PRACE Final Report Form – Project Access

1. General information

1.1. Proposal ID

2012060992

1.2. Type of proposal granted: Project Access

1.3. Period of access to the PRACE facilities

Originally granted: 1st November 2012 – 31st October 2013
Actual access to PrACE system started on 1st January 2013

1.4. Name of the PRACE facility assigned

MareNostrum III Barcelona Supercomputing Center, Spain

1.5. Name of the Principal Investigator

Dr Colin Jones, Swedish Meteorological and Hydrological Institute

2. Project information

2.1. Project name

HiResClim: High Resolution Ensemble Climate Modeling

2.2. Main research field (in brackets the corresponding ERC fields)

- ☐ Economics, Finance and Management (SH1, SH2)
- ☐ Linguistics, Cognition and Culture (SH3, SH4, SH5, SH6)
- ☐ Biochemistry, Bioinformatics and Life sciences (LS1, LS2, LS8, LS9)
- ☐ Physiology and Medicine (LS3, LS4, LS5 LS6, LS7)
- ☐ Mathematics and Computer Sciences (PE1, PE6)
- ☐ Fundamental Physics (PE2, PE3)
- ☐ Chemical Sciences and Materials (PE4, PE5)
- ☐ Engineering (PE7, PE8)
- ☐ Universe Sciences (PE9)
- ☒ Earth System Sciences (PE10)

2.3. Institutions and research team members

Laurent Terray, CERFACS, France, Sophie Valcke, CERFACS, France, Eric Maisonnave, CERFACS, France, Christophe Cassou, CERFACS, France, Klaus Wyser, Swedish Meterological and Hydrological Institute (SMHI), Sweden, Uwe Fladrich, Swedish Meterological and Hydrological Institute (SMHI), Sweden, Muhammad Asif, Catalan Institute of Climate Sciences, Spain, Domingo Manubens, Catalan Institute of Climate Sciences, Spain, Francisco Doblas-Reyes, Catalan Institute of Climate Sciences, Spain, Chandan Basu, Linkoping University, Sweden, Torgny Faxen, Linkoping University, Sweden, Wilco Hazeleger, Royal Netherlands Meteorological Institute (KNMI), Netherlands, Richard Bintanja, Royal Netherlands Meteorological Institute (KNMI), Netherlands, Camiel Severijns, Royal Netherlands Meteorological Institute (KNMI), Netherlands

2.4. Summary of the project *(Maximum 500 words)*
Please fill in the field with the same text used in the application form.

HiResClim aims to make major advances in the science of climate change modelling . This will be achieved by addressing the dual requirements of; increased climate model resolution and increased number of ensemble realizations of future climate conditions for a range of plausible socio-economic development pathways. Increased model resolution aims to deliver a significant improvement in our ability to simulate key modes of climate and weather variability and thereby provide reliable estimates of future changes in this variability. A large ensemble approach acknowledges the inherent uncertainty in estimating long-term changes in climate, particularly in phenomena that are highly variable and, of which, changes in the occurrence of the rare but intense events are those impacting society and nature most strongly. To provide credible risk assessment statistics on future change in phenomena such as; extra-tropical and tropical cyclones, heatwaves, droughts and flood events, the combination of high climate model resolution and a large ensemble approach is unavoidable. In HiResClim we attack both of these requirements in a balanced approach, which, as well as being the most efficient way to utilise the most advanced HPC systems of today, is also the only path to providing more robust and actionable estimates of future climate change.

2.5. Description of the results obtained from the scientific point of view, future perspectives, benefits to our society, and the benefits of using computer resources. *(Maximum of 1000 words)*

The success obtained at this stage did not allow analysing any of the simulations as they were not complete. Most of the effort invested in the last three months has focused on providing feedback to the BSC user support team to stabilize the machine to perform our simulations.

2.6. Expected future work in the area

A proposal for the second year of the project has just been submitted. It reformulates the original proposal for the second year framing a more ambitious set of simulations in the seamless climate paradigm.

2.7. Images of the results including description or caption *(Minimum resolution of 300 dpi)*
Please attach the images to this form.

2.8. Publications or reports regarding the developed project
(Format: Author(s). "Title". Publication, volume, issue, page, month year)

n.a.

2.9. Patents registered in relation with the developed project
(Format: patent identifier, title and description) (Maximum of 850 characters)

n.a.

2.10. Name and Surname of the students that have deployed their thesis collaborating in the developed project and title of the thesis.

n.a.

2.11. Talks given in the area of the project

- U. Fladrich "EC-Earth 3" at the Second Workshop on Coupling Technologies for Earth System Models (CW2013)

2.12. Other information (URLs, logos, photographies, etc.)
Please fill in the box with the regarding information and attach the photographies or logos to this form.

n.a.

2.13. Any further funding obtained as a result of the developed project

No further funding has been obtained, although most of the EU projects linked to this proposal (IS-ENES2, SPECS) have already started or are about to start. Full details are now provided in the second-year proposal.

3. Feedback and technical deployment

3.1. Feedback on the centers/PRACE mechanism *(Maximum 500 words)*

- A uniform utilisation of the resource over the allocation time, as preferred by PrACE, does not seem a realistic assumption. This demand, though understandable, conflicts with the computational requirements of the project's scientific experiments (spin-up runs followed by ensemble experiments), the need for a technical start-up phase (e.g. analysis and tuning of computational performance with respect to the specific experiments) and the unavoidable appearance of initial technical problems. Please, note that HiResClim was prepared to run on a mature machine like Curie, while the resources were finally offered on a machine unavailable at the time of starting the project like MareNostrum III.
- A preparatory phase of some weeks prior to the actual project accounting would be appreciated. This (low resource) phase could be optional and adjusted in time corresponding to the expected technical difficulties. In almost any case, some time is needed for porting the model as well as establishing the work-flow on the PrACE platform. In case of new system installations, an extended preparatory phase is considered appropriate. Overall, such a phase could facilitate a smooth utilisation of the PrACE resources.
- The complete work-flow of the use of coupled climate models should be taken into account and be reflected in technical considerations. This includes the interaction of computationally intensive experiments with post-processing phases and data storage.
- A better appreciation of the technical needs for coupled climate models is needed. These models constitute complex software and an involved development process. Hence, it is not desirable to adapt the models to a single purpose or a single computing platform. Scientific flexibility is a key issue for climate models, which implies rather bulky software, complex runtime environment, and diverse work-flow.

3.2. Explanation of how the computer time was used compared with the work plan presented in the proposal. Justification of discrepancies, especially if the computer time was not completely used. *(Maximum 500 words)*

The HighResClim project was granted starting 1st November 2012, however, access to the MareNostrum III machine could not be provided by BSC until the beginning of 2013, which delayed the project start for about 10 weeks.
A number of issues have appeared in the course of the current activities. All of them are being worked on with the help of the BSC PrACE support, but some still remain and hinder to forge ahead with the actual experiments of the project.

- **LSF job description headers**
  The length of job description headers is apparently limited by LSF to a certain number of bytes. This conflicts with the needs of autosubmit. A simple workaround has been applied, but a lifting of the limit would be advantageous.
- **Intel MPI**
  It turned out that Intel MPI, in its default configuration, was de-functional when used for the coupled model (multi executable) on a number of cores higher than 500. Standalone atmosphere experiments (single executable) were not affected. BSC has proposed a modified configuration, by means of environment variables, which allows to run some of the intended

experiments. Nevertheless, the high-resolution configuration of EC-Earth 3 in particular can still only be run on a very limited number of cores (615 cores shared by the atmosphere, ocean, and coupler components). This is the main obstacle for running the high-resolution spin-up runs as planned in the proposal.

- **OpenMPI**
  Initial tests suggest that the computational performance is substantially reduced (by a factor of two) by using OpenMPI as compared to Intel MPI. Some optimisation has been done by BSC, but the inferior performance as compared to InteMPI remains a fact and penalizes the experiment.
- **GPFS issues**
  A number of runs experience occasional (random) patterns of severe performance degradation, which are, according to BSC information, associated with GPFS problems. This affects the experiment work-flow as jobs are running out of time and fail before completion. In particular, the number of files created by the climate model seems to negatively affect performance and stability of the file system for the planned model runs.
- **Post-processing facilities**
  A number of nodes or a dedicated machine with large memory is needed to deal with post-processing of the output, specially for rebuilding the high resolution ocean data.

As a result of the listed system issues, the spin-up runs have been further delayed (beyond the access delay) according to the HiResClim application. Consequently, substantially less computer time than planned has been consumed at this time.
As a result of delays in developing an efficient performance of the high-resolution configuration of EC-Earth on MareNostrum III and the consequent likelihood of only have ~7-8 months of wall-clock time to perform the high-resolution EC-Earth spin-up run planned for year 1 of the project, we have modified our compute plans for year 1 so as to maximize the allocated cpu hours and scientific outcomes of the first year. Our original 2-year proposal envisaged the first year primarily consisting of 2 pre-industrial spin-up runs of EC-Earth (T511L91/ORCA025L75) and ARPEGE (T359L31/ORCA025L75) for ~700 simulated years. Due to the significantly higher resolution of the EC-Earth atmosphere, it will likely now not be feasible to make a single 700 year spin up run during the remainder of year 1 with this configuration. To maximize use of resources we are now pursuing 4 spin-up simulations during year 1 that can all be run in parallel, the 2 original planned runs, plus 2 further configurations of EC-Earth. The latter 2 configurations both utilise a lower resolution atmosphere in EC-Earth (T255, ~75km instead of 39km at T511) with this atmosphere coupled to the same NEMO ocean model, run respectively at 0.25° and 1° global resolution. Analysis of these 3 EC-Earth spin-up runs will allow us to assess the benefits of (i) increasing atmospheric model resolution with a common ocean resolution, (ii) increasing ocean model resolution with a common atmosphere and (iii) higher resolution in both model components, in the context of simulating multi-decadal climate variability. In year 2 of HiResClim *(requested through the 7th PrACE call)* only results from the EC-Earth high resolution and ARPEGE spin up runs will be used to initialize 20th century historical climate simulations and subsequent climate projections and predictions at the same resolutions. Output from the 2 lower resolution EC-Earth spin-up runs will be used to initialize equivalent-resolution and time period 20th century and projection/prediction experiments on either partner national HPC systems or Tier 1 systems available through competitive national PrACE instruments. Following this strategy we maximize the outcomes of year 1 resources and through synergy with use of national and Tier 1 platforms are able to address a number of important questions with respect to the role of both ocean and atmospheric model resolution with respect to simulating future,

anthropgenic-induced, climate change and in the context of seasonal and decadal climate prediction.

3.3. Please, let us know if you plan to apply for a regular PRACE project? If not, explain us why. *(Maximum 500 words)*

A second-year proposal has been submitted to the 7th PrACE tier-0 open call.