Barcelona Supercomputing Center
Centro Nacional de Supercomputación

RES
RED ESPAÑOLA DE SUPERCOMPUTACIÓN

AECT-2019-2-0003 The Decadal Climate Prediction Project - DCPP

1. General Information

Activity Id

AECT-2019-2-0003

a) Activity Title

The Decadal Climate Prediction Project - DCPP

b) Area

Astronomy, Space and Earth Sciences

2. Research Project Description

a) Is this a Test Activity?

No

b) Is this a Long Term Activity that will extend over two application periods?

Yes

c) Brief description of the 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. 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. Near-term climate prediction (NTCP) is concerned with the climate evolution and the 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, inland flooding and 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 and century-scale climate projections performed as part of the sixth phase of the Climate Model Intercomparison Project, also known as CMIP6, for the elaboration of the Intergovernmental Panel on Climate Change (IPCC) reports.

The goal of NTCP 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, for surface air temperature and to some extent precipitation, but also for extreme events such as tropical storms. 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. 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. 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 (impact of volcanoes), including the study of the mechanisms that determine these behaviours.

d) Grants and funded projects related to this activity

Reference code	Project title		
H2020-SC5-2016-776613	European Climate Prediction system (EUCP)		
	Starting date 2017-12-01	Ending date 2021-11-30	Total financing (in EUR) 12.999.515,00
Financing source European			

e) Brief description of the Project (if this Activity takes place in the context of a Technology or Industrial Project)

f) Specific Activity proposed

The experiment proposed here focuses on the production of the DCPP component A and part of component C re-forcasts and will constitute part of BSC's contribution to the international CMIP6 exercise. The experiment will be run with the newly operational version 3.3.1 of EC-Earth using its standard resolution: T255-ORCA1, which corresponds to 75 km and 100 km in the atmosphere and ocean respectively. 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.

DCPP-A

Component A consists of 10-member ensembles of 10-year long simulations initialised once a year on the 1 st of November over the period 1960-2018, which corresponds to approximately 6,000 years of simulations. The initial conditions for each of these simulations are taken from the ERA-40/ERA-Interim atmospheric and ORA-S5 ocean re-analyses, which are either produced at or interpolated to the same resolution as the EC-Earth model. 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. 10 member has been estimated by the community as the minimum number of ensemble members sufficient to extract the signal from the noise.

59 start dates x 10 members x 10 year simulations = 5900 simulated years (1 year corresponding to 1 job in this case)

DCPP-C

Component C consists in targeted simulations and predictions intended to investigate the origins, mechanisms and predictability of long timescale variations in climate as well as their regional imprints and to investigate the influence and consequences of volcanic eruptions on decadal prediction and predictability. In the first case, we will produce a twin set of experiments wherein the sea surface temperature over the Northern Atlantic is restored towards either Atlantic Multidecadal Variability (AMV) positive and negative anomalies. The AMV is a slow (~20-30 year), naturally occuring oscillation in Atlantic ocean temperature and has been shown to modulate the North American and European climate as well as hurricane variability, amongst others. More precisely, we restore the SST using a restoring coefficient of 40 W m-2 K-1 over the Atlantic region limited by 10N and 65N and using a 8 degree buffer. Comparing the two sets of experiments will help understand the impact of the AMV on European climate and the pathway through which the responses are expressed throughout the ocean and the atmosphere. For this twin set of experiments, the DCPP protocol requires a minimum of 25 ensemble members of 10-year long simulations. The atmospheric forcings remain constant at pre-industrial level. It should be noted that the surface restoring has been implemented by members of the PI's team as part of the European H2020 PRIMAVERA project and is thus readily available in EC-Earth3.

25 members x 10 year simulations x 2 experiments = 500 simulated year

The last component of DCPP-C consists in a set of experiments aimed at understanding the prediction skill associated with volcanic eruptions. In this case, we will rerun the hindcasts for the periods covering the last three large volcanic eruptions but this time excluding the volcanic forcing. This consists in simulations initialised near the years surrounding the 1963 Agung, the 1982 el Chichón and the 1991 Pinatubo eruptions. In total, we will have fifteen 10-year and 10-member experiments, reaching a total of 1500 years of simulation. This set of experiments will allow investigating the atmospheric and oceanic responses to the last eruptions in decadal hindcasts to diagnose the forecast system skill associated with the volcanic forcing.

15 start dates x 10 members x 10 year simulations = 1500

So the total of simulated years 5900 (DCPP-A) + 500 (DCPP-C AMV) + 1500 (DCPP-C volcano) = 7900 years

All the simulations produced will be made publicly available through the Earth System Grid Federation (ESGF) node hosted at the BSC.

g) Computational algorithms and codes outline

As mentioned above, EC-Earth3 includes three major components: the atmospheric model IFS, the ocean model NEMO 3.6 and 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 and has been employed since 2015 in more than 100 projects. The NEMO version used in EC-Earth 3.3 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.

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).

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. The Earth Science department of the BSC 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, we have recently developed different optimizations for the EC-Earth coupling and for the NEMO model.

As the experiment will have a complicated workflow in certain phases, the Autosubmit software 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 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.

3. Software and Numerical Libraries

Software components that the project team requires for the activity.

a) Applications + Libraries

BLAS, HDF5, LAPACK, NETCDF, R, SCALAPACK, UDUNITS, SZIP

b) Compilers and Development Tools

GCC, TOTALVIEW, INTEL, MVAPICH2

c) Utilities + Parallel Debuggers and Performance Analysis Tools

CMAKE, PERL, PYTHON, GNUPLOT, NCVIEW, NCL, AUTOCONF

d) Other requested software

GRIB_API, GRIBEX, CDO, CDFTOOLS

e) Proprietary software

4. Research Team Description

a) Personal Data

Name of Team Leader	Louis-Philippe Caron
Gender	Male
Institution	Barcelona Supercomputing Center
e-mail	louis-philippe.caron@bsc.es
Phone	934054290
Nationality	Canada

b) The employment contract of the activity leader with the research organisation is valid at least 3 months after the end of the allocation period.

Yes

c) Curriculum Vitae of the Team Leader

Louis-Philippe Caron has obtained his Ph.D. in 2011 from the Université du Québec à Montréal. He then worked as a post-doctoral fellow in the meteorology department of Stockholm University where, in collaboration with the Swedish Meteorological and Hydrological Institute (SMHI), he produced a series of decadal experiments in the context of the IPCC CMIP5 project using the EC-Earth2.3 CGCM (these data are now freely available through the ESGF portal).

Since 2013, LPC has been a member of the Climate Prediction Group in the Earth Science Department of the Barcelona Supercomputing Center (previously affiliated with the Catalan Institute of Climate Sciences), working in the field of seasonal and decadal forecasting of tropical cyclone activity and on the different climate influences driving cyclone high and low frequency variability. He is also a member of the seasonal Tropical Cyclone Prediction Panel of the Working Group on Tropical Meteorology Research (WGTMR) and a member of the Decadal Climate Prediction Project (DCPP) of the Working Group on Group on Subseasonal to Interdecadal Prediction (WGSIP) of the World Meteorological Organization (WMO).

In recent years, he has collaborated and published with scientists from international public institutions such as the National Center for Atmospheric Research (NCAR; USA), Princeton University (USA), Columbia University (USA), Met Office (UK) and the University of Melbourne (Australia), but also from the private sector (Risk Management Solutions (UK). He has contributed to 22 peer-reviewed manuscript (9 as first author), all in Q-1 journals. He has participated in one FP7 project (SPECS) and is currently participating in 2 H2020 (PRIMAVERA, EUCP) projects. He is also the co-coordinator of the ERA4CS MEDSCOPE project and has been the PI of two MINECO projects (HIATUS, RESPONS). He was a Juan de la Cierva Incorporacion fellow from January 2016 to December 2017.

d) Names of other researchers involved in this activity

All from the Barcelona Supercomputing Center.

Pablo Echevarria (pablo.echevarria@bsc.es)
Arthur Amaral Ramos (arthur.amaral@bsc.es)
Simon Wild (simon.wild@bsc.es)
Roberto Bilbao (roberto.bilbao@bsc.es)
Vladimir Lapin (vladimir.lapin@bsc.es)
Julian Rodrigo Berlin (julian.berlin@bsc.es)
Yohan Ruprich Robert (yohan.ruprich@bsc.es)

e) Relevant publications

Smith, DM, R Eade, AA Scaife, L.-P. Caron et al. (2019) Robust skill of decadal climate predictions. npj Climate and Atmospheric Science, accepted.

Smith, DM, AA Scaife, E Hawkins, R Bilbao, GJ Boer, M Caian, L-P Caron et al. (2018) Predicted chance that global warming will temporarily exceed 1.5C. Geophys. Res. Lett.

45, 11,895-11,903. doi: 10.1029/2018GL079362.

Lavender, SL, KJE Walsh, L-P Caron, M King, S Monkiewicz, M Guishard, Q Zhang and B Hunt (2018) Estimation of the maximum annual number of North Atlantic tropical cyclones using climate models. Science Advances. 4(8), doi: 10.1126/sciadv.aat6509.

Caron, L-P, L Hermanson, A Dobbin, J Imbers, L Lledó and GA Vecchi (2018) How skilful are the multi-annual forecasts of Atlantic hurricane activity? Bull Amer Meteor Soc. 99, 403-413.

Caron, L-P, L Hermanson and FJ Doblas-Reyes (2015) Multi-annual forecasts of Atlantic U.S. tropical cyclone wind damage potential. Geophys Res Lett, 42, 2417-2425.

5. Resources

a) Estimated resources required for the Activity for the current Application Period

Requested machine	MareNostrum 4 ((Intel(R) Xeon(R) Platinum 8160, 2.10GHz with Intel(R) Omni-Path / 165888 cores)
Interprocess communication	Null

Typical Job Run	
Number of processors needed for each job	768.00
Estimated number of jobs to submit	8640.00
Average job durations (hours) per job	1.42
Total memory used by the job (GBytes)	200.00
Largest Job Run	
Number of processors needed for each job	768.00
Estimated number of jobs to submit	8640.00
Average job durations (hours) per job	1.42

Total memory used by the job (GBytes)		200.00		
Total disk space (Gigabytes)	**Minimum**	100.00	**Desirable**	150.00
Total scratch space (Gigabytes)	**Minimum**	20.00	**Desirable**	50.00
Total tape space (Gigabytes) (*)	**Minimum**	0.00	**Desirable**	0.00
Total Requested time (Thousands of hours)		9500.00		

If this activity is asking for more than 10Million CPU hours, you need to justify the amount of resources requested for the activity. (max 1000 characters)

7,900 jobs x 768 (# of processors) x 1.42 (avg duration) 8,615,421 hours. To this 8.6M hours, we add a 10% overhead to account for the cost of the output post-processing and the data management on the HPC platform (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 additional 10% also accounts for repeated jobs due to hardware and software (numerical instabilities) failures. This additional 10% of resources brings the total to 9.5 M hours.

INFORMATION: The estimated cost of the requested hours, considering only the electricity cost, is 10165 euros.

The required resources have to be executed in the selected machines, the other architectures do not fit the requirements to execute the proposal.
** this option implies that if no hours in this machine/these machines are available, the acces committee will reject the full application.

b) Estimate of the total resources that the Activity will require until it is completed (including the present and all the following Application Periods)

Number of application periods expected to complete this Activity

2

Total Requested Time (thousands of hours) expected to complete this Activity (sum of both periods)

9500.00

6. Abstract for publication

The experiment consists in running the CMIP6 DCPP A and C experiments using the EC-Earth3 climate model. DCPP-A consists of 10-member ensemble simulations initialised once a year over the period 1960-2018. Component C consists in targeted simulations intended to investigate the origins, mechanisms and predictability of long timescale variations in climate as well as their regional imprints. In DCPP-C, we produce a twin set of experiment wherein the sea surface temperature is restored towards either positive or negative Atlantic anomalies in order to understand its impact on European climate. The second component of DCPP-C consists in rerunning the some DCPP-A hindcast for periods covering three recent volcanic eruptions but excluding the volcanic forcing in order to estimate the prediction skill associated with volcanic eruptions.

7. Contact with CURES during last year

Information about the RES Users Committee (CURES).

a) User has contacted the CURES during last year

No

b) If not, indicate why you have not contacted the CURES

Because this is my first application to RES.

Usage Terms & Conditions

- The Usage Terms & Conditions have been already accepted.