REQUEST FOR A SPECIAL PROJECT 2016–2018

MEMBER STATE: Germany

Principal Investigator[1]: Stefan Kollet

Affiliation: Forschungszentrum Jülich, Agrosphere (IBG-3)

Address: Wilhelm-Johnen-Straße

52425 Jülich

Germany

E-mail: s.kollet@fz-juelich.de

Other researchers: Jessica Keune, Meteorological Institute, Bonn University, Germany
Klaus Görgen, Jülich Supercomputing Centre, Germany
Florian Pappenberger, ECMWF

Project Title: Integrated Simulations of the Terrestrial System over the European CORDEX Domain

If this is a continuation of an existing project, please state the computer project account assigned previously.	**SP** DEKOLL_________________	
Starting year: (Each project will have a well defined duration, up to a maximum of 3 years, agreed at the beginning of the project.)	2016	
Would you accept support for 1 year only, if necessary?	YES ☒	NO ☐

Computer resources required for 2016-2018: (The maximum project duration is 3 years, therefore a continuation project cannot request resources for 2018.)	**2016**	**2017**	**2018**
High Performance Computing Facility (units)	30.000.000	30.000.000	-
Data storage capacity (total archive volume) (gigabytes)	60.000	60.000	-

*An electronic copy of this form **must be sent** via e-mail to:* *special_projects@ecmwf.int*

Electronic copy of the form sent on (please specify date):

June 29th 2015……………………

Continue overleaf

[1] The Principal Investigator will act as contact person for this Special Project and, in particular, will be asked to register the project, provide an annual progress report of the project's activities, etc.

Principal Investigator: Stefan Kollet……………………………………………………

Project Title: Integrated Simulations of the Terrestrial System over the European CORDEX Domain………………………………………………

Extended abstract

It is expected that Special Projects requesting large amounts of computing resources (500,000 SBU or more) should provide a more detailed abstract/project description (3-5 pages) including a scientific plan, a justification of the computer resources requested and the technical characteristics of the code to be used. The Scientific Advisory Committee and the Technical Advisory Committee review the scientific and technical aspects of each Special Project application. The review process takes into account the resources available, the quality of the scientific and technical proposals, the use of ECMWF software and data infrastructure, and their relevance to ECMWF's objectives. - Descriptions of all accepted projects will be published on the ECMWF website.

Abstract

The objective of this study is to perform high-resolution fully coupled aquifer-to-atmosphere simulations over the European CORDEX domain. The simulations will be performed with the integrated Terrestrial Systems Modeling Platform, TerrSysMP, consisting of the three-dimensional surface-subsurface model ParFlow, the Community Land Model CLM3.5 and the numerical weather prediction model COSMO of the German Weather Service (Shrestha et al., 2014, Gasper et al., 2014). At the ECMWF, the system will be set up with an initial spatial resolution of 0.11° (12.5km), which will be increased to 3km over the course of the project as a proof-of-concept study. The simulations will be used to interrogate the two-way feedbacks of groundwater and soil moisture dynamics with essential climate variables, such as air temperature and precipitation, at continental scales. Additionally, since TerrSysMP provides simulation results of *all* hydrologic states and fluxes, the potential for a fresh water monitoring system including e.g., stream discharge and groundwater storages will be assessed with observations. In the beginning, it is planned to perform event-based simulations, focusing on the floods in the early summer of 2013 and 2014, and the heat wave in 2003. These simulations will be followed by time slice simulations covering multiple decades between 1986 to 2005 for validation, and near-term and far-term scenario simulations between 2041 to 2060 and 2081 to 2100, respectively. The simulations are compared to in-situ and remote sensing observations, such as FLUXNET, gauge measurements, SMOS and GRACE, and reanalysis products, such as the recently published high-resolution European reanalysis (Bollmeyer et al., 2015) and ERA Interim as well as ERA Interim Land (Balsamo et al., 2015). Later on, for the proof-of-concept study at 3km resolution up to 6 months of simulation are planned. From a technical perspective it is important to note that extensive preliminary work has been performed i.e. the TerrSysMP model over the European CORDEX domain has been implemented and tested at the aforementioned resolutions. Table 1 gives an overview of the model setup.

Table 1. Model setup of TerrSysMP over the European CORDEX domain with a rotated north pole at (39.25,-162). The vertical resolutions of each model varies with height and depth. The atmospheric model COSMO has a extra sponge of 8 grid cells on each side.

		Horizontal dimensions	Horizontal resolution	Vertical dimensions	time step
CRODEX EUR-11	**COSMO**	444 x 432	0.11°	50	60s
	CLM3.5	436 x 424	0.11°	10	1h
	ParFlow	436 x 424	0.11°	15	1h
CORDEX EUR-3	**COSMO**	1600 x 1552	0.0275°	50	15s
	CLM3.5	1492 x 1544	0.0275°	10	1h
	ParFlow	1492 x 1544	0.0275°	15	1h

Particularly, the impact of the groundwater parameterization on atmospheric processes through land surface-atmosphere feedbacks, such as the soil moisture-temperature and the soil moisture-precipitation feedback, are analyzed. Therefore two different configurations of TerrSysMP are run, where only the parameterization of the subsurface differs. In this case, a physics-based 3D groundwater implementation is compared to a one-dimensional free-drainage approach, which is commonly used in the land surface scheme of regional climate models. Figure 1 illustrates the different configurations of TerrSysMP. In the former, 3D variably saturated groundwater flow is simulated from the bedrock to the land surface (TerrSysMP (3D), figure 1a) honouring a dynamic water table and groundwater reservoir. This leads to distinct groundwater convergence zones along e.g. river corridors and regionally wetter (or drier) soil moisture conditions.

In the latter, a free drainage boundary condition is implemented reflecting a kind of open-bucket parameterization commonly applied in land surface models leading to uniform drying of the subsurface. Convergence of groundwater and baseflow conditions along river corridors cannot be simulated. Here, a sensitivity study over at least 5 years is planned in order to analyze the significance of the impact of groundwater on atmospheric climate on longer time scales and vice versa.

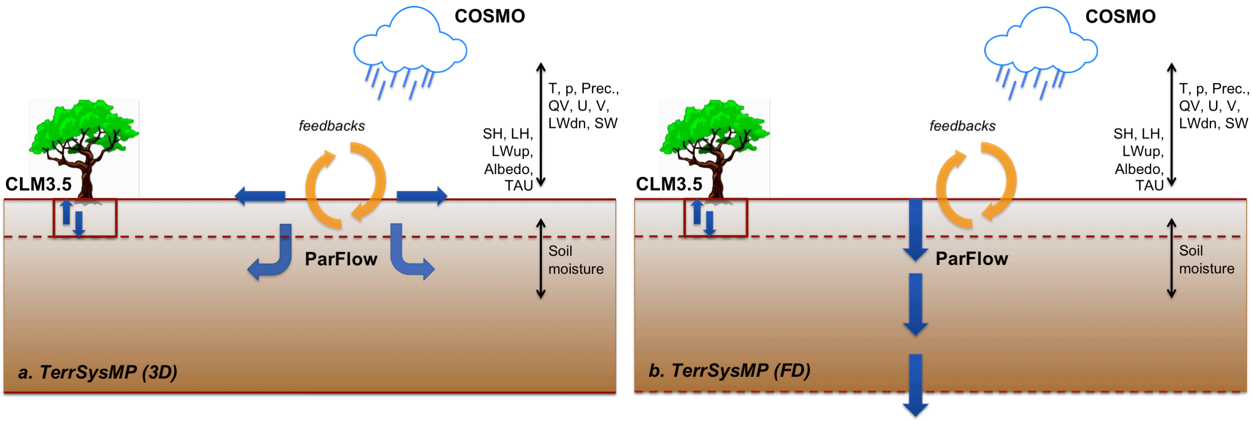

Figure 1. Vertical sketch of the setup of TerrSysMP with the atmospheric weather prediction model COSMO, the land surface model CLM3.5 and the subsurface-surface model ParFlow. Figures 1a and 1b illustrate the different configurations of TerrSysMP with a physics-based 3D subsurface and a 1D free drainage approach, respectively.

Figure 2 shows the model domain and a snapshot of soil moisture and cloud liquid and ice water content at the beginning of June 2013.

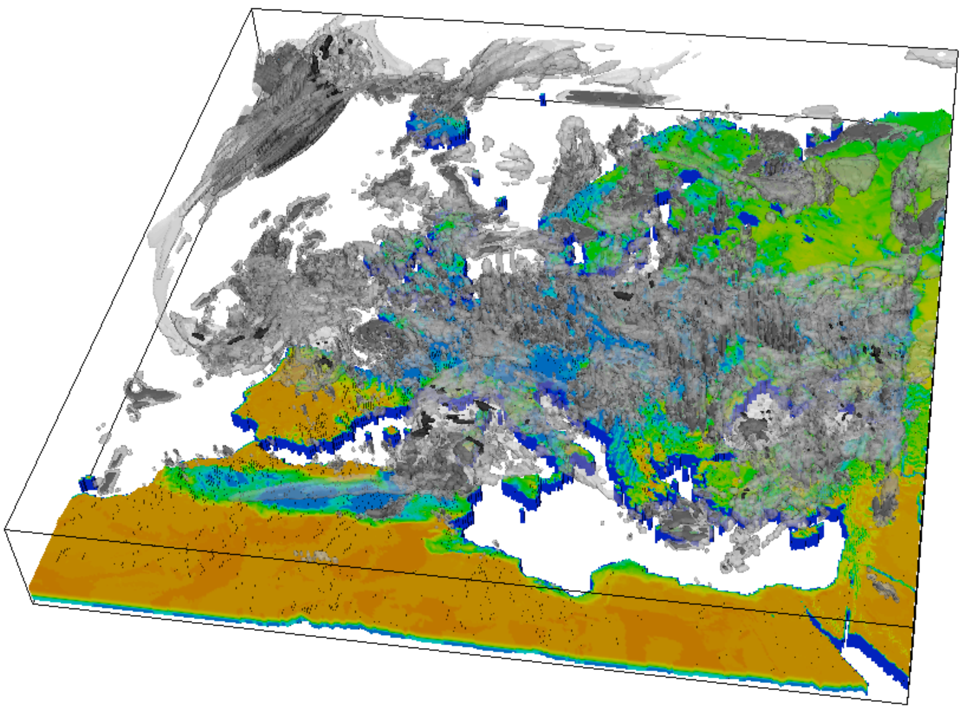

Figure 2. Snapshot of fully coupled simulations with TerrSysMP of 3D groundwater and soil moisture (orange: dry; blue: wet) and liquid/ice cloud water content in the summer of 2013 over the European CORDEX domain at 12 km lateral resolution (courtesy of Jessica Keune, Forschungszentrum Jülich and Meteorological Institute, Bonn University)

Work plan

The workplan for the simulations of TerrSysMP over the CORDEX domain can be divided into three parts:

1. successfully implementing TerrSysMP over CORDEX in the ECMWF supercomputer environment including the workflow in ecFlow.
2. reaching dynamic equilibrium (spinup) of hydrological compartments of TerrSysMP
3. starting production runs for events (such as 2003 drought, 2013 flood) and over years for the fully-coupled system with two different subsurface configurations.

Additionally, the proof-of-concept with 3km resolution over the European CORDEX domain is run. Table 2 gives an overview of the upcoming tasks and simulations.

Task 1 has been almost completed and will not require much more time in future. The modelling system is completely integrated into ecFlow in order to enable automated restarts. Furthermore, this setup makes use of the ECMWF database by automatically fetching ERA Interim data as lateral boundary conditions for TerrSysMP. Task 2 deals with the spinup runs, which are a basic necessity in order to have consistent initial conditions for the hydrologic compartment in the fully coupled simulations. Some work in this context has already been done but a dynamic equilibrium of the hydrological compartments in TerrSysMP has not been reached yet. Nevertheless the spinup process should be completed within the first months of project start. Task 3 constitutes the actual production runs over events and over multiple years and decades using two configurations of TerrSysMP with a 3D setup and a 1D free drainage approach as explained above.

Table 2. Workplan for the simulations of TerrSysMP over the European CORDEX domain.

Year		2016												2017											
Month		J	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J	A	S	O	N	D
Implementation																									
Spinup																									
Production run (12 km)	TerrSysMP (3D)																								
Production run (12 km)	TerrSysMP (FD)																								
Proof-of-concept (3 km)	TerrSysMP (3D)																								
Analysis																									

Justification of the requested computer resources

The current fully coupled setup of TerrSysMP at ECMWF with 12 km horizontal resolution uses 4 nodes, each with 24 CPUs and runs 3 hours in 30 minutes wall clock time. This is about 960 SBUs in total or 320 SBUs per simulated hour. Extrapolating this performance to a one-year simulation of the fully coupled system, computational resources of about 3.000.000 SBUs are required. For decadal simulations 30.000.000 SBUs are necessary. This requested compute time covers *one* simulation of the fully coupled system (ParFlow-CLM-COSMO) for 10 years or two simulations of TerrSysMP with different configurations, i.e. 3D groundwater and a 1D free-drainage approach, over 5 years each.

The required model output of the three models is about 17GB / day. For one year simulations, data storage of 6 TB is needed. Decadal simulations require about 60 TB of storage. The requested storage covers hourly output of the minimum basic output of each model.

Technical characteristics of Terrestrial Systems Monitoring Platform, TerrSysMP

The technical characteristics of TerrSysMP are taken from Gasper et al. 2014 with slight modifications. "The parallel Terrestrial Systems Modeling Platform (v1.0) consists of the numerical weather prediction system (COSMO, v4.11) of the German Weather Service (Baldauf et al., 2011), the Community Land Model (CLM, v3.5) (Oleson et al., 2008), and the variably saturated surface–subsurface flow code ParFlow (v3.1) (Kollet and Maxwell, 2006). For details with regard to the different component models, the reader is referred to the aforementioned publications. In TerrSysMP, these component models were integrated in a scale-consistent way conserving moisture and energy from the subsurface across the land surface into the atmosphere (Fig. 1). The interested reader is referred to Shrestha et al. (2014) for a more detailed description of the modeling system. Each component model is itself parallel and has been demonstrated to scale efficiently to a large number of parallel tasks (e.g., Kollet et al., 2010). In order to couple differently structured component models to simulate complex systems, it is necessary to match a specified interface to exchange fluxes and states. Tailoring this interface exclusively for a certain model environment does not provide the flexibility and compatibility that is needed for various scientific modeling platforms. The obvious solution is a coupling strategy that abstracts that interface via synchronous data exchange, time step management, grid transformation and interpolation methods, and I/O with a low cost and strong stability on different computing environments. In TerrSysMP, the interface abstraction relies on the Multiple Program Multiple Data (MPMD) execution model, which forms the basis of the external Ocean-Atmosphere-Sea-Ice- Soil coupler, OASIS (Valcke, 2013). With the MPMD functionality, which is offered by most Message Passing Interface (MPI) implementations, it is possible to run several executables within the same global MPI_COMM_WORLD communicator. This functionality enables a coupler that has an external "view" of all component models reflecting the key requirement of high modularity and is especially useful in coupling of component models with fast development cycles and heterogeneous computation loads (Chang et al., 1997). The implementation of the coupler is almost non-invasive. Therefore component models remain independent, which allows for interchangeable executables as a major advantage. Thus, OASIS links the aforementioned component models as independent executables, and can be implemented in two different versions: OASIS3 and OASIS3-MCT (OASIS3 including the Model Coupling Toolkit libraries). In case of OASIS3, the coupler is implemented as an additional independent executable, while in case of OASIS3-MCT the coupler is attached to each individual component model as a library."

References

Baldauf, M., Seifert, A., Forstner, J., Majewski, D., Raschendorfer, M., and Reinhardt, T., 2011, Operational Convective-Scale Numerical Weather Prediction with the COSMO Model: Description and Sensitivities, Mon. Weather Rev., 139, 3887–3905, doi:10.1175/Mwr-D-10-05013.1.

Balsamo, G., Albergel, C., Beljaars, A., Boussetta, S., Brun, E., Cloke, H., Dee, D., Dutra, E., Muñoz-Sabater, J., Pappenberger, F., de Rosnay, P., Stockdale, T., and Vitart, F., 2015, ERA-Interim/Land: a global land surface reanalysis data set, Hydrol. Earth Syst. Sci., 19, 389-407, doi:10.5194/hess-19-389-2015.

Bollmeyer, C., Keller, J. D., Ohlwein, C., Wahl, S., Crewell, S., Friederichs, P., Hense, A., Keune, J., Kneifel, S., Pscheidt, I., Redl, S. and Steinke, S., 2015, Towards a high-resolution regional reanalysis for the European CORDEX domain. Q.J.R. Meteorol. Soc., 141: 1–15. doi: 10.1002/qj.2486.

Chang, C.-C., Czajkowski, G., Eicken, T. V., and Kesselman, C., 1997, Evaluating the performance limitations of MPMD communication, Proceedings of the 1997 ACM/IEEE conference on Supercomputing, San Jose, CA.

Kollet, S. J. and Maxwell, R. M., 2006, Integrated surface-groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model, Adv. Water Resour., 29, 945–958, doi:10.1016/J.Advwatres.2005.08.006.

Kollet, S. J., Maxwell, R. M., Woodward, C. S., Smith, S., Vanderborght, J., Vereecken, H., and Simmer, C., 2010, Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources, Water Resour. Res., 46, W04201, doi:10.1029/2009wr008730.

Gasper, F., K. Goergen, P. Shrestha, M. Sulis, J. Rihani, M. Geimer and S.J. Kollet, 2014, 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.

Oleson, K. W., Niu, G. Y., Yang, Z. L., Lawrence, D. M., Thornton,P. E., Lawrence, P. J., Stockli, R., Dickinson, R. E., Bonan, G. B., Levis, S., Dai, A., and Qian, T., 2008, Improvements to the Community Land Model and their impact on the hydrological cycle, J. Geophys. Res.-Biogeosci., 113, G01021, doi:10.1029/2007jg000563.

Shrestha, P., M. Sulis, M. Masbou, S.J. Kollet, C. Simmer, 2014, A scale-consistent Terrestrial Systems Modeling Platform based on COSMO, CLM and ParFlow, Monthly Weather Review, 142, 3466–3483.

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