



The on-line coupled atmospheric chemistry model system MECO(n) – Part 5: Expanding the Multi-Model-Driver (MMD v2.0) for 2-way data exchange including data interpolation via GRID (v1.0)

Astrid Kerkweg^{1,a}, Christiane Hofmann^{1,a}, Patrick Jöckel², Mariano Mertens², and Gregor Pante^{1,b}

¹Institut für Physik der Atmosphäre, Johannes Gutenberg-Universität Mainz, 55099 Mainz, Germany

²Deutsches Zentrum für Luft- und Raumfahrt (DLR), Institut für Physik der Atmosphäre, 82234 Oberpfaffenhofen, Germany

^anow at: Meteorologisches Institut, Rheinische Friedrich-Wilhelms-Universität Bonn, Auf dem Hügel 20,
53121 Bonn, Germany

^bnow at: Institute of Meteorology and Climate Research, Department Troposphere Research (IMK-TRO), Karlsruhe Institute
of Technology (KIT), 76131 Karlsruhe, Germany

Correspondence: Astrid Kerkweg (kerkweg@uni-bonn.de)

Received: 31 March 2017 – Discussion started: 25 April 2017

Revised: 17 November 2017 – Accepted: 21 December 2017 – Published: 16 March 2018

Abstract. As part of the Modular Earth Submodel System (MESSy), the Multi-Model-Driver (MMD v1.0) was developed to couple online the regional Consortium for Small-scale Modeling (COSMO) model into a driving model, which can be either the regional COSMO model or the global European Centre Hamburg general circulation model (ECHAM) (see Part 2 of the model documentation). The coupled system is called MECO(n), i.e., MESSy-fied ECHAM and COSMO models nested n times. In this article, which is part of the model documentation of the MECO(n) system, the second generation of MMD is introduced. MMD comprises the message-passing infrastructure required for the parallel execution (multiple programme multiple data, MPMD) of different models and the communication of the individual model instances, i.e. between the driving and the driven models. Initially, the MMD library was developed for a one-way coupling between the global chemistry–climate ECHAM/MESSy atmospheric chemistry (EMAC) model and an arbitrary number of (optionally cascaded) instances of the regional chemistry–climate model COSMO/MESSy. Thus, MMD (v1.0) provided only functions for unidirectional data transfer, i.e. from the larger-scale to the smaller-scale models.

Soon, extended applications requiring data transfer from the small-scale model back to the larger-scale model became of interest. For instance, the original fields of the larger-scale model can directly be compared to the upscaled small-scale

fields to analyse the improvements gained through the small-scale calculations, after the results are upscaled. Moreover, the fields originating from the two different models might be fed into the same diagnostic tool, e.g. the online calculation of the radiative forcing calculated consistently with the same radiation scheme. Last but not least, enabling the two-way data transfer between two models is the first important step on the way to a fully dynamical and chemical two-way coupling of the various model instances.

In MMD (v1.0), interpolation between the base model grids is performed via the COSMO preprocessing tool INT2LM, which was implemented into the MMD submodel for online interpolation, specifically for mapping onto the rotated COSMO grid. A more flexible algorithm is required for the backward mapping. Thus, MMD (v2.0) uses the new MESSy submodel GRID for the generalised definition of arbitrary grids and for the transformation of data between them.

In this article, we explain the basics of the MMD expansion and the newly developed generic MESSy submodel GRID (v1.0) and show some examples of the abovementioned applications.

1 Introduction

As the fifth part of a paper series about the MECO(n) – the Modular Earth Submodel System (MESSy)-fied European Centre Hamburg general circulation model (ECHAM) and Consortium for Small-scale Modeling (COSMO) models nested n times – system, and as such a component of the ACP/GMD special issue¹ about MESSy, this article documents the progress of the MESSy code development. More specifically, the second-generation of the Multi-Model-Driver (MMD v2.0) is introduced, which enables the two-way online nesting between different model instances (base models). Online nesting means that the coupled models exchange their data via the computer memory, in contrast to the data exchange via files on disks in common offline nesting procedures. Thus, this article describes further development of the one-way online nesting system presented in the second part of the paper series (Kerkweg and Jöckel, 2012b).

We achieve the nesting by coupling different models; thus, our two-way nesting is implemented as two-way coupling of global and regional atmospheric models. Usually the term “two-way coupling” is used in the context of different Earth system compartment models, such as land, ocean or atmospheric models being connected within a comprehensive Earth system model. Here, two-way nesting through two-way coupling is used to distinguish it from fundamentally different other nesting techniques, as, for instance, regional static grid refinement. For a number of atmospheric models, grid refinement features exist. Usually, the grid resolution needs to be subdivided by a fixed factor: e.g. a factor of 3 for the Weather Research and Forecasting (WRF) model (Moeng et al., 2007; Harris and Durran, 2010) or 2 for the icosahedral non-hydrostatic (ICON) model (Zaengl et al., 2015). These constraints minimise the interpolation error, especially for the horizontal interpolation. At least these two models deal with the different grid refinement areas within the same executable; i.e. they are coupled “internally” (see Kerkweg and Jöckel, 2012b for a discussion of internal and external coupling). The individual grid refinement areas within one model configuration are usually called “patches”. In contrast to this, the MECO(n) system is implemented as an external coupling, i.e. a real two-way nesting of the same or different base models (here, COSMO/MESSy and EMAC).

In the MECO(n) system, we follow this second approach for the following reasons:

- It is necessary to couple the model instances externally, since different base models, EMAC and COSMO/MESSy, are nested into each other. This in itself prevents the “patch approach”, as the internal coupling or patch approach is usually a feature of regional grid refinements, which is directly embedded in (or part of) the model code, as, for instance, in WRF or ICON, in which the user can specify the number of patches and

their corresponding domains flexibly at runtime. For such a feature, however, the entire model code needs to be “aware” of a(n arbitrary) number of grid refined patches. To equip legacy code (as COSMO or ECHAM) supplementarily with such a feature would basically mean to rewrite the entire code from scratch. The reason is that all prognostic (and diagnostic) variables need to exist on each patch technically independent of each other.

- Different COSMO/MESSy model instances are nested into each other using the same algorithms as for the EMAC–COSMO/MESSy nesting. The external coupling approach was favoured here due to limitations of the Fortran95 name space: in fully object-oriented languages, overloaded “sets” or “instances” of the same variable(s) could be defined; however, the Fortran95 language standard does not allow to have the same variable with the same name in the same name space more than once. Thus, a complete recoding of the base model, e.g. replacing arrays by structures of arrays, would be required for the patch approach.
- A nesting of COSMO/MESSy model instances employing different grids (e.g. rotated differently) is possible. This also includes the possibility to realise an arbitrary resolution jump; i.e. the factor for the grid refinement is freely selectable, in contrast to the fixed factors of 2 or 3 as required by the ICON or the WRF model, respectively. Especially for air quality applications, a higher-resolution jump is necessary to reduce computational costs. Here, a global instance providing consistent boundary data is required, while the scientific focus is on a much finer-resolved model instance.
- Due to the external coupling, prognostic variables are not necessarily all coupled back to the coarser model. Thus, two-way nesting does not necessarily imply (full) feedback of the smaller- to the coarse-scale model. Consequently, the coupling can also be used to couple back diagnostic fields only. Additionally, testing of the influence of the coupling of different (individual) variables is easier to accomplish by external coupling.

Apart from these advantages, the external coupling proves to be more challenging than the internal coupling. Horizontal and vertical interpolation errors are expected to be larger, depending on the relations between the different grids and differences in their orographies. From these, the adaption to the higher-resolved orography of the nested simulation causes the largest error. An additional disadvantage of all external coupling approaches is the need for the user to optimise the distribution of the available parallel tasks among the different model instances in order to achieve an optimum runtime performance with minimised waiting times between the model instances.

¹http://www.geosci-model-dev.net/special_issue10_22.html

As far as we know, the only other two-way online nested modelling system using external coupling is an MPI-ESM (Giorgetta et al., 2013) – COSMO climate limited area model (CLM) coupling via OASIS3-MCT (Model Coupling Toolkit) (Will et al., 2017). This was developed in parallel to our two-way coupling approach within the same German Ministry of Education and Research (BMBF)-funded MiKlip project², as different approaches had to be assessed. In contrast to our system, the MPI-ESM–COSMO coupling via OASIS3-MCT is restricted to the coupling of one COSMO instance only; i.e. no further online COSMO–COSMO coupling is possible in the system of Will et al. (2017). Technically, this COSMO–COSMO coupling would of course also be possible, but it is not implemented. In the rest of the article, we will use the terms “two-way coupling” and “two-way nesting” synonymously for the approach chosen in the MECO(n) system.

This article documents a major achievement in the development of the online coupled MECO(n) system, the central part of which is the MESSy software. As described by Jöckel et al. (2016), Baumgaertner et al. (2016) and the MESSy home page (<http://messy-interface.org>; last access: 15 November 2017):

“The Modular Earth Submodel System (MESSy) is a software providing a framework for a standardized, bottom-up implementation of Earth System Models (or parts of those) with flexible complexity. “Bottom-up” means the MESSy software provides an infrastructure with generalized interfaces for the standardized control and interconnection (=coupling) of “low-level ESM components” (dynamic cores, physical parameterizations, chemistry packages, diagnostics etc.), which are called submodels. MESSy comprises currently about 60 submodels (i.e., coded MESSy conform):

- infrastructure (=the framework) submodels (sometimes called generic submodels),
- diagnostic, atmospheric chemistry and model physics related submodels.

The main design concept of MESSy is the strict separation of process description (=process and diagnostic submodels) from model infrastructure (e.g., memory management, input/output, flow control, ...).

Within MESSy, the operator splitting is formalized as the fundamental concept. Model codes are organized in 4 conceptual software layers: a basemodel of any level of complexity is complemented by a basemodel interface layer (BMIL). A further interface layer to the submodels (SMIL, submodel interface layer) makes it possible to keep process

submodels as distinct as possible in the submodel core layer (SMCL).”

MESSy currently employs the programming language Fortran90/95 with some rare exceptions linking libraries containing C or C++ code.

Furthermore, different base models, e.g. the global model ECHAM³, the regional COSMO model⁴ and the coupled global climate model CESM1⁵ have been expanded by the MESSy middleware (i.e. the MESSy infrastructure components) to enable a standardised expansion by additional or alternative process components (e.g. for physics or chemistry) and diagnostic components, which we call MESSy submodels.

In Part 2 of the MECO(n) model documentation, the one-way online coupled model system MECO(n), for which MMD was developed initially, was described in detail (Kerkweg and Jöckel, 2012b). In the online coupled system MECO(n), an arbitrary number of COSMO/MESSy model instances are nested online into one master model. This driving model can be either the global EMAC or a coarser COSMO/MESSy model instance. The data exchange is implemented as client–server system, where the driving model acts as server providing the client model with the data required for the calculation of the initial and boundary fields used to drive the regional model.

The MMD v1.0 provides the software necessary for the data exchange from the server to the client model and for the calculation of the initial and boundary data. MMD consists of two parts: (1) a library which performs the data exchange between the model instances and (2) MESSy submodels, which organise and process these data.

In addition to the functionalities provided by MMD (v1.0), the update of MMD presented here provides the possibility to exchange data in both directions during the time in-

³The core atmospheric model is the fifth-generation European Centre Hamburg general circulation model (ECHAM5; Roeckner et al., 2006). The ECHAM/MESSy atmospheric chemistry (EMAC) model is a numerical chemistry and climate simulation system that includes submodels describing tropospheric and middle atmosphere processes and their interaction with ocean, land and human influences (Jöckel et al., 2010).

⁴COSMO is the regional weather prediction model of the Consortium for Small-scale Modeling (Steppeler et al., 2003; Doms and Schättler, 1999) and the community model of the German regional climate research (Rockel et al., 2008). By implementing the MESSy interface, the COSMO model was expanded to a regional chemistry (climate) model (Kerkweg and Jöckel, 2012a; Mertens et al., 2016).

⁵The Community Earth System Model version 1.2.1 (CESM1) is a fully coupled climate model. “CESM is sponsored by the National Science Foundation (NSF) and the US Department of Energy (DOE). Administration of the CESM is maintained by the Climate and Global Dynamics Laboratory (CGD) at the National Center for Atmospheric Research (NCAR)” (cited from <http://www.cesm.ucar.edu/models/cesm1.0/>, last access date: 27 September 2016) (Hurrell et al., 2013; Baumgaertner et al., 2016).

²<https://www.fona-miklip.de/>

tegration phase of a simulation. For the unidirectional data exchange, the expanded INT2LM⁶ software was used to interpolate the data from the driving model grid to the target model grid. This software is a specialised software for the calculation of the initial and boundary data of the COSMO model. Therefore, a different software is required to interpolate the data from the finer to the coarser grid for the data sent from the client model to the server model. According to the MESSy philosophy of strict separation and generalisation, we therefore developed the new generic submodel GRID (v1.0), which is also documented in the present article. GRID can be used for all grid mapping operations required during a simulation. Nevertheless, in MMD, it is currently used for the parent-to-client coupling only.

In the next section, we describe the new developments within MMD. As the data mapping between the different grids is central to this further development of MMD, Sect. 3 introduces the newly developed GRID submodel, which provides the required mapping functionalities used for the remapping from the finer to the coarser model instance. Some examples for two-way data exchange are shown in Sect. 4. A brief runtime performance analysis of the model is presented in Sect. 5.

2 MMD v2.0

MMD is the coupling software performing the data transfer between two independent base models running within the same MPI environment. Appendix A of Kerkweg and Jöckel (2012b) provides an overview about different coupling approaches; especially the differences between internal and external coupling are discussed. Furthermore, Sect. 4 of Kerkweg and Jöckel (2012b) explains why MMD was chosen as coupling software between different MESSy base models. In summary, apart from the reasons named already in the introduction, MMD provides the best balance between fast data transfer and the possibility to integrate model-specific software, such as INT2LM, into the coupling procedure. INT2LM is the software provided by the German Weather Service (Deutscher Wetterdienst, DWD) for the calculation of the initial and boundary data for the regional COSMO model. This software was included in MMD (v1.0) as sub-submodel INT2COSMO.

INT2LM, and thus INT2COSMO, does not only include the interpolation routines to map the driving model fields to the regional model grid. It furthermore processes the external data used as input to the COSMO model and provides the calculation of additional fields required by the COSMO model, which are not necessarily provided directly by the driving model.

⁶See part V of the COSMO model documentation: <http://www2.cosmo-model.org/content/model/documentation/core/default.htm>; last access date: 29 September 2016.

The coupling was implemented following a client–server approach. Therefore, in MMD (v1.0), all routines and modules have been named server (serv) or client (clnt) in accordance to the model using them. In MMD (v2.0), the routines and modules have been renamed to parent and child instead of server and client. This was required, as the term server implies that this model is sending the data. As in MMD (v2.0), data are sent in both directions; the terms parent and child for the coarser and the finer model, respectively, are better suited.

MMD consists of two parts:

1. a library performing the data transfer, which is independent from the coupled models and
2. the part for data provision and processing implemented as MESSy submodels.

The library was extended by a few subroutines enabling the data transfer in both directions. The larger changes occurred in the MESSy submodels, as the data processing routines for the back transfer of the data had to be implemented. In the following subsections, an overview about the changes and additions made within these two parts of MMD are described. The MMD library manual and the MMD user manual in the Supplement provide all technical details about the implementation.

2.1 The MMD (v2.0) Library⁷

The MMD library manages the two-way data exchange between the different tasks of one EMAC and/or an arbitrary number of COSMO/MESSy instances as illustrated in Fig. 1. The configuration of the client–server system is defined in the Fortran95 namelist file `MMD_layout.nml` (which is written automatically by the run script). This namelist file contains the information about the overall number of model instances within the current MECO(n) setup (i.e. $n + 1$), the number of MPI tasks assigned to each model and the definition of the parent model of the respective model (for further details, see the MMD (v2.0) library manual in the Supplement). The library contains a high-level API for the data exchange between the different models. Figure 2 illustrates the functional principle of the MMD library.

During the initialisation phase, the exchange of information required by the parent from the child model, and vice versa, is accomplished by utilising the MPI routines `MPI_send` and `MPI_recv`. During the integration phase, data can be exchanged in both directions, i.e. from the parent to the child model, and vice versa. Point-to-point, single-sided, non-blocking communication is applied to exchange the required data between the different MPI tasks. “Checkpointing” (the technical term for “restarting”) is required (not

⁷The text of this section is adopted from the initial publication of the MMD library in Kerkweg and Jöckel (2012b).

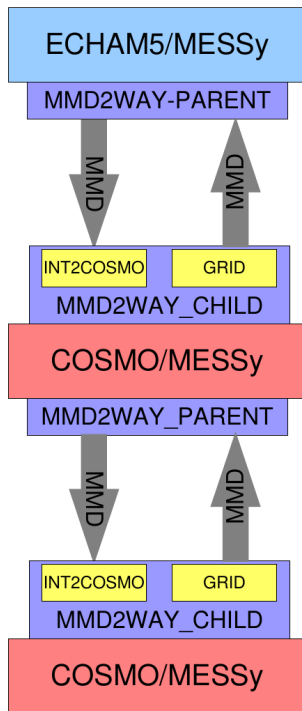


Figure 1. Illustration of the connection between the different MMD parts. The example is for a MECO(2) setup, i.e. EMAC with a cascade of two nested COSMO/MESSy instances.

only for climate simulations) to be able to continue a simulation after hardware failures and for branching off sensitivity studies, and last but not least, it is required to split a simulation into parts, fitting into the typical time limits of a job scheduler on a supercomputer. To enable check-pointing, one additional communication step occurs during the integration phase: for the synchronisation of the models with respect to the check-pointing, the parent model has to send the information regarding whether the simulation will be interrupted after the current time step. This data exchange is implemented as direct MPI communication using `MPI_send` and `MPI_recv`.

As the routine `MPI_alloc_mem`, used to allocate the memory (buffer) required for the data exchange, can only be used in C (and not in Fortran95), some parts of the MMD library are written in C; however, most parts are written in Fortran95 for consistency with the POINTER arithmetic used for the MESSy memory management (see Jöckel et al., 2010). The MMD library routines and their usage are described in detail in the MMD (v2.0) library manual (see the Supplement).

2.2 The MESSy submodel MMD2WAY

In addition to the library, the MMD software comprises a regular MESSy submodel as the “wrapper”. This submodel provides and processes the data transferred by the library.

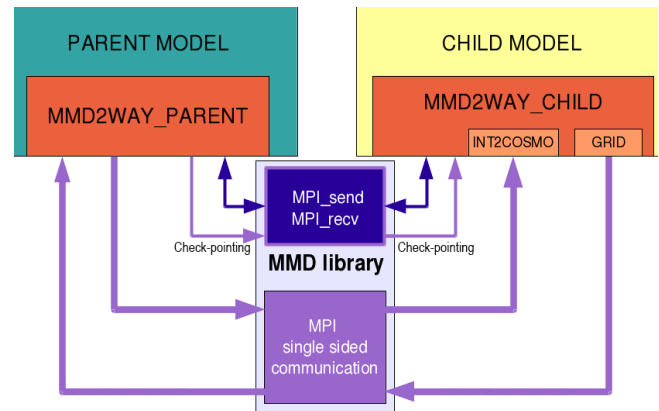


Figure 2. Illustration of the communication managed by MMD between a parent and a child model. Dark violet colours indicate data flow during the initial phase, while purple indicates the data flow during the integration phase.

MMD (v1.0) contained two MESSy submodels: one for the server (MMDSERV) and one for the client (MMDCLNT). Here, the server controls the timing of the client model and “serves” the data, which is processed by the client. In the new MMD version, the client also provides data to the server model. The only remaining difference between the models with respect to the data transfer is the time control of the models. Therefore, the server and client models have been renamed to parent and child models, omitting the impression that only the server acts as “data server”. Consequently, the new MESSy submodel consists of two submodels: MMD2WAY_PARENT and MMD2WAY_CHILD. These sub-submodels provide the same functionalities for the one-way online coupling as MMDSERV and MMDCLNT in MMD (v1.0), respectively, as described in detail in Part 2 of the MECO(n) model documentation (Kerkweg and Jöckel, 2012b). In the initial phase of a model simulation, the following measures take place.

- The parent imprints its time settings on the child model; these are end date, restart trigger date and, at the very beginning of a model instance, the (re-)start date as start date of this instance.
- The field names required from the parent are read from the `&CPL_CHILD_ECHAM` or `&CPL_CHILD_COSMO` namelist in the `mmd2way.nml` namelist file for ECHAM or COSMO as parent models, respectively. The names of the parent fields are sent to the parent, and in both models pointers to the respective data fields and dimension information are set.
- The exchange matrix, the so-called “index list”, is set up. This index list provides the information for which grid box (index pair (i_p, j_p) on which parent parallel task (PE_p) exchanges data with which child grid box (i_c, j_c) on which child parallel tasks (PE_c). For this, the

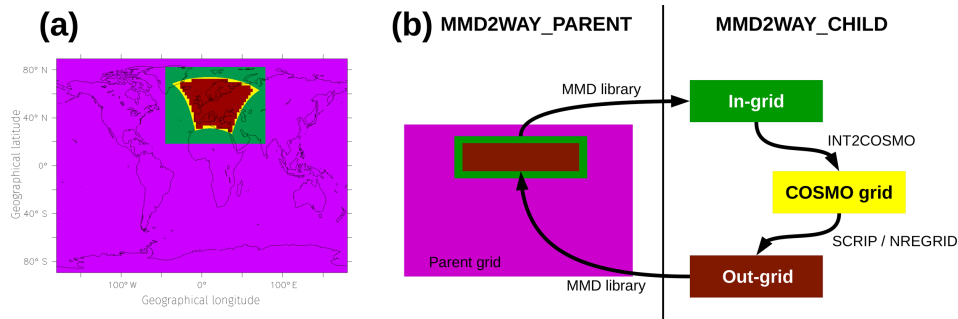


Figure 3. Relation of the different grids: EMAC grid in pink, the in-grid and out-grid defined by MMD2WAY_CHILD in green and brown, respectively, and the COSMO/MESSy model grid in yellow. Panel (a): position of the different grids relative to each other, for the example of a European domain. Panel (b) illustration of data flow between the different grids.

child model has to define an “in-grid”. This is a sub-area of the parent grid (i.e., it has the same rotation and the same mesh size) and completely overlays the child grid. Figure 3 illustrates the relation between the different grids. Afterwards, the data are transformed from the in-grid to the child grid using the expanded version of the preprocessing software INT2LM for the COSMO model (INT2COSMO; see Kerkweg and Jöckel, 2012b, for further explanations).

During the integration phase, the MMD library sends the data from the parent model to the child model. The child model calculates the required initial and boundary conditions from the parent data and transforms additional data to the child grid.

This functionality for the one-way coupling is kept the same in MMD (v2.0). In addition, MMD2WAY_PARENT and MMD2WAY_CHILD have been expanded for the data transfer from the child to the parent model. For most functionalities of the one-way coupling, a counterpart for the data transfer in the other direction could be implemented by keeping the same logic. Thus, a namelist (&CPL_PAR_CHILD) in the parent model namelist file `mmd2way.nml` determines which fields are exchanged between the child and the parent models. In the initial phase of a model simulation, this information is transferred to the child model. Both models set pointers to their corresponding data objects. Again, the child model has to define a grid, which is a subpart of the parent model grid (called “out-grid”), and it has to perform the data transformation from the child model grid to the out-grid. The decision to transform the data within the child model was taken in order to minimise the amount of data to be transferred between the models: as the parent model grid will usually be coarser resolved, data on this grid are exchanged via MMD. For the transformation from the child to the out-grid, the newly written MESSy infrastructure submodel GRID is used (see Sect. 3). First, the data are remapped horizontally, before the vertically remapping proceeds in an extra step. For the time being, only conservative remapping, as provided by GRID, is utilised as a horizontal transforma-

tion method in MMD2WAY_CHILD. As the COSMO model uses a staggered Arakawa-C grid, the wind components need to be interpolated to the grid midpoints prior to the horizontal remapping for the COSMO–EMAC coupling, as the EMAC wind components are defined on the grid midpoints. For the COSMO–COSMO coupling, the wind components are interpolated directly between the staggered grids; i.e. they are always defined on the box edges.

The vertical remapping depends on the parent model. If EMAC is the parent model, NREGRID is used for the vertical remapping of the fields. In this case, data of a non-hydrostatic model with a fixed vertical geometry need to be converted for a hydrostatic model using hybrid pressure coordinates. The vertical coordinate in the COSMO model is defined as a pseudo-hybrid pressure axis. For this, the hybrid coefficient calculation as provided by INT2LM is used as the input vertical axis to the vertical interpolation via NREGRID. Furthermore, the new surface pressure in the EMAC model is approximated by an iterative calculation of the pressure, temperature and humidity (vapour, liquid water and cloud ice) vertical profiles. For the vertical interpolation of the COSMO–COSMO coupling, the INT2COSMO spline interpolation is used.

The interpolated data are sent to the parent model, where they are subsequently weighted (if requested) and assigned to the target parent model fields. For the utilisation of the child data by the parent model, two methods are distinguished:

“0”: For purely diagnostic applications, the field is only used as input to the parent model; i.e. this field is created by the parent coupling submodel and thus independent of other model data objects. In this case, the memory is allocated by MMD2WAY_PARENT, and the transferred field is copied to this memory without any further modifications.

“1”: For feedback from the finer- to the coarser-resolved model instance, the exchanged field is used to directly modify a parent model field. Therefore,

no additional memory needs to be allocated by MMD2WAY_PARENT.

Using method 1, there are two options for modifying a prognostic variable of the parent model:

- the value of the variable can be changed directly or
- the tendency of the variable can be modified.

For all non-prognostic variables, only option (a) is possible.

For both options, a weighting between the original value of the parent field (P) and the child model field (C) is applied:

$$P(i, j, k, tlev) = P(i, j, k, tlev) \cdot (1.0 - f_{mn} \cdot f_{vw}(k) \cdot f_w(i, j)) + C(i, j, k, 1) \cdot f_{mn} \cdot f_{vw}(k) \cdot f_w(i, j). \quad (1)$$

Here, i and j are the indices along the horizontal dimensions, k is the index along the vertical dimension, and $tlev$ indicates (if applicable) the respective time level.

The different weight coefficients are as follows:

- f_{mn} is the relaxation strength. Its value is set in the parent namelist individually for each field.
- f_{vw} is a vertical weight function. It depends on the vertical index k . In most cases, the domain coupled back from the child model does not cover the full height of the parent model⁸. To avoid artificial discontinuities in the data fields, a weight function is required, which gradually decreases from 1 in the core domain to 0 towards the edge of the domain. The weight function is implemented as a cosine function:

$$\begin{aligned} f_{vw}(k) &= 0.0 \quad \text{for } k \leq k_{\min} \\ f_{vw}(k) &= \cos\left(\frac{\pi}{2.0} \cdot \left(1 - \frac{k - k_{\min} - 1}{n_k - 1}\right)\right)^2 \\ &\quad \text{for } k_{\min} < k \leq k_{\min} + n_k \\ f_{vw}(k) &= 1.0 \quad \text{for } k > k_{\min} + n_k. \end{aligned} \quad (2)$$

In Eq. (2), it is assumed that the vertical index k increases from top to bottom; k_{\min} is the height index of the top of the child domain; n_k is the number of vertical layers the cosine function should cover.

- f_w is the horizontal weight function. This weight function is required to avoid artificial discontinuities at the borders of the area, where the fields are relaxed to the child variables. Currently, the user can choose between three different implementations by namelist:

“0”: f_w is set to 1 everywhere in the child domain. This option is for testing only, as it may lead to artificial discontinuities in the data.

“1”: f_w is implemented as the sum of two cosine functions:

$$f_w(i, j) = 1.0 - (\cos(x)^e + \cos(y)^e) \quad (3)$$

with $x = \pi \cdot \frac{i}{i_{\max}}$; $y = \pi \cdot \frac{j}{j_{\max}}$.

i_{\max} and j_{\max} are the number of grid points in the two horizontal directions, respectively. The exponent e is set by namelist. Its default value is 14.

“2”: f_w decreases in the form of a cosine from 1 in the domain inner part to 0 at the borders of the coupled domain. The width of the damping zone is determined by a namelist parameter `damprel`. Its valid range is $[0, 0.5]$. This number determines the relative width of the damping zone. If, for example, a model’s domain consists of 100 grid boxes in the x direction (index i) and of 50 grid boxes in the y direction (index j), for `damprel` = 0.2, the damping zone is in the x and y directions 20 and 10 grid boxes wide, respectively.

All these weight functions are defined on the child grid. They are once, during the initialisation phase, transformed in the same way as the data and sent to the parent model for application during the integration phase. Figure 4 displays the different weight functions for a domain over Europe. The upper row shows the weight functions as defined on the child model grid. Note that the coupled domain is smaller than the child domain (with the exception of $f_w = 0$). This is because the damping zone of the regional model itself should not be coupled back to the parent model, as this is, for two-way coupled variables, directly influenced by the parent model, and thus spurious damping or amplifications could occur. The lower row of Fig. 4 shows the same weight functions after the transformation to the parent grid.

If the tendency is subject to change (i.e. method 1, option b is used), first the current value of the parent field (P) needs to be calculated from the values at the previous time step plus the tendencies of the current time step. This field is modified according to Eq. (1) and an additional tendency is calculated from the difference between the parent fields before and after the modification.

3 The generic MESSy submodel GRID (v1.0)

Due to the increasing complexity of Earth system models, grid transformations at runtime of the model (e.g. remapping from an atmosphere grid to a higher-resolved land grid, and vice versa) are more and more commonly required. To avoid several implementations throughout the code, the MESSy in-

⁸For example, in the case of the COSMO/MESSy model, only the data below the damping layer should be coupled back.

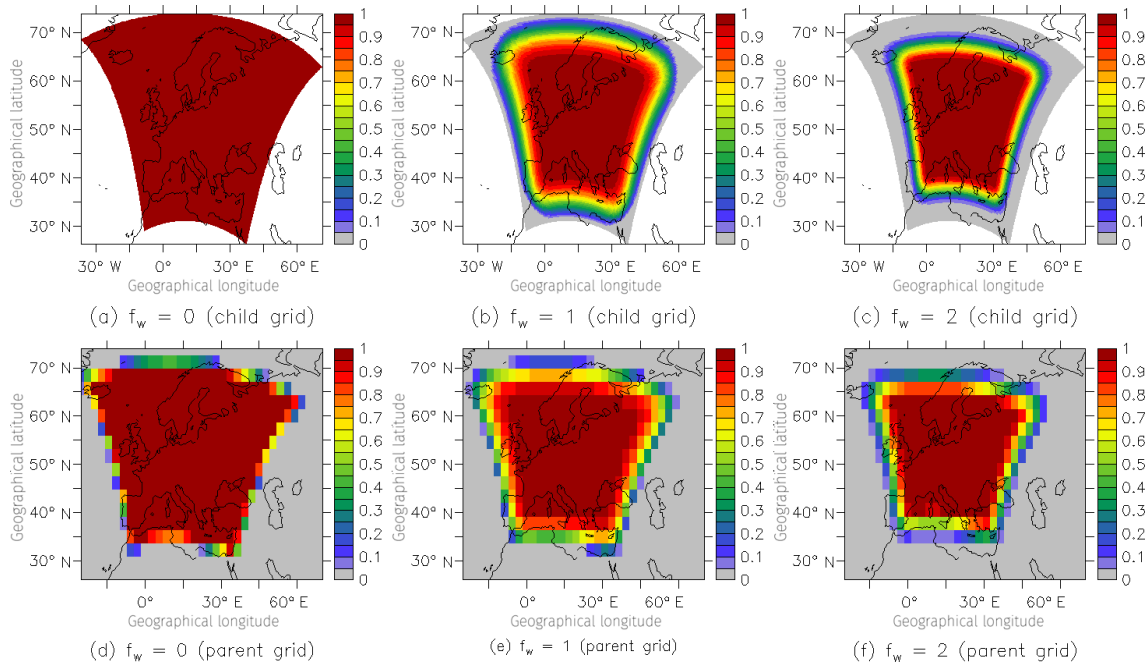


Figure 4. Weight functions (f_w ; see Sect. 2.2, p. 9) for the different weight types. Upper row: weight functions as calculated on the child model grid. Lower row: weight functions after transformation to the parent model grid.

infrastructure submodel GRID⁹ was implemented providing an online transformation (remapping) functionality. The common grid processing functionality provided by GRID includes the routines for grid definition, grid modification and the transformation between different grids. Implementation as one important part of the model infrastructure simplifies the maintenance and expansion of the functionality, because it is utilised jointly by all model components. As the infrastructure module is written in a general way, performance optimisation or additional grid types, transformation algorithms, etc. can be implemented straightforwardly.

Currently, two horizontal grid types are treated by GRID:

- rectangular grids which are orthogonal in geocoordinates (rectilinear grids, e.g. the ECHAM grid) and
- rectangular grids which are orthogonal in another reference system, i.e. curvilinear grids (e.g. the COSMO grid).

The 3-D spatial grids consist of one of the abovementioned horizontal grids and a vertical dimension. The vertical axis can be defined in different ways, e.g. as a height- or pressure-based coordinate.

For the transformation between different grids, different methods are provided for conservative remapping and (not necessarily conservative) interpolation.

⁹The names of MESSy submodels are written in capital letters throughout the article, even though they are not necessarily acronyms.

Moreover, the MESSy infrastructure submodel GRID code was written to be

- well structured to flexibly support expansions,
- as simple as possible, to keep it maintainable,
- efficient, i.e. to show good runtime performance (therefore, it must work in a parallel environment and scale appropriately), and it is
- designed to cause an as small as possible memory footprint during operation.

For a grid transformation, first the source and the target grid need to be defined. Second, the remapping of data between these grids can be calculated. The following two subsections give an overview of the functionalities provided by GRID (v1.0). Their implementation in GRID is organised as follows:

1. The SMCL of GRID provides a unified interface for the definition of all grids required in all MESSy submodels. It is implemented as a Fortran95 structure, which contains all required information of a grid in a generalised way.
2. The sub-submodel GRID_TRAFO provides the interface routines to use this grid information for the transformation between the different grids. GRID_TRAFO

utilises third-party grid transformation codes: currently NREGRID¹⁰ (Jöckel, 2006) and SCRIP (Jones, 1999).

The GRID user manual in the Supplement provides detailed information about the usage of the GRID submodel.

3.1 The submodel core layer of GRID

Earth system models usually define grids in spherical geometry. Two different grid types are distinguished in GRID v1.0: (1) rectangular grids, which are orthogonal in geocoordinates and (2) curvilinear grids. The implementation of non-rectangular, unstructured grids (e.g. the ICON grid) is ongoing.

Most of the internal data types of GRID follow the netCDF data model. The hierarchical data structures follow mainly those of NCREGRID (Jöckel, 2006). The definition of the Fortran structure, which contains all components required for the definition of a geohybrid grid, was extended and generalised for the usage in GRID. The geohybrid grid, as defined by Jöckel (2006), consists of a horizontal grid space, which comprises geographical latitude and longitude of the mesh vertices and/or centres. For different types of grids, different structure components for the definition of the horizontal grid are specified. The vertical grid space is defined in analogy to the hybrid pressure level definition. Depending on the setting of the coefficients and of the reference and surface pressure, the vertical axis can be defined as one of (1) hybrid pressure axes, (2) constant pressure axes, (3) constant height axes or (4) sigma levels. More details can be found in the GRID user manual in the Supplement.

The GRID SMCL routines also comprise subroutines for the handling of the grid structures, i.e. routines for initialising, copying, importing, exporting and printing a variable of the grid structure type. Beyond that, routines necessary for defining a grid, storing it in a concatenated list, locating an already defined grid within this list and comparing grids are part of the GRID SMCL. During a model simulation, the definition of an arbitrary number of georeferencing grids and the transformations between those grids is possible.

3.1.1 GRID_TRAFO

The main intention of the GRID_TRAFO submodel is to provide routines for the transformation of gridded geolocated data. GRID_TRAFO comprises NREGRID (Jöckel, 2006), the standard remapping tool in EMAC and the SCRIP¹¹ software (Jones, 1999). While NREGRID is restricted to mapping between orthogonal 2-D or 3-D grids, SCRIP provides transformations to/from curvilinear or unstructured grids. Here, we use grid “transformation” as a generic term for both

conservative remapping (or “regridding”) as well as for (not necessarily conservative) interpolation.

The geohybrid grid structure provides all information required for the grid conversion. As each remapping software (NREGRID and SCRIP) relies on its specific grid information structure, GRID_TRAFO additionally provides routines to extract these as required by the respective mapping software; i.e. it provides the “middleware” or acts as “wrapper” for the established mapping software. The remapping algorithms automatically apply the correct conversion routines, depending on the associated structure components. While the core mapping algorithms differ, GRID_TRAFO provides unified interfaces for the conversion between different grids. Additional interpolation schemes can be easily added, if required in the future.

The details are explained in the GRID user manual, which is part of the Supplement.

3.1.2 NREGRID

The remapping algorithm NREGRID is a recursive algorithm, which is applicable to arbitrary orthogonal (including curvilinear) grids of any dimension. It is used for the re-discretisation of “gridded” geoscientific data between n -dimensional (usually $n = 2$ or 3) orthogonal grids. The conservative re-discretisation of extensive or intensive variables is based on the calculation of the overlap (area or volume) matrix between source and destination grid boxes. For orthogonal grids, these overlap matrices can nicely be calculated recursively, since the overlap area/volume is zero as soon as at least the overlap interval along one axis (dimension) is zero. Since the recursive nature of this algorithm limits its application to orthogonal grids, it cannot be applied for re-discretisations between the (in geographical coordinates orthogonal) Gaussian grid of ECHAM5 and the rotated (in geographical coordinates non-orthogonal) COSMO grid.

Details about the NREGRID algorithm have been published by Jöckel (2006).

3.1.3 SCRIP

As NREGRID is limited to the remapping between equally oriented orthogonal grids, the implementation of an algorithm able to transform between different curvilinear or even unstructured grids became necessary. For this, the SCRIP software¹² (Jones, 1999) version 1.4 provided by the Los Alamos National Laboratory has been utilised. SCRIP “is a software package used to generate interpolation weights for remapping fields from one grid to another in spherical geometry. The package currently supports four types of remappings. The first is a conservative remapping scheme that is ideally suited to a coupled model context where the area-

¹⁰Note: the infrastructure submodel previously used in EMAC is named NCREGRID, while the remapping algorithm itself is called NREGRID.

¹¹Spherical Coordinate Remapping and Interpolation Package.

¹²<http://oceans11.lanl.gov/trac/SCRIP> (last access: 5 February 2018). The official link named in the SCRIP user guide (<http://climate.acl.lanl.gov/software/SCRIP>) is not available anymore.

integrated field (e.g. water or heat flux) must be conserved. The second type of mapping is a basic bilinear interpolation which has been slightly generalised to perform a local bilinear interpolation. A third method is a bicubic interpolation similar to the bilinear method. The last type of remapping is a distance-weighted average of nearest-neighbour points. The bilinear and bicubic schemes can only be used with logically rectangular grids; the other two methods can be used for any grid in spherical coordinates.” (Quoted from the SCRIP user guide, Introduction; Jones, 1998.)

3.1.4 Application of GRID in MMD2WAY_CHILD

The COSMO model uses a rotated grid and the orientation between two COSMO model instances or the COSMO and the EMAC model is arbitrary. As NREGRID requires equally oriented orthogonal grids, it is not applicable in MMD2WAY. Sadly, SCRIP provides only algorithms for horizontal grid transformation. Thus, two steps are required for the remapping of 3-D data fields.

Usually the biggest challenge in two-way nesting of two atmospheric models is the height correction required due to the differently resolved orographies of the child and parent models. Thus, it seems to be a natural choice to first re-grid horizontally and to perform the vertical regridding intertwined with the height adjustment as a second step. For child-to-parent coupling, first horizontal remapping via SCRIP is conducted. In a second step, the vertical remapping is performed using NREGRID for COSMO–EMAC coupling. For the COSMO–COSMO coupling, it was decided to use the INT2COSMO spline interpolation.

Offhandedly, one might expect that INT2LM could be completely replaced by GRID, but this is not trivial. INT2LM provides much more functionalities than just remapping. It reads and processes the external data required as input for the COSMO model (especially for the initialisation of the model). Moreover, it performs some field adjustments concerning inconsistencies between the land–sea mask of the COSMO model and the incoming data. Therefore, it is not possible to completely eliminate INT2LM. One could, however, indeed exchange the horizontal and vertical interpolation routines. We started to test this, but in the first place performance (with respect to the results) of the child model was downgraded. The main reason is that INT2LM does not only perform a vertical remapping but preserves structure and characteristics of the boundary layer (i.e. up to 850 hPa) by moving it to the height of the target orography and remapping only the remaining part of the vertical column. This procedure, as it is implemented in INT2LM, is not reversible and thus introduces spurious effects for different orographies, which are always present because of the different horizontal resolutions of the nested models. Anyhow, for the offline nested COSMO model, this is the preferred way, as this makes physically more sense compared to a simple vertical interpolation. Unfortunately, it causes inconsistencies, as it

is not reversible. The Supplement contains an example illustrating the deviations caused in the tracer profiles due to this height adjustment procedure.

3.2 The base model interface layer of GRID

The backbone of each model is its grid; e.g. for an atmospheric model, the horizontal domain is given by a definition of the geographical longitudes and latitudes of the model’s grid midpoints and the grid corners. The vertical dimension is usually defined by a height- or pressure-based coordinate. As this grid (hereafter denoted as “basegrid”) is the reference for most submodels and processes, the basegrid is defined in the base model interface layer (BMIL) for usage in all MESSy submodels. In the case of MMD2WAY, MMD2WAY_CHILD utilises the basegrid as source grid for the mapping to the out-grid as target grid.

4 Example applications using the two-way coupled MECO(n) system

Keeping the remaining issues in mind (Sect. 3.1.4), the current implementation, nevertheless, allows already for some useful applications. For instance, data can be transferred online from the finer to the coarser grid to be compared on the coarse grid. A simple example is shown in Sect. 4.1. Additionally, diagnostic tools can be used to interpret global and regional model results consistently, e.g. for radiative forcing, which is consistently determined only if calculated with the same radiation code. Section 4.2 illustrates this by utilising the radiative forcing calculations of EMAC.

As discussed in Sect. 2.2, the two-way coupling of prognostic variables is technically implemented in the MMD2WAY submodel. Thus, Sect. 4.3.1 gives an example for an EMAC–COSMO/MESSy coupling, where dust tracers are two-way coupled. Finally, Sect. 4.3.2 shows the full dynamical two-way coupling of two COSMO/MESSy model instances located over the Atlantic Ocean (i.e. over flat terrain) using the same height coordinates.

4.1 Simple examples of added value through aggregated subgrid-scale information

Depending on their resolution, only certain processes can be resolved by atmospheric models, whereas others have to be parameterised. Naturally, smaller-scale models can resolve more processes explicitly. It is still under debate whether or not the aggregation of the subgrid-scale information provided by the smaller-scale model to the larger-scale model constitutes an added value for the larger-scale model. This issue might be answered with the help of two-way coupled applications. Most probably, the answer will differ for different processes. For some dynamical processes, e.g. the generation of Rossby waves or hurricanes (see Sect. 4.3.2), the upscaling

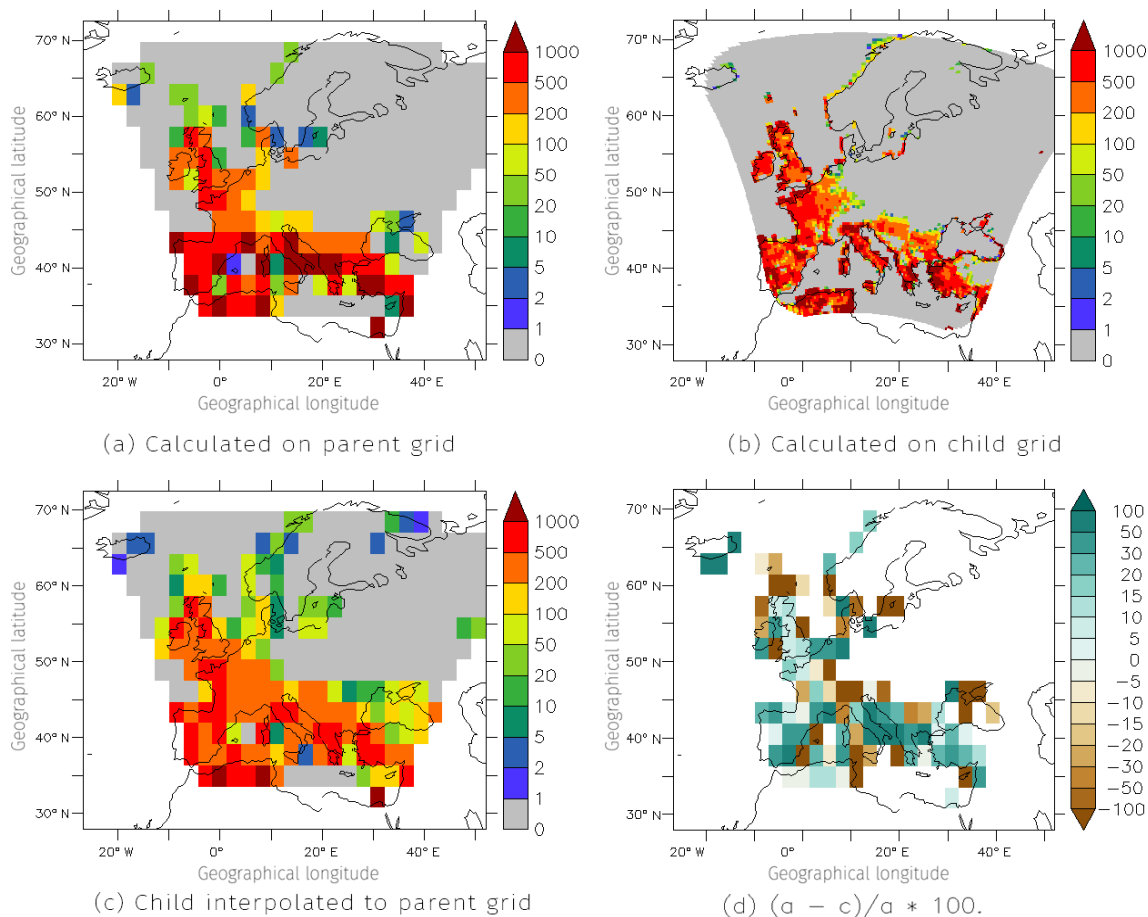


Figure 5. Biogenic NO emissions flux in $\text{pg m}^{-2} \text{s}^{-1}$ for one distinct date in January 2003. The data have been masked to show only the coupled region.

might result in an added value, as these phenomena originate from smaller-scale perturbations.

For chemistry models, especially the treatment of emissions is of interest. On the one hand, emissions, which depend on soil properties and/or on prognostic variables in the model (the so-called online emissions, because they are calculated during the simulation), can substantially differ between models with different resolution. One example is dust emissions, which depend on the 10 m wind speed, soil properties and soil moisture (see Sect. 4.3.1). On the other hand, it is normally assumed that even point and line emissions are instantly mixed within the grid box into which they are emitted. This leads to a higher dilution in larger-scale models. Especially in highly polluted regions, or more generally near emission sources, this might influence the simulated chemical regime, as atmospheric chemistry is highly non-linear.

Figure 5 illustrates, as a simple example, the resolution effect on online calculated nitrogen oxide (NO) soil emissions (Kerkweg et al., 2006b). These emissions strongly depend on the soil properties and thus differ substantially between the models.

Figure 5a depicts the NO emission flux as calculated on a global EMAC model grid of T42 ($\approx 2.8^\circ$) resolution. Figure 5b shows how these emission fluxes look on a COSMO/MESSy grid with 0.36° horizontal resolution. If COSMO/MESSy was two-way coupled into EMAC, and EMAC was using the NO emissions coupled from COSMO/MESSy instead of calculating them itself, the emissions aggregated from COSMO/MESSy to the EMAC grid would be as in Fig. 5c. Figure 5d depicts the difference in percent between the emissions directly calculated by EMAC (Fig. 5a) and coupled back from COSMO/MESSy (Fig. 5c). Naturally, the emission fluxes on the COSMO/MESSy grid show much finer structures as a result of the finer grid and therefore finer distributed soil properties. However, the largest differences between the upscaled (Fig. 5c) and EMAC-calculated (Fig. 5a) emission fluxes occur at the coast lines (Fig. 5d), which is mostly due to the much finer-resolved land–sea mask in the smaller-scale model. The NO emission flux integrated over the coupled domain is 3.29 and $2.63 \text{ kg(NO) s}^{-1}$ for the parent and the child model, respectively. Thus, the differences in the soil properties of

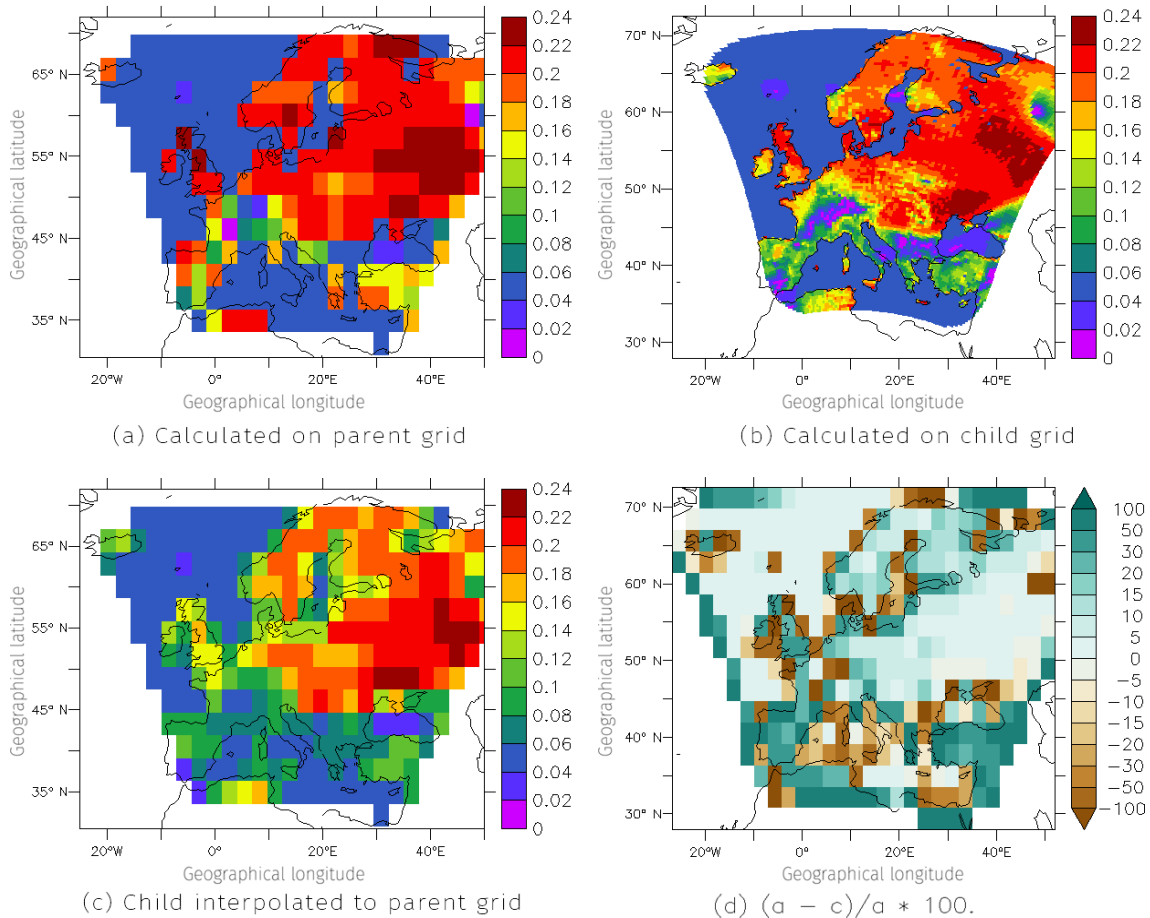


Figure 6. Ozone dry deposition velocities in cm s^{-1} for one distinct date in January 2003. The data have been masked to show only the coupled region.

the two models account for a difference of $0.66 \text{ kg (NO) s}^{-1}$. The integrated NO emission flux regridded from the child to the parent grid is $2.78 \text{ kg (NO) s}^{-1}$, providing an emission flux lower by $0.51 \text{ kg (NO) s}^{-1}$ compared to the directly calculated integrated emission flux. The difference of $0.15 \text{ kg (NO) s}^{-1}$ between the flux in the regional domain and its integral over the global domain simply results from the not fully congruent areas, over which the integrals are taken in the rotated domain and the global domain, respectively.

As a second example, the dry deposition velocities for ozone are displayed in Fig. 6 (Kerkweg et al., 2006a). The features discussed for the previous example appear here as well. Additionally, the ozone dry deposition velocities calculated by COSMO/MESSy are much more evenly distributed in the Mediterranean region, while they are slightly but systematically smaller over eastern Europe, which is most probably due to different soil properties and also due to the different turbulence schemes employed by the two base models.

4.2 Use of specific diagnostic tools: radiative forcing

To evaluate the radiative forcing for two MECO(n) instances consistently, the MESSy submodel RAD (Dietmüller et al., 2016) is used for the calculation of the radiative forcing of a COSMO/MESSy instance online coupled to the EMAC model. As COSMO/MESSy and EMAC use different radiation schemes, this is one way of conducting a consistent comparison.

Here, results are shown from simulations using a setup as published by Mertens et al. (2016). A COSMO/MESSy instance over Europe (0.44° resolution) was coupled to the global EMAC domain. The ozone field calculated by COSMO/MESSy was sent back to EMAC using MMD. However, the ozone field coupled back from COSMO/MESSy is zero or undefined outside of the coupled region, i.e. the horizontal and vertical relaxation areas as well as those parts of the globe, which are not covered by the COSMO/MESSy instance. Therefore, the uncovered points are filled with the ozone field calculated by EMAC; as for the calculation of the radiative flux in EMAC, global non-




Figure 7. Differences (COSMO/MESSy minus EMAC) of the radiative fluxes (shortwave (sw), longwave (lw) and net) averaged over 35–60° N and –10–30° E for July 2008. (a) REF simulation; (b) SENS simulation.

zero fields must be fed into the diagnostic routine. With this ozone field, a second, diagnostic radiation call is performed using RAD. Two simulations are investigated:

- REF: EMAC and COSMO/MESSy are using the same emission data set (MACCity, Granier et al., 2011).
- SENS: EMAC uses the MACCity inventory, while COSMO/MESSy applies a Deutsches Zentrum für Luft- und Raumfahrt (DLR) specific inventory.

Here, the differences (COSMO/MESSy minus EMAC) of the radiative fluxes, area-averaged over central Europe (35–60° N; –10–30° E) for July 2008, are compared. Figure 7a shows the vertical profiles of the differences of the clear-sky radiative fluxes applying the same emissions (REF). The larger ozone values as simulated by COSMO/MESSy compared to EMAC near the tropopause lead to a positive radiative flux difference in the longwave as well as in the shortwave bands around 200 hPa. If the emissions in COSMO/MESSy are changed (SENS; Fig. 7b), lower ozone values are simulated by COSMO/MESSy compared to EMAC up to around 800 hPa. These lower values lead to a negative difference of the longwave radiative fluxes compared to EMAC.

4.3 Two-way coupling of prognostic variables

Next, we show two examples for the coupling of 3-D prognostic variables. First, the two-way coupling of dust tracers between EMAC and one COSMO/MESSy instance is shown. Secondly, all dynamical variables of two COSMO/MESSy instances are two-way coupled to demonstrate the potential of the two-way coupling. To avoid the adjustment of the orographies, the smaller COSMO/MESSy instance is predominantly located over the ocean.




Figure 8. Time series of dust emission (Gg) in the EMAC (black) and in the COSMO/MESSy (red) instance.

4.3.1 Dust

Dust emissions are very sensitive to the model resolution, as they depend on the soil type and the wind velocity. Typically, dust emission schemes are developed for a specific model resolution. They include scaling factors to adapt them as well as possible to other resolutions.

In our example, we use a MECO(1) setup, coupling the dust tracers of a COSMO/MESSy instance with 0.36° horizontal resolution back to the EMAC model in T63 spectral resolution.

Figure 8 shows the dust emission fluxes integrated over a domain ranging from 60° W to 60° E and from 45° S to 45° N. Due to different soil type distributions and higher wind maxima in the COSMO/MESSy instance, the latter produces much higher dust emission fluxes, as the simulation was performed without any resolution-dependent tuning of the emission scheme. These higher emissions are reflected also in the horizontal dust column mass (mg m^{-2}) distribution. Figure 9 displays the dust column mass for 6 March 2004. Figure 9a shows the dust column mass (in mg m^{-2}) in the COSMO/MESSy instance, Fig. 9b the EMAC dust column mass in the one-way coupled simulation and Fig. 9c the result of the two-way coupled simulation. Obviously, the COSMO/MESSy instance exhibits much finer structures than both EMAC instances. However, the maximum present in the COSMO/MESSy instance is much better represented in the EMAC two-way simulation, as intended.

This simulation contains still a small error with respect to the vertical distribution of the dust, as the height adjustment for the orography is not yet consistent for the two coupling directions (see Sect. 3.1.4). Nevertheless, this example illustrates the potential of the two-way coupling to improve the

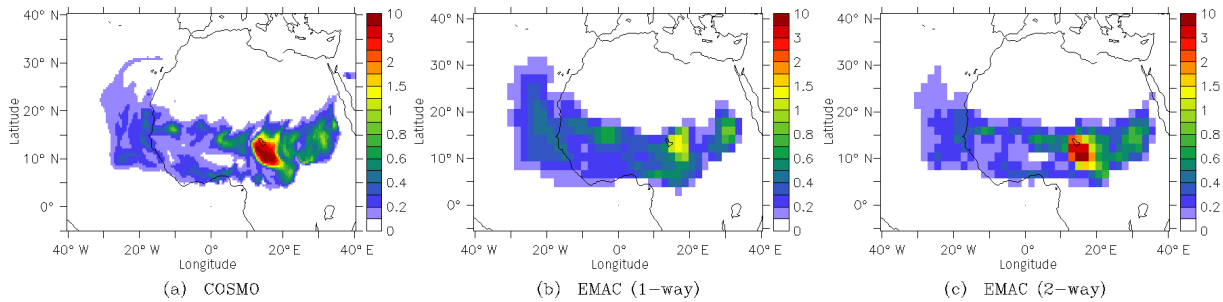


Figure 9. Dust column mass (mg m^{-2}). Shown is an instantaneous value for 6 March 2004, 00:00 UTC. (a) For COSMO/MESSy; (b) for EMAC; (c) for two-way coupled EMAC, thus influenced by the COSMO/MESSy instance.

coarse representation of quantities, which are determined by smaller-scale features.

4.3.2 Hurricanes

Tropical cyclones (TCs), developing on the west coast of Africa over the tropical east Atlantic, are known to be precursors for hurricanes causing damages in the US (e.g. Ike in 2008; Dean in 2009) or over Europe (e.g. Helene in 2006; Katia in 2011). Those TCs often originate as disturbances of the African jet stream, so-called African easterly waves (AEWs) over the African continent. In the case of suitable conditions over the Atlantic, these AEWs have the potential to develop into TCs and finally into hurricanes. Therefore, forecasting the development, track and intensity of hurricanes requires both a high model resolution to capture all the multi-scale interactions, prerequisite for the development of the initial TC, and a huge model domain capturing the African continent as well as the Atlantic Ocean (e.g. Rappaport et al., 2009; Schwendike and Jones, 2010).

As an example, the development of a hurricane named Isaac is analysed here. It originated in September 2000 as a TC from an AEW on the west coast of Africa. The National Hurricane Center classified it as a hurricane for the first time on 23 September, 12:00 UTC, before it reached maximum intensity as a category 4 hurricane on 28 September, 18:00 UTC (<https://coast.noaa.gov/hurricanes/>). Afterwards, its track turned to the north-east and after extratropical transition the system reached Great Britain readily identifiable by strong wind gusts 2 days later.

To demonstrate the potential of the dynamical two-way coupling between two COSMO/MESSy instances, a MECO(2) setup is applied to simulate the development of Isaac. This means that two COSMO/MESSy instances, varying in horizontal resolution, time step length and model domain, are coupled to the global EMAC model. The finer-resolved ($0.11^\circ \approx 12 \text{ km}$) COSMO/MESSy instance is driven by the coarser-resolved ($0.22^\circ \approx 25 \text{ km}$) COSMO/MESSy instance. Initial and boundary data for the coarser COSMO/MESSy instance are transformed from EMAC ($T106 \approx 120 \text{ km}$). Since this study focuses on the spe-

cific development of Isaac, a weak nudging of four prognostic variables (temperature, divergence, vorticity and the logarithm of surface pressure) towards ECMWF analysis data is applied for EMAC (as described by Jöckel, 2006) during the first 2 weeks after the start of the simulation (15 September, 00:00 UTC). Once the hurricane leaves the model domain of the finer-resolved COSMO/MESSy instance (29 September, 00:00 UTC; Fig. 11), the nudging is switched off and the EMAC instance is completely unconstrained afterwards. In the event that the COSMO/MESSy instances are two-way coupled, the dynamical information from the finer-resolved instance, comprising the temperature (T), the wind velocities (U, V, W), the pressure deviation from the reference atmosphere (PP) and moisture (QV, QC, QI) is fed back to the coarser instance. The domains covered by the model instances are shown in Fig. 11 (grey and blue areas).

In Figs. 10 and 11, the results obtained with the two-way coupled COSMO/MESSy instances (Figs. 10b and 11b) are compared to those of the one-way coupled system (Figs. 10a and 11a). Deviations are validated using the National Hurricane Centers HURricane DATabases (HURDAT) data set (Landsea and Franklin, 2013), which is part of the International Best Track Archive for Climate Stewardship (Knapp et al., 2010, IBTrACS (v03r04), available under <https://doi.org/10.7289/V5NK3BZP>).

Although EMAC is nudged during the genesis phase, the development of Isaac is not captured by EMAC, as there is no pressure decrease visible in the time series of minimum sea level pressure (SLP_{\min} ; Fig. 10, black contour). In contrast, Isaac initially originates in both COSMO/MESSy instances (Fig. 10, blue and red contours), independent of the coupling strategy and horizontal resolution and approximately at the correct time (23 September) compared to the best track estimate (Fig. 10, green contour). However, there are strong differences comparing the one-way and two-way coupled instances during the ongoing development of SLP_{\min} : while the final intensification of Isaac simulated with the one-way coupled instances does not start before 26 September, the two-way coupled instances are able to capture the initial decrease. Even though the intensity of Isaac in the one-way coupled simulations coincides better with the reference on




Figure 10. Time series of SLP_{min} (in the area of 10° – 50° N and 65° – 25° W) in the one-way (a) and two-way (b) coupled simulations for EMAC (black), COSMO/MESSy $_{0.22^{\circ}}$ (red) and COSMO/MESSy $_{0.11^{\circ}}$ (blue) based on 6-hourly data. The best track intensity from the National Hurricane Centers HURricane DATABases (HURDAT) is shown as a reference (green). EMAC is nudged until 29 September (dashed line; see text for details).

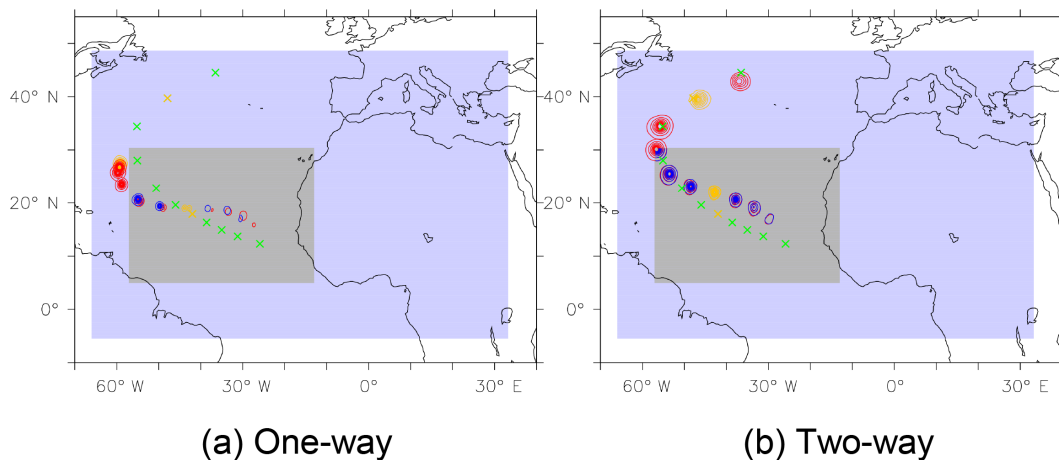


Figure 11. Track position for hurricane Isaac in the one-way (a) and two-way (b) coupled simulations. The daily SLP (less than 1005 hPa, 5 hPa intervals, in the area of 0° – 50° N and 65° – 25° W, starting on 23 September, 00:00 UTC) is shown as contours for the COSMO/MESSy instances (red: COSMO/MESSy $_{0.22^{\circ}}$, blue: COSMO/MESSy $_{0.11^{\circ}}$). The best track position from HURDAT is marked as reference (green crosses). To allow for a temporal comparison, the positions on 26 September and 1 October are yellow coloured for all tracks. The different model domains of the COSMO/MESSy instances are shaded (blue: COSMO/MESSy $_{0.22^{\circ}}$, grey: COSMO/MESSy $_{0.11^{\circ}}$).

27 September, the position and further track of Isaac differs distinctly from the best track position (Fig. 11) from this time on. In contrast, the dynamical two-way coupling between the COSMO/MESSy instances leads to a correct representation of the track and intensity of Isaac in the coarser-resolved model instance, even after the system has left the model domain of the finer-resolved instance.

By simulating the development of Isaac with the MECO(2) setup, the potential of the dynamical two-way coupling between two COSMO/MESSy instances is demonstrated: to capture the multi-scale interactions, prerequisite for the development of the initial TC of Isaac, in this case, a horizontal model resolution of 0.11° is required. The model domain of this fine-resolved instance, however, can be kept

small, if the dynamical information is fed back to the coarser-resolved model instance in the two-way coupled mode.

Overall, the results of the examples shown here indicate that the two-way coupling has the potential to improve the representation of hurricanes in the coarser COSMO/MESSy model instance.

5 Model performance

Due to technical reasons, the frequency of data exchange between the child and the parent models must be the same as for the parent-to-child data transfer. During the latter, the two time slices of the boundary fields, between which COSMO




Figure 12. Hourly averaged wall-clock time spent for the processing of the coupling data in the submodels (a) MMD2WAY_CHILD and (b) MMD2WAY_PARENT for a MECO(2) setup, shown for different couplings between the two COSMO/MESSy instances: one-way coupling (black) and two-way dynamical coupling (red). The dashed lines are for the additional coupling of 139 tracers.

usually performs a linear time interpolation, are now filled with the data of the actual time step. This was required to enable a non-sequential two-way coupling. However, it limits the choice of the coupling frequency, which should be chosen to be as small as possible, i.e. as the smallest common multiple of the parent and the child model time steps. For this reason, a sensitivity analysis of different coupling frequencies is not provided here.

Usually people are aware of other couplings (e.g. ocean–atmosphere coupling) in which, for instance, for mass conservation, fluxes need to be accumulated/averaged over the coupling interval. In contrast to this, our two-way coupling of two atmosphere models utilises a relaxation technique at the lateral boundaries for the parent-to-child exchange. For the child-to-parent coupling, a relaxation for the entire coupling domain modifies the coarser model results according to the finer-resolved fields. Thus, since we do not couple fluxes for which mass conservation would be required but correct the results directly; accumulation or averaging over time is not feasible.

The runtime performance depends first and foremost on the specific model setups (e.g. on the complexity of the chosen chemistry representation). However, in the end, the overall performance is mostly determined by the “degree of balance” of the distribution of parallel tasks among the different model instances. We discussed this in detail in Part 2 of our series (Kerkweg and Jöckel, 2012b) for the one-way nesting case. The same principles hold for the two-way exchange, except for the complication that communication waiting times depend now on bidirectional data exchange. Thus, it is up to the user to find (experimentally) the optimum task distribution to minimise communication waiting times. The Supplement contains an example, showing the additional costs of two-way coupling, for one specific MECO(2) setup without chemistry.

Figure 12 exemplarily sketches the costs of the coupling. A MECO(2) setup similar to the hurricane case was integrated for 1 day and the residence time in the respective routines transforming the data has been measured. Because the child model does all the data transformations between the

two grids, it consumes much more computing time than the parent model. The difference between the one-way coupled (black) and the solely dynamically two-way coupled (red) simulation is small, as only six additional fields need to be interpolated. Already for the one-way coupled simulation, adding 139 chemical tracers (black dashed line) triples the processing time in the child model, while it requires 6-fold time, if they are two-way coupled (red dashed line).

In contrast to this, the number of coupling fields provokes no systematic increase of computing time in the parent model. Although the coupling time increases significantly in the child model, the time consumption for the coupling is still negligible in comparison to the computing time required for the calculation of the chemistry of these 139 chemical tracers.

6 Conclusions

In this article, we present the next generation of the MMD (v2.0). While MMD (v1.0) provided all the tools required for the one-way coupling of dynamical and chemical models (e.g. EMAC and an arbitrary number of COSMO/MESSy instances) following a client–server approach, version v2.0 was further developed to allow for data exchange from the client to the server model.

To reach this goal, the MMD library was expanded by the respective subroutines for the data exchange via MPI. The new submodel MMD2WAY includes the features of the previous MMD (v1.0) submodels MMDCLNT and MMDSERV plus all functionalities for the data transfer from the child to the parent model.

The new MESSy infrastructure submodel GRID is used for the transformation from the child model grid to the subpart of the parent model grid overlapped by the child model. For the horizontal data remapping, the SCRIP software implemented in the GRID submodel is used. For the vertical regridding for COSMO–EMAC coupling, NREGRID, which is also part of the GRID submodel, is utilised. In contrast to this, the vertical remapping for the COSMO–COSMO cou-

pling is performed using the spline interpolation as provided by INT2COSMO.

Currently, an inconsistency between the vertical regrid- ing routines in INT2COSMO and GRID, and in the adap- tion of the resolution-dependent orography, limits a fully consistent dynamical coupling. This is especially the case for the EMAC–COSMO/MESSy coupling, as the vertical coor- dinate of EMAC is pressure based, while the COSMO model uses a height-based vertical coordinate. Nevertheless, over flat terrain and between COSMO/MESSy instances, coupling of 3-D prognostic variables is possible.

The capabilities of the two-way coupling have been demonstrated on the basis of four examples: (a) a compar- ison of fields upscaled from the regional model to the global model grid with the global model fields, (b) the compar- ison of radiative forcing calculated consistently with the same radiation scheme, (c) the two-way coupling of dust tracers, whose emission fluxes are highly grid resolution dependent and (d) the dynamical coupling of two COSMO/MESSy in- stances influencing the development of a hurricane within the coarse COSMO/MESSy model domain.

The Supplement contains the manuals for the MESSy infrastructure submodel GRID, the MMD library and the MMD user manual.

To develop a fully dynamical two-way coupling for the MECO(n) system, the INT2COSMO routines need to be re- placed by more generic routines to enable a fully consistent coupling in both directions.

Code availability. The submodel GRID and the two-way coupling code are part of the official MESSy distribution (code release v2.53 and younger). The code as described is part of MESSy, which is continuously further developed and applied by a consortium of in- stitutions. The usage of MESSy and access to the source code is licensed to all affiliates of institutions which are members of the MESSy Consortium. Institutions can be a member of the MESSy Consortium by signing the MESSy Memorandum of Understand- ing. More information can be found on the MESSy Consortium website (www.messy-interface.org).

The Supplement related to this article is available online at <https://doi.org/10.5194/gmd-11-1059-2018-supplement>.

Author contributions. AK was the leader of the FLAGSHIP project. She developed the largest part of MMD (library and MMD2WAY submodel) and the generic submodel GRID. CH and GP have been part of the FLAGSHIP project and considerably con- tributed to the MMD2WAY development. CH performed the hur- ricane study. GP provided the results of the dust tracer example. MM was the first “FLAGSHIP external user” of the two-way cou- pled system and contributed the radiative forcing example. PJ was one of the masterminds of the FLAGSHIP team. Additionally, the GRID software is based on the NCREGRID software developed by

him. Thus, he considerably contributed to the implementation of the GRID submodel.

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

Acknowledgements. The work was financed by the German Ministry of Education and Research (BMBF) in the framework of the MiKlip (Mittelfristige Klimaprognose/Decadal Prediction) subproject FLAGSHIP (Feedback of a Limited-Area model to the Global-Scale implemented for Hindcasts and Projections). Mariano Mertens acknowledges funding by the DLR project “Verkehr in Europa”. We thank Bastian Kern and Andrea Pozzer for fruitful discussions concerning the application of SCRIP for interpolation between the EMAC and the MPIOM grid. The authors wish to acknowledge use of the Ferret programme for the graphics in this paper. Ferret is a product of NOAA’s Pacific Marine Environmental Laboratory (information is available at <http://www.ferret.noaa.gov>). Additionally, we thank the members of the MiKlip subproject MesoTel (Andreas Will, Ingo Kirchner, Markus Thürkow, Stefan Weiher and Mareike Schuster) for fruitful discussion of the dynamical two-way coupling between the EMAC and the COSMO/MESSy model.

Edited by: Sophie Valcke

Reviewed by: two anonymous referees

References

- Baumgaertner, A. J. G., Jöckel, P., Kerkweg, A., Sander, R., and Tost, H.: Implementation of the Community Earth System Model (CESM) version 1.2.1 as a new base model into version 2.50 of the MESSy framework, *Geosci. Model Dev.*, 9, 125–135, <https://doi.org/10.5194/gmd-9-125-2016>, 2016.
- Dietmüller, S., Jöckel, P., Tost, H., Kunze, M., Gellhorn, C., Brinkop, S., Frömming, C., Ponater, M., Steil, B., Lauer, A., and Hendricks, J.: A new radiation infrastructure for the Mod- ular Earth Submodel System (MESSy, based on version 2.51), *Geosci. Model Dev.*, 9, 2209–2222, <https://doi.org/10.5194/gmd-9-2209-2016>, 2016.
- Doms, G. and Schättler, U.: The nonhydrostatic limited-area model LM of DWD. Part 1: Scientific documentation, Deutscher Wet- terdienst, Offenbach, available at: www.cosmo-model.org (last access: 5 February 2018), 1999.
- Giorgetta, M. A., Jungclaus, J., Reick, C. H., Legutke, S., Bader, J., Böttinger, 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., Mikolajew- icz, U., Mueller, W., Notz, D., Pithan, F., Raddatz, T., Rast, S., Redler, R., Roeckner, E., Schmidt, H., Schnur, R., Segsnei- der, 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 Intercom- parison Project phase 5, *J. Adv. Model. Earth Sy.*, 5, 572–597, <https://doi.org/10.1002/jame.20038>, 2013.

- Granier, C., Bessagnet, B., Bond, T., D'Angiola, A., van der Gon, H. D., Frost, G., Heil, A., Kaiser, J., Kinne, S., Klimont, Z., Kloster, S., Lamarque, J.-F., Liousse, C., Masui, T., Meleux, F., Mieville, A., Ohara, T., Raut, J.-C., Riahi, K., Schultz, M., Smith, S., Thompson, A., Aardenne, J., Werf, G., and Vuuren, D.: Evolution of anthropogenic and biomass burning emissions of air pollutants at global and regional scales during the 1980–2010 period, *Climatic Change*, 109, 163–190, 2011.
- Harris, L. M. and Durran, D. R.: An Idealized Comparison of One-Way and Two-Way Grid Nesting, *Mon. Weather Rev.*, 138, 2174–2187, <https://doi.org/10.1175/2010MWR3080.1>, 2010.
- Hurrell, J. W., Holland, M. M., Gent, P. R., Ghan, S., Kay, J. E., Kushner, P. J., Lamarque, J. F., Large, W. G., Lawrence, D., Lindsay, K., Lipscomb, W. H., Long, M. C., Mahowald, N., Marsh, D. R., Neale, R. B., Rasch, P., Vavrus, S., Vertenstein, M., Bader, D., Collins, W. D., Hack, J. J., Kiehl, J., and Marshall, S.: The Community Earth System Model A Framework for Collaborative Research, *B. Am. Meteorol. Soc.*, 94, 1339–1360, <https://doi.org/10.1175/BAMS-D-12-00121.1>, 2013.
- Jöckel, P.: Technical note: Recursive rediscritisation of geoscientific data in the Modular Earth Submodel System (MESSy), *Atmos. Chem. Phys.*, 6, 3557–3562, <https://doi.org/10.5194/acp-6-3557-2006>, 2006.
- Jöckel, P., Kerkweg, A., Pozzer, A., Sander, R., Tost, H., Riede, H., Baumgaertner, A., Gromov, S., and Kern, B.: Development cycle 2 of the Modular Earth Submodel System (MESSy2), *Geosci. Model Dev.*, 3, 717–752, <https://doi.org/10.5194/gmd-3-717-2010>, 2010.
- Jöckel, P., Tost, H., Pozzer, A., Kunze, M., Kirner, O., Brenninkmeijer, C. A. M., Brinkop, S., Cai, D. S., Dyroff, C., Eckstein, J., Frank, F., Garny, H., Gottschaldt, K.-D., Graf, P., Grewe, V., Kerkweg, A., Kern, B., Matthes, S., Mertens, M., Meul, S., Neumaier, M., Nützel, M., Oberländer-Hayn, S., Ruhnke, R., Runde, T., Sander, R., Scharffe, D., and Zahn, A.: Earth System Chemistry integrated Modelling (ESCI-Mo) with the Modular Earth Submodel System (MESSy) version 2.51, *Geosci. Model Dev.*, 9, 1153–1200, <https://doi.org/10.5194/gmd-9-1153-2016>, 2016.
- Jones, P.: First- and Second-Order Conservative Remapping Schemes for Grids in Spherical Coordinates, *Mon. Weather Rev.*, 127, 2204–2210, 1999.
- Jones, P. W.: A User's Guide for SCRIP: A Spherical Coordinate Remapping and Interpolation Package, Version 1.4, Theoretical Division Los Alamos National Laboratory, The User's Guide is part of the downloadable SCRIP code distribution, available at: <http://oceans11.lanl.gov/trac/SCRIP> (last access: 5 February 2018), 1998.
- 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, <https://doi.org/10.5194/gmd-5-87-2012>, 2012a.
- Kerkweg, A. and Jöckel, P.: The 1-way on-line coupled atmospheric chemistry model system MECO(n) – Part 2: On-line coupling with the Multi-Model-Driver (MMD), *Geosci. Model Dev.*, 5, 111–128, <https://doi.org/10.5194/gmd-5-111-2012>, 2012b.
- Kerkweg, A., Buchholz, J., Ganzeveld, L., Pozzer, A., Tost, H., and Jöckel, P.: Technical Note: An implementation of the dry removal processes DRY DEPosition and SEDimentation in the Modular Earth Submodel System (MESSy), *Atmos. Chem. Phys.*, 6, 4617–4632, <https://doi.org/10.5194/acp-6-4617-2006>, 2006a.
- Kerkweg, A., Sander, R., Tost, H., and Jöckel, P.: Technical note: Implementation of prescribed (OFFLEM), calculated (ONLEM), and pseudo-emissions (TNUDGE) of chemical species in the Modular Earth Submodel System (MESSy), *Atmos. Chem. Phys.*, 6, 3603–3609, <https://doi.org/10.5194/acp-6-3603-2006>, 2006b.
- Knapp, K. R., Kruk, M. C., Levinson, D. H., Diamond, H. J., and Neumann, C. J.: The International Best Track Archive for Climate Stewardship (IBTrACS), *B. Am. Meteorol. Soc.*, 91, 363–376, <https://doi.org/10.1175/2009BAMS2755.1>, 2010.
- Landsea, C. W. and Franklin, J. L.: Atlantic Hurricane Database Uncertainty and Presentation of a New Database Format, *Mon. Weather Rev.*, 141, 3576–3592, <https://doi.org/10.1175/MWR-D-12-00254.1>, 2013.
- Mertens, M., Kerkweg, A., Jöckel, P., Tost, H., and Hofmann, C.: The 1-way on-line coupled model system MECO(n) – Part 4: Chemical evaluation (based on MESSy v2.52), *Geosci. Model Dev.*, 9, 3545–3567, <https://doi.org/10.5194/gmd-9-3545-2016>, 2016.
- Moeng, C.-H., Dudhia, J., Klemp, J., and Sullivan, P.: Examining two-way grid nesting for large eddy simulation of the PBL using the WRF model, *Mon. Weather Rev.*, 135, 2295–2311, <https://doi.org/10.1175/MWR3406.1>, 2007.
- Rappaport, E. N., Franklin, J. L., Avila, L. A., Baig, S. R., Beven, II, J. L., Blake, E. S., Burr, C. A., Jiing, J.-G., Juckins, C. A., Knabb, R. D., Landsea, C. W., Mainelli, M., Mayfield, M., McAdie, C. J., Pasch, R. J., Sisko, C., Stewart, S. R., and Tribble, A. N.: Advances and Challenges at the National Hurricane Center, *Weather Forecast.*, 24, 395–419, <https://doi.org/10.1175/2008WAF2222128.1>, 2009.
- Rockel, B., Will, A., and Hense, A.: The Regional Climate Model COSMO-CLM (CCLM), *Meteorol. Z.*, 17, 347–348, 2008.
- Roeckner, E., Brokopf, R., Esch, M., Giorgetta, M., Hagemann, S., Kornblüeh, L., Manzini, E., Schlese, U., and Schulzweida, U.: Sensitivity of simulated climate to horizontal and vertical resolution in the ECHAM5 atmosphere model, *J. Climate*, 19, 3771–3791, 2006.
- Schwendike, J. and Jones, S. C.: Convection in an African Easterly Wave over West Africa and the eastern Atlantic: A model case study of Helene (2006), *Q. J. Roy. Meteor. Soc.*, 136, 364–396, <https://doi.org/10.1002/qj.566>, 2010.
- Stappeler, J., Doms, G., Schättler, U., Bitzer, H. W., Gassmann, A., Damrath, U., and Gregoric, G.: Meso-gamma scale forecasts using the non-hydrostatic model LM, *Meteorol. Atmos. Phys.*, 82, 75–96, <https://doi.org/10.1007/s00703-001-0592-9>, 2003.
- Will, A., Akhtar, N., Brauch, J., Breil, M., Davin, E., Hagemann, H. T. M., Maisonnave, E., Thürkow, M., and Weiher, S.: 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, *Geosci. Model Dev.*, 10, 1549–1586, <https://doi.org/10.5194/gmd-10-1549-2017>, 2017.
- Zaengl, G., Reinert, D., Ripodas, P., and Baldauf, M.: The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core, *Q. J. Roy. Meteor. Soc.*, 141, 563–579, <https://doi.org/10.1002/qj.2378>, 2015.