Project: 1313
Project title: Horizon EU-project EPOC [**E**xplaining and **P**redicting the **O**cean **C**onveyor]
Project lead: Jochem Marotzke (MPI-M), Dian Putrasahan (MPI-M)
Allocation period: 2023-07-01 to 2024-06-30

Achievements in the previous allocation period

The general objective of EPOC is aimed at generating a conceptual framework of the Atlantic Meridional Overturning Circulation (AMOC), and MPI-M involvement in EPOC is mainly in WP 2, 3, and 4. More specifically:

- WP2 for understanding key processes responsible for maintaining or breaking meridional connectivity of ocean transport
- WP3 to identify processes and drivers of recent change in AMOC and infer role of external forcing,
- WP4 for assessing key processes and feedback for future changes in AMOC

In WP2, MPI-M uses existing MPI-ESM runs and will configure and run very high-resolution coupled simulation (ICON) for the period 1993-2022, with an additional 30-year control simulation to account for model drift. For WP3 and WP4, MPI-M will be producing MPI-ESM-HR and MPI-ESM-ER simulations (historical with fixed GHGs) and use them in conjunction with those from PRIMAVERA to address the respective objectives.

In the first year of EPOC (July 2022- April 2023), we were porting and running experiments with MPI-ESM-HR and -ER so that analysis with their PRIMAVERA counterpart can be made (WP3 and 4). Oliver Gutjahr started on the project in Jan 2023, and has succeeded in porting and running MPI-ESM with HR and ER configuration on Levante. He optimized and obtained a throughput of ~2.4SYPD on 181 nodes for the ER configuration. It is anticipated that the runs will be completed by the end of this allocation period (June 2023). Additionally, we took a look at the effect of resolution (NR=1°, HR=0.4°, ER=0.1°) on AMOC connectivity in MPI-ESM runs (WP2), and found that connectivity of subpolar AMOC (40°N-60°N) to AMOC in lower latitudes decreases with increasing resolution (Fig. 1).

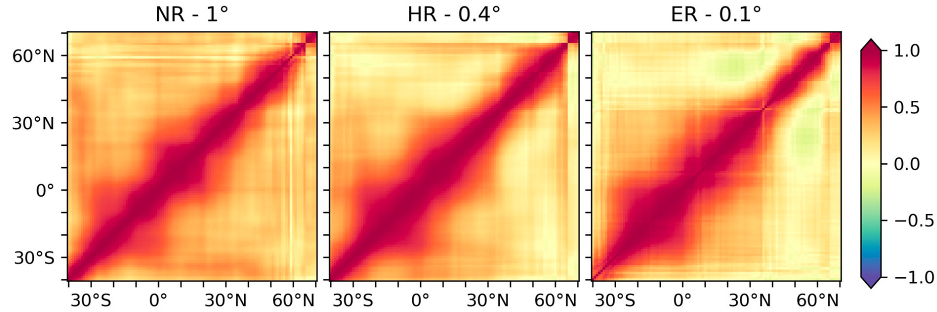

Figure 1: Cross-correlation of AMOC at each latitude for 3 different resolutions

We have been working on a new coupled ICON configuration for WP2, constantly weighing out the trade-off between computational load vs resolving processes of interest. Originally, we were going to use the coupled 5-km ICON-Sapphire configuration, akin to NextGEMS. However, the resolution for the North Atlantic is not as satisfying, particularly when looking to processes near the Flemish Cap where observational arrays are placed and can be compared with. We therefore tried and opted for a telescoping ocean grid to get resolution as high as ~1.5km around the Flemish Cap and submesoscale eddy-permitting resolution in the North Atlantic and GIN Seas. We considered between using a base 2.5km vs 5km grid and concluded on a telescope grid with base 5km resolution due to computational cost (Fig. 2).

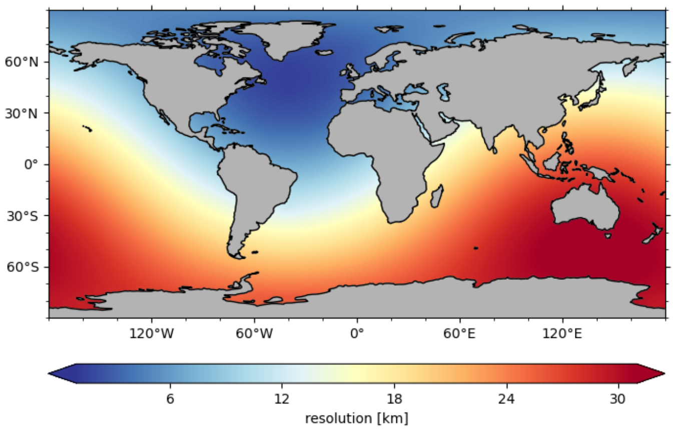

Figure 2: Horizontal resolution of telescopic ocean grid

To reduce data load and increase throughput while keeping the processes we want to capture, we played with reducing the vertical resolution of the ocean. Here are some of the scientific considerations of the vertical resolution at various depths of the ocean in relation to capturing certain scientific phenomena:
1) thickness of first layer for air-sea interactions (we kept it at ~2m for consistency with NextGEMS)

2) high resolution (~2-10m in upper 100m) to capture mixed layer depth dynamics
3) decent resolution (~50m) in upper 1000m for overflows in the North Atlantic, e.g. Denmark Strait (~700m depth), Faroe Banks (~900m depth), Flemish Cap (~1000m)
4) 200m resolution near mid-Atlantic ridge (~3000m deep), Charlie-Gibbs fracture zone (~3000m, where deep Labrador Sea water flow) or places where internal tides/waves are generated

Originally, 128 levels were used, and we finally settled on 72 levels while satisfying all the considerations above (Fig. 3). Further details on how we arrived on the 72 levels can be found on https://pad.gwdg.de/s/agipmdesM Based on the stand-alone ICON-O test simulations, we found that on 251 nodes, the original R2B9L128 had a throughput of 360 SDPD, while with the new R2B9L72 configuration on 126 nodes provided a throughput of 720 SDPD, which is an equivalent to a factor 4 improvement in performance. The 72-level vertical configuration will also be employed in EERIE, NextGEMS and DestinE for consistency and ease of comparability. The atmospheric component (ICON-A) is computationally 3-4 times more expensive compared to the same ocean resolution, so we re-evaluated the necessity of a 5-km atmosphere and concluded that for the purposes of EPOC, a 10-km atmosphere would suffice in capturing the resolved processes necessary and relevant to our studies while significantly reducing the computational requirements, making it very similar to the EERIE configuration and throughput. In summary, the new EPOC configuration is a telescoping oceanic grid with a 5-km base resolution that focuses in the North Atlantic, has a resolution as high as ~1.8km around the Flemish Cap and 72 vertical levels, coupled with a uniform 10-km atmosphere with 90 levels.

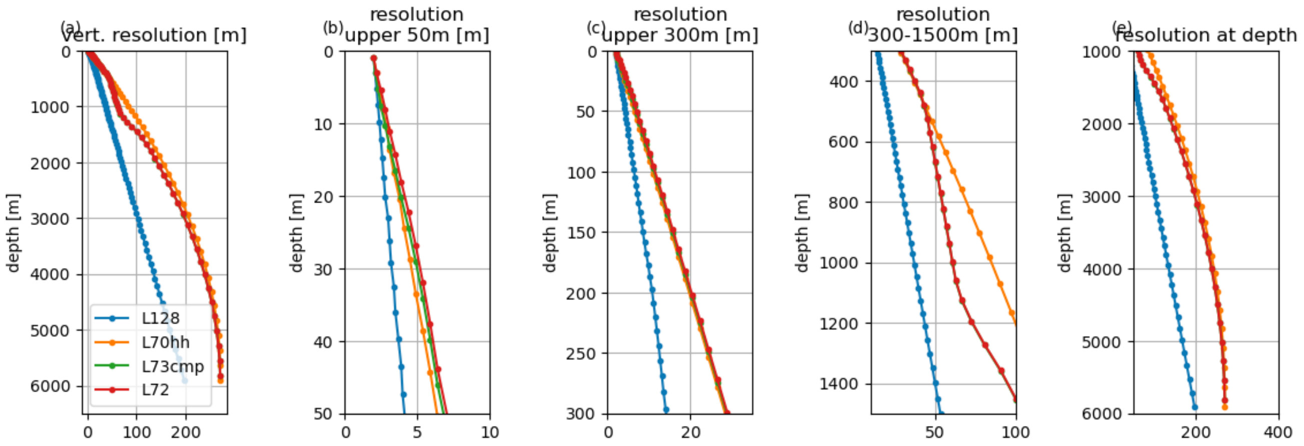

Figure 3: Various vertical resolution. Originally, we used L128 (blue) and we settled on L72 (red)

Additional benefits in the form of throughput and disk storage space come with the new EPOC configuration. On 400 nodes, the node hours per simulated year decreased from 43,800 down to 12,000, while the throughput improved from 0.2 SYPD to 0.8 SYPD. The quadrupling of throughput is a combination of reduced ocean vertical resolution as well as adjustment in load balancing of nodes between ocean and atmosphere components of the model. Original load balance was 24:8 for atmosphere to ocean, and for new configuration, we use 16:16. Here, the load balance accounted for ratio of the atmosphere-to-ocean resolution, as well as the telescopic ocean grid where we had to use a conservative time step (~150sec) relative to a typical time step for R2B9 uniform grid (~300sec). Previously, with the coupled 5-km configuration, job integration lengths were less than a month, so we could not output on monthly intervals, rather on daily intervals which amounted to at least 16TB/simulated year for just daily 3D atm and ocean output. With the new EPOC configuration, our integration job length is now ~1 month/job, allowing for monthly output such that with daily 2D outputs and monthly 3D output, this amounts to ~2TB/year. High frequency output is desirable but to account for storage space, this will be done in time slices, which would be an additional ~16TB/year for the period of the time slice (~5years).

The table below summarizes the differences between the original configuration and the new EPOC configuration.

	Original configuration	**New EPOC configuration**
Ocean resolution	uniform 5km	telescopic 5km, focal point around Flemish Cap at ~1.5-km resolution
Atmosphere resolution	uniform 5km	uniform 10km
Ocean vertical levels	128	72
Throughput on 400 nodes	0.2 SYPD	0.8 SYPD
Node hours per simulated year	43,800	12,000
Integration length	30 years transient	30 + 30 years (transient & control)
Storage per simulated year	~5.8TB/year	~2TB/year + ~16TB/year for 5 years

It took some time to finalize the new EPOC configuration. We are currently performing the coupled spin-up, which based on HighResMIP, may require 30-50 simulated years and we anticipate to complete this by the end of the current 2022/2023 allocation period. We thus expect to start production runs (30-years control & 30-years transient) at the beginning of the next allocation period (2023-2024).