SPECIAL PROJECT PROGRESS REPORT

Progress Reports should be 2 to 10 pages in length, depending on importance of the project. All the following mandatory information needs to be provided.

Reporting year	2015-2016
Project Title:	Optimization of the OceanVar oceanographic data assimilation system for high-resolution applications …………………………………………………………….……
Computer Project Account:	spitstor
Principal Investigator(s):	Andrea Storto …………………………………………………………….……
Affiliation:	CMCC
Name of ECMWF scientist(s) collaborating to the project (if applicable)	Not Applicable
Start date of the project:	27/06/2015
Expected end date:	30/06/2017

Computer resources allocated/used for the current year and the previous one
(if applicable)
Please answer for all project resources

		Previous year		**Current year**	
		Allocated	Used	Allocated	Used
High Performance Computing Facility	(units)	150000	3320	150000	0
Data storage capacity	(Gbytes)	500	1	500	0

Summary of project objectives
(10 lines max)

The project aims at optimizing the OceanVar variational ocean data assimilation scheme for high-resolution applications. We aim at improving the memory consumption and computational resources required by OceanVar and increasing its scalability focussing on the optimization of the horizontal operator that model horizontal background-error correlations. An improved version of the recursive filter operator that model horizontal correlations will be implemented in global high-resolution (eddy-resolving) configuration. Such an improved operator will take advantage of a rigorous newly formulated mathematical framework, thanks to which we are able to avoid the use of extension zones in the west-east wrapping and inland ghost points to impose cyclic and land-sea lateral boundary conditions, respectively.

Summary of problems encountered (if any)
(20 lines max)

Contrary to what was expected, we have been using local computing resources for testing the new filter formulation that led to the scientific publication of Mirouze and Storto (2016) (see below the reference). This is because we focus on simple test that did not involve massive parallelization. We plan to rely on the ECMWF HPC facility in the follow-up of the project.

Summary of results of the current year (from July of previous year to June of current year)

A new formulation of the recursive filter has been set up and implemented in OceanVar, and successfully tested in a ¼ degree configuration of the global ocean data assimilation system used at CMCC. The scientific achievements are documented in details in the article reported below. Cyclic and land/sea boundary conditions have been reformulated analytically through the inclusion of a corrective term corresponding to the use of Neumann or Dirichlet boundary conditions. This strategy replaced the former strategy, where an extension zone with duplicated observations was used to impose cyclic conditions, and “ghost points” (i.e. inland extension of the computational domain close to shoreline) were used to handle the land/sea boundary conditions. The new strategy, implemented and tested in the ¼ degree global configuration of OceanVar, not only proves more accurate with respect to the previous one, but also improves the memory consumption and the computational time required by OceanVar.

List of publications/reports from the project with complete references

Mirouze, I. and Storto, A. (2016), Handling boundaries with the one-dimensional first-order recursive filter. Q.J.R. Meteorol. Soc.. doi: 10.1002/qj.2840

Summary of plans for the continuation of the project
(10 lines max)

During the continuation of the project, we plan to extend the approach used for specifying exact analytical boundary conditions for the recursive filter to the problem of parallelizing the recursive filter.
Indeed, it is possible to define boundary conditions between different domains (processes) that the data assimilation problem is divided into, in a way similar to what has been done for the lateral boundary conditions. Such a strategy would require a relatively small MPI communication compared to parallelization strategies usually implemented in grid-point recursive filter (e.g. transposition of the field before and after 1D application of the filter).
This approach will be compared with the one currently implemented in OceanVar, which consists in a domain decomposition with large overlap areas (several halo points) and no MPI communication. It is expected that the new approach will improve the flexibility of the scheme and its scalability (reducing the halo points), and the memory consumption, which is crucial for very high-resolution global ocean applications.