Project name	HiResNTCP: High-resolution near-term climate predictions
Research field	PE10_3 Climatology and climate change

Project leader

Title	Prof
Last name	Doblas-Reyes
First name	Francisco
Organisation name*	Barcelona Supercomputing Center
Department*	Earth Sciences
Group*	NA
Country	Spain

Abstract (max. half a page)

The future evolution of climate in the near term, from one year to a decade or so ahead, is of significant importance to our changing society. Near-term climate prediction (NTCP) is concerned with the climate evolution in those time scales, where climate is impacted by internal low frequency variability, in addition to changes due to anthropogenic or natural forcing. NTCP, by contrast with long-term climate projections, aims to produce a skilful and reliable prediction of the future evolution of climate considering both external forcing and internal climate variability, as well as their interaction. To achieve this, NTCP systems start from estimates of the observed climate system as the model's initial condition. The ability of current climate models to make skilful NTCP is usually tested by performing re-forecasts, ensembles of predictions that start from initial states in the past. The sample of re-forecasts can be compared with the observed evolution of the climate system over the same period, a validation that is essential if users are to have confidence in the forecasts. One of the key open science questions around NTCP is motivated by recent studies establishing that the typical atmospheric and oceanic resolutions used for global climate experimentation are a serious limiting factor to correctly reproduce both climate mean state and variability. Seasonal climate prediction experiments performed with PRACE resources have shown the positive impact of using unprecedented resolution increases in global climate models upon forecast quality and climate reproducibility. However, the impact of equivalent increases in resolution has not yet been tested in NTCP because the technical and scientific challenges have been too important until now. Taking advantage of the lessons being learned from multi-model coordinated global climate experiments at high resolution, **the HiResNTCP proposal aims at investigating the potential improvement of NTCP in terms of model climate, including the reduction of the drift problem, and forecast quality associated with an unprecedented increase in global climate model resolution**. The experiment proposed consists of a continuous run that will generate initial conditions for a set of 10-member ensemble re-forecasts produced over the period 1960-2018. Such an experiment would not be possible without appropriate access to tier-0 computing resources and the associated support offered by PRACE. A total of **33 Mcore-hours are requested in this proposal**. An ambitious set of analyses will be performed. The project will be carried out by members of the Earth Sciences Department of the Barcelona Supercomputing Center with the EC-Earth global climate model. The data produced will be made publicly accessible via the Earth System Grid Federation data node hosted by the Department.

IMPORTANT NOTICE

Upload a <u>single</u> document, using this template, <u>in PDF format</u>. All of the sections below **MUST BE COMPLETED** and <u>cannot exceed 8 MB</u>. The minimum font size allowed is 12 point; the minimum margin is 2 cm. The **maximum number of pages allowed is 20 pages, including cover page and references**. Reviewers will be instructed not to consider any pages out of the limit. Applicants are requested to **attach their full CV and list of publications** at the end of the document. The pages of this attachment will not be counted in the page limit.

1. **Research project**

The future evolution of climate in the near term, from one year to a decade or so ahead, is of significant importance to our modern society (Hewitt et al., 2017; Buontempo and Hewitt, 2017). Decision makers in many sectors of the economy and those concerned with human climate resilience can benefit greatly from authoritative, skilful and reliable predictions of the near-term climate. On these time scales, climate is impacted by internal low frequency variability from the oceans, atmosphere, land, and cryosphere, in addition to changes due to anthropogenic or natural forcing (Kirtman et al., 2013). Near-term climate prediction (NTCP) is concerned with the climate evolution and the implied changes in regional conditions that determine the probability of extreme events, such as frequency and intensity of heat waves and cold snaps, tropical and extratropical storms, downpours, inland and coastal flooding as well as droughts. Recent research has revealed considerable potential for NTCP when coupled climate models are initialized by the contemporaneous climate state, particularly in the oceans. Progress in these areas will bridge an important gap between seasonal forecasting (Doblas-Reyes et al., 2013a) and century-scale climate projections performed as part of the different phases of the Climate Model Intercomparison Project, which enters now its sixth phase, also known as CMIP6 (Eyring et al., 2016), for the elaboration of the Intergovernmental Panel on Climate Change (IPCC) reports. This progress will also provide a valuable building block to the seamless climate services delivery chain.

The goal of NTCP, by contrast with long-term climate projections, is to produce a skilful and reliable prediction of the future evolution of climate, considering external forcing, like for instance the increase in greenhouse gas concentrations or volcanic eruptions that have just happened, and internal climate variability, as well as their interaction. NTCP systems thus use the present and projected anthropogenic forcing in the same way as climate projections do, but, in addition, start from the observed state of the coupled climate system prescribed as the model's initial condition. Recent studies have shown that such models achieve skilful predictions over a period of several years (Smith et al., 2007; Meehl et al., 2014; Doblas-Reyes et al., 2013b). This holds primarily for surface air temperature and to some extent precipitation, but also for the frequency of extreme events such as tropical storms (Caron et al., 2015).

The premise of NTCP is that the coupled atmosphere-ocean-land-cryosphere climate system contains predictable elements on interannual to decadal timescales. For example, if the behaviour of coupled atmosphere-ocean processes in the North Atlantic Ocean can be adequately modelled and is susceptible of being predicted, their influence will spread over surrounding land areas via the bridging effect of the atmosphere. The addition of other elements to the initial climate conditions and their evolving response, such as those of the land surface and cryosphere, has the potential to add further predictability to the system.

The ability of state-of-the-art climate models to make skilful NTCP is usually tested by performing retrospective predictions or re-forecasts. These are ensembles of predictions of up to ten years of forecast time that start from initial states in the past. The set of initial states of a specific ensemble are equi-probable and different within a range that aims to represent best estimates of the observational uncertainty. This process is repeated for enough start times over the past decades (subject to the availability of sufficient climate observations) to produce an assessment of the forecast quality over past decades. The sample of re-forecasts can be compared with the observed evolution of the climate system over the same period, allowing for forecast quality assessment and the unavoidable bias adjustment. Such re-forecast-based validation of near-term climate predictions is essential if users are to have confidence in the forecasts and make suitable use of its results.

To make progress in the field, a large international NTCP exercise has recently started as a contribution of the climate-prediction community to CMIP6 and the elaboration of the next IPCC report. The Decadal Climate Prediction Project (DCPP) is a coordinated multi-model investigation into NTCP, predictability, and variability (Boer et al., 2016). The DCPP builds on recent improvements in models, in the reanalysis of climate data, in methods of initialization and ensemble generation, and in data treatment and analysis. It is closely linked to the World Climate Research Programme's (WCRP) Grand Challenge on Near Term Climate Prediction. The DCPP consists of three components. Component A comprises the production and analysis of an extensive archive of retrospective forecasts to be used to assess and understand historical NTCP forecast quality as a basis for forecasting on annual to decadal timescales. Component B undertakes ongoing production, analysis and dissemination of experimental quasi-real-time multi-model forecasts as a basis for potential operational forecast production. Component C involves the organisation and coordination of case studies of particular climate shifts and variations, both internal and naturally forced (e.g. global warming slowdown, impact of volcanoes), including the study of the mechanisms that determine these behaviours.

Among the range of open science questions around NTCP, recent studies have established that the typical atmospheric and oceanic resolutions used for global climate experimentation are a limiting factor to correctly reproduce climate mean state and variability. Noteworthy, improvements in the representation of climate processes when increasing resolutions in global models have been obtained in the simulation of El Niño Southern Oscillation (ENSO) (Masson et al., 2012), the global water cycle (Demory et al., 2014), tropical-extratropical interactions (Baatsen et al., 2014), or heat waves and droughts (Van Haren et al. 2015). In the climate prediction context, some of the simulations of the PRACE project HiResClim led to a seminal paper on the impact of resolution increase in seasonal (up to four months into the future) climate modelling and prediction (Prodhomme et al., 2016). In a companion study, the conclusion of the beneficial impact of the climate model resolution in the exploitation of the land-surface initial condition information was documented in Siegert et al. (2017). However, this hypothesis has not yet been tested in NTCP because the technical and scientific challenges have been too important until now.

Project scope and plan

The main obstacle to running global climate models at resolutions approaching 25 km is computational. Up to now, few research centres have carried out such highly-demanding simulations. The usefulness of the experiments has been limited by 1) the inability to perform enough simulations to learn how to address a true model development at those resolutions in an iterative process, as it is done with the same models when run at standard resolutions, and 2) the lack of robustness of the results obtained, which tend to be model dependent in some cases. HighResMIP is another CMIP6 coordinated exercise that explores the challenges and solutions to build a large multi-model ensemble of high-resolution global simulations with a low resolution counterpart following a common experimental protocol, i.e. a common integration period, forcing and sets of boundary conditions. This coordinated exercise, of which the European contribution is ensured by the H2020 PRIMAVERA project, will allow for an optimisation of the computing resources usage toward a common goal of identifying the robust benefits of increased model resolution based on a multi-model ensemble of long-term climate simulations. Following the example of HiResMIP and aiming at learning from its lessons, **the HiResNTCP proposal aims at investigating the potential improvement of NTCP in terms of model climate, including the reduction of the drift problem, and forecast quality associated with an unprecedented increase in global climate model resolution**. The experiment proposed is self-contained and follows a simplification of the component A of DCPP. It is addressed using a single climate model as an initial exploratory study, but the decision has been discussed with and agreed upon relevant members of the community. It is clear to us that such an ambitious experiment would not be possible without appropriate access to tier-0 computing resources.

The impact of the resolution increase in NTCP will be measured in three particular aspects: 1) the understanding of fundamental climate mechanisms, in particular those related to climate predictability such as the simulation of the North Atlantic and North Pacific multiannual variability, 2) fundamental aspects of climate modelling, in particular the origin of problems like shock, drift and bias, and 3) the preparation of initial conditions consistent with the model climate through the assimilation of observational information. A number of diagnostics will be used to address these aspects. The diagnostics, which include forecast time-dependent bias estimates, forecast quality assessments of temperature and precipitation fields for different forecast horizons, the analysis of the main modes of multiannual variability in the tropical and northern hemisphere basins (e.g. the Atlantic Multidecadal Variability mode) and their impact on meteorological variables over land, as well as the prediction capability of a key diagnostic for the ongoing global environmental change as is the global-mean temperature, will be performed with the set of diagnostics and metrics tools developed by the proposing team. In this context, the analyses in HiResNTCP will be partly driven by the progress of the PRIMAVERA project, which focuses, using process-based assessments, on the representation of mean state, variability and teleconnections over a wide range of timescales.

The project will be carried out by the Earth Sciences Department of the Barcelona Supercomputing Center (BSC-ES henceforth) using the EC-Earth global climate model, which is described in detail in the following sections. BSC-ES members are involved in both HiResMIP and the PRIMAVERA project, for which they have already obtained computing resources from PRACE in previous calls, while the PI is a member of the DCPP and the Grand Challenge on Near-Term Climate Prediction panels.

2. Methodology

EC-Earth3 comprises three major components: the atmospheric model IFS (Integrated Forecasting System) Cy36r4, the ocean model NEMO 3.6, which also includes the LIM3 sea-ice model, and OASIS3 that couples the main components.

IFS is an operational global meteorological forecasting model developed and maintained by the European Centre of Medium-Range Weather Forecasts (ECMWF). The dynamical core of IFS is hydrostatic, two-time-level, semi-implicit, semi-Lagrangian and applies spectral transformations between grid-point space and spectral space. Vertically, the model is discretized using a finite-element scheme. A reduced Gaussian grid is used in the horizontal. The IFS cycle used in EC-Earth 3.2 CY36r4.

NEMO is a state-of-the-art modelling framework for the ocean, based on the Navier-Stokes equations, used for oceanographic research, operational oceanography, seasonal forecasting and climate research studies. It is used by a large community with more than 1,000 registered users (half in Europe, and half elsewhere) and has been employed since 2015 in more than 100 projects. The NEMO version used in EC-Earth 3.2 is known as “v3.6 stable”. The core of the NEMO model is OPA, a primitive equation model adapted to regional and global ocean circulation problems down to kilometric scale. Prognostic variables are the three-dimensional velocity field, a linear or non-linear sea surface height, the temperature and the salinity. In the horizontal direction, the model uses a curvilinear orthogonal grid and in the vertical direction, a full or partial step z-coordinate, or s-coordinate, or a mixture of the two. The distribution of variables is a three-dimensional Arakawa C-type grid. Various physical choices are available to describe the ocean physics as well as various HPC functionalities to improve the model performance.

A separate executable should be configured and built for each one of them. The resolution proposed (T511L91: ~535,000 horizontal grid points with 91 vertical levels, ORCA025L75: 1,475,000 horizontal grid points with 75 vertical levels) can efficiently perform the calculations using around 2,000 processes, with the optimal number chosen with the range of efficient compute-core usage per model executable (see section 4 for details). IFS and NEMO fully support a parallel environment, while OASIS3 supports a pseudo-parallel environment. OASIS3 requires Cray pointers. For IFS there is a possibility to activate an OpenMP switch but, in this case, the implemented MPI should be thread-safe. IFS generates the output in GRIB format (which is then thinned and written in NetCDF) and NEMO in NetCDF, while OASIS3 does not generate any output. At the end of a simulation the three components always generate restarts (checkpoints) separately (IFS in binary, and NEMO and OASIS3 in NetCDF format).

For configuring and building the model executables, GNU make 3.81 or 3.81+, FORTRAN 77/90/95 compliant compiler with pre-processing capabilities and NetCDF4 deployed with HDF5 and SZIP are needed. A tool for automatic build configuration called “ec-conf” can be used. This useful tool requires Python 2.4.3 or 2.4.3+ (although it does not work yet with Python 3.0+). For NEMO, the FCM bash and perl mechanism is essential, as it is the I/O GRIB_API 1.9.9 or 1.9.9+ and GRIBEX 370 mechanism that are needed for IFS. To test the model with the run scripts, GNU date (64-bit) is also required.

The simulations will require MPI libraries and runtime facilities (MPICH2, MPICH-MX, HP-MPI, OpenMPI), optimization and data handling tools, such as BLAS, LAPACK, HDF4, HDF5, NETCDF, PARMETIS, SCALAPACK, P-NETCDF, UDUNITS, GRIB_API, CDFTOOLS v2, CDO, NCO and general configurations tools, such as PERL, PYTHON, AUTOCONF and AUTOMAKE.

The hardware that best fits the needs of the model is made of nodes with a general purpose core. Due to lack of source code adaptations, any usage with an accelerator or any other computing device, would not take full advantage of these resources. This is the main reason why the request is formulated for the Marenostrum4 platform.

Related to the performance of the model, all the model components (IFS, NEMO and OASIS), are regularly benchmarked and analysed using a methodology based on extracting traces using Extrae from real executions. These traces are displayed using the Paraver software and processed to discover possible bottlenecks (Acosta et al., 2016). BSC-ES actively contributes to improve the performance of EC-Earth, having presence both in the EC-Earth Technical Working Group and in the NEMO System Team HPC group. In particular, BSC-ES has recently developed different optimizations for the EC-Earth coupling (Acosta, M.C. et al., 2016) and for the NEMO model (Tintó Prims et al., 2015).

As the experiment will have a complicated workflow in certain phases, the Autosubmit software (Manubens-Gil et al., 2016) will be adopted to manage the workflow and ensure a uniform and optimal use of the resources. The jobs will be managed, and packed in groups in a single big job whenever required, by Autosubmit to better manage the I/O system while maximising the use of the machine. The data storage and data transfer can be organised to only require a space of 8 TB in the working file system (more details are provided in section 7). This required scratch space is motivated by the large amount of output to be generated. The output will be transferred to the BSC-ES local storage as soon as each chunk of simulation is completed, which implies that data transfer will be part of the experiment workflow and be continuously active. To complement the requirements, around 500 GB of “home” space will be required to host the code, its modified versions, and ancillary files, while a “work” space of around 15 TB should host the initial conditions to start the re-forecasts from and the reference files required by the assimilation run to be kept as close as possible to an observational reference.

The BSC-ES department has developed a suite of tools to evaluate and validate the model outputs against observational data from different observational sources (the specific diagnostics and metrics are listed in section 1). Part of the validation suite runs as the results are generated to monitor the experiments from a scientific point of view. These tools include R packages like s2dverification or Python-based solutions like Earth Diagnostics, both developed at BSC-ES and including several levels of parallelisation. This evaluation is completely integrated in the workflow, allowing early detection of any possible computational issues and, therefore, reducing waste of resources.

3. Resource management

The experiment proposed focuses on the production of a subset of the DCPP component A re-forecasts. Component A consists of 10-member ensembles of 10-year long simulations initialised once a year over the period 1960-2018 on the 1st of November. However, this is an extremely demanding exercise for the model resolution considered in this proposal. Instead, an alternative experimental design has been discussed in the community as a minimum experiment size to obtain robust results in a comparison between experiments. It consists in producing 10-member ensembles of 5-year long simulations initialised once a year over the period 1960-2018. This reduces the experiment size four times, to a total of 1,500 years of simulation. It was considered important to keep the ensemble size as large as possible given the results obtained from both seasonal and decadal forecast systems that suggest that ensemble size is key to obtain robust improvements in forecast quality estimates (Scaife et al., 2014; Marotzke et al., 2016). At the same time, keeping a long evaluation period (1960-2018) is a critical factor for robustness, although as the sampling of start dates is decreased to once every two years the analysis should take into account the possible aliasing introduced by some phenomena having a similar frequency like ENSO (García-Serrano and Doblas-Reyes, 2012). Finally, the reduction in the forecast time from 10 to 5 years has been chosen on the basis that most of the additional skill in decadal prediction over standard methods is found in the first few forecast years. Should substantial improvements in the decadal prediction forecast quality be found in the experiments proposed in HiResNTCP, an extension of the simulations in the forecast time dimension could be considered in a future proposal. In this way, which agrees with previous discussions held by the DCPP panel members, an expensive experiment like the DCPP component A can be addressed, but still requiring a substantial amount of resources.

The generation of initial conditions consistent with the model configuration chosen requires a continuous simulation over the period 1960-2018. Experience has shown that a spin-up period prior to 1960 where the model is nudged to observational estimates prior to 1960 in a periodic way is needed, the length of the spin up requiring up to 50 years (Müller et al., 2014). The continuous simulation, or assimilation run (Marotzke et al., 2016), is performed by nudging the model components (typically the ocean and the atmosphere) to observational estimates. For EC-Earth those estimates are taken from the ERA-40/ERA-Interim atmospheric and ORA-S5 ocean re-analyses, which are either produced or interpolated at the same resolution of the EC-Earth model. These are the reference files that will be stored in the “work” space for a few weeks until the assimilation run is completed. Checkpoints of the assimilation run will be stored once a year to serve as initial conditions of the re-forecasts. The ensemble of initial conditions for the re-forecasts will be generated by introducing random perturbations in the temperature fields of both the ocean and the atmosphere.

Table 1 summarises the experiments planned and the resources requested in this proposal.

	Simulation (years)	**Computing time (Mcore-hours)**	**Output (TB)**
Assimilation run	50+59=119	2.2	16.8+3=19.8
Re-forecasts	10x5x30=1,500	27.8	212.5
Total	1,619	30.0	232.3

Table 1: Resources requested. The estimates have been made on the basis of a cost of 5.3h wallclock time (using 3,456 cores) and 142 GB of output stored per year of simulation with the T511L91-ORCA025L75 configuration. The checkpoints of the assimilation run kept to generate the initial conditions of the re-forecasts have a size of 100 GB and have been added in the corresponding line. These estimates have been obtained by running EC-Earth3.2 on MareNostrum4.

Project scope and plan

These experiments will follow the schedule indicated in Table 2, a Gantt diagram with the schedule of the proposal. The first month will be devoted to the model setup, upload of the ancillary and reference files, and the creation of the test suite that ensures that the whole workflow produces test results according to plan.

Months since start	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Assimilation run			M1									
Re-forecasts									M3			
Analysis and data dissemination						M2						M4

Table 2: Schedule of the activities to be undertaken. The top row indicates the time since the start of the project. Milestones are indicated in the corresponding month with an "M".

Using 3,456 cores, the basis taken to prepare Table 1, one year of simulation will run in about 5.3h. Therefore, subject to an adequate performance of the queues, the assimilation run should be completed in about 6 weeks. Assuming that the re-forecasts can start as soon as the first checkpoint to be used as initial condition is available, the re-forecasts could be started at the end of the spin up of 50 years, which should occur at around 3 weeks after the start of the experiment. The simulations will be performed in sections, referred to as chunks henceforth, of six months.

However, prudence recommends checking that the assimilation run produces results in the long term, notably multiannual and decadal variability, matching well a set of independent observational references. For this reason, it has been considered that at least 30 years of the assimilation run will be monitored before accepting the initial conditions generated. This brings the start of the production of the re-forecasts not before one month after the beginning of the spin up. For this reason, the re-forecast production has been scheduled to start in month 3, although it could also start slightly before that time. The assimilation run will be carefully monitored from both the computational and scientific points of view as it progresses; a set of monitoring tools that will be embedded in the workflow have been described in the previous section.

Should the assimilation run show indications of spurious variability in the main climate processes that act as sources of the decadal predictability, the problem will be diagnosed over a period of up to a month and the spin up and assimilation run restarted. This potential delay of two months will allow for a contingency plan to ensure that the full resource allocated is used with the guarantee of obtaining sensible scientific results. Extra computing resources of 10% of the total, i.e. 3 Mcore-hours, are requested to cover for this eventual contingency, making a total of **33 Mcore-hours requested in this proposal**. The buffer described should also account for failing jobs (both for the assimilation run and the re-forecasts) that will need to be repeated, an aspect that our previous experience with PRACE projects highly recommends.

The re-forecasts will start in month 3. With a throughput similar to the spin up and the assimilation run of around 3 years simulated per day, and assuming that at least five ensemble members can be run simultaneously (we remind the reader that the simulation of each member and start date is completely independent), the re-forecasts will take four months to complete. Of course, the capability of running more ensemble members could speed up the completion of the re-forecasts but there are some constraints related to the data transfer described below that might prevent this from happening. The re-forecasts will be run with chunks of one year. Given that there are scheduled and non-scheduled interruptions in either the service of the HPC platform or the BSC local storage, a buffer of an additional month could accommodate the delays associated in the experiment production. In this scenario the re-forecasts should be ready by month 7, which could become month 9 if the contingency plan mentioned in the previous paragraph has to be activated.

In the previous paragraph and sections the possibility of running several ensemble members simultaneously has been mentioned. Autosubmit offers the possibility of wrapping a number of tasks (including simulations, diagnostics, and data transfers with all their dependencies) in a single job that can increase arbitrarily the size of the submitted job and accommodate the experiment progress to the recommendations of the platform operators depending on the machine load. The wrapping can be activated or deactivated at any point of any of the experiments.

The EC-Earth workflow includes a transfer job that moves the data from the computing platform to the BSC storage as soon as a chunk of simulation is complete. In spite of the wide range of transfer rates recorded between the HPC platforms used by the BSC-ES scientists to this date and the fact that the computing time could be allocated on a machine different to Marenostrum4, it is expected that the transfer jobs will move the data to the BSC local storage at a sufficient speed for not more than 10% of the total output data to reside at the same time in the computing platform.

The analysis of the simulations will start as soon as the first simulations are available. The analysis will not only take into account the simulations performed in HiResNTCP, but also a set of historical simulations performed at the same resolution as part of the HiResMIP experiment, which is currently underway, and the equivalent DCPP exercise, historical simulations and climate projections at standard resolution (T255L91/ORCA1L75) that will be performed using tier-1 computing resources allocated in Marenostrum4 and with the same model version as in HiResNTCP.

Should the HiResNTCP computing time be offered on a platform different to Marenostrum4, the schedule could be delayed by one month, the time necessary to port both the model and the running environment, perform the corresponding scalability exercise and successfully run the test suite. Besides, to ensure the compatibility of the experiments performed in different platforms, a reproducibility procedure has been established by the BSC-ES for the whole EC-Earth consortium. This procedure requires running 10-member ensembles of 5-year long simulations on the two platforms to be compared, where the simulations are started from the same initial conditions. A statistical test on the most relevant climate variables is applied to reach a decision of whether the experiments belong to the same population of climate simulations. These tests can be performed with the model at the standard resolution, which implies that their computational cost is marginal compared to the requested allocation of HiResNTCP.

To monitor the resource management and set targets for the dissemination of the results, a number of milestones have been considered and added to Table 2:

- M1: Delivery of the assimilation run in the BSC-ES ESGF node
- M2: Documentation of the scientific characteristics of the assimilation run as a BSC-ES technical memorandum prior to submission to a peer-review journal
- M3: Delivery of the decadal re-forecasts in the BSC-ES ESGF node
- M4: Documentation of the forecast quality and the predictability in the re-forecasts in comparison with standard resolution re-forecasts as a BSC-ES technical memorandum prior to submission to a peer-review journal and final report

4. **Justification of Tier-0 needs**

The most recent version of EC-Earth will be used for the high-resolution experiments planned in this proposal. This is the same version that will be used in the CMIP6 exercise. Running the high-resolution configuration of EC-Earth proposed in HiResNTCP requires a system with high-level resources. EC-Earth has already been deployed on Marenostrum4 (one of the PRACE tier-0 systems) and run on similar platforms as part of the previous PRACE projects HiResClim, HiResClim2 and LSIHP.

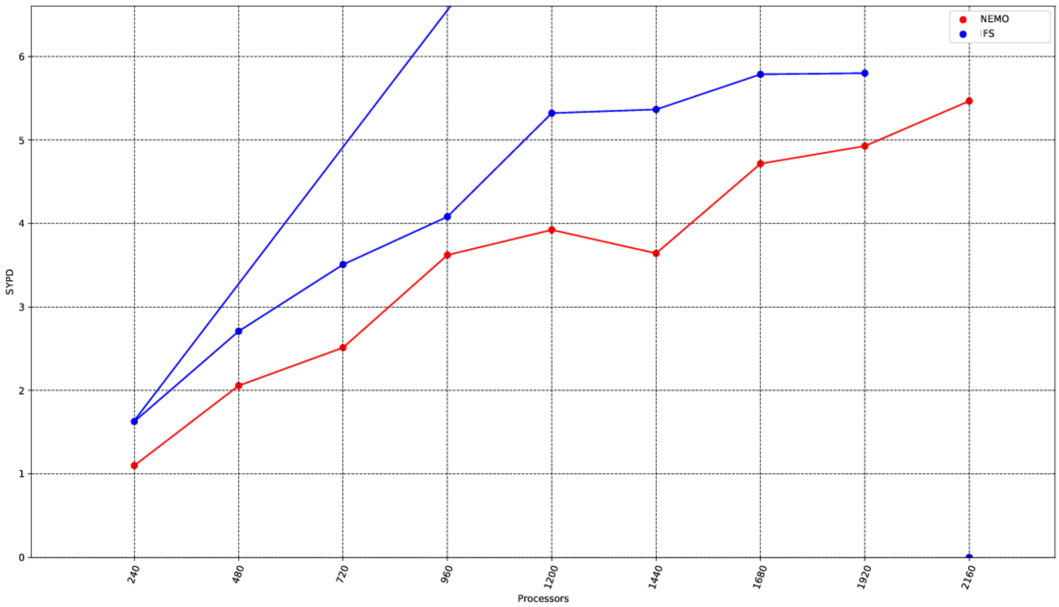

Figure 1. NEMO (red) and IFS (blue) scalability in the EC-Earth3.2 climate model. The throughput is expressed in simulated years per day (SYPD) of wallclock time, an efficiency measure typically used in climate studies as a function of the number of processes used for each main model component. The tests have been performed on Marenostrum4 with samples of five one-month runs for each processor combination, the average of which is shown in the figure. The horizontal axis corresponds to the number of cores used.

The proposed configuration uses 3,456 cores, devoting 1,680 cores to IFS, 1,680 to NEMO and 95 cores for the I/O server XIOS, plus one for the element known as runoff mapper, which uses a separate binary, all in a single job. This number might change slightly once the model is installed for the experiment production because the model performance changes with the level of I/O selected. The final list of variables will follow the requirements set by DCPP for decadal prediction experiments, which is still not closed. In this scenario, a year of simulation would have an approximate cost of 18,500 core-hours. This is the factor used in the preparation of Table 1.

The factor has been obtained from scalability analyses of each component, as shown in Figure 1. Acosta et al. (2016) showed that, while for coupled application the load balance between components has to be taken into account in the scalability process, the process needs to start with the scalability analysis of each individual component. However, showing that the optimization of one component (e.g. the reduction of the execution time of IFS) does not reduce the execution time of the coupled application if there are other slower components also requires performing a load balance analysis. The final choice depends on the needs of the specific problem, where either time or energy to solution can be minimised. For the problem we describe in this proposal the load balance between the two components is achieved when NEMO is the slowest component (Acosta et al., 2016).

#Cores	NEMO SYPD	IFS SYPD	NEMO wallclock	IFS wallclock
240	1.1	1.6	21.8	14.7
480	2.1	2.7	11.7	8.9
720	2.5	3.5	9.6	6.8
960	3.6	4.1	6.6	5.9
1,200	3.9	5.3	6.1	4.5
1,440	3.6	5.3	6.6	4.5
1,680	4.7	5.8	5.1	4.1
1,920	4.9	5.8	4.9	4.1

Table 3: NEMO and IFS performance in the EC-Earth3.2 global climate model. The model speed is expressed in simulated years per day (SYPD) and wallclock time (in hours) proportional to a one-year chunk. The tests have been performed on Marenostrum4 with samples of five one-month runs for each processor combination, the average of which is shown in the figure. The numbers provided do not include the time for initialisation and finalisation of the job.

From both model components, the slowest one is NEMO. During the scaling exercises of EC-Earth3.2 the oceanic component was targeted to eliminate a number of bottlenecks that limited the scalability of the coupled model. For instance, in the coupled executions the ELPiN (Exclude Land Processes in NEMO) method has been implemented by the BSC-ES members to find an optimal domain decomposition for NEMO with computation of only ocean subdomains. This improves substantially both the throughput and efficiency.

Figure 1 and Table 3 show that IFS scales up to 1,680 cores, the limitation being mainly due to the collective communications to produce the output every few model time steps. On the other hand, although NEMO scales beyond 1,680 cores, a number of processes for NEMO achieving a longer wallclock time than for IFS is required, which is achieved with 1,680 processes too. In this case, the coupled model will be unable to run with a shorter wallclock time than NEMO's. Since the coupling represents around 10% of the execution, the final value for the wallclock time is 5.3 hours for a year of simulation. This value could be decreased using more processes for NEMO but at the expense of a waste of resources for IFS.

The investigation to obtain an EC-Earth3.2 configuration that provides the best parallel model performance is ongoing, in close collaboration with the BSC Computer Sciences department, in two directions to reach an optimum scalability. These actions include not only an adjustment of the model configuration and a balance of the number of cores devoted to each one of its components, but also modifications of the code itself and work on the parallel programming models adopted in the different components.

5. Justification of the amount of resources requested

Table 4 lists the experiments described in section 3, which comprise one spin up simulation and an assimilation run performed in chunks of six months plus 300 re-forecasts (10 members times 30 start dates) to be performed in chunks of one year. They have to be run successively, which implies that the use of the allocation will grow with time, picking at around month 3. A simulation of one year is expected to require a wallclock time of 5.3 hours using 3,456 cores, following the benchmark performed on Marenostrum4 described in the previous section.

Run type	# Runs	# Steps/Run	Walltime/Step	#CPUcores	Total core hours/Type Run
Spin up	1	100	2.65h	3,456	915,840
Assimilation run	1	118	2.65h	3,456	1,080,691
Re-forecasts	300	5	5.3h	3,456	27,475,200
Total					29,471,731

Table 4: Detailed cost of the experiments proposed.

Table 4 does not include the cost of the output post-processing and the data management on the HPC platform because its cost is marginal compared to performing the simulations (the workflow contains three jobs performing these tasks at the end of each model chunk; these jobs do not run for longer than the time taken by the model to run the next chunk takes). Each one of these jobs uses a handful of cores.

The final estimate is for a total request of 33 million core-hours, which includes the numbers described in Table 4 plus a buffer of 10% to account for repeated jobs due to hardware (a problem experienced in Marenostrum4 in the last few months) and software (numerical instabilities) failures plus the contingency plan associated with the potential repetition of the spin up and part of the assimilation run.

6. Data management plan

Each year of simulation is estimated to produce about 142 GB of output in about 120 files (one per each ocean, sea ice or atmospheric variable) to be stored long term. The size of each file varies because some files contain 2D and others 3D variables. This output will be post-processed on the HPC platform to compute a small number of 1D climate indices for the experiment scientific monitoring and to format the files according to international standards. To make sure that any chunk can be repeated should an output file discovered to be corrupted, the checkpoint files (containing a snapshot of the atmosphere, ocean, size and coupler) of a size approaching 100 GB will be also kept in the long-term repository. In particular, the checkpoints of the assimilation run will act as initial conditions for the re-forecasts. Chunks can be repeated at any time by the workflow manager for as long as the PRACE accounts remain open to solve any contingency that corrupted the data.

Thanks to the well-defined workflow structure and the use of an adapted workflow manager like Autosubmit, the data produced by each simulation will be downloaded to the local BSC storage as soon as the post-processing of a chunk is completed, which should happen around 11 hours after the chunk started running, although this time depends on the queue capacity for the different types of jobs. In this context, a standard two-week delay between the end of the project and the closing of the PRACE accounts is more than enough for the team members to clean the HPC repositories.

Assuming that at the production peak there will be five ensemble members of the re-forecasts running simultaneously, each year of five-member ensemble simulation will produce 710 GB of output and 500 GB of the corresponding checkpoints. Taking into account that the next group of ensemble members will start running as the post-processing of the previous one is underway, the typical amount of data that will reside simultaneously in the working file system should be at least three times that produced by the five members mentioned: one for the set of simulations running, another one for the set being post-processed and, finally, a third one for the set being transferred to the BSC local archive. This approach requires around 4 TB of working space. Transfer rates between Marenostrum4 and the BSC local archive of up to 5 TB per day have been reached, which should fit the plan mentioned. However, as the transfer rate is not yet known because it depends on the platform where the request is allocated, the production rate will have to be adjusted to this factor. To cater for maintenance activities in the BSC local storage or an increase in the line traffic that might temporarily lower the transfer rate, the space requested for handling the output data is increased to 8 TB, which offers a buffer of one day of theoretical maximum production that cannot be transferred immediately.

The data stored in the BSC local archive will be managed and curated by the Data and Diagnostics Team of BSC-ES. They have developed a framework to store all the simulations and the observational data required by the BSC-ES researchers that offers access to all the data with a strict documentation, organisation and, of course, formatting. The data produced in HiResNTCP will be made accessible via the Earth System Grid Federation (ESGF) node hosted by the BSC-ES as part of the CMIP6 experiments. The terms of access will be fully public and accept commercial use of the data to better link to the future activities in decadal prediction of the Copernicus programme, which is arguably one of the most efficient ways to link to a wide range of climate data users.

Both the data disseminated and the publications originated from them will appropriately acknowledge PRACE and will follow the criteria set by the EC for open data access.

7. Past Experience

The BSC-ES has been involved in the IS-ENES2/PRACE-1IP working group focusing on the EC-Earth3 adaptation to Tier-0 machines. It has tested a range of EC-Earth3 configurations, in the atmospheric resolutions T255/511/799/1579 on several HPC systems: SGI Altix 3500, NEC-SX6, Linux cluster with Intel Xeon, Dell PowerEdge 2900, IBM pSeries 575 Power6, IBM Power PC and BlueGen. Finally, members of the BSC Earth Sciences Department coordinated the recent HiResClim and HiResClim2 projects supported by PRACE.

The Department has the privilege to have an expert team on parallel model's performance, with the ability to analyse parallel programming model codes using cutting-edge performance tools. This advantage in resources and expertise allows the department to always have the latest performance results and therefore be able to determine the optimum configurations for EC-Earth, IFS and NEMO. The BSC-ES is a collaborator of the NEMO development team, and member of the NEMO HPC working group, providing the consortium with performance reports and code optimisations. As a result of this collaboration two optimisation branches were created in the NEMO code repository, both branches being merged into the NEMO3.6 stable version and trunk. These developments are taken as a basis for further improvements in the communications within the model. Following the success of this cooperation, a similar collaboration has been established with the OpenIFS and IFS developers at ECMWF. To set up this partnership, a two month visit from team members to ECMWF has been done in 2017. During this interaction, the BSC performance tools have been applied for the analysis of the IFS model, leading to identify and apply two optimisations in the code. Moreover, the collaboration with the Research Department at ECMWF continues with the profiling analysis and possible optimizations for IFS and OpenIFS. Besides, led by the BSC-ES, ECMWF and the EC-Earth community are integrating an IO server for OpenIFS, which will increase the computational performance of the atmospheric model considerably. In parallel, the BSC-ES is also member of the EC-Earth Technical Working Group, where several profiling analysis have been done to increase the performane of the coupled version of the model.

Figures 2 illustrates an example of the performance tool's output. This is a view provided by the Paraver tool from one single EC-Earth model execution, focused on its two main components, NEMO and IFS. The number of resources used is less than the configuration recommended due to the memory needed to generate the traces. The image displays the communications pattern along time, representing time on the horizontal axis, while the vertical axis corresponds to the different processes executing the model, the first part for NEMO and the second one for IFS. The different colours correspond to different MPI communication functions, except the light blue, which corresponds to no communication. This Paraver view is very useful to determine the communications within the model and explain where the possible bottlenecks come from, especially when dealing with coupled models which have to communicate between components. For example, the green color for IFS represents the waiting time needed to synchronize the coupled model for the next time step, which means an unloaded balance in the execution. The first part of Figure 2 focuses on the NEMO component model, which in this configuration takes the first 512 processes. The image shows the importance of the communications in NEMO, which for a large percentage of time is not performing effective computations (red colour). It also shows that a large portion of this time is devoted to global communications, which appear in pink. Those collective communications belong to the horizontal diffusion routine, inside the ice model (LIM) used in NEMO, and they were one of the main targets of the optimisations, so this problem was solved to a large extent. In the second part of the figure the IFS component model reserves most of the processors used for the execution (1,024 processes). Simplifying, it can be said that the first half has less MPI communication, with more computation-only regions, while the second half contains a section performing lots of broadcasting messages (yellow and white colour block), which impacts the scalability of the code dramatically. However, these broadcasts have been removed thanks to the new optimised configuration recommended (Acosta et al., 2016) for the new version of EC-Earth, reducing this block by 90%.

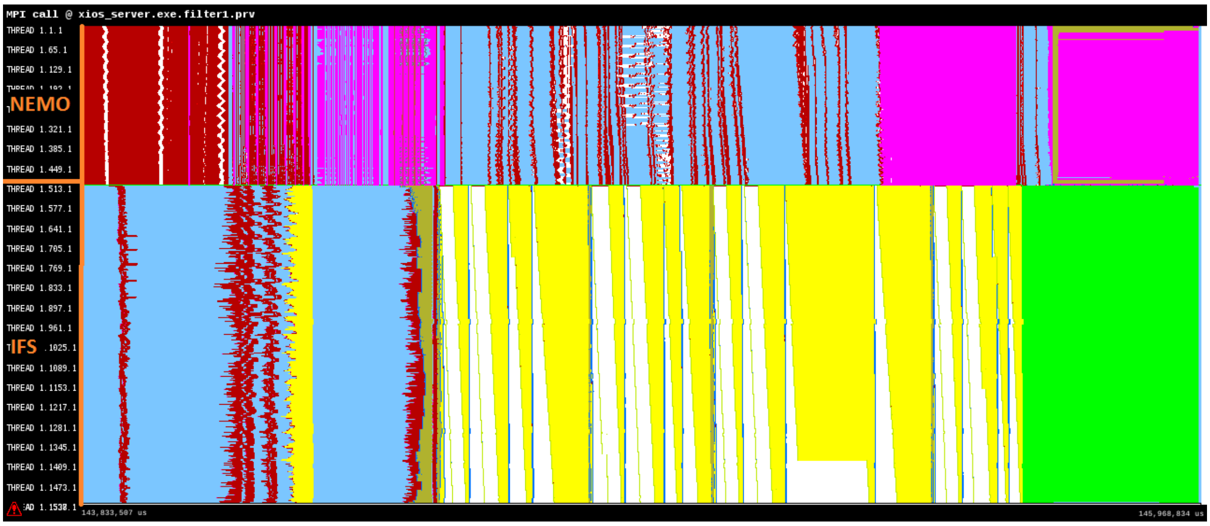

Figure 2: Paraver view of the NEMO and IFS components in an EC-Earth3 model execution. The horizontal lines give the behaviour of the different processes (1 to 512 for IFS and 513 to 1,536 for NEMO) along time. Each colour corresponds to one different MPI communication function.

The proposal is a new project that has not been submitted to any other HPC platform and that is not planned to be submitted elsewhere.

8. Previous results and communication plan

BSC-ES members are involved in a range of projects that assess the impact of model resolution increase, as explained in previous sections. These efforts have been supported by PRACE computing resources in previous calls. For instance, the PI has led another PRACE project, HiResClim (2013081637), granted in the seventh call with 65 million hours in Marenostrum3. HiResClim led to a seminal paper on the impact of resolution increase in climate modelling and prediction, Prodhomme et al. (2016), and the data are still being used by a number of scientists, particularly in the context of the PRIMAVERA project. Siegert et al. (2017) is another paper that used these simulations to address fundamental issues in the robustness of the conclusions achieved when comparing twin climate experiments. In addition, the PI has obtained a large amount of Tier-1 computing resources via the competitive calls of the Red Española de Supercomputación (RES), the US INCITE programme and the calls for special projects of the European Centre for Medium-Range Weather Forecasts (ECMWF). These resources resulted in more than 30 scientific publications where the EC-Earth model has been used.

Members of the proposing team are regularly involved in a number of PRACE activities. For instance, Mr Serradell gave the presentation “Software stack deployment for Earth System Modelling using Spack” at the most recent PRACEDays event in Barcelona. Besides, the BSC-ES is responsible for the annual organisation of the PRACE Advanced Training Course “Earth Sciences Simulation Environments”, which is held at the BSC and coordinated by the PI.

The PI is a member of the DCPP and the Grand Challenge on Near-Term Climate Prediction panels, which offers a unique platform to advertise the outcome of this project and discuss potential links to other initiatives, be it a future multi-model European effort as part of the H2020 EUCP project (due to start in December 2017) where the PI leads the NTCP research or an international coordinated initiative bringing together groups from Japan, North America and Europe that also builds on top of the DCPP effort.

Beyond DCPP, the BSC-ES participates in one of the most interesting tasks of the Grand Challenge on Near-Term Climate Prediction: the provision of decadal predictions in real time for the elaboration of WMO’s Annual-to-Decadal Climate Outlook. Should the results of the experiment proposed in HiResNTCP suggest an increased forecast quality of the high-resolution EC-Earth global climate forecast system, the system will contribute regularly to this exchange in real time, for which tier-1 resources regularly allocated to the BSC-ES can be used.

Data transferred to the BSC local storage will be disseminated following the CMIP6 conventions for variables, files and directory names through the BSC-ES ESGF data node. These data will be exploited as part of the analyses planned in the DCPP activities and in the H2020 EUCP project by as many partners as possible taking into account that the data access will be open. To favour this initiative, the project will be described in a section of the BSC-ES wiki and the BSC website summarising the experimental setup, the way to access the data and a number of key figures illustrating preliminary results.

At least one publication in an international peer-reviewed journal will be prepared summarising the results. It is important that this publication is submitted soon after the results are obtained for the results to be on time to be considered in the preparation of the Sixth IPCC Assessment Report. This article will summarise the thorough assessment of the added-value of increasing the EC-Earth3 resolution for the representation of the model mean state, variability and teleconnections, as well as for the increase in forecast quality. To speed up the impact of the experiment, BSC technical memoranda will be released with preliminary results, a process that offers a faster turnaround for dissemination until peer-reviewed material can be made available.

The results will be widely presented in scientific conferences, be submitted for presentation in the PRACEDays event of the years in which the project is active and will feature at the meetings of international programmatic initiatives of the World Climate Research Programme, the World Meteorological Organisation (for instance, in the Second Operational Climate Prediction conference to be organised by the BSC-ES in June 2018) and the EC (for instance, in the increasingly pervasive climate services events) in which the PI or the team members lead or participate.

9. **Additional needs** (optional)

The use of NCO, CDO, TOTALVIEW and NCVIEW is required to allow debugging in case a problem arises on the HPC. CDFTOOLS and Python (with a set of identified modules) will be essential for post-processing and data reduction before transfer to the local storage space (in this sense bandwidth will be required to move data from the HPC facility to our storage facilities).

10. **Track record**

This proposal is relevant to two climate modelling H2020 projects in which the PI is heavily involved as work package leader: PRIMAVERA and EUCP. PRIMAVERA is a large project that aims at developing a new generation of advanced and well-evaluated high-resolution global climate models, with which simulations of global climate are performed with unprecedented resolution in a multi-model environment. EUCP is a project that aims at bringing climate modelling closer to climate services by developing a completely new global and regional European climate multi-model system in which the impact of unprecedented resolution plays a central role. EUCP will start in next few weeks and, as PRIMAVERA, will represent a key contribution to the Sixth IPCC Assessment Report. Besides, the PI has just engaged in a collaborative bilateral project with the EC Joint Research Centre to investigate the applicability of decadal climate prediction to develop crop-yield information relevant to the Commission.

The PI has led another PRACE project, HiResClim (2013081637), granted in the seventh call with 65 million hours in Marenostrum3. HiResClim led to a seminal paper on the impact of resolution increase in climate modelling and prediction, Prodhomme et al. (2016), and the data are still being used by a number of scientists, particularly in the context of the PRIMAVERA project. Siegert et al. (2017) is another paper that used these simulations to address fundamental issues in the robustness of the conclusions achieved when comparing twin climate experiments. In addition, the PI has obtained a large amount of Tier-1 computing resources via the competitive calls of the Red Española de Supercomputación (RES), the US INCITE programme and the calls for special projects of the European Centre for Medium-Range Weather Forecasts (ECMWF).

The PI is also involved as a team member in the tier-0 PRACE project HRERS that finishes in March 2018.

The members of the team, in addition to the PI, are listed below:

- Dr Juan Camilo Acosta Navarro obtained a PhD in environmental science from Stockholm University. He has a background in atmospheric dynamics, cloud and aerosol microphysics. During his Ph.D. he studied the link between remote climatic changes induced by regional radiative forcing with a particular emphasis on Arctic climate and its teleconnections. He is currently a research scientist at the Climate Prediction group at BSC-ES and his current work focuses on understanding the drivers of sea-ice variability and how it may affect climate locally and remotely. He has been involved in several large projects: CRAICC (Norden, NordForsk), PEGASOS (EU FP7 Large Scale Integrating Project), APPLICATE (H2020) and PRIMAVERA (H2020). He has published nine scientific peer-reviewed articles, being first author on three of them. He has also supervised two MSc theses at Stockholm University.
- Dr Pablo Ortega is currently co-leading the BSC-ES Climate Prediction group. He has a broad background on ocean dynamics and decadal variability and prediction, in particular in the North Atlantic. He accumulated six years of experience as postdoctoral researcher in recognised European research centres, which have allowed him to establish a solid heterogeneous network of collaborations with research groups across Europe and to actively participate in nine research projects. He oversees the contribution of the group to two national (HIATUS and DANAE, both MINECO-

funded) and four European projects: FP7 (PREFACE), H2020 (INTAROS and APPLICATE) and COPERNICUS (QA4Seas). He is author of 20 publications in journals of the first quartile (9 as a first author, 5 as second), from which 6 were published in high-impact journals (Nature as first author, Nature Geosciences (x3), Nature Communications and BAMS), receiving a total of 186 citations with an h-index of 9 (November 2017, from Scopus). He has previously supervised two undergraduate and one master student and is currently supervising one PhD student and six postdoctoral scientists.

- Dr Yohan Ruprich-Robert is a postdoctoral scientist in the BSC-ES Climate Prediction group. He has a background in ocean and atmosphere dynamics and experience in climate prediction at decadal timescale. During his postdoctoral research at Princeton University, he set up an innovative experimental protocol that allows the investigation of the climate impacts of the Atlantic Multidecadal Variability using coupled climate models. This protocol has been adopted by DCPP as component C for CMIP6 and will be used for internationally coordinated experiments. He is currently involved in several projects (PREACE, PRIMAVERA, Hiatus, MEDSCOPE). He has a h-index (i-index) of 4 (3), with 5 articles in peer-review journals, 2 articles under review, and 57 citations. He presented his results in more than 20 international conferences and meetings, including 2 invited talks.
- Mr Kim Serradell is currently the leader of the Computational Earth Sciences group of the BSC-ES. The group is a multidisciplinary team of 19 members with different IT profiles that interacts closely with all the other groups of the BSC-ES. The group has among its tasks providing help and guidance to the scientists with the technical issues related to their work and developing a framework for the most efficient use of HPC resources. Kim Serradell is the PI of the ESIWACE H2020 project that focuses on porting and optimising codes like EC-Earth at the highest-resolution possible on a range of platforms.
- Mr Miguel Castrillo is an expert in performance profiling and optimisation applied to NEMO and EC-Earth with a strong experience in software development. He is coordinator of the Models and Workflows team in the BSC-ES. He has experience in deploying and running the NEMO ocean model on several platforms. He is a member of the Technical Issues Working Group of EC-Earth and a member of the NEMO model HPC Working Group.
- Mr Pablo Echevarria has a Masters’ degree in computer sciences. His thesis was about the implementation of a data assimilation cycle using LETKF and the NOAA WW3 wave model. He has worked since 2012 on modelling on HPC platforms in the national weather service of Argentina. He is currently member of the BSC-ES Computational Earth Sciences group. He is the person responsible of the coordination of the technical aspects of the EC-Earth experiments in the department.
- Mr Pierre-Antoine Bretonnière holds a Masters’ Degree in Mathematical and Mechanical Modelling from the Matmeca engineer school in Bordeaux (France). Graduated in 2010, he has worked in several climate research institutes (CERFACS - Toulouse - France, Catalan Institute of Climate Sciences - Barcelona - Spain and the BSC-ES). His work focuses on climate model outputs and diagnostics, data management and model coupling. He is in charge of the Data and Diagnostics team of

the BSC-ES. He is the person in charge of the data management and data conventions definitions and has participated in a number of European projects. He is also involved in the Research Data Alliance (RDA) framework as chairman of the "Weather, climate and air quality" interest group.

11. References

Acosta, M.C., X. Yepes-Arbós, S. Valcke, E. Maisonnave, K. Serradell, O. Mula-Valls and F.J. Doblas-Reyes (2016). Performance analysis of EC-Earth 3.2: Coupling. BSC-CES Technical Memorandum 2016-006, 38 pp.

Baatsen, M., R.J. Haarsma, A.J. Van Delden and H. de Vries (2014). Severe autumn storms in future western Europe with a warmer Atlantic ocean. Climate Dyn., 45, 949-964, doi:10.1007/s00382-014-2329-8.

Boer, G.J, et al. (2016). The Decadal Climate Prediction Project (DCPP) contribution to CMIP6. Geosci. Model Dev., 9, 3751-3777, doi:10.5194/gmd-9-3751-2016.

Buontempo, C. and C. Hewitt (2017). EUPORIAS and the development of climate services. Clim. Ser., doi: 10.1016/j.cliser.2017.06.011.

Caron, L.-P., L. Hermanson and F.J. Doblas-Reyes (2015). Multi-annual forecasts of Atlantic U.S. tropical cyclone wind damage potential. Geophys. Res. Letters, 42, 2417-2425, doi:10.1002/2015GL063303.

Demory, M.-E., et al. (2014). The role of horizontal resolution in simulating drivers of the global hydrological cycle. Climate Dyn., 42, 2201-2225.

Doblas-Reyes, F.J., J. García-Serrano, F. Lienert, A. Pintó Biescas and L.R.L. Rodrigues (2013a). Seasonal climate predictability and forecasting: status and prospects. WIREs Climate Chan., 4, 245-268, doi:10.1002/WCC.217.

Doblas-Reyes, F.J., et al. (2013b). Initialized near-term regional climate change prediction. Nature comm., 4, 1715, doi: 10.1038/ncomms2704.

Eyring, V., et al. (2016) Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. Geosci. Model Dev., 9, 1937-1958, doi:10.5194/gmd-9-1937-2016.

García-Serrano, J. and F.J. Doblas-Reyes (2012). On the assessment of near-surface global temperature and North Atlantic multi-decadal variability in the ENSEMBLES decadal hindcast. Climate Dyn., 39, 2025-2040, doi:10.1007/s00382-012-1413-1.

Hewitt, C., C. Buontempo, P. Newton, F.J. Doblas-Reyes, K. Jochumsen and D. Quadfasel (2017). Climate observations, climate modelling and climate services. Bull. Amer. Meteorol. Soc., 98, 1503-1506, doi: 10.1175/BAMS-D-17-0012.1.

Kirtman, B., et al. (2013). Near-term Climate Change: Projections and Predictability. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

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. In proceedings of the International Conference on High Performance Computing & Simulation (HPCS), Innsbruck, pp. 895-900, doi:10.1109/HPCSim.2016.7568429.

Marotzke, J., et al. (2016). MiKlip - A national research project on decadal climate prediction. Bull. Amer. Meteor. Soc., doi:10.1175/BAMS-D-15-00184.1.

Masson, S., P. Terray, G. Madec, J.-J. Luo, T. Yamagata and K. Takahashi (2012). Impact of intra-daily SST variability on ENSO characteristics in a coupled model. Climate Dyn., 39, 681-707.

Müller, W.A., H. Pohlmann, F. Sienz and D. Smith (2014). Decadal climate prediction for the period 1901-2010 with a coupled climate model. Geophys. Res. Lett., 41, 2100-2107.

Prodhomme, C., L. Batté, F. Massonnet, P. Davini, O. Bellprat, V. Guemas, and F.J. Doblas-Reyes (2016). Benefits of increasing the model resolution for the seasonal forecast quality in EC-Earth. J. Climate, 29, 9141-9162, doi:10.1175/JCLI-D-16-0117.1.

Scaife, A.A., et al. (2014). Skilful long range prediction of European and North American winters. Geophys. Res. Letters, 41, 2514-2519, doi: 10.1002/2014GL059637.

Siegert, S., D. Stephenson, O. Bellprat, M. Ménégoz and F.J. Doblas-Reyes (2017). Detecting improvements in forecast correlation skill: Statistical testing and power analysis. Mon. Weather Rev., 145, 437-450, doi:10.1175/MWR-D-16-0037.1.

Tintó Prims, O., M. Castrillo, K. Serradell, O. Mula-Valls and F.J. Doblas-Reyes (2015). Optimization of an ocean model using performance tools. BSC-CES Technical Memorandum 2015-002, 16 pp.

Van Haren, R., R.J. Haarsma, H. de Vries, G.J. van Oldenborgh and W. Hazeleger (2015). Resolution dependence of circulation forced future central European summer drying. Env. Res. Letters, 10, 055002, doi:10.1088/1748-9326/10/5/055002.

Appendix – CV of the Principal Investigator with list of publications (<= 5 years)

Short profile

Francisco J. Doblas-Reyes is research professor of the Institució Catalana de Recerca i Estudis Avançats (ICREA), the most prestigious Catalan institution employing senior scientists in all research fields. He develops his research activity as director of the Earth Sciences Department of the Barcelona Supercomputing Center, position that he took in November 2014 after leading the Climate Forecasting Unit of the Institut Català de Ciències del Clima (IC3) from 2010 to 2015, also in Barcelona.

Prof Doblas-Reyes has a long research experience in the understanding, simulation and prediction of climate variability and change. He has participated and coordinated a number of recognised international research projects and initiatives related to his research interests. He is the author of more than 100 peer-review journal articles and has a long record of obtaining competitive computing and financial resources to undertake this research. He is currently co-chair of the World Climate Research Programme's Working Group on Seasonal-to-Interannual Prediction and a member of a number of other international panels.

He has been a lead author of the most recent assessment report of the Intergovernmental Panel on Climate Change and, since then, been heavily involved in the design of both technical and scientific aspects of the Sixth Coupled Model Intercomparison Project, which should be the modelling basis for the next assessment report.

Education

- MSc degree in Physics at the Universidad Complutense of Madrid (Spain), June 1991.
- Post-graduate courses at the Universidad Complutense of Madrid from October 1992 to October 1996, and at Météo-France in September-October 1994.
- PhD in Physics with honours at the Universidad Complutense of Madrid (Spain). Thesis entitled "Atmospheric blocking: GCM simulation and associated precipitation patterns" (in Spanish). Date of dissertation: 22nd May 1996.
- BSc degree in Mathematics at the Universidad Complutense of Madrid, June 1997.

Professional Experience and Employment History

- Since November 2014, ICREA research professor at the Barcelona Supercomputing Center (BSC, Barcelona, Spain), working as head of the Earth Sciences Department (http://www.bsc.es), a group of more than 50 people (scientists, engineers and students).
- From December 2009 to November 2015, ICREA research professor at the Institut Català de Ciències del Clima (IC3, Barcelona, Spain), working as senior scientist and head of the Climate Forecasting Unit (http://ic3.cat/wikicfu), a group of more than 20 people (scientists, technicians and PhD students).
- From March 2000 to November 2009 at the European Centre for Medium-Range Weather Forecasts (ECMWF, Reading, UK), as research scientist.
- Visiting scientist at CINECA (Bologna, Italy) in March 2000.
- From February 1999 to February 2000 at the Centro de Astrobiología, Instituto Nacional de Técnicas Aeroespaciales (Madrid, Spain) as research assistant.
- Visiting scientist at the Institute of Atmospheric Physics (Academy of Sciences, Prague, Czech Republic) in July and August 1997.
- From January 1997 to January 1999 at CNRM (Météo-France, Toulouse, France) as research assistant.
- Visiting scientist at the CNRM (Météo-France, Toulouse, France) in September-November 1994.
- Visiting scientist at the Instituto Meteorológico Nacional (Asunción, Paraguay) in August 1994.
- From October 1992 to September 1996 at the Departamento de Física de la Tierra, Astronomía y Astrofísica II (Universidad Complutense, Madrid, Spain) as PhD student.

Membership in Committees and Panels

- CLIVAR's Decadal Climate Variability and Predictability panel, member since 2015.
- Working Group on Seasonal-to-Interannual Prediction (WGSIP) of the World Climate Research Programme (WCRP), member since 2011, co-chair since 2012 (http://www.wcrp-climate.org/index.php/wgsip-overview).
- Modelling Advisory Council (WMAC) of the WCRP, member since 2012 (http://www.wcrp-climate.org/WMAC.shtml).
- Decadal Climate Prediction Panel (DCPP) of the WCRP, member since 2012 (http://www.wcrp-climate.org/decadal/cmip5.shtml).
- European Network for Earth System Modelling (ENES) High-Performance Computing Task Force, member since 2012 (https://verc.enes.org/ISENES2/project/na1-wp2-enes-strategy).
- European Climate, Observations and Modelling for Services (ECOMS) panel of the European Commission, member 2012-2016 (http://www.eu-ecoms.eu).
- Polar Prediction Project (PPP) of the World Weather Research Programme (WWRP), member 2011-2015 (http://polarprediction.net).
- Fifth Assessment Report (AR5) of the Intergovernmental Panel on Climate Change (IPCC), lead author 2010-2013.

Selected Peer-Reviewed Publications (out of more than 130)

1. Massonnet, F., O. Bellprat, V. Guemas and F.J. Doblas-Reyes (2016). Using climate models to estimate the quality of global observational data sets. *Science*, **6311**, 452-455, doi:10.1126/science.aaf6369.
2. Doblas-Reyes, F.J., I. Andreu-Burillo, Y. Chikamoto, J. García-Serrano, V. Guemas, M. Kimoto, T. Mochizuki, L.R.L. Rodrigues and G.J. van Oldenborgh (2013). Initialized near-term regional climate change prediction. *Nature Communications*, **4**, 1715, doi:10.1038/ncomms2704.
3. Guemas, V., F.J. Doblas-Reyes, I. Andreu-Burillo and M. Asif (2013). Retrospective prediction of the global warming slowdown in the past decade. *Nature Climate Change*, **3**, 649-653, doi:10.1038/nclimate1863.
4. Doblas-Reyes, F.J., J. García-Serrano, F. Lienert, A. Pintó Biescas and L. R. L. Rodrigues (2013). Seasonal climate predictability and forecasting: status and prospects. *WIREs Climate Change*, **4**, 245-268, doi:10.1002/WCC.217.
5. Kirtman, B., S. Power, J.A. Adedoyin, G.J. Boer, R. Bojariu, I. Camilloni, F.J. Doblas-Reyes, A.M. Fiore, M. Kimoto, G.A. Meehl, M. Prather, A. Sarr, C. Schär, R. Sutton, G.J. van Oldenborgh, G. Vecchi and H.J. Wang (2013)*. Near-term climate change: Projections and predictability. Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (Eds), 953-1028, Cambridge University Press.
6. Thomson, M.C., F.J. Doblas-Reyes, S.J. Mason, R. Hagedorn, S.J. Connor, T. Phindela, A.P. Morse and T.N. Palmer (2006). Malaria early warnings based on seasonal climate forecasts from multi-model ensembles. *Nature*, **439**, 576-579.

Competitive Resources (more than 7 million euros obtained since 2012)

- EUCP (EUropean Climate Prediction system), European Commission H2020, 2017-2021, contract 776613, 1,115,125 euros.
- PRIMAVERA (PRocess-based climate sIMulation: AdVances in high resolution modelling and European climate Risk Assessment), European Commission H2020, 2015-2019, contract 641727, 1,277,425 euros.
- RESILIENCE (Strengthening the European Energy Network using Climate Services), coordinator, Spanish Ministry of Economy and Competitiveness (MINECO), 2014-2016, 270,000 euros.

- SPECS (Seasonal-to-decadal climate Prediction for the improvement of European Climate Services), coordinator, European Commission FP7, 2012-2017, contract 3038378, 1,615,305 euros.
- QA4Seas (Quality Assessment Strategies for Multi-model Seasonal Forecasts), Copernicus Climate Change Service, coordinator, 2016-2018, 731,214 euros.

Awards and Recognitions

- Recipient of the Mumm-Gerbier Prize in 2006 (http://www.wmo.int/pages/about/awards/winners_mumm_en.html).