PRACE application form

7th PRACE Project Access Call

Detailed Project Document

Project name	HiResClim: High Resolution Ensemble Climate Modeling
Research field	PE10_3 Climatology and climate change

Project leader

Title	Professor
Last name	Doblas-Reyes
First name	Francisco
Organisation name*	Institut Català de Ciències del Clima (IC3)
Department*	Climate Forecasting Unit
Group*	
Country	Spain

Research project: Science aims and description

Over the past few years the international climate modelling community has collaborated around the Fifth Coupled Model Intercomparison Project (CMIP5, http://cmip-pcmdi.llnl.gov/cmip5, Taylor et al. 2012), producing a coordinated set of climate simulations to support the Fifth Assessment Report of the Intergovernmental Panel for Climate Change. For the first time, CMIP5 promotes two experiment protocols: climate projections for the rest of the century and initialized climate predictions.

In the climate projection protocol a coupled general circulation model (GCM) is integrated for many hundreds of years, using estimates of pre-industrial greenhouse gas (GHG) and aerosol concentrations, solar intensity, volcanic aerosols and land-use. Once the model attains a stable equilibrium with the forcing, a number of model states are saved that sample the simulated climate variability. Each of these states is used as an initial condition to extend the pre-industrial simulation through the historical past (in CMIP5, 1850-2005), using observed estimates of GHGs, anthropogenic, natural and volcanic aerosols, land-use and solar variability. These simulations are evaluated against observations to determine the quality of present-day GCMs. The historical runs are then extended into the future, using different GHG, aerosol and land-use scenarios that sample a range of plausible socio-economic futures up to ~2100, referred to as representative concentration pathways (RCPs, Moss et al. 2010). Application of the same RCPs to different GCMs for the period 2006-2100 supports a multi-model ensemble of climate projections.

Climate prediction differs from the projections by using observation-based initial conditions to start an ensemble of GCM simulations from a known, recent date, with the initialized climate system providing predictability beyond typical atmospheric time scales. Climate prediction, which typically involves large ensembles of up to 10-year long simulations, has recently emerged as a possibility due to the convergence of important improvements in (i) ocean observational coverage (ii) GCM simulation quality (iii) our understanding of processes providing inter-annual predictability (Doblas-Reyes et al. 2013).

Climate prediction is, however, a wider concept than that advanced by CMIP5 and includes, among others, seasonal forecasting which aims to provide probabilistic climate information at time scales of up to one year. In the wake of the World Modelling Summit in 2008, the seamless concept (Shukla et al. 2009) was coined where the use of numerical weather forecasting and climate prediction methods can help quantify and reduce uncertainty in climate predictions. The seamless concept is used as a paradigm to enable climate simulations addressing a range of different time scales to contribute to an overall improvement of climate modelling. In particular, there is an increasing use of seasonal predictions to diagnose systematic problems in climate modelling (Vanniere et al. 2013).

The EC-Earth (Hazeleger et al. 2012, http://ecearth.knmi.nl, https://dev.ec-earth.org) and ARPEGE-NEMO GCMs contributed to the projection and prediction protocols of CMIP5 and regularly perform seasonal predictions, using resolutions of ~125-150 km and ~1° in the atmosphere and ocean, respectively. While such resolution compares favourably with other CMIP5 models, it is poor in terms of resolutions required for an accurate simulation of important modes of climate and weather variability. Higher model resolution has been suggested to improve many aspects of climate simulations. Jung et al. (2012) and Kinter et al. (2013) found significant improvements in the simulation and prediction of many atmospheric features such as tropical precipitation and the frequency/intensity of both tropical and mid-latitude cyclones in IFS-only simulations (the atmospheric component of EC-Earth) after using dedicated high-end computing resources to enable high spatial resolution. Caron et al. (2011) using a variable resolution AGCM found significant improvements in the representation of Atlantic tropical cyclones from 150 to 30 km resolution. Resolving mesoscale ocean eddies allows for a more realistic representation of ice drift and deformation (Zhang et al. 1999) and frontal scale air-sea interaction (Bryan et al. 2010). Recently, a couple of studies reported the first use of coupled climate models in the resolution range ~50 km (atmosphere) and ~10-30 km (ocean), identifying large improvements in a number of systematic errors seen in lower resolution versions of the same models. These include a reduction in the double ITCZ problem (GFDL CGCM; Delworth et al. 2012), improved intensity distribution of tropical rainfall (MIROC4h CGCM, Sakamoto et al. 2012), improved coastal upwelling and ENSO variability (Sakamoto et al. 2012) and improved structure of the North Atlantic Ocean circulation (Delworth et al 2012).

Potential changes in these modes of climate variability in response to near future GHG and aerosol forcing will likely impact society and ecosystems far more than a slow change in the time-mean climate. It is therefore important that GCMs realistically simulate such variability. In HiResClim we address this and other problems through experiments based on CMIP5 and standard seasonal prediction protocols, employing significantly increased resolution. Specifically, the atmosphere will use T359 (~55 km) and 31 vertical levels and T511 (~39 km) and 91 levels in ARPEGE-NEMO HR and EC-Earth, respectively, with 0.25° horizontal resolution and 75 vertical levels for the NEMO ocean model in both cases. Compared to the original CMIP5 simulations, horizontal resolution is increased by a factor of up to 3.5, with increased number of vertical levels in both ocean and atmosphere.

Following the production of the spin-up, pre-industrial simulations carried out during the first year of the project, three different model states from the GCM pre-industrial run will initialize a three-member ensemble of historical simulations. Each of the three historical runs realized with both models will spawn a future projection following the RCP4.5 scenario covering the period 2006-2100; for EC-Earth, additional future projections following the RCP8.5 scenario will also be performed. Such an ensemble is necessary to establish probabilistic estimates of future changes in the important, but highly variable, weather and climate phenomena listed above. This ensemble will be augmented over time using tier-1 HPC resources available to the partner institutes.

Both the seasonal and decadal predictions, which contribute to the EU FP7 project SPECS led by the HiResClim project PI, will be performed throughout the duration of the project taking advantage of the fact they are made up of a large number of independent simulations.

Climate prediction has to deal with the dual requirement of increasing model resolution while also performing enough predictions over the observed past to formulate a robust estimate of forecast quality. The ensemble method used acknowledges the inherent uncertainty in estimating future climate states. These are the two independent dimensions, ensemble size and number of start dates, which additionally increases the cost of this part of the experiment. The partners can adjust the production of the predictions to the progress of the historical and RCP simulations, ensuring a uniform use of resources across the year. However, the historical simulations are an important prerequisite to some of the high-resolution predictions. In particular, a present-day coupled climatology is required for the observation-based anomaly initialization procedure developed for EC-Earth (Hazeleger et al. 2013). The anomaly-initialized EC-Earth predictions will be compared to predictions using full-field initialization with the same model. The historical CMIP5 simulations also provide a set of uninitialized model runs against which predictability improvements in the initialized forecasts can be evaluated. The high-resolution predictions made in HiResClim for the different time scales will also be compared to lower resolution predictions, as planned in SPECS, with the latter simulations being performed on national HPC allocations This experimental set up is the most ambitious high-resolution, multi-model ensemble of simulations, predictions and projections planned to this date to explore the seamless concept HiResClim results will be saved on the Earth System Grid, thus becoming openly available to the wider research community.

This experimental set up is the most ambitious high-resolution, multi-model ensemble of simulations, predictions and projections planned to this date to explore the seamless concept.

Management of the resources requested

The schedule of the experiment integration is detailed in response to Question 6. Figure 1 illustrates the approximate schedule for the second-year HiResClim experiments.

As the models have already been run on the requested system during the first year of the project most of the planned simulations can begin straightaway. The first project year is dominated by the core 700-year pre-industrial spin-up experiments that are necessary for the subsequent historical and future simulations. In the case of EC-Earth, the spin-up experiments are carried out using three different resolutions of EC-Earth, T255L91-ORCA1L46, T255L91-ORCA025L75 and T511L91-ORCA025L75. This decision was taken for several reasons: 1) to make a more efficient use of the first-year resources given the problems encountered to date on Marenostrum III, 2) to create restarts to perform historical simulations at the three resolutions and 3) to assess the impact of the resolution in the spin-up process and subsequent prediction/projection experiments. Due to the extreme cost of running the T511L91-ORCA025L75 version of EC-Earth, it has been planned to use the lower resolution spin-up runs to investigate a range of approaches to accelerate the production of balanced, spun-up states for subsequent use as initial conditions for high-resolution historical simulations. The second-year retains this basic strategy, except only the high-resolution experiments will be run under HiResClim and experiments using the two lower resolutions will be performed on other platforms using the partners’ or competitive resources, potentially including the PRACE DECI instrument. These in-kind contributions will allow an assessment of the impact of resolution at all stages of the climate modelling process. For ARPEGE-NEMO, the spin-up experiment is done with the T359L31-ORCA025L75 resolution only.

For each model, the three historical simulations will be made as soon as the spin-up runs are complete. Model states to start the historical runs will be drawn from the last ~150 years of the pre-industrial runs, hence the first historical runs may start towards the end of the first year, December 2013. These simulations are independent of each other and will be run in parallel to make efficient use of available compute nodes and a uniform use of the resources through the year. Each historical run requires ~60-90 wall-clock days.

The RCP scenarios will start between months 3 and 5 of the second year. As these start from the end of the historical simulations, the complete set of historical and RCP simulations may require ~6 to 8 months of continuous machine use. The additional months to the end of the year will be used as a buffer should any unforeseen model or system problems arise. This is a possibility given that such an ambitious experiment has never been performed before.

The climate predictions, both seasonal and decadal, will be performed throughout the year taking advantage of the fact that they are made of a large number of independent simulations. For instance, the multi-annual predictions are made of ten-member ensemble, five-year long simulations started once a year (on the first of November) over the period 1993-2009 using two (full-field and anomaly) initialization methods for EC-Earth but only the full-field initialization method for ARPEGE-NEMO HR. The selected period is when high-resolution ocean initial conditions from MERCATOR are available. This makes a total of 170 individual simulations (10 members times 17 start dates over 1993-2009) per type of initialization, which can all run in parallel. The anomaly-initialization predictions will be run last as they require the historical simulations to be available to estimate the mean model climate. The Autosubmit software, briefly described below, will be used to manage the work flow to adapt it to the progress of the historical and RCP simulations and ensure a uniform use of resources. The extension of some of the multi-annual simulations to ten forecast years planned for the EC-Earth model has been scheduled in months 8 to 10 as confirmation that the first five forecast years provide satisfactory results is needed first. Finally, the seasonal predictions also performed with EC-Earth only offer the largest computing flexibility as they are a set of more than 1,000 simulations that will be managed, and packed in groups in a single executable if required, by Autosubmit to optimize the use of the machine and avoid collapsing the I/O system. It has to be noted that seasonal predictions are currently being realized with ARPEGE-NEMO HR in the framework of the SPRUCE PRACE project; the experiments planned here with EC-Earth at even higher resolution for the atmosphere (T511L91 versus T359L31) will extend the SPRUCE project allowing us to quantify the scientific added value of increased resolution.

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
Historical runs (3 members, 1850-2005)	🟦	🟦	🟦	🟦	🟦							
Scenarios RCP4.5 and 8.5 (3 members, 2006-2100)					🟦	🟦	🟦	🟦				
Multi-annual predictions of both full and anomaly initialization with annual restarts with 5-year forecast length (10 members, 1993-2009)	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦		
Seven start dates of the multi-annual predictions of both full and anomaly initialization extended to 10-year forecast length (10 members, 1993-2009)								🟦	🟦	🟦		
Seasonal predictions with three start dates (Feb, May, Aug) per year with seven-month forecast length (20 members, 1993-2009)										🟦	🟦	🟦

Figure 1. Approximate schedule of the second-year HiResClim experiments.

Numerical methods and algorithms

The main concern for optimum use on Tier-0 machines is the problem size, combined with the ability to run numerous independent ensemble members in parallel.

The atmosphere-land component of ARPEGE-NEMO HR is the operational weather forecast model of Météo-France, while EC-Earth is based on cycle 36R4 of the ECMWF Integrated Forecast System (IFS). The basic code infrastructure of IFS and ARPEGE has been shared for the past ~15 years allowing close collaboration on development targeting optimum performance on a range of HPC systems. The resolution proposed for HiResClim (ARPEGE: T359, ~300.000 horizontal grid points and EC-Earth: T511, ~425,000 points) will help efficiently share calculations between 1000-1500 sub-domains, increasing the range of efficient compute-core usage per model executable. The need to make an ensemble of simulations (~3 members for the historical and scenario experiments and 10-20 for the climate predictions), each independent of the other, also increases the efficient use of the architecture and ensures a smooth use of the resources through the year.

NEMO, developed by a European consortium, is the ocean component in both models. NEMO uses the OASIS coupler, developed and maintained by CERFACS, for coupling atmosphere-land models. NEMO has already been adapted for Tier-0 machines by several ENES (European Network for Earth System Modeling) laboratories, including all four in this consortium, in collaboration with the PRACE-1IP support team. This group optimized a very high-resolution (T799/ORCA025) version of EC-Earth. It undertook the core mapping of MPI tasks (lead by John Donners, SARA), the implementation of dynamical memory allocation and load-balanced domain decomposition (lead by Andrew Porter, STFC) and the design of optimal coupling interface and strategy (lead by CERFACS) on different architectures. The chosen problem size for NEMO (ORCA025, 1442 x 1021 horizontal grid points) enables a distribution to 1000 sub-domains. An important development for HiResClim is that the models now uses the latest version of the OASIS coupler (OASIS-MCT), which enables a direct and parallel exchange of coupling fields between every model process, avoiding the previous coupler-centric bottleneck. HiResClim provides an opportunity to measure OASIS-MCT performance running at high-resolution. This will provide information to other European modelling groups also using the OASIS software.

The HiResClim 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 minimum amount of necessary diagnostics will be stored on disk and immediately post-processed for transfer to final storage. The partners expect to test the newly developed XIOS IO server, including an option to save diagnostic data on a reduced grid, although the development required might prevent the ongoing and planned simulations to take advantage of this facility. Experience from this effort will inform other European groups investigating the XIOS IO server within IS-ENES2.

Since output transfer will be performed as the simulations are produced, the partners will organize the data storage and data transfer with a disk space of 20 Tb in the “scratch” file system, when the total amount of data to be transferred amounts to ~200 Tb. The required scratch space is motivated by the requirement to perform many simulations at the same time.

This is a desirable feature of our experimental set up and is made possible because there is no dependency between the members of an ensemble or the predictions for different start dates. As a high-resolution one-year forecast with the high-resolution configuration can produce up to 100 Gb of raw data, a realistic number of maximum simulations that can be performed simultaneously on the chosen platform such as five ensemble members for three independent start dates require about 30,000 cores, will produce ~1.5 Tb per year of simulation (or every six hours of wall clock time). These data will be held in the work space before it being post-processed to reduce the data volume to less than a quarter of the size, which will be transferred to the "archive" file system and then to the local institutions.

Around 500 Gb of home space will be required to host the code and its modified versions.

Explain why this project needs to run on a Tier-0 system

This proposal is an important step in furthering the ENES preparations for efficient use of PRACE Tier-0 and other large-scale HPC systems. The EU FP7 project IS-ENES2 is due to start in April 2013 and has as a central goal the technical development of European climate and Earth system models for efficient application on high-end HPC systems. In HiResClim, the two models employed address this aim concentrating on the two leading science requirements for accurate and reliable climate forecasting, namely increased model resolution and the use of a large ensemble. The former is critical for models to realistically represent the observed power spectrum of climate and weather variability, while the latter is necessary for developing robust estimates of future climate changes and for determining both seasonal and near-term climate predictability. The substantial increase in model resolution in HiResClim, combined with the required reduction in model time step, increases the computational cost of one simulated year by ~20 compared to that experienced with CMIP5 lower-resolution coupled models. Combined with the aim to produce a suitable ensemble size and the need for either long runs or large sets of past predictions to generate robust climate estimates, the proposed set of core experiments are not feasible on Tier-1 or 2 facilities. They are hence most efficiently addressed on a major Tier 0 system, particularly Marenostrum III as it is the machine where the first year of the project is being carried out. Lower resolution versions of these experiments will be performed on machines other than the Tier-0 platform requested here, as explained below.

It is important to emphasize that in HiResClim many of the planned simulations are independent of each other, allowing an ensemble of simulations to be bundled into a single job for efficient use of parallel architectures. IC3 recently tested EC-Earth3 managed by the Autosubmit system (Johnson et al. 2013) on the Cray XE Jaguar in collaboration with the INCITE programme of the US Department of Energy, achieving an efficient use of ~60,000 cores with a single executable wrapping an ensemble of high-resolution integrations. The Autosubmit option will be extensively used in the project, where either more than 20 seasonal or decadal predictions or 6 RCP future projections (2 RCPs times 3 members) can all run in parallel.

Splitting the project across multiple platforms will increase significantly the amount of enabling work for the project. A good example of how important this is can be read in the progress report of the first year of HiResClim, where important delays and technical problems experienced on Marenostrum III are described. The amount of work required to port the code and to set up the monitoring and post-processing systems will likely be multiplied by the number of platforms, and therefore might compromise the achievement of such an ambitious experiment in a reasonable time. We hence suggest avoiding such a splitting at all costs.

While in the first year of the project 38 million core hours were obtained, the second year an increase is requested. However, the capability of running more jobs at the same time, up to 15-20 (depending on the total number of cores available and the final choice of number of cores per job, but with the possibility of easily using up to 40,000 cores) compared to just three jobs for each EC-Earth resolution and for ARPEGE-NEMO HR in the first year (see progress report), is an important difference between the plans for each year. Obviously, the management system should be able to create different I/O patterns for the simulations corresponding to the different members and/or start dates.

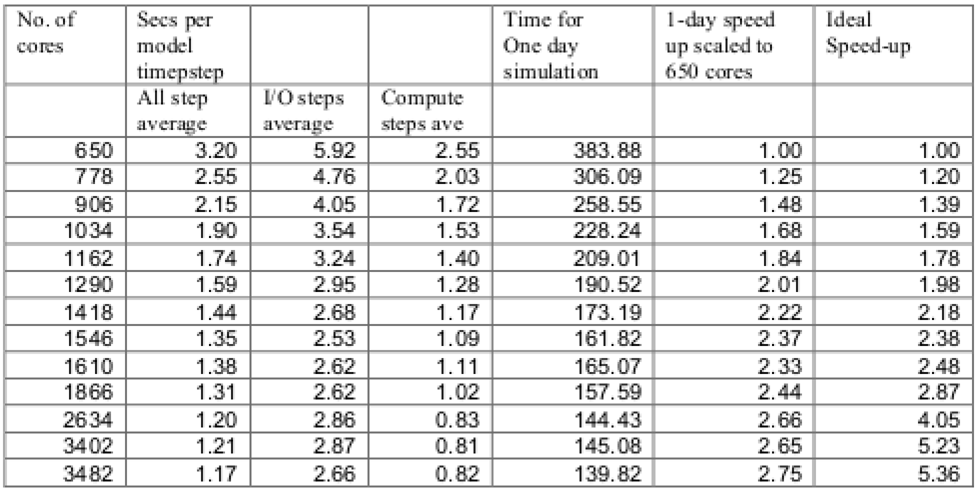

No. of cores	Secs per model timestep			Time for One day simulation	1-day speed up scaled to 650 cores	Ideal Speed-up
	All step average	I/O steps average	Compute steps ave			
650	3.20	5.92	2.55	383.88	1.00	1.00
778	2.55	4.76	2.03	306.09	1.25	1.20
906	2.15	4.05	1.72	258.55	1.48	1.39
1034	1.90	3.54	1.53	228.24	1.68	1.59
1162	1.74	3.24	1.40	209.01	1.84	1.78
1290	1.59	2.95	1.28	190.52	2.01	1.98
1418	1.44	2.68	1.17	173.19	2.22	2.18
1546	1.35	2.53	1.09	161.82	2.37	2.38
1610	1.38	2.62	1.11	165.07	2.33	2.48
1866	1.31	2.62	1.02	157.59	2.44	2.87
2634	1.20	2.86	0.83	144.43	2.66	4.05
3402	1.21	2.87	0.81	145.08	2.65	5.23
3482	1.17	2.66	0.82	139.82	2.75	5.36

Table 1. Timings of EC-Earth T799/0.25° on the Curie Fat Node system, relative to 650 cores. Also shown are the average time (in secs) for a single model time step (All step) and for only I/O steps and only compute steps (with no I/O occurring).

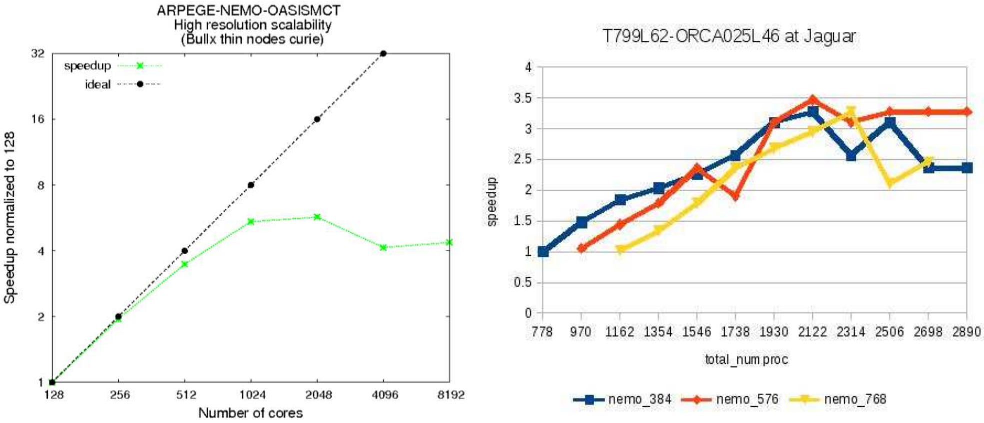

Figure 2. (Left) Speed up of ARPEGE-NEMO T359/0.25º configuration relative to 128 core performance on the Curie thin node system versus a perfect scaling line. (Right) Speed up of the EC-Earth T799L62/025L46 configuration on Jaguar (Cray XE6) as a function of the number of cores expressed relative to 778 cores for different number of NEMO cores: 384 (blue), 576 (orange) and 768 (yellow). The best efficiency was obtained with 2122 cores (1536 for IFS, 576 for NEMO and 10 for OASIS).

An important delay has been experienced up to now in the first year, as described in the progress report made available to the PRACE office. This delay, due to causes beyond the partners’control, has prevented an extensive technical performance (short integrations with different numbers of cores) of the model on Marenostrum III to this date. However, the models have been tested in other HPCs like Curie and Jaguar. For instance, scaling work on a PrACE Tier 1 Cray XT6 in Sweden with the EC-Earth T511/ORCA025 configuration suggests the potential for a doubling of the current throughput on Marenostrum III to reach the estimates used in this proposal, once sufficient scaling analysis has been performed. Figure 2 and table 1 show scaling details of EC-Earth T799/ORCA025 from the IS-ENES/PRACE1IP evaluation and the collaboration with the INCITE programme, and for ARPEGE-NEMO T359/ORCA025.

The experience gained in making such a coordinated, multi-model high-resolution ensemble will be of huge importance to the ENES community. In particular, the experience will feed directly into a major JRA in IS-ENES2, that brings together the majority of European climate models around a common multi-GCM, multi-member integration protocol to address bottlenecks to efficient GCM use on high-end HPC systems. A second EU FP7 large-scale project, SPECS, has a primary focus on improving European seasonal-to-decadal climate prediction capabilities. SPECS also has a work package addressing the development and application of a high-resolution multi-GCM, multi-member ensemble of seasonal-to-decadal predictions. All institutes in HiResClim are partners in both IS-ENES2 and SPECS.

Describe your experience of using HPC resources in the past

The partners have more than two decades experience developing and running GCMs on a range of HPC systems. Both project models made extensive contributions to the projection and prediction protocols in CMIP5. The ARPEGE-NEMO simulations were a collaboration between CERFACS and Météo-France, while KNMI, SMHI and IC3 were three of the most active EC-Earth institutes. The partners also have contributed extensively to a large number of past and on-going EU projects, including DEMETER, PREDICATE, DYNAMITE, PRUDENCE, ENSEMBLES, COMBINE, IS-ENES, EMBRACE (coordinator Jones, SMHI), as well as two large-scale projects started in November 2012 and focused on the development of climate services: SPECS (coordinator Doblas-Reyes, IC3) and EUPORIAS. CERFACS is the leader of the PRACE SPRUCE project, which targets suites of seasonal predictions on Curie using ARPEGE-NEMO HR.

CERFACS is currently leading major improvements to the OASIS code (Valcke 2013) to make optimal use of the Tier-0 machines, currently available in the new OASIS3-MCT version. Maisonnave and Valcke, within the IS-ENES project, have worked extensively on GCM coupling on HPC resources and provided support on OASIS-based coupling to several European laboratories. In collaboration with PRACE-2IP, a study on NEMO model fault-tolerance compliance has been initiated (Maisonnave et al. 2011).

Every HiResClim institute has been involved in the IS-ENES/PRACE-1IP working group focusing on the EC-Earth adaptation to Tier-0 machines. The KNMI team developed the first climate version of EC-Earth (version 2.3) that was subsequently used for all CMIP5 simulations. Since then, SMHI led the development of the latest EC-Earth (version 3), the one used in this proposal, and have tested a range of EC-Earth configurations, in the atmospheric resolutions T255/511/799 on several HPC systems: SGI Altix 3500, NEC-SX6, Linux cluster with Intel Xeon, Dell PowerEdge 2900, IBM pSeries 575 Power6 and IBM Power PC.

IC3 leads the development of Autosubmit, a python-based wrapper software that can manage any type of climate simulation workflow. It can also bundle numerous jobs into a single executable, performing job-control on parallel sets of simulations throughout the execution.

Justify the number of core hours requested

Table 2 lists the core and optional runs planned for the second year of the project. These assume a restart frequency (single job run length) of 3 simulated months for both models. Depending on progress in year 1, some year 2 simulations may be able to start towards the end of year 1. In year 2, all simulations can run in parallel: the 3 historical simulations can operate in parallel, as can the subsequent 6 RCP projection runs and the different climate predictions, for which each start date and member is, again, an independent simulation.

The cost of ARPEGE-NEMO HR is 0.6 times that of EC-Earth, for which a three-month simulation requires a wall-clock time of 1.5 hours. Hence, the wall time per step for those experiments planned with both models is 2.4 hours, the sum of 1.5 for EC-Earth and 0.9 for ARPEGE-NEMO HR. ARPEGE-NEMO HR will only perform a subset of the experiments as others, such as the seasonal predictions, are performed in PRACE SPRUCE.

Run type	# Runs	# Steps/Run	Walltime/Step	# CPU cores	Total core hours/Type Run
Historical runs (2 models, 3 members, 1850-2005)	3	660	2.4 h	2048	9.7 Mhours
Scenario RCP4.5 (2 models, 3 members, 2006-2100)	3	380	2.4 h	2048	5.6 Mhours
Scenario RCP8.5 (1 model, 3 members, 2006-2100)	3	380	1.5 h	2048	3.5 Mhours
Multi-annual predictions of full initialization with annual restarts with 5-year forecast length (2 models, 10 members, 1993-2009)	10x17=170	20	2.4 h	2048	16.7 Mhours
Multi-annual predictions of anomaly initialization with annual restarts with 5-year forecast length (1 model, 10 members, 1993-2009)	10x17=170	20	1.5 h	2048	10.5 Mhours
Seven start dates of the multi-annual predictions extended to 10-year forecast length (full initialization, 1 model, 10 members, 1993-2009)	10x7=70	20	1.5 h	2048	4.3 Mhours
Seven start dates of the multi-annual predictions extended to 10-year forecast length (anomaly initialization, 1 model, 10 members, 1993-2009)	10x7=70	20	1.5 h	2048	4.3 Mhours
Seasonal predictions with three start dates (Feb, May, Aug) per year with seven-month forecast length (1 model, 20 members, 1993-2009)	20x17x3=1020	2	1.7 h	2048	7.1 Mhours
Total					**61.7 Mhours**

Table 2. Cost of the experiments planned for the second year of HiResClim.

The final estimate is for a total request of 65 Mhours, which includes the numbers described in table 2 plus a small buffer of 5% to account for failing jobs that will need to be repeated.

The cost of the low-resolution EC-Earth version of the experiments is 54% (EC-Earth T255-ORCA025) and 8% (EC-Earth T255-ORCA1) of those described in table 2, which illustrates how important the in-kind contributions to HiResClim are.

References

Bryan, F.O., R. Tomas, J.M. Dennis, D.B. Chelton, N.G. Loeb and J.L. McClean, 2010. Frontal scale air-sea interaction in high-resolution coupled climate models. J. Climate, 23, 6277-6291, doi: 10.1175/2010JCLI3665.1.

Caron L-P, C.G.Jones and K.Winger, 2011. Impact of resolution and downscaling technique in simulating recent Atlantic tropical cyclone activity. Climate Dyn., 37, 869-892.

Delworth, T.L., A. Rosati, W. Anderson, A.J. Adcroft, V. Balaji, R. Benson, K. Dixon, S.M. Griffies, H.-C. Lee, R.C. Pacanowski, G.A. Vecchi, A.T. Wittenberg, F. Zeng and R. Zhang, 2012. Simulated climate and climate change in the GFDL CM2.5 High-resolution coupled climate model. J. Climate, 25, 2755-2781, doi:10.1175/JCLI-D-11-00316.1.

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 Comm., doi:10.1038/ncomms2704.

Hazeleger, W., X. Wang, C. Severijns, S. Ştefănescu, R. Bintanja, A. Sterl, K. Wyser, T. Semmler, S. Yang, B. van den Hurk, T. van Noije, E. van der Linden and K. van der Wiel, 2012. EC-Earth V2.2: description and validation of a new seamless earth system prediction model. Climate Dyn., 39, 2611-2629, doi:10.1007/s00382-011-1228-5.

Hazeleger, W., V. Guemas, B. Wouters, S. Corti, I. Andreu-Burillo, F. J. Doblas-Reyes, K. Wyser and M. Caian, 2013. Multiyear climate predictions using two initialisation strategies. Geophys. Res. Letters, doi:10.1002/grl.50355.

Johnson, C., A. Carter, I. Bethune, K. Statford, M. Alava, V. Cardoso, M. Asif, B.S.A. Schuberth, T. Weinzierl, 2013. PRACE DECI (Distributed European Computing Initiative) Minisymposium. Applied Parallel and Scientific Computing - Lecture Notes in Computer Science, 7782, 43-60, doi:10.1007/978-3-642-36803-5_3.

Jung, T., M.J. Miller, T.N. Palmer, P. Towers, N. Wedi, D. Achuthavarier, J.D. Adams, E.L. Altshuler, B.A. Cash, J.L. Kinter III, L. Marx, C. Stan and K.I. Hodges, 2012. High-resolution global climate simulations with the ECMWF model in Project Athena: Experimental design, model climate and seasonal forecast skill. J. Climate, 25, 3155-3172, doi:10.1175/JCLI-D-11-00265.1.

Kinter III, J.L. , B.A. Cash, D. Achuthavarier, J.D. Adams, E.L. Altshuler, P.A. Dirmeyer, B. Doty, B. Huang, L. Marx, J. Manganello, C. Stan, T. Wakefield, E.K. Jin, T.N. Palmer, M. Hamrud, T. Jung, M.J. Miller, P. Towers, N. Wedi, M. Satoh, H. Tomita, C. Kodama, H. Taniguchi, P. Andrews, T. Baer, M. Ezell, C. Halloy, D. John, B. Loftis, R. Mohr, K. Wong, T. Nasuno, K. Oouchi and Y. Yamada, 2013. Revolutionizing climate modeling with Project Athena: A multi-institutional, international collaboration. Bull. Amer. Meteor. Soc., 94, 231-245, doi:10.1175/BAMS-D-11-00043.1.

Maisonnave, E., L. Coquart and S. Valcke, 2011. PRACE preparatory access for high resolution ARPEGE-NEMO porting on Bullx TGCC platform. Working Note, WN/CMGC/11/77, SUC au CERFACS, URA CERFACS/CNRS No1875, France.

Moss, R.H, J.A. Edmonds, K.A. Hibbard, M.R. Manning, S.K. Rose, D.P. van Vuuren, T.R. Carter, S. Emori, M. Kainuma, T. Kram, G.A. Meehl, J.F.B. Mitchell, N. Nakicenovic, K. Riahi, S.J. Smith, R.J. Stouffer, A.M. Thomson, J.P. Weyant and T.J. Wilbanks, 2010. The next generation of scenarios for climate change research and assessment. Nature, 463, 747-756, doi:10.1038/nature08823.

Sakamoto T.T. Y. Komuro, T. Nishimura, M. Ishii, H. Tatebe, H. Shiogama, A. Hasegawa, T. Toyoda, M. Mori, T. Suzuki, Y. Imada, T. Nozawa, K. Takata, T. Mochizuki, K. Ogochi, S. Emori, H. Hasumi and M. Kimoto, 2012. A new high-resolution atmosphere-ocean coupled general circulation model. J Japan Met Soc, 90, 3, 325-359, doi:10.2151/jmsj.2012-301.

Shukla, J., R. Hagedorn, M. Miller, T.N. Palmer, B. Hoskins, J. Kinter III, J. Marotzke and J. Slingo, 2009. Strategies: Revolution in climate prediction is both necessary and possible: A declaration at the World Modeling Summit for Climate Prediction. Bull. Amer. Meteor. Soc., 90, 16-19, doi:10.1175/2008BAMS2759.1.

Taylor, K.E., R.J. Stouffer and G.A. Meehl, 2012. An overview of CMIP5 and the experiment design. Bull. Amer. Meteor. Soc., 93, 485-498, doi:10.1175/BAMS-D-11-00094.1.

Valcke, S., 2013. The OASIS3 coupler: a European climate modelling community software. Geosci. Model Dev., 6, 373-388, doi:10.5194/gmd-6-373-2013.

Vannière, B., E. Guilyardi, G. Madec, F.J. Doblas-Reyes and S. Woolnough, 2013. Using seasonal hindcasts to understand the origin of the equatorial cold tongue bias in CGCMs and its impact on ENSO. Climate Dyn., 40, 963-981, doi:10.1007/s00382-012-1429-6.

Zhang, Y., W. Maslowski and A. J. Semtner, 1999. Impact of mesoscale ocean currents on sea ice in high-resolution Arctic ice and ocean simulations. J. Geophys. Res., 104, 18409-18429, doi:10.1029/1999JC900158.