Report: ESiWACE: Scalability of Earth System Models

Project: 1040
Project title: ESiWACE: Scalability of Earth System Models
Principal Investigators: Florian Ziemen (DKRZ), Daniel Klocke (DWD), Bernadette Fritzsch (AWI), Bjorn Stevens (MPI-M), Julian Kunkel (University of Reading), Joachim Biercamp (DKRZ)
Allocation Period: 2020-01-01 to 2020-12-31

ICON (~340 000 nh + MPI-M contribution)

We performed several simulations with the ESiWACE2 flagship model ICON. We contributed to the DYAMOND Winter model intercomparison with four major simulations. Two atmosphere-only simulations (at 2.5 km and 5 km resolution) were performed with the NWP physics of ICON. Two simulations at 5 km resolution were performed with the SAPPHIRE physics of ICON, one of those atmosphere-only, one coupled with the ocean. The two simulations at 5 km resolution will allow us to analyze the influence of the different physics packages on the simulation results. By comparing the 2.5 and the 5 km resolution simulations, we can study the effects of resolution on the representation of clouds and convection. The 5 km coupled simulation is paving the way for coupled high-resolution modeling and allows us to study the influence of high-resolution atmosphere-ocean interactions.
Based on the results of these first DYAMOND Winter simulations, we performed an experiment of one full seasonal cycle with the coupled 5 km ICON setup. To compensate a warming of the simulated climate, due to a lack of low tropical clouds, we increased the ocean albedo for the longer integration. The simulation is already in the 11th months, and we plan to complete the first year of this first-ever fully coupled global storm and ocean-eddy simulation soon. As such, it served as a laboratory for analyzing the model’s behavior during the different weather states of the annual cycle and the stability of the simulated climate system (Figure 1). While the required changes in the model configuration during the simulation impair an analysis of the full annual cycle, these experiments are necessary to develop the new flagship climate model generation, and still provide scientific insights into the climate system and atmosphere-ocean interactions.

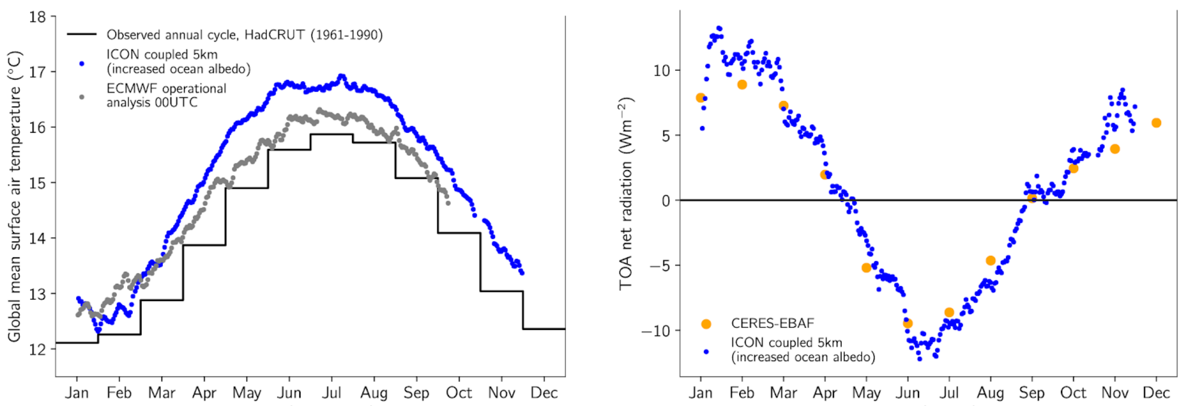

Figure 1: Global mean surface temperature (left) and top of the atmosphere net radiation (right) of the coupled 5 km ICON simulations in comparison to observations for nearly the complete seasonal cycle.

First results indicate very promising improvements when compared to conventional climate models with coarser horizontal grids. One example is the tropical precipitation in the southeastern Pacific exhibits a dry area, which is expected when compared to observations (see Figure 2). Conventional climate models struggle to reproduce this dry area feature west of Peru (Fiedler et al. 2020), which is often referred to as the double ITCZ bias. Past studies were able to alleviate this bias only by pinpoint and costly model tuning (Woelfle et al., 2019), thereby introducing model biases in other domains. In contrast to that, the 5km storm-resolving setup allowed for an instant improvement without pinpoint model tuning. This out-of-the-box-improvment is probably related to the explicit resolving of atmospheric convection of the order of 5-10km horizontal scale.

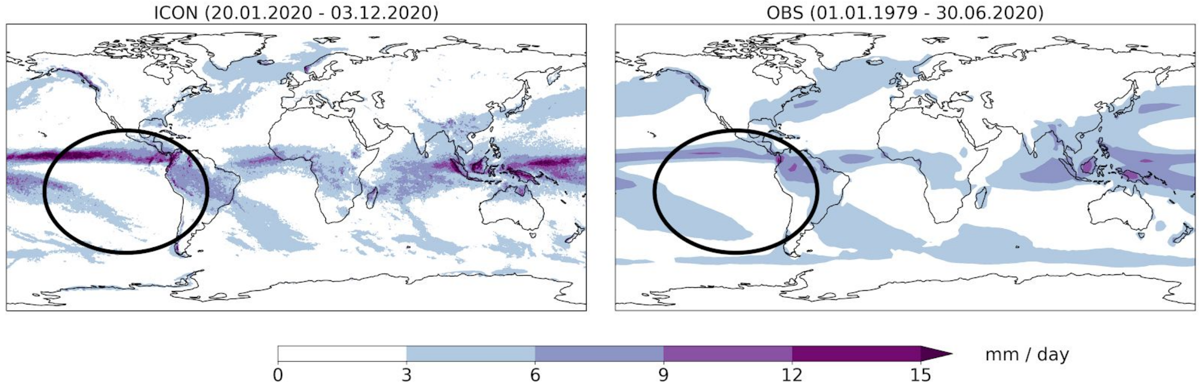

Figure 2: Global precipitation (mm/day) in the (left) ICON DYAMOND 5km coupled storm-resolving simulation and (right) observations (GPCP). Highlighted is the southeastern tropical Pacific featuring realistic dry conditions.

Owing to 5km horizontal resolution, small-scale air-sea interactions are represented. For example, the Gulf Stream in the North Atlantic is captured with all its fine scale structures such as meandering of the flow and its eddies separating from the mean flow (Fig. 3, left).

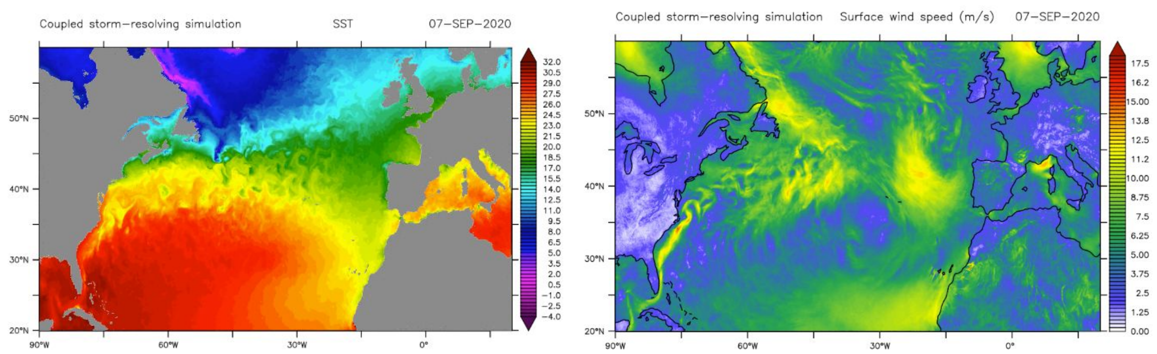

Figure 3: Snapshot of North Atlantic (left) sea surface temperature and (right) surface wind speed in the ICON DYAMOND 5km coupled storm-resolving simulation.

This introduces a low-level atmospheric response on scales as seen from surface wind speeds (Figure 3 right). Those fine-scale features can not be simulated by coarse resolution models and one important challenge is to assess whether these features are realistic.

Overall it must be stressed that this simulation is only of 1 year simulation length and thus these results must be viewed with care. However, these promising indications of mean-state improvements and the resolving of fine-scale ocean-atmosphere interactions serve as an important and motivating basis for performing a longer simulation of this kind. This would provide a more robust basis for assessing the added value of such high-cost simulations and will eventually improve the confidence of climate projections.

One simulation with particularly very high-frequent data output was performed with respect to data visualization, in particular to create a dome-projection video that was then realized in collaboration with the developer teams of ParaView and Intel OSPray[1]. This video was produced for the Eurovis 2020 WISDOME contest (later cancelled due to covid), and pioneered new ray tracing techniques for rendering high-resolution climate data. DKRZ also used the creation of this video to develop capabilities in rendering 3D animations via batch jobs, a feature that will be crucial in dealing with the enormous amounts of data created by the next-generation climate models.
The video will be presented at the Supercomputing Conference (SC2020) as a Visualization Showcase with a peer-reviewed short paper contribution (roughly equivalent to a short paper in geosciences). The data and the rendering techniques were also used for animations that will be shown in Bjorn Stevens' SC 2020 Keynote, enhancing the visibility of DKRZ and German climate modeling within the supercomputing community.
MPI-M contributed compute resources used via user k203123 in project mh0287. In August, September, and October, these totaled at 76 000 node hours. We expect to continue at a high rate in November and December, as the resources available via bk1040 were fully consumed during the usual phase of relatively low demand in summer.

The ICON DYAMOND winter data is at DKRZ and we are in the process of figuring out the best way of transfering the IFS data from ECMWF to DKRZ (last time we used hard disks, but due to the home-office policy of ECMWF this is not an option this time). We have published a data request document on the page of the DYAMOND winter intercomparison[2] asking modelers to provide the data in chunks that facilitate providing only subsets of the data on disk, and fetching reasonably sized chunks of data from tape as required.
In this context, we will build on the recently released packems tools that MPI-M and DKRZ have developed for a more convenient access to the data on the tape system, and with the installation of the new HSM storage system to be installed in 2021[3], we expect further workflow improvements.

FESOM (~110 000 nh + AWI contribution)

During the reporting period, the Finite-volumE Seaice–Ocean circulation Model, Version 2.0 (FESOM2), was further optimized. The new global eddy resolving ocean setup with 33 Million surface nodes and ¼ Rossby radius resolution (Figure 4) was run. The model resolution varies from 1 km in High Latitudes to 6-8 km on the Equator and ca. 14 km in the Tropics. This setup represents the first one globally eddy-resolving ocean model which is

[1] https://youtu.be/4SO2VhmeI4I
[2] https://www.esiwace.eu/services/dyamond/phase-2
[3] https://www.strongboxdata.com/blog/stronglink-wins-150pb-data-management-project-for-the-german-climate-computing-center-dkrz/

capable of running for decades. For correct representation of the water masses exchange through the straits we increase the resolution there down to several hundreds meters (Figure 5).

Because of the huge amount of model nodes and amounts of processors that have to be used, some additional work, focused on scalability and MPI parallelization was done. The main bottleneck finally was I/O. We continued to work on parallized I/O for further improvement of the model performance.

The model demonstrated its capability to carry out long term eddy-resolving climate simulations (100 years and more). In case of sufficient computational resources it can make up to 4 years per day. Preliminary results can be seen on Figure 5 and Figure 6.

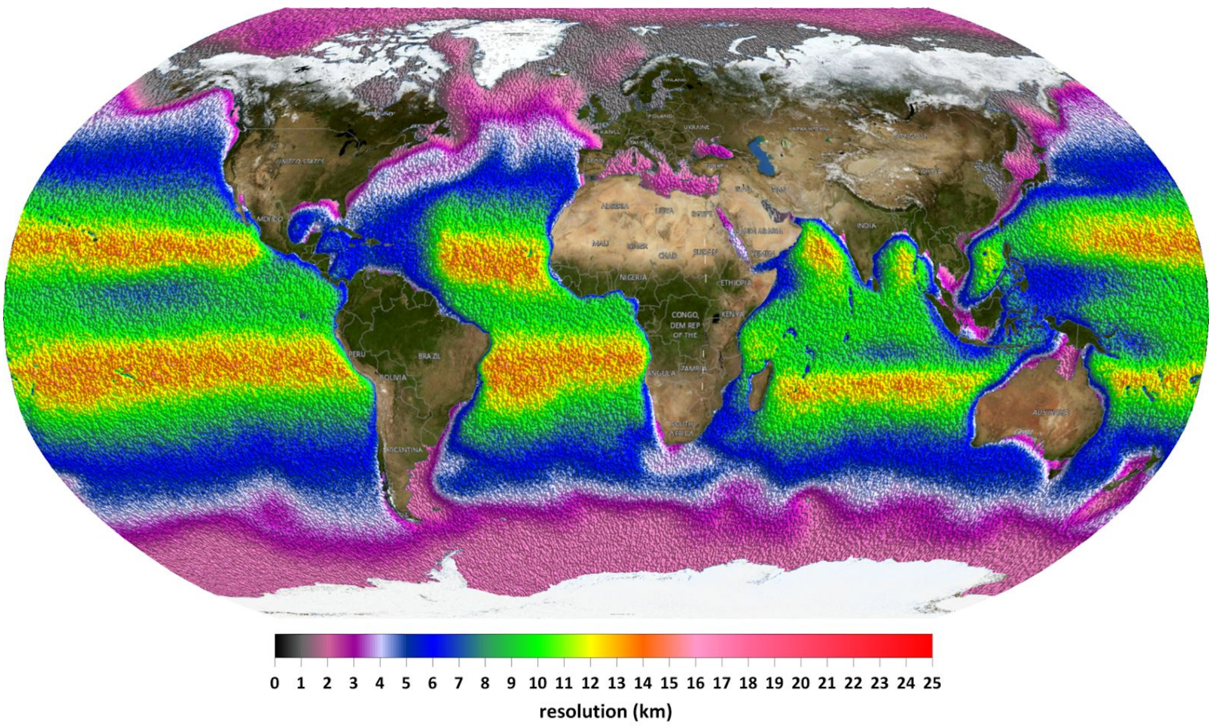

Figure 4: Ocean resolution

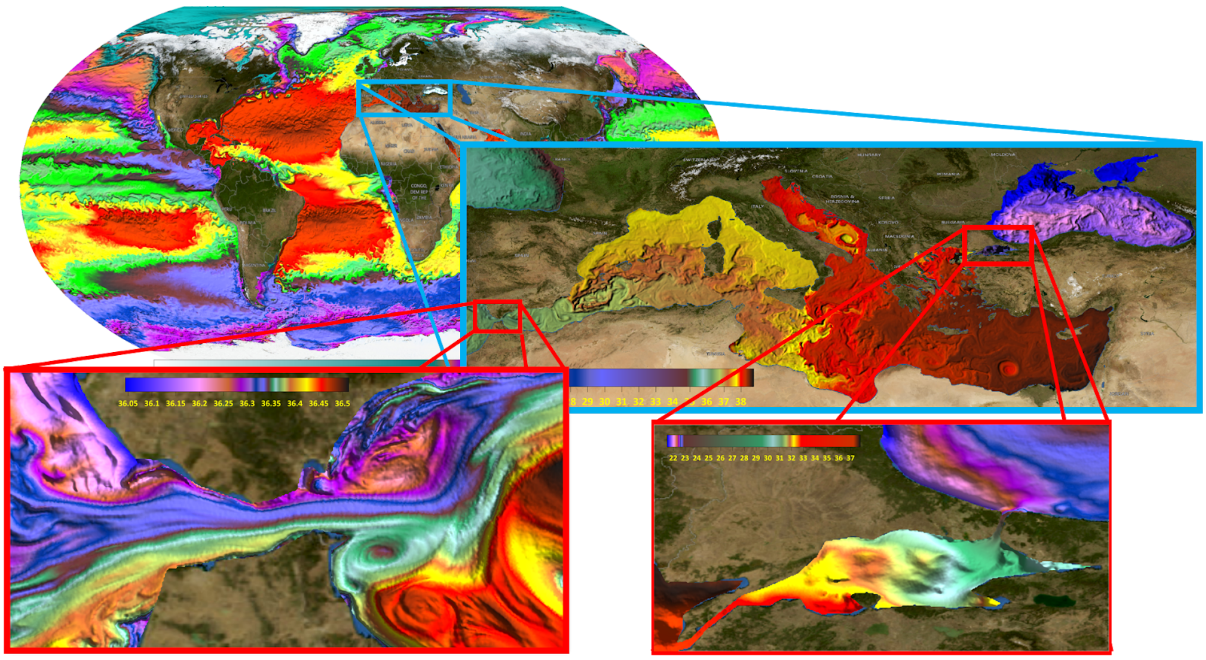

Figure 5: Sea surface salinity snapshots after 10 simulated years. The front panels show dynamics in the Strait of Gibraltar, Bosphorus and Dardanelles, where the model resolution comes down to 300-400 meters.

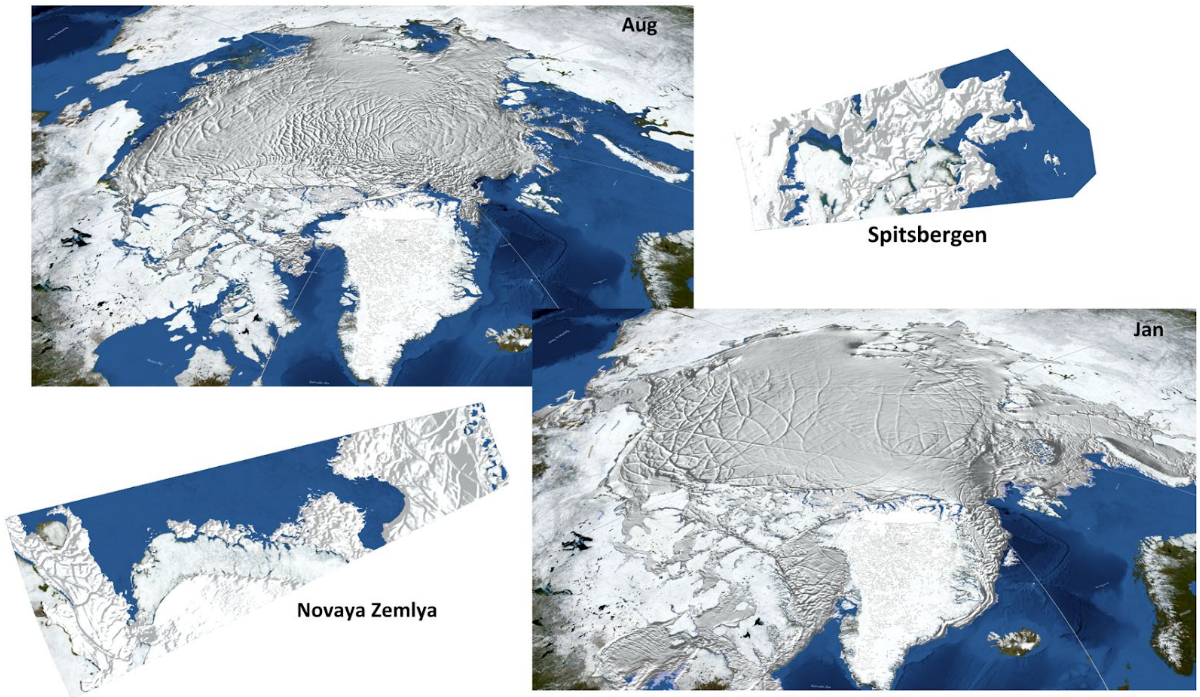

Figure 6: Simulated sea ice thickness in the Arctic after 10 years of the model run. Upper left panel August, lower right - January. Spitsbergen and Novaya Zemlya plots show the leads with open water.