Project: **1040**

Project title: **ESiWACE: Scalability of Earth System Models**

Principal investigator: **Joachim Biercamp**

Report period: **2018-01-01 to 2018-12-31**

## ICON – 2.5km and 5km global high-resolution simulations

In the reporting period, ICON was successfully prepared to carry out 40-day DYAMOND runs, corresponding to 2.5km and 5km-resolving global simulations. Due to various technical issues, debugging consumed severe amounts of the compute time, along with the actual DYAMOND experiments. To our knowledge, the 2.5km simulation is the first within the DYAMOND initiative and, generally, one of the first simulations that has ever been carried out at this unprecedented global resolution. Figure 1 shows an early state of the 40-day 2.5km DYAMOND run.

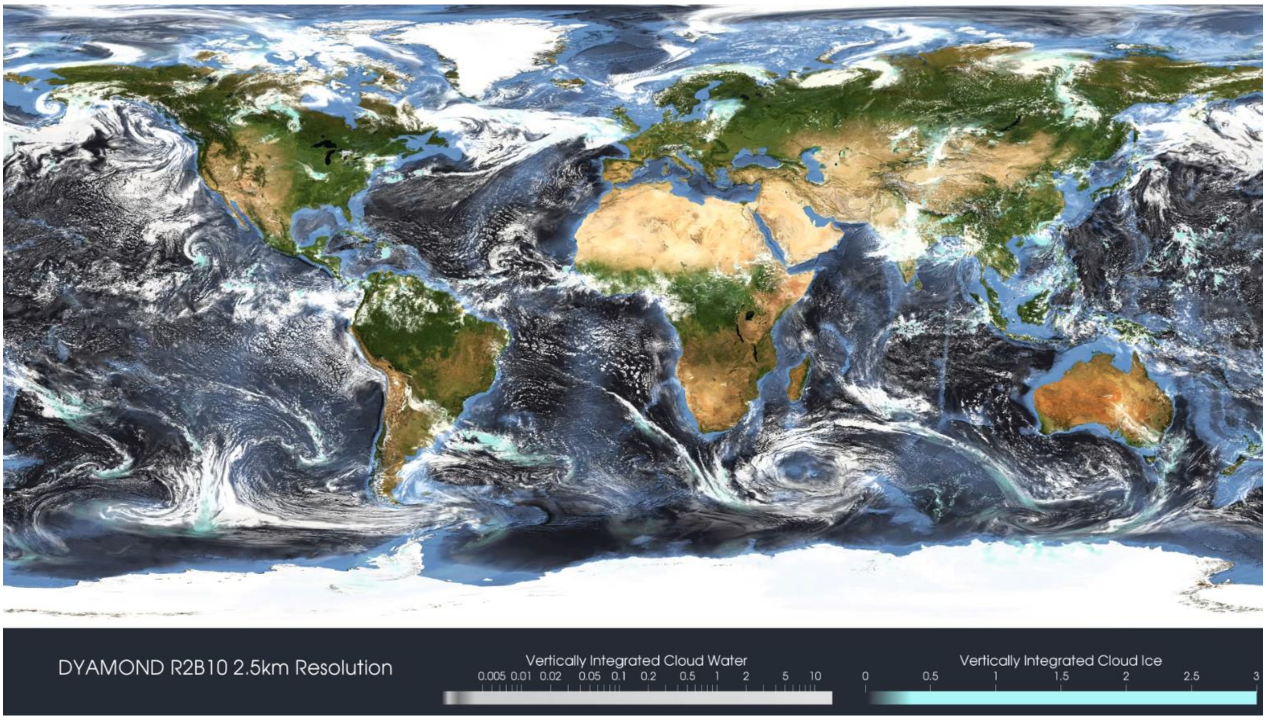


**Figure 1. DYAMOND 2.5km-resolving simulation with ICON. The set up corresponds to ca. 84 million horizontal cells and 90 vertical levels. The time step size was chosen as 22.5 seconds, yielding a total of 153,600 time steps. The throughput of the 2.5km set up was ca. 3 simulated days per day on 540 compute nodes of Mistral, compute2.**

Investigation of the high-resolution features, in particular with regard to resolving some of the convective physics, is in progress. Preliminary analyses suggest amongst others that prediction of mean tropical precipitation rates is significantly improved in the high-resolution, explicit convection-resolving ICON-DYAMOND simulation.

We have further improved performance of the DYAMOND set up in collaboration with senior researchers at MPI-M and Deutscher Wetterdienst, including hardware-aware macro definitions, loop re-orderings and I/O configuration. In particular, compared to previous high-resolution simulations in this project bk1040, DYAMOND runs feature full 2D and 3D I/O. The asynchronous I/O configuration has for this purpose been tuned, severely increasing I/O scalability on up to 900 compute nodes (Mistral, partition compute2, considering the 5km DYAMOND set up). All these improvements accumulated to an increase in throughput from original 19 simulated days per day (SDPD) to 26 SDPD on 300 nodes in the 5km simulation set up.

Excerpts of the scientific DYAMOND results and scalability measurements have been submitted for publication [NEU18]. In this context, we have further – based on the 5km DYAMOND set up – developed performance models, capable to predict and extrapolate scalability of the model at a different (e.g., 1km) resolution and at higher node counts. This is expected to be of great relevance to predict performance on upcoming (pre-)exascale supercomputers.

## FESOM2 – Performance optimisations and scalability improvements

During the reporting period, our main aim (parallel to the ongoing work on the validation of the new model) has been the performance and scalability optimisation. The main bottleneck identified in the previous round of scalability runs (early 2018) was the sea ice model that demonstrated rapidly deteriorating scalability starting from 2000-3000 2D vertices/compute core. By adopting a carefully chosen numerical and algorithmic remedy in the sea ice sub-model (elasto-visco-plastic (EVP) subcycling), we have been able to substantially improve the scalability of the sea ice within FESOM2. In addition, our hierarchic parallel partitioning of the computational meshes originally created to improve the mapping of the parallel partitions to the cluster topology appears to have beneficial effects on the performance of the sea ice model part – especially on systems with slower communication (e.g. Emmy at RRZ Erlangen that uses QDR fat tree Infiniband with 40 Gbit/s network).

The runs on Mistral in the framework of the previous phase of ESiWACE required much more CPU time than originally planned and requested due to various problems with the code scalability. This work helped to optimise the performance of the sea ice sub-model in FESOM2 and allowed to extend a nearly linear scaling trend for the ocean dynamics part of FESOM2 to even greater numbers of compute nodes (currently up to 320 Mistral nodes). The latest simulation runs also played a key role in identifying a new bottleneck that started to manifest itself at large numbers of nodes: the 2D linear solver for the sea surface height. Any further performance optimisation and scalability improvement must focus on this model component.

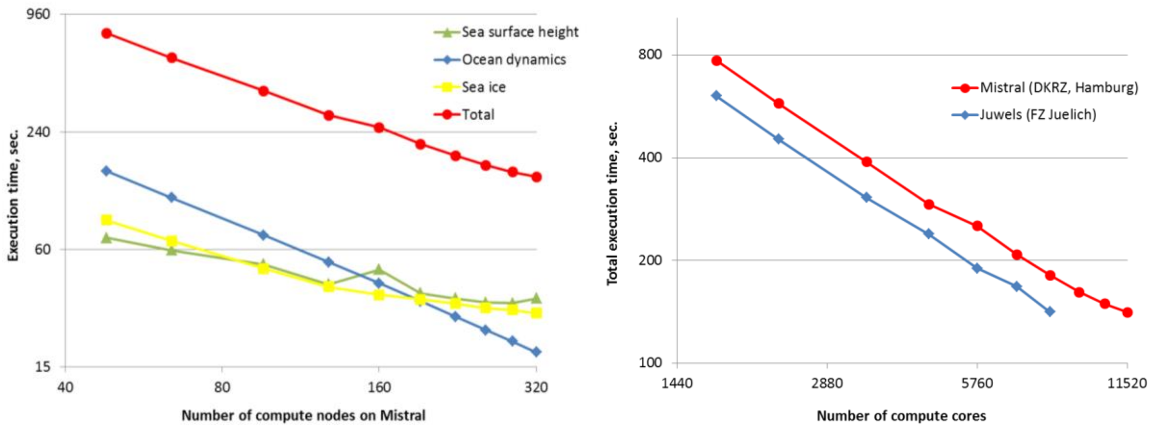


**Figure 2. Left: Mistral scaling for Storm mesh (5.6 mio 2D vertices). Right: Mistral vs. Juwels scaling for Storm mesh. 5 simulation days at 4 min time step (1800 time steps in total) were considered.**

## References

[NEU18] P. Neumann et al. *Assessing the Scales in Numerical Weather and Climate Predictions: Will Exascale be the Rescue?* Submitted, 2018