Report on project bk1341

Project: **1341**
Project title: **WarmWorld**
Principal investigator: **Claudia Frauen**
Allocation period: **2023-01-01 to 2023-12-31**

1 ICON SR-ESM baseline 1.25 km coupled

As part of the BMBF funded WarmWorld project, we performed the first ever global coupled climate model simulation with 1.25 km resolution using ICON. Results were shown on the German news "Tagesthemen" and reported on the websites of DKRZ[1] and MPIM[2]. We refer to these simulations as Apollo simulations, as we recreated the iconic Blue Marble picture, which the Apollo17 astronauts took on their way to the moon on the 7th of December 1972. We simulated 7 days, with saving full output for the 7th day. The data is made available to the research community. Based on this simulation output, we also explored new data and visualisation strategies, which allow the interactive exploration of the data[3].
Within this reporting period, we used bk1314 to extend this simulations for a high profile publication we have requested Nature at Climate Change. We used less compute time (36000 node hours, about 38%) for this project than anticipated. The reason is, that we did the ocean spin-up and debugging under a different compute project and that the model's compute efficiency was improved significantly. As we had never performed a simulation at this scale (number of grid points), compute time estimates were extrapolated from coarser resolutions, which where not fully optimised at the time.

2 ICON SR-ESM 5.0 km coupled hybrid

To explore areas for possible performance improvements, an ICON coupled baseline run with spatial resolution of 5 km (R02B09) for atmosphere and ocean has been performed. It uses the nextGEMS Cycle 3 setup, run on 270 compute nodes, with 248 nodes dedicated to the atmosphere component and 22 nodes to the ocean component and gained a throughput of 54.254 SDPD (0.201 SH/node hour). To investigate the scaling of this setup and find possible performance limitations, this run has been repeated with 265, 398, 518, 602, 707 and 1000 compute nodes, always using the same relation of nodes between atmosphere and ocean component (248:22). Looking closely at the timers showed a broken scaling for runs using more than 600 nodes which is mainly caused by the waiting time for the asynchronous output of the atmosphere. This behaviour is now under investigation.

Unfortunately, for most of the year it was not possible to run heterogenous jobs on Levante. Only recently, the problem has been fixed by the vendor. Further, a bug in the coupled 5km model setup had to be fixed before we could start working on a hybrid coupled setup. Thus, we were not able to use any of the requested GPU ressources for the ICON model until the end of October when we managed to do at least the technical setup of the hybrid coupled model. We used a slightly modified setup from above which works with latest version of the ICON model. Further, we disabled mixed precision because it has not yet been validated on GPUs. The reference simulation on CPUs again uses 270 CPU nodes of which 22 are used for the ocean model. With this setup we achieve a throughput of 42.541 SDPD. We run the same setup as a hetjob with the atmosphere running on 40 GPU nodes and the ocean on 27 CPU nodes and additionaly 7 CPU nodes for I/O and get a throughput of 59.3 SDPD. Further optimisation and validation is needed, but can be done until the end of the year.

3 FESOM standalone baseline 4.0 km

The 16326 CPU node hours have been used to run two FESOM standalone simulations with the D3 mesh in order to test the FESOM refactoring branch performances on a longer run using two different ocean vertical mixing schemes (`cvmix_TKE` and KPP). 45 compute nodes have been used to simulate 58 years for each of the two experiments (6.4 SY/day). In Figure 1 the sea ice extent time series show that no major differences were retrieved from the comparison.

Computational resources were also utilized to evaluate the development configuration of IFS-FESOM. The IFS is based on cycle 48r1 with the atmospheric grid Tco1279, featuring a horizontal grid spacing of approximately 9

[1] https://tinyurl.com/mrza9yu
[2] https://tinyurl.com/2p9bbbhy
[3] https://tinyurl.com/4ewbu7jz

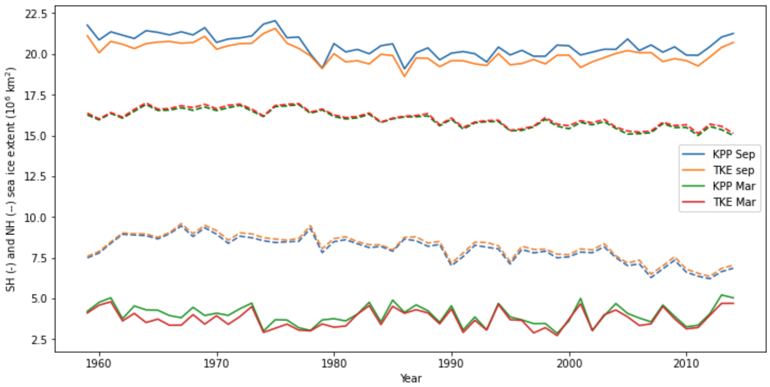

Figure 1: Sea ice extent values for the entire simulated period.

km and 137 vertical levels in the atmosphere. The ocean model, FESOM version 2.5, operates on an unstructured NG5 mesh with a horizontal resolution of approximately 12 km in the tropics, which gradually decreases to about 4–5 km in polar regions, and it employs 70 vertical levels. The validation process followed the protocol outlined below:

1. the ocean-only spin-up was initiated, from EN4 data for January 1950 and January 1985 for temperature and salinity. Forcing was provided by ERA5 reanalysis data for the periods 1950–1954 and 1985–1989, respectively.

2. the coupled spin-up was executed for the time intervals of 1950–1951 and 1990–1991.

In Figure 2 we can see that IFS-FESOM exhibits a moderate level of performance in capturing the evolution of globally averaged monthly surface air temperatures, suggesting a reasonable outcome.

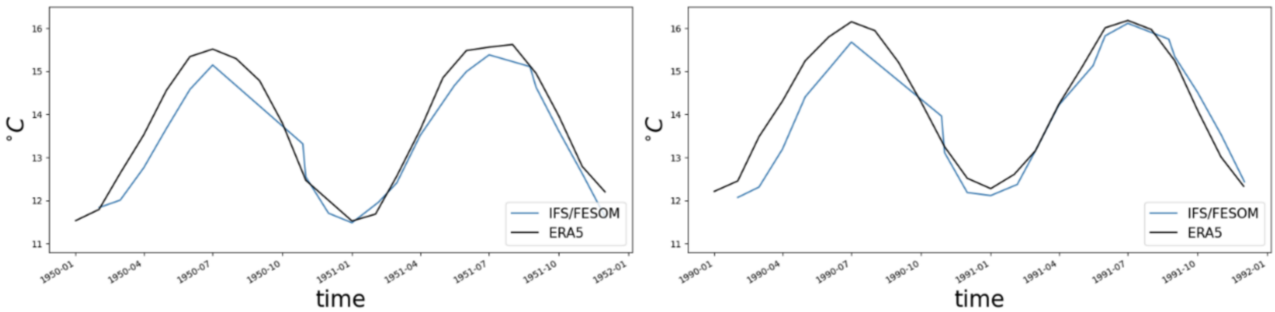

Figure 2: (Left) Depicts the evolution of globally averaged monthly surface air temperatures in IFS-FESOM over a 2-year period in the 1950s during the coupled spin-up, compared to ERA5 surface air temperature data from 1950 to 1951. (Right) Presents the same analysis, but for the years 1990 to 1991.

In the summer of 2023, our initial plan was to carry out several development experiments involving concurrent model runs. However, due to significant queuing times on Levante, we temporarily shifted our work to the JSC machine for these tasks. As a result, we were only able to make use of 30% of the resources we had originally requested at DKRZ. For the next phase of the project, we will resume our development tasks on Levante.

4 FESOM GPU port development

We were unable to utilize the GPU resources as we initially intended. A noticeable disparity between the expected and actual throughput became evident during our testing of FESOM with a GPU on the LUMI machine. Despite having ported a part of the model to the GPU, any performance improvements in kernel computations were consistently nullified by the overhead resulting from frequent data copying between the CPU and GPU. Consequently, we have increased our efforts to migrate additional parts of the FESOM code to the GPU. Additionally, GPU support for IFS became available only recently and has not yet been tested on Levante. In the next phase of the project, we intend to concentrate on using more GPU.