



The COSMO-CLM 4.8 regional climate model coupled to regional ocean, land surface and global earth system models using OASIS3-MCT: description and performance

Andreas Will¹, Naveed Akhtar², Jennifer Brauch³, Marcus Breil⁴, Edouard Davin⁵, Ha T. M. Ho-Hagemann⁶, Eric Maisonnave⁷, Markus Thürkow⁸, and Stefan Weiher¹

¹Institute for Environmental Sciences, BTU Cottbus-Senftenberg, Germany

²Institute for Atmospheric and Environmental Sciences, Goethe University, Frankfurt am Main, Germany

³Deutscher Wetterdienst, Offenbach am Main, Germany

⁴Institute of Meteorology and Climate Research, KIT, Karlsruhe, Germany

⁵Institute for Atmospheric and Climate Science, ETH, Zurich, Switzerland

⁶Institute of Coastal Research, Helmholtz-Center Geesthacht, Germany

⁷Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, CERFACS-CNRS, Toulouse, France

⁸Institute of Meteorology, FU Berlin, Germany

Correspondence to: Andreas Will (will@b-tu.de)

Received: 28 February 2016 – Discussion started: 20 April 2016

Revised: 20 January 2017 – Accepted: 4 February 2017 – Published: 13 April 2017

Abstract. We developed a coupled regional climate system model based on the CCLM regional climate model. Within this model system, using OASIS3-MCT as a coupler, CCLM can be coupled to two land surface models (the Community Land Model (CLM) and VEG3D), the NEMO-MED12 regional ocean model for the Mediterranean Sea, two ocean models for the North and Baltic seas (NEMO-NORDIC and TRIMNP+CICE) and the MPI-ESM Earth system model.

We first present the different model components and the unified OASIS3-MCT interface which handles all couplings in a consistent way, minimising the model source code modifications and defining the physical and numerical aspects of the couplings. We also address specific coupling issues like the handling of different domains, multiple usage of the MCT library and exchange of 3-D fields.

We analyse and compare the computational performance of the different couplings based on real-case simulations over Europe. The usage of the LUCIA tool implemented in OASIS3-MCT enables the quantification of the contributions of the coupled components to the overall coupling cost. These individual contributions are (1) cost of the model(s) coupled, (2) direct cost of coupling including horizontal interpolation and communication between the components,

(3) load imbalance, (4) cost of different usage of processors by CCLM in coupled and stand-alone mode and (5) residual cost including i.a. CCLM additional computations.

Finally a procedure for finding an optimum processor configuration for each of the couplings was developed considering the time to solution, computing cost and parallel efficiency of the simulation. The optimum configurations are presented for sequential, concurrent and mixed (sequential+concurrent) coupling layouts. The procedure applied can be regarded as independent of the specific coupling layout and coupling details.

We found that the direct cost of coupling, i.e. communications and horizontal interpolation, in OASIS3-MCT remains below 7 % of the CCLM stand-alone cost for all couplings investigated. This is in particular true for the exchange of 450 2-D fields between CCLM and MPI-ESM. We identified remaining limitations in the coupling strategies and discuss possible future improvements of the computational efficiency.

1 Introduction

The aim of regional climate models is to represent the meso-scale dynamics within a limited area by using appropriate physical parameters describing the region and solving a system of equations derived from first principles of physics describing the dynamics. Most of the current regional climate models (RCMs) are atmosphere–land models and are computationally demanding. They aim to represent the meso-scale dynamics within the atmosphere and between the atmosphere and the land surface and to suppress parts of the interactivity between the atmosphere and the other components of the climate system. The interactivity is either altered by the use of a simplified component model (e.g. over land) or even suppressed when top, lateral and/or ocean surface boundary conditions of the atmospheric component model of the RCM are prescribed by reanalysis or large-scale Earth system model (ESM) outputs.

The neglected meso-scale feedbacks and inconsistencies of the boundary conditions (Laprise et al., 2008; Becker et al., 2015) might be well accountable for a substantial part of large- and regional-scale biases found in RCM simulations at 10–50 km horizontal resolution (see e.g. Kotlarski et al., 2014 for Europe). This hypothesis gains further evidence from the results of convection-permitting simulations, in which these processes are not regarded either. These simulations provide more regional-scale information and improve e.g. the precipitation distribution in mountainous regions, but they usually do not show a reduction of the large-scale biases (see e.g. Prein et al., 2013).

The potential of explicit simulation of the processes neglected or prescribed in land–atmosphere RCMs has been investigated using ESMs with variable horizontal resolution (Hertwig et al., 2015; Hagos et al., 2013) and RCMs two-way coupled with global ESMs (Lorenz and Jacob, 2005; Inatsu and Kimoto, 2009), with regional oceans (Döscher et al., 2002; Gualdi et al., 2013; Zou and Zhou, 2013; Bülow et al., 2014; Akhtar et al., 2014; Pham et al., 2014; Ho-Hagemann et al., 2013, 2015) and/or with more sophisticated land surface models (Wilhelm et al., 2014; Davin et al., 2011).

A significant increase in the climate change signal was found by Somot et al. (2008) in the ARPEGE model with the horizontal grid refined over Europe and two-way coupled with a regional ocean for the Mediterranean Sea. This suggests that building regional climate system models (RC-SMs) with explicit modelling of the interaction between meso scales in the atmosphere, ocean and land surface (by ocean–atmosphere and atmosphere–land couplings) and between meso scales and large scales in the atmosphere (and ocean) (by coupling of regional with global models) might be relevant for an improved representation of regional climate and climate change. Furthermore, the large-scale dynamics can be significantly improved by two-way coupling with meso scales if upscaling is a relevant process.

However, a decision to use the growing computational resources for an explicit simulation of interactions suppressed otherwise does not depend only on its physical impact on the simulation quality, but also on the extra cost in comparison with e.g. a further increase in the model's grid resolution.

In this paper we present a prototype of a RCM, a concept of finding an optimum configuration of computational resources, and discuss the extra cost of coupling in comparison with an RCM solution. The RCM prototype is based on the COSMO-CLM (CCLM) non-hydrostatic regional climate model (Rockel et al., 2008), which belongs to the class of land–atmosphere RCMs. We present couplings of CCLM with one other model applied successfully over Europe on climatological timescales.

The coupling of CCLM with a land surface scheme replaces the TERRA land surface scheme of CCLM. One scheme coupled is the VEG3D soil and vegetation model. It is extensively tested in central Europe and western Africa on regional scales and has, in comparison with TERRA, an implemented vegetation layer. The other scheme coupled is the Community Land Model (CLM) (version 4.0). It is a state-of-the-art land surface scheme developed for all climate zones and global applications.

The couplings with the regional ocean models replace the prescribed SSTs over regional ocean surfaces and allow for meso-scale interaction. High-resolution configurations for the regional oceans in the European domain are available for the NEMO community ocean model. We use the configurations for the Mediterranean (with NEMO version 3.2) and for the Baltic and North seas (with NEMO version 3.3, including the LIM3 sea ice model). A second high-resolution configuration for the Baltic and North seas is available for the TRIMNP regional ocean model along with the CICE sea ice model.

The coupling with the Earth system model replaces the atmospheric lateral and top boundary condition and the lower boundary condition over the oceans (SST) and allows for a common solution between the RCM and ESM at the RCM boundaries, thus reducing the boundary effect of one-way RCM solutions. Furthermore, it extends the opportunities of multi-scale modelling. We couple the state-of-the-art MPI-ESM Earth system model (version 6.1), which is widely used in regional climate applications of CCLM in one-way mode.

Additional models, which can be coupled with CCLM in the same way but which are not discussed in this article, are the ROMS ocean model (Byrne et al., 2015) and the ParFLOW hydrological model (Gasper et al., 2014) together with CLM.

Each coupling is using the OASIS3-MCT (Valcke et al., 2013) coupler, a fully parallelised version of the widely used OASIS3 coupler and a unified OASIS3 interface in CCLM. The solutions found for particular problems of coupling of a regional climate model using features of OASIS3-MCT will be presented in this paper as well.

An alternative coupling strategy is available for CCLM. It is based on an internal coupling of the models of interest with the master routine MESSy resulting in the compilation of one executable (Kerkweg and Joeckel, 2012). This coupling strategy is not investigated in this study.

The climate system models, either global (ESMs) or regional (RCSMs), are computationally demanding. Keeping the computing cost small contributes substantially to the climate system models' usability. For this reason the present paper also focuses on the coupled systems' computational efficiency, which greatly relies on the parallelisation of the OASIS3-MCT coupler.

An optimisation of the computational performance is considered to be highly dependent on the model system and/or the computational machine used. However, several studies show transferability of optimisation strategies and universality of certain aspects of the performance. Worley et al. (2011) analysed the performance of the Community Earth System Model (CESM) and found a good scalability of the concurrently running CLM and sequentially running CICE down to approximately 100 grid points per processor for two different resolutions and computing architectures. Furthermore, they found the CICE scalability to be limited by a domain decomposition, which follows that of the ocean model, resulting in a very low number of ice grid points in subdomains. Lin-Jiong et al. (2012) investigated a weak scaling (discussed in Sect. 4.3) of the FAMIL model (IAP, Beijing) and found a performance similar to that of the optimised configuration of the CESM (Worley et al., 2011). This result indicates that a careful investigation of the model performance leads to similar results for similar computational problems. An analysis of the CESM at very high resolutions by Dennis et al. (2012) showed that a cost reduction by a factor of 3 or so can be achieved using an optimal layout of model components. Later Alexeev et al. (2014) presented an algorithm for finding an optimum model coupling layout (concurrent, sequential) and processor distribution between the model components minimising the load imbalance in the CESM.

These results indicate that the optimised computational performance is weakly dependent on the computing architecture or on the individual model components but depends on the coupling method. Furthermore, the application of an optimisation procedure was found to be beneficial.

In this study we present a detailed analysis of the performances of CCLM+X (X: another model) coupled model systems on IBM POWER6 machine *Blizzard* located at DKRZ, Hamburg, for a real climate simulation configuration over Europe. We calculate the speed and cost of the individual models in coupled mode and of the coupler itself. We identify the reasons for reduced speed or increased cost for each coupling and reasonable processor configurations and suggest an optimum processor configuration for each coupling considering the cost and speed of the simulation. Particularities of the performance of a coupled RCM are highlighted together with the potential of the OASIS3-MCT coupling soft-

ware. We suggest a procedure of optimisation of an RCSM processor configuration, which can be generalised. However, we show that some relevant optimisations are possible only due to features available with the OASIS3-MCT coupler.

Finally we present an analysis of the extra cost of coupling at optimum configuration. We separate the cost of (i) components of the model system coupled, (ii) the OASIS3-MCT coupler including horizontal interpolation and communication between the components, (iii) load imbalance, (iv) different usage of processors by CCLM in coupled and stand-alone mode and (v) residual cost including additional computations in CCLM. This allows one to identify the unavoidable cost of coupling and the bottlenecks.

The paper is organised as follows. The models coupled are described in Sect. 2. Section 3 focuses on the OASIS3-MCT coupling method and its interfaces for the individual couplings. The coupling method description encompasses the OASIS3-MCT functionality, method of the coupling optimisation and particularities of coupling of a regional climate model system. The model interface description gives a summary of the physics and numerics of the individual couplings. In Sect. 4 the computational efficiency of individual couplings is presented and discussed. Finally, the conclusions and an outlook are given in Sect. 5. For improved readability, Tables 1 and 2 provide an overview of the acronyms frequently used throughout the paper and of the investigated couplings.

2 Description of regional climate model system components

The further development of the COSMO model in Climate Mode (COSMO-CLM or CCLM) presented here aims at overcoming the limitations of the regional soil–atmosphere climate model, as discussed in the introduction, by replacing prescribed vegetation, lower boundary condition over sea surfaces and the lateral and top boundary conditions with interactions between dynamical models.

The models selected for coupling with CCLM need to fulfil the requirements of the intended range of application, which are (1) the simulation at varying scales from convection-resolving up to 50 km grid spacing, (2) local-scale up to continental-scale simulation domains and (3) full capability at least for European model domains. We decided to couple the NEMO ocean model for the Mediterranean Sea (NEMO-MED12) and the Baltic and North seas (NEMO-NORDIC), alternatively the TRIMNP regional ocean model together with the CICE sea ice model for the Baltic and North seas (TRIMNP+CICE), the Community Land Model (CLM) of soil and vegetation (replacing the TERRA multi-layer soil model), or alternatively the VEG3D soil and vegetation model and the MPI-ESM global earth system model for two-way coupling with the regional atmosphere. Table 2 gives an overview of all model systems investigated, their components

Table 1. List of abbreviations used throughout the paper.

Acronym	Meaning
COSMO	Limited-area model of the COnsortium for Small-scale MOdeling
COSMO-CLM	COSMO model in CLimate Mode
CCLM	Abbreviation of COSMO-CLM
CCLM _{OC}	CCLM in coupled mode using the mapping of optimum processor configuration
CCLM _{sa}	CCLM stand-alone, not in coupled mode
CCLM _{sa,sc}	CCLM _{sa} using the same mapping as in coupled mode
CCLM _{sa,OC}	CCLM _{sa} using the mapping of optimum processor configuration
CLM	Community Land Model
VEG3D	Soil and vegetation model of KIT
NEMO	Community model “Nucleus for European Modeling of the Ocean”
NEMO-MED12	NEMO 3.2 for the Mediterranean sea
NEMO-NORDIC	NEMO 3.3 for the North and Baltic seas
TRIMNP	Tidal, Residual, Intertidal mudflat Model Nested parallel Processing regional ocean model
CICE	Sea ice model of LANL
MPI-ESM	Global Earth System Model of MPIfM Hamburg
ECHAM	Atmosphere model (ECMWF dynamics and MPIfM Hamburg physics) of MPI-ESM
MPIOM	MPIfM Hamburg Ocean Model of MPI-ESM
OASIS3-MCT	Coupling software for Earth System Models of CERFACS
CESM	Community Earth System Model
Institutions	
MPIfM	Max-Planck-Institut für Meteorologie Hamburg, Germany
LANL	Los Alamos National Laboratory, USA
CERFACS	Centre Europeen de Recherche et de Formation Avancee en Calcul Scientifique, Toulouse, France
CLM-Community	Climate Limited-area Modeling (CLM-)Community
ECMWF	European Center for Medium Range Weather Forecast, Reading, Great Britain
NCAR	National Center for Atmospheric Research, Boulder, USA
CNRS	Centre National de Recherche Scientifique, Paris, France
ETH	Eidgenössische Technische Hochschule, Zürich, Switzerland
KIT	Karlsruher Institut für Technologie, Germany
GUF	Goethe-Universität Frankfurt am Main, Germany
HZG	Helmholtz-Zentrum Geesthacht, Germany
BTU	Brandenburgische Technische Universität Cottbus-Senftenberg, Cottbus, Germany
FUB	Freie Universität Berlin, Germany
Model domains	
CORDEX-EU	CORDEX domain for regional climate simulations over Europe

Table 2. Coupled model systems, their components and the institution at which they are maintained. For the meaning of the acronyms see Table 1.

Coupled model system	Institution	First coupled component	Second coupled component
CCLM+CLM	ETH	CLM	–
CCLM+VEG3D	KIT	VEG3D	–
CCLM+NEMO-MED12	GUF	NEMO-MED12	–
CCLM+TRIMNP+CICE	HZG	TRIMNP	CICE
CCLM+MPI-ESM	BTU and FUB	ECHAM	MPIOM

and institutions at which they are maintained. An overview of the models selected for coupling with CCLM is given in Table 3 together with the main model developer, configuration details of high relevance for computational performance, the

model complexity (see Balaji et al., 2017) and a reference in which a detailed model description can be found. The model domains are plotted in Fig. 1. More information on the avail-




Figure 1. Map of coupled-system components. The horizontal domains of all components are bounded by the CCLM domain (CORDEX-EU), except MPI-ESM (= ECHAM+MPI-OM), which is solved on the global domain. The CLM and VEG3D domains cover CCLM (land). TRIMNP, CICE and NEMO-NORDIC share area 1. Additionally, CICE covers area 4, NEMO-NORDIC area 3 and TRIMNP areas 2, 3 and 4.

ability of the CCLM coupled model systems can be found in Appendix A.

In the following, the models used are briefly described with respect to model history, space–time scales of applicability and model physics and dynamics relevant for the coupling.

2.1 COSMO-CLM

COSMO-CLM (CCLM) is the COSMO model in climate mode. COSMO model is a non-hydrostatic limited-area atmosphere–soil model originally developed by the Deutscher Wetterdienst for operational numerical weather prediction (NWP). Additionally, it is used for climate, environmental (Vogel et al., 2009) and idealised studies (Baldauf et al., 2011).

The COSMO physics and dynamics are designed for operational applications at horizontal resolutions of 1 to 50 km for NWP and RCM applications. The basis of this capability is a stable and efficient solution of the non-hydrostatic system of equations for the moist, deep atmosphere on a spherical, rotated, terrain-following, staggered Arakawa C grid with a hybrid z level coordinate. The model physics and dynamics are described in Doms et al. (2011) and Doms and Baldauf (2015) respectively. The features of the model are discussed in Baldauf et al. (2011).

The COSMO model’s climate mode (Rockel et al., 2008) is a technical extension for long-time simulations and all related developments are unified with COSMO regularly. The important aspects of the climate mode are time dependency of the vegetation parameters and of the prescribed SSTs and usability of the output of several global and regional climate models as initial and boundary conditions. All other aspects

related to the climate mode, e.g. the restart option for soil and atmosphere, the NetCDF model input and output, online computation of climate quantities, and the sea ice module or spectral nudging, can be used in other modes of the COSMO model as well.

The `cosmo_4.8_clm19` model version is the recommended version of the CLM-Community (Kotlarski et al., 2014) and it is used for the couplings, but for CCLM+CLM and for stand-alone simulations. CCLM as part of the CCLM+CLM coupled system is used in a slightly different version (`cosmo_5.0_clm1`). The way this affects the performance results is presented in Sect. 4.4.

2.2 MPI-ESM

The global Earth System Model of the Max Planck Institute for Meteorology Hamburg (MPI-ESM; Stevens et al., 2013) consists of subsystem models for ocean and atmosphere, cryo-, pedo- and bio-sphere. The ECHAM6 hydrostatic general circulation model uses the transform method for horizontal computations. The derivatives are computed in spectral space, and the transports and physics tendencies on a regular grid in physical space. A pressure-based sigma coordinate is used for vertical discretisation. The MPIOM ocean model (Jungclaus et al., 2013) is a regular grid model with the option of local grid refinement. The terrestrial bio- and pedosphere component model is JSBACH (Reick et al., 2013; Schneck et al., 2013). The marine biogeochemistry model used is HAMOCC5 (Ilyina et al., 2013). A key aspect is the implementation of the bio-geo-chemistry of the carbon cycle, which allows e.g. investigation of the dynamics of the greenhouse gas concentrations (Giorgetta et al., 2013). The subsystem models are coupled via the OASIS3-MCT coupler (Valcke et al., 2013) which was implemented recently by I. Fast of DKRZ in the CMIP5 model version. This allows parallelised and efficient coupling of a huge amount of data, which is a requirement of atmosphere–atmosphere coupling.

The MPI-ESM reference configuration uses a spectral resolution of T63, which is equivalent to a spatial resolution of about 320 km for atmospheric dynamics and 200 km for model physics. Vertically the atmosphere is resolved by 47 hybrid sigma-pressure levels, with the top level at 0.01 hPa. The MPIOM reference configuration uses the GR15L40 resolution which corresponds to a bipolar grid with a horizontal resolution of approximately 165 km near the Equator and 40 vertical levels, most of them within the upper 400 m. The North Pole and the South Pole are located over Greenland and Antarctica in order to avoid the “pole problem” and to achieve a higher resolution in the Atlantic region (Jungclaus et al., 2013).

2.3 NEMO

The Nucleus for European Modeling of the Ocean (NEMO) is based on the primitive equations. It can be adapted for re-

Table 3. Properties of the models coupled. For the meanings of the acronyms, see Table 1. The configuration used is a coarse-grid regional climate simulation configuration used for sensitivity studies, tests and continental-scale climate simulations. Model complexity is measured as the number of prognostic variables. For a comprehensive definition, see Balaji et al. (2017).

Model	CCLM	CLM	VEG3D	MPI-ESM
Full name	COSMO model in mode	Community Land Model	Vegetation model	Max Planck Institute Earth System Model
Institution	CLM-Community	NCAR and other institutions	KIT	MPIfM Hamburg
Coupling area	CORDEX-EU	CORDEX-EU land	CORDEX-EU land	CORDEX-EU
Horizontal res. (km)	50	50	50	330
No. of levels	40/45	15	10	47
Time step (s)	300	3600	300	600
Grid points (10 ³)	766	142	95	3118
Complexity	35	<1	<1	58
Reference	Baldauf et al. (2011)	Oleson et al. (2010)	Schädler (1990)	Stevens et al. (2013)
Model	NEMO-MED12	NEMO-NORDIC	TRIMNP	CICE
Full name	Nucleus for European Modeling of the Ocean – Mediterranean Sea	Nucleus for European Modeling of the Ocean – North and Baltic seas	Tidal, Residual, Intertidal mudflat Model Nested parallel Processing	Sea Ice Model
Institution	CNRS	CNRS	Univ. Trento, HZG	LANL
Coupling area	Mediterranean Sea (without Black Sea)	North and Baltic seas	North and Baltic Sea	Baltic Sea and Kattegat
Horizontal res. (km)	6–8	3.7	12.8	12.8
No. of levels	50	56	50	5
Time step (s)	720	300	240	240
Grid points (10 ³)	2767	4187	877	28
Complexity	8	8	11	<1
Reference	Madec (2008), Lebeaupin et al. (2011); Akhar et al. (2014)	Hordeir et al. (2013), Dieterich et al. (2013), Pham et al. (2014)	Casulli and Cattani (1994), Casulli and Stelling (1998); Hagemann et al. (2013)	Hunke et al. (2013); Ho-Hagemann et al. (2013)

gional and global applications. The sea ice (LIM3) or the marine biogeochemistry module with passive tracers (TOP) can be used optionally. NEMO uses staggered variable positions together with a geographic or Mercator horizontal grid and a terrain-following σ coordinate (curvilinear grid) or a z coordinate with full or partial bathymetry steps (orthogonal grid). A hybrid vertical coordinate (z coordinate near the top and σ coordinate near the bottom boundary) is possible as well (for details see Madec, 2011).

CCLM is coupled to two different regional versions of the NEMO model, adapted to specific conditions of the region of application. For the North and Baltic seas, the sea ice module (LIM3) of NEMO is activated and the model is applied with a free surface to enable the tidal forcing, whereas in the Mediterranean Sea, the ocean model runs with a classical rigid-lid formulation in which the sea surface height is simulated via pressure differences. Both model set-ups are briefly introduced in the following two sub-sections.

2.3.1 Mediterranean Sea

Lebeaupin et al. (2011), Beuvier et al. (2012) and Akhtar et al. (2014) adapted NEMO version 3.2 (Madec, 2008) to the regional ocean conditions of the Mediterranean Sea, hereafter called *NEMO-MED12*. It covers the whole Mediterranean Sea excluding the Black Sea. The NEMO-MED12 grid is a section of the standard irregular ORCA12 grid (Madec, 2008) with an eddy-resolving $1/12^\circ$ horizontal resolution, stretched in the latitudinal direction, equivalent to 6–8 km horizontal resolution. In the vertical, 50 unevenly spaced levels are used with 23 levels in the top layer of 100 m depth. A time step of 12 min is used.

The initial conditions for potential temperature and salinity are taken from the Medatlas (MEDAR-Group, 2002). The freshwater inflow from rivers is prescribed by a climatology taken from the RivDis database (Vörösmarty et al., 1996) with seasonal variations calibrated for each river by Beuvier et al. (2010) based on Ludwig et al. (2009). In this context, the Black Sea is considered as a river for which climatological monthly values are calculated from a dataset of Stanev and Peneva (2002). The water exchange with the Atlantic Ocean is parameterised using a buffer zone west of the Strait of Gibraltar with a thermohaline relaxation to the World Ocean Atlas data of Levitus et al. (2005).

2.3.2 North and Baltic seas

Hordoir et al. (2013), Dieterich et al. (2013) and Pham et al. (2014) adapted the NEMO version 3.3 to the regional ocean conditions of the North and Baltic seas, hereafter called *NEMO-NORDIC*. Part of NEMO 3.3 is the LIM3 sea ice model including a representation of dynamic and thermodynamic processes (for details see Vancoppenolle et al., 2009). The NEMO-NORDIC domain covers the whole Baltic and North Sea area with two open boundaries to the Atlantic

Ocean: the southern, meridional boundary in the English Channel and the northern, zonal boundary between the Hebrides and Norway. The horizontal resolution is 2 nautical miles (about 3.7 km) with 56 stretched vertical levels. The time step used is 5 min. No freshwater flux correction for the ocean surface is applied. NEMO-NORDIC uses a free top surface to include the tidal forcing in the dynamics. Thus, the tidal potential has to be prescribed at the open boundaries in the North Sea. Here, we use the output of the global tidal model of Egbert and Erofeeva (2002).

The lateral freshwater inflow from rivers plays a crucial role for the salinity budget of the North and Baltic seas. It is taken from the daily time series of river runoff from the E-HYPE model output operated at SMHI (Lindström et al., 2010). The World Ocean Atlas data (Levitus et al., 2005) are used for the initial and lateral boundary conditions of potential temperature and salinity.

2.4 TRIMNP and CICE

TRIMNP (Tidal, Residual, Intertidal Mudflat Model Nested Parallel Processing) is the regional ocean model of the University of Trento, Italy (Casulli and Cattani, 1994; Casulli and Stelling, 1998). The domain of TRIMNP covers the Baltic Sea, the North Sea and a part of the north-eastern Atlantic Ocean, with the north-western corner over Iceland and the south-western corner over Spain at the Bay of Biscay. TRIMNP is designed with a horizontal grid mesh size of 12.8 km and 50 vertical layers. The thickness of the top 20 layers is 1 m each and increases with depth up to 600 m for the remaining layers. The model time step is 240 s. Initial states and boundary conditions of water temperature, salinity, and velocity components for the ocean layers are determined using the monthly ORAS-4 reanalysis data of ECMWF (Balmaseda et al., 2013). The daily Advanced Very High Resolution Radiometer AVHRR2 data of the National Oceanic and Atmospheric Administration of the USA are used for surface temperature and the World Ocean Atlas data (Levitus and Boyer, 1994) for surface salinity. No tide is taken into account in the current version of TRIMNP. Monthly river inflows of 33 rivers to the North Sea and the Baltic Sea are rough estimates based on climatological annual mean, minimum and maximum values (H. Kapitza, HZG Geesthacht, Germany, personal communication, 2012).

The CICE sea ice model version 5.0 is developed at the Los Alamos National Laboratory, USA (<http://oceans11.lanl.gov/trac/CICE/wiki>), to represent dynamic and thermodynamic processes of sea ice in global climate models (for more details, see Hunke et al., 2013). In this study CICE is adapted to the region of the Baltic Sea and Kattegat, a part of the North Sea, on a 12.8 km grid with five ice categories. Initial conditions of CICE are determined using the AVHRR2 SST.

2.5 VEG3D

VEG3D is a multi-layer soil–vegetation–atmosphere transfer model (Schädler, 1990) designed for regional climate applications and maintained by the Institute of Meteorology and Climate Research at the Karlsruhe Institute of Technology. VEG3D considers radiation interactions with vegetation and soil, and calculates the turbulent heat fluxes between the soil, the vegetation and the atmosphere, as well as the thermal transport and hydrological processes in soil, snow and canopy.

The radiation interaction and the moisture and turbulent fluxes between soil surface and the atmosphere are regulated by a massless vegetation layer located between the lowest atmospheric level and the soil surface, having its own canopy temperature, specific humidity and energy balance. The multi-layer soil model solves the heat conduction equation for temperature and the Richardson equation for soil water content. Thereby, vertically differing soil types can be considered within one soil column, comprising 10 stretched layers with its bottom at a depth of 15.34 m. The heat conductivity depends on the soil type and the water content. In case of soil freezing the ice phase is taken into account. The soil texture has 17 classes. Three classes are reserved for water, rock and ice. The remaining 14 classes are taken from the USDA Textural Soil Classification (Staff, 1999).

Ten different landuse classes are considered: water, bare soil, urban area and seven vegetation types. Vegetation parameters like the leaf area index or the plant cover follow a prescribed annual cycle.

Up to two additional snow layers on top are created, if the snow cover is higher than 0.01 m. The physical properties of the snow depend on its age, metamorphosis, melting and freezing. A snow layer on a vegetated grid cell changes the vegetation albedo, emissivity and turbulent transfer coefficients for heat as well.

An evaluation of VEG3D in comparison with TERRA in western Africa is presented by Köhler et al. (2012).

2.6 Community Land Model

The Community Land Model (CLM) is a state-of-the-art land surface model designed for climate applications. Biogeophysical processes represented by CLM include radiation interactions with vegetation and soil, the fluxes of momentum, sensible and latent heat from vegetation and soil and the heat transfer in soil and snow. Snow and canopy hydrology, stomatal physiology and photosynthesis are modelled as well.

Subgrid-scale surface heterogeneity is represented using a tile approach allowing five different land units (vegetated, urban, lake, glacier, wetland). The vegetated land unit is itself subdivided into 17 different plant-functional types (or more when the crop module is active). Temperature, energy and water fluxes are determined separately for the canopy

layer and the soil. This allows a more realistic representation of canopy effects than in bulk schemes, which have a single surface temperature and energy balance. The soil column has 15 layers, the deepest layer reaching 42 m in depth. Thermal calculations explicitly account for the effect of soil texture (vertically varying), soil liquid water, soil ice and freezing/melting. CLM includes a prognostic water table depth and groundwater reservoir allowing for a dynamic bottom-boundary condition for hydrological calculations rather than a free drainage condition. A snow model with up to five layers enables the representation of snow accumulation and compaction, melt/freezing cycles in the snowpack and the effect of snow aging on surface albedo.

CLM also includes processes such as carbon and nitrogen dynamics, biogenic emissions, crop dynamics, transient land cover change and ecosystem dynamics. These processes are activated optionally and are not considered in the present study. A full description of the model equations and input datasets is provided in Oleson et al. (2010) (for CLM4.0) and Oleson et al. (2013) (for CLM4.5). An offline evaluation of CLM4.0 surface fluxes and hydrology at the global scale is provided by Lawrence et al. (2011).

CLM is developed as part of the Community Earth System Model (CESM) (Collins et al., 2006; Dickinson et al., 2006) but it has been also coupled to other global (NorESM) or regional (Steiner et al., 2005, 2009; Kumar et al., 2008) climate models. In particular, an earlier version of CLM (CLM3.5) has been coupled to CCLM (Davin et al., 2011; Davin and Seneviratne, 2012) using a “sub-routine” approach for the coupling. Here we use a more recent version of CLM (CLM4.0 as part of the CESM1_2.0 package) coupled to CCLM via OASIS3-MCT rather than through a sub-routine call. A scientific evaluation of this coupled system, also referred to as COSMO-CLM², is provided in Davin et al. (2016). Note that CLM4.5 is also included in CESM1_2.0 and can be also coupled to CCLM using the same framework.

3 Description and optimisation of CCLM couplings via OASIS3-MCT

The computational performance, usability and maintainability of a complex model system depend on the coupling method used, the ability of the coupler to run efficiently in the computing architecture, and on the flexibility of the coupler to deal with different requirements of the coupling depending on model physics and numerics.

In the following, the physics and numerics of the coupling of CCLM with different models (or *components* of the coupled system) via OASIS3-MCT are discussed and the different aspects of optimisation of the computational performance of the individual couplings are highlighted. In Sect. 3.1.1 the main differences between coupling methods are discussed, the main properties of the OASIS3-MCT coupling method are described, the new OASIS3-MCT features

are highlighted and the steps of optimisation of the computational performance of a regional coupled model system are discussed considering different coupling layouts (concurrent/sequential). In Sects. 3.2 to 3.5 the physics and numerics of the couplings are described. In these sections a list of the exchanged variables, the additional computations and the interpolation methods is presented. The time step organisation of each model coupled is given in Appendix B.

3.1 Efficient coupling of a regional climate model

The complexity of the climate system leads to developments of independent models for different components of the climate system. Software solutions are widely used to organise the interaction between the models in order to simulate the development of the climate system. However, the solutions should be accurate, the simulation computationally efficient and the model system easy to maintain. Appropriate software solutions have been developed mainly for global earth system models. As will be shown in the following, the specific features of regional climate system models lead to new requirements which can be met using OASIS3-MCT.

In this section the OASIS3-MCT coupling method is described with a focus on the new features of the Model Coupling Toolkit (MCT) and the solutions found for the particular requirements of regional climate system modelling. Furthermore, a concept for finding of an optimum processor configuration is presented.

3.1.1 Choice of the coupling method

Lateral-, top- and/or bottom-boundary conditions for regional geophysical models are traditionally read from files and updated regularly at runtime. We call this approach *off-line (one-way) coupling*. For various reasons, one could decide to calculate these boundary conditions with another geophysical model – at runtime – in an *online (one-way) coupling*. If this additional model in return receives information from the first model modifying the boundary conditions provided by the first to the second, an *online two-way coupling* is established. In any of these cases, model exchanges must be synchronised. This could be done by (1) reading data from file, (2) calling one model as a subroutine of the other or (3) using a coupler which is software that enables online data exchanges between models.

Communicating information from model to model boundaries via reading from and writing to a file is known to be quite simple to implement but computationally inefficient, particularly in the case of non-parallelised I/O and high frequencies of disc access. In contrast, calling component models as subroutines exhibits much better performances because the information is exchanged directly in memory. Nevertheless, the inclusion of an additional model in a “subroutine style” requires comprehensive modifications of the source code. Furthermore, the modifications need to be updated for

every new source code version. Since the early 90s, software solutions have been developed which allow coupling between geophysical models in a non-intrusive, flexible and computationally efficient way. This facilitates use of the last released model versions in couplings of models developed and maintained by different communities.

One of the software solutions for coupling of geophysical models is the OASIS coupler, which is widely used in the climate modelling community (see for example Valcke, 2013, and Maisonnave et al., 2013). Its latest version, OASIS3-MCT version 2.0 (Valcke et al., 2013), is fully parallelised. Masson et al. (2012) proved its efficiency for high-resolution quasi-global models on top-end supercomputers. A second proof is presented in this paper in Sect. 4.5. This shows that the parallelisation is required for the coupling between a regional climate and global earth system model.

3.1.2 Features of the OASIS3 Model Coupling Toolkit (OASIS3-MCT)

A separate executable (coupler) was necessary to the former version of OASIS. OASIS3-MCT consists of a FORTRAN application programming interface (API). Its subroutines have to be added in all coupled-system component models. The part of the program in which the OASIS3-MCT API routines are located is called the *component interface*. There is no independent OASIS executable anymore, as was the case with OASIS3. With OASIS3-MCT, every communication between the component models is directly executed via the Model Coupling Toolkit (MCT, in Jacob et al., 2005) based on the Message Passing Interface (MPI). This significantly improves the performance over OASIS3, because the bottleneck due to the sequential separate coupler is entirely removed as shown e.g. in Gasper et al. (2014).

In the following, we point out the potential of the new OASIS3-MCT coupler and discuss the peculiarities of its application for coupling in the COSMO model in CLimate Mode (COSMO-CLM or CCLM). If there is no difference between the OASIS versions, we use the acronym OASIS; otherwise, the OASIS version is specified.

In the OASIS coupling paradigm, each model is a *component* of a *coupled system*. Each *component* is included as a separate executable up to OASIS3-MCT version 2.0. Using version 3.0 this is not a constraint anymore. Now a *component* can be an externally coupled component model or an internally coupled model component. This e.g. facilitates the use of the same physics of coupling for internally and externally coupled components, e.g. different land surface schemes.

At runtime, all components are launched together in a single MPI context. The parameters defining the properties of a coupled system are provided to OASIS via an ASCII file called *namcouple*. By means of this file the *component's* coupling fields and coupling intervals are associated. Specific calls of the *OASIS3-MCT Application Programming Inter-*

face (API) in a component interface described in Sects. 3.2 to 3.5 define a component's coupling characteristics, that is, (1) the name of incoming and outgoing coupling fields, (2) the grids on which each of the coupling fields are discretised, (3) a mask (binary-sparse array) describing where coupling fields are described on the grids and (4) the partitioning (MPI-parallel decomposition into subdomains) of the grids. The component partitioning and grid do not have to be the same for each component as OASIS3-MCT is able to scatter and gather the arrays of coupling fields if they are exchanged with a component that is decomposed differently. Similarly, OASIS is able to perform interpolations between different grids. OASIS is also able to perform time averaging or accumulation for exchanges at a coupling time step, e.g. if the components' time steps differ. In total, six to eight API routines have to be called by each component to start MPI communications, declare the component's name, possibly get back the MPI local communicator for internal communications, declare the grid partitioning and variable names, finalise the component's coupling characteristics declaration, send and receive the coupling fields and, finally, close the MPI context at the component's runtime end. The number of routines, whose arguments require easily identifiable model quantities, is the most important feature of the OASIS3-MCT coupling library that contributes to its non-intrusiveness. In addition, each component can be modified separately or another component can be added later. This facilitates a shared maintenance between the users of the coupled-model system: when a new development or a version upgrade is done in one component, the modification scarcely affects the other components. This ensures the modularity and interoperability of any OASIS-coupled system.

As previously mentioned, OASIS3-MCT includes the MCT library, based on MPI, for direct parallel communications between components. To ensure that calculations are delayed only by receiving of coupling fields or interpolation of these fields, MPI non-blocking sending is used by OASIS3-MCT so that sending coupling fields is a quasi-instantaneous operation. The SCRIP library (Jones, 1997) included in OASIS3-MCT provides a set of standard operations (for example bilinear and bicubic interpolation, Gaussian-weighted N-nearest-neighbour averages) to calculate, for each source grid point, an interpolation weight that is used to derive an interpolated value at each (non-masked) target grid point. OASIS3-MCT can also (re-)use interpolation weights calculated offline. Intensively tested for demanding configurations (Craig et al., 2012), the MCT library performs the definition of the parallel communication pattern needed to optimise exchanges of coupling fields between each component's MPI subdomain. It is important to note that unlike the "subroutine coupling" each component coupled via OASIS3-MCT can keep its parallel decomposition so that each of them can be used at its optimum scalability. In some cases, this optimum can be adjusted to ensure a good load balance between components. The two optimisa-

tion aims that strongly matter for computational performance are discussed in the next section.

3.1.3 Synchronisation and optimisation of a regional coupled system

A component receiving information from one or several other component has to wait for the information before it can perform its own calculations. In case of a two-way coupling this component provides information needed by the other coupled-system component(s). As mentioned earlier, the information exchange is quasi-instantaneously performed, if the time needed to perform interpolations can be neglected which is the case even for 3-D-field couplings (as discussed in Sect. 4.6). Therefore, the total duration of a coupled-system simulation can be separated into two parts for each component: (1) a waiting time in which a component waits for boundary conditions and (2) a computing time in which a component's calculations are performed. The duration of a stand-alone, that is, un-coupled component simulation approximates the coupled-component's computing time. In a coupled system this time can be shorter than in the uncoupled mode, since the reading of boundary conditions from file (in stand-alone mode) is partially or entirely replaced by the coupling. It is also important to note that components can perform their calculations sequentially or concurrently.

The coupled-system's total sequential simulation time can be expected to be equal to the sum of the individual component's calculation times, potentially increased by the time needed to interpolate and communicate coupling fields between the components. The computational constraint induced by a sequential coupling algorithm depends on the computing architecture. If one process can be started on each core, the cores allocated for one model system component are idle while others are performing calculations and vice versa. In such a case the performance optimisation strategy needs to consider the component's waiting time. If more than one process can be started on each core, each component can use all cores sequentially and an allocation of the same number of cores to each component can avoid any waiting time. This is discussed in more detail in the following paragraphs.

The constraints of sequential coupling are often alleviated if calculations of a coupled-system component can be performed with coupling fields of another component's previous coupling time step. This concurrent coupling strategy is possible if one of the two sets of exchanged quantities is slowly changing in comparison to the other set. For example, sea surface temperatures of an ocean model are slowly changing in comparison to fluxes coming from an atmosphere model. However, now the time to solution of each component can be substantially different and an optimisation strategy needs to minimise the waiting time.

Thus, the strategy of synchronisation of the components depends on the layout of the coupling (sequential or concurrent) in order to reduce the waiting time as much as possible.

It is important to note that huge differences in computational performance can be found for different coupling layouts due to different scalability of the modular *component*.

Since computational efficiency is one of the key aspects of any coupled system, the various aspects affecting it are discussed. These are the performances of the *component*, of the coupling library and of the coupled system. Hereby the design of the interface and the OASIS3-MCT coupling parameters, which enable optimisation of the efficiency, are described.

The *component*'s performance depends on its scalability. The optimum partitioning has to be set for each parallel component by means of a strong scaling analysis (discussed in Sect. 4.1). This analysis, which results in finding the scalability limit (the maximum speed) or the scalability optimum (the acceptable level of parallel efficiency), can be difficult to obtain for each *component* in a multi-component context. In this article, we propose to simply consider the previously defined concept of the computing time (excluding the waiting time from the total time to solution). In Sect. 4 we will describe our strategy to separate the measurement of computing and waiting times for each *component* and how to deduce the optimum MPI partitioning from the scaling analysis.

The optimisation of OASIS3-MCT coupling library performance is relevant for the efficiency of the data exchange between *components* discretised on different grids. The parallelised interpolations are performed by the OASIS3-MCT library routines called by the source or by the target *component*. An interpolation will be faster if performed (1) by the model with the larger number of MPI processes available (up to the OASIS3-MCT interpolation scalability limit) and/or (2) by the fastest model (until the OASIS3-MCT interpolation together with the fastest model's calculations last longer than the calculations of the slowest model).

A significant improvement of interpolation and communication performances can be achieved by coupling of multiple variables that share the same coupling characteristics via a single communication, that is, by using the technique called *pseudo-3-D coupling*. Via this option, a single interpolation and a single send/receive instruction are executed for a whole group of coupling fields, for example, all levels and variables in an atmosphere–atmosphere coupling at one time instead of all coupling fields and levels separately. The option groups several small MPI messages into a big one and, thus, reduces communications. Furthermore, the number of matrix multiplications is reduced because it is performed on big arrays. This functionality can easily be set via the “namcouple” parameter file (see Sect. B2.4 in Valcke et al., 2013). The impact on the performance of CCLM atmosphere–atmosphere coupling is discussed in Sect. 4.6). See also Maisonave et al. (2013).

The optimisation of the performance of a coupled system relies on the allocation of an optimum number of computing resources to each model. If the *components*' calculations are performed concurrently, the waiting time needs to be min-

imised. This can be achieved by balancing the load of the two (or more) *components* between the available computing resources: the slower *component* is granted more resources, leading to an increase in its parallelism and a decrease in its computing time. The opposite is done for the fastest *component* until an equilibrium is reached. Section 4 gives examples of this operation and describes the strategy to find a compromise between each *component*'s optimum scalability and the load balance between all *components*.

On all high-performance operating systems it is possible to run one process of a parallel application on one core in a so-called *single-threading* (ST) mode (Fig. 2a). Should the core of the operating system feature the so-called *simultaneous multi-threading* (SMT) mode, two (or more) processes/threads of the same (in a *non-alternating process distribution*; Fig. 2b) or of different (in an *alternating process distribution*; Fig. 2c) applications can be executed simultaneously on the same core. Applying SMT mode is more efficient for well-scaling parallel applications, leading to an increase in speed of the order of magnitude of 10% compared to the ST mode. Usually it is possible to specify which process is executed on which core (see Fig. 2). In these cases the SMT mode with alternating distribution of *component* processes can be used, and the waiting time of sequentially coupled *components* can be avoided. Starting each model *component* on each core is usually the optimum configuration, since the reduction of the waiting time of cores outperforms the increase in the time to solution by using ST mode instead of SMT mode (at each time one process is executed on each core). In the case of concurrent couplings, however, it is possible to use SMT mode with a non-alternating process distribution.

The optimisation procedure applied is described in more detail in Sect. 4.3 for the couplings considered. The results are discussed in Sect. 4.6.

3.1.4 Regional climate model coupling particularities

In addition to the standard OASIS functionalities, some adaptation of the OASIS3-MCT API routines were necessary to fit special requirements of the regional-to-regional and regional-to-global couplings presented in this article.

A regional model covers only a portion of earth's sphere and requires boundary conditions at its domain boundaries. This has two immediate consequences for coupling: first, two regional models do not necessarily cover exactly the same part of earth's sphere. This implies that the geographic boundaries of the model's computational domains and of coupled variables may not be the same in the source and target *components* of a coupled system. Second, a regional model can be coupled with a global model or another limited-area model, and some of the variables which need to be exchanged are 3-D, as in the case of atmosphere-to-atmosphere or ocean-to-ocean coupling.

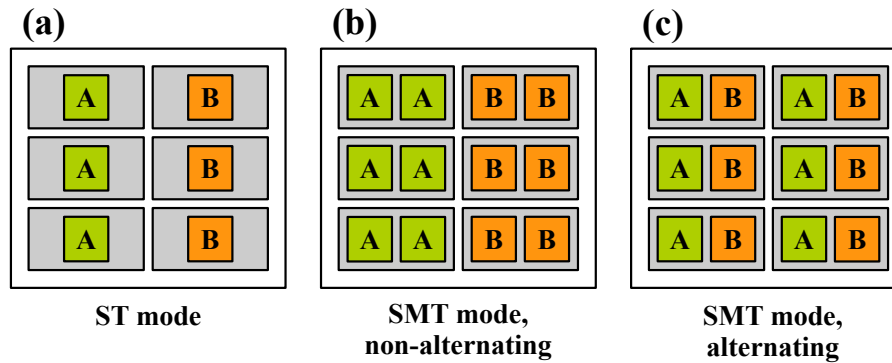


Figure 2. Schematic processes distribution on a hypothetical computing node with six cores (grey-shaded areas) in (a) ST mode, (b) SMT mode with non-alternating processes distribution and (c) SMT mode with alternating processes distribution. “A” and “B” are processes belonging to two different *components* of the model system sharing the same node. In (b) and (c) two processes of the same (b) or different (c) *component* share one core using the simultaneous multi-threading (SMT) technique, while in (a) only one process per core is launched in the single-threading (ST) mode.

A major part of the OASIS community uses global models. Therefore, OASIS standard features fit global model coupling requirements. Consequently, the coupling library must be adapted or used in an unconventional way, described in the following, to be able to cope with the extra demands mentioned.

Limited-area field exchange has to deal with a mismatch of the domains of the models coupled. Differences between the (land and ocean) models coupled to CCLM lead to two solutions for the mismatch of the model domains. For coupling with the Community Land Model (CLM) the CLM domain is extended in such a way that at least all land points of the CCLM domain are covered. Then, all CLM grid points located outside of the CCLM domain are masked. To achieve this, a uniform array on the CCLM grid is interpolated by OASIS3-MCT to the CLM grid using the same interpolation method as for the coupling fields. On the CLM grid the uniform array contains the projection weights of the CCLM on the CLM grid points. This field is used to construct a new CLM domain containing all grid points necessary for interpolation. However, this solution is not applicable to all coupled-system *components*. In ocean models, a domain modification would complicate the definition of ocean boundary conditions or even lead to numerical instabilities at the new boundaries. Thus, the original ocean domain, which must be smaller than the CCLM domain, is interpolated to the CCLM grid. At runtime, all CCLM ocean grid points located inside the interpolated area are filled with values interpolated from the ocean model and all CCLM ocean grid points located outside the interpolated area are filled with external forcing data.

Multiple usage of the MCT library occurred in the CCLM+CLM coupled system implementation making some modifications of the OASIS3-MCT version 2.0 necessary. Since the MCT library has no re-entrancy properties, a duplication of the MCT library and a renaming of the OASIS3-

MCT calling instruction were necessary. This modification ensures the capability of coupling any other CESM *component* via OASIS3-MCT. The additional usage of the MCT library occurred in the CESM framework of CLM version 4.0. More precisely, the DATM model interface in the CESM module is using the CPL7 coupler including the MCT library for data exchange.

Interpolation of 3-D fields is necessary in an atmosphere-to-atmosphere coupling. The OASIS3-MCT library is used to provide 3-D boundary conditions to the regional model and a 3-D feedback to the global coarse-grid model. OASIS is not able to interpolate the 3-D fields vertically, mainly because of the complexity of vertical interpolations in geophysical models (different orographies, level numbers and formulations of the vertical grid). However, it is possible to decompose the operation into two steps: (1) horizontal interpolation with OASIS3-MCT and (2) model-specific vertical interpolation performed in the source or target *component's* interface. The first operation does not require any adaption of the OASIS3-MCT library and can be solved in the most efficient manner by the pseudo-3-D coupling option described in Sect. 3.1.3. The second operation requires a case-dependent algorithm addressing aspects such as interpolation and extrapolation of the boundary layer over different orographies, change in the coordinate variable, conservation properties as well as interpolation efficiency and accuracy.

An exchange of 3-D fields, which occurs in the CCLM+MPI-ESM coupling, requires a more intensive usage of the OASIS3-MCT library functionalities than observed so far in the climate modelling community. The 3-D regional-to-global coupling is even more computationally demanding than its global-to-regional opposite. Now, all grid points of the CCLM domain have to be interpolated instead of just the grid points of a global domain that are covered by the regional domain. The amount of data exchanged is rarely reached by any other coupled system of the community due

to (1) the high number of exchanged 2-D fields, (2) the high number of exchanged grid points (full CCLM domain) and (3) the high exchange frequency at every ECHAM time step. In addition, as will be explained in Sect. 3.2, the coupling between CCLM and MPI-ESM needs to be sequential and, thus, the exchange speed has a direct impact on the simulation's total time to solution.

Interpolation methods used in OASIS3-MCT are the SCRIP standard interpolations: bilinear, bicubic, first- and second-order conservative. However, the interpolation accuracy might not be sufficient and/or the method is inappropriate for certain applications. This is for example the case with the atmosphere-to-atmosphere coupling CCLM+MPI-ESM. The linear methods turned out to be of low accuracy and the second-order conservative method requires the availability of the spatial derivatives on the source grid. Up to now, the latter cannot be calculated efficiently in ECHAM (see Sect. 3.2 for details). Other higher-order interpolation methods can be applied by providing weights of the source grid points at the target grid points. This method was successfully applied in the CCLM+MPI-ESM coupling by application of a bicubic interpolation using a 16-point stencil. In Sect. 3.2 to 3.5 the interpolation methods recommended for the individual couplings are given.

3.2 CCLM+MPI-ESM

The CCLM+MPI-ESM two-way coupled system presented here provides a stable solution over climatological timescales. In the CCLM+MPI-ESM two-way coupled system the 3-D atmospheric fields are exchanged between the non-hydrostatic atmosphere model of CCLM and the ECHAM hydrostatic atmosphere model of MPI-ESM. In MPI-ESM the CCLM solution is replacing the ECHAM solution within the coupled (limited-area) domain of the global atmosphere. In CCLM the MPI-ESM solution is used as a boundary condition at the top, lateral and ocean bottom boundaries in the same way as in standard one-way nesting. Both models, CCLM and MPI-ESM, run sequentially (see also Appendix B).

CCLM recalculates the ECHAM time step in dependence on the boundary conditions provided by MPI-ESM. In MPI-ESM the ECHAM solution is updated within the coupled domain of the globe using the solution provided by CCLM. The CCLM is solving the equations in physical space. ECHAM is using the transform method between the physical and the spectral space. For computational-efficiency reasons the data exchange in ECHAM is done in grid point space. This avoids costly transformations between grid point and spectral space. Since the simulation results of CCLM need to become effective in ECHAM dynamics, the two-way coupling is implemented in ECHAM after the transformation from spectral to grid point space and before the computation of advection (see Figs. 8 and DKRZ, 1993 for details).

Table 4. Variables exchanged between CCLM and the MPI-ESM global model. The CF standard-names convention is used. Units are given as defined in CCLM. \otimes : information is sent by CCLM; \boxtimes : information is received by CCLM. 3-D indicates that a three-dimensional field is sent/received.

Variable (unit)	CCLM+MPI-ESM
Temperature (K)	$\boxtimes\otimes$ 3-D
U component of wind (m s^{-1})	$\boxtimes\otimes$ 3-D
V component of wind (m s^{-1})	$\boxtimes\otimes$ 3-D
Specific humidity (kg kg^{-1})	$\boxtimes\otimes$ 3-D
Specific cloud liquid water content (kg kg^{-1})	$\boxtimes\otimes$ 3-D
Specific cloud ice content (kg kg^{-1})	$\boxtimes\otimes$ 3-D
Surface pressure (Pa)	$\boxtimes\otimes$
Sea surface temperature SST (K)	\boxtimes
Surface snow amount (m)	\boxtimes
Surface geopotential (m s^{-2})	\boxtimes

$$\text{SST} = (\text{sea_ice_area_fraction} \cdot T_{\text{sea ice}}) + (\text{SST} \cdot (1 - \text{sea_ice_area_fraction}))$$

ECHAM provides the boundary conditions for CCLM at time level $t = t_n$ of the three time levels $t_n - (\Delta t)_E$, t_n and $t_n + (\Delta t)_E$ of ECHAM's leap frog time integration scheme. However, the second part of the Assilin time filtering in ECHAM for this time level has to be executed after the advection calculation in `dyn` (see Fig. 8) in which the tendency due to two-way coupling needs to be included. Thus, the fields sent to CCLM as boundary conditions do not undergo the second part of the Assilin time filtering. The CCLM is integrated over j time steps between the ECHAM time level t_{n-1} and t_n . However, the coupling time may also be a multiple of an ECHAM time step $(\Delta t)_E$.

A complete list of variables exchanged between ECHAM and CCLM is given in Table 4. The time step organisation is described in Appendix B and shown in Fig. 7 for CCLM and in Fig. 8 for ECHAM. The data sent in routine `couple_put_e2c` of ECHAM to OASIS3-MCT are the 3-D variables temperature, u and v components of the wind velocity, specific humidity, cloud liquid and ice water content and the 2-D fields surface pressure, surface temperature and surface snow amount. At initial time the surface geopotential is sent for calculation of the orography differences between the model grids. After horizontal interpolation to the CCLM grid via the bilinear SCRIP interpolation¹ by OASIS3-MCT, the 3-D variables are received in CCLM by the routine `receive fld` and vertically interpolated to the CCLM grid keeping the height of the 300 hPa level constant and using the hydrostatic approximation. Afterwards, the horizontal wind vector velocity components of ECHAM are rotated from the geographical (lon, lat) ECHAM to the rotated (rlon, rlat) CCLM coordinate system. Here the

¹This interpolation is used for the performance tests only. For physical coupling the conservative interpolation second order (CO2) is used, which requires an additional computation of derivatives. Alternatively, a bicubic interpolation can be used which has the same accuracy as CO2.

receive_fld routine and the additional computations of online coupling ECHAM_2_CCLM in CCLM end and the interpolated data are used to initialise the bound lines at the next CCLM time levels $t_m = t_{n-1} + k \cdot (\Delta t)_C \leq t_n$, with $k \leq j = (\Delta t)_E / (\Delta t)_C$. However, the final time of CCLM integration $t_{m+j} = t_m + j \cdot (\Delta t)_C = t_n$ is equal to the time t_n of the ECHAM data received.

After integrating between $t_n - i \cdot (\Delta t)_E$ and t_n , the 3-D fields of temperature, u and v velocity components, specific humidity and cloud liquid and ice water content of CCLM are vertically interpolated to the ECHAM vertical grid in the send_fld routine following the same procedure as in the CCLM receive interface and keeping the height of the 300 hPa level of the CCLM pressure constant. The wind velocity vector components are rotated back to the geographical directions of the ECHAM grid. The 3-D fields and the hydrostatically approximated surface pressure are sent to OASIS3-MCT, horizontally interpolated to the ECHAM grid by OASIS3-MCT² and received in ECHAM grid space in routine couple_get_c2e. In ECHAM the CCLM solution is relaxed at the lateral and top boundaries of the CCLM domain by means of a cosine weight function over a range of 5 to 10 ECHAM grid boxes using a weight between zero at the outer boundary and one in the central part of the CCLM domain. Additional fields are calculated and relaxed in the CCLM domain for a consistent update of the ECHAM prognostic variables. These are the horizontal derivatives of temperature, surface pressure, u and v wind velocity, divergence and vorticity.

A strong initialisation perturbation is avoided by slowly increasing the maximum coupling weight to 1 with time, following the function $\text{weight} = \text{weight}_{\max} \cdot (\sin((t/t_{\text{end}}) \cdot \pi/2))$, with t_{end} equal to 1 month.

3.3 CCLM+NEMO-MED12

CCLM and the NEMO ocean model are coupled concurrently for the Mediterranean Sea (NEMO-MED12) and for the North and Baltic seas (NEMO-NORDIC). Table 5 gives an overview of the variables exchanged. Bicubic interpolation between the horizontal grids is used for all variables.

At the beginning of the NEMO time integration (see Fig. 7) the CCLM receives the sea surface temperature (SST) and – only in the case of coupling with the North and Baltic seas – also the sea ice fraction from the ocean model. At the end of each NEMO time step CCLM sends average water, heat and momentum fluxes to OASIS3-MCT. In the NEMO-NORDIC set-up CCLM additionally sends the averaged sea level pressure (SLP)_C needed in NEMO to link the exchange of water between the North and Baltic seas directly to the atmospheric pressure. The sea ice fraction affects the radia-

tive and turbulent fluxes due to different albedo and roughness length of ice. In both coupling set-ups SST is the lower boundary condition for CCLM and is used to calculate the heat budget in the lowest atmospheric layer. The averaged wind stress is a direct momentum flux for NEMO to calculate the water motion. Solar and non-solar radiation are needed by NEMO to calculate the heat fluxes. $E-P$ (evaporation minus precipitation) is the net gain ($E - P < 0$) or loss ($E - P > 0$) of freshwater at the water surface. This water flux adjusts the salinity of the uppermost ocean layer.

In all CCLM grid cells where there is no active ocean model underneath, the lower boundary condition (SST) is taken from ERA-Interim re-analyses. The sea ice fraction in the Atlantic Ocean is derived from the ERA-Interim SST where $\text{SST} < -1.7^\circ\text{C}$, which is a salinity-dependent freezing temperature.

On the NEMO side, the coupling interface is included similarly to CCLM, as can be seen in Fig. 9. There is a set-up of the coupling interface at the beginning of the NEMO simulation. At the beginning of the time loop NEMO receives the upper boundary conditions from OASIS3-MCT and, before the time loop ends, it sends the coupling fields (average SST and sea ice fraction for NEMO-NORDIC) to OASIS3-MCT.

3.4 CCLM+TRIMNP+CICE

In the CCLM+TRIMNP+CICE coupled system (denoted as COSTRICE; Ho-Hagemann et al., 2013), all fields are exchanged every hour between the three models CCLM, TRIMNP and CICE running concurrently. An overview of variables exchanged among the three models is given in Table 5. The “surface temperature over sea/ocean” is sent to CCLM instead of “SST” to avoid a potential inconsistency in case of sea ice existence. As shown in Fig. 7, CCLM receives the skin temperature (T_{skin}) at the beginning of each CCLM time step over the coupling areas, the North and Baltic seas. The skin temperature T_{skin} is a weighted average of sea ice and sea surface temperature. It is not a linear combination of skin temperatures over water and over ice weighted by the sea ice fraction. Instead, the skin temperature over ice T_{Ice} and the sea ice fraction A_{Ice} of CICE are sent to TRIMNP, where they are used to compute the heat flux HFL, that is, the net outgoing long-wave radiation. HFL is used to compute the skin temperature of each grid cell via the Stefan–Boltzmann law.

At the end of the time step, after the physics and dynamics computations and output writing, CCLM sends the variables listed in Table 5 to TRIMNP and CICE for calculation of wind stress, freshwater, momentum and heat flux. TRIMNP can either directly use the sensible and latent heat fluxes from CCLM (considered as the flux coupling method; see e.g. Döscher et al., 2002) or compute the turbulent fluxes using the temperature and humidity density differences between air and sea as well as the wind speed (considered as the coupling method via state variables; see e.g.

²The bilinear interpolation is used. The usage of a second-order conservative interpolation requires horizontal derivatives of the variables exchanged. This is not implemented in this version of the CCLM send interface.

Table 5. As Table 4 but variables exchanged between CCLM and the NEMO, TRIMNP and CICE ocean models.

Variable (unit)	CCLM+ NEMO-MED12	CCLM+ NEMO-NORDIC	CCLM+ TRIMNP+ CICE
Surface temperature over sea/ocean (K)	☐	☐	☐
2 m temperature (K)	–	–	⊗
Potential temperature NSL (K)	–	–	⊗
Temperature NSL (K)	–	–	⊗
Sea ice area fraction (1)	–	☐	–
Surface pressure (Pa)	–	⊗	–
Mean sea level pressure (Pa)	–	–	⊗
Surface downward eastward and northward stress (Pa)	⊗	⊗	–
Surface net downward short-wave flux ($W m^{-2}$)	⊗	⊗	⊗
Surface net downward long-wave flux ($W m^{-2}$)	–	–	⊗
Non-solar radiation NSR ($W m^{-2}$)	⊗	⊗	–
Surface downward latent heat flux ($W m^{-2}$)	–	–	⊗
Surface downward heat flux HFL ($W m^{-2}$)	–	–	⊗
Evaporation–precipitation $E-P$ ($kg m^{-2}$)	⊗	⊗	–
Total precipitation flux TPF ($kg m^{-2} s^{-1}$)	–	–	⊗
Rain flux RF ($kg m^{-2} s^{-1}$)	–	–	⊗
Snow flux SF ($kg m^{-2} s^{-1}$)	–	–	⊗
U and V component of 10 m wind ($m s^{-1}$)	–	–	⊗
2 m relative humidity (%)	–	–	⊗
Specific humidity NSL ($kg kg^{-1}$)	–	–	⊗
Total cloud cover (1)	–	–	⊗
Half height of lowest CCLM level (m)	–	–	⊗
Air density NSL ($kg m^{-3}$)	–	–	⊗

NSL = lowest (near-surface) level of the 3-D variable;

NSR = surface net downward long-wave flux + surface downward latent and sensible heat flux;

HFL = surface net downward short-wave flux + surface downward long-wave flux + surface downward latent and sensible heat flux;

TPF = RF + SF = convective and large-scale rainfall flux + convective and large-scale snowfall flux;

$E-P$ = $-(\text{surface downward latent heat flux/LHV}) - \text{TPF}$; LHV = latent heat of vapourisation = $2.501 \times 10^6 J kg^{-1}$.

Rummukainen et al., 2001). The method used is specified in the subroutine `heat_flux` of TRIMNP.

In addition to the fields received from CCLM, the CICE sea ice model requires from TRIMNP the SST, salinity, water velocity components, ocean surface slope, and freezing/melting potential energy. CICE sends to TRIMNP the water and ice temperature, sea ice fraction, freshwater flux, ice-to-ocean heat flux, short-wave flux through ice to ocean and ice stress components. The horizontal interpolation method applied in CCLM+TRIMNP+CICE is the SCRIP nearest-neighbour inverse-distance-weighting fourth-order interpolation (DISTWGT).

Note that the coupling method differs between CCLM+TRIMNP+CICE and CCLM+NEMO-NORDIC (see Sect. 3.3). In the latter, SSTs and sea ice fraction from NEMO are sent to CCLM so that the sea ice fraction from NEMO affects the radiative and turbulent fluxes of CCLM due to different albedo and roughness length of ice. But in CCLM+TRIMNP+CICE, only SSTs are passed to CCLM. Although these SSTs implicitly contain information of sea ice fraction, which is sent from CICE to TRIMNP, the albedo of sea ice in CCLM is not taken from CICE but calculated in the atmospheric model independently. The reason for this inconsistent calculation of albedo between these two coupled systems originates from a fact that a

tile-approach has not been applied for the CCLM version used in the present study. Here, partial covers within a grid box are not accounted for, hence, partial fluxes, i.e. the partial sea ice cover, snow on sea ice and water on sea ice are not considered. In a water grid box of this CCLM version, the albedo parameterisation switches from ocean to sea ice if the surface temperature is below a freezing temperature threshold of $-1.7^\circ C$. Coupled to NEMO-NORDIC, CCLM obtains the sea ice fraction, but the albedo and roughness length of a grid box in CCLM are calculated as a weighted average of water and sea ice portions which is a parameter aggregation approach.

Moreover, even if the sea ice fraction from CICE would be sent to CCLM, such as done for NEMO-NORDIC, the latent and sensible heat fluxes in CCLM would still be different to those in CICE due to different turbulence schemes of the two models CCLM and CICE. This different calculation of heat fluxes in the two models leads to another inconsistency in the current set-up which can only be removed if all models coupled use the same radiation and turbulent energy fluxes. These fluxes should preferably be calculated in one of the models at the highest resolution, for example in the CICE model for fluxes over sea ice. Such a strategy shall be applied in future studies, but is beyond the scope of the CCLM version used in this study.

3.5 CCLM+VEG3D and CCLM+CLM

The two-way couplings between CCLM and VEG3D and between CCLM and CLM are implemented in a similar way. First, the call to the LSM (OASIS send and receive; see Fig. 7) is placed at the same location in the code as the call to CCLM's native land surface scheme, TERRA_ML, which is switched off when either VEG3D or CLM is used. This ensures that the sequence of calls in CCLM remains the same regardless of whether TERRA_ML, VEG3D or CLM is used. In the default configuration used here CCLM and CLM (or VEG3D) are executed sequentially, thus mimicking the “sub-routine” type of coupling used with TERRA_ML. Note that it is also possible to run CCLM and the LSM concurrently, but this is not discussed here. Details of the time step organisation of VEG3D and CLM are described in the Appendix and shown in Figs. 12 and 13.

VEG3D runs at the same time step and on the same horizontal rotated grid (0.44° here) as CCLM with no need for any horizontal interpolations. CLM uses a regular lat–lon grid and the coupling fields are interpolated using bilinear interpolation (atmosphere to LSM) and distance-weighted interpolation (LSM to atmosphere). The time step of CLM is synchronised with the CCLM radiative transfer scheme time step (1 h in this application) with the idea that the frequency of the radiation update determines the radiative forcing at the surface.

The LSMs need to receive the following atmospheric forcing fields (see also Table 6): the total amount of precipitation, the short- and long-wave downward radiation, the surface pressure, the wind speed, the temperature and the specific humidity of the lowest atmospheric model layer.

VEG3D additionally needs information about the time-dependent composition of the vegetation to describe its influence on radiation interactions and turbulent fluxes correctly. This includes the leaf area index, the plant cover and a vegetation function which describes the annual cycle of vegetation parameters based on a simple cosine function depending on latitude and day. They are exchanged at the beginning of each simulated day.

One specificity of the coupling concerns the turbulent fluxes of latent and sensible heat. In its turbulence scheme, CCLM does not directly use surface fluxes. It uses surface states (surface temperature and humidity) together with turbulent diffusion coefficients of heat, moisture and momentum. Therefore, the diffusion coefficients need to be calculated from the surface fluxes received by CCLM. This is done by deriving, in a first step, the coefficient for heat (assumed to be the same as the one for moisture in CCLM) based on the sensible heat flux. In a second step an effective surface humidity is calculated using the latent heat flux and the derived diffusion coefficient for heat.

4 Computational efficiency

Computational efficiency is an important property of a numerical model's usability and applicability and has many aspects. A particular coupled model system can be very inefficient even if each *component* has a high computational efficiency in stand-alone mode and in other couplings. Thus, optimising the computational performance of a coupled model system can save a substantial amount of resources in terms of simulation time and cost. We focus here on aspects of computational efficiency related directly to coupling of different models overall tested in other applications and use real case model configurations for each *component* of a coupled system.

We use a three step approach. First, the scalability of different coupled model systems and of its *components* is investigated. Second, an optimum configuration of resources is derived and third, different components of extra cost of coupling at optimum configuration are quantified. For this purpose the *Load-balancing Utility and Coupling Implementation Appraisal* (LUCIA), developed at CERFACS, Toulouse, France (Maisonnavé and Caubel, 2014) is used, which is available together with the OASIS3-MCT coupler.

More precisely, we investigate the scalability of each coupled system's *component* in terms of simulation speed, computational cost and parallel efficiency, the time needed for horizontal interpolations by OASIS3-MCT and the load balance in the case of concurrently running *components*. Based on these results, an optimum configuration for all couplings is suggested. Finally, the cost of all *components* at optimum configurations are compared with the cost of CCLM stand-alone at configuration used in coupled system and at optimum configuration ($CCLM_{sa,OC}$) of the stand-alone simulation.

4.1 Simulation set-up and methodology

A parallel program's runtime $T(n, R)$ mainly depends on two variables: the problem size n and the number of cores R , that is, the resources. In scaling theory, a *weak scaling* is performed with the notion of solving an increasing problem size in the same time, while as in a *strong scaling* a fixed problem size is solved more quickly with an increasing amount of resources. Due to resource limits on the common high-performance computer we chose to conduct a strong-scaling analysis with a common model set-up allowing for an easier comparability of the results. By means of the scalability study we identified an optimum configuration for each coupling which served as a basis to address two central questions. (1) How much does it cost to add one (or more) *component(s)* to CCLM? (2) How big are the costs of different *components* and of OASIS3-MCT to transform the information between the *components'* grids? The first question can only be answered by a comparison to a reference which is, in this study, a CCLM stand-alone simulation. The second question

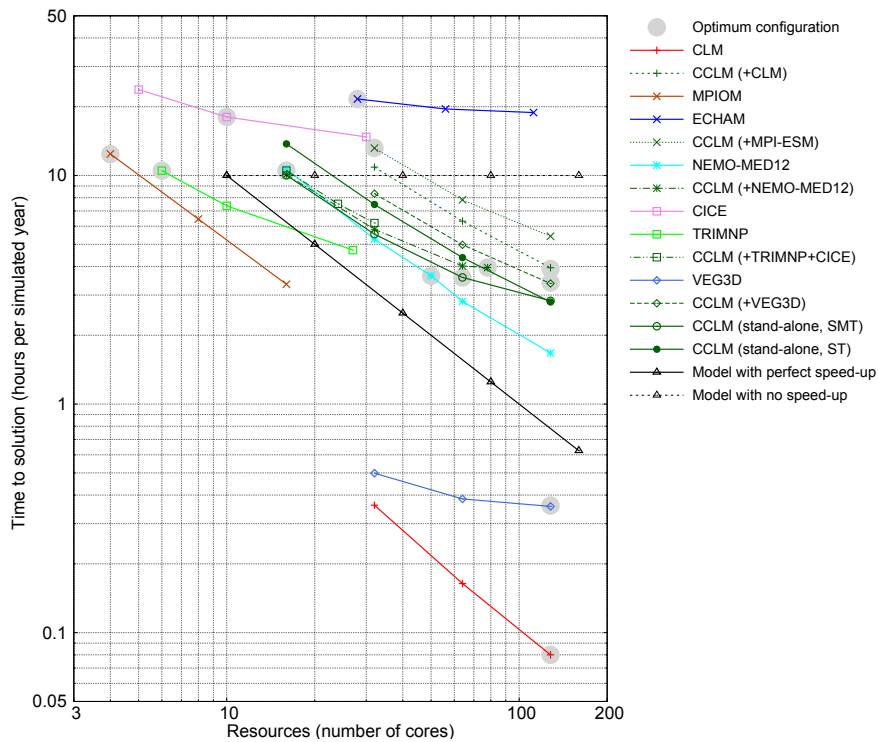


Figure 3. Time to solution of model components of the coupled systems (indicated for CCLM in brackets) and for CCLM stand-alone ($CCLM_{sa}$) in hours per simulated year (HPSY) in dependence on the computational resources (number of cores) in single-threading (ST) and multi-threading (SMT) mode. The times for model components ECHAM and MPIOM of MPI-ESM are given separately. The optimum configuration of each component is highlighted by a grey dot. The hypothetical result for a model with perfect and no speed-up is given as well.

can directly be answered by the measurements of LUCIA. We used this OASIS3-MCT tool to measure the computing and waiting time of each *component* in a coupled model system (see Sect. 3.1.3) as well as the time needed for interpolation of fields before and after sending or receiving.

A recommended configuration was chosen for the COSMO-CLM reference model at 0.44 horizontal resolution. The other *components'* set-ups are those used by the developers of the particular coupling (see Sect. 2 for more details) for climate modelling applications in the CORDEX-EU domain. This means that I/O, model physics and dynamics are chosen in the same way as for climate applications in order to obtain a realistic estimate of the performance of the couplings. The simulated period is 1 month; the horizontal grid has 132 by 129 grid points and 0.44° (ca. 50 km) horizontal grid spacing. In the vertical, 45 levels are used for the CCLM+MPI-ESM and CCLM+VEG3D couplings as well as for the $CCLM_{sa}$ simulations. All other couplings use 40 levels. The impact of this difference on the numerical performance is compensated for by a simple post-processing scaling of the measured CCLM computing time $T_{CCLM,45}$ of the CCLM component that employs 45 levels assuming a linear scaling of the CCLM computing time with the number of

levels as $T_{CCLM} = 0.8 \cdot T_{CCLM,45} \cdot \frac{40}{45} + 0.2 \cdot T_{CCLM,45}$.³ The usage of a real-case configuration allows one to provide realistic computing times.

The computing architecture used is *Blizzard* at *Deutsches Klimarechenzentrum* (DKRZ) in Hamburg, Germany. It is an IBM Power6 machine with nodes consisting of 16 dual-core CPUs (16 processors, 32 cores). Simultaneous multi-threading (SMT; see Sect. 3.1.3) allows one to launch two processes on each core. A maximum of 64 threads can be launched on one node.

The measures used in this paper to present and discuss the computational performance are well known in scalability analyses: (1) *time to solution* in Hours Per Simulated Year (HPSY), (2) *cost* in Core Hours Per Simulated Year (CHPSY) and (3) *parallel efficiency* (PE) (see Table 7 for details).

Usually, $HPSY_1$ is the time to solution of a *component* executed serially, that is, using one process ($R = 1$) and $HPSY_2$ is the time to solution if executed using $R_2 > R_1$ parallel processes. Some *components*, like ECHAM, cannot be executed serially. This is why the reference number of threads is $R_1 \geq 2$ for all coupled-system *components*.

³The estimation that 80 % of CCLM's computations depend on the number of model levels is based on CCLM's internal time measurements. $T_{CCLM,45}$ is the time measured by LUCIA.




Figure 4. As Fig. 3 but for the cost of the *components* in core hours per simulated year.




Figure 5. As Fig. 3 but for the parallel efficiency of the *components* in % of the reference configuration.

Table 6. As Table 4 but variables exchanged between CCLM and the VEG3D and CLM land surface models.

Variable (unit)	CCLM+VEG3D	CCLM+CLM
Leaf area index (1)	⊗	–
Plant cover (1)	⊗	–
Vegetation function (1)	⊗	–
Surface albedo (1)	□	□
Height of lowest level (m)	–	⊗
Surface pressure (Pa)	⊗	–
Pressure NSL (Pa)	⊗	⊗
Snow flux SF ($\text{kg m}^{-2} \text{s}^{-1}$)	⊗	⊗
Rain flux RF ($\text{kg m}^{-2} \text{s}^{-1}$)	⊗	⊗
Temperature NSL (K)	⊗	⊗
Grid-mean surface temperature (K)	□	□
Soil surface temperature (K)	□	–
Snow surface temperature (K)	□	–
Surface snow amount (m)	□	–
Density of snow (kg m^{-3})	□	–
Thickness of snow (m)	□	–
Canopy water amount (m)	□	–
Specific humidity NSL (kg kg^{-1})	⊗	⊗
Surface specific humidity (kg kg^{-1})	□	–
Subsurface runoff (kg m^{-2})	□	–
Surface runoff (kg m^{-2})	□	–
Wind speed $ \mathbf{v} $ NSL (m s^{-1})	⊗	–
U and V component of wind NSL (m s^{-1})	–	⊗
Surface downward sensible heat flux (W m^{-2})	□	□
Surface downward latent heat flux (W m^{-2})	–	□
Surface direct and diffuse downwelling short-wave flux in air (W m^{-2})	⊗	⊗
Surface net downward long-wave flux (W m^{-2})	⊗	⊗
Surface flux of water vapour ($\text{s}^{-1} \text{m}^{-2}$)	□	–
Surface downward eastward and northward flux (U/V momentum flux, Pa)	–	□

NSL: lowest (near-surface) level of the 3-D variable
 RF: convective and large-scale rainfall flux; SF: convective and large-scale snowfall flux;
 SWD_S: surface diffuse and direct downwelling short-wave flux in air.

Table 7. Measures of computational performance used for computational performance analysis.

Measure (unit)	Acronym	Description
simulated years (1)	sy	Number of simulated physical years
number of cores (1)	n	Number of computational cores used in a simulation per model component
number of threads (1)	R	Number of parallel processes or threads configured in a simulation per model component. On <i>Blizzard</i> at DKRZ one or two threads can be started on one core.
time to solution (HPSY)	T	Simulation time of a model component measured by LUCIA per simulated year
speed (HPSY ⁻¹)	s	$= T^{-1}$ is the number of simulated years per simulated hour by a model component
costs (CHPSY)	–	$= T \cdot n$ is the core hours used by a model component running on n cores per simulated year
speed-up (%)	SU	$= \frac{\text{HPSY}_1(R_1)}{\text{HPSY}_2(R_2)} \cdot 100$ is the ratio of time to solution of a model component configured for reference and actual number of threads
parallel efficiency (%)	PE	$= \frac{\text{CHPSY}_1}{\text{CHPSY}_2} \cdot 100$ is the ratio of core hours per simulated year for reference (CHPSY ₁) and actual (CHPSY ₂) number of cores




Figure 6. Time to solution and cost of *components* of the coupled systems at optimum configuration of couplings investigated and of stand-alone CCLM. The boxes' widths correspond to the number of cores used per component. The area of each box is equal to the costs (the amount of core hours per simulated year) consumed by each component calculations, including coupling interpolations. The white areas indicate the load imbalance between concurrently running components. See Table 8 for details.

If the resources of a perfectly scaling parallel application are doubled, the speed would be doubled and therefore the cost would remain constant, the parallel efficiency would be 100 %, and the speed-up would be 200 %. A parallel efficiency of 50 % is reached if the costs of CHPSY₂ are twice as big as those of the reference configuration CHPSY₁.

Inconsistencies of the time to solution of approximately 10 % were found between measurements obtained from simulations conducted at two different physical times. This gives a measure of the dependency of the time to solution on the status of the machine used, particularly originating from the I/O. Nevertheless, the time to solution and cost are given with higher accuracy to highlight the consistency of the numbers.

4.2 Scalability results

Figure 3 shows the results of the performance measurement *time to solution* for all *components* individually in coupled mode and for CCLM_{sa} (in ST and SMT mode). As reference, the slopes of a model at no speed-up and at perfect speed-up are shown. Three groups can be identified. CLM and VEG3D have the shortest times to solution and, thus, they are the fastest *components*. The three models of regional oceans coupling with CCLM and the CCLM models in coupled as well as in stand-alone mode need about 2–10 HPSY. The overall slowest *components* are CICE and ECHAM which need about 20 HPSY at reference configuration. Within the range

of resources investigated CICE, ECHAM and VEG3D exhibit almost no speed-up in coupled mode (i.e. including additional computations). On the contrary, MPIOM, NEMO-MED12 and CLM have a very good scalability up to the tested limit of 128 cores.

Figure 4 shows the second relevant performance measure, the absolute cost of computation in core hours per simulated year for the same couplings together with the perfect and no speed-up slopes. The aforementioned three groups slightly change their composition. VEG3D and CLM are not only the fastest, but also the cheapest *components*, the latter becoming even cheaper with increasing resources. A little bit more expensive but mostly of the same order of magnitude as the land surface *components* are the regional ocean *components* MPIOM and TRIMNP followed by CICE, NEMO-MED12 and all the different coupled CCLMs. The NEMO model is approximately 2 times more expensive than TRIMNP. The configuration of the CICE model is as expensive as the CCLM regional climate model. The cost of CCLM differs by a factor of 2 between the stand-alone and different coupled versions. The most expensive one is coupled to ECHAM, which is also the most expensive *component*.

In order to analyse the performance of the couplings in more detail, we took measurements of the stand-alone CCLM in single-threading (ST) and multi-threading (SMT) mode. The direct comparison provides the information on how much CCLM's speed and cost benefit from switching from ST to SMT mode. As shown in Fig. 3 at 16 cores the CCLM in SMT mode is 27 % faster. When allocating 128 cores both modes arrive at about the same speed. This can be explained by increasing cost of MPI communications with decreasing number of grid points/thread. Since the number of threads in SMT mode is twice for the same core number and thus the number of grid points per thread is half, the scalability limit of approximately 1.5 points exchanged per computational grid point is reached at approximately 100 points/thread (if three bound lines are exchanged), resulting in a scalability limit at approximately 80 cores in SMT mode and 160 cores in ST mode (see also the CCLM+NEMO-MED12 coupling in Sect. 4.4).

4.3 Strategy for finding an optimum configuration

The optimisation strategy that we pursue is empirical rather than strictly mathematical, which is why we understand “optimum” more as “near-optimum”. Due to the heterogeneity of our coupled systems, a single algorithm cannot be proposed (as in Balaprakash et al., 2014). Nonetheless, our results show that these empirical methods are sufficient, regarding the complexity of the couplings investigated here, and lead to satisfying results.

Obviously, “optimum” has to be a compromise between cost and time to solution. In order to find a unique configuration we suggest the optimum to have a parallel efficiency higher than 50 % of the cost of the reference configuration,

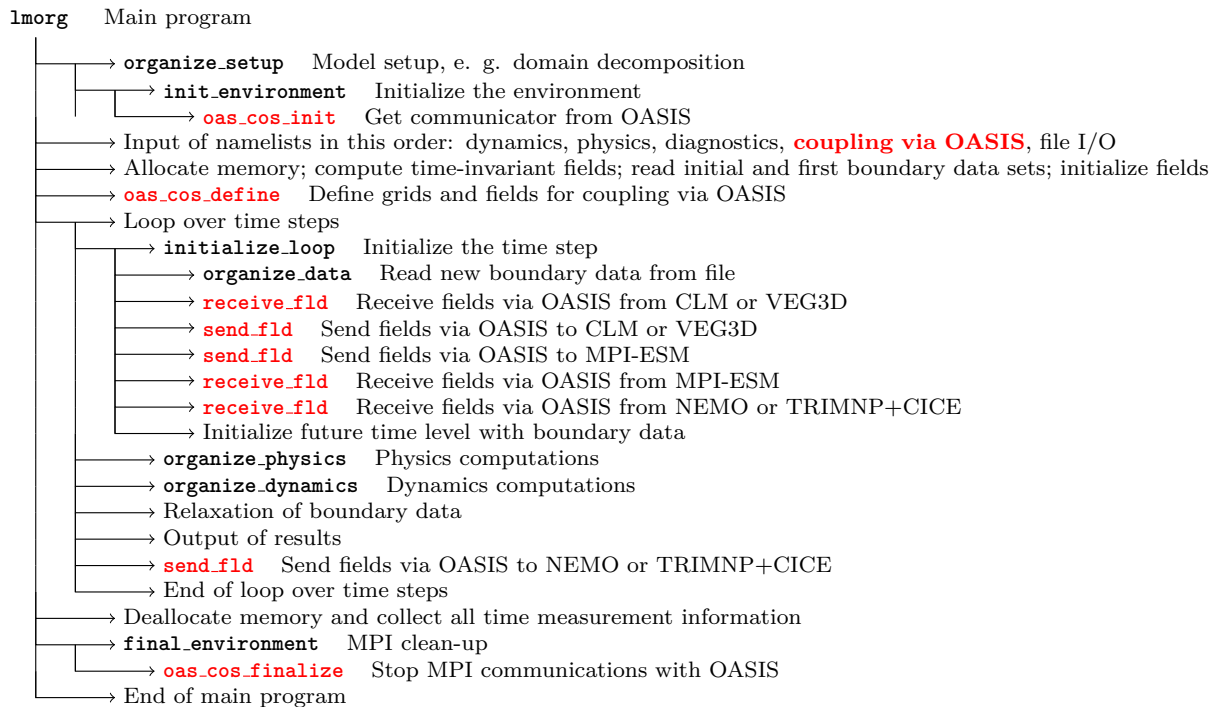


Figure 7. Simplified flow diagram of the main program of the COSMO model in Climate Mode (CCLM), version 4.8_clm19_uoi. The red highlighted parts indicate the locations at which the additional computations necessary for coupling are executed and the calls to the OASIS interface take place. Where applicable, the component models to which the respective calls apply are given.

until which increasing cost can be regarded as still acceptable. In the case of scalability of all *components* and no substantial cost of necessary additional calculations, this guarantees that the coupled-system’s time to solution is only slightly bigger than that of the *component* with the highest cost.

However, such “optimum” configuration depends on the reference configuration. In this study for all couplings the one-node configuration is regarded to have 100 % parallel efficiency.

An additional constraint is sometimes given by the CPU accounting policy of the computing centre, if consumption is measured “per node” and not “per core”. This leads to a restriction of the “optimum” configuration (r_1, r_2, \dots, r_n) of cores r_i for each *component* of the coupled system to those, for which the total number of cores $R = \sum_i r_i$ is a multiplex of the number of cores r_n per node: $R = \#nodes \cdot r_n$.

An exception is the case of very low scalability of a *component* which has a time to solution similar to the time to solution of the coupled model system. In this case an increase in the number of cores results in an increase in cost and in no decrease in time to solution. In such a case the optimum configuration is the one with the lower cost, even if the limit of 50 % parallel efficiency is fulfilled for the configuration with the higher cost.

The strategies of identifying an optimum configuration are different for sequential and concurrent couplings due to the

possible waiting time, which needs to be considered with concurrent couplings.

For sequential couplings (CCLM+CLM, CCLM+VEG3D and CCLM+MPI-ESM) the SMT mode and an alternating distribution of processes (ADP) is used to keep all cores busy at all times. The possible *component*-internal load imbalances, which occurs when parts of the code are not executed in parallel, are neglected. The effect of ADP has been investigated for CCLM+MPI-ESM coupling on one node ($n = 1$) in more detail and the results are presented in Sect. 4.6.

The optimum configuration is found by starting the measuring of the computing time on one node for all *components*, doubling the resources and measuring the computing time again and again as long as all *components*’ parallel efficiencies remain above 50 %. One could decide to stop at a higher parallel efficiency if cost are a limiting factor.

For concurrent couplings (CCLM+NEMO-MED12 and CCLM+TRIMNP+CICE) the SMT mode with non-alternating processes distribution is used aiming to speed up all *components* in comparison to the ST mode and to reduce the inter-node communication.

The optimisation process of a concurrently coupled model system additionally needs to consider minimising the load imbalance between all *components*. For a given total number of cores (cost) used, the time to solution is minimised if all *components* have the same time to solution (no load



Figure 8. As Fig. 7 but for the ECHAM global atmosphere model of MPI-ESM.

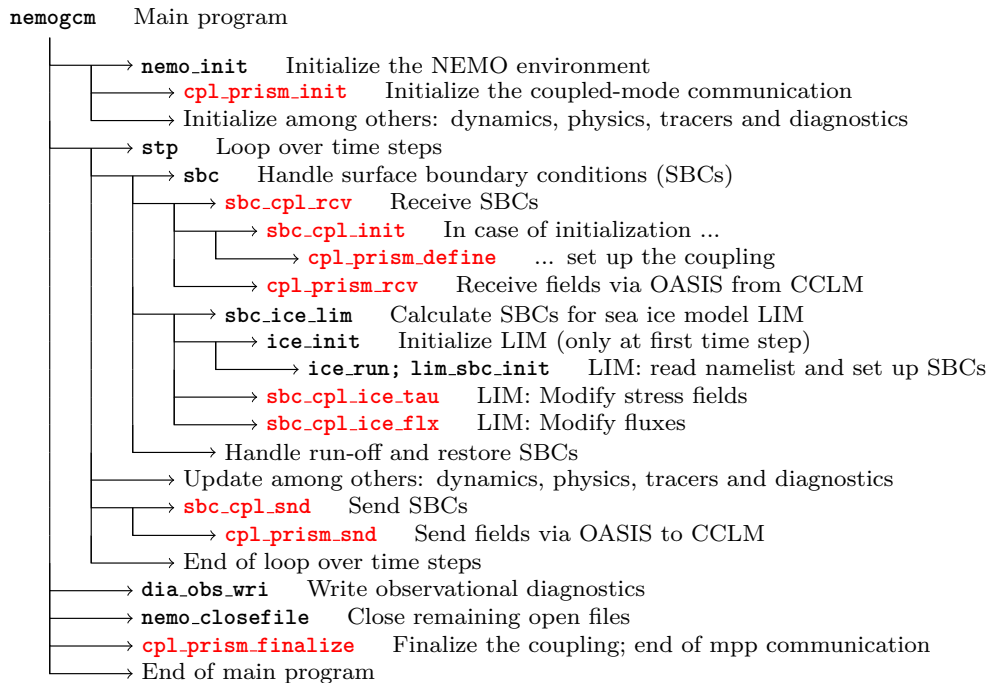


Figure 9. As Fig. 8 but for the NEMO version 3.3 ocean model.

imbalance) and thus no cores are idle during the simulation. Practically speaking, one starts with a first-guess distribution of processes between all *components* on one node, measures each *component's* computing and waiting time and adjusts the process distribution between the *components* if the waiting time of at least one *component* is larger than 5 % of the total runtime. If, finally, the waiting times of all *components* are small, the following chain of action is repeated several times: doubling resources for each *component*, measuring computing times, and adjusting and re-distributing the processes if necessary. If cost is a limiting factor, this is repeated until the cost reaches a pre-defined limit. If cost is not a limiting factor, the procedure should be repeated until the model with the highest time to solution reaches the proposed parallel-efficiency limit of 50 %.

4.4 The optimum configurations

We applied the strategy for finding an optimum configuration described in Sect. 4.3 to the CCLM couplings with a regional ocean (TRIMNP+CICE or NEMO-MED12), an alternative land surface scheme (CLM or VEG3D) or the atmosphere of a global earth system model (MPI-ESM). The optimum configurations found for CCLM_{sa} and all coupled systems are shown in Fig. 6 and in more detail in Table 8. The parallel efficiency used as criterion of finding the optimum configuration is shown in Fig. 5.

The minimum number of cores which should be used is 32 (one node). For sequential coupling an alternating distribution of processes is used and thus one CCLM and one cou-

pled *component* (VEG3D, CLM) process are started on each core. For CCLM+VEG3D and CCLM+CLM the CCLM is more expensive and thus the scalability limit of CCLM determines the optimum configuration. In this case the fair reference for CCLM is CCLM stand-alone (CCLM_{sa}) on 32 cores in single-threading (ST) mode. As shown in Fig. 5 the parallel efficiency of 50 % for COSMO stand-alone in ST mode is reached at 128 cores or four nodes, and thus the 128-core configuration is selected as the optimum.

For concurrent coupling the SMT mode with non-alternating distribution of processes is used, which is more efficient than the alternating SMT and the ST modes. The cores are shared between CCLM and the coupled *components* (NEMO-MED12 and TRIMNP+CICE). For these couplings CCLM is the most expensive *component* as well, and thus the reference for CCLM is CCLM_{sa} on 16 cores (0.5 nodes) in SMT mode. As shown in Fig. 5 the parallel efficiency of 50 % for COSMO stand-alone in SMT mode using 16 cores as a reference is reached at approximately 100 cores. For CCLM+NEMO-MED12 coupling a two-node configuration with 78 cores for CCLM and 50 cores for NEMO-MED12 resulted in an overall decrease in load imbalance to an acceptable 3.1 % of the total cost. Increasing the number of cores beyond 80 for CCLM did not change the time to solution much, because CCLM already approaches the parallel-efficiency limit by using 78 cores. This prevented one from finding the optimum configuration using three nodes. The corresponding NEMO-MED12 measurements at 50 cores are a bit out of scaling as well. This is probably caused by the I/O

Table 8. Analysis of optimum configurations of the coupled systems (CS) given in the table header (see also Fig. 6 and Tables 2 and 3). *seq* refers to sequential and *con* to concurrent couplings. *Thread mode* is either the ST or the SMT mode (see Fig. 2). *APD* indicates whether an alternating processes distribution was used or not. *levels in CCLM* gives the simulated number of levels and *CCLM version* is the CCLM model version used for coupling. Relative *Time to solution (%)* and *Cost (%)* are calculated with respect to the reference, which is the CCLM stand-alone configuration CCLM_{sa} using 64 cores and non-alternating SMT mode. The time to solution includes the time needed for OASIS interpolations. All relative quantities in lines 2.2–2.3 and 3.2–3.3.5 are given in percent of CCLM_{sa} time to solution (line 8) and cost (all others). CS-CCLM_{sa} gives the differences between CS and the optimum CCLM_{sa} configuration. This difference is separated in 5 components of cost: *coupled component* models coupled with CCLM. *OASIS hor: interp.* all horizontal interpolations computed by OASIS. *load imbalance* load imbalance between the concurrently running models. CCLM_{sa,sc} – CCLM_{sa} difference between stand-alone CCLM process mappings used in the particular coupling and for optimum configuration. CCLM – CCLM_{sa,sc} difference between coupled and stand-alone CCLM using process mapping of the coupling

	CCLM stand-alone	CCLM+ CLM	CCLM+ VEG3D	CCLM+ NEMO-MED12	CCLM+ TRIMNP +CICE	CCLM+ ECHAM+ MPIOM	
1.1	Type of coupling	–	seq	seq	con	seq + con	
1.2	Thread mode	SMT	SMT	SMT	SMT	SMT	
1.3	APD used	–	yes	yes	no	yes	
1.4	# nodes	2	4	4	1	1	
1.5	# cores per component	64	128, 128	128, 128	16, 6, 10	32, 28, 4	
1.6	levels in CCLM	45	40	45	40	45	
1.7	CCLM version	4.8	5.0	4.8	4.8	4.8	
2.1	Time to solution (HPSY)	3.6	4.0	3.7	4.0	18.0	34.8
2.2	Time to solution (%)	100.0	111.1	102.8	111.1	450.0	866.7
2.3	CS – CCLM _{sa} (%)	–	11.1	2.8	11.1	350.0	766.7
3.1	CS Cost (CHPSY)	230.4	512.0	473.6	512.0	576.0	1113.6
3.2	CS Cost (%)	100.0	222.2	205.6	222.2	250.0	483.3
3.3	CS – CCLM _{sa} (%)	–	122.2	105.6	122.2	150.0	383.3
3.3.1	coupled component (%)	–	4.3	19.7	79.9	27.2+77.9	261+20.1
3.3.2	OASIS hor: interp. (%)	–	6.3	0.0	0.05	0.76	3.3
3.3.3	load imbalance (%)	–	–	–	6.9	71.5	17.2
3.3.4	CCLM _{sa,sc} – CCLM _{sa} (%)	–	56.2	56.2	16.3	–30.0	4.3
3.3.5	CCLM – CCLM _{sa,sc} (%)	–	55.4	29.7	19.0	2.6	77.4

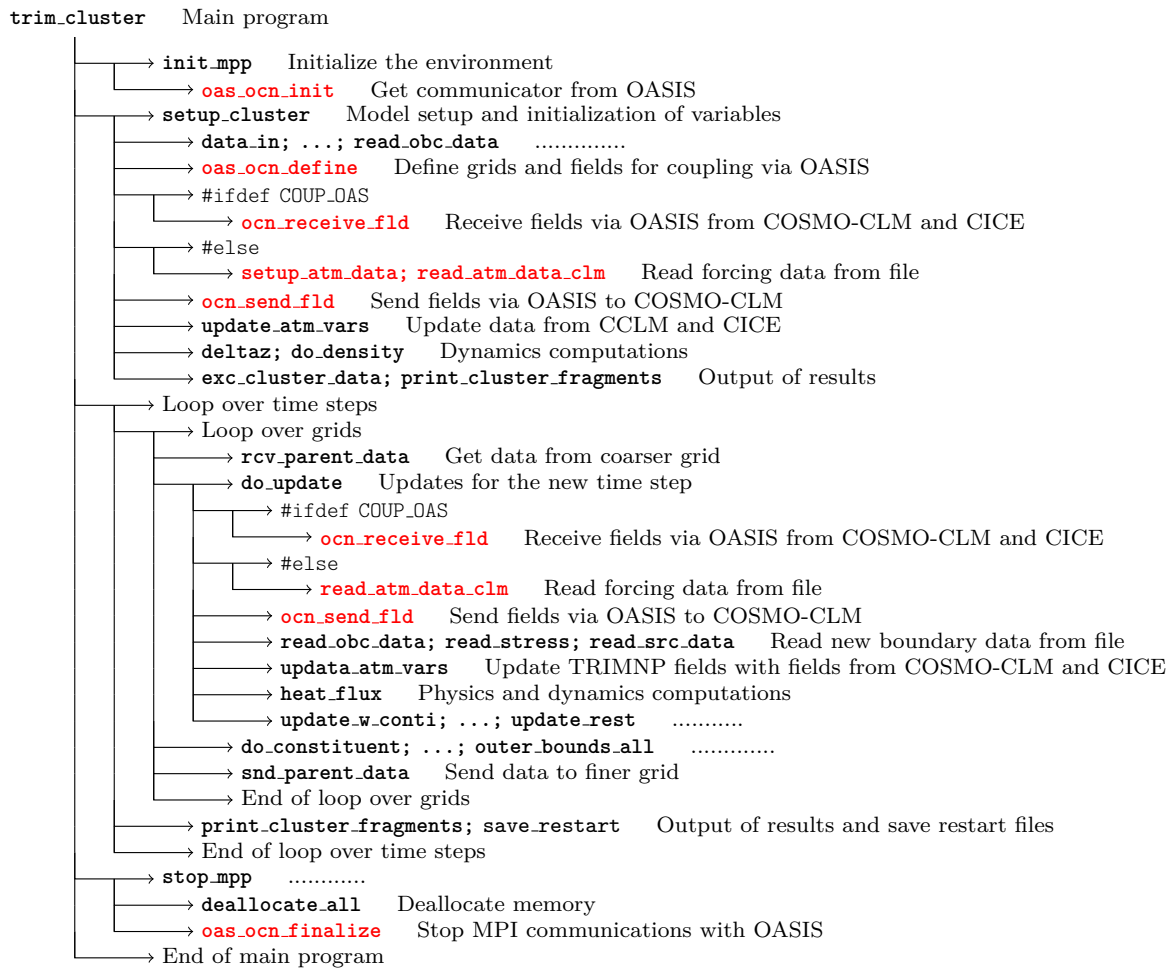


Figure 10. As Fig. 8 but for the TRIMNP ocean model.

which increased for unknown reasons on the machine used between the time of conduction of the first series of simulations and of the optimised simulations.

For CCLM+TRIMNP+CICE no scalability is found for CICE. As shown in Fig. 5 a parallel efficiency smaller than 50 % is found for CICE at approximately 15 cores. As shown in Fig. 3 the time to solution for all core numbers investigated is higher for CICE than for CCLM in SMT mode. Thus, a load imbalance smaller than 5 % can hardly be found using one node. The optimum configuration found is thus a one-node configuration using the CCLM reference configuration (16 cores).

The CCLM+MPI-ESM coupling is a combination of sequential coupling between CCLM and ECHAM and concurrent coupling between ECHAM and the MPIOM ocean model. As shown in Fig. 4 MPIOM is much cheaper than ECHAM and, thus, the coupling is dominated by the sequential coupling between CCLM and ECHAM. As shown in Fig. 3, ECHAM is the most expensive *component* and it exhibits no decrease in time to solution by increasing the

number of cores from 28 to 56, i.e. it exhibits a very low scalability. Thus, as described in the strategy for finding the optimum configuration, even if a parallel efficiency higher than 50 % for up to 64 cores (see Fig. 5) is found, the optimum configuration is the 32-core (one-node) configuration, since no significant reduction of the time to solution can be achieved by further increasing the number of cores.

An analysis of additional cost of coupling requires a definition of a reference. We use the cost of CCLM stand-alone at optimum configuration (CCLM_{sa,OC}). We found the SMT mode with non-alternating distribution of processes and 64 cores to be the optimum configuration for CCLM resulting in a time to solution of 3.6 HPSY and cost of 230.4 CHPSY. As shown in Sect. 4.2, SMT mode with non-alternating processes distribution is the most efficient and the scalability limit is reached at approximately 80 cores in SMT mode due to limited number of grid points used. The double of 64 cores is beyond the scalability limit of this particular model grid.

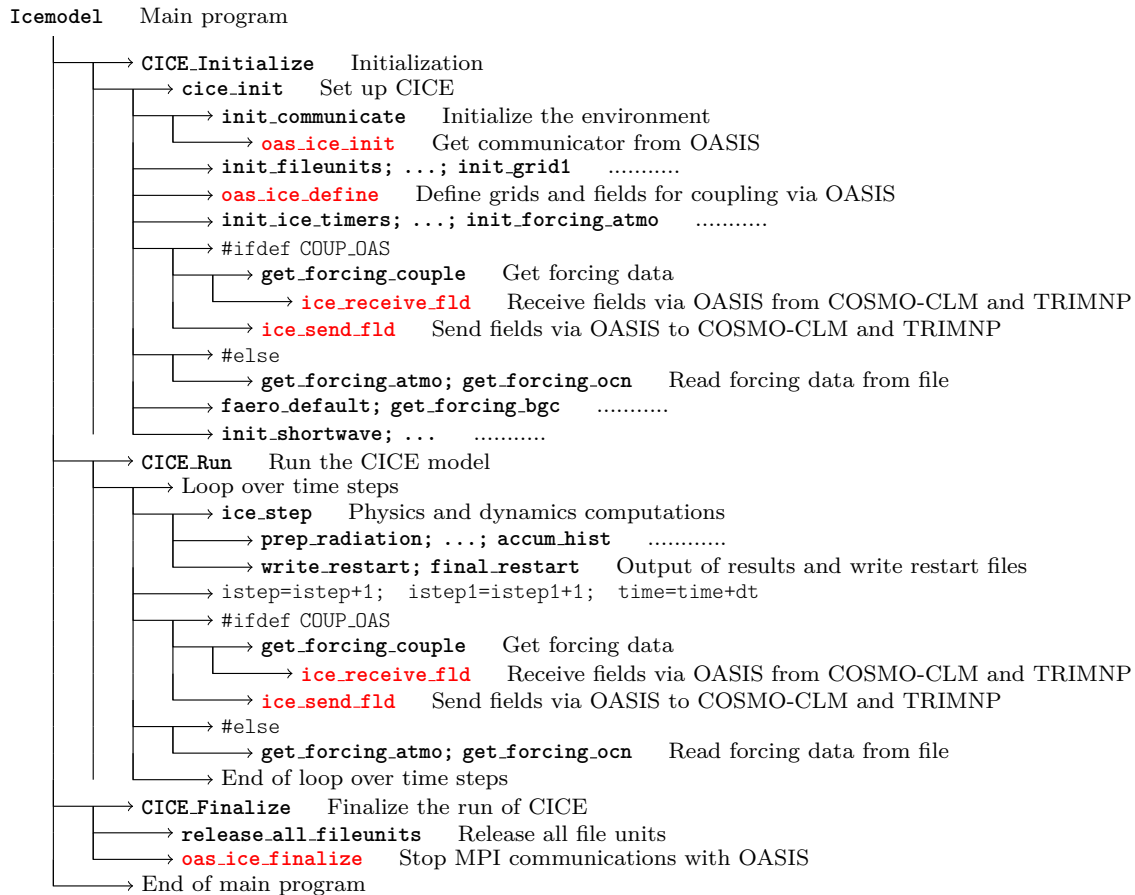


Figure 11. As Fig. 8 but for the CICE sea ice model.

4.5 Extra time and cost

Figure 6 shows the times to solution (vertical axis) and cost (box area) of the *components* of the coupled systems at optimum configurations together with the load imbalance. It exhibits significant differences between the coupled model systems, CCLM_{OC} and CCLM_{sa,OC}. The direct coupling cost of the OASIS3-MCT coupler are not shown. This is due to the fact that they are negligible in comparison with the cost of the coupled models. This is not necessarily the case, in particular when a huge amount of fields is exchanged. The relevant steps to reduce these direct coupling cost are described in Sect. 4.6.

Table 8 gives a summary of an analysis of each optimum configuration (line 3.1 and 3.2) using the opportunities provided by LUCIA and by additional internal measurements of timing. It focuses on the cost analysis of the relative difference between the cost of CS and CCLM_{sa} (line 3.3) and provides its separation into 5 components:

1. *coupled component(s)*: cost of the *component(s)* coupled to CCLM

2. *OASIS hor. interp.*: cost of OASIS horizontal interpolations between the grids and communication between the *components*
3. *load imbalance*: cost of waiting time of the *component* with the shorter time to solution in case of concurrent coupling
4. CCLM_{sa,sc} – CCLM_{sa}: cost difference due to usage of another CCLM process mapping (alternating/non alternating SMT or ST mode and a different number of cores).
5. CCLM–CCLM_{sa,sc}: extra cost of CCLM in coupled mode. It contains additional computations in the coupling interface, differences due to different model versions (as in CCLM+CLM), differences in performance of CCLM by using the core and memory together with other *components* and uncertainties of measurement due to variability in performance of the computing system.

The optimum configurations of sequential couplings CCLM+CLM and CCLM+VEG3D can be identified as the configurations with the smallest extra time (11.1 and 2.8 %)

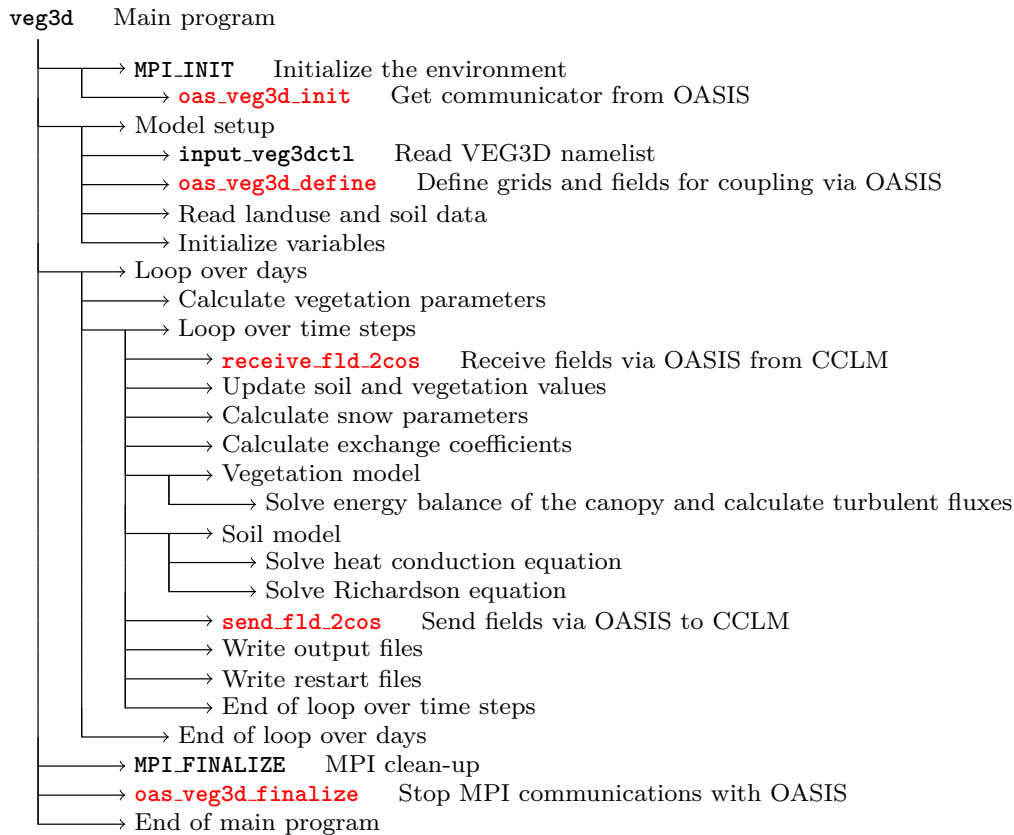


Figure 12. As Fig. 8 but for the VEG3D soil–vegetation model.

and extra cost (122.2 and 105.6 %) respectively (see line 3.3 in Table 8). They use 128 cores for each *component* in SMT mode with alternating processes distribution (line 1.5 in Table 8). A substantial part (56.2 %) of the extra cost in CCLM+CLM and CCLM+VEG3D can be explained by a different mapping of CCLM (line 3.3.4 in Table 8). The 128 CCLM processes of our reference optimum configuration are mapped on 64 cores (CCLM_{sa,OC} mapping). The 128 CCLM processes in optimum configuration of the coupled mode are mapped on 128 cores (CCLM_{OC} mapping) but, in each core, memory, bandwidth and disk access are shared with a land surface model process. These higher cost can be regarded as the price for keeping the time to solution only marginally bigger than that of CCLM_{sa,OC} (see line 2.1 in Table 8) and avoiding of 50 % idle time in sequential mode. The replacement of the CCLM model component TERRA (1 % of CCLM_{sa} cost) by a land surface *component* is the second important part of extra cost with 4.3 % for CLM and 19.3 % for VEG3D (line 3.3.1 in Table 8). The 5 times higher cost of VEG3D in comparison with CCLM is due to low scalability of VEG3D (see Fig. 3). The OASIS horizontal interpolations (line 3.3.2 in Table 8) produce 6.3 % extra cost in CCLM+CLM. No extra cost occurs due to horizontal interpolation in CCLM+VEG3D coupling, since the same grid

is used in CCLM and VEG3D, and due to load imbalance, which is obsolete in sequential coupling. The remaining extra cost are assumed to be the cost difference between the coupled CCLM and CCLM_{sa,OC}. They are found to be 55.4 and 29.7 % for CLM and VEG3D coupling respectively. A substantial part of the relatively high extra cost of CCLM in coupled mode of CCLM+CLM can be explained by higher cost of `cosmo_5.0_clm1`, used in CCLM+CLM, in comparison with `cosmo_4.8_clm19`, used in all other couplings (see line 1.7 in Table 8). CCLM_{sa} performance measurements with both versions (but on a different machine than *Blizzard*) reveal a `cosmo_5.0_clm1` time to solution 45 % longer than for `cosmo_4.8_clm19`.

The concurrent coupling of CCLM with NEMO for Mediterranean Sea (CCLM+NEMO-MED12) is as expensive as CCLM+CLM and exhibits at the systems' optimum configuration 4.0 HPSY time to solution and 512.0 CHPSY cost (line 3.1 and 3.2 in Table 8). The extra cost of 122 % are dominated by the cost of the coupled *component*, which are 79.9 % of the CCLM_{sa,OC} cost. The second important cost of 16.3 % can be explained by the higher number of cores used by CCLM_{OC} than CCLM_{sa,OC} at optimum configurations (line 1.5 and 3.3.4 in Table 8). The load imbalance of 6.9 % of CCLM_{sa,OC} is below the intended limit of 5 % of

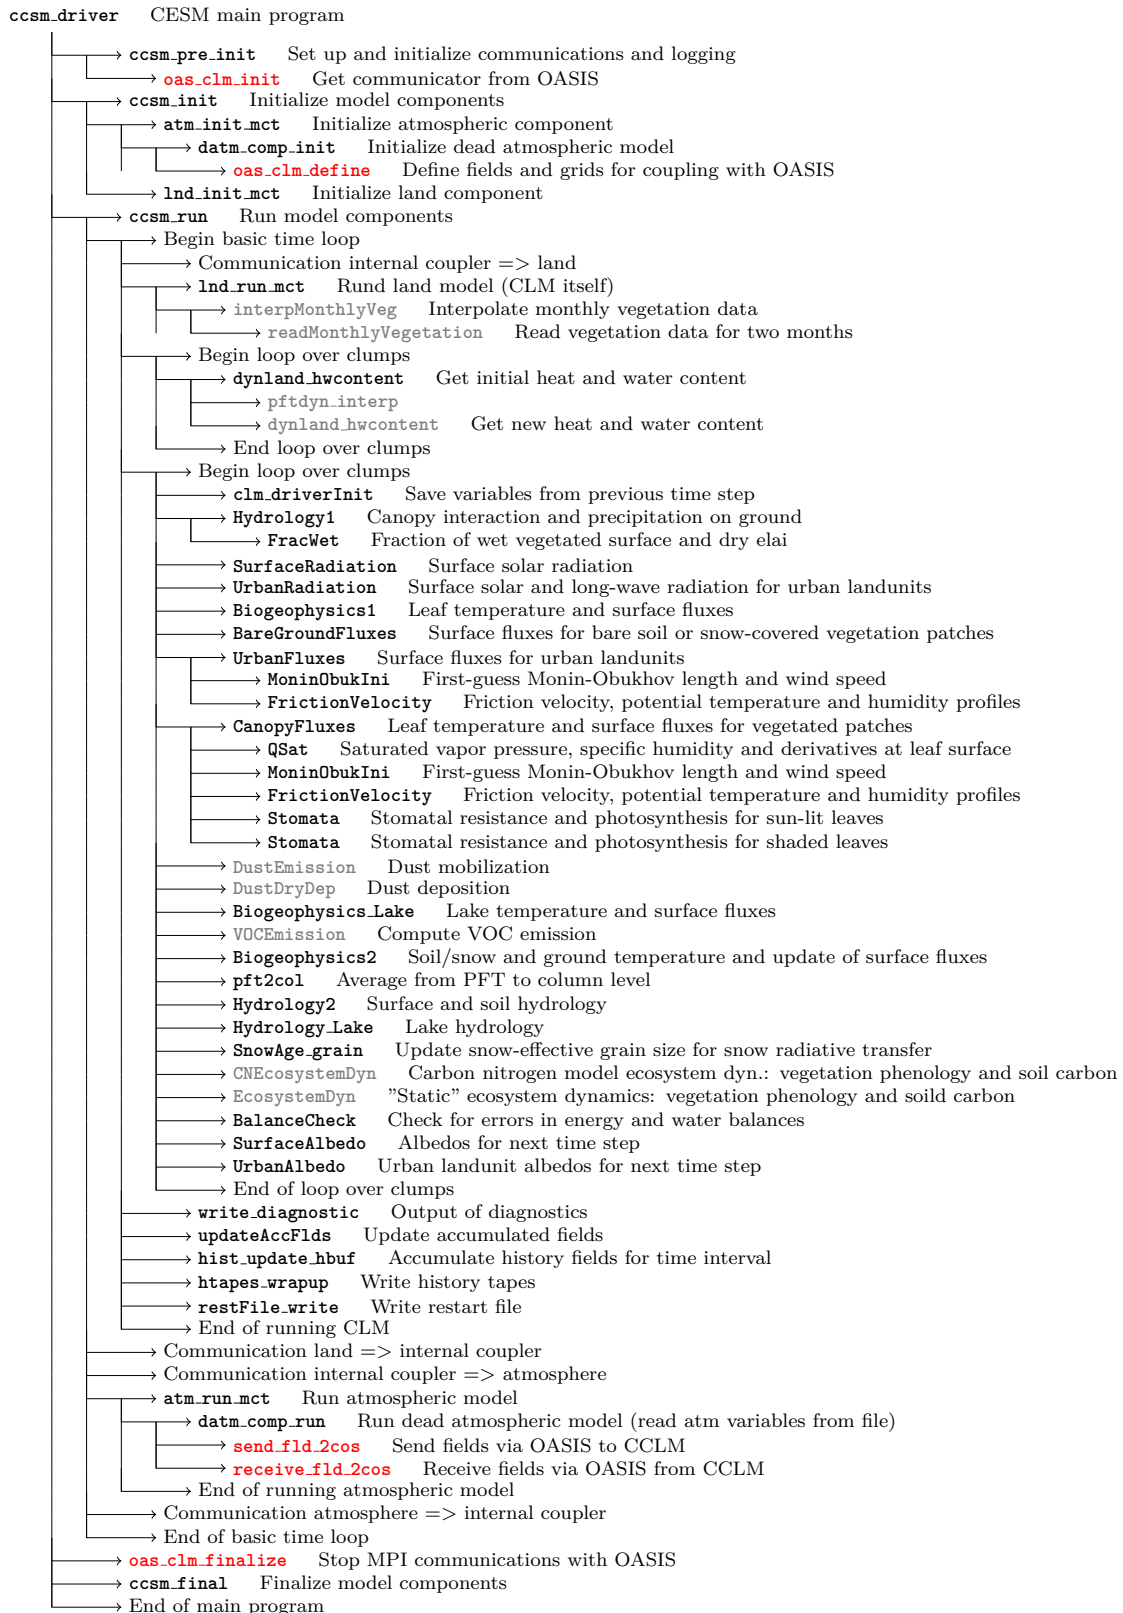


Figure 13. As Fig. 8 but for the Community Land Model (CLM). The grey highlighted routines are optional.

the cost of the coupled system. The extra cost of CCLM_{OC} of 19 % are smaller than for the land surface scheme couplings.

The optimum configuration of the coupling with TRIMNP+CICE for the North and Baltic seas (CCLM+TRIMNP+CICE) has a time to solution of 18 HPSY and a cost of 576 CHPSY. This is 3.5 times longer than CCLM_{sa,OC} due to lack of scalability of the CICE sea ice model and 1.5 times more expensive than CCLM_{sa,OC} (lines 2.3 and 3.3 of Table 8). The dominating components of the extra cost are the costs of the *components* coupled with CCLM. The TRIMNP ocean model cost 27.2 % and the CICE ice model 77.9 % of the CCLM_{sa,OC} cost. The second important component of the extra cost is the load imbalance. Due to CICE's low speed-up and the fact that the time to solution of CICE is generally significantly higher than that of TRIMNP and CCLM, there is no common speed of all three *components*. The load imbalance at optimum configuration is 71.5 % of the CCLM_{sa,OC} cost. However, a further decrease in CCLM and TRIMNP cores reduces the load imbalance but not the cost of coupling, since the time to solution of CICE decreases very slowly with the number of processors. The CCLM mapping used in the coupled system is 30 % cheaper than CCLM_{sa,OC}. This reduces the extra cost without increasing the time to solution. The OASIS3-MCT interpolation cost of 0.8 % of the CCLM_{sa,OC} cost is negligible. The extra cost of CCLM in coupled mode is found to be 2.6 % of the CCLM_{sa,OC} cost only.

The most complex (see the definition in Balaji et al., 2017) and most expensive coupling presented here is the sequential coupling of CCLM with the MPI-ESM global earth system model. The model components directly coupled are the non-hydrostatic atmosphere model of CCLM and the ECHAM hydrostatic atmosphere model, which is a *component* of MPI-ESM. The complexity of the coupling is increased by an additional MPI-ESM internal concurrent coupling via OASIS3-MCT between the ECHAM global atmosphere model and the MPIOM global ocean model. From the point of view of OASIS, the CCLM+MPI-ESM coupling is a CCLM+ECHAM+MPIOM coupling. In this list ECHAM has a similar complexity to CCLM but on a global scale. At optimum configuration the time to solution of CCLM+ECHAM+MPIOM is 34.8 HPSY and the cost is 1113.6 CHPSY (lines 2.1 and 3.3.1 in Table 8). It takes 7.67 times longer than CCLM_{sa,OC} due to lack of scalability of ECHAM in coupled mode. A model-internal timing measurement revealed no scalability and high cost of a necessary additional computation of horizontal derivatives executed in the ECHAM coupling interface using a spline method. Connected herewith, the cost of ECHAM, which is 261 % of the CCLM_{sa,OC} cost, is the major part of the total extra cost of 383 %. In stand-alone mode the cost of MPI-ESM at optimum processor configuration (one node) is 64% of the CCLM_{sa,OC} cost, and thus 197% of CCLM_{sa,OC} is the extra costs of coupling of MPI-ESM. The second *component* MPIOM cost 20.1 % of CCLM_{sa,OC}. The load imbalance

using 4 cores for MPIOM and 28 for ECHAM is 17.2 %. However, a further reduction of the number of MPIOM cores (and increase in the number of ECHAM cores) can reduce the load imbalance but not the time to solution and cost of MPI-ESM. The cost of CCLM stand-alone using the same mapping (CCLM_{sa,sc}) as for CCLM coupled to MPI-ESM is 4.3 % higher than the cost of CCLM_{sa,OC} (line 3.3.4 in Table 8). Interestingly, the cost of OASIS horizontal interpolations is 3.3 % only. This achievement is discussed in more detail in the next section. Finally, the extra cost of CCLM in the coupled mode of CCLM+ECHAM+MPIOM is 77.4 %. They are the highest of all couplings. Additional internal measurements allowed one to identify additional computations in the CCLM coupling interface as being responsible for a substantial part of this cost. The vertical spline interpolation of the 3-D fields exchanged between the models was found to consume 51.8 % of the CCLM_{sa,OC} cost, which is 2/3 of the extra cost of CCLM_{OC}.

Interestingly, a direct comparison of complexity and grid point number G (see the definition in Balaji et al., 2017) given in Table 3 with the extra cost of coupling given in Table 8 shows that the couplings with short time to solution and lowest extra cost are those of low complexity. On the other hand, the most expensive coupling with the longest time to solution is that of the highest complexity and with the largest number of grid points.

4.6 Coupling cost reduction

The CCLM+MPI-ESM coupling is one of the most intensive couplings that has up to now been realised with OASIS3(-MCT) in terms of number of coupling fields and coupling time steps: 450 2-D fields are exchanged every ECHAM coupling time step, that is, every 10 simulated minutes (see Sect. 3.2). Most of these 2-D fields are levels of 3-D atmospheric fields. We show in this section that a conscious choice of coupling software and computing platform features can have a significant impact on time to solution and cost.

To make the CCLM+MPI-ESM coupling more efficient, all levels of a 3-D variable are sent and received in a single MPI message using the concept of *pseudo-3-D coupling*, as described in Sect. 3.1.3, thus reducing the number of sent and received fields (see Table 4). The change from 2-D to pseudo-3-D coupling leads to a decrease in the cost of the coupled system running on 32 cores by 3.7 % of the coupled system, which corresponds to 25 % of the CCLM_{sa,OC} cost. At the same time the cost of the OASIS3-MCT interpolations is reduced by 76 %, which corresponds to an additional reduction of cost by 12 % of the CCLM_{sa,OC} cost. The total reduction of cost by exchanging one 3-D field is 34 % of the CCLM_{sa,OC} cost.

The second optimisation step is a change in mapping of running processes on cores. Instead of non-alternating, an alternating distribution of processes of sequentially running *components* is used such that on each core one process of

each component model is started. This reduced the time to solution and cost of the coupled system running on 32 cores and using *pseudo-3-D coupling* by 35.8 %, which is 226 % of $CCLM_{sa,OC}$. The expected reduction of time to solution is 25.5 %. It is a combined effect of increasing the time to solution by changing the mapping from 16 cores in SMT mode to 32 cores in ST mode (here $CCLM_{sa}$ measurements are used) and of reducing it by making 50 % of the idle time of the cores in sequential coupling available for computations. A separate investigation of CCLM, ECHAM and MPIOM time to solution and cost revealed strong deviations from the expectation for the individual *components*. A higher relative decrease of 46.4 % was found for ECHAM due to a dramatic reduction of the time to solution of the inefficient calculation of the derivatives (needed for coupling with CCLM only) by one process. The CCLM's time to solution in coupled mode was reduced by 9.2 % only. Additional internal measurements of CCLM revealed that the discrepancy of 16.3 % originates from reduced scalability of some subroutines of CCLM in coupled mode, which is probably related to sharing of memory between CCLM and ECHAM when running on the same core in coupled mode. In particular the CCLM interface and the physics computations show almost no speed-up.

The combined effect of usage of 3-D-field exchange and of an alternating process distribution lead to an overall reduction of the total time to solution and cost of the coupled system CCLM+MPI-ESM by 39 %, which corresponds to 261 % of the $CCLM_{sa,OC}$ cost.

5 Conclusions

We presented a prototype of a regional climate system model based on the non-hydrostatic, limited-area COSMO model in CLimate Mode (CCLM) coupled to regional ocean, land surface and global earth system models using the fully parallelised OASIS3-MCT coupler. We showed how particularities of regional coupling can be solved using the features of OASIS3-MCT and how an optimum configuration of computational resources can be found. Finally we analysed the extra cost of coupling and identified the unavoidable cost and the bottlenecks.

We showed that the measures *time to solution*, *cost* and *parallel efficiency* of each *component* and of the coupled system, provided by OASIS3-MCT tool LUCIA, are sufficient to find an optimum processor configuration for sequential, concurrent and mixed regional coupling with CCLM. Thus, it could be applicable to other regional coupled model systems as well.

The analysis of the extra cost of individual couplings at optimum configuration, presented here, was found to be a useful step of development of a regional climate system model. The results reveal that the regional climate system model at optimum configuration can have a similar time to solution as the RCM, but at extra costs which are approximately the

cost of the RCM for each coupling if (i) scalability problems can be avoided and (ii) the extra cost of additional computations can be kept small. This is found for concurrent and sequential coupling layouts for different reasons (see Table 8 for details).

The prototype of the regional climate system model consists of two-way couplings between the COSMO model in Climate Mode (COSMO-CLM or CCLM), which is an atmosphere–land model, two alternative land surface schemes (VEG3D, CLM) replacing TERRA, a regional ocean model (NEMO-MED12) for the Mediterranean Sea and two alternative regional ocean models (NEMO-NORDIC, TRIMNP+CICE) for the North and Baltic seas and the MPI-ESM earth system model. A unified OASIS3-MCT interface (UOI) was developed and successfully applied for all couplings. All couplings are organised in a least intrusive way such that the modifications of all *components* of the coupled systems are mainly limited to the call of two subroutines receiving and sending the exchanged fields (as shown in Figs. 7 to 13) and performing the necessary additional computations.

The features of the fully parallelised OASIS3-MCT coupler have been used to address the particularities the couplings investigated. We presented solutions for (i) using the OASIS coupling library for an exchange of data between different domains, (ii) for multiple usage of the MCT library (in different couplings), (iii) an efficient exchange of more than 450 2-D fields and (iv) usage of higher order (than linear) interpolation methods.

A series of simulations has been conducted with an aim to analyse the computational performance of the couplings. The CORDEX-EU grid configuration of CCLM on a common computing system (*Blizzard* at DKRZ) has been used in order to keep the results comparable.

The LUCIA tool of OASIS3-MCT has been used to measure the computing time used by each *component* and by the coupler for communication and horizontal interpolation in dependence on the computing resources used. This allows an estimation of the computing time for intermediate computing resources and thus determination of an optimum configuration based on a limited number of measurements. Furthermore, the scaling of each *component* of the coupled system can be analysed and compared with that of the model in stand-alone mode. Thus, the extra cost of coupling is measured and the origins of the relevant extra cost can be analysed.

The scaling of CCLM was found to be very similar in stand-alone and in coupled mode. The weaker scaling, which occurred in some configurations, was found to originate from additional computations which do not scale but are necessary for coupling. In some cases the model physics or the I/O routines exhibited a weaker scaling, most probably due to limited memory.

The results confirm that parallel efficiency is decreasing substantially if the number of grid points per core is below

80. For the configuration used (132×129 grid points), this limits the number of cores, which can be used efficiently to 80 in SMT mode and 160 in ST mode.

For the first time a sequential coupling of approximately 450 2-D fields using the OASIS3-MCT parallelised coupler was investigated. It was shown that the direct costs of coupling by OASIS3-MCT (interpolation and communication) are negligible in comparison with the cost of the coupled atmosphere–atmosphere model system. We showed that the exchange of one (pseudo-)3-D field instead of many 2-D fields reduces the cost of communication drastically.

The idling of cores due to sequential coupling could be avoided by a dedicated launching of one process of each of the two sequentially running models on each core making use of the multi-threading mode available on the machine *Blizzard*. This feature is available on other machines as well.

A strategy for finding an optimum configuration was developed. Optimum configurations were identified for all investigated couplings considering three aspects of climate modelling performance: time to solution, cost and parallel efficiency. The optimum configuration of a coupled system, which involves a *component* not scaling well with available resources, is suggested to be used at minimum cost, if time to solution cannot be decreased significantly. This is the case for CCLM+MPI-ESM and CCLM+TRIMNP+CICE couplings. An exception is the CCLM+VEG3D coupling. VEG3D was found to have a weak scaling but a small workload in comparison to CCLM. Thus, it has a negligible impact on the performance of the coupled system.

The analysis of the extra cost of coupling at optimum configuration using LUCIA and CCLM stand-alone performance measurements allowed one to distinguish five components (lines 3.3.1–3.3.5 in Table 8): (i) cost of coupled *components*, (ii) OASIS horizontal interpolation and communication (direct coupling cost), (iii) load imbalance (if concurrently coupled), (iv) additional/minor cost of different usage of processors by CCLM in coupled and stand-alone mode and (v) residual cost including i.a. CCLM additional computations and extraordinary behaviour of the *components* in coupled mode due to e.g. sharing of the memory. This allowed one to identify the unavoidable cost and the bottlenecks of each coupling.

The analysis of the extra cost of coupling in comparison with CCLM stand-alone (see Table 8) at optimum processor configuration can be summarised as follows.

- The land surface scheme (CCLM+CLM) exhibits the same speed and 122% extra cost and it can hardly be further improved. Probably up to 20% extra cost is avoidable. Approximately 100% extra cost is unavoidable: (1) extra cost of keeping the speed of the coupled system high by using a higher number of cores, (2) the need to use the single-threading mode to avoid idle time of cores in sequential coupling and (3) the

higher cost of `cosmo_5.0_clm1` in comparison with `cosmo_4.8_clm19`.

- The soil and vegetation model (CCLM+VEG3D) exhibits the same speed and 105.6% extra cost, and it can hardly be further improved as well. Probably up to 50% extra cost is avoidable. These are (1) the higher cost of VEG3D in comparison with TERRA and (2) of CCLM in coupled mode. Approximately 56% extra cost (same as for CCLM+CLM) is unavoidable: (1) extra cost of keeping the speed of the coupled system high by using a higher number of cores and (2) the need to use the single-threading mode to avoid idle time of cores in sequential coupling.
- The Mediterranean ocean model (CCLM+NEMO-MED12) exhibits same speed and 122% extra cost. It hardly can be further improved as well. Probably 20% extra cost of CCLM in coupled mode are avoidable. Approximately 100% extra cost are unavoidable: (1) cost of NEMO-MED12, (2) extra cost of keeping the speed of the coupled system high by using a higher number of cores and (3) small extra cost of load imbalance due to concurrent coupling.
- The North and Baltic seas model (CCLM+TRIMNP+CICE) exhibits a much longer time to solution (+350%) and 150% extra cost. The longer time to solution and 70% extra cost of load imbalance are due to the lack of scalability of the CICE model.
- The global earth system model (CCLM+MPI-ESM) exhibits a very long time to solution (+766%) and high extra cost (+383%). The longer time to solution and approximately 235% extra cost are due to a lack of scalability of the ECHAM model. Additionally, 77% extra cost is due to vertical interpolation of 3-D fields in CCLM.

We found bottlenecks of coupling in the CCLM+TRIMNP+CCLM and CCLM+MPI-ESM couplings.

A direct comparison between NEMO and TRIMNP+CICE is not possible because the cost of NEMO-NORDIC has not been measured on the same machine and for the same configuration. The lower cost of TRIMNP in comparison with NEMO-MED12 can be more than explained by the difference in the number of grid points and time steps. The surface of the North and Baltic seas is approximately half of the Mediterranean surface. Furthermore, approximately a double horizontal resolution is used in the NEMO-MED12 coupling, resulting in a factor of 16.

Code availability. The source code availability is described in Appendix A.

Appendix A: Source code availability

The COSMO model in Climate Mode (COSMO-CLM or CCLM) is an atmosphere model coupled to the TERRA soil-vegetation model. Other regional processes in the climate system like ocean and ice sheet dynamics, plant responses, aerosol–cloud interaction, and the feedback to the GCM driving the RCM are made available by coupling COSMO-CLM via OASIS3-MCT with other models.

The CCLM model source code is freely available for scientific usage by members of the CLM-Community. The CLM-Community (www.clm-community.eu) is a network of scientists who accept the CLM-Community agreement. To become a member, please contact the CLM-Community coordination office at DWD, Germany (clm-coordination@dwd.de).

The current recommended version of CCLM is COSMO_131108_5.0_clm⁹. It comes together with a recommendation for the configurations for the European domain.

The development of fully coupled CCLM is an ongoing research project within the CLM-Community. The unified OASIS3-MCT coupling interface, necessary to ensure coupling of CCLM with any other component, is available by contacting one of the authors and will be part of a future official CCLM version. All other components, including OASIS3-MCT interface for the component, are available by contacting the authors. The OASIS3-MCT coupling library can be downloaded at <https://verc.enes.org/oasis/>.

The CCLM+MPIESM two-way coupled system was developed at BTU Cottbus and FU Berlin. Please contact Andreas Will (will@b-tu.de) for more information about the source codes.

The Community Land Model (CLM) is freely available as part of the Community Earth System Model (CESM) package and can be obtained through a SVN server after registration. Registration and access: <http://www.cesm.ucar.edu/models/cesm1.2>.

For information about a possible usage of VEG3D, please contact Marcus Breil at KIT (marcus.breil@kit.edu).

The Nucleus for European Modeling of the Ocean (NEMO) is a community model. It can be adapted for regional and global applications. To access NEMO, please visit the webpage <http://www.nemo-ocean.eu/> and register there with signing the CeCILL licence agreement. Please contact Jennifer Brauch (jennifer.brauch@dwd.de) to get more information about the employed NEMO configurations.

For information about the modified version of TRIMNP, please contact Ha Hagemann at HZG (ha.hagemann@hzg.de). The CICE sea ice model version 5.0 is developed at the Los Alamos National Laboratory, USA (<http://oceans11.lanl.gov/trac/CICE/wiki>). Please contact Ha Hagemann at HZG

for more details to set up CICE for the North Sea and Baltic Sea.

Appendix B: Model time step organisation

In the following, the time step organisation within the models coupled is described. This aims at providing a basis of understanding of the coupling between the models.

B1 COSMO model in Climate Mode (COSMO-CLM or CCLM)

Figure 7 gives an overview of the model initialisation procedure, of the *Runge–Kutta* time step loop and of final calculations. The subroutines that contain all modifications of the model necessary for coupling are highlighted in red.

At the beginning ($t = t_m$) of the CCLM time step $(\Delta t)_c$ in `initialize_loop` the lateral, top and the ocean surface boundary conditions are updated. In `organize_data` the future boundary conditions at $t_f \geq t_m + \Delta t_c$ on the COSMO grid are read from a file (if necessary). As next `send_fld` and `receive_fld` routines are executed sending the CCLM fields to or receiving them from OASIS3-MCT in coupled simulations (if necessary). The details including the positioning of the `send_fld` routines are explained in Sect. 3.2 to 3.5.

At the end of the `initialize_loop` routine the model variables available at previous $t_p \leq t_m$ and next time $t_m < t_f$ of the boundary update are interpolated linearly in time (if necessary) and used to initialise the bound lines of the CCLM model grid at the next model time level $t_m + (\Delta t)_c$ for the variables u and v wind, temperature and pressure deviation from a reference atmosphere profile, specific humidity, cloud liquid and ice water content, surface temperature over water surfaces and – in the bound lines only – surface specific humidity, snow surface temperature and surface snow amount.

In `organize_physics` all tendencies due to physical parameterisations between the current t_m and the next time level $t_m + (\Delta t)_c$ are computed in dependence on the model variables at time t_m . Thus, they are not part of the *Runge–Kutta* time stepping. In `organize_dynamics` the terms of the Euler equation are computed.

The solution at the next time level $t_m + (\Delta t)_c$ is relaxed to the solution prescribed at the boundaries using an exponential function for the lateral boundary relaxation and a cosine function for the top boundary Rayleigh damping (Doms and Baldauf, 2015). At the lower boundary a slip boundary condition is used together with a boundary layer parameterisation scheme (Doms et al., 2011).

B2 MPI-ESM

Figure 8 gives an overview of the ECHAM leapfrog time step (see DKRZ, 1993 for details). Here the fields at time level

⁹Status of October 2016

t_{n+1} are computed by updating the time level t_{n-1} using tendencies computed at time level t_n .

After model initialisation in `initialize` and `init_memory` and reading of initial conditions in `iorestart` or `ioinitial` the time step begins in `stepon` by reading the boundary conditions for the models coupled in `bc_list_read` if necessary, in this case for the MPIOM ocean model. In `couple_get_o2a` the fields sent by MPIOM to ECHAM (SSTs, SICs) for time level t_n are received if necessary.

The time loop (`stepon`) has three main parts. It begins with the computations in spectral space, followed by grid space and spectral-space computations. In `scanl` the spatial derivatives (`sym2`, `ewd`, `fft1`) are computed for time level t_n in Fourier space followed by the transformation into grid-space variables on the lon/lat grid. Now, the computations needed for two-way coupling with CCLM (`twc`) are done for time level t_n variables followed by advection (`dyn`, `ldo_advection`) at t_n , the second part of the time filtering of the variables at time t_n (`tf2`), the calculation of the advection tendencies and update of fields for t_{n+1} (`ldo_advection`). Now, the first part of the time filtering of the time level t_{n+1} (`tf1`) is done followed by the computation of physical tendencies at t_n (`physc`). The remaining spectral-space computations in `scanl` begin with the reverse Fourier transformation (`fft2`).

B3 NEMO-MED12

In Fig. 9 the flow diagram of NEMO 3.3 is shown. At the beginning the mpp communication is initialised by `cpl_prism_init`. This is followed by the general initialisation of the NEMO model. All OASIS3-MCT fields are defined inside the time loop, when `sbc` (surface boundary conditions) is called the first time. In `sbc_cpl_init` the variables which are sent and received are defined over ocean and sea ice if applicable. At the end of `sbc_cpl_init` the grid is initialised on which the fields are exchanged. In `cpl_prism_rcv` NEMO receives from OASIS3-MCT the fields necessary as initial and upper boundary conditions. NEMO-MED12 and NEMO-NORDIC follow the time lag procedure of OASIS3-MCT appropriate for concurrent coupling. NEMO receives the restart files provided by OASIS3-MCT containing the CCLM fields at restart time. At all following coupling times the fields received are not the CCLM fields at the coupling time but at a previous time, which is the coupling time minus a specified time lag. If a sea ice model is used, the fluxes from CCLM to NEMO have to be modified over surfaces containing sea ice. Hereafter, NEMO is integrated forward in time. At the end of the time loop in `sbc_cpl_snd` the surface boundary conditions are sent to CCLM. After the time loop integration the mpp communication is finished in `cpl_prism_finalize`.

B4 TRIMNP+CICE

Figures 10 and 11 show the flow diagrams of TRIMNP and CICE in which red parts are modifications of the models and blue parts are additional computations necessary for coupling. First, initialisation is done by calling `init_mpp` and `cice_init` in TRIMNP and CICE respectively. In `cice_init`, the model configuration and the initial values of variables are set up for CICE, while for TRIMNP `setup_cluster` is used for the same purpose. In both models the receiving (`ocn_receive_fld`, `ice_receive_fld`) and sending (`ocn_send_fld`, `ice_send_fld`) subroutines are used in the first time step ($t=0$) prior to the time loop to provide the initial forcing. The time loop of TRIMNP covers a grid loop in which several grids at higher resolutions are potentially *one-way* nested for specific sub-regions with rather complex bathymetry, e.g. Kattegat of the North Sea. Note that for the coupling, only the first/main grid is applied. The grid loop begins with `rcv_parent_data` that sends data from the coarser grid to the nested grid. Then, `do_update` updates the forcing data passed from CCLM and CICE as well as the lateral boundary data are read from files. After updating, the physics and dynamics computations are mainly done in `heat_flux`, `turbo_adv`, `turbo_gotm`, `do_constituent`, `do_explicit` and `do_implicit`. At the end of the grid loop, the main grid sends data to the finer grid by calling `snd_parent_data` if necessary. At the end of each time step, output and restart data are written to files. Eventually, `stop_mpp` is called at the end of the main program to de-allocate the memory of all variables and finalise the program.

The time loop of CICE has two main parts. In the first part `ice_step`, physical, dynamical and thermo-dynamical processes of the time step $t=t_n$ are mainly computed in `step_therm1`, `step_therm2`, `step_radiation`, `biogeochemistry` and `step_dynamics`, followed by `write_restart` and `final_restart` for writing the output and restart files. Then, the time step is increased to a new time step $t=t_{n+1}$, followed by an update of forcing data from CCLM and TRIMNP via `ice_receive_fld` if necessary and a sending of fields to CCLM and TRIMNP via `ice_send_fld`. At the end of the time loop, all file units are released in `release_all_fileunits` and `oas_ice_finalize` concludes the main program.

B5 VEG3D

Figure 12 shows the flow diagram of VEG3D for the coupled system. In a first step the `oas_veg3d_init` sub-routine is called in order to initialise the MPI communication for the coupling. Afterwards, the model set-up is specified by reading the VEG3D namelist and by loading external landuse and soil datasets. The definition of the grid and the coupling fields is done in `oas_veg3d_define`.

The main program includes two time loops. In the first time loop vegetation parameters are calculated for every simulated day. In the second loop (over the model time steps) the coupling fields from CCLM are received via OASIS3-MCT in `receive_fld_2cos` at every coupling time step. Using these updated fields the energy balance of the canopy for the current time level t_n is solved iteratively and based on this the latent and sensible heat fluxes are calculated. The heat conduction and the Richardson equation for the time level t_{n+1} are solved by a semi-implicit Crank–Nicholson method. After these calculations the simulated coupling fields from VEG3D are sent to CCLM in `send_fld_2cos`. At the end, output and restart files are written for selected time steps. The `oas_veg3d_finalize` subroutine stops the coupling via OASIS3-MCT.

B6 CLM

CLM is embedded within the CESM modelling system and its multiple components. In the case of land-only simulations, the active components are the driver/internal coupler (CPL7), CLM and a data atmosphere component. The latter is substituted to the atmospheric component used in coupled mode and provides the atmospheric forcing usually read from a file. In the framework of the OASIS3-MCT coupling, however, the file reading is deactivated and replaced by the coupling fields received from OASIS3-MCT (`receive_field_2cos`). The send operation (`send_field_2cos`) is also positioned in the data atmosphere component in order to enforce the same sequence of calls as in CESM. The definition of coupling fields and grids for the OASIS3-MCT coupling is also done in the data atmosphere component during initialisation before the time loop. Additionally, the initialisation (`oas_clm_init`) and finalisation (`oas_clm_finalize`) of the MPI communicator for the OASIS3-MCT coupling is positioned in the CESM driver respectively before and after the time loop. The sequence of hydrological and biogeophysical calculations during the time loop is given in black and the calls to optional modules are marked in grey.

The Supplement related to this article is available online at doi:10.5194/gmd-10-1549-2017-supplement.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The development of CCLM couplings would have not been possible without the continuous work done by OASIS, COSMO and CLM-Community colleagues and provision of computing time and support by computing centers. In particular we would like to thank Ulrich Schaettler (DWD) and Hans-Jürgen Panitz (KIT Karlsruhe) for source code maintenance of COSMO and COSMO-CLM. The OASIS support leading to our results received funding from the European Union Seventh Framework program under the IS-ENES2 project (grant agreement no. 312979). The overall support and provision of computing time by DKRZ Hamburg and the hosting of a developers' workshop by the Swiss National Supercomputing Centre (CSCS) are gratefully acknowledged.

We would like to highlight the contributions to the work presented here by further colleagues. First of all, Irina Fast provided the solution for dedicated distribution of model tasks on cores and the MPI-ESM version using the OASIS3-MCT coupler. Andreas Dobler (FU Berlin) made the pioneering work in coupling of COSMO-CLM using OASIS3. Sophie Valcke (CERFACS) provided the OASIS3-MCT support necessary to solve the problems of coupling with a regional model. We also thank Matthieu Leclair (ETH Zürich) for his critical comments which helped improve the manuscript.

Finally, we would like to kindly acknowledge the support provided by the German Ministry for Education and Research, in particular of CCLM+TRIMNP+CICE (in REKLIM), CCLM+NEMO (MiKliP II grant, FKZ 01LP1518C), CCLM+VEG3D (MiKliP I+II grant) and CCLM+MPI-ESM (MiKliP I grant, FKZ 01LP1165).

Edited by: J. Fyke

Reviewed by: S. Valcke and one anonymous referee

References

- Akhtar, N., Brauch, J., Dobler, A., Béranger, K., and Ahrens, B.: Medicanes in an ocean-atmosphere coupled regional climate model, *Nat. Hazards Earth Syst. Sci.*, 14, 2189–2201, doi:10.5194/nhess-14-2189-2014, 2014.
- Alexeev, Y., Mickelson, S., Leyffer, S., Jacob, R., and Craig, A.: The Heuristic Static Load-Balancing Algorithm Applied to the Community Earth System Model, in: 28th IEEE International Parallel and Distributed Processing Symposium Workshops, IEEE, 1581–1590, doi:10.1109/IPDPSW.2014.177, 2014.
- Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Caubel, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., and Wright, G.: CPMIP: measurements of real computational performance of Earth system models in CMIP6, *Geosci. Model Dev.*, 10, 19–34, doi:10.5194/gmd-10-19-2017, 2017.
- Balaprakash, P., Alexeev, Y., Mickelson, S. A., Leyffer, S., Jacob, R., and Craig, A.: Machine-learning-based load balancing for Community Ice CodE component in CESM, in: International Conference on High Performance Computing for Computational Science, Springer, 79–91, 2014.
- Baldauf, M., Seifert, A., Foerstner, J., Majewski, D., Raschendorfer, M., and Reinhardt, T.: Operational convective-scale numerical weather prediction with the COSMO model: description and sensitivities, *Mon. Weather Rev.*, 139, 3887–3905, 2011.
- Balmaseda, M. A., Mogensen, K., and Weaver, A. T.: Evaluation of the ECMWF ocean reanalysis system ORAS4, *Q. J. Roy. Meteor. Soc.*, 139, 1132–1161, doi:10.1002/qj.2063, 2013.
- Becker, N., Ulbrich, U., and Klein, R.: Systematic large-scale secondary circulations in a regional climate model, *Geophys. Res. Lett.*, 42, 1944–8007, doi:10.1002/2015GL063955, 2015.
- Beuviel, J., Sevault, F., Herrmann, M., Kontoyiannis, H., Ludwig, W., Rixen, M., Stanev, E., Béranger, K., and Somot, S.: Modelling the Mediterranean sea interannual variability during 1961–2000: focus on the Eastern Mediterranean Transient (EMT), *J. Geophys. Res.*, 115, C08517, doi:10.1029/2009JC005950, 2010.
- Beuviel, J., Lebeau-pin-Brossier, C., Beranger, K., Arsouze, T., Bourdalle-Badie, R., Deltel, C., Drillet, Y., Drobinski, P., Ferry, N., Lyard, F., Sevault, F., and Somot, S.: MED12, Oceanic Component for the Modeling of the Regional Mediterranean Earth System, *Mercator Ocean Quarterly Newsletter*, 46, 60–66, 2012.
- Bülow, K., Dietrich, C., Elizalde, A., Gröger, M., Heinrich, H., Hüttl-Kabos, S., Klein, B., Mayer, B., Meier, H. M., Mikolajewicz, U., Narayan, N., Pohlmann, T., Rosenhagen, G., Schimanke, S., Sein, D., and Su, J.: Comparison of three regional coupled ocean atmosphere models for the North Sea under today's and future climate conditions, *KLIWAS Schriftenreihe KLIWAS-27/2014*, Koblenz, Bundesanstalt für Gewässerkunde, doi:10.5675/Kliwas_27/2014, 2014.
- Byrne, D., Papritz, L., Frenger, I., Munnich, M., and Gruber, N.: Atmospheric Response to Mesoscale Sea Surface Temperature Anomalies: Assessment of Mechanisms and Coupling Strength in a High-Resolution Coupled Model over the South Atlantic, *J. Atmos. Sci.*, 72, 1872–1890, doi:10.1175/JAS-D-14-0195.1, 2015.
- Casulli, V. and Cattani, E.: Stability, Accuracy and Efficiency of a Semi-Implicit Method for Three-Dimensional Shallow Water Flow, *Comput. Math. Appl.*, 27, 99–112, 1994.
- Casulli, V. and Stelling, G. S.: Numerical Simulation of 3D Quasi-Hydrostatic, Free-Surface Flows, *J. Hydrol. Eng.*, 124, 678–686, 1998.
- Collins, W. D., Bitz, C. M., Blackmon, M. L., Bonan, G. B., Bretherton, C. S., Carton, J. A., Chang, P., Doney, S. C., Hack, J. J., Henderson, T. B., Kiehl, J. T., Large, W. G., McKenna, D. S., Santer, B. D., and Smith, R. D.: The Community Climate System Model version 3 (CCSM3), *J. Climate*, 19, 2122–2143, 2006.
- Craig, A., Vertenstein, M., and Jacob, R.: A new flexible coupler for earth system modeling developed for CCSM4 and CESM1, *Int. J. High Perform. C.*, 26, 31–42, doi:10.1177/1094342011428141, 2012.
- Davin, E. L. and Seneviratne, S. I.: Role of land surface processes and diffuse/direct radiation partitioning in simulating the European climate, *Biogeosciences*, 9, 1695–1707, doi:10.5194/bg-9-1695-2012, 2012.

- Davin, E. L., Stoeckli, R., Jaeger, E. B., Levis, S., and Seneviratne, S. I.: COSMO-CLM2: a new version of the COSMO-CLM model coupled to the Community Land Model, *Clim. Dynam.*, 37, 1889–1907, doi:10.1007/s00382-011-1019-z, 2011.
- Davin, E. L., Maisonnave, E., and Seneviratne, S. I.: Is land surface processes representation a possible weak link in current Regional Climate Models?, *Environ. Res. Lett.*, 11, 074027, doi:10.1088/1748-9326/11/7/074027, 2016.
- Dennis, J. M., Vertenstein, M., Worley, P. H., Mirin, A. A., Craig, A. P., Jacob, R., and Mickelson, S.: Computational performance of ultra-high-resolution capability in the Community Earth System Model, *Int. J. High Perform. C.*, 26, 5–16, doi:10.1177/1094342012436965, 2012.
- Dickinson, R., Oleson, K., Bonan, G., Hoffman, F., Thornton, P., Vertenstein, M., Yang, Z., and Zeng, X.: The Community Land Model and its climate statistics as a component of the Community Climate System Model, *J. Climate*, 19, 2302–2324, 2006.
- Dieterich, C., Schimanke, S., Wang, S., Väli, G., Liu, Y., Hordoir, R., Axell, L., and Meier, H.: Evaluation of the SMHI coupled atmosphere-ice-ocean model RCA4-NEMO, Tech. Rep. 47, Sveriges Meteorologiska och Hydrologiska Institut (SMHI), Sweden, 2013.
- DKRZ: The ECHAM3 Atmospheric General Circulation Model, Report no. 6, 2nd revision, Deutsches Klimarechenzentrum, Hamburg, 1993.
- Doms, G. and Baldauf, M.: A Description of the nonhydrostatic regional model LM, Part I: Dynamics and Numerics – COSMO V5.1, Tech. rep., Deutscher Wetterdienst, P.O. Box 100465, 63004 Offenbach, Germany, 2015.
- Doms, G., Förstner, J., Heise, E., Herzog, H.-J., Mironov, D., Raschendorfer, M., Reinhardt, T., Ritter, B., Schrodin, R., Schulz, J.-P., and Vogel, G.: A Description of the nonhydrostatic regional model LM, Part II: Physical Parameterization – LM_F90 4.20, Tech. rep., Deutscher Wetterdienst, P.O. Box 100465, 63004 Offenbach, Germany, 2011.
- Döscher, R., Will'en, U., Jones, C., Rutgersson, A., Meier, H.E.M., Hansson, U., and Graham, L.P.D.: The development of the regional coupled ocean-atmosphere model RCAO, *Boreal Environ. Res.*, 7, 183–192, 2002.
- Egbert, G. D. and Erofeeva, S. Y.: Efficient Inverse Modeling of Barotropic Ocean Tides, *J. Atmos. Ocean. Tech.*, 19, 183–204, 2002.
- Gasper, F., Goergen, K., Shrestha, P., Sulis, M., Rihani, J., Geimer, M., and Kollet, S.: Implementation and scaling of the fully coupled Terrestrial Systems Modeling Platform (TerrSysMP v1.0) in a massively parallel supercomputing environment – a case study on JUQUEEN (IBM Blue Gene/Q), *Geosci. Model Dev.*, 7, 2531–2543, doi:10.5194/gmd-7-2531-2014, 2014.
- Giorgetta, M. A., Jungclaus, J., Reick, C. H., Legutke, S., Bader, J., Boettinger, M., Brovkin, V., Crueger, T., Esch, M., Fieg, K., Glushak, K., Gayler, V., Haak, H., Hollweg, H.-D., Ilyina, T., Kinne, S., Kornbluh, L., Matei, D., Mauritsen, T., Mikolajewicz, U., Mueller, W., Notz, D., Pithan, F., Raddatz, T., Rast, S., Redler, R., Roeckner, E., Schmidt, H., Schnur, R., Segschneider, J., Six, K. D., Stockhause, M., Timmreck, C., Wegner, J., Widmann, H., Wieners, K.-H., Claussen, M., Marotzke, J., and Stevens, B.: Climate and carbon cycle changes from 1850 to 2100 in MPI-ESM simulations for the Coupled Model Intercomparison Project phase 5, *J. Adv. Model. Earth Syst.*, 5, 572–597, doi:10.1002/jame.20038, 2013.
- Gualdi, S., Somot, S., Li, L., Artale, V., Adani, M., Bellucci, A., Braun, A., Calmanti, S., Carillo, A., Dell'Aquila, A., Déqué, M., Dubois, C., Elizalde, A., Harzallah, A., Jacob, D., L'Hévéder, B., May, W., Oddo, P., Ruti, P., Sanna, A., Sannino, G., Scocimarro, E., Sevault, F., and Navarra, A.: THE CIRCE simulations: Regional climate change projections with realistic representation of the mediterranean sea, *B. Am. Meteorol. Soc.*, 94, 65–81, doi:10.1175/BAMS-D-11-00136.1, 2013.
- Hagos, S., Leung, R., Rauscher, S. A., and Ringler, T.: Error Characteristics of Two Grid Refinement Approaches in Aquaplanet Simulations: MPAS-A and WRF, *Mon. Weather Rev.*, 141, 3022–3030, doi:10.1175/MWR-D-12-00338.1, 2013.
- Hertwig, E., Storch, J. V., Handorf, D., Dethloff, K., Fast, I., and Krismer, T.: Effect of horizontal resolution on ECHAM6 AMIP performance, *Clim. Dynam.*, 45, 185–211, doi:10.1007/s00382-014-2396-x, 2015.
- Ho-Hagemann, H. T. M., Rockel, B., Kapitzka, H., Geyer, B., and Meyer, E.: COSTRICE – an atmosphere-ocean-sea ice model coupled system using OASIS3, HZG Report 2013-5, Tech. rep., Helmholtz-Zentrum Geesthacht, Geesthacht, Germany, 2013.
- Ho-Hagemann, H. T. M., Hagemann, S., and Rockel, B.: On the role of soil moisture in the generation of heavy rainfall during the Oder flood event in July 1997, *Tellus A*, 67, 1–17, doi:10.3402/tellusa.v67.28661, 2015.
- Hordoir, R., Dieterich, C., Basu, C., Dietze, H., and Meier, H. E. M.: Freshwater outflow of the Baltic Sea and transport in the Norwegian current: A statistical correlation analysis based on a numerical experiment, *Cont. Shelf Res.*, 64, 1–9, doi:10.1016/j.csr.2013.05.006, 2013.
- Hunke, E. C., Lipscomb, W. H., Turner, A. K., Jeffery, N., and Elliott, S.: CICE: The Los Alamos Sea Ice Model. Documentation and Software User's Manual. Version 5.0, Tech. Rep. LA-CC-06-012, T-3 Fluid Dynamics Group, Los Alamos National Laboratory, 2013.
- Ilyina, T., Six, K. D., Segschneider, J., Maier-Reimer, E., Li, H., and Nunez-Riboni, I.: Global ocean biogeochemistry model HAMOCC: Model architecture and performance as component of the MPI-Earth System Model in different CMIP5 experimental realizations, *J. Adv. Model. Earth Syst.*, 5, 287–315, doi:10.1029/2012MS000178, 2013.
- Inatsu, M. and Kimoto, M.: A scale interaction study on East Asian cyclogenesis using a general circulation model with an interactively nested regional model, *Mon. Weather Rev.*, 137, 2851–2868, doi:10.1175/2009MWR2825.1, 2009.
- Jacob, R., Larson, J., and Ong, E.: M × N communication and parallel interpolation in Community Climate System Model Version 3 using the model coupling toolkit, *Int. J. High Perform. C.*, 19, 293–307, 2005.
- Jones, P.: A user's guide for SCRIP: A spherical coordinate remapping and interpolation package, Tech. rep., Los Alamos National Laboratory, 1997.
- Jungclaus, J. H., Fischer, N., Haak, H., Lohmann, K., Marotzke, J., Matei, D., Mikolajewicz, U., Notz, D., and von Storch, J.-S.: Characteristics of the ocean simulations in MPIOM, the ocean component of the MPI Earth System Model, *J. Adv. Model. Earth Syst.*, 5, 422–446, doi:10.1002/jame.20023, 2013.

- Kerkweg, A. and Jöckel, P.: The 1-way on-line coupled atmospheric chemistry model system MECO(n) – Part 1: Description of the limited-area atmospheric chemistry model COSMO/MESSy, *Geosci. Model Dev.*, 5, 87–110, doi:10.5194/gmd-5-87-2012, 2012.
- Köhler, M., Schädler, G., Gantner, L., Kalthoff, N., Königer, F., and Kottmeier, C.: Validation of two SVAT models for different periods during the West African monsoon, *Meteorol. Z.*, 21, 509–524, 2012.
- Kotlarski, S., Keuler, K., Christensen, O. B., Colette, A., Déqué, M., Gobiet, A., Goergen, K., Jacob, D., Lüthi, D., van Meijgaard, E., Nikulin, G., Schär, C., Teichmann, C., Vautard, R., Warrach-Sagi, K., and Wulfmeyer, V.: Regional climate modeling on European scales: a joint standard evaluation of the EURO-CORDEX RCM ensemble, *Geosci. Model Dev.*, 7, 1297–1333, doi:10.5194/gmd-7-1297-2014, 2014.
- Kumar, S. V., Peters-Lidard, C. D., Eastman, J. L., and Tao, W.-K.: An integrated high-resolution hydrometeorological modeling testbed using LIS and WRF, *Environ. Model. Softw.*, 23, 169–181, doi:10.1016/j.envsoft.2007.05.012, 2008.
- Laprise, R., de Elia, R., Caya, D., Biner, S., Lucas-Picher, P., Diaconescu, E., Leduc, M., Alexandru, A., Separovic, L., and Climate, C. N. R.: Challenging some tenets of Regional Climate Modelling, *Meteorol. Atmos. Phys.*, 100, 3–22, doi:10.1007/s00703-008-0292-9, 2008.
- Lawrence, D. M., Oleson, K. W., Flanner, M. G., Thornton, P. E., Swenson, S. C., Lawrence, P. J., Zeng, X., Yang, Z.-L., Levis, S., Sakaguchi, K., Bonan, G. B., and Slater, A. G.: Parameterization Improvements and Functional and Structural Advances in Version 4 of the Community Land Model, *J. Adv. Model. Earth Syst.*, 3, M03001, doi:10.1029/2011MS000045, 2011.
- Lebeaupin, C., Béranger, K., Deltel, C., and Drobinski, P.: The Mediterranean response to different space-time resolution atmospheric forcings using perpetual mode sensitivity simulations, *Ocean Model.*, 36, 1–25, doi:10.1016/j.ocemod.2010.10.008, 2011.
- Levitus, S. and Boyer, T. P.: World Ocean Atlas, vol. 4: Temperature, number 4, NOAA/OAR/ESRL PSD, Boulder, Colorado, USA, 1994.
- Levitus, S., Antonov, J. I., and Boyer, T. P.: Warming of the world ocean, *Geophys. Res. Lett.*, 32, L02604, doi:10.1029/2004GL021592, 2005.
- Lindström, G., Pers, C. P., Rosberg, R., Strömqvist, J., and Arheimer, B.: Development and test of the HYPE (Hydrological Predictions for the Environment) model – A water quality model for different spatial scales, *Hydrol. Res.*, 41, 295–319, 2010.
- Lin-Jiong, Z., Yi-Min, L., Qing, B., Hai-Yang, Y., and Guo-Xiong, W.: Computational Performance of the High-Resolution Atmospheric Model FAMIL, *Atmospheric and Oceanic Science Letters*, 5, 355–359, 2012.
- Lorenz, P. and Jacob, D.: Influence of regional scale information on the global circulation: A two-way nesting climate simulation, *Geophys. Res. Lett.*, 32, L18706, doi:10.1029/2005GL023351, 2005.
- Ludwig, W., Dumont, E., Meybeck, M., and Heussner, S.: River discharges of water and nutrients to the Mediterranean and Black Sea: Major drivers for ecosystem changes during past and future decades?, *Prog. Oceanog.*, 80, 199–217, doi:10.1016/j.pocean.2009.02.001, 2009.
- Madec, G.: NEMO ocean engine, Tech. Rep. 27, Note du Pole de modélisation, Institut Pierre-Simon Laplace (IPSL), France, 2008.
- Madec, G.: NEMO ocean engine (version 3.3), Tech. Rep. 27, Note du Pole de modélisation, Institut Pierre-Simon Laplace (IPSL), France, 2011.
- Maisonnavé, E. and Caubel, A.: LUCIA, load balancing tool for OASIS coupled systems, TR-CMGC 14-63, CERFACS, 2014.
- Maisonnavé, E., Valcke, S., and Foujols, M.-A.: OASIS Dedicated User Support 2009-2012, Synthesis, Tech. rep., TR/CMGC/13/19, SUC au CERFACS, URA CERFACS/CNRS No1875, Toulouse, France, 2013.
- Masson, S., Hourdin, C., Benshila, R., Maisonnavé, E., Meurdesoif, Y., Mazauric, C., Samson, G., Colas, F., Madec, G., Bourdall Badie, R., Valcke, S., and Coquart, L.: Tropical Channel NEMO-OASIS-WRF Coupled simulations at very high resolution, in: 13th WRF Users' Workshop, 25–29 June 2012, Boulder, CO, USA, 2012.
- MEDAR-Group: Mediterranean and Black Sea database of temperature, salinity and biochemical parameters and climatological atlas, 4 CD-ROM, available at: www.ifremer.fr/medar/ (last access: 9 April 2017), European Commission Marine Science and Technology Programme (MAST), 2002.
- Oleson, K., Lawrence, D., Bonan, G., Flanner, M., Kluzek, E., Lawrence, P., Levis, S., Swenson, S., Thornton, P., Dai, A., Decker, M., Dickinson, R., Feddema, J., Heald, C., Hoffman, F., Lamarque, J.-F., Mahowald, N., Niu, G.-Y., Qian, T., Randerson, J., Running, S., Sakaguchi, K., Slater, A., Stockli, R., Wang, A., Yang, Z.-L., Zeng, X., and Zeng, X.: Technical description of version 4.0 of the Community Land Model (CLM), NCAR Tech. Note NCAR/TN-478+STR, Nat. Cent. for Atmos. Res., Boulder, CO, 2010.
- Oleson, K., Lawrence, D., Bonan, G., Drewniak, B., Huang, M., Koven, C., Levis, S., Li, F., Riley, W., Subin, Z., Swenson, S., Thornton, P., Bozbiyik, A., Fisher, R., Kluzek, E., Lamarque, J.-F., Lawrence, P., Leung, L., Lipscomb, W., Muszala, S., Ricciuto, D., Sacks, W., Sun, Y., Tang, J., and Yang, Z.-L.: Technical description of version 4.5 of the Community Land Model (CLM), NCAR Tech. Note NCAR/TN-503+STR, Natl. Cent. for Atmos. Res., Boulder, CO, doi:10.5065/D6RR1W7M, 2013.
- Pham, T., Brauch, J., Dieterich, D., Früh, B., and Ahrens, B.: New coupled atmosphere-ocean-ice system COSMO-CLM/NEMO: On the air temperature sensitivity on the North and Baltic Seas, *Oceanologia*, 56, 167–189, doi:10.5697/oc.56-2.167, 2014.
- Prein, A. F., Gobiet, A., Suklitsch, M., Truhetz, H., Awan, N. K., Keuler, K., and Georgievski, G.: Added value of convection permitting seasonal simulations, *Clim. Dynam.*, 41, 2655–2677, doi:10.1007/s00382-013-1744-6, 2013.
- Reick, C., Raddatz, T., Brovkin, V., and Gayler, V.: Representation of natural and anthropogenic land cover change in MPI-ESM, *J. Adv. Model. Earth Syst.*, 5, 459–482, doi:10.1002/jame.20022, 2013.
- Rockel, B., Will, A., and Hense, A.: The Regional Climate Model CLM, *Meteorol. Z.*, 17, 347–348, 2008.
- Rummukainen, M., Raesaenen, J., Bringfelt, B., Ullerstig, A., Omstedt, A., Willen, U., Hansson, U., and Jones, C.: A regional climate model for northern Europe: model description and results from the downscaling of two GCM control simulations, *Clim. Dynam.*, 17, 339–359, 2001.

- Schädler, G.: Numerische Simulationen zur Wechselwirkung zwischen Landoberfläche und atmosphärischer Grenzschicht, PhD thesis, Karlsruher Institut für Technologie, Institut für Meteorologie und Klimaforschung, 1990.
- Schneck, R., Reick, C., and Raddatz, T.: Land contributions to natural CO₂ variability on time scales of centuries, *J. Adv. Model. Earth Syst.*, 5, 354–365, doi:10.1002/jame.20029, 2013.
- Somot, S., Sevault, F., Déqué, M., and Crépon, M.: 21st century climate change scenario for the Mediterranean using a coupled Atmosphere–Ocean Regional Climate Model, *Global Planet. Change*, 63, 112–126, doi:10.1016/j.gloplacha.2007.10.003, 2008.
- Staff, S. S.: Soil taxonomy: A basic system of soil classification for making and interpreting soil surveys, 2nd Edn., Tech. rep., Natural Resources Conservation Service, U.S. Department of Agriculture Handbook 436, 1999.
- Stanev, E. and Peneva, E.: Regional sea level response to global forcing. Black Sea examples, *J. Global Planet. Change*, 32, 33–47, 2002.
- Steiner, A., Pal, J., Giorgi, F., Dickinson, R., and Chameides, W.: The coupling of the Common Land Model (CLM0) to a regional climate model (RegCM), *Theor. Appl. Climatol.*, 82, 225–243, doi:10.1007/s00704-005-0132-5, 2005.
- Steiner, A. L., Pal, J. S., Rauscher, S. A., Bell, J. L., Diffenbaugh, N. S., Boone, A., Sloan, L. C., and Giorgi, F.: Land surface coupling in regional climate simulations of the West African monsoon, *Clim. Dynam.*, 33, 869–892, doi:10.1007/s00382-009-0543-6, 2009.
- Stevens, B., Giorgetta, M. A., Esch, M., Mauritsen, T., Crueger, T., Rast, S., Salzmann, M., Schmidt, H., Bader, J., Block, K., Brokopf, R., Fast, I., Kinne, S., Kornbluh, L., Lohmann, U., Pincus, R., Reichler, T., and Roeckner, E.: Atmospheric component of the MPI-M Earth System Model: ECHAM6, *J. Adv. Model. Earth Syst.*, 5, 146–172, doi:10.1002/jame.20015, 2013.
- Valcke, S.: The OASIS3 coupler: a European climate modelling community software, *Geosci. Model Dev.*, 6, 373–388, doi:10.5194/gmd-6-373-2013, 2013.
- Valcke, S., Craig, T., and Coquart, L.: OASIS3-MCT User Guide, OASIS3-MCT 2.0, Tech. rep., TR/CMGC/13/17, CERFACS/CNRS SUC URA No 1875, Toulouse, France, 2013.
- Vancoppenolle, M., Fichefet, T., Goosse, H., Bouillon, S., Madec, G., and Maqueda, M.: Simulating the mass balance and salinity of arctic and antarctic sea ice, *Ocean Model.*, 27, 33–53, 2009.
- Vogel, B., Vogel, H., Bäumer, D., Bangert, M., Lundgren, K., Rinke, R., and Stanelle, T.: The comprehensive model system COSMO-ART – Radiative impact of aerosol on the state of the atmosphere on the regional scale, *Atmos. Chem. Phys.*, 9, 8661–8680, doi:10.5194/acp-9-8661-2009, 2009.
- Vörösmarty, C. J., Fekete, B. M., and Tucker, B. A.: Global River Discharge Database, Version 1.0 (RivDIS V1.0), Volumes 0 through 6, A contribution to IHP-V Theme 1, Technical Documents in Hydrology Series, UNESCO, Paris, 1996.
- Wilhelm, C., Rechid, D., and Jacob, D.: Interactive coupling of regional atmosphere with biosphere in the new generation regional climate system model REMO-iMOVE, *Geosci. Model Dev.*, 7, 1093–1114, doi:10.5194/gmd-7-1093-2014, 2014.
- Worley, P. H., Mirin, A. A., Craig, A. P., Taylor, M. A., Dennis, J. M., and Vertenstein, M.: Performance of the Community Earth System Model, in: SC'11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE, ACM, New York, NY, USA, doi:10.1145/2063384.2063457, 2011.
- Zou, L. and Zhou, T.: Can a Regional Ocean Atmosphere Coupled Model Improve the Simulation of the Interannual Variability of the Western North Pacific Summer Monsoon?, *J. Climate*, 26, 2353–2367, doi:10.1175/JCLI-D-11-00722.1, 2013.