Overview of the first year of the NEMO global 1/36° configuration (ORCA36) development

Clément Bricaud(1), Miguel Castrillo(2)
(1): Mercator Ocean International, (2): Barcelona Supercomputer Center

Contact: cbricaud@mercator-ocean.fr

Global 1/36° (ORCA36): context

- Model configuration for future CMEMS/MOI global forecasting and reanalysis systems

- Based on NEMO 4

- Projects:

IMMERSE (EU H2020)

ESIWACE2 (EU H2020)

ESiWACE2 - for future exascale climate and weather predictions

Projet objectives:

- improve efficiency and productivity of numerical weather and climate simulation and prepare them for future exascale systems
- prepares the European weather and climate community to make use of future exascale systems in a co-design effort involving modelling groups, computer scientists and HPC industry

MOI involved in WP1.Task 1.1:

- WP1. : « Production runs at unprecedented resolution on pre-exascale supercomputers »
- Task 1.1: « Develop infrastructure for production-mode configurations »

« To enable production-mode simulations at the highest resolution possible, to be able to fill a significant fraction of a pre- exascale EuroHPC system, and to allow a scientific comparison of results some infrastructure needs to be developed. »

⇒ **Provide NEMO-based global 1/36° ORCA36 configuration**

EU Horizon 2020 grant agreements 823988

IMMERSE: world-class marine modelling tools for its next generation systems

Projet objectives:

- *Develop a new, efficient, stable and scalable NEMO reference code with improved performances adapted to exploit future HPC technologies in the context of CMEMS systems.*
- *Develop NEMO for the challenges of delivering ocean state estimates and forecasts describing ocean dynamics and biogeochemistry at kilometric scale with improved accuracy*
- *Prepare the exploitation of the next generation of high resolution observing networks within CMEMS systems and in detailed, downstream modelling systems.*

IMMERSE-WP6: ORCA36= high resolution configuration used as a bench
For developped code in WP3 (numerics), WP4 (HPC) and WP5 (physics)

- Short simulations (operationnal objective: 7 days) to assess developments
- A first simulation (several months) to validate the configuration
- A long simulation with NEMO4/IMMERSE code to highlight IMMERSE developments

Global 1/36° (ORCA36): context

- Collaborations:

CMEMS contract with BSC:

« 87-GLOBAL-CMEMS-NEMO: EVOLUTION AND OPTIMISATION OF THE NEMO CODE USED FOR THE MFC-GLO IN CMEMS » :

NEMO HPC performances, especially with global 1/36°

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CMEMS contract with CNRS/IGE/MEOM team:

« 2-GLO-HR Evolution of CMEMS Global High Resolution MFC »

IGE Institut des Géosciences de l'Environnement

➢ sensitivity of NEMO solutions to numerical and parametric choices in realistic configurations

an Atlantic (20S-81N) 1/12° configuration with AGRIF zooms (1/12° to 1/48° and 75 to 200 vertical levels)

➢ Definition of metrics to assess resolved fine-scale structures

Small scale vorticity variance, KE wavenumber spectra, regularity of resolved fields at the grid scale,submesoscale vertical buoyancy flux, fine scale horizontal gradient of surface buoyancy

Model configuration

- Horizontal: tripolar ORCA grid, 1/36° résolution (2-3km)
- Vertical: 75 Z-levels, 1 meter at surface
- Bathymetrie: based on ETOPO08
- Runoff: climatology
- Atmospheric forcing: Era-interim (on-line interpolation)
- Initial condition: from MOI ¼° reanalysis (shorten model spinup)

Model parametrizations

NEMO 4

Variable volume

Forcing:

Erainterim with NCAR bulk formulae and analytical diurnal cycle

Surface frequency frequency: every time-step

Atmospheric pressure gradient added in ocean & ice Eqs.

2 bands light penetration scheme

Sea Ice model :

SI3

Levitating sea ice

5 categories

EOS80 for equation of state

Hydrostatic pressure gradient: s-coordinate (standard jacobian formulation)

BC:

Lateral friction: free slip

logarithmic top/bottom drag coefficient

Tracers transport:

TVD advection scheme 4th order on horizontal and vertical

Explicite diffusion with triad iso-neutral operator

No damping

Dynamic:

Advection: flux form - 3rd order UBS

No explicit diffusion

EEN energy & enstrophy scheme (with masked averaging of e3t divided by the sum of mask)

Vertical physic:

Vertical mixing: GLS

adaptive-implicit vertical advection (Shchepetkin 2015)

Activate NEMO observations operator

SST:
ODYSEA
L3S 0,1°

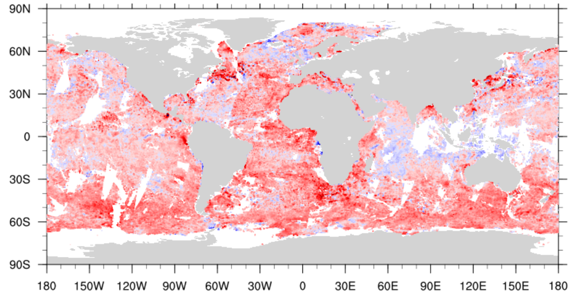

In situ :
Coriolis

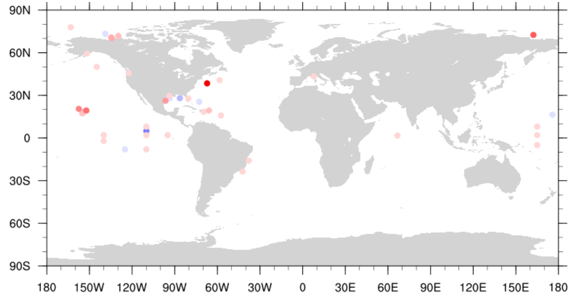

SLA:
S3A,C2,J3,
ALT,H2Y

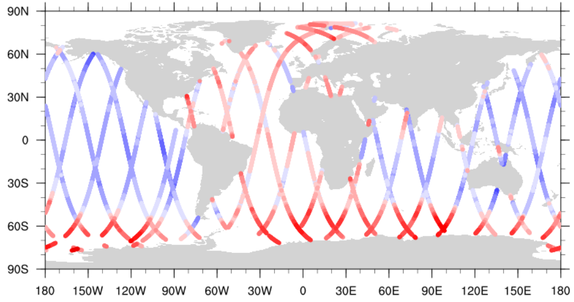

Good way to provide model-observations comparison

But… out of memory….

- **Why?** Each processor:

load the **data on all the domain**: -
select which observations are inside its MPI sub-domain
make the colocalisation

An example of data volume for the first week:
SST (1Gb) + SLA (100Mb) + INSITU (3.5 Gb) : 4.5G for each processor

- **Consequence:** Peak of memory during initialisation => Need to depopulate : 1 proc on 2

- **Solution to avoid this problem ?**

✓ Don't need to read observation data itself, positions are enough (but need to concatenate after)
✓ Preprocess obs dataset: split on MPP sub-domains
✓ Read obs day per day instead loading all the full run period dataset at initialisation

MERCATOR OCEAN INTERNATIONAL

Run setup

- Run starts at 20170101
- 3 full months of simulation

- Time step: 120s

- Twin runs with global ¼° (ORCA025) and global 1/12° (ORCA12) performed

- Computer: ECMWF CRAY CCA (Lustre file system, One node: 36 cores , 128 Gb memory)

11783 NEMO subdomains => **18 processus per node (depopulate)** => 655 nodes
240 Xios servers => 2 processus per node (depopulate) => 120 nodes
Total => 775 nodes

Results

1 hour sea surface velocities after 3 months

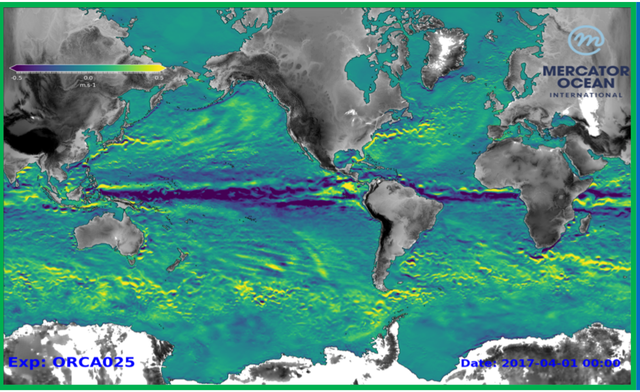

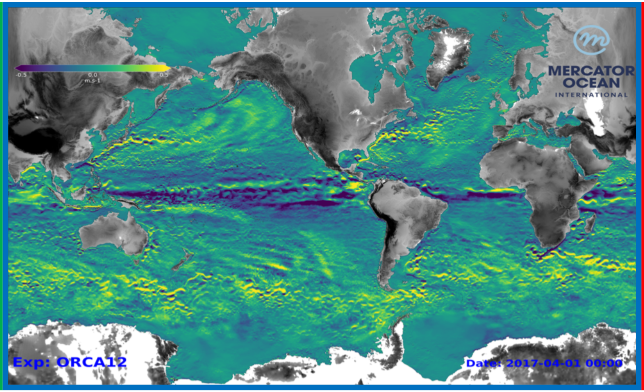

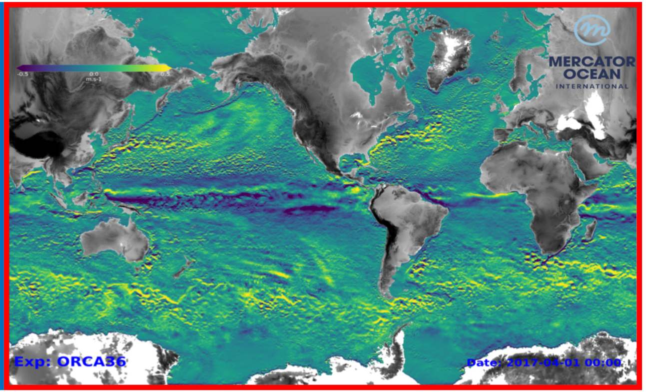

¼° 1/12° 1/36°

Results

1 hour sea surface relative vorticity after 3 months

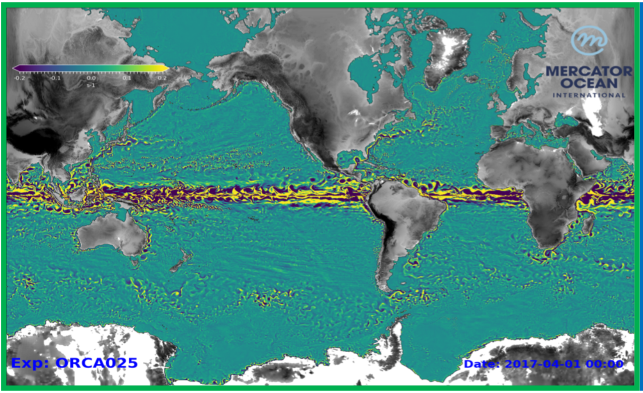

¼°

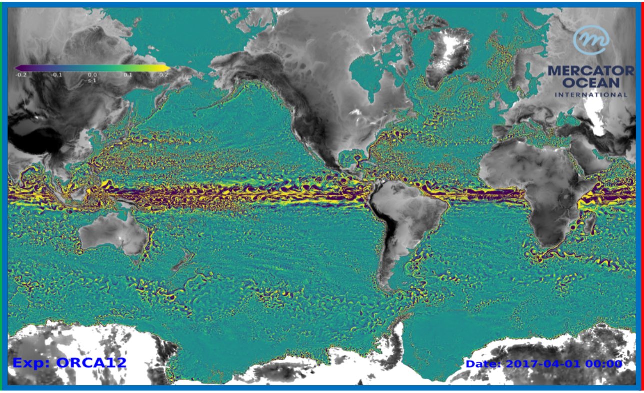

1/12°

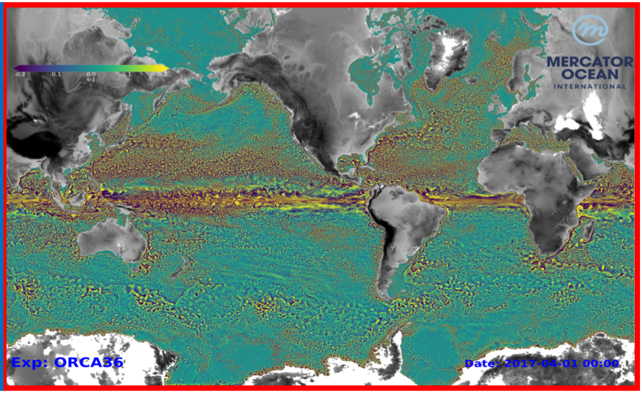

1/36°

Results

1 hour sea surface relative vorticity after 3 months

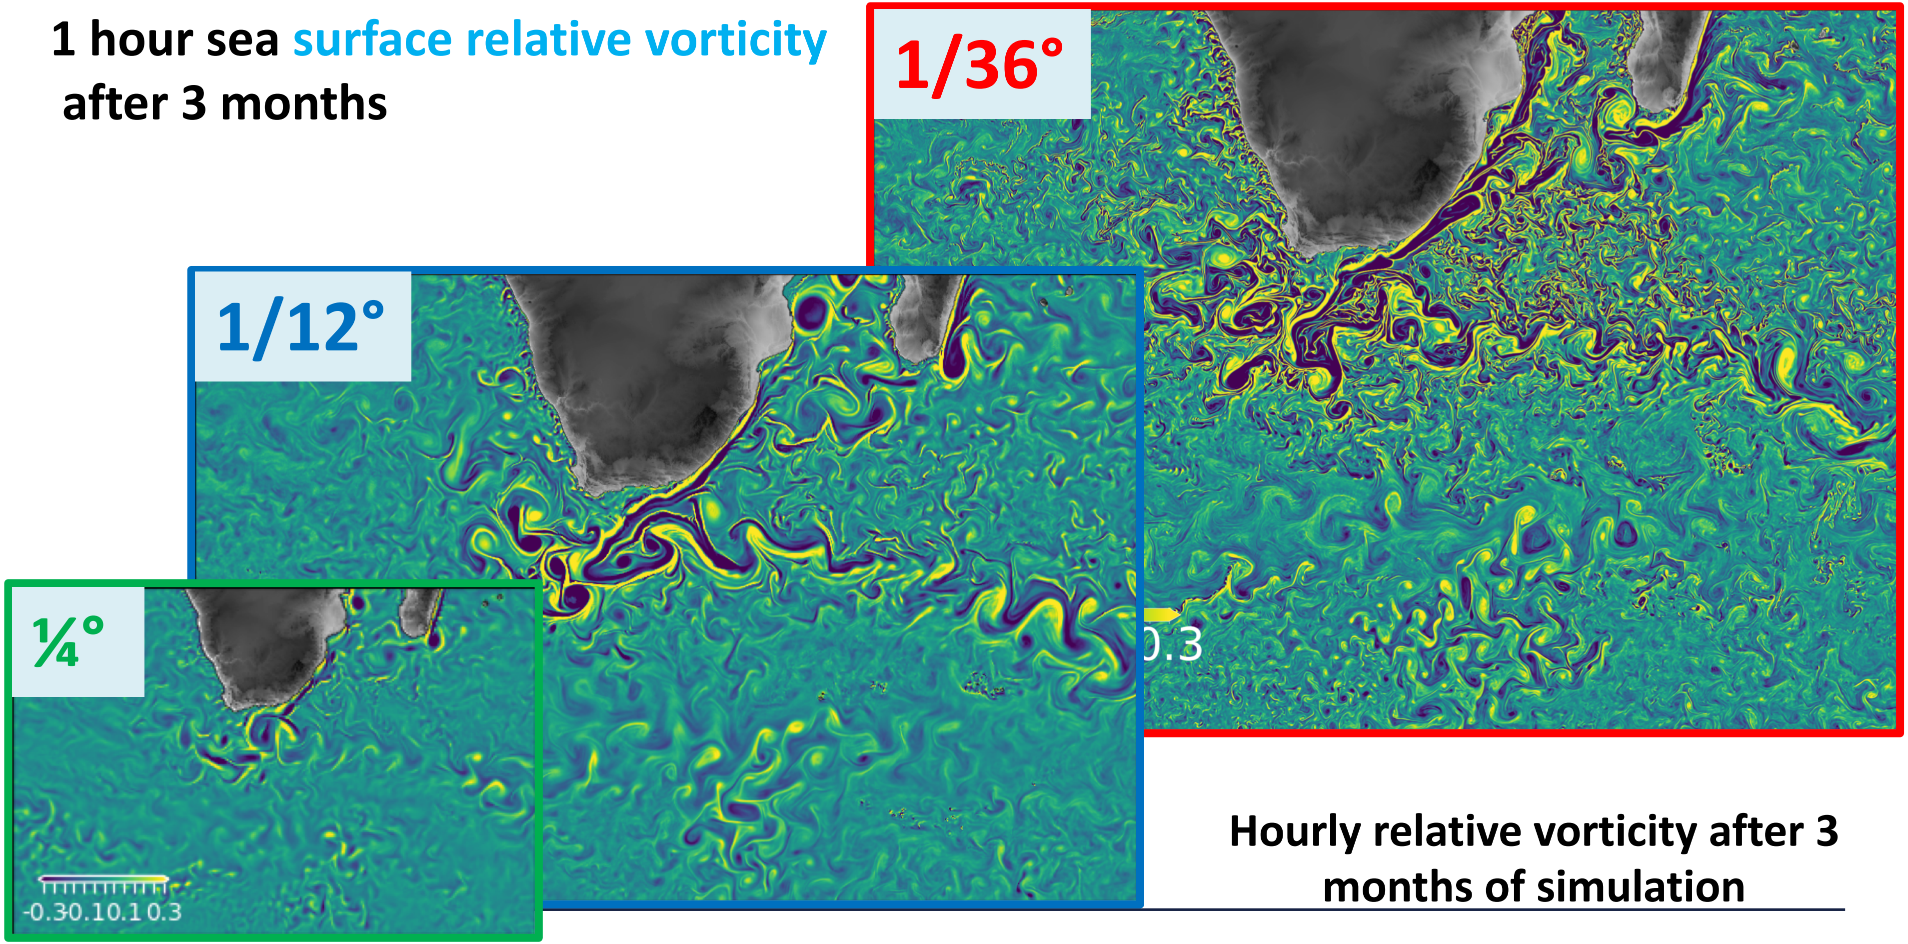

Hourly relative vorticity after 3 months of simulation

Results

Monthly SSH variance from hourly fields after 3 months

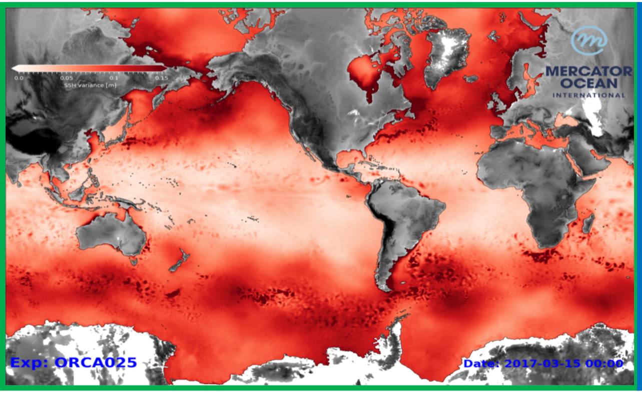

¼°

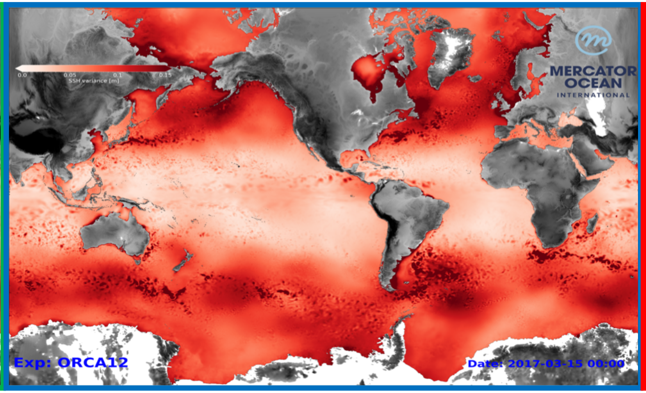

1/12°

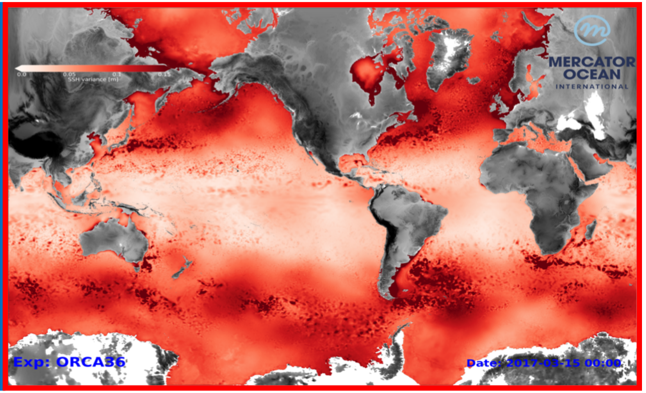

1/36°

Results

2017-01-15 to 2017-04-08
KE power spectra

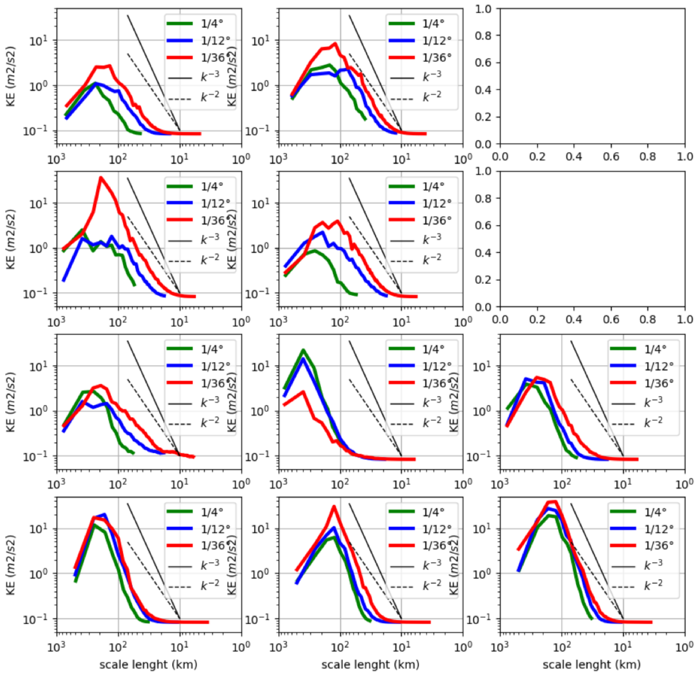

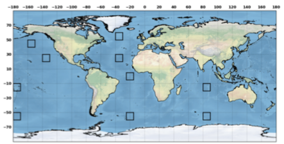

2017-01-15 to 2017-04-08
SSH power spectra
comparison to Jason3

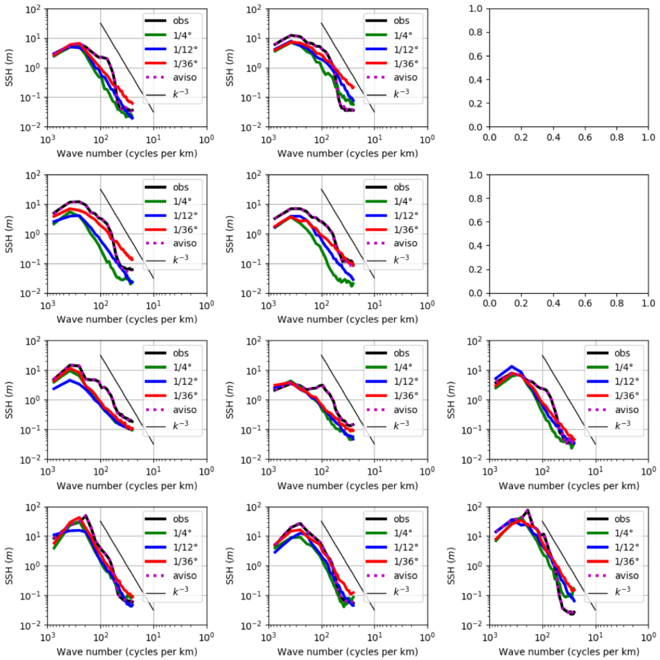

Global 1/36° HPC performances analysis

Work realized by BSC with CMEMS contract:

« 87-GLOBAL-CMEMS-NEMO: EVOLUTION AND OPTIMISATION OF THE NEMO CODE USED FOR THE MFC-GLO IN CMEMS »

ORCA36 scalability
(no forcing, no sea-ice, no outputs)

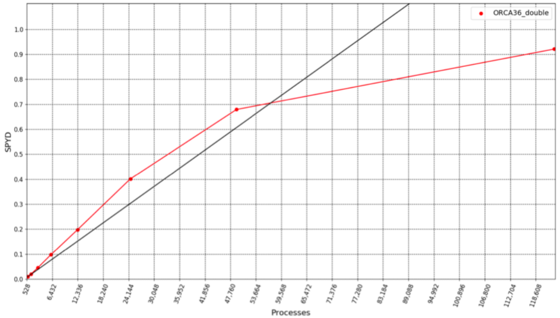

Scalability tested up to 122 000 cores
Good scalability up to 50 000 cores

ORCA36 scalability
(no forcing, no sea-ice, no outputs)
Simple vs double precision

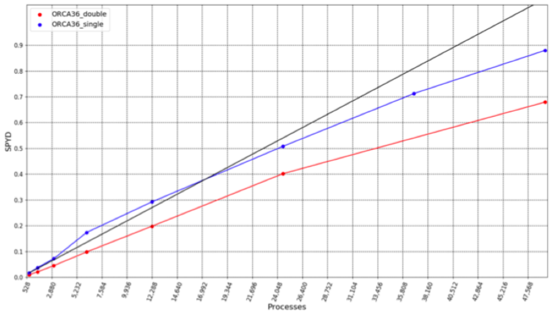

Simple precision improve scalability compared to double precision

ORCA36 scalability
(no outputs)
Sea-ice model impact

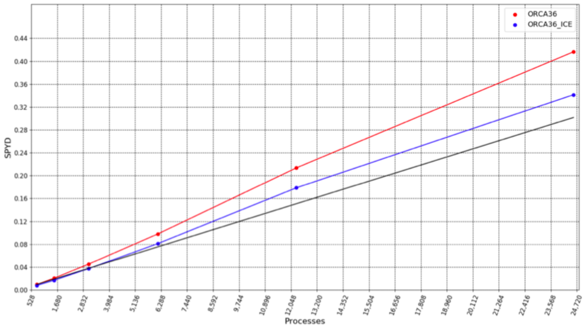

Global 1/36° HPC performances analysis

performance metrics

Number of processes	**768**	**3,072**	**6,144**
Parallel efficiency	93.72	85.74	82.65
Load balance	97.42	92.4	92.74
Communication efficiency	96.2	92.79	89.12
Computation scalability*	100	130.25	110.98
Global efficiency*	**93.72**	**111.67**	**91.72**
IPC scalability*	100	123.47	118.24
Instruction scalability*	100	102.63	94.69
Frequency scalability*	100	102.78	99.12
Speedup*	**1.00**	**4.77**	**2.14**
Average IPC	0.29	0.35	0.42
Average frequency (GHz)	2.09	2.15	2.13

** These values use the column on their left as reference.*

- parallel efficiency decreases with the scale
- gains in computational efficiency due to the increase in instructions per cycle (IPC).
  better exploitation of the shared resources
  faster memory operations

Proportion of useful instructions
(those not involved in MPI communication)

Function	768	3,072	6,144
divhor_m..div_hor_	0.60%	0.61%	0.61%
step_mp_stp_	0.13%	0.16%	0.19%
sbcmod_mp_sbc_	0.05%	0.05%	0.04%
usrdef_s..sbc_oce_	0.49%	0.48%	0.45%
lib_fort.._sum_2d_	0.23%	0.23%	0.23%
eosbn2_mp_rab_3d_	3.52%	3.43%	3.21%
eosbn2_mp_bn2_	0.82%	0.82%	0.84%
zdfphy_m..zdf_phy_	0.20%	0.27%	0.35%
zdfdrg_m..zdf_drg_	0.12%	0.11%	0.10%
zdfsh2_m..zdf_sh2_	0.79%	0.79%	0.78%
zdfgls_m..zdf_gls_	**6.08%**	**6.02%**	**5.86%**
zdfmxl_m..zdf_mxl_	0.27%	0.29%	0.30%
sshwzv_m..ssh_nxt_	0.09%	0.09%	0.11%
domvvl_m.._sf_nxt_	1.86%	1.98%	2.12%
sshwzv_mp_wzv_	0.35%	0.37%	0.45%
eosbn2_m..itu_pot_	1.37%	1.34%	1.26%
zpshde_m..zps_hde_	0.18%	0.17%	0.17%
eosbn2_m..situ_2d_	0.04%	0.04%	0.04%
dynadv_u..adv_ubs_	3.95%	4.07%	4.34%
dynvor_m..vor_een_	1.35%	1.34%	1.42%
dynhpg_m..hpg_sco_	0.75%	0.73%	0.70%
dynspg_t.._spg_ts_	**20.34%**	**21.21%**	**22.63%**
dynzdf_m..dyn_zdf_	2.00%	2.11%	2.24%
trasbc_m..tra_sbc_	0.01%	0.01%	0.01%
traqsr_m..tra_qsr_	**16.59%**	**15.84%**	**14.61%**
traadv_m..tra_adv_	0.26%	0.28%	0.35%
traadv_f..adv_fct_	2.95%	3.05%	3.18%
traadv_f.._nonosc_	4.54%	4.47%	4.40%
traldf_l..ldf_lap_	0.90%	0.97%	1.05%
trazdf_m..tra_zdf_	0.08%	0.09%	0.11%
trazdf_m..zdf_imp_	1.06%	1.14%	1.22%
tranxt_m..tra_nxt_	**18.50%**	**17.71%**	**16.63%**
dynnxt_m..dyn_nxt_	2.66%	2.86%	3.24%
sshwzv_m..ssh_swp_	0.01%	0.01%	0.01%
domvvl_m.._sf_swp_	2.33%	2.42%	2.49%
stpctl_m..stp_ctl_	4.52%	4.41%	4.24%

- most of these instructions are Load and Stores and not floating point operations.
- Impact of code writing or compiler optimization?

Summary

- A configuration is existing

- It is running with NEMO 4…

  with reasonable performance (for development phase)

  with NEMO observations operator (but need more memory)

- Right way to provide a configuration running on NEMO 4 for IMMERSE and ESIWACE2

- Good feedback for BSC study on NEMO4/ORCA36 HPC performances

MERCATOR OCEAN INTERNATIONAL

What happens next?

- NEMO 4 version upgrade (NEMO 4.1 « IMMERSE » version)
- Longer run (10/2012 => 2020)
- Forcing: ERAinterim => IFS

- Extension of domain southward: add under ice shelf seas
- Add tidal forcing (for tidal internal waves)
- Uses Atmospheric Boundary Layer: dynamical downscalling of atmo data to model resolution
- Improve model parametrization tunning

- Increase ouput volume (1 hour 3D outputs)
- Increase MPI domain splitting
- Switch to new Meteo France BULL and/or ECMWF computers
- Ask for a PRACE project with IGE/Ocean-Next/CMCC/BSC

End