



# REPRODUCIBILITY OF AN EARTH SYSTEM MODEL UNDER A CHANGE IN COMPUTING ENVIRONMENT

TECHNICAL NOTE

BSC-CPG-2018-002

F. Massonnet<sup>1,2</sup>, M. Ménégoz<sup>2,3</sup>, M. C. Acosta<sup>2</sup>,  
X. Yepes-Arbós<sup>2</sup>, E. Exarchou<sup>2</sup>, F.J. Doblas-Reyes<sup>2,4</sup>

<sup>1</sup>Georges Lemaître Centre for Earth and Climate Research  
*Earth and Life Institute, Université catholique de Louvain*  
3, Place Louis Pasteur, 1348 Louvain-la-Neuve, Belgium

<sup>2</sup>Earth Sciences Department  
*Barcelona Supercomputing Center - Centro Nacional de  
Supercomputación (BSC-CNS)*  
C/Jordi Girona, 29, 08034 Barcelona, Spain

<sup>3</sup>Institut des Géosciences de l'Environnement  
*Université Grenoble-Alpes CNRS*  
CS 40700 38 058 Grenoble Cedex 9, France

<sup>4</sup>ICREA  
Pg. Lluís Companys 23, 08010 Barcelona, Spain

08 October 2018

## *Series: Earth Sciences (ES) Technical Report*

A full list of ES Publications can be found on our website under:

[https://earth.bsc.es/wiki/doku.php?id=library:external:technical\\_memoranda](https://earth.bsc.es/wiki/doku.php?id=library:external:technical_memoranda)

® Copyright 2018

Barcelona Supercomputing Center-Centro Nacional de  
Supercomputación (BSC-CNS)

C/Jordi Girona, 31 | 08034 Barcelona (Spain)

Library and scientific copyrights belong to BSC and are reserved in all countries. This publication is not to be reprinted or translated in whole or in part without the written permission of the Director. Appropriate non-commercial use will normally be granted under the condition that reference is made to BSC. The information within this publication is given in good faith and considered to be true, but BSC accepts no liability for error, omission and for loss or damage arising from its use.

## Summary

Most Earth System Models (ESMs) are running under different high performance computing (HPC) environments. This has several advantages, from allowing different groups to work with the same tool in parallel, to leveraging the burden of ensemble climate simulations but also offering alternative solutions in case of shutdown (expected or not) of any of the environments. However, for obvious scientific reasons, it is critical to ensure that ESMs provide reproducible results under changes in computing environment. While reproducibility in a strong sense (bit-for-bit) is in general unfeasible, it can be hoped that results obtained under one computing environment are statistically indistinguishable from those obtained under another environment (reproducibility in a weak sense). Here, we develop a protocol to assess the reproducibility (in a weak sense) of the EC-Earth ESM. Using two versions of that ESM, we present one case of non-reproducibility and one case of reproducibility. The non-reproducible case occurs with the older version of the model and likely finds its origin in the treatment of river runoffs along Antarctic coasts. However, it was not possible to nail the problem further down. By contrast, the more recent version of the model provides reproducible results. The protocol introduced in this note can help users of EC-Earth and, by extension other ESMs, assessing whether their model can be ported from one HPC environment to another. Our results and experience in this work suggest that the default assumption should be that an ESM is not reproducible under changes in the HPC environment, until proven otherwise.



## Contents

1.	Introduction .....	2
2.	Issues of reproducibility of Earth System Models .....	3
2.1.	Origins of non-reproducibility .....	3
2.2.	State of the art .....	4
3.	Methods.....	5
3.1.	Earth System Model .....	5
3.1.1.	Code information and revisions used .....	5
3.1.2.	Atmosphere component .....	5
3.1.3.	Ocean and sea ice components.....	5
3.1.4.	Land .....	6
3.1.5.	Coupling.....	6
3.2.	Protocol for testing reproducibility.....	6
3.2.1.	Generation of simulations.....	7
3.2.2.	Calculation of standard indices.....	7
3.2.3.	Statistical testing.....	8
3.2.4.	Experimental setup.....	9
4.	Results .....	11
5.	Concluding remarks and recommendations.....	15
6.	Additional information .....	17
7.	Acknowledgements.....	18
	References.....	19

# 1. Introduction

Numerical models of the climate system are an essential pillar of modern climate research. These models are the primary source of information for understanding the functioning of the Earth's climate, for attributing observed changes to specific drivers, and for the development of mitigation and adaptation policies, among others (IPCC, 2014). Over the years, these models have become more and more complex. Today's Earth System Models (ESMs) consist of several components of the climate system (ocean, atmosphere, cryosphere, biosphere) coupled together, and often feature several millions of lines of code. Due to their high computational requirements, ESMs usually run on high performance computing (HPC) facilities, or supercomputers. As such, climate science now fully entails an important computational component. Climate scientists - developers, users, and now computer scientists, are typically facing three questions regarding computational aspects of the ESMs:

1. How can better performance be achieved for a given model configuration?
2. What is the accuracy of the solution returned?
3. Are the results reproducible under changes in hardware or software?

These three aspects (performance, accuracy and reproducibility) usually conflict and cannot all be achieved at the same time (Corden and Kreitzer, 2015). In this note, we are primarily interested in the third aspect. While the reproducible character of scientific results appears as a natural requirement for any branch of sciences (Berg, 2018), this is a challenge in the particular context of Earth System Modeling due to external constraints imposed to the researchers, such as the fact that ESMs are run on supercomputers usually not owned by the researcher himself nor by his institution. Therefore, the meaning of reproducibility for climate sciences, how to test it and whether ESMs provide reproducible results are all legitimate questions to ask. Our goal with this note is threefold:

- We aim at establishing a protocol for detecting the possible non-reproducibility of an ESM under changes in the HPC environment (compiler, distribution of processors, compilation options, flags, ...);
- We aim at reporting examples of non-reproducibility that highlight the need to run ESMs in full awareness of possible existence of bugs in its code;
- We aim at alerting the climate community about the underestimated role of hardware or software errors in the final model solution.

This note first reviews the existing literature on the topic, which is relatively meager compared to the importance it has. Then, we present EC-Earth, the ESM used in this work, as well as the protocol for checking its reproducibility across environments. Finally we report instances of reproducibility and non-reproducibility in EC-Earth before formulating recommendations for potential users of climate models.

## 2. Issues of reproducibility of Earth System Models

It has been long established that output from computer codes of weather, atmospheric and by extension climate models would inevitably face reproducibility issues. The reason is fundamental. On the one hand, the dynamics underlying the evolution of the atmosphere is highly sensitive to initial conditions as first pointed out by Lorenz (1963). That is, two integrations started from arbitrarily close initial conditions will quickly depart from one another, with doubling time of errors of 2-3 days, due to the intrinsic internal noise of the climate system. On the other hand, computer codes are based on finite-precision arithmetics, and the representation of numbers or operations can (slightly) change whenever a different compiler, optimization level or Message Passing Interface (MPI) configuration is used. Combining these two effects inevitably leads to issues of reproducibility when codes of ESMs are executed in different computing environments, unless specific precautions are taken.

We review in more detail the reasons behind the non-reproducibility of ESMs and the literature published on the topic so far.

### 2.1. Origins of non-reproducibility

The governing equations solved by computational models are represented using floating-point variables in binary base. The general representation of a variable consists in several bits to represent the value, exponent and sign of a number. This means that we have a finite number of bits used to represent each real number. This limits our capacity to represent data with enough fidelity, but also determines the magnitude of the numerical errors that will be added to the results of an algorithm due to round-off and other sources of numerical errors. Although there are different reasons for the origin of this kind of errors, most of them are related to 1) how the compiler does the translation of floating point calculations to assembler code and 2) how the calculations are done during parallel computation. Both of them will be explained briefly below.

Results produced by a simulation using a specific compiling setup (version, target hardware, flag compilations ...) may be irreproducible under a different compiling setup because trivial round-off errors introduced in the compiled code can potentially trigger significant changes in simulation results. A compiler not only translates the code from a high-level programming language to a low-level language but also tries to improve computational performance of the codes with compiler optimization schemes. The optimizations (or simply, the translation to assembler code) done by the compiler may introduce round-off errors (or even bugs) that are easily overlooked due to the uncertainties or unknowns in ESMs.

Another source of non-reproducibility is the non-deterministic nature of parallel applications. When global collective communications are used, all the resources working in parallel have to send and receive some data. These data, which will be collected by a master process, may arrive in random order (due to delays in message passing between processes). When data is processed following the order of arrival, the results can end up being non-deterministic because of round-off errors produced by the different order to collect the final result. (We

recall that the commutative and associative properties of mathematics do not hold in finite-precision arithmetics). There are several techniques to avoid round-off errors during parallel computation but this implies, in some way, to degrade the computational performance of the execution. This happens, for example, when requiring the collective communications to be sequenced in a prescribed order. Other techniques can be used to reduce the impact of maintaining a particular order to do the operations, such as a binary tree process to calculate the collective communications, avoiding a single sequential order but yet depending on the load balance of the parallel execution to achieve peak performance. All these options can be implemented into the code by the developer, others are inserted directly by the compiler or the library used for the parallel execution, where one more time again, things depend on the compilation setup chosen and can differ (the final round-off errors) from one to another.

## 2.2. State of the art

Rosinski and Williamson (1997) were the first to raise the concern of reproducibility in a global atmospheric model, and formulated several criteria to validate the porting of such models from one computing environment to another. Recognizing already that bit-for-bit reproducibility would be impossible to meet, they proposed that the long-term statistics of the model solution in the new computing environment should match those in a trusted environment. Subsequent studies tested the sensitivity of results to domain decomposition, change in compiler, or usage of different libraries. They all came to the same conclusion that changes in behavior induced by hardware or software differences were not negligible compared to other sources of error such as uncertainty in model parameters or initial conditions. These conclusions were found to hold from weather (Thomas et al., 2002) to seasonal (Hong et al., 2013) and even climate change (Knight et al., 2008) time scales.

Arguably, the most comprehensive and complete study on the topic is that from Baker et al. (2015). Recognizing that the atmosphere exhibits coherency across variables and across space, Baker and colleagues proposed a protocol to identify possible non-reproducibility in standard atmospheric fields, accounting for the strong covariance that may exist between these fields. While useful, the Baker et al. study addresses only short (1-yr) time scales and is only concerned by atmospheric variables. It is important to acknowledge that variations in hardware or software can potentially impact slower components of the climate system, that the time of emergence of the differences may exceed one year, and that long runs might be needed to disentangle internal climate variability from a true signal. As an example, Servonnat et al. (2013) investigated the reproducibility of the IPSL-CM5A-LR climate model to changes in HPC environments. They found that for dynamical variables like pressure or precipitation, at least ~70 years would be needed to ensure that one given signal lies within the bounds of the reference signal.

## 3. Methods

### 3.1. Earth System Model

EC-Earth is a state-of-the-art ESM developed by the EC-Earth consortium, counting close to 20 European institutions (Hazeleger et al., 2012 and Doescher et al., in prep). EC-Earth is a community model developed in a collaborative and decentralized way. EC-Earth consists of coupled component models for atmosphere, ocean, land and sea ice, as described hereunder.

In this study, two versions of the EC-Earth ESM are used. The first one, denoted EC-Earth 3.1 hereinafter, is the “interim” version that was developed between the fifth and sixth stages of the Coupled Model Intercomparison Project (CMIP5 and CMIP6). The second one, denoted EC-Earth 3.2, is the “near-CMIP6” version that was used during the two years preceding the official release of EC-Earth for CMIP6.

#### 3.1.1. Code information and revisions used

The EC-Earth source codes used for this study were managed through the Subversion (SVN) version control system. For EC-Earth 3.1, the revision r1722 (EC-Earth3.1 official release) of the code was used. For EC-Earth 3.2, the revision r3906 of the code was used.

#### 3.1.2. Atmosphere component

The atmosphere component of EC-Earth 3.1 is the Integrated Forecasting System (IFS), cycle 36r4, of the European Centre for Medium-Range Weather Forecasts (ECMWF). IFS is a primitive equation model with fully interactive cloud and radiation physics. The T255 (~80 km) spectral resolution features 91 vertical levels (up to 1 Pa). The time step is 2700 seconds. IFS is adapted to High Performance Computing (HPC) using the distributed memory paradigm with the standard MPI. It uses domain decomposition to distribute the workload among MPI processes in the horizontal plane, increasing the complexity and overhead of the execution to satisfy the requirements for parallel execution. The atmosphere component of EC-Earth3.2 is the same (IFS cycle 36r4). With respect to the model version used for CMIP5 (Hazeleger et al., 2012), the main updates and improvements in EC-Earth 3.1 include an improved radiation scheme (Morcrette et al., 2008), and a new cloud microphysics scheme (Forbes et al., 2011) in the atmosphere.

#### 3.1.3. Ocean and sea ice components

The ocean component of EC-Earth 3.1 is the version 3.3.1 of NEMO2 (Madec, 2008). NEMO uses the so-called ORCA1 configurations, which consists of a tripolar grid with poles over northern North America, Siberia and Antarctica at a resolution of about 1°. Higher resolution, by roughly a factor 3, is applied close to the equator in order to better resolve tropical instability waves. 46 vertical levels are used, and the vertical grid thickness ranges between 6m and 250m. The effects of the subgrid scale processes (mainly the mesoscale eddies) are represented by an isopycnal mixing/advection parameterization as proposed by Gent and

McWilliams (1990) while the vertical mixing is parameterized according to a local turbulent kinetic energy (TKE) closure scheme (Blanke and Delecluse, 1993). A bottom boundary layer scheme, similar to that of Beckmann and Döscher (1997), is used to improve the representation of dense water spreading. The ocean component is coupled to the Louvain-la-Neuve sea Ice Model, version 3 (LIM3, Vancoppenolle et al., 2009) which is a dynamic-thermodynamic model explicitly accounting for subgrid scale variations in ice thickness. However, in EC-Earth 3.1, only one ice thickness category was used due to numerical instabilities when the default configuration was used with five thickness categories.

EC-Earth 3.2 uses the version 3.6 of the NEMO model and an updated version of the LIM3 model, which this time runs with five ice thickness categories. The ocean grid is identical except that it has 75 vertical levels.

NEMO is adapted to HPC using the shared memory paradigm with the standard MPI. Similar to IFS, It uses domain decomposition to distribute the workload among MPI processes.

#### **3.1.4. Land**

Both EC-Earth versions 3.1 and 3.2 use the H-TESEL (TESEL for Tiled ECMWF Scheme for Surface Exchanges over Land) scheme for the land surface (van den Hurk et al. 2000).

#### **3.1.5. Coupling**

The atmosphere and ocean–sea ice components of EC-Earth are coupled every time step with the Ocean Atmosphere Sea Ice Soil coupler version 3 (OASIS3; Valcke, 2012). OASIS allows exchanging different fields among components (such as IFS or NEMO) during the execution of EC-Earth. The coupling process involves the transformation of the fields from the source grid to the target grid (including interpolation and conservative operations when it is needed) and the explicit communication among components using MPI communications. OASIS is able to work using MPI to exchange fields between the source and target grids. For EC-Earth 3.1 OASIS3 is used as an independent application, while with EC-Earth 3.2 OASIS3-MCT is called using library functions (thus not requiring dedicated processors).

### **3.2. Protocol for testing reproducibility**

Our protocol is designed to test whether a given version of EC-Earth (either 3.1 or 3.2) gives reproducible results under two computing environments, named “A” and “B” for the sake of illustration (proper names will be given in the next section). Before designing the protocol for reproducibility itself, it was checked and confirmed that both EC-Earth 3.1 and 3.2 are fully deterministic. This was done using appropriate keys that force the parallel code to be executed in the same conditions, at the expense of an increase in computing time: for each version, two one-year integrations were conducted consecutively under the same computing environment (same executable, same machine, same domain decomposition, same MPI ordering). The results were found to be bit-for-bit identical in both cases.

Our protocol for testing the reproducibility of EC-Earth implies the generation of ensemble simulations. Ensemble simulations are necessary to gauge the imprint of internal variability on climate indices like time averages, and disentangle that source of variability from actual changes caused by hardware or software differences. In an attempt to reach reasonable statistical power while keeping a low level of significance (see below), and constrained by limited computational resources, we run 5-member, 20-year simulations for both “A” and “B” computing environments. In the following, each of these 5-member ensemble simulations are termed “integrations”.

### 3.2.1. Generation of simulations

The five members of integrations conducted on environments A and B always start from unique atmospheric and sea ice restarts, obtained from a long equilibrium simulation conducted at the Italian National Research Council (CNR) (personal communication, Paolo Davini: [http://sansone.to.isac.cnr.it/eearth/init/year1850\\_tome/15010101/](http://sansone.to.isac.cnr.it/eearth/init/year1850_tome/15010101/)). An ocean restart was also obtained from this equilibrium simulation, and five random but deterministic perturbations were added to the sea surface temperature of this restart at all grid points (gaussian perturbation, standard deviation:  $10^{-4}$  K). The introduction of these tiny perturbations allows ensemble spread to develop in integrations A and B and to eventually sample the model’s own internal climate variability. Note that by the deterministic nature of the perturbations, pairs of members always start from the same triplet of atmospheric, oceanic and sea ice restarts: the first member of integration A and the first member of integration B start from identical initial conditions, and so for the second member, the third one, etc.

Integrations A and B are conducted under an annually repeating pre-industrial constant forcing. The integrations are 20-year long. The initial year is arbitrarily set to 1850, thus the period covers 1850-1869.

### 3.2.2. Calculation of standard indices

Due to the large amount of output produced by each simulation, it is impossible in practice to compare exhaustively all aspects of integrations A and B to one another. Therefore, the outputs from integrations A and B are first post-processed in an identical way. The code used to post-process the outputs is available publicly (see “Additional Information” at the end of this document). Based on the list of standard metrics proposed by Reichler and Kim (2008), we record for each integration standard ocean, atmosphere and sea ice output variables: 3-D air temperature, humidity and components of the wind; 2-D total precipitation, mean sea-level pressure, air surface temperature, wind stress and surface thermal radiation; 2-D sea surface temperature and salinity, and sea ice concentration. These fields are averaged monthly (240 time steps over 20 years) and the grand time mean is also saved (1 time step over 20 years).

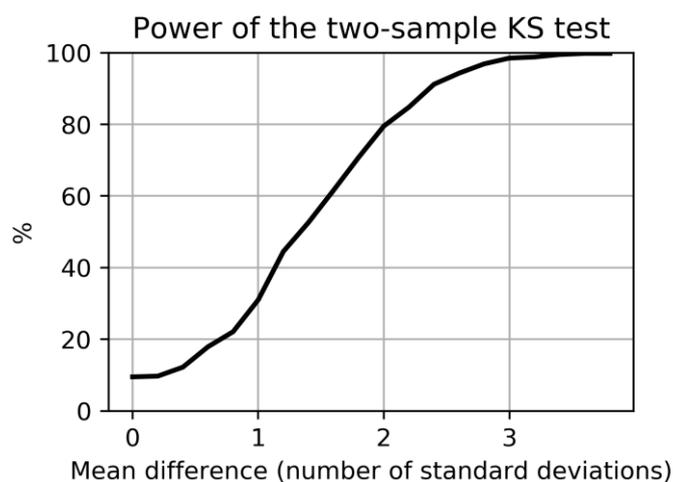
Then, several options can be used to compare the integrations A and B. A first approach would be to directly compare the model outputs from A and B. A second one that we have

considered here consists in comparing the model fields to a reference data set as described in Reichler and Kim (2008), using as reference observational and reanalysis datasets. The goal here is not to make a proper model evaluation, but rather to come up with a set of scalars characterizing the distance between a model run and a reference (whatever it may be). Therefore, the intrinsic quality of the reference data sets does not matter for our question. For each field, a grid-cell area weighted average of the model departure from the corresponding reference is evaluated and then normalized by the variance of that field in the reference data set. Thus, for each field, one number is retained that describes the mismatch between that field in the integration, and the reference. Five such numbers are available for each integration, since each integration uses five ensemble members.

### 3.2.3. Statistical testing

For each index derived in Sect. 3.2.2, we need to compare two five-member ensembles and determine whether the two ensembles are statistically indistinguishable from one another. Since no prior assumption can be made on the underlying statistical distribution of the samples, we use a two-sample Kolmogorov-Smirnov (KS) test. KS tests are non-parametric, which makes them suitable for our application.

A Monte-Carlo analysis reveals that for a prescribed level of significance of 5%, the power of the two-sample KS test exceeds 80% (a standard in research) when the means from the two samples are separated by at least two standard deviations, in case of Gaussian distributions (Fig. 1). Stated otherwise, when the means of two ensembles are separated by less than 2 standard deviations, there is a non-negligible chance (>20% at least) that the difference is not detected by the KS test.



**Figure 1. Power of the statistical test used.** Probability that a two-sample Kolmogorov-Smirnov (KS) test returns a  $p$ -value below a prescribed significance level of 5%, for two 5-member normal samples with equal standard deviations, and means separated by 0, 0.1, 0.2, ...4.0 standard deviations ( $x$  axis). The probabilities were estimated using 1,000 Monte-Carlo runs for each effect size. Gaussian distributions with equal variance were assumed for the samples.

### 3.2.4. Experimental setup

For the purpose of this paper, which is to introduce a protocol for reproducibility and illustrate cases of (non)-reproducibility in ESMs, only two computing environments were considered (Table 1). Each version of EC-Earth was used to produce one integration in each computing environment, resulting in four experiments (Table 2). The experiments were deployed and run using the Autosubmit scheduler (Manubens-Gil et al., 2016), which ensures an identical treatment of source code, namelist, compilation flags and input files management throughout. It should be noted that each experiment runs under a different domain decomposition, but sensitivity experiments conducted under the same computing environment and where only the domain decomposition was changed, indicated that this did not induce detectable changes in the results (in the sense of the KS statistic described in section 3.2.3).

	<i>Computing Environment</i>	
	<i>ECMWF-CCA</i>	<i>MareNostrum3</i>
<i>Location</i>	Reading, UK	Barcelona, Spain
<i>Motherboard</i>	Cray XC30 system	IBM dx360 M4
<i>Processor</i>	Dual 12-core E5-2697 v2 (Ivy Bridge) series processors (2.7 GHz), 24 cores per node	2x Intel SandyBridge-EP E5-2670/1600 20M 8-core at 2.6 GHz, 16 cores per node
<i>Operating system</i>	Cray Linux Environment (CLE) 5.2	Linux - SuSe Distribution 11 SP2
<i>Compiler</i>	Intel(R) 64 Compiler XE for applications running on Intel(R) 64, Version 14.0.1.106 Build 20131008	Intel(R) 64 Compiler XE for applications running on Intel(R) 64, Version 13.0.1.117 Build 20121010
<i>MPI version</i>	Cray mpich2 v6.2.0	Intel MPI v4.1.3.049
<i>LAPACK version</i>	Cray libsci v12.2.0	Intel MKL v11.0.1
<i>SZIP, HDF5, NetCDF4</i>	v2.1, v1.8.11, v4.3.0	v2.1, v1.8.14, v4.2
<i>GribAPI, GribEX</i>	v1.13.0, v000395	v1.14.0, v000370

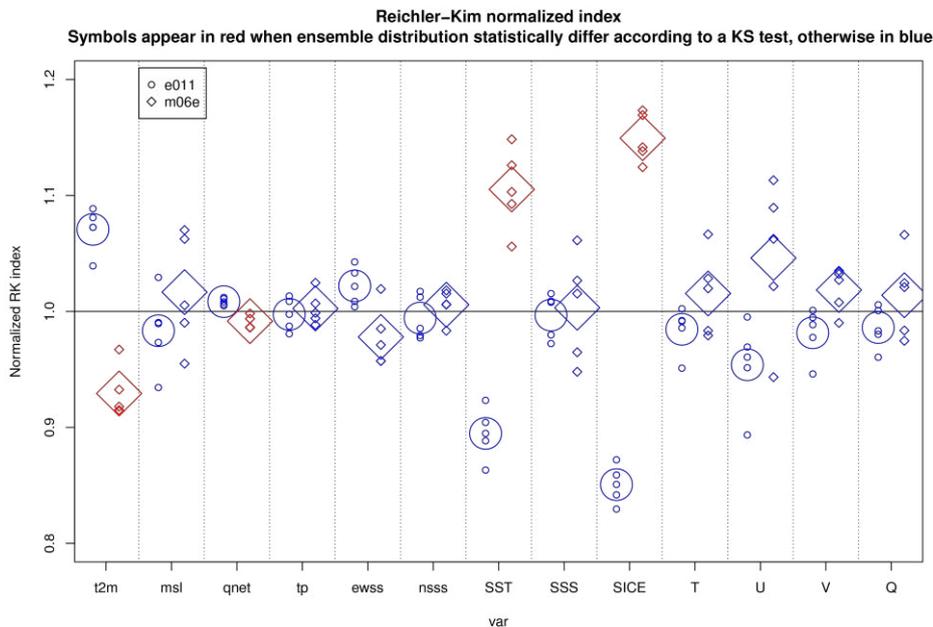
*Table 1. The two computing environments considered in this study.*

	<i>Experiments</i>			
<b>Experiment name</b>	E3.1_CCA	E3.1_MN3	E3.2_CCA	E3.2_MN3
<b>Internal ID (BSC convention)</b>	e011	m06e	a0gi	a0go
<b>Computing Environment</b>	ECMWF-CCA	MareNostrum3	ECMWF-CCA	MareNostrum3
<b>EC-Earth version</b>	3.1	3.1	3.2	3.2
<b>Processors (IFS +NEMO+OASIS)</b>	598 (480+96+22)	512 (384+96+22)	432 (288+144) (OASIS: library)	416 (288+128) (OASIS:library)
<b>F Flags</b>	-O2 -g - traceback -vec- report0 -r8	-O2 -g - traceback - vec-report0 -r8	-O2 -g - traceback -r8 - fp-model strict -xHost	-O2 -fp-model precise -xHost -g -traceback -r8
<b>C Flags</b>	-O2 -g - traceback	-O2 -g - traceback	-O2 -g - traceback -fp- model strict - xHost	-O2 -fp-model precise -xHost -g -traceback
<b>LD Flags</b>	-O2 -g - traceback	-O2 -g - traceback	-O2 -g - traceback -fp- model strict - xHost	-O2 -fp-model precise -xHost -g -traceback

*Table 2. The four experiments considered in this study*

## 4. Results

We first ran two integrations of EC-Earth 3.1 under the ECMWF-CCA and MareNostrum3 computing environments, respectively. Results revealed that for four of the parameters considered (out of thirteen, i.e. about 30%), an incompatibility was detected (Fig. 2). Since the probability of making a type I error is set to 5%, the incompatibility might not be explainable by chance only - though we should recognize that all the parameters considered display covariances that make the thirteen tests not fully independent. Differences in sea ice concentration and sea surface temperature appear very large, hinting that more investigation should be devoted to the models behavior at high latitudes.



**Figure 2.** Distribution of the normalized Reichler and Kim (2008) indices for the two simulations e011 (EC-Earth 3.1 on ECMWF-CCA, circles) and m06e (EC-Earth 3.1 on MareNostrum3, squares), for thirteen parameters: 2-m air temperature (t2m), mean sea level pressure (msl), net thermal radiation (qnet), total precipitation (tp), zonal wind stress (ewss), meridional wind stress (nsss), sea surface temperature (SST), sea surface salinity (SSS), sea ice concentration (SICE), 3-D air temperature (T), 3-D zonal wind (U), 3-D meridional wind (V), specific humidity (Q). The indices appear in red when the distribution of m06e is statistically incompatible (in a two-sample Kolmogorov-Smirnov sense) with the e011 distribution. The significance level of the KS test is set to 5%.

A spatial analysis of the difference in near surface air temperature (Fig. 3) points the Southern Ocean as the possible region of origin for the discrepancies. From the map, it appears that differences arising from this region could be responsible for the difference seen in all other parameters of Fig. 2. We further nail down the origin of the differences to winter Antarctic sea ice (Fig. 4): September ice extent departs significantly between the two integrations, and the difference in the mean values exceeds by more than a factor of two the

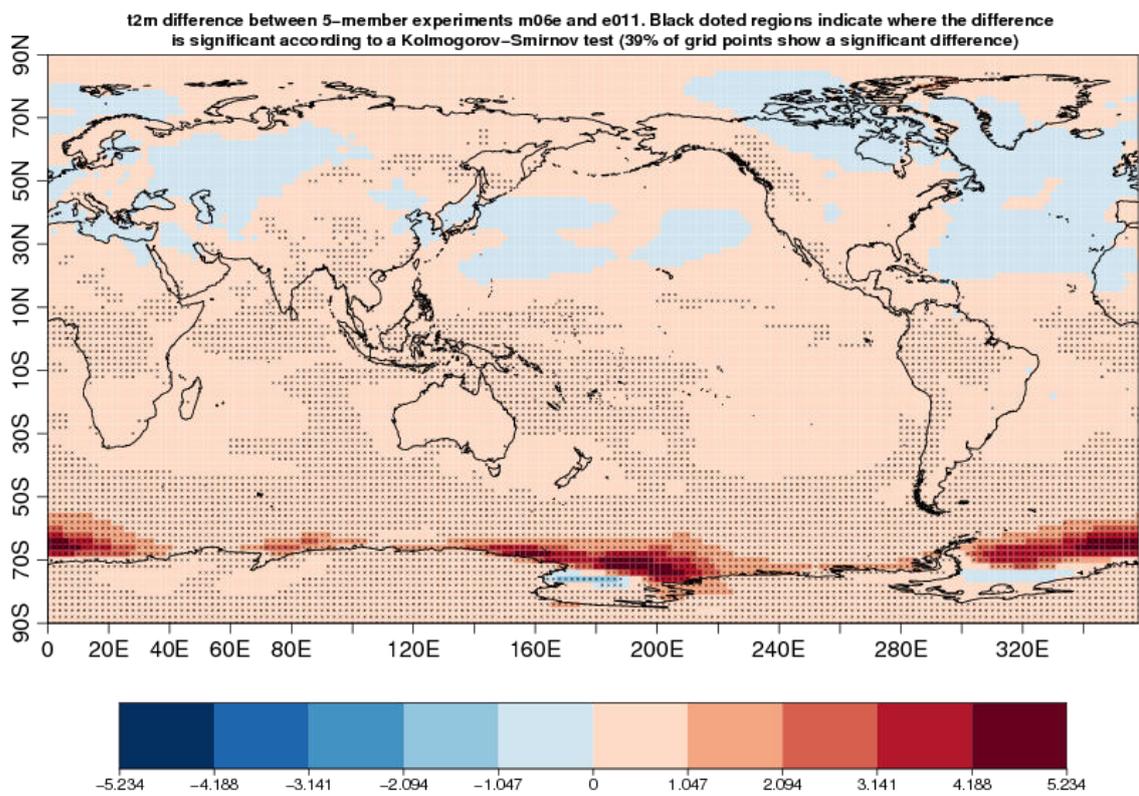


Figure 3. Difference in 20-year mean, ensemble mean near surface air temperature between the experiments m06e and e011 (red means m06e is warmer). Dots indicate pixels where the two 5-member samples are statistically incompatible in a Kolmogorov-Smirnov sense.

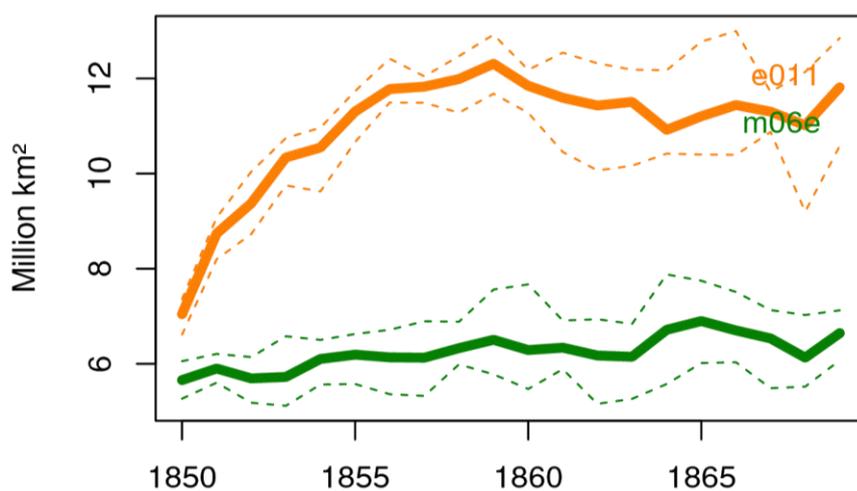
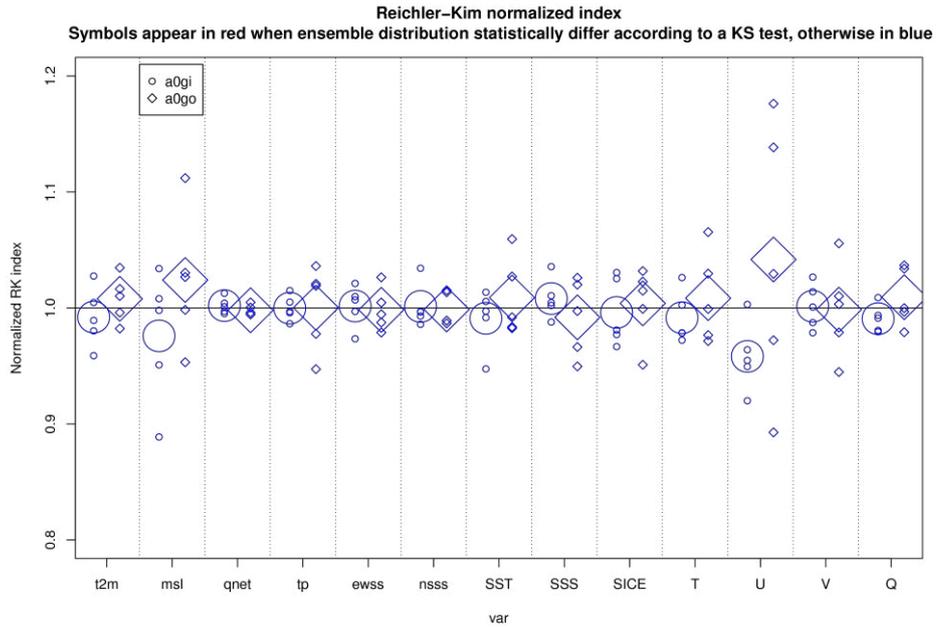


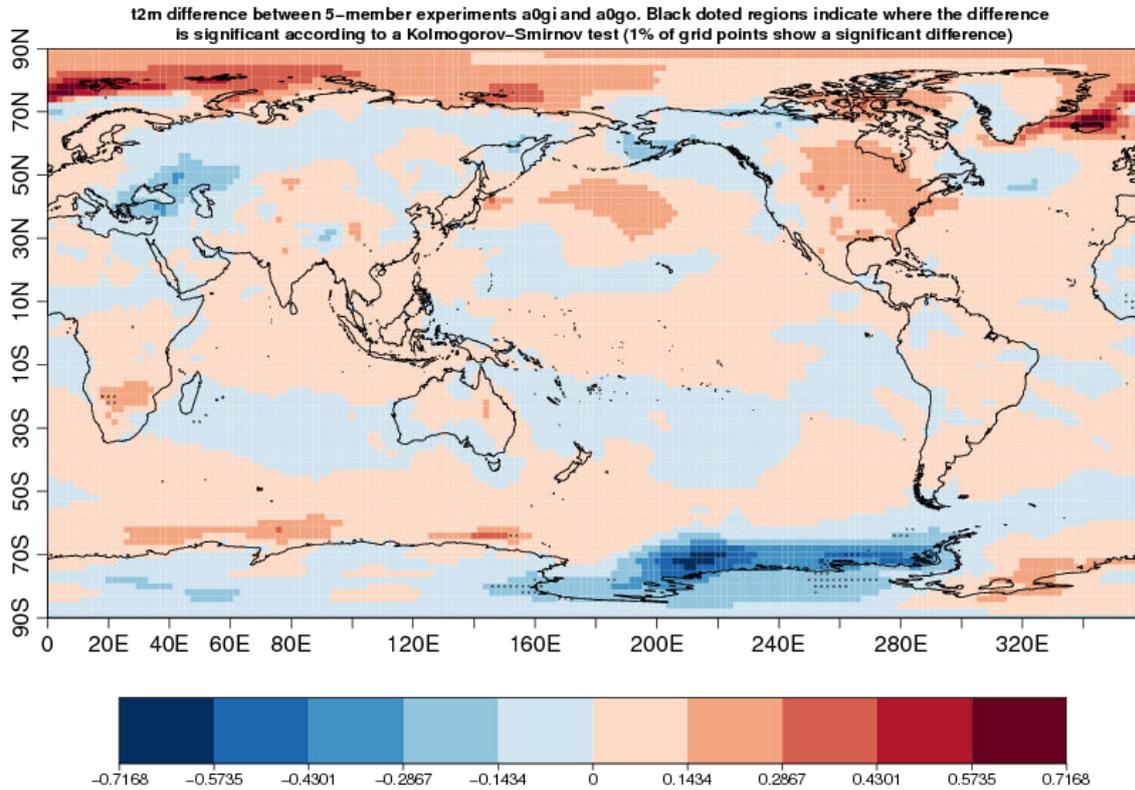
Figure 4. Time series of ensemble mean Antarctic September sea ice extent in the two experiments e011 (orange) and m06e (green). The dashed lines indicate the ensemble minimum and maximum (i.e., the range). Sea ice extent is calculated as the sum of areas of ocean pixels containing more than 15% of sea ice.

inter-member range of each model. Thus, we can be suspicious about the reproducible character of one experiment with respect to another. We then attempted to seek possible physical reasons behind this non-reproducibility. After investigation (not detailed here), one hypothesis is the following: a Fortran array involved in the river runoff routines of the NEMO model is not declared in the header of the routine (as it should). When this is the case, the compiler fills the arrays with some default values. However, which default values are set (0.0, 9999.0, NaN...) depends on the compiler itself. We noted that river runoff values off the coast of Antarctica were significantly different from one experiment to another, resulting in sea surface salinity differences by more than 1 standard deviation of internal variability. If these sea surface salinity differences at the coast spread further to the open ocean, they can eventually cause strong changes in the ocean column stratification (Kjellson et al., 2015). If vertical ocean mixing is sufficiently high in one simulation due to large positive sea surface salinity anomalies, it can even prevent sea ice formation in fall and winter. From Fig. 4, the problematic simulation seems to be the one carried out on MareNostrum3, although the ECMWF-CCA simulation is also on the low side (the observed wintertime Antarctic sea ice extent is in the range 15-20 million km<sup>2</sup>).

The analysis was then repeated with the newer version of the model, EC-Earth3.2 (experiments a0gi and a0go on ECMWF-CCA and MareNostrum3, respectively). In that case, we found no instance of incompatibility between any of the 13 parameters considered (Fig. 5). A spatial analysis (Fig. 6) suggests that only 1% of the grid points display an incompatibility for 2-m air temperature. (We recall that under the null hypothesis of no difference, significant differences are expected to occur 5% of the time). There is no sufficient evidence to reject the hypothesis that the two simulations are producing the same climate.



*Figure 5. Same as Fig. 2 but for the pair of experiments carried out with EC-Earth 3.2, namely a0gi and a0go.*



*Figure 6. Same as Fig. 3 but for the pair of experiments carried out with EC-Earth 3.2, namely a0gi and a0go.*

## 5. Concluding remarks and recommendations

Two different versions of the EC-Earth ESM were run under two different computing environments. In one case (model version 3.1), the change of environment implied a significant difference in simulated climates finding its origin in the Southern Ocean, while in the other (the model version 3.2), it did not. What can explain these different outcomes? Our protocol, like others, cannot inform on the source of non-reproducibility, but can inform whether there may exist one (Baker et al., 2015), so in-depth analyses that go beyond the scope of this note would be necessary to trace the origin of non-reproducibility with version 3.1. However, we suspect that the presence of a bug, present in EC-Earth 3.1 but no longer in EC-Earth 3.2, could explain this result. This flag allows stopping the execution when floating point exceptions, such as division by zero, are encountered. In fact, we were never able to run the EC-Earth 3.1 model with the “-fpe0” flag activated during compilation (but could well run the model if this flag was disabled). Our guess is that by disabling this flag, the model still encounters bugs (probably linked to the array initialization mentioned in Sect. 4), but these bugs give different outcomes depending on the computing environment. This worrying error that we obtained by porting a climate model from one HPC environment to another one highlights the necessity to choose adequate options of compilation when developing a model, without giving way to the temptation of excluding safe compilation options to bypass a compilation error in a new HPC. Such a result highlights also the fact that porting a code from one HPC to another might be an opportunity to detect errors in model codes.

We finally formulate a set of practical recommendations, gained during the realization of this work:

- The default assumption should be that ESMs are not reproducible under changes in computing environments. Climate scientists often assume that a model code would give identical climates regardless of where this code is executed. Our experience indicates that the picture is more complicated, and that codes (especially when they are bugged, as they often inevitably are) interfere with computing environments in sometimes unpredictable ways. Thus, it is safer to always assume that a model code will give different results from one computing environment to another, and to try proving the opposite - i.e., that the model executed in the two computing environments gives results that cannot be deemed incompatible. Our protocol fulfills this goal.
- Bugs in models are likely to be interpreted differently depending on the computing environment, and therefore cause significant changes in the simulated climates. In order to herd oneself from this inconvenient situation, and since bugs are by definition hidden to model developers, we formulate the recommendation to (1) systematically compile the model code with the -fpe0 flag, so that the model would not be able to run in case of severe bugs, and (2) run systematically the reproducibility protocol each time the ESM has to be ported to a new machine.
- Besides the frequently quoted sources of prediction uncertainty like climate model

error, initial condition errors and climate forcing uncertainty (Hawkins and Sutton, 2009), hardware and software potentially affect the ESM climate (mean state and perhaps response, though this latter point was not investigated here) in a way that deserves more attention. Users of climate models have not always the background to appreciate the importance of these impacts. Changes that may appear unimportant, like the reordering of the call to physical routines, could profoundly affect the model results (Donahue and Caldwell, 2018). For climate model users, a better understanding of the conditions that guarantee the reproducibility of ESMs is a necessary step to bring more trust in these central tools in their work.

## 6. Additional information

The code for testing the reproducibility of EC-Earth is available at <https://github.com/plesager/ece3-postproc.git>

## 7. Acknowledgements

We are grateful to Paolo Davini, Uwe Fladrich, Philippe Le Sager, Klaus Wyser, Ralf Döscher, Etienne Tourigny, Virginie Guemas, Kim Serradell, Oriol Mula -Valls, Miguel Castrillo, Oriol Tinto-Prims, Omar Bellprat and Javier Garcia-Serrano for helpful discussions. We acknowledge the ECMWF and MareNostrum3 for providing computational resources to carry out the experiments. This work was partly supported by the following projects: IS-ENES2, ESiWACE, PRIMAVERA, APPLICATE.

## References

- Baker, A. H., Hammerling, D. M., Levy, M. N., Xu, H., Dennis, J. M., Eaton, B. E., ... Williamson, D. (2015). A new ensemble-based consistency test for the Community Earth System Model (pyCECT v1.0). *Geoscientific Model Development*, 8(9), 2829-2840. <https://doi.org/10.5194/gmd-8-2829-2015>
- Beckmann, A. and R. Döscher (1997). A method for improved representation of dense water spreading over topography in geopotential-coordinate models. *J. Phys. Oceanogr.*,27, 581-591
- Berg, J. (2018). Progress on reproducibility, *Science*, Editorial, [doi:10.1126/science.aar8654](https://doi.org/10.1126/science.aar8654)
- Blanke, B. and P. Delecluse (1993). Variability of the tropical atlantic ocean simulated by a general circulation model with two different mixed-layer physics. *J. Phys. Oceanogr.*,23, 1363-1388
- Corden, M. J., and Kreitzer, M. (2015). Consistency of Floating-Point Results using the Intel Compiler or Why doesn't my application always give the same answer? <https://software.intel.com/en-us/articles/consistency-of-floating-point-results-using-the-intel-compiler> (last access: April 2018)
- Donahue, A. S., and Caldwell, P. M. (2018). Impact of Physics Parameterization Ordering in a Global Atmosphere Model. *Journal of Advances in Modeling Earth Systems*, 10(2), 481-499. <https://doi.org/10.1002/2017ms001067>
- Döscher, R. et al. (2018). EC-Earth V3: Description and evaluation of an Earth System Model (in preparation)
- Forbes, R. M., A. M. Tompkins, and A. Untch (2011). A New Prognostic Bulk Microphysics Scheme for the IFS. European Centre for Medium-Range Weather Forecasts.
- Gent, P. R. and J. C. McWilliams (1990). Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.*, 20, 150-155
- Hawkins, E., and Sutton, R. (2009). The Potential to Narrow Uncertainty in Regional Climate Predictions. *Bulletin of the American Meteorological Society*, 90(8), 1095-1108. <https://doi.org/10.1175/2009bams2607.1>
- Hazeleger, W., et al. (2012). EC-earth V2.2: Description and validation of a new seamless earth system prediction model. *Climate Dynamics*, 39 (11), 2611-2629.
- Hong, S.-Y., Koo, M.-S., Jang, J., Esther Kim, J.-E., Park, H., Joh, M.-S., ... Oh, T.-J. (2013). An Evaluation of the Software System Dependency of a Global Atmospheric Model. *Monthly Weather Review*, 141(11), 4165-4172. <https://doi.org/10.1175/mwr-d-12-00352.1>

IPCC, 2014: Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)]. IPCC, Geneva, Switzerland, 151 pp.

Kjellsson, J., Holland, P. R., Marshall, G. J., Mathiot, P., Aksenov, Y., Coward, A. C., and Ridley, J. (2015). Model sensitivity of the Weddell and Ross seas, Antarctica, to vertical mixing and freshwater forcing. *Ocean Modelling*, 94, 141-152

Knight, C. G., Knight, S. H. E., Massey, N., Aina, T., Christensen, C., Frame, D. J., ... Allen, M. R. (2007). Association of parameter, software, and hardware variation with large-scale behavior across 57,000 climate models. *Proceedings of the National Academy of Sciences*, 104(30), 12259-12264. <https://doi.org/10.1073/pnas.0608144104>

Lorenz, E. N. (1963). Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 20(2), 130-14

Madec, G. (2008). NEMO ocean engine. Note du Pôle de modélisation. Institut Pierre-958 Simon Laplace (IPSL) France 27. ISSN No 1288-1619

Manubens-Gil, D., J. Vegas-Regidor, C. Prodhomme, O. Mula-Valls and F. J. Doblas-Reyes (2016). Seamless management of ensemble climate prediction experiments on HPC platforms, 2016 International Conference on High Performance Computing & Simulation (HPCS), Innsbruck, 2016, pp. 895-900. doi: 10.1109/HPCSim.2016.7568429

Morcrette, J., H. Barker, J. Cole, M. Iacono, and R. Pincus (2008). Impact of a new radiation package, mcrad, in the ecmwf integrated forecasting system. *Monthly Weather Review*, 136 (12), 4773-4798.

Reichler, T., and Kim, J. (2008). How Well Do Coupled Models Simulate Today's Climate? *Bulletin of the American Meteorological Society*, 89(3), 303-312. <https://doi.org/10.1175/bams-89-3-303>

Rosinski, J. M., and Williamson, D. L (1997). The accumulation of rounding errors and port validations for global atmospheric models. *J. Sci. Comput.* 18, 552-564

Servonnat, J., Foujols, M. A., Hourdin, F., Caubel, A., Terray, P., and Marti, O. (2013). Comparaison du climat préindustriel du modèle IPSL-CM5A-LR sur différents calculateurs utilisés à l'IPSL. *Bulletin d'Information ORAP 77* (<http://orap.irisa.fr/wp-content/uploads/2016/03/Biorap-77.pdf>) (last accessed : April 2018)

Song, Z., Qiao, F., Lei, X., and Wang, C. (2012). Influence of parallel computational uncertainty on simulations of the Coupled General Climate Model. *Geoscientific Model Development*, 5(2), 313-319. <https://doi.org/10.5194/gmd-5-313-2012>

Thomas, S. J., Hacker, J. P., Desgagné, M., and Stull, R. B. (2002). An Ensemble Analysis of Forecast Errors Related to Floating Point Performance. *Weather and Forecasting*,

17(4), 898-906. [https://doi.org/10.1175/1520-0434\(2002\)017<0898:aeaofe>2.0.co;2](https://doi.org/10.1175/1520-0434(2002)017<0898:aeaofe>2.0.co;2)

Valcke, S. (2012). The OASIS3 coupler: a European climate modelling community software. *Geoscientific Model Development Discussions*, 5(3), 2139-2178. <https://doi.org/10.5194/gmdd-5-2139-2012>

van den Hurk, B. J. J. M., A. C. M. Beljaars, and A. K. Betts (2000). Offline validation of the ERA40 surface scheme. ECMWF Tech. Memo. 295, 42 pp.

Vancoppenolle, M., Fichefet, T., Goosse, H., Bouillon, S., Madec, G., and Maqueda, M. A. M. (2009). Simulating the mass balance and salinity of Arctic and Antarctic sea ice. 1. Model description and validation. *Ocean Modelling*, 27(1-2), 33-53. <https://doi.org/10.1016/j.ocemod.2008.10.005>

## Other references consulted throughout this work

Janakiraman, S., Ratnam, J. V., and Kaginalkar, A. (2000). Study of machine round-off response on weather forecasting simulations using high performance computing systems. In *Proceedings Fourth International Conference/Exhibition on High Performance Computing in the Asia-Pacific Region*. IEEE. <https://doi.org/10.1109/hpc.2000.843604>

Lawrence, M. G., Landgraf, J., Jöckel, P. and Eaton, B. (1999). Artifacts in Global Atmospheric Modeling: Two Recent Examples. *Eos* 80, 11 123-128

Li, R., Liu, L., Yang, G., Zhang, C., and Wang, B. (2016). Bitwise identical compiling setup: prospective for reproducibility and reliability of Earth system modeling. *Geoscientific Model Development*, 9(2), 731-748. <https://doi.org/10.5194/gmd-9-731-2016>

Senoner, J.-M., García, M., Mendez, S., Staffelbach, G., Vermorel, O., and Poinso, T. (2008). Growth of Rounding Errors and Repetitiveness of Large Eddy Simulations. *AIAA Journal*, 46(7), 1773-1781. <https://doi.org/10.2514/1.34862>

Steyer, M. (2015). Intel MPI Library Conditional Reproducibility, in *The Parallel Universe*, Intel Software Mag Issue 21, [https://shop.sos-software.com/download-sos/intel-parallel-universe\\_mag\\_issue21.pdf](https://shop.sos-software.com/download-sos/intel-parallel-universe_mag_issue21.pdf) (last accessed: April 2018)

Stodden, V., Baily, D. H., Borweing, J., LeVeque, R. J., Rider, W., and Stein, W. (2013). Setting the Default to Reproducible. Report from “Reproducibility in Computational and Experimental Mathematics” workshop ([https://icerm.brown.edu/tw12-5-rcem/icerm\\_report.pdf](https://icerm.brown.edu/tw12-5-rcem/icerm_report.pdf)) (last accessed: April 2018)