



C-Coupler2: a flexible and user-friendly community coupler for model coupling and nesting

Li Liu^{1,2}, Cheng Zhang^{1,2}, Ruizhe Li^{1,2}, Bin Wang^{1,2,3}, and Guangwen Yang^{1,2}

¹Ministry of Education Key Laboratory for Earth System Modeling, Department of Earth System Science, Tsinghua University, Beijing, China

²Joint Center for Global Change Studies (JCGCS), Beijing, China

³State Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics (LASG), Institute of Atmospheric Physics, Chinese Academy of Sciences, Beijing, China

Correspondence: Li Liu (liuli-cess@tsinghua.edu.cn) and Cheng Zhang (zhangc-cess@tsinghua.edu.cn)

Received: 1 February 2018 – Discussion started: 12 February 2018

Revised: 29 July 2018 – Accepted: 3 August 2018 – Published: 31 August 2018

Abstract. The Chinese C-Coupler (Community Coupler) family aims primarily to develop coupled models for weather forecasting and climate simulation and prediction. It is targeted to serve various coupled models with flexibility, user-friendliness, and extensive coupling functions. C-Coupler2, the latest version, includes a series of new features in addition to those of C-Coupler1 – including a common, flexible, and user-friendly coupling configuration interface that combines a set of application programming interfaces and a set of XML-formatted configuration files; the capability of coupling within one executable or the same subset of MPI (message passing interface) processes; flexible and automatic coupling procedure generation for any subset of component models; dynamic 3-D coupling that enables convenient coupling of fields on 3-D grids with time-evolving vertical coordinate values; non-blocking data transfer; facilitation for model nesting; facilitation for increment coupling; adaptive restart capability; and finally a debugging capability. C-Coupler2 is ready for use to develop various coupled or nested models. It has passed a number of test cases involving model coupling and nesting, and with various MPI process layouts between component models, and has already been used in several real coupled models.

1 Introduction

Couplers, which can handle data interpolation and data transfer between different models and different grids, have been widely used to develop coupled models for fields such as weather forecasting and climate simulation and prediction. The Community Coupler (C-Coupler) family was initiated in 2010 in China. It aims to flexibly serve various coupled models with a user-friendly interface and to provide extensive coupling functions to complement the ever increasing number of coupled models being developed and used in China. To this end, the first version (C-Coupler1; Liu et al., 2014) included features such as flexible coupling configuration based on configuration files and a 3-D coupling capability. Two coupled models were built using C-Coupler1. The first is a coupled climate system model version FGOALS-gc (Liu et al., 2014) that was built by replacing the original coupler CPL6 (Craig et al., 2005) used in the CMIP5 (Coupled Model Intercomparison Project, Phase 5) model FGOALS-g2 (a grid point version of the Flexible Global Ocean–Atmosphere–Land System model; Li et al., 2013a) by C-Coupler1. FGOALS-gc can achieve exactly the same (bit-wise identical) simulation results as FGOALS-g2, and was further used in several Chinese research projects. The second is a regional coupled model FIO-AOW (Zhao et al., 2017) that consists of an atmosphere model WRF (Weather Research and Forecasting model; Wang et al., 2014), an ocean model POM (Princeton Ocean Model; Wang et al., 2010), and a wave model MASNUM (MARine Science and NUMerical Modeling; Yang et al., 2005). FIO-AOW employs both 2-

D and 3-D coupling, where the coupling from MASNUM to POM includes a 3-D field, the wave-induced mixing coefficient (Qiao et al., 2004). FIO-AOW has been used in research to improve typhoon forecasting (Zhao et al., 2017). These coupled models demonstrate that C-Coupler1 can be used for different coupling configurations. C-Coupler1 demonstrates the feasibility of the general design of C-Coupler, although as the first version, it does not fully achieve all our targets and still has some limitations. A significant limitation is that C-Coupler1 is not sufficiently user-friendly. To construct a coupled model with C-Coupler1 requires much effort to prepare the configuration files in ASCII format. For example, there are about 2500 and 700 lines in the configuration files for FGOALS-gc and FIO-AOW, respectively. With such a limitation, C-Coupler1 did not achieve wide usage.

The next version, C-Coupler2, includes not only increased coupling function support but also flexibility and user-friendliness. It is now ready for use and has passed hundreds of test cases based on a sample coupled model and several real coupled models. Compared with C-Coupler1 or even other existing couplers, C-Coupler2 has the following new features.

1. *A common, flexible, and user-friendly coupling configuration interface that combines a set of application programming interfaces (APIs) and a set of configuration files in XML format.* This interface enables the user to flexibly and conveniently specify or change coupling configurations including the component models used in a coupled model, the time step of each component model, model grids, parallel decompositions on model grids, coupling frequencies, coupling lags between component models (given two component models that are coupled, they can have different model times in a coupling exchange), the fields to be coupled, the data type of coupling fields, coupling connections between component models, and coupling procedure generations. Remapping configurations are also modifiable: the user can either use the remapping weights that are read from an existing remapping weight file produced by software tools such as SCRIP (Jones, 1999), ESMF (Hill et al., 2004), YAC (Hanke et al., 2016), and CoR1 (Liu et al., 2013) or use remapping weights that are automatically generated by C-Coupler2 in parallel.
2. *Coupling within one executable or the same subset of MPI (message passing interface) processes.* The component models that are coupled together can be in either multiple executables or in one, and can fully or partially share the same subset of MPI processes. Different processes in the same component model can also be coupled with C-Coupler2.
3. *Flexible and automatic coupling procedure generation.* In a coupling procedure generation, the coupling generator can automatically detect existing component models in the coupled model, detect possible coupling connections within a subset of component models, and generate a coupling procedure for each coupling connection. A coupling procedure can include a set of operations such as data transfer, data interpolation, data type transformation, and data averaging when necessary (please refer to Sect. 4.3 for details). Multiple coupling procedure generations can be performed for a coupled model, and a coupling procedure generation can be performed for any subset of component models.
4. *Dynamic 3-D coupling capability.* This allows for convenient coupling of fields between two 3-D grids, either of which has variable vertical coordinate values that change in time integration.
5. *Non-blocking data transfer.* It is implemented with two-sided MPI communication (i.e., *MPI_Isend* and *MPI_Irecv*) by default, and with one-sided MPI communication (i.e., *MPI_put* and *MPI_get*) as an additional option to enable flexible setting of coupling lags and to minimize potential deadlocks.
6. *Facilitation for model nesting.* C-Coupler2 facilitates a regional model (either a component model or a coupled model constructed with a coupler) to be nested (either one way or two ways) into itself or another model without significant changes to the model codes, and can enable different grid domains in a nested system to be integrated simultaneously for better parallel performance.
7. *Facilitation for incremental coupling.* An existing coupled model using any coupler can be used as a component model by C-Coupler2, which employs the component models of the existing coupled model as its child components. Thus, an existing coupled model can be coupled with another model (either itself as a single component model or a coupled model) to make a bigger coupled model, where only the new model coupling through C-Coupler2 is required to be newly developed, without changes to the original coupling in the existing coupled model.
8. *Adaptive restart capability.* It can adaptively achieve exact (bitwise identical) restart for coupling fields no matter the setting of coupling lags and no matter the implementation of coupling. It also provides support for model fields and can automatically determine a right model time for restarting a “*continue*” run.
9. *Debugging capability.* A series of implementations in C-Coupler2 enable it to alert the user, as early as possible, to the potential risks of constructing or using a coupled model, and guide the user to fix the errors in model codes and configuration files.

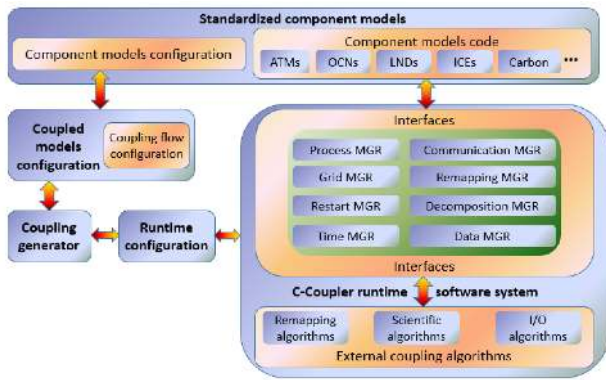


Figure 1. Original software design of C-Coupler (from Liu et al., 2014).

The remainder of this paper is organized as follows. We briefly review C-Coupler and C-Coupler1 in Sect. 2, introduce the motivation for the development of C-Coupler2 in Sect. 3, describe the implementation of C-Coupler2 in Sect. 4, evaluate C-Coupler2 in Sect. 5, and briefly summarize this paper and discuss the future work in Sect. 6.

2 Brief review of C-Coupler and C-Coupler1

The general goal of C-Coupler includes two key aspects. First, C-Coupler can serve various coupling configurations in different coupled models. For example, C-Coupler works as a library and can be used to develop a centralized coupler component for a coupled model and can also achieve model coupling between component models without a separate coupler component. Second, a component model only tells C-Coupler which fields it can provide and which fields it wants to obtain from a coupled model, without caring about where the fields it wants are from (i.e., from specific data files or from specific component models) and where the fields it can provide will be transferred to. Thus, a component model can have an identical code version in coupled models with different coupling configurations. This general goal can be achieved through C-Coupler's original software design (Fig. 1), which consists of a coupling configuration system, a coupling generator, and a runtime software system. The runtime software system works a common, flexible, and extendable library that includes various coupling functions or can even integrate external calculation routines to serve various coupling configurations. The configuration system defines common rules for describing various coupling configurations. Besides the coupling configuration of component models and coupled models, the configuration system includes a runtime configuration, which describes detailed coupling procedures corresponding to a coupling configuration. This is the input of the runtime software system. The coupling generator can automatically generate the run-

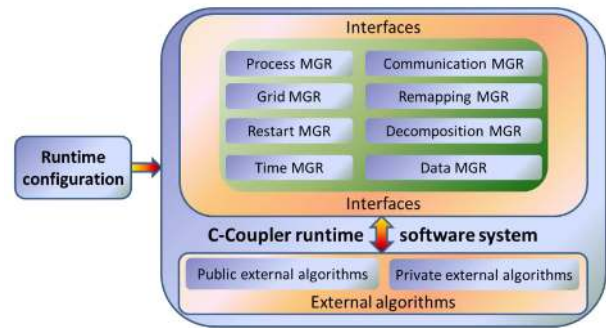


Figure 2. C-Coupler1 software structure.

time configuration, facilitating the construction of a coupled model.

The first version, C-Coupler1, was not developed with the aim of achieving the full software design in Fig. 1, but only with a focus on the runtime software system and the runtime configuration (Fig. 2). To describe the runtime configuration, a set of ASCII configuration files were designed. Besides the traditional coupling functions of data transfer and data interpolation for 2-D coupling, the runtime software system of C-Coupler1 can integrate external calculation routines and has a 3-D coupling capability that enables convenient coupling of fields between different 3-D grids. To achieve simultaneous 2-D and 3-D coupling, remapping software CoR1 was developed and included in C-Coupler1. CoR1 can effectively manage 1-D, 2-D, and 3-D grids, and can interpolate the fields on such grids, where the 3-D interpolation is performed in a “2-D + 1-D” manner (“2-D” corresponds to interpolation between horizontal sub-grids, and “1-D” corresponds to interpolation between vertical sub-grids).

3 Motivation

We considered the following motivations when designing and developing C-Coupler2.

3.1 Coupling configuration

The runtime configuration in C-Coupler1 is almost fully based on configuration files, which can improve flexibility in specifying or changing the coupling configuration, but their overuse may significantly lower user-friendliness. For example, many changes to configuration files are required for changing coupling frequencies based on C-Coupler1. Moreover, overuse of configuration files can cause problems. Configuration files can be inconsistent with the configuration information determined by the model codes. For example, C-Coupler1 will read in the time step of each component model from the runtime configuration files, while the user can change the time step of a component model through the *namelist* file or the model codes. C-Coupler1 must read in

each model grid through a grid data file managed by the runtime configuration, while a model grid of a component model can be generated by the model code or read from a grid data file that is not managed by the runtime configuration. To avoid problems resulting from such inconsistencies, extra effort is required to develop code to detect them, and the user will have to fix the corresponding configuration files when an inconsistency is detected.

Therefore, C-Coupler2 should not allow configuration files to include any configuration information determined by component models, and it should provide flexible APIs to enable component models to specify various coupling configuration information flexibly. Considering various kinds of component models, the configuration system should have commonality in, for example, supporting various kinds of component models and model grids. Considering the low readability of the ASCII format, another format with better readability should be used to design the configuration files.

3.2 Coupling within one executable or a subset of MPI processes

Similar to CPL6, C-Coupler1 requires each component model to have its own executable. However, there are increased requirements for coupling within one executable or a subset of MPI processes. For example, CESM (Community Earth System Model; Hurrell et al., 2013) has the component models and driver containing the coupler CPL7 (Craig et al., 2012) enclosed in a unique executable, and any two different component models can run on non-overlapping, partially overlapping, or overlapping MPI processes. The rapid expansion of model codes requires modularization to guarantee the quality of the models' software, and a coupler can be used to achieve this when it can support coupling between different procedures in the same component model.

3.3 Dynamic 3-D coupling

Atmospheric chemistry modeling is becoming increasingly important for simulating air quality and climate. Such modeling strongly depends on meteorological fields, and can be included as an internal package in an atmosphere model where the atmospheric chemistry package generally uses the same 3-D grid as the model. The rapid development of atmospheric chemistry modeling has led to stand-alone atmospheric chemistry models, such as GEOS-Chem (Long et al., 2015), which read in meteorological fields from data files that can be produced by various atmosphere models. As increasing numbers of atmosphere models require the time-variant aerosol concentration, which can be produced by atmospheric chemistry modeling, there is increasing demand for two-way coupling between an atmosphere model and a stand-alone atmospheric chemistry model. Even if an atmosphere model includes an atmospheric chemistry package, considering that atmospheric chemistry modeling generally

is very time consuming, it might run with a lower resolution. Overall, there is increasing demand for 3-D coupling between atmosphere models and atmospheric chemistry models (or packages) with different 3-D grids.

Despite its 3-D coupling capability, C-Coupler1 might fail to handle the 3-D coupling between an atmosphere model and an atmospheric chemistry model, because it requires the 3-D grids to be constant throughout the whole simulation, whereas the terrain-following pressure coordinates that are widely used in atmosphere models and atmospheric chemistry models make the vertical coordinate values of 3-D grids change with the surface pressure in time integration. In this paper, we call 3-D coupling on constant grids "static 3-D coupling" and 3-D coupling on the 3-D grids with time-variant vertical coordinate values "dynamic 3-D coupling". A coupler having a dynamic 3-D coupling capability will be much more capable of achieving coupling between an atmosphere model and an atmospheric chemistry model (or package).

3.4 Coupling procedure generation

Model coupling is generally achieved through coupling procedures that consist of operations such as data transfer, data interpolation, data averaging, data type transformation, etc. Some existing couplers require the user to explicitly develop all coupling procedures. This is inflexible and not user-friendly enough, because the user must modify the model code, perhaps even significantly, when developing a new coupled model or changing coupling configurations. The coupler OASIS (Redler et al., 2010; Valcke, 2013; Craig et al., 2017) is more flexible and user-friendly in this regard, because it can automatically generate coupling procedures.

To make C-Coupler2 flexible and user-friendly, it should also include a coupling generator capable of automatically generating coupling procedures. In fact, a coupling generator has already been considered in the original software design of C-Coupler.

3.5 Non-blocking data transfer

Data transfer enables a sender to transfer a set of coupling fields to a receiver. A data send/receive operation is blocking when it does not return until the communication is finished (i.e., the receiver has successfully received the data), while a non-blocking operation can return immediately before the communication is finished. In a two-way coupled model, a component model always executes both data send and receive operations. As mentioned above, C-Coupler aims to enable a model to have identical code versions in different coupled models, so the order of data send and data receive operations in a component model can remain the same in different coupled models. To avoid potential deadlocks, we recommend to execute data send operations as early as possible and execute data receive operations as late as possi-




Figure 3. Example of model coupling between two component models: deadlocks occur when blocking data send/receive operations are used.

ble. Specifically, in the initialization stage or at a time step, data send operations should occur before data receive operations. Figure 3 shows an example of model coupling between two component models, in each of which the data send operation is executed before the data receive operation at each time step. During blocking data send operations, the data send operations in both component models cannot return, because the corresponding data receiving operations subsequent to the data send operations will never be executed, leading to a deadlock. Similarly, blocking data send operations can also introduce deadlocks to model coupling within the same component model. Therefore, non-blocking data transfer is highly desirable for developing C-Coupler2.

3.6 Model nesting

Model nesting generally involves nesting a small grid domain with finer resolution into a larger grid domain with coarser resolution. This approach has been widely used in weather forecasting and climate simulation to achieve higher-resolution simulations in key grid domains, without significantly increasing the computational cost. Generally, a regional model can be nested into another model so that different grid domains are simulated by different models, while some models such as WRF have a self-nesting capability, where different grid domains are simulated by the same model. Although WRF and its self-nesting capability have been widely used, the corresponding software implementation has a number of limitations. First, a data structure that can simultaneously manage the fields on different grid domains and a driver that orders initialization and integration among different grid domains are implemented in WRF. For a regional model without a self-nesting capability, significant code changes in the data structure and driver are required to achieve the self-nesting capability. Second, all grid domains must use the same set of MPI processes for integration, so that grid domains must run one by one, not simultaneously. Such an implementation can limit parallelism as well as scalability to the grid domains with fewest grid cells, and will also waste the parallelism between different grid domains.

Model nesting will introduce field exchange between the same type of component models on different grid domains.

As such field exchange generally includes data transfer and data interpolation that are the fundamental functions of a coupler, model nesting can potentially benefit from couplers. If each domain in model nesting can be treated as a component model in model coupling, a regional model can easily achieve self-nesting with its original data structure only managing the fields on one grid domain, and different domains can be integrated simultaneously on different sets of MPI processes for higher parallelism and better parallel efficiency. To aid in the nesting of a regional coupled model (e.g., a regional ocean–atmosphere coupled model) to itself or another coupled model, couplers can serve the field exchanges both between the same type of component models on different grid domains and between different types of component models on the same grid domain.

3.7 Incremental coupling

Building a new coupled model version involves either directly coupling a set of component models together or updating an existing coupled model through coupling external component models or replacing some component models. Such updating of an existing coupled model is here called “incremental coupling”. Directly coupling many component models together is difficult and possibly unwise, because it requires much effort in software implementation, software testing, scientific testing, etc., while incremental coupling is always better when a suitable existing coupled model is available. However, incremental coupling may still face some technical challenges when the existing coupled model and the component models to be coupled have different software frameworks. For example, He et al. (2013) successfully nested WRF into CESM, where both the main driver of CESM and the driver of WRF were modified to enable CESM to drive the integration of WRF and so achieve effective nesting. Successful incremental coupling can give a new coupled model that may become a new code version corresponding to the original coupled model. Further developing the original and the new coupled model in separate code version branches can lead to conflicts when trying to merge the two branches. For example, following the work of He et al. (2013), the main driver of CESM in the original code version branch (managed and maintained by the National Center for Atmospheric Research, NCAR) was significantly changed without considering WRF nesting, leading to much further work being required to re-nest the latest version of WRF into the latest version of CESM.

As C-Coupler aims to enable a model (either a component model or a coupled model) to have an identical code version in different coupled models (i.e., a model can have the same code in different coupled models after incremental coupling), C-Coupler should be able to facilitate incremental coupling.

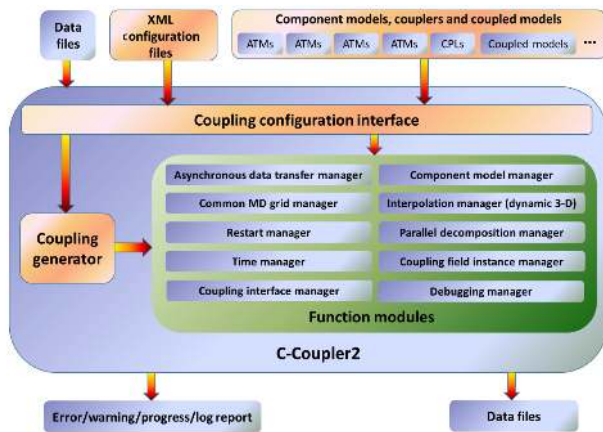


Figure 4. C-Coupler2 software structure.

3.8 Restart capability

A component model generally has the capability of exactly restarting a simulation run from a checkpoint that was produced in a previous run. To make a coupled model achieve an exact restart capability, besides each component model, all coupling procedures should also have an exact restart capability. A lot of effort will be required for users to directly develop an exact restart capability of coupling procedures and new effort will be further required after updating coupling procedures. To improve user-friendliness and to enable a model to have an identical code version in different coupled models, C-Coupler2 should be able to automatically achieve an exact restart capability of coupling procedures, without requiring users to develop specific code for different coupled models.

3.9 Debugging capability

Models can behave anomalously where their run exits due to an error but without giving a report. In such a case, the corresponding simulation setting might be abandoned and another tried, or much effort might be expended locating and fixing the model code segment corresponding to the abnormal exit. Fixing an error is not easy, because it can easily and quickly propagate throughout a component model and from one component model to another through a coupler.

C-Coupler2 aims to facilitate software debugging for model coupling. Specifically, C-Coupler2 should promptly report an error after an abnormal exit, and the error report should effectively help to locate the code segment or configuration file that requires fixing. Moreover, C-Coupler2 should thoroughly examine its inputs to avoid the propagation of errors.

4 Design and implementation of C-Coupler2

Similar with OASIS3-MCT (Craig et al., 2017), C-Coupler2 also works as a library without a driver layer and is driven by calls from the models. As the second version of C-Coupler, C-Coupler2 is guided by the family's general coupling architecture (Fig. 1), so it should be applicable to various coupled models and enable a model to have an identical code version in different coupled models. These considerations influenced the design of the software structure of C-Coupler2 (Fig. 4), which consists of a coupling configuration interface, a coupling generator, and a set of function modules. This software structure is similar to the original software design of C-Coupler (Fig. 1), but has the following differences.

1. The original software design of C-Coupler has the coupling generator as a stand-alone tool that produces the runtime configuration files that drive the runtime software system. However, C-Coupler2 works as a common and flexible library (which can be viewed as the runtime software system), and the coupling generator is an internal program of the library. The coupling generator does not produce runtime configuration files, but directly uses the function modules to generate coupling procedures. Such a design can save redundant code development related to runtime configuration files.
2. Coupling procedure generation in the original software design of C-Coupler fully depends on the offline configuration files that are managed by the configuration system. In C-Coupler2, coupling procedure generation depends on the coupling configuration information obtained by the coupling configuration interface via online API calls and offline configuration files.
3. C-Coupler2 does not include functions to support integrating external algorithms. This will be further discussed in Sect. 6.

In detail, the function modules of C-Coupler2 include managers for non-blocking data transfer, component models, grids, remapping functions, restart capability, parallel decomposition, model time, coupling field instances, coupling interfaces, and debugging capability. The non-blocking data transfer manager manages a set of runtime data transfer algorithms, each of which is responsible for the non-blocking transfer of a set of coupling fields within a component model or between two different component models. The component model manager handles basic information (e.g., name, type, MPI processes) about the component models registered to C-Coupler2. The grid manager manages model grids registered to C-Coupler2; similar to the grid manager in C-Coupler1, it also utilizes CoR1 to support various types of grid with dimensions from 1-D to 4-D. The remapping manager controls a set of runtime remapping algorithms, each of which interpolates a set of coupling fields from one grid to another.

Similar to the remapping manager in C-Coupler1, it also utilizes CoR1 to achieve data interpolation between any kind of grid with dimensions from 1-D to 4-D. It has been further upgraded to support dynamic 3-D interpolation. The restart manager achieves the adaptive restart capability that enables each component model as well as the whole coupled model to conveniently achieve exact restart. The parallel decomposition manager oversees parallel decompositions on model grids. Similar to C-Coupler1, each parallel decomposition must be on a 2-D horizontal grid that has been registered to C-Coupler2, while the parallel decomposition on vertical grids remains unsupported. The coupling field instance manager supervises a set of coupling field instances registered by component models or used by C-Coupler2 internally. The coupling interface manager operates a set of coupling interfaces, each of which imports, exports, or remaps a set of coupling fields. The time manager manages the model time of each component model and manages a set of timers. A timer can be used to control the time to execute a coupling interface and to control lag in model coupling. The debugging manager enables C-Coupler2 as well as component models to flexibly report log information or errors.

We will further introduce here the design and implementation related to each main feature of C-Coupler2, including the common, flexible, and user-friendly coupling configuration interface; coupling within one executable or a subset of MPI processes; flexible and automatic coupling procedure generation; dynamic 3-D coupling capability; non-blocking data transfer; facilitation for model nesting; facilitation for incremental coupling; and a debugging capability. Moreover, we will list out some default options for using C-Coupler2.

4.1 Common, flexible, and user-friendly coupling configuration interface

The coupling generator of C-Coupler2 can automatically generate coupling procedures for model coupling and nesting. As it takes coupling configuration information as its input, the coupling configuration interface should be able to obtain sufficient information for successful coupling procedure generation. Moreover, the constitution of the coupling configuration information determines the flexibility of specifying coupling configurations, and more types of coupling configuration information generally means greater flexibility of C-Coupler2. Finally, the coupling configuration information of C-Coupler2 includes the following.

1. *Basic information about each component model, including the model name, model type, the parallel setting (i.e., the MPI processes that are involved in running the component model), and the relationship with other component models.* To facilitate incremental coupling, an existing coupled model with any coupler can be referred to C-Coupler2 as a component model, and a component model of the existing coupled model can be further referred to C-Coupler2 as a child component model. For

a component model with self-nesting capability within one executable, one grid domain can be employed as a component model and a smaller grid domain directly nested to it can be used as its child component model. Different component models can share common MPI processes.

2. *Coupling connections.* Model coupling by C-Coupler2 can be viewed as a set of data flows, each of which couples a set of coupling fields provided by a component model to a component model that uses these fields – possibly the same component model, as C-Coupler2 supports coupling within one component model. Here we call such a data flow a “coupling connection”. The coupling generator can automatically detect all possible coupling connections, while the user can also specify some coupling connections with higher priority.
3. *Attributes of coupling fields.* Coupling fields are distinguished using field names. All component models in C-Coupler2 share the same name space of the coupling fields as well as the default attributes corresponding to each field name.
4. *Model grids.* A coupling field is either a scalar variable or is associated with a model grid. A model grid may be vertical or horizontal, or a 3-D grid consisting of a horizontal grid and a vertical grid. There might be some relationship between two grids, e.g., a horizontal or vertical grid can be a sub-grid of a 3-D grid.
5. *Decomposition of grid domain for parallelization.* To accelerate modeling on a modern high-performance computer with many processor cores, a grid domain in a component model is generally decomposed into a number of subdomains, each of which is assigned to an MPI process for parallel integration. We call this “parallel decomposition”.
6. *Coupling field instances.* A coupling field generally has multiple instances in a coupled model. First, different component models can produce or use the same coupling field. For example, when all grid domains in self-nesting WRF are registered as component models in C-Coupler2, they can produce the same coupling fields (e.g., precipitation), where each component model has its own coupling field instances. Second, a given component model can have different instances of the same coupling field due to different model grids or different parallel decompositions on the same model grid. For example, a component model can interpolate a coupling field from a source grid to a target grid, which means that this coupling field has two separate instances: one on the source grid and one on the target grid.
7. *Conducting coupling field instances.* A component model can export coupling field instances to the coupled

model, import coupling field instances from the coupled model, or remap its own coupling field instances on a source grid to the coupling field instances on a target grid.

8. *Coupling frequencies.* A component model can specify the frequency at which it exports, imports, or remaps the coupling field instances. Different coupling frequencies might be needed in different simulations, e.g., model coupling can be more frequent when the resolution increases.
9. *Model time.* C-Coupler2 manages model time information for each component model to control model coupling in the time integration of the whole coupled model. It uses a separate and unique time manager for each active component model. The coupling frequencies should be consistent with the model time. For example, a coupling frequency should be a positive integer multiple of the time step of the corresponding model.
10. *Remapping configurations.* Most existing couplers, including C-Coupler1, enables the user to specify how to remap a set of fields from a source grid to a target grid, e.g., using the offline remapping weights read from an input data file produced by a remapping software tool or using the online remapping weights produced by the coupler (if supported).
11. *Shared input parameters for a model run.* C-Coupler will require shared input parameters for a model run, such as case information of the model run, the start time of the model run, how to stop the model run, and the frequency at which to write restart files.

The above coupling configuration information can be classified into two categories: private coupling configuration information of a component model (including information about the component model, model grids, parallel decompositions, coupling field instances, conducting coupling field instances, coupling frequencies, and model time) and public coupling configuration information shared by component models (including coupling connections, attributes of coupling fields, and shared input parameters for a model run). Considering the motivation for coupling configuration (Sect. 3.1), we design a set of C-Coupler2 APIs to enable a component model to flexibly specify its private coupling configuration information through model codes, and design a set of configuration files for flexibly specifying the public coupling configuration information. Although the remapping configurations can be either private or public (the source and target grid used in data remapping may belong to the same or different component models), we only design the corresponding configuration file to guarantee a unique way to specify remapping configurations.

We introduce in this section the implementation of the coupling configuration interface for each kind of configuration information.

4.1.1 C-Coupler2 APIs

This subsection will briefly introduce each kind of C-Coupler2 API. For examples of using C-Coupler2 APIs to implement coupled models, please refer to Sect. S1 of the Supplement.

APIs for component model management

To couple component models running on non-overlapping, partially overlapping, or overlapping subsets of MPI processes, C-Coupler2 allows a component model to run on any subset of MPI processes. Therefore, the coupler can support almost any kind of MPI process layout among the component models. Figure 5 shows an example of a complex MPI process layout: *comp1*, *comp2*, and *comp3* do not share any MPI process; *comp4* runs on a proper subset of the MPI processes of *comp1*; *comp8* runs on all MPI processes of *comp2*; and *comp4* and *comp5* partially share some MPI processes. Moreover, there are relationships between the component models in Fig. 5: *comp1* is the parent of *comp4* and *comp5*; *comp5* is the parent of *comp6* and *comp7*; and *comp2* is the parent of *comp8*. In C-Coupler2, a component model must cover all MPI processes of its children (e.g., *comp1* in Fig. 5 includes all processes of *comp4* and *comp5*). A component model without a parent is a root component model (e.g., *comp1*, *comp2*, and *comp3* in Fig. 5 are root component models). Each MPI process must belong to a unique root component model (e.g., each process in Fig. 5 only belongs to one of *comp1*, *comp2*, or *comp3*), i.e., all root component models cover all MPI processes without sharing any MPI process with each other. This constraint seems contradictory to the target of supporting shared MPI processes among component models, and may make C-Coupler2 unable to support some MPI process layouts. For example, given that a component model consists of two component models that run on partially overlapping subsets of MPI processes, both component models cannot be root component models. To support this kind of MPI process layout, a coupled model can be registered as a root component model of C-Coupler2, and its component models can be further registered as children of the root component model.

Coupler2 provides 10 APIs for component model management, including *CCPL_register_component*, *CCPL_end_coupling_configuration*, etc. *CPL_register_component* is responsible for registering a component model to C-Coupler2. C-Coupler2 only serves component models registered to it. (Almost any model can be registered to C-Coupler2, while it is unnecessary to register a model whose model coupling is fully served by other couplers to C-Coupler2.) The arguments of this API include the ID of the parent compo-

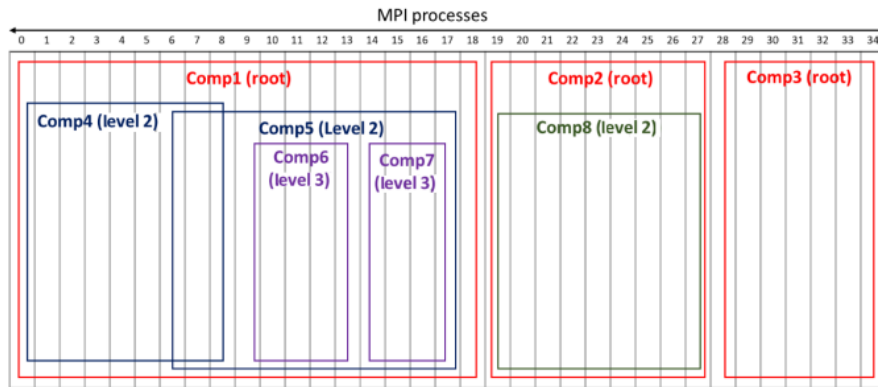


Figure 5. Sample process layout of component models (comp1–comp8).

component model, model name, model type, and MPI communicator. Any component model except a root component model must have a parent. C-Coupler2 will allocate an ID and generate a unique full name for each component model that is formatted as “*parent_full_name@model_name*”, where “*model_name*” means the name of the current component model and “*parent_full_name*” is the full name of the parent component model. (For a root component model, “*parent_full_name*” corresponds to an empty string.) A component model is either “active” or “pseudo” (inactive), as specified by the model type. A pseudo component model can be the parent of some component models, while its name will not be included in the full name of any component model. Moreover, coupling configurations cannot be further specified to a pseudo component model. Table 1 lists the model types currently supported by C-Coupler2. Note that “*active_coupled_system*” and “*pseudo_coupled_system*” indicate that an existing coupled model can be registered as a component model of C-Coupler2. This API can create the MPI communicator of the component model when required. It will start the stage of coupling configuration of the component model, while the API *CCPL_end_coupling_configuration* will finalize the stage of coupling configuration. A component model can successfully call *CCPL_end_coupling_configuration* only when all its children component models have already called this API.

For more details of the APIs for component model management, please refer to the user guide (<https://gitlab.com/c-coupler-group/c-coupler-doc/raw/master/C-Coupler2UserGuide.pdf>, last access: 24 August 2018).

APIs for time management

C-Coupler2 provides 26 APIs for time management, including *CCPL_set_normal_time_step*, *CCPL_check_current_time*, *CCPL_define_single_timer*,

etc. These APIs enable C-Coupler2 to manage the model time information for each active component model. Detailed time information of a component model can also be accessed through C-Coupler2, and thus a component model can employ C-Coupler2 for its model time management. A component model with its own model time management must keep its model time constantly consistent with C-Coupler2. The API *CCPL_check_current_time* can be used to check such consistency. An active component model can have a unique time manager that is not activated until a unique time step has been set through the API *CCPL_set_normal_time_step*. After a time manager is activated, the user can access detailed information on the model time, define timers, advance the model time, and use timers to control model coupling.

C-Coupler2 currently only provides the API *CCPL_define_single_timer* to define a periodic timer that is an alarm for specifying coupling period. The arguments of this API include the ID of the corresponding component model, a period unit, a period count, a local lag count, and an optional remote lag count. The period unit and period count specify the period of the timer. The local lag count corresponds to the period unit, which is used to specify a local lag (it can be viewed as a time offset from the start time) that influences when the timer is on. For example, a timer set with $\langle \text{period_unit} = \text{“steps”}, \text{period_count} = \text{“5”}, \text{local_lag_count} = \text{“2”} \rangle$ will be on at the 2nd, 7th, 12th, etc. (i.e., $5i + 2$, where i is a non-negative integer) time steps of the corresponding component model. The remote lag count also corresponds to the period unit. It can be used to specify a lag on a coupling connection between two component models or within one component model. Its default value is 0 (i.e., no lag). Note that the lag for a coupling connection is determined by the timer from the receiver component model. The lag corresponding to a coupling connection can be viewed as the model time difference from the receiver component model to the sender component model, which can control the time sequence

Table 1. Model types currently supported by C-Coupler2.

Model type	Description	Remark
cpl	Coupler	Active component model
atm	Atmosphere model	Active component model
glc	Glacier model	Active component model
atm_chem	Atmospheric chemistry model	Active component model
ocn	Ocean model	Active component model
lnd	Land surface model	Active component model
sea_ice	Sea ice model	Active component model
wave	Wave model	Active component model
roff	Runoff model	Active component model
active_coupled_system	Coupled model that consists of a set of component models	Active component model
pseudo_coupled_system	Coupled model that consists of a set of component models	Pseudo component model

between the two component models. For example, given a lag of 1/ – 1 h, the coupling fields produced by the sender component model at the sender’s 0th / 1st hour will be obtained by the receiver component model at the receiver’s 1st / 0th hour. Thus, the user can flexibly achieve concurrent runs or sequential runs between component models. Incorrectly setting “*remote_lag_count*” may introduce deadlocks between component models.

For more details about the APIs for time management, please refer to the user guide.

APIs for grid management

Each grid managed by C-Coupler2 belongs to a unique active component model. A grid shared by multiple component models should be registered to each component model separately. The keyword for a grid can be expressed as *<ID of the component model, grid name>*. Therefore, different grids in the same component model cannot have the same grid name, while grids in different component models can have the same grid name.

C-Coupler2 provides 15 APIs for grid management. A horizontal grid can be registered via global grid data (through the API *CCPL_register_H2D_grid_via_global_data*), local grid data (through the API *CCPL_register_H2D_grid_via_local_data*), or a grid data file (through the API *CCPL_register_H2D_grid_via_file*). Considering that a horizontal grid in a component model may be determined by another component model (e.g., the horizontal grid of a land surface model will be determined by an atmosphere model when both models require the same horizontal grid), we designed the API *CCPL_register_H2D_grid_from_another_component*.

A vertical grid can be registered via global grid data. The coordinate of a vertical grid can be registered as a Z coordinate (through the API *CCPL_register_VID_Z_grid_via_model_data*),

a sigma coordinate (through the API *CCPL_register_VID_sigma_grid_via_model_data*), or a hybrid coordinate (through the API *CCPL_register_VID_hybrid_grid_via_model_data*).

A 3-D grid can be registered by combining a horizontal grid and a vertical grid (through the API *CCPL_register_MD_grid_via_multi_grids*). Thus, C-Coupler2 can know the relationship between a 3-D grid and its sub-grids. A 3-D grid can be either an interface-level grid or a middle-point grid. A middle-point grid can be generated from an interface-level grid through the API *CCPL_register_mid_point_grid*. Thus, C-Coupler2 can know the relationship between an interface-level 3-D grid and a middle-point 3-D grid.

For a 3-D grid that consists of a horizontal grid and a vertical grid with sigma or hybrid coordinate, C-Coupler2 can set its unique surface field on the horizontal grid in order to calculate the vertical coordinates at each horizontal grid point. The surface field of a 3-D grid can be static (through the API *CCPL_set_3D_grid_constant_surface_field*), dynamic (through the API *CCPL_set_3D_grid_variable_surface_field*), or external (through the API *CCPL_set_3D_grid_external_surface_field*). A static surface field has constant values with time integration, so the vertical coordinate values in the corresponding 3-D grid are constant. A dynamic surface field has changing values with time integration, so that the vertical coordinate values in the corresponding 3-D grid vary. An external surface field has values determined by the surface field of another 3-D grid.

For more details about the APIs for grid management, please refer to the user guide.

API for parallel decomposition management

To accelerate modeling by taking advantage of a high-performance computer with many processor cores, the model

needs to be parallelized with MPI, whereby the domains of the model grids are decomposed into separate subdomains for parallel integration. To accommodate the parallel integration of component models and to allow model coupling to be handled in parallel, C-Coupler2 both manages parallel decompositions and provides APIs to enable active component models to register their parallel decompositions to the coupler. C-Coupler2 currently supports parallel decompositions only on horizontal grids, and further parallel decomposition on vertical grids is not supported yet. Therefore, parallel decomposition is associated with a horizontal grid, and thus with the component model corresponding to the horizontal grid. The keyword for parallel decomposition is expressed as *<ID of component model, parallel decomposition name>*. Therefore, different parallel decompositions in the same component model cannot have the same name, while parallel decompositions in different component models can have the same name. Multiple parallel decompositions on the same horizontal grid are allowed.

Parallel decomposition on a horizontal grid is described through enumerating global grid cell indexes of the local grid cells assigned to each MPI process of the corresponding component model. A valid global grid cell index should be between 1 and the size of the horizontal grid. For any local grid cells that need not be considered in model coupling (e.g., land-only grid cells in an ocean model), the corresponding values of the global grid cell index can be set to a C-Coupler2 pre-defined variable *CCPL_NULL_INT*, to save some overheads in model coupling.

Currently, C-Coupler2 provides only one API for parallel decomposition management, i.e., *CCPL_register_normal_parallel_decomp*. Please refer to the user guide for more details.

API for coupling field instance management

A coupling field instance includes a set of meta-information and a memory buffer that keeps the data values of an instance of a coupling field. A coupling field instance is associated with a unique component model, a unique grid, and a unique parallel decomposition. An attribute of “*buf_mark*”, which is a non-negative integer mark given by users, is employed in each coupling field instance to separate multiple coupling field instances in the same component model, on the same grid, and on the same parallel decomposition. For example, as the land surface, oceans, and sea ice lie under the atmosphere, an atmosphere model may receive multiple coupling field instances of surface temperature from land surface, ocean, or sea ice models. Therefore, the keyword for a coupling field instance is expressed as *<field name, ID of component model, ID of grid, ID of parallel decomposition, buf_mark>*. For a scalar coupling field instance that is not on a grid, the corresponding grid ID and parallel decomposition ID should be set to -1 .

Currently, C-Coupler2 provides only one API for field instance management, i.e., *CCPL_register_field_instance*. This API allows a component model to register a coupling field instance to the coupler to provide, obtain, and remap coupling field instances in model coupling. An internal model field instance that will not be used in model coupling can also be registered to C-Coupler2 for exact restart capability. For more details of this API, please refer to the user guide for more details.

APIs for coupling interface management

In C-Coupler2, an active component model can handle coupling field instances through coupling interfaces. The keyword of a coupling interface is expressed as *<ID of the component model, interface name>*. Therefore, different coupling interfaces in the same component model cannot have the same interface name, while coupling interfaces in different component models can have the same interface name.

Coupling interfaces are classified into three categories: import, export, and remap. An import interface enables a component model to obtain coupling field instances from itself (for example, C-Coupler2 can achieve coupling between the physical package and the dynamic core in the same component model) or other component models. Specifically, it can be specified to obtain instantaneous or averaged coupling field instances. An export interface enables a component model to provide a number of coupling field instances to the coupled model. A remap interface enables a component model to remap its coupling fields from a source grid to a target grid. There are two detailed kinds of remap interface: normal and fraction-based. A normal remap interface directly interpolates coupling field instances from the source grid to the target grid, while a fraction-based remap interface additionally takes a source fraction (for example, the area fraction of atmosphere, ocean, land surface, or sea ice in each cell of the source grid) as input, will first adjust the values of coupling field instances on the source grid based on the source fraction before remapping, and will finally adjust the values of coupling field instances on the target grid based on the target fraction after remapping (the source fraction is also remapped from the source grid to the target grid to produce the target fraction at the same time). Fraction-based remap interfaces are generally necessary to guarantee conservation in model coupling between different horizontal grids.

There are three steps taken to utilize a coupling interface. The coupling interface is first registered, whereby a timer is required to be specified to control the timing of coupling interface execution. Coupling procedures are next generated for the coupling interface, which is then executed in the third step. Although the API to execute a coupling interface can be called at each time step, a coupling interface will be truly executed only when its timer is bypassed or its timer is on. C-Coupler2 allows the timer to be bypassed when executing a coupling interface, in order to achieve flexible coupling at the

initialization stage of the coupled model. Note that the timer of a coupling interface cannot be bypassed again if this coupling interface has already been executed with the timer on, and when the timer of a coupling interface is not bypassed, the coupling interface will be truly executed at most once each time step, which means that any additional API calls for executing the coupling interface at a time step will be ignored.

For a remap interface that does not refer to coupling between different coupling interfaces or different component models, its coupling procedures are generated implicitly by the coupling generator when registering it. Coupling procedures of an export/import interface are also generated automatically by the coupling generator, but will not be generated when registering the interface, because an export/import interface refers to coupling between different coupling interfaces in the same or different component models. To generate coupling procedures for export or import interfaces, the coupling generator will analyze possible connections from export interfaces to import interfaces based on the field name of each coupling field instance. A coupling connection from an export interface to an import interface can be generated only when these two coupling interfaces have common field names. Regarding a field name, C-Coupler2 allows an export interface to be connected to any number of import interfaces, while forcing an import interface to be connected from a unique export interface. In other words, each coupling field instance in an import interface must have only one provider. If there are multiple providers for a coupling field instance in an import interface, the user must select only one provider through the corresponding configuration file (see subsection of Sect. 4.1.2). Different coupling field instances in an import interface can have different providers. The coupling procedures for import and export interfaces are generated through explicitly calling the APIs for coupling procedure generation.

An export interface or a remap interface can always be executed successfully without error, while the execution of import interfaces can fail and lead to an error report, if the coupling procedures of some necessary coupling field instances have not been generated (i.e., if the providers of some necessary coupling field instances have not been found). When registering an import interface, each import coupling field instance can be specified as necessary or optional. No error will be reported if the providers of some optional coupling field instances have not been found.

C-Coupler2 provides 8 APIs for coupling interface management, including *CCPL_register_port_interface*, *CCPL_register_import_interface*, *CCPL_register_normal_remap_interface*, *CCPL_execute_interface_using_id*, etc. Please refer to the user guide for more details.

APIs for coupling procedure generation

The development of a coupling procedure generation capability in C-Coupler2 has experienced two designs, and the second design can be viewed as an upgrade of the first design and is currently implemented in C-Coupler2. The first design enforces only one global coupling procedure generation for the whole coupled model: it is performed when finalizing the coupling configuration stages of the whole coupled model (when all root component models are calling the API *CCPL_end_coupling_configuration*), and involves all component models that have already been registered to C-Coupler2. Along with the development of C-Coupler2, the first design shows the following limitations:

1. It assumes that the whole coupled model as well as each component model is organized as a unique three-stage flowchart consisting of a coupling configuration stage, coupling procedure generation, and a model coupling run. However, the main drivers of many existing models, such as the CESM driver, consist of not just a unique but multiple three-stage flowcharts, indicating that multiple coupling procedure generations are necessary for wide usage in real cases.
2. As a global coupling procedure generation will involve the participation of all component models, and will require global synchronization of all the MPI processes in the whole coupled model, it will be costly, inconvenient, and unnecessary to conduct multiple global coupling procedure generations. For example, an increment coupling case that seeks to nest a regional atmosphere model into an existing climate system model through C-Coupler2 requires only partial coupling procedure generation between the regional atmosphere model and the global atmosphere model.
3. Besides the global coupling procedure generation, the second design achieves partial coupling procedure generation for any subset of component models through the APIs for coupling procedure generation. The coupling procedure generation related to a component model is classified as either individual or family. Individual mode considers only the given component model in coupling procedure generation, while family mode considers the given component model and its descendant component models in the same coupling procedure generation. When registering a component model through the API *CCPL_register_component*, it can be specified to enable or disable the given component model in the family coupling procedure generation of its parent or any ancestor. The API *CCPL_do_external_coupling_generation* can do coupling procedure generation regarding any subset of component models, where either individual or family coupling procedure generation can be specified for each given component model.

The API `CCPL_get_configurable_comps_full_names` allows flexible specification of a subset of component models in an XML configuration file; it can cooperate with the API `CCPL_do_external_coupling_generation` to further improve the flexibility of coupling procedure generation. Besides partial coupling procedure generations, a global coupling procedure generation will still be performed when root component models are calling the API `CCPL_end_coupling_configuration`, while a root component model that has been disabled in the family coupling procedure generation will not be involved in the global coupling procedure generation.

4. Coupling procedure generation requires the synchronization of all MPI processes of the involved component models. Please refer to the user guide for more details of the APIs for coupling procedure generation.

Other APIs

Like most component models, C-Coupler2 can restart model simulation from a checkpoint. It does so through 6 APIs, including `CCPL_do_restart_write_IO`, `CCPL_start_restart_read_IO`, etc. More details of these APIs can be found in the user guide. The restart management not only serves the variables or data involved in the model coupling handled by C-Coupler2 but can also serve the internal field instances of any component model that have been registered to C-Coupler2. To achieve the restart capability of a model coupled using C-Coupler2, all active component models should separately call the corresponding APIs. Besides the “initial” run, C-Coupler2 supports three types of model run: “continue”, “branch”, and “hybrid”, which are related to the restart capability. Detailed implementation of the restart capability will be further introduced in Sect. 4.8.

C-Coupler2 enables each MPI process in each component model to have a separate log file, thus improving the parallel debugging capability. Several APIs (`CCPL_report_log`, `CCPL_report_progress`, `CCPL_report_error`, `CCPL_get_comp_log_file_name`, and `CCPL_get_comp_log_file_device`) allow component models to benefit from such a capability. For more details of these APIs, please refer to the user guide.

4.1.2 C-Coupler2 configuration files

As mentioned above, the C-Coupler2 configuration files allow flexible specification of public coupling configuration information including shared input parameters for a model run, attributes of coupling fields, remapping configurations, and coupling connections. In order to achieve good readability, all configuration files are in XML format. This subsection briefly introduces the four kinds of configuration files, i.e., input parameter configuration file, field attribute configuration file, remapping configuration file, and coupling connec-

tion configuration file. Additional details can be found in the user guide.

Input parameter configuration file

The input parameter configuration file specifies a set of global input parameters shared by all component models. The input parameters include simulation times (e.g., start and stop times), the type of simulation run (i.e., initial, continue, branch, or hybrid), whether leap years are considered, and the frequency of writing restart data files. Note that C-Coupler2 requires all component models to use the same start and stop times, and the user should guarantee that the input parameter configuration file is consistent with the corresponding modeling settings.

Field attribute configuration file

When registering a coupling field instance to C-Coupler2, the field name should be specified as an input parameter. A field name is only legal when there is a corresponding entry in the field attribute configuration file that is shared by all component models in a coupled model. When the coupling generator automatically tries to generate coupling procedures, field names are used to detect possible coupling connections between coupling interfaces: an import interface and an export interface can have a coupling connection only when their coupling field instances have common field names. The attributes of each coupling field include “*long_name*”, “*default_unit*”, “*dimensions*”, and “*type*”. The attribute “*dimensions*” means a label of grid dimensions. It can be set to “*0D*”, “*H2D*”, “*V1D*”, or “*V3D*”, denoting that a field is a scalar variable that is not on any grid, is on a horizontal grid, is on a vertical grid, or is on a 3-D grid that consists of a horizontal grid and a vertical grid, respectively. The attribute “*type*” gives the type of coupling field, either “*state*” or “*flux*”.

Remapping configuration file

Remapping configuration files can flexibly and conveniently specify the remapping of coupling fields between grids, as follows.

1. For remapping from a source horizontal grid to a target horizontal grid, the user can either employ the remapping weights that are automatically generated by C-Coupler2 in parallel or read from an existing remapping weight file produced by external software such as SCRIP, ESMF, YAC, or CoR1. Unstructured horizontal grids such as cube-sphere grid or non-quadrilateral grids are supported in the online remapping weight generation.
2. Like C-Coupler1, C-Coupler2 uses the “2-D + 1-D” approach to achieve 3-D remapping. Regarding 3-D remapping, the remapping configuration for the 2-D (horizontal) remapping and for the 1-D (vertical)

remapping can be specified separately; the 2-D remapping can also use the remapping weights loaded from a remapping weight file.

3. Different coupling fields in the same component model can have different remapping configurations, and the same coupling field in different component models can have different remapping configurations.
4. Given a coupling field, a component model can either use its own remapping configuration or use that inherited from its parent (if its own remapping configuration is not specified); a root component model (i.e., one without a parent) without a specified remapping configuration can use the specified overall remapping configuration or use the default remapping configuration set by C-Coupler2 (if the overall remapping configuration is not specified). In the default remapping configuration, the bilinear remapping algorithm is used to remap the “state” fields between horizontal grids, the conservative remapping algorithm is used to remap the “flux” fields between horizontal grids, and the linear remapping algorithm is used to remap the vertical and time dimensions. Note that all remapping weights in the default remapping configuration are generated automatically by C-Coupler2.
5. A remapping configuration file consists of a set of remapping settings, each of which can specify the remapping configuration for all coupling fields, coupling fields of the same type (“flux” or “state”), or a specific set of coupling fields (possibly even only one field). A prioritization strategy is designed accordingly: a remapping setting corresponding to all coupling fields is at the lowest priority, a remapping setting corresponding to a type of coupling field is at medium priority, and a remapping setting corresponding to specific coupling field is at the highest priority.

A procedure with data remapping for a given coupling field on a coupling connection between two different component models will be generated when the component models use different corresponding grids. It is possible that the remapping configuration of this coupling field is not the same in the two component models. In such a case, C-Coupler2 will only use the remapping configuration in the source component model (the component model that exports the coupling field). In general, given a coupling field on a coupling connection, C-Coupler2 uses only the remapping configuration in the source component model for coupling procedure generation. Therefore, it is meaningless to specify remapping configurations for the imported coupling fields of a component model.

Figure 6 shows an example of a remapping configuration file that consists of three active remapping settings (corresponding to the XML node of “remapping_setting” with the

attribute “status” of value “on”). The first remapping setting (from L1 to L15 in Fig. 8) corresponds to all coupling fields, the second remapping setting (from L16 to L27 in Fig. 8) corresponds to the “flux” coupling fields, and the third (from L28 to L39 in Fig. 8) corresponds to two specific coupling fields “t_atm_3D” and “ghs_atm_3D” that should be 3-D “state” fields. Specific to the remapping configuration of these two fields, the vertical remapping configuration is determined by the third remapping setting and the horizontal remapping configuration is determined by the first remapping setting. Both the first and second remapping settings specify an online horizontal remapping algorithm (corresponding to the XML node of “H2D_algorithm”) and offline horizontal remapping weight files (corresponding to the XML node of “H2D_weights”). Note that offline remapping weight files have higher priority than the online remapping algorithm. To generate the coupling procedures for data remapping, a remapping weight file in the corresponding remapping configuration (if present) will be used if it matches both the source grid and the target grid of this data remapping.

Coupling connection configuration files

A coupling connection configuration file can be used to specify coupling connections for a component model. The connections are classified into three types: (1) for import interfaces, (2) for model grids (corresponding to the API *CCPL_register_H2D_grid_from_another_component*), and (3) sets of component model full names used for external coupling procedure generation (corresponding to the API *CCPL_do_external_coupling_generation*). The coupling connections for an import interface specify the providers (a provider is a component model as well as its export interface) of imported coupling fields, while a distinct subset of imported coupling fields can have a distinct provider. When the coupling generator tries to generate coupling procedures for an imported field of an import interface, it will first check the corresponding coupling connection configuration file. If the configuration file contains a coupling connection for the imported field, the coupling generator will only use the corresponding provider in coupling procedure generation; otherwise, it will try to detect a provider and will report an error if no provider or multiple providers are detected.

4.2 Coupling within one executable or a subset of MPI processes

As mentioned in Sect. 4.1, C-Coupler2 allows coupling between two component models that run on non-overlapping, partially overlapping, or overlapping MPI processes, and can also automatically generate coupling connections between the export and import interfaces of the same component model. Each component model registered to C-Coupler2 can have its own model coupling resources, including time step,

```

<root>
L1:  <remapping_setting status="on">
L2:    <remapping_algorithms status="on">
L3:      <H2D_algorithm status="on" name="bilinear">
L4:        <parameter name="enable_extrapolate" value="true" />
L5:      </H2D_algorithm>
L6:      <V1D_algorithm status="on" name="linear">
L7:        <parameter name="enable_extrapolate" value="true" />
L8:      </V1D_algorithm>
L9:      <H2D_weights status="on">
L10:        <file name="map_to_global_grid1_default.nc" />
L11:        <file name="map_to_regional_grid1_default.nc" />
L12:      </H2D_weights>
L13:    </remapping_algorithms>
L14:    <fields status="on" specification="default" />
L15:  </remapping_setting>

L16: <remapping_setting status="on">
L17:   <remapping_algorithms status="on">
L18:     <H2D_algorithm status="on" name="conserv_2D" />
L19:     <H2D_weights status="on">
L20:       <file name="map_to_global_grid1_conserv.nc" />
L21:       <file name="map_to_regional_grid1_conserv.nc" />
L22:     </H2D_weights>
L23:   </remapping_algorithms>
L24:   <fields status="on" specification="type">
L25:     <entry value="flux" />
L26:   </fields>
L27: </remapping_setting>
L28: <remapping_setting status="on">
L29:   <remapping_algorithms status="on">
L30:     <V1D_algorithm status="on" name="linear">
L31:       <parameter name="enable_extrapolate" value="true" />
L32:       <parameter name="use_logarithmic_coordinate" value="true" />
L33:     </V1D_algorithm>
L34:   </remapping_algorithms>
L35:   <fields status="on" specification="name">
L36:     <entry value="t_atm_3D" />
L37:     <entry value="ghs_atm_3D" />
L38:   </fields>
L39: </remapping_setting>
</root>

```

Figure 6. Sample of a remapping configuration file.

timers, model grids, parallel decompositions, coupling field instances, and coupling interfaces. In other words, a model coupling resource must be associated with a unique component model. Most model coupling resources, including the time step, model grids, parallel decompositions, coupling field instances, and coupling interfaces, are public to a component model and shared by all its MPI processes. When registering a public model coupling resource of a component model, all MPI processes of the component model are required to call the corresponding API simultaneously, with consistent parameters. To manage different component models and model coupling resources effectively, each component model, as well as each model coupling resource, has a unique ID.

4.3 Flexible and automatic coupling procedure generation

4.3.1 Creation of MPI communicators

As mentioned in Sect. 4.1.1, C-Coupler2 can achieve a partial coupling procedure generation for any subset of component models. A partial coupling procedure generation should only involve the MPI processes of the corresponding subset of component models but not all MPI processes of the whole coupled model. Therefore, the first step in C-Coupler2 is to create a unified MPI communicator that only includes the MPI communicator of each component model involved in the same coupling procedure generation. To create a unified MPI communicator based on a set of sub MPI communicators,

the root process in a sub MPI communicator should know the root process (for example, the ID of the process in the global communicator MPI_COMM_WORLD) of each other MPI communicator. In other words, the component models involved in the same coupling procedure generation should know the MPI processes of each other. As the registration of root component models generally involves all MPI processes of the whole coupled model, each root component model can easily know the MPI processes of any other root component model through global synchronization. However, as the registration of a non-root component model only involves a subset of MPI processes (the MPI processes of itself in general and the MPI processes of its parent component model at most), it is possible that a component model does not know the MPI processes of a non-root component model. For an example based on Fig. 5, when registering the component model *comp4* that is a child of the root component model *comp1*, the root component model *comp2* and its child *comp8* will not be notified by this event, so that *comp2* and *comp8* do not know which MPI processes belong to *comp4*. Similarly, *comp1* and its children *comp4* and *comp5* do not know which MPI processes belong to *comp8*. Given that *comp4* and *comp8* are involved in the same coupling procedure generation, they should be synchronized for creating a unified communicator, but they do not know the MPI processes of each other.

To overcome the above challenge, a new solution should be implemented to enable one component model to know the MPI processes of any other component model. One possible solution is that the root process of *comp4* broadcasts a message to all processes of the whole coupled model and then the processes of *comp8* reply. However, that will introduce global synchronization. To avoid global synchronization, we implemented a file-based solution in C-Coupler2 as follows. C-Coupler2 will write information about the MPI processes of each component model into an internal XML file. Thus, in a coupling procedure generation, a component model can know the MPI processes of any other component model through reading the corresponding XML file. The XML files of all component models are put under the same directory and this directory will be automatically emptied when initializing C-Coupler2 in a new coupled model run. The XML file corresponding to a component model will be produced as early as when registering the component model. When a component model wants to read an XML file that does not exist currently, it will keep waiting until the corresponding component model has been registered. Only one MPI process of a component model will write or read the XML file, to minimize the overhead of this solution.

4.3.2 Timers matching and lags

The coupling generator will generate a coupling procedure for each coupling connection from an export interface to an import interface. As introduced in a subsection of Sect. 4.1.1,

the API for registering an import/export interface takes as an input parameter a timer, which specifies when a component model must import/export coupling fields. The second challenge during the coupling procedure generation is achieving effective coupling with the timers of the import and export interfaces, which can be set independently, are different, or even do not “match” in periods. For example, how to achieve effective coupling when the periods of the import and export timers are 900 and 200 s, respectively (assuming that no lags are specified in the two timers)? At the model time of 0 s, both timers are activated, and the import interface will obtain the coupling fields from the export interface. Before the second activation of the import timer (at 900 s), the export timer will have been on four times (i.e., at 200, 400, 600, and 800 s). The import interface at 900 s can obtain either the average values of the coupling fields from the four times at which the export timer was on, or the instantaneous coupling fields at its last activation (i.e., at 800 s). The choice regarding average values or instantaneous values is made by the user when registering the import interface. Similarly, before the third activation of the import timer (at 1800 s), the export timer will have been on a further five times, at model times of 1000, 1200, 1400, 1600, and 1800 s. Thus, the import interface (at 1800 s) will obtain from the export interface either the average values of the coupling fields at its five intervening occasions or the instantaneous coupling fields at its last activation (at 1800 s) based on the user’s choice. As an additional example, suppose that periods of the import and export timers are now 200 and 900 s, respectively, still with no lags specified in the timers. As before, the import interface will obtain the coupling fields from the export interface at the model time of 0 s when both timers are on. Although the import timer will be further activated at model times of 200, 400, 600 and 800 s, the import interface will not obtain new coupling fields from the export interface, but will use those previously obtained at 0 s, because the export interface will not export coupling fields again until a model time of 900 s. In summary, regardless of the difference in periods between the import and export timers, C-Coupler2 can adapt to conduct model coupling in a suitable manner.

As introduced in the subsections of Sect. 4.1.1, a coupling lag can be specified through the input parameter of the timer when registering an import interface. Given a lag of m seconds ($m \neq 0$), the coupling fields obtained by a receiver component model at its model time of $N + m$ seconds are exported by a sender component model at its model time of N seconds. To support coupling lags, existing coupled models such as CCSM3 (Collins et al., 2006) or FGOALS-g2 essentially extend the simulation stop time of some component models, so that component models have different simulation periods. However, C-Coupler2 only supports a uniform stop time among all component models. To support coupling lags in C-Coupler2 effectively, an import interface will be bypassed if it would request coupling fields produced after the simulation stop time, and an export interface will be by-

passed if its coupling fields would be used after the simulation stop time. Such an implementation may introduce errors into the model states at the last steps of simulation run. We therefore propose to extend the simulation period properly to guarantee correct simulation of the model states in the concerned simulation period. For example, given a coupled model consisting of an atmosphere model and an ocean model and that the coupling lag from the atmosphere model to the ocean model is -600 s (i.e., the ocean model at its model time of 0 s will receive the coupling fields of the atmosphere model at its model time of 600 s), when the stop time of a simulation run is 86 400 s, the ocean model after the model time of 85 800 s will not obtain the coupling fields from the atmosphere model because the corresponding model time of the atmosphere model will be later than the stop time. Therefore, incorrect states of the ocean model may be obtained after its model time of 85 800 s. If users want to guarantee correct states of the ocean model before the model time of 86 400 s, the stop time should be extended to 87 000 s or later.

4.3.3 Steps for coupling procedure generation

The coupling procedure generation for a subset of component models follows the steps outlined below.

1. No matter which API is used to start the coupling procedure generation (subsection in Sect. 4.1.1), the coupling generator first confirms the subset of component models participating in the coupling procedure generation and confirms their MPI processes.
2. Determine all coupling connections. An export interface and an import interface will be connected for model coupling only when they have common coupling fields (with the same field names). As a component model manages its own coupling interfaces as well as coupling fields, an MPI communicator that includes all MPI processes in the subset of component models will be generated for aggregating the information of all coupling interfaces among different component models. As the user can also specify coupling connections through configuration files, file reading is required for analyzing possible coupling connections. To minimize the cost of reading, only one MPI process analyzes possible coupling connections, while other MPI processes await its results. An error will be reported if a coupling field in an import interface has multiple providers in this coupling procedure generation.
3. Generate a coupling procedure for each coupling connection. A coupling connection aims to couple a set of fields from an export interface to an import interface. When these interfaces belong to different component models, their models will exchange information about the corresponding timers, model grids, parallel decompositions, remapping configurations, data

types, etc. If a coupling field has different data types in the two coupling interfaces, an operation of data type transformation will be generated. The coupling generator adaptively selects a component model to execute the data type transformation for improved model coupling. For example, given that the data type in the export/import interface is *double/float*, the sender component model will transform the data type from *double* (8 bytes) to *float* (4 bytes), so that *float* values but not *double* values will be transferred from the export interface to the import interface. If a coupling field has different grids in the two coupling interfaces, a runtime algorithm for parallel data interpolation will be generated following the corresponding remapping configuration, where existing remapping weights will be used or new remapping weights will be read from an external data file or calculated by C-Coupler2 online and in parallel. Currently, only the receiver component model executes the parallel data interpolation. In the future, the coupling generator will adaptively select a component model to process data interpolation calculation for better coupling performance. When the import interface has been specified to import time-averaged coupling fields through the corresponding input parameter of the API *CCPL_register_import_interface* (please refer to the user guide for more details), operations for data averaging will be generated. To transfer the coupling fields from the export interface to the import interface, a runtime algorithm of non-block data transfer will be generated. In summary, a coupling procedure can include as necessary a runtime algorithm for data transfer, a runtime algorithm for data interpolation, operations for data type transformation, and operations for data averaging.

As a coupling procedure generation can be performed for any subset of component models, a component model can participate in multiple coupling procedure generations. In other words, the coupling procedures of a component model or even an import/export interface can be incrementally generated through multiple coupling procedure generations. For an import interface in a coupling procedure generation, only the import fields whose coupling procedures have not been generated will be considered in the coupling procedure generation, while the import fields whose coupling procedures have already been generated will be neglected.

4.4 Dynamic 3-D coupling capability

Given a 3-D grid that consists of a horizontal grid and a vertical grid with sigma or hybrid coordinates, the vertical coordinate values at each horizontal grid point are determined by a unique surface field on the horizontal grid. For example, the 3-D grid of an atmosphere model with sigma or hybrid coordinates will have constant vertical coordinate values when the surface field is terrain height, but the

values will be variable and change in time integration when the surface field is surface pressure, because the terrain height generally remains constant while the surface pressure changes in time integration. C-Coupler2 therefore provides two APIs, *CCPL_set_3D_grid_constant_surface_field* and *CCPL_set_3D_grid_variable_surface_field* (subsection in Sect. 4.1.1), for specifying constant and variable surface fields, respectively. Given a 3-D grid of an atmospheric chemistry model, the API *CCPL_set_3D_grid_variable_surface_field* can be used to specify the surface pressure as the surface field. As an atmospheric chemistry model generally does not produce the surface pressure, additional implementation will be required to enable an atmospheric chemistry model to obtain external surface pressure (e.g., from an atmosphere model in online model coupling). To facilitate coupling implementation for such a case, C-Coupler2 provides the API *CCPL_set_3D_grid_external_surface_field*, which, rather than specifying a surface field, states that the surface field of a 3-D grid is externally determined by the surface field of another 3-D grid. Moreover, the external surface field of a 3-D grid will be obtained automatically and implicitly by C-Coupler2 in model coupling.

As mentioned above, the 3-D interpolation involved in 3-D coupling is still performed in the “2-D + 1-D” manner in C-Coupler2, where 2-D interpolation between the horizontal sub-grids is performed first, followed by 1-D interpolation between the vertical sub-grids. Given a 3-D interpolation from a source 3-D grid (expressed as $H2D_s + V1D_s$) to a target grid ($H2D_t + V1D_t$), the 2-D interpolation between the horizontal sub-grids $H2D_s$ and $H2D_t$ eventually interpolates coupling fields from the source 3-D grid to an intermediate 3-D grid consisting of $H2D_t$ and $V1D_s$, and thus the 1-D interpolation between the vertical sub-grids $V1D_s$ and $V1D_t$ eventually interpolates coupling fields from the intermediate 3-D grid to the target 3-D grid. Specifically, the 2-D interpolation will be performed a number of times, each of which corresponds to a horizontal level of the source and intermediate 3-D grids, and the 1-D interpolation will also be performed for a number of times, each corresponding to a column in the intermediate and target 3-D grids. For dynamic 3-D coupling, 2-D interpolation can use pre-calculated remapping weights, because the horizontal sub-grids do not change throughout a simulation, while 1-D interpolation cannot use pre-calculated remapping weights, and instead must dynamically calculate the remapping weights according to the changes in vertical coordinate values in the source or target 3-D grid in time integration. To complement dynamic 3-D interpolation based on the implementation of static 3-D interpolation in C-Coupler1, dynamic calculation for 1-D remapping weights is implemented with the following steps in C-Coupler2.

1. If the source 3-D grid has a variable surface field, the import interface first receives the source surface field

transferred from the export interface, and next uses the pre-calculated horizontal remapping weights to interpolate the source surface field from the source horizontal grid (the horizontal sub-grid of the source 3-D grid) to the target horizontal grid (data interpolation will be bypassed if the two horizontal grids are the same). The source surface field on the target horizontal grid will be used as the surface field of the intermediate 3-D grid, and will be further used as the target surface field when the target 3-D grid has an external surface field.

2. If the source 3-D grid has an external surface field, the import interface uses the target surface field as the surface field of the intermediate 3-D grid (the target 3-D grid must have a non-external surface field in this case).
3. The import interface calculates the vertical coordinate values of the intermediate/target 3-D grid when the 3-D grid has a surface field (the import interface can obtain all constant information of the source 3-D grid in coupling procedure generation before the first execution of the corresponding export and import interfaces).
4. For each column in the intermediate or target 3-D grid, the import interface calculates the 1-D remapping weights.

As dynamic 3-D interpolation cannot fully utilize pre-calculated remapping weights and must update 1-D remapping weights at almost all coupling steps, it has a higher computational cost than static 3-D interpolation. To minimize the impact of the increased computation cost, all of the above steps, including data transfer for the source surface field, 2-D interpolation for the source surface field, calculation of vertical coordinates of the intermediate/target 3-D grid, and calculation of 1-D remapping weights, are parallelized based on the MPI processes and parallel decompositions in the corresponding component models. Moreover, the implementations of static 3-D interpolation and dynamic 3-D interpolation are unified. In detail, static 3-D interpolation will be treated as dynamic 3-D interpolation at the first step of coupling; the 1-D remapping weights will thus be calculated online during the first step of coupling; and they will be treated as static 3-D interpolation and use the existing remapping weights in the subsequent coupling steps.

4.5 Non-blocking data transfer

Non-blocking data transfer is a necessary function for achieving coupling within one executable or a subset of MPI processes. Two-sided MPI communication (e.g., *MPI_Send*, *MPI_Recv*, and their non-blocking modes *MPI_Isend*, *MPI_Irecv*) has been widely used in existing couplers for data transfer. It is still used in C-Coupler2 as a default option of data transfer. To achieve non-blocking data transfer based on two-sided MPI communication, the import interfaces of C-Coupler2 are blocking while the export interfaces use the

non-blocking mode; a mechanism that afterwards checks and waits for the completion of the last execution of export interfaces. Similar non-blocking data transfer has already been implemented in OASIS3-MCT.

Unpredictable “deadlocks” in non-blocking two-sided MPI-communication-based data transfer can occur when an excessive number of messages sent to a single MPI process exhaust the message passing buffer space allocated by the MPI library (Dennis et al., 2012). We therefore implemented an additional option of non-blocking data transfer based on one-sided MPI communication (i.e., *MPI_Put* and *MPI_Get*), which enables C-Coupler2 to manage the message passing buffer space, so as to ensure a “safe” implementation of non-blocking data transfer. If unpredictable “deadlocks” happen when using the default two-sided MPI-communication-based data transfer, users will be advised to try the one-sided MPI-communication-based implementation. In fact, we have encountered such a case when using our sample coupled model for software testing (Sect. 5).

The wait for two-sided MPI communication is easily achieved through the MPI function *MPI_Wait*. However, there is no corresponding MPI function available for one-sided MPI communication, and thus extra effort is required to implement such a mechanism. In detail, non-blocking data transfer based on one-sided MPI communication is implemented as follows in C-Coupler2.

1. The export interface calls *MPI_Put* to send coupling field values to the message passing buffer managed by the import interface. Before sending coupling field values, the export interface examines whether the message passing buffer of the import interface is available. The message passing buffer remains unavailable until the import interface has received coupling field values from the last data transfer.
2. Before the import interface copies out coupling field values from its message passing buffer, it first checks whether new coupling field values have been received. Besides coupling field values, the export interface will send a model time tag to the import interface at the same time. If coupling field values are only associated with tags later than the tag of the last receive, it means that new coupling field values have been received. After copying out coupling field values, the message passing buffer is set as available. Another model time tag is used for identifying the status of the message passing buffer (i.e., available or unavailable). The export interface uses *MPI_Get/MPI_Put* to query/update the status of the message passing buffer.
3. The data receive command issued by the C-Coupler2 API calls is blocking. Beyond the API calls, C-Coupler2 issues additional non-blocking data receives for import interfaces, in order to make the data receive finish and the data send execute as early as possible. In a non-

blocking data receive, if new coupling field values have been received, C-Coupler2 will copy coupling field values and set the message passing buffer as available; otherwise, C-Coupler2 will do nothing.

As noted above, C-Coupler2 provides flexibility in setting a lag on a coupling connection via the “*remote_lag_count*” in the timer of the import interface. One challenge associated with this function is that a deadlock can occur if the lag is greater than the corresponding coupling period. For example, Fig. 7 includes two component models (*comp1* and *comp2*) coupled with two connections. The first coupling connection is from the export interface *exp1* of *comp1* to the import interface *imp2* of *comp2*, while the second is from the export interface *exp2* of *comp2* to the import interface *imp1* of *comp1*. Both connections have equal coupling periods of 600 s. At each coupling step of each model, the export interface is executed before the import interface. In Fig. 7a, there is no lag on each coupling connection, and thus *comp1* and *comp2* can run concurrently. In Fig. 7b, there is no lag on the second coupling connection, while the first coupling connection has a lag of 600 s, which means that *imp2* at the current coupling step will receive the coupling fields from *exp1* at the previous coupling step. At the first coupling step, *exp1_1* (meaning *exp1* executed at the first coupling step) tries to send coupling fields to *imp2_2*. As non-blocking data transfer is used, *exp1_1* can successfully put the coupling fields into the message passing buffer of *imp2*, and thus *comp1* can finish *imp1_1*, and so finish the first coupling step. At the same time, *comp2* can finish the first coupling step (*imp2_1* will not be executed, because it corresponds to *exp1* executed at the model time of -600 s that is earlier than the start time of the model run; the coupling field values imported by *imp2* can be initialized via data files or coupling in the initialization stage of the coupled model). At the second coupling step, *exp1_2* will first await the message passing buffer of *imp2* that still keeps the coupling fields from *exp1_1*. After *comp2* finishes *exp2_2*, *imp2_2* is executed, and the message passing buffer of *imp2* will be set as available. Next, *exp1_2* can successfully put the coupling fields into the message passing buffer of *imp2*. Although the lag of 600 s on the second coupling connection does not introduce a deadlock, a problem in the sequencing of *comp1* and *comp2* is introduced: *exp1* in *comp1* must await the completion of *imp2* in *comp2* at the same coupling step. In Fig. 7c, there is no lag on the second coupling connection, while the first has a lag of 1200 s. Similarly to Fig. 7b, both *comp1* and *comp2* can finish the first coupling step. At the second coupling step, *exp1_2* first awaits the message passing buffer of *imp2* that will not be set as available before *comp2* finishes *imp2_3* at the third coupling step, while *exp2_2* can successfully put the coupling fields into the message passing buffer of *imp1*, which has been set as available by *imp1_1* at the first coupling step. At the third coupling step of *comp2* (*exp1_2* in *comp1* is still waiting at the second coupling step), *exp2_3* first awaits the

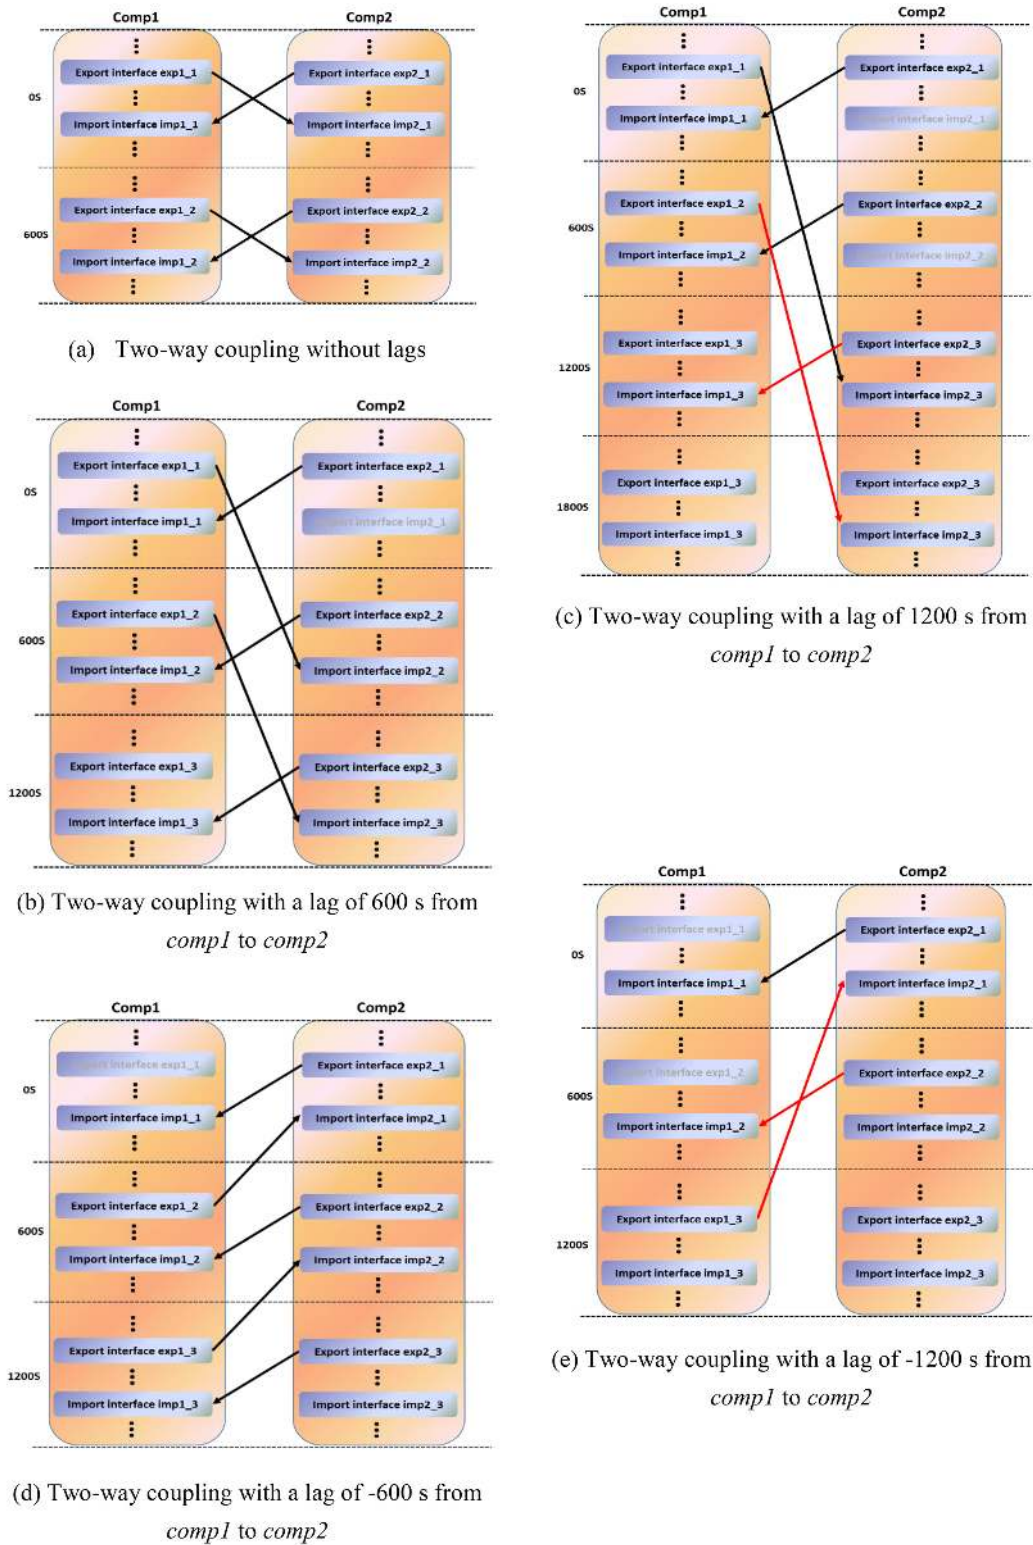


Figure 7. Sample two-way couplings with different lag settings. A coupling interface with light grey words means that it will not be executed, because it corresponds to the coupling interface executed at the model time earlier than the start time (0 s) of the model run.

message passing buffer of *imp1* that will not be set as available before *comp1* finishes *imp1_2*. As a result, both *comp1* and *comp2* wait for each other, causing a deadlock.

The sequencing problem in Fig. 7b and the deadlock in Fig. 7c result from the unavailability of message passing buffers. They are almost unsolvable regarding to two-sided MPI communication, because the message passing buffer space is allocated by the MPI library. However, they are not unbreakable when regarding one-sided MPI communication, and can be avoided by increasing the corresponding message passing buffers. In C-Coupler2, the message passing buffers for one-sided MPI-communication-based data transfer can be increased adaptively. At each time step of a component model, C-Coupler2 checks each import interface and will adaptively increase the message passing buffers when required. As a result, a positive lag on a coupling connection will never result in a sequencing problem or a deadlock between component models, when the one-sided MPI-communication-based non-blocking data transfer is used. The extra memory usage due to the increased message passing buffers would be affordable in most cases, because generally only a small proportion of model fields are involved in model coupling.

A negative lag can also be specified for a coupling connection, but can result in sequencing problems between component models or even an unbreakable deadlock. For example, Fig. 7d shows no lag on the second coupling connection and a lag of -600 s on the first, which means that *imp2* at the current coupling step will receive the coupling fields from *exp1* at the next coupling step. This lag setting will not introduce a deadlock, but will introduce a sequencing problem between *comp1* and *comp2*: *imp1* is coupled with *exp2* at the same step, while *imp2* at the current coupling step waits for *exp1* at the next coupling step. In Fig. 7e, there is no lag on the second coupling connection, while the first has a lag of -1200 s, which introduces an unsolvable deadlock (corresponding to the red arrows in the figure) even with one-sided MPI communication, where import interfaces are awaiting the export interfaces that cannot be executed until the import interfaces return.

4.6 Facilitation for model nesting

C-Coupler2 can help achieve self-nesting in a regional model that does not originally support this possibility, as follows:

1. The code of the regional model can still only manage a unique grid domain, but multiple grid domains for self-nesting can be achieved through running multiple copies of the executable of the regional model on non-overlapping MPI processes, while each copy can have separate input parameters and input data files for a unique grid domain and can be registered as a separate component model of C-Coupler2. The different grid domains should have different component model

names, but they can use the same names for the model grids, parallel decompositions, coupling fields, coupling interfaces, etc. Therefore, C-Coupler2 only requires the regional model to obtain a few additional input parameters. In other words, slight modification of the *namelist* file and the corresponding model code of the regional model can enable C-Coupler2 to recognize multiple grid domains.

2. Given that a small grid domain is nested in a larger grid domain, C-Coupler2 can recognize the relationship between the two grid domains through the coupling connection configuration files. As all grid domains can correspond to identical code in the regional model, the coupling connection configuration files of different grid domains can be similar, differing only in terms of the full names of component models in the file contents. Thus, the coupling connection configuration files of all grid domains can be generated easily or even automatically by a script.
3. The self-nesting capability requires the exchange of model fields that are generally 3-D between grid domains. Implementation of this exchanging can benefit from the 3-D coupling capability, especially the dynamic 3-D coupling capability of C-Coupler2. Moreover, given that a small grid domain is nested in a larger grid domain, the coupling procedures for exchanging model fields between them can be automatically generated in their partial coupling procedure generation.

C-Coupler2 does not provide any lateral boundary condition scheme. This is not a problem, because a regional model generally includes lateral boundary condition schemes that can also be used in self-nesting. To achieve two-way self-nesting, schemes for using the feedback from smaller grid domains should be added to the regional model.

As each grid domain corresponds to a separate copy of the executable, each grid domain can easily use a distinct subset of MPI processes, which allows for simultaneous integration of grid domains for better parallel performance. Scientists may want to integrate a grid domain earlier than its nested grid domains. For example, after a grid domain finishes integration from 0 to 90 s, its nested grid domains can start the integration from 0 s. This can be achieved by adjusting coupling lags among grid domains based on the timers of the corresponding import interfaces. In one-way nesting, the coupling lags generally do not affect the parallelism among grid domains. In other words, different grid domains can always be integrated simultaneously in one-way nesting, regardless of the coupling lags. In two-way nesting, even when a specific setting of coupling lags forces sequential running between a grid domain and its nested grid domains, multiple nested grid domains of the same grid domain can also run simultaneously, so that C-Coupler can also help improve the parallel performance of self-nesting. One challenge resulting

from sequential running is that the corresponding processors will be essentially idle, and therefore wasted, when a grid domain is waiting for another grid domain. In the future, we will investigate technical solutions to overcome this inefficiency. Moreover, it may be an interesting topic to investigate the scientific impact of different settings of coupling lags in model nesting.

Similarly, C-Coupler2 can be used to nest a regional model into a different model. For a model that already has self-nesting capability (such as WRF), all grid domains and the field instances on each grid domain are allocated and managed by the regional model itself, while each grid domain as well as the field instances on it can be registered to C-Coupler2 as a component model. Therefore, a grid domain in a self-nesting model can be further coupled with another kind of grid domain or component model through C-Coupler2.

4.7 Facilitation for incremental coupling

Incremental coupling can be viewed as coupling external component models with an existing coupled model. A straightforward implementation is to treat the external component models as internal component models of the existing coupled model, and use the coupler of the existing coupled model to handle the corresponding incremental coupling. For example, regarding the work of nesting WRF into CESM done by He et al. (2013) that has been introduced in Section 3.7, WRF is treated as an internal component model of CESM, and the incremental coupling for its nesting is handled by CPL7, the coupler of CESM. A major challenge in this kind of implementation is that the independence might need to be broken between external component models and the existing coupled model that may have been developed independently by different groups for a number of years. This introduces significant code changes to the models (even including the coupler), and results in inconsistent code versions of the same model among different model groups. For example, all component models of CESM share the same driver and are compiled into a unique executable, while WRF has its own driver, different from the others. When treating WRF as an internal component model of CESM, WRF will have to use the driver of CESM, and will also be compiled into the unique executable. Thus, WRF's original driver and compiling scripts as well as CESM need to be modified. Moreover, as the original driver and coupler of CESM do not consider the existence of a regional atmosphere model, the driver and coupler codes of CESM also need to be modified.

Incremental coupling faces the fundamental problem of guaranteeing independence between external component models and the existing coupled model, so as to minimize code changes to the models or the coupler. To help in this regard, C-Coupler2 should minimize the constraints on using external component models and existing coupled models that are already coupled with other couples; it should also work as an additional coupler specifically for incremental coupling as

part of coupling in a new coupled model, thus letting developers focus only on the coupling between external component models and the corresponding component models in the existing coupled model. In response to these requirements, C-Coupler2 includes the following implementations for incremental coupling.

1. An existing coupled model can be registered to C-Coupler2 as a component model, and its component models involved in incremental coupling can be further registered as its children, while other component models that are irrelevant to incremental coupling can be neglected. Generally, several API calls are enough for the model registration, which only introduces slight code changes to the existing coupled model. As C-Coupler2 can support almost any MPI process layout among component models, a component model in any existing coupled model can be easily registered to C-Coupler2. Similarly, an external component model can be easily registered to C-Coupler2.
2. As C-Coupler2 allows coupling procedure generation to be performed for any subset of component models, partial coupling procedure generations for only the component models relevant to incremental coupling can be performed flexibly. Similarly, several API calls are enough for partial coupling procedure generations, which only introduces a slight code change to the external component models and the existing coupled model.

4.8 Adaptive restart capability

In a coupled model, all component models generally share a unique restart timer that notifies all component models to prepare restart data corresponding to the same model time (called restart write model time hereafter). It does not mean that the restart data files should only include the data values at the restart write model time. Examples 1–3 in Sect. S2 indicate that the restart data files corresponding to a restart write model time may need to include the coupling field instance values at different model time corresponding to a positive lag on a coupling connection or an order between writing restart data files and advancing model time, and that it may need to bypass the data export of some export interfaces at some model time after restarting the coupled model run.

To conveniently achieve exact restart for coupling fields under any setting of coupling lags and any order between writing restart data files and advancing model time, the restart manager of C-Coupler2 provides an adaptive restart capability implemented as follows (For further illustration of the implementation, please refer to example 4 in Sect. S2).

– *The restart manager conducts restart writing as follows:*

1. The restart data corresponding to a component model are classified into two categories: management information that will be written into a binary

formatted data file at once, and field instance values that will be written into a NetCDF formatted data file incrementally.

2. Each active component model that has been registered to C-Coupler2 should call the API `CCPL_do_restart_write_IO` at each time when the restart timer is on (Generally, this API will be bypassed when the restart timer is not on).
 3. `CCPL_do_restart_write_IO` will set the restart writing model time as the current model time, optionally write the current values of all restart-related field instances into the NetCDF restart data file based on the user's configuration, and prepare the restart management information that will be written into the binary restart data file later.
 4. When a receiver component model is importing the values of a coupling field instance from a sender component model, if the model time at the sender or receiver component model is not later than the restart writing model time, the values obtained by the receiver component model will be written into the corresponding NetCDF restart data file automatically and incrementally.
 5. Restart management information will be written into a binary restart data file only when the corresponding component model is advancing its model time and has already received all values with the sender's model time earlier than the restart writing model time.
- *The restart manager conducts restart reading as follows:*
6. When users want to restart a coupled model run, each active component model should firstly call the API `CCPL_start_restart_read_IO` (can be called at most once and will be bypassed in an “initial” run) which will read in the restart management information from the corresponding binary restart data file, and the model time corresponding to the restart management information is marked as restarted model time.
 7. A component model can call the API `CCPL_restart_read_fields_all` or `CCPL_restart_read_fields_interface` (both APIs can be called multiple times and redundant restart read of the same values will be avoided in multiple calls; they will be bypassed in an “initial” run) to read in the values of some field instances at the restarted model time from the corresponding NetCDF restart data file, if required.
 8. In a “continue” run or a “hybrid” run, when a sender component model tries to export a coupling field instance to a receiver component model, if the corresponding model time at the sender or receiver component model is not later than the restarted model time, the sender will bypass the field instance export; when a receiver component model tries to import a coupling field instance from a sender component model, if the corresponding model time at the sender or receiver component model is not later than the restarted model time, the values of the import field instance will be automatically read from the corresponding NetCDF restart data file but not imported from the sender component model (an error will be reported if the corresponding restart data file does not contain the import field instance values).
- *The restart manager provides the following supports for the “continue” run that will automatically and exactly restart the model run from a previous restart write model time:*
9. C-Coupler2 automatically records the latest restart write model time of each component model in an implicit file called “*rpointer*” file that will be updated only when the corresponding binary restart data file has been produced.
 10. Besides the latest restart write model time, C-Coupler2 will also automatically and implicitly record the previous restart write model time before the latest restart write model time for each component model.
 11. When starting a “continue” run, C-Coupler2 will adaptively determine the right restarted model time (i.e., the latest restart write model time or the previous restart write model time) and then restart the model run.
 12. A component model can also get the right restarted model time determined by C-Coupler2 through the API `CCPL_get_restart_setting` (please refer to the user guide for details).
- A coupling lag can be adaptively achieved through setting the remote lag count in the timer of an import interface when the import interface is executed without bypassing its timer. However, even when constantly bypassing the timer of an import interface, a coupling lag can also be achieved through adjusting the execution of the import interface and the corresponding export interfaces. Please note that, C-Coupler2 currently does not guarantee exact restart capability under such kind of coupling lag specification. We therefore highly recommend users to enable the timer when executing a coupling interface as soon as possible, especially in the main loop of the time integration.
- For a summary, when the timer of each coupling interface is enabled in the main loop of the time integration, C-Coupler2 will guarantee exact restart for coupling field

values no matter the setting of coupling lags and no matter the implementation of coupling. In a “*continue*” run, C-Coupler2 will adaptively determine the right model time to exactly restart the simulation. In a “*branch*” run or “*hybrid*” run, C-Coupler2 will stop the model simulation with an error report if the restarted model time specified by users is wrong (e.g., some corresponding restart data files are missing or do not contain necessary coupling field instance values). Besides coupling fields, C-Coupler2 also provides the support of exact restart capability for the model fields irrelevant to coupling.

4.9 Debugging capability

The following aspects enhance the debugging capability of C-Coupler2.

1. C-Coupler2 performs a series of checks for almost all API calls. For example, when registering a component model, model grid, parallel decomposition, coupling field, or coupling interface, when setting the time step of a component model, and when executing a coupling interface, C-Coupler2 can check whether all MPI processes of the component model call the API at the same time and with consistent parameters. For example, when registering a horizontal grid with global grid data or registering a vertical grid, C-Coupler2 can check whether the grid data are the same among MPI processes, and when registering a coupling interface, C-Coupler2 can check whether the timer, coupling field instances, and other parameters are consistent among MPI processes. When an API call includes an array as a parameter, C-Coupler2 can check the size of the array. For example, when registering a coupling field instance, C-Coupler2 can check whether the array size of the memory buffer of the coupling field instance matches that required. When an API call includes the ID of a coupling resource as a parameter, C-Coupler2 can check whether the ID is legal. When an API call reads information in configuration files, C-Coupler2 can check whether the files are in the correct XML format and check the correctness of the required information. Given the additional overheads in computation and communication introduced by performing such checks, most of them can be disabled in a model run. We strongly recommend that the user enables the checks fully when developing a coupled model.
2. When an error or a warning is detected, it will be reported, including a suggestion for fixing the relevant model codes or configuration files. Almost all APIs include an optional input parameter “*annotation*”, which is a string giving a hint for locating the model code of the API call corresponding to an error or warning. There are around 1000 error reports throughout the code of C-Coupler2.

3. C-Coupler2 can report many kinds of log information about coupling configurations, progress in handling coupling configurations, coupling procedure generations, coupling runs, and the values of coupling field instances. Each process in a coupled model can have a separate log file for reporting log information, errors, and warnings, which can facilitate parallel debugging. Moreover, C-Coupler2 also enables the internal code of a component model to report log information through the C-Coupler2 log files. The user can enable or disable the reporting of log information.

5 Evaluation

This section evaluates C-Coupler2 in several aspects, including software testing, scaling of initialization, data transfer, memory use, and dynamic 3-D coupling.

5.1 Software testing

To improve the reliability of C-Coupler2 in various areas of application, we first designed a sample coupled model that includes coupling between several sample component models and self-nesting component models. Next, we developed hundreds of use cases based on the sample coupled model, to evaluate whether C-Coupler2 properly detects and reports errors in various cases of incorrect use, properly generates coupling procedures and handles model coupling and nesting in correct cases.

Besides the sample coupled model, existing real coupled models were used to test C-Coupler2, including FIO-AOW, BCC_CSM (Beijing Climate Center Climate System Model; Xin et al., 2013), and CESM. For each coupled model, C-Coupler2 was used to replace some coupling functions from other couplers, while trying to keep exactly the same (bit-wise identical) simulation results. Specifically, C-Coupler2 was used to replace C-Coupler1 employed in FIO-AOW, MCT-based data transfer (Larson et al., 2005) and interpolation functions of the CPL5 coupler (the version immediately prior to CPL6) employed in BCC_CSM, and MCT-based data transfer and interpolation functions of the CPL7 coupler employed in CESM. Various MPI process numbers of component models and various run types (i.e., “*initial*”, “*continue*”, “*branch*”, and “*hybrid*” run) of each real coupled model were used for testing using nearly 2000 processor cores. As CESM enables the user to flexibly set MPI process layouts among component models, the tests considered non-overlapping, partially overlapping, and overlapping MPI processes among component models of CESM.

We further coupled an atmosphere model GAMIL2 (Grid-point Atmospheric Model of IAP LASG version 2; Li et al., 2013b) and the CESM version with C-Coupler2 (the original atmosphere model was disabled) via incremental coupling, which generates the executables of both GAMIL2 and

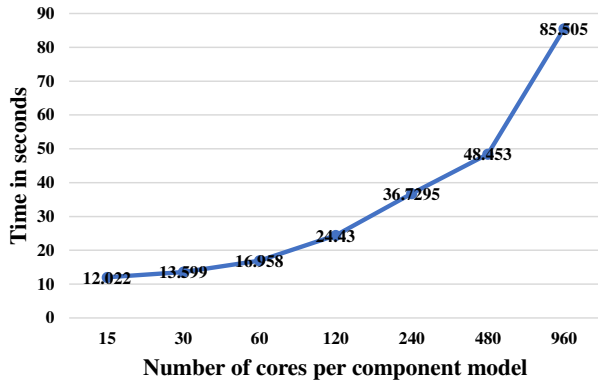


Figure 8. Initialization cost for coupling two trial models with C-Coupler2 on a supercomputer with Intel Xeon CPUs and an InfiniBand network.

CESM, while GAMIL2 keeps its original driver unchanged. C-Coupler2 thus successfully demonstrated incremental coupling. The dynamic 3-D coupling capability of C-Coupler2 was evaluated when coupling GAMIL2 and GEOS-Chem. A coupled model with the atmosphere model in MPAS (Model for Prediction Across Scales) (<https://mpas-dev.github.io/>, last access: 24 August 2018) and a wave model WaveWatch (polar.ncep.noaa.gov/waves/wavewatch/, last access: 24 August 2018) has been newly developed, which also contributes to software testing of C-Coupler2.

Moreover, various processors (i.e., Intel X86, IBM Power, and the SW26010 processors employed in the Sunway TaihuLight system), various operating systems (i.e., Linux and IBM AIX), various compilers (i.e., Intel compilers, GNU compilers, and IBM compilers), and various MPI libraries (Intel MPI, MPICH, Open MPI, and IBM MPI) were involved in testing C-Coupler2.

5.2 Initialization cost

The evaluation of initialization cost considered coupling 10 2-D fields between two trial component models that define horizontal grids but do not have real model initialization. The component models' horizontal grids were a regular longitude–latitude grid with 1440×720 grid points and a tripolar grid with 1440×1021 grid points. Therefore, model coupling required data remapping, and we used the corresponding remapping weights file generated by CoR1. The two component models ran concurrently on a supercomputer, with the same number of processor cores (MPI processes). Each computing node on the supercomputer included two Intel Xeon CPUs, with 20 processor cores in total, and all computing nodes were connected with an InfiniBand network. The codes were compiled by an Intel Fortran and C++ compiler at the optimization level O2, using an Intel MPI library.

Figure 8 shows the initialization cost of C-Coupler2 when scaling the number of processor cores of each component

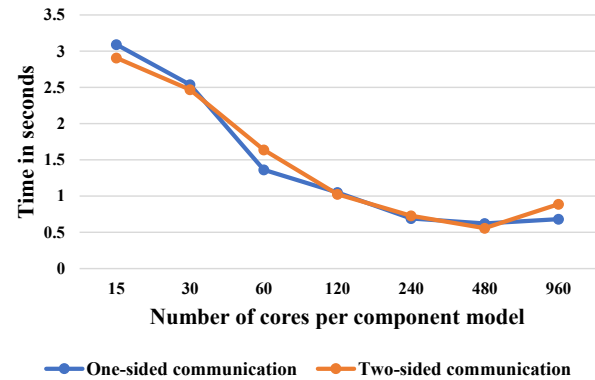


Figure 9. Comparison of data transfer times (for 100 ping-pong couplings) between a one-sided and a two-sided implementation, with the same configuration as Fig. 8.

model from 15 to 960. It increases with increasing core number. C-Coupler2 initialization consists of several steps, including registering component models, setting time steps, registering model grids, registering parallel decompositions, registering coupling field instances, registering coupling interfaces, generating coupling procedures, reading and then distributing the remapping weights from a file, and computing routing information for data transfer between component models. Most of these steps include non-scalable operations, i.e., MPI collective communications or I/O. Increasing the core number increases the overhead of the non-scalable operations, and thus increases initialization cost. The results in Fig. 8 indicate that the initialization cost of C-Coupler2 may not be negligible, but it would be affordable in most cases, especially for long-term climate simulations.

5.3 Data transfer

As introduced above, C-Coupler2 can employ one-sided MPI communication in data transfer, while other existing couplers only use two-sided MPI communication in data transfer. We evaluated the performance of our one-sided communication in comparison to a two-sided implementation, based on a ping-pong coupling for the same configuration used in Sect. 5.2. Figure 9 shows the cost of the two implementations when scaling the number of processor cores of each component model from 15 to 960. The times are per 100 ping-pong couplings. Overall, the one-sided communication achieves similar performance to the two-sided communication. In other words, the option of one-sided MPI communication does not obviously degrade the performance of data transfer.

5.4 Memory usage

Figure 10 shows the memory use per core for the coupled model configuration used in Sect. 5.2, as measured using

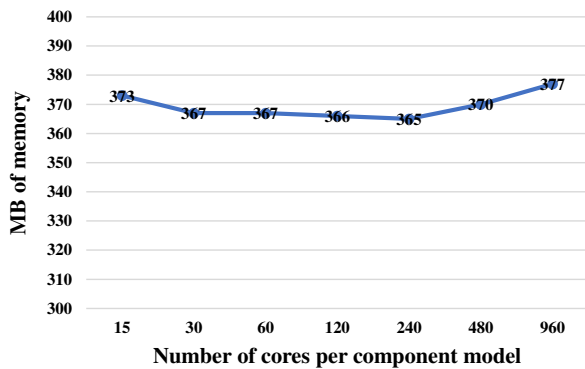


Figure 10. Memory use of C-Coupler2 for the trial coupled model considered in Sect. 5.2.

the gptl (<https://github.com/jmrosinski/GPTL>, last access: 24 August 2018) interface (similar to Craig et al., 2017). The memory usage remains around 360–380 MB regardless of the core number from 15 to 960. Lacking computing resources, we were unable to evaluate the memory usage at a much higher number of processor cores. When more processor cores are used, the memory use per core for the whole coupled model configuration will get higher owing to the MPI memory footprint (Balaji et al., 2009). However, the memory use per core due to C-Coupler2 itself would not get higher because the number of local grid cells per core will get lower.

5.5 Dynamic 3-D coupling

The capability of dynamic 3-D coupling was tested based on the coupled model consisting of GAMIL2 and GEOS-Chem that has been mentioned before. In this coupled model, the 3-D grids of GAMIL2 and GEOS-Chem are different in both horizontal direction and vertical direction. Specifically, GAMIL2 includes a 2.8° horizontal grid with the uniform grid in the low and middle latitudes region and a weighted even-area grid in the high latitudes and polar region, and 26-sigma vertical levels (pressure normalized by surface pressure) with the model top at 2.194 hPa, while GEOS-Chem uses a $4^\circ \times 5^\circ$ uniform horizontal grid and 72-hybrid vertical levels with the model top at 0.01 hPa. As GEOS-Chem does not calculate the atmospheric pressure, the surface pressure corresponding to its 3-D grid is declared as an external surface field through the API `CCPL_set_3D_grid_external_surface_field` and will be determined by the surface field of the 3-D grid of GAMIL in dynamic 3-D coupling. For the first step of the evaluation, we examined the vertical coordination values of the two 3-D grids and the corresponding vertical remapping weights that change in time integration. Next, we examined each coupling field from one component model to the other. For example, Figs. 11 and 12 show the values of the coupling fields

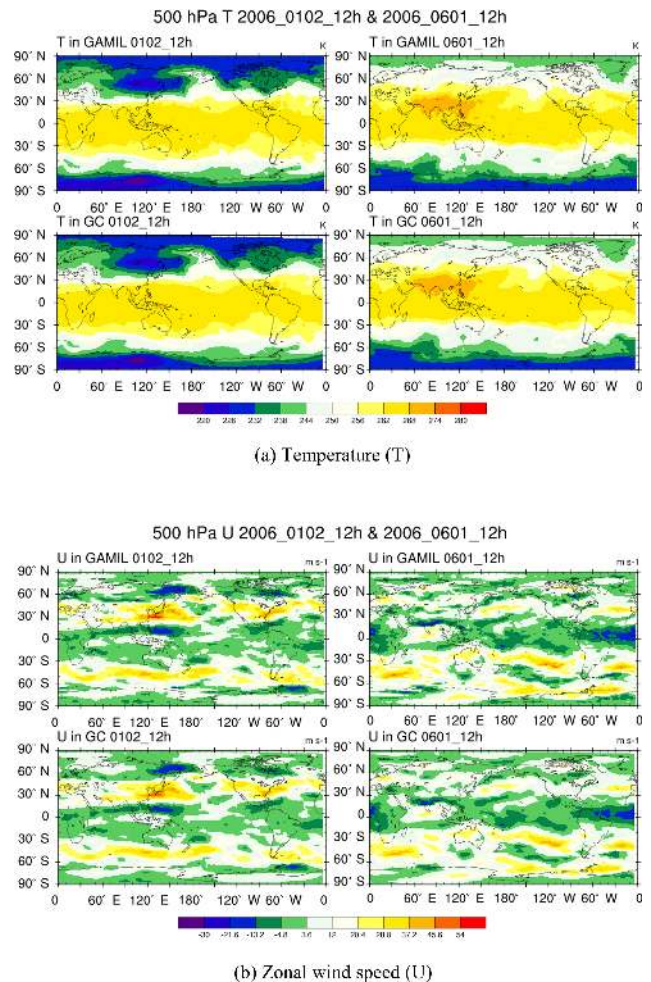


Figure 11. The temperature, T (a), and zonal wind speed, U (b), from GAMIL2 to GEOS-Chem (GC) at the 500 hPa level at two different model times.

of atmospheric temperature (T) and the zonal wind speed (U) from GAMIL2 and GEOS-Chem, which reveal that C-Coupler2 achieves consistent results between the source grid and the target grid in dynamic 3-D coupling.

To evaluate the parallel scalability of dynamic 3-D coupling, we derived a new configuration from the configuration used in Sect. 5.2, where the 2-D grid corresponding to the 2-D coupling fields in each component model is extended to be a 3-D grid with a 50-level vertical sub-grid of sigma coordinate. The sigma coordinate values are different between the two component models, and the surface field corresponding to the 3-D grid of one component model changes at every time step, while the surface field corresponding to the other component model has been set to be external. Therefore, the vertical coordinate values of all 3-D grids and the corresponding vertical remapping weights are updated in each time of 3-D coupling. Figure 13 shows the parallel speedup of dynamic 3-D coupling when scaling the number of processor cores of each component model from 15 to 960, measured

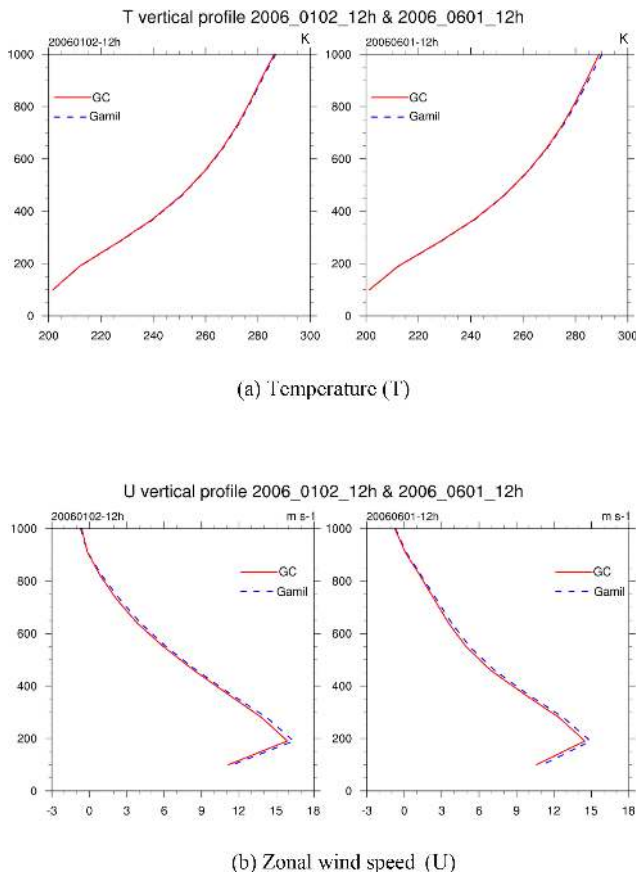


Figure 12. The global vertical profile of the temperature T (a), and zonal wind speed U (b), from GAMIL2 to GEOS-Chem (GC) at the 500 hPa level at two different model times.

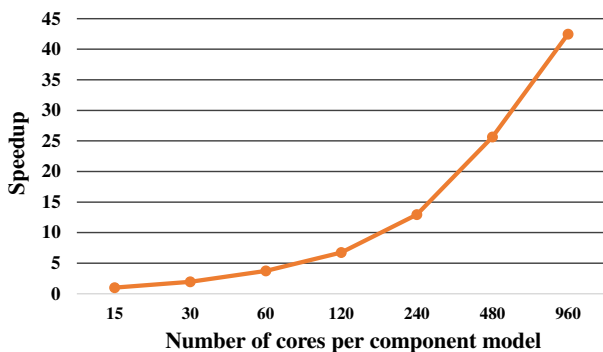


Figure 13. The parallel speedup of dynamic 3-D coupling (for 100 ping-pong couplings) between the two component models, with a new configuration derived from the configuration used in Fig. 8. The speedup is normalized to the time at 15 cores per component model (1583 s).

from 100 ping-pong couplings. The results in Fig. 13 reveal that C-Coupler2 can achieve acceptable parallel scalability in dynamic 3-D coupling.

6 Summary and future work

As a new version of C-Coupler, C-Coupler2 follows the family's targets and the main design, but is significantly different from C-Coupler1 with advancements in many aspects, e.g., coupling configuration, 3-D coupling capability, coupling procedure generation, support for incremental coupling, coupling lags, etc. (please refer to Table 2 for the differences between C-Coupler1 and C-Coupler2 in all aspects):

1. The coupling configuration of C-Coupler2 properly combines nearly 80 APIs and several kinds of XML-formatted configuration files, while the coupling configuration of C-Coupler1 strongly depends on ASCII formatted configuration files. Therefore, the coupling configuration of C-Coupler2 is much more flexible and user-friendly.
2. C-Coupler2 supports dynamic 3-D coupling that enables vertical coordinate values of 3-D grids to change in time integration, while C-Coupler1 only supports static 3-D coupling where coordinate values of 3-D grids must remain constant throughout time integration.
3. C-Coupler2 has a coupling procedure generation capability that achieves automatic and incremental coupling procedure generation for any subset of component models, while C-Coupler1 does not achieve such a capability.
4. C-Coupler2 can facilitate the implementation of incremental coupling because an existing coupled model as well as its component models coupled with any coupler can be used as component models of C-Coupler2, while C-Coupler1 does not support coupled component models using other couplers.
5. The APIs non-blocking data transfer and adaptive restart capability of C-Coupler2 enable users to flexibly set coupling lags that can be greater than the corresponding coupling periods, while a coupling lag in C-Coupler1 cannot exceed the corresponding coupling period.

As shown in Table 2, C-Coupler1 has the capability of integrating external calculation routines, which means that a private subroutine of a component model or a common algorithm such as a flux calculation algorithm can be registered as an external calculation routine. An external calculation routine cannot have any explicit argument while its inputs and outputs are implicitly specified through the corresponding configuration files. The integration of a Fortran external calculation routine generally requires an additional C interface. An external calculation routine can be further used as a runtime algorithm in a coupling procedure by specifying it in the corresponding configuration files. C-Coupler2 does not inherit this capability from C-Coupler1, because configuration files for implicitly specifying the inputs and outputs

Table 2. Differences between C-Coupler1 and C-Coupler2.

No.	Technical aspects	C-Coupler1	C-Coupler2
1	Coupling configuration	Strongly depends on ASCII formatted configuration files	Properly combines APIs and XML-formatted configuration files
2	MPI process layout of coupling	Only root component models are supported; cannot handle model coupling within a subset of MPI processes or the same component model	Can support almost any kind of MPI process layout among component models (including existing coupled models); can handle model coupling within a subset of MPI processes or the same component model
3	3-D coupling capability	Static 3-D coupling only	Both static and dynamic 3-D coupling
4	Coupling procedure generation	No coupling procedure generation function	Automatic and incremental coupling procedure generation for any subset of component models
5	Data transfer	Blocking data transfer based on two-sided MPI communication	Non-blocking data transfer is implemented with two-sided MPI communication by default, and with one-sided MPI communication as an additional option
6	Support for model nesting	No specific support	Can facilitate nesting a regional model into itself or into another model
7	Support for incremental coupling	No specific support	Can facilitate coupling external component models with an existing coupled model
8	Debugging capability	Not prioritized; little support	Implemented, with support provided
9	Coupling lags	Supported, but lag cannot be greater than the corresponding coupling period	Supported, with lag able to be greater than the corresponding coupling period
10	Coupling procedures	Runtime algorithms in a coupling procedure are explicitly specified in a configuration file	Runtime algorithms in a coupling procedure are implicitly generated by the coupling generator
11	Restart capability	Can achieve exact restart for model field instances and coupling field instances, without supporting the case with a coupling lag greater than the corresponding coupling period.	Can achieve exact restart for model field instances and coupling field instances, no matter the setting of coupling lag. Can automatically determine a right model time for restarting a “ <i>continue</i> ” run
12	Online remapping weight generation	Sequential online remapping weight generation	Parallelized online remapping weight generation
13	Integration of external calculation routines	Can integrate an external calculation routine as a runtime algorithm and then further use it in a coupling procedure	Cannot integrate an external calculation routine

of an external calculation routine and for specifying the runtime algorithms in a coupling procedure do not exist in C-Coupler2. We intend to recover this capability in future versions of C-Coupler.

As C-Coupler1’s coupling configuration interface has been significantly changed in C-Coupler2, C-Coupler2 does not achieve backwards compatibility. However, keeping backwards compatibility will be a primary goal for future C-Coupler versions. Therefore, we will try to make C-Coupler versions support all existing C-Coupler2 APIs, so that the coupled models using C-Coupler2 can be conveniently upgraded with future C-Coupler versions without code changes.

Although the results in Sect. 5 indicate that the initialization cost and memory use of C-Coupler2 may be affordable

in most cases, a problem might arise when the model resolution or the number of processor cores is extremely high. When developing future C-Coupler versions, we will investigate ways to decrease initialization cost and memory use.

Code availability. C-Coupler2 is an open-source coupler that is always free for non-commercial activities. The latest version and future updates of the source code, user guide, and examples can be downloaded from <https://github.com/C-Coupler-Group/c-coupler-lib> (last access: 24 August 2018). We highly recommend users to watch this project, so as to be notified of the future updates and bug fixes. Anyone who detects a new bug of C-Coupler2 and contributes to the corresponding test cases will be awarded a bonus.

We also applied a DOI for a specific code version of C-Coupler2, that can be accessed via <https://doi.org/10.5281/zenodo.1283512>.

The Supplement related to this article is available online at <https://doi.org/10.5194/gmd-11-3557-2018-supplement>.

Author contributions. LL was responsible for most aspects of C-Coupler2, including the design, code development, software testing, user manual, and this paper. CZ designed and developed the non-blocking data transfer functionality, conducted software testing based on several real coupled models, and co-led a training. RL conducted software testing based on several real coupled models, co-led a training, and manages the code release. BW contributed to the design of the main target functions. GY contributed to the evaluation of the dynamic 3-D coupling capability. All authors contributed to paper writing.

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

Acknowledgements. This work was jointly supported in part by the National Grand Fundamental Research 973 Program of China (grant no. 2014CB441302) and the National Key Research Project of China (grant nos. 2017YFC1501903 and 2016YFA0602200).

Edited by: Sophie Valcke

Reviewed by: two anonymous referees

References

- Balaji, P., Buntinas, D., Goodell, D., Gropp, W. D., Kumar, S., Lusk, E. L., Thakur, R., and Traff, J. L.: MPI on a Million Processors. Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lect. Notes Comp. Sci., 5759, 20–30, https://doi.org/10.1007/978-3-642-03770-2_9, 2009.
- Collins, W. D., Bitz, C. M., Blackmon, M. L., Bonan, G. B., Bretherton, C. S., Carton, J. A., Chang, P., Doney, S. C., Hack, J. J., Henderson, T. B., Kiehl, J. T., Large, W. G., McKenna, D. S., Santer, B. D., and Smith, R. D.: The Community Climate System Model Version 3 (CCSM3), *J. Climate*, 19, 2122–2143, 2006.
- Craig, A. P., Jacob, R. L., Kauffman, B., Bettge, T., Larson, J. W., Ong, E. T., Ding, C. H. Q., and He, Y.: CPL6: The New Extensible, High Performance Parallel Coupler for the Community Climate System Model, *Int. J. High Perform. C.*, 19, 309–327, 2005.
- Craig, A. P., Vertenstein, M., and Jacob, R.: A New Flexible Coupler for Earth System Modeling developed for CCSM4 and CESM1, *Int. J. High Perform. C.*, 26, 31–42, <https://doi.org/10.1177/1094342011428141>, 2012.
- Craig, A., Valcke, S., and Coquart, L.: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, *Geosci. Model Dev.*, 10, 3297–3308, <https://doi.org/10.5194/gmd-10-3297-2017>, 2017.
- Dennis, J. M., Edwards, J., Loy, R., Jacob, R., Mirin, A. A., Craig, A. P., and Vertenstein, M.: An application-level parallel I/O library for Earth system models, *Int. J. High Perform. C.*, 26, 43–53, 2012.
- Hanke, M., Redler, R., Holfeld, T., and Yastremsky, M.: YAC 1.2.0: new aspects for coupling software in Earth system modelling, *Geosci. Model Dev.*, 9, 2755–2769, <https://doi.org/10.5194/gmd-9-2755-2016>, 2016.
- He, J., Zhang, M., Lin, W., Colle B., Liu, P., and Vogelmann, A. M.: The WRF nested within the CESM: Simulations of a midlatitude cyclone over the southern great plains, *J. Adv. Model Earth Sy.*, 5, 611–622, 2013.
- Hill, C., DeLuca, C., Balaji, C. V., Suarez, M., and Silva, A. D.: Architecture of the Earth System Modeling Framework, *Comput. Sci. Eng.*, 6, 18–28, 2004.
- 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, 2013.
- Larson, J., Jacob, R., and Ong, E.: The Model Coupling Toolkit: A New Fortran90 Toolkit for Building Multiphysics Parallel Coupled Models, *Int. J. High Perf. Comp. App.*, 19, 277–292, <https://doi.org/10.1177/1094342005056116>, 2005.
- Li, L. J., Lin, P. F., Yu, Y. Q., Wang, B., Zhou, T. J., Liu, L., Liu, J. P., Bao, Q., Xu, S. M., Huang, W. Y., Xia, K., Pu, Y., Dong, L., Shen, S., Liu, Y. M., Hu, N., Liu, M. M., Sun, W. Q., Shi, X. J., Zheng, W. P., Wu, B., Song, M.-R., Liu, H. L., Zhang, X. H., Wu, G. X., Xue, W., Huang, X. M., Yang, G. W., Song, Z. Y., and Qiao, F. L.: The Flexible Global Ocean-Atmosphere-Land System Model: Grid-point Version 2: FGOALS-g2, *Adv. Atmos. Sci.*, 30, 543–560, 2013a.
- Li, L. J., Wang, B., Dong, L., Liu, L., Shen, S., Hu, N., Sun, W., Wang, Y., Huang, W., Shi, X., Pu, Y., and Yang, G.: Evaluation of Grid-point Atmospheric Model of IAP LASG version 2 (GAMIL2), *Adv. Atmos. Sci.*, 30, 855–867, <https://doi.org/10.1007/s00376013-2157-5>, 2013b.
- Liu, L., Yang, G., and Wang, B.: CoR: a multi-dimensional common remapping software for Earth System Models, in: The Second Workshop on Coupling Technologies for Earth System Models (CW2013), available at: <https://wiki.cc.gatech.edu/CW2013/index.php/Program> (last access: 8 May 2014), 2013.
- Liu, L., Yang, G., Wang, B., Zhang, C., Li, R., Zhang, Z., Ji, Y., and Wang, L.: C-Coupler1: a Chinese community coupler for Earth system modeling, *Geosci. Model Dev.*, 7, 2281–2302, <https://doi.org/10.5194/gmd-7-2281-2014>, 2014.
- Long, M. S., Yantosca, R., Nielsen, J. E., Keller, C. A., da Silva, A., Sulprizio, M. P., Pawson, S., and Jacob, D. J.: Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry module for Earth system models, *Geosci. Model Dev.*, 8, 595–602, <https://doi.org/10.5194/gmd-8-595-2015>, 2015.
- Jones, P.: Conservative remapping: First- and second-order conservative remapping, *Mon. Weather Rev.*, 127, 2204–2210, [https://doi.org/10.1175/1520-0493\(1999\)127<2204:FASOCR>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<2204:FASOCR>2.0.CO;2), 1999.

- Qiao, F., Yuan, Y., Yang, Y., Zheng, Q., Xia, C., and Ma, J.: Wave-induced mixing in the upper ocean: Distribution and application in a global ocean circulation model, *Geophys. Res. Lett.*, 31, L11303, <https://doi.org/10.1029/2004GL019824>, 2004.
- Redler, R., Valcke, S., and Ritzdorf, H.: OASIS4 – a coupling software for next generation earth system modelling, *Geosci. Model Dev.*, 3, 87–104, <https://doi.org/10.5194/gmd-3-87-2010>, 2010.
- Valcke, S.: The OASIS3 coupler: a European climate modelling community software, *Geosci. Model Dev.*, 6, 373–388, <https://doi.org/10.5194/gmd-6-373-2013>, 2013.
- Wang, G., Qiao, F., and Xia, C.: Parallelization of a coupled wave-circulation model and its application, *Ocean Dynam.*, 60, 331–339, 2010.
- Wang, W., Barker, D., Bray, J., Bruyère, C., Duda, M., Dudhia, J., Gill, D., and Michalakes, J.: WRF Version 3 Modeling System User's Guide, available at: http://www2.mmm.ucar.edu/wrf/users/docs/user_guide_V3/contents.html (last access: 24 August 2018), 2014.
- Xin, X.-G., Wu, T.-W., and Zhang, J.: Introduction of CMIP5 experiments carried out with the climate system models of Beijing Climate Center, *Adv. Clim. Change Res.*, 4, 41–49, <https://doi.org/10.3724/SP.J.1248.2013.041>, 2013.
- Yang, Y., Qiao, F., Zhao, W., Teng, Y., and Yuan, Y.: MASNUM ocean wave numerical model in spherical coordinates and its application, *Acta Oceanol. Sin.*, 27, 1–7, 2005.
- Zhao, B., Qiao, F., Cavaleri, L., Wang, G., Bertotti, L., and Liu, L.: Sensitivity of typhoon modeling to surface waves and rainfall, *J. Geophys. Res.-Oceans*, 122, 1702–1723, <https://doi.org/10.1002/2016JC012262>, 2017.