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
Project Title:	High-resolution climate prediction with EC-Earth
Computer Project Account:	SPESICCF
Principal Investigator(s):	Francisco J. Doblas-Reyes
	……………………………………………………….……
Affiliation:	Institut Català de Ciències del Clima (IC3)
Name of ECMWF scientist(s) collaborating to the project (if applicable)	……………………………………………………….…… ……………………………………………………….……
Start date of the project:	01/01/2015
Expected end date:	31/12/2016

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)			37,050,000	38,499,421
Data storage capacity	(Gbytes)			12,150	3,431

Summary of project objectives

(10 lines max)

The main project objective is to study the impact of increased resolution in both the ocean and the atmosphere on seasonal prediction quality. To this end, we plan to compare seasonal predictions performed with the high- and standard-resolution configurations of EC-Earth. Motivated by the important role of the ocean in seasonal prediction, we also intend to investigate the impact of ocean initialisation by using initial conditions from three different ocean reanalysis datasets. An important computational cost is associated with each of these experiments and, more generally, to all types of seasonal-to-decadal climate prediction simulations. It would be therefore desirable to distribute all simulations across several HPC platforms, but this can only be done if differences in hardware/software do not introduce differences in the results. No one had ever investigated the reproducibility of climate simulations across different platforms. Hence, we made it a side objective of this project.

Summary of problems encountered (if any)

(20 lines max)

In order to run our simulations at ECMWF we use Autosubmit, which is a Python-based tool developed at IC3 to submit, manage and monitor climate simulations on HPC platforms. Autosubmit uses ecaccess in order to interactively access the ECMWF (or any other) computing facilities. While using Autosubmit to launch the simulations, we encountered a problem in making the ecaccess submission command to successfully deal with the remote submission of the simulation job script in cca. We contacted by email the ECMWF support, who tracked down the problem and found out that a user alias for the "ls" command on cca caused the ecaccess submission commands to fail. ECMWF support informed us that they plan to implement a more permanent solution to allow users having aliases of sensitive commands without affecting ecaccess. We greatly acknowledge the assistance of ECMWF support.

In the original proposal we estimated that the cost of each month of the simulation is 28,000 SBU. The initial estimate was based on the official EC-Earth3.1 with the T511-ORCA025-LIM3 configuration run on the first stages of cca. The actual cost in SBU of each month of simulation we have run lately, however, turns out to be 60,000 SBU. This means that to complete 25% of the ORCA025L75 set of forecast with GloSea5 we used all the resources scheduled for 2015 (38,499,421 SBU). Therefore, to complete the remaining 75% of this experiment, which is scheduled for 2015, we estimate that an additional 115,500,000 SBU would be needed, which is not an acceptable solution. We are working with user support to find out what could explain this unexpected behaviour, find a solution with a cost similar to the one obtained during the scaling exercise at the time of submitting the proposal and starting the experiments, and proceed with the request of a sensible amount of additional resources.

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

At the time of writing this report (end of June 2015), seven start dates of seasonal hindcasts have been run (experiment a00p in the table below). These hindcasts have been compared to existing simulations to assess the benefits of resolution and initial conditions on seasonal forecast quality. Note that the standard resolution configuration (T255L91-ORCA1L46) is initialized with GLORYS2V1 interpolated to the ORCA1L46 grid (experiment m04c).

	Four-month hindcasts, 1993-2000, initialized 1st May, five members		
Experiment name	a00p	m04p	m04c
Oceanic resolution	**ORCA025L75**		**ORCA1L46**
Atmospheric resolution	**T511L91**		**T255L91**
Ocean initial conditions	**GloSea5**	**GLORYS2V1**	
Sea ice initial conditions	**GLORYS2V1**		

Figure 1a shows the model SST bias for boreal summer (JJA) for the available hindcasts, i.e. 1993-2000, for the two sets of high-resolution hindcasts, a00p and m04p. The bias has been computed with respect to HadISST observational data. Figure 1b shows the differences in SST between a00p and m04p to highlight the impact of the different ocean reanalyses taken as initial conditions. The main patterns of long-standing SST EC-Earth3 biases are present in both sets of forecasts, namely the cold equatorial Pacific bias, the warm bias in the equatorial and south east Atlantic and the cold bias in the northern hemisphere. There are two notable improvements in a00p: the equatorial biases, i.e. the cold bias in central and eastern equatorial Pacific and the warm bias in tropical and south east Atlantic, which are present in most climate models including EC-Earth, are reduced, thereby improving the simulations of the equatorial climate. Further improvements in the west coast of Australia and in the Southern Ocean also indicate some first promising results of using the GloSea5 initial conditions.

Figure 2 shows the model bias for the sets of forecasts a00p and m04c exactly as in Figure 1. The two sets differ in the resolution of atmosphere and ocean components, and in the ocean initial conditions. Similarly as in the case above, the high-resolution experiment a00p initialized by GloSea5 shows smaller SST biases in the equatorial Pacific and Atlantic Oceans. The pattern of changes is very similar to the one shown in the comparison with m04p (i.e. the patterns in Figs. 1b and 2b are close to each other).

Figure 3 shows the skill of the SST averaged in the Niño3.4 region for the boreal summer for the hindcasts initialized over 1993-2000 for the three different experiments. The correlation of the ensemble mean has been computed with respect to three different observational datasets (HadISST, ERAint and ERSST). Overall, the experiment a00p has the best skill when compared to the other two experiments, followed by the high-resolution GLORYS2V1 experiment m04p, whereas the standard-resolution experiment m04c has the worst performance among the thee experiments. The figure also illustrates the large observational uncertainty in the skill estimates as the correlation changes substantially depending on the observational reference used.

The three experiments described above were conducted on two different machines: ECMWF for a04p and Marenostrum3 at Barcelona Supercomputing Center (BSC) for m04p and m04c. To ensure that the experiments are comparable, it is required that all experimental conditions are identical except the initial conditions and/or the resolution. However, it is known that difference in motherboards, compilers, number of processors or optimization flags do induce round-off errors and make bitwise reproducibility impossible (Thomas et al., 2002; Senoner et al., 2008; Hong et al., 2013). Because climate is defined through statistical distributions rather than single realizations, this non-bitwise reproducibility is not an issue as long as ensembles that are run on different machines are undistinguishable from each other, statistically speaking. To validate this hypothesis, we ran the same 5-member, 20-yr control simulation on three different platforms with EC-Earth3 in its standard configuration (T255L91-ORCA1L46). One of these platforms was cca; the two others were run on Ithaca (the IC3 cluster) and BSC’s Marenostrum3. All other parameters than HPC settings were forced to be identical (initial conditions, forcings, model code) thanks to the possibilities offered by the Autosubmit software described above. The results obtained so far show differences in the Southern Ocean (Fig. 4) that exceed the inter-member spread. In other words, in some regions, running the same code on two machines gives systematically more differences than running two members on the same machine. Note though that no significant differences were found in the tropics in these simulations. It is therefore unlikely that these machine-to-machine differences affect the conclusions presented above on the role of resolution and initial conditions on skill in the tropics, especially at short time scales.

These preliminary results show the potential benefit of alternative initial conditions to reduce the development of biases in EC-Earth. They also show the relevance of running high-resolution simulations to increase correlation skill, although the validity of these two results should be re-assessed when more hindcasts will become available. In summary, they suggest that improvements in the initial conditions provide a better mean climate estimate, while increases in resolution provide an improvement in Niño3.4 SST skill. This motivates the continuation of existing simulations and the design of additional simulations that will use the ORAS5 and ORAP5 ocean reanalysis as ocean initial conditions.

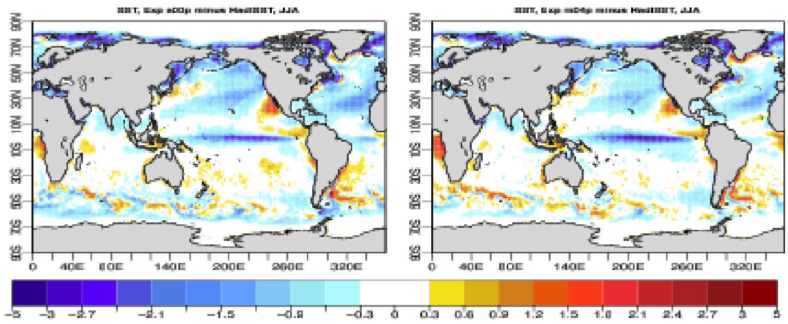

(a) Model bias in SST in forecasts using the HR configuration, initialized by GloSea5 (left) and GLORYS (right).

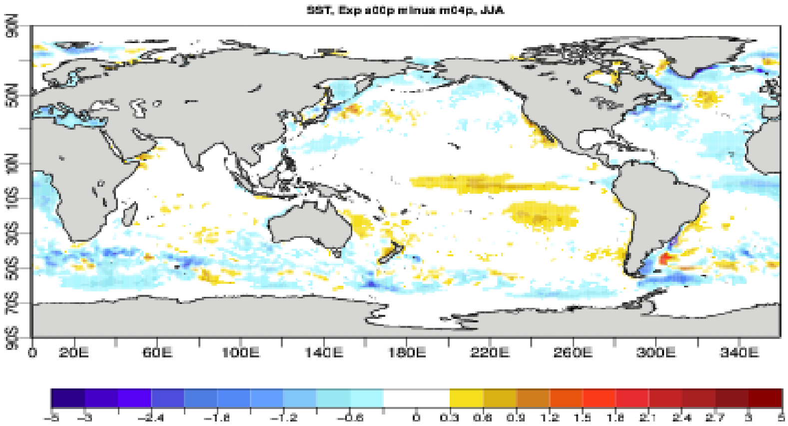

(b) Difference between the two panel plots in (a).

Figure 1: (a) Model bias in SST during summer (JJA), defined as the difference between the seasonal prediction hindcasts and observations from HadISST (Reyner et al., 2003) for 1993-2000, for simulations using the same model setup but are initialized by GloSea5 (left) and GLORYS2V1 (right). (b) Difference in SSTs between hindcasts initialized by GloSea5 and hindcasts initialized by GLORYS2V1.

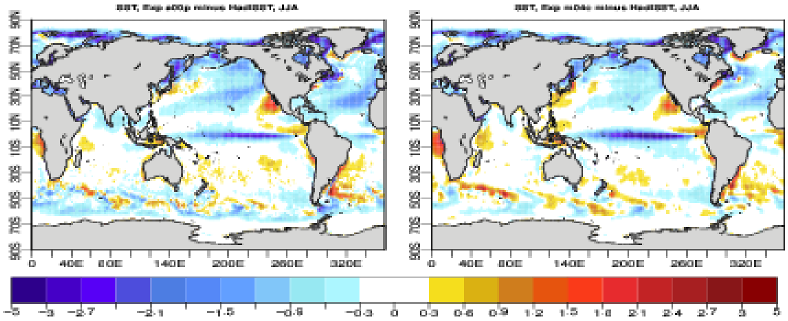

(a) Model bias in SST in forecasts using the HR configuration and initialized by GloSea5 (left), and forecasts using the SR configuration and initialized by GLORYS (right).

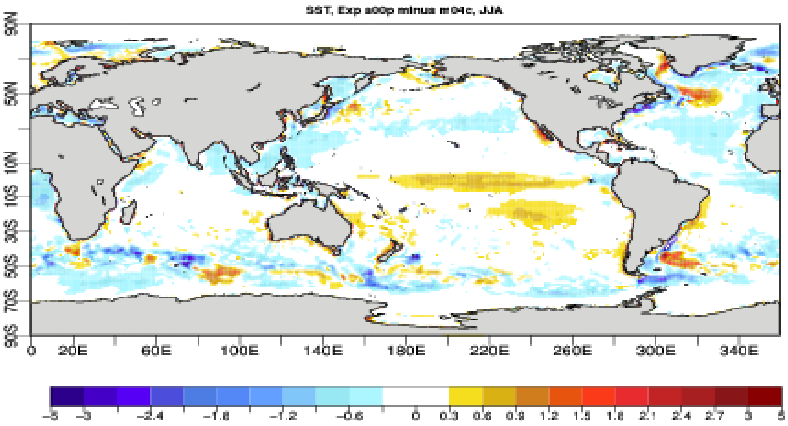

(b) Difference between the two panel plots in (a).

Figure 2: (a) Model bias in SST during summer (JJA), defined as the difference between the seasonal prediction hindcasts and observations from HadISST (Reyner et al., 2003) for 1993-2000, for simulations that use the HR configuration and are initialized by GloSea5 (left), and for simulations that use the SR and are initialized by GLORYS2V1 (right). (b) Difference in SSTs between HR hindcasts initialized by GloSea5 and SR hindcasts initialized by GLORYS2V1.

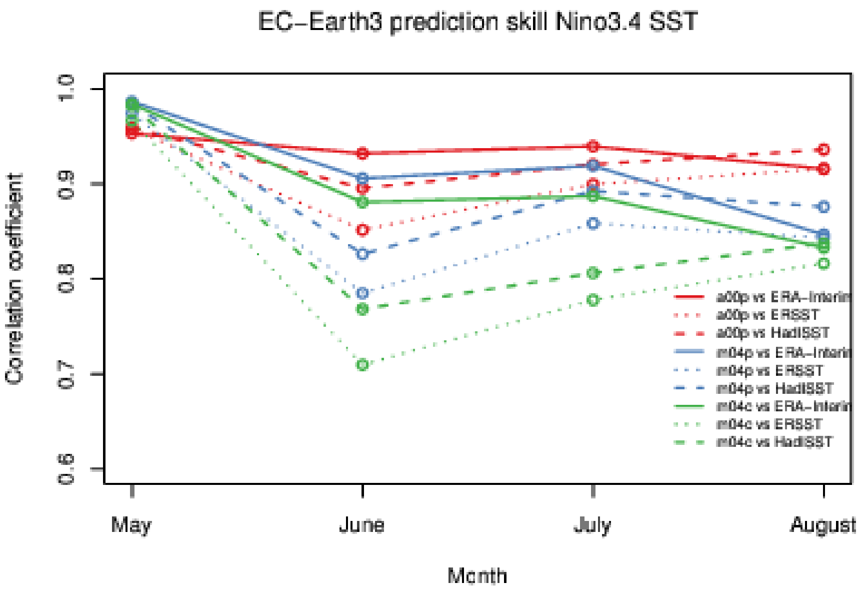

Figure 3: Prediction skill in the SSTs in the Niño3.4 region for JJA shown as the correlation of the ensemble mean prediction with three different observational datasets (HadISST, ERAinterim and ERSST) for the three different experiments a00p, m04p and m04c.

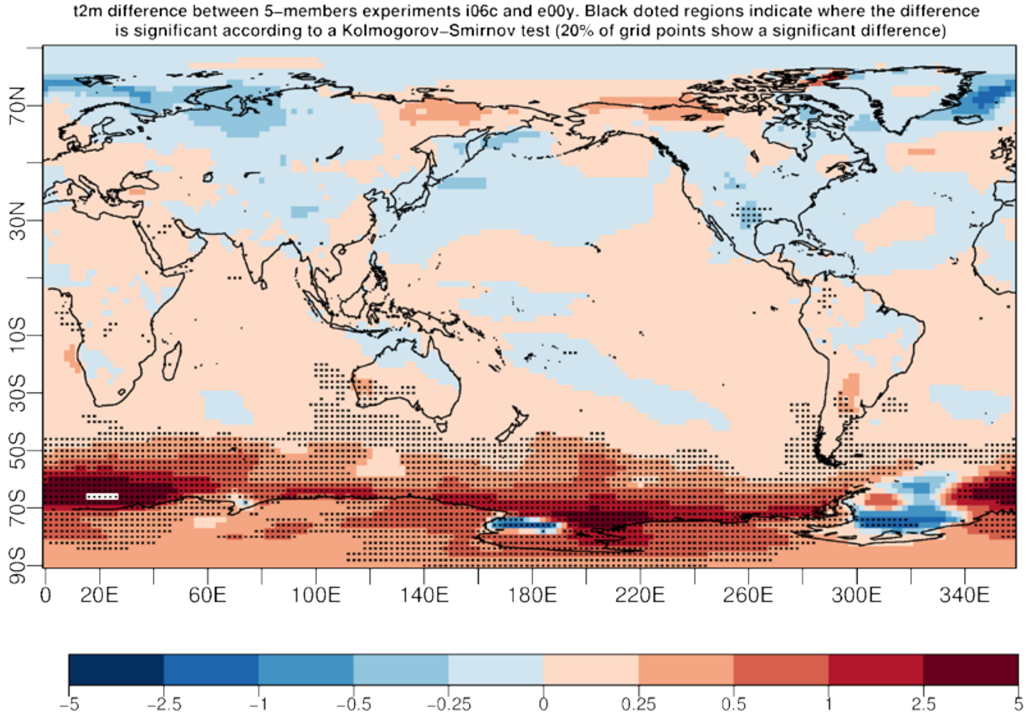

Figure 4: Mean difference in average 2m air temperature (°C) over a 20-year, 5-member simulation performed on Ithaca (IC3) and cca (ECMWF) under pre-industrial conditions. The dots indicate regions where the two 5-member ensembles were found to be statistically different from each other following a Kolmogorov-Smirnov test at a significance level of 5%.

List of publications/reports from the project with complete references

None yet.

Summary of plans for the continuation of the project

(10 lines max)

About 25% of the ORCA025L75 set of hindcasts with GloSea5 is completed at the time of writing. We plan to continue and complete this set of forecast in 2015. An additional two sets of forecasts are still planned in this special project:

1. ORAP5 ORCA025L75, to be performed between 2015 and 2016
2. ORAS5 ORCA025L75, which will be available in 2016 and should allow the simulation to be performed in 2016

For completing 25% of the ORCA025L75 set of forecast with GloSea5 we used all the resources scheduled for 2015 (38,499,421 SBU). Therefore, to complete the remaining 75% of this experiment, which is scheduled for 2015, we estimate that an additional 115,500,000 SBU are needed.

References

Hong, S. –Y., Koo, M. –S., Jang, J., Kim., J. –E., Park, H., Joh., M. –S., Kang, J. –H., Oh, T. –J. (2013). An Evaluation of the software system dependency on a global atmospheric model. *Monthly Weather Review*, **141**, 4165-4172.

Rayner, N. A., D. E. Parker, E. B. Horton, C. K. Folland, L. V. Alexander, D. P. Rowell, E. C. Kent, and A. Kaplan (2003). Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century. *J. Geophys. Res.*, **108**, 4407.

Senoner, J.M., García, M., Mendez, S., Staffelbach, G., Vermorel, O. and T. Poinsot (2008). Growth of rounding errors and repetitivity of large eddy simulations. *AIAA Journal*, **46**, 1773-1781.

Thomas, S. J., Hacker, J . P., Desgagné and M., Stull, R. B. (2001). An ensemble analysis of forecast errors related to floating point performance. *Weather and Forecasting*, **17,** 898-906.