Project name

Project name	CMIP6 BSC contribution to HighResMIP (HighResMIP_BSC)
Research field	PE10_3 Climatology and climate change

Project leader

Title	Dr
Last name	Guemas
First name	Virginie
Organisation name*	Barcelona Supercomputing Center (BSC)
Department*	Earth Sciences Department
Group*	Climate Prediction Group
Country	Spain

IMPORTANT NOTICE

Please **upload a single document**, using this template, in PDF format, which must not exceed 8 MB and addresses **ALL of the points below**. The minimum font size that you can use is 12 point and you must adhere to the page limits indicated below. The **maximum number of pages you may submit, including references is 10 pages** (pages out of the limit are not considered).

1. Describe your research project. Include discussion of the scientific questions that you are planning to address and the overall scientific goals of the project. It is important that you describe the novelty, impact and timeliness of the proposal. (This section must be no longer than 2 pages).

Recent studies have established that the typical atmospheric and oceanic resolutions used for the CMIP5 coordinated exercise (Coupled Model Intercomparison Project, phase 5) are a limiting factor to correctly reproduce climate mean state and variability. Noteworthy improvements when increasing resolutions have been obtained in the simulation of El Niño Southern Oscillation (ENSO) (Shaffrey et al. 2009, Masson et al. 2012), Tropical Instability Waves (Roberts et al. 2009), the Gulf Stream and its influence on the atmosphere (Chassignet and Marshall 2008; Kuwano-Yoshida et al. 2010), the global water cycle (Demory et al. 2014), jet stream (Lu et al., 2015), storm tracks (Hodges et al. 2011), Euro-Atlantic blockings (Jung et al 2012), tropical cyclones (Zhao et al. 2009, Bengtsson et al. 2007), tropical-extratropical interactions (Baatsen et al. 2014, Haarsma et al. 2013), monsoons (Sperber et al. 1994; Lal et al. 1997; Martin 1999), heat waves and droughts (Van Haren et al. 2015) and sea ice drift and deformation (Zhang et al 1999; Gent et al, 2010).

The main obstacle to running climate models at higher resolution is computational. Up to now, relatively few research centres have carried out such highly-demanding simulations. These were relatively short experiments, with very few or only one member and not performed in a coordinated way. The HighResMIP coordinated exercise, as part the Sixth Phase of the Coupled Model Intercomparison Project (CMIP6), offers a framework for increasing synergies and building a large multi-model ensemble of high resolution simulations with a low resolution counterpart following a common experimental protocol, i.e. a common integration period, forcing and boundary conditions. This coordinated exercise will allow for an optimization of the computing resources usage toward a common goal of identifying the robust benefits of increased model resolution based on multi-model ensemble simulations. This is the first time ever such a largely demanding computing exercise will be implemented and the largest ensemble ever built at such high resolution. The ensemble spread will allow diagnosing uncertainties and attributing them to model formulation. The CMIP6 framework ensures a high level of participation to this project.

Process-based assessment will be central to HighResMIP. The main goal is to pin down physical and dynamical reasons behind differences in model representation induced by resolution change. Such a process-based assessment is supported by the PRIMAVERA project, funded by the European commission under the H2020 framework programme in which BSC is highly involved. New diagnostics and metrics will be developed which fulfil essential criteria such as: 1) focusing on processes or feedbacks rather than large-scale performance metrics that can hide compensation of errors; 2) relying on observable variables; 3) being stable with respect to internal variability on different timescales, i.e. insensitive to the time period selected; 4) being interpretable physically; 5) being distributed in publicly released tools to ensure reproducibility. The HighResMIP simulations will focus on present-day and near-term future climate. Hence, process-based assessment will not focus on climate sensitivity but on the representation of mean state, variability and teleconnections on a wide range of timescales. The primary goal will be to determine which processes can be represented reliably at typical CMIP5 resolutions and what is the minimum resolution required for an adequate representation of other processes as well as what are the limitations of representing such processes in lower resolution models. A wide variety of processes will be considered from global to regional scales and from the upper atmosphere down to mesoscale eddies.

Through the HighResMIP_BSC PRACE-funded project, BSC plans to take part in HighResMIP using EC-Earth3, which corresponds to the latest available model version of the community model EC-Earth. A spectral truncation of the atmospheric model (IFS) at T511 (approx. 40 km globally) and 91 vertical levels will be used as well as grid resolution of the ocean model (NEMO3.6) of 0.25° globally (approximately 25 km) with 75 vertical levels which thickness increases from 1m below surface up to 500m in the deep ocean. Compared to the resolution of the model used in CMIP5 simulations, the horizontal resolution is increased by a factor of up to 4, with an increased number of vertical levels in both ocean and atmosphere by factors of 1.5 to 2. Accounting for the reduced time step necessary for stability, this corresponds to an increase in computing resources by a factor of about 30.

2. Describe how you will manage the resources requested? Use a <u>Gantt chart or equivalent</u> to illustrate this (mandatory). (1 page).

BSC will be involved in the Tier 2 of the HighResMIP project. Its experimental protocol consists in running a 50-year spinup under 1950 conditions followed by: 1) a historical simulations covering the 1950-2050 period, 2) a control simulation under 1950 conditions run for 100 years. This experimental protocol has been chosen as a compromise between limiting the computational cost to ensure that a maximum number of participating institutes can afford it and allowing for a minimization (thanks to the spinup) and subtraction (thanks to the control) of the model drift. The spinup will be initialised from the EN4 ocean climatology. BSC aims at running 3 members by introducing random perturbations at the end of the spinup. This makes 50 years of spinup and 200 years of simulations for each member, i.e. a total of 650 years of simulation. These experiments will follow the schedule indicated in Table 1. These simulations will be run with EC-Earth 3 in configuration ORCA025L75 – T511L91 (described in section 1). The low resolution counterpart will be performed using national resources.

	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12
Spinup												
Member 1												
Member 2												
Member 3												

Table 1: Approximate schedule of the experiments to be performed. M5 stands for Month 5.

Table 2 summarises the resources requested and detailed in section 6.

	Simulations (years)	Total Core-hours (millions)	Total archive (TB)
Spinup	50	0.9	50
Member 1	200	3.5	200
Member 2	200	3.5	200
Member 3	200	3.5	200
Total	650	11.4	650

Table 2: Resources requested. The estimates have been made on the basis of a cost of 0.9 wallclock hours (by using 1,616 cores) and 1TB of output per year of simulation with the T511L91-ORCA025L75 configuration. These estimates have been obtained by running EC-Earth3.1 on MareNostrum III.

The Autosubmit software (Asif et al., 2014), briefly described below, will be used to manage the workflow and ensure a uniform and optimal use of the resources. Our simulations offer flexibility as they comprise 3 independent 200-year long simulations which can run in parallel for an optimal use of the PRACE computing resources. They will be managed, and packed in groups in a single big job if required, by Autosubmit to optimize the use of the machine and avoid collapsing the I/O system. The data storage and data transfer can be organized with a disk space of 10 TB in the "scratch" file system. This required scratch space is motivated by the large amount of output to be generated. These output data will be transferred immediately locally. Around 500 GB of "home" space will be required to host the code and its modified versions.

3. **Describe the numerical methods and algorithms that you are planning to use, improve, or develop, the codes, packages or libraries that you need to undertake the project, and how these will enable the research to be achieved. (1 page)**

EC-Earth3 comprises three major components: IFS, NEMO and OASIS3. It is essential to configure and build a separate executable for each one of them. The resolution proposed (T511: ~425,000 grid points, ORCA025: ~1,475,000 grid points) will help efficiently share calculations between 1000-1500 sub-domains, increasing the range of efficient compute-core usage per model executable. 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 and NEMO in NetCDF, while OASIS3 does not generate any output. At the end of a simulation the three components always generate restarts 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 newly designed 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.

4. **Explain why this project needs to run on a Tier-0 system, why the machine you have requested is suitable for the project and how the use of the system will enable the science proposed. You should describe the architecture, machine/system name and the problem sizes that have been used to test for scaling and provide supporting evidence.**

   <u>Provide both a table and scaling plot</u> (mandatory) such as the ones shown below with example data to illustrate the information requested (2 pages).

Members of the BSC Earth Science Deparment recently tested EC-Earth3 managed by the Autosubmit system (Asif et al., 2014) on the Cray XE Jaguar in collaboration with the IN-CITE program of the US Department of Energy, achieving an efficient use of ~60,000 cores with a single big job wrapping an ensemble of high-resolution integrations.

Total number of EC-Earth3 cores	Elapsed time (in seconds)			
	NEMO (384)	NEMO (576)	NEMO (768)	NEMO (960)
583	2,042.67			
775	1,387.67	2,042.67		
967	1,417.00	1,203.67	2,041.67	

1159	1,432.00	1,118.67	1,208.00	2,126.00
1351	1,434.33	1,104.67	1,095.67	1,234.67
1543	1,405.67	1,122.00	1,099.33	1,114.33
1735	1,434.67	1,091.67	1,119.33	1,142.00
1927	1,421.33	1,097.67	1,087.33	1,126.00
2119	1,432.33	1,086.33	1,120.67	1,130.00
2311	1,480.67	1,170.00	1,076.33	1,127.67
2503		1,131.00	1,063.33	1,184.00
2695			1,109.67	1,169.67
2887				1,296.33

Table 3: EC-Earth3 (T511-ORCA025) ran for 10 simulation days to obtain the elapsed time (average of 3 sets). The total number of cores used for EC-Earth3 is indicated in the first column while the number of cores used for NEMO (384, 576, 768 or 960) is indicated in columns 2 to 5 and OASIS always use 7 (fixed). The number of cores used for IFS is the difference between that used for EC-Earth and the sum of those used for NEMO and OASIS.

The most recent version of EC-Earth and its highest resolution configuration currently available have been chosen for HighResMIP_BSC. Running EC-Earth with such configuration demands a system with high-level resources, and building a large multi-model ensemble of these characteristics reduces the candidates to a few machines in the world. In this sense, EC-Earth has already been deployed at MareNostrum III (a tier-0 system; having 3056 nodes with 2x Intel SandyBridge-EP E5-2670/1600 8-core processors at 2.6 GHz, having 8x4GB DDR3-1600 DIMMS (2GB/core) per node and Infiniband FDR10 network) and used within the previous PRACE projects "HiResClim" and "HiResClim2". Members of the BSC Earth Sciences Department therefore have experience on running EC-Earth3 on MareNostrum III, and have invested a lot of efforts and resources in finding the best configuration for this model in this particular environment.

During these previous projects, along with scientific experimentation, an extensive scaling exercise was carried out (Asif et al., 2014) on EC-Earth3, where the model was analyzed, not only in coupled mode, but also its individual components in uncoupled/stand-alone mode. To illustrate the problem size, Table 3 and Figure 1 depict the elapsed time obtained by running the model in coupled mode. Figure 1 also shows the respective speedup and efficiency plots.

During this EC-Earth scaling exercise, the oceanic component of EC-Earth3 was identified as a clear bottleneck which limits the scalability of the coupled model. To enhance the speed of this oceanic component, the north-fold optimization technique was tested and adopted for EC-Earth3 (according to suggestions from NEMO experts). A 30% gain in performance in NEMO stand-alone mode and a 20% gain in EC-Earth coupled mode were obtained. Thanks to these improvements, the model could complete one month of simulation in 0.9 hours with IFS+NEMO+OASIS: 640+960+7=1607~1616 cores (according to MareNostrum III support, for a better optimization, the total number of cores should be divisible by the number of cores per node so instead of 1607, a total number of cores of 1616 was chosen).

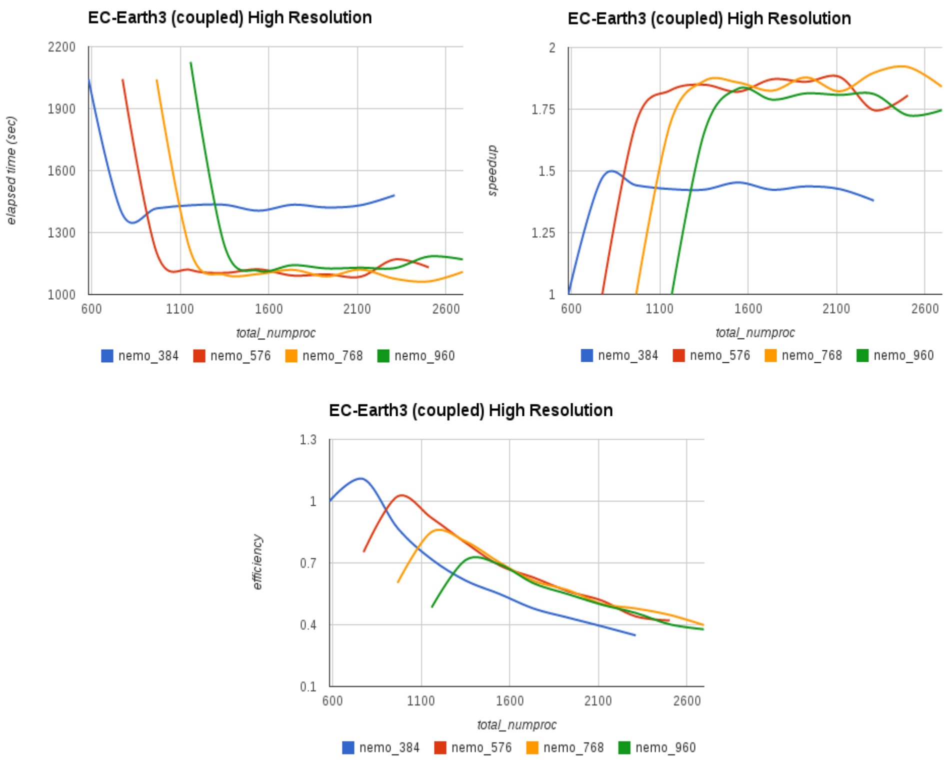

Figure 1: EC-Earth3 (T511-ORCA025-LIM3) ran for 10 simulation days to estimate the elapsed time (average of 3 sets). The blue, red, orange and green curves correspond to results for tests with NEMO using 384, 576, 768 and 960 processors while the speedup and efficiency values are relative to the reference test with 583 processors for EC-Earth.

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

5. Describe your experience of using HPC resources in the past and how you will manage using a Tier-0 system. What other experience do you and your team bring to this project? (1 page).

The BSC Earth Sciences Department has been involved in the IS-ENES/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 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. Finally, members of the BSC Earth Sciences Department coordinated the recent HiResClim and HiResClim2 projects supported by PRACE

The BSC Earth Sciences Department has the privilege to have an expert team on parallel model's performance, which is specialized in analyzing 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 and NEMO. Currently the BSC Earth Sciences Department is collaborator of the NEMO development team, and member of the NEMO HPC working group, providing the Consortium with performance reports and code optimizations. As a result of this collaboration two optimization branches were created in the NEMO code repository by members of the BSC Earth Sciences Department; the first became part of the NEMO 3.6 stable version, and the latter is scheduled to be merged shortly. Following the success of this cooperation, a similar collaboration has been established with the OpenIFS developers at ECMWF.

Figures 2 & 3 illustrate an example of the performance tool's output. These are two different views provided by the Paraver tool from one single EC-Earth model execution, focused on its different components. Both images display the communications pattern along time, representing time on the horizontal axis, while the vertical correspond to the different processes executing the model. 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. Figure 2 focuses on the IFS component model, which has reserved most of the processors used for the execution. 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 message, which can impact the scalability of the code. Figure 3 focuses on the NEMO component model, which in this configuration takes the last 32 processors. The image clearly shows the importance of the communications in NEMO, which for a large percentage of time is not performing effective computations. It also shows that a big portion of this time is devoted to global communications, which are painted in pink colour. We know that 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 our optimizations, so this problem was solved to a large extent.

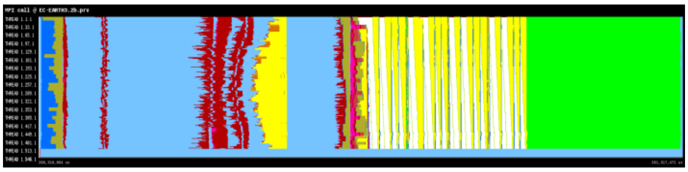

Figure 2: MPI call Paraver view of the IFS component model in an EC-Earth's model execution. The horizontal lines give the behaviour of the different processes (1 to 512) along time. Each colour corresponds to one different MPI communication function.

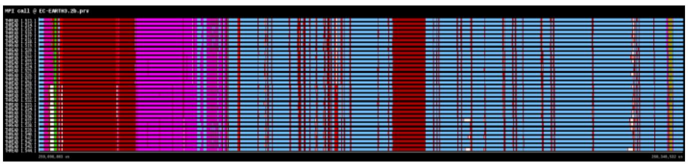

Figure 3: MPI call Paraver view of the NEMO component model in an EC-Earth's model execution. The horizontal lines give the behaviour of the different processes (513 to 544) along time. Each colour corresponds to one different MPI communication function.

6. **Justify the number of core hours requested. This should include information such as: run type, wall clock time per step, number of jobs per run type, the number of CPU cores and the total core hours per run type. This information should take the form of a table like the one shown below with example data (mandatory). Explain how the**

core hours requested will be used (1 page). In case of Multi-year access, please provide an estimation of the core hours requested for the second and third years in the table.

Table 4 lists the experiments described in section 2, which comprise between 150 and 600 chunks of 4 month length, to reach 50 or 200 years of simulations respectively. The three 200-year long simulations can be run independently and simultaneously depending on the machine load by Autosubmit. These experiments use the HR EC-Earth3 configuration. Its cost, for a four-month simulation requires a wall-clock time of 3.6 hours. A benchmarking performed within the framework of the HiResClim project suggests that, taking into account the average load of the MareNostrum III queues, optimum performance is obtained using 1,616 procs.

Run type	# Runs	# Steps/Run	Walltime/Step	# CPU cores / Step	Total core hours
Spinup	1	150	3.6	1,616	872,640
Member 1	1	600	3.6	1,616	3,490,560
Member 2	1	600	3.6	1,616	3,490,560
Member 3	1	600	3.6	1,616	3,490,560
Total					11,344,320

Table 4: Cost of the experiments proposed.

The final estimate is for a total request of 12 million core-hours, which includes the numbers described in Table 4 plus a small buffer of 5% to account for failing jobs that will need to be repeated.

The experiments will be run using Autosubmit, the launching and monitoring solution developed by the group of the applicant that allows the remote submission of EC-Earth and NEMO experiments. Autosubmit includes in the workflow of the experiments a job that retrieves the data back to the Department data storage as soon as each chunk of simulation has completed. This means that the estimates for the archive are an absolute upper value in case the automatic download does not perform as expected. It is very likely that this number will approach a figure ten times smaller.

7. References

Asif, M., A. Cencerrado, O. Mula-Valls, D. Manubens, F.J. Doblas-Reyes and A. Cortés (2014). Impact of I/O and data management in ensemble large scale climate forecasting using EC-Earth3. Procedia Computer Science, 29, 2370-2379, doi:10.1016/j.procs.2014.05.221.

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

Bengtsson, L., K. I. Hodges, M. Esch, N. Keenlyside, L. Kornblueh, J.-J. Luo, and T. Yamagata, 2007: How may tropical cyclones change in a warmer climate? Tellus, 59A, 539-561.

Chassignet, E.P., and D. P. Marshall, 2008: Gulf stream separation in numerical ocean models. In: Hecht, M., Hasumi, H. (eds.), Eddy-Resolving Ocean Modeling, AGU Monog. Ser., 39-62.

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

Gent, P. R., S. G. Yeager, R. B. Neale, S. Levis, and D. A. Bailey, 2010: Improvements in a half degree atmosphere/land version of the CCSM. Clim. Dyn., 34, 819–833, doi:10.1007/s00382-009-0614-8.

Haarsma, R.J., W. Hazeleger, C. Severijns, H. de Vries, A. Sterl, R. Bintanja, G.J. van Oldenborgh and H.W. van den Brink, 2013: More hurricanes to hit Western Europe due to global warming. Geophys. Res. Lett. doi:10.1002/grl.50360.

Hodges, K. I., et al., 2011: A comparison of extratropical cyclones in recent re-analyses ERA-Interim, NASA MERRA, NCEP CFSR, JRA-25. J. Clim., 24, 4888-4906.

Jung, T., et al., 2012: High-resolution global climate simulation with the ECMWF model in project Athena: Experimental design, model climate, and seasonal forecast skill. J. Clim., 25, 3155-3172.

Kuwano-Yoshida A, Minobe S, Xie S-P, 2010: Precipitation response to the Gulf Stream in an Atmospheric GCM. J. Clim 23, 3676–3698.

Lal, M., U. Cubasch, J.P. Perlwitz, and J. Waszkewitz, 1997: Simulation of the Indian monsoon climatology in ECHAM3 climate model: Sensitivity to horizontal resolution. Int. J. Climatol., 17, 847-858

Martin, G. M., 1999: The simulation of the Asian summer monsoon, and its sensitivity to horizontal resolution, in the UK meteorological office unified model. Q.J.R. Meteorol. Soc., 125: 1499–1525. doi:10.1002/qj.49712555703

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 Dynamics, 39, 681-707

Roberts, M. J., et al., 2009: Impact of Resolution on the Tropical Pacific Circulation in a Matrix of Coupled Models. J. Clim., 22, 2541-2556.

Zhao, M., et al., 2009: Simulations of Global Hurricane Climatology, Interannual Variability, and Response to Global Warming Using a 50km Resolution GCM. J. Climate, 33, 6653-6678.

Shaffrey, L. C., et al., 2009: U.K. HiGEM: the new U.K. High-Resolution Global Environment Model-model description and basic evaluation. J. Clim., 22,1861-1896.

Sperber, K. R., H. Sultan, G. L. Potter, and J. S. Boyle, 1994: Simulation of the Northern Summer Monsoon in the ECMWF Model: Sensitivity to Horizontal Resolution. Mon. Wea. Rev., 122, 2461–2481. doi:10.1175/1520-0493(1994)122<2461:SOTNSM>2.0.CO;2

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. Lett., 10, 055002. DOI: 10.1088/1748-9326/10/5/055002.

Zhang, Y., Maslowski W., Semtner, A.J, 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.