Implementing a parallel version of a variational scheme in a global assimilation system at eddy-resolving resolution

Andrea Cipollone[1] *(andrea.cipollone@cmcc.it)*, Andrea Storto[1,2], Simona Masina[1]

1. ODA Division, Centro Euro-Mediterraneo sui Cambiamenti Climatici, Bologna, Italy; 2. CNR-ISMAR, Rome, Italy

Assimilation of mesososcale variability

The representation of mesoscale variability at global scale, represents a long-standing problem in the ocean forecast and reanalysis community. Besides the improvements of ocean general circulation models, a day-by-day realistic representation of such variability is triggered by the capability of assimilation systems to correct the ocean states at the same scale. This work extends a 3dvar system designed for oceanic applications (Storto et al 2011), to cope with global eddy-resolving grid and dense observational datasets in a hybrid-parallelized environment (Cipollone et al., 2020). The scheme is included in the global eddy-resolving ocean forecasting system at the resolution of 1/16 that is in place at CMCC.

OceanVar assimilation scheme

OceanVar is an assimilation scheme that belongs to the family of three-dimensional variational (3DVAR) analysis systems specifically designed to deal with coastlines and bathymetry for oceanography applications (Dobricic & Pinardi, 2008; *Storto et al 2011)*. At regional level, the system is used in Mediterranean and Black Sea operational forecasting systems available from the CMEMS catalogue (Clementi et al 2017). At global level, OceanVar has been employed in several global Reanalysis products (C-GLORS series: Storto 2016,2019). The scheme has been parallelized to address high-res grid, the scalability is favoured by the implementation of a weak-constrained formulation of the continuity requirement across artificial boundaries (Fig. 1). Such formulation limits possible boundary discontinuities to be less than a prescribed value ($\boldsymbol{\varepsilon}$)

Strong constraint (exact solution): $\boldsymbol{v_1}[\cap_{1j}] - \boldsymbol{v_j}[\cap_{j1}] = 0$ → Weak constraint (approximate solution): $\boldsymbol{v_1}[\cap_{1j}] - \boldsymbol{v_j}[\cap_{j1}] = \varepsilon_{1j}$

Minimized by a new term in the cost function

$$\boldsymbol{J}_1[\boldsymbol{v_1}] = \boldsymbol{J}[\boldsymbol{v_1}] + \boldsymbol{J}_{\mathrm{bd}}[\boldsymbol{v_1}] = \boldsymbol{J}[\boldsymbol{v_1}] + \frac{1}{2} \sum_{\mathrm{j\in Neigh.of\ 1}} (\boldsymbol{v_1}[\cap_{1j}] - \boldsymbol{v_j}[\cap_{j1}]) \boldsymbol{Q}_{1j}^{-1} (\boldsymbol{v_1}[\cap_{1j}] - \boldsymbol{v_j}[\cap_{j1}])$$

Standard formulation

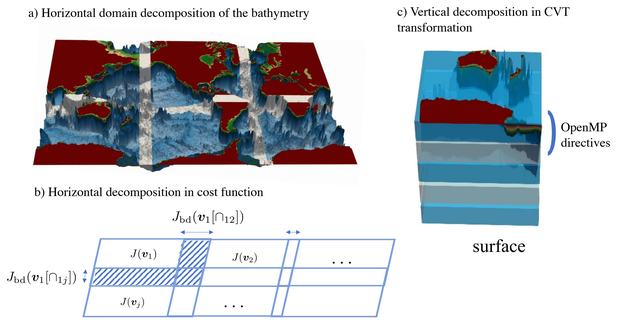

Figure 1: Schematic representation of the hybrid domain decomposition in OceanVar: Panel a) shows the horizontal domain decomposition with overlapping boundaries (gray color) where communications are driven by MPI. The corresponding cost function decomposition is seen in Panel b). Panel c) represents the vertical decomposition applied in the Control Vector Transformation (Dobricic & Pinardi, 2008) controlled by shared-memory OpenMP directives.

Results from the new parallelized scheme

Accuracy: The new boundary term generates a "crossing flow" throughout the boundaries of horizontal subdomains that reduces the discontinuity in the overlapping area. The magnitude of such flow mainly depends on the ratio between the discontinuity and error covariance matrix (**Q**). By imposing that such error is well below a physical threshold (i.e. background error), the **Q** becomes of diagonal form in the control space :

$$\boldsymbol{Q}_{ij} = q\boldsymbol{I}_{ij} \quad q < 1$$

Where q can be tuned in order to promote/penalize such flow: reducing q, the solution is driven toward the exact one (Fig.2-3)

Scalability: The new boundary term allows to move the MPI communication outside the CVT transformation and in the calculation of cost function. This drastically reduces the number of total MPI calls. The reduction of execution time for one iteration in the global configuration is shown Fig. 4, the computational cost of MPI-IO processes and the speed-up of the second layer of parallelization (OpenMP directives) is shown

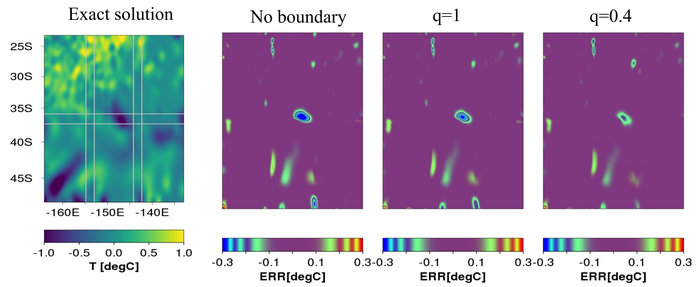

Figure 2: The exact incremental solution in one specific region (left panel) is compared w.r.t. the solution from a 6 subdomains and different values of q (difference from the the exact solution is shown).

Figure 3: Percentual gain (w.r.t. the exact solution) for different decomposition and boundary widths (in number of points).

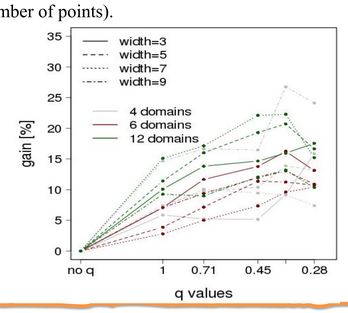

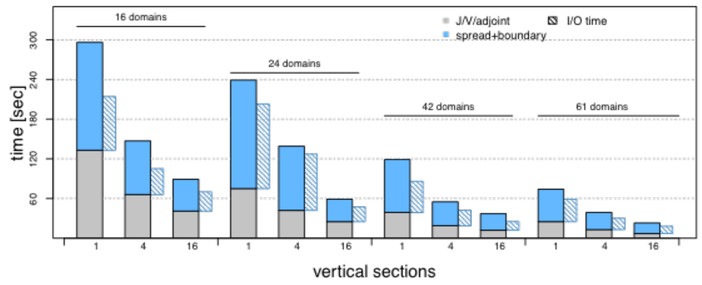

Figure 4: Execution time for one iteration and different hybrid domain decompositions (global configuration). The blue bars refer to overall communication time (MPI + unbalance), the grey bar shows the computing time for calculation of J and CVT transformation. Shaded narrow bars correspond to an estimate of the MPI I/O time only

The role of resolution at assimilation level

The present availability of dense satellite datasets suggests that the use of coarse grid at assimilation level can be a suboptimal choice for eddy-resolving forecasting systems. To quantify the impact of high-res grid at assimilation level, a twin experiment is performed that considers two independent cycles of model integration and assimilation for one month. The model configuration is identical with 1/16 deg resolution globally and 98 vertical levels (Iovino et al. 2016) while at the assimilation step, the first experiment exploits a coarser grid (¼ deg and 50 vertical level) while the second uses the same grid of the model. Same observational dataset is used: CMEMS NRT INSITU and NRT along-track L3 SLA data dataset, L3 Advanced Very High-Resolution Radiometer (AVHRR) and L3 Advanced Microwave Scanning Radiometer 2 (AMSR2).

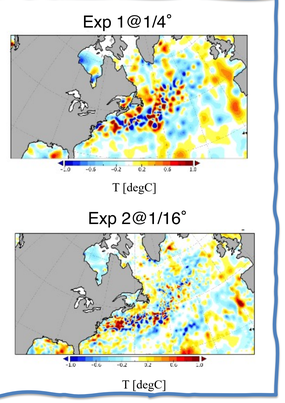

Figure 5: Typical SST increments at surface for the two experiments

Results from high-res assimilation

Experiment 2 (assimilation in high-res grid) performs significantly better in the case of satellite data (altimetry and sea surface temperature from infrared sensors) with an improvement up to 15% in the RMSE. This is true especially at mid- and high latitudes where the Rossby radius is no longer resolved by ¼° (not shown). Improvements around 8-10% can be seen in the RMSE against SLA observation while no significant impact is shown when compared sparse in situ data (Figure 6).

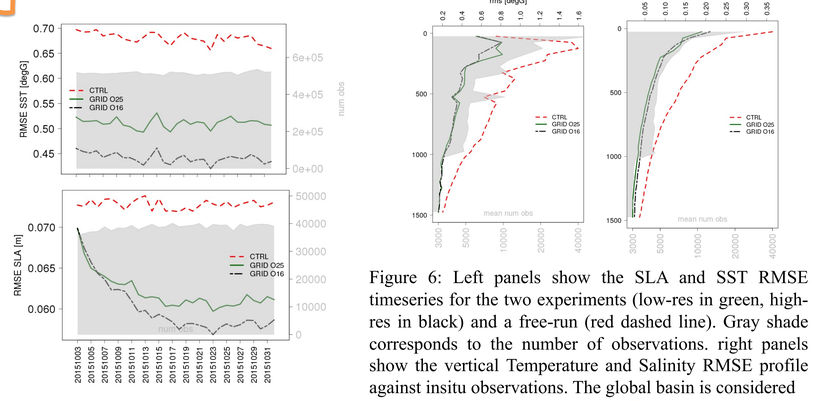

Figure 6: Left panels show the SLA and SST RMSE timeseries for the two experiments (low-res in green, high-res in black) and a free-run (red dashed line). Gray shade corresponds to the number of observations. right panels show the vertical Temperature and Salinity RMSE profile against insitu observations. The global basin is considered

MAIN REFERENCES: *Cipollone, A., Storto, A., & Masina, S. (2020), JTECH, 37(10), 1865-1876 ,DOI:10.1175/JTECH-D-19-0099.1. Storto, A., Dobricic, S., Masina, S., and Di Pietro, P. (2011), Mon. Weather Rev., 139, 738–754; Iovino D., Storto A., Masina S., Cipollone A. & Stepanov V. (2016), Geosci. Model Dev. , 9, 2665.; Clementi, E., P. Oddo, M. Drudi, N. Pinardi, G. Korres, and A. Grandi, (2017), 67 (10), 1293–1312, doi:10.1007/s10236-017-1087-7; Storto, A., and Masina S. (2016b), ESSD, 8 (2), 679–696. Storto, A., Masina S., Simoncelli S. et al. (2019), Clim Dyn 53: 287.*