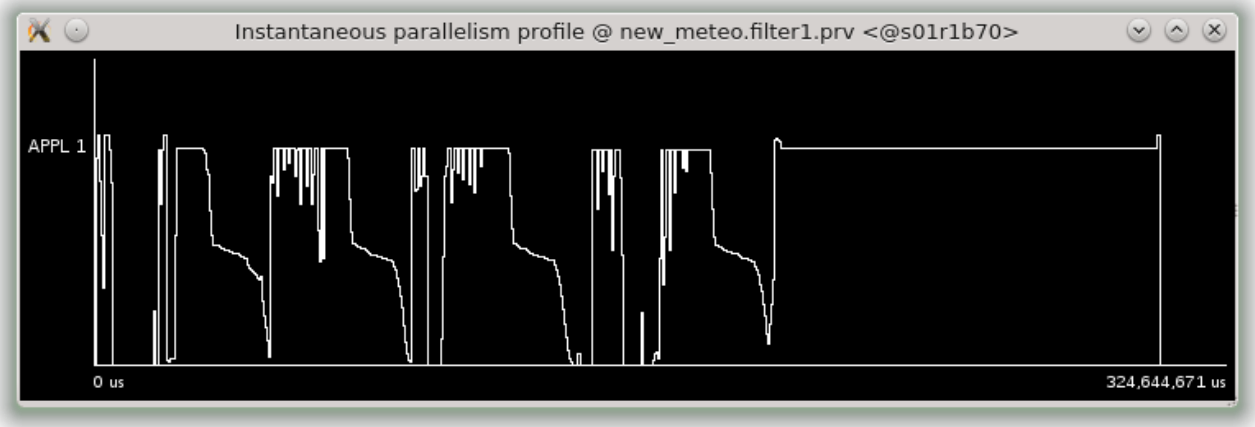

Paraver - Information about functions

Three hours simulation of NMMB, global 24km

	Meteo	
Functions	Percentage	IPC
rrtm	13.7% - 52% (31.3%)	2.18 - 2.38
gather_ layers	8.26% - 13.7% (11.1%)	X
scatter_ layers	10.6% - 14.1% (12.1%)	X

	Meteo + aerosols	
Functions	Percentage	IPC
rrtm	8.8% - 33.4% (20.33%)	2.2 – 2.4
gather_layers	11.9% - 22% (17.4%)	x
scatter_layers	14.4% - 26.6% (19.5%)	X

	Meteo + aerosols + chemistry	
Functions	Percentage	IPC
run_ebi	14% - 20.3% (16.55%)	0.71-1.11
rrtm	3.97% - 15.07% (9.05%)	2.17 – 2.37
gather_ layers	12.37% - 24.55% (16.93%)	X
scatter_ layers	14.65% - 26.58% (19%)	X

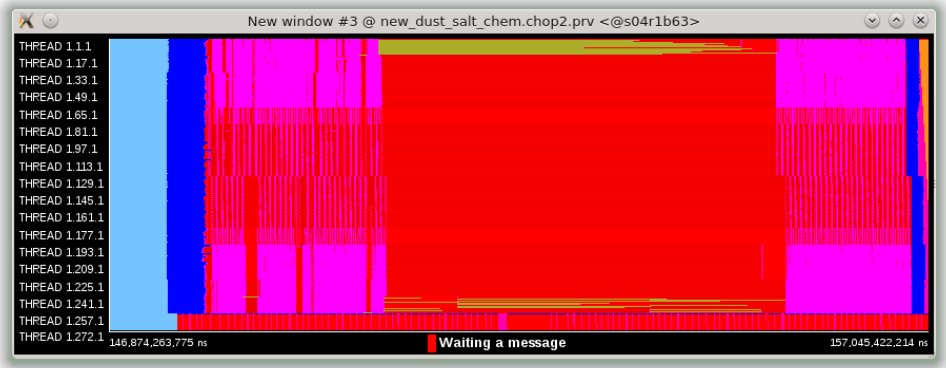

Paraver – Issues/Requests

- Instrumenting functions
  - Dyninst consumes a lot of memory, we had to use 4-8 cores per node
  - The flag –finstrument-functions seems to be more efficient

- We implemented a pseudo throttling mechanism in order to decrease the size of the trace files

- The merging procedure of the intermediate trace files should consume less memory

- We need I/O instrumentation to understand better the future versions of the NMMB model

Roadmap to OmpSs

- NMMB is based on the Earth System Modelling Framework (ESMF)
- The current ESMF release (v3.1) is not supporting threads
- However, the development version of NMMB uses ESMF v6.3
- Post-process broke because of some other issues (which will be fixed)
- The new version of NMMB with OmpSs support has been compiled
- Ready to apply and test OmpSs

Improved I/O (future work)

- Parallel NetCDF written to single files by all MPI tasks

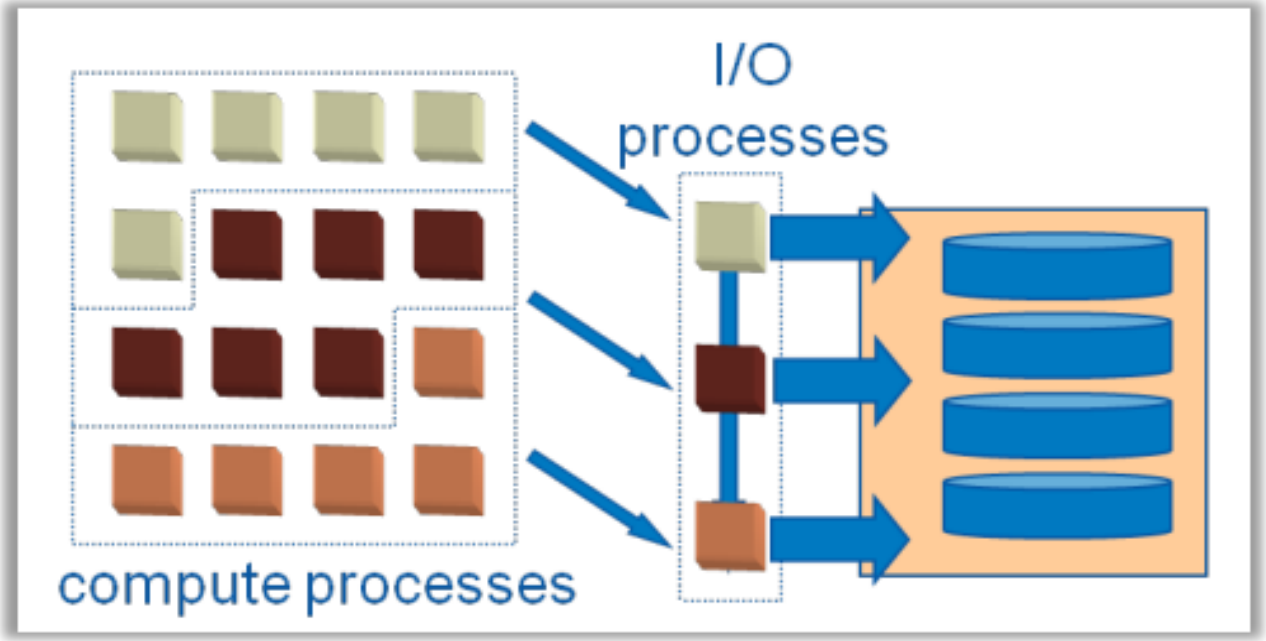

Future work

- Add parallel I/O (writing and reading)
- Use OmpSs programming model
  - Study GPU case
  - Explore Xeon Phi
- Prepare NMMB model for higher resolutions, first milestone is the global model for 12km
- Improve performance and scale NMMB for thousands of cores
- Collaboration with the Computer Science department to prepare a submission to PRACE Scientific and Industrial Conference 2014 (Judit Gimenez, Julian Morillo)

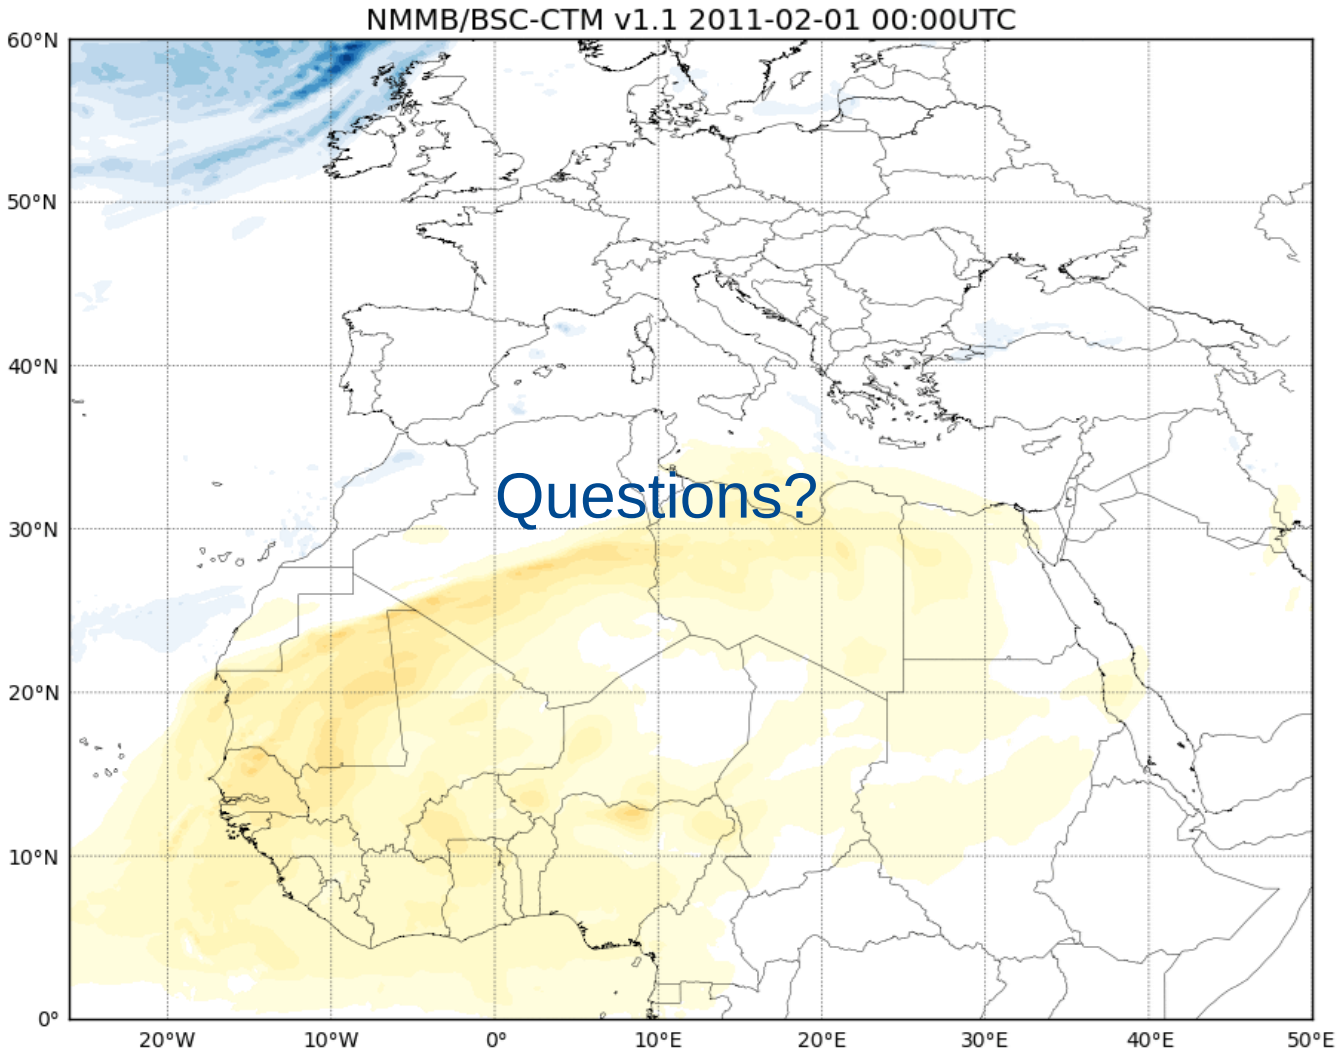
NMMB/BSC-CTM v1.1 2011-02-01 00:00UTC
60°N
50°N
40°N
30°N
20°N
10°N
0°
20°W
10°W
0°
10°E
20°E
30°E
40°E
50°E
Questions?